



THÈSE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE

Délivré par :

Institut Supérieur de l'Aéronautique et de l'Espace

Présentée et soutenue par :

Paul ESCANDE

le lundi 26 septembre 2016

Titre :

Compression et inférence des opérateurs intégraux : applications à la
restauration d'images dégradées par des flous variables

Approximation and estimation of integral operators : applications to the
restoration of images degraded by spatially varying blurs

École doctorale et discipline ou spécialité :

ED MITT : Mathématiques appliquées

Unité de recherche :

Hors équipe d'accueil ISAE-ONERA et ICA, précisez : ITAV

Directeur(s) de Thèse :

M. Jérémie BIGOT (directeur de thèse)

M. Pierre WEISS (co-directeur de thèse)

Jury :

Jean-François AUJOL Professeur des Universités - Président du Jury

Jérémie BIGOT Professeur des Universités - Directeur

Rémi GRIBONVAL Directeur de Recherche - Rapporteur

Marc HOFFMANN Professeur des Universités - Membre du Jury

Denis MATIGNON Professeur des Universités - Membre du Jury

Jean-Luc STARCK Directeur de Recherche - Membre du Jury

Bruno TORRÉSANI Professeur des Universités - Rapporteur

Pierre WEISS Chargé de Recherche - Co-directeur

À Christiane,
À Jacqueline,
mes Mamies.

Remerciements

Chères toutes, chers tous,

Il est maintenant temps de tourner la page (pas celle-ci !). En trois années, un nombre incroyable de personnes, elles aussi incroyables, ont défilé à mes cotés, au labo et en dehors. Toutes ces personnes ont contribué à mon épanouissement et à l'aboutissement de ce travail dans les meilleures conditions. Il est grand temps de les remercier !

Mes tout premiers remerciements vont à mes deux directeurs Jérémie et Pierre. Ils m'ont tout d'abord fait confiance pour tenter cette aventure. Leur talent, dynamisme, fraîcheur d'esprit et bienveillance m'ont été précieux ! Pierre a dû me supporter quotidiennement (et vice versa) mais il m'a toujours reçu dans la bonne humeur (et bonne musique), il m'a apporté son aide et ses conseils avisés (non vestimentaires). Au delà de l'aspect scientifique, il me tient à cœur de souligner la générosité et la sincérité de Pierre et Jérémie qui se sont toujours démenés pour moi. Pierre et Jérémie sont deux personnes merveilleuses. Merci pour tout ce que vous m'avez apporté, scientifiquement et humainement !

Je remercie le PRES (Pôle de Recherche et d'Enseignement Supérieur) et la Région de Toulouse qui ont financé cette thèse au travers du projet MODIM (Modélisation de données en grande dimension pour la déconvolution à flou variable des images) vaillamment défendu par Pierre et Jérémie, encore eux !

J'aimerais ensuite exprimer ma profonde reconnaissance à Rémi Gribonval et Bruno Torrèsani qui ont accepté très volontiers de rapporter ma thèse. Je souhaite aussi les remercier pour la lecture très attentive du manuscrit ainsi que pour leurs questions intéressées, commentaires avisés et suggestions passionnantes. Je suis honoré que Jean-François Aujol, Marc Hoffmann, Denis Matignon et Jean-Luc Starck aient accepté de se joindre à eux pour constituer un jury fantastique !

Un grand merci à tout ceux et celles qui ont pris parti dans l'aventure PRIMO team qui avait la malédiction de se prémunir d'un préfixe malicieux 'dé'. En particulier, Jérôme, Wenxing, Jonas et Frédéric qui animaient régulièrement les journées par leur enthousiasme, leur grande culture et humour ; et mes aînés, la grande sœur Claire et le demi-grand-frère Nicolas pour leur présence stimulante et les moments partagés en conf ou ailleurs !

J'en viens maintenant aux collègues de l'ITAV, mon laboratoire d'accueil ! Merci à Bernard de m'avoir accepté en ses murs, à Corinne et Aurore pour cette collaboration. Aurore a montré un talent pour fournir des exercices de débogage inédits. Merci à Manon

qui a réalisé une étude numérique préliminaire d'un des chapitres de cette thèse. Merci à Jacques pour l'organisation des sessions de Rugby et celles d'un autre genre ! Et si je ne parlais pas de l'ultimate, qui animait le chef et tous ses stagiaires (qui se voyaient apposer une close de participation exclusive dans leur convention), je ne pourrais pas quitter le labo sans peur de représailles ! L'ultimate c'est génial ! Merci à Camille et à Baptistine pour avoir toujours gérer avec brio, les missions déclarées au dernier moment. Merci aux copines chics, gourmandes, joviales, zens: Ludivine et Nathalie, pour avoir agrémenter la vie quotidienne au labo. Bon courage à Sophie, Aurélie, Ludivine et Fabien, pour leur thèse ! Valentin, je te souhaite beaucoup de courage avec Pierre, mais tu vas te régaler ! Merci à tous ceux que j'oublie !

Je remercie également tous mes amis du BSC ! Les entraînements, matchs et autres festivités ont animé avec grand bonheur toutes ces années. Trop d'anecdotes sont à résumer, mais on retiendra les tops modèles de chez Tatïe: Dams et Dums, Bertrand de chez Boss et les pédiatres sans frontières Akram et Aurel. Merci à Romain pour les dégustations de vins et à Jean-Marc pour ses vanes d'un autre monde ! Merci aux petits nouveaux Sélim et Hugo qui prennent toujours le temps pour une dernière. Les copines, je vous oublie pas: Julie, Anne-cé, Gab, Sophie, Vivie et Carole. Vous allez me manquer !

Maintenant, le creeew: n.m. association de ~~malfaiteurs~~ à visée la déconnade ayant pour membres: Flo, Lulu, Lisa, Coco, Thomas, Manu, KK, Grugru, Fitia et Gérard. Merci pour tous les weekends/vacances passés ensemble depuis tant d'années ! C'est toujours un plaisir immense d'être avec vous. Merci à Flo pour avoir pris soin de ma condition physique, Coco et Thomas pour les fameuses pâtes aux palourdes, Lulu pour son roulé de makki légendaire, Manu pour son accueil à Nice, Grugru pour ses retards, KK pour la cuisson parfaite des écrevisses, et Fitia et Gérard parce qu'il begins to well done !

Merci à mes parents pour leur soutien sans faille, de m'avoir donné goût à la curiosité et à la persévérance. A mon petit frère grand, Romain, qui nous aura régales avec ses petits plats ! À Valérie pour m'avoir accueilli (avec de succulents nectars) de nombreuses fois quand le vent m'amenait à Paname. Merci à Philippe pour l'accueil à Bordeaux, mais la prochaine fois évite le dimanche soir !!

Finalement, les ultimes vont à Lisa, qui, à elle seule embellit toutes mes journées, mêmes les plus dures. Mille mercis !

The restoration of images degraded by spatially varying blurs is a problem of increasing importance. It is encountered in many applications such as astronomy, computer vision and fluorescence microscopy where images can be of size $1000 \times 1000 \times 1000$ pixels. Variable blurs can be modelled by linear integral operators H that map a sharp image u to its blurred version Hu defined by

$$Hu(x) = \int_{\Omega} K(x, y)u(y)dy, \quad \forall x \in \Omega = [0, 1]^d$$

where $K : \Omega \times \Omega \rightarrow \mathbb{R}$ is called the kernel. After discretization of the image on a grid of N pixels, H can be viewed as a matrix of size $N \times N$. For targeted applications, matrices contain 10^{18} coefficients. This simple observation illustrates the difficulties associated to this problem: i) the storage of a huge amount of data, ii) the prohibitive computation costs of matrix-vector products. This problem suffers from the challenging curse of dimensionality. In addition, in many applications, the operator is usually unknown or only partially known. There are therefore two different problems, the approximation and the estimation of blurring operators. They are intricate *and* have to be addressed with a global overview.

Most of the work of this thesis is dedicated to the development of new models and computational methods to address those issues.

In a first part, this work studied the approximation methods of integral operators. Existing approaches in the literature can be separated in two classes. The most developed approaches consists in constructing a low rank decomposition of the kernel K . Similarly to the seminal work of Beylkin, Coifman and Rokhlin, we studied the representation of blurring operators in wavelet bases. We showed that matrix-vector products can be rapidly computed in $O(N\epsilon^{-M/d})$ with a precision ϵ in spectral norm where M is a scalar describing the regularity of the kernel. This type of approximations can be used to design restoration algorithms that are two orders of magnitude faster than current methods.

The second class contains methods commonly used in the imaging community. They consist in constructing a low rank decomposition of the Time Varying Impulse Response (TVIR) $T : \Omega \times \Omega \rightarrow \mathbb{R}$ defined by $T(x, y) = K(x + y, y)$. Until now, those methods were partially studied and this work bridges the gap in the comprehension of their performances. Moreover, it allowed the identification of a representation that “super”

compresses the operator. This representation appeals for the development of new identification strategies.

In a second part, this work addresses the challenging problem of the estimation of operators. Recent theoretical works studied this problem but none of them can be implemented in targeted applications. In the specific case where some scattered impulse responses are observed, this work proposes the construction of an estimator of the operator that can be evaluated numerically in large dimensions. Theoretical guarantees on the performance of the estimator are also provided.

Finally, this thesis studied other imaging problems. Images from light-sheet microscopy are degraded by stripe shaped attenuations. These phenomena can be modelled by multiplicative structured noises. This work proposes to solve a convex optimization and convincing results are obtained on real data.

The development of quantitative indices that measure the similarity of images is a challenging question in imaging. The illumination of a scene can vary between two moments. Most indices are non invariant to these variations and will fail to assess the similarity of the same scene between the two instants. We proposed a similarity measure that is invariant to illumination changes.

Le problème de restauration d'images dégradées par des flous variables connaît un attrait croissant et touche plusieurs domaines tels que l'astronomie, la vision par ordinateur et la microscopie à feuille de lumière où les images sont de taille $1000 \times 1000 \times 1000$ pixels. Les flous variables peuvent être modélisés par des opérateurs intégraux qui associent à une image nette u , une image floue Hu définie par

$$Hu(x) = \int_{\Omega} K(x, y)u(y)dy, \quad \forall x \in \Omega = [0, 1]^d$$

où $K : \Omega \times \Omega \rightarrow \mathbb{R}$ est appelé le noyau. Une fois discrétisé pour être appliqué sur des images de N pixels, l'opérateur H peut être vu comme une matrice de taille $N \times N$. Pour les applications visées, la matrice contient 10^{18} coefficients. On voit apparaître ici les difficultés liées à ce problème de restauration des images qui sont i) le stockage de ce grand volume de données, ii) les coûts de calculs prohibitifs des produits matrice-vecteur. Ce problème souffre du fléau de la dimension. D'autre part, dans beaucoup d'applications, l'opérateur de flou n'est pas ou que partialement connu. Il y a donc deux problèmes complémentaires mais étroitement liés qui sont l'approximation et l'estimation des opérateurs de flou.

Cette thèse a consisté à développer des nouveaux modèles et méthodes numériques permettant de traiter ces problèmes.

Dans une première partie, ce travail a étudié les méthodes d'approximation des opérateurs intégraux. Elles peuvent être distinguées en deux groupes. Le plus étoffé contient les méthodes qui construisent une décomposition de rang faible du noyau K . Similairement aux travaux fondateurs du groupe de Beylkin, Coifman et Rokhlin, nous avons étudié les représentations des opérateurs de flou dans des bases d'ondelettes. Nous avons montré que les produits matrice-vecteur Hu peuvent être appliqués en $O(N\epsilon^{-M/d})$ avec une précision ϵ où M est un scalaire décrivant la régularité du noyau. Ces approximations peuvent donner lieu à des temps de restauration des images réduits de deux ordres de grandeurs.

Les méthodes du deuxième groupe, couramment utilisées dans la communauté du traitement d'image, construisent une décomposition de rang faible de la Time Varying Impulse Response (TVIR) $T : \Omega \times \Omega \rightarrow \mathbb{R}$ définie par $T(x, y) = K(x + y, y)$. Elles sont restées jusqu'à présent non-étudiées et ce travail a permis de combler un manque dans

la compréhension de leurs performances. De plus, il a permis d'identifier des représentations qui compressent considérablement les opérateurs et rendent possible l'estimation de l'opérateur.

Dans une deuxième partie, ce travail aborde le délicat problème d'estimation des opérateurs. Récemment, de nouveaux travaux théoriques s'y sont intéressés, cependant aucun d'eux ne peut être implémenté en pratique. Dans le cas où quelques réponses impulsives arbitrairement réparties dans l'espace sont connues, ce travail propose une construction d'un estimateur de l'opérateur qui soit numériquement applicable en grande dimension. Nous donnons aussi des garanties théoriques sur ses performances.

Dans un troisième temps, cette thèse étudie d'autres types de dégradation. Les images issues de la microscopie à feuille de lumière sont altérées par des atténuations en forme de raies. Ce phénomène peut-être modélisé par un bruit multiplicatif structuré. Ce travail propose un modèle de résolution convexe qui rend le problème soluble aisément et qui donne des résultats probants sur des données réelles.

Les indices de comparaison de deux images consistent une question délicate. Entre deux instants, une scène peut être illuminée différemment. Beaucoup d'indices sont sensibles à ces variations et échoueront à reconnaître la similarité de la scène entre deux instants. Nous avons proposé une mesure de similarité des images qui est invariante aux changements d'illumination.

Contents

1	Introduction	1
I	Approximation of operators	49
2	Approximation of integral operators using convolution-product expansions	51
3	Sparse wavelet representations of spatially varying blurring operators	77
4	Accelerating $\ell^1 - \ell^2$ deblurring using wavelet expansions of operators	121
II	Estimation of operators	155
5	Identification of linear operators from scattered impulse responses	157
III	Imaging problems	181
6	A Variational Model for Multiplicative Structured Noise Removal	183
7	Contrast Invariant SNR	211
8	High-resolution in-depth imaging of optically cleared thick samples using an adaptive SPIM	231
	Activités Scientifiques	247
	Bibliographie	251

1

Introduction

Les problématiques de ma thèse sont motivées par une nouvelle technique d'imagerie à fluorescence appelée Selective Plane Illumination Microscope (SPIM) [234]. Ce nouveau microscope permet d'imager des échantillons vivants en 3D et en temps réel. Ceci permet de mieux comprendre les mécanismes cellulaires mis en œuvre dans la croissance des tumeurs, l'embryologie, l'action de médicaments, etc...

Les images produites par ce microscope sont très grandes : une taille typique est $1000 \times 1000 \times 1000$ voxels. Des phénomènes complexes de propagation de la lumière (diffusion, réfraction et atténuation) dégradent les images. Celles-ci souffrent notamment de flous variables et de bruits multiplicatifs structurés. Voir Figure 1.1 pour quelques exemples.




Figure 1.1 – Trois exemples d'images obtenues grâce au SPIM. A gauche, un plan XZ d'un sphéroïde HCT116 dégradé par un flou variable. Au centre, un plan XY d'un derme de peau humaine dégradé par un bruit multiplicatif structuré (raies d'atténuation). A droite, un plan XZ d'un sphéroïde HCT116 dégradé par un flou variable et un bruit multiplicatif structuré.

Ma thèse comporte des contributions dans quatre domaines :

Approximation d'opérateurs linéaires intégraux : Les opérateurs de flou peuvent être modélisés par des opérateurs intégraux H qui, à une image nette $u \in L^2(\Omega)$, associent une image floue $Hu \in L^2(\Omega)$, définie pour tout point x d'un domaine

d -dimensionnel $\Omega = [0, 1]^d$ par :

$$Hu(x) = \int_{\Omega} K(x, y)u(y)dy, \quad (1.1)$$

où $K : \Omega \times \Omega \rightarrow \mathbb{R}$ est appelé le noyau de l'opérateur intégral. Pour être appliqué sur des images de N pixels l'opérateur doit être discrétisé. Après cette étape, l'opérateur H peut être vu comme une matrice de taille $N \times N$. Pour les applications ciblées, la matrice contient 10^{18} coefficients et un produit matrice-vecteur peut être calculé en $O(N^2)$ opérations, ce qui est gigantesque ! On dit que le problème souffre du *fléau de la dimension*.

Dans cette thèse, nous cherchons à apporter des solutions pour faire face à ce déluge de données et proposer des algorithmes de restauration rapides.

Estimation de opérateurs intégraux : Dans de nombreuses applications, l'opérateur de flou est inconnu. Dans le cas de la microscopie à feuille de lumière, les dégradations sont très souvent dépendantes de l'échantillon. En astronomie, les turbulences atmosphériques varient au cours du temps. Avant toute tentative de restauration, l'opérateur doit être estimé à partir de quelques réponses impulsionnelles réparties de manière arbitraire. A notre connaissance, la plupart des algorithmes existants sont inutilisables du fait de la dimensionnalité du problème. Nous proposons une approche fonctionnelle avec de solides garanties théoriques.

Restauration d'images dégradées par des bruits multiplicatifs et structurés :

Les images issues du SPIM sont dégradées par des raies d'absorption qui peuvent être modélisées comme des bruits multiplicatifs structurés. Jusqu'à présent, certains auteurs avaient proposé des méthodes de restauration de bruits i.i.d. multiplicatifs, tandis que d'autres traitaient le cas de bruits structurés additifs. En collaboration avec Pierre Weiss et Wenxing Zhang, nous avons proposé un nouveau modèle de restauration convexe qui traite les deux problèmes simultanément et offre des résultats convaincants sur données réelles.

Mesure de similarité d'images invariante aux contrastes locaux : Les mesures de similarité entre images couramment utilisées (SNR, SSIM, ...) sont sensibles aux variations d'illumination. En collaboration avec Pierre Weiss et Yiqiu Dong, nous avons élaboré une mesure invariante à ces changements. Elle repose sur un problème d'optimisation convexe permettant de trouver une image possédant un arbre de composantes fixé, la plus proche d'une image de référence. L'arbre des composantes permet de coder la relation d'inclusion entre les composantes connexes des ensembles de niveau d'une image et de définir mathématiquement la notion de changement de contraste local.

Les problématiques d'approximation et d'estimation d'opérateurs constituent le cœur de cette thèse. Ces deux problèmes peuvent paraître orthogonaux mais ils doivent être pensés et résolus de manière unifiée.

Dans cette introduction, nous présenterons d’abord quelques applications qui pourraient bénéficier de mes travaux de thèse. Ceci nous permettra aussi de justifier nos choix de modélisation. Dans une deuxième phase, nous verrons quelles sont les méthodes utilisées dans la littérature, leurs défauts et pourquoi nous avons élaboré de nouvelles approches. Finalement, nous présenterons les contributions que cette thèse a apporté aux différents domaines. Les buts de cette introduction sont de donner une vue globale - avec suffisamment de détails - des travaux menés et d’orienter le lecteur vers le chapitre adéquat pour une description plus fine du propos.

1.1 Motivations de mes travaux : Imagerie, microscopie à feuille de lumière

Les origines de ma thèse sont motivées par les problématiques associées à une nouvelle technologie d’imagerie. Dans cette section, je détaille ces motivations.

1.1.1 Contexte

Pendant une longue période, les biologistes travaillaient avec des modèles cellulaires bi-dimensionnels, les boîtes de pétri par exemple. Ces modèles arrivent à leurs limites lors de l’étude d’ensembles plus complexes comme les tumeurs. Dans ces corps on observe notamment un gradient de diffusion des nutriments et une hétérogénéité de l’état cellulaire (les cellules centrales sont en nécrose). Ces phénomènes ne peuvent pas être reproduits en 2D. Ainsi, à l’ITAV ¹ les biologistes ont mis au point un modèle 3D de tumeur, appelé sphéroïde, qui rend compte de ces particularités.

Les modèles 3D sont plus délicats à imager car leur intérieur est difficilement accessible. Pour y remédier, plusieurs type de microscopes à fluorescence ont été mis au point. Le principe général est d’exciter des fluorophores intégrés dans les cellules à l’aide d’un laser et d’observer la réponse lumineuse de l’échantillon. Dans les années 1980, les microscopes confocaux [174] ont fait leur apparition commerciale. Ces microscopes reconstruisent une image à partir de l’excitation point par point de l’échantillon. Ils ont plusieurs inconvénients, les plus importants étant le fait que le rayon laser traverse l’échantillon jusqu’au point de focalisation et que les voies d’émission et d’excitation sont identiques. Cela implique une photo-toxicité importante, une faible profondeur de pénétration et un grand étalement de la réponse impulsionnelle dans la direction axiale. Ces effets peuvent être partiellement réduits en utilisant une voie d’émission bi-photonique [86]. Deux rayons lasers de faible énergie sont combinés pour exciter le fluorophore, il en résulte une photo-toxicité plus faible, une profondeur de pénétration plus grande et une réponse impulsionnelle d’étalement axial plus petit.

En 1993, l’idée du microscope à feuille de lumière émerge [234]. Dans cette modalité d’imagerie, les voies d’émission et de détection sont perpendiculaires et une feuille de lumière créée à l’aide d’une lentille cylindrique est utilisée pour exciter une coupe entière

¹Institut des Technologies Avancées en sciences du Vivant

de l'échantillon. La photo-toxicité s'en trouve réduite et les temps de capture d'une image sont accélérés par des facteurs allant de 100 à 1000. Ce microscope est donc bien plus adapté au suivi en temps réel d'un amas de cellules se développant sous l'action de médicaments, contraintes mécaniques, etc... Les images produites souffrent cependant de plusieurs défauts. Premièrement, la feuille de lumière générée par la lentille cylindrique n'est pas plane, elle est hyperbolique.

En son centre, la feuille de lumière est fine alors que quand la distance au foyer augmente, elle devient de plus en plus épaisse, voir Figure 1.2. L'excitation d'un volume de plus en plus grand de fluorophores contribue à créer un phénomène de flou variable sur le plan XZ illustré Figure 1.1. Deuxièmement, la complexité des échantillons observés et l'hétérogénéité des indices de réfractions induisent des dégradations de la feuille de lumière qui sont dépendantes de l'échantillon observé.

Ces deux phénomènes font qu'il est impossible de connaître l'opérateur de dégradation de façon précise. Des stratégies peuvent être élaborées pour obtenir des informations locales sur l'opérateur. Par exemple, des micro-billes (pouvant être assimilées à des diracs) peuvent être ajoutées aux échantillons. Leur image permet alors d'obtenir la réponse impulsionnelle de l'opérateur de flou à la position de bille [142] voir Figure 1.3. L'insertion de micro-billes est néanmoins difficile à mettre en œuvre et celles-ci peuvent perturber le développement de l'échantillon. D'autres méthodes plus avancées ont été mises au point et consistent à exciter avec un laser certains fluorophores émettant dans une fréquence différente. Ceci permet d'observer la réponse impulsionnelle à des positions arbitraires [11].

Ces problématiques se retrouvent aussi en astronomie. Les turbulences atmosphériques perturbent la propagation de la lumière et varient au cours du temps. Certaines étoiles dites "étoiles guides" peuvent être assimilées à des diracs et donnent une information locale sur le flou, voir Figure 1.4.

D'autres phénomènes apparaissent dans le cas particulier de la microscopie à feuille de lumière : les raies d'absorption. Lors de son trajet, la lumière peut être absorbée par des éléments de l'échantillon. Une ombre portée se crée sur le contenu suivant voir Figure 1.1 pour une illustration. Ces raies brident les performances des méthodes d'analyse des images. Il convient alors de développer des méthodes permettant de corriger ces phénomènes.

1.1.2 Modélisation

Flous variables

Les flous variables sont modélisés par des opérateurs linéaires continus qui - à une image nette $u \in L^2(\Omega)$ - associent une image floue $Hu \in L^2(\Omega)$. Le théorème de Schwartz implique que n'importe quel opérateur linéaire peut être représenté par une équation intégrale de la forme (1.1). Il est donc important de spécifier la classe de noyaux K rencontrés en pratique, afin de développer des méthodes adaptées. Une grande partie des flous de la microscopie ont des réponses impulsionnelles (abrégées RI dans la suite) satisfaisant les propriétés suivantes :




Figure 1.2 – Deux images illustrant le fonctionnement du SPIM. A gauche, une schématisation de la feuille de lumière en bleu se propageant dans le sphéroïde (schéma de Aurore Masson). A droite, une reconstruction 3D d'un intestin imagé à l'aide d'un SPIM. Cette image est issue d'une collaboration entre StromaLab et la plateforme imagerie de l'ITAV (Jacques Rouquette et Lise Tesseydre).

Décroissance spatiale Bien souvent les réponses impulsionnelles ont un support borné (indicatrice des capteurs CCD) ou au moins ont une décroissance rapide (tâche de Airy, flou gaussien, ...)

Régularité des RI Les réponses impulsionnelles sont lisses, i.e. les fonctions $K(\cdot, y)$ sont régulières pour tout $y \in \Omega$.

Régularité des variations des RI Les réponses impulsionnelles varient de manière régulière. La taille de la réponse impulsionnelle dépend de la profondeur de l'échantillon que la lumière a dû traverser, cette quantité est intrinsèquement continue.

Ces trois propriétés permettent de proposer une classe de noyaux $\mathcal{A}(M, f)$ d'intérêt. On note $W^{M,p}(\Omega)$ l'espace de Sobolev des fonctions aux M premières dérivées dans $L^p(\Omega)$, voir section 3.2 p.81.

Définition 1.1.1. Soient $M \in \mathbb{N}$ et $f : [0, 1] \rightarrow \mathbb{R}_+$ une fonction bornée décroissante. La classe $\mathcal{A}(M, f)$ est constituée des noyaux $K \in W^{M,\infty}(\Omega \times \Omega)$ qui satisfont les propriétés de décroissance suivantes

$$\begin{aligned} \forall |\alpha| \leq M, \forall (x, y) \in \Omega \times \Omega, & \quad |\partial_x^\alpha K(x, y)| \leq f(\|x - y\|_\infty) \\ \forall |\alpha| \leq M, \forall (x, y) \in \Omega \times \Omega, & \quad |\partial_y^\alpha K(x, y)| \leq f(\|x - y\|_\infty). \end{aligned}$$




Figure 1.3 – Des micro-billes sont insérées dans un sphéroïde [142] et permettent d’obtenir une information sur la réponse impulsionnelle à différentes positions.

Tous les opérateurs intégraux admettant un noyau $K \in \mathcal{A}(M, f)$ sont des opérateurs linéaires continus de $L^2(\Omega) \rightarrow W^{M,\infty}(\Omega)$. En effet, la définition de K permet d’appliquer la convergence dominée puis l’inégalité de Hölder qui donnent

$$\begin{aligned} |\partial_x^\alpha H u(x)| &\leq \int_{\Omega} |\partial_x^\alpha K(x, y)| |u(y)| dy \\ &\leq \|\partial_x^\alpha K(x, \cdot)\|_{L^\infty(\Omega)} \|u\|_{L^1(\Omega)}, \end{aligned}$$

pour $|\alpha| \leq M$. Ainsi, $\|\partial^\alpha H u\|_{L^\infty(\Omega)} \leq |\Omega|^{1/2} \|K\|_{W^{M,\infty}(\Omega \times \Omega)} \|u\|_{L^2(\Omega)}$. Cette propriété décrit bien les opérateurs de flous qui sont des opérateurs régularisants et qui - à une image quelconque de $L^2(\Omega)$ - associent une image dans $W^{M,\infty}(\Omega)$. Cet espace décrit des fonctions très régulières : il est formé de fonctions dont les dérivées sont lipschitziennes jusqu’à l’ordre $M - 1$.

Raies d’absorption

Les raies d’absorption peuvent être modélisées comme des bruits à la fois *multiplicatifs* et *structurés*. L’image observée u_0 peut être modélisée comme $u_0 = u \odot \eta$ où u est l’image “idéale”, $\odot$ est la multiplication entre deux fonctions et η le bruit. Ce dernier peut - dans le cas de la microscopie - être modélisé comme un processus stochastique stationnaire $\eta = \psi \star \lambda$ où λ est un bruit i.i.d. d’une certaine loi et ψ est un filtre de convolution tenant compte de la structure du bruit.

Le problème de la restauration des images qui sont à la fois dégradées par des flous variables et par des raies d’absorption est un problème inverse complexe. Idéalement, il faudrait développer des méthodes qui unifient ces deux problématiques. Dans le temps imparti de cette thèse, nous n’avons pu étudier ces problèmes que séparément.




Figure 1.4 – Simulation d'étoiles observées à l'aide du télescope Hubble [180]. Chaque tache encadrée d'un carré est une étoile observée. On y observe la variation de la réponse impulsionnelle.

1.1.3 Difficultés numériques

Pour être appliqué à des images de N pixels l'opérateur doit être discrétisé. Après cette étape, l'opérateur H peut être vu comme une matrice de taille $N \times N$. Pour les applications ciblées, la matrice peut contenir de l'ordre de 10^{18} coefficients et un produit matrice-vecteur peut être calculé de façon naïve en $O(N^2)$ opérations.

La restauration des images floues est un problème inverse mal posé. Les méthodes de résolution sont souvent itératives et nécessitent d'appliquer l'opérateur H et son adjoint H^* à chaque itération. Le nombre d'opérations nécessaire est donc complètement prohibitif avec des approches naïves. Pour illustrer le propos, considérons une hypothétique machine ayant un processeur calculant 3 milliards d'opérations par seconde (3GFlops). Dans ces conditions, un produit matrice-vecteur est calculé en un temps de l'ordre de la dizaine d'années. L'utilisation d'une architecture parallèle permet - dans une certaine mesure - de réduire ce temps de calcul, mais pas suffisamment. Une des contributions principale de cette thèse est le développement de nouvelles approches de compression et de calcul de l'opérateur.

1.2 Applications similaires

Les problèmes liés à l'approximation et à l'estimation des opérateurs intégraux en grande dimension se retrouvent dans des applications très variées. L'une des principales raisons de cette omniprésence est le théorème de noyaux de Schwartz. Soit $D(\Omega)$ l'espace des fonctions $\mathcal{C}^\infty(\Omega)$ à support compact et $D'(\Omega)$ son dual, l'espace des distributions. Ce

théorème affirme que toute application linéaire continue de $D(\Omega) \rightarrow D'(\Omega)$ peut s'écrire sous la forme (1.1), du moins si on autorise le noyau K à être une distributions sur $\Omega \times \Omega$. De plus, ce noyau K est déterminé de manière unique [208]. En utilisant un argument de densité, on peut étendre ce résultat à tous les espaces rencontrés en pratique : les espaces de distributions, les espaces de Lebesgue, de Sobolev, etc... Un opérateur linéaire intégral est l'analogue continu de la représentation matricielle des applications linéaires.

1.2.1 Systèmes de communications

Les opérateurs linéaires intégraux de type (1.1) apparaissent aussi dans les systèmes de communications. L'estimation de filtres linéaires variants temporellement est étudiée depuis les années 1950 (voir e.g. [143]). Ce problème connaît un regain d'intérêt avec l'OFDM (Orthogonal frequency-division multiplexing) [253], une technologie de transfert de données qui est utilisée dans beaucoup de systèmes de communications tels que l'ADSL, certaines normes Wi-Fi [159] ou les réseaux téléphoniques 4G [220].

Dans les réseaux de communications sans-fil, les informations sont propagées sous forme d'ondes électro-magnétiques dans l'environnement ambiant et sont sujettes à beaucoup de dégradations :

Affaiblissement par trajets multiples les environnements complexes incluant les bâtiments, montagnes ou océans réfléchissent les signaux et une antenne peut recevoir plusieurs versions retardées du signal émis. L'information reçue est donc une somme du signal d'intérêt et de différentes versions retardées par les réflexions.

Parasites atmosphériques En grande partie, ce sont les décharges électriques produites par la foudre et en moindre mesure les rayonnements électromagnétiques naturellement présents dans l'atmosphère.

Effet Doppler Lorsque les émetteurs et récepteurs sont en mouvement relatif, le signal subit une distorsion fréquentielle.

Avant l'arrivée de l'OFDM, les signaux étaient transportés sur une seule onde porteuse d'une unique fréquence. Ce fonctionnement était très sensible aux dégradations évoquées précédemment. L'OFDM est un système qui transporte les données sur plusieurs sous-porteuses et permet de réduire les effets dus aux dégradations. Pour ces raisons, cette technologie est un domaine de recherche en effervescence, notamment dans le domaine des mathématiques appliquées [219, 17, 136]. L'idée de base de l'OFDM est de construire un ensemble de fonctions sous-porteuses qui sont des translations et modulations d'une seule fonction de référence ψ :

$$\psi_{j,m}(t) = \psi(t - jT)e^{2i\pi mFt}.$$

Dans cette équation, T est la période du symbole et F est la séparation fréquentielle. L'orthogonalité du système de fonctions $\psi_{j,m}(t)$ dépend des choix de ψ , de T et de F . On voit apparaître un lien entre l'OFDM et les trames de Gabor et on en déduit que

pour obtenir une reconstruction parfaite, la condition $TF \geq 1$ doit être satisfaite [148]. Le signal envoyé par le système est alors décomposé comme

$$s(t) = \sum_{j,m} c_{j,m} \psi_{j,m}, \quad (1.2)$$

et le signal reçu est $r = Hs + \epsilon$ où ϵ est une perturbation stochastique de moyenne nulle i.i.d. et H est l'opérateur décrivant le système de communication. Il a été modélisé comme étant un opérateur dispersif temporellement et fréquentiellement, voir [22] pour différentes modélisations. Ces opérateurs ont des noyaux qui sont la somme de diracs et d'autres fonctions bornées ayant une décroissance hors-diagonale [17]. Ce sont donc des noyaux très similaires à ceux rencontrés en microscopie à la différence qu'une somme de diracs peut apparaître, ceci est exprimé par la condition

$$\operatorname{ess\,sup}_{x,y \in \mathbb{R}^d} |K(x,y)|v(x-y) < +\infty \quad (1.3)$$

où $v : \mathbb{R}^d \rightarrow \mathbb{R}_+$ est une fonction de poids. Ces fonctions de poids sont choisies avec précaution de sorte que l'espace des opérateurs de noyaux K satisfaisant (1.3), muni d'une norme appropriée, forme une algèbre unitaire d'opérateurs continus de $L^2(\mathbb{R}^d)$ dans $L^2(\mathbb{R}^d)$.

1.2.2 Opérateurs pseudo-différentiels

Des opérateurs similaires apparaissent dans l'étude des équations aux dérivées partielles : les opérateurs pseudo-différentiels. Ces opérateurs ont été initialement introduits dans l'étude des EDPs elliptiques. Un exemple typique est l'équation de Poisson qui décrit le potentiel ϕ d'un champ causé par une densité de charge ou de masse f connue : $\Delta\phi = f$ où Δ est l'opérateur Laplacien. Elle apparaît par exemple en électrostatique et en astronomie.

Une EDP peut s'écrire sous la forme

$$Pu = \sum_{|\alpha| \leq m} a_\alpha \cdot \partial^\alpha u = f, \quad (1.4)$$

où les $a_\alpha : \mathbb{R}^d \rightarrow \mathbb{R}$ et P est un opérateur différentiel d'ordre m . Sur $\mathbb{R}^d$ et sous certaines conditions de régularité sur les fonctions a_α , il est possible d'écrire P sous la forme :

$$Pu(x) = \int_{\mathbb{R}^d} a(x, \xi) \hat{u}(\xi) e^{2i\pi\xi \cdot x} d\xi, \quad (1.5)$$

où $\hat{u}$ est la transformée de Fourier de u définie par

$$\hat{u}(\xi) = \int_{\mathbb{R}^d} u(x) e^{-2i\pi\xi \cdot x} dx, \quad (1.6)$$

et a est le symbole de l'opérateur différentiel P défini par

$$a(x, \xi) = \sum_{|\alpha| \leq m} a_\alpha(x) (2i\pi\xi)^\alpha. \quad (1.7)$$

qui est déduit de l'expression des dérivées dans le domaine de Fourier. L'idée directrice de la théorie des opérateurs différentiels est d'étudier les propriétés de P à partir des propriétés de son symbole a , qui est polynomial en la variable de Fourier $\xi \in \mathbb{R}^d$. Par exemple, on dit que l'opérateur est elliptique [227] si le cône

$$C(x) = \left\{ \xi \in \mathbb{R}^d \mid \sum_{|\alpha|=m} a_\alpha(x) (2i\pi\xi)^\alpha = 0 \right\} \quad (1.8)$$

contient au plus un point, l'origine. C'est-à-dire que P est elliptique si tous les opérateurs à coefficients constants obtenus en fixant une valeur de x sont elliptiques. On remarque que dans le cadre d'une EDP elliptique d'ordre 2, de forme canonique

$$\sum_{i,j} a_{i,j}(x) \frac{\partial^2 u(x)}{\partial x_i \partial x_j} = f, \quad (1.9)$$

on retrouve bien la condition que les matrices $A(x) = (a_{i,j}(x))_{i,j}$ soient définies positives pour tout x .

Les opérateurs pseudo-différentiels sont des outils qui ont été introduits pour généraliser les opérateurs différentiels pour plusieurs raisons :

- Les opérateurs différentiels avec des coefficients $a_\alpha \in C^\infty$ bornés peuvent être vus comme des opérateurs pseudo-différentiels.
- L'inverse d'un opérateur différentiel elliptique d'ordre m est un opérateur pseudo-différentiel d'ordre $-m$ [227]. Ces outils permettent donc de traiter les opérateurs et leurs inverses sous le même point de vue.
- Ils généralisent les opérateurs singuliers.

Par analogie à (1.7), on définit la classe des opérateurs pseudo-différentiels d'ordre m - classiquement notée par $\Psi_{\rho,\delta}^m$ avec $0 \leq \rho, \delta \leq 1$ - qui contient tous les opérateurs de la forme (1.5) qui ont des symboles $a \in C^\infty$ tels que :

$$|\partial_\xi^\alpha \partial_x^\beta a(x, \xi)| \leq C_{\alpha,\beta} (1 + |\xi|)^{m-\rho|\alpha|-\delta|\beta|}, \quad (1.10)$$

pour tout multi-indice α, β , et tout couple (x, ξ) . La constante $C_{\alpha,\beta}$ est indépendante de x et ξ .

En vertu du théorème de Schwartz, les opérateurs pseudo-différentiels de classe $\Psi_{\rho,\delta}^m$ ont un lien très étroit avec les opérateurs intégraux de la forme (1.1). On peut montrer [217, 227] que les distributions K associées aux opérateurs de $\Psi_{1,0}^m$ coïncident avec une fonction qui est C^∞ en dehors de la diagonale et qui a une décroissance de type :

$$|\partial_x^\beta \partial_y^\alpha K(x, y)| \leq A_{\alpha,\beta,d} \|x - y\|^{-d-m-\rho|\alpha|-\delta|\beta|}, \quad (1.11)$$

pour tout multi-index α, β . On remarque alors que l'opérateur est un opérateur singulier avec une singularité d'ordre m sur la diagonale. Ainsi, les opérateurs pseudo-différentiels généralisent les opérateurs intégraux singuliers.

Le formalisme d'opérateurs pseudo-différentiels permet d'étudier plus facilement la continuité d'un opérateur. En effet, les opérateurs de $\Psi_{\rho,\delta}^m$ peuvent être prolongés en des opérateurs linéaires continus de $H^s(\mathbb{R}^d) \rightarrow H^{s-m}(\mathbb{R}^d)$ pour tout $x \in \mathbb{R}$ [227]. Un autre résultat permet de donner des conditions d'extension des $\Psi_{\rho,\delta}^m$ par des opérateurs continus de $L^2(\mathbb{R}^d) \rightarrow L^2(\mathbb{R}^d)$ [225]. Les indices doivent satisfaire $m \leq 0$, $0 \leq \rho \leq 1$, $0 \leq \delta < 1$ et $\rho - \delta - \frac{m}{d} \geq 0$. Ce résultat est en défaut dès que $\delta = 1$, en effet la classe $\Psi_{1,1}^m$ contient des membres qui ne peuvent pas être étendus par des opérateurs continus $L^2(\mathbb{R}^d) \rightarrow L^2(\mathbb{R}^d)$ [134].

Cette classe très particulière $\Psi_{1,1}^0$ contient la classe des opérateurs dit de Calderón-Zygmund [41, 42, 171] qui ont un noyau satisfaisant l'inégalité :

$$|\partial_x^\alpha K(x, y)| + |\partial_y^\alpha K(x, y)| \lesssim \|x - y\|^{-d-|\alpha|}, \quad (1.12)$$

pour tout multi-indice $|\alpha| \leq 1$. Pour assurer la continuité des opérateurs de cette classe, le théorème $T(1)$ [80] requiert trois hypothèses supplémentaires : i) l'opérateur est faiblement continu, ii) la distribution $H(1)$ doit être dans BMO , l'espace des fonctions à oscillations moyennes bornées, iii) $H^*(1)$ doit être aussi dans BMO . Ce résultat peut aussi se généraliser par le $T(b)$ théorème [81]. L'étude de cette classe d'opérateurs a motivé Yves Meyer à construire des bases orthonormales d'ondelettes. Ces travaux sont un outil central de cette thèse.

1.2.3 Équations aux dérivées partielles hyperboliques – Géophysique

Les phénomènes de propagations sont modélisés généralement par des EDPs hyperboliques. L'exemple récurrent concerne la propagation des ondes. L'étude des sols est un problème de grande importance pour toutes les sociétés pétrolières en recherche d'un nouveau gisement à exploiter, et dans une moindre mesure pour les études scientifiques nécessaires à la compréhension des mécanismes internes de la Terre [66, 250]. La surface de la planète est stratifiée en couches géologiques de différentes natures (sable, argile, calcaire, marnes, granite,...) qui ont chacune une densité spécifique et donc une vitesse de propagation des ondes différentes. Les réservoirs de pétrole potentiels sont des volumes piégés par une couche imperméable (argile par exemple) et qui contiennent une roche poreuse. Les géophysiciens s'intéressent à la localisation de ces gisements en étudiant la réponse du sol à des ondes de chocs. Suivant la géométrie du sol, les ondes vont se réfléchir et se réfracter aux changements de couche et seront recueillies à la surface par des capteurs sismiques, voir Figure 1.5. A partir de ces informations, les géophysiciens sont en mesure de repérer les gisements potentiels.

La physique de propagation des ondes est modélisée par l'équation différentielle

$$\begin{cases} \frac{\partial^2 u}{\partial t^2}(t, x) - c^2(x)\Delta u(t, x) + f(t, x), \\ u(0, x) = u_0(x), \\ \frac{\partial u}{\partial t}(0, x) = u_1(x), \end{cases} \quad (1.13)$$




Figure 1.5 – Schématisation d'une prospection géophysique sur Terre. Les couches de différentes couleurs représentent la stratification de la croûte terrestre. Des chocs ou des vibrations sonores sont générés (étoile rouge) et se propagent dans les couches géophysiques. Les ondes peuvent être réfractées ou réfléchies et sont recueillies en surface par les capteurs en vert.

que l'on regarde par simplicité dans un domaine sans bord. L'équation des ondes est une EDP linéaire hyperbolique. Quand c est une fonction lisse, les solutions de cette équation sont des ondes se propageant à la vitesse $c(x)$ au point x . L'existence et l'unicité des solutions sont garanties dans le cadre très général où $\log c \in L^\infty$ [156].

Lors de la prospection, des capteurs sont placés à la surface du sol, disons en $(x_i)_{i=1}^n$, et mesurent la solution u après la sollicitation du sol à une stimulation. Afin d'identifier les gisements, l'idée est de retrouver la fonction c à partir des fonctions $(u(\cdot, x_i))_{i=1}^n$ qui peuvent être bruitées. On note $\mathcal{F}$ l'opérateur $c \mapsto (u(\cdot, x_i))_{i=1}^n$. Les géophysiciens sont donc en présence d'un problème inverse mal posé et non linéaire [222]. Les méthodes de résolution sont itératives et mettent en œuvre les calculs avec $\mathcal{F}$ et son adjoint $\mathcal{F}^*$. En sismique 3D, les fonctions c sont typiquement de taille 10^9 et les données collectées $(u(\cdot, x_i))_{i=1}^n$ de taille 10^{12} . On voit apparaître ici aussi le grand nombre d'opérations nécessaires à la résolution du problème inverse, puisque chaque calcul $\mathcal{F}(c)$ ou $\mathcal{F}^*(c)$ nécessite la résolution d'une EDP (1.13) de très grande dimension. Il devient alors crucial de trouver des méthodes de calcul qui permettent de déterminer la solution de l'EDP rapidement. Le calcul de ces solutions font intervenir des opérateurs Fourier intégraux (FIO) de la forme :

$$Hf(x) = \int_{\mathbb{R}^d} a(x, \xi) e^{2i\pi\Phi(x, \xi)} \hat{f}(\xi) d\xi \quad (1.14)$$

où Φ est la fonction de phase et a le symbole de l'opérateur. Ces opérateurs sont fréquemment rencontrés dans la théorie des EDP, l'analyse et le calcul numérique. Ils

ont été largement étudié théoriquement [133], mais des algorithmes rapides de produits matrice-vecteur ont été développés que récemment [84, 44]. Les FIOs apparaissent dans une grande classe d'EDP hyperbolique, les conditions standards sur les phases et symboles sont [217] :

- Φ est C^∞ , 1-homogène en ξ et satisfait la condition de non dégénérescence $|\det \nabla_x \nabla_\xi \Phi| > 0$ uniformément sur x et ξ .
- a est un symbole d'ordre M i.e. a est C^∞ et satisfait $|\partial_\xi^\alpha \partial_x^\beta a(x, \xi)| \leq C_{\alpha, \beta} (1 + |\xi|)^{M - |\alpha|}$.

Les FIO généralisent les opérateurs pseudo-différentiels et intégraux : pour obtenir un opérateur pseudo-différentiel il suffit de choisir la fonction de phase $\Phi(x, \xi) = x \cdot \xi$.

1.2.4 Résumé

Les opérateurs linéaires intégraux considérés dans cette thèse sont des opérateurs assez simples faisant partie de classes d'opérateurs largement étudiées (Pseudo-différentiels, Calderón-Zygmund, Fourier intégraux). De nombreux résultats existent pour ces classes très larges. Dans cette thèse, nous exploiterons certaines propriétés spécifiques des noyaux de la classe $\mathcal{A}(M, f)$ pour construire des méthodes plus adaptées.

1.3 Un horizon des méthodes de calcul rapides des opérateurs intégraux

Beaucoup de méthodes ont été développées pour évaluer des intégrales de type (1.1). Après discrétisation le problème revient à évaluer des produits matrice-vecteur Hu . L'idée générale est de s'autoriser une petite erreur contrôlée sur le résultat, qui permet de réduire grandement le nombre d'opérations nécessaires. Ces méthodes de calculs sont donc très liées à la théorie de l'approximation. On peut catégoriser ces méthodes en trois types :

- celles qui sont trop spécifiques et qui ne peuvent pas être utilisées dans nos applications pratiques.
- celles qui construisent des approximations de rang faible des noyaux K ,
- celles qui construisent des approximations de rang faible de la "Time Varying Impulse Response" (TVIRs) $T : \Omega \times \Omega \rightarrow \mathbb{R}$ définie par $T(x, y) = K(x + y, y)$,

1.3.1 Discrétisation de l'opérateur

Pour pouvoir appliquer l'opérateur numériquement il faut au préalable le discrétiser. Les méthodes de Galerkin forment une classe de méthodes qui permettent d'approcher un problème en dimension infinie par un problème en dimension finie. Elles sont utilisées

notamment dans les méthodes d'éléments finis et éléments finis de frontière. L'idée est de choisir une base $(\phi_i)_{1 \leq i \leq N}$ d'un sous espace vectoriel V_N de l'espace de $L^2(\Omega)$. Un opérateur $H : L^2(\Omega) \rightarrow L^2(\Omega)$ peut alors se discrétiser en une matrice $\mathbf{H}^N \in \mathbb{R}^{N \times N}$ définie par $\mathbf{H}_{i,j}^N = \langle H\phi_j, \phi_i \rangle$. Notons $S^N : H \rightarrow \mathbf{H}^N$ l'application qui discrétise l'opérateur. A partir d'une matrice $\mathbf{H}^N$ on peut reconstruire un opérateur H^N en utilisant par exemple la pseudo-inverse $S^{N,+}$ de S^N . On note $H^N = S^{N,+}(\mathbf{H}^N)$. Dans le cas particulier d'une base orthonormale $(\phi_i)_{1 \leq i \leq N}$, l'opérateur H^N est donné par :

$$H^N u = \sum_{1 \leq i,j \leq N} \mathbf{H}_{i,j}^N \langle \phi_i, u \rangle \phi_j. \quad (1.15)$$

La base $(\phi_i)_{1 \leq i \leq N}$ peut-être choisie de manière à obtenir de bonnes propriétés d'approximation, ou pour décrire le système d'acquisition. Par exemple, les images sont souvent acquises sur une grille de pixels et dans ce cas les ϕ_i seront les indicatrices de carrés pavant le domaine Ω . Tout l'enjeu est maintenant d'être capable de calculer rapidement un produit matrice-vecteur $\mathbf{H}\mathbf{u}$ où $\mathbf{u} \in \mathbb{R}^N$ est un vecteur arbitraire.

1.3.2 Opérateurs invariants par translation : convolutions

Dans le cas particulier où l'opérateur est invariant par translation, i.e. $K(x, y) = k(x - y)$, on peut utiliser la transformée de Fourier pour évaluer rapidement des produits matrice-vecteur $\mathbf{H}\mathbf{u}$. En effet, $\mathbf{H}\mathbf{u} = \mathbf{k} \star \mathbf{u} = \mathbf{F}^*(\hat{\mathbf{k}} \cdot \hat{\mathbf{u}})$, où $\mathbf{F} \in \mathbb{C}^{N \times N}$ est la transformée de Fourier discrète et $\hat{\mathbf{k}} = \mathbf{F}\mathbf{k}$. En utilisant la transformée de Fourier rapide (FFT), la matrice $\mathbf{H}$ pourra donc être appliquée en $O(N \log N)$ opérations [74].

Dès que l'opérateur a des réponses impulsionnelles qui varient dans le champ, cette méthode n'est plus applicable et il faut trouver d'autres stratégies.

1.3.3 Quelques méthodes d'approximation

Dans cette partie, je décris quelques méthodes couramment utilisées en imagerie pour approcher des produits matrice-vecteur. J'ai étudié ces méthodes en début de thèse et elles se sont avérées trop peu précises ou trop spécifiques pour être utiles dans les problèmes qui nous intéressaient.

Matrices creuses et largeur de bande

Les opérateurs de flous de la classe $\mathcal{A}(M, f)$ ont des réponses impulsionnelles qui ont une décroissance connue ou des supports compacts. Cette dernière propriété peut être modélisée en choisissant la fonction $f = \mathbb{1}_{[0, \kappa]}$ où $\kappa \in (0, 1]$ est le rayon maximal des réponses impulsionnelles de l'opérateur et de l'adjoint.

L'approche la plus simple qui permet de réduire les coûts de calculs des produits matrice-vecteur est d'exploiter cette propriété de support compact en utilisant une matrice creuse. Cela permet d'obtenir une complexité de $O(\kappa^d N)$ opérations. C'est une méthode qui est très simple à mettre en place, mais elle devient très vite inutilisable quand la taille de l'image et des réponses impulsionnelles deviennent grandes, c'est à dire dans beaucoup de cas pratiques.

Composition d'un difféomorphisme et d'une convolution

L'une des premières alternatives proposée permettant de réduire la complexité des produits matrice-vecteur consiste à appliquer un difféomorphisme au domaine de l'image Ω suivi d'une convolution implémentée avec des FFTs et le difféomorphisme inverse [203, 204, 168, 223]. Le difféomorphisme est choisi de sorte à transformer l'opérateur de flou variable en une convolution. Cette approche souffre de deux inconvénients importants :

- premièrement, il a été montré que tous les opérateurs de flou variable ne pouvaient être approchés de cette façon [168],
- deuxièmement, cette méthode nécessite des techniques d'interpolations précises sur des grilles euclidiennes très fines pour correctement coder les difféomorphismes.

Approximations séparables des noyaux

Les approximations séparables des noyaux ont été décrites par exemple dans [5, 144]. L'idée est d'approcher le noyau K par un noyau séparable $\tilde{K}$ de la forme :

$$\tilde{K}(x, y) = \prod_{k=1}^d \tilde{K}_k(x_k, y_k) \quad (1.16)$$

où chaque $\tilde{K}_k : [0, 1] \times [0, 1] \rightarrow \mathbb{R}$ est un noyau agissant dans une seule dimension. Sous cette hypothèse, l'opérateur approchant $\tilde{H}$ peut-être décomposé comme le produit de d opérateurs agissant sur une dimension $\tilde{H}_k$ comme $\tilde{H} : \tilde{H}_1 \circ \dots \circ \tilde{H}_d$ avec

$$\tilde{H}_k u(x) = \int_{y_k \in [0, 1]} \tilde{K}_k(x_k, y_k) u(x_1, \dots, y_k, \dots, x_d) dy_k, \quad (1.17)$$

pour tout $x \in \Omega$. En discret maintenant, cette structure séparable des noyaux permet de calculer $\tilde{\mathbf{H}}\mathbf{u}$ en $O(d\kappa N^{1+1/d})$.

L'hypothèse de séparabilité implique que les réponses impulsionnelles ainsi que leurs variations sont séparables. Malheureusement, la plupart des réponses impulsionnelles rencontrées en pratique ne sont pas séparables voir Figure 3.3. Néanmoins, en microscopie à fluorescence des travaux utilisent ce type d'approximation [251, 195, 158, 24].

1.3.4 Changement dans une base fixe

Le principe des méthodes décrites ici est d'exprimer la matrice $\mathbf{H}$ dans une autre base, ce qui permet de concentrer l'information sur peu de coefficients. Un seuillage adapté permet alors de coder une approximation de $\mathbf{H}$ sous forme creuse.

Principes et enjeux

Soit $\Psi \in \mathbb{R}^{N \times N}$ une matrice orthonormale. La matrice $\mathbf{H}$ peut s'écrire sous la forme $\mathbf{H} = \Psi \Theta \Psi^*$ où $\Theta \in \mathbb{R}^{N \times N}$ est le représentant de $\mathbf{H}$ dans la base Ψ . Si la base Ψ est

bien choisie, la majorité de l'information peut être concentrée sur un petit nombre de coefficients de Θ . On rend de cette manière le noyau compressible. Par exemple, il est bien connu que les ondelettes ont de très bonnes propriétés d'approximation linéaires ou non linéaires de grandes classes de fonctions.

Ce changement de base peut aussi s'interpréter comme une approximation de rang faible du noyau. En effet, $\mathbf{H} = \Psi\Theta\Psi^*$ se traduit par

$$\mathbf{H}[i, j] = \sum_{(\lambda, \mu) \in \Lambda^2} \Theta[\lambda, \mu] \Psi_\lambda[i] \Psi_\mu[j] \quad (1.18)$$

où Ψ_λ est la colonne indexée par λ de la matrice Ψ et Λ est un ensemble d'indices pouvant être identifié à $\{1, \dots, N\}$. Construire une approximation de rang faible du noyau consiste à trouver un ensemble $\Xi \subset \Lambda^2$ de faible cardinalité, de sorte que

$$\tilde{\mathbf{H}}[i, j] = \sum_{(\lambda, \mu) \in \Xi} \Theta[\lambda, \mu] \Psi_\lambda[i] \Psi_\mu[j]. \quad (1.19)$$

La matrice $\Psi_\lambda \Psi_\mu^T$ est une matrice de rang un et si le cardinal de Ξ est faible, la forme (1.19) correspond à une approximation de rang faible de $\mathbf{H}$. De plus, si les fonctions Ψ_λ sont à support compact, cette décomposition peut être interprétée comme une approximation locale de rang faible. Cette idée - illustrée sur la Figure 1.6 - constitue le fondement des méthodes parmi les plus efficaces à l'heure actuelle.

Quelques questions interviennent naturellement lors de l'étude de ces décompositions :

1. Quel est le nombre d'opérations à réaliser pour obtenir une précision ϵ sur une norme matricielle $\|\mathbf{H} - \tilde{\mathbf{H}}\|$? Peut-on relier cette erreur d'approximation à une norme d'opérateur $\|H - S^{N,+}(\tilde{\mathbf{H}})\|$?
2. Quelle est la meilleure base Ψ pour approcher un opérateur donné ? Une classe d'opérateurs ?
3. Comment choisir la base Ψ pour obtenir simultanément une bonne compression et pouvoir effectuer des produits matrice-vecteur rapidement ?

Toutes ces questions sont étroitement liées et doivent être prises en compte simultanément. La difficulté est d'obtenir un bon compromis entre temps de calcul et capacité de compression.

Décomposition en bases d'ondelettes

Les bases d'ondelettes sont des bases inconditionnelles de nombreux espaces fonctionnels. Elles permettent d'approcher efficacement les éléments de nombreuses classes de fonctions, telles que les espaces de Sobolev, l'espace BV, les espaces de Besov, ... Elles sont par exemple au cœur du standard de compression Jpeg2000. Bien que leur utilisation

dans le domaine de l'imagerie se limite actuellement à l'analyse des images, nous montrons dans cette thèse qu'elles présentent aussi un grand intérêt pour la représentation des opérateurs.

Les idées explorées remontent aux travaux fondateurs de Yves Meyer [169]. Dans ce livre, l'auteur montre de façon élégante la continuité des opérateurs de Calderón-Zygmund grâce aux bases d'ondelettes orthonormales. En 1D, un opérateur de Calderón-Zygmund est un opérateur intégral dont le noyau K satisfait les conditions

$$|\partial_x^M K(x, y)| + |\partial_y^M K(x, y)| \lesssim \|x - y\|^{-(M+1)}, \quad (1.20)$$

pour tout $x \neq y$. Soit $(\psi_{j,m})_{j,m \in \mathbb{Z}}$ une base d'ondelettes à M moments nuls et à support compact de $L^2(\mathbb{R})$. Les coefficients $\Theta_{j,m,k,n} = \langle H\psi_{j,m}, \psi_{k,n} \rangle$ de la matrice Θ satisfont - sous la condition (1.20) et sous les hypothèses du Théorème $T(1)$ - l'inégalité suivante :

$$|\Theta_{j,m,k,n}| \lesssim 2^{-|j-k|(M+1/2)} \left(\frac{2^{-j} + 2^{-k}}{2^{-j} + 2^{-k} + |2^{-j}m + 2^{-k}n|} \right)^{M+1}. \quad (1.21)$$

La continuité des opérateurs de Calderón-Zygmund peut être démontrée assez facilement en utilisant cette inégalité et le test de Schur [207].

Les travaux de Gregory Beylkin, Ronald Coifman et Vladimir Rokhlin [29] ont ensuite ouvert la voie sur le développement de méthodes numériques efficaces reposant sur la décomposition d'opérateurs dans des bases d'ondelettes. L'inégalité (1.21) permet de construire une ϵ -approximation $\tilde{\mathbf{H}}$ de $\mathbf{H}$ - au sens où $\|\mathbf{H} - \tilde{\mathbf{H}}\|_{2 \rightarrow 2} \leq \epsilon$ - pour laquelle des produits matrice-vecteur peuvent être évalués en $O(N \log N \epsilon^{-d/M})$ opérations. Cette approximation est obtenue en seillant de manière astucieuse les coefficients de Θ .

Ces résultats ont été utilisés dans beaucoup de domaines d'applications où les opérateurs singuliers et pseudo-différentiels sont rencontrés. Ils ont été affinés dans de nombreux ouvrages [140, 3, 77, 69, 67]. En plus de la parcimonie, les décompositions des opérateurs dans les bases d'ondelettes permettent de construire des préconditionneurs diagonaux efficaces qui sont très importants pour la résolution des systèmes linéaires par méthodes itératives. Ces deux avantages ont contribué au succès des ondelettes dans la compression des opérateurs et dans la résolution de problèmes essentiellement rencontrés dans le domaine des EDP. Elles n'ont pas connu un tel succès dans la communauté du traitement des images.

Dans notre cas d'opérateurs de flous de la classe $\mathcal{A}(M, f)$ les ondelettes paraissent avoir un bon potentiel : les images sont fortement compressibles dans le domaine des ondelettes et elles sont en général définies sur des grilles cartésiennes qui sont le cadre idéal pour la discrétisation des ondelettes. Plusieurs points doivent être investigués :

- Les noyaux des opérateurs de $\mathcal{A}(M, f)$ ne sont pas singuliers le long de la diagonal, est-ce que l'estimation (1.21) peut être améliorée ?
- Quelle est le nombre d'opérations nécessaires pour un produit matrice-vecteur de précision ϵ ?




Figure 1.6 – Illustration de l’interprétation en approximation multi-échelle de rang faible dans des bases d’ondelettes. Pour générer cette image, le noyau d’un opérateur a été décomposé dans une base d’ondelettes, puis les coefficients ont été seuillés à un certain niveau. Les rectangles rouges représentent les supports des ondelettes $\psi_\lambda \otimes \psi_\mu$ correspondant aux coefficients $\Theta_{\lambda,\mu}$ non-seuillés. On remarque que les zones de variations lentes sont approchées par un grand bloc de rang faible et que la diagonale est fortement décomposée.

1.3.5 Approximations locales de rang faible des noyaux : méthodes SVD

Soit $\mathbf{H} \in \mathbb{R}^{N \times N}$ une matrice. On peut effectuer une décomposition en valeurs singulières de celle-ci de la forme $\mathbf{H} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^*$ où $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_N]$ et $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_N]$ sont deux matrices orthogonales et $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_N)$ est une matrice diagonale. Les valeurs $\sigma_1 \geq \sigma_2 \geq \dots \geq 0$ sont les valeurs singulières de $\mathbf{H}$.

Il est bien connu que la meilleure approximation de rang faible d'une matrice - au sens de la norme spectrale - peut être obtenue grâce à une décomposition en valeurs singulières. Le minimiseur du problème variationnel suivant :

$$\min_{\substack{\tilde{\mathbf{H}} \in \mathbb{R}^{N \times N} \\ \text{rang}(\tilde{\mathbf{H}}) \leq m}} \|\mathbf{H} - \tilde{\mathbf{H}}\|_{2 \rightarrow 2}. \quad (1.22)$$

est

$$\tilde{\mathbf{H}} = \sum_{k=1}^m \sigma_k \mathbf{u}_k \otimes \mathbf{v}_k. \quad (1.23)$$

Bien qu'elle soit optimale au sens décrit ci-dessus, cette décomposition souffre de plusieurs problèmes pour l'analyse numérique. Premièrement, même si les opérateurs étudiés ici sont compacts, leur rang peut être grand et il faut donc choisir de grandes valeurs de m pour obtenir une bonne approximation. Deuxièmement, les vecteurs singuliers ne sont généralement pas localisés en espace et un produit matrice-vecteur est donc calculé en $O(mN)$ opérations, ce qui peut être prohibitif. Pour faire face à ces difficultés, plusieurs méthodes ont été développées.

L'idée de base est de faire des approximations *locales* de rang faible du noyau. Une des premières méthodes de ce type est probablement les méthodes multipolaires de Leslie Greengard et Vladimir Rokhlin [199, 119]. C'est l'un des dix algorithmes les plus importants du xx^e siècle selon la Society for Industrial and Applied Mathematics (SIAM). Initialement développée dans le cas du problème à n -corps et vite généralisée aux équation d'Helmholtz, où des matrices pleines apparaissent, son principe est de grouper les contributions des champs lointains et de mieux détailler les contributions voisines. Cette construction hiérarchique permet de réduire les coûts de stockage et des produits matrice-vecteur qui peuvent se calculer en $O(N \log(1/\epsilon))$ opérations pour une précision ϵ sur un vecteur de $\mathbb{R}^N$, au lieu de $O(N^2)$.

Cette idée de hiérarchisation de l'information a été exploitée ensuite dans les approches du type *panel clustering* [125], *hierarchical matrices* $\mathcal{H}$ -matrices [124], *mosaic-skeleton method* [230] et d'autres [44].

L'hypothèse exploitée dans ces approches est la notion de régularité asymptotique du noyau de l'opérateur intégral (1.1) qui est souvent exprimée sous la forme

$$|\partial_x^\alpha \partial_y^\beta K(x, y)| \lesssim \|x - y\|^{-|\alpha| - |\beta| - m} \quad (1.24)$$

pour une constante $m > 0$, pour tout multi-indices α et β et pour tout $x, y \in \mathbb{R}^d$ en dehors de la diagonale $x \neq y$. C'est une condition similaire à la décroissance des opérateurs pseudo-différentiels $\Psi_{1,1}^m$.

Les méthodes de *panel clustering* exploitent cette propriété de la façon suivante : pour un $x \in \Omega$ et un sous-ensemble $\omega \subset \Omega$ donnés qui satisfont $\text{dist}(x, \omega) \geq C \text{diam}(\omega)$, on remplace $K(x, y)$ par son développement de Taylor en un point $y_0 \in \omega$ tronqué à l'ordre k :

$$\widetilde{K}(x, y) = \sum_{\nu=1}^k p_{\nu}(x) q_{\nu}(y). \quad (1.25)$$

Cela se traduit par une approximation de rang k au point x . Les ensembles ω peuvent être construits de manière hiérarchique. Il a été remarqué que les coefficients $p_{\nu}(x)$ peuvent être utilisés pour un voisinage du point x . On construit ainsi des approximations de rang faible des blocs de la matrice.

La force de ces méthodes est de construire un arbre de blocs de la matrice discrétisée par Galerkin $\mathbf{H}_{i,j} = \langle H\phi_i, \phi_j \rangle$, où $(\phi_i)_{i \in I}$ est une base d'un sous-espace vectoriel de $L^2(\Omega)$, bien souvent éléments finis ou éléments de frontière. On cherche à identifier les blocs d'indices $\sigma \times \tau$ tels que les matrices $(\mathbf{H}[i, j])|_{i \in \sigma, j \in \tau}$ soient de rangs faibles. Dans le cas d'une régularité asymptotique (1.24), la condition d'admissibilité $\min(\text{diam}\Omega_{\tau}, \text{diam}\Omega_{\sigma}) \leq \eta \text{dist}(\Omega_{\tau}, \Omega_{\sigma})$ où $\Omega_{\tau} = \cup_{i \in \tau} \text{supp } \phi_i$, permet de "prédire" quels seront les blocs susceptibles d'avoir une approximation de rang faible. Ce sont les blocs qui sont suffisamment séparés. Les approximations de rang faible peuvent ensuite être calculées par différentes méthodes comme les SVD tronquées, Adaptive Cross Approximation (ACA), etc.. Finalement, on obtient des arbres de décomposition en blocs de rangs faibles très similaires à celui présenté Figure 1.6.

Les méthodes de $\mathcal{H}$ -matrices permettent de réduire les coûts de stockage des matrices à $O(Nk \log N)$ où k est le rang maximal de chaque blocs. Pour obtenir une précision ϵ sur les produits matrice-vecteur, le rang doit être de l'ordre de $\log \epsilon^{-1}$. Le produit matrice-vecteur avec une ϵ -approximation s'obtient donc en $O(\log(\epsilon^{-1})N \log N)$. Ces taux sont meilleurs que ceux présentés dans cette thèse, mais reposent sur une régularité C^{∞} du noyau - en dehors de la diagonale.

De plus, même si la dépendance de la complexité en $\log \epsilon^{-1}$ est intéressante, elle nécessite de connaître le noyau K de façon analytique, et ne semble donc pas adaptée à l'identification d'opérateurs.

1.3.6 Approximations de rang faible des TVIRs

On va maintenant s'intéresser à des méthodes qui sont largement utilisées dans la communauté du traitement d'images et qui consistent à décomposer la TVIR $T : \Omega \times \Omega$ définie par $T(x, y) = K(x + y, y)$ plutôt que le noyau K . On se place dans le cadre d'un opérateur Hilbert-Schmidt, i.e. $K \in L^2(\Omega \times \Omega)$.

Sous des hypothèses de petites variations des réponses impulsives, il est naturel d'approcher localement H par une convolution. Cette idée est - de loin - la plus exploitée en imagerie [228, 180, 112, 114, 135, 132, 175, 85] et elle apparaît sous beaucoup de noms différents : *sectional methods*, les méthodes *overlap-add* et *overlap-save*, *piecewise convolutions*, *anisoplanatic convolutions*, *filter flow*, *windowed-convolutions*,...

Le principe de ces méthodes est de décomposer le domaine de l'image en sous-régions et de considérer que l'opérateur restreint à chaque région est approximativement une convolution. Les résultats sur chaque sous-domaine sont ensuite assemblés pour obtenir l'image floue. Dans sa version la plus simple, on forme une partition du domaine en carrés mais des discontinuités apparaîtront aux changements de sous-régions lors de la restauration de l'image. Les sous-régions peuvent se recouvrir pour permettre des interpolations entre elles.

Mathématiquement, ces méthodes consistent à construire une approximation H_m de H de la forme

$$H_m u = \sum_{k=1}^m h_k \star (w_k \odot u) \quad (1.26)$$

que nous appellerons produit-convolution en référence à [39] ou de la forme

$$H_m u = \sum_{k=1}^m w_k \odot (h_k \star u). \quad (1.27)$$

Les questions qui nous sont venues en étudiant ces approches sont :

1. Comment choisir les fonctions h_k et w_k ?
2. Quelle est la complexité du calcul de $H_m u$?
3. Quelle est l'erreur d'approximation $\|H - H_m\|$? Et combien d'opérations sont nécessaires pour appliquer cette approximation ?

Les deux premières questions ont été bien étudiées dans la littérature. A notre connaissance, il n'existe pas de travaux abordant le troisième point, et un chapitre de la thèse y sera dédié.

Une observation faite dans [112] permet d'interpréter les méthodes de produit-convolution comme une approximation de rang-faible de la TVIR de H . En effet, la TVIR de H_m que l'on note T_m s'écrit :

$$T_m(x, y) = \sum_{k=1}^m h_k(x) w_k(y). \quad (1.28)$$

Dans ces méthodes on cherche à construire une approximation de rang faible de T qui est en général de rang numérique plus faible que le noyau K voir Figure 1.7. Tout comme les approximations creuses dans des bases, on peut choisir h_k et w_k en utilisant une SVD de manière à minimiser la norme de Hilbert-Schmidt $\|H - H_m\|_{HS}$. Cependant, les fonctions w_k n'ont aucune raison d'être à support compact et le coût d'un produit matrice-vecteur devient trop grand. Comme décrit précédemment, nous verrons qu'il est préférable de faire des décompositions multi-niveau de la TVIR. La Figure 1.8 illustre une décomposition multi-niveau de la TVIR. Cette illustration permet d'interpréter géométriquement le comportement de la méthode et de la comparer au décomposition multi-niveau des noyaux.




Figure 1.7 – Différents noyaux K (ligne du haut) avec leur TVIR T associées (ligne du bas). On observe que le rang de la TVIR est plus faible que celui du noyau K .

1.3.7 Résumé

Le problème de l'approximation d'opérateurs a été étudié depuis de nombreuses années dans les domaines de l'analyse numérique, des EDP (ondes, diffusion, ...) ou de la physique computationnelle. De façon surprenante, ces techniques ne semblent pas avoir été exploitées dans le domaine du traitement des images - et en particulier des opérateurs de flou variable.

Les méthodes les plus efficaces pour effectuer rapidement des produits matrice-vecteur sont probablement les $\mathcal{H}$ -matrices qui reposent sur une connaissance parfaite de l'opérateur. Dans cette thèse, nous avons délibérément évité ce type de méthode, dans l'optique de pouvoir traiter des problèmes de défloutage aveugle où un problème central est l'identification de l'opérateur.

Les méthodes reposant sur la décomposition d'une *classe d'opérateurs* dans une *base fixe* (e.g. les ondelettes) semblent plus pertinentes dans les problèmes d'identification et c'est donc l'approche que nous avons privilégiée. Nous verrons de plus que cette approche permet d'exploiter le fait que l'opérateur et les images sont parcimonieux dans la même base, ce qui aura des conséquences importantes pour l'analyse numérique.

1.4 Méthodes d'estimation

Nous nous sommes intéressés dans cette thèse à des problèmes inverses dits aveugles ou semi-aveugles : on souhaite reconstruire un opérateur $\mathbf{H}$ et une image $\mathbf{u}$ à partir d'une




Figure 1.8 – Illustration d’une décomposition multi-échelle de rang faible de la TVIR visualisée sur le noyau. Ici, une bases d’ondelettes a été utilisée pour décomposer les lignes de la TVIR. Les rectangles rouges représentent les supports des fonctions $h_k \otimes w_k$. qui ont été pivotés pour obtenir une représentation similaire à la Figure 1.6

observation

$$\mathbf{u}_0 = \mathbf{H}\mathbf{u} + \mathbf{b}, \quad (1.29)$$

où $\mathbf{b}$ est une perturbation inconnue.

Nous dirons que le problème est aveugle, lorsque l'on n'a aucune information sur l'opérateur $\mathbf{H}$ en dehors d'hypothèses de régularité. Il sera dit semi-aveugle lorsqu'une informations partielle sur l'opérateur est disponible. Par exemple, on peut dans certaines applications avoir accès à son action sur quelques fonctions test.

1.4.1 Méthodes aveugles

La majorité des méthodes de résolution de problèmes inverses aveugles reposent sur la résolution d'un problème variationnel non convexe de la forme suivante [21, 55, 154, 118] :

$$(\mathbf{H}^*, \mathbf{u}^*) \in \arg \min_{(\mathbf{H}^*, \mathbf{u}^*) \in \mathbb{R}^{N \times N} \times \mathbb{R}^N} \frac{1}{2} \|\mathbf{H}\mathbf{u} - \mathbf{u}_0\|_2^2 + R_1(\mathbf{H}) + R_2(\mathbf{u}), \quad (1.30)$$

où R_1 et R_2 sont des fonctions convexes qui codent les connaissances a priori sur l'opérateur et l'image. C'est un problème non-convexe (il est dit bi-convexe, car convexe dans chaque variable séparément). Il est généralement résolu de façon approximative en utilisant des algorithmes de type minimisations alternées. Les récents progrès en optimisation non convexe donnent des garanties de convergence vers des points critiques [8, 33].

Malgré les difficultés liées à la non-convexité et au coût calculatoire important, ces méthodes peuvent donner de bons résultats d'estimation et de restauration dans certains cas comme la déconvolution avec des noyaux spécifiques (dont la transformée de Fourier ne décroît pas trop rapidement). Nous avons utilisé une technique de ce type pour résoudre un problème concret dans le dernier chapitre de cette thèse. Nous pensons cependant que ces techniques sont encore très mal comprises au niveau théorique. Des travaux ont néanmoins étudié ce problème et montrent en particulier que les formules variationnelles couramment choisies ont des solutions triviales d'aucun intérêt pratique [25].

Nous n'avons pas vraiment essayé d'exploiter cette stratégie dans le cas des flous variables, car il nous a semblé plus pertinent de d'abord améliorer l'identifiabilité du problème, en réduisant au maximum le nombre de coefficients à estimer. Nous pensons que les contributions de cette thèse devraient permettre de faire avancer ce domaine de recherche.

1.4.2 Méthodes semi-aveugles

Dans cette thèse, nous étudions le problème d'estimation d'un opérateur à partir de quelques réponses impulsionnelles réparties arbitrairement dans l'espace. C'est un problème de longue date qui a commencé dans le domaine des télécommunications [143]. Quand l'opérateur de dégradation dépend de très peu de paramètres (e.g. la variance d'une gaussienne), une méthode d'estimation paramétrique peut être mise en œuvre pour identifier le noyau K . Ces méthodes s'appliquent cependant dans des cadres assez limités. Elles peuvent difficilement être utilisées en microscopie à feuille de lumière où les

dégradations ont de grandes variations suivant les échantillons. De nombreuses méthodes exploitent plutôt une hypothèse de compressibilité (e.g. symbole de Kohn-Nirenberg à bande limitée, parcimonie dans une base,...) [149, 192, 45, 62, 191] et c’est cette approche que nous suivrons dans cette thèse.

Sondage de matrices

Les travaux [192, 62] considèrent l’estimation d’une matrice $\mathbf{A} \in \mathbb{R}^{M \times N}$ représentée par une combinaison linéaire de peu d’atomes d’un dictionnaire de matrices $\mathbf{B} = (\mathbf{B}_j)_{j=1}^l$:

$$\mathbf{A} = \sum_{j=1}^l c_j \mathbf{B}_j. \quad (1.31)$$

où $\mathbf{c} \in \mathbb{R}^l$ est un vecteur p -parcimonieux. Afin d’estimer la matrice, il “suffit” d’estimer le vecteur $\mathbf{c}$. Ces deux travaux développent des méthodologies différentes mais partent tous les deux de la connaissance d’un vecteur $\mathbf{u} \in \mathbb{R}^N$ et de l’action de $\mathbf{A}$ sur ce vecteur $\mathbf{A}\mathbf{u} \in \mathbb{R}^M$. Ils tentent de résoudre le système linéaire :

$$\mathbf{A}\mathbf{u} = \sum_{j=1}^p c_j (\mathbf{B}_j \mathbf{u}) = \mathbf{L}\mathbf{c}, \quad (1.32)$$

où $\mathbf{L} \in \mathbb{R}^{M \times l}$ a pour colonnes les vecteurs $(\mathbf{B}_j \mathbf{u})$.

Dans [62], les auteurs étudient le conditionnement de $\mathbf{L}$ dans le cas où les vecteurs $\mathbf{u}$ sont aléatoires Gaussiens i.i.d. Ils montrent notamment que si la matrice de Gram $(\langle \mathbf{B}_i, \mathbf{B}_j \rangle_F)_{1 \leq i, j \leq l}$ est bien conditionnée et que les $\mathbf{B}_j$ ont un haut rang numérique, alors choisir $N \propto p \log^2(N)$ permet d’avoir une matrice $\mathbf{L}$ bien conditionnée avec grande probabilité. Les résultats peuvent être améliorés en considérant l’action de $\mathbf{A}$ sur plusieurs vecteur $(\mathbf{u}_i)_{1 \leq i \leq q}$.

Dans [192], les auteurs utilisent d’autres idées issues de l’échantillonnage compressé [46, 92, 31, 37] et proposent de résoudre le problème d’optimisation

$$\min_{\substack{\mathbf{c} \in \mathbb{R}^l \\ \mathbf{L}\mathbf{c} = \mathbf{A}\mathbf{u}}} \|\mathbf{c}\|_1. \quad (1.33)$$

Ce problème est résolu avec un algorithme de *basis pursuit* [60] et des garanties de reconstruction exactes sont données cette fois-ci sous des conditions de cohérence et pour différents choix de matrices $\mathbf{B}_j$. Nous pensons que cette stratégie est probablement mal adaptée aux dimensions auxquelles nous sommes confrontés. De plus, nous verrons que l’on peut souvent avoir une bonne connaissance du support a priori, et son identification n’est donc pas nécessaire.

Approches de type Shannon

Les travaux [149] considèrent l’estimation d’opérateurs d’Hilbert-Schmidt qui ont un symbole de Kohn-Nirenberg à une largeur de bande limitée. Ce symbole $N : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$

est défini :

$$N(x, \xi) = \int_{\mathbb{R}^d} T(x, y) e^{-2i\pi y \cdot \xi} dy, \quad (1.34)$$

qui est la transformée de Fourier de la TVIR T selon la variable y . Cette transformation est aussi appelée la transformée de Fourier symplectique du noyau K . L'hypothèse de travail est une largeur de bande $B > 0$ telle que $\forall |\xi| > B, N(x, \xi) = 0$. Cette hypothèse implique que les fonctions $T(x, \cdot)$ sont C^∞ . Sous cette hypothèse, les variations des réponses impulsionnelles sont lisses. Des théorèmes de reconstruction exacte sont développés lorsque l'action de l'opérateur sur un peigne de diracs est connue. Ces résultats sont généralisés dans [191] à des opérateurs qui ne sont plus nécessairement Hilbert-Schmidt.

Dans ce travail, les auteurs supposent connues la somme de réponses impulsionnelles, ce qui est particulièrement pertinent pour les applications (e.g. on observe simultanément plusieurs étoiles). Cependant, les méthodes développées dans [149, 191] semblent particulièrement lourdes termes de temps de calculs pour être exploitées en très grande dimension. De plus les hypothèses de régularité sont inadaptées aux problèmes d'imagerie.

Interpolation dans des bases

Dans le cas où quelques réponses impulsionnelles sont connues il est possible d'interpoler les coefficients des réponses impulsionnelles dans une base, ce qui permet de réduire le nombre de paramètres à estimer. Cette méthode d'estimation est très courante en astronomie, voir [113] qui propose un survol de toutes ces méthodes.

Soit $(\phi_\lambda)_{\lambda \in \Lambda}$ une base orthonormale de dimension N d'un sous espace de $L^2(\Omega)$. Le principe de ces méthodes est d'estimer la fonction $F : \Lambda \times \Omega \rightarrow \mathbb{R}$ définie par

$$F(\lambda, y) = \langle T(\cdot, y) \phi_\lambda \rangle, \quad (1.35)$$

à partir de la donnée de quelques réponses impulsionnelles connues $\mathbf{F}_i \in \mathbb{R}^N$ définies par $\mathbf{F}_i[\lambda] = F(\lambda, y_i) + \epsilon_i[\lambda]$ aux positions $(y_i)_{1 \leq i \leq n}$. Les vecteurs ϵ_i sont aléatoires i.i.d. de moyenne nulle. Une fois l'estimateur $\hat{F}$ de F établi, un estimateur $\hat{T}$ de la TVIR peut-être construit par :

$$\hat{T}(x, y) = \sum_{\lambda \in \Lambda} \hat{F}(\lambda, y) \phi_\lambda(x). \quad (1.36)$$

Le point épineux de ces méthodes est de construire un estimateur de $\hat{F}$. De nombreuses méthodes ont été proposées. Beaucoup d'entre elles fonctionnent par interpolation ligne par ligne de $\hat{F}$ par différentes méthodes : interpolation polynomiale, splines, fonctions radiales, etc... [113]. Ces méthodes ont de nombreux avantages : elles peuvent être appliquées numériquement en grande dimension (contrairement aux méthodes citées plus haut) et fournissent des résultats convaincants. Cependant, beaucoup de questions sont actuellement ouvertes :

- Comment gérer le bruit dans la méthode d'estimation tout en gardant des méthodes rapides ?

- Quels sont les propriétés de l'estimateur de $\hat{T}$ ainsi reconstruit ?
- Pouvons nous donner une vitesse d'approximation du genre $\mathbb{E}\|H - \hat{H}\|_{2 \rightarrow 2}^2$? où H est l'opérateur exact et $\hat{H}$ l'opérateur de TVIR $\hat{H}$.
- Quelle est l'influence de la distribution des positions $(y_i)_{1 \leq i \leq n}$ dans l'estimation ?

Nous fournirons des réponses à ces questions dans cette thèse.

1.5 Contributions

Dans cette section, nous allons aborder les contributions de cette thèse.

1.5.1 Approximation parcimonieuse des flous variables

Modélisation des flous – Chapitre 3

On considère des images d -dimensionnelles définies sur l'hypercube $\Omega = [0, 1]^d$, vues comme des fonctions $u : \Omega \rightarrow \mathbb{R}$. Les flous variables peuvent être modélisés comme des opérateurs linéaires intégraux qui, à une image nette $u \in L^2(\Omega)$ associent une image floue $Hu \in L^2(\Omega)$ de la forme

$$Hu(x) = \int_{\Omega} K(x, y)u(y)dy, \quad (1.37)$$

où $K : \Omega \times \Omega \rightarrow \mathbb{R}$ est appelé le noyau de l'opérateur intégral. En microscopie ou plus généralement dans les systèmes d'acquisition des images les flous apparaissant peuvent être caractérisés comme des opérateurs de la classe $\mathcal{A}(M, f)$, voir définition 1.1.1.

Erreur d'approximation et complexité – Chapitre 3

La classe d'opérateurs $\mathcal{A}(M, f)$ partage des similarités avec les opérateurs de Calderón-Zygmund. Nous avons donc eu l'idée d'étudier la décomposition des opérateurs de flou dans des bases d'ondelettes.

On considère une base orthonormale d'ondelettes de $L^2(\Omega)$ que l'on note $(\psi_{\lambda})_{\lambda \in \Lambda}$ où Λ est un ensemble d'indices générique. Comme H est un opérateur linéaire entre deux espaces de Hilbert, il peut s'écrire $H = \Psi\Theta\Psi^*$ où Ψ^* est la transformée en ondelettes, Ψ son inverse et Θ est une matrice infinie, qui représente l'opérateur dans le domaine des ondelettes. Ses coefficients sont définis par :

$$\theta_{\lambda, \mu} = \langle H\psi_{\lambda}, \psi_{\mu} \rangle, \quad \forall \lambda, \mu \in \Lambda. \quad (1.38)$$

Nous pouvons aussi obtenir une borne supérieure sur les coefficients de $\theta_{\lambda, \mu}$, tout comme ceux des opérateurs de Calderón-Zygmund, voir Lemme 1.5.1.

Lemma 1.5.1 (Décroissance des $\theta_{\lambda,\mu}$). *Supposons que H soit un opérateur de noyau $K \in \mathcal{A}(M, f)$. Supposons de plus que l'ondelette mère de la base $(\psi_\lambda)_{\lambda \in \Lambda}$ soit à support compact et possède M moments nuls. Alors*

$$|\theta_{\lambda,\mu}| \lesssim 2^{-(M+d/2)||\lambda|-|\mu||} 2^{-\min(|\lambda|,|\mu|)(M+d)} f_{\lambda,\mu}, \quad (1.39)$$

où $|\lambda|$ désigne l'échelle de l'ondelette ψ_λ et $f_{\lambda,\mu} = f(\text{dist}(\text{supp } \psi_\lambda, \text{supp } \psi_\mu))$.

L'estimation proposée dans le lemme 1.5.1 est plus fine que celle proposée par Yves Meyer (1.21). En effet, elle exploite la non-singularité des noyaux et leur caractère $W^{M,\infty}$ afin d'obtenir un terme de décroissance supplémentaire $2^{-\min(|\lambda|,|\mu|)(M+d)}$. Ce terme assure une décroissance des coefficients dans les hautes échelles. C'est un résultat nécessaire à l'obtention des résultats d'approximation et de complexité, ainsi qu'au développement de préconditionneurs diagonaux.

Theorem 1.5.1. *Soit Θ la matrice de taille $N \times N$ obtenue par troncature de Θ à l'échelle $\log_2(N)/d-1$. Notons la matrice $\mathbf{H} = \Psi\Theta\Psi^*$ où Ψ^* et Ψ sont les transformées en ondelettes discrètes.*

En utilisant une règle de seuillage appropriée sur Θ , on peut construire une matrice $\tilde{\mathbf{H}}_\epsilon$ qui est une ϵ -approximation i.e. $\|\mathbf{H} - \tilde{\mathbf{H}}_\epsilon\|_{2 \rightarrow 2} \leq \epsilon$ qui peut être appliquée à un vecteur de $\mathbb{R}^N$ arbitraire en $O(N\epsilon^{-d/M})$ opérations.

Si l'on suppose en plus que les réponses impulsionnelles sont à support compact de rayon maximal κ , i.e. que f a un support compact sur $[0, \kappa]$, le nombre d'opérations devient $O(N\kappa^d\epsilon^{-d/M})$.

Ce théorème permet de distinguer l'influence de chaque paramètre dans la vitesse d'approximation, on remarque notamment que :

- La méthode d'approximation se comporte en $\epsilon^{-d/M}$ où M est la régularité de l'opérateur. Elle est capable de capter cette régularité automatiquement.
- La dépendance en κ - la taille du support des réponses impulsionnelles - apparaît dans ce théorème, ce qui est particulièrement important en traitement d'images. On pourrait obtenir des résultats similaires en utilisant la transformée de Fourier, mais la complexité serait indépendante de κ .

Structure de parcimonie – Chapitre 3

Les résultats d'approximation et de complexité du Théorème 1.5.1 sont obtenus en utilisant une règle de seuillage particulière exploitant le lemme de Schur, qui est utile pour contrôler la norme spectrale. En pratique, cette règle de seuillage ne donne pas des résultats d'approximation satisfaisants dans le domaine du traitement des images. Ceci explique probablement le fait que les décompositions des opérateurs dans des bases d'ondelettes ne soient pas plus populaires en imagerie ². Nous nous sommes donc intéressés plus précisément à la sélection des coefficients dans les matrices Θ .

²De nombreux chercheurs nous ont rapporté avoir essayé de décomposer les opérateurs dans des bases d'ondelette, sans succès.




Figure 1.9 – Une illustration de la compression d’un opérateur de flou variable dans le domaine des ondelettes. Nous considérons un opérateur agissant sur des signaux 1D avec un noyau défini sur $[0, 1]$ par $K(x, y) = \frac{1}{\sigma(y)\sqrt{2\pi}} \exp\left(-\frac{(x-y)^2}{2\sigma^2(y)}\right)$ où $\sigma(y) = 4 + 10y$. Toutes les réponses impulsionnelles sont des gaussiennes de variances croissant linéairement. En haut : à gauche $\mathbf{H}$, à droite $\mathbf{H}$ en échelle $\log_{10}$. En bas : à gauche Θ obtenue en utilisant des ondelettes de Daubechies à 10 moments nuls et un niveau de décomposition $J = 7$; à droite Θ en échelle $\log_{10}$. On remarque que la représentation en ondelettes est bien plus parcimonieuse que dans le domaine spatial. On observe de plus la décroissance capturée par le Lemme 1.5.1

Les images peuvent être décrites comme étant des éléments d'un espace de Besov $B_{1,1}^1(\Omega)$ ou de $BV(\Omega)$ l'espace des fonctions à variations bornées. Cette observation nous a motivé à introduire une norme d'opérateur de la forme suivante :

$$\|\mathbf{H}\|_{X \rightarrow 2} = \sup_{\|\mathbf{u}\|_X \leq 1} \|\mathbf{H}\mathbf{u}\|_2, \quad (1.40)$$

qui permet de mesurer l'erreur d'approximation uniformément sur une classe d'intérêt décrite par $\|\cdot\|_X$.

On souhaite donc trouver une matrice $\tilde{\mathbf{H}}_L = \Psi \mathbf{S}_L \Psi^*$ qui approche $\mathbf{H}$ au sens $\|\cdot\|_{X \rightarrow 2}$ et où $\mathbf{S}_L$ contient au pire L coefficients non-nuls. La matrice est obtenue en résolvant :

$$\min_{\substack{\mathbf{S}_L \in \mathbb{R}^{N \times N} \\ L\text{-parcimonieuse}}} \left\| \mathbf{H} - \tilde{\mathbf{H}}_L \right\|_{X \rightarrow 2}. \quad (1.41)$$

La solution de ce problème donne la meilleure matrice L -parcimonieuse $\mathbf{S}_L$ sur la boule unité définie par la norme $\|\cdot\|_X$. Les espaces $BV(\Omega)$ et de Besov $B_{1,1}^1$ peuvent être décrits par des normes ℓ^1 à poids sur les coefficients d'ondelettes [72, 164]. Ceci nous a amené à considérer une norme $\|\mathbf{u}\|_X = \|\Sigma \Psi^* \mathbf{u}\|_1$ où $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_N)$ est une matrice diagonale. En utilisant le fait que la norme $\|\cdot\|_X$ définie sur les coefficients en ondelettes et que les opérateurs sont représentés eux aussi en ondelettes, nous avons proposé un algorithme glouton qui permet de calculer la solution de (1.41) en $O(N^2 \log N)$ opérations. Moralement, les coefficients des basses échelles ont plus de poids que ceux des hautes fréquences.

Cette complexité est trop importante pour pouvoir utiliser ce résultat en pratique. Il permet néanmoins d'obtenir des résultats bien plus intéressants. Comme on connaît a priori la décroissance des coefficients $\theta_{\lambda,\mu}$ pour la classe de noyaux dans $\mathcal{A}(M, f)$, on peut trouver la meilleure structure de parcimonie uniformément sur la classe $\mathcal{A}(M, f)$. On définit $\mathcal{H}$ l'ensemble des matrices $\mathbf{H} = \Psi \Theta \Psi^*$ où Θ vérifie la décroissance correspondante, voir Lemme 1.5.1. Notons Ξ un ensemble d'indices de $N \times N$ de cardinalité L , et $\mathcal{H}(\Xi)$ l'ensemble des matrices de $\mathcal{H}$ qui ont pour support Ξ . Nous avons développé un algorithme glouton de complexité $O(N \log N)$ qui permet d'obtenir la solution de

$$\min_{|\Xi|=L} \sup_{\mathbf{H} \in \mathcal{H}} \inf_{\tilde{\mathbf{H}}_L \in \mathcal{H}(\Xi)} \left\| \mathbf{H} - \tilde{\mathbf{H}}_L \right\|_{X \rightarrow 2}. \quad (1.42)$$

La réduction de la complexité de cet algorithme vient de l'exploitation des invariances dans l'estimation de la décroissance des coefficients du Lemme 1.5.1. On note ici que l'obtention d'estimations plus fines permet d'obtenir des structures de parcimonie plus adaptées aux matrices issues des noyaux de la classe $\mathcal{A}(M, f)$. Notons que la sortie de notre algorithme dépend explicitement de M , la régularité du noyau et de f qui code sa décroissance. Le résultat est optimal, uniformément sur la classe $\mathcal{A}(M, f)$. De plus, les coefficients qui seront significatifs dans le codage des matrices sont identifiés a priori et cela permet de réduire la dimension du problème d'estimation des matrices de flous.

Ces travaux ont permis d'améliorer significativement les résultats d'approximation sur le problème direct et sur le problème de restauration, voir Figure 1.10.




Figure 1.10 – Un exemple de déconvolution (1.45) montrant l'importance de la méthode de seuillage. En haut l'image dégradée avec un flou de mouvement et un bruit Gaussien i.i.d. de niveau 5.10^{-3} le $\text{pSNR} = 17.85\text{dB}$. Les deux matrices Θ ont été construites avec le même nombre de coefficients (qui correspond à 57 opérations par pixels). Au milieu : la méthode de seuillage simple $\text{pSNR} = 23.71\text{dB}$. En bas : le seuillage pondéré (1.41), $\text{pSNR} = 24.07\text{dB}$. Au delà de l'augmentation du pSNR , on observe que les artefacts de restauration sont réduits significativement.

Résolution temps-réel de problèmes inverses – Chapitre 4

En utilisant les différents ingrédients présentés précédemment, nous avons ensuite illustré l'intérêt de ces méthodes d'approximation dans le cas de la restauration des images. Dans beaucoup d'applications, les images sont observées suivant le modèle $\mathbf{u}_0 = \mathbf{H}\mathbf{u} + \boldsymbol{\eta}$ où $\boldsymbol{\eta}$ est un bruit i.i.d. Gaussien de moyenne nulle. Dans ce contexte, une procédure de restauration usuelle consiste à trouver le minimiseur de

$$\mathbf{u}^* = \Psi \arg \min_{\mathbf{x} \in \mathbb{R}^N} \frac{1}{2} \|\mathbf{H}\Psi\mathbf{x} - \mathbf{u}_0\|_2^2 + \lambda \|\mathbf{x}\|_1, \quad (1.43)$$

où Ψ^* est la transformée en ondelette et $\lambda > 0$ est un paramètre de régularisation. Beaucoup de méthodes ont été développées pour résoudre ce genre de problème. Un exemple typique est la descente de gradient proximale accélérée qui - dans le cas de la norme $\|\cdot\|_1$ - est plus connue sous le nom de FISTA (Fast Iterative Soft Thresholding Algorithm). Ces méthodes nécessitent d'appliquer $\Psi^*\mathbf{H}^*\mathbf{H}\Psi$ à chaque itération.

Dans le cas d'une convolution $\mathbf{H} = \mathbf{h} \star \cdot$, l'idée la plus répandue est d'utiliser la transformée de Fourier afin de diagonaliser la matrice $\mathbf{H} = \mathbf{F}\mathbf{D}\mathbf{F}^*$ où $\mathbf{F}^*$ est la transformée de Fourier et $\mathbf{D} = \text{diag}(\hat{\mathbf{h}})$ est la matrice diagonale formée avec la transformée de Fourier de $\mathbf{h}$. Dans ce cas $\Psi^*\mathbf{H}^*\mathbf{H}\Psi = \Psi^*\mathbf{F}|\mathbf{D}|^2\mathbf{F}^*\Psi$. Numériquement, pendant la restauration environ 95% du temps de calcul est utilisé pour faire des aller-retours dans les domaines Fourier/ondelettes.

Ces aller-retours entre les domaines peuvent être évités en exprimant (1.43) entièrement dans le domaine des ondelettes :

$$\mathbf{u}^* = \Psi \arg \min_{\mathbf{x} \in \mathbb{R}^N} \frac{1}{2} \|\boldsymbol{\Theta}\mathbf{x} - \mathbf{x}_0\|_2^2 + \lambda \|\mathbf{x}\|_1, \quad (1.44)$$

où $\mathbf{x}_0 = \Psi^*\mathbf{u}_0$. Comme présenté précédemment, les matrices $\boldsymbol{\Theta}$ ont beaucoup de coefficients négligeables. En utilisant une règle de seuillage adaptée, (1.41) ou (1.42), on peut construire une matrice $\tilde{\boldsymbol{\Theta}}$ approchant $\boldsymbol{\Theta}$. Il devient alors raisonnable de résoudre le problème approché suivant

$$\tilde{\mathbf{u}}^* = \Psi \arg \min_{\mathbf{x} \in \mathbb{R}^N} \frac{1}{2} \|\tilde{\boldsymbol{\Theta}}\mathbf{x} - \mathbf{x}_0\|_2^2 + \lambda \|\mathbf{x}\|_1, \quad (1.45)$$

qui peut être résolu uniquement avec des produits matrice-vecteur avec $\tilde{\boldsymbol{\Theta}}^*\tilde{\boldsymbol{\Theta}}$ qui est très creuse.

Nous avons montré que la résolution de (1.45), au lieu des implémentations utilisant des FFTs, accélérerait les temps de restauration par des facteurs allant de 30 à 200 pour un résultat équivalent. A titre d'exemple, une image 1024×1024 peut être restaurée en 0.15 secondes, ce qui correspond à du temps-réel, voir Figure 1.11. Ces résultats, qui de notre point de vue constituent une petite révolution dans ce domaine, ont été rendus possibles par deux autres éléments :

La structure $\boldsymbol{\Theta}^*\boldsymbol{\Theta}$: Comme $\mathbf{H}$ est une convolution, nous avons montré que $\boldsymbol{\Theta}^*\boldsymbol{\Theta}$ avait des blocs circulants. La matrice $\boldsymbol{\Theta}^*\boldsymbol{\Theta}$ peut donc être calculée en $O((2^d - 1)JN)$

opérations, où J est le niveau de décomposition de la transformée en ondelettes, au lieu de $O(N^3)$ pour une matrice $\mathbf{H}$ quelconque.

Les préconditionneurs Les paramètres influant dans la vitesse de convergence des algorithmes comme FISTA sont encore peu connus - mais des travaux intéressants proposent des explications [155]. De la même manière que pour les méthodes itératives de type gradient conjugué, nous avons proposé des préconditionneurs qui améliorent la répartition des valeurs propres de la matrice $\Theta^* \Theta$. Les préconditionneurs doivent être diagonaux pour préserver le calcul explicite de l'opérateur proximal associé à la norme ℓ^1 . Ils permettent d'accélérer la convergence de l'algorithme par des facteurs allant de 2 à 6.

Finalement, tous les ingrédients qui rendent ces avancées possibles sont :

- la structure particulière de $\Theta^* \Theta$: parcimonie et blocs circulants,
- la parcimonie des opérateurs *et* des images dans le *même* domaine,
- des stratégies de seuillage adaptées,
- des préconditionneurs diagonaux efficaces.

1.5.2 Étude des approximations des *produit-convolutions* – Chapitre 2

Les méthodes d'approximation par *produit-convolutions* permettent d'approcher H par un opérateur H_m de la forme

$$H_m u = \sum_{k=1}^m h_k \star (w_k \odot u). \quad (1.46)$$

La condition de régularité pour pouvoir décomposer un noyau K sous la forme d'une $\mathcal{H}$ -matrice est qu'il soit asymptotiquement régulier (voir définition (1.24)). Cette hypothèse implique que $K(x, \cdot)$ et $K(\cdot, y)$ soient réguliers pour tout x et y . Nous avons montré que la condition de régularité utile pour étudier les erreurs d'approximation des méthodes de produit-convolution porte uniquement sur les fonctions $T(x, \cdot)$. Intuitivement, cette hypothèse se traduit par le fait que réponses impulsionnelles varient régulièrement.

Dans notre travail, nous avons considéré la classe suivante d'opérateurs :

Definition 1.5.1 (Classe de TVIR $\mathcal{T}^s$). Soit $\mathcal{T}^s$ l'ensemble des fonctions $T : \Omega \times \Omega \rightarrow \mathbb{R}$ satisfaisant la condition de régularité : $T(x, \cdot) \in H^s(\Omega)$, $\forall x \in \Omega$ et $\|T(x, \cdot)\|_{H^s(\Omega)}$ est uniformément bornée en x i.e. :

$$\sup_{x \in \Omega} \|T(x, \cdot)\|_{H^s(\Omega)} \leq C < +\infty. \quad (1.47)$$




Figure 1.11 – Un exemple de déconvolution d’une image issue d’un microscope confocal. En haut : l’image dégradée par un flou gaussien anisotrope et par un bruit Gaussien i.i.d. de niveau $5 \cdot 10^{-3}$, $\text{pSNR} = 23.94\text{dB}$. En bas : l’image restaurée $\text{pSNR} = 26.33\text{dB}$. La méthode classique utilisant la FFT obtient l’image en 25 secondes, la méthode proposée en 0.15 secondes (accélération $\times 166$).

Une première analyse théorique de ces méthodes

Sous ces faibles hypothèses de régularité, nous avons montré que la plupart des méthodes décrites dans la littérature ont des vitesses d'approximation de type $\|H - H_m\|_{HS} = O(\kappa^{d/2} m^{-s/d})$ où $0 < \kappa \leq 1$ désigne le côté du plus petit hyper-cube contenant le support des réponses impulsionnelles. Cette erreur d'approximation ne peut pas être améliorée uniformément sur la classe de fonctions $\mathcal{T}^s$. Nous avons de plus montré que la complexité de ces méthodes pour atteindre une précision ϵ sur la norme d'Hilbert-Schmidt d'opérateur était :

- $O\left(\kappa^{\frac{d}{2s}} N \log(N) \epsilon^{-d/s}\right)$ si les fonctions w_k n'avaient pas de contrainte de support. Par exemple, SVD, Fourier, ...
- $O\left(\kappa^{d(\frac{2s+1}{2s})} N \log(\kappa^d N) \epsilon^{-d/s}\right)$ si les fonctions w_k ont un support compact. Par exemple, B-splines, ondelettes, ... On remarque que le nombre d'opérations théoriques est plus faible d'un facteur κ^d .

Pour illustrer ces méthodes, considérons que les fonctions w_k sont choisies parmi les éléments d'une base d'ondelettes. La TVIR de l'opérateur est décomposée sur cette base :

$$T(x, y) = \sum_{\mu \in \Lambda} c_\mu(x) \psi_\mu(y), \quad (1.48)$$

où c_μ est appelé le *symbole ondelette* et est défini par $c_\mu(x) = \langle T(x, \cdot), \psi_\mu \rangle$. L'approximation de rang faible T_m est obtenue en choisissant m indices μ . On obtient des représentations de ces symboles Figure 1.12.

Proposition d'une représentation indépendante de la discrétisation

Pendant l'étude des approximations par *produit-convolution* nous avons étudié une autre représentation de la TVIR qui permet de représenter un opérateur avec très peu de coefficients. Cette idée avait été suggérée dans [28], sans analyse théorique précise. Dans le paragraphe précédent, nous n'avions supposé aucune régularité des fonctions $T(\cdot, y)$. Pour appliquer cette stratégie, nous avons introduit l'hypothèse $T \in H^{r,s}(\Omega \times \Omega)$, où

$$H^{r,s}(\Omega \times \Omega) = \left\{ T : \Omega \times \Omega \rightarrow \mathbb{R}, \partial_x^\alpha \partial_y^\beta T \in L^2(\Omega \times \Omega), \forall |\alpha| \leq r, \forall |\beta| \leq s \right\}. \quad (1.49)$$

Cette espace semble naturel pour décrire les différences de régularité entre les réponses impulsionnelles et leurs variations. La TVIR peut se décomposer sous la forme :

$$T(x, y) = \sum_{\lambda, \mu \in \Lambda} c_{\lambda, \mu} \psi_\lambda(x) \psi_\mu(y), \quad (1.50)$$

où ψ est une ondelette mère générant une base orthonormale de $L^2(\Omega)$. En utilisant cette représentation, nous avons montré qu'il était possible de construire une méthode d'approximation H_m de H s'approchant en $\|H - H_m\|_{HS} = O(m^{-s/d})$ et qui contient seulement $O(m^2)$ coefficients - au lieu de $O(Nm)$ pour les méthodes de *produit-convolution*.




Figure 1.12 – Représentation des “symboles ondelettes”. Ligne du haut les TVIR T . Ligne du bas : les c_μ associés. Les barres rouges indiquent le changement entre les échelles. Pour la TVIR de gauche, on remarque que les coefficients en ondelettes décroissent très rapidement quand les échelles augmentent. La décroissance est plus lente pour les deux autres TVIR, car moins régulières. L’adaptabilité des ondelettes peut se visualiser sur la TVIR de droite où l’on voit quelques coefficients non nuls dans les hautes échelles, ils sont concentrés autour des discontinuités. Il n’y aura que quelques couples (c_μ, ψ_μ) nécessaires pour coder cette discontinuité, ce qui n’est pas le cas pour les décomposition en Fourier ou B-splines.

existantes. Ce nombre de coefficients est *indépendant* de la dimension de la discrétisation N . Cette représentation très compacte ouvre la voie à de nouvelles méthodes innovantes d'estimation de l'opérateur, cependant, nous n'avons pas eu le temps de les explorer dans cette thèse. Voir Figure 1.13 pour une illustration de la compacité de la représentation.




Figure 1.13 – Illustration de la représentation compacte (en bas) de l'opérateur en comparaison avec les symboles ondelettes (en haut). Les nombres de coefficients non-nuls sont affichés (Nnz) pour chaque combinaison TVIR/représentation. On remarque que la représentation compacte contient très peu de coefficients significatifs.

1.5.3 Estimation des opérateurs à partir de réponses impulsionnelles éparpillées – Chapitre 5

Le cadre d'estimation choisi est le suivant. Soit $(\phi_\lambda)_{\lambda \in \Lambda}$ une base orthonormale d'un sous-espace de $L^2(\Omega)$ de dimension N . Nous avons mis au point une méthode d'interpolation de la fonction $F : \Lambda \times \Omega \rightarrow \mathbb{R}$ définie par

$$F(\lambda, y) = \langle T(\cdot, y), \phi_\lambda \rangle, \quad (1.51)$$

à partir de la donnée de quelques réponses impulsionnelles connues $\mathbf{F}_i \in \mathbb{R}^N$ définies par $\mathbf{F}_i[\lambda] = F(\lambda, y_i) + \epsilon_i[\lambda]$ aux positions $(y_i)_{1 \leq i \leq n}$. En notant $\hat{T}$ l'estimateur obtenu, on construit un estimateur de T comme suit :

$$\hat{T}(x, y) = \sum_{\lambda \in \Lambda} \hat{F}(\lambda, y) \phi_\lambda(x). \quad (1.52)$$

Le problème d'interpolation des coefficients d'un opérateur a été formulé avec la TVIR car - comme nous l'avons vu - cette représentation permet de coder les régularités des variations des réponses impulsionnelles de manière très simple.

Dans la suite, nous allons supposer que la TVIR T appartient à l'espace $\mathcal{T}^{r,s}$ défini par toutes les fonctions $T : \Omega \times \Omega \rightarrow \mathbb{R}$ telles que $T(x, \cdot) \in H^s(\Omega)$ uniformément pour tout $x \in \Omega$ et $T(\cdot, y) \in H^r(\Omega)$ uniformément pour tout $y \in \Omega$.

Algorithme numériquement réalisable en grande dimension

Une observation récurrente qui explique l'utilisation des ondelettes est leur capacité de compresser une grande classe de signaux. Notamment, les ondelettes compressent les réponses impulsionnelles en très peu de coefficients. C'est une des raisons qui nous a amené à choisir de décomposer les réponses impulsionnelles dans une bases d'ondelettes $(\psi_\lambda)_{\lambda \in \Lambda}$.

Nous avons proposé un estimateur $\hat{T}$ de la forme

$$\hat{T} = \arg \min_{T: \Omega \times \Omega \rightarrow \mathbb{R}} \frac{1}{n} \sum_{i=1}^n \| \langle T(\cdot, y_i), \psi_\lambda \rangle_{\lambda \in \Lambda} - \mathbf{F}_i \|_2^2 + \mu \left((1 - \alpha) \int_{\Omega} |T(x, \cdot)|_{BL^s(\Omega)}^2 dx + \alpha \int_{\Omega} \|T(\cdot, y)\|_{H^r(\Omega)}^2 dy \right). \quad (1.53)$$

Le terme d'attache aux données permet de trouver une TVIR qui est proche des réponses impulsionnelles connues. Le premier terme de régularisation fait intervenir une semi-norme de l'espace de Beppo-Levi BL^s et permet d'interpoler les données manquantes, tandis que le deuxième terme lisse le bruit introduit dans les données. Le paramètre $\alpha \in [0, 1)$ permet de pondérer les deux régularisations. Ce problème d'estimation peut être vu comme un problème d'apprentissage multi-tâches qui fait intervenir des Reproducing Kernel Hilbert Spaces (RKHS) à valeurs vectorielles.

Les espaces de Sobolev peuvent être caractérisés par une norme à poids sur les coefficients d'ondelettes. Cette propriété permet d'écrire le problème (1.53) comme

$$\hat{F} = \arg \min_{F: \Lambda \times \Omega \rightarrow \mathbb{R}} \frac{1}{n} \sum_{i=1}^n \|F(\cdot, y_i) - \mathbf{F}_i\|_2^2 + \mu \|D^s \mathbf{F}\|_{\mathbb{R}^N \times L^2(\Omega)} + \gamma \|\mathbf{L} F\|_{\mathbb{R}^N \times L^2(\Omega)}, \quad (1.54)$$

où $D^s = \sum_{|\beta|=s} \partial^\beta$ est l'opérateur différentiel codant la semi-norme $H^s(\Omega)$, $\mathbf{L}$ est une matrice diagonale $N \times N$ de coefficients $\mathbf{L}[\lambda, \lambda] = 2^{r|\lambda|}$ et $\|\cdot\|_{\mathbb{R}^N \times L^2(\Omega)}$ décrit la norme

$$\|F\|_{\mathbb{R}^N \times L^2(\Omega)} = \sum_{\lambda \in \Lambda} \|F(\lambda, \cdot)\|_{L^2(\Omega)}^2. \quad (1.55)$$

La solution de ce problème $\hat{F}$ peut être obtenue en construisant $\hat{F}(\lambda, \cdot) = \hat{f}_\lambda$ où $\hat{f}_\lambda : \Omega \rightarrow \mathbb{R}$ est solution du problème

$$\hat{f}_\lambda = \arg \min_{f: \Omega \rightarrow \mathbb{R}} \frac{1}{n} \sum_{i=1}^n (\mathbf{F}_i[\lambda] - f(y_i))^2 + \mu \|D^s f\|_{L^2(\Omega)}^2 + \gamma \|2^{r|\lambda|} f\|_{L^2(\Omega)}^2. \quad (1.56)$$

Chacun de ces sous problèmes peuvent être résolu à l'aide de fonctions radiales, en résolvant un système linéaire de petite taille $n \times n$. Le noyau des fonctions radiales est déterminé de manière explicite.

Les ondelettes sont encore ici cruciales pour construire une méthode d'estimation rapide. Elles sont utiles pour plusieurs raisons :

- Elles compressent les réponses impulsionnelles : l'ensemble Λ est de petite dimension.
- Elles permettent de décomposer le problème d'estimation sous la forme d'une interpolation indépendante ligne par ligne. Chacune de ces interpolations se fait rapidement par la résolution d'un système linéaire de petite taille.

La figure 1.14 illustre la méthode.




Figure 1.14 – Un résultat d'interpolation d'un opérateur défini sur des images 256×256 . Les opérateurs exact et interpolé sont appliqués à un peigne de dirac supporté sur une grille cartésienne. Le noyau de l'opérateur est une gaussienne de matrice de covariance $\Sigma(y_1, y_2) = \text{diag}(\sigma^2(y_1, y_2), \sigma^2(y_1, y_2))$ où $\sigma(y_1, y_2) = 1 + 2 \max(1 - y_1, y_1)$ for $(y_1, y_2) \in [0, 1]^2$. Dans cette expérience 64 PSFs ont été utilisées et dégradées par un bruit additif gaussien d'écart-type 10^{-3} . En bas, une image a été dégradée en utilisant l'opérateur exact et dégradée en un bruit additif gaussien d'écart-type 10^{-3} . L'image restaurée a été obtenue en utilisant l'opérateur interpolé.

Taux de convergence

La vitesse de convergence de l'estimateur a été étudiée théoriquement. Sous l'hypothèse de quasi-uniformité de l'ensemble $\{y_1, \dots, y_n\}$, nous montrons que l'estimateur $\hat{H}$ construit par la méthode décrite ci-dessus satisfait :

$$\mathbb{E}\|H - \hat{H}\|_{2 \rightarrow 2}^2 \lesssim N^{-\frac{2r}{d}} + (Nn^{-1})^{\frac{2s}{2s+d}} \quad (1.57)$$

pour un choix du paramètre $\mu \propto (Nn^{-1})^{\frac{2s}{2s+d}}$.

Résumé

Au final, nous avons proposé une méthode d'estimation qui :

- permet de lisser le bruit introduit par les données tout en maintenant des algorithmes d'interpolation ligne à ligne qui sont résolus par des petits systèmes linéaires,
- permet d'obtenir un estimateur avec une interprétation variationnelle dans le domaine spatial.

De plus, nous avons étudié la vitesse d'estimation en risque quadratique de l'opérateur ainsi estimé. Il nous semble que celle-ci est optimale, bien que ce point reste à élucider.

Les résultats développés dans le Chapitre 5 s'appliquent à un choix plus large de base $(\phi_\lambda)_{\lambda \in \Lambda}$. Ils permettent de comprendre le comportement de certaines méthodes d'interpolation des coefficients dans des bases utilisées couramment en astronomie.

1.5.4 Bruits multiplicatifs structurés – Chapitre 6

J'ai eu la chance de collaborer avec Pierre Weiss et Wenxing Zhang sur le problème de restauration d'images dégradées par des bruits multiplicatifs structurés.

En microscopie à fluorescence, les images sont dégradées par des bruits qui sont à la fois structurés et multiplicatifs. Ceci est principalement dû aux effets d'absorption de la lumière. Les bruits structurés additifs ont été étudiés dans [106], et les bruits indépendants multiplicatifs dans [210, 9, 216, 95]. Nous avons proposé une méthode originale combinant ces deux aspects.

En microscopie à fluorescence, le bruit $\boldsymbol{\eta}$ peut être modélisé comme un processus stochastique stationnaire de type $\boldsymbol{\eta} = \boldsymbol{\psi} \star \boldsymbol{\lambda}$ où $\boldsymbol{\lambda} \in \mathbb{R}^N$ est un vecteur positif indépendant et $\boldsymbol{\psi}$ est un filtre de convolution qui dépend de la structure du bruit. L'image observée est alors reliée à l'image "idéale" par $\mathbf{u}_0 = \mathbf{u} \odot \boldsymbol{\eta}$. Utiliser une division (au lieu d'une multiplication standard) est essentiel afin d'obtenir un problème d'optimisation convexe. Nous avons proposé de résoudre le problème d'optimisation :

$$\min_{\boldsymbol{\lambda} \in \mathbb{R}_+^N} \|\nabla(\mathbf{u}_0 \odot (\boldsymbol{\psi} \star \boldsymbol{\lambda}))\|_1 + \alpha \langle \boldsymbol{\lambda} - \log \boldsymbol{\lambda}, \mathbf{1} \rangle. \quad (1.58)$$

Ce problème peut être justifié d'un point de vue probabiliste en utilisant le principe du maximum a posteriori. Nous le résolvons à l'aide d'un algorithme primal-dual et nous

l'avons implémenté sur GPU. Les résultats numériques sur des images réelles démontrent l'utilité de la méthode qui est maintenant utilisée en routine par les biologistes, voir Figure 1.15.




Figure 1.15 – Expérience de débruitage sur données réelles 1040×1390 . A gauche : l'image dégradée est un plan XY d'un derme de peau imagé à l'aide du SPIM. On observe les raies d'atténuation sur l'image. A droite : l'image restaurée par notre méthode, la plupart des raies ont été supprimées.

1.5.5 SNRs invariants aux contrastes locaux – Chapitre 7

Dans le but de comparer efficacement des images issues de différents algorithmes de débruitage, nous avons proposé une mesure de qualité qui est invariante aux changements de contrastes locaux. Ce travail est issu d'une collaboration avec Pierre Weiss et Yiqiu Dong (DTU Denmark).

La performance des algorithmes de traitement d'images repose souvent sur l'invariance aux changements de contrastes locaux, qui constituent un modèle simple de changement d'illumination. La plupart des travaux traitant de ce problème consistent à décomposer l'image en petites zones, et à normaliser les moyennes et les variances sur chaque zone. Nous avons développé une nouvelle approche très différente.

C'est une méthode d'optimisation convexe permettant de trouver une image possédant un arbre de composantes fixé, la plus proche d'une image de référence. L'arbre des composantes permet de coder la relation d'inclusion entre les composantes connexes des ensembles de niveau d'une image et de définir mathématiquement la notion de changement de contraste local. A partir d'une image de référence u_0 et d'une image cible u_1 , notre algorithme permet de trouver l'image u la plus proche de u_0 , qui diffère uniquement d'un changement de contraste local de u_1 . Ce problème s'exprime sous la forme :

$$\min_{u \in \mathcal{U}_1} \frac{1}{2} \|u - u_0\|_2^2 \quad (1.59)$$

où $\mathcal{U}_1$ est l'ensemble des images qui partagent le même arbre de composantes que u_1 . Il peut être décrit par un arbre ou de façon alternative par un graphe dirigé acyclique. Notre algorithme repose sur cette deuxième représentation qui permet d'identifier notre

problème à celui d'une régression isotonique. Ce problème a reçu une attention considérable dans la littérature, mais les approches actuelles ne permettent pas de traiter les graphes à plusieurs millions de sommets apparaissant en traitement des images. Nous avons mis en œuvre un algorithme d'optimisation du premier ordre que nous avons analysé théoriquement. Notre algorithme a de nombreuses applications potentielles telles que la mesure de qualité image, la détection de changement, le traitement d'images couleurs ou le recalage d'images. Voir Figure 1.16 pour une illustration.




Figure 1.16 – Un exemple de correction de changement d'illumination. On remarque que la méthode est capable d'identifier la similarité de la scène.

1.6 Perspectives

Des questions ont émergé au cours de cette thèse, mais n'ont pas été traitées - en partie par manque de temps. Je détaille quelques problèmes que j'aimerais continuer à explorer ci-dessous.

1.6.1 Approximation des opérateurs

Décomposition rapides des opérateurs dans des bases d'ondelettes

Comme nous l'avons vu, les ondelettes sont des outils pertinents pour représenter des opérateurs intégraux dans le but de restaurer des images dégradées par des flous variables. Une des difficultés à surmonter pour les rendre plus populaires est la réduction du coût de calcul de la représentation Θ . Une approche naïve permet de calculer cette représentation en $O(N^3)$ opérations, ce qui est prohibitif!

Nous avons développé quelques algorithmes pour calculer cette matrice efficacement, cependant, leurs performances pratiques sont encore trop mauvaises et nous avons choisi de ne pas les présenter dans cette thèse. Notons tout de même que le problème est résolu de façon efficace dans le cadre des convolutions, où la structure de la matrice Θ nous a permis d'élaborer des algorithmes en $O(N)$. Nous pensons continuer à étudier ce problème par la suite.

Structure de parcimonie adaptée au problème inverse

Les méthodes développées pour déterminer a priori une structure de parcimonie (1.41) ou (1.42) sont optimales pour le problème direct i.e. le calcul de $\mathbf{H}\mathbf{u}$. Notre but est en fait de résoudre un problème inverse du type

$$\hat{\mathbf{u}}_\lambda \in \arg \min_{\mathbf{u} \in \mathbb{R}^N} \frac{1}{2} \|\mathbf{H}\mathbf{u} - \mathbf{u}_0\|_2^2 + \lambda R(\mathbf{u}). \quad (1.60)$$

Les résultats numériques obtenus au Chapitre 3 laissent penser que déterminer une structure optimale pour ce problème inverse permettrait d'obtenir des résultats de restauration meilleurs. Une telle structure de parcimonie Ξ de cardinalité L pourrait être obtenue en résolvant le problème :

$$\min_{|\Xi|=L} \sup_{\|\mathbf{u}\|_X \leq 1} \mathbb{E} \|\hat{\mathbf{u}}_\lambda - \mathbf{u}\|_2^2, \quad (1.61)$$

où l'espérance est prise sur la perturbation stochastique et $\hat{\mathbf{u}}_\lambda$ est la solution du problème d'optimisation obtenue avec l'opérateur $\tilde{\mathbf{H}} = \Psi \Theta_L \Psi^*$, la matrice Θ_L est supportée sur Ξ . La structure de parcimonie Ξ serait plus adaptée au problème inverse et au paramètre de régularisation λ . Nous avons à peine effleuré ce problème pendant la thèse. Il semble particulièrement complexe à résoudre, mais aussi particulièrement pertinent.

Décomposition des opérateurs dans des trames

La restauration du problème (1.43) utilise une régularisation promouvant la parcimonie des images dans une base d'ondelettes orthogonales. Numériquement, ces bases sont connues pour ne pas être invariantes aux translations. Des artefacts peuvent apparaître sur l'image restaurée. L'utilisation de trame d'ondelettes invariantes par translations permet de réduire significativement ces artefacts. Cependant, les temps de transformation sont beaucoup plus longs et la méthodologie développée dans le Chapitre 4 pourrait être d'un grand intérêt. Des questions doivent néanmoins être étudiées préalablement :

- La représentation d'un opérateur dans une trame n'est pas unique, comment choisir une représentation adaptée ? Sur quels critères ? On peut imaginer tenter de trouver la représentation la plus parcimonieuse de sorte à minimiser le coût d'un produit matrice-vecteur. Des résultats numériques préliminaires obtenus dans le cadre d'opérateurs de convolution sont particulièrement encourageants : on peut obtenir des représentations d'opérateurs beaucoup plus parcimonieuses dans des trames d'ondelettes que dans des bases orthogonales.
- La méthode développée dans le Chapitre 4 repose sur le fait que $\Psi\Psi^* = \Psi^*\Psi = I$, ce qui n'est plus le cas pour une trame. Est-il tout de même possible d'élaborer un algorithme rapide de restauration ?

L'idée de représenter des opérateurs dans des trames n'est pas nouvelle [109, 108, 93]. Ces travaux étudient l'approximation d'opérateur par des multiplicateurs de Gabor, qui sont des opérateurs diagonaux dans la trame du même nom.

Optimalité de la décomposition en ondelettes

Comme nous l'avons vu dans ces travaux, le choix de la base qui permet de décomposer les opérateurs est crucial. Elle doit permettre de compresser une classe d'opérateurs tout en fournissant des produits matrice-vecteur rapides. Nous avons fait le choix des bases d'ondelettes pour beaucoup de raisons expliquées précédemment. On peut se demander si ce choix est optimal et se rapprocher de la théorie des Kolmogorov n -widths pour le montrer de façon formelle [193].

Apprentissage de représentation

Enfin, notons que la tendance actuelle est plus l'apprentissage de représentations adaptées à des classes restreintes de problèmes [161, 50, 160]. Ces idées ont un bon potentiel dans le domaine de l'approximation d'opérateurs et sont probablement une bonne piste de recherche. Nous avons choisi des bases "universelles" dans cette thèse, qui permettent de représenter simultanément beaucoup d'objets différents (matrices et vecteurs). Il est cependant probable que de bien meilleures représentations puissent être apprises en prenant en compte simultanément la structure de l'opérateur et celle de la classe d'images à traiter.

1.6.2 Méthodes de produit-convolution

Nous avons étudié les performances théoriques des approximations par produit-convolutions sous l'hypothèse de régularité $T(x, \cdot) \in H^s(\Omega)$ pour tout $x \in \Omega$. Il est intéressant pour certaines applications d'étudier les vitesses de convergence de ces méthodes dans le cas d'espace plus généraux. Par exemple l'hypothèse $T \in BV(\Omega \times \Omega)$ pourrait être traitée en utilisant les résultats de la théorie de l'approximation non-linéaire. Ce cadre est pertinent par exemple, dans le cadre de la vision par ordinateur, où des objets peuvent

se déplacer à différentes vitesses. Dans ce cas, une modélisation approximative du flou est constant par morceaux ³.

Les qualités d'approximation des ces méthodes ont été mesurées en utilisant la norme Hilbert-Schmidt d'opérateur $\|\cdot\|_{HS}$. Comme nous l'avons vu, il est souvent plus intéressant de considérer des normes de type

$$\|H\|_{X \rightarrow Y} = \sup_{\|u\|_X \leq 1} \|Hu\|_Y, \quad (1.62)$$

où $\|\cdot\|_X$ et $\|\cdot\|_Y$ sont des normes caractérisant des espaces fonctionnels. Pour le moment, la démarche à adopter pour contrôler ce genre de norme n'est pas claire. C'est d'ailleurs un des grands avantages des représentations en ondelettes proposées dans le Chapitre 3 : l'adaptation à différentes normes est bien plus facile.

1.6.3 Estimation des opérateurs

Dans le chapitre 5, nous avons considéré le cadre de l'estimation d'un opérateur à partir de quelques réponses impulsionnelles. Dans quelques applications, comme la prospection pétrolière, l'action de l'opérateur est connu sur un ensemble de fonctions test $(u_i)_{1 \leq i \leq n}$ quelconques. Il me semble intéressant d'étudier une méthode d'estimation plus générale que celle proposée dans ces travaux. Plusieurs questions sont à étudier :

- Comment élaborer une méthode d'estimation réalisable en très grande dimension ?
- Il y a-t-il une interprétation spatiale de l'estimateur ?
- Quel est le risque de l'estimateur $\mathbb{E}\|H - \hat{H}\|_{X \rightarrow Y}$?
- Comment choisir les n meilleures fonctions qui permettent de minimiser l'erreur d'estimation ?

Ces questions sont assez liées à la théorie récente de l'échantillonnage compressé. De mon point de vue, la dimensionnalité du problème rend cependant l'identification du support difficile. Il me semble plus pertinent de fixer le support en faisant des hypothèses de régularité du noyau, puis d'utiliser des techniques d'estimation sur un ensemble de petite dimension. Les techniques de super-compression évoquées dans le Chapitre 2 semblent être un outil particulièrement adapté à ce problème.

1.7 Organisation de la thèse

Cette thèse est divisée en trois parties. La première est consacrée à l'étude des approximations des opérateurs de flou et à leur utilisation dans la résolution des problèmes inverses. La deuxième concerne le problème d'estimation des opérateurs. Enfin, la troisième regroupe les différents problèmes d'imagerie traités au cours de cette thèse.

³En fait cette approximation est imparfaite, car elle ne tient pas compte des occlusions

Partie I : Approximation of operators

Chapitre 2 : *Approximation of integral operators using convolution-product expansions*

Ce chapitre est issu de [101], soumis en mai 2016 en collaboration avec Pierre Weiss. On y considère une classe d'opérateurs linéaires intégraux de réponses impulsionnelles variant régulièrement dans le temps ou l'espace. Ces opérateurs apparaissent dans un grand nombre d'applications allant du traitement des signaux/images à la biologie. Évaluer leurs actions sur des fonctions est un problème intensif numériquement. Il est cependant nécessaire pour beaucoup d'applications. On analyse une technique d'approximation appelée développement en produit-convolution dans laquelle l'opérateur est approché localement par une convolution. Cette hypothèse permet de construire des algorithmes rapides reposant sur la transformée de Fourier. Nous avons construit différents types de développements, étudié leur vitesse d'approximation et leur complexité en fonction de la régularité des variations des réponses impulsionnelles. Cette analyse suggère de nouvelles implémentations reposant sur des ondelettes. Elles présentent plusieurs atouts : une vitesse d'approximation optimale, de faibles coûts de calculs et de stockage et une adaptivité à la régularité du noyau. Les méthodes proposées dans ce chapitre sont des alternatives aux méthodes plus standards comme *panel clustering*, *cross approximations*, *développement en ondelettes*, $\mathcal{H}$ -matrices.

Chapitre 3 : *Sparse wavelet representation of spatially varying blurring operators*

Ce chapitre est issu de [100] publié dans *SIAM Journal on Imaging Science* 2015 en collaboration avec Pierre Weiss. La restauration d'images dégradées par des flous variables est un problème rencontré dans beaucoup de disciplines comme l'astrophysique, la vision par ordinateur et l'imagerie médicale. L'un des plus gros challenges dans la résolution de ce problème est d'élaborer des algorithmes numériquement efficaces qui permettent d'approcher des opérateurs intégraux. Nous introduisons une nouvelle méthode reposant sur des approximations parcimonieuses des opérateurs de flou dans des bases d'ondelettes. Cette méthode requiert $O(N\epsilon^{-d/M})$ opérations pour produire une ϵ -approximation où N est le nombre de pixels d'une image d -dimensionnelle et $M \geq 1$ est une scalaire décrivant la régularité du flou. De plus, nous proposons une idée originale qui permet de définir a priori une structure de parcimonie quand seulement la régularité de l'opérateur est connue. Les expériences numériques nous montrent que notre algorithme fournit des résultats significativement meilleurs que les méthodes de l'état-de-l'art.

Chapitre 4 : *Accelerating $\ell^1 - \ell^2$ deblurring using wavelet expansions of operators*

Ce chapitre est issu de [99] soumis en décembre 2015 en collaboration avec Pierre Weiss. La restauration d'images dégradées par des flous est un problème fondamental en traitement d'image. Il est généralement résolu en utilisant des procé-

dures d’optimisation gourmandes en calculs. Le but de ce chapitre est de proposer de nouvelles stratégies efficaces permettant de réduire les temps de calculs pour des modèles de restauration simple utilisant des régularisations basées sur des transformées en ondelettes orthogonales. Nous montrons que la minimisation peut être accélérée significativement en utilisant le fait que les images et les opérateurs sont compressibles dans la même base orthonormale d’ondelettes. La méthodologie proposée repose sur trois ingrédients : i) une approximation parcimonieuse de l’opérateur dans une base d’ondelettes, ii) un préconditionneur diagonal et iii) une implémentation massivement parallèle. La combinaison de ces trois ingrédients donne lieu à des facteurs d’accélération variant de 30 à 250 sur une station de travail standard. Par exemple, la restauration d’une image 1024×1024 est effectuée en 0.15 secondes.

Partie II : Estimation of operators

Chapitre 5 : *Identification of linear operators from scattered impulse responses*

Ce chapitre est issu de [98], un article en préparation et en collaboration avec Jérémie Bigot et Pierre Weiss. Préalablement à la restauration d’images dégradées par des flous variables, les opérateurs de flou doivent être estimés. Nous considérons donc le problème de l’estimation d’un opérateur intégral régulier à partir de la connaissance de n réponses impulsionnelles éparpillées dans le domaine. Nous proposons un estimateur $\hat{H}$ de l’opérateur H calculable en très grande dimension. La méthode repose sur une interpolation ligne par ligne des coefficients en ondelettes de la Time Varying Impulse Response (TVIR). L’estimateur a une interprétation spatiale sous forme variationnelle. Nous avons obtenu la vitesse d’estimation $\mathbb{E}\|H - \hat{H}\|_{2 \rightarrow 2}^2 \lesssim N^{-2r/d} + (Nn^{-1})^{\frac{2s}{2s+d}}$ où N est la dimension de la discrétisation et r et s sont des scalaires décrivant la régularité des réponses impulsionnelles et de leurs variations. Cette méthode est illustrée numériquement sur des problèmes d’interpolation d’opérateurs.

Partie III : Imaging problems

Chapitre 6 : *A variational model for multiplicative structured noise removal*

Ce chapitre est issu de [104], accepté en Juin 2016 pour publication dans *Journal of Mathematical Imaging and Vision (JMIV)*. Il est le fruit d’une collaboration avec Pierre Weiss et Wenxing Zhang. Nous considérons le problème d’estimation d’images dégradées par des bruits qui sont à la fois structurés et multiplicatifs. Notre première motivation a été les images issues du Selective Plane Illumination Microscope (SPIM) qui souffrent souvent d’inhomogénéités d’intensité dues à l’absorption et la réfraction de la lumière lors de sa propagation dans l’échantillon. Ce type de dégradations apparaît aussi dans d’autres modalités d’imagerie comme

l'échographie. Nous avons modélisé les bruits multiplicatifs comme des processus stationnaires avec une distribution connue. Cette modélisation donne lieu à un nouveau problème de restauration convexe reposant sur le principe du maximum a posteriori. Après l'étude de quelques propriétés analytiques des minimiseurs, nous proposons finalement des méthodes d'optimisation rapides implémentées sur GPU. Les expériences numériques sur des images 2D issues de la microscopie à fluorescence démontrent l'utilité pratique de ce modèle.

Chapitre 7 : *Contrast Invariant SNR*

Ce chapitre est issu de [244], en cours de préparation et en collaboration avec Pierre Weiss et Yiqiu Dong. Dans ce chapitre nous concevons une mesure de similarité d'images qui est invariante aux changements de contrastes locaux, qui constituent un modèle simple de changement d'illumination. A partir de deux images, l'algorithme construit une image qui est la plus proche de la première mais qui partage le même arbre de composantes que la deuxième. Ce problème peut-être vu comme un problème d'optimisation spécifique appelé régression isotonique. Nous donnons quelques propriétés analytiques de la solution de ce problème. Nous avons aussi élaboré une procédure d'optimisation de premier ordre adaptée ainsi qu'une analyse complète de sa complexité. La méthode proposée se montre en pratique plus efficace et fiable que les méthodes existantes basées sur la méthode des points intérieurs. L'algorithme a aussi des applications potentielles en détection de changements, traitement des images couleurs et fusion d'images.

Chapitre 8 : *High-resolution in-depth imaging of optically cleared thick samples using an adaptive SPIM*

Cette article est issu de [165] publié dans Nature Scientific Reports en Novembre 2015. Il est le fruit d'une collaboration les équipes d'imagerie et de biologie de l'ITAV : Aurore Masson et Corinne Lorenzo ; Céline Frongia et Bernard Ducommun ainsi que la société d'Imagine Optic représentée ici par Grégory Clouvel. De nos jours, les microscopes à feuille de lumière permettent d'imager des échantillons sur plusieurs centaines de microns de profondeur. Les images produites par ce microscope sont dégradées par de la diffusion, des absorptions et des aberrations optiques. Les variations spatiales des indices de réfraction dans l'échantillon perturbent grandement le parcours de la lumière. Cela conduit à une perte de signal et de contraste dans les régions profondes. Ces effets sont particulièrement marqués lors de l'étude d'échantillons complexes inhomogènes. Récemment, des traitements chimiques ont été développés et permettent d'homogénéiser les indices de réfraction. Bien que ce traitement réduise les dégradations liées aux variations d'indices de réfraction dans l'échantillon, il crée une discontinuité entre les indices de l'échantillon traité et le milieu utilisé pour les objectifs. Une nouvelle dégradation apparaît que nous proposons de corriger avec un système d'optique adaptative qui peut être couplée à un algorithme de déconvolution aveugle.

Part I

Approximation of operators

2

Approximation of integral operators using convolution-product expansions

Résumé : *Ce chapitre est issu de [101], soumis en mai 2016 en collaboration avec Pierre Weiss. On y considère une classe d'opérateurs linéaires intégraux de réponses impulsionnelles variant régulièrement dans le temps ou l'espace. Ces opérateurs apparaissent dans un grand nombre d'applications allant du traitement des signaux/images à la biologie. Évaluer leurs actions sur des fonctions est un problème intensif numériquement. Il est cependant nécessaire pour beaucoup d'applications. On analyse une technique d'approximation appelée développement en produit-convolution dans laquelle l'opérateur est approché localement par une convolution. Cette hypothèse permet de construire des algorithmes rapides reposant sur la transformée de Fourier. Nous avons construit différents types de développements, étudié leur vitesse d'approximation et leur complexité en fonction de la régularité des variations des réponses impulsionnelles. Cette analyse suggère de nouvelles implémentations reposant sur des ondelettes. Elles présentent plusieurs atouts : une vitesse d'approximation optimale, de faibles coûts de calculs et de stockage et une adaptivité à la régularité du noyau. Les méthodes proposées dans ce chapitre sont des alternatives aux méthodes plus standards comme panel clustering, cross approximations, développement en ondelettes, $\mathcal{H}$ -matrices.*

Abstract: *We consider a class of linear integral operators with impulse responses varying regularly in time or space. These operators appear in a large number of applications ranging from signal/image processing to biology. Evaluating their action on functions is a computationally intensive problem necessary for many practical problems. We analyze a technique called product-convolution expansion: the operator is locally approximated by a convolution, allowing to design fast numerical algorithms based on the fast Fourier transform. We design various types of expansions, provide their explicit rates of approximation and their complexity depending on the time varying impulse response smoothness. This analysis suggests novel wavelet based implementations of the method with numerous assets such as optimal approximation rates, low complexity and storage requirements as well as adaptivity to the kernels regularity. The proposed methods are an alternative to more standard procedures such as panel clustering, cross approximations, wavelet expansions or hierarchical matrices.*

Contents

2.1	Introduction	53
2.2	Notation	55
2.3	Preliminary facts	57
2.3.1	Assumptions on the operator and examples	57
2.3.2	Examples	58
2.3.3	Product-convolution expansions as low-rank approximations	60
2.3.4	Discretization	60
2.3.5	Implementation and complexity	62
2.4	Projections on linear subspaces	63
2.4.1	Fourier decompositions	63
2.4.2	Spline decompositions	64
2.4.3	Wavelet decompositions	66
2.4.4	Interpolation VS approximation	68
2.4.5	On Meyer's operator representation	68
2.5	Adaptive decompositions	70
2.5.1	Singular value decompositions	71
2.5.2	The optimization approach in [85]	73
2.6	Summary and extensions	74
2.6.1	A summary of all results	74
2.6.2	Extensions to higher dimensions	75
2.6.3	Extensions to less regular spaces	75
2.6.4	Controls in other norms	75
2.7	Conclusion	76

2.1 Introduction

We are interested in the compact representation and fast evaluation of a class of space or time varying linear integral operators with regular variations. Such operators appear in a large number of applications ranging from wireless communications [198, 135] to seismic data analysis [121], biology [114] and image processing [203].

In all these applications, a key numerical problem is to efficiently evaluate the action of the operator and its adjoint on given functions. This is necessary - for instance - to design fast inverse problems solvers. The main objective of this paper is to analyze the complexity of a set of approximation techniques coined *product-convolution expansions*.

We are interested in bounded linear integral operator $H : L^2(\Omega) \rightarrow L^2(\Omega)$ defined from a kernel K by:

$$Hu(x) = \int_{\Omega} K(x, y)u(y) dy. \quad (2.1)$$

for all $u \in L^2(\Omega)$, where $\Omega \subset \mathbb{R}^d$. Evaluating integrals of type (2.1) is a major challenge in numerical analysis and many methods have been developed in the literature. Nearly all methods share the same basic principle: decompose the operator kernel as a sum of low rank matrices with a multi-scale structure. This is the case in panel clustering methods [125], hierarchical matrices [34], cross approximations [188] or wavelet expansions [29]. The method proposed in this paper basically shares the same idea, except that the time varying impulse response T of the operator is decomposed instead of the kernel K . The *time varying impulse response* (TVIR) T of H is defined by:

$$T(x, y) = K(x + y, y). \quad (2.2)$$

The TVIR representation of H allows formalizing the notion of regularly varying integral operator: the functions $T(x, \cdot)$ should be “smooth” for all $x \in \Omega$. Intuitively, the smoothness assumption means that two neighboring impulse responses should only differ slightly. Under this assumption, it is tempting to approximate H locally by a convolution. Two different approaches have been proposed in the literature to achieve this. The first one is called *convolution-product expansion of order m* and consists of approximating H by an operator H_m of type:

$$H_m u = \sum_{k=1}^m w_k \odot (h_k \star u), \quad (2.3)$$

where $\odot$ denotes the standard multiplication for functions and the Hadamard product for vectors, and $\star$ denotes the convolution operator. The second one, called *product-convolution expansion of order m* , is at the core of this paper and consists of using an expansion of type:

$$H_m u = \sum_{k=1}^m h_k \star (w_k \odot u). \quad (2.4)$$

These two types of approximations have been used for a long time in the field of imaging (and to a lesser extent mobile communications and biology) and progressively became

more and more refined [228, 180, 112, 114, 135, 7, 132, 175, 85]. In particular, the recent work [85] provides a nice overview of existing choices for the functions h_k and w_k as well as new ideas leading to significant improvements. Many different names have been used in the literature to describe expansions of type (2.3) and (2.4) depending on the communities: sectional methods, overlap-add and overlap-save methods, piecewise convolutions, anisoplanatic convolutions, parallel product-convolution, filter flow, windowed-convolutions... The term product-convolution comes from the field of mathematics [39]¹. We believe that it precisely describes the set of expansions of type (2.3) and therefore chose this naming. It was already used in the field of imaging by [7]. Now that product-convolution expansions have been described, natural questions arise:

- i) How to choose the functions h_k and w_k ?
- ii) What is the numerical complexity of evaluating products of type $H_m u$?
- iii) What is the resulting approximation error $\|H_m - H\|$, where $\|\cdot\|$ is a norm over the space of operators?
- iv) How many operations are needed in order to obtain an approximation H_m such that $\|H_m - H\| \leq \epsilon$?

Elements i) and ii) have been studied thoroughly and improved over the years in the mentioned papers. The main questions addressed herein are points iii) and iv). To the best of our knowledge, they have been ignored until now. They are however necessary in order to evaluate the theoretical performance of different product-convolution expansions and to compare their respective advantages precisely.

The main outcome of this paper is the following: under smoothness assumptions of type $T(x, \cdot) \in H^s(\Omega)$ for all $x \in \Omega$ (the Hilbert space of functions in $L^2(\Omega)$ with s derivatives in $L^2(\Omega)$), most methods proposed in the literature - if implemented correctly - ensure a decay of type $\|H_m - H\|_{HS} = O(m^{-s})$, where $\|\cdot\|_{HS}$ is the Hilbert-Schmidt norm. Moreover, this bound cannot be improved uniformly on the considered smoothness class. By adding a support condition of type $\text{supp}(T(x, \cdot)) \subseteq [-\kappa/2, \kappa/2]$, the bound becomes $\|H_m - H\|_{HS} = O(\sqrt{\kappa} m^{-s})$. More importantly, bounded supports allow reducing the computational burden. After discretization on n time points, we show that the number of operations required to satisfy $\|H_m - H\|_{HS} \leq \epsilon$ vary from $O\left(\kappa^{\frac{1}{2s}} n \log_2(n) \epsilon^{-1/s}\right)$ to $O\left(\kappa^{\frac{2s+1}{2s}} n \log_2(\kappa n) \epsilon^{-1/s}\right)$ depending on the choices of w_k and h_k . We also show that the compressed operator representations of Meyer [170] can be used under additional regularity assumptions.

An important difference of product-convolution expansions compared to most methods in the literature [125, 34, 188, 29, 100] is that they are insensitive to the smoothness

¹With the terminology of [39], the name product-convolution would have been convolution-product and vice-versa. The name product-convolution seems more appropriate to describe a product followed by a convolution.

of $T(\cdot, y)$. The smoothness in the x direction is a useful property to control the discretization error, but not the approximation rate. The proposed methodology might therefore be particularly competitive in applications with irregular impulse responses.

The paper is organized as follows. In section 2.2, we describe the notation and introduce a few standard results of approximation theory. In section 2.3, we precisely describe the class of operators studied in this paper, show how to discretize them and provide the numerical complexity of evaluating product-convolution expansions of type (2.4). Sections 2.4 and 2.5 contain the full approximation analysis for two different kinds of approaches called linear or adaptive methods. Section 2.6 contains a summary and a few additional comments.

2.2 Notation

Let a and b denote functions depending on some parameters. The relationship $a \asymp b$ means that a and b are equivalent, i.e. that there exists $0 < c_1 \leq c_2$ such that $c_1 a \leq b \leq c_2 a$. Constants appearing in inequalities will be denoted by C and may vary at each occurrence. If a dependence on a parameter exists (e.g. ϵ), we will use the notation $C(\epsilon)$.

In most of the paper, we work on the unit circle $\Omega = \mathbb{R} \setminus \mathbb{Z}$ sometimes identified with the interval $[-\frac{1}{2}, \frac{1}{2}]$. This choice is driven by simplicity of exposition and the results can be extended to bounded domains such as $\Omega = [0, 1]^d$ (see section 2.6.2). Let $L^2(\Omega)$ denote the space of square integrable functions on Ω . The Sobolev space $H^s(\Omega)$ is defined as the set of functions in $L^2(\Omega)$ with weak derivatives up to order s in $L^2(\Omega)$. The k -th weak derivative of $u \in H^s(\Omega)$ is denoted $u^{(s)}$. The norm and semi-norm of $u \in H^s(\Omega)$ are defined by:

$$\|u\|_{H^s(\Omega)} = \sum_{k=0}^s \|u^{(k)}\|_{L^2(\Omega)} \quad \text{and} \quad |u|_{H^s(\Omega)} = \|u^{(s)}\|_{L^2(\Omega)}. \quad (2.5)$$

The sequence of functions $(e_k)_{k \in \mathbb{Z}}$ where $e_k : x \mapsto \exp(2i\pi kx)$ is a Hilbert basis of $L^2(\Omega)$ (see e.g. [145]).

Definition 2.2.1. Let $u \in L^2(\Omega)$ and $e_k : x \mapsto \exp(2i\pi kx)$ denote the k -th Fourier atom. The Fourier series coefficients $\hat{u}[k]$ of u are defined for all $k \in \mathbb{Z}$ by:

$$\hat{u}[k] = \int_{\Omega} u(x) \overline{e_k(x)} dx = \int_{\Omega} u(x) \exp(-2i\pi kx) dx. \quad (2.6)$$

The space $H^s(\Omega)$ can be characterized through Fourier series.

Lemma 2.2.1 (Fourier characterization of Sobolev norms).

$$\|u\|_{H^s(\Omega)}^2 \asymp \sum_{k \in \mathbb{Z}} |\hat{u}[k]|^2 (1 + |k|^2)^s. \quad (2.7)$$

Definition 2.2.2 (B-spline of order α). Let $\alpha \in \mathbb{N}$ and $m \geq \alpha + 2$ be two integers. The B-spline of order 0 is defined by

$$B_{0,m} = \mathbb{1}_{[-1/(2m), 1/(2m)]}. \quad (2.8)$$

The B-spline of order $\alpha \in \mathbb{N}^*$ is defined by recurrence by:

$$B_{\alpha,m} = m B_{0,m} \star B_{\alpha-1,m} = m^\alpha \underbrace{B_{0,m} \star \dots \star B_{0,m}}_{\alpha \text{ times}}. \quad (2.9)$$

The set of cardinal B-splines of order α is denoted $\mathcal{B}_{\alpha,m}$ and defined by:

$$\mathcal{B}_{\alpha,m} = \left\{ f(\cdot) = \sum_{k=0}^{m-1} c_k B_{\alpha,m}(\cdot - k/m), \ c_k \in \mathbb{R}, \ 0 \leq k \leq m-1 \right\}. \quad (2.10)$$

In this work, we use the Daubechies wavelet basis for $L^2(\mathbb{R})$ [78]. We let ϕ and ψ denote the scaling and mother wavelets and assume that the mother wavelet ψ has α vanishing moments, i.e.

$$\forall 0 \leq m < \alpha, \quad \int_{[0,1]} t^m \psi(t) dt = 0. \quad (2.11)$$

Daubechies wavelets satisfy $\text{supp}(\psi) = [-\alpha + 1, \alpha]$, see [163, Theorem 7.9, p. 294]. Translated and dilated versions of the wavelets are defined, for all $j > 0$ by

$$\psi_{j,l}(x) = 2^{j/2} \psi(2^j x - l). \quad (2.12)$$

The set of functions $(\psi_{j,l})_{j \in \mathbb{N}, l \in \mathbb{Z}}$, is an orthonormal basis of $L^2(\mathbb{R})$ with the convention $\psi_{0,l} = \phi(x - l)$. There are different ways to construct a wavelet basis on the interval $[-1/2, 1/2]$ from a wavelet basis on $L^2(\mathbb{R})$. Here, we use boundary wavelets defined in [71]. We refer to [79, 163] for more details on the construction of wavelet bases. This yields an orthonormal basis $(\psi_\lambda)_{\lambda \in \Lambda}$ of $L^2(\Omega)$, where

$$\Lambda = \{(j, l), j \in \mathbb{N}, 0 \leq l \leq 2^j\}. \quad (2.13)$$

We let $I_\lambda = \text{supp}(\psi_\lambda)$ and for $\lambda \in \Lambda$, we use the notation $|\lambda| = j$.

Let u and v be two functions in $L^2(\Omega)$, the notation $u \otimes v$ will be used both to indicate the function $w \in L^2(\Omega \times \Omega)$ defined by

$$w(x, y) = (u \otimes v)(x, y) = u(x)v(y), \quad (2.14)$$

or the Hilbert-Schmidt operator $w : L^2(\Omega) \rightarrow L^2(\Omega)$ defined for all $f \in L^2(\Omega)$ by:

$$w(f) = (u \otimes v)f = \langle u, f \rangle v. \quad (2.15)$$

The meaning can be inferred depending on the context. Let $H : L^2(\Omega) \rightarrow L^2(\Omega)$ denote a linear integral operators. Its kernel will always be denoted K and its time varying impulse response T . The linear integral operator with kernel T will be denoted J .

The following result is an extension of the singular value decomposition to operators.

Lemma 2.2.2 (Schmidt decomposition [193, Theorem 2.2] or [129, Theorem 1 p. 215]). *Let $H : L^2(\Omega) \rightarrow L^2(\Omega)$ denote a compact operator. There exists two finite or countable orthonormal systems $\{e_1, \dots\}$, $\{f_1, \dots\}$ of $L^2(\Omega)$ and a finite or infinite sequence $\sigma_1 \geq \sigma_2 \geq \dots$ of positive numbers (tending to zero if it is infinite), such that H can be decomposed as:*

$$H = \sum_{k \geq 1} \sigma_k \cdot e_k \otimes f_k. \quad (2.16)$$

A function $u \in L^2(\Omega)$ is denoted in regular font whereas its discretized version $\mathbf{u} \in \mathbb{R}^n$ is denoted in bold font. The value of function u at $x \in \Omega$ is denoted $u(x)$, while the i -th coefficient of vector $\mathbf{u} \in \mathbb{R}^n$ is denoted $\mathbf{u}[i]$. Similarly, an operator $H : L^2(\Omega) \rightarrow L^2(\Omega)$ is denoted in upper-case regular font whereas its discretized version $\mathbf{H} \in \mathbb{R}^{n \times n}$ is denoted in upper-case bold font.

2.3 Preliminary facts

In this section, we gather a few basic results necessary to derive approximation results.

2.3.1 Assumptions on the operator and examples

All the results stated in this paper rely on the assumption that the TVIR T of H is a sufficiently simple function. By simple, we mean that i) the functions $T(x, \cdot)$ are smooth for all $x \in \Omega$ and ii) the impulse responses $T(\cdot, y)$ have a bounded support or a fast decay for all $y \in \Omega$.

There are numerous ways to capture the regularity of a function. In this paper, we assume that $T(x, \cdot)$ lives in the Hilbert spaces $H^s(\Omega)$ for all $x \in \Omega$. This hypothesis is deliberately simple to clarify the proofs and the main ideas.

Definition 2.3.1 (Class $\mathcal{T}^s$). We let $\mathcal{T}^s$ denote the class of functions $T : \Omega \times \Omega \rightarrow \mathbb{R}$ satisfying the smoothness condition: $T(x, \cdot) \in H^s(\Omega)$, $\forall x \in \Omega$ and $\|T(x, \cdot)\|_{H^s(\Omega)}$ is uniformly bounded in x , i.e:

$$\sup_{x \in \Omega} \|T(x, \cdot)\|_{H^s(\Omega)} \leq C < +\infty. \quad (2.17)$$

Note that if $T \in \mathcal{T}^s$, then H is a Hilbert-Schmidt operator since:

$$\|H\|_{HS}^2 = \int_{\Omega} \int_{\Omega} K(x, y)^2 dx dy \quad (2.18)$$

$$= \int_{\Omega} \int_{\Omega} T(x, y)^2 dx dy \quad (2.19)$$

$$= \int_{\Omega} \|T(x, \cdot)\|_{L^2(\Omega)}^2 dx < +\infty. \quad (2.20)$$

We will often use the following regularity assumption.

Assumption 2.3.1. *The TVIR T of H belongs to $\mathcal{T}^s$.*

In many applications, the impulse responses have a bounded support, or at least a fast spatial decay allowing to neglect the tails. This property will be exploited to design faster algorithms. This hypothesis can be expressed by the following assumption.

Assumption 2.3.2. $T(x, y) = 0, \forall |x| > \kappa/2$.

2.3.2 Examples

We provide 3 examples of kernels that may appear in applications. Figure 2.1 shows each kernel as a 2D image, the associated TVIR and the spectrum of the operator J (the linear integral operator with kernel T) computed with an SVD.

Example 2.3.1. A typical kernel that motivates our study is defined by:

$$K(x, y) = \frac{1}{\sqrt{2\pi}\sigma(y)} \exp\left(-\frac{(x-y)^2}{2\sigma^2(y)}\right). \quad (2.21)$$

The impulse responses $K(\cdot, y)$ are Gaussian for all $y \in \Omega$. Their variance $\sigma(y) > 0$ varies depending on the position y . The TVIR of K is defined by:

$$T(x, y) = \frac{1}{\sqrt{2\pi}\sigma(y)} \exp\left(-\frac{x^2}{2\sigma^2(y)}\right). \quad (2.22)$$

The impulse responses $T(\cdot, y)$ are not compactly supported, therefore, $\kappa = 1$ in Assumption 2.3.2. However, it is possible to truncate them by setting $\kappa = 3 \sup_{y \in \Omega} \sigma(y)$ for instance. This kernel satisfies Assumption 2.3.1 only if $\sigma : \Omega \rightarrow \mathbb{R}$ is sufficiently smooth. In Figure 2.1, left column, we set $\sigma(y) = 0.08 + 0.02 \cos(2\pi y)$.

Example 2.3.2. The second example is given by:

$$T(x, y) = \frac{2}{\sigma(y)} \max(1 - 2\sigma(y)|x|, 0). \quad (2.23)$$

The impulse responses $T(\cdot, y)$ are cardinal B-splines of degree 1 and width $\sigma(y) > 0$. They are compactly supported with $\kappa = \sup_{y \in \Omega} \sigma(y)$. This kernel satisfies Assumption 2.3.2 only if $\sigma : \Omega \rightarrow \mathbb{R}$ is sufficiently smooth. In Figure 2.1, central column, we set $\sigma(y) = 0.1 + 0.3(1 - |y|)$. This kernel satisfies Assumption 2.3.1 with $s = 1$.

Example 2.3.3. The last example is a discontinuous TVIR. We set:

$$T(x, y) = g_{\sigma_1}(x) \mathbb{1}_{[-1/4, 1/4]}(y) + g_{\sigma_2}(x) (1 - \mathbb{1}_{[-1/4, 1/4]}(y)), \quad (2.24)$$

where $g_\sigma(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{\sigma^2}\right)$. This corresponds to the last column in Figure 2.1, with $\sigma_1 = 0.05$ and $\sigma_2 = 0.1$. For this kernel, both Assumptions 2.3.1 and 2.3.2 are violated. Notice however that T is the sum of two tensor products and can therefore be represented using only four 1D functions. The spectrum of J should have only 2 non zero elements. This is verified in Figure 2.1i up to numerical errors.




Figure 2.1 – Different kernels K , the associated TVIR T and the spectrum of the operator J . Left column corresponds to Example 2.3.1. Central column corresponds to Example 2.3.2. Right column corresponds to Example 2.3.3.

2.3.3 Product-convolution expansions as low-rank approximations

Though similar in spirit, convolution-product (2.3) and product-convolution (2.4) expansions have a quite different interpretation captured by the following lemma.

Lemma 2.3.1. *The TVIR of the convolution-product expansion T_m in (2.3) is given by:*

$$T_m(x, y) = \sum_{k=1}^m h_k(x) w_k(x + y). \quad (2.25)$$

The TVIR of the product-convolution expansion T_m in (2.4) is given by:

$$T_m(x, y) = \sum_{k=1}^m h_k(x) w_k(y). \quad (2.26)$$

Proof. We only prove (2.26) since the proof of (2.25) relies on the same arguments. By definition:

$$(H_m u)(x) = \left(\sum_{k=1}^m h_k \star (w_k \odot u) \right) (x) \quad (2.27)$$

$$= \int_{\Omega} \sum_{k=1}^m h_k(x - y) w_k(y) u(y) dy. \quad (2.28)$$

By identification, this yields:

$$K_m(x, y) = \sum_{k=1}^m h_k(x - y) w_k(y), \quad (2.29)$$

so that

$$T_m(x, y) = \sum_{k=1}^m h_k(x) w_k(y). \quad (2.30)$$

□

As can be seen in (2.26), product-convolution expansions consist of finding low-rank approximations of the TVIR. This interpretation was already proposed in [85] for instance and is the key observation to derive the forthcoming results. The expansion (2.25) does not share this simple interpretation and we do not investigate it further in this paper.

2.3.4 Discretization


In order to implement a product-convolution expansion of type 2.4, the problem first needs to be discretized. We address this problem with a Galerkin formalism. Let $(\varphi_1, \dots, \varphi_n)$ be a basis of a finite dimensional vector space V^n of $L^2(\Omega)$. Given an operator $H : L^2(\Omega) \rightarrow L^2(\Omega)$, we can construct a matrix $\mathbf{H}^n \in \mathbb{R}^{n \times n}$ defined for all

$1 \leq i, j \leq n$ by $\mathbf{H}^n[i, j] = \langle H\varphi_j, \varphi_i \rangle$. Let $S^n : H \mapsto \mathbf{H}^n$ denote the discretization operator. From a matrix $\mathbf{H}^n$, an operator H^n can be reconstructed using, for instance, the pseudo-inverse $S^{n,+}$ of S^n . We let $H^n = S^{n,+}(\mathbf{H}^n)$. For instance, if $(\varphi_1, \dots, \varphi_n)$ is an *orthonormal* basis of V^n , the operator H^n is given by:


$$H^n = S^{n,+}(\mathbf{H}^n) = \sum_{1 \leq i, j \leq n} \mathbf{H}^n[i, j] \varphi_i \otimes \varphi_j. \quad (2.31)$$

This paper is dedicated to analyzing methods denoted $\mathcal{A}_m$ that provide an approximation $H_m = \mathcal{A}_m(H)$ of type (2.4), given an input operator H . Our analysis provides guarantees on the distance $\|H - H_m\|_{HS}$ depending on m and the regularity properties of the input operator H , for different methods. Depending on the context, two different approaches can be used to implement $\mathcal{A}_m$.

- Compute the matrix $\mathbf{H}_m^n = S^n(H_m)$ using numerical integration procedures. Then create an operator $H_m^n = S^{n,+}(\mathbf{H}_m^n)$. This approach suffers from two defects. First, it is only possible by assuming that the kernel of H is given analytically. Moreover it might be computationally intractable. It is illustrated below.



- In many applications, the operator H is not given explicitly. Instead, we only have access to its discretization $\mathbf{H}^n$. Then it is possible to construct a discrete approximation algorithm $\mathcal{A}_m$ yielding a discrete approximation $\mathbf{H}_m^n = \mathcal{A}_m(\mathbf{H}^n)$. This matrix can then be mapped back to the continuous world using the pseudo-inverse: $H_m^n = S^{n,+}(\mathbf{H}_m^n)$. This is illustrated below. In this paper, we will analyze the construction complexity of $\mathbf{H}_m^n$ using this second approach.



Ideally, we would like to provide guarantees on $\|H - H_m^n\|_{HS}$ depending on m and n . In the first approach, this is possible by using the following inequality:

$$\|H - H_m^n\|_{HS} \leq \underbrace{\|H - H_m\|_{HS}}_{\epsilon_a(m)} + \underbrace{\|H_m - H_m^n\|_{HS}}_{\epsilon_d(n)}, \quad (2.32)$$

where $\epsilon_a(m)$ is the approximation error studied in this paper and $\epsilon_d(n)$ is the discretization error. Under mild regularity assumptions on K , it is possible to obtain results of type $\epsilon_d(n) = O(n^{-\gamma})$, where γ depends on the smoothness of K . For instance, if $K \in H^r(\Omega \times \Omega)$, the error satisfies $\epsilon_d(n) = O(n^{-r/2})$ for many bases including Fourier, wavelets and B-splines [65]. For $K \in BV(\Omega \times \Omega)$, the space of functions with bounded variations, $\epsilon_d(n) = O(n^{-1/4})$, see [163, Theorem 9.3]. As will be seen later, the approximation error $\epsilon_a(m)$ behaves like $O(m^{-s})$. The proposed approximation technique

is valuable only in the case $m \ll n$, so that the discretization error can be considered negligible compared to the approximation error most practical applications. In all the paper, we assume that $\epsilon_d(n)$ is negligible compared to $\epsilon_a(m)$ without mention.

In the second approach, the error analysis is more complex since there is an additional bias due to the algorithm discretization. This bias is captured by the following inequality:

$$\|H - H_m^n\|_{HS} \leq \underbrace{\|H - H^n\|_{HS}}_{\epsilon_d(n)} + \underbrace{\|H^n - \mathcal{A}_m(H^n)\|_{HS}}_{\epsilon_a(m)} + \underbrace{\|\mathcal{A}_m(H^n) - H_m^n\|_{HS}}_{\epsilon_b(m,n)}. \quad (2.33)$$

The bias $\epsilon_b(m, n) = \|\mathcal{A}_m(S^{n,+}(S^n(H))) - S^{n,+}(\mathcal{A}_m(S^n(H)))\|_{HS}$ accounts for the difference between using the discrete or continuous approximation algorithm. In this paper, we do not study this bias error and assume that it is negligible compared to the approximation error ϵ_a .

2.3.5 Implementation and complexity

Let $\mathbf{F}_n \in \mathbb{C}^{n \times n}$ denote the discrete inverse Fourier transform and $\mathbf{F}_n^*$ denote the discrete Fourier transform. Matrix-vector products $\mathbf{F}_n \mathbf{u}$ or $\mathbf{F}_n^* \mathbf{u}$ can be evaluated in $O(n \log_2(n))$ operations using the fast Fourier transform (FFT). The discrete convolution product $\mathbf{v} = \mathbf{h} \star \mathbf{u}$ is defined for all $i \in \mathbb{Z}$ by $v[i] = \sum_{j=1}^n u[i-j]h[j]$, with circular boundary conditions.

Discrete convolution products can be evaluated in $O(n \log_2(n))$ operations by using the following fundamental identity:

$$\mathbf{v} = \mathbf{F}_n \cdot ((\mathbf{F}_n^* \mathbf{h}) \odot (\mathbf{F}_n^* \mathbf{u})). \quad (2.34)$$

Hence a convolution can be implemented using three FFTs ($O(n \log_2(n))$ operations) and a point-wise multiplication ($O(n)$ operations). This being said, it is straightforward to implement formula (2.4) with an $O(mn \log_2(n))$ algorithm.

Under the additional assumption that w_k and h_k are supported on bounded intervals, the complexity can be improved. We assume that, after discretization, $\mathbf{h}_k$ and $\mathbf{w}_k$ are compactly supported, with support length $q_k \leq n$ and $p_k \leq n$ respectively.

Lemma 2.3.2. *A matrix-vector product of type (2.4) can be implemented with a complexity that does not exceed $O(\sum_{k=1}^m (p_k + q_k) \log_2(\min(p_k, q_k)))$ operations.*

Proof. A convolution product of type $\mathbf{h}_k \star (\mathbf{w}_k \odot \mathbf{u})$ can be evaluated in $O((p_k + q_k) \log(p_k + q_k))$ operations. Indeed, the support of $\mathbf{h}_k \star (\mathbf{w}_k \odot \mathbf{u})$ has no more than $p_k + q_k$ contiguous non-zeros elements. Using the Stockham sectioning algorithm [218], the complexity can be further decreased to $O((p_k + q_k) \log_2(\min(p_k, q_k)))$ operations. This idea was proposed in [132]. $\square$

2.4 Projections on linear subspaces

We now turn to the problem of choosing the functions h_k and w_k in equation (2.4). The idea studied in this section is to fix a subspace $E_m = \text{span}(e_k, k \in \{1, \dots, m\})$ of $L^2(\Omega)$ and to approximate $T(x, \cdot)$ as:

$$T_m(x, y) = \sum_{k=1}^m c_k(x) e_k(y). \quad (2.35)$$

For instance, the coefficients c_k can be chosen so that $T_m(x, \cdot)$ is a projection of $T(x, \cdot)$ onto E_m . We propose to analyze three different family of functions e_k : Fourier atoms, wavelets atoms and B-splines. We analyze their complexity and approximation properties as well as their respective advantages.

2.4.1 Fourier decompositions

It is well known that functions in $H^s(\Omega)$ can be well approximated by linear combination of low-frequency Fourier atoms. This loose statement is captured by the following lemma.

Lemma 2.4.1 ([89, 88]). *Let $f \in H^s(\Omega)$ and f_m denote its partial Fourier series:*

$$f_m = \sum_{k=-m}^m \hat{f}[k] e_k, \quad (2.36)$$

where $e_k(y) = \exp(2i\pi ky)$. Then

$$\|f_m - f\|_{L^2(\Omega)} \leq C m^{-s} \|f\|_{H^s(\Omega)}. \quad (2.37)$$

The so-called Kohn-Nirenberg symbol N of H is defined for all $(x, k) \in \Omega \times \mathbb{Z}$ by

$$N(x, k) = \int_{\Omega} T(x, y) \exp(-2i\pi ky) dy. \quad (2.38)$$

Illustrations of different Kohn-Nirenberg symbols are provided in Figure 2.2.

Corollary 2.4.1. *Set $e_k(y) = \exp(2i\pi ky)$ and define T_m by:*

$$T_m(x, y) = \sum_{|k| \leq m} N(x, k) e_k(y). \quad (2.39)$$

Then, under Assumptions 2.3.1 and 2.3.2

$$\|H_m - H\|_{HS} \leq C \sqrt{\kappa} m^{-s}. \quad (2.40)$$

Proof. By Lemma 2.4.1 and Assumption 2.3.1, $\|T_m(x, \cdot) - T(x, \cdot)\|_{L^2(\Omega)} \leq Cm^{-s}$ for some constant C and for all $x \in \Omega$. In addition, by Assumption 2.3.2, $\|T_m(x, \cdot) - T(x, \cdot)\|_{L^2(\Omega)} = 0$ for $|x| > \kappa/2$. Therefore:

$$\|H_m - H\|_{HS}^2 = \int_{\Omega} \int_{\Omega} (T_m(x, y) - T(x, y))^2 dx dy \quad (2.41)$$

$$= \int_{\Omega} \|T_m(x, \cdot) - T(x, \cdot)\|_{L^2(\Omega)}^2 dx \quad (2.42)$$

$$\leq \kappa C^2 m^{-2s} \quad (2.43)$$

□

As will be seen later, the convergence rate (2.40) is optimal in the sense that no product-convolution expansion of order m can achieve a better rate under the sole Assumptions 2.3.1 and 2.3.2.

Corollary 2.4.2. *Under Assumptions 2.3.1 and 2.3.2, let $\epsilon > 0$ and set $m = \lceil C\epsilon^{-1/s}\kappa^{1/2s} \rceil$. Then H_m satisfies $\|H - H_m\|_{HS} \leq \epsilon$ and products with H_m and H_m^* can be evaluated with no more than $O(\kappa^{1/2s}n \log n \epsilon^{-1/s})$ operations.*

Proof. Since Fourier atoms are not localized in the time domain, the modulation functions $\mathbf{w}_k$ are supported on intervals of size $p = n$. The complexity of computing a matrix vector product is therefore $O(mn \log(n))$ operations by Lemma 2.3.2. □

Finally, let us mention that computing the discrete Kohn-Nirenberg $\mathbf{N}$ costs $O(\kappa n^2 \log_2(n))$ operations (κn discrete Fourier transforms of size n). The storage cost of this Fourier representation is $O(m\kappa n)$ since one has to store κn coefficients for each of the m vectors $\mathbf{h}_k$.

In the next two sections, we show that replacing Fourier atoms by wavelet atoms or B-splines preserves the optimal rate of convergence in $O(\sqrt{\kappa}m^{-s})$, but has the additional advantage of being localized in space, thereby reducing complexity.

2.4.2 Spline decompositions

B-Splines form a Riesz basis with dual Riesz basis of form [63]:

$$(\tilde{B}_{\alpha,m}(\cdot - k/m))_{0 \leq k \leq m-1}. \quad (2.44)$$

The projection f_m of any $f \in L^2(\Omega)$ onto $\mathcal{B}_{\alpha,m}$ can be expressed as:

$$f_m = \arg \min_{\tilde{f} \in \mathcal{B}_{\alpha,m}} \|\tilde{f} - f\|_2^2 \quad (2.45)$$

$$= \sum_{k=0}^{m-1} \langle f, \tilde{B}_{\alpha,m}(\cdot - k/m) \rangle B_{\alpha,m}(\cdot - k/m). \quad (2.46)$$




Figure 2.2 – Kohn-Nirenberg symbols of the kernels given in Examples 2.3.1, 2.3.2 and 2.3.3 in $\log_{10}$ scale. Observe how the decay speed from the center (low frequencies) to the outer parts (high frequencies) changes depending on the TVIR smoothness. Note: the lowest values of the Kohn-Nirenberg symbol have been set to 10^{-4} for visualization purposes.

Theorem 2.4.1 ([30, p. 87] or [89, p. 420]). *Let $f \in H^s(\Omega)$ and $\alpha \geq s$, then*

$$\|f - f_m\|_2 \leq C\sqrt{\kappa}m^{-s}\|f\|_{H^s(\Omega)}. \quad (2.47)$$

The following result directly follows.

Corollary 2.4.3. *Set $\alpha \geq s$. For each $x \in \Omega$, let $(c_k(x))_{0 \leq k \leq m-1}$ be defined as*

$$c_k(x) = \langle T(x, \cdot), \tilde{B}_{\alpha,m}(\cdot - k/m) \rangle. \quad (2.48)$$

Define T_m by:

$$T_m(x, y) = \sum_{k=0}^{m-1} c_k(x) B_{\alpha,m}(y - k/m). \quad (2.49)$$

If $\alpha \geq s$, then, under Assumptions 2.3.1 and 2.3.2,

$$\|H_m - H\|_{HS} \leq C\sqrt{\kappa}m^{-s}. \quad (2.50)$$

Proof. The proof is similar to that of Corollary (2.4.1). □

Corollary 2.4.4. *Under Assumptions 2.3.1 and 2.3.2, let $\epsilon > 0$ and set $m = \lceil C\epsilon^{-1/s}\kappa^{1/2s} \rceil$. Then H_m satisfies $\|H - H_m\|_{HS} \leq \epsilon$ and products with H_m and H_m^* can be evaluated with no more than*

$$O\left(\left(s + \kappa^{1+1/2s}\epsilon^{-1/s}\right)n \log_2(\kappa n)\right) \quad (2.51)$$

operations. For small ϵ and large n , the complexity behaves like

$$O\left(\kappa^{1+1/2s}n \log_2(\kappa n)\epsilon^{-1/s}\right). \quad (2.52)$$

Proof. In this approximation, m B-splines are used to cover Ω . B-splines have a compact support of size $(\alpha + 1)/m$. This property leads to windowing vector $\mathbf{w}_k$ with support of size $p = \lceil (\alpha + 1) \frac{n}{m} \rceil$. Furthermore the vectors $(\mathbf{h}_k)$ have a support of size $q = \kappa n$. Combining these two results with Lemma 2.3.2 and Corollary 2.4.3 yields the result for the choice $\alpha = s$. $\square$

The complexity of computing the vectors $\mathbf{c}_k$ is $O(\kappa n^2 \log(n))$ (κn projections with complexity $n \log(n)$, see e.g. [231]).

As can be seen in Corollary (2.4.4), B-splines approximations are preferable over Fourier decompositions whenever the support size κ is small.

2.4.3 Wavelet decompositions

Lemma 2.4.2 ([163, Theorem 9.5]). *Let $f \in H^s(\Omega)$ and f_m denote its partial wavelet series:*

$$f_m = \sum_{|\mu| \leq \lceil \log_2(m) \rceil} c_\mu \psi_\mu, \quad (2.53)$$

where ψ is a Daubechies wavelet with $\alpha > s$ vanishing moments and $c_\mu = \langle \psi_\mu, f \rangle$. Then

$$\|f_m - f\|_{L^2(\Omega)} \leq C m^{-s} \|f\|_{H^s(\Omega)}. \quad (2.54)$$

A direct consequence is the following corollary.

Corollary 2.4.5. *Let ψ be a Daubechies wavelet with $\alpha = s + 1$ vanishing moments. Define T_m by:*

$$T_m(x, y) = \sum_{|\mu| \leq \lceil \log_2(m) \rceil} c_\mu(x) \psi_\mu(y), \quad (2.55)$$

where $c_\mu(x) = \langle \psi_\mu, T(x, \cdot) \rangle$. Then, under Assumptions 2.3.1 and 2.3.2

$$\|H_m - H\|_{HS} \leq C \sqrt{\kappa} m^{-s}. \quad (2.56)$$

Proof. The proof is identical to that of Corollary (2.4.1). $\square$

Proposition 2.4.1. *Let $\epsilon > 0$ and set $m = \lceil C \epsilon^{-1/s} \kappa^{1/2s} \rceil$. Under Assumptions 2.3.1 and 2.3.2 H_m satisfies $\|H - H_m\|_{HS} \leq \epsilon$ and products with H_m and H_m^* can be evaluated with no more than*

$$O\left(\left(sn \log_2\left(\epsilon^{-1/s} \kappa^{1/2s}\right) + \kappa^{1+1/2s} n \epsilon^{-1/s}\right) \log_2(\kappa n)\right) \quad (2.57)$$

operations. For small ϵ , the complexity behaves like

$$O\left(\kappa^{1+1/2s} n \log_2(\kappa n) \epsilon^{-1/s}\right). \quad (2.58)$$




Figure 2.3 – “Wavelet symbols” of the operators given in Examples 2.3.1, 2.3.2 and 2.3.3 in $\log_{10}$ scale. The red bars indicate separations between scales. Notice that the wavelet coefficients in kernel 1 rapidly decay as scales increase. The decay is slower for kernels 2 and 3 which are less regular. The adaptivity of wavelets can be visualized in Kernel 3: some wavelet coefficients are non zero at large scales, but they are all concentrated around discontinuities. Therefore only a few number of couples (c_μ, ψ_μ) will be necessary to encode the discontinuities. This was not the case with Fourier or B-spline atoms.

Proof. In (2.55), the windowing vectors $\mathbf{w}_k$ are wavelets ψ_μ of support of size $\min((2s+1)n2^{-|\mu|}, n)$. Therefore each convolution has to be performed on intervals of size $|\psi_\mu| + \kappa n$. Since there are 2^j wavelets at scale j , the total number of operations is:

$$\sum_{\mu \mid |\mu| < \log_2(m)} (|\psi_\mu| + \kappa n) \log_2(\min(|\psi_\mu|, \kappa n)) \quad (2.59)$$

$$\leq \sum_{\mu \mid |\mu| < \log_2(m)} ((2s+1)n2^{-|\mu|} + \kappa n) \log_2(\kappa n) \quad (2.60)$$

$$= \sum_{j=0}^{\log_2(m)-1} 2^j \left((2s+1)n2^{-j} + \kappa n \right) \log_2(\kappa n) \quad (2.61)$$

$$= \sum_{j=0}^{\log_2(m)-1} \left((2s+1)n + 2^j \kappa n \right) \log_2(\kappa n) \quad (2.62)$$

$$\leq ((2s+1)n \log_2(m) + m \kappa n) \log_2(\kappa n) \quad (2.63)$$

$$= \left((2s+1)n \log_2(\epsilon^{-1/s} \kappa^{1/2s}) + \epsilon^{-1/s} \kappa^{1+1/2s} n \right) \log_2(\kappa n). \quad (2.64)$$

□

Computing the vectors $\mathbf{c}_\mu$ costs $O(\kappa s n^2)$ operations (κn discrete wavelet transforms of size n). The storage cost of this wavelet representation is $O(m \kappa n)$ since one has to store κn coefficients for each of the m functions $\mathbf{h}_k$.

As can be seen from this analysis, wavelet and B-spline approximations roughly have the same complexity over the class $\mathcal{T}^s$. The first advantage of wavelets compared to B-splines is that the coefficients $c_\mu(x)$ have a simple analytic expression, while B-splines

coefficients c_k are found by solving a linear system. This is slightly more complicated to implement.

The second significant advantage of wavelets compared to B-splines with fixed knots is that they are known to characterize much more general function spaces than $H^s(\Omega)$. For instance, if all functions $T(x, \cdot)$ have a single discontinuity at a given $y \in \Omega$, only a few coefficients $c_\mu(x)$ will remain of large amplitude. Wavelets will be able to efficiently encode the discontinuity, while B-splines with fixed knots - which are not localized in nature - will fail to well approximate the TVIR. It is therefore possible to use wavelets in an adaptive way. This effect is visible on Figure 2.3c: despite discontinuities, only wavelets localized around the discontinuities yield large coefficients. In the next section, we propose two other adaptive methods, in the sense that they are able to automatically adapt to the TVIR regularity.

2.4.4 Interpolation VS approximation

In all previous results, we constructed the functions w_k and h_k in 2.4 by projecting $T(x, \cdot)$ onto linear subspaces. This is only possible if the whole TVIR T is available. In very large scale applications, this assumption is unrealistic, since the TVIR contains n^2 coefficients, which cannot even be stored. Instead of assuming a full knowledge of T , some authors (e.g. [180]) assume that the impulse responses $T(\cdot, y)$ are available only at a discrete set of points $y_i = i/m$ for $1 \leq i \leq m$.

In that case, it is possible to *interpolate* the impulse responses instead of approximating them. Given a linear subspace $E_m = \text{span}(e_k, k \in \{1, \dots, m\})$, where the atoms e_k are assumed to be linearly independent, the functions $c_k(x)$ in (2.35) are chosen by solving the set of linear systems:

$$\sum_{k=1}^m c_k(x) e_k(y_i) = T_m(x, y_i) \quad \text{for } 1 \leq i \leq m. \quad (2.65)$$

In the discrete setting, under Assumption 2.3.2, this amounts to solving $[\kappa n]$ linear systems of size $m \times m$. We do not discuss the rates of approximation for this interpolation technique since they are usually expressed in the L^∞ -norm under more stringent smoothness assumptions than $T \in \mathcal{T}^s$. We refer the interested reader to [206, 89, 91] for results on spline and wavelet interpolants.

2.4.5 On Meyer's operator representation

Up to now, we only assumed a regularity of T in the y direction, meaning that the impulse responses vary smoothly in space. In many applications, the impulse responses themselves are smooth. In this section, we show that this additional regularity assumption can be used to further compress the operator. Finding a compact operator representation is a key to treat identification problems (e.g. blind deblurring in imaging).

Since $(\psi_\lambda)_{\lambda \in \Lambda}$ is a Hilbert basis of $L^2(\Omega)$, the set of tensor product functions $(\psi_\lambda \otimes \psi_\mu)_{\lambda \in \Lambda, \mu \in \Lambda}$ is a Hilbert basis of $L^2(\Omega \times \Omega)$. Therefore, any $T \in L^2(\Omega \times \Omega)$ can be

expanded as:

$$T(x, y) = \sum_{\lambda \in \Lambda} \sum_{\mu \in \Lambda} c_{\lambda, \mu} \psi_{\lambda}(x) \psi_{\mu}(y). \quad (2.66)$$

The main idea of the construction in this section consists of keeping only the coefficients $c_{\lambda, \mu}$ of large amplitude. A similar idea was proposed in the BCR paper [29]², except that the kernel K was expanded instead of the TVIR T . Decomposing T was suggested by Beylkin at the end of [28] without a precise analysis.

In this section, we assume that $T \in H^{r, s}(\Omega \times \Omega)$, where

$$H^{r, s}(\Omega \times \Omega) = \{T : \Omega \times \Omega \rightarrow \mathbb{R}, \partial_x^{\alpha_1} \partial_y^{\alpha_2} T \in L^2(\Omega \times \Omega), \forall \alpha_1 \in \{0, \dots, r\}, \forall \alpha_2 \in \{0, \dots, s\}\}. \quad (2.67)$$

This space arises naturally in applications, where the impulse response regularity r might differ from the regularity s of their variations. Notice that $H^{2s}(\Omega \times \Omega) \subset H^{s, s}(\Omega \times \Omega) \subset H^s(\Omega)$.

Theorem 2.4.2. *Assume that $T \in H^{r, s}(\Omega \times \Omega)$ and satisfies Assumption 2.3.2. Assume that ψ has $\max(r, s) + 1$ vanishing moments. Let $c_{\lambda, \mu} = \langle T, \psi_{\lambda} \otimes \psi_{\mu} \rangle$. Define*

$$H_{m_1, m_2} = \sum_{|\lambda| \leq \log_2(m_1)} \sum_{|\mu| \leq \log_2(m_2)} c_{\lambda, \mu} \psi_{\lambda} \otimes \psi_{\mu}. \quad (2.68)$$

Let $m \in \mathbb{N}$, set $m_1 = \lceil m^{s/(r+s)} \rceil$, $m_2 = \lceil m^{r/(r+s)} \rceil$ and $H_m = H_{m_1, m_2}$. Then

$$\|H - H_m\|_{HS} \leq C \sqrt{\kappa} m^{-\frac{rs}{r+s}}. \quad (2.69)$$

Proof. First notice that

$$T_{\infty, m_2} = \sum_{|\mu| \leq \lceil \log_2(m_2) \rceil} c_{\mu} \otimes \psi_{\mu}, \quad (2.70)$$

where $c_{\mu}(x) = \langle T(x, \cdot), \psi_{\mu} \rangle$. From Corollary 2.4.5, we get:

$$\|T_{\infty, m_2} - T\|_{L^2(\Omega \times \Omega)} \leq C \sqrt{\kappa} m_2^{-s}. \quad (2.71)$$

Now, notice that $c_{\mu} \in H^r(\Omega)$. Indeed, for all $0 \leq k \leq r$, we get:

$$\int_{\Omega} (\partial_x^k c_{\mu}(x))^2 dx \quad (2.72)$$

$$= \int_{\Omega} \left(\partial_x^k \int_{\Omega} T(x, y) \psi_{\mu}(y) dy \right)^2 dx \quad (2.73)$$

$$= \int_{\Omega} \left(\int_{\Omega} (\partial_x^k T)(x, y) \psi_{\mu}(y) dy \right)^2 dx \quad (2.74)$$

$$\leq \int_{\Omega} \|(\partial_x^k T)(x, \cdot)\|_{L^2(\Omega)}^2 \|\psi_{\mu}\|_{L^2(\Omega)}^2 dx \quad (2.75)$$

$$= \|(\partial_x^k T)\|_{L^2(\Omega \times \Omega)}^2 < +\infty. \quad (2.76)$$

²This was also the basic idea in our recent paper [100].

Therefore, we can use Lemma 2.4.2 again to show:

$$\|T_{\infty, m_2} - T_{m_1, m_2}\|_{L^2(\Omega \times \Omega)} \leq C\sqrt{\kappa}m_1^{-r}. \quad (2.77)$$

Finally, using the triangle inequality, we get:

$$\|T - T_{m_1, m_2}\|_{HS} \leq C\sqrt{\kappa}(m_1^{-r} + m_2^{-s}). \quad (2.78)$$

By setting $m_1 = m_2^{s/r}$, the two approximation errors in the right-hand side of (2.78) are balanced. This motivates the choice of m_1 and m_2 indicated in the theorem. $\square$

The approximation result in inequality (2.69) is worse than the previous ones. For instance if $r = s$, then the bound becomes $O(\sqrt{\kappa}m^{-s/2})$ instead of $O(\sqrt{\kappa}m^{-s})$ in all previous theorems. The great advantage of this representation is the operator storage: until now, the whole set of vectors $(\mathbf{c}_\mu)$ had to be stored ($O(\kappa nm)$ values), while now, only m coefficients $c_{\lambda, \mu}$ are required. For instance, in the case $r = s$, for an equivalent precision, the storage cost of the new representation is $O(\kappa m^2)$ instead of $O(\kappa nm)$.

In addition, evaluating matrix-vector products can be achieved rapidly by using the following trick:

$$\mathbf{H}_m \mathbf{u} = \sum_{|\lambda| \leq \log_2(m_1)} \sum_{|\mu| \leq \log_2(m_2)} c_{\lambda, \mu} \psi_\lambda \star (\psi_\mu \odot \mathbf{u}) \quad (2.79)$$

$$= \sum_{|\mu| \leq \log_2(m_2)} \left(\sum_{|\lambda| \leq \log_2(m_1)} c_{\lambda, \mu} \psi_\lambda \right) \star (\psi_\mu \odot \mathbf{u}). \quad (2.80)$$

By letting $\tilde{\mathbf{c}}_\mu = \sum_{|\lambda| \leq \log_2(m_1)} c_{\lambda, \mu} \psi_\lambda$, we get

$$\mathbf{H}_m \mathbf{u} = \sum_{|\mu| \leq \log_2(m_2)} \tilde{\mathbf{c}}_\mu \star (\psi_\mu \odot \mathbf{u}). \quad (2.81)$$

which can be computed in $O(m_2 \kappa n \log_2(\kappa n))$ operations. This remark leads to the following proposition.

Proposition 2.4.2. *Assume that $T \in H^{r,s}(\Omega \times \Omega)$ and that it satisfies Assumption 2.3.2. Set $m = \left\lceil \left(\frac{\epsilon}{C\sqrt{\kappa}} \right)^{-(r+s)/rs} \right\rceil$. Then the operator H_m defined in Theorem 2.4.2 satisfies $\|H - H_m\|_{HS} \leq \epsilon$ and the number of operations necessary to evaluate a product with H_m or H_m^* is bounded above by $O\left(\epsilon^{-1/s} \kappa^{\frac{2s+1}{2s}} n \log_2(n)\right)$.*

Notice that the complexity of matrix-vector products is unchanged compared to the wavelet or spline approaches with a much better compression ability. However, this method requires a preprocessing to compute $\tilde{\mathbf{c}}_\mu$ with complexity $\epsilon^{-1/s} \kappa^{1/2s} n$.

2.5 Adaptive decompositions

In the last section, all methods shared the same principle: project $T(x, \cdot)$ on a fixed basis for each $x \in \Omega$. Instead of fixing a basis, one can try to find a basis adapted to the operator at hand. This idea was proposed in [112] and [85].




Figure 2.4 – Meyer’s representations of the operators in Examples 2.3.1, 2.3.2 and 2.3.3 in $\log_{10}$ scale.

2.5.1 Singular value decompositions

The authors of [112] proposed to use a singular value decomposition (SVD) of the TVIR in order to construct the functions h_k and w_k . In this section we first detail this idea and then analyze it from an approximation theoretic point of view. Let $J : L^2(\Omega) \rightarrow L^2(\Omega)$ denote the linear integral operator with kernel $T \in \mathcal{T}^s$. First notice that J is a Hilbert-Schmidt operator since $\|J\|_{HS} = \|H\|_{HS}$. By Lemma 2.2.2 and since Hilbert-Schmidt operators are compact, there exists two Hilbert bases (e_k) and (f_k) of $L^2(\Omega)$ such that J can be decomposed as

$$J = \sum_{k \geq 1} \sigma_k \cdot e_k \otimes f_k, \quad (2.82)$$

leading to

$$T(x, y) = \sum_{k=1}^{+\infty} \sigma_k f_k(x) e_k(y). \quad (2.83)$$

The following result is a standard.

Theorem 2.5.1. *For a given m , a set of functions $(h_k)_{1 \leq k \leq m}$ and $(w_k)_{1 \leq k \leq m}$ that minimizes $\|H_m - H\|_{HS}$ is given by:*

$$h_k = \sigma_k f_k \quad \text{and} \quad w_k = e_k. \quad (2.84)$$

Moreover, if $T(x, \cdot)$ satisfies Assumptions 2.3.1 and 2.3.2, we get:

$$\|H_m - H\|_{HS} = O(\sqrt{\kappa} m^{-s}). \quad (2.85)$$

Proof. The proof of optimality (2.85) is standard. Since T_m is the best rank m approximation of T , it is necessarily better than bound (2.40), yielding (2.85). $\square$

Theorem 2.5.2. *For all $\epsilon > 0$ and $m < n$, there exists an operator H with TVIR satisfying 2.3.1 and 2.3.2 such that:*

$$\|H_m - H\|_{HS} \geq C \sqrt{\kappa} m^{-(s+\epsilon)}. \quad (2.86)$$

Proof. In order to prove, (2.86), we construct a “worst case” TVIR T . We first begin by constructing a kernel T with $\kappa = 1$ to show a simple pathological TVIR. Define T by:

$$T(x, y) = \sum_{k \in \mathbb{Z}} \sigma_k f_k(x) f_k(y), \quad (2.87)$$

where $f_k(x) = \exp(2i\pi kx)$ is the k -th Fourier atom, $\sigma_0 = 0$ and $\sigma_k = \sigma_{-k} = \frac{1}{|k|^{s+1/2+\epsilon/2}}$ for $|k| \geq 1$. With this choice,

$$T(x, y) = \sum_{|k| \leq N} 2\sigma_k \cos(2\pi(x + y)) \quad (2.88)$$

is real for all (x, y) . We now prove that $T \in \mathcal{T}^s$. The k -th Fourier coefficient of $T(x, \cdot)$ is given by $\sigma_k f_k(x)$ which is bounded by σ_k for all x . By Lemma 2.2.1, $T(x, \cdot)$ therefore belongs to $H^s(\Omega)$ for all $x \in \Omega$. By construction, the spectrum of T is $(|\sigma_k|)_{k \in \mathbb{N}}$, therefore for any rank $2m + 1$ approximation of T , we get:

$$\|T - T_{2m+1}\|_{HS}^2 \geq \sum_{|k| \geq m+1} \frac{1}{|k|^{2s+1+\epsilon}} \quad (2.89)$$

$$\geq \int_{m+1}^{\infty} \frac{2}{t^{2s+1+\epsilon}} dt \quad (2.90)$$

$$= \frac{1}{2s + \epsilon} \frac{2}{(m + 1)^{2s+\epsilon}} \quad (2.91)$$

$$= O(m^{-2s-\epsilon}), \quad (2.92)$$

proving the result for $\kappa = 1$. Notice that the kernel K of the operator with TVIR T only depends on x :

$$K(x, y) = \sum_{|k| \leq N} 2\sigma_k \cos(2\pi x). \quad (2.93)$$

Therefore the worst case TVIR exhibited here is that of a *rank 1* operator H . Obviously, it cannot be well approximated by product-convolution expansions.

Let us now construct a TVIR satisfying Assumption 2.3.2. For this, we first construct an orthonormal basis $(\tilde{f}_k)_{k \in \mathbb{Z}}$ of $L^2([-\kappa/2, \kappa/2])$ defined by:

$$\tilde{f}_k(x) = \begin{cases} \frac{1}{\sqrt{\kappa}} f_k\left(\frac{x}{\kappa}\right) & \text{if } |x| \leq \frac{\kappa}{2}, \\ 0 & \text{otherwise.} \end{cases} \quad (2.94)$$

The worst case operator considered now is defined by:

$$T(x, y) = \sum_{k \in \mathbb{Z}} \tilde{\sigma}_k \tilde{f}_k(x) f_k(y). \quad (2.95)$$

Its spectrum is $(|\tilde{\sigma}_k|)_{k \in \mathbb{Z}}$, and we get

$$|\langle T(x, \cdot), f_k \rangle| = |\tilde{\sigma}_k \tilde{f}_k(x)| = \frac{1}{\kappa} |\tilde{\sigma}_k|. \quad (2.96)$$

By Lemma 2.4.1, if $\tilde{\sigma}_k = \frac{\kappa}{(1+|k|2)^s|k|^{1+\epsilon}}$, then $\|T(x, \cdot)\|_{H^s(\Omega)}$ is uniformly bounded by a constant independent of κ . Moreover, by reproducing the reasoning in (2.89), we get:

$$\|T - T_{2m+1}\|_{HS}^2 = O(\kappa m^{-2s-\epsilon}). \quad (2.97)$$

□

Even if the SVD provides an optimal decomposition, there is no guarantee that functions e_k are supported on an interval of small size. As an example, it suffices to consider the “worst case” TVIR given in equation (2.87). Therefore, vectors $\mathbf{w}_k$ are generically supported on intervals of size $p = n$. This yields the following proposition.

Corollary 2.5.1. *Let $\epsilon > 0$ and set $m = \lceil C\epsilon^{-1/s}\kappa^{1/2s} \rceil$. Then H_m satisfies $\|H - H_m\|_{HS} \leq \epsilon$ and a product with H_m and H_m^* can be evaluated with no more than $O(\kappa^{1/2s}n \log n \epsilon^{-1/s})$ operations.*

Computing the first m singular vectors in (2.83) can be achieved in roughly $O(\kappa n^2 \log(m))$ operations thanks to recent advances in randomized algorithms [127]. The storage cost for this approach is $O(mn)$ since the vectors $\mathbf{e}_k$ have no reason to be compactly supported.

2.5.2 The optimization approach in [85]

In [85], the authors propose to construct the windowing functions w_k and the filters h_k using constrained optimization procedures. For a fixed m , they propose solving:

$$\min_{(h_k, w_k)_{1 \leq k \leq m}} \left\| T - \sum_{k=1}^m h_k \otimes w_k \right\|_{HS}^2 \quad (2.98)$$

under an additional constraint that $\text{supp}(w_k) \subset \omega_k$ with ω_k chosen so that $\cup_{k=1}^m \omega_k = \Omega$. A decomposition of type (2.98) is known as structured low rank approximation [64]. This problem is non convex and to the best of our knowledge, there currently exists no algorithm running in a reasonable time to find its global minimizer. It can however be solved approximately using alternating minimization like algorithms.

Depending on the choice of the supports ω_k , different convergence rates can be expected. However, by using the results for B-splines in section 2.4.2, we obtain the following proposition.

Proposition 2.5.1. *Set $\omega_k = [(k-1)/m, k/m + s/m]$ and let $(h_k, w_k)_{1 \leq k \leq m}$ denote the global minimizer of (2.98). Define T_m by $T_m(x, y) = \sum_{k=1}^m h_k(x)w_k(y)$. Then:*

$$\|T - T_m\|_{HS}^2 \leq C\sqrt{\kappa}m^{-s}. \quad (2.99)$$

Set $m = \lceil \kappa^{1/2s}C\epsilon^{-1/s} \rceil$, then $\|H_m - H\|_{HS} \leq \epsilon$ and the evaluation of a product with H_m or H_m^ is of order*

$$O(\kappa^{1+1/2s}n \log(n)\epsilon^{-1/s}). \quad (2.100)$$

Proof. First notice that cardinal B-Splines are also supported on $[(k-1)/m, k/m + s/m]$. Since the method in [85] provides the best choices for (h_k, w_k) , the distance $\|H_m - H\|_{HS}$ is necessarily lower than that obtained using B-splines in Corollary 2.4.3. $\square$

Finally, let us mention that - owing to Corollary 2.4.5 - it might be interesting to use the optimization approach (2.98) with windows of varying sizes.

2.6 Summary and extensions

2.6.1 A summary of all results

Table 2.1 summarizes the results derived so far under Assumptions 2.3.1 and 2.3.2. In the particular case of Meyer's methods, we assume that $T \in H^{r,s}(\Omega \times \Omega)$ instead of Assumption 2.3.1. As can be seen in this table, different methods should be used depending on the application. The best methods are:

- Wavelets: they are adaptive, have a relatively low construction complexity, and matrix-vector products also have the best complexity.
- Meyer: this method has a big advantage in terms of storage. The operator can be represented very compactly with this approach. It has a good potential for problems where the operator should be inferred (e.g. blind deblurring). It however requires stronger regularity assumptions.
- The SVD and the method proposed in [85] both share an optimal adaptivity. The representation however depends on the operator and it is more costly to evaluate it.

<i>Method</i>	<i>Approximation</i>	<i>Product</i>	<i>Construction</i>	<i>Storage</i>	<i>Adaptivity</i>
Fourier 2.4.1	$O\left(\kappa^{\frac{1}{2}} m^{-s}\right)$	$O\left(\kappa^{\frac{1}{2s}} n \log(n) \epsilon^{-\frac{1}{s}}\right)$	$O(\kappa n^2 \log(n))$	$O(m\kappa n)$	$\times$
B-Splines 2.4.2	$O\left(\kappa^{\frac{1}{2}} m^{-s}\right)$	$O\left(\kappa^{\frac{2s+1}{2s}} n \log(n) \epsilon^{-\frac{1}{s}}\right)$	$O(\kappa n^2 \log(n))$	$O(m\kappa n)$	$\times$
Wavelets 2.4.3	$O\left(\kappa^{\frac{1}{2}} m^{-s}\right)$	$O\left(\kappa^{\frac{2s+1}{2s}} n \log(n) \epsilon^{-\frac{1}{s}}\right)$	$O(\kappa s n^2)$	$O(m\kappa n)$	$\checkmark$
Meyer 2.4.5	$O\left(\kappa^{\frac{1}{2}} m^{-\frac{rs}{r+s}}\right)$	$O\left(\kappa^{\frac{2s+1}{2s}} n \log(n) \epsilon^{-\frac{1}{s}}\right)$	$O(s n^2)$	$O(m)$	$\checkmark$
SVD 2.5.1	$O\left(\kappa^{\frac{1}{2}} m^{-s}\right)$	$O\left(\kappa^{\frac{1}{2s}} n \log(n) \epsilon^{-\frac{1}{s}}\right)$	$O(\kappa n^2 \log(m))$	$O(mn)$	$\checkmark$
[85] 2.5.2	$O\left(\kappa^{\frac{1}{2}} m^{-s}\right)$	$O\left(\kappa^{\frac{2s+1}{2s}} n \log(n) \epsilon^{-\frac{1}{s}}\right)$	High (iterative)	$O(m\kappa n)$	$\checkmark$

Table 2.1 – Summary of the properties of different constructions. Approximation $\equiv$ approximation rates in terms of m . Product $\equiv$ matrix-vector product complexity to get an ϵ approximation. Construction $\equiv$ complexity of the construction of order m representation. Storage $\equiv$ cost of storage of a given representation. Adaptivity $\equiv$ ability to automatically adapt to different input operators.

2.6.2 Extensions to higher dimensions

Most of the results provided in this paper are based on standard approximation results in 1D, such as Lemmas 2.4.1, 2.4.2 and 2.4.1. All these lemmas can be extended to higher dimension and we refer the interested reader to [89, 163, 88, 193] for more details.

We now assume that $\Omega = [0, 1]^d$ and that the diameter of the impulse responses is bounded by $\kappa \in [0, 1]$. Using the mentioned results, it is straightforward to show that the approximation rate of all methods now becomes

$$\|H - H_m\|_{HS} = O(\kappa^{d/2} m^{-s/d}). \quad (2.101)$$

The space Ω can be discretized on a finite dimensional space of size n^d . Similarly, all complexity results given in Table 2.1 are still valid by replacing n by n^d , $\epsilon^{-1/s}$ by $\epsilon^{-d/s}$ and κ by κ^d .

2.6.3 Extensions to less regular spaces

Until now, we assumed that the TVIR T belongs to Hilbert spaces (see e.g. Assumption 2.3.1). This assumption was deliberately chosen easy to clarify the presentation. The results can most likely be extended to much more general spaces using nonlinear approximation theory results [88].

For instance, assume that $T \in BV(\Omega \times \Omega)$, the space of functions with bounded variations. Then, it is well known (see e.g. [68]) that T can be expressed compactly on a Hilbert basis of tensor-product wavelets. Therefore, the product-convolution expansion 2.4 could be used by using the trick proposed in 2.81.

Similarly, most of the kernels found in partial differential equations (e.g. Calderón-Zygmund operators) are singular at the origin. Once again, it is well known [170] that wavelets are able to capture the singularities and the proposed methods can most likely be applied to this setting too.

A precise setting useful for applications requires more work and we leave this issue open for future work.

2.6.4 Controls in other norms

In all the paper we only controlled the Hilbert-Schmidt norm $\|\cdot\|_{HS}$. This choice simplifies the analysis and also allows getting bounds for the spectral norm

$$\|H\|_{2 \rightarrow 2} = \sup_{\|u\|_{L^2(\Omega)} \leq 1} \|Hu\|_{L^2(\Omega)}, \quad (2.102)$$

since $\|H\|_{2 \rightarrow 2} \leq \|H\|_{HS}$. In applications, it often makes sense to consider other operator norms defined by

$$\|H\|_{X \rightarrow Y} = \sup_{\|u\|_X \leq 1} \|Hu\|_Y, \quad (2.103)$$

where $\|\cdot\|_X$ and $\|\cdot\|_Y$ are norms characterizing some function spaces. We showed in [100] that this idea could highly improve practical approximation results.

Unfortunately, it is not clear yet how to extend the proposed results and algorithms to such a setting and we leave this question open for the future. Let us mention that our previous experience shows that this idea can highly change the method's efficiency.

2.7 Conclusion

In this paper, we analyzed the approximation rates and numerical complexity of product-convolution expansions. This approach was shown to be efficient whenever the time or space varying impulse response of the operator is well approximated by a low rank tensor. We showed that this situation occurs under mild regularity assumptions, making the approach relevant for a large class of applications. We also proposed a few original implementations of this methods based on orthogonal wavelet decompositions and analyzed their respective advantages precisely. Finally, we suggested a few ideas to further improve the practical efficiency of the method.

3

Sparse wavelet representations of spatially varying blurring operators

Résumé : *Ce chapitre est issu de [100] publié dans SIAM Journal on Imaging Science 2015 en collaboration avec Pierre Weiss. La restauration d'images dégradées par des flous variables est un problème rencontré dans beaucoup de disciplines comme l'astrophysique, la vision par ordinateur et l'imagerie médicale. L'un des plus gros challenges dans la résolution de ce problème est d'élaborer des algorithmes numériquement efficaces qui permettent d'approcher des opérateurs intégraux. Nous introduisons une nouvelle méthode reposant sur des approximations parcimonieuses des opérateurs de flou dans des bases d'ondelettes. Cette méthode requiert $O(N\epsilon^{-d/M})$ opérations pour produire une ϵ -approximation où N est le nombre de pixels d'une image d -dimensionnelle et $M \geq 1$ est une scalaire décrivant la régularité du flou. De plus, nous proposons une idée originale qui permet de définir a priori une structure de parcimonie quand seulement la régularité de l'opérateur est connue. Les expériences numériques nous montrent que notre algorithme fournit des résultats significativement meilleurs que les méthodes de l'état-de-l'art elles que les méthodes de produit-convolution.*

Abstract: *Restoring images degraded by spatially varying blur is a problem encountered in many disciplines such as astrophysics, computer vision or biomedical imaging. One of the main challenges to perform this task is to design efficient numerical algorithms to approximate integral operators. We introduce a new method based on a sparse approximation of the blurring operator in the wavelet domain. This method requires $\mathcal{O}(N\epsilon^{-d/M})$ operations to provide ϵ -approximations, where N is the number of pixels of a d -dimensional image and $M \geq 1$ is a scalar describing the regularity of the blur kernel. In addition, we propose original methods to define sparsity patterns when only the operators regularity is known. Numerical experiments reveal that our algorithm provides a significant improvement compared to standard methods based on windowed convolutions.*

Contents

3.1	Introduction	79
3.1.1	Contributions of the paper	81
3.1.2	Outline of the paper	81
3.2	Notation	81
3.3	Blurring operators and their mathematical properties	84
3.3.1	A mathematical description of blurring operators	84
3.4	Wavelet representation of the blurring operator	86
3.4.1	Discretization of the operator by projection	86
3.4.2	Discretization by projection on a wavelet basis	87
3.4.3	Theoretical guarantees with sparse approximations	87
3.5	Identification of sparsity patterns	89
3.5.1	Problem formalization	89
3.5.2	Link with the approach in [241]	91
3.5.3	An algorithm when Θ is known	92
3.5.4	An algorithm when Θ is unknown	92
3.6	Numerical experiments	96
3.6.1	Preliminaries	98
3.6.2	Application to direct problems	100
3.6.3	Application to inverse problems	109
3.7	Conclusion	110
3.7.1	Brief summary	110
3.7.2	Outlook	110
3.8	Proof of Lemma 3.4.1	113
3.9	Proof of Theorem 3.4.1	115

3.1 Introduction

The problem of image restoration in the presence of spatially varying blur appears in many domains. Examples of applications in computer vision, biomedical imaging and astronomy are shown in Figures 3.1 and 3.2 respectively. In this paper, we propose new solutions to address one of the main difficulties associated to this problem: the computational evaluation of matrix-vector products.

A spatially variant blurring operator can be modelled as a linear operator and therefore be represented by a matrix $\mathbf{H}$ of size $N \times N$, where N represents the number of pixels of a d -dimensional image. Sizes of typical images range from $N = 10^6$ for small 2D images, to $N = 10^{10}$ for large 2D or 3D images. Storing matrices and computing matrix-vector products using the standard representation is impossible for such sizes: it amounts to tera or exabytes of data/operations. In cases where the Point Spread Functions (PSF) supports are sufficiently small in average over the image domain, the operator can be coded as a sparse matrix and be applied using traditional approaches. However, in many practical applications this method turns out to be too intensive and cannot be applied with decent computing times. This may be due to i) large PSFs supports or ii) the need for super-resolution applications where the PSFs sizes increase with the resolution. Spatially varying blurring matrices therefore require the development of computational tools to compress them and evaluate them in an efficient way.

Existing approaches

To the best of our knowledge, the first attempts to address this issue appeared at the beginning of the seventies (see e.g. [203]). Since then, many techniques were proposed. We describe them briefly below

Composition of diffeomorphisms and convolutions One of the first method proposed to reduce the computational complexity, is based on first applying a diffeomorphism to the image domain [203, 204, 168, 223, 105] followed by a convolution using FFTs and an inverse diffeomorphism. The diffeomorphism is chosen in order to transform the spatially varying blur into an invariant one. This approach suffers from two important drawbacks:

- first it was shown that not all spatially varying kernel can be approximated by this approach [168],
- second, this method requires good interpolation methods and the use of Euclidean grids with small grid size in order to correctly estimate integrals.

Separable approximations Another common idea is to approximate the kernel of the operator by a separable one that operates in only one dimension. The computational complexity of a product is thus reduced to d applications of one-dimensional operators. It drastically improves the performance of algorithms. For instance, in 3D fluorescence microscopy, the authors of [195, 158, 23, 251] proposed to approximate PSFs by anisotropic

Gaussians and assumed that the Gaussian variances only vary along one direction (e.g., the direction of light propagation). The separability assumption implies that both the PSF and its variations are separable. Unfortunately, most physically realistic PSFs are not separable and do not vary in a separable manner (see e.g., Figure 3.3). This method is therefore usually too crude.

Wavelet or Gabor multipliers Some works [57, 103, 109, 136] proposed to approximate blurring operators $\mathbf{H}$ using operators diagonal in wavelet bases, wavelet packet or Gabor frames. This idea consists of defining an approximation $\tilde{\mathbf{H}}$ of kind $\tilde{\mathbf{H}} = \Psi \Sigma \Psi^*$, where Ψ^* and Ψ are wavelet or Gabor transforms and Σ is a diagonal matrix. These diagonal approximations mimic the fact that shift-invariant operators are diagonal in the Fourier domain. These approaches lead to fast $\mathcal{O}(N)$ or $\mathcal{O}(N \log(N))$ algorithms to compute matrix-vector products. In [103], we proposed to deblur images using diagonal approximations of the blurring operators in redundant wavelet packet bases. This approximation was shown to be fast and efficient in deblurring images when the exact operator was scarcely known or in high noise levels. It is however too coarse for applications with low noise levels. This approach seems however promising. Gabor multipliers are considered the state-of-the-art for 1D signals in OFDM systems for instance (slowly varying smoothing operators).

Weighted convolutions Probably the most commonly used approaches consist of approximating the integral kernel by spatially weighted sum of convolutions. Among these approaches two different ideas have been explored. The first one will be called *windowed convolutions* in this paper and appeared in [179, 180, 128, 132, 85]. The second one was proposed in [112] and consists of expanding the PSFs in a common basis of small dimensionality.

Windowed convolutions consists of locally stationary approximations of the kernel. We advise the reading of [85] for an up-to-date description of this approach and its numerous refinements. The main idea is to decompose the image domain into subregions and perform a convolution on each subregion. The results are then gathered together to obtain the blurred image. In its simplest form, this approach consists in partitioning the domain Ω in squares of equal sizes. More advanced strategies consist in decomposing the domain with overlapping subregions. The blurred image can then be obtained by using windowing functions that interpolate the kernel between subregions (see, e.g., [179, 132, 85]). Various methods have been proposed to interpolate the PSF. In [132], a linear interpolation is performed, and in [85] higher order interpolation of the PSF are handled.

Sparse wavelet approximations The approach studied in this paper was proposed recently and independently in [241, 242, 102]. The main idea is to represent the operator in the wavelet domain by using a change of basis. This change of basis, followed by a thresholding operation allows sparsifying the operator and use sparse matrix-vector

products. The main objective of this work is to provide solid theoretical foundations to these approaches.

3.1.1 Contributions of the paper

Our first contribution is the design of a new approach based on sparse approximation of $\mathbf{H}$ in the wavelet domain. Using techniques initially developed for pseudo-differential operators [29, 170], we show that approximations $\tilde{\mathbf{H}}$ satisfying $\|\mathbf{H} - \tilde{\mathbf{H}}\|_{2 \rightarrow 2} \leq \epsilon$, can be obtained with this new technique, in no more than $\mathcal{O}(N\epsilon^{-d/M})$ operations. In this complexity bound, $M \geq 1$ is an integer that describes the smoothness of the blur kernel.

Controlling the spectral norm is usually of little relevance in image processing. Our second contribution is the design of algorithms that iteratively construct sparse matrix patterns adapted to the structure of images. These algorithms rely on the fact that both natural images and operators can be compressed simultaneously in the same wavelet basis.

As a third contribution, we propose an algorithm to design a generic sparsity structure when only the operators regularity is known. This paves the way to the use of wavelet based approaches in blind deblurring problems where operators need to be inferred from the data.

We finish the paper by numerical experiments. We show that the proposed algorithms allow significant speed ups compared to some windowed convolutions based methods.

Let us emphasize that the present paper is a continuation of our recent contribution [102]. The main evolution is that i) we provide all the theoretical foundations of the approach with precise hypotheses, ii) we propose a method to automatically generate adequate sparsity patterns and iii) we conduct a thorough numerical analysis of the method.

3.1.2 Outline of the paper

The outline of this paper is as follows. We introduce the notation used throughout the paper in Section 3.2. We propose an original mathematical description of blurring operators appearing in image processing in Section 3.3. We introduce the proposed method and analyze its theoretical efficiency Section 3.4. We then propose various algorithms to design good sparsity patterns in Section 3.5. Finally, we perform numerical tests to analyze the proposed method and compare it to the standard windowed convolutions based methods in Section 3.6.

3.2 Notation

In this paper, we consider d dimensional images defined on a domain $\Omega = [0, 1]^d$. The space $L^2(\Omega)$ will denote the space of squared integrable functions defined on Ω . The set of compactly supported $\mathcal{C}^\infty(\Omega)$ functions is denoted $D(\Omega)$. Its dual space (the set of distributions) is denoted $D'(\Omega)$.




Figure 3.1 – An example in computer vision. Image degraded by spatially varying blur due to a camera shake. Images are from [131] and used here by courtesy of Michael Hirsch.




Figure 3.2 – An example in biology. Image of a multicellular tumor spheroid imaged in 3D using Selective Plane Illumination Microscope (SPIM). Fluorescence beads (in green) are inserted in the tumor model and allow the observation of the PSF at different locations. Nuclei are stained in red. On the left-hand-side, 3D PSFs outside the sample are observed. On the right-hand-side, 3D PSFs inside the sample are observed. This image is from [142] and used here by courtesy of Corinne Lorenzo.




Figure 3.3 – Three PSFs displayed in a XZ plan at different z depths: $-20\mu\text{m}$, $0\mu\text{m}$ and $20\mu\text{m}$. PSFs are generated using Gibson and Lanni 3D optical model from the PSF Generator [147]. The parameters used are $n_i = 1.5$, $n_s = 1.33$, $t_i = 150\mu\text{m}$, $\text{NA} = 1.4$ and a wavelength of 610nm .

Let $\alpha = (\alpha_1, \dots, \alpha_d)$ denote a multi-index. The sum of its components is denoted $|\alpha| = \sum_{i=1}^d \alpha_i$. The Sobolev spaces $W^{M,p}$ are defined as the set of functions $f \in L^p$ with partial derivatives up to order M in L^p where $p \in [1, +\infty]$ and $M \in \mathbb{N}$. These spaces, equipped with the following norm are Banach spaces

$$\|f\|_{W^{M,p}} = \|f\|_{L^p} + |f|_{W^{M,p}}, \quad \text{where,} \quad |f|_{W^{M,p}} = \sum_{|\alpha|=M} \|\partial^\alpha f\|_{L^p}. \quad (3.1)$$

In this notation, $\partial^\alpha f = \frac{\partial^{\alpha_1}}{\partial x_1^{\alpha_1}} \dots \frac{\partial^{\alpha_d}}{\partial x_d^{\alpha_d}} f$.

Let X and Y denote two metric spaces endowed with their respective norms $\|\cdot\|_X$ and $\|\cdot\|_Y$. In all the paper $H : X \rightarrow Y$ will denote a linear operator and H^* its adjoint operator. The subordinate operator norm is defined by

$$\|H\|_{X \rightarrow Y} = \sup_{x \in X, \|x\|_X = 1} \|Hx\|_Y.$$

The notation $\|H\|_{p \rightarrow q}$ corresponds to the case where X and Y are endowed with the standard L^p and L^q norms. In all the paper, operators acting in a continuous domain are written in plain text format H . Finite dimensional matrices are written in bold fonts $\mathbf{H}$. Approximation operators will be denoted $\tilde{H}$ in the continuous domain or $\tilde{\mathbf{H}}$ in the discrete domain.

In this paper we consider a compactly supported wavelet basis of $L^2(\Omega)$. We first introduce wavelet basis of $L^2([0, 1])$. We let ϕ and ψ denote the scaling and mother wavelets. We assume that the mother-wavelet ψ has M vanishing moments, i.e.

$$\text{for all } 0 \leq m < M, \quad \int_{[0,1]} t^m \psi(t) dt = 0.$$

We assume that $\text{supp}(\psi) = [-c(M)/2, c(M)/2]$. Note that $c(M) \geq 2M - 1$, with equality for Daubechies wavelets, see, e.g., [164, Theorem 7.9, p. 294].

We define translated and dilated versions of the wavelets for $j \geq 0$ as follows

$$\begin{aligned}\phi_{j,l} &= 2^{j/2} \phi(2^j \cdot - l), \\ \psi_{j,l} &= 2^{j/2} \psi(2^j \cdot - l),\end{aligned}\tag{3.2}$$

with $l \in \mathcal{T}_j$ and $\mathcal{T}_j = \{0, \dots, 2^j - 1\}$.

In dimension d , we use separable wavelet bases, see, e.g., [164, Theorem 7.26, p. 348]. Let $m = (m_1, \dots, m_d)$. Define $\rho_{j,l}^0 = \phi_{j,l}$ and $\rho_{j,l}^1 = \psi_{j,l}$. Let $e = (e_1, \dots, e_d) \in \{0, 1\}^d$. For ease of reading, we will use the shorthand notation $\lambda = (j, m, e)$. We also denote

$$\Lambda_0 = \left\{ (j, m, e) \mid j \in \mathbb{Z}, m \in \mathcal{T}_j, e \in \{0, 1\}^d \right\}$$

and

$$\Lambda = \left\{ (j, m, e) \mid j \in \mathbb{Z}, m \in \mathcal{T}_j, e \in \{0, 1\}^d \setminus \{0\} \right\}.$$

Wavelet ψ_λ is defined by $\psi_\lambda(x_1, \dots, x_d) = \psi_{j,m}^e(x_1, \dots, x_d) = \rho_{j,m_1}^{e_1}(x_1) \dots \rho_{j,m_d}^{e_d}(x_d)$. Elements of the separable wavelet basis consist of tensor products of scaling and mother wavelets at the same scale. Note that if $e \neq 0$ wavelet $\psi_{j,m}^e$ has M vanishing moments in $\mathbb{R}^d$. We let $I_{j,m} = \cup_e \text{supp } \psi_{j,m}^e$ and $I_\lambda = \text{supp } \psi_\lambda$.

We assume that every function $f \in L^2(\Omega)$ can be written as

$$\begin{aligned}u &= \langle u, \psi_{0,0}^0 \rangle \psi_{0,0}^0 + \sum_{e \in \{0,1\}^d \setminus \{0\}} \sum_{j=0}^{+\infty} \sum_{m \in \mathcal{T}_j} \langle u, \psi_{j,m}^e \rangle \psi_{j,m}^e \\ &= \langle u, \psi_{0,0}^0 \rangle \psi_{0,0}^0 + \sum_{\lambda \in \Lambda} \langle u, \psi_\lambda \rangle \psi_\lambda \\ &= \sum_{\lambda \in \Lambda_0} \langle u, \psi_\lambda \rangle \psi_\lambda\end{aligned}$$

This is a slight abuse since wavelets defined in (3.2) do not define a Hilbert basis of $L^2([0, 1]^d)$. There are various ways to define wavelet bases on the interval [71] and wavelets having a support intersecting the boundary should be given a different definition. We stick to these definitions to keep the proofs simple.

We let $\Psi^* : L^2(\Omega) \rightarrow l^2(\mathbb{Z})$ denote the wavelet decomposition operator and $\Psi : l^2(\mathbb{Z}) \rightarrow L^2(\Omega)$ its associated reconstruction operator. The discrete wavelet transform is denoted $\Psi : \mathbb{R}^N \rightarrow \mathbb{R}^N$. We refer to [164, 79, 71] for more details on the construction of wavelet bases.

3.3 Blurring operators and their mathematical properties

3.3.1 A mathematical description of blurring operators

In this paper, we consider d -dimensional real-valued images defined on a domain $\Omega = [0, 1]^d$, where d denotes the space dimension. We consider a blurring operator $H :$

$L^2(\Omega) \rightarrow L^2(\Omega)$ defined for any $u \in L^2(\Omega)$ by the following integral operator:

$$\forall x \in \Omega, \quad Hu(x) = \int_{y \in \Omega} K(x, y)u(y)dy. \quad (3.3)$$

The function $K : \Omega \times \Omega \rightarrow \mathbb{R}$ is a kernel that defines the Point Spread Function (PSF) $K(\cdot, y)$ at each location $y \in \Omega$. The image Hu is the blurred version of u . By the Schwartz kernel theorem, any linear operator which is continuous from $D(\Omega) \rightarrow D'(\Omega)$, can be represented by (3.3) where K is a generalized function. We thus need to determine properties of K specific to blurring operators that will allow to design efficient numerical algorithms to approximate the integral (3.3).

We propose a definition of the class of blurring operators below.

Definition 3.3.1 (Blurring operators). Let $M \in \mathbb{N}$ and $f : [0, 1] \rightarrow \mathbb{R}_+$ denote a non-increasing bounded function. An integral operator is called a blurring operator in the class $\mathcal{A}(M, f)$ if it satisfies the following properties:

1. Its kernel $K \in W^{M, \infty}(\Omega \times \Omega)$;
2. The partial derivatives of K satisfy:

$$(a) \quad \forall |\alpha| \leq M, \quad \forall (x, y) \in \Omega \times \Omega, \quad |\partial_x^\alpha K(x, y)| \leq f(\|x - y\|_\infty). \quad (3.4)$$

$$(b) \quad \forall |\alpha| \leq M, \quad \forall (x, y) \in \Omega \times \Omega, \quad \left| \partial_y^\alpha K(x, y) \right| \leq f(\|x - y\|_\infty). \quad (3.5)$$

Let us justify this model from a physical point of view. Most imaging systems satisfy the following properties:

Spatial decay

The PSFs usually have a bounded support (e.g. motion blurs, convolution with the CCD sensors support) or at least a fast spatial decay (Airy pattern, Gaussian blurs,...). This property can be modelled as property 2a. For instance, the 2D Airy disk describing the PSF due to diffraction of light in a circular aperture satisfies 2a with $f(r) = \frac{1}{(1+r)^4}$ (see e.g. [35]).

PSF smoothness

In most imaging applications, the PSF at $y \in \Omega$, $K(\cdot, y)$ is smooth. Indeed it is the result of a convolution with the acquisition device impulse response which is smooth (e.g. Airy disk). This assumption motivates inequality (3.4).

PSFs variations are smooth

We assume that the PSF does not vary abruptly on the image domain. This property can be modelled by inequality (3.5). It does not hold true in all applications. For instance, when sharp discontinuities occur in the depth maps, the PSFs can only be considered as piecewise regular. This assumption simplifies the analysis of

numerical procedures to approximate H . Moreover, it seems reasonable in many settings. For instance, in fluorescence microscopy, the PSF width (or Strehl ratio) mostly depends on the optical thickness, i.e. the quantity of matter laser light has to go through, and this quantity is intrinsically continuous. Even in cases where the PSFs variations are not smooth, the discontinuities locations are usually known only approximately and it seems important to smooth the transitions in order to avoid reconstruction artifacts [14].

Remark 3.3.1. A standard assumption in image processing is that the constant functions are preserved by the operator H . This hypothesis ensures that brightness is preserved on the image domain. In this paper we do not make this assumption and thus encompass image formation models comprising blur and attenuation. Handling attenuation is crucial in domains such as fluorescence microscopy.

Remark 3.3.2. The above properties are important to derive mathematical theories, but only represent an approximation of real systems. The methods proposed in this paper may be applied even if the above properties are not satisfied and are likely to perform well. It is notably possible to relax the boundedness assumption.

3.4 Wavelet representation of the blurring operator

In this section, we show that blurring operators can be well approximated by sparse representations in the wavelet domain. Since H is a linear operator in a Hilbert space, it can be written as $H = \Psi\Theta\Psi^*$, where $\Theta : l^2(\mathbb{Z}) \rightarrow l^2(\mathbb{Z})$ is the (infinite dimensional) matrix representation of the blur operator in the wavelet domain. Matrix Θ is characterized by the coefficients:

$$\theta_{\lambda,\mu} = \langle H\psi_\lambda, \psi_\mu \rangle, \quad \forall \lambda, \mu \in \Lambda. \quad (3.6)$$

In their seminal papers [170, 171, 29], Y. Meyer, R. Coifman, G. Beylkin and V. Rokhlin prove that the coefficients of Θ decrease fastly away from its diagonal for a large class of pseudo-differential operators. They also show that this property allows to design fast numerical algorithms to approximate H , by thresholding Θ to obtain a sparse matrix. In this section, we detail this approach precisely and adapt it to the class of blurring operators.

This section is organized as follows: first, we discuss the interest of approximating H in a wavelet basis rather than using the standard discretization. Second, we provide various theoretical results concerning the number of coefficients necessary to obtain an ϵ -approximation of H .

3.4.1 Discretization of the operator by projection

The proposed method relies on a Galerkin discretization of H . The main idea is to use a projection on a finite dimensional linear subspace $V_q = \text{Span}(\varphi_1, \dots, \varphi_q)$ of $L^2(\Omega)$ where $(\varphi_1, \varphi_2, \dots)$ is an orthonormal basis of $L^2(\Omega)$. We define a projected operator H_q by

$H_q u = P_{V_q} H P_{V_q} u$. where P_{V_q} is the projector on V_q . We can associate a $q \times q$ matrix Θ to this operator defined by $\Theta = (\langle H \varphi_i, \varphi_j \rangle)_{1 \leq i, j \leq q}$.

It is very common in image processing to assume that natural images belong to functional spaces containing functions with some degree of regularity. For instance, images are often assumed to be of bounded total variation [201]. This hypothesis implies that

$$\|u - P_{V_q} u\|_2 = \mathcal{O}(q^{-\alpha}) \quad (3.7)$$

for a certain $\alpha > 0$. For instance, in 1D, if $(\varphi_1, \varphi_2, \dots)$ is a wavelet or a Fourier basis and $u \in H^1(\Omega)$ then $\alpha = 2$. For $u \in BV(\Omega)$ (the space of bounded variation functions), $\alpha = 1$ in 1D and $\alpha = 1/2$ in 2D [164, 72].

Moreover, if we assume that H is a regularizing operator, meaning that $\|Hu - P_{V_q} Hu\|_2 = \mathcal{O}(q^{-\beta})$ with $\beta \geq \alpha$ for all u satisfying (3.7), then we have:

$$\begin{aligned} & \|Hu - H_q u\|_2 \\ &= \|Hu - P_{V_q} H(u + P_{V_q} u - u)\|_2 \\ &\leq \|Hu - P_{V_q} Hu\|_2 + \|P_{V_q} H\|_{2 \rightarrow 2} \|P_{V_q} u - u\|_2 \\ &= \mathcal{O}(q^{-\alpha}). \end{aligned}$$

This simple analysis shows that under mild assumptions, the Galerkin approximation of the operator converges and that the convergence rate can be controlled. The situation is not as easy for standard discretization using finite elements for instance (see, e.g., [238, 16] where a value $\alpha = 1/6$ is obtained in 2D for BV functions, while the simple analysis above leads to $\alpha = 1/2$).

3.4.2 Discretization by projection on a wavelet basis

In order to get a representation of the operator in a finite dimensional setting, we truncate the wavelet representation at scale J . This way, we obtain an operator $\tilde{H}$ acting on a space of dimension N , where $N = 1 + \sum_{j=0}^{J-1} (2^d - 1)2^{dj}$ denotes the numbers of wavelets kept to represent images.

After discretization, it can be written in the following convenient form:

$$\mathbf{H} = \Psi \Theta \Psi^* \quad (3.8)$$

where $\Psi : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is the discrete separable wavelet transform. Matrix Θ is an $N \times N$ matrix which corresponds to a truncated version (also called finite section) of the matrix Θ defined in (3.6).

3.4.3 Theoretical guarantees with sparse approximations

Sparse approximations of integral operators have been studied theoretically in [29, 170]. They then have been successfully used in the numerical analysis of PDEs [76, 70, 67]. Surprisingly, they have been scarcely applied to image processing. The two exceptions we are aware of are the paper [57], where the authors show that wavelet multipliers can

be useful to approximate foveation operators. More recently, [241] proposed an approach that is very much related to that of our paper.

Let us provide a typical result that motivates the proposed approach.

Lemma 3.4.1 (Decay of $\theta_{\lambda,\mu}$). *Assume that H is a blurring operator (see Definition 3.3.1) in the class $\mathcal{A}(M, f)$. Assume that the mother wavelet is compactly supported with M vanishing moments.*

Then, the coefficients of Θ satisfy the following inequality for all $\lambda = (j, m, e) \in \Lambda$ and $\mu = (k, n, e') \in \Lambda$:

$$|\theta_{\lambda,\mu}| \leq C_M 2^{-(M+\frac{d}{2})|j-k|} 2^{-\min(j,k)(M+d)} f_{\lambda,\mu} \quad (3.9)$$

where $f_{\lambda,\mu} = f(\text{dist}(I_\lambda, I_\mu))$, C_M is a constant that does not depend on λ and μ and

$$\begin{aligned} \text{dist}(I_\lambda, I_\mu) &= \inf_{x \in I_\lambda, y \in I_\mu} \|x - y\|_\infty \\ &= \max \left(0, \left\| 2^{-j}m - 2^{-k}n \right\|_\infty - (2^{-j} + 2^{-k}) \frac{c(M)}{2} \right). \end{aligned} \quad (3.10)$$

Proof. See Appendix 3.8. □

Lemma 3.4.1 is the key to obtain all subsequent complexity estimates.

Theorem 3.4.1. *Let Θ_η be the matrix obtained by zeroing all coefficients in Θ such that*

$$2^{-\min(j,k)(M+d)} f_{\lambda,\mu} \leq \eta,$$

with $\lambda = (j, m, e) \in \Lambda$ and $\mu = (k, n, e') \in \Lambda$.

Let $\tilde{\mathbf{H}}_\eta = \Psi \Theta_\eta \Psi^$ denote the resulting operator. Suppose that f is compactly supported in $[0, \kappa]$ and that $\eta \leq \log_2(N)^{-(M+d)/d}$. Then:*

i) *The number of non zero coefficients in Θ_η is bounded above by*

$$C'_M N \kappa^d \eta^{-\frac{d}{M+d}} \quad (3.11)$$

where $C'_M > 0$ is independent of N .

ii) *The approximation $\tilde{\mathbf{H}}_\eta$ satisfies $\left\| \mathbf{H} - \tilde{\mathbf{H}}_\eta \right\|_{2 \rightarrow 2} \lesssim \eta^{\frac{M}{M+d}}$.*

iii) *The number of coefficients needed to satisfy $\left\| \mathbf{H} - \tilde{\mathbf{H}}_\eta \right\|_{2 \rightarrow 2} \leq \epsilon$ is bounded above by*

$$C''_M N \kappa^d \epsilon^{-\frac{d}{M}} \quad (3.12)$$

where $C''_M > 0$ is independent of N .

Proof. See Appendix 3.9. □

Let us summarize the main conclusions drawn from this section:

- A discretization in the wavelet domain provides better theoretical guarantees than the standard quadrature rules (see Section 3.4.1).
- The method is capable of handling *automatically* the degree of smoothness of the integral kernel K since there is a dependency in $\epsilon^{-\frac{d}{M}}$ where M is the smoothness of the integral operator.
- We will see in the next section that the method is quite versatile since different sparsity patterns can be chosen depending on the knowledge of the blur kernel and on the regularity of the signals that are to be processed.
- The method can also handle more general singular operators as was shown in the seminal papers [170, 171, 29].

Remark 3.4.1. Similar bounds as (3.9) can be derived with less stringent assumptions. First, the domain can be unbounded, given that kernels have a sufficiently fast decay at infinity. Second, the kernel can blow up on its diagonal, which is the key to study Calderon-Zygmund operators (see [170, 171, 29] for more details). We stucked to this simpler setting to simplify the proofs.

3.5 Identification of sparsity patterns

A key step to control the approximation quality is the selection of the coefficients in the matrix Θ that should be kept. For instance, a simple thresholding of Θ leads to sub-optimal and somewhat disappointing results. In this section we propose algorithms to select the most relevant coefficients for images belonging to functional spaces such as that of bounded variation functions. We study the case where Θ is known completely and the case where only an upper-bound such as (3.9) is available.

3.5.1 Problem formalization

Let $\mathbf{H}$ be the $N^d \times N^d$ matrix defined in equation (3.8). We wish to approximate $\mathbf{H}$ by a matrix $\tilde{\mathbf{H}}_K$ of kind $\Psi \mathbf{S}_K \Psi^*$ where $\mathbf{S}_K$ is a matrix with at most K non-zero coefficients. Let $\mathbb{S}_K$ denote the space of $N \times N$ matrices with at most K non-zero coefficients. The problem we address in this paragraph reads

$$\begin{aligned} & \min_{\mathbf{S}_K \in \mathbb{S}_K} \left\| \mathbf{H} - \tilde{\mathbf{H}}_K \right\|_{X \rightarrow 2} \\ &= \min_{\mathbf{S}_K \in \mathbb{S}_K} \max_{\|\mathbf{u}\|_X \leq 1} \left\| \mathbf{H}\mathbf{u} - \Psi \mathbf{S}_K \Psi^* \mathbf{u} \right\|_2. \end{aligned}$$

The solution of this problem provides the best K -sparse matrix $\mathbf{S}_K$, in the sense that no other choice provides a better SNR uniformly on the unit-ball $\{\mathbf{u} \in \mathbb{R}^N, \|\mathbf{u}\|_X \leq 1\}$.

Theoretical choice of the space X

The norm $\|\cdot\|_X$ should be chosen depending on the type of images that have to be blurred. For instance, it is well-known that natural images are highly compressible in the wavelet domain [164, 211]. This observation is the basis of JPEG2000 compression standard. Therefore, a natural choice could be to set $\|\mathbf{u}\|_X = \|\Psi^* \mathbf{u}\|_1$. This choice will ensure a good reconstruction of images that have a wavelet decomposition with a low ℓ^1 -norm.

Another very common assumption in image processing is that images have a bounded total variation. The space of functions with bounded total variation [10] contains images discontinuous along edges with finite length. It is one of the most successful tools for image processing tasks such as denoising, segmentation, reconstruction, ... Functions in $BV(\Omega)$ can be characterized by their wavelet coefficients [72, 164]. For instance, if $u \in BV(\Omega)$, then

$$\sum_{\lambda \in \Lambda_0} 2^{j(1-\frac{d}{2})} |\langle u, \psi_\lambda \rangle| < +\infty \quad (3.13)$$

for all wavelet bases. This result is due to embeddings of BV space in Besov spaces which are characterized by their wavelet coefficients (see [67] for more details on Besov spaces). This result motivated us to consider norms defined by

$$\|\mathbf{u}\|_X = \|\Sigma \Psi^* \mathbf{u}\|_1$$

where $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_N)$ is a diagonal matrix. Depending on the regularity level of the images considered, different diagonal coefficients can be used. For instance, for BV signals in 1D, one could set $\sigma_i = 2^{j(i)/2}$ where $j(i)$ is the scale of the i -th wavelet, owing to (3.13).

Practical choice of the space X

More generally, it is possible to adapt the weights σ_i depending on the images to recover. Most images exhibit a similar decay of wavelet coefficients across subbands. This decay is a characteristic of the functions regularity (see e.g. [130]). To illustrate this fact, we conducted a simple experiment in Figure 3.4. We evaluate the maximal value of the amplitude of wavelet coefficients of three images with different contents across scales. The wavelet transform is decomposed at level 4 and we normalize the images so that their maximum wavelet coefficient is 1. As can be seen even though the maximal values differ from one image to the next, their overall behavior is the same: amplitudes decay nearly dyadically from one scale to the next. The same phenomenon can be observed with the mean value.

This experiment suggests setting $\sigma_i = 2^{j(i)}$ in order to normalize the wavelet coefficients amplitude in each subband. Once again, the same idea was explored in [241].

An optimization problem

We can now take advantage of the fact that images and operators are sparse in the same wavelet basis. Let $\mathbf{z} = \Psi^* \mathbf{u}$ and $\Delta = \Theta - \mathbf{S}_K$. Since we consider orthogonal wavelet




Figure 3.4 – Three pictures and the mean amplitude of their wavelet coefficients at each scale of the wavelet transform.

transforms, we have $\|\Psi \mathbf{u}\|_2 = \|\mathbf{u}\|_2$, for any $\mathbf{u} \in \mathbb{R}^N$ therefore:

$$\begin{aligned}
 \|\mathbf{H} - \tilde{\mathbf{H}}_K\|_{X \rightarrow 2} &= \max_{\|\mathbf{u}\|_X \leq 1} \|\Psi(\Theta - \mathbf{S}_K)\Psi^* \mathbf{u}\|_2 \\
 &= \max_{\|\Sigma \mathbf{z}\|_1 \leq 1} \|(\Theta - \mathbf{S}_K)\mathbf{z}\|_2 \\
 &= \max_{\|\mathbf{z}\|_1 \leq 1} \|\Delta \Sigma^{-1} \mathbf{z}\|_2.
 \end{aligned}$$

Since the operator norm $\|\mathbf{A}\|_{1 \rightarrow 2} = \max_{1 \leq i \leq N} \|\mathbf{A}^{(i)}\|_2$, where $\mathbf{A}^{(i)}$ denote the i -th column of the $N \times N$ matrix $\mathbf{A}$ and by remarking that $(\Delta \Sigma^{-1})^{(i)} = \Delta^{(i)} \sigma_i^{-1}$, we finally get the following simple expression for the operator norm:

$$\|\mathbf{H} - \tilde{\mathbf{H}}\|_{X \rightarrow 2} = \max_{1 \leq i \leq N} \frac{1}{\sigma_i} \|\Delta^{(i)}\|_2. \quad (3.14)$$

Our goal is thus to find the solution of:

$$\min_{\mathbf{S}_K \in \mathbb{S}_K} \max_{1 \leq i \leq N} \frac{1}{\sigma_i} \|\Delta^{(i)}\|_2. \quad (3.15)$$

3.5.2 Link with the approach in [241]

In this paragraph, we show that the method proposed in [242, 241], can be interpreted with the formalism given above. In those papers, Θ is approximated by $\tilde{\Theta}$ using the following rule:

$$\tilde{\Theta}_{i,j} = \begin{cases} \Theta_{i,j} & \text{if } \frac{\Theta_{i,j}}{w_j} \text{ is in the } K \text{ largest values of } \Theta \mathbf{W}^{-1} \\ 0 & \text{otherwise.} \end{cases} \quad (3.16)$$

The weights w_i are set as constant by subbands and learned as described in paragraph 3.5.1.

The thresholding rule (3.16) can be interpreted as the solution of the following problem:

$$\min_{\tilde{\Theta} \in \mathbb{S}_K} \left\| \Theta - \tilde{\Theta} \right\|_{\mathbf{W} \rightarrow \infty},$$

where here $\|x\|_{\mathbf{W}} = \|\mathbf{W}x\|_1$ with $\mathbf{W} = \text{diag}(w_i)$ a diagonal matrix. Indeed, the above problem is equivalent to:

$$\min_{\tilde{\Theta} \in \mathbb{S}_K} \max_{1 \leq i, j \leq N} \left| \frac{1}{w_j} (\Theta - \tilde{\Theta})_{i,j} \right|.$$

In other words, the method proposed in [242, 241] constructs a K best-term approximation of Θ in the metric $\|\cdot\|_{\mathbf{W} \rightarrow \infty}$.

Overall, the problem is very similar to (3.15), except that the image quality is evaluated through an infinite norm in the wavelet domain, while we propose using a Euclidean norm in the spatial domain. We believe that this choice is more relevant for image processing since the SNR is the most common measure of image quality. In practice, we will see in the numerical experiments that both methods lead to very similar practical results.

Finally, let us mention that the authors in [241] have an additional concern of storing the matrix representation with the least memory. They therefore *quantize* the coefficients in Θ . Since the main goal in this paper is the design of fast algorithms for matrix-vector products, we do not consider this extra refinement.

3.5.3 An algorithm when Θ is known

Finding the minimizer of problem (3.15) can be achieved using a simple greedy algorithm: the matrix $\mathbf{S}_{k+1}$ is obtained by adding the largest coefficient of the column Δ_i with largest Euclidean norm to $\mathbf{S}_k$. This procedure can be implemented efficiently by using quick sort algorithms. The complete procedure is described in Algorithm 1. The overall complexity of this algorithm is $\mathcal{O}(N^2 \log(N))$. The most computationally intensive step is the sorting procedure in the initialisation. The loop on k can be accelerated by first sorting the set $(\gamma_j)_{1 \leq j \leq N}$, but the algorithm's complexity remains essentially unchanged.

3.5.4 An algorithm when Θ is unknown

In the previous paragraph, we assumed that the full matrix Θ was known. There are at least two reasons that make this assumption irrelevant. First, computing Θ is very computationally intensive and it is not even possible to store this matrix in RAM for medium sized images (e.g. 512×512). Second, in blind deblurring problems, the operator $\mathbf{H}$ needs to be inferred from the data and adding priors on the sparsity pattern of $\mathbf{S}_K$ might be an efficient choice to improve the problem identifiability.

When Θ is unknown, we may take advantage of equation (3.9) to define sparsity patterns. A naive approach would consist in applying Algorithm (1) directly on the upper-bound (3.9). However, this matrix cannot be stored and this approach is applicable only for small images. In order to reduce the computational burden, one may

Algorithm 1: An algorithm to find the minimizer of (3.15).

Input:

Θ : $N \times N$ matrix;

Σ : Diagonal matrix;

K : the number of elements in the thresholded matrix;

Output:

$\mathbf{S}_K$: Matrix minimizing (3.15)

Initialization:

Set $\mathbf{S}_K = \mathbf{0} \in \mathbb{R}^{N \times N}$;

Sort the coefficients of each column $\Theta^{(j)}$ of Θ in decreasing order;

Obtain $\mathbf{A}^{(j)}$ the sorted columns $\Theta^{(j)}$ and index sets I_j ;

The sorted columns $\mathbf{A}^{(j)}$ and index set I_j satisfy $\mathbf{A}^{(j)}(i) = \Theta^{(j)}(I_j(i))$;

Compute the norms $\gamma_j = \frac{\|\Theta^{(j)}\|_2^2}{\sigma_j^2}$;

Define $\mathbf{O} = (1, \dots, 1) \in \mathbb{R}^N$;

$\mathbf{O}(j)$ is the index of the largest coefficient in $\mathbf{A}^{(j)}$ not yet added to $\mathbf{S}_K$;

begin

for $k = 1$ **to** K **do**

 Find $l = \arg \max_{j=1 \dots N} \gamma_j$;

(Find the column l with largest Euclidean norm)

 Set $\mathbf{S}_K(I_l(\mathbf{O}(l)), l) = \Theta(I_l(\mathbf{O}(l)), l)$;

(Add the coefficient in the l -th column at index $I_l(\mathbf{O}(l))$)

 Update $\gamma_l = \gamma_l - \left(\frac{\mathbf{A}^{(l)}(\mathbf{O}(l))}{\sigma_l} \right)^2$;

(Update norms vector)

 Set $\mathbf{O}(l) = \mathbf{O}(l) + 1$;

(The next value to add in l -th column will be at index $\mathbf{O}(l) + 1$)

end

end

take advantage of the special structure of the upper-bound: equation (3.9) indicates that the coefficients $\theta_{\lambda,\mu}$ can be discarded for sufficiently large $|j - k|$ and sufficiently large distance between the wavelet supports. Equation (3.9) thus means that for a given wavelet ψ_λ , only its spatial neighbours in neighbouring scales have significant correlation coefficients $\langle H\psi_\lambda, \psi_\mu \rangle$. We may thus construct sparsity patterns using the notion of multiscale neighbourhoods defined below.

Definition 3.5.1 (Multiscale shift). The multiscale shift $s_{\lambda,\mu} \in \mathbb{Z}^d$ between two wavelets ψ_λ and ψ_μ is defined by

$$s_{\lambda,\mu} = \left\lfloor \frac{n}{2^{\max(k-j,0)}} \right\rfloor - \left\lfloor \frac{m}{2^{\max(j-k,0)}} \right\rfloor. \quad (3.17)$$

We recall that $\lambda = (j, m, e) \in \Lambda$ and $\mu = (k, n, e') \in \Lambda$. Note that for $k = j$, the multiscale shift is just $s_{\lambda,\mu} = n - m$ and corresponds to the standard shift between wavelets, measured as a multiple of the characteristic size 2^{-j} . The divisions by $2^{\max(k-j,0)}$ and $2^{\max(j-k,0)}$ allow to rescale the shifts at the coarsest level. This definition is illustrated in Figure 3.5.

Definition 3.5.2 (Multiscale neighborhood). Let

$$\mathcal{N} = \left\{ (j, (k, s)), (j, k) \in \{0, \dots, \log_2(N) - 1\}^2, s \in \{0, \dots, 2^{\min(j,k)} - 1\}^d \right\}$$

denote the set of all neighborhood relationships, i.e. the set of all possible couples of type (scale, (scale, shift)). A multiscale neighborhood $\mathcal{N}$ is an element of the powerset $\mathcal{P}(\mathcal{N})$.

Definition 3.5.3 (Multiscale neighbors). Given a multiscale neighborhood $\mathcal{N}$, two wavelets ψ_λ and ψ_μ will be said to be $\mathcal{N}$ -neighbors if $(j, (k, s_{\lambda,\mu})) \in \mathcal{N}$ where $s_{\lambda,\mu}$ is defined in equation (3.17).

The problem of finding a sparsity pattern is now reduced to finding a good multiscale neighborhood. In what follows, we let $\mathcal{N}(j) = \{(k, s), (j, (k, s)) \in \mathcal{N}\}$ denote the set of all possible neighborhood relationships at scale j . This is illustrated in Figure 3.6. Let $\mathcal{N} \in \mathcal{P}(\mathcal{N})$ denote a multiscale neighborhood. We define the matrix $\mathbf{S}_{\mathcal{N}}$ as follows:

$$\mathbf{S}_{\mathcal{N}}(\lambda, \mu) = \begin{cases} \theta_{\lambda,\mu} & \text{if } \psi_\lambda \text{ is an } \mathcal{N}\text{-neighbor of } \psi_\mu \\ 0 & \text{otherwise.} \end{cases}$$

Equation (3.9) indicates that

$$|\theta_{\lambda,\mu}| \leq u(j, k, s)$$

with

$$u(j, k, s) = C_M 2^{-(M + \frac{d}{2})|j-k| - (M+d)\min(j,k)} f_{j,k,s} \quad (3.18)$$




Figure 3.5 – Illustration of a multiscale shift on a 1D signal of size 8 with the Haar basis. The shifts are computed with respect to wavelet $\psi_{1,1}$. Wavelets $\psi_{0,0}$, $\psi_{2,2}$ and $\psi_{2,3}$ have a multiscale shift $s = 0$ with $\psi_{1,1}$ since their support intersects that of $\psi_{1,1}$. Wavelets $\psi_{1,0}$, $\psi_{2,0}$ and $\psi_{2,1}$ have a multiscale shift $s = -1$ with $\psi_{1,1}$ since their support intersects that of $\psi_{1,0}$.




Figure 3.6 – Illustration of a multiscale neighborhood on a 1D signal. In this example, the neighborhood at scale 1 is $\mathcal{N}(1) = \{(-1, 0), (0, -1), (0, 0), (0, 1), (1, -1), (1, 0), (1, 1), (2, 0)\}$. Notice that the two red wavelets at scale 2 are neighbors of the orange wavelet at scale 1 and that this relationship is described through only one shift.

and $f_{j,k,s} = f\left(\max\left(0, 2^{-\min(j,k)} \|s\|_\infty - (2^{-j} + 2^{-k}) \frac{c(M)}{2}\right)\right)$. Let $\mathbf{U}$ be the matrix defined by $\mathbf{U}(\lambda, \mu) = u(j, k, s_{\lambda, \mu})$. Finding a good sparsity pattern can now be achieved by solving the following problem:

$$\min_{\substack{\mathcal{N} \in \mathcal{P}(\mathcal{N}) \\ |\mathcal{N}|=K}} \max_{1 \leq i \leq N} \frac{1}{\sigma_i} \left\| (\mathbf{U} - \mathbf{S}_{\mathcal{N}})^{(i)} \right\|_2 \quad (3.19)$$

where $(\mathbf{U} - \mathbf{S}_{\mathcal{N}})^{(i)}$ denotes the i -th column of $(\mathbf{U} - \mathbf{S}_{\mathcal{N}})$.

In what follows, we assume that σ_i only depends on the scale $j(i)$ of the i -th wavelet. Similarly to the previous section, finding the optimal sparsity pattern can be performed using a greedy algorithm. A multiscale neighborhood is constructed by iteratively adding the couple (scale, (scale, shift)) that minimizes a residual. This technique is described in Algorithm 2.

Note that the norms γ_k only depend on the scale $j(k)$, so that the initialisation step only requires $\mathcal{O}(N \log_2(N))$ operations. Similarly to Algorithm 1, this algorithm can be accelerated by first sorting the elements of $u(j, k, s)$ in decreasing order. The overall complexity for this algorithm is $\mathcal{O}(N \log(N)^2)$ operations.

3.6 Numerical experiments

In this section we perform various numerical experiments in order to illustrate the theory proposed in the previous sections and to compare the practical efficiency of wavelet based methods against windowed convolutions (WC) based approaches. We first describe the operators and images used in our experiments. Second, we provide numerical

Algorithm 2: An algorithm to find the minimizer of (3.19).

Input:

u : Upper-bound defined in (3.18);

Σ : Diagonal matrix;

K : the number of elements of the neighborhood;

Output:

$\mathcal{N}$: multiscale neighborhood minimizing (3.19)

Initialization:

Set $\mathcal{N} = \emptyset$;

Compute the norms $\gamma_k = \frac{\|\mathbf{U}^{(k)}\|_2^2}{\sigma_k^2}$ using the upper-bound u ;

begin

for $k = 1$ **to** K **do**

 Find $j^* = \arg \max_{j=1 \dots N} \gamma_j$;

(The column with largest norm)

 Find $(k^*, s^*) = \arg \max_{(k,s) \in \mathcal{N}(j^*)} u^2(j^*, k, s) 2^{\max(j^*-k, 0)}$;

(The best scale and shift for this column is (k^, s^*))*

(The number of elements in the neighborhood relationship $(j^, (k, s))$ is $2^{\max(j^*-k, 0)}$)*

 Update $\mathcal{N} = \mathcal{N} \cup \{(j^*, (k^*, s^*))\}$;

 Set $\gamma_k = \gamma_k - u^2(j^*, k^*, s^*) \cdot 2^{\max(j^*-k, 0)}$

end

end

experiments for the direct problems. Finally, we provide numerical comparisons for the inverse problem.

3.6.1 Preliminaries

Test images

We consider a set of 16 images of different natures: standard image processing images (the boat, the house, Lena, Mandrill (see Figure 3.7a), peppers, cameraman), two satellite images, three medical images, three buildings images, and two test pattern images (see Figure 3.7b). Due to memory limitations, we only consider images of size $N = 256 \times 256$. Note that a full matrix of size $N \times N$ stored in double precision weights around 32 gigabytes.

Test operators

Three different blur kernels of different complexities are considered, see Figure 3.8. The PSFs in Figure 3.8a and 3.8b modeled for all $x \in [0, 1]^2$ by 2D Gaussians. Therefore the associated kernel is defined for all $(x, y) \in [0, 1]^2 \times [0, 1]^2$ by

$$K(x, y) = \frac{1}{2\pi |C(y)|} \exp \left[\frac{1}{2} (y - x)^T C^{-1}(y) (y - x) \right].$$

The covariance matrices C are defined as:

- In Figure 3.8a: $C(y) = \begin{pmatrix} f(y_1) & 0 \\ 0 & f(y_1) \end{pmatrix}$ with $f(t) = 2t$, for $t \in [0, 1]$. The PSFs are truncated out of a 11×11 support.
- In Figure 3.8b: $C(y) = R(y)^T D(y) R(y)$ where $R(y)$ is a rotation matrix of angle $\theta = \arctan \left(\frac{y_1 - 0.5}{y_2 - 0.5} \right)$ and $D(y) = \begin{pmatrix} g(y) & 0 \\ 0 & h(y) \end{pmatrix}$ with $g(y) = 10 \|y - (0.5, 0.5)^T\|_2$ and $h(y) = 2 \|y - (0.5, 0.5)^T\|_2$. The PSFs are truncated out of a 21×21 support.

The PSFs in Figure 3.8c were proposed in [212] as an approximation of real spatially optical blurs.

Computation of the full Θ matrix

Before applying our approximation methods, matrix Θ needs to be computed explicitly. The coefficients $\langle H\psi_\lambda, \psi_\mu \rangle$ are approximated by their discrete counterparts. If ψ_λ and ψ_μ denote discrete wavelets, we simply compute the wavelet transform of $\mathbf{H}\psi_\lambda$ and store it into the λ -th column of Θ . This computation scheme is summarized in Algorithm 3. This algorithm corresponds to the use of rectangle methods to evaluate the dot-products:

$$\int_{\Omega} \int_{\Omega} K(x, y) \psi_\lambda(y) \psi_\mu(x) dy dx \simeq \frac{1}{N^{2d}} \sum_{x \in X} \sum_{y \in X} K(x, y) \psi_\lambda(y) \psi_\mu(x). \quad (3.20)$$




Figure 3.7 – The two images of size 256×256 used in these numerical experiments




Figure 3.8 – PSFs maps used in the paper. The PSFs in Figure 3.8a are Gaussians with equal variances increasing in the vertical direction. The PSFs in Figure 3.8b are anisotropic Gaussians with covariance matrices that depend on the polar coordinates. The PSFs in Figure 3.8c are based on paper [212].

Algorithm 3: An algorithm to compute Θ

Output:

Θ : the full matrix of $\mathbf{H}$

begin

forall the λ do

 Compute the wavelet ψ_λ using an inverse wavelet transform

 Compute the blurred wavelet $\mathbf{H}\psi_\lambda$

 Compute $\left(\langle \mathbf{H}\psi_\lambda, \psi_\mu \rangle\right)_\mu$ using one forward wavelet transform

 Set $\left(\langle \mathbf{H}\psi_\lambda, \psi_\mu \rangle\right)_\mu$ in the λ -th column of Θ .

end

end

3.6.2 Application to direct problems

In this section, we investigate the approximation properties of the proposed approaches in the aim of computing matrix-vector products. In all numerical experiments, we use an orthogonal wavelet transform with 4 decomposition levels. We always use Daubechies wavelets.

Influence of vanishing moments

First of all we demonstrate the influence of vanishing moments on the quality of approximations. For each number of vanishing moments $M \in \{1, 2, 4, 6, 10\}$, a sparse approximation $\tilde{\mathbf{H}}$ is constructed by thresholding Θ , keeping the $K = l \times N$ largest coefficients with $l \in \{0 \dots 40\}$. Then for each $\mathbf{u}$ in the set of 16 images, we compare $\tilde{\mathbf{H}}\mathbf{u}$ to $\mathbf{H}\mathbf{u}$ computing the pSNR. We then plot the average of pSNRs over the set of images with respect to the number of operations needed for a matrix-vector product. The results of this experiment are displayed in Figure 3.9. It appears that for the considered operators, using as many vanishing moments as possible was preferable. Using more than 10 vanishing moments however led to insignificant performance increase while making the numerical complexity higher. Therefore, in all the following numerical experiments we will use Daubechies wavelets with 10 vanishing moments. Note that paper [241] only explored the use of Haar wavelets. This experiment shows that very significant improvements can be obtained by leveraging regularity of the integral kernel using vanishing moments. The behavior was predicted by Theorem 3.4.1.

Comparison of different methods

Wavelets VS windowed convolutions. In this first numerical experiment, we evaluate $\|\mathbf{H} - \tilde{\mathbf{H}}\|_{2 \rightarrow 2}$ where $\tilde{\mathbf{H}}$ is obtained by windowed convolutions method or sparse approximations in the wavelet domain.




Figure 3.9 – pSNR of the blurred image using the approximated operator $\tilde{\mathbf{H}}\mathbf{u}$ with respect to the blurred image using the exact operator $\mathbf{H}\mathbf{u}$. pSNRs have been averaged over the set of test images. Daubechies wavelets have been used with different number vanishing moments $M \in \{1, 2, 4, 6, 10\}$. The case $M = 1$ corresponds to Haar wavelets.

The sparse approximation of the operator is constructed by thresholding the matrix Θ in order to keep the K largest coefficients. We have set $K = 2^l \times N$ with $l \in \{0 \dots 2 \log_2 N\}$. This way K is a multiple of the number of pixels in the image. The windowed convolutions method is constructed by partitioning the image into $2^l \times 2^l$ sub-images where $l \in \{1 \dots \log_2 N\}$. We also studied the case where sub-images overlap and linearly interpolated the blur between sub-images as proposed in [179, 132]. The overlap has been fixed to 50% of the sub-images sizes.

For each sub-image size, and each overlap, the norm $\|\mathbf{H} - \tilde{\mathbf{H}}\|_{2 \rightarrow 2}$ is approximated using a power method [116]. We stop the iterative process when the difference between the eigenvalues of two successive iterations is smaller than $10^{-8} \|\mathbf{H}\|_{2 \rightarrow 2}$. The number of operations associated to each type of approximation is computed using theoretical complexities. For sparse matrix-vector product the number of operations is proportional to the number of non-zero coefficients in the matrix. For windowed convolutions methods, the number of operations is proportional to the number of windows ($2^l \times 2^l$) multiplied by the cost of a discrete convolution over a window $\left(\frac{N}{2^l} + N\kappa\right)^2 \log_2 \left(\frac{N}{2^l} + N\kappa\right)$.

Figure 3.10 shows the results of this experiment. The wavelet based method seems to perform much better than windowed convolutions methods for both operators. The gap is however significantly larger for the rotation blur in Figure 3.8b. This experiment therefore suggests that the advantage of wavelet based approaches will depend on the type of blur considered.




Figure 3.10 – The operator norms $\|\mathbf{H} - \tilde{\mathbf{H}}\|_{2 \rightarrow 2}$ are displayed for the three proposed kernels. (Left: kernel Figure 3.8a, middle: kernel in Figure 3.8b, right: kernel in Figure 3.8c). Norms are plotted with respect to the number of operations needed to compute $\tilde{\mathbf{H}}\mathbf{u}$. The abscissas are in log scale.

The influence of sparsity patterns. In this numerical experiment, we obtain a K -sparse matrix Θ_K using either a simple thresholding strategy, Algorithm 1 or Algorithm 2. We evaluate the error $\|\mathbf{H} - \tilde{\mathbf{H}}\|_{X \rightarrow 2}$ defined in (3.14) for each methods. We set $\sigma_i = 2^{j(i)}$, where $j(i)$ corresponds to the scale of the i -th wavelet. As can be seen from Figure 3.11, Algorithm 1 provides a much better error decay for each operator than the simple thresholding strategy. This fact will be verified for real images in next paragraph. Algorithm 2 has a much slower decay than both thresholding algorithm. Notice that this algorithm is essentially blind, in the sense that it does not require knowing the exact matrix Θ to select the pattern. It would therefore work for a whole class of blur kernels, whereas the simple thresholding and Algorithm 1 work only for a specific matrix.




Figure 3.11 – The operator norms $\|\mathbf{H} - \tilde{\mathbf{H}}\|_{X \rightarrow 2}$ are displayed for kernels Figure 3.8a (left) and Figure 3.8b (right); and with respect to the number of operations needed to compute $\tilde{\mathbf{H}}u$. The abscissas are in log scale. Daubechies wavelets with 10 vanishing moments have been used.

Figure 3.12 shows the sparsity patterns of matrices obtained with Algorithms 1 and 2 for $K = 30N$ and $K = 128N$ coefficients. The sparsity patterns look quite similar. However, Algorithm 1 selects subbands that are not selected by Algorithm 2, which might

explain the significant performance differences. Similarly, Algorithm 2 select subbands that would probably be crucial for some blur kernels, but which are not significant for this particular blur kernel.

Quality of matrix-vector products for real images

In this section, we evaluate the performance of wavelet based methods for matrix-vector products with real images.

Quality VS complexity. We compare $\tilde{\mathbf{H}}\mathbf{u}$ to $\mathbf{H}\mathbf{u}$, where $\mathbf{u}$ is the image in Figure 3.7b and where $\tilde{\mathbf{H}}$ is obtained either by windowed convolutions methods or by sparse wavelet approximations. We plot the pSNR between the exact blurred image $\mathbf{H}\mathbf{u}$ and the blurred image using the approximated operator $\tilde{\mathbf{H}}\mathbf{u}$ in Figure 3.13. Different approximation methods are tested:

Thresholded matrix: This corresponds to a simple thresholding of the wavelet matrix Θ .

Σ n°1: This corresponds to applying Algorithm 1 with $\sigma_i = 1$, $\forall i$ where $j(i)$ corresponds to the scale of the i -th wavelet.

Σ n°2: This corresponds to applying Algorithm 1 with $\sigma_i = 2^{j(i)} \forall i$.


[241]: The method presented in [241] with $K = l \times N$ coefficients in the matrix, with $l \in \{1, \dots, 100\}$.

WC, Overlap 50%: This corresponds to the windowed convolution with 50% overlap. We use this overlap since it produces better pSNRs.


Algo 2: The algorithm finds multi-scale neighbourhoods until $K = l \times N$ coefficients populate the matrix, with $l \in \{1, \dots, 100\}$. In this experiment, we set $M = 1$, $f(t) = \frac{1}{1+t}$ and $\sigma_i = 2^{j(i)}$, $\forall i$.

The pSNRs are averaged over the set of 16 images. The results of this experiment are displayed in Figure 3.13 for the two kernels from Figures 3.8b and 3.8a. Let us summarize the conclusions from this experiment:


- A clear fact is that windowed convolution methods are significantly outperformed by wavelet based methods for all blur kernels. Moreover, the differences between wavelet and windowed convolution based methods get larger as the blurs regularity decreases.
- A second result is that wavelet based methods with fixed sparsity patterns (Algo 2) are quite satisfactory for very sparse patterns (i.e. less than $20N$ operations) and kernels 3.8a and 3.8b. We believe that the most important regime for applications is in the range $[N, 20N]$, so that this result is rather positive. However, Algo 2 suffers from two important drawbacks: first, the increase in SNR after a certain value




(a) Algorithm 1 – $K = 30N$



(b) Algorithm 2 – $K = 30N$



(c) Algorithm 1 – $K = 128N$



(d) Algorithm 2 – $K = 128N$

Figure 3.12 – The structure of the wavelet matrices of Θ_K are displayed for Algorithms 1 and 2 and for $K = 30N$ and $K = 128N$ coefficients. Algorithm 1 has been applied using the second $\Sigma = \text{diag}(2^{j(i)})_i$ matrix. This experiment corresponds to the blur in Figure 3.8b

becomes very slow. Second, this algorithm provides very disappointing results for the last blur map 3.8c. These results suggest that this method should be used with caution if one aims at obtaining very good approximations. In particular, the algorithm is dependent on the bound (3.9) which itself depends on user given parameters such as function f in (2a). Modifying those parameters might result in better results, but is usually hard to tweak manually.

- The methods Σ n° 1, Σ n° 2, *Thresholded matrix* all behave similarly. Method Σ n° 2 is however significantly better, showing the importance of choosing the weights σ_i in equation (3.15) carefully.
- The methods Σ n° 1, Σ n° 2, *Thresholded matrix* outperform the method proposed in [241] for very sparse patterns ($< 20N$) and get outperformed for mid-range sparsification $> 40N$. The main difference between algorithm [241] and the methods proposed in this paper is the number of vanishing moments. In [241], the authors propose using the Haar wavelet (i.e. 1 vanishing moment), while we use Daubechies wavelets with 10 vanishing moments. In practice, this results in better approximation properties in the very sparse regime, which might be the most important in applications. For mid-range sparsification, the Haar wavelet provides better results. Two reasons might explain this phenomenon: first, Haar wavelets have a small spatial support, therefore matrix Θ contains less non-zero coefficients when expressed with Haar wavelets than Daubechies wavelets. Second, the constants C'_M and C''_M in Theorem (3.4.1) are increasing functions of the number of vanishing moments.

Illustration of artefacts. Figure 3.14 provides a comparison of the windowed convolutions methods and the wavelet based approach in terms of approximation quality and computing times. The following conclusions can be drawn from this experiment:

- The residual artefacts appearing in the windowed convolutions approach and wavelet based approach are different. They are localized at the interfaces between sub-images for the windowed convolutions approach while they span the whole image domain for the wavelet based approach. It is likely that using translation and/or rotation invariant wavelet would improve the result substantially.
- The approximation using the second Σ matrix produces the best results and should be preferred over more simple approaches.
- In our implementation, the windowed convolutions approach (implemented in C) is outperformed by the wavelet based method (implemented in Matlab with C-mex files). For instance, for a precision of 45dBs, the wavelet based approach is about 10 times faster.
- The computing time of 1.21 seconds for the windowed convolutions approach with a 2×2 partition might look awkward since the computing times are significantly

lower for finer partitions. This is because the efficiency of FFT methods depend greatly on the image size. The time needed to compute an FFT is usually lower for sizes that have a prime factorization comprising only small primes (e.g. less than 7). This phenomenon explains the fact that the practical complexity of windowed convolutions algorithms may increase in a chaotic manner with respect to m .




Figure 3.13 – pSNR of the blurred image using the approximated operators $\tilde{\mathbf{H}}\mathbf{u}$ with respect to the blurred image using the exact operator $\mathbf{H}\mathbf{u}$. The results have been obtained with blur Figure 3.8a for top-left graph, blur Figure 3.8b for top-right graph and blur Figure 3.8c for the bottom. pSNR are averaged over the set of 16 images.



	Piece. Conv.	Difference	Algorithm 1	Difference	$l =$
2×2	31.90 dB		36.66 dB		5
1.21 sec					0.039 sec
4×4	38.49 dB		45.87 dB		30
0.17 sec					0.040 sec
8×8	44.51 dB		50.26 dB		50
0.36 sec					0.048 sec
16×16	53.75 dB		57.79 dB		100
0.39 sec					0.058 sec

Figure 3.14 – Blurred images and the differences $\mathbf{Hu} - \tilde{\mathbf{H}}\mathbf{u}$ for the kernel Figure 3.8b. Results on the left are obtained using windowed convolutions approximations with 2×2 , 4×4 , 8×8 and 16×16 partitionings all with 50% overlap. Results on the right are obtained using Algorithm 1 with the second $\Sigma = \text{diag}(2^{j(i)})_i$ matrix keeping $K = lN$ coefficients. The pSNR and the time needed for the computation for the matrix-vector product are shown.

3.6.3 Application to inverse problems

In this experiment we compare the methods efficiency in deblurring problems. We assume the following classical image degradation model

$$\mathbf{v} = \mathbf{H}\mathbf{u} + \boldsymbol{\eta}, \quad \boldsymbol{\eta} \sim \mathcal{N}(0, \sigma^2 \text{Id}), \quad (3.21)$$

where $\mathbf{v}$ is the degraded image observed, $\mathbf{u}$ is the image to restore, $\mathbf{H}$ is the blurring operator and σ^2 is the noise variance. A standard TV-L2 optimization problem is solved to restore the image $\mathbf{u}$:

$$\text{Find } \mathbf{u}^* \in \arg \min_{\mathbf{u} \in \mathbb{R}^N, \|\tilde{\mathbf{H}}\mathbf{u} - \mathbf{v}\|_2^2 \leq \alpha} TV(\mathbf{u}), \quad (3.22)$$

where $\tilde{\mathbf{H}}$ is an approximating operator and TV is the isotropic total variation of $\mathbf{u}$. The optimization problem is solved using the primal-dual algorithm proposed in [54]. We do not detail the resolution method since it is now well documented in the literature.

An important remark is that the interest of the total variation term is not only to regularize the ill-posed inverse problem, but also to handle the errors in the operator approximation. In practice we found that setting $\alpha = (1 + \epsilon)\sigma^2 N$ where $\epsilon > 0$ is a small parameter provides good experimental results.

In Figures 3.15 and 3.16, we present deblurring results using Figure 3.7b with kernel 3.8b.

In both the noisy and noiseless cases, the 4×4 windowed convolutions method performs worst reconstructions than wavelet approaches with $30N$. Moreover, they are between 4 and 6 times significantly slower. Surprisingly even the implementation in the space domain is faster. The reason for that is probably a difference in the quality of implementation: we use Matlab sparse matrix-vector products for space and wavelet methods. This routine is cautiously optimized while our c implementation of windowed convolutions can probably be improved. In addition, let us mention that two wavelet transforms need to be computed at each iteration with the wavelet based methods, while this is not necessary with the space implementation. It is likely that the acceleration factor would have been significantly higher if wavelet based regularizations had been used.

In the noiseless case, the simple thresholding approach provides significantly better SNRs than the more advanced proposed in this paper and in [241]. Note however that it produces more significant visual artefacts. This result might come as a surprise at first sight. However, as was explained in section 3.5, our aim to design sparsity patterns was to minimize an operator norm $\|\mathbf{H} - \tilde{\mathbf{H}}\|_{X \rightarrow 2}$. When dealing with an inverse problem, approximating the direct operator is not as relevant as approximating its inverse. This calls for new methods specific to inverse problems.

In the noisy case, all three thresholding strategies produce results of a similar quality. The Haar wavelet transform is however about twice faster since the Haar wavelet support is smaller. Moreover, the results obtained with the approximated matrices are nearly as good as the ones with the true operator. It suggests that it is not necessary to construct

accurate approximations of the operators in practical problems. This observation is also supported by the experiment in Figure 3.17. In this experiment, we plot the pSNR of the deblurred image in presence of noise with respect to the number of elements in Θ_K . Interestingly, a matrix containing only $20N$ coefficients leads to deblurred images close to the results obtained with the exact operator. In this experiment, a total of $K = 5N$ coefficients in Θ_K is enough to retrieve satisfactory results. This is a very encouraging result for blind deblurring problems.

3.7 Conclusion

3.7.1 Brief summary

In this paper, we introduced an original method to represent spatially varying blur operators in the wavelet domain. We showed that this new technique has a great adaptivity to the smoothness of the operator and exhibit an $\mathcal{O}(N\epsilon^{-d/M})$ complexity, where M denotes the kernel regularity. This method is versatile since it is possible to adapt it to the kind of images that have to be treated. We showed that much better performance to approximate the direct operator can be obtained by leveraging the fact that natural signals exhibit some structure in the wavelet domain. Moreover, we proposed an original method to design sparsity patterns for class of blurring operators when only the operator regularity is known. These theoretical results were confirmed by practical experiments on real images. Even though our conclusions are still preliminary since we tested only small 256×256 images, the wavelet based methods seem to significantly outperform standard windowed convolutions based approaches. Moreover, they seem to provide satisfactory deblurring results on practical problems with a complexity no greater than $5N$ operations, where N denotes the pixels number.

3.7.2 Outlook

We provided a simple complexity analysis based solely on the *global* regularity of the kernel function. It is well known that wavelets are able to adapt locally to the structures of images or operators [70]. The method should thus provide an efficient tool for piecewise regular blurs appearing in computer vision for instance. It could be interesting to evaluate precisely the complexity of wavelet based approximations for piecewise regular blurs.

A key problem of the wavelet based approach is the need to project the operator on a wavelet basis. In this paper we performed this operation using the computationally intensive Algorithm 3. It could be interesting to derive fast projection methods. Let us note that such methods already exist in the literature [29]. A similar procedure was used in the specific context of spatially varying blur in [241].

Moreover, the proposed method can already be applied to situations where the blur mostly depends on the instrument: the wavelet representation has to be computed once for all off-line, and then all deblurring operations can be handled much faster. This



(a) Degraded image
21.85dB



(b) Exact operator
34.53dB – 64.87 sec



(c) Simple thresh
31.68dB – 21.68 sec



(d) Algorithm 1
30.57dB – 21.16 sec




(e) WC 4×4
28.37dB – 85.60 sec



(f) [241]
30.53dB – 14.12 sec

Figure 3.15 – Deblurring results for kernel Figure 3.8b and without noise. Top-left: degraded image. Top-right: deblurred using the exact operator. Middle-left: deblurred by the wavelet based method and a simple thresholding. Middle-right: deblurred by the wavelet based method and Algorithm 2 with the second $\Sigma = \text{diag}(2^{j(i)})_i$ matrix. Bottom: deblurred using a 4×4 windowed convolutions algorithm with 50% overlap. For wavelet methods $K = 30N$ coefficients are kept in matrices. pSNR are displayed for each restoration.



(a) Degraded image
21.62dB



(b) Exact operator
29.09dB – 64.87 sec



(c) Simple thresh
28.64dB – 21.68 sec



(d) Algorithm 1
28.24dB – 21.16 sec



(e) WC 4×4
27.62dB – 85.60 sec



(f) [241]
28.37dB – 14.12 sec

Figure 3.16 – Deblurring results for kernel Figure 3.8b and with $\sigma = 0.02$ noise. Top-left: degraded image. Top-right: deblurred using the exact operator. Middle-left: deblurred by the wavelet based method and a simple thresholding. Middle-right: deblurred by the wavelet based method and Algorithm 2 with the second $\Sigma = \text{diag}(2^{j(i)})_i$ matrix. Bottom: deblurred using a 4×4 windowed convolutions algorithm with 50% overlap. For wavelet methods $K = 30N$ coefficients are kept in matrices. pSNR are displayed for each restoration.




Figure 3.17 – pSNR of the deblurred image with respect to the number of coefficients in the matrix divided by N for the image Figure 3.7a and the kernel Figure 3.8a. The matrix is constructed using Algorithm 1 with the second $\Sigma = \text{diag}(2^{j(i)})_i$ matrix with $K = lN$ coefficients for l from 1 to 30. Deblurred images using these matrices are compared with the one obtained with the exact operator.

situation occurs in satellite imaging or for some fluorescence microscopes (see e.g. [126, 226, 158]).

The design of good sparsity patterns is an open and promising research avenue. In particular, designing patterns adapted to specific inverse problems could have some impact as was illustrated in section 3.6.3.

Another exciting research perspective is the problem of blind deconvolution. Expressing the unknown operator as a sparse matrix in the wavelet domain is a good way to improve the problem identifiability. This is however far from being sufficient since the blind deconvolution problem has far more unknowns (a full operator and an image) than data (a single image). Further assumptions should thus be made on the wavelet coefficients regularity, and we plan to study this problem in a forthcoming work.

Finally let us mention that we observed some artefacts when using the wavelet based methods with high sparsity levels. This is probably due to their non translation and rotation invariance. It could be interesting to study sparse approximations in redundant wavelet bases or other time-frequency bases. It was shown for instance in [43] that curvelets are near optimal to represent Fourier integral operators. Similarly, Gabor frames are known to be very efficient to describe smoothly varying integral operators in the 1D setting [136].

3.8 Proof of Lemma 3.4.1

We let Π_M denote the set of polynomials of degree less or equal to M .

Lemma 3.8.1 below is a common result in numerical analysis [87] (see also Theorem

3.2.1 in [67]). It ensures that the approximation error of a function by a polynomial of degree M is bounded by the Sobolev semi-norm $W^{M,p}$.

Lemma 3.8.1 (Polynomial approximation). *For $1 \leq p \leq +\infty$, $M \in \mathbb{N}^*$ and $\Omega \subset \mathbb{R}^d$ a bounded domain, the following bound holds*

$$\inf_{g \in \Pi_M} \|f - g\|_{L^p(\Omega)} \leq C |f|_{W^{M+1,p}(\Omega)}, \quad (3.23)$$

where C is a constant that depends on d, M, p and Ω only.

Moreover, if $I_h \subset \Omega \subset \mathbb{R}^d$ is a cube of sidelength h , the following estimate holds

$$\inf_{g \in \Pi_M} \|f - g\|_{L^p(I_h)} \leq Ch^{M+1} |f|_{W^{M+1,p}(I_h)}, \quad (3.24)$$

where C is a constant only depending on d, M, p and Ω .

Let $I_\lambda = \text{supp}(\psi_\lambda)$. From the wavelets definition, we get

$$I_\lambda = 2^{-j}(m + [-c(M)/2, c(M)/2]^d)$$

therefore $|I_\lambda| = c(M)^d \cdot 2^{-jd}$. We will now prove Lemma 3.4.1.

Proof of Lemma 3.4.1. Since the mapping $(x, y) \mapsto K(x, y)\psi_\lambda(y)\psi_\mu(x)$ is bounded, it is also absolutely integrable on compact domains. Therefore $\langle H\psi_\lambda, \psi_\mu \rangle$ is well-defined for all (λ, μ) . Recall that $\lambda = (j, m, e) \in \Lambda$ and $\mu = (k, n, e') \in \Lambda$. Moreover Fubini's theorem can be applied and we get

$$\begin{aligned} \langle H\psi_\lambda, \psi_\mu \rangle &= \int_{I_\mu} \int_{I_\lambda} K(x, y)\psi_\lambda(y)\psi_\mu(x) dy dx \\ &= \int_{I_\lambda} \int_{I_\mu} K(x, y)\psi_\lambda(y)\psi_\mu(x) dx dy. \end{aligned}$$

To prove the result, we distinguish the cases $j \leq k$ and $j > k$. In this proof, we focus on the case $j \leq k$. The other one can be obtained by symmetry, using the facts that $\langle H\psi_\lambda, \psi_\mu \rangle = \langle \psi_\lambda, H^*\psi_\mu \rangle$ and that H and H^* are both blurring operators in the same class.

To exploit the regularity of K and ψ , note that for all $g \in \Pi_{M-1}$, $\int_{I_\mu} g(x)\psi_\mu(x) dx = 0$ since ψ has M vanishing moments. Therefore,

$$\langle H\psi_\lambda, \psi_\mu \rangle = \int_{I_\lambda} \int_{I_\mu} (K(x, y) - g(x)) \psi_\lambda(y) \psi_\mu(x) dx dy, \quad \forall g \in \Pi_{M-1}$$

and

$$|\langle H\psi_\lambda, \psi_\mu \rangle| \leq \int_{I_\lambda} \int_{I_\mu} |K(x, y) - g(x)| |\psi_\lambda(y)| |\psi_\mu(x)| dx dy,$$

for all $g \in \Pi_{M-1}$. Therefore it implies that,

$$\begin{aligned} |\langle H\psi_\lambda, \psi_\mu \rangle| &\leq \int_{I_\lambda} \inf_{g \in \Pi_{M-1}} \int_{I_\mu} |K(x, y) - g(x)| |\psi_\lambda(y)| |\psi_\mu(x)| dx dy, \\ &\leq \int_{I_\lambda} \inf_{g \in \Pi_{M-1}} \|K(\cdot, y) - g\|_{L^\infty(I_\mu)} \|\psi_\mu\|_{L^1(I_\mu)} |\psi_\lambda(y)| dy. \end{aligned}$$

By Lemma 3.8.1, $\inf_{g \in \Pi_{M-1}} \|K(\cdot, y) - g\|_{L^\infty(I_\mu)} \lesssim 2^{-kM} |K(\cdot, y)|_{W^{M, \infty}(I_\mu)}$ since I_μ is a cube of sidelength $c(M) \cdot 2^{-k}$. We thus obtain

$$\begin{aligned} |\langle H\psi_\lambda, \psi_\mu \rangle| &\lesssim 2^{-kM} \|\psi_\mu\|_{L^1(I_\mu)} \|\psi_\lambda\|_{L^1(I_\lambda)} \operatorname{ess\,sup}_{y \in I_{j,m}} |K(\cdot, y)|_{W^{M, \infty}(I_\mu)} \\ &\lesssim 2^{-kM} 2^{-\frac{dj}{2}} 2^{-\frac{dk}{2}} \operatorname{ess\,sup}_{y \in I_\lambda} |K(\cdot, y)|_{W^{M, \infty}(I_\mu)} \end{aligned}$$

since $\|\psi_\lambda\|_{L^1} = 2^{-\frac{dj}{2}} \|\psi\|_{L^1}$.

Since $H \in \mathcal{A}(M, f)$

$$\begin{aligned} \operatorname{ess\,sup}_{y \in I_\lambda} |K(\cdot, y)|_{W^{M, \infty}(I_\mu)} &= \operatorname{ess\,sup}_{y \in I_\lambda} \sum_{|\alpha|=M} \operatorname{ess\,sup}_{x \in I_\mu} |\partial_x^\alpha K(x, y)| \\ &\leq \sum_{|\alpha|=M} \operatorname{ess\,sup}_{(x, y) \in I_\lambda \times I_\mu} f(\|x - y\|_\infty) \\ &\lesssim \operatorname{ess\,sup}_{(x, y) \in I_\lambda \times I_\mu} f(\|x - y\|_\infty). \end{aligned}$$

Because f is a non-increasing function, $f(\|x - y\|_\infty) \leq f(\operatorname{dist}(I_\lambda, I_\mu))$ since $\operatorname{dist}(I_\lambda, I_\mu) = \inf_{(x, y) \in I_\lambda \times I_\mu} \|x - y\|_\infty$. Therefore

$$\begin{aligned} |\langle H\psi_\lambda, \psi_\mu \rangle| &\lesssim 2^{-kM} 2^{-\frac{dj}{2}} 2^{-\frac{dk}{2}} f(\operatorname{dist}(I_\lambda, I_\mu)) \\ &= 2^{-(M+\frac{d}{2})|j-k|} 2^{-j(M+d)} f(\operatorname{dist}(I_\lambda, I_\mu)). \end{aligned}$$

The case $k < j$ gives

$$|\langle H\psi_\lambda, \psi_\mu \rangle| \lesssim 2^{-(M+\frac{d}{2})|j-k|} 2^{-k(M+d)} f(\operatorname{dist}(I_\lambda, I_\mu)),$$

which allows to conclude that

$$|\langle H\psi_\lambda, \psi_\mu \rangle| \lesssim 2^{-(M+\frac{d}{2})|j-k|} 2^{-\min(j, k)(M+d)} f(\operatorname{dist}(I_\lambda, I_\mu)),$$

□

3.9 Proof of Theorem 3.4.1

Let us begin with some preliminary results. Recall that $\lambda = (j, m, e) \in \Lambda$ and $\mu = (k, n, e') \in \Lambda$. Since f is compactly supported on $[0, \kappa]$ and bounded by c_f , we have $f_{\lambda, \mu} = f(\operatorname{dist}(I_\lambda, I_\mu)) \leq c_f \mathbb{1}_{\operatorname{dist}(I_\lambda, I_\mu) \leq \kappa}$. By equation (3.10), $\operatorname{dist}(I_\mu, I_\lambda) \leq \kappa$ if $\|2^{-j}m - 2^{-k}n\|_\infty \leq R_{j, k}^\kappa$, where $R_{j, k}^\kappa = (2^{-j} + 2^{-k})c(M)/2 + \kappa$.

Lemma 3.9.1. *Define*

$$\mathcal{G}_{j,k}^{e,e'} = \left\{ (m, n) \in \mathcal{T}_j \times \mathcal{T}_k \mid \mathbb{1}_{\text{dist}(I_\lambda, I_\mu) \leq \kappa} = 1 \right\}.$$

Then $|\mathcal{G}_{j,k}^{e,e'}| \leq (2^j 2^{k+1} R_{j,k}^\kappa)^d$.

Proof. First note that

$$\mathcal{G}_{j,k}^{e,e'} = \left\{ (m, n) \in \mathcal{T}_j \times \mathcal{T}_k \mid \left| 2^{-j} m_i - 2^{-k} n_i \right| \leq R_{j,k}^\kappa, \quad \forall i \in \{1, \dots, d\} \right\}.$$

Now, define $\mathcal{G}_{j,k,m}^{e,e'} = \left\{ n \in \mathcal{T}_k \mid (m, n) \in \mathcal{G}_{j,k}^{e,e'} \right\}$. For a fixed (j, k, m, e, e') the set $\mathcal{G}_{j,k,m}^{e,e'}$ is a discrete hyper-cube of sidelength bounded above by $2^{k+1} R_{j,k}^\kappa$. Therefore $|\mathcal{G}_{j,k,m}^{e,e'}| \leq (2^{k+1} R_{j,k}^\kappa)^d$ coefficients. Moreover, $|\mathcal{T}_j| = 2^{jd}$, hence the number of coefficients in $\mathcal{G}_{j,k}^{e,e'}$ is bounded above by $(2^j 2^k R_{j,k}^\kappa)^d$. $\square$

Proof of i). We denote $J_{\max} = \log_2(N)/d$ the highest scale of decomposition. First note that a sufficient condition for $2^{-\min(j,k)(M+d)} f_{\lambda,\mu} \leq \eta$ is that $\min(j, k) \geq J(\eta)$ with $J(\eta) = \frac{-\log_2(\eta/c_f)}{M+d}$. In the following, we let $\tilde{J}(\eta) = \min(J(\eta), J_{\max})$ and define

$$\mathcal{G} = \bigcup_{\min(j,k) < J(\eta)} \bigcup_{e,e' \in \{0,1\}^d \setminus \{0\}} \mathcal{G}_{j,k}^{e,e'}.$$

The overall number of non zero coefficients $|\mathcal{G}|$ in Θ_η satisfies

$$\begin{aligned} |\mathcal{G}| &= \sum_{j=0}^{J_{\max}-1} \sum_{k=0}^{J_{\max}-1} \sum_{e,e' \in \{0,1\}^d} \# \mathcal{G}_{j,k}^{e,e'} \mathbb{1}_{\min(j,k) < J(\eta)} \\ &\lesssim (2^d - 1)^2 \sum_{j=0}^{J_{\max}-1} \sum_{k=0}^{J_{\max}-1} \mathbb{1}_{\min(j,k) < J(\eta)} 2^{jd} 2^{kd} \left(\frac{c(M)}{2} (2^{-j} + 2^{-k}) + \kappa \right)^d \\ &\lesssim \sum_{j=0}^{J_{\max}-1} \sum_{k=0}^{J_{\max}-1} \mathbb{1}_{\min(j,k) < J(\eta)} 2^{jd} 2^{kd} \left(\frac{c(M)^d}{2^d} 2^{-dj} + \frac{c(M)^d}{2^d} 2^{-dk} + \kappa^d \right) \\ &\lesssim \sum_{j=0}^{J_{\max}-1} \sum_{k=0}^{J_{\max}-1} \mathbb{1}_{\min(j,k) < J(\eta)} 2^{kd} + \sum_{j=0}^{J_{\max}-1} \sum_{k=0}^{J_{\max}-1} \mathbb{1}_{\min(j,k) < J(\eta)} 2^{jd} \\ &\quad + \sum_{j=0}^{J_{\max}-1} \sum_{k=0}^{J_{\max}-1} \mathbb{1}_{\min(j,k) < J(\eta)} 2^{kd} 2^{jd} \kappa^d. \end{aligned}$$

The first sum yields

$$\begin{aligned}
& \sum_{j=0}^{J_{\max}-1} \sum_{k=0}^{J_{\max}-1} \mathbb{1}_{\min(j,k) < J(\eta)} 2^{kd} \\
&= \left(\sum_{j=0}^{\tilde{J}(\eta)-1} \sum_{k=j}^{J_{\max}-1} 2^{kd} + \sum_{k=0}^{\tilde{J}(\eta)-1} 2^{kd} \sum_{j=k}^{J_{\max}-1} 1 \right) \\
&\lesssim \tilde{J}(\eta)N + 2^{d\tilde{J}(\eta)} \log_2(N) \lesssim \log_2(N)N.
\end{aligned}$$

The second sum is handled similarly and the third sum gives

$$\begin{aligned}
& \sum_{j=0}^{J_{\max}-1} \sum_{k=0}^{J_{\max}-1} \mathbb{1}_{\min(j,k) < J(\eta)} 2^{kd} 2^{kd} \kappa^d \\
&= \kappa^d \sum_{j=0}^{\tilde{J}(\eta)-1} 2^{jd} \sum_{k=j}^{J_{\max}-1} 2^{kd} + \sum_{k=0}^{\tilde{J}(\eta)-1} 2^{kd} \sum_{j=k}^{J_{\max}-1} 2^{jd} \\
&\lesssim \kappa^d N 2^{d\tilde{J}(\eta)}.
\end{aligned}$$

Overall $|\mathcal{G}| \lesssim \log_2(N)N + \eta^{-\frac{d}{M+d}}N$. For $\eta \leq \log_2(N)^{-(M+d)/d}$, the dominating terms are of kind $\eta^{-\frac{d}{M+d}}$, hence $|\mathcal{G}| \lesssim \eta^{-\frac{d}{M+d}}N\kappa^d$. $\square$

Proof of ii). Since Ψ is an orthogonal wavelet transform

$$\|\mathbf{H} - \tilde{\mathbf{H}}_\eta\|_{2 \rightarrow 2} = \|\Theta - \Theta_\eta\|_{2 \rightarrow 2}.$$

Let $\Delta_\eta = \Theta - \Theta_\eta$. We will make use of the following version of Shur inequality

$$\|\Delta_\eta\|_{2 \rightarrow 2}^2 \leq \|\Delta_\eta\|_{1 \rightarrow 1} \|\Delta_\eta\|_{\infty \rightarrow \infty}. \quad (3.25)$$

Since the upper-bound (3.9) is symmetric,

$$\|\Delta_\eta\|_{\infty \rightarrow \infty} = \|\Delta_\eta\|_{1 \rightarrow 1} = \max_{\lambda \in \Lambda} \sum_{\mu \in \Lambda} |\Delta_{\lambda, \mu}|$$

By definition of Θ_η we get that

$$\begin{aligned}
\sum_{\mu \in \Lambda} |\Delta_{\lambda, \mu}| &= \sum_{k=0}^{J_{\max}-1} \sum_{e' \in \{0,1\}^d \setminus \{0\}} \sum_{n \in \mathcal{G}_{j,k,m}^{e,e'}} |\theta_{\lambda, \mu}| \mathbb{1}_{\min(j,k) > J(\eta)} \\
&\lesssim \sum_{k=0}^{J_{\max}-1} \sum_{e' \in \{0,1\}^d \setminus \{0\}} \sum_{n \in \mathcal{G}_{j,k,m}^{e,e'}} 2^{-(M+\frac{d}{2})|j-k|} 2^{-\min(j,k)(M+d)} \mathbb{1}_{\min(j,k) > J(\eta)}.
\end{aligned}$$

Then

$$\begin{aligned}
\sum_{\mu \in \Lambda} |\Delta_{\lambda, \mu}| &\lesssim \sum_{k=0}^{J_{\max}-1} 2^{-(M+\frac{d}{2})|j-k|} 2^{-\min(j,k)(M+d)} \mathbb{1}_{\min(j,k) > J(\eta)} \left| \mathcal{G}_{j,k}^{e,e'} \right| \\
&\lesssim \sum_{k=0}^{j-1} (2^k R_{j,k}^\kappa)^d 2^{(k-j)(M+d/2)} 2^{-k(M+d)} \mathbb{1}_{k > J(\eta)} \\
&\quad + \sum_{k=j}^{J_{\max}-1} (2^k R_{j,k}^\kappa)^d 2^{(j-k)(M+d/2)} 2^{-j(M+d)} \mathbb{1}_{j > J(\eta)}.
\end{aligned}$$

The first sum on $k < j$ is equal to

$$\begin{aligned}
A_1 &= 2^{-jM} 2^{-jd/2} \sum_{k=0}^{j-1} (2^{k/2} R_{j,k}^\kappa)^d \mathbb{1}_{k > J(\eta)} \\
&= 2^{-jM} 2^{-jd/2} \mathbb{1}_{j > J(\eta)} \sum_{k=J(\eta)}^{j-1} (2^{k/2} R_{j,k}^\kappa)^d.
\end{aligned}$$

The second sum on $k \geq j$ is:

$$A_2 = \mathbb{1}_{j > J(\eta)} 2^{-jd/2} \sum_{k=j}^{J_{\max}-1} (R_{j,k}^\kappa)^d 2^{-k(M-d/2)}.$$

Now, notice that $(R_{j,k}^\kappa)^d \lesssim 2^{-jd} + 2^{-kd} + \kappa^d$. Thus

$$\begin{aligned}
A_1 &\lesssim 2^{-jM} 2^{-jd/2} \mathbb{1}_{j > J(\eta)} \sum_{k=J(\eta)}^{j-1} \left(2^{dk/2} 2^{-jd} + 2^{-dk/2} + 2^{kd/2} \kappa^d \right) \\
&\lesssim 2^{-jM} 2^{-jd/2} \mathbb{1}_{j > J(\eta)} \left(2^{-jd} 2^{jd/2} + 2^{-\frac{d}{2}J(\eta)} + \kappa^d 2^{jd/2} \right) \\
&= 2^{-jM} \mathbb{1}_{j > J(\eta)} \left(2^{-jd} + 2^{-\frac{d}{2}(J(\eta)+j)} + \kappa^d \right).
\end{aligned}$$

And

$$\begin{aligned}
A_2 &\lesssim \mathbb{1}_{j > J(\eta)} 2^{-jd/2} \sum_{k=j}^{J_{\max}-1} \left(2^{-jd} + 2^{-kd} + \kappa^d \right) 2^{-k(M-d/2)} \\
&\lesssim \mathbb{1}_{j > J(\eta)} 2^{-jd/2} \left(2^{-jd} 2^{-j(M-d/2)} + 2^{-j(M+d/2)} + \kappa^d 2^{-j(M-d/2)} \right) \\
&\lesssim \mathbb{1}_{j > J(\eta)} 2^{-jM} \left(2^{-jd} + \kappa^d \right).
\end{aligned}$$

Hence

$$\sum_{\mu \in \Lambda} |\Delta_{\lambda, \mu}| \lesssim \mathbb{1}_{j > J(\eta)} 2^{-jM} \left(2^{-jd} + \kappa^d + 2^{-\frac{d}{2}(J(\eta)+j)} \right).$$

Therefore

$$\begin{aligned}
\|\Delta_\eta\|_{1 \rightarrow 1} &\lesssim 2^{-J(\eta)M} \left(2^{-J(\eta)d} + \kappa^d + 2^{-dJ(\eta)} \right) \\
&\lesssim 2^{-J(\eta)M} \left(2^{-J(\eta)d} + \kappa^d \right) \\
&\lesssim \eta + \kappa^d \eta^{\frac{M}{M+d}} \\
&\lesssim \kappa^d \eta^{\frac{M}{M+d}} \quad \text{for small } \eta.
\end{aligned}$$

Finally, we can see that there exists a constant C_M independent of N such that

$$\|\Delta_\eta\|_{1 \rightarrow 1} \leq C_M \kappa^d \eta^{\frac{M}{M+d}} \quad \text{and} \quad \|\Delta_\eta\|_{\infty \rightarrow \infty} \leq C_M \kappa^d \eta^{\frac{M}{M+d}}.$$

It suffices to use inequality (3.25) to conclude. □

Proof of iii). This is a direct consequence of point i) and ii). □

4

Accelerating $\ell^1 - \ell^2$ deblurring using wavelet expansions of operators

Résumé : *Ce chapitre est issu de [99] soumis en décembre 2015 en collaboration avec Pierre Weiss. La restauration d'images dégradées par des flous est un problème fondamental en traitement d'image. Il est généralement résolu en utilisant des procédures d'optimisation gourmandes en calculs. Le but de ce chapitre est de proposer de nouvelles stratégies efficaces permettant de réduire les temps de calculs pour des modèles de restauration simple utilisant des régularisations basées sur des transformées en ondelettes orthogonales. Nous montrons que la minimisation peut être accélérée significativement en utilisant le fait que les images et les opérateurs sont compressibles dans la même base orthonormale d'ondelettes. La méthodologie proposée repose sur trois ingrédients: i) une approximation parcimonieuse de l'opérateur dans une base d'ondelettes, ii) un préconditionneur diagonal et iii) une implémentation massivement parallèle. La combinaison de ces trois ingrédients donne lieu à des facteurs d'accélération variant de 30 à 250 sur une station de travail standard. Par exemple, la restauration d'une image 1024×1024 est effectuée en 0.15 secondes.*

Abstract: *Image deblurring is a fundamental problem in imaging, usually solved with computationally intensive optimization procedures. The goal of this paper is to provide new efficient strategies to reduce computing times for simple deblurring models regularized using orthogonal wavelet transforms. We show that the minimization can be significantly accelerated by leveraging the fact that images and blur operators are compressible in the same orthogonal wavelet basis. The proposed methodology consists of three ingredients: i) a sparse approximation of the blur operator in wavelet bases, ii) a diagonal preconditioner and iii) an implementation on massively parallel architectures. Combining the three ingredients leads to acceleration factors ranging from 30 to 250 on a typical workstation. For instance, a 1024×1024 image can be deblurred in 0.15 seconds.*

Contents

4.1	Introduction	123
4.1.1	Image formation and image restoration models	123
4.1.2	Standard optimization algorithms	124
4.2	Main ideas	125
4.2.1	Related works	126
4.2.2	Contributions	126
4.2.3	Paper outline	127
4.3	Notation	127
4.4	Wavelet decompositions of blurring operators	129
4.4.1	Definition of blurring operators	129
4.4.2	Decomposition in wavelet bases	129
4.4.3	Approximation in wavelet bases	130
4.4.4	Discretization	131
4.4.5	Illustration	131
4.4.6	Decomposition of convolutions	131
4.4.7	Thresholding strategies	134
4.5	On the design of preconditioners	136
4.5.1	What governs convergence rates?	137
4.5.2	Jacobi preconditioner	138
4.5.3	SPAI preconditioner	138
4.5.4	GPU implementation	139
4.6	Numerical experiments	139
4.6.1	On the role of thresholding strategies	140
4.6.2	Approximation in wavelet bases	140
4.6.3	A competitor: the ADMM	142
4.6.4	Comparing preconditioners and ADMM	147
4.6.5	Computing times	147
4.6.6	Dependency on the blur kernel	149
4.6.7	Dependency on resolution	151
4.6.8	Spatially varying blurs	151

4.1 Introduction

Most imaging devices produce blurry images. This degradation very often prevents the correct interpretation of image contents and sometimes ruins expensive experiments. One of the most advertised examples of that type is Hubble space telescope ¹, which was discovered to suffer from severe optical aberrations after being launched. Such situations occur on a daily basis in fields such as biomedical imaging, astronomy or conventional photography.

Starting from the seventies, a large number of numerical methods to deblur images was therefore developed. The first methods were based on linear estimators such as the Wiener filter [249]. They were progressively replaced by more complicated nonlinear methods, incorporating prior knowledge on the image contents. We refer the interested reader to the following review papers [215, 239, 196] to get an overview of the available techniques.

Despite providing better reconstruction results, the most efficient methods are often disregarded in practice, due to their high computational complexity, especially for large 2D or 3D images. The goal of this paper is to develop new numerical strategies that significantly reduce the computational burden of image deblurring. The proposed ideas yield a *fast deblurring method*, compatible with *large data* and *routine use*. They allow handling both *stationary and spatially varying blurs*. The proposed algorithm does not reach the state-of-the-art in terms of image quality, because the prior is too simple, but still performs well in short computing times.

4.1.1 Image formation and image restoration models

In this paper, we assume that the observed image u_0 reads:

$$u_0 = Hu + b, \quad (4.1)$$

where $u \in \mathbb{R}^N$ is the clean image we wish to recover, $b \sim \mathcal{N}(0, \sigma^2 I_N)$ is a white Gaussian noise of standard deviation σ and $H \in \mathbb{R}^{N \times N}$ is a known blurring operator. Loosely speaking, a blurring operator replaces the value of a pixel by a mean of its neighbors. A precise definition will be given in Section 4.4.

Let $\Psi \in \mathbb{R}^{N \times N}$ denote an orthogonal wavelet transform and let $A = H\Psi$. A standard variational formulation to restore u consists of solving:

$$\min_{x \in \mathbb{R}^N} E(x) = F(x) + G(x). \quad (4.2)$$

In this equation, $F(x)$ is a quadratic data fidelity term defined by

$$F(x) = \frac{1}{2} \|Ax - u_0\|_2^2. \quad (4.3)$$

¹The total cost of Hubble telescope is estimated at 10 billions US Dollars [13].

The regularization term $G(x)$ is defined by:

$$G(x) = \|x\|_{1,w} = \sum_{i=1}^N w[i] |x[i]| \quad (4.4)$$

The vector of weights $w \in \mathbb{R}_+^N$ is a regularization parameter that may vary across subbands of the wavelet transform. The weighted ℓ^1 -norm is well known to promote sparse vectors. This is usually advantageous since images are compressible in the wavelet domain. Overall, problem (4.2) consists of finding an image Ψx consistent with the observed data u_0 with a sparse representation x in the wavelet domain.

Many groups worldwide have proposed minimizing similar cost functions in the literature, see e.g. [110, 182, 235, 236]. The current trend is to use redundant dictionaries Ψ such as the undecimated wavelet transforms or learned transforms instead of orthogonal transforms [214, 59, 51, 40]. This usually allows reducing reconstruction artifacts. We focus here on the case where Ψ is orthogonal. This property will help designing much faster algorithms traded for some image quality.

4.1.2 Standard optimization algorithms

A lot of algorithms based on proximity operators were designed in the last decade to solve convex problems of type (4.2). We refer the reader to the review papers [19, 73] to get an overview of the available techniques. A typical method is the accelerated proximal gradient descent, also known as FISTA (Fast Iterative Soft Thresholding Algorithm) [18]. By letting $\|A\|_2$ denote the largest singular value of A , it takes the form described in Algorithm 4. This method got very popular lately due to its ease of implementation and relatively fast convergence.

Algorithm 4: Accelerated proximal gradient descent

Input: Initial guess $x^{(0)} = y^{(1)}$, $\tau = 1/\ A\ _2^2$ and Nit for $k = 1$ to Nit do	
Compute $\nabla F(y^{(k)}) = A^*(Ay^{(k)} - u_0)$.	// 99.35''
$x^{(k)} = \text{Prox}_{\tau G} \left(y^{(k)} - \tau \nabla F(y^{(k)}) \right)$.	// 2.7''
$y^{(k+1)} = x^{(k)} + \frac{k-1}{k+2} (x^{(k)} - x^{(k-1)})$.	// 1.1''
end	

Let us illustrate this method on a practical deconvolution experiment. We use a 1024×1024 image and assume that $Hu = h \star u$, where $\star$ denotes the discrete convolution product and h is a motion blur described on Figure 4.4b. In practice, the PSNR of the deblurred image stabilizes after 500 iterations. The computing times on a workstation with Matlab and mex-files is around 103''15. The result is shown on Figure 4.5. Profiling the code leads to the computing times shown on the right-hand-side of Algorithm 4. As can be seen, 96% of the computing time is spent in the gradient evaluation. This

requires computing two wavelet transforms and two fast Fourier transforms. This simple experiment reveals that two approaches can be used to reduce computing times:

- *Accelerate gradients computation.*
- *Use more sophisticated minimization algorithms to accelerate convergence.*

4.2 Main ideas

The method proposed in this paper relies on three ideas. First, function F in equation (4.3) can be approximated by another function F_K such that ∇F_K is inexpensive to compute. Second, we characterize precisely the structure of the Hessian of F_K , allowing to design efficient preconditioners. Finally, we implement the iterative algorithm on a GPU. The first two ideas, which constitute the main contribution of this paper, are motivated by our recent observation that spatially varying blur operators are compressible and have a well characterized structure in the wavelet domain [100]. We showed that matrix

$$\Theta = \Psi^* H \Psi, \quad (4.5)$$

which differs from H by a change of basis, has a particular banded structure, with many negligible entries. Therefore, one can construct a K -sparse matrix Θ_K (i.e. a matrix with at most K non zero entries) such that $\Theta_K \simeq \Theta$.

Problem approximation Using the fact that Ψ is an orthogonal transform allows writing that:

$$\min_{x \in \mathbb{R}^N} \frac{1}{2} \|H \Psi x - u_0\|_2^2 + \|x\|_{1,w} \quad (4.6)$$

$$= \min_{x \in \mathbb{R}^N} \frac{1}{2} \|\Psi^* (H \Psi x - u_0)\|_2^2 + \|x\|_{1,w} \quad (4.7)$$

$$= \min_{x \in \mathbb{R}^N} \frac{1}{2} \|\Theta x - x_0\|_2^2 + \|x\|_{1,w}, \quad (4.8)$$

where $x_0 = \Psi^* u_0$ is the wavelet decomposition of u_0 . Problem (4.8) is expressed entirely in the wavelet domain, contrarily to problem (4.2). However, matrix-vector products with Θ might be computationally expensive. We therefore approximate the variational problem (4.8) by:

$$\min_{x \in \mathbb{R}^N} \frac{1}{2} \|\Theta_K x - x_0\|_2^2 + \|x\|_{1,w}. \quad (4.9)$$

Now, let $F_K(x) = \frac{1}{2} \|\Theta_K x - x_0\|_2^2$. The gradient of F_K reads:

$$\nabla F_K(x) = \Theta_K^* (\Theta_K x - x_0), \quad (4.10)$$

and therefore requires computing two matrix-vector products with sparse matrices. Computing the approximate gradient (4.10) is usually much cheaper than computing $\nabla F(x)$

exactly using Fast Fourier transforms and fast wavelet transforms. This may come as a surprise since they are respectively of complexity $O(N \log(N))$ and $O(N)$. In fact, we will see that in favorable cases, the evaluation of $\nabla F_K(x)$ may require about 2 operations per pixel!

Preconditioning The second ingredient of our method relies on the observation that the Hessian matrix $H_{F_K}(x) = \Theta_K^* \Theta_K$ has a near diagonal structure with decreasing entries on the diagonal. This allows designing *efficient preconditioners*, which reduces the number of iterations necessary to reach a satisfactory precision. In practice preconditioning leads to acceleration factors ranging from 2 to 5.

GPU implementation Finally, using massively parallel programming on graphic cards still leads to an acceleration factor of order 10 on an NVIDIA K20c. Of course, this factor could be improved further by using more powerful graphic cards. Combining the three ingredients leads to algorithms that are from 4 to 250 times faster than FISTA or the ADMM algorithms applied to (4.2), which arguably constitute the current state-of-the-art.

4.2.1 Related works

The idea of characterizing integral operators in the wavelet domain appeared nearly at the same time as wavelets, at the end of the eighties. Y. Meyer characterized many properties of Calderón-Zygmund operators in his seminal book [169]. Later, Beylkin, Coifman and Rokhlin [29], showed that those theoretical results may have important consequences for the fast resolution of partial differential equations and the compression of matrices. Since then, the idea of using multi-scale representations has been used extensively in numerical simulation of physical phenomena. The interested reader can refer to [67] for some applications.

Quite surprisingly, it seems that very few researchers attempted to apply them in imaging. In [57, 162], the authors proposed to approximate integral operators by matrices diagonal in the wavelet domain. Our experience is that diagonal approximations are too crude to provide sufficiently good approximations. More recently the authors of [241, 100] proposed independently to compress operators in the wavelet domain. However they did not explore its implications for the fast resolution of inverse problems.

On the side of preconditioning, the two references [235, 236] are closely related to our work. The authors designed a few preconditioners to accelerate the convergence of the proximal gradient descent (also called thresholded Landweber algorithm or Iterative Soft Thresholding Algorithm). Overall, the idea of preconditioning is therefore not new. To the best of our knowledge, our contribution is however the first that is based on a precise understanding of the structure of Θ .

4.2.2 Contributions

The main contributions of this paper are as follows.

- We show that the ideas proposed in [100] make the new formulation (4.8) particularly attractive for the design of optimization algorithms.
- We provide a fast $O(N \log(N)^2)$ algorithm to compute matrix Θ in the case where H is a convolution operator.
- We design preconditioners based on clear understanding of the problem geometry. To the best of our knowledge, existing attempts to precondition the problem are rather heuristic [235, 236].
- We provide a gallery of numerical comparisons and examples to show the well-foundedness of the proposed approach.

4.2.3 Paper outline

The paper is structured as follows. We first provide some notation and definitions in section 4.3. We then provide a few results characterizing the structure of blurring operators in section 4.4. This section is mostly based on our previous work [100]. We propose two simple explicit preconditioners in section 4.5. Finally, we perform numerical experiments and comparisons in section 4.6.

4.3 Notation

In this paper, we consider d dimensional images. To simplify the discussion, we use periodic boundary conditions and work on the d -dimensional torus $\Omega = \mathbb{T}^d$, where $\mathbb{T}^d = \mathbb{R}^d / \mathbb{Z}^d$. The space $L^2(\Omega)$ denotes the space of squared integrable functions defined on Ω .

Let $\alpha = (\alpha_1, \dots, \alpha_d)$ denote a multi-index. The sum of its components is denoted $|\alpha| = \sum_{i=1}^d \alpha_i$. The Sobolev spaces $W^{M,p}$ are defined as the set of functions $f \in L^p$ with partial derivatives up to order M in L^p where $p \in [1, +\infty]$ and $M \in \mathbb{N}$. These spaces, equipped with the following norm are Banach spaces

$$\|f\|_{W^{M,p}} = \sum_{|\alpha| \leq M} \|\partial^\alpha f\|_{L^p}, \quad (4.11)$$

where $\partial^\alpha f = \frac{\partial^{\alpha_1}}{\partial x_1^{\alpha_1}} \dots \frac{\partial^{\alpha_d}}{\partial x_d^{\alpha_d}} f$.

Let us now define a wavelet basis on $L^2(\Omega)$. To this end, we first introduce a 1D wavelet basis on $\mathbb{T}$. Let ϕ and ψ denote the scaling and mother wavelets and assume that the mother wavelet ψ has M vanishing moments, i.e.

$$\text{for all } 0 \leq m < M, \quad \int_{[0,1]} t^m \psi(t) dt = 0. \quad (4.12)$$

We assume that $\text{supp}(\psi) = [-c(M)/2, c(M)/2]$. Note that $c(M) \geq 2M - 1$, with equality for Daubechies wavelets, see, e.g., [164, Theorem 7.9, p. 294].

Translated and dilated versions of the wavelets are defined, for $j \geq 0$, as follows

$$\phi_{j,l} = 2^{j/2} \phi(2^j \cdot - l), \quad (4.13)$$

$$\psi_{j,l} = 2^{j/2} \psi(2^j \cdot - l), \quad (4.14)$$

with $l \in \mathcal{T}_j$ and $\mathcal{T}_j = \{0, \dots, 2^j - 1\}$.

In dimension d , we use isotropic separable wavelet bases, see, e.g., [164, Theorem 7.26, p. 348]. Let $m = (m_1, \dots, m_d)$. Define $\rho_{j,l}^0 = \phi_{j,l}$ and $\rho_{j,l}^1 = \psi_{j,l}$. Let $e = (e_1, \dots, e_d) \in \{0, 1\}^d$. For the ease of reading, we will use the shorthand notation $\lambda = (j, m, e)$ and $|\lambda| = j$. We also let

$$\Lambda_0 = \left\{ (j, m, e) \mid j \in \mathbb{Z}, m \in \mathcal{T}_j, e \in \{0, 1\}^d \right\} \quad (4.15)$$

and

$$\Lambda = \left\{ (j, m, e) \mid j \in \mathbb{Z}, m \in \mathcal{T}_j, e \in \{0, 1\}^d \setminus \{0_{\mathbb{R}^d}\} \right\}. \quad (4.16)$$

Wavelet ψ_λ is defined by $\psi_\lambda(x_1, \dots, x_d) = \psi_{j,m}^e(x_1, \dots, x_d) = \rho_{j,m_1}^{e_1}(x_1) \dots \rho_{j,m_d}^{e_d}(x_d)$. Elements of the separable wavelet basis consist of tensor products of scaling and mother wavelets at the same scale. Note that if $e \neq 0$ wavelet $\psi_{j,m}^e$ has M vanishing moments in $\mathbb{R}^d$. Let $I_{j,m} = \cup_e \text{supp } \psi_{j,m}^e$ and $I_\lambda = \text{supp } \psi_\lambda$. The distance between the supports of ψ_λ and ψ_μ is defined by

$$\text{dist}(I_\lambda, I_\mu) = \inf_{x \in I_\lambda, y \in I_\mu} \|x - y\|_\infty \quad (4.17)$$

$$= \max \left(0, \left\| 2^{-j}m - 2^{-k}n \right\|_\infty - (2^{-j} + 2^{-k}) \frac{c(M)}{2} \right). \quad (4.18)$$

This distance will play an important role to describe the structure of matrix Θ .

With these definitions, every function $f \in L^2(\Omega)$ can be written as

$$u = \langle u, \psi_{0,0}^0 \rangle \psi_{0,0}^0 + \sum_{e \in \{0,1\}^d \setminus \{0\}} \sum_{j=0}^{+\infty} \sum_{m \in \mathcal{T}_j} \langle u, \psi_{j,m}^e \rangle \psi_{j,m}^e \quad (4.19)$$

$$= \langle u, \psi_{0,0}^0 \rangle \psi_{0,0}^0 + \sum_{\lambda \in \Lambda} \langle u, \psi_\lambda \rangle \psi_\lambda \quad (4.20)$$

$$= \sum_{\lambda \in \Lambda_0} \langle u, \psi_\lambda \rangle \psi_\lambda. \quad (4.21)$$

Finally, we let $\Psi^* : L^2(\Omega) \rightarrow l^2(\mathbb{Z})$ denote the wavelet decomposition operator and $\Psi : l^2(\mathbb{Z}) \rightarrow L^2(\Omega)$ its associated reconstruction operator. The discrete wavelet transform is also denoted $\Psi : \mathbb{R}^N \rightarrow \mathbb{R}^N$. We refer to [164, 79, 71] for more details on the construction of wavelet bases.

4.4 Wavelet decompositions of blurring operators

In this section we remind some results on the decomposition of blurring operators in the wavelet domain.

4.4.1 Definition of blurring operators

A blurring operator H can be modeled by a linear integral operator $H : L^2(\Omega) \rightarrow L^2(\Omega)$:

$$\forall x \in \Omega, \quad Hu(x) = \int_{\Omega} K(x, y)u(y)dy. \quad (4.22)$$

The function $K : \Omega \times \Omega \rightarrow \mathbb{R}$ is called kernel of the integral operator and defines the Point Spread Function (PSF) $K(\cdot, y)$ at location $y \in \Omega$. The image Hu is the blurred version of u . Following our recent paper [100], we define blurring operators as follows.

Definition 4.4.1 (Blurring operators [100]). Let $M \in \mathbb{N}$ and $f : [0, 1] \rightarrow \mathbb{R}_+$ denote a non-increasing bounded function. An integral operator is called a blurring operator in the class $\mathcal{A}(M, f)$ if it satisfies the following properties:

1. Its kernel $K \in W^{M, \infty}(\Omega \times \Omega)$;
2. The partial derivatives of K satisfy:

$$(a) \quad \forall |\alpha| \leq M, \forall (x, y) \in \Omega \times \Omega, \quad |\partial_x^\alpha K(x, y)| \leq f(\|x - y\|_\infty), \quad (4.23)$$

$$(b) \quad \forall |\alpha| \leq M, \forall (x, y) \in \Omega \times \Omega, \quad \left| \partial_y^\alpha K(x, y) \right| \leq f(\|x - y\|_\infty). \quad (4.24)$$

Condition (4.23) means that the PSF is smooth, while condition (4.24) indicates that the PSFs vary smoothly. These regularity assumptions are met in a large number of practical problems. In addition, they allow deriving theorems similar to those of the seminal papers of Y. Meyer, R. Coifman, G. Beylkin and V. Rokhlin [171, 29]. Those results basically state that an operator in the class $\mathcal{A}(M, f)$ can be represented and computed efficiently when decomposed in a wavelet basis. We make this key idea precise in the next paragraph.

4.4.2 Decomposition in wavelet bases

Since H is defined on the Hilbert space $L^2(\Omega)$, it can be written as $H = \Psi\Theta\Psi^*$, where $\Theta : \ell^2(\mathbb{Z}) \rightarrow \ell^2(\mathbb{Z})$ is the infinite matrix representation of the operator H in the wavelet domain. Matrix Θ is characterized by the coefficients:

$$\theta_{\lambda, \mu} = \Theta[\lambda, \mu] = \langle H\psi_\lambda, \psi_\mu \rangle. \quad (4.25)$$

The following result provides a good upper-bound on their amplitude.

Theorem 4.4.1 (Representation of blurring operator in wavelet bases [100]). *Let $f_{\lambda,\mu} = f(\text{dist}(I_\lambda, I_\mu))$ and assume that:*

- *Operator H belongs to the class $\mathcal{A}(M, f)$ (see Definition 4.4.1).*
- *The mother wavelet is compactly supported with M vanishing moments.*

Then for all $\lambda = (j, m, e) \in \Lambda$ and $\mu = (k, n, e') \in \Lambda$, with $e, e' \neq 0$:

$$|\theta_{\lambda,\mu}| \leq C_M 2^{-(M+\frac{d}{2})|j-k|} 2^{-\min(j,k)(M+d)} f_{\lambda,\mu}, \quad (4.26)$$

where C_M is a constant that does not depend on λ and μ .

The coefficients of Θ decay exponentially with respect to the scale difference and also as a function of the distance between the two wavelets supports.

4.4.3 Approximation in wavelet bases

In order to get a representation of the operator in a finite dimensional setting, the wavelet representation can be truncated at scale J . Let $\Theta^{(J)}$ denote the infinite matrix defined by:

$$\Theta^{(J)}[\lambda, \mu] = \begin{cases} \theta_{\lambda,\mu} & \text{if } |\lambda| \leq J \text{ and } |\mu| \leq J, \\ 0 & \text{otherwise.} \end{cases} \quad (4.27)$$

This matrix contains at most N^2 non-zero coefficients, where $N = 1 + \sum_{j=0}^{J-1} (2^d - 1)2^{dj}$ denotes the numbers of wavelets kept to represent functions. The operator $H^{(J)} = \Psi \Theta^{(J)} \Psi^*$ is a Galerkin approximation of H . It is possible to control the discretization error $\|H^{(J)} - H\|_{X \rightarrow Y}$, where $\|\cdot\|_{X \rightarrow Y}$ is some operator norm between two function spaces X and Y . We refer to [100] for more details.

The following theorem is a variation of [29]. Loosely speaking, it states that $H^{(J)}$ can be well approximated by a matrix containing only $O(N)$ coefficients.

Theorem 4.4.2 (Computation of blurring operators in wavelet bases [100]). *Set $0 \leq \eta \leq \log_2(N)^{-(M+d)/d}$. Let $\Theta_\eta^{(J)}$ be the matrix obtained by zeroing all coefficients in $\Theta^{(J)}$ such that*

$$2^{-\min(j,k)(M+d)} f_{\lambda,\mu} \leq \eta. \quad (4.28)$$

Define $H_\eta^{(J)} = \Psi \Theta_\eta^{(J)} \Psi^$. Under the same hypotheses as Theorem (4.4.1), the number of coefficients needed to satisfy $\|H^{(J)} - H_\eta^{(J)}\|_{2 \rightarrow 2} \leq \epsilon$ is bounded above by*

$$C'_M N \epsilon^{-\frac{d}{M}} \quad (4.29)$$

where $C'_M > 0$ is independent of N .

This theorem has important consequences for numerical analysis. It states that evaluations of H can be obtained with an ϵ accuracy using only $O(N \epsilon^{-\frac{d}{M}})$ operations. Note that the smoothness M of the kernel is handled automatically. Of interest, let us mention that [171, 29] proposed similar results under less stringent conditions. In particular, similar inequalities may still hold true if the kernel blows up on the diagonal.

4.4.4 Discretization

In the discrete setting, the above results can be exploited as follows. Given a matrix $H \in \mathbb{R}^{N \times N}$ that represents a discretized version of H , we perform the change of basis:

$$\Theta = \Psi^* H \Psi, \quad (4.30)$$

where $\Psi \in \mathbb{R}^{N \times N}$ is the discrete isotropic separable wavelet transform. Similarly to the continuous setting, matrix Θ is essentially concentrated along the diagonals of the wavelet sub-bands (see Figure 4.1).

Probably the main reason why this decomposition has very seldom been used in practice is that it is very computationally demanding. Computing the whole set of coefficients $(\langle H\psi_\lambda, \psi_\mu \rangle)_{\lambda, \mu}$ is an $O(N^3)$ algorithm! To evaluate this set, one has to apply H to each of the N discrete wavelets. The cost of evaluating one vector $H\psi_\lambda$ is $O(N^2)$ operations since it is a matrix-vector product with an $N \times N$ matrix. Then, evaluating $(\langle H\psi_\lambda, \psi_\mu \rangle)_{\lambda, \mu}$ for all μ and a fixed λ costs $O(N)$ operations (a discrete wavelet transform). The overall complexity is therefore dominated by the evaluation of N vectors $H\psi_\lambda$, with complexity $O(N^3)$.

This computational burden is hardly tractable for large signals, unless it can be computed once for all. This happens when the blur is device dependent but not sample dependent. We will see that computing times are drastically reduced for *convolution operators* in section 4.4.6.

4.4.5 Illustration

In order to illustrate the various results provided so far, let us consider an operator acting on 1D signals, with kernel defined by

$$K(x, y) = \frac{1}{\sigma(y)\sqrt{2\pi}} \exp\left(-\frac{(x-y)^2}{2\sigma^2(y)}\right), \quad (4.31)$$

where $\sigma(y) = 4 + 10y$. All PSFs are Gaussian with a variance that increases linearly. The matrix is displayed in linear scale (resp. log scale) in Figure 4.1 top-left (resp. top-right). The wavelet representation of the matrix is displayed in linear scale (resp. log scale) in Figure 4.1 bottom-left (resp. bottom-right). As can be seen, the matrix expressed in the wavelet basis is sparser than in the space domain. It has a particular banded structure captured by Theorem 4.4.1.

4.4.6 Decomposition of convolutions

From now on, we assume that H is a convolution with a kernel h . The results below hold both in the continuous and the discrete setting. We establish them in the discrete setting to ease the implementation. Matrix Θ can be decomposed into its wavelet sub-bands:

$$\Theta = \left(\Theta_{j,k}^{e,e'} \right)_{j,k}, \quad \text{with } \Theta_{j,k}^{e,e'} = \left(\left\langle H\psi_{j,m}^e, \psi_{k,n}^{e'} \right\rangle \right)_{m \in \mathcal{T}_j, n \in \mathcal{T}_k}. \quad (4.32)$$




Figure 4.1 – An illustration of the compression of a spatially varying blur in the wavelet domain. Top-left: H . Top-right: H in $\log_{10}$ -scale. Bottom-left: Θ obtained using Daubechies wavelets with 10 vanishing moments and a decomposition level $J = 7$. Bottom-right: Θ in $\log_{10}$ -scale.

For instance, on 1D signals with $J = 2$, matrix Θ can be decomposed as shown in Figure 4.2, left. Let us now describe the specific structure of the sub-bands $\Theta_{j,k}^{e,e'}$ for convolutions. We will need the following definitions.

Definition 4.4.2 (Translation operator). Let $a \in \mathbb{R}^N$ denote a d -dimensional image and $m \in \mathbb{Z}^d$ denote a shift. The translated image $b = \tau_m(a)$ is defined for all $i_1, \dots, i_d$ by:

$$b[i_1, \dots, i_d] = a[i_1 - m_1, \dots, i_d - m_d] \quad (4.33)$$

with circular boundary conditions.

Definition 4.4.3 (Rectangular circulant matrices). Let $A \in \mathbb{R}^{2^j \times 2^k}$ denote a rectangular matrix. It is called circulant if and only if:

- When $k \geq j$: there exists $a \in \mathbb{R}^{2^k}$, such that, for all $0 \leq l \leq 2^j - 1$,

$$A[l, :] = \tau_{2^{k-j}l}(a).$$

- When $k < j$: there exists $a \in \mathbb{R}^{2^j}$, such that, for all $0 \leq l \leq 2^k - 1$,

$$A[:, l] = \tau_{2^{j-k}l}(a).$$

As an example, a 4×8 circulant matrix is of the form:

$$A = \begin{pmatrix} a_1 & a_2 & a_3 & a_4 & a_5 & a_6 & a_7 & a_8 \\ a_7 & a_8 & a_1 & a_2 & a_3 & a_4 & a_5 & a_6 \\ a_5 & a_6 & a_7 & a_8 & a_1 & a_2 & a_3 & a_4 \\ a_3 & a_4 & a_5 & a_6 & a_7 & a_8 & a_1 & a_2 \end{pmatrix}.$$




Figure 4.2 – Left: structure of Θ . Right: Θ in $\log_{10}$ -scale when H is a convolution with a Gaussian kernel.

Theorem 4.4.3 states that all the wavelet sub-bands of Θ are circulant for convolution operators. This is illustrated on Figure 4.2, right.

Theorem 4.4.3 (Circulant structure of Θ). *Let H be a convolution matrix and define $\Theta = \Psi^* H \Psi$ be its wavelet representation. Then, for all $j, k \in [0, J]$ and $e, e' \in \{0, 1\}^d$, the sub-matrices $\Theta_{j,k}^{e,e'}$ are circulant.*

Proof. We only treat the case $j \geq k$, since the case $j < k$ is similar. Let $a \in \mathbb{R}^{2^j}$ be defined by:

$$a[m] = \langle h \star \psi_{j,m}^e, \psi_{k,0}^{e'} \rangle. \quad (4.34)$$

We have:

$$\langle h \star \psi_{j,m}^e, \psi_{k,n}^{e'} \rangle = \langle h \star \psi_{j,m}^e, \tau_{2^k n}(\psi_{k,0}^{e'}) \rangle \quad (4.35)$$

$$= \langle \tau_{-2^k n}(h \star \psi_{j,m}^e), \psi_{k,0}^{e'} \rangle \quad (4.36)$$

$$= \langle h \star \tau_{-2^k n}(\psi_{j,m}^e), \psi_{k,0}^{e'} \rangle \quad (4.37)$$

$$= \langle h \star \psi_{j,m-2^{j-k}n}^e, \psi_{k,0}^{e'} \rangle \quad (4.38)$$

$$= a[m - 2^{j-k}n]. \quad (4.39)$$

The sub-matrix $\Theta_{j,k}^{e,e'}$ is therefore circulant. In this list of identities, we only used the fact that the adjoint of $\tau_{2^k n}$ is $\tau_{-2^k n}$ and the fact that translations and convolution commute. $\square$

The main consequence of Theorem 4.4.3 is that the computation of matrix Θ reduces to computing one column or one row of each matrix $\Theta_{j,k}^{e,e'}$. This can be achieved by computing $(2^d - 1)J$ wavelet transforms (see Algorithm 5). The complexity of computing Θ therefore reduces to $O((2^d - 1)JN)$ operations instead of $O(N^3)$ operations for spatially varying operators.

4.4.7 Thresholding strategies

Theorem 4.4.2 ensures that one can construct good sparse approximations of Θ . However, the thresholding strategy suggested by the theorem turns out to be impractical.

In this section, we propose efficient thresholding strategies. Most of the proposed ideas come from our recent work [100], and we refer to this paper for more details. The algorithm specific to convolution operators is new.

Let us define the operator norm:

$$\|H\|_{X \rightarrow Y} = \sup_{\|u\|_X \leq 1} \|Hu\|_Y, \quad (4.40)$$

where $\|\cdot\|_X$ and $\|\cdot\|_Y$ denote two norms on $\mathbb{R}^N$. A natural way to obtain a K -sparse approximation Θ_K of Θ consists of finding the minimizer of:

$$\min_{\Theta_K, K\text{-sparse}} \|\Theta_K - \Theta\|_{X \rightarrow Y}, \quad (4.41)$$

Algorithm 5: An algorithm to compute Θ for convolution operator

Input: $h \in \mathbb{R}^N$, the convolution kernel of H
Output: Θ , the wavelet representation of H
for $(j, e) \in [0, J] \times \{0, 1\}^d$ **do**
 Compute the wavelet ψ_λ with $\lambda = (j, e, 0)$.
 Compute the blurred wavelets $H\psi_\lambda$ and $H^*\psi_\lambda$.
 Compute $(\langle H\psi_\lambda, \psi_\mu \rangle)_\mu$ using one forward wavelet transform.
 Compute $(\langle H^*\psi_\lambda, \psi_\mu \rangle)_\mu$ using one forward wavelet transform.
 for $(k, e') \in [0, J] \times \{0, 1\}^d$ **do**
 if $k \geq j$ **then**
 $\Theta_{j,k}^{e,e'}$ is the circulant matrix with column: $(\langle H\psi_\lambda, \psi_{k,n}^{e'} \rangle)_n$
 else
 $\Theta_{j,k}^{e,e'}$ is the circulant matrix with row:
 $(\langle H^*\psi_\lambda, \psi_{k,n}^{e'} \rangle)_n = (\langle \psi_\lambda, H\psi_{k,n}^{e'} \rangle)_n$

where $H_K = \Psi\Theta_K\Psi^*$. The most naive thresholding strategy consists of constructing a matrix Θ_K such that:

$$\Theta_K[\lambda, \mu] = \begin{cases} \Theta[\lambda, \mu] & \text{if } |\Theta[\lambda, \mu]| \text{ is among the } K \text{ largest values of } |\Theta|, \\ 0 & \text{otherwise.} \end{cases} \quad (4.42)$$

This thresholding strategy can be understood as the solution of the minimization problem (4.41), by setting $\|\cdot\|_X = \|\Psi^* \cdot\|_1$ and $\|\cdot\|_Y = \|\Psi^* \cdot\|_\infty$.

The ℓ^1 -norm of the wavelet coefficients is not adapted to the description of images. Natural images are often modeled as elements of Besov spaces or the space of bounded variation functions [72, 10]. These spaces can be characterized by the decay of their wavelet coefficients [67] across sub-bands. This motivates to set $\|\cdot\|_X = \|\Sigma\Psi^* \cdot\|_1$ where $\Sigma = \text{diag}(\sigma) \in \mathbb{R}^{N \times N}$ is a diagonal matrix and where $\sigma \in \mathbb{R}^N$ is constant by levels. The thresholding strategy using this metric can be expressed as the minimization problem:

$$\min_{\Theta_K, K\text{-sparse}} \sup_{\|\Sigma x\|_1 \leq 1} \|(\Theta - \Theta_K)x\|_\infty. \quad (4.43)$$

Its solution is given in closed-form by:

$$\Theta_K[\lambda, \mu] = \begin{cases} \Theta[\lambda, \mu] & \text{if } |\sigma_\mu \Theta[\lambda, \mu]| \text{ is among the } K \text{ largest values of } |\Theta\Sigma|, \\ 0 & \text{otherwise.} \end{cases} \quad (4.44)$$

The weights σ_μ must be adapted to the class of images to recover. In practice we found that setting $\sigma_\mu = 2^{-k}$ for $\mu = (k, e', n)$ is a good choice. These weights can also be trained from a set of images belonging the class of interest. Finally let us mention that

we also proposed greedy algorithms when setting $\|\cdot\|_Y = \|\cdot\|_2$ in [100]. In practice, it turns out that both approaches yield similar results.

We illustrate the importance of the thresholding strategy in section 4.6.1, Figure 4.7.

4.5 On the design of preconditioners

Let $P \in \mathbb{R}^{N \times N}$ denote a Symmetric Positive Definite (SPD) matrix. There are two equivalent ways to understand preconditioning: one is based on a change of variable, while the other is based on a metric change.

For the change of variable, let z be defined by $x = P^{1/2}z$. A solution x^* of problem (4.9) reads $x^* = P^{1/2}z^*$, where z^* is a solution of:

$$\min_{z \in \mathbb{R}^N} \frac{1}{2} \|\Theta_K P^{1/2} z - x_0\|_2^2 + \|P^{1/2} z\|_{1,w}. \quad (4.45)$$

The convergence rate of iterative methods applied to (4.45) is now driven by the properties of matrix $\Theta_K P^{1/2}$ instead of Θ_K . By choosing P adequately, one can expect to significantly accelerate convergence rates. For instance, if Θ_K is invertible and $P^{1/2} = \Theta_K^{-1}$, one iteration of a proximal gradient descent provides the exact solution of the problem.

For the metric change, the idea is to define a new scalar product defined by

$$\langle x, y \rangle_P = \langle Px, y \rangle, \quad (4.46)$$

and to consider the associated norm $\|x\|_P = \sqrt{\langle Px, x \rangle}$. By doing so, the gradient and proximal operators are modified, which leads to different dynamics. The preconditioned FISTA algorithm is given in Algorithm 6.

Algorithm 6: Preconditioned accelerated proximal gradient descent

Input: Initial guess $x^{(0)} = y^{(1)}$, $\tau = 1/\|\Theta_K^* \Theta_K P^{-1}\|_2$ and Nit
for $k = 1$ **to** Nit **do**
 Compute $\nabla F(y^{(k)}) = \Theta_K^* (\Theta_K y^{(k)} - u_0)$.
 $x^{(k)} = \text{Prox}_{\tau G}^P \left(y^{(k)} - \tau P^{-1} \nabla F(y^{(k)}) \right)$.
 $y^{(k+1)} = x^{(k)} + \frac{k-1}{k+2} (x^{(k)} - x^{(k-1)})$.
end

In this algorithm

$$\text{Prox}_{\tau G}^P(z_0) = \arg \min_{z \in \mathbb{R}^N} \frac{1}{2} \|z - z_0\|_P^2 + \tau G(z). \quad (4.47)$$

Unfortunately, it is impossible to provide a closed-form expression of (4.47), unless matrix P has a very simple structure (e.g. diagonal). Finding an efficient preconditioner therefore requires: i) defining a structure for P compatible with fast evaluations of the proximal operator (4.47) and ii) improving some “properties” of $\Theta_K P^{1/2}$ using this structure.

4.5.1 What governs convergence rates?

Good preconditioners are often heuristic. The following sentence is taken from a reference textbook about the resolution of linear systems by Y. Saad [202]: “Finding a good preconditioner to solve a given sparse linear system is often viewed as a combination of art and science. Theoretical results are rare and some methods work surprisingly well, often despite expectation.” In what follows, we will first show that existing convergence results are indeed of little help. We then provide two simple diagonal preconditioners.

Let us look at the convergence rate of Algorithm 4 applied to problem (4.45). The following theorem appears in [184, 18] for instance.

Theorem 4.5.1. *Let $A = P^{-1/2}\Theta_K^*\Theta_KP^{-1/2}$ and set $L = \lambda_{\max}(A^*A)$. The iterates in Algorithm 6 satisfy:*

$$E(x^{(k)}) - E(x^*) \leq L\|x - x_0\|_2^2 \cdot \min \left(\frac{1}{k^2}, \frac{1}{2} \left(\frac{\sqrt{\kappa(A)} - 1}{\sqrt{\kappa(A)} + 1} \right)^{2k} \right), \quad (4.48)$$

where $\kappa(A)$ designs the condition number of A :

$$\kappa(A) = \begin{cases} \sqrt{\frac{\lambda_{\max}(A^*A)}{\lambda_{\min}(A^*A)}} & \text{if } \lambda_{\min}(A^*A) > 0, \\ +\infty & \text{otherwise.} \end{cases} \quad (4.49)$$

When dealing with ill-posed inverse problems, the condition number $\kappa(A)$ is huge or infinite and bound (4.48) therefore reduces to

$$E(x^{(k)}) - E(x^*) \leq \frac{L\|x - x_0\|_2^2}{k^2}, \quad (4.50)$$

even for a very large number of iterations. Unfortunately, this bound tells very little about which properties of A characterize the convergence rate. Only the largest singular value of A seems to matter. The rest of the spectrum does not appear, while it obviously plays a key role.

Recently, more subtle results were proposed in [224] for the specific $\ell^1 - \ell^2$ problem and in [155] for a broad class of problems. These theoretical results were shown to fit some experiments very well. Let us state a typical result.

Theorem 4.5.2. *Assume that problem (4.45) admits a unique minimizer x^* . Let $S^* = \text{supp}(x^*)$ denote the solution’s support. Then Algorithm 6 generates a sequence $(x^{(k)})_{k \in \mathbb{N}}$ satisfying the following properties:*

- The sequence $(x^{(k)})_{k \in \mathbb{N}}$ converges to x^* .
- There exists an iteration number k^* such that, for $k \geq k^*$, $\text{supp}(x^{(k)}) = \text{supp}(x^*)$.
- If in addition

$$\langle Ax, Ax \rangle \geq \alpha\|x\|_2^2, \quad \forall x \text{ s.t. } \text{supp}(x) \subseteq S^*, \quad (4.51)$$

then the sequence of iterates $(x^{(k)})_{k \in \mathbb{N}}$ converges linearly to x^* : there exists $0 \leq \rho < 1$ s.t.

$$\|x^{(k)} - x^*\|_2 = O(\rho^k). \quad (4.52)$$

Remark 4.5.1. The uniqueness of a solution x^* is not required if the algorithm converges. This can be ensured if the algorithm is slightly modified [53].

The main consequence of Theorem 4.5.2 is that good preconditioners should depend on the support S^* of the solution. Obviously, this support is unknown at the start of the algorithm. Moreover, for compact operators, the value α in equation (4.51) is null or very small. Therefore - once again - Theorem 4.5.2 seems to be of little help to find well founded preconditioners.

In this paper we therefore restrict our attention to two standard heuristic preconditioners: Jacobi and Sparse Approximate Inverses (SPAI) [122, 202]. The overall idea is to cluster the eigenvalues of A^*A .

4.5.2 Jacobi preconditioner

The Jacobi preconditioner is one of the most popular diagonal preconditioner, it consists of setting

$$P = \max(\text{diag}(\Theta_K^* \Theta_K), \epsilon), \quad (4.53)$$

where ϵ is a small constant and the maximum should be understood component-wise. The parameter ϵ guarantees the invertibility of P .

The idea of this preconditioner is to make the Hessian matrix $P^{-1/2} \Theta_K^* \Theta_K P^{-1/2}$ “close” to the identity. This preconditioner has a simple analytic expression and is known to perform well for diagonally dominant matrices. Blurring matrices expressed in the wavelet domain have a fast decay away from the diagonal, but are usually not diagonally dominant. Moreover, the parameter ϵ has to be tuned. In our numerical experiments, this parameter was hand-tuned so as to ensure the fastest convergence.

4.5.3 SPAI preconditioner

The preconditioned gradient in Algorithm 6 involves matrix $P^{-1} \Theta_K^* \Theta_K$. The idea of sparse approximate inverses is to cluster the eigenvalues of $P^{-1} \Theta_K^* \Theta_K$ around 1. To improve the clustering, a possibility is to solve the following optimization problem:

$$\arg \min_{P, \text{ diagonal}} \left\| \text{Id} - P^{-1} \Theta_K^* \Theta_K \right\|_F^2, \quad (4.54)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. This formulation is standard in the numerical analysis community [122, 202].

Lemma 4.5.1. *Let $M = \Theta_K^* \Theta_K$. The set of solutions of (4.54) reads:*

$$P[i, i] = \begin{cases} \frac{M^2[i, i]}{M[i, i]} & \text{if } M[i, i] \neq 0, \\ \text{an arbitrary positive value otherwise.} \end{cases} \quad (4.55)$$

Proof. First notice that problem (4.54) can be rewritten as

$$\arg \min_{P, \text{ diagonal}} \left\| \text{Id} - MP^{-1} \right\|_F^2, \quad (4.56)$$

by taking the transpose of the matrices, since M is symmetric and P diagonal. The Karush-Kuhn-Tucker optimality conditions for problem (4.56) yield the existence of a Lagrange multiplier $\mu \in \mathbb{R}^{N \times N}$ such that:

$$M(MP^{-1} - I) + \mu = 0_{\mathbb{R}^N}, \quad (4.57)$$

with

$$\mu[i, j] = \begin{cases} 0 & \text{if } i = j, \\ \text{an arbitrary value} & \text{otherwise.} \end{cases} \quad (4.58)$$

Therefore, for all i ,

$$(M^2 P^{-1})[i, i] = M[i, i], \quad (4.59)$$

which can be rewritten as

$$M^2[i, i]P^{-1}[i, i] = M[i, i], \quad (4.60)$$

since P is diagonal. If $M^2[i, i] = 0$, then $M[i, i] = 0$ since $M^2[i, i]$ is the squared norm of the i -th column of M . In that case, $P^{-1}[i, i]$ can take an arbitrary value. Otherwise $P^{-1}[i, i] = M[i, i]/M^2[i, i]$, finishing the proof. $\square$

4.5.4 GPU implementation

Algorithm 6 has a structure that is near ideal for parallelization on GPU. It consists in iteratively:

- Evaluating matrix-vector products with sparse matrices. This was achieved with the cuSPARSE library from CUDA.
- Evaluating the action of shrinkage operators. This is a simple operation for a GPU since it suffices to unroll a *for* loop.

All computations are done on the GPU: the only exchanges with the CPU are i) loading the wavelet coefficients and the sparse matrix at the beginning of the iterative process and ii) returning the restored wavelet coefficients at the end. The reported computing times do not include the transfer time.

4.6 Numerical experiments

In this section we propose a set of numerical experiments to illustrate the proposed methodology and to compare its efficiency with respect to state-of-the-art approaches. In all experiments, we assume that the blur kernel and the noise variance are known (non blind deblurring).

The numerical experiments are performed on two 1024×1024 images with values rescaled in $[0, 1]$, see Figure 4.3. We also consider two different blurs, see Figure 4.4. The PSF in Figure 4.4a is an anisotropic 2D Gaussian with kernel defined for all $(t_1, t_2) \in [0, 1]^2$ by

$$k(t_1, t_2) = \begin{cases} \exp\left(-\frac{t_1^2}{2\sigma^2} - \frac{t_2^2}{2\sigma^2}\right) & \text{if } t_1 \geq 0, \\ \exp\left(-\frac{4t_1^2}{2\sigma^2} - \frac{t_2^2}{2\sigma^2}\right) & \text{otherwise,} \end{cases}$$

with $\sigma = 5$. This PSF is smooth, which is a favorable situation for our method, see Theorem 4.4.2. The PSF in Figure 4.4b is a simulation of motion blur. This PSF is probably one of the worst for the proposed technique since it is singular. The PSF is generated from a set of $l = 5$ points drawn at random from a Gaussian distribution with standard deviation $\sigma_1 = 8$. Then the points are joined using a cubic spline and the resulting curve is blurred using a Gaussian kernel of standard deviation $\sigma_2 = 1$.

All our numerical experiments are based on Symmlet 6 wavelets decomposed $J = 6$ times. This choice offers a good compromise between computing times and visual quality of the results. The weights w in function G in (4.4) were defined by $w[i] = j(i)$, where $j(i)$ denotes the scale of the i -th wavelet coefficient. This choice was hand tuned so as to produce the best deblurring results. Figures 4.5 and 4.6 display two typical deconvolution results using these parameters.

The numerical experiments were performed on Matlab2014b on an Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz with 200Gb RAM in double precision. Automatic multi-threading was disabled by launching Matlab with:

```
>> matlab -singleCompThread
```

For the experiments led on GPU, we use a NVIDIA Tesla K20c containing 2496 CUDA cores and 5GB internal memory. All computations were performed in double precision.

4.6.1 On the role of thresholding strategies

We first illustrate the influence of the thresholding strategy discussed in Section 4.4.7. We construct two matrices having the same number of coefficients but built using two different thresholding rules: the naive thresholding given in equation (4.42) and the weighted thresholding given in equation (4.44). Figure 4.7 displays the images restored with each of these two matrices. It is clear that the weighted thresholding strategy significantly outperforms the simple one: it produces less artifacts and a higher pSNR. In all the following experiments, this thresholding scheme will be used.

4.6.2 Approximation in wavelet bases

In this paragraph, we illustrate the influence of the approximation on the deblurring quality. We compare the solution of the original problem (4.2) with the solution of the approximated problem (4.9) for different numbers of coefficients K . Computing the exact gradient $\nabla F = \Psi^* H^* H \Psi$ requires two fast Fourier transforms, two fast wavelet transforms and a multiplication by a diagonal matrix. Its complexity is therefore:




Figure 4.3 – Original images 1024×1024 .




Figure 4.4 – PSFs used in the paper. (4.4a) is a skewed Gaussian. (4.4b) is a motion blur.

$2N \log_2(N) + 2lN + N$, with l denoting the wavelet filter size. The number of operations per pixel is therefore $2 \log_2(N) + 2l + 1$. The approximate gradient ∇F_K requires two matrix-vector products with a K -sparse matrix. Its complexity is therefore $2 \frac{K}{N}$ operations per pixel. Figure (4.8) displays the restoration quality with respect to the number of operations per pixel.

For the smooth PSF in Figure 4.4a, the standard approach requires 89 operations per pixel, while the wavelet method requires 20 operations per pixel to obtain the same pSNR. This represents an acceleration of a factor 4.5. For users ready to accept a decrease of pSNR of 0.2dB, K can be chosen even significantly lower, leading to an acceleration factor of 40 and around 2.2 operations per pixels! For the less regular PSF 4.4b, the gain is less important. To obtain a similar pSNR, the proposed approach is in fact slower with 138 operations per pixel instead of 89 for the standard approach. However, accepting a decrease of pSNR of 0.2dB, our method leads to an acceleration by a factor 1.1. To summarize, the proposed approximation does not really lead to interesting acceleration factors for motion blurs. Note however that the preconditioners can be used even if the operator is not expanded in the blur domain.

The different behavior between the two blurs was predicted by Theorem 4.4.2, since the compressibility of operators in wavelet bases strongly depends on the regularity M of the PSF.

4.6.3 A competitor: the ADMM

FISTA is probably one of the most popular approaches to solve the ℓ^1 problem (4.2). Another very effective and popular approach is the alternating direction of multipliers (ADMM) [115, 187]. We therefore compare the proposed approach to the ADMM too.

The problem we wish to solve reads:

$$\min_{u \in \mathbb{R}^N} \frac{1}{2} \|Hu - u_0\|_2^2 + \|\Psi^T u\|_{1,w}. \quad (4.61)$$




Figure 4.5 – A deconvolution example. The book image in Figure 4.3 is blurred with the motion blur Figure 4.4b and degraded with a noise level of $5 \cdot 10^{-3}$. The pSNR of the degraded image (on top) is 17.85dB. Problem (4.2) is solved using the exact operator, $\lambda = 10^{-4}$, 500 iterations and Symmlet 6 wavelets decomposed 6 times. The pSNR of the restored image is 24.14dB.




Figure 4.6 – A deconvolution example. The confocal image Figure 4.3 has been blurred with blur Figure 4.4a and degraded with a noise level of $5 \cdot 10^{-3}$. The pSNR of the degraded image (on top) is 23.94dB. Problem (4.2) is solved using the exact operator, $\lambda = 10^{-4}$, 500 iterations and Symmlet 6 wavelets decomposed 6 times. The pSNR of the restored image is 26.33dB.




Figure 4.7 – A deconvolution example showing the importance of the thresholding strategy. The book image on Figure 4.3 is blurred with the kernel in Figure 4.4b and degraded with a noise level of $5 \cdot 10^{-3}$ (see Figure 4.5). Matrices have been constructed with the same number of coefficients that corresponds to 57 operations per pixel. Top: the result for the simple thresholding strategy, $\text{pSNR} = 23.71\text{dB}$. Bottom: the weighted strategy $\text{pSNR} = 24.07\text{dB}$.




Figure 4.8 – Evolution of the pSNR of the deconvolved image w.r.t. the number of operations per pixel per iteration. The gray vertical line gives the number of operations per pixel per iteration to solve the exact ℓ^1 -problem 4.2 using FFTs and FWTs. The horizontal line gives the pSNR obtained using the exact operator.

The ADMM can be applied as follows. First use the following splitting:

$$\min_{u \in \mathbb{R}^N, x \in \mathbb{R}^N, x = \Psi^T u} \frac{1}{2} \|Hu - u_0\|_2^2 + \|x\|_{1,w}. \quad (4.62)$$

The augmented Lagrangian associated to problem (4.62) is:

$$\mathcal{L}(x, u, \lambda) = \frac{1}{2} \|Hu - u_0\|_2^2 + \|x\|_{1,w} + \langle \lambda, x - \Psi^T u \rangle + \frac{\beta}{2} \|x - \Psi^T u\|_2^2. \quad (4.63)$$

The ADMM then takes the following algorithmic form:

Algorithm 7: ADMM to solve (4.2)

Input: Initial guesses $x^{(0)}$, $u^{(0)}$, $\lambda^{(0)}$ and Nit
for $k = 1$ **to** Nit **do**
 Solve $x^{(k+1)} = \arg \min_{x \in \mathbb{R}^N} \mathcal{L}(x, u^{(k)}, \lambda^{(k)})$. Solve
 $u^{(k+1)} = \arg \min_{u \in \mathbb{R}^N} \mathcal{L}(x^{(k+1)}, u, \lambda^{(k)})$. $\lambda^{(k+1)} = \lambda^{(k)} + \beta (\Psi^T u^{(k+1)} - x^{(k+1)})$.
end

The first step is a soft-thresholding operation on x . The second step consists of solving a linear system of type:

$$(H^T H + \beta I)u = c, \quad (4.64)$$

with $c = H^T u_0 - \Psi \lambda + \beta \Psi x$. Solving this linear system can be achieved in $O(N \log(N))$ operations for convolution operators, since $(H^T H + \beta I)$ can be diagonalized using the

discrete Fourier transform. This is no longer the case for spatially varying blurs. In that case, iterative methods such as a conjugate gradient descent can be used.

Algorithm 7 can be implemented on the GPU. The linear system is solved with the parallel version of the FFT called cuFFT. We used the open-source parallel implementation of the wavelet transform available here <https://github.com/pierrelepaleo/PDWT>.

4.6.4 Comparing preconditioners and ADMM

We now illustrate the interest of using the preconditioners described in Section 4.5. We compare the cost function w.r.t. the iterations number for different methods: ISTA, FISTA, FISTA with a Jacobi preconditioner (see (4.53)), FISTA with a SPAI preconditioner (see (4.55)) and the ADMM described in paragraph 4.6.3. For the Jacobi preconditioner, we optimized ϵ by trial and error in order to get a fixed accuracy in a minimum number of iterations.

As can be seen in Figure 4.9, the Jacobi and SPAI preconditioners allow reducing the iterations number significantly. We observed that the SPAI preconditioner outperformed the Jacobi preconditioner for all blurs we tested and thus recommend SPAI in general. From a practical point view, a speed-up of a factor 3 is obtained for both blurs.

As can be seen, ADMM is significantly faster than preconditioned FISTA for both blurs, at least in the first iterations. The number of iterations necessary to reach a relative accuracy of 10^{-3} is roughly divided by 4. Notice that ADMM and FISTA handle the quadratic term $\frac{1}{2}\|Hu - u_0\|_2^2$ in a very different way: the minimization is basically performed exactly with the ADMM, while it is solved iteratively in FISTA algorithm. This being said, the proposed approach is still significantly faster for smooth kernels. For the skewed Gaussian blur, the proposed algorithm is 4 times faster, since each iteration of the proposed FISTA approach has a significantly lower cost. For the motion blur, the ADMM is twice faster since the number of non-zero components in Θ_K is much higher.


4.6.5 Computing times

In this paragraph, we time precisely what can be gained using the proposed approach on the two examples of Figure 4.6 and 4.5. The proposed approach consists of:


- Finding a number K such that deconvolving the image with matrix Θ_K instead of Θ leads to a decrease of pSNR of less than 0.2dB.
- For each optimization method, finding a number of iterations Nit leading to a precision

$$E(x^{(Nit)}) - E(x^*) \leq 10^{-3}E(x^{(0)}). \quad (4.65)$$


In all experiments, matrix Θ_K is computed off-line, meaning that we assume it is known beforehand. In all the paper, the “speed-up” is defined as the ratio between the time needed by the standard FISTA algorithm to reach a given accuracy, divided by the time needed by the proposed algorithm to reach the same accuracy.




(a) Skewed Gaussian



(b) Skewed Gaussian



(c) Motion blur



(d) Motion blur

Figure 4.9 – Cost function with respect to iterations and time for different preconditioners.

	Exact	GPU	FISTA	GPU	Jacobi	GPU	SPAI	GPU	ADMM	GPU
Iterations number		117		127		55		43		16
Time (in seconds)	24.30	2.47	2.57	0.43	1.31	0.19	1.16	0.15	4.22	0.59

Table 4.1 – Timing and iterations number depending on the method. The number of operations per pixel is 2.46. This experiment corresponds to the Skewed Gaussian blur in Figure 4.6.

	Exact	GPU	FISTA	GPU	Jacobi	GPU	SPAI	GPU	ADMM	GPU
Iterations number		107		107		52		36		8
Time (in seconds)	20.03	1.67	16.7	1.82	7.48	0.89	6.54	0.62	1.95	0.28

Table 4.2 – Timing and iterations number depending on the method. The number of operations per pixel is 39.7. This experiment corresponds to the motion blur in Figure 4.5.

The results are displayed in Table 4.1 for the skewed Gaussian blur and in Table 4.2 for the motion blur. For the skewed Gaussian, the total speed-up is roughly 162, which can be decomposed as: sparsification = 7.8, preconditioning = 2.7, GPU = 7.7. For the motion blur, the total speed-up is roughly 32, which can be decomposed as: sparsification = 1.01, preconditioning = 3, GPU = 10.5.

As can be seen from this example, the proposed sparsification may accelerate computations significantly for smooth enough blurs. On these two examples, the preconditioning led to an acceleration of a factor 3. Finally, GPU programming allows accelerations of a factor 7-8, which is on par with what is usually reported in the literature.

Note that for the smooth blurs encountered in microscopy, the total computing time is 0.17 seconds for a 1024×1024 image, which can be considered as real-time.

4.6.6 Dependency on the blur kernel

In this paragraph, we analyze the method behavior with respect to different blur kernels. We consider 5 different types of kernels commonly encountered in applications: Gaussian blur, skewed Gaussian blur, motion blur, Airy pattern and defocus blur. For each type, we consider two different widths ($\sigma = 2.5$ and $\sigma = 5$). The blurs are shown in Figure 4.10. Table 4.3 summarizes the acceleration provided by using simultaneously the sparse wavelet approximation, SPAI preconditioner and GPU programming. We used the same protocol as Section 4.6.5. The acceleration varies from 218 (large Airy pattern) to 19 (large motion blur). As expected, the speed-up strongly depends on the kernel smoothness. Of interest, let us mention that the blurs encountered in applications such as astronomy or microscopy (Airy, Gaussian, defocus) all benefit greatly from the proposed approach. The acceleration factor for the least smooth blur, corresponding to the motion blur, still leads to a significant acceleration, showing that the proposed methodology can be used in nearly all types of deblurring applications.




Figure 4.10 – The different blurs used to analyze the method’s efficiency.

Blur	Time (Fourier)	Time (Proposed)	Speed-up	# ops per pixels
Gaussian (small)	14.8	0.15	99	4
Gaussian (large)	17.8	0.14	127	2
Skewed (small)	11.2	0.11	100	10
Skewed (large)	10.5	0.1	102	4
Motion (small)	6.0	0.26	23	80
Motion (large)	9.7	0.51	19	80
Airy (small)	15.16	0.081	187	4
Airy (large)	18.6	0.085	218	2
Defocus (small)	20.2	0.23	87	20
Defocus (large)	21.89	0.20	110	10

Table 4.3 – Speed-up of ℓ^1 - ℓ^2 deconvolution with respect to the different blur kernels, see Figure 4.10.

Resolution	512	1024	2048	4096
Time (Fourier)	3.19	17.19	76	352
Time Wavelet + GPU + SPAI	0.07	0.25	0.55	1.35
Total Speed-up	44	70	141	260
Speed-up sparse	4.1	4.5	9.6	9.7
Speed-up SPAI	2.4	2.2	2.1	2.7
Speed-up GPU	4.5	7.1	7.0	10.0

Table 4.4 – Speed-up of ℓ^1 - ℓ^2 deconvolution with respect to the image resolution.

4.6.7 Dependency on resolution

In this paragraph, we aim at illustrating that the method efficiency increases with resolution. To this end, we deconvolve the phantom in [123] with resolutions ranging from 512×512 to 4096×4096 . The convolution kernel is a Gaussian. Its standard deviation is chosen as $\sigma = 2^L/200$, where 2^L is the number of pixels in each direction. This choice is the natural scaling that ensures resolution invariance. We then reproduce the experiment of the previous section to evaluate the speed-up for each resolution. The results are displayed in Table 4.4. As can be seen, the speed-up increases significantly with the resolution, which could be expected, since as resolution increases, the kernel’s smoothness increases. Of interest, note that 1.35 seconds is enough to restore a 4096×4096 image.

4.6.8 Spatially varying blurs

Finally, we show that the proposed ideas can be successfully applied to spatially varying blurs. We work on 512×512 images. The confocal image -Figure 4.3 - has been degraded with a spatially varying blur presented in Figure 4.11. This blur models realistic degradations appearing in optical systems [212]. The matrix H associated to this blur contains about $537N$ non-zero coefficients. The aim of this section is to compare the performance of the following methods:

- Problem (4.6) solved using FISTA algorithm and the exact operator implemented with a sparse matrix, see Algorithm 4.
- Problem (4.6) solved using ADMM algorithm and the exact operator implemented with a sparse matrix. The linear system is solved with a conjugate gradient descent initialized with an initial guess corresponding to the solution at the previous iteration and a stopping criterion optimized so as to get the fastest convergence.
- Problem (4.45) solved using FISTA algorithm, with an operator approximated in the wavelet domain and a Jacobi preconditioner, i.e. Algorithm 6.

Similarly as the previous section, the number of coefficients K is chosen such that deblurring the image with matrix Θ_K instead of Θ leads to a decrease of pSNR of less




Figure 4.11 – An illustration of the spatially varying blur used in the numerical experiments. The blurring operator is applied to a Dirac comb to obtain the PSFs at various locations. This blur models realistic degradations appearing in optical systems [212].

	GPU	GPU	GPU
	FISTA Spatial	FISTA Wavelet	ADMM
Iterations number	110	44	6
Time (in seconds)	47.8 / 9.54	0.73 / 0.13	107.3 / 21.5

Table 4.5 – Timing and iterations number depending on the method.

than 0.2dB. To reach this we had to set $K = 15N$ which is about 36 times smaller than the number of coefficients in H .

The deblurred image is displayed in Figure 4.12. The performance of each algorithm is provided on Figure 4.13. The ADMM algorithm has a faster convergence than FISTA. However each iteration involves the resolution of a linear system with a conjugate gradient descent which considerably increases computing times. FISTA algorithm implemented in the spatial domain involves matrix-vector products with H , which - again - leads to slow computing times.




Figure 4.12 – A restoration of images degraded by spatially varying blurs. The confocal image - 512×512 - in Figure 4.3 has been blurred with the operator in Figure 4.11 and degraded with a noise level of $5 \cdot 10^{-3}$. The pSNR of the degraded image (on top) is 22.3dB. Problem (4.2) is solved using the exact operator, $\lambda = 10^{-4}$, 500 iterations and Symmlet 6 wavelets decomposed 6 times. The pSNR of the restored image is 27.29dB.




Figure 4.13 – Illustration of the performance of the three deblurring methods. The cost functions are displayed with respect to the number of iterations on the left and the time on the right.

Part II

Estimation of operators

5

Identification of linear operators from scattered impulse responses

Résumé : Ce chapitre est issu de [98], un article en préparation et en collaboration avec Jérémie Bigot et Pierre Weiss. Préalablement à la restauration d'images dégradées par des flous variables, les opérateurs de flou doivent être estimés. Nous considérons donc le problème de l'estimation d'un opérateur intégral régulier à partir de la connaissance de n réponses impulsionnelles éparpillées dans le domaine. Nous proposons un estimateur $\hat{H}$ de l'opérateur H calculable en très grande dimension. La méthode repose sur une interpolation ligne par ligne des coefficients en ondelettes de la Time Varying Impulse Response (TVIR). L'estimateur a une interprétation spatiale sous forme variationnelle. Nous avons obtenu la vitesse d'estimation $\mathbb{E}\|H - \hat{H}\|_{2 \rightarrow 2}^2 \lesssim N^{-r/d} + N^{\frac{2s}{2s+d}} n^{-\frac{2s}{2s+d}}$ où N est la dimension de la discrétisation, r et s sont des scalaires décrivant la régularité des réponses impulsionnelles et de leurs variations. Cette méthode est illustrée numériquement sur des problèmes d'interpolation d'opérateurs de flou.

Abstract: Restoring images degraded by spatially varying blurs often requires the estimation of an operator that is partially unknown. We consider the problem of estimating smooth integral operators from n scattered impulse responses. We propose an estimator $\hat{H}$ of the operator H with a method that is numerically tractable in high-dimension. The approach is based on a row by row interpolation of the Time Varying Impulse Response (TVIR). The provided estimator has a variational interpretation in the spatial domain. We also propose a theoretical analysis of the performance of the estimator. We found a convergence rate $\mathbb{E}\|H - \hat{H}\|_{2 \rightarrow 2}^2 \lesssim N^{-r/d} + N^{\frac{2s}{2s+d}} n^{-\frac{2s}{2s+d}}$ where N is discretization dimension, r and s are scalars describing the regularity of the impulse responses and their variations. This method is numerically illustrated with blurring operators.

Contents

5.1	Introduction	159
5.2	Problem setting	161
5.2.1	Space varying impulse response regularity	161
5.2.2	Smooth variations of the impulse responses	162
5.2.3	Sampling model	163
5.3	Main results	164
5.3.1	Construction of an estimator	164
5.3.2	Mixed-Sobolev space interpretation	165
5.3.3	Numerical complexity	165
5.3.4	Convergence rates	167
5.3.5	Extensions	168
5.4	Radial basis functions implementation	168
5.4.1	Standard approximation results in RKHS	168
5.4.2	Application to our problem	170
5.5	Proofs of the main results	170
5.5.1	Operator norm risk	172
5.5.2	Discretization error ϵ_d	172
5.5.3	Estimation error ϵ_e	173
5.5.4	Proof of the main results	179

5.1 Introduction

Let Ω denote a subset of $\mathbb{R}^d$ and $H : L^2(\Omega) \rightarrow L^2(\Omega)$ denote a linear integral operator defined for all $u \in L^2(\Omega)$ by:

$$Hu(x) = \int_{\Omega} K(x, y)u(y)dy, \quad (5.1)$$

where $K : \Omega \times \Omega \rightarrow \mathbb{R}$, is the operator's kernel. Given a set of functions $(u_i)_{1 \leq i \leq n}$, the problem of operator identification consists of recovering H from the knowledge of $f_i = Hu_i + \epsilon_i$, where ϵ_i is an unknown perturbation.

This problem arises in many fields of science and engineering such as mobile communication [143], imaging [113] and geophysics [20]. Many different reconstruction approaches have been developed, depending on the operator's regularity and the set of test functions (u_i) . Assuming that H has a bandlimited Kohn-Nirenberg symbol and that its action on a dirac comb is known, a few authors proposed extensions of Shannon's sampling theorem [143, 149, 191]. Another recent trend is to assume that H can be decomposed as a linear combination of a small number of elementary operators. When the operators are known, recovering H amounts to solving a linear system. The work [62] analyzes the conditioning of this linear system when H is a matrix applied to a random Gaussian vector. When the operator can be sparsely represented in a dictionary of elementary matrices, compressed sensing theories can be developed [192]. Finally, in astrophysics, a few authors considered interpolating the coefficients of a few known impulse responses (also called Point Spread Functions, PSF) in a well chosen basis [113, 167, 56]. This strategy corresponds to assuming that $u_i = \delta_{y_i}$ and is often used when the PSFs are compactly supported and have smooth variations. Notice that in this setting, each PSFs is known *independently* of the others, contrarily to the work [191].

This last approach is particularly effective in large scale imaging applications due to two useful facts. First, representing the impulse responses in a small dimensional basis allows reducing the number of parameters to identify. Second, there now exists efficient interpolation schemes based on radial basis functions. Despite its wide use and empirical success, this method still lacks of solid mathematical foundations and many practical questions remain open:

- Under what hypotheses on the operator H can this method be applied?
- What is the influence of the geometry of the set $(y_i)_{1 \leq i \leq n}$?
- Is the reconstruction stable to the perturbations $(\epsilon_i)_{1 \leq i \leq n}$? If not, how to make robust reconstructions, tractable in very large scale problems?
- What theoretical guarantees can be provided in this challenging setting?

The objective of this work is to provide answers to the above mentioned questions. We design a robust algorithm applicable in large scale applications. It yields a finite dimensional operator estimator of H allowing for fast matrix-vector products, which are

essential for further processing. The theoretical convergence rate of the estimator as the number of observations increases is studied thoroughly. An illustration of the problem and the output of the proposed algorithm is provided in Figure 5.1.




Figure 5.1 – Top: Reconstruction of a spatially varying blur operator. (a): The exact operator applied to a 2D dirac comb. (b): 64 impulse responses corrupted by additive Gaussian noise. (c): reconstructed operator. Bottom: a deconvolution experiment. (d): original image. (e): blurry and noisy image. (f): deblurred image using the operator reconstructed in Figure (c). The operator’s impulse responses are Gaussians with covariance matrices $\Sigma(y_1, y_2) = \text{diag}(\sigma^2(y_1, y_2), \sigma^2(y_1, y_2))$ where $\sigma(y_1, y_2) = 1 + 2 \max(1 - y_1, y_1)$ for $(y_1, y_2) \in [0, 1]^2$.

The outline of this paper is as follows. We first specify the problem setting precisely in Section 5.2. We then describe the main outcomes of our study in Section 5.3. We provide a detailed explanation of the numerical algorithm in Section 5.4. Finally, the proofs of the main results are given in Section 5.5.

5.2 Problem setting

In this section, we precisely describe the problem setting. We assume that $\Omega \subset \mathbb{R}^d$ is a bounded, open, and connected set, with Lipschitz continuous boundary. The value of a function f at x is denoted $f(x)$, while the i -th value of a vector $v \in \mathbb{R}^N$ is denoted $v[i]$. The (i, j) -th element of a matrix A is denoted $A[i, j]$. The Sobolev space $H^s(\Omega)$ is defined for $s \in \mathbb{N}$ by

$$H^s(\Omega) = \left\{ u \in L^2(\Omega), \partial^\alpha u \in L^2(\Omega), \text{ for all multi-index } \alpha \in \mathbb{N}^d \text{ s.t. } |\alpha| = \sum_{i=1}^d \alpha[i] \leq s \right\}. \quad (5.2)$$

The space $H^s(\Omega)$ can be endowed with a norm $\|u\|_{H^s(\Omega)} = \left(\sum_{|\alpha| \leq s} \|\partial^\alpha u\|_{L^2(\Omega)}^2 \right)^{1/2}$ and the semi-norm $|u|_{H^s(\Omega)} = \left(\sum_{|\alpha|=s} \|\partial^\alpha u\|_{L^2(\Omega)}^2 \right)^{1/2}$. In addition, we will use the equivalent Beppo-Levi semi-norm defined by $|u|_{BL^s(\Omega)}^2 = \sum_{|\alpha|=s} \frac{s!}{\alpha_1! \alpha_2! \dots \alpha_d!} \|\partial^\alpha u\|_{L^2(\Omega)}^2$ and the Beppo-Levi semi-inner product defined by

$$\langle f, g \rangle_{BL^s(\Omega)} = \sum_{|\alpha|=s} \frac{s!}{\alpha_1! \alpha_2! \dots \alpha_d!} \langle \partial^\alpha f, \partial^\alpha g \rangle_{L^2(\Omega)}. \quad (5.3)$$

5.2.1 Space varying impulse response regularity

An integral operator can be represented in many different ways. A key representation in this paper is the Space Varying Impulse Response (SVIR) $S : \Omega \times \Omega \rightarrow \mathbb{R}$ defined by:

$$S(x, y) = K(x + y, y). \quad (5.4)$$

The impulse response or Point Spread Function (PSF) at location $y \in \Omega$ is defined by $S(\cdot, y)$. The SVIR encodes the impulse response variations in the y direction, instead of the $(x - y)$ direction for the kernel representation, see Figure 5.2 for a 1D example. It is convenient, since in many applications, the smoothness of S in the x and y directions is driven by completely different physical phenomena. For instance, in astrophysics, the regularity of $S(\cdot, y)$ depends on the optical system, while the regularity of $S(x, \cdot)$ may depend on exterior factors such as a weak gravitational lensing [56]. This property will be expressed through specific regularity assumptions of S defined hereafter.

Let $(\phi_k)_{k \in \mathbb{N}}$ denote a Hilbert basis of $L^2(\Omega)$ and $\mathcal{E}^r(\Omega)$ denote the following Hilbert space.

Definition 5.2.1. The space $\mathcal{E}^r(\Omega)$ is defined, for all $r \in \mathbb{R}$ and $r > \frac{d}{2}$, as the set of functions $u \in L^2(\Omega)$ such that:

$$\|u\|_{\mathcal{E}^r(\Omega)}^2 = \sum_{k \in \mathbb{N}} w[k] |\langle u, \phi_k \rangle|^2 < +\infty, \quad (5.5)$$

where $w : \mathbb{N} \rightarrow \mathbb{R}_+^*$ is a weight function satisfying $w[k] \gtrsim k^{2r/d}$.




Figure 5.2 – Two representations of an integral operator H acting on $L^2([0, 1])$. The kernel is defined by $K(x, y) = \frac{1}{\sqrt{2\pi\sigma(y)}} \exp\left(-\frac{1}{2\sigma(y)^2}|x - y|^2\right)$, where $\sigma(y) = 0.05(1 + 2\min(y, 1 - y))$. Left: kernel representation (see equation (5.1)). Right: SVIR representation (see equation (5.4)).

This definition is introduced in reference to the Sobolev spaces $H^m(\Omega)$ which can be defined - alternatively to equation (5.2) - by:

$$H^m(\Omega) = \left\{ u \in L^2(\Omega), \sum_{k \in \mathbb{N}} (1 + k^2)^{m/d} |\langle u, \phi_k \rangle|^2 < +\infty \right\}, \quad (5.6)$$

where the basis (ϕ_k) is either the Fourier basis or a wavelet basis with at least $m + 1$ vanishing moments (see e.g. [164, Chapter 9]). Definition 5.2.1 encompasses many other spaces. For instance, it allows choosing a basis $(\phi_k)_{k \in \mathbb{N}}$ that is best adapted to the impulse responses at hand, by using principal component analysis, as was proposed in a few applied papers [141, 26].

Assumption 5.2.1 (Impulse response regularity). *In all the paper, we will assume the following regularity condition.*

$$\sup_{y \in \Omega} \|S(\cdot, y)\|_{\mathcal{E}^r(\Omega)} dy < +\infty. \quad (5.7)$$

When $(\phi_k)_{k \in \mathbb{N}}$ is a Fourier or a wavelet basis, condition (5.7) simplifies to $\partial_x^\alpha S(\cdot, y) \in L^2(\Omega)$ for all $y \in \Omega$ and for all multi-index $\alpha \in \mathbb{N}^d$ such that $|\alpha| \leq r$.

5.2.2 Smooth variations of the impulse responses

In addition to the impulse responses regularity Assumption 5.2.1, we need to state a regularity condition for the impulse responses variations. In order to use fine approximation results based on radial basis functions [6], we will use the following regularity condition.

Assumption 5.2.2 (Smooth variations). *Throughout the paper, we assume that*

$$\partial_y^\beta S(x, \cdot) \in L^2(\Omega), \quad \forall x \in \Omega \text{ and for all multi-index } \beta \text{ s.t. } |\beta| \leq s. \quad (5.8)$$

In the particular case $\mathcal{E}^r(\Omega) = H^r(\Omega)$, the two assumptions 5.2.2 and 5.2.1 implies that S belongs to the *mixed-Sobolev space* $H^{r,s}(\Omega \times \Omega)$ consisting of functions with $\partial_x^\alpha S \in L^2(\Omega \times \Omega)$ for all multi-index $|\alpha| \leq r$ and $\partial_y^\beta S \in L^2(\Omega \times \Omega)$ for all multi-index $|\beta| \leq s$.

5.2.3 Sampling model

The main purpose of this paper is the reconstruction of the SVIR of an operator from the observation of a few impulse responses $S(\cdot, y_i)$ at scattered (but known) locations $(y_i)_{1 \leq i \leq n}$ in Ω . In applications, the PSFs $S(\cdot, y_i)$ can only be observed through a projection onto an N dimensional linear subspace V^N . In this paper, we assume that the linear subspace V_N reads

$$V_N = \text{span}(\phi_k, 1 \leq k \leq N). \quad (5.9)$$

In addition, the data is often corrupted by noise and we therefore observe a set of N dimensional vectors $(F_i^\epsilon)_{1 \leq i \leq n}$ defined for all $k \in \{1, \dots, N\}$ by

$$F_i^\epsilon[k] = \langle S(\cdot, y_i), \phi_k \rangle + \epsilon_i[k], \quad 1 \leq i \leq n, \quad (5.10)$$

where ϵ_i is a random vector with independent and identically distributed (iid) components with zero mean and variance σ^2 .

The assumption that V_N is defined using basis (ϕ_k) simplifies the analysis, since the representation and observation bases coincide. It would be interesting for applications to consider cases where V_N is defined using another Hilbert basis, but we would then need to use the theory of generalized sampling, which is significantly more involved (see e.g. [1]). We therefore leave this question aside in this paper.

Finally, we will show that the approximation efficiency of our method depends on the geometry of the set of data locations, and - in particular - on the fill and separation distances defined below.

Definition 5.2.2 (Fill distance). The fill distance of $Y = \{y_1, \dots, y_n\} \subset \Omega$ is defined as:

$$h_{Y,\Omega} = \sup_{y \in \Omega} \min_{1 \leq j \leq n} \|y - y_j\|_2. \quad (5.11)$$

This is the distance for which any $y \in \Omega$ is at most at a distance $h_{Y,\Omega}$ of Y . It can also be interpreted as the radius of the largest ball which is completely contained in Ω without intersecting Y .

Definition 5.2.3 (Separation distance). The separation distance of $Y = \{y_1, \dots, y_n\} \subset \Omega$ is defined as:

$$q_{Y,\Omega} = \frac{1}{2} \min_{i \neq j} \|y_i - y_j\|_2. \quad (5.12)$$

This quantity gives the maximal radius $r > 0$ such that all balls $\{y \in \mathbb{R}^d : \|y - y_j\|_2 \leq r\}$ are disjoint.

Finally, the following condition will be shown to play a key role in our analysis.

Definition 5.2.4 (Quasi-uniformity condition). A set of data locations $Y = \{y_1, \dots, y_n\} \subset \Omega$ is said to be quasi-uniform with respect to a constant $B > 0$ if

$$q_{Y,\Omega} \leq h_{Y,\Omega} \leq Bq_{Y,\Omega}. \quad (5.13)$$

5.3 Main results

5.3.1 Construction of an estimator

Let $F : \Omega \rightarrow \mathbb{R}^N$ denote the vector-valued function representing the impulse responses coefficients (IRC) in basis $(\phi_k)_{k \in \mathbb{N}}$:

$$F(y)[k] = \langle S(\cdot, y), \phi_k \rangle. \quad (5.14)$$

Based on the observation model (5.10), a natural approach to estimate the SVIR, consists in constructing an IRC $\hat{F} : \Omega \rightarrow \mathbb{R}^N$ such that $\hat{F}$ is an estimate of F . The estimated SVIR is then defined as

$$\hat{S}(x, y) = \sum_{k=1}^N \hat{F}(y)[k] \phi_k(x), \text{ for } (x, y) \in \Omega \times \Omega. \quad (5.15)$$

The two assumptions 5.2.1 and 5.2.2, motivate the introduction of the following space.

Definition 5.3.1 (Space $\mathcal{H}$ of IRC). The space $\mathcal{H}(\Omega)$ of admissible IRC is defined as the set of vector-valued functions $G : \Omega \rightarrow \mathbb{R}^N$ such that

$$\|G\|_{\mathcal{H}(\Omega)}^2 = \alpha \int_{y \in \Omega} \sum_{k=1}^N w[k] |G(y)[k]|^2 dy + (1 - \alpha) \sum_{k=1}^N |G(\cdot)[k]|_{BL^s(\Omega)}^2 < +\infty, \quad (5.16)$$

where $\alpha \in [0, 1)$ allows to balance the smoothness in each direction.

Lemma 5.3.1. *The SVIRs satisfying assumptions 5.2.1 and 5.2.2 have an IRC belonging to $\mathcal{H}(\Omega)$.*

Lemma 5.3.2. *The space $\mathcal{H}(\mathbb{R}^d)$ is a vector-valued Reproducing Kernel Hilbert Space (RKHS).*

To construct an estimator of F , we use the formalism of regression and smoothing in vector-valued Reproducing Kernel Hilbert Spaces (RKHS) [172, 173] which leads to define $\hat{F}_\mu$ as the minimizer of the optimisation problem

$$\hat{F}_\mu = \arg \min_{F \in \mathcal{H}(\mathbb{R}^d)} \frac{1}{n} \sum_{i=1}^n \|F_i^e - F(y_i)\|_{\mathbb{R}^N}^2 + \mu \|F\|_{\mathcal{H}(\mathbb{R}^d)}^2, \quad (5.17)$$

where $\mu > 0$ is a regularization parameter. Notice that the optimization is performed on $\mathcal{H}(\mathbb{R}^d)$ and not $\mathcal{H}(\Omega)$, for technical reasons related to the use of radial basis functions. The formalism of vector-valued RKHS has been developed for the purpose of multi-task learning, and its application to operator estimator appears to be novel.

5.3.2 Mixed-Sobolev space interpretation

In the specific case where $(\phi_k)_{k \in \mathbb{N}}$ is a wavelet or a Fourier basis and $N = +\infty$, the proposed methodology can be interpreted in terms of SVIR instead of IRC.

Lemma 5.3.3. *The cost function in problem 5.17 is equivalent to*

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \|F_i^\epsilon - (\langle S(\cdot, y_i), \phi_k \rangle)_{k \in \mathbb{N}}\|_2^2 + \mu \left(\alpha \int_{\mathbb{R}^d} |S(\cdot, y)|_{H^r(\mathbb{R}^d)}^2 dy \right. \\ \left. + (1 - \alpha) \int_{\mathbb{R}^d} |S(x, \cdot)|_{BL^s(\mathbb{R}^d)}^2 dx \right). \end{aligned} \quad (5.18)$$

In this formulation, the data fidelity term allows finding a SVIR that is close to the observed data, the first regularization term allows smoothing the additive noise on the acquired PSFs and the second interpolates the missing data.

5.3.3 Numerical complexity

Thanks to the results in [173], computing $\hat{F}_\mu$ amounts to solving a *finite-dimensional* system of linear equations. However, for an arbitrary orthonormal basis $(\phi_\lambda)_{\lambda \in \Lambda}$, and without any further assumptions on the kernel of the RKHS $\mathcal{H}(\mathbb{R}^d)$, evaluating $\hat{F}_\mu$ leads to the resolution of a *full* $nN \times nN$ linear system, which is untractable for large N and n .

With the specific choice of norm introduced in Definition 5.3.1, the problem simplifies to the resolution of N systems of equations of size $n \times n$. This yields the following proposition:

Proposition 5.3.1. *The solution of (5.17) can be computed in no more than $O(Nn^3)$ operations.*

In addition, if the weight function w is piecewise constant, some $n \times n$ matrices are identical, allowing to compute an LU factorization once for all and using it to solve many systems. In the specific case where $(\phi_k)_{k \in \mathbb{N}}$ is a wavelet basis, it is natural to set function w as a constant over each wavelet subband [164, Thm. 9.4]. This yields the following result.

Proposition 5.3.2. *If w is set as constant over each subband of a wavelet basis, the solution of (5.17) can be computed in no more than $O\left(\frac{\log(N)}{d}n^3 + Nn^2\right)$ operations.*

Finally for well chosen bases $(\phi_k)_{k \in \mathbb{N}}$ - including wavelets - the impulse responses can be well approximated using a small number N of atoms, making the method tractable even in very large scale applications.

The proposed ideas are illustrated on Figure 5.3. As can be seen on Figure 5.3 (e), computing the IRC in a wavelet basis allows expressing most of the information contained in the SVIR on a few lines only. Given the noisy dataset, the proposed




Figure 5.3 – An interpolation result. The exact SVIR corresponds to Gaussian PSFs with standard deviation $\sigma(t) = 0.05(1 + 2\min(t, 1 - t))$ for $t \in [0, 1]$. In this setting 7 PSFs are observed and corrupted by an additive iid Gaussian noise of standard deviation 5.10^{-3} .

algorithm simultaneously interpolates along lines and denoises along rows to obtain the results in Figure 5.3 (c) and (f).

To conclude this paragraph, let us mention that the representation of an operator of type (5.15) can be used to evaluate matrix-vector products rapidly. We refer the interested reader to [101] for more details.

5.3.4 Convergence rates

The convergence of the proposed estimator with respect to the number n of observations is captured by the following theorem.

Theorem 5.3.1. *Assume that S satisfies assumptions 5.2.1 and 5.2.2 and that it is sampled using model (5.10) under the quasi-uniformity condition given in Definition 5.2.4. Then the estimating operator $\hat{H}$ with SVIR $\hat{S}$ defined in equation (5.15) satisfies the following inequality*

$$\mathbb{E} \left(\|H - \hat{H}\|_{2 \rightarrow 2}^2 \right) \lesssim N^{-\frac{2r}{d}} + (N\sigma^2 n^{-1})^{\frac{2s}{2s+d}}, \quad (5.19)$$

for $\mu \propto (N\sigma^2 n^{-1})^{\frac{2s}{2s+d}}$.

Theorem 5.3.2. *Assume that S satisfies assumptions 5.2.1 and 5.2.2 and that it is sampled using model (5.10) under the quasi-uniformity condition given in Definition 5.2.4. Then the estimating operator $\hat{H}$ with SVIR $\hat{S}$ defined in equation (5.15) satisfies the following inequality*

$$\mathbb{E} \left(\|H - \hat{H}\|_{2 \rightarrow 2}^2 \right) \lesssim (\sigma^2 n^{-1})^{\frac{2q}{2q+d}}, \quad (5.20)$$

with the relation $1/q = 1/r + 1/s$, for $\mu \propto (\sigma^2 n^{-1})^{\frac{2q}{2q+d}}$ and $N \propto (\sigma^{-2} n)^{\frac{dq}{r(2q+d)}}$.

This last theorem gives some insights on the estimator behavior. In particular:

- It provides an explicit way of choosing the value of the regularization parameter μ : it should decrease as the number of observations increases.
- If the number of observations n is small, it is unnecessary to project the impulse responses on a high dimensional basis (i.e. N large). The basic reason is that not enough information has been collected to reconstruct the fine details of the kernel.
- A close inspection at the proof shows that only the fill distance $h_{Y,\Omega}$ plays a role in the convergence rate. The quasi-uniformity condition is unnecessary. It is however very likely that this hypothesis allows showing that the upper-bound (5.20) is optimal. This explains why we preferred stating the theorems under a quasi-uniformity condition and not explicitly provide the dependency on $h_{Y,\Omega}$. This aspect of the work might change before publication.

5.3.5 Extensions

The results announced above are still preliminary and there are at least two additional aspects that we plan to investigate.

Optimality Though we did not have time to show optimality results in this thesis yet, we do believe that the convergence rate (5.19), cannot be improved uniformly on the considered class of SVIR.

Impulse responses regularization An attentive reader might have realized that there is no explicit dependency on the balance parameter α in equation (5.19). This means that the effect of the smoothing term $\int_{y \in \Omega} \sum_{k=1}^N w[k] |G(y)[k]|^2 dy$ is not clear yet from a theoretical point of view. From a practical point of view, this term allows reducing noise on the data. We observed experimentally that it provides significant improvements over a simple line by line interpolation. To the best of our knowledge, this is the first time that such a term is introduced and it therefore represents a significant contribution of this work. Clarifying this regularizing aspect requires a better control of the eigenvalues of a matrix $\tilde{\Gamma}$ appearing in the proofs. We are rather optimistic about this control and expect to have a better understanding of it by the end of the PhD period.

5.4 Radial basis functions implementation

The objective of this section is to provide a fast algorithm to solve problem (5.17) and to prove Lemma 5.3.1 and 5.3.2. A few tools related to radial basis functions and useful for the subsequent proofs are also introduced.

A key observation is provided below.

Lemma 5.4.1. *For $k \in \{1, \dots, N\}$, function $\hat{F}(\cdot)[k]$ is the solution of the following variational problem:*

$$\min_{f \in H^s(\mathbb{R}^d)} \frac{1}{n} \sum_{i=1}^n (F_i^\epsilon[k] - f(y_i))^2 + \mu \left(\alpha w[k] \|f\|_{L^2(\mathbb{R}^d)}^2 + (1 - \alpha) \|f\|_{BL^s(\mathbb{R}^d)}^2 \right). \quad (5.21)$$

Proof. It suffices to remark that problem (5.17) consists of solving N independent sub-problems. $\square$

We now focus on the resolution of sub-problem (5.21) which completely fits the framework of radial basis function approximation.

5.4.1 Standard approximation results in RKHS

Let us begin with a few results about RKHS. Most of the results can be found in the book of Wendland [246].

Definition 5.4.1 (Positive definite function). A continuous function $\Phi : \mathbb{R}^d \rightarrow \mathbb{C}$ is called positive semi-definite if, for all $n \in \mathbb{N}$, all sets of pairwise distinct centers $Y = \{y_1, \dots, y_n\} \subset \mathbb{R}^d$, and all $\alpha \in \mathbb{C}^n$, the quadratic form

$$\sum_{j=1}^n \sum_{k=1}^n \alpha_j \bar{\alpha}_k \Phi(x_j - x_k) \quad (5.22)$$

is nonnegative. Function Φ is called positive definite if the quadratic form is positive for all $\alpha \in \mathbb{C}^n \setminus \{0\}$.

Definition 5.4.2 (Reproducing kernel). Let $\mathcal{G}$ denote a Hilbert space of real-valued functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ endowed with a scalar product $\langle \cdot, \cdot \rangle_{\mathcal{G}}$. A function $\Phi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is called reproducing kernel for $\mathcal{G}$ if

1. $\Phi(\cdot, y) \in \mathcal{G}, \quad \forall y \in \mathbb{R}^d,$
2. $f(y) = \langle f, \Phi(\cdot, y) \rangle_{\mathcal{G}},$ for all $f \in \mathcal{G}$ and all $y \in \mathbb{R}^d$.

Theorem 5.4.1 (RKHS). Suppose that $\mathcal{G}$ is a Hilbert space of functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$. Then the following statements are equivalent:

1. the point evaluations functionals are continuous for all $y \in \mathbb{R}^d$.
2. $\mathcal{G}$ has a reproducing kernel.

A Hilbert space satisfying the properties above is called a Reproducing Kernel Hilbert Space (RKHS).

The Fourier transform of a function $f \in L^1(\mathbb{R}^d)$ is defined by

$$\widehat{f}(\xi) = \mathcal{F}[f](\xi) = \int_{x \in \mathbb{R}^d} f(x) e^{-i\langle x, \xi \rangle} dx, \quad (5.23)$$

and the inverse transform by

$$\mathcal{F}^{-1}[\widehat{f}](x) = \int_{\xi \in \mathbb{R}^d} \widehat{f}(\xi) e^{i\langle x, \xi \rangle} d\xi. \quad (5.24)$$

The Fourier transform can be extended to $L^2(\mathbb{R}^d)$ and to $\mathcal{S}'(\mathbb{R}^d)$ the space of tempered distributions.

Theorem 5.4.2 ([246, Theorem 10.12]). Suppose that $\Phi \in C(\mathbb{R}^d) \cap L^1(\mathbb{R}^d)$ is a real-valued positive definite function. Define $\mathcal{G} = \{f \in L^2(\mathbb{R}^d) \cap C(\mathbb{R}^d) : \widehat{f}/\sqrt{\widehat{\Phi}} \in L^2(\mathbb{R}^d)\}$ equipped with

$$\langle f, g \rangle_{\mathcal{G}} = (2\pi)^{-d/2} \int_{\mathbb{R}^d} \frac{\widehat{f}(\xi) \overline{\widehat{g}(\xi)}}{\widehat{\Phi}(\xi)} d\xi. \quad (5.25)$$

Then $\mathcal{G}$ is a real Hilbert space with inner-product $\langle \cdot, \cdot \rangle_{\mathcal{G}}$ and reproducing kernel $\Phi(\cdot - \cdot)$.

Theorem 5.4.3. *Let $\mathcal{G}$ be an RKHS with positive definite reproducing kernel Φ . Let $(y_1, \dots, y_n)$ denote a set of points in $\mathbb{R}^d$ and $z \in \mathbb{R}^n$ denote a set of altitudes. The solution of the following approximation problem*

$$\min_{u \in \mathcal{G}} \frac{1}{n} \sum_{i=1}^n (u(y_i) - z[i])^2 + \frac{\mu}{2} \|u\|_{\mathcal{G}}^2 \quad (5.26)$$

can be written as:

$$u(x) = \sum_{i=1}^n c[i] \Phi(x - y_i), \quad (5.27)$$

where vector $c \in \mathbb{R}^n$ is the unique solution of the following linear system of equations

$$(G + n\mu \text{Id})c = z \text{ with } G[i, j] = \Phi(y_i - y_j). \quad (5.28)$$

5.4.2 Application to our problem

Let us now show how the above results help solving problem (5.21).

Proposition 5.4.1. *Let $\mathcal{G}$ be the Hilbert space of functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ such that $|f|_{BL^s(\mathbb{R}^d)}^2 + \|f\|_{L^2(\Omega)}^2 < +\infty$, equipped with the inner product:*

$$\langle f, g \rangle_{\mathcal{G}} = (1 - \alpha) \langle f, g \rangle_{BL^s(\mathbb{R}^d)} + \alpha w[k] \langle f, g \rangle_{L^2(\mathbb{R}^d)}^2. \quad (5.29)$$

Then $\mathcal{G}$ is an RKHS and its scalar product reads

$$\langle f, g \rangle_{\mathcal{G}} = (2\pi)^{-d/2} \int_{\mathbb{R}^d} \frac{\widehat{f}(\xi) \overline{\widehat{g}(\xi)}}{\widehat{\Phi}(\xi)} d\xi, \quad (5.30)$$

where the reproducing kernel Φ , is defined by:

$$\widehat{\Phi}_k(\xi) = \left((1 - \alpha) \|\xi\|^{2s} + \alpha w[k] \right)^{-1}. \quad (5.31)$$

Proof. The proof is a direct application of the different results stated previously. $\square$

The Fourier transform $\widehat{\Phi}_k$ is radial, so that Φ_k is radial too and the resolution of (5.21) fits the formalism of radial basis functions interpolation/approximation [38].

Remark 5.4.1. For some applications, it makes sense to set $w[k] = 0$ for some values of k . For instance, if $(\phi_k)_{k \in \mathbb{N}}$ is a wavelet basis, then it is usually good to set $w[k] = 0$ when k is the index of a scaling wavelet. In that case, the theory of conditionally positive definite kernels should be used instead of the one above. We do not detail this aspect since it is well described in standard textbooks [246, 38].

5.5 Proofs of the main results

In this section, we prove Theorem 5.3.1 about the convergence rate of the quadratic risk $\mathbb{E} \|S - \hat{S}\|_{HS}^2$.

Algorithm 8: Estimation algorithm

Input: Weight vector $w \in \mathbb{R}^N$. Regularity $s \in \mathbb{N}$.

PSF locations $Y = \{y_1, \dots, y_n\} \in \mathbb{R}^{d \times n}$.

Observed data $(F_i^\epsilon)_{1 \leq i \leq n}$, where $F_i^\epsilon \in \mathbb{R}^N$.

begin

 Identify the $m \leq N$ weights of identical values in $w \in \mathbb{R}^N$ *// $O(N \log(N))$*
 for *Each unique weight ω* **do** *// $O(mn^3)$*

 Compute matrix G from formula (5.28) with Φ defined in (5.31).

 Compute a LU decomposition of $M_\omega = (G + n\mu \text{Id}) = L_\omega U_\omega$.

end

for $k = 1$ **to** N **do** *// $O(Nn^2)$*

 Identify the value ω such that $w[k] = \omega$.

 Set $z = (F_i^\epsilon[k])_{1 \leq i \leq n}$.

 Solve the linear system $L_\omega U_\omega c_k = z$.

 Possibly reconstruct $\hat{F}$ by (see equation (5.27))

$$\hat{F}(x)[k] = \sum_{i=1}^n c_k[i] \Phi(x - y_i).$$

end

end

5.5.1 Operator norm risk

To analyse the theoretical properties of a given estimator of the operator H , we introduce the quadratic risk defined as:

$$R(\hat{H}, H) = \mathbb{E} \left\| \hat{H} - H \right\|_{2 \rightarrow 2}^2, \quad (5.32)$$

where $\hat{H}$ is the operator associated to the SVIR $\hat{S}$ defined in (5.15). The above expectation is taken with respect to the distribution of the observations in (5.10). The operator norm appearing in (5.32) is defined by:

$$\|H\|_{2 \rightarrow 2} = \sup_{\|u\|_{L^2(\Omega)}=1} \|Hu\|_2, \quad (5.33)$$

and can be bounded by the Hilbert-Schmidt norm of H , that is: $\|H\|_{2 \rightarrow 2} \leq \|H\|_{HS} = \|K\|_{L^2(\Omega \times \Omega)} = \|S\|_{L^2(\Omega \times \Omega)}$. From this observation we get that:

$$\begin{aligned} R(\hat{H}, H) &\leq \mathbb{E} \left\| \hat{H} - H \right\|_{HS}^2 \\ &\leq 2 \left(\|H - H_N\|_{HS}^2 + \mathbb{E} \left\| H_N - \hat{H} \right\|_{HS}^2 \right) \\ &= 2 \left(\underbrace{\|S - S_N\|_{L^2(\Omega \times \Omega)}^2}_{\epsilon_d(N)} + \underbrace{\mathbb{E} \left\| S_N - \hat{S} \right\|_{L^2(\Omega \times \Omega)}^2}_{\epsilon_e(n)} \right), \end{aligned} \quad (5.34)$$

where H_N is the operator associated to the SVIR S_N defined by

$$S_N(x, y) = \sum_{k=1}^N F(y)[k] \phi_k(x)$$

and $\hat{H}$ the estimating operator associated to the SVIR $\hat{S}$ as in (5.15).

In equation (5.34), the risk is decomposed as the sum of two terms $\epsilon_e(n)$ and $\epsilon_d(N)$. The first one $\epsilon_d(N)$ is the error introduced by the discretization step. The second term $\epsilon_e(N)$ is the quadratic risk between S_N and the estimator $\hat{S}$. In the next sections, we provide upper-bounds for $\epsilon_d(N)$ and $\epsilon_e(n)$.

5.5.2 Discretization error ϵ_d

The discretization error $\epsilon_d(N)$ can be controlled using the standard approximation result below (see e.g. [164, Theorem 9.1, p. 437]).

Theorem 5.5.1. *Let $f \in \mathcal{E}^r(\Omega)$ and let $f_N = \sum_{k=1}^N \langle f, \phi_k \rangle \phi_k$. Then*

$$\|f - f_N\|_2^2 \leq c \|f\|_{\mathcal{E}^r(\Omega)}^2 N^{-2r/d}, \quad (5.35)$$

where c is a universal constant.

Corollary 5.5.1. *Under assumptions 5.2.1 and 5.2.2, the discretization error satisfies:*

$$\epsilon_d(N) \lesssim N^{-2r/d}. \quad (5.36)$$

Proof. By assumption 5.2.1, $S(\cdot, y) \in \mathcal{E}^r(\Omega)$ for all $y \in \Omega$. Therefore, by Theorem 5.5.1:

$$\|S(\cdot, y) - S_N(\cdot, y)\|_{L^2(\Omega)}^2 \leq cN^{-2r/d}. \quad (5.37)$$

Finally:

$$\begin{aligned} \|S - S_N\|_{L^2(\Omega \times \Omega)}^2 &= \int_{y \in \Omega} \|S(\cdot, y) - S_N(\cdot, y)\|_{L^2(\Omega)}^2 dy \\ &\leq |\Omega|c \left(\sup_{y \in \Omega} \|S(\cdot, y)\|_{\mathcal{E}^r(\Omega)}^2 \right) N^{-2r/d}. \end{aligned}$$

□

5.5.3 Estimation error ϵ_e

This section provides an upper-bound on the estimation error

$$\epsilon_e(n) = \mathbb{E} \left\| S_N - \hat{S} \right\|_{L^2(\Omega \times \Omega)}^2. \quad (5.38)$$

This part is significantly harder than the rest of the paper. Let us begin with a simple remark.

Lemma 5.5.1. *The estimation error satisfies*

$$\epsilon_e(n) = \|F - \hat{F}\|_{\mathbb{R}^N \times L^2(\Omega)}^2. \quad (5.39)$$

Proof. Since $(\phi_k)_{1 \leq k \leq N}$ is an orthonormal basis, Parseval's theorem gives

$$\begin{aligned} \|S_N - \hat{S}\|_{L^2(\Omega \times \Omega)}^2 &= \int_{\Omega} \int_{\Omega} \left(S_N(x, y) - \hat{S}(x, y) \right)^2 dx dy \\ &= \int_{\Omega} \int_{\Omega} \left(\sum_{k=1}^N (F(y)[k] - \hat{F}(y)[k]) \phi_k(x) \right)^2 dx dy \\ &= \int_{\Omega} \sum_{k=1}^N (F(y)[k] - \hat{F}(y)[k])^2 dy \\ &= \sum_{k=1}^N \|F(\cdot)[k] - \hat{F}(\cdot)[k]\|_{L^2(\Omega)}^2 =: \|F - \hat{F}\|_{\mathbb{R}^N \times L^2(\Omega)}^2. \end{aligned} \quad (5.40)$$

□

By Lemma 5.4.1 the estimator defined in (5.17) can be decomposed as N independent estimators. Lemma 5.5.2 below provides a convergence rate for each of them. This result is strongly related to the work in [233] on smoothing splines. Unfortunately, we cannot directly apply the results in [233] to our setting since the kernel defined in (5.31) is not that of a thin-plate smoothing spline.

Lemma 5.5.2. *Suppose that $\Omega \subset \mathbb{R}^d$ is a bounded connected open set in $\mathbb{R}^d$ with Lipschitz continuous boundary and that the set $Y = \{y_1, \dots, y_n\} \subset \Omega$ of PSF locations satisfies a quasi-uniformity condition in the sense of Definition 5.2.4. Then, each function $\hat{F}(\cdot)[k]$ solution of problem (5.21) satisfies:*

$$\mathbb{E}\|\hat{F}(\cdot)[k] - F(\cdot)[k]\|_{L^2(\Omega)}^2 \lesssim \mu\|F(\cdot)[k]\|_{H^s(\Omega)}^2 + n^{-1}\sigma^2\mu^{-\frac{d}{2s}}, \quad (5.41)$$

provided that $n\mu^{d/2s} \geq 1$.

Proof. In order to prove the upper-bound (5.41), we first decompose the expected squared error $\mathbb{E}\|\hat{F}(\cdot)[k] - F(\cdot)[k]\|_{L^2(\Omega)}^2$ into bias and variance terms:

$$\mathbb{E}\|\hat{F}(\cdot)[k] - F(\cdot)[k]\|_{L^2(\Omega)}^2 \lesssim \underbrace{\|\hat{F}^0(\cdot)[k] - F(\cdot)[k]\|_{L^2(\Omega)}^2}_{\text{Bias term}} + \underbrace{\mathbb{E}\|\hat{F}^0(\cdot)[k] - \hat{F}(\cdot)[k]\|_{L^2(\Omega)}^2}_{\text{Variance term}}, \quad (5.42)$$

where $\hat{F}^0(\cdot)[k]$ is the solution of the noise-free problem

$$\hat{F}^0(\cdot)[k] = \arg \min_{f \in H^s(\mathbb{R}^d)} \frac{1}{n} \sum_{i=1}^n (F(y_i)[k] - f(y_i))^2 + \mu \left(\alpha w[k] \|f\|_{L^2(\mathbb{R}^d)}^2 + (1 - \alpha) |f|_{BL^s(\mathbb{R}^d)}^2 \right). \quad (5.43)$$

We then treat the bias and variance terms separately.

Control of the bias The bias control relies on sampling inequalities in Sobolev spaces. They first appeared in [94] to control the norm of functions in Sobolev spaces with scattered zeros. They have been generalized in different ways, see e.g. [247] and [6]. In this paper, we will use the following result from [6].

Theorem 5.5.2 ([6, Theorem 4.1]). *Let $\Omega \subset \mathbb{R}^d$ be a bounded connected open set with Lipschitz continuous boundary.*

Let $p, q, x \in [1, +\infty]$. Let s be a real number such that $s \leq d$ if $p = 1$, $s > d/p$ if $1 < p < \infty$ or $s \in \mathbb{N}^$ if $p = \infty$. Furthermore, let $l_0 = s - d(1/p - 1/q)_+$ and $\gamma = \max(p, q, x)$ where $(\cdot)_+ = \max(0, \cdot)$.*

Then, there exist two positive constants η_s (depending on Ω and s) and C (depending on Ω , n, s, p, q and x) satisfying the following property: for any finite set $Y \subset \bar{\Omega}$ (or $Y \subset \Omega$ if $p = 1$ and $s = d$) such that $h_{Y,\Omega} \leq \eta_s$, for any $u \in W^{s,p}(\Omega)$ and for any $l = 0, \dots, \lceil l_0 \rceil - 1$, we have

$$\|u\|_{W^{l,q}(\Omega)} \leq C \left(h_{Y,\Omega}^{s-l-d(1/p-1/q)_+} |u|_{W^{s,p}(\Omega)} + h_{Y,\Omega}^{d/\gamma-l} \|u\|_Y \right), \quad (5.44)$$

where $\|u|_Y\|_x = (\sum_{i=1}^n u(y_i)^x)^{1/x}$. If $s \in \mathbb{N}^*$ this bound also holds with $l = l_0$ when either $p < q < \infty$ and $l_0 \in \mathbb{N}$ or $(p, q) = (1, \infty)$ or $p \geq q$.

The above theorem is the key to obtain Proposition 5.5.1 below.

Proposition 5.5.1. *Set $a > 0$ and let $\mathcal{G}(\Omega)$ be the RKHS with norm defined by $\|\cdot\|_{\mathcal{G}(\Omega)}^2 = |\cdot|_{BL^s(\Omega)}^2 + a\|\cdot\|_{L^2(\Omega)}^2$. Let $u \in H^s(\Omega)$ denote a target function and $Y = \{y_1, \dots, y_n\} \subset \Omega$ a data site set. Let f_μ denote the solution of the following variational problem*

$$f_\mu = \arg \min_{f \in \mathcal{G}(\mathbb{R}^d)} \frac{1}{n} \sum_{i=1}^n (u(y_i) - f(y_i))^2 + \mu \|f\|_{\mathcal{G}(\mathbb{R}^d)}^2. \quad (5.45)$$

Then

$$\|f_\mu - u\|_{L^2(\Omega)} \leq C \left(h_{Y,\Omega}^s + h_{Y,\Omega}^{d/2} \sqrt{n\mu} \right) \|u\|_{H^s(\Omega)}, \quad (5.46)$$

where C is a constant depending only on Ω and s and $h_{Y,\Omega}$ is the fill distance defined in 5.2.2.

Proof. By applying the Sobolev sampling inequality of Theorem 5.5.2 for $p = q = x = 2$, $l = 0$, we get

$$\|v\|_{L^2(\Omega)} \leq C \left(h_{Y,\Omega}^s |v|_{H^s(\Omega)} + h_{Y,\Omega}^{d/2} \left(\sum_{i=1}^n v(y_i)^2 \right)^{1/2} \right), \quad (5.47)$$

for all $v \in H^s$. This inequality applied to function $v = f_\mu - u$ yields

$$\|f_\mu - u\|_{L^2(\Omega)} \leq C \left(h_{Y,\Omega}^s |f_\mu - u|_{H^s(\Omega)} + h_{Y,\Omega}^{d/2} \left(\sum_{i=1}^n (f_\mu(y_i) - u(y_i))^2 \right)^{1/2} \right), \quad (5.48)$$

The remaining task is to bound $|f_\mu - u|_{H^s(\Omega)}$ and $(\sum_{i=1}^n (f_\mu(y_i) - u(y_i))^2)^{1/2}$ by $\|u\|_{H^s(\Omega)}$. One part of the difficulty lies in the fact that f_μ minimizes the semi-norm over $\mathbb{R}^n$ and that we wish a control over the domain Ω . This motivates the introduction of an extension operator $P : H^s(\Omega) \rightarrow H^s(\mathbb{R}^d)$ defined by

$$Pu = \arg \min_{\substack{f \in H^s(\mathbb{R}^d) \\ f|_\Omega = u}} \|f\|_{H^s(\mathbb{R}^d)}^2. \quad (5.49)$$

By [94, Lemma 3.1], this operator is continuous, i.e. there exists a constant $K > 0$ (depending on Ω and s) such that for all $u \in H^s(\Omega)$, $\|Pu\|_{H^s(\mathbb{R}^d)} \leq K\|u\|_{H^s(\Omega)}$. Now, let $f_0 : \Omega \rightarrow \mathbb{R}$ denote the solution of

$$f_0 = \arg \min_{\substack{f \in \mathcal{G}(\mathbb{R}^d) \\ f(y_j) = u(y_j)}} \|f\|_{\mathcal{G}(\mathbb{R}^d)}^2, \quad (5.50)$$

By strict convexity of the squared norm $\|\cdot\|_{\mathcal{G}(\mathbb{R}^d)}^2$, function f_0 is uniquely determined. Moreover, it satisfies $\|f_0\|_{\mathcal{G}(\Omega)} \leq \|f_0\|_{\mathcal{G}(\mathbb{R}^d)} \leq \|Pu\|_{\mathcal{G}(\mathbb{R}^d)} \leq \|Pu\|_{H^s(\mathbb{R}^d)} \leq K\|u\|_{H^s(\Omega)}$.

Now, let us define two functionals $f \mapsto E(f) = \frac{1}{n} \sum_{i=1}^n (u(y_i) - f(y_i))^2$ and $f \mapsto J(f) = \|f\|_{\mathcal{G}(\mathbb{R}^d)}^2$. Since f_μ is the minimizer of (5.45), it satisfies

$$E(f_\mu) + \mu J(f_\mu) \leq E(f_0) + \mu J(f_0). \quad (5.51)$$

In addition $E(f_0) = 0$ and $J(f_0) \leq \|Pu\|_{\mathcal{G}(\mathbb{R}^d)}^2 \leq \|Pu\|_{H^s(\mathbb{R}^d)}^2 \leq K^2 \|u\|_{H^s(\Omega)}^2$. Hence,

$$E(f_\mu) + \mu J(f_\mu) \leq K\mu \|u\|_{H^s(\Omega)}^2 \quad (5.52)$$

which yields

$$E(f_\mu) = \frac{1}{n} \sum_{i=1}^n (u(y_i) - f_\mu(y_i))^2 \leq K^2 \mu \|u\|_{H^s(\Omega)}^2. \quad (5.53)$$

To finish, the triangle inequality yields $|f_\mu - u|_{H^s(\Omega)} \leq |f_\mu|_{H^s(\Omega)} + |u|_{H^s(\Omega)}$. Then, $|u|_{H^s(\Omega)} \leq \|u\|_{H^s(\Omega)}$ and bound (5.52) leads to:

$$|f_\mu|_{H^s(\Omega)} \leq \|f_\mu\|_{\mathcal{G}(\mathbb{R}^d)} \leq \|f_0\|_{\mathcal{G}(\mathbb{R}^d)} \leq K \|u\|_{H^s(\Omega)}. \quad (5.54)$$

Therefore

$$|f_\mu - u|_{H^s(\Omega)} \leq (K + 1) \|u\|_{H^s(\Omega)}. \quad (5.55)$$

Replacing bounds (5.55) and (5.53) in the sampling inequality (5.48) completes the proof of Proposition 5.5.1. $\square$

Applying Proposition 5.5.1 to $\hat{F}^0(\cdot)[k]$, we get

$$\|\hat{F}^0(\cdot)[k] - F(\cdot)[k]\|_{L^2(\Omega)}^2 \leq C \left(h_{Y,\Omega}^s + \sqrt{\mu n} h_{Y,\Omega}^{d/2} \right)^2 \|F(\cdot)[k]\|_{H^s(\Omega)}^2. \quad (5.56)$$

The trick is now to use the quasi-uniformity condition given in Definition 5.2.4 to control $h_{Y,\Omega}^s$ and $\sqrt{\mu n} h_{Y,\Omega}^{d/2}$. This is achieved using the following proposition.

Proposition 5.5.2 ([246, Proposition 14.1] or [233]). *Let $Y = \{y_1, \dots, y_n\} \subset \Omega$ be a quasi-uniform set with respect to B . Then, there exist constants $c > 0$ and $C > 0$ depending only on d , Ω and B such that,*

$$cn^{-1} \leq h_{Y,\Omega}^d \leq Cn^{-1}. \quad (5.57)$$

Condition $n\mu^{d/2s} \geq 1$ combined with the right-hand-side of (5.57) yields $h_{Y,\Omega}^d \leq C\mu^{d/2s}$, so that $h_{Y,\Omega}^s \lesssim \sqrt{\mu}$. Similarly, the right-hand-side of (5.57) yields $\sqrt{\mu n} h_{Y,\Omega}^{d/2} \lesssim \sqrt{\mu}$. Hence

$$\|\hat{F}^0(\cdot)[k] - F(\cdot)[k]\|_{L^2(\Omega)}^2 \lesssim \mu \|F(\cdot)[k]\|_{H^s(\Omega)}^2. \quad (5.58)$$

Control of the variance The variance term is treated following arguments similar to those in [233]. However, the change of kernel needs additional treatments. First of all, note that due to the linearity of the estimators (5.21), we have $\hat{F}_\mu^0(\cdot)[k] - \hat{F}_\mu(\cdot)[k] = f_k^\eta$ with $\eta \in \mathbb{R}^n$ defined as $\eta[i] = \epsilon_i[k]$ and

$$f_k^\eta = \arg \min_{f \in H^s(\mathbb{R}^d)} \frac{1}{n} \sum_{i=1}^n (f(y_i) - \eta[i])^2 + \mu \left(\alpha w[k] \|f\|_{L^2(\mathbb{R}^d)}^2 + (1 - \alpha) \|f\|_{BL^s(\mathbb{R}^d)}^2 \right). \quad (5.59)$$

We therefore need to estimate $\mathbb{E} \|f_k^\eta\|_{L^2(\Omega)}^2$. From Theorem 5.5.2 applied with $p = q = x = 2$ and $l = 0$ we obtain that for $u \in H^s(\Omega)$

$$\|u\|_{L^2(\Omega)} \leq C \left(h_{Y,\Omega}^s |u|_{H^s(\Omega)} + h_{Y,\Omega}^{d/2} \|u\|_Y \right).$$

Using the above inequality together with Proposition 5.5.2, we get that

$$\|f_k^\eta\|_{L^2(\Omega)}^2 \leq 2C \left(h_{Y,\Omega}^{2s} |f_k^\eta|_{H^s(\Omega)}^2 + n^{-1} \sum_{i=1}^n f_k^\eta(y_i)^2 \right).$$

As in [233], let us define the $n \times n$ matrix $\tilde{\Gamma}$ such that

$$\langle \tilde{\Gamma} z, z \rangle = \min_{\substack{u \in BL_s(\mathbb{R}^d) \\ u(y_i) = z[i]}} (1 - \alpha) |u|_{H^s(\mathbb{R}^d)}^2 + \alpha w[k] \|u\|_{L^2(\mathbb{R}^d)}^2. \quad (5.60)$$

The solution of problem (5.60) is a spline interpolating the data $(y_i, z[i])_{i=1}^n$. Using this matrix, we can write (5.59) as:

$$\min_{z \in \mathbb{R}^n} \frac{1}{n} \sum_{i=1}^n (z[i] - \eta[i])^2 + \mu \langle \tilde{\Gamma} z, z \rangle,$$

see [233, 232, 237] for details. Thus, the solution $\hat{z} = (f_k^\eta(y_i))_{i=1}^n$ is obtained by:

$$\hat{z} = (\text{Id} + n\mu\tilde{\Gamma})^{-1}\eta.$$

By letting $E_\mu = (\text{Id} + n\mu\tilde{\Gamma})^{-1}$, we obtain

$$n^{-1} \sum_{i=1}^n f_k^\eta(y_i)^2 = n^{-1} \sum_{i=1}^n \hat{z}[i]^2 = n^{-1} \eta^T E_\mu^2 \eta$$

and

$$\begin{aligned} (1 - \alpha) |f_k^\eta|_{H^s(\mathbb{R}^d)}^2 + \alpha w[k] \|f_k^\eta\|_{L^2(\mathbb{R}^d)}^2 &= \hat{z}^T \tilde{\Gamma} \hat{z} = \eta^T E_\mu \tilde{\Gamma} E_\mu \eta \\ &= (n\mu)^{-1} \eta^T E_\mu (E_\mu^{-1} - \text{Id}) E_\mu \eta \\ &= (n\mu)^{-1} \eta^T (E_\mu - E_\mu^2) \eta. \end{aligned}$$

Thus

$$|f_k^\eta|_{H^s(\Omega)}^2 \leq |f_k^\eta|_{H^s(\mathbb{R}^d)}^2 \leq (n\mu)^{-1} \eta^T (E_\mu - E_\mu^2) \eta.$$

Using the fact that η has i.i.d. components with zero mean and variance σ^2 , we get that,

$$\mathbb{E} \left[n^{-1} \sum_{i=1}^n f_{\lambda}^{\eta}(y_i)^2 \right] = n^{-1} \sigma^2 \text{Tr}(E_{\mu}^2),$$

and on the other hand

$$\begin{aligned} \mathbb{E} |f_k^{\eta}|_{H^s(\Omega)}^2 &\leq (n\mu)^{-1} \sigma^2 (\text{Tr}(E_{\mu}) - \text{Tr}(E_{\mu}^2)) \\ &\leq (n\mu)^{-1} \sigma^2 \text{Tr}(E_{\mu}). \end{aligned}$$

We now have to focus on the estimation of both $\text{Tr}(E_{\mu}) = \sum_{i=1}^n (1 + n\mu\lambda_i(\tilde{\Gamma}))^{-1}$ and $\text{Tr}(E_{\mu}^2) = \sum_{i=1}^n (1 + n\mu\lambda_i(\tilde{\Gamma}))^{-2}$, where $\lambda_i(\tilde{\Gamma})$ is the i -th eigenvalue of $\tilde{\Gamma}$. This will be achieved by analyzing the eigenvalues of the matrix $\tilde{\Gamma}$. This step is uneasy. Hopefully, Utreras analyzed the eigenvalues of the matrix Γ associated to thin-plate splines in [233]. Matrix Γ is defined in a similar way as (5.60):

$$\langle \Gamma z, z \rangle = \min_{\substack{u \in BL_s(\mathbb{R}^d) \\ u(y_i) = z[i]}} |u|_{H^s(\mathbb{R}^d)}^2. \quad (5.61)$$

One therefore has that $(1 - \alpha)z^T \Gamma z \leq z^T \tilde{\Gamma} z$ for all $z \in \mathbb{R}^N$. Therefore the matrix $\tilde{\Gamma} - (1 - \alpha)\Gamma$ is semi-definite positive. By virtue of Weyl's monotonicity theorem [248], we get that $(1 - \alpha)\lambda_i(\Gamma) \leq \lambda_i(\tilde{\Gamma})$. Hence we can bound the traces of the matrices E_{μ} and E_{μ}^2 as follows

$$\begin{aligned} \text{Tr}(E_{\mu}) &\leq \sum_{i=1}^n (1 + (1 - \alpha)n\mu\lambda_i(\Gamma))^{-1}, \\ \text{Tr}(E_{\mu}^2) &\leq \sum_{i=1}^n (1 + (1 - \alpha)n\mu\lambda_i(\Gamma))^{-2}. \end{aligned}$$

It is shown in [233], that $\gamma = \binom{s-1+d}{s-1}$ eigenvalues $\lambda_i(\Gamma)$ are null and the others satisfy $i^{2m/d}n^{-1} \lesssim \lambda_i(\Gamma) \lesssim i^{2m/d}n^{-1}$ for $\gamma + 1 \leq i \leq n$. Following [233], it can be shown that both traces are bounded by quantities proportional to $\mu^{-d/2s}$. Thus, one has that

$$\mathbb{E} \|f_k^{\eta}\|_{L^2(\Omega)}^2 \lesssim \sigma^2 (n^{-1} \mu^{-d/2s} + n^{-1} h_{Y,\Omega}^{2s} \mu^{-1} \mu^{-d/2s}).$$

Since $\mu^{d/2s}n \geq 1$ and using Proposition 5.5.2 that gives $n \lesssim h_{Y,\Omega}^{-d}$ we obtain that $h_{Y,\Omega}^{2s} \mu \lesssim 1$. Hence

$$\mathbb{E} \|f_k^{\eta}\|_{L^2(\Omega)}^2 \lesssim \sigma^2 n^{-1} \mu^{-d/2s} (1 + h_{Y,\Omega}^{2s} \mu) \lesssim \sigma^2 n^{-1} \mu^{-d/2s},$$

which completes the proof of Lemma 5.5.2. $\square$

5.5.4 Proof of the main results

Proof of Theorem 5.3.1

Proof. By equation (5.34):

$$\mathbb{E}\|\hat{H} - H\|_{2 \rightarrow 2}^2 \leq 2(\epsilon_d(N) + \epsilon_e(n)). \quad (5.62)$$

By Corollary (5.5.1)

$$\epsilon_d(N) \lesssim N^{-2r/d}. \quad (5.63)$$

Now, let us control ϵ_e . Let $F_N = F(\cdot)[1 : N]$.

$$\epsilon_e(n) = \mathbb{E}\|\hat{F} - F_N\|_{\mathbb{R}^N \times L^2(\Omega)}^2 \quad (5.64)$$

$$= \sum_{k=1}^N \mathbb{E}\|\hat{F}(\cdot)[k] - F_N(\cdot)[k]\|_{L^2(\Omega)}^2 \quad (5.65)$$

$$\stackrel{(5.41)}{\lesssim} \sum_{k=1}^N \left(\mu \|F_N(\cdot)[k]\|_{H^s(\Omega)}^2 + n^{-1} \sigma^2 \mu^{-d/2s} \right) \quad (5.66)$$

$$= \mu \|F_N\|_{\mathbb{R}^N \times H^s(\Omega)}^2 + N n^{-1} \sigma^2 \mu^{-d/2s}. \quad (5.67)$$

This upper-bound allows to set the value of the regularization parameter μ by balancing the two terms $\mu \|F_N\|_{\mathbb{R}^N \times H^s(\Omega)}^2$ and $N n^{-1} \sigma^2 \mu^{-d/2s}$:

$$\mu \|F\|_{\mathbb{R}^N \times H^s(\Omega)}^2 \propto N n^{-1} \sigma^2 \mu^{-d/2s}. \quad (5.68)$$

This yields

$$\mu \propto \left(N \sigma^2 n^{-1} \|F\|_{\mathbb{R}^N \times H^s(\Omega)}^{-2} \right)^{\frac{2s}{2s+d}}. \quad (5.69)$$

Plugging this value in the upper-bound of $\epsilon_e(n)$ gives

$$\begin{aligned} \mu \|F\|_{\mathbb{R}^N \times H^s(\Omega)}^2 &\propto \left(N \sigma^2 n^{-1} \|F\|_{\mathbb{R}^N \times H^s(\Omega)}^{-2} \right)^{\frac{2s}{2s+d}} \|F\|_{\mathbb{R}^N \times H^s(\Omega)}^2, \\ &= (N \sigma^2 n^{-1})^{\frac{2s}{2s+d}} \|F\|_{\mathbb{R}^N \times H^s(\Omega)}^{\frac{2d}{2s+d}} \end{aligned} \quad (5.70)$$

and

$$\epsilon_e(n) \lesssim (N \sigma^2 n^{-1})^{\frac{2s}{2s+d}}. \quad (5.71)$$

□

Proof of Theorem 5.3.2

Proof. To obtain this bound we use Theorem 5.3.1 and we balance the two terms so that:

$$N^{-2r/d} \propto (N \sigma^2 n^{-1})^{\frac{2s}{2s+d}} \quad (5.72)$$

This gives the choice $N \propto (\sigma^{-2}n)^{\frac{2sd}{4rs+2rd+2sd}}$. Replacing N by this value in bound (5.19) gives

$$\begin{aligned} N^{-2r/d} &\propto (\sigma^2 n^{-1})^{\frac{4rs}{4rs+2rd+2sd}}, \\ &= (\sigma^2 n^{-1})^{\frac{2q}{2q+d}}. \end{aligned} \tag{5.73}$$

□

Part III

Imaging problems

6

A Variational Model for Multiplicative Structured Noise Removal

Résumé : *Ce chapitre est issu de [104], accepté en Juin 2016 pour publication dans Journal of Mathematical Imaging and Vision (JMIV). Il est le fruit d'une collaboration avec Pierre Weiss et Wenxing Zhang. Nous considérons le problème d'estimation d'images dégradées par des bruits qui sont à la fois structurés et multiplicatifs. Notre première motivation a été les images issues du Selective Plane Illumination Microscope (SPIM) qui souffrent souvent d'inhomogénéités d'intensité dues à l'absorption et la réfraction de la lumière lors de sa propagation dans l'échantillon. Ce type de dégradations apparaît aussi dans d'autres modalités d'imagerie comme l'échographie. Nous avons modélisé les bruits multiplicatifs comme des processus stationnaires avec une distribution connue. Cette modélisation donne lieu à un nouveau problème de restauration convexe reposant sur le principe du maximum a posteriori. Après l'étude de quelques propriétés analytiques des minimiseurs, nous proposons finalement des méthodes d'optimisation rapides implémentées sur GPU. Les expériences numériques sur des images 2D issues de la microscopie à fluorescence démontrent l'utilité pratique de ce modèle.*

Abstract: *We consider the problem of restoring images impaired by noise that is simultaneously structured and multiplicative. Our primary motivation for this setting is the Selective Plane Illumination Microscope (SPIM) which often suffers from severe inhomogeneities due to light absorption and scattering. This type of degradation arises in other imaging devices such as ultrasonic imaging. We model the multiplicative noise as a stationary process with known distribution. This leads to a novel convex image restoration model based on a maximum a posteriori estimator. After establishing some analytical properties of the minimizers, we finally propose a fast optimization method on GPU. Numerical experiments on 2D fluorescence microscopy images demonstrate the usefulness of the proposed models in practical applications.*

Contents

6.1	Introduction	185
6.1.1	Related works	186
6.1.2	Our contribution	187
6.1.3	Paper structure	187
6.2	Preliminaries	188
6.2.1	Notation	188
6.2.2	Convex sets and functions	188
6.2.3	Discretization of differential operators	189
6.2.4	The (inverse)-gamma distribution	190
6.3	Noise model	190
6.3.1	Modeling the noise	190
6.3.2	A few noise properties	192
6.3.3	Noise simulations	193
6.4	Restoration model	193
6.4.1	MAP based restoration models	193
6.4.2	Simplifying the model	196
6.4.3	Duality and mean preservation	198
6.4.4	Relationships to other models in the white noise regime	199
6.5	Solvers for the proposed model (6.44)	199
6.6	Numerical experiments	200
6.6.1	The case of white noise	200
6.6.2	The case of structured noise	201
6.6.3	A few results on real images	204




Figure 6.1 – SPIM acquisition of a HCT116 tumor spheroid. Large stripes strongly impair image quality.

6.1 Introduction

Noise reduction is a long-standing problem which plays a fundamental role to simplify image analysis. Denoising algorithms should be designed by accounting carefully for the noise properties. Many works focus on additive white noise (i.e. noise that is independent pixel-wise). This problem is arguably close to have reached a mature state, at least in the Gaussian white noise setting [153]. In this paper, we consider a setting where the noise is made of spatially correlated patterns that affect the image multiplicatively. To the best of our knowledge, this problem is left unexplored until now.

Our initial motivation is an instance of fluorescence microscopy called Selective Plane Illumination Microscopy (SPIM) [139]. In this modality, samples are excited by a light sheet that coincides with the focal plane of a camera. While passing through the sample, the light sheet is diffracted and absorbed. As a result, images suffer from stripes parallel to the direction of light propagation. A typical image obtained with a SPIM is displayed in Figure 6.1. As can be seen, large stripes make automatic image analysis difficult. The proposed ideas, though initially developed for the SPIM, are likely to be useful for other imaging devices such as ultrasound imaging (speckle noise) [194], hyperspectral remote sensing imaging (waterfall effects) [117, 139] or nanotomography and synchrotron based X-ray microtomography (“ripple” effect) [137, 189, 178, 229, 150, 32].

6.1.1 Related works

In this section, we provide a brief and incomplete review of existing methods to treat multiplicative noise and additive structured noise.

Multiplicative noise

Probably the simplest way to treat multiplicative noise simply consists of applying a logarithm to the image and then use standard techniques for additive white noise. This however introduces a bias that needs to be corrected with dedicated methods [95]. The first variational methods that attempted to take the multiplicative nature of noise into account more finely are probably due to Rudin, Lions, Osher in [200] and later by Shi and Osher [210]. These methods were however not based on a clear statistical framework. Roughly at the same time, Aubert and Aujol [9] proposed a nonconvex variational model based on a Maximum A Posteriori (MAP) estimator. The nonconvexity usually makes it impossible to find global minimizers. This motivated Teuber and Steidl [216] to develop convex variational methods based on I -divergences. More recently, some authors extended the ideas from patch-based and dictionary learning restoration to the multiplicative case [82, 138], with significant improvements. It is however unclear how to apply patch based restoration with correlated noise. Finally, let us mention that a few authors also studied deconvolution problems with multiplicative noise [200, 90, 252]. These models are closely related to the proposed approach, since deconvolution can be seen as a denoising problem with structured noise after a direct deconvolution. The main difference is that the noise structure is completely different in the proposed approach, since its Fourier transform does not blow up in high frequencies.

Structured noise

In recent years, treating structured noise gained importance in fields such as life sciences, astronomy or seismology. Many works are dedicated to a particular imaging modality, since noise structure is highly dependent on the physical acquisition principle. To name a few, Münch et al [178] devised a wavelet-Fourier based filter for removing stripe and ring artifacts in ion beam nanotomography and computed tomography (CT) X-Ray microtomography. Aizenberg and Butakoff [2] proposed to use median filters in the Fourier domain to reduce quasi-periodic noise. Boas and Fleischmann [32] reviewed variant structured artifacts in computed tomography (CT) and developed an iterative reconstruction approach. Cornelis et al [75] designed double filters (a smoothing Wiener filter and an adaptive filter) approach for digitally removing periodic canvas patterns superimposed on painting artworks. Chen and Pellequer [61] developed a “divide-and-conquer” approach, where the Fourier spectrum of the image is divided into central and off-center regions for noisy pixels detection and intensity restoration, to remove heavy and fine stripes in atomic force microscopy bio-molecular images. Anas et al [4] provided a detailed statistical analysis including classification, detection and correction of ring artifacts in flat panel detectors based micro-CT images. Kim et al [146] suggested to

reduce ring artifacts in sinogram data by calculating the line-ratio and equalizing detector element in sinogram space. Chang et al [58] developed a variational destriping model by combining the unidirectional total variation and framelet regularization. Fitschen et al [111] proposed similar ideas to remove the curtaining effect and stripes in Ion Beam nanotomography images. Finally, Sur and Grédecac [221] proposed an automatized algorithm based on natural image statistics to remove quasi-periodic noise from strain images.

To finish, let us mention a previous works by one of the authors [107, 106] (which played a role in the emergence of subsequent methods such as [58, 111, 221]). Therein, we proposed a generic framework to treat additive structured noise within a variational framework. The noise was modeled as the convolution of a given filter with a random vector with i.i.d. components. The motivation behind this modeling is that in many applications, the noise is stationary: its distribution should be invariant to translations. The MAP principle then leads to a convex restoration model solved with efficient first order methods. The proposed methodology was proved to be very useful in different image modalities (SPIM, nanotomography, atomic force microscopy, bathymetry, satellite imaging...). It is used on a daily basis in the imaging facility of our laboratory. However, it leads to some artifacts and the restoration is imperfect when images suffer from with very large or dark stripes.

6.1.2 Our contribution

This work is built upon the previous paper [107]. The first contribution is a phenomenological model to describe the random patterns that appear in Figure 6.1. The proposed statistical model reproduces the degradations observed in practice rather faithfully. Its second important feature, is that the maximum a posteriori principle leads to convex restoration models. Convexity allows us to determine some analytical properties of the minimizers and to design an efficient minimization procedure. We finally apply the proposed algorithm to synthetic and real images, to demonstrate the interest of the proposed approach in biological imaging. The proposed algorithm is implemented on GPU, interfaced with Fiji [205] and distributed freely on our webpage.

6.1.3 Paper structure

We first provide some background information in section 6.2. In section 6.3, we design a first image formation model. The MAP leads to a convex variational problem that contains four parameters. In section 6.4, we analyze the proposed model and show that it can be simplified and contain only one regularization parameter. We then provide some analytical properties of the simplified model minimizers and design a minimization algorithm. Finally, we propose some numerical results on synthetic and real images in section 6.6.

6.2 Preliminaries

6.2.1 Notation

Let E denote the vector space of images defined on $\Omega = \{1, \dots, n_1\} \times \dots \times \{1, \dots, n_d\}$. The total number of pixels is therefore $n = n_1 \cdot \dots \cdot n_d$. Let E_+ denote the space of images on Ω with non-negative entries. The pixels of the image are identified by a multi-index $\mathbf{i} = (i_1, \dots, i_d) \in \Omega$. For any image $\mathbf{u} \in E$ and scalar $p \in [1, +\infty)$, we let $\|\mathbf{u}\|_p := (\sum_{\mathbf{i} \in \Omega} |\mathbf{u}[\mathbf{i}]|^p)^{1/p}$ and $\|\mathbf{u}\|_\infty := \max_{\mathbf{i} \in \Omega} |\mathbf{u}[\mathbf{i}]|$ denote the standard l^p and l^∞ -norm, respectively. The dot product of $\mathbf{u}_1$ and $\mathbf{u}_2 \in E$ is defined by $\langle \mathbf{u}_1, \mathbf{u}_2 \rangle := \sum_{\mathbf{i} \in \Omega} \mathbf{u}_1[\mathbf{i}] \mathbf{u}_2[\mathbf{i}]$. The symbol $\mathbf{1}$ stands for the element of E with all values equal to 1. The identity operator is denoted $\mathbf{I}$. The notation $\text{diag}(\mathbf{u})$ indicates a diagonal operator with diagonal elements equal to the entries of $\mathbf{u}$. Let $V = E^d$ denote the space of discrete vector fields on Ω . For any $\mathbf{q} = (\mathbf{q}_1, \dots, \mathbf{q}_d) \in V$, $|\mathbf{q}|$ denotes an element of E with $\mathbf{i}$ -th entry defined by

$$|\mathbf{q}|[\mathbf{i}] := \sqrt{\sum_{l=1}^d \mathbf{q}_l[\mathbf{i}]^2}. \quad (6.1)$$

6.2.2 Convex sets and functions

For simplicity, we identify E with its dual E^* , the space of linear forms on E . Let $f : E \rightarrow (-\infty, +\infty]$ be a function. The domain of f is defined by $\text{dom} f := \{\mathbf{x} \in E \mid f(\mathbf{x}) < +\infty\}$. Function f is proper if it is not identically equal to $+\infty$. It is closed if its epigraph $\text{epi} f := \{(\mathbf{x}, z) \in E \times \mathbb{R} \mid f(\mathbf{x}) \leq z\}$ is closed in $E \times \mathbb{R}$. The indicator of a set $D \subset E$ is defined by

$$\chi_D(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \in D, \\ +\infty & \text{otherwise.} \end{cases} \quad (6.2)$$

The proximity operator of the proper convex function f is defined by

$$\text{prox}_f(\mathbf{x}) := \arg \min_{\mathbf{y} \in E} \left\{ f(\mathbf{y}) + \frac{1}{2} \|\mathbf{y} - \mathbf{x}\|_2^2 \right\}, \quad \forall \mathbf{x} \in E. \quad (6.3)$$

The conjugate of f is the function $f^* : E \rightarrow (-\infty, +\infty]$ defined by

$$f^*(\mathbf{y}) := \sup_{\mathbf{x} \in E} \{ \langle \mathbf{x}, \mathbf{y} \rangle - f(\mathbf{x}) \}. \quad (6.4)$$

Let $D \subseteq E$ be a convex and closed subset of E . The normal cone to D at point $\boldsymbol{\lambda}$ is denoted $\mathcal{N}_D(\boldsymbol{\lambda})$. It is defined for all $\boldsymbol{\lambda} \in D$ by:

$$\mathcal{N}_D(\boldsymbol{\lambda}) = \{ \boldsymbol{\eta} \in E, \langle \boldsymbol{\eta}, \boldsymbol{\lambda}' - \boldsymbol{\lambda} \rangle \leq 0, \forall \boldsymbol{\lambda}' \in D \}. \quad (6.5)$$

Notice that since $\mathcal{N}_D(\boldsymbol{\lambda})$ is a cone:

$$t\mathcal{N}_D(\boldsymbol{\lambda}) = \mathcal{N}_D(\boldsymbol{\lambda}), \quad \forall t > 0, \forall \boldsymbol{\lambda} \in E. \quad (6.6)$$

6.2.3 Discretization of differential operators

The discrete gradient of $\mathbf{u} \in E$ is defined by:

$$\nabla \mathbf{u} = (\partial_1 \mathbf{u}, \dots, \partial_d \mathbf{u}) \in V. \quad (6.7)$$

The partial derivatives are defined as in [52] by:

$$(\partial_l \mathbf{u})[\mathbf{i}] = \begin{cases} \mathbf{u}[\dots, i_l + 1, \dots] - \mathbf{u}[\dots, i_l, \dots] & \text{if } i_l < n_l, \\ 0 & \text{otherwise.} \end{cases} \quad (6.8)$$

The adjoint operator of ∂_l^* is the unique operator satisfying:

$$\langle \partial_l \mathbf{u}, \mathbf{q}_l \rangle = \langle \mathbf{u}, \partial_l^* \mathbf{q}_l \rangle, \quad \forall \mathbf{u}, \mathbf{q}_l \in E. \quad (6.9)$$

It is easy to establish that:

$$(\partial_l^* \mathbf{q}_l)[\mathbf{i}] = \begin{cases} \mathbf{u}[\dots, i_l, \dots] - \mathbf{u}[\dots, i_l - 1, \dots] & \text{if } 1 < i_l < n_l, \\ \mathbf{u}[\dots, i_l, \dots] & \text{if } i_l = 1, \\ -\mathbf{u}[\dots, i_l - 1, \dots] & \text{if } i_l = n_l. \end{cases} \quad (6.10)$$

The adjoint of the gradient operator is defined by:

$$\begin{aligned} \nabla^* : V &\rightarrow E \\ \mathbf{q} &\mapsto \sum_{l=1}^d \partial_l^* \mathbf{q}_l. \end{aligned} \quad (6.11)$$

Let $\mathbf{A} : E \rightarrow V$ denote a linear operator. Its spectral norm is defined by

$$\|\mathbf{A}\| = \sup_{\mathbf{x} \in E, \|\mathbf{x}\|_2 \leq 1} \|\mathbf{A}\mathbf{x}\|_2. \quad (6.12)$$

The point-wise product between two elements $\mathbf{u}$ and $\mathbf{v}$ of E is denoted $\mathbf{u} \odot \mathbf{v}$:

$$(\mathbf{u} \odot \mathbf{v})[\mathbf{i}] = \mathbf{u}[\mathbf{i}] \mathbf{v}[\mathbf{i}]. \quad (6.13)$$

The point-wise division between two elements $\mathbf{u}$ and $\mathbf{v}$ of E is denoted $\mathbf{u} \oslash \mathbf{v}$:

$$(\mathbf{u} \oslash \mathbf{v})[\mathbf{i}] = \mathbf{u}[\mathbf{i}] / \mathbf{v}[\mathbf{i}]. \quad (6.14)$$

The convolution product $\mathbf{w} = \mathbf{u} \star \mathbf{v}$ between $\mathbf{u}$ and $\mathbf{v}$ is defined by:

$$\mathbf{w}[\mathbf{i}] = \sum_{\mathbf{j} \in \Omega} \mathbf{u}[\mathbf{i} - \mathbf{j}] \mathbf{v}[\mathbf{j}]. \quad (6.15)$$

In this definition, we assumed periodic boundary conditions. The discrete Fourier transform operator on E is denoted $\mathbf{F}$. It satisfies the fundamental relationship:

$$\mathbf{F}(\mathbf{u} \star \mathbf{v}) = \mathbf{F}(\mathbf{u}) \odot \mathbf{F}(\mathbf{v}). \quad (6.16)$$

The discrete Dirac mass is denoted $\boldsymbol{\delta}$. It satisfies $\mathbf{u} \star \boldsymbol{\delta} = \mathbf{u}$ for all $\mathbf{u} \in E$.

6.2.4 The (inverse)-gamma distribution

A random variable X that is gamma-distributed is denoted by $X \sim \text{Gamma}(a, b)$. Its probability density function (p.d.f.) is defined by:

$$\mathbb{P}(x) = \frac{b^a}{\Gamma(a)} x^{a-1} \exp(-bx), \quad \forall x > 0,$$

where $\Gamma(\cdot)$ is the gamma function. The parameter $a > 0$ is called shape parameter, while $b > 0$ is called inverse scale parameter. If $X \sim \text{Gamma}(a, b)$, then $\mathbb{E}(X) = a/b$ and $\text{var}(X) = a/(b^2)$. For any $c > 0$, if $X \sim \text{Gamma}(a, b)$, then $cX \sim \text{Gamma}(a, cb)$. Let

$$X_i \sim \text{Gamma}(a_i, b), \quad i \in \{1, \dots, n\} \quad (6.17)$$

be independent random variables, then

$$\sum_{i=1}^n X_i \sim \text{Gamma}\left(\sum_{i=1}^n a_i, b\right). \quad (6.18)$$

A random variable that is inverse-gamma distributed is denoted by $X \sim \text{InvGamma}(a, b)$. Its p.d.f. is:

$$\mathbb{P}(x) = \frac{b^a}{\Gamma(a)} x^{-a-1} \exp(-b/x), \quad \forall x > 0.$$

If $X \sim \text{Gamma}(a, b)$ then $1/X \sim \text{InvGamma}(a, 1/b)$.

6.3 Noise model

6.3.1 Modeling the noise

Designing a precise image formation model in the examples given in Figure 6.1 is a hard task. Many different physical phenomena such as diffraction, absorption and scattering with unknown physical quantities are involved. It is in fact likely that good forward models of light propagation are too complicated to use in the frame of inverse problems. This observation leads us to develop a simple phenomenological model. From a practical point of view, a good model should:

- i) Reproduce approximately what is observed in real experiments.
 - ii) Lead to optimization problems that can be solved with a reasonable complexity.
- In this paper, we will be particularly interested in designing *convex problems*.

Let $\mathbf{u}$ denote the ideal image we wish to recover and $\mathbf{u}_0$ denote the degraded image. The first observation is that the noise is *multiplicative* since it is due to attenuation of the excitation light through the sample: the light emitted by a fluorescent sample is approximately proportional to the excitation intensity. The standard modeling of such

noise consists of writing $\mathbf{u}_0 = \mathbf{u} \odot \boldsymbol{\xi}$, where $\boldsymbol{\xi}$ is some random vector. This equation however leads to serious numerical troubles since the maximum a posteriori principle leads to problems of the form:

$$\min_{(\mathbf{u}, \boldsymbol{\xi}) \in E_+ \times E_+, \mathbf{u}_0 = \mathbf{u} \odot \boldsymbol{\xi}} f(\mathbf{u}) + g(\boldsymbol{\xi}), \quad (6.19)$$

where functions f and g are priors on the image and the noise respectively. The constraint set $\{(\mathbf{u}, \boldsymbol{\xi}) \in E_+ \times E_+, \mathbf{u}_0 = \mathbf{u} \odot \boldsymbol{\xi}\}$ is nonconvex and - except for specific f and g - finding the global minimizer of problem (6.19) is therefore out of reach.

To avoid this pitfall, we instead write that:

$$\mathbf{u}_0 = \mathbf{u} \oslash \boldsymbol{\eta}, \quad (6.20)$$

where $\boldsymbol{\eta}$ is the realization of some random vector. The constraint set obtained by using a division instead of a multiplication becomes the linear subspace $\{(\mathbf{u}, \boldsymbol{\eta}) \in E_+ \times E_+, \boldsymbol{\eta} \odot \mathbf{u}_0 = \mathbf{u}\}$. The whole minimization problem

$$\min_{(\mathbf{u}, \boldsymbol{\eta}) \in E_+ \times E_+, \boldsymbol{\eta} \odot \mathbf{u}_0 = \mathbf{u}} f(\mathbf{u}) + g(\boldsymbol{\eta}), \quad (6.21)$$

therefore becomes convex as soon as f and g are convex.

The main difficulty is now to construct a probability distribution function for the random vector $\boldsymbol{\eta}$. Similarly to [107, 106], we assume that the noise is *stationary*, meaning that all translated versions of $\boldsymbol{\eta}$ have the same likelihood as $\boldsymbol{\eta}$. A simple way to generate a stationary random vector consists in writing

$$\boldsymbol{\eta} = \boldsymbol{\psi} \star \boldsymbol{\lambda}, \quad (6.22)$$

where $\boldsymbol{\psi}$ is a convolution filter that depends on the noise structure and $\boldsymbol{\lambda}$ is the realization of a random vector in E with *independent* entries.

In order to specify our model completely, we still need to define a distribution for $\boldsymbol{\lambda}$. Since fluorescence images have nonnegative values, all components of $\boldsymbol{\eta}$ need to be positive. A simple way to ensure this is to set a convolution filter $\boldsymbol{\psi} \neq 0$ with nonnegative entries and to define $\boldsymbol{\lambda}$ as a positive random vector. By doing so, the event $\boldsymbol{\eta}[\mathbf{i}] = 0$ occurs with null probability. Among the positive, log-concave distributions, the gamma distribution has two parameters, allowing to control its mean and variance. This is important to be able varying the attenuation amplitude.

We now specified every piece of our image formation model. Let us recollect everything below.

Proposed image formation model.

We assume that:

$$\mathbf{u}_0 = \mathbf{u} \odot (\boldsymbol{\psi} \star \boldsymbol{\lambda}), \quad \text{where:} \quad (6.23)$$

- $\mathbf{u}_0 \in E_+$ is the observed, noisy image.
- $\mathbf{u} \in E_+$ is the clean image we wish to recover.
- $\boldsymbol{\psi} \in E_+$ is a convolution filter.
- $\boldsymbol{\lambda} \in E_+$ is the realization of a random vector with i.i.d. components following a $\text{Gamma}(a, b)$ distribution.

6.3.2 A few noise properties

The case $\boldsymbol{\psi} = \boldsymbol{\delta}$.

In the specific case $\boldsymbol{\psi} = \boldsymbol{\delta}$, our model is adapted to i.i.d. multiplicative noise. The generative model (6.20) can be rewritten as

$$\mathbf{u}_0 = \mathbf{u} \odot \boldsymbol{\varrho}, \quad (6.24)$$

where $\boldsymbol{\varrho}$ follows an *inverse-gamma* distribution of parameters a and b^{-1} . If we follow a Bayesian point of view, our model is therefore well suited to multiplicative inverse-gamma noise. This departs from the usual modeling used in SAR imaging [9]. However, we will see that this model also performs well in the case $\mathbf{u}_0 = \mathbf{u} \odot \boldsymbol{\eta}$, where $\boldsymbol{\eta}$ is a white gamma noise.

Indicator functions

Let $\omega \subseteq \Omega$ denote a subset of cardinality $|\omega|$. Set

$$\boldsymbol{\psi}[\mathbf{i}] = \begin{cases} 1/|\omega| & \text{if } \mathbf{i} \in \omega, \\ 0 & \text{otherwise.} \end{cases} \quad (6.25)$$

Then, a marginal $\boldsymbol{\eta}[\mathbf{i}]$ is just a sum of i.i.d. random variables with distribution $\text{Gamma}(a, b/|\omega|)$. Therefore $\boldsymbol{\eta}[\mathbf{i}] \sim \text{Gamma}(|\omega|a, b/|\omega|)$ and

$$1/\boldsymbol{\eta}[\mathbf{i}] \sim \text{InvGamma}(|\omega|a, |\omega|/b). \quad (6.26)$$

This simple observation allows to evaluate the mean and variance of the marginals. We get:

$$\mathbb{E} \left(\frac{1}{\boldsymbol{\eta}[\mathbf{i}]} \right) = \frac{|\omega|}{b(|\omega|a - 1)} \quad (6.27)$$

and

$$\text{var} \left(\frac{1}{\boldsymbol{\eta}[\mathbf{i}]} \right) = \frac{|\omega|^2}{b^2(|\omega|a - 1)^2(|\omega|a - 2)}. \quad (6.28)$$

Since we have to set $a > 1$ in order to preserve convexity, we get that for large $|\omega|$, $\mathbb{E}\left(\frac{1}{\eta[\mathbf{i}]}\right) \simeq \frac{1}{ba}$ and $\text{var}\left(\frac{1}{\eta[\mathbf{i}]}\right) \simeq 0$. As a conclusion, the proposed model cannot reproduce large attenuation dynamics, whenever the patterns have a large support. This is a limitation of the model.

The generic case.

When ψ is an arbitrary nonnegative filter, studying the statistical properties of $\mathbf{1} \odot \eta$ is significantly more complicated. Let $\eta[\mathbf{i}] = \sum_{\mathbf{j} \in \Omega} \lambda[\mathbf{j}] \psi[\mathbf{i} - \mathbf{j}]$ be a sum of independent random variables. Each term $\lambda[\mathbf{j}] \psi[\mathbf{i} - \mathbf{j}]$ in the sum follows a $\text{Gamma}(a, b \psi[\mathbf{i} - \mathbf{j}])$ distribution. There is no simple analytical expression for such a sum, see e.g. [177]. We therefore do not investigate further the statistical properties of our model in the general case.

6.3.3 Noise simulations

Figure 6.2 illustrates some realizations of a random vector of type $\mathbf{1} \odot (\psi \star \lambda)$ for various ψ and various a with $b = a$. As can be seen, the noise amplitude cannot vary too much for $a > 1$, when the filter size is large. This property was explained in the previous section. For $0 < a < 1$, the variance of the inverse-gamma distribution is undefined: the stochastic process λ can take isolated high values that dominate all the others. For instance, the top left realization took a single huge value exceeding the largest value tolerated by the computer. This explains why it looks completely gray.

Figure 6.3 illustrates different examples of noisy images that can be generated by model (6.20) for $a > 1$. As can be seen, various degradations resembling what is observed in SPIM imaging can be reproduced.

6.4 Restoration model

6.4.1 MAP based restoration models

In this paragraph, we propose a MAP restoration model. We aim at solving:

$$\max_{\mathbf{u}, \lambda} \mathbb{P}(\mathbf{u}, \lambda | \mathbf{u}_0). \quad (6.29)$$

By using Bayes rule, this amounts to:

$$\max_{\mathbf{u}, \lambda} \frac{\mathbb{P}(\mathbf{u}_0 | \mathbf{u}, \lambda) \mathbb{P}(\mathbf{u}, \lambda)}{\mathbb{P}(\mathbf{u}_0)}. \quad (6.30)$$

Assuming that $\mathbf{u}$ and λ are *independent*, we get that $\mathbb{P}(\mathbf{u}, \lambda) = \mathbb{P}(\mathbf{u})\mathbb{P}(\lambda)$. Moreover, by setting

$$\Xi = \{(\mathbf{u}, \lambda), (\psi \star \lambda) \odot \mathbf{u}_0 = \mathbf{u}\}, \quad (6.31)$$




Figure 6.2 – Examples of a random process of type $1 \otimes (\psi \star \lambda)$, where $\lambda[i] \sim \text{Gamma}(a, a)$ are independent random variables. First column: filter ψ . From the second to the last column: realizations with different values of a : 0.2, 1.1 and 2.




Figure 6.3 – Examples of noisy images generated by model (6.20). For all images, $\lambda[i] \sim \text{Gamma}(a, a)$. Top-Left: original image. Top-Right: $\psi = \delta$ and $a = 2$. Bottom-Left: ψ is a segment and $a = 1.1$. Bottom-Right: ψ is an indicator function of a 4×4 square and $a = 1.1$.

we get

$$\mathbb{P}(\mathbf{u}_0|\mathbf{u}, \boldsymbol{\lambda}) = \begin{cases} 0 & \text{if } (\mathbf{u}, \boldsymbol{\lambda}) \notin \Xi, \\ 1 & \text{otherwise.} \end{cases} \quad (6.32)$$

By taking the negative log in (6.30), we therefore get that the minimizers of (6.30) coincide with the minimizers of:

$$\min_{(\mathbf{u}, \boldsymbol{\lambda}) \in \Xi} -\log(\mathbb{P}(\mathbf{u})) - \log(\mathbb{P}(\boldsymbol{\lambda})). \quad (6.33)$$

In this paper, we use the standard hypothesis that images have a low total variation, which can be expressed as:

$$\mathbb{P}(\mathbf{u}) \propto \exp(-c\|\nabla \mathbf{u}\|_1), \quad (6.34)$$

for some $c > 0$. Since $\boldsymbol{\lambda}$ follows a gamma distribution, the variational problem we end with reads:

$$\min_{(\mathbf{u}, \boldsymbol{\lambda}) \in \Xi} c\|\nabla \mathbf{u}\|_1 + \langle b\boldsymbol{\lambda} - (a-1)\log \boldsymbol{\lambda}, \mathbf{1} \rangle. \quad (6.35)$$

Finally, by expressing $\mathbf{u}$ in terms of $\boldsymbol{\lambda}$, we obtain:

$$\min_{\boldsymbol{\lambda} \in E_+} c\|\nabla(\mathbf{u}_0 \odot (\boldsymbol{\psi} \star \boldsymbol{\lambda}))\|_1 + \langle b\boldsymbol{\lambda} - (a-1)\log \boldsymbol{\lambda}, \mathbf{1} \rangle. \quad (6.36)$$

This functional is convex for $b > 0$ and $a \geq 1$.

6.4.2 Simplifying the model

Model (6.36) contains three explicit parameters (a, b, c) and one implicit parameter $\|\boldsymbol{\psi}\|_1$: the ℓ^1 -norm of $\boldsymbol{\psi}$. Tuning 4 parameters is very hard in practice and we propose to simplify parameter estimation in this section.

Proposition 6.4.1. *Problem (6.36) admits a unique minimizer if $c \geq 0$, $a > 1$ and $b > 0$.*

Proof. Set

$$g(\boldsymbol{\lambda}) = c\|\nabla(\mathbf{u}_0 \odot (\boldsymbol{\psi} \star \boldsymbol{\lambda}))\|_1 + \langle b\boldsymbol{\lambda} - (a-1)\log \boldsymbol{\lambda}, \mathbf{1} \rangle. \quad (6.37)$$

First notice that the set E_+ is convex closed. Moreover, $g(\boldsymbol{\lambda})$ is strictly convex and closed. Indeed, for $a > 1$, the function defined on E_+ by $\boldsymbol{\lambda} \mapsto -(a-1)\langle \log \boldsymbol{\lambda}, \mathbf{1} \rangle$ is strictly convex since its Hessian $(a-1)\text{diag}(\mathbf{1} \oslash (\boldsymbol{\lambda}^2))$ is positive definite for all $\boldsymbol{\lambda} \in E_+$. Problem (6.36) consists of minimizing a strictly convex function over a convex closed set. This is sufficient to ensure uniqueness of a minimizer if it exists. The existence of a solution is granted by noticing that $g(\boldsymbol{\lambda})$ is continuous on the interior of its domain since it is convex. Moreover, it is coercive since the 1D function $t \mapsto -(a-1)\log(t) + bt$ goes to $+\infty$ as $t \rightarrow +\infty$. $\square$

In order to reduce the number of parameters, it is possible to divide the cost function in (6.36) by $c\|\boldsymbol{\psi}\|_1$. By doing so, it is straightforward to see that the minimizer of (6.36) is equal to the minimizer of:

$$\min_{\boldsymbol{\lambda} \in E_+} \|\nabla(\mathbf{u}_0 \odot (\tilde{\boldsymbol{\psi}} \star \boldsymbol{\lambda}))\|_1 + \langle \beta \boldsymbol{\lambda} - \alpha \log \boldsymbol{\lambda}, \mathbf{1} \rangle. \quad (6.38)$$

where $\tilde{\boldsymbol{\psi}} = \frac{\boldsymbol{\psi}}{\|\boldsymbol{\psi}\|_1}$, $\alpha = (a-1)/(c\|\boldsymbol{\psi}\|_1)$ and $\beta = b/c\|\boldsymbol{\psi}\|_1$. Model (6.38) still contains two parameters. The following proposition shows that it can be further reduced to 1 parameter.

Proposition 6.4.2. *Let $\boldsymbol{\lambda}(\alpha, \beta)$ denote the minimizer of (6.38). Then, for all $t > 0$*

$$\boldsymbol{\lambda}(t\alpha, \beta) = t\boldsymbol{\lambda}(\alpha, \beta). \quad (6.39)$$

Proof. Let $\tilde{\boldsymbol{\lambda}} = t\boldsymbol{\lambda}(\alpha, \beta)$. The minimizer $\boldsymbol{\lambda}(\alpha, \beta)$ of (6.38) satisfies:

$$0 \in \tilde{\mathbf{A}}^* \partial_{\|\cdot\|_1}(\tilde{\mathbf{A}}\boldsymbol{\lambda}(\alpha, \beta)) + \beta \mathbf{1} - \alpha \mathbf{1} \odot \boldsymbol{\lambda} + \mathcal{N}_{E_+}(\boldsymbol{\lambda}), \quad (6.40)$$

where $\tilde{\mathbf{A}}$ is the linear operator defined for all $\boldsymbol{\lambda} \in E$ by

$$\tilde{\mathbf{A}}\boldsymbol{\lambda} = \nabla(\mathbf{u}_0 \odot (\tilde{\boldsymbol{\psi}} \star \boldsymbol{\lambda})). \quad (6.41)$$

Therefore $\tilde{\boldsymbol{\lambda}}$ obeys:

$$0 \in \tilde{\mathbf{A}}^* \partial_{\|\cdot\|_1}(\tilde{\mathbf{A}}\tilde{\boldsymbol{\lambda}}/t) + \beta \mathbf{1} - \alpha t \mathbf{1} \odot \tilde{\boldsymbol{\lambda}} + \mathcal{N}_{E_+}(\tilde{\boldsymbol{\lambda}}). \quad (6.42)$$

Since $\partial_{\|\cdot\|_1}(\tilde{\mathbf{A}}\tilde{\boldsymbol{\lambda}}/t) = \partial_{\|\cdot\|_1}(\tilde{\mathbf{A}}\tilde{\boldsymbol{\lambda}})$, this is still equivalent to:

$$0 \in \tilde{\mathbf{A}}^* \partial_{\|\cdot\|_1}(\tilde{\mathbf{A}}\tilde{\boldsymbol{\lambda}}) + \beta \mathbf{1} - \alpha t \mathbf{1} \odot \tilde{\boldsymbol{\lambda}} + \mathcal{N}_{E_+}(\tilde{\boldsymbol{\lambda}}). \quad (6.43)$$

This inclusion characterizes the minimizer of problem (6.38), where α is replaced by $t\alpha$. $\square$

Proposition 6.4.2 indicates that only β really matters. The value α only allows to set the amplitude of the solution. In fluorescence microscopy, this amplitude depends on the excitation intensity which is usually unknown. Fluorescence images are therefore not used for quantitative measures, and only contrasts matter. As a conclusion, the model could be further simplified by setting $\alpha = 1$. We will see in Proposition 6.4.3 that setting $\beta = \alpha$ allows preserving the harmonic mean of $\mathbf{u}_0$ and this choice therefore seems more natural. The following model summarizes the simplifications we have proposed in this section.

Proposed variational formulation.

The variational problem we study in this paper reads:

$$\min_{\boldsymbol{\lambda} \in E_+} \|\nabla(\mathbf{u}_0 \odot (\boldsymbol{\psi} \star \boldsymbol{\lambda}))\|_1 + \alpha \langle \boldsymbol{\lambda} - \log(\boldsymbol{\lambda}), \mathbf{1} \rangle, \quad (6.44)$$

where $\boldsymbol{\psi} \in E_+$, $\|\boldsymbol{\psi}\|_1 = 1$ and $\alpha > 0$.

6.4.3 Duality and mean preservation

In this paragraph, we prove that model (6.44) preserves the *harmonic* mean of the input image $\mathbf{u}_0$ in the specific case $\boldsymbol{\psi} = \boldsymbol{\delta}$. In the generic case, we prove that the weighted harmonic mean of $\boldsymbol{\lambda}$ is constant.

Proposition 6.4.3. *Let $\boldsymbol{\lambda}$ be the solution of (6.44). Assume that $\mathbf{u}_0 > 0$ and let $\bar{\boldsymbol{\lambda}}$ be any vector such that $\boldsymbol{\psi} \star \bar{\boldsymbol{\lambda}} = \mathbf{1} \oslash \mathbf{u}_0$ ¹. Then*

$$\sum_{\mathbf{i} \in \Omega} \frac{\bar{\boldsymbol{\lambda}}[\mathbf{i}]}{\boldsymbol{\lambda}[\mathbf{i}]} = \sum_{\mathbf{i} \in \Omega} \bar{\boldsymbol{\lambda}}[\mathbf{i}], \quad (6.45)$$

In the particular case $\boldsymbol{\psi} = \boldsymbol{\delta}$ - thus $\mathbf{u} = \mathbf{u}_0 \odot \boldsymbol{\lambda}$ - this yields:

$$\sum_{\mathbf{i} \in \Omega} \frac{1}{\mathbf{u}[\mathbf{i}]} = \sum_{\mathbf{i} \in \Omega} \frac{1}{\mathbf{u}_0[\mathbf{i}]} \quad (6.46)$$

Proof. Let us first use Fenchel-Rockafellar duality [197], in order to derive the dual problem of (6.44). Let $\mathbf{A}$ be the linear operator defined, for all $\boldsymbol{\lambda} \in E$, by $\mathbf{A}\boldsymbol{\lambda} = \nabla(\mathbf{u}_0 \odot (\boldsymbol{\psi} \star \boldsymbol{\lambda}))$. We have:

$$\begin{aligned} & \min_{\boldsymbol{\lambda} \in E_+} \|\mathbf{A}\boldsymbol{\lambda}\|_1 + \alpha \langle \boldsymbol{\lambda} - \log(\boldsymbol{\lambda}), \mathbf{1} \rangle \\ &= \min_{\boldsymbol{\lambda} \in E_+} \sup_{\mathbf{q} \in V, \|\mathbf{q}\|_\infty \leq 1} \langle \mathbf{A}\boldsymbol{\lambda}, \mathbf{q} \rangle + \alpha \langle \boldsymbol{\lambda} - \log(\boldsymbol{\lambda}), \mathbf{1} \rangle \\ &= \sup_{\mathbf{q} \in V, \|\mathbf{q}\|_\infty \leq 1} \inf_{\boldsymbol{\lambda} \in E_+} \langle \mathbf{A}\boldsymbol{\lambda}, \mathbf{q} \rangle + \alpha \langle \boldsymbol{\lambda} - \log(\boldsymbol{\lambda}), \mathbf{1} \rangle. \end{aligned}$$

The minimizer of the inner-problem satisfies:

$$\mathbf{A}^* \mathbf{q} + \alpha \mathbf{1} - \alpha \mathbf{1} \oslash \boldsymbol{\lambda} + \mathcal{N}_{E_+}(\boldsymbol{\lambda}) \ni 0. \quad (6.47)$$

Since the logarithm is a barrier function for the set E_+ , the minimizer belongs to the interior of E_+ and the optimality conditions simplify to:

$$\mathbf{A}^* \mathbf{q} + \alpha \mathbf{1} - \alpha \mathbf{1} \oslash \boldsymbol{\lambda} = 0. \quad (6.48)$$

Therefore, the primal-dual pair $(\boldsymbol{\lambda}, \mathbf{q})$ satisfies:

$$\boldsymbol{\lambda} = \alpha \mathbf{1} \oslash (\mathbf{A}^* \mathbf{q} + \alpha \mathbf{1}). \quad (6.49)$$

By injecting this result in the inner-problem, we obtain the dual problem:

$$\sup_{\|\mathbf{q}\|_\infty \leq 1} \langle \log(\alpha \mathbf{1} + \mathbf{A}^* \mathbf{q}), \mathbf{1} \rangle + \alpha (1 - \langle \log(\alpha \mathbf{1}), \mathbf{1} \rangle), \quad (6.50)$$

¹A sufficient condition for existence of such a $\bar{\boldsymbol{\lambda}}$ is that the Fourier transform $\hat{\boldsymbol{\psi}}$ does not vanish.

which share the minimizers of

$$\sup_{\|\mathbf{q}\|_\infty \leq 1} \langle \log(\alpha \mathbf{1} + \mathbf{A}^* \mathbf{q}), \mathbf{1} \rangle. \quad (6.51)$$

Then, remark that:

$$\mathbf{1} \oslash \boldsymbol{\lambda} = \frac{\mathbf{A}^* \mathbf{q}}{\alpha} + \mathbf{1}. \quad (6.52)$$

By construction, $\bar{\boldsymbol{\lambda}} \in \text{Ker}(\mathbf{A})$. Since $\text{Ker}(\mathbf{A}) = \text{Im}(\mathbf{A}^*)^\perp$, we get that $\langle \mathbf{A}^* \mathbf{q}, \bar{\boldsymbol{\lambda}} \rangle = 0$ for all $\mathbf{q} \in V$. Therefore:

$$\langle \mathbf{1} \oslash \boldsymbol{\lambda}, \bar{\boldsymbol{\lambda}} \rangle = \langle \mathbf{1}, \bar{\boldsymbol{\lambda}} \rangle. \quad (6.53)$$

If $\boldsymbol{\psi} = \boldsymbol{\delta}$, then $\bar{\boldsymbol{\lambda}} = \mathbf{1} \oslash \mathbf{u}_0$ and $\boldsymbol{\lambda} = \mathbf{u} \oslash \mathbf{u}_0$. This yields (6.46). $\square$

6.4.4 Relationships to other models in the white noise regime

When $\boldsymbol{\psi} = \boldsymbol{\delta}$ is a Dirac mass, our model is adapted to white noise denoising. Problem (6.44) can be rephrased in terms of $\mathbf{u}$ as:

$$\min_{\mathbf{u} \in E_+} \|\nabla \mathbf{u}\|_1 + \alpha \langle \mathbf{u} \oslash \mathbf{u}_0 - \log(\mathbf{u} \oslash \mathbf{u}_0), \mathbf{1} \rangle. \quad (6.54)$$

Comparatively, the convex Shi-Osher or Teuber-Steidl model proposed in [216, 210] is:

$$\min_{\mathbf{u} \in E_+} \|\nabla \mathbf{u}\|_1 + \alpha (\langle \mathbf{u}, \mathbf{1} \rangle - \langle \mathbf{u}_0, \log(\mathbf{u}) \rangle). \quad (6.55)$$

The nonconvex Aubert-Aujol model proposed in [9] is:

$$\min_{\mathbf{u} \in E_+} \|\nabla \mathbf{u}\|_1 + \alpha \langle \mathbf{u}_0 \oslash \mathbf{u} + \log(\mathbf{u}), \mathbf{1} \rangle. \quad (6.56)$$

All models are proved to preserve the geometric mean, while our preserves the harmonic mean. It is unclear to us what property is better. We will however see in the numerical experiments that our model better preserves contrasts.

6.5 Solvers for the proposed model (6.44)

In this section we provide an algorithm to solve problem (6.44). A large amount of numerical approaches were developed recently to solve non-smooth convex problems of that type. Problem (6.44) can be rewritten as:

$$\min_{\boldsymbol{\lambda} \in E} J(\boldsymbol{\lambda}) = F(\mathbf{A}\boldsymbol{\lambda}) + G(\boldsymbol{\lambda}), \quad (6.57)$$

where $\mathbf{A} : E \rightarrow V$ defined by $\mathbf{A}\boldsymbol{\lambda} = \nabla(\mathbf{u}_0 \odot (\boldsymbol{\psi} \star \boldsymbol{\lambda}))$, $F : V \rightarrow \mathbb{R}$ is defined as $F(\mathbf{q}) = \|\mathbf{q}\|_1$ and $G : E \rightarrow \mathbb{R} \cup \{+\infty\}$ is defined as $G(\boldsymbol{\lambda}) = \alpha \langle \boldsymbol{\lambda} - \log \boldsymbol{\lambda}, \mathbf{1} \rangle + \chi_{E_+}(\boldsymbol{\lambda})$. This reformulation perfectly fits the framework of the first-order primal-dual algorithm proposed in [54]. It is described in Algorithm 9.

Algorithm 9: First order primal-dual algorithm [54]

Input : ϵ : the desired precision,
 $(\mathbf{x}_0, \mathbf{y}_0) \in E_+ \times V$: a starting point.
Output: an approximate solution of problem (6.57).
Init. : Choose $\tau, \sigma > 0$ such that $\sigma\tau\|\mathbf{A}\|^2 < 1$,
set $\theta = 1$, $k = 0$ and $\bar{\mathbf{x}}_0 = \mathbf{x}_0$.
while *Convergence criterion* $> \epsilon$ **do**
 $\mathbf{y}_{k+1} = \text{prox}_{\sigma F^*}(\mathbf{y}_k + \sigma \mathbf{A} \bar{\mathbf{x}}_k)$
 $\mathbf{x}_{k+1} = \text{prox}_{\tau G}(\mathbf{x}_k - \tau \mathbf{A}^* \mathbf{y}_{k+1})$
 $\bar{\mathbf{x}}_{k+1} = \mathbf{x}_{k+1} + \theta(\mathbf{x}_{k+1} - \mathbf{x}_k)$
 $k = k + 1$

Algorithm 9 generates a feasible sequence of iterates $(\mathbf{x}_k)_{k \in \mathbb{N}}$ that converges to the minimizer $\mathbf{x}^*$ of (6.57). Moreover the iterates satisfy $J(\mathbf{x}_k) - J(\mathbf{x}^*) = O\left(\frac{1}{k}\right)$ which is arguably optimal. The proximity operators of F^* and G have a simple closed form expression:

$$\text{prox}_{\sigma F^*}(\mathbf{q})[\mathbf{i}] = \frac{\mathbf{q}[\mathbf{i}]}{\max(1, |\mathbf{q}[\mathbf{i}]|)}, \quad (6.58)$$

where the division is meant component-wise on each coordinate of $\mathbf{q}$ and

$$\text{prox}_{\tau G}(\mathbf{x}) = \frac{1}{2} \left(\mathbf{x} - \tau \alpha \mathbf{1} + \sqrt{(\mathbf{x} - \tau \alpha \mathbf{1})^2 + 4\tau \alpha \mathbf{1}} \right). \quad (6.59)$$

This algorithm was implemented on a GPU using CUDA. Except convolutions, all operators appearing in this algorithm are pixel-wise operations which are the perfect setting for GPU. To implement convolutions we used the cuFFT library. In all experiments, the GPU implementation has been run in double precision on a NVIDIA Tesla K20c containing 2496 CUDA cores and 5GB internal memory.

6.6 Numerical experiments

In this section, we first provide some numerical results on synthetic images, where ground-truth is available and then turn to real images.

6.6.1 The case of white noise

In this paragraph, we first generate noisy images by the formula:

$$\mathbf{u}_0 = \mathbf{u} \odot \boldsymbol{\eta}, \quad (6.60)$$

where $\boldsymbol{\eta} \in E$ is a random vector with i.i.d. components. The distribution of the marginals $\boldsymbol{\eta}[\mathbf{i}]$ is set as a gamma distribution $\text{Gamma}(a, b)$ in Figure 6.4 and as an inverse-gamma

	a	10	20	50
Degraded	SNRr	10.6	13.2	17.1
	SSIMr	0.69	0.78	0.87
[216, 210]	SNRr	15.5	16	16.6
	SSIMr	0.78	0.81	0.83
MSNR	SNRr	17.2	18.7	20.7
	SSIMr	0.79	0.85	0.90

Table 6.1 – SNRr and SSIMr for image formation model $\mathbf{u}_0 = \mathbf{u} \odot \boldsymbol{\eta}$, where $\boldsymbol{\eta}$ is a gamma distributed white noise with $\boldsymbol{\eta}[1] \sim \text{Gamma}(a, a)$

distribution in Figure 6.5. Our model is supposed to be well adapted only to the inverse-gamma distribution. We provide comparisons with the denoising model (6.55) proposed in [216, 210]. The minimizer of (6.44) is referred to as MSNR, for Multiplicative Stationary Noise Removal.

In order to compare the restoration results, we evaluate the rescaled Signal to Noise Ratio, denoted SNRr and defined by:

$$\text{SNRr}(\mathbf{u}, \mathbf{u}_0) = -\min_{a \in \mathbb{R}} \log_{10} \left(\frac{\|a\mathbf{u} - \mathbf{u}_0\|_2^2}{\|\mathbf{u}_0\|_2^2} \right). \quad (6.61)$$

The reason to use this image quality measure is that the models only yield results valid up to a multiplicative constant. We also compute another image quality measure called SSIMr, defined as the Standard Structural Similarity Index Measure [240] applied to the rescaled image $a\mathbf{u}$ where a is optimized in (6.61). In all experiments, the model parameters are set so as to maximize the SNRr.

As can be seen in Figures 6.4 and 6.5 and Tables 6.1 and 6.2, our model provides significantly higher rescaled SNRs and SSIMs than model (6.55). When having a close look at the images, it can be seen however that the details are preserved similarly in both approaches. The main difference lies in the fact that model (6.55) does not preserve the contrasts as well as the proposed model. Moreover one can remark that the proposed restoration model - that is designed for inverse-Gamma distribution - performs well for Gamma-distributed. An explanation of this phenomenon can be found in [120, 157]. Morally, data-fitting terms derived from noise distributions can be neglected against the TV regularization term which strongly promotes a very specific type of images.

6.6.2 The case of structured noise

We now investigate the efficiency of the method proposed in this paper versus VSNR (Variational Stationary Noise Remover) [107, 106]. We use the 256×256 drosophila image see Figure 6.3 top left, corrupted with noise generated using two different filters $\boldsymbol{\psi}$. The first one is an anisotropic Gaussian and the second one is the indicator of a square. The data term parameter α is set so as to maximize the SNRr. Results are displayed Figures 6.6 and 6.7. The proposed method provides higher rescaled SNR and




Figure 6.4 – Denoising experiment for image formation model $\mathbf{u}_0 = \mathbf{u} \odot \boldsymbol{\eta}$, where $\boldsymbol{\eta}$ is a gamma distributed white noise with $\boldsymbol{\eta}[1] \sim \text{Gamma}(50, 50)$. The SNRr for MSNR is stable after 50 iterations. Computing time MATLAB - CPU: 0.15s.

	a	10	20	50
Degraded	SNRr	9.84	12.8	16.9
	SSIMr	0.66	0.77	0.87
[216, 210]	SNRr	15.2	16	16.8
	SSIMr	0.79	0.81	0.85
MSNR	SNRr	17.4	18.9	20.8
	SSIMr	0.79	0.85	0.91

Table 6.2 – SNRr and SSIMr for image formation model $\mathbf{u}_0 = \mathbf{u} \odot \boldsymbol{\eta}$, where $\boldsymbol{\eta}$ is a gamma distributed white noise with $\boldsymbol{\eta}[1] \sim \text{Gamma}(a, a)$




Figure 6.5 – Denoising experiment for image formation model $\mathbf{u}_0 = \mathbf{u} \oslash \boldsymbol{\eta}$, where $\boldsymbol{\eta}$ is a gamma distributed white noise with $\boldsymbol{\eta}[1] \sim \text{Gamma}(20, 20)$. The SNRr for MSNR is stable after 50 iterations. Computing time MATLAB - CPU: 0.15s.

images containing finer details and more natural contrasts. The computing times are also very appealing.

6.6.3 A few results on real images

To end up with our numerical validation, we perform experiments on real images. We consider two different images and compare the results on these images obtained with VSNR and the method proposed in this paper, see Figures 6.8 and 6.10. Following the previous works [107, 106], we set ψ as an anisotropic Gaussian filter elongated in the stripes direction. This choice allows to accurately describe the noise in the frequency domain. In particular, the filter's smoothness ensures that only low frequencies will be changed by the variational model (6.44). This is a key feature to preserve fine details in the image.

The results can only be compared qualitatively since no ground truth is available. A quick inspection of the results advocates for the method proposed in this paper. Images restored with VSNR suffer from smear artifacts at the locations of stripes, especially in dark region. Images proposed by the MSNR do not suffer from these artifacts and overall have a better contrast see Figure 6.9.




Figure 6.6 – Denoising simulation with structured noise. The image is corrupted with ψ a segment and $a = 1.1$ see Figure 6.3 bottom-left. The SNRr is stable after 50 iterations. Computing time GPU: 0.03s.




Figure 6.7 – Denoising simulation with structured noise. The image is corrupted with ψ the indicator function 4×4 square and $a = 1.1$ see Figure 6.3 bottom-right. The SNRr is stable after 40 iterations. Computing time GPU: 0.025s.




Figure 6.8 – Denoising experiment on real data: 1040×1390 . The degraded image is a XY plane of an skin dermis acquired using a SPIM. The image is used here by courtesy of Claire Jardet from Genoskin. The iterates stabilize visually after 300 iterations. Computing time GPU: 3.3s (about 11ms per iteration)




Figure 6.9 – Zoom on Figure 6.8. Observe in the top figure that VSNR algorithm smears the black region. In the bottom figure, observe that the contrasts seem slightly more pronounced in MSNR.




Figure 6.10 – Denoising experiment on real data : 499×561 . The degraded image is a XY plane of spheroid acquired using a SPIM. The iterates stabilize visually after 500 iterations. Computing time GPU: 3s (about 6ms per iteration)

7

Contrast Invariant SNR

Résumé : *Ce chapitre est issu de [244], en cours de préparation et en collaboration avec Pierre Weiss et Yiqiu Dong. Dans ce chapitre nous concevons une mesure de similarité d'images qui est invariante aux changements de contrastes locaux, qui constituent un modèle simple de changement d'illumination. A partir de deux images, l'algorithme construit une image qui est la plus proche de la première mais qui partage le même arbre de composantes que la deuxième. Ce problème peut-être vu comme un problème d'optimisation spécifique appelé régression isotonique. Nous donnons quelques propriétés analytiques de la solution de ce problème. Nous avons aussi élaboré une procédure d'optimisation de premier ordre adaptée ainsi qu'une analyse complète de sa complexité. La méthode proposée se montre en pratique plus efficace et fiable que les méthodes existantes basées sur la méthode des points intérieurs. L'algorithme a aussi des applications potentielles en détection de changements, traitement des images couleurs et fusion d'images.*

Abstract: *We design an image quality measure independent of local contrast changes, which constitute simple models of illumination changes. Given two images, the algorithm provides the image closest to the first one that has the component tree of the second. This problem can be cast as a specific convex program called isotonic regression. We provide a few analytic properties of the solutions to this problem. We also design a tailored first order optimization procedure together with a full complexity analysis. The proposed method turns out to be practically more efficient and reliable than the best existing algorithms based on interior point methods. The algorithm has potential applications in change detection, color image processing or image fusion.*

Contents

7.1	Introduction	213
7.2	Existing approaches	213
7.2.1	Contrast equalization	213
7.2.2	Optimal linear and affine maps	214
7.2.3	Optimal global contrast change	214
7.3	Optimal local contrast changes	215
7.3.1	Topographic maps and local contrast changes	215
7.3.2	A compact representation of local contrast changes	215
7.3.3	A new convex minimization algorithm	217
7.3.4	Complexity analysis	220
7.3.5	Analytical properties of the minimizer	222
7.3.6	Analytic properties of SNR_{loc}	223
7.4	Robust contrast changes	224
7.4.1	The principle	224
7.4.2	A minimization algorithm	224
7.5	Numerical results	225
7.5.1	Image comparison and change detection	225
7.5.2	A large panel of illumination changes	228
7.5.3	Robust contrast change	228

7.1 Introduction

Invariance to illumination conditions is often a key element for the success of image processing algorithms. The whole field of mathematical morphology is based on contrast invariance [209]. The structural similarity index [240] - one of the most popular image quality measures - also strongly relies on a partial invariance to illumination changes.

In this paper, we introduce a novel algorithm that allows comparing two images in a way independent of *local contrast changes*. Most of the works dedicated to illumination invariance consist in decomposing the image into patches and normalizing means and variances on the patches. Here, we follow a different trail and consider that two images differ by a local contrast change if they share the same *component tree* [47, 49]. Given a reference image $u_0 : \Omega \rightarrow \mathbb{R}$ and another image $u_1 : \Omega \rightarrow \mathbb{R}$, we propose to measure their similarity by computing the following value:

$$\Delta_{loc}(u_1, u_0) = \min_{T \in \mathcal{T}} \|u_0 - T(u_1)\|_2^2, \quad (7.1)$$

where $\mathcal{T}$ is the set of local contrast changes. This amounts to finding the best match with u_0 among all images that have the same component tree as u_1 . The locally contrast invariant signal-to-noise-ratio is defined by:

$$SNR_{loc}(u_1, u_0) = -10 \log_{10}(\Delta_{loc}(u_1, u_0) / \|u_0\|_2^2). \quad (7.2)$$

Let T^* denote the optimal contrast change in problem (7.1). The image $u^* = T^*(u_1)$ has a geometry identical to u_1 , with the contrast of u_0 .

We introduce an efficient algorithm based on convex programming to solve (7.1) and provide a full analysis of its complexity. The proposed method has potential applications in different fields such as image quality measure, change detection [245], image fusion or color image processing [48, 12, 97].

7.2 Existing approaches

Various approaches are commonly used to compare two images u_0 and u_1 independently of illumination variations. We briefly describe a few of them below. A Matlab implementation of each method is provided here <http://www.math.univ-toulouse.fr/~weiss/PageCodes.html>.

7.2.1 Contrast equalization

Probably the most common approach consists in equalizing histograms, i.e. changing the gray-values of u_1 in such a way that the resulting histogram matches approximately that of u (see e.g. [83, 36]). This approach suffers from the fact that the image geometry is completely forgotten: histograms only account for gray-level distributions and not geometrical features such as edges, textures,... An example of two images with identical histogram and different geometrical contents is provided in Fig. 7.1.




Figure 7.1 – Two images with different geometrical features can have an identical histogram.

7.2.2 Optimal linear and affine maps

The set $\mathcal{T}$ in problem (7.1) can be replaced by any class of transformations that describe changes of illuminations. Probably the simplest classes $\mathcal{T}$ are the set of linear maps $T(u) = au$ or the set of affine maps $T(u) = au + b$, where a and b are scalars. The solution of both problems can be computed explicitly in terms of u_0 and u_1 . The same approach can be used locally and the L^2 -norm can be replaced by a weighted L^2 -norm. This idea is the basis of the Structural Similarity Index Measure (SSIM).

7.2.3 Optimal global contrast change

A richer set of transformations $\mathcal{T}$ is that of global contrast changes. Two images u_0 and u_1 are said to differ by a global contrast change if there exists a non increasing function $T : \mathbb{R} \rightarrow \mathbb{R}$ such that $T(u_1) = u_0$. Finding the best global contrast change amounts to solving:

$$\Delta_{glo}(u_1, u_0) = \min_{T: \mathbb{R} \rightarrow \mathbb{R}, \text{ non increasing}} \frac{1}{2} \|T(u_1) - u_0\|_2^2. \quad (7.3)$$

We let

$$SNR_{glo}(u_1, u_0) = -10 \log_{10}(\Delta_{glo}(u_1, u_0) / \|u_0\|_2^2)$$

denote the globally contrast invariant SNR. Problem (7.3) is a simple case of isotonic regression [96, 27]. Due to the simple structure of the constraint set, this problem can be solved in $O(n)$ operations using active sets type methods.

Unfortunately, global contrast changes do not capture all the complexity of illumination changes: in most applications, the variations are local.

7.3 Optimal local contrast changes

7.3.1 Topographic maps and local contrast changes

Mathematical morphology emerged with the works of Matheron [166]. Therein, he proposed to analyze an image u by using geometric operations on its (upper)-level sets $\chi_\lambda = \{x \in \Omega, u(x) \geq \lambda\}$ or its level lines $LL_\lambda = \{x \in \Omega, u(x) = \lambda\}$ ¹. The level-sets and level lines are geometrical features invariant to global contrast changes. In order to obtain a representation invariant to local contrast changes, it is possible to consider the *connected components* of these objects. This idea was proposed and detailed thoroughly in [47, 176]. The connected components, thanks to the inclusion relationships, can be encoded in a tree structure called component tree, useful in many applications [49, 181]. This tree structure allows formalizing the notion of local contrast change: two images differ by a local contrast change if they share the same component tree.

This definition has many assets. In particular, it allows working both in the discrete and continuous settings. For our purposes however, we will use a simpler definition valid for discrete images only. Let Ω denote a discrete set of n vertices (pixels) endowed with a neighborhood relationship $\mathcal{N} : \Omega \rightarrow \mathcal{P}(\Omega)$, the power set of Ω . For each pixel $x \in \Omega$, the set $\mathcal{N}(x)$ is the set of all neighbors of x . We assume that the number of elements of $\mathcal{N}(x)$ denoted $|\mathcal{N}(x)|$ is bounded independently of x by a constant $c_{\max}$. In our implementation, we simply use 4-connectivity and therefore $c_{\max} = 4$. We can now define local contrast changes.

Definition 7.3.1. Let u and u_1 denote two images. They are said to differ by a local contrast change if the order relationship between adjacent pixels is the same for each image. More precisely, $\forall x \in \Omega$ and $\forall y \in \mathcal{N}(x)$:

$$\text{sign}(u(x) - u(y)) = \text{sign}(u_1(x) - u_1(y)), \quad (7.4)$$

with the convention $\text{sign}(0) = 0$.

7.3.2 A compact representation of local contrast changes

Problem (7.1) can be rephrased as follows:

$$\min_{u \in \mathcal{U}_1} \frac{1}{2} \|u_0 - u\|_2^2, \quad (7.5)$$

where $\mathcal{U}_1$ is the set of images satisfying condition (7.4). Let $(\omega_i)_{1 \leq i \leq p}$ denote the connected components of the level lines of u_1 . These $p \leq n$ components form a partition of Ω . The first step of our algorithm is to construct the sets $(\omega_i)_{1 \leq i \leq p}$ and a directed acyclic graph $G = (V, E)$ from the image u_1 . The set $V = (v_1, \dots, v_p)$ are the vertices of this graph and represent the sets $(\omega_i)_{1 \leq i \leq p}$. The set $E = (e_1, \dots, e_m)$ are the edges of the graph. Edge $e_k \in E$ is an ordered pair of vertices written $e_k = (I(k), J(k))$ going from

¹We deliberately simplify the definition of level lines here.




Figure 7.2 – Illustration of the graph construction

vertex $I(k)$ to vertex $J(k)$. Such an edge exists if the sets $\omega_{I(k)}$ and $\omega_{J(k)}$ are connected via the neighborhood relationship $\mathcal{N}$ and if $u_1(\omega_{J(k)}) > u_1(\omega_{I(k)})$. The graph G can be encoded through an incidence matrix (or more precisely its transpose) $A \in \mathbb{R}^{m \times p}$. Each row of this matrix describes an edge with the convention $A(k, I(k)) = -1$ and $A(k, J(k)) = 1$ and all the other coefficients of row k are null. A simple 3×3 image u_1 , the associated regions $(\omega_i)_{1 \leq i \leq 4}$, graph and incidence matrix are represented in Figure 7.2.

The list of regions $(\omega_i)_{1 \leq i \leq p}$ can be constructed in $O(n)$ operations using flood fill algorithms. The graph or matrix A can be constructed in $O(n \log(n))$ operations. The idea is to first scan all edges in $\mathcal{N}$ to construct a preliminary matrix $\tilde{A}$ with repetitions. For instance, region ω_1 is connected three times to ω_3 (see arrows in Figure 7.2a). The complexity of this part is $O(c_{\max} n)$. Then, the repetitions can be suppressed in $O(n \log(n))$ operations using a quicksort algorithm.

The set $\mathcal{U}_1$ can now be described compactly as follows:

$$\mathcal{U}_1 = \{u : \Omega \rightarrow \mathbb{R}, u|_{\omega_i} = \alpha_i, 1 \leq i \leq p, A\alpha > 0\}.$$

This set is convex but not closed. The solution of (7.5) might therefore not exist. To ensure existence, we replace (7.5) by

$$\min_{u \in \mathcal{U}_1} \frac{1}{2} \|u_0 - u\|_2^2, \quad \text{where} \tag{7.6}$$

$$\bar{\mathcal{U}}_1 = \{u : \Omega \rightarrow \mathbb{R}, u|_{\omega_i} = \alpha_i, 1 \leq i \leq p, A\alpha \geq 0\}.$$

7.3.3 A new convex minimization algorithm

In this section, we concentrate on the numerical resolution of problem (7.6). We first simplify it as follows

$$\begin{aligned} \min_{u \in \mathcal{U}_1} \frac{1}{2} \sum_{x \in \Omega} (u_0(x) - u(x))^2 \\ = \min_{u \in \mathcal{U}_1} \frac{1}{2} \sum_{i=1}^p \sum_{x \in \omega_i} (u_0(x) - \alpha_i)^2. \end{aligned}$$

Next, let $\beta_i = \frac{1}{|\omega_i|} \sum_{x \in \omega_i} u_0(x)$ denote the mean of u_0 over region ω_i . We get:

$$\begin{aligned} \sum_{x \in \omega_i} (u_0(x) - \alpha_i)^2 &= \sum_{x \in \omega_i} (u_0(x) - \beta_i + \beta_i - \alpha_i)^2 \\ &= \sum_{x \in \omega_i} (u_0(x) - \beta_i)^2 + (\beta_i - \alpha_i)^2 + 2(u_0(x) - \beta_i)(\beta_i - \alpha_i) \\ &= |\omega_i| \text{Var}_{\omega_i}(u_0) + |\omega_i|(\beta_i - \alpha_i)^2, \end{aligned}$$

where $\text{Var}_{\omega_i}(u_0)$ is the variance u_0 over region ω_i . Therefore, the problem reads

$$\min_{A\alpha \geq 0} \frac{1}{2} \sum_{i=1}^p |\omega_i|(\beta_i - \alpha_i)^2 + |\Omega| \text{Var}_{\Omega}(u_0).$$

By letting $w \in \mathbb{R}^p$ denote the vector with components $w_i = |\omega_i|$, $W = \text{diag}(w)$ and by skipping the constant term $|\Omega| \text{Var}_{\Omega}(u_0)$, problem (7.1) finally simplifies to:

$$\min_{A\alpha \geq 0} \frac{1}{2} \langle W(\alpha - \beta), \alpha - \beta \rangle. \quad (7.7)$$

Problem (7.7) is - once again - an isotonic regression problem. It is however much more complicated than problem (7.3), due to the near arbitrary structure of matrix A . In the rest of the paper, we let α^* denote the unique minimizer of (7.7). Uniqueness is due to the fact that (7.7) amounts to projecting a vector onto a *closed convex set*.

Isotonic regression problems received a considerable attention in the optimization literature (see e.g. [15, 96, 27, 190] for a few algorithms). To the best of our knowledge, the best existing complexity estimates to solve them - in the case of an arbitrary directed acyclic graph - were provided recently in [151]. Therein, the authors propose an interior point algorithm [186] exploiting the special graph structure of matrix A [213]. Let α^* denote the unique solution of problem (7.7). Their tailored algorithm provides a feasible estimate $\alpha^{(\epsilon)}$ of α^* satisfying $A\alpha^{(\epsilon)} \geq 0$ with $\|\alpha^{(\epsilon)} - \alpha^*\|_2^2 \leq \epsilon$ in no more than

$$O(m^{1.5} \log^2 p \log(p/\epsilon)) \quad (7.8)$$

operations. In practice, this algorithm works nicely for small m , but in our experience, the method fails to converge when dealing with large graphs. In what follows, we therefore design a simpler first order algorithm.

The main idea is to exploit strong convexity of the squared l^2 -norm to design a first order algorithm on the dual. Nesterov type accelerations [183] make this method particularly relevant for large scale problems [243]. Proposition 7.3.1 summarizes the nice properties of the dual problem.

Proposition 7.3.1. *The dual problem of (7.7) reads:*

$$\sup_{\lambda \leq 0} D(\lambda) = -\frac{1}{2} \|W^{-1/2} A^T \lambda\|_2^2 + \langle \lambda, A\beta \rangle. \quad (7.9)$$

Let $\alpha(\lambda) = \beta - W^{-1} A^T \lambda$, then any primal-dual solution (α^*, λ^*) satisfies $\alpha^* = \alpha(\lambda^*)$. Function D is differentiable with L -Lipschitz continuous gradient and $L = \lambda_{\max}(AW^{-1}A^T)$. Finally, the following inequality holds for any $\lambda \in \mathbb{R}^m$:

$$\|\alpha(\lambda) - \alpha^*\|_2^2 \leq 2(D(\lambda^*) - D(\lambda)). \quad (7.10)$$

Proof. We only sketch the proof. The idea is to use Fenchel-Rockafellar duality for convex optimization:

$$\begin{aligned} & \min_{A\alpha \geq 0} \frac{1}{2} \langle W(\alpha - \beta), \alpha - \beta \rangle \\ &= \min_{\alpha \in \mathbb{R}^m} \sup_{\lambda \leq 0} \frac{1}{2} \langle W(\alpha - \beta), \alpha - \beta \rangle + \langle A\alpha, \lambda \rangle \\ &= \sup_{\lambda \leq 0} \min_{\alpha \in \mathbb{R}^m} \frac{1}{2} \langle W(\alpha - \beta), \alpha - \beta \rangle + \langle A\alpha, \lambda \rangle. \end{aligned}$$

The primal-dual relationship $\alpha(\lambda)$ is obtained by finding the minimizer of the inner-problem in the last equation. The dual problem is found by replacing α by $\alpha(\lambda)$ in the inner-problem.

Function D is obviously differentiable with $\nabla D(\lambda) = -AW^{-1}A^T\lambda + A\beta$. Therefore, $\forall(\lambda_1, \lambda_2)$, we get:

$$\begin{aligned} \|\nabla D(\lambda_1) - \nabla D(\lambda_2)\|_2 &= \|AW^{-1}A^T(\lambda_1 - \lambda_2)\|_2 \\ &\leq \lambda_{\max}(AW^{-1}A^T) \|\lambda_1 - \lambda_2\|_2. \end{aligned}$$

Inequality (7.10) is a direct consequence of a little known result about the Fenchel-Rockafellar dual of problems involving a strongly convex function. We refer the reader to lemma D.1 in [37] for more details. $\square$

Proposition 7.3.2. *The Lipschitz constant L satisfies $L \leq 4c_{\max}$.*

Proof. Notice that $\lambda_{\max}(AW^{-1}A^T) = \sigma_{\max}^2(AW^{-1/2})$, where $\sigma_{\max}$ stands for the largest

singular value. Moreover

$$\begin{aligned}
\|AW^{-1/2}\alpha\|_2^2 &= \sum_{k=1}^m \left(\frac{\alpha_{\mathbf{I}(k)}}{\sqrt{w_{\mathbf{I}(k)}}} - \frac{\alpha_{\mathbf{J}(k)}}{\sqrt{w_{\mathbf{J}(k)}}} \right)^2 \\
&\leq \sum_{k=1}^m 2 \left(\frac{\alpha_{\mathbf{I}(k)}^2}{w_{\mathbf{I}(k)}} + \frac{\alpha_{\mathbf{J}(k)}^2}{w_{\mathbf{J}(k)}} \right) \\
&= 4 \sum_{k=1}^m \frac{\alpha_{\mathbf{I}(k)}^2}{w_{\mathbf{I}(k)}} \\
&= 4 \sum_{i=1}^p n_i \frac{\alpha_i^2}{w_i},
\end{aligned}$$

where n_i denotes the number of edges starting from vertex i (the outdegree). To conclude, notice that each pixel in region ω_j has at most $c_{\max}$ neighbors. Therefore $n_i \leq w_i c_{\max}$ and we finally get:

$$\|AW^{-1/2}\alpha\|_2^2 \leq 4c_{\max} \sum_{i=1}^p \alpha_i^2 = 4c_{\max} \|\alpha\|_2^2. \quad (7.11)$$

□

Problem (7.9) has a simple structure, compatible with the use of accelerated projected gradient ascents methods [185] described in Algorithm 10.

Algorithm 10: Accelerated proximal gradient ascent method.

Input : Initial guess $\mu^{(1)} \in \mathbb{R}^m$, $\tau = 1/L$ and Nit

for $k = 1$ *to* Nit **do**

$$\left[\begin{array}{l} \lambda^{(k)} = \min \left(\mu^{(k)} + \tau \nabla D(\mu^{(k)}), 0 \right) \\ \mu^{(k+1)} = \lambda^{(k)} + \frac{k-1}{k+2} (\lambda^{(k)} - \lambda^{(k-1)}) \\ \alpha^{(k)} = \alpha(\lambda^{(k)}) \end{array} \right.$$

Proposition 7.3.3. *Algorithm 10 provides the following guarantees:*

$$\|\alpha^{(k)} - \alpha^*\|_2^2 \leq \frac{2L \|\lambda^{(0)} - \lambda^*\|_2^2}{k^2}, \quad (7.12)$$

where λ^* is any solution of the dual problem (7.9). The complexity to obtain an estimate $\alpha^{(k)}$ satisfying $\|\alpha^{(k)} - \alpha^*\|_2 \leq \epsilon$ is bounded above by

$$O\left(\frac{m}{\epsilon} \|\lambda^{(0)} - \lambda^*\|_2\right) \text{ operations.} \quad (7.13)$$

Proof. Standard convergence results [185] state that:

$$D(\lambda^{(k)}) - D(\lambda^*) \leq \frac{L\|\lambda^{(0)} - \lambda^*\|_2^2}{k^2}.$$

Combining this result with inequality (7.10) directly yields (7.12).

To obtain bound (7.13), first remark that each iteration of Algorithm 10 requires two matrix-vector products with A and A^T of complexity $O(m)$. The bound is then a direct consequence of bound (7.12) and Proposition 7.3.2. $\square$

At this point, the convergence analysis is not complete since $\|\lambda^{(0)} - \lambda^*\|_2$ can be arbitrarily large. In order to compare the proposed first order method with the interior point method from [151], we need to upper-bound this quantity.

7.3.4 Complexity analysis

In this paragraph, we analyze the theoretical efficiency of Algorithm 10. We consider the special case $W = \text{Id}$ for the ease of exposition. In practice, controlling the *absolute error* $\|\alpha^{(k)} - \alpha^*\|_2$ is probably less relevant than the *relative error* $\frac{\|\alpha^{(k)} - \alpha^*\|_2}{\|\alpha^{(0)} - \alpha^*\|_2}$. This motivates setting $\epsilon = \eta\|\alpha^{(0)} - \alpha^*\|_2$ in equation (7.13), where $\eta \in [0, 1]$ is a parameter describing the relative precision of the solution.

Setting $\lambda^{(0)} = 0$ and noticing that:

$$\begin{aligned} \|\alpha^{(0)} - \alpha^*\|_2 &= \|\beta - \alpha^*\|_2 \\ &= \|A^T \lambda^*\|_2, \end{aligned}$$

the complexity in terms of η becomes:

$$O\left(\frac{m}{\eta} \frac{\|\lambda^*\|_2}{\|A^T \lambda^*\|_2}\right). \quad (7.14)$$

In what follows, we show that - unfortunately - the ratio $\frac{\|\lambda^*\|_2}{\|A^T \lambda^*\|_2}$ can behave like m and the complexity therefore becomes $O\left(\frac{m^2}{\eta}\right)$, which is significantly worst than interior point methods [151], both in terms of dimension and precision. In all practical examples that we treated, the ratio $\frac{\|\lambda^*\|_2}{\|A^T \lambda^*\|_2}$ however remained bounded by values never exceeding 10, explaining the practical success of the proposed method.

Example of a hard problem An example of hard graph (a simple line graph) is provided in Figure 7.3. For this graph, Algorithm 10 can be interpreted as a *diffusion process*, which is known to be extremely slow. In particular, Nesterov shows that diffusions are the worst case problems for first order methods in [185, p.59].

Proposition 7.3.4. *Consider a simple line graph as depicted in Figure 7.3, with p even and $W = \text{Id}$. Set*

$$\beta_i = \begin{cases} 1 & \text{if } i \leq p/2, \\ -1 & \text{otherwise.} \end{cases} \quad (7.15)$$

Then the primal-dual solution (α^*, λ^*) of the isotonic regression problem (7.7) is given by $\alpha^* = 0$ and

$$\lambda_k^* = \begin{cases} -k & \text{if } 1 \leq k \leq p/2, \\ -n + k & \text{if } p/2 + 1 \leq k \leq p. \end{cases} \quad (7.16)$$

This implies that

$$\frac{\|\lambda^*\|_2}{\|A^T \lambda^*\|_2} \sim m. \quad (7.17)$$

Proof. For this simple graph, $m = p - 1$. To check that (7.16) is a solution, it suffices to verify that the three conditions (7.25), (7.26) and (7.27) are satisfied. This is done by direct inspection, using the fact that for this graph:

$$(A^T \lambda)_i = \begin{cases} -\lambda_1 & \text{if } i = 1 \\ -\lambda_i + \lambda_{i-1} & \text{if } 2 \leq i \leq p-1 \\ \lambda_{p-1} & \text{if } i = p. \end{cases} \quad (7.18)$$

Relationship (7.17) is due to the fact that the sum of squares $\sum_{k=1}^m k^2 = m(m+1)(2m+1)/6 \sim m^3$ so that $\|\lambda^*\|_2^2 \sim m^3$ and $\|A^T \lambda^*\|_2^2 = \|\beta\|_2^2 = m$. $\square$




Figure 7.3 – Worst case graph




Figure 7.4 – First 20000 iterations of the primal-dual pair $(\alpha^{(k)}, \lambda^{(k)})$. Top: β is displayed in red while $\alpha^{(k)}$ varies from green to blue with iterations. Bottom: $\lambda^{(k)}$ varies from green to blue with iterations. A new curve is displayed every 1000 iterations. As can be seen, the convergence is very slow.

Example of a nice problem In order to rehabilitate our approach, let us show that the ratio $\frac{\|\lambda^*\|_2}{\|A^T \lambda^*\|_2}$ can be bounded independently of m for “nice” graphs.

Proposition 7.3.5. *For any $\lambda \leq 0$ and the graph depicted in Figure 7.5 satisfies:*

$$\frac{1}{2} \leq \frac{\|\lambda^*\|_2}{\|A^T \lambda^*\|_2} \leq \frac{1}{\sqrt{2}}. \quad (7.19)$$

Proof. For this graph, we get:

$$(A^T \lambda)_i = \begin{pmatrix} -\lambda_1 \\ \lambda_1 + \lambda_2 \\ -\lambda_2 - \lambda_3 \\ \vdots \\ \lambda_{n-2} + \lambda_{n-1} \\ -\lambda_{n-1} \end{pmatrix}. \quad (7.20)$$

Therefore:

$$\begin{aligned} \|A^T \lambda\|_2^2 &= \lambda_1^2 + \lambda_{n-1}^2 + \sum_{k=1}^{n-2} (\lambda_k + \lambda_{k+1})^2 \\ &= 2 \sum_{k=1}^{n-1} \lambda_k^2 + 2 \sum_{k=1}^{n-2} \lambda_k \lambda_{k+1}, \end{aligned}$$

and

$$2\|\lambda\|_2^2 \leq \|A^T \lambda\|_2^2 \leq 4\|\lambda\|_2^2. \quad (7.21)$$

□

In the general case, we conjecture that the method's complexity depends on the length of the longest path in the graph.




Figure 7.5 – A nice graph

7.3.5 Analytical properties of the minimizer

Let us now provide two analytical properties of the minimizer. The following proposition states that the mean of u^* is equal to that of u_0 .

Proposition 7.3.6. *The solution α^* of problem (7.7) satisfies the following property.*

$$\sum_{i=1}^p w_i \alpha_i^* = \sum_{i=1}^p w_i \beta_i^*. \quad (7.22)$$

Proof. Let $\mathbf{1}$ denote the vector with all components equal to 1. Notice that $\text{Ker}(A) = \text{span}(\mathbf{1})$. Therefore, $\text{Im}(A^T) = \text{span}(\mathbf{1})^\perp$ and we get: $\langle \mathbf{1}, A^T \lambda^* \rangle = \langle \mathbf{1}, W(\alpha^* - \beta) \rangle = 0$. □

The second proposition states that u^* satisfies the *maximum principle* in the sense that:

$$\min_{x \in \Omega} u_0(x) \leq u^*(y) \leq \max_{x \in \Omega} u_0(x) \quad (7.23)$$

for all $y \in \Omega$.

Proposition 7.3.7. *Let $\beta_+ = \max_{1 \leq i \leq p} \beta_i$ and $\beta_- = \min_{1 \leq i \leq p} \beta_i$, then:*

$$\beta_- \leq \alpha_i^* \leq \beta_+, \quad \forall i \in \{1, \dots, p\}. \quad (7.24)$$

Proof. The Karush-Kuhn-Tucker optimality conditions for λ^* read:

$$A^T \lambda^* = W(\beta - \alpha^*), \quad (7.25)$$

$$\lambda^* \leq 0, \quad (7.26)$$

$$\lambda_i^* = 0 \quad \text{if } (A\alpha^*)_i > 0. \quad (7.27)$$

Now, let $E_{out}(i) = \mathcal{I}^{-1}(i)$ and $E_{in}(j) = \mathcal{J}^{-1}(i)$ denote the sets of edges that get in and out of vertex i respectively. With this notation, we get:

$$(A^T \lambda)_i = \sum_{k \in E_{out}(i)} \lambda_k - \sum_{k \in E_{in}(i)} \lambda_k. \quad (7.28)$$

Let i denote the index of a region corresponding to a local maximum of u_1 . This implies that $E_{out}(i) = \emptyset$. Therefore

$$(A^T \lambda^*)_i = - \sum_{k \in E_{in}(i)} \lambda_k^* \geq 0, \quad (7.29)$$

by equation (7.26). By equation (7.25), this implies that $\alpha_i^* \leq \beta_i$.

Therefore, for all positions i corresponding to local maxima of α , we get $\alpha_i^* \leq \beta_i \leq \beta_+$. The fact that all local maxima of α^* are below β_+ implies, in particular, that the global maximum of α^* is below β_+ , so $\alpha_i^* \leq \beta_+$ for all $i \in \{1, \dots, p\}$.

A similar reasoning shows that $\alpha_i^* \geq \beta_-$ for all i . $\square$

7.3.6 Analytic properties of SNR_{loc}

To finish this theoretical study, let us mention a few properties of the value $\Delta_{loc}(u_1, u_0)$:

- It is invariant to linear and affine transforms with a coefficient $a \geq 0$, i.e.

$$SNR_{loc}(au_1 + b, u_0) = SNR_{loc}(u_1, u_0), \quad \forall a \geq 0. \quad (7.30)$$

- It is also invariant to global contrast changes, since global contrast changes are specific instances of local contrast changes. For all non decreasing functions $\phi : \mathbb{R} \rightarrow \mathbb{R}$, we get:

$$SNR_{loc}(\phi(u_1), u_0) = SNR_{loc}(u_1, u_0). \quad (7.31)$$

- In general, it is not symmetric: $\Delta_{loc}(u_1, u_0) \neq \Delta_{loc}(u_0, u_1)$. However, it is possible to make it symmetric by computing

$$\max(\Delta_{loc}(u_1, u_0), \Delta_{loc}(u_0, u_1)) \quad \text{or} \quad \min(\Delta_{loc}(u_0, u_1), \Delta_{loc}(u_1, u_0)). \quad (7.32)$$

7.4 Robust contrast changes

The proposed methodology has a few properties that may be considered as drawbacks in applications. First, a constant image differs by a local contrast change from *any* other image:

$$SNR_{loc}(constant, u) = +\infty, \forall u. \quad (7.33)$$

This is somehow natural: a photograph taken in a lightless environment could possibly represent any scene. However, it can be a hindrance in some applications.

Another drawback of the proposed approach is illustrated in Fig. 7.7. In this experiment, we compare two images given in Fig. 7.7a and 7.7b of different scenes taken under similar lighting conditions. The minimizer of problem (7.1) is displayed in Fig. 7.7c. This image differs from Fig. 7.7b only by a local contrast change, but is definitely very different. In addition, a very faint gray level variation in Fig. 7.7b can result in a huge distortion, see red arrows.

All these observations come from the same underlying cause: local contrast changes are a class of transformations that is too wide to correctly model usual lighting variations. In this section, we propose simple restrictions to avoid the mentioned flaws.

7.4.1 The principle

The idea is to accept only local contrast variations that belong to a range defined from the contrasts of the input image u_1 . This leads to the following definition:

Definition 7.4.1. Let $\theta \leq \Theta$ denote two constants and u and u_1 denote two images. They are said to differ by a robust (θ, Θ) contrast change if the following inequalities are satisfied for all $x \in \Omega$ and for all $y \in \mathcal{N}(x)$ such that $u_1(x) \geq u_1(y)$:

$$\theta(u_1(x) - u_1(y)) \leq u(x) - u(y) \leq \Theta(u_1(x) - u_1(y)). \quad (7.34)$$

Notice that θ and Θ are not assumed to be nonnegative, so that this definition also allows for local contrast inversions.

Similarly to the previous sections, let us now consider the following value:

$$\Delta_{rob}(u_1, u_0) = \min_{T \in \mathcal{T}} \|u_0 - T(u_1)\|_2^2, \quad (7.35)$$

where $\mathcal{T}$ is the set of robust (θ, Θ) contrast changes and define:

$$SNR_{rob}(u_1, u_0) = -10 \log_{10}(\Delta_{rob}(u_1, u_0) / \|u_0\|_2^2). \quad (7.36)$$

7.4.2 A minimization algorithm

Let us express this problem into a form more suitable for numerical computations. Let

$$\gamma_i = \frac{1}{|\omega_i|} \sum_{x \in \omega_i} u_1(x) \quad (7.37)$$

denote the mean of u_1 over region ω_i . Define $c = \theta A\gamma$ and $C = \Theta A\gamma$. By construction, $A\gamma > 0$ so that $0 \leq c \leq C$, where the inequality is meant component-wise. Instead of looking for the minimizer of (7.7), we propose to evaluate:

$$\min_{c \leq A\alpha \leq C} \frac{1}{2} \langle W(\alpha - \beta), \alpha - \beta \rangle. \quad (7.38)$$

In general, the solution of (7.38) does not satisfy the analytical properties given in Section 7.3.5, but the dual algorithm described in the previous section can be generalized quite easily. The main algorithmic facts are summarized in Proposition 7.4.1 and Algorithm 11.

Proposition 7.4.1. *The dual problem of (7.38) reads:*

$$\sup_{\lambda \in \mathbb{R}^m} D(\lambda) + \sum_{k=1}^m (c_k \cdot \min(\lambda_k, 0) + C_k \cdot \max(\lambda_k, 0)). \quad (7.39)$$

Let $\alpha(\lambda) = \beta - W^{-1}A^T\lambda$, then the sequence $\alpha^{(k)}$ generated by Algorithm 11 satisfies:

$$\|\alpha^{(k)} - \alpha^*\|_2^2 \leq \frac{2L\|\lambda^{(0)} - \lambda^*\|_2^2}{k^2}. \quad (7.40)$$

Algorithm 11: Accelerated proximal gradient ascent method.

Input : Initial guess $\mu^{(1)} \in \mathbb{R}^m$, $\tau = 1/L$ and *Nit*

for $k = 1$ *to* *Nit* **do**

$$\left[\begin{array}{l} \tilde{\lambda}^{(k)} = \mu^{(k)} + \tau \nabla D(\mu^{(k)}) \\ \lambda^{(k)} = \tilde{\lambda}^{(k)} - \tau \min \left(\max \left(\tilde{\lambda}^{(k)} / \tau, c \right), C \right) \\ \mu^{(k+1)} = \lambda^{(k)} + \frac{k-1}{k+2} (\lambda^{(k)} - \lambda^{(k-1)}) \\ \alpha^{(k)} = \alpha(\lambda^{(k)}) \end{array} \right.$$

7.5 Numerical results

7.5.1 Image comparison and change detection

In order to assess the relevance of the proposed approach for image comparison and change detection, we took pictures of two scenes - denoted F and G - under different lighting conditions (window shutter closed or open), see Figures 7.6 and 7.7.

As can be seen in Figures 7.6 and 7.7, the proposed local contrast change algorithm is able to correctly distinguish changes of illuminations (Fig. 7.6e is nearly gray) from changes of scene configuration (Fig. 7.7e contains only the differences from u to u_1).




Figure 7.6 – Comparing local and global contrast changes




Figure 7.7 – Comparing local and global contrast changes

	F_1	F_2	F_3	F_4	F_5	G_1	G_2	G_3	G_4	G_5
F_1	Inf	18.11	14.60	16.05	17.24	11.65	11.21	11.14	10.26	12.47
F_2	18.88	Inf	23.91	18.47	15.52	11.51	11.32	11.26	11.00	12.79
F_3	17.05	25.54	Inf	18.45	14.74	11.56	11.40	11.33	11.12	12.87
F_4	14.94	17.38	16.83	Inf	18.04	11.98	11.81	11.75	11.60	13.17
F_5	17.14	15.01	13.03	18.82	Inf	12.24	11.75	11.69	10.64	12.79
G_1	10.54	9.49	8.57	11.44	11.21	Inf	29.56	29.20	18.08	20.20
G_2	10.58	9.68	8.83	11.59	11.24	29.34	Inf	36.34	19.61	21.07
G_3	10.58	9.68	8.82	11.60	11.24	28.81	36.28	Inf	19.76	21.12
G_4	10.41	10.34	9.71	12.46	10.97	19.51	21.59	21.63	Inf	21.49
G_5	10.06	9.48	8.69	11.42	10.58	18.28	19.30	19.30	19.03	Inf

Table 7.1 – Results global contrasts

	F_1	F_2	F_3	F_4	F_5	G_1	G_2	G_3	G_4	G_5
F_1	Inf	32.12	26.13	25.61	24.64	18.36	18.16	18.10	18.07	19.14
F_2	33.03	Inf	39.26	26.26	23.84	18.36	18.28	18.22	18.34	19.38
F_3	31.28	43.49	Inf	25.86	23.22	17.98	17.80	17.75	17.68	18.93
F_4	23.89	26.00	26.31	Inf	31.17	18.97	18.68	18.65	18.23	19.57
F_5	25.06	25.37	23.42	31.22	Inf	18.89	18.63	18.61	18.52	19.45
G_1	16.91	16.05	14.95	17.98	17.50	Inf	42.15	39.66	31.35	30.53
G_2	16.99	16.36	15.32	18.26	17.65	41.19	Inf	45.04	31.77	31.93
G_3	16.69	16.19	15.20	18.09	17.27	39.07	47.71	Inf	32.56	32.41
G_4	16.10	15.86	15.13	18.01	16.68	31.85	34.76	34.83	Inf	35.04
G_5	15.85	15.45	14.48	17.49	16.33	28.32	31.57	31.67	30.10	Inf

Table 7.2 – Results local contrasts

7.5.2 A large panel of illumination changes

We also evaluated the global and local SNR between all pairs of images in Figure 7.8. The results are displayed in Tables 7.1 and 7.2 respectively.

As can be seen in these tables, the SNR between images corresponding to identical scenes is higher than that of images corresponding to different scenes, both for the local and global contrast changes. However this experiment shows that the local SNR is *much more discriminative*. The mean difference between scenes corresponding to different scenes is about 5dB for the global SNR, while it is about 10dBs for the local SNR.

This experiment suggests that local contrast changes provide a much more accurate model of illumination changes than global contrast changes.

7.5.3 Robust contrast change

Finally, we show how the robust contrast change described in Section 7.4 behaves in Fig. 7.9. We set $\theta = 0.5$ and $\Theta = 2$. As can be seen on Fig. 7.9d, the robust contrast



(a) F_1



(b) G_1



(c) F_2



(d) G_2



(e) F_3



(f) G_3



(g) F_4



(h) G_4



(i) F_5



(j) G_5

Figure 7.8 – Different images used for comparison

change provides an image that could really be due to lighting variations: it much better resembles the original picture 7.9b than the local contrast change in Fig. 7.9c.




Figure 7.9 – Comparing local and robust contrast changes

8

High-resolution in-depth imaging of optically cleared thick samples using an adaptive SPIM

Durant ma thèse, j'ai eu le privilège de prendre part à une collaboration avec les équipes d'imagerie et de biologie de l'ITAV représentées par Aurore Masson, Corinne Lorenzo ; Céline Frongia et Bernard Ducommun ainsi que la société d'Imagine Optic représentée par Grégory Clouvel. Cette collaboration a donné lieu à une publication dans la revue Nature Scientific Report en Novembre 2015 [165].

L'article issu de cette collaboration est *étroitement* lié aux problématiques évoquées dans cette thèse bien que la plupart de son contenu présente des détails techniques éloignés d'une thèse de mathématiques appliquées. Ainsi, mes contributions relatives à cette collaboration sont résumées ici et l'article complet est présenté à titre d'annexe.

La microscopie à feuille de lumière est une nouvelle technique permettant d'imager des échantillons sur plusieurs centaines de microns de profondeur. Les images produites par ce microscope sont dégradées par différents phénomènes :

- la diffusion et l'absorption qui sont intrinsèques à l'échantillon,
- des aberrations optiques produites à la fois par le microscope et l'échantillon. La contribution de l'échantillon est due à sa complexité structurelle qui conduit à des variations locales des indices de réfraction dans l'échantillon.

Récemment, des traitements chimiques ont été développés et permettent d'homogénéiser les indices de réfraction des échantillons pour réduire les dégradations optiques induites. Toutefois, le traitement crée une discontinuité entre les indices de l'échantillon traité et le milieu utilisé pour les objectifs. Ce saut d'indice induit une dégradation simplifiée qui peut être simulée partiellement.

La publication contient de nombreux détails sur les méthodes de traitements chimiques des échantillons et sur le système d'optique adaptative. Mes contributions consistent en deux points :

Définir une mesure de qualité d'image L'optique adaptative est utilisée pour corriger les dégradations liées aux changements des indices de réfraction du milieu. Pour quantifier l'apport lié à ce système, nous avons mis au point une mesure de

qualité des images. Cette méthode repose sur l'étude de la décroissance des fréquences des images. Les images floues ont une décroissance plus rapide que les images nettes. La comparaison relative d'un profil radial des coefficients de Fourier a montré que les images corrigées par optique adaptative étaient plus nettes. Cette indicateur pourrait probablement être amélioré en exploitant la structure des phases des images [152].

Proposer un algorithme de déconvolution aveugle J'ai aussi proposé un algorithme de déconvolution aveugle qui permet de restaurer les images tout en estimant le noyau de dégradation. La déconvolution a permis d'obtenir des images bien plus nettes et l'étude des noyaux estimés a confirmé que l'optique adaptative améliorait nettement la qualité des images. Le couplage d'un système d'optique adaptative et de cet algorithme offre des images qui deviennent exploitables, voir Figure 8.1.




Figure 8.1 – Illustration de l'amélioration de la qualité des images obtenues grâce au couplage optique adaptative/déconvolution aveugle. A gauche : un plan d'un modèle tumeur imagé sans optique adaptative et sans post-traitement. A droite : l'image a été acquise après correction par le système d'optique adaptative et restaurée à l'aide de l'algorithme de déconvolution aveugle.

SCIENTIFIC REPORTS

OPEN

High-resolution in-depth imaging of optically cleared thick samples using an adaptive SPIM

Aurore Masson^{1,2}, Paul Escande^{1,2,3}, Céline Frongia^{1,2}, Grégory Clouvel⁵, Bernard Ducommun^{1,2,4} & Corinne Lorenzo^{1,2}

Received: 07 July 2015

Accepted: 21 October 2015

Published: 18 November 2015

Today, Light Sheet Fluorescence Microscopy (LSFM) makes it possible to image fluorescent samples through depths of several hundreds of microns. However, LSFM also suffers from scattering, absorption and optical aberrations. Spatial variations in the refractive index inside the samples cause major changes to the light path resulting in loss of signal and contrast in the deepest regions, thus impairing in-depth imaging capability. These effects are particularly marked when inhomogeneous, complex biological samples are under study. Recently, chemical treatments have been developed to render a sample transparent by homogenizing its refractive index (RI), consequently enabling a reduction of scattering phenomena and a simplification of optical aberration patterns. One drawback of these methods is that the resulting RI of cleared samples does not match the working RI medium generally used for LSFM lenses. This RI mismatch leads to the presence of low-order aberrations and therefore to a significant degradation of image quality. In this paper, we introduce an original optical-chemical combined method based on an adaptive SPIM and a water-based clearing protocol enabling compensation for aberrations arising from RI mismatches induced by optical clearing methods and acquisition of high-resolution in-depth images of optically cleared complex thick samples such as Multi-Cellular Tumour Spheroids.

Light sheet fluorescence microscopy (LSFM), also known as Selective plane illumination microscopy (SPIM), represents a universal and versatile technique for three-dimensional (3D) imaging of live tissues and organisms with subcellular resolution^{1–4} and, undoubtedly, is emerging as a useful tool for performing 3D imaging of complex thick biological samples^{3,5}. Nevertheless, as for all fluorescence microscopes, it still remains limited for in-depth imaging of scattering and of heterogeneous samples. Indeed, optical aberrations, absorption and scattering of both excitation and emission result in a loss of signal and contrast, limiting practical use for imaging up to a few hundred μm deep. In complex thick samples, scattering and optical aberrations arising from refractive index (RI) discontinuities between and within cells are the main processes which contribute to degradation of image quality⁶ and which limit the resolving power of optical imaging techniques. To overcome these obstacles, LSFM can be combined with an optical clearing method which chemically treats tissues to render them transparent^{7–11}. Recent progress in tissue clearing methods has facilitated microscopic analysis of whole embryos, tissues and intact organisms. These methods work by minimizing RI mismatches in tissues so that photons undergo less, or almost no, scattering. Furthermore, by homogenizing the RI in fixed samples, optical aberrations induced by the sample itself are reduced or eliminated.

However, achieving high transparency in the sample is not enough to acquire high-resolution 3D images. Indeed, a common problem in imaging optically cleared samples is the immersion media of objectives. The latter are designed to work with a specific RI medium (n_i) that often differs from the RI of

¹Université de Toulouse, ITAV-USR3505, F-31106 Toulouse, France. ²CNRS, ITAV-USR3505, F-31106 Toulouse, France. ³ISAE, DISC, F-31106 Toulouse, France. ⁴CHU de Toulouse, F-31059 Toulouse, France. ⁵Imagine Optic Orsay F-91400, France. Correspondence and requests for materials should be addressed to C.L. (email: corinne.lorenzo@itav.fr)

the clearing agent (n_2), therefore leading to a mismatch between RI n_1 and n_2 . Consequently, strong defocus and spherical aberrations are imparted into the system, causing significant light perturbation such as defocus and focal elongation, together with a significant decrease in signal-to-noise ratio and resolution¹². Furthermore, aberration increases in proportion to imaging depth, thus hindering high-resolution imaging of large volumetric samples. Recently, however, the availability of customised objectives specifically designed to work with given clearing agents has made it possible to solve this problem¹³. But since their development is still in its infancy, few such affordable objectives are on the market. An alternative solution has been suggested by Silvestri & co-worker¹⁴, consisting of correcting aberrations by means of adaptive optics (AO). A method originally developed for astronomic telescopes, AO measures and corrects aberration, thus restoring optimum performance of the imaging system¹⁵. Over the past decade, AO has been implemented in microscopy¹⁶ and has proven to be a valuable tool for correcting aberration and restoring diffraction-limited resolution in different kinds of fluorescence microscopes^{17–20}. Recently, we successfully applied a Wavefront sensor Adaptive Optics scheme to use with SPIM ($_{\text{WAO}}$ SPIM) to correct complex aberration induced by the sample itself²¹. We demonstrated the ability of $_{\text{WAO}}$ SPIM to improve in-depth imaging of deep complex samples such as Multi Cellular Tumour Spheroids (MCTS). These are tissue-mimic models that are highly useful in investigating the influence of malignant cell interactions during cell proliferation²². Due to their opacity and density, however, they raise significant challenges for imaging by light microscopy. We have shown that spatial variation of aberrations within such thick, non-transparent biological samples causes major changes to the light path and limits the quality of AO correction in an extended field of view²¹.

To date, AO has not yet been combined with optical clearing imaging methods. In the current study, we propose a straightforward method for aberration correction relying on the use of $_{\text{WAO}}$ SPIM in an open-loop mode and on a single set of corrections, enabling us to improve imaging across the entire volume of cleared samples studied. Furthermore, we demonstrate that this method combined with deconvolution helps recover more details and suffer less artifacts. This efficient, user-friendly and versatile strategy which makes it possible to compensate for aberrations arising from RI mismatches induced by optical clearing methods was applied to 3D imaging of cleared complex thick tissue mimics represented here by MCTS.

Results

The source of aberration: theoretical analysis. To test the proposed method, MCTS either labelled with the EdU Click chemistry method (MCTS_EdU) and enabling DNA detection in proliferating cells or MCTS stably expressing a histone H2B nuclear protein (H2B) fused to the mCherry fluorescent protein (MCTS_mCherryH2B) were cleared by using the most recently developed method, CUBIC²³. CUBIC is a water-based optical clearing method that works by (1) breaking down lipids in fixed tissues and (2) replacing the surrounding intracellular media by a higher RI solution. In our system, detection was performed using water-dipping objective lenses. The sample was embedded in a cylinder with a mixture of 2% agarose and clearing agent that was immersed in PBS. CUBIC is a reversible method, which means that after several hours of incubation in PBS/water, cleared MCTS nearly fully recovered their appearance. Thus, embedding the cleared MCTS in a mixture of agarose and clearing agent enabled us to efficiently clear the sample and to maintain the sample clarity by limiting reversibility in PBS/water. A schematic illustration of the sample mounting geometry is shown in Fig. 1A. Applying CUBIC led to a refractive index mismatch between the clearing agent (RI: 1.48–1.49, Fig. 1A) and the objective immersion medium (herein RI: 1.33, Fig. 1A), thereby introducing defocus and spherical aberration. This delayed a fraction of ballistic photons and contributed to point spread function (PSF) enlargement and therefore to resolution degradation.

Defocus and spherical aberrations are well known^{24,25} and can be understood by considering the shape of the wavefront in the pupil plane (P) of an objective lens with a given RI n_1 . Since optical clearing methods result in RI homogenisation within the sample, we assumed the cleared MCTS to be a homogeneous medium of uniform RI n_2 . In a first approximation, we also assumed that the interface between the mismatched mediums was flat. In this case, the pupil function was modelled by

$$P(\rho, d) = e^{-j\Psi(\rho, d)} \quad (1)$$

where $\rho \in [0, 1]$ is the normalized pupil radius and d is the nominal focusing depth in a perfectly matched system (absence of a RI boundary). In an aberrated system, the pupil function is modified by the wave aberration function Ψ (or phase error). This function can be decomposed as a weighted sum of Zernike polynomials (Z) and can be expanded into a series of radially symmetric Zernike polynomials of zero azimuthal order ($Z_{j,0}$, where j and i are respectively the radial and azimuthal orders) with aberration coefficients $A_{j,0}$

$$\Psi(\rho, d) = \sum_{\substack{j=2 \\ j \text{ even}}}^{\infty} \Psi_j(\rho, d) = d \frac{2\pi}{\lambda} \text{NA} \sum_{\substack{j=2 \\ j \text{ even}}}^{\infty} A_{j,0} Z_{j,0}(\rho) \quad (2)$$




Figure 1. (A) Schematic drawing (top view and side view) of the SPIM physiological chamber and the sample-mounting configuration with a zoom view of the geometry of agarose/clearing agent cylinder in which the MCTS was included. (B) Magnitude of aberration coefficients Eqs (3 and 4) as a function of NA. (C,D) Aberration function $\Psi(\rho, d)$ as function of d for $\rho = 1$ (B) or as function of ρ for $d = 200 \mu\text{m}$ (D) Eq (2). Different aberration terms Ψ calculated for an axial distance $d = 200 \mu\text{m}$ considering only the 1st spherical order (d , solid line) and composed by the sum of the first three spherical terms (d , dashed line). (E) Evolution of Ψ for different d values. The edge of MCTS was estimated to be at $d = 50 \mu\text{m}$. The red line corresponds to the compensated optical section ($d = 200 \mu\text{m}$), where average amplitude of Ψ is predicted.

Here j represents the orders of defocus ($j = 2$) and all higher order spherical aberrations ($j = 4, 6, 8, \dots$), NA is the numerical aperture, and λ the fluorescence emission wavelength. Considering only defocus and spherical aberrations, the aberration coefficients can be calculated using eqs (3) and (4)

$$A_{j,0} = B_j(n_1) - B_j(n_2) \quad (3)$$

where

$$B_j(n_k) = \left[1 - \frac{j-1}{j+3} \tan^4 \left\{ \frac{\text{asin}\{NA/n_k\}}{2} \right\} \right] \times \frac{\tan^{j-1} \left\{ \frac{\text{asin}\{NA/n_k\}}{2} \right\}}{2(j-1)\sqrt{j+1}} \quad (4)$$

The numerical computation of these equations shows the strong influence of the NA objective, the nominal focusing depth, and the RI mismatch on the extent of defocus and spherical aberration (Fig. 1B–D). The aberration coefficients $A_{j,0}$ were plotted as a function of NA, for $j=2, 4, 6$ and 8 (Fig. 1B). As shown, aberration coefficients $A_{j,0}$ are dominated by defocus ($j=2$) and 1st order spherical aberration ($j=4$) and drop for high orders of j . High NA objectives are more sensitive to spherical aberrations than low NA objectives. Considering the detection objective (X20, NA 0.5) and the emission wavelength used in this study (510 nm), we calculated the corresponding aberration function ($\Psi_j(\rho, d)$) (Fig. 1C). The latter rises proportionally to the nominal focusing depth corresponding to the distance between the edge of the sample and the depth of the focal plane. In this current optical configuration, a nominal focusing depth of 400 μm , for example, produced a focusing error of 50 μm (Supplementary Fig. S1a,b). Second and 3rd spherical orders for an objective with a NA of 0.5 can be disregarded (respectively 10^{-6} and 10^{-7} , Fig. 1B). However, the addition of oscillating terms, even low-amplitude ones such as spherical aberrations, might induce a major wavefront modification. For this reason, it was important to ensure that the sum of the 1st, 2nd and 3rd spherical orders gave rise to the same result as when considering only the 1st order spherical aberration (Fig. 1D). We calculated the root mean square error (RMSE) to compare the difference between the aberration function Ψ corresponding to the sum of the first 3 spherical orders ($j=4,6,8$) and the aberration function relative to the 1st spherical order ($j=4$).

$$\text{RMSE} = \sqrt{\frac{\sum_{\rho} \left| \left(\sum_{\substack{j=4 \\ j \text{ even}}}^8 \Psi_j(\rho, d) \right) - \Psi_4(\rho, d) \right|^2}{\sum_{\rho} \rho^2}} \quad (5)$$

This deviation was estimated at 0.023%, meaning that using a model considering only the 1st spherical order was well adapted in our case.

From the numerical simulation, we observed that the amplitude of the aberration function increased proportionally with depth and yielded an intermediate weight contribution at around 200 μm deep within the cleared MCTS (Fig. 1E). This imaging plane, hereafter known as the compensated optical section, corresponded to the MCTS centre minus the distance (usually 50–100 μm) between the immersion medium (n_1)/mounting cylinder (n_2) interface and the cleared MCTS boundary (Fig. 1A). The distance variation was essentially due to the sample mounting procedure. The samples were most often found to be located in one of the cylinder quarters, rarely centered (Fig. 1A). This current depth-dependent spherical aberration model assumes a flat interface between the mismatched mediums. It is well known that a cylindrical-shaped sample holder (Fig. 1A) introduces new aberrations, which must also be taken into consideration. As described by Bourgenot and colleagues²⁶, the primary source of aberrations resulting from the cylindrical interface with a sample holder is astigmatism. In addition, it has been shown by Turaga and Holy²⁷ that the magnitude of aberrations such as astigmatism, coma and trefoil, arising by imaging at a tilt through a planar surface between mediums of mismatched refractive index, increased with the tilt angle. In this particular case, the detection axis was not perpendicular to the sample interface and the rays travelling in the propagation direction were subject to different sample thicknesses. A similar effect occurs with cylindrical interfaces and lead to the same nature of aberration pattern. However, in our case the sample position could not be precisely controlled inside the agarose/clearing agent cylinder, thus leading to—high variability of these aberration magnitudes.

Experimental Results

In order to minimize the aberrations due to sample positioning, it was important to systematically position the samples in the same way in the $_{\text{WAO}}$ SPIM. Therefore, sample holders were carefully rotated until the quarter cylinder where the MCTS was closest to the cylinder edge was located in front of the objective lens (Fig. 1A). After that, the light sheet was refocused manually on the cleared MCTS compensated optical section. This step, consisting of translating the light sheet in the z axis (approximately 1–2 μm) and in the x axis (approximately $6.5 \pm 2.5 \mu\text{m}$, mean $\pm$ SD) until obtaining a clear image, enabled us to compensate for focus error (Supplementary Fig. S1c). Then, in order to achieve high-resolution images of cleared MCTS, aberration correction was performed by using an open-loop strategy, which consisted in first manually adjusting spherical aberration and defocus (residual errors) and then the other modes of higher amplitude accountable for major phase errors such as astigmatism, coma and trefoil. Prior to

Haso index	Equation	Name
3	$\sqrt{3} (2\rho^2 - 1)$	Focus
4	$\sqrt{6} \rho^2 \cos 2\theta$	Astigmatism 0°
5	$\sqrt{6} \rho^2 \sin 2\theta$	Astigmatism 45°
6	$\sqrt{8} (3\rho^2 - 2)\rho \cos \theta$	Coma 0°
7	$\sqrt{8} (3\rho^2 - 2)\rho \sin \theta$	Coma 45°
8	$\sqrt{5} (6\rho^4 - 6\rho^2 + 1)$	3rd order Spherical
9	$\sqrt{8} \rho^3 \cos 3\theta$	Trefoil 0°
10	$\sqrt{8} \rho^3 \sin 3\theta$	Trefoil 30°

Table 1. List of the Zernike terms with their corresponding equation used in this study.

this step the deformable mirror command matrix was computed during a calibration process based on the characterisation of each actuator mirror using a reference sample (see Methods). Table 1 gives the list of Zernike terms with their corresponding equation used for the subsequent experiments (indices given by the supplier). The first 10 Zernike azimuthal orders were used, excluding the lowest two values corresponding to tilts. The latter did not affect the image quality. Overall, the user successively adjusted the coefficient value for each Zernike mode with reference to the compensated optical section image output until a clear image with maximum possible brightness and sharpness levels was obtained. This procedure can easily be performed in the same way that objective correction collars can be adjusted. Once the set of Zernike coefficients was adjusted and applied a z-stack of the entire cleared MCTS was then acquired either with or without correction (respectively, AO-on or AO-off) and compared. Supplementary Figure S2 shows the Zernike coefficients of two individual sample corrections and the average magnitude of Zernike terms corrected for different experiments. As expected, the RI mismatch introduced spherical aberrations. Results show that predicted values of spherical aberrations ($0.18\mu\text{m}$) were close to the experimental data ($0.10\mu\text{m} \pm 0.05$, mean $\pm$ SD) (Supplementary Fig. S2). We also observed that defocus was almost fully compensated for after light sheet refocusing (residual focus error $0.02\mu\text{m} \pm 0.01$, mean $\pm$ SD). In addition, we also encountered astigmatism, coma and trefoil errors. Although the samples were carefully positioned in the same way for each experiment, the magnitude of these aberrations was significantly different from one sample to other (Supplementary Fig. S2) and was thus difficult to predict. In terms of image quality, the corrected images displayed improved sharpness and attenuated blur (Fig. 2; Supplementary Movie S1). The gain in information was clearly illustrated by the apparition of thin structures unresolved without correction. Moreover, improvement was visible throughout all the optical section images of cleared MCTS (Fig. 2A,B). The apparition of structural details could be seen in the first planes ($30\mu\text{m}$ in depth) and beyond $100\mu\text{m}$ deep to the compensated optical section ($190\mu\text{m}$ in depth) (Fig. 2B). In order to qualitatively evaluate improvement in image quality after correction, fluorescence intensity profiles along different lines were extracted from Fig. 2B and analysed (Fig. 2C). After correction (AO-on), we could clearly distinguish distinctly separated and higher intensity peaks in the profiles of both lines (1) and (2) corresponding, respectively, to condensed DNA of a prometaphase figure and to poorly resolved nuclei distinguishable without correction (AO-off).

Image quality assessment. In addition to fluorescence intensity profiles, improvement of image quality was then also quantitatively evaluated in terms of local contrast and spatial frequencies. The former were measured according to Jorand *et al.*²¹, using a gradient intensity map representing the local total variance computed for square Regions of Interest (ROIs). Intensity gradient map ratios (IG_{ratio}) of different optical sections of cleared MCTS_EdU are shown in Fig. 3. The AO-on images represent those acquired at different focal depths after performing aberration correction on the compensated optical section. Local contrast enhancement can be clearly seen for all optical section images of cleared MCTS_EdU. We measured an average improvement of 10–20%, as high as 40% in some subregions (Fig. 3). Up to $200\mu\text{m}$ in depth, correction was uniformly efficient across the field of view. Beyond $250\mu\text{m}$ in depth, contrast improvement was less noticeable than in the first optical section images of cleared MCTS. As expected, correction was more efficient and uniformly close to the compensated optical section (located in this case at $150\mu\text{m}$ deep), but nonetheless, a significant improvement (approximately 10%) was maintained even at a depth of $300\mu\text{m}$.

To get an accurate assessment of the improvement provided by correction, we decided to introduce a novel metric based on spatial frequencies allowing us to quantify and compare each spatial frequency component of images which were corrected or not. The aim was to evaluate the gain in spatial resolution for each object size in the image. Since spatial frequency components carry information about the image in great detail the more contrasted the latter become after aberration correction, the more the frequency weight should increase. A ratio factor was calculated corresponding to the relative difference of each frequency (RD) by comparing each component to the reference measured in the AO-off image.




Figure 2. Cleared MCTS labelled with EDU-click Alexa 594 images were acquired with $_{\text{WAO}}\text{SPIM}$ either without (AO-off) or with (AO-on) correction at a fixed excitation intensity at 594 nm and a 200 ms exposure time. Light sheet illumination comes from the right. (a) Top: cleared MCTS-Edu images acquired at a focal depth of 100 μm before (AO-off) and after (AO-on) aberration correction. Bottom: Magnified views of top inset views (*) and (**). Scale bar 10 μm. (b) Subregions of (AO-off) and (AO-on) images taken at different focal depths (30, 50, 100, 150, 190 μm). Scale Bar 10 μm. (c) Intensity profiles along lines 1 (50 μm in depth) and 2 (100 μm in depth) indicated in (b). Scale bar = 10 μm.

For an $M \times N$ image u its 2-D discrete Fourier Transform was defined as follows:

$$\hat{u}(f_x, f_y) = \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} u(x, y) e^{-2j\pi \left(\frac{xf_x}{M} + \frac{yf_y}{N} \right)} \quad (6)$$

with $(f_x, f_y) \in [0, M-1] \times [0, N-1]$ being the coordinates in the frequency domain. This representation enabled us to localize the spatial frequency corresponding to each object with a specific size defined by d_r . The latter was radially spread out around the zero frequency and was distributed in a ring Ω_r which comprised $|f_x| = M/2d_r$ in the x direction and $|f_y| = N/2d_r$ in the y direction. Spatial frequency components contained in each ring Ω_r were summed to create a radial component associated to the radial frequency (f_r). We constructed $\hat{U}(f_r)$, the normalized radial representation of $\hat{u}$

$$\hat{U}(f_r) = \frac{\sum_{\Omega_r} |\hat{u}(f_x, f_y)|}{|\hat{u}(0, 0)|} \quad (7)$$

To evaluate the improvement of images, we then calculated the relative differences (RD) of radial frequency between images acquired before, U_{off} , and after, U_{on} , correction. RD was defined by

$$RD(f_r) = \frac{\hat{U}_{\text{on}}(f_r) - \hat{U}_{\text{off}}(f_r)}{\hat{U}_{\text{off}}(f_r)} \quad (8)$$

RD values were plotted in function of spatial frequencies (Figs 4 and 5) for different optical section images of cleared MCTS_Edu (Fig. 4A) and cleared MCTS_mCherryH2B (Figs 4B and 5). The maximum value of spatial frequencies referred to in x coordinates ($1.03 \mu\text{m}^{-1}$) corresponds to the frequency cut-off of our optical system (resolving power limit). Consequently, values above this limit had no




Figure 3. IG_{ROI} Ratio of cleared MCTS_EdU image of subregions taken at different focal depths (50, 100, 150, 200, 250, 300 μm), calculated as $IG_{ROI}Ratio = \frac{IG_{ROI}^{AOon} - IG_{ROI}^{AOoff}}{IG_{ROI}^{AOoff}}$ where IG_{ROI}^{AOon} is IG_{ROI} mapping image calculated from images obtained with correction and IG_{ROI}^{AOoff} is IG_{ROI} mapping image calculated from images obtained without correction. The square size of ROI corresponds to a typical diameter of a nucleus (10 μm). Scale bar = 25 μm .

significance and are not reported. Positive values of RD indicated an effective correction for a given frequency (f_r) thus enabling us to retrieve certain details characterized by specific size, while negative RD values indicated an inefficient correction leading instead to degradation of information. Regardless of the imaging depth, we observed a similar distribution of RD with a peak for spatial frequencies




Figure 4. RD curves corresponding to the relative difference between Fourier contents of AO-on and AO-off images obtained for different optical section images of cleared MCTS_EdU **(a)** and cleared MCTS_mCherryH2B **(b)** immersed for 5 days in CUBIC reagent 1. Mean curves are estimated from $n = 4$ experiments and error bars represent the standard deviation.

ranging from 0.25 to $0.35 \mu\text{m}^{-1}$ and a full width at half maximum for frequencies ranging from 0.15 to $0.55 \mu\text{m}^{-1}$ corresponding to object size typically matching that of nuclei and subnuclear structures (ranging from 6.7 to $1.8 \mu\text{m}$). Maximal values of RD were found for the optical section images acquired in the first 100 microns and decreased beyond that up to a factor two. Interestingly, we noted that the frequency weight varied less (standard deviation) around the compensated optical section (100 – $150 \mu\text{m}$




Figure 5. RD curves corresponding to the relative difference between Fourier contents of AO-on and AO-off images obtained for different optical section images of cleared MCTS_mCherryH2B immersed for 2 days in CUBIC reagent 1. Mean curves are estimated from $n = 4$ experiments, and error bars represent standard deviation.

in depth), indicating that image quality enhancement was robust around the compensated optical section and somewhat less beyond. Furthermore, we observed weaker values of RD and greater variability for cleared MCTS_mCherryH2B (Fig. 4B) than for cleared MCTS_EdU (Fig. 4A). Although the enhancement in image quality was still quantifiable, its efficiency was noticeably lower. Fluorescence preservation and clearing processes are known to be important trade-offs since the chemical treatments generally used in clearing methods increase transparency but at the same time degrade fluorescence of proteinaceous fluorophores such as mCherry. CUBIC is well known for preserving fluorescence in this context, but we nonetheless observed a drop in mCherry fluorescence levels after 5 days with the first CUBIC reagent (R1) which led to a low level of fluorescence. This was verified by reducing R1 incubation time to 2 days in order to better preserve the mCherry fluorescence (Fig. 5). This change, however, was accompanied by inefficiently cleared processes. Maximum RD values were higher (0.5 vs. 0.3) in this case than for cleared MCTS_mCherryH2B incubated for 5 days in R1 and were equivalent to those obtained for cleared MCTS_EdU. Nevertheless, the correction was less reliable (greater variability), and the improvement in image quality was observable only within a restricted range of frequencies, from 0.1 to 0.3 μm^{-1} . Finally, despite weak signal intensity levels, which could have been a limiting factor for correction, we were still able to achieve a certain degree of correction and thus significantly improve images of cleared MCTS_mCherryH2B. The signal intensity level must therefore be recognized as an important factor in correction ability and, consequently, in achieving high image quality.

AO improve the deconvolution process. To deconvolve the images, we used blind-image deconvolution making it possible to recover a sharp image by researching the unknown kernel k blur²⁸ (Supplementary information S1). Figure 6 shows the resulting deconvolved images acquired with (AO-on) or without (AO-off) correction, which we called respectively “Deblurred AO-on” and “Deblurred AO-off”. In the case of an initially sharp image, this algorithm makes it possible to expect an estimated PSF tending to a Dirac function. We observed that for the uncorrected images, the corresponding deblurred images were enhanced with artificially added high frequencies. Furthermore, they were associated with a larger kernel. In contrast, deblurred AO-on images were well defined. As shown in Supplementary Figure S3 and Movie S2, improvements could be observed in the whole deblurred AO-on z-stack images compared to deblurred AO-off sample. In conclusion, image correction enabled us to reveal details not visible without correction and contributed to successful deconvolution since the deconvolution process was not efficient on uncorrected images. Consequently, this regularisation approach can be effectively combined with correction in order to obtain a final sharp deblurred image.




Figure 6. Images of cleared MCTS_EdU images (corresponding to the compensated optical section, depth = 100 μm) acquired with (AO-on) or without (AO-off) correction and obtained after blind deconvolution. k_{off} and k_{on} represent the PSF function between the deblurred and the raw AO-off and AO-on images calculated by the deconvolution algorithm. Scale bar = 25 μm .

Versatility of the method. To further demonstrate the versatility of our strategy and its compatibility with any water-based clearing method and any kind of staining, we then used the CLARITY method^{29,30} with our samples. This method, compatible with immunostaining applications, is based on transforming tissue to a nanoporous hydrogel and removing lipids with an active or passive diffusion process based on ionic detergent solubilisation. MCTS were cleared with CLARITY (RI of approximately 1.45) and subsequently immunostained with an antibody directed against the α tubulin protein, a constituent of microtubule structure. We chose this particular structure in order to test correction ability in

sparsely labelled samples. In contrast to CUBIC, CLARITY is an irreversible method making it possible to immerse the cleared samples directly in water/PBS while maintaining their transparency. Therefore the cleared MCTS were embedded in a cylinder of 1% agarose which was then immersed in PBS. In this case, the RI of sample mounting holders was close to water, which enabled us to disregard aberration arising from sample positioning. Indeed astigmatism, coma and trefoil terms were not present and only defocus and spherical aberration had to be compensated for to improve image quality (Supplementary Fig. S4a). Supplementary Figure S4 shows the resulting images of microtubules, corrected (AO-on) or not (AO-off), with the corresponding local contrast map (IG) and RD plots. The microtubules were imaged at a depth of approximately 60 μm . After correction, contrast improvement was significant and enabled us to clearly distinguish particular fine structures such as the cytokinetic bridge (Supplementary Fig. S4b) and the mitotic spindle (Supplementary Fig. S4c).

Discussion

SPIM has emerged as a versatile and indispensable instrument for long-term live 3D imaging of organism models. Furthermore, recent progress in optical clearing methods and their combination with LSFM have, without any doubt, improved optical in-depth imaging capabilities in large-scale tissue samples. Our results show that it is indeed possible to simplify the aberration pattern of tissue mimics, such as MCTS, by optically clearing them. In this context we chose water-based clearing methods such as, for example, CUBIC or CLARITY for their compatibility with SPIM configuration¹⁰. However, even if these methods guarantee high transparency and simplify MCTS aberration patterns, it was apparent that even a small RI mismatch between immersion medium and cleared samples can cause a significant loss of resolution due to spherical aberrations. RI mismatch is thus a key consideration for high-resolution in-depth imaging of water-based cleared samples. However, the correction procedure is greatly simplified in this context since the optical properties of the sample can be approximated, thus enabling predictive correction. This fact motivated us to exploit the potential of $_{\text{WAO}}$ SPIM for correcting this kind of aberration.

To address this issue, we first provided a comprehensive and simple model based on RI mismatch in $_{\text{WAO}}$ SPIM. The numerical simulation study of phase errors showed that low-order aberration correction driven by the radially symmetric Zernike terms should be sufficient to provide significant image quality improvement. We then defined a single focal depth, the compensated optical section, where the contribution weight of phase errors was of an intermediate level, and where we subsequently carried out the open loop. However, depending on the geometry and nature of the sample mounting holder, sample-positioning aberrations such as astigmatism, coma and trefoil can be found and must be taken into consideration. Our results closely reflected the simulation, and it was experimentally demonstrated that correcting from 4 to 10 Zernike terms provides significant recovery of signal and resolution. In addition, we showed that even for small initial aberrations, low-order correction is still beneficial. Image quality improvement was evaluated using two image quality metrics based, respectively, on gradient intensity and spatial frequency content. In general, the corrected images show considerably improved quality in terms of peak intensities and contrast. However, we were not able to reveal image quality improvement for cleared MCTS_mCherryH2B by using the gradient intensity map, whereas the spatial frequencies content metric did reflect significant enhancement. In other words, the intensity gradient map does not provide consistently reliable information on images with low signal-to-noise ratio such as cleared MCTS_mCherryH2B images. In contrast, RD is suitable when signal-to-noise ratio is low. Overall, our results demonstrate that this “semi-empirical” strategy is robust with respect to changes in experimental conditions. In particular, it makes it possible to reach a satisfactory compromise in image quality improvement over all focal depths and can also be easily used in different kinds of staining and clearing protocols. Finally, we have shown that this method, when combined with fast space-invariant deconvolution algorithms, improves the quality of raw corrected images and ensures that the PSF is closer to the theoretical model.

However, when imaging at different depths, our strategy does not currently provide plane-by-plane correction. An approach using an image-based sensorless adaptive scheme³¹ would merit exploration in the future on, for example, (i) methods relying on the theoretical model described above (Supplementary information S2), or (ii) iterative algorithms based on trial-and-error to optimally converge within an image quality metric.

In conclusion, we propose a simple and user-friendly procedure enabling image quality improvement in light sheet microscopy for cleared samples capable of dealing with all existing and forthcoming water-based clearing methods. With the rapidly growing list of clearing protocols, there has never been a greater need for solutions making it possible to compensate for a range of RI for different kinds of cleared samples. Ongoing advances in the combination of these methods with sophisticated techniques such as $_{\text{WAO}}$ SPIM significantly contribute to a better understanding of the practical benefits and the trade-offs of optical-chemical combined methods, and will undoubtedly open up new experimental avenues for exploring thick complex samples such as tissue mimics.

Methods

$_{\text{WAO}}$ SPIM setup. The layout of the set-up is shown in Supplementary Figure S6. The illumination path (orange and blue lines) is similar to that of a conventional SPIM. The main difference concerns

the detection path (red and green lines) where an adaptive optic loop has been implemented. For more details see reference²¹.

Calibration procedure. The AO loop correction was driven by CASAOTM software (Imagine Optic) which enables communication with the Hartman Shack Wavefront sensor (HSWF) and the Deformable Mirror (DM). The calibration procedure began by positioning a fluorescent bead in the field of view. An interaction matrix was then computed between HSWF and DM. In the algorithm implemented in the CASAOTM software, information from the HSWF sensor is decomposed into the orthogonal modes of the DM using a standard algorithm of singular value decomposition. A command matrix was then calculated making it possible to correct up to 40 modes. Prior to the calibration procedure, the “static” aberrations due to the non-common path between the HSWF and the camera were measured and were compensated for. The purpose of this step was to set the DM shape so that it would correct imaging path aberrations. The corresponding DM shape was referred to as “AO-off”. For more details see reference²¹.

Cell culture. HCT116 wild type and HCT116-mCherry_H2B cells were cultured in DMEM + GlutaMAX (Dulbecco's Modified Eagle Medium; Gibco), supplemented with 10% fetal bovine serum and 1% penicillin-streptomycin (Pen Strep; Gibco) and maintained at 37 °C with 5% CO₂ in an incubator.

MCTS production. To produce the MCTS, we used the centrifugation method described in³. MCTS were prepared in 96-well plates that were coated with 20 mg/ml polyHEMA (Sigma). Cells were plated at a density of 600 cells/well in 100 µl cell culture medium per well then centrifuged to enable MCTS formation. After 4–5 days of growth, MCTS of 400 µm in diameter were collected, washed three times with PBS and then fixed with 10% neutral buffered formalin (Sigma-Aldrich) at room temperature over 2 hours.

EdU labelling of MCTS. EdU labelling was performed by using the Click EdU Alexa Fluor imaging kit (Molecular Probes). Briefly, EdU (5-ethynyl-2'-deoxyuridine) is a thymidine analogue which can be incorporated into the DNA during the DNA synthesis phase (phase S) in proliferating cells. EdU detection relies on a simple and quick click reaction, a copper-catalyzed covalent reaction between an azide and an alkyne, that does not necessitate a DNA denaturation step such as BrdU detection. EdU was added to the culture medium to obtain a final concentration of 10 µM. After 24 hours of incubation at 37 °C, MCTS were washed in PBS and then fixed. EdU detection, based on the click reaction between EdU and Alexa Fluor 594 dye, was performed following the manufacturer's instructions.

Immunofluorescence. Following the clearing protocol, immunohistological studies can be successfully carried out by allowing for penetration of biomolecules. Cleared-MCTS were washed in PBS/0.5% BSA/0.1% Triton v/v. Incubation was carried out with antibodies against αTubulin (mouse polyclonal, Sigma, 1/2000 at 4 °C, for 4 days). After being washed in PBS/0.1% Triton v/v, the secondary antibody was added (anti-mouse conjugated with Alexa594, Molecular Probes, 1/800, at room temperature, over a 4-day period).

CUBIC protocol. CUBIC (Clear, Unobstructed Brain Imaging Cocktail) is a chemical reversible method adapted for clearing large structure like mouse organs⁹. Two main reagents are used: Reagent 1 (R1) was prepared as a mixture of 25 wt% urea (U5378, Sigma Aldrich) 25 wt% N,N,N',N'-tetrakis (2 hydroxypropyl) ethylenediamine (122262, Sigma Aldrich) and 15 wt% polyethylene glycol mono-*p*-isooctylphenyl ether/Triton X-100 (2000, EUROMEDEX). Reagent 2 (R2) was prepared with 50 wt% sucrose, 25 wt% urea, 10 wt% 2,2',2''-nitrilotriethanol (90279, Sigma Aldrich) and 0.1 v/v% Triton X-100 (2000, EUROMEDEX).

For MCTS clearing, samples were immersed in R1 at 37 °C with gentle shaking for either 2 or 4 days, after which the solution was changed and the sample immersed in the same volume of fresh R1 for an additional 3–4 days. Then different washing steps were carried out at room temperature to gently adapt the refractive index inside the MCTS: 1) with a PBS solution, several times, 2) with a PBS/sucrose 20% solution for a minimum of 2 hours, and 3) with a PBS/glycerol 50% solution for a minimum of 2 hours. MCTS were then immersed in R2 at 37 °C with gentle shaking for 1–2 days. For storage, the samples were placed in PBS at 4 °C after R1 step. Washing and R2 immersion steps were pursued before imaging experiments. For _{WAO}SPIM imaging, CUBIC-cleared MCTS were embedded in a cylinder made of a mixture of low-melting 2% agarose and clearing R2 agent (usually one third for two thirds of final volume).

CLARITY protocol. CLARITY is an irreversible clearing technology developed around non-reversible chemical transformation²⁹. This method brings transparency and permeability to macromolecules by removing lipids. Biological samples are transformed into a stable macromolecule-permeable structure. MCTS are embedded for 4 days at 4 °C in a hydrogel mixture of 3% Acrylamide /BisAcrylamide (A4058, Sigma Aldrich), 0.15% N,N,N',N'-Tetramethylethylenediamine (T9281, Sigma Aldrich), 0.007% Ammonium persulfate (Sigma Aldrich), 3.6% Formaldehyde (F8775, Sigma Aldrich). The second solution is the clearing reagent made with 4% SDS (Ambion), boric acid (Sigma Aldrich), with sufficient H₂O, to reach PH 8.5 NaOH (Sigma Aldrich). After a polymerization step of at least 3 hours at 43 °C, the clearing solution was added and maintained for 7 days at 43 °C. CLARITY-cleared MCTS can be stored

at 4°C in PBS. Transparency of the MCTS was sufficiently stable in PBS to either embed them in a 1% agarose cylinder or to put them into a sample chamber inside a phytigel container filled with PBS as described in Desmaison & *al.*, 2012.

Imaging processing. Images were processed with the open-source image-processing Fiji package and with MATLAB. SPIM images were registered with the StackReg plugin available in the Fiji software. For image quality assessment (spatial frequency comparison and intensity gradient maps), Matlab codes were developed and are available on request.

References

- Huisken, J., Swoger, J., Del Bene, F., Wittbrodt, J. & Stelzer, E. H. Optical sectioning deep inside live embryos by selective plane illumination microscopy. *Science* **305**, 1007–1009 (2004).
- Keller, P. J., Schmidt, A. D., Wittbrodt, J. & Stelzer, E. H. Reconstruction of zebrafish early embryonic development by scanned light sheet microscopy. *Science* **322**, 1065–1069 (2008).
- Lorenzo, C. *et al.* Live cell division dynamics monitoring in 3D large spheroid tumor models using light sheet microscopy. *Cell Div* **6**, 22. (2011)
- Chen, B. C. *et al.* Lattice light-sheet microscopy: imaging molecules to embryos at high spatiotemporal resolution. *Science* **346**, 1257998 (2014).
- Dodt, H. U. *et al.* Ultramicroscopy: three-dimensional visualization of neuronal networks in the whole mouse brain. *Nat Methods* **4**, 331–336 (2007).
- Chylek, P. Absorption and scattering of light by small particles. By C. F. Bohren and d. R. Huffman. *Appl Opt* **25**, 3166 (1986).
- Becker, K., Jahrling, N., Saghafi, S. & Dodt, H. U. Immunostaining, dehydration, and clearing of mouse embryos for ultramicroscopy. *Cold Spring Harb Protoc* 2013, 743–744 (2013).
- Ertürk, A. *et al.* Three-dimensional imaging of solvent-cleared organs using 3DISCO. *Nat Protoc* **7**, 1983–1995 (2012).
- Tainaka, K. *et al.* Whole-body imaging with single-cell resolution by tissue decolorization. *Cell* **159**, 911–924 (2014).
- Tomer, R., Ye, L., Hsueh, B. & Deisseroth, K. Advanced CLARITY for rapid and high-resolution imaging of intact tissues. *Nat Protoc* **9**, 1682–1697 (2014).
- Costantini, I. *et al.* A versatile clearing agent for multi-modal brain imaging. *Sci Rep* **5**, 9808 (2015).
- Diaspro, A., Federici, F. & Robello, M. Influence of refractive-index mismatch in high-resolution three-dimensional confocal microscopy. *Appl Opt* **41**, 685–690 (2002).
- Marx, V. Microscopy: seeing through tissue. *Nat Methods* **11**, 1209–1214 (2014).
- Silvestri, L., Sacconi, L. & Pavone, F. S. Correcting spherical aberrations in confocal light sheet microscopy: a theoretical study. *Microsc. Res. Tech.* **77**, 483–491 (2014).
- Tyson, R. K. *Principles of Adaptive Optics*. Academic press, San Diego CA (1991).
- Booth, M. J. Adaptive optics in microscopy. *Philos. Transac. A Math. Phys. Eng. Sci.* **365**, 2829–2843 (2007).
- Kner, P., Sedat, J. W., Agard, D. A. & Kam, Z. High-resolution wide-field microscopy with adaptive optics for spherical aberration correction and motionless focusing. *J. Microsc* **237**, 136–147 (2010).
- Tao, X. *et al.* Adaptive optics confocal microscopy using direct wavefront sensing. *Opt. Lett* **36**, 1062–1064 (2011).
- Wang, K. *et al.* Rapid adaptive optical recovery of optimal resolution over large volumes. *Nat methods* **11**, 625–628 (2014).
- Fraiser, V. *et al.* Adaptive optics in spinning disk microscopy: improved contrast and brightness by a simple and fast method. *J. Microsc* **00**, 1–9 (2015).
- Jorand, R. *et al.* Deep and clear optical imaging of thick inhomogeneous samples. *PLoS One* **7**, e35795.
- Sutherland, R. M. Cell and environment interactions in tumor microregions: the multicell spheroid model. *Science* **240**, 177–184 (1988).
- Susaki, E. A. *et al.* Whole-brain imaging with single-cell resolution using chemical cocktails and computational analysis. *Cell* **157**, 726–739 (2014).
- Török, P., V., P. & Németh, G. Analytical solution of the diffraction integrals and interpretation of wave-front distortion when light is focused through a planar interface between materials of mismatched refractive indices. *J. Opt. Soc. Am* **12**, 2660–2671 (1995).
- Booth, M. J. & W., T. Refractive-index-mismatch induced aberrations in single-photon and two-photon microscopy and the use of aberration correction. *J. Biomed. Opt.* **6**, 266–272 (2001).
- Bourgenot, C., Saunter, C. D., Taylor, J. M., Girkin, J. M. & Love, G. D. 3D adaptive optics in a light sheet microscope. *Opt Express* **20**, 13252–13261 (2012).
- Turaga, D. & Holy, T. E. Aberrations and their correction in light-sheet microscopy: a low-dimensional parametrization. *Biomed Opt Express* **4**, 1654–1661 (2013).
- Krishnan, D., Tay, T. & Fergus, R. Blind deconvolution using a normalized sparsity measure. In: 2011 IEEE conference on Computer Vision and Pattern Recognition (CVPR), 233–240 (June 2011).
- Chung, K. & Deisseroth, K. CLARITY for mapping the nervous system. *Nat Methods* **10**, 508–513 (2013).
- Lee, H., Park, J. H., Seo, I., Park, S. H. & Kim, S. Improved application of the electrophoretic tissue clearing technology, CLARITY, to intact solid organs including brain, pancreas, liver, kidney, lung, and intestine. *BMC Dev Biol* **14**, 48 (2014).
- Booth, M. J., Neil, M. A., Juskaitis, R. & Wilson, T. Adaptive aberration correction in a confocal microscope. *Proc Natl Acad Sci USA* **99**, 5788–5792 (2002).

Acknowledgements

This work was supported by the CNRS, l'Université Paul Sabatier, l'Agence Nationale de la Recherche (ANR12-BSV5-0008-01) and the Toulouse University IDEX Emergence programme. The authors wish to acknowledge the TRI-Genotoul and ITAV imaging facilities. We would like to thank G. Taillefer for manuscript proofreading and editing suggestions as well as Pierre Weiss, Jordi Andilla and Xavier Levêque for helpful comments on the manuscript.

Author Contributions

A.M. and C.L. conceived of and designed the experiments. A.M., C.F. and G.C. performed the experiments. A.M. and P.E. conceived of and designed the RD image quality metric. A.M., P.E. and C.L. analysed the data. A.M., B.D. and C.L. wrote the manuscript with the input of all authors. C.L. supervised the project.

Additional Information

Supplementary information accompanies this paper at <http://www.nature.com/srep>

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Masson, A. *et al.* High-resolution in-depth imaging of optically cleared thick samples using an adaptive SPIM. *Sci. Rep.* 5, 16898; doi: 10.1038/srep16898 (2015).



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>

Publications

Journaux

- P. Escande, P. Weiss, *Sparse Wavelet Representations of Spatially Varying Blurring Operators*, SIAM Journal on Imaging Sciences (SIIMS), 2015, vol. 8, no 4, p. 2976-3014.
- P. Escande, P. Weiss, W. Zhang, *A Variational Model for Multiplicative Structured Noise Removal*, Accepted for publication in JMIV, 2016
- A. Masson, P. Escande, C. Frongia, G. Clouvel, B. Ducommun, C. Lorenzo, *High-resolution in-depth imaging of optically cleared thick samples using an adaptive SPIM*, Nature Scientific Reports, 2015.
- P. Escande, P. Weiss, *Real-time $l1$ - $l2$ deblurring using wavelet expansions of operators*, Submitted 2015.
- P. Escande, P. Weiss, *Approximation of integral operators using product-convolution expansions*, Submitted 2016.
- P. Escande, J. Bigot, P. Weiss, *Identification of linear operators from scattered impulse responses*, In preparation.
- P. Weiss, P. Escande, Y. Dong, *Contrast Invariant SNR*, Submitted 2016.

Conférences

- P. Escande, P. Weiss, F. Malgouyres, *Spatially Varying Blur Recovery. Diagonal Approximations in the Wavelet Domain*, Proc. ICPRAM, 2013.
- P. Escande, P. Weiss, F. Malgouyres, *Image Restoration Using Sparse Approximations of Spatially Varying Blur Operators in the Wavelet Domain.*, Proc. NCMIP, 2013.
- P. Escande, P. Weiss, *Approximation et Estimation des Opérateurs de Flou Variable*, Proc. GRETSI 2015.

Relecteur pour des revues internationales

- JMIV : Journal of Mathematical Imaging and Vision
- SIAM Journal on Imaging Sciences
- Nature and Methods
- Sensors

Exposés

- September 2016 : MAIA 2016, Multivariate Approximation and Interpolation with Applications, CIRM Marseille (France)
- June 2016 : GDR ISIS, Telecom ParisTech, Paris (France)
Changement de contraste local optimal
- May 2016 : CANUM 2016, Obernai (France)
Approximation and estimation of spatially varying operators
- November 2015 : BioImage Informatics 2015, GDR 2588, Institut Pasteur, Paris (France)
On wavelet expansion of blurring operators and its application to deconvolution
- September 2015 : COSMOSTAT Journal Club 2015, CEA Saclay (France)
On wavelet expansion of blurring operators
- September 2015 : GRETSI 2015, Lyon (France)
Approximation and estimation of spatially varying blurring operators
- June 2015 : SMAI 2015, Les Karellis, Savoie (France)
Approximation and estimation of spatially varying blurring operators
- November 2014 : GDR ISIS, Telecom ParisTech, Paris (France)
Approximation and estimation of spatially varying blurring operators
- May 2014 : ICCHA 5, Vanderbilt University, Nashville TN (USA)
Sparse Approximations of Spatially Varying Blur Operators in the Wavelet Domain
- May 2014 : SIAM Conference on IMAGING SCIENCE, Baptist University, Hong-Kong
Sparse Approximations of Spatially Varying Blur Operators in the Wavelet Domain
- May 2013 : NCMIP, ENS Cachan (France)
Image restoration using sparse approximations of spatially varying blur operators in the wavelet domain
- February 2013 : ICPRAM 2013, Barcelona (Spain)
Spatially Varying Blur Recovery. Diagonal Approximations in the Wavelet Domain.

Enseignement

2015-2016

- L2 SPIEEA, Université Paul Sabatier Toulouse
Géométries 2D et 3D, Algèbre linéaire, Calcul vectoriel (24h)
- L3 SID Probabilités, Université Paul Sabatier Toulouse
Variables aléatoires, Convergence, Théorème central limite (14h)
- L1 Outils Maths, Université Paul Sabatier Toulouse
Trigonométrie, Algèbre, EDO (22h)

2014-2015

- L2 SPIEEA, Université Paul Sabatier Toulouse
Géométries 2D et 3D, Algèbre linéaire, Calcul vectoriel (11h)
- L1 SFA Math 4, Université Paul Sabatier Toulouse
Outils d'analyse des suites (48h)
- L2 Math, Université Paul Sabatier Toulouse
Développement du projet professionnel (8h)

2013-2014

- L1 SFA Math 4, Université Paul Sabatier Toulouse
Outils d'analyse des suites (48h)
- L3 SID Inferential Statistics, Université Paul Sabatier Toulouse
Théorie de l'estimation statistique et tests gaussiens (16h)
- Projet 4A GMM, INSA de Toulouse
Une première approche de la déconvolution aveugle (12h)

Bibliographie

- [1] Ben Adcock and Anders C Hansen. A generalized sampling theorem for stable reconstructions in arbitrary bases. *Journal of Fourier Analysis and Applications*, 18(4) :685–716, 2012.
- [2] Igor Aizenberg and Constantine Butakoff. A windowed Gaussian notch filter for quasi-periodic noise removal. *Image and Vision Computing*, 26(10) :1347–1353, 2008.
- [3] Bradley K Alpert. A class of bases in L^2 for the sparse representation of integral operators. *SIAM journal on Mathematical Analysis*, 24(1) :246–262, 1993.
- [4] Emran Mohammad Abu Anas, Soo Yeol Lee, and Md Kamrul Hasan. Classification of ring artifacts for their effective removal using type adaptive correction schemes. *Computers in biology and medicine*, 41(6) :390–401, 2011.
- [5] Edward Angel and Anil K. Jain. Restoration of images degraded by spatially varying pointspread functions by a conjugate gradient method. *Applied Optics*, 17(14) :2186–2190, 1978.
- [6] Rémi Arcangéli, María Cruz López de Silanes, and Juan José Torrens. An extension of a bound for functions in Sobolev spaces, with applications to (m, s)-spline interpolation and smoothing. *Numerische Mathematik*, 107(2) :181–211, 2007.
- [7] Muthuvel Arigovindan, Joshua Shaevitz, John McGowan, John W Sedat, and David A. Agard. A parallel product-convolution approach for representing the depth varying point spread functions in 3d widefield microscopy based on principal component analysis. *Optics express*, 18(7) :6461–6476, 2010.
- [8] Hédý Attouch, Jérôme Bolte, Patrick Redont, and Antoine Soubeyran. Proximal alternating minimization and projection methods for nonconvex problems : an approach based on the Kurdyka-Lojasiewicz inequality. *Mathematics of Operations Research*, 35(2) :438–457, 2010.
- [9] Gilles Aubert and Jean-Francois Aujol. A variational approach to removing multiplicative noise. *SIAM Journal on Applied Mathematics*, 68(4) :925–946, 2008.

- [10] Gilles Aubert and Pierre Kornprobst. *Mathematical problems in image processing : partial differential equations and the calculus of variations*, volume 147. Springer Science & Business Media, 2006.
- [11] Rodrigo Aviles-Espinosa, Jordi Andilla, Rafael Porcar-Guezenec, Omar E Olarte, Marta Nieto, Xavier Levecq, David Artigas, and Pablo Loza-Alvarez. Measurement and correction of in vivo sample aberrations employing a nonlinear guide-star in two-photon excited fluorescence microscopy. *Biomedical optics express*, 2(11) :3135–3149, 2011.
- [12] Coloma Ballester, Vicent Caselles, Laura Igual, Joan Verdera, and Bernard Rougé. A variational model for P+XS image fusion. *International Journal of Computer Vision*, 69(1) :43–58, 2006.
- [13] W. Ballhaus, John Casani, Steven Dorfman, David Gallagher, Garth Illingworth, John Klineberg, and David Schurr. James Webb Space Telescope (JWST) Independent Comprehensive Review Panel (ICRP). 2010.
- [14] Leah Bar, Nir Sochen, and Nahum Kiryati. Restoration of images with piecewise space-variant blur. In *Scale Space and Variational Methods in Computer Vision*, pages 533–544. Springer, 2007.
- [15] RE Barlow and HD Brunk. The isotonic regression problem and its dual. *Journal of the American Statistical Association*, 67(337) :140–147, 1972.
- [16] Sören Bartels. Total variation minimization with finite elements : Convergence and iterative solution. *SIAM Journal on Numerical Analysis*, 50 :1162–1180, 2012.
- [17] Scott Beaver. Banach algebras of integral operators, off-diagonal decay, and applications in wireless communications. *arXiv preprint math/0406198*, 2004.
- [18] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1) :183–202, 2009.
- [19] Amir Beck and Marc Teboulle. Gradient-based algorithms with applications to signal recovery. *Convex Optimization in Signal Processing and Communications*, 2009.
- [20] Rosalie Bélanger-Rioux and Laurent Demanet. Compressed absorbing boundary conditions via matrix probing. *SIAM Journal on Numerical Analysis*, 53(5) :2441–2471, 2015.
- [21] Anthony J Bell and Terrence J Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6) :1129–1159, 1995.

- [22] Philip A Bello. Characterization of randomly time-variant linear channels. *Communications Systems, IEEE transactions on*, 11(4) :360–393, 1963.
- [23] Saïma Ben Hadj and Laure Blanc-Féraud. Restoration method for spatially variant blurred images. Research report, INRIA, June 2011.
- [24] Saïma Ben Hadj and Laure Blanc Féraud. *Restoration method for spatially variant blurred images*. PhD thesis, INRIA, 2011.
- [25] Alexis Benichoux, Emmanuel Vincent, and Rémi Gribonval. A fundamental pitfall in blind deconvolution with sparse and shift-invariant priors. In *ICASSP-38th International Conference on Acoustics, Speech, and Signal Processing-2013*, 2013.
- [26] Joel Bergé, Sedona Price, Adam Amara, and Jason Rhodes. On point spread function modelling : towards optimal interpolation. *Monthly Notices of the Royal Astronomical Society*, 419(3) :2356–2368, 2012.
- [27] Michael J Best and Nilotpal Chakravarti. Active set algorithms for isotonic regression ; a unifying framework. *Mathematical Programming*, 47(1-3) :425–439, 1990.
- [28] Gregory Beylkin. On the representation of operators in bases of compactly supported wavelets. *SIAM Journal on Numerical Analysis*, 29(6) :1716–1740, 1992.
- [29] Gregory Beylkin, Ronald Coifman, and Vladimir Rokhlin. Fast wavelet transforms and numerical algorithms I. *Communications on pure and applied mathematics*, 44(2) :141–183, 1991.
- [30] Anatoly Yu Bezhaev and Vladimir Aleksandrovich Vasilenko. *Variational theory of splines*. Springer, 2001.
- [31] Jérémie Bigot, Claire Boyer, and Pierre Weiss. An analysis of block sampling strategies in compressed sensing. *arXiv preprint arXiv :1305.4446*, 2013.
- [32] F Edward Boas and Dominik Fleischmann. CT artifacts : causes and reduction techniques. *Imaging in Medicine*, 4(2) :229–240, 2012.
- [33] Jérôme Bolte, Shoham Sabach, and Marc Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1-2) :459–494, 2014.
- [34] Steffen Börm, Lars Grasedyck, and Wolfgang Hackbusch. Introduction to hierarchical matrices with applications. *Engineering Analysis with Boundary Elements*, 27(5) :405–422, 2003.
- [35] Max Born and Emil Wolf. *Principles of optics : electromagnetic theory of propagation, interference and diffraction of light*. CUP Archive, 1999.
- [36] Alan C Bovik. *Handbook of image and video processing*. Academic press, 2010.

- [37] Claire Boyer, Pierre Weiss, and Jérémie Bigot. An algorithm for variable density sampling with block-constrained acquisition. *SIAM Journal on Imaging Sciences*, 7(2) :1080–1107, 2014.
- [38] Martin D Buhmann. Radial basis functions : theory and implementations. *Cambridge Monographs on Applied and Computational Mathematics*, 12 :147–165, 2003.
- [39] Robert Busby and Harvey Smith. Product-convolution operators and mixed-norm spaces. *Transactions of the American Mathematical Society*, 263(2) :309–341, 1981.
- [40] Jian-Feng Cai and Zuowei Shen. Framelet based deconvolution. *J. Comput. Math*, 28(3) :289–308, 2010.
- [41] Alberto P Calderon and Antoni Zygmund. On the existence of certain singular integrals. *Acta Mathematica*, 88(1) :85–139, 1952.
- [42] Alberto P Calderón and Antoni Zygmund. On singular integrals. *American Journal of Mathematics*, 78(2) :289–309, 1956.
- [43] Emmanuel Candes and Laurent Demanet. Curvelets and Fourier integral operators. *Comptes Rendus Mathématique*, 336 :395–398, 2003.
- [44] Emmanuel Candes, Laurent Demanet, and Lexing Ying. Fast computation of Fourier integral operators. *SIAM Journal on Scientific Computing*, 29(6) :2464–2493, 2007.
- [45] Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6) :717–772, 2009.
- [46] Emmanuel J Candes, Justin K Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on pure and applied mathematics*, 59(8) :1207–1223, 2006.
- [47] Vicent Caselles, Bartomeu Coll, and Jean-Michel Morel. Topographic maps and local contrast changes in natural images. *International Journal of Computer Vision*, 33(1) :5–27, 1999.
- [48] Vicent Caselles, Bartomeu Coll, and Jean-Michel Morel. Geometry and color in natural images. *Journal of Mathematical Imaging and Vision*, 16(2) :89–105, 2002.
- [49] Vicent Caselles and Pascal Monasse. *Geometric description of images as topographic maps*. Springer, 2009.
- [50] Olivier Chabiron, Francois Malgouyres, Jean-Yves Tournet, and Nicolas Dobigeon. Toward fast transform learning. *International Journal of Computer Vision*, 114(2-3) :195–216, 2015.
- [51] Anwei Chai and Zuowei Shen. Deconvolution : A wavelet frame approach. *Numerische Mathematik*, 106(4) :529–587, 2007.

- [52] Antonin Chambolle. An algorithm for total variation minimization and applications. *Journal of Mathematical imaging and vision*, 20(1-2) :89–97, 2004.
- [53] Antonin Chambolle and Charles Dossal. On the convergence of the iterates of FISTA. *Preprint hal-01060130*, September, 2014.
- [54] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40 :120–145, 2011.
- [55] Tony F Chan and Chiu-Kwong Wong. Total variation blind deconvolution. *Image Processing, IEEE Transactions on*, 7(3) :370–375, 1998.
- [56] Chihway Chang, PJ Marshall, JG Jernigan, JR Peterson, SM Kahn, SF Gull, Y AlSayyad, Z Ahmad, J Bankert, D Bard, et al. Atmospheric point spread function interpolation for weak lensing in short exposure imaging data. *Monthly Notices of the Royal Astronomical Society*, 427(3) :2572–2587, 2012.
- [57] Ee-Chien Chang, Stephane Mallat, and Chee Yap. Wavelet foveation. *Applied and Computational Harmonic Analysis*, 9(3) :312–335, 2000.
- [58] Yi Chang, Houzhang Fang, Luxin Yan, and Hai Liu. Robust destriping method with unidirectional total variation and framelet regularization. *Optics express*, 21(20) :23307–23323, 2013.
- [59] Caroline Chaux, Patrick L Combettes, Jean-Christophe Pesquet, and Valérie R Wajs. A variational formulation for frame-based inverse problems. *Inverse Problems*, 23(4) :1495, 2007.
- [60] Scott Shaobing Chen, David L Donoho, and Michael A Saunders. Atomic decomposition by basis pursuit. *SIAM review*, 43(1) :129–159, 2001.
- [61] Shu-W Chen and Jean-Luc Pellequer. Destripe : frequency-based algorithm for removing stripe noises from AFM images. *BMC structural biology*, 11(1) :7, 2011.
- [62] Jiawei Chiu and Laurent Demanet. Matrix probing and its conditioning. *SIAM Journal on Numerical Analysis*, 50(1) :171–193, 2012.
- [63] Ole Christensen. *An introduction to frames and Riesz bases*. Springer Science & Business Media, 2013.
- [64] Moody Chu, Robert Funderlic, and Robert Plemmons. Structured low rank approximation. *Linear algebra and its applications*, 366 :157–172, 2003.
- [65] Philippe G. Ciarlet. *The finite element method for elliptic problems*, volume 40. Siam, 2002.
- [66] Jon F Claerbout. Fundamentals of geophysical data processing. 1985.

- [67] Albert Cohen. *Numerical analysis of wavelet methods*, volume 32. Elsevier, 2003.
- [68] Albert Cohen, Wolfgang Dahmen, Ingrid Daubechies, and Ronald DeVore. Harmonic analysis of the space BV . *Revista Matemática Iberoamericana*, 19(1) :235–263, 2003.
- [69] Albert Cohen, Wolfgang Dahmen, and Ronald DeVore. Adaptive wavelet methods for elliptic operator equations : convergence rates. *Mathematics of Computation*, 70(233) :27–75, 2001.
- [70] Albert Cohen, Wolfgang Dahmen, and Ronald A DeVore. Adaptive wavelet methods II-beyond the elliptic case. *Foundations of Computational Mathematics*, 2 :203–245, 2002.
- [71] Albert Cohen, Ingrid Daubechies, and Pierre Vial. Wavelets on the interval and fast wavelet transforms. *Applied and Computational Harmonic Analysis*, 1(1) :54–81, 1993.
- [72] Albert Cohen, Ronald DeVore, Pencho Petrushev, and Hong Xu. Nonlinear approximation and the space $BV(\mathbb{R}^2)$. *American Journal of Mathematics*, pages 587–628, 1999.
- [73] Patrick L Combettes and Jean-Christophe Pesquet. Proximal splitting methods in signal processing. In *Fixed-point algorithms for inverse problems in science and engineering*, pages 185–212. Springer, 2011.
- [74] James W Cooley and John W Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of computation*, 19(90) :297–301, 1965.
- [75] Bruno Cornelis, Ann Doods, Jan Cornelis, and Peter Schelkens. Digital canvas removal in paintings. *Signal Processing*, 92(4) :1166–1171, 2012.
- [76] W. Dahmen, S. Prössdorf, and R. Schneider. Wavelet approximation methods for pseudodifferential equations II : matrix compression and fast solution. *Advances in Computational Mathematics*, 1 :259–335, October 1993.
- [77] Wolfgang Dahmen. Wavelet and multiscale methods for operator equations. *Acta numerica*, 6 :55–228, 1997.
- [78] Ingrid Daubechies. Orthonormal bases of compactly supported wavelets. *Communications on pure and applied mathematics*, 41(7) :909–996, 1988.
- [79] Ingrid Daubechies. *Ten lectures on wavelets*, volume 61. SIAM, 1992.
- [80] Guy David and Jean-Lin Journé. A boundedness criterion for generalized Calderón-Zygmund operators. *The Annals of Mathematics*, 120 :371–397, 1984.

- [81] Guy David, Jean-Lin Journé, and Stephen Semmes. Opérateurs de Calderón-Zygmund, fonctions para-accrétives et interpolation. *Revista Matemática Iberoamericana*, 1(4) :1–56, 1985.
- [82] C.A. Deledalle, L. Denis, and F. Tupin. Iterative Weighted Maximum Likelihood Denoising with Probabilistic Patch-Based Weights. *IEEE Transactions on Image Processing*, 18(12) :2661–2672, 2009.
- [83] Julie Delon. Midway image equalization. *Journal of Mathematical Imaging and Vision*, 21(2) :119–134, 2004.
- [84] Laurent Demanet. *Curvelets, wave atoms, and wave equations*. PhD thesis, California Institute of Technology, 2006.
- [85] Loïc Denis, Eric Thiébaud, Ferréol Soulez, Jean-Marie Becker, and Rahul Mourya. Fast approximations of shift-variant blur. *International Journal of Computer Vision*, 115(3) :253–278, 2015.
- [86] Winfried Denk and Karel Svoboda. Photon upmanship : why multiphoton imaging is more than a gimmick. *Neuron*, 18(3) :351–357, 1997.
- [87] Jacques Deny and Jacques-Louis Lions. Les espaces du type de Beppo Levi. In *Annales de l’institut Fourier*, volume 5, pages 305–370. Institut Fourier, 1954.
- [88] Ronald A. DeVore. Nonlinear approximation. *Acta numerica*, 7 :51–150, 1998.
- [89] Ronald A. DeVore and George G. Lorentz. *Constructive approximation*, volume 303. Springer Science & Business Media, 1993.
- [90] Yiqiu Dong and Tieyong Zeng. A convex variational model for restoring blurred images with multiplicative noise. *SIAM Journal on Imaging Sciences*, 6(3) :1598–1625, 2013.
- [91] David L. Donoho. Interpolating wavelet transforms. *Preprint, Department of Statistics, Stanford University*, 2(3), 1992.
- [92] David L Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4) :1289–1306, 2006.
- [93] Monika Dörfler and Bruno Torrèsani. Representation of operators in the time-frequency domain and generalized Gabor multipliers. *Journal of Fourier Analysis and Applications*, 16(2) :261–293, 2010.
- [94] Jean Duchon. Sur l’erreur d’interpolation des fonctions de plusieurs variables par les D^m -splines. *RAIRO-Analyse numérique*, 12(4) :325–334, 1978.
- [95] Sylvain Durand, Jalal Fadili, and Mila Nikolova. Multiplicative noise removal using l1 fidelity on frame coefficients. *Journal of Mathematical Imaging and Vision*, 36(3) :201–226, 2010.

- [96] Richard L Dykstra, Tim Robertson, et al. An algorithm for isotonic regression for two or more independent variables. *The Annals of Statistics*, 10(3) :708–716, 1982.
- [97] Matthias Joachim Ehrhardt and Simon R Arridge. Vector-valued image processing by parallel level sets. *Image Processing, IEEE Transactions on*, 23(1) :9–18, 2014.
- [98] Paul Escande, Jérémie Bigot, and Pierre Weiss. Identification of linear operators from scattered impulse responses. 2016.
- [99] Paul Escande and Pierre Weiss. Real-time ℓ^1 – ℓ^2 deblurring using wavelet expansions of operators. *arXiv preprint arXiv :1512.08401*, 2015.
- [100] Paul Escande and Pierre Weiss. Sparse wavelet representations of spatially varying blurring operators. *SIAM Journal on Imaging Sciences*, 8(4) :2976–3014, 2015.
- [101] Paul Escande and Pierre Weiss. Approximation of integral operators using convolution-product expansions. *arXiv preprint arXiv :1604.04374*, 2016.
- [102] Paul Escande, Pierre Weiss, and François Malgouyres. Image restoration using sparse approximations of spatially varying blur operators in the wavelet domain. *Proc. NCMIP*, 2013.
- [103] Paul Escande, Pierre Weiss, and François Malgouyres. Spatially varying blur recovery. diagonal approximations in the wavelet domain. *Proceedings of ICPRAM*, 2013.
- [104] Paul Escande, Pierre Weiss, and Wenxing Zhang. A variational model for multiplicative structured noise removal. 2016.
- [105] Claudio Estatico and Fabio Di Benedetto. Shift-invariant approximations of structured shift-variant blurring matrices. *Numerical Algorithms*, 62(4) :615–635, 2013.
- [106] Jérôme Fehrenbach and Pierre Weiss. Processing stationary noise : Model and parameter selection in variational methods. *SIAM Journal on Imaging Sciences*, 7(2) :613–640, 2014.
- [107] Jérôme Fehrenbach, Pierre Weiss, and Corinne Lorenzo. Variational algorithms to remove stationary noise : applications to microscopy imaging. *Image Processing, IEEE Transactions on*, 21(10) :4420–4430, 2012.
- [108] Hans G Feichtinger, Mario Hampejs, and Günther Kracher. Approximation of matrices by Gabor multipliers. *Signal Processing Letters, IEEE*, 11(11) :883–886, 2004.
- [109] Hans G Feichtinger and Krzysztof Nowak. A first survey of Gabor multipliers. In *Advances in Gabor analysis*, pages 99–128. Springer, 2003.
- [110] Mário AT Figueiredo and Robert D Nowak. An EM algorithm for wavelet-based image restoration. *Image Processing, IEEE Transactions on*, 12(8) :906–916, 2003.

- [111] Jan Henrik Fitschen, Jianwei Ma, and Sebastian Schuff. Removal of curtaining effects by a variational model with directional first and second order differences. *arXiv preprint arXiv :1507.00112*, 2015.
- [112] Ralf C Flicker and François J Rigaut. Anisoplanatic deconvolution of adaptive optics images. *Journal of the Optical Society of America*, 22(3) :504–513, 2005.
- [113] M Gentile, F Courbin, and G Meylan. Interpolating point spread function anisotropy. *Astronomy & Astrophysics*, 549 :A1, 2013.
- [114] Erez Gilad and Jost Von Hardenberg. A fast algorithm for convolution integrals with space and time variant kernels. *Journal of Computational Physics*, 216(1) :326–336, 2006.
- [115] Roland Glowinski. *Numerical methods for non-linear variational problems*. Springer, 1984.
- [116] Gene H Golub and Charles F Van Loan. *Matrix computations*, volume 3. JHU Press, 2012.
- [117] Luis Gómez-Chova, Luis Alonso, Luis Guanter, Gustavo Camps-Valls, Javier Calpe, and José Moreno. Correction of systematic spatial noise in push-broom hyperspectral sensors : application to chris/proba images. *Applied optics*, 47(28) :F46–F60, 2008.
- [118] Adriana González, Véronique Delouille, and Laurent Jacques. Non-parametric PSF estimation from celestial transit solar images using blind deconvolution. *Journal of Space Weather and Space Climate*, 6 :A1, 2016.
- [119] Leslie Greengard and Vladimir Rokhlin. A fast algorithm for particle simulations. *Journal of computational physics*, 73(2) :325–348, 1987.
- [120] Rémi Gribonval. Should penalized least squares regression be interpreted as maximum a posteriori estimation? *IEEE Transactions on Signal Processing*, 59(5) :2405–2410, 2011.
- [121] L.J. Griffiths, F.R. Smolka, and L.D. Trembly. Adaptive deconvolution : A new technique for processing time-varying seismic data. *Geophysics*, 42(4) :742–759, 1977.
- [122] Marcus J Grote and Thomas Huckle. Parallel preconditioning with sparse approximate inverses. *SIAM Journal on Scientific Computing*, 18(3) :838–853, 1997.
- [123] Matthieu Guerquin-Kern, L Lejeune, Klaas P Pruessmann, and Michael Unser. Realistic analytical phantoms for parallel magnetic resonance imaging. *Medical Imaging, IEEE Transactions on*, 31(3) :626–636, 2012.
- [124] Wolfgang Hackbusch. A sparse matrix arithmetic based on $\mathcal{H}$ -matrices. part I : Introduction to $\mathcal{H}$ -matrices. *Computing*, 62(2) :89–108, 1999.

- [125] Wolfgang Hackbusch and Zenon Paul Nowak. On the fast matrix multiplication in the boundary element method by panel clustering. *Numerische Mathematik*, 54(4) :463–491, 1989.
- [126] Nasreddine Hajlaoui, Caroline Chaux, Guillaume Perrin, Frédéric Falzon, and Amel Benazza-Benyahia. Satellite image restoration in the context of a spatially varying point spread function. *JOSA A*, 27(6) :1473–1481, 2010.
- [127] Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp. Finding structure with randomness : Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2) :217–288, 2011.
- [128] Per Christian Hansen, James G. Nagy, and Dianne P. O’leary. *Deblurring images : matrices, spectra, and filtering*. Siam, 2006.
- [129] Aleksandr Yakovlevich Helemskii. *Lectures and exercises on functional analysis*, volume 233. American mathematical society Providence, RI, 2006.
- [130] Eugenio Hernández and Guido Weiss. *A first course on wavelets*. CRC press, 1996.
- [131] Michael Hirsch, Christian J Schuler, Stefan Harmeling, and Bernhard Scholkopf. Fast removal of non-uniform camera shake. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 463–470. IEEE, 2011.
- [132] Michael Hirsch, Suvrit Sra, Bernhard Scholkopf, and Stefan Harmeling. Efficient filter flow for space-variant multiframe blind deconvolution. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2010.
- [133] Lars Hörmander. Fourier integral operators. I. *Acta mathematica*, 127(1) :79–183, 1971.
- [134] Lars Hörmander. Pseudo-differential operators of type 1, 1. *Communications in partial differential equations*, 13(9) :1085–1111, 1988.
- [135] Tomasz Hrycak, Saptarshi Das, Gerald Matz, and Hans G. Feichtinger. Low complexity equalization for doubly selective channels modeled by a basis expansion. *Signal Processing, IEEE Transactions on*, 58(11) :5706–5719, 2010.
- [136] Tomasz Hrycak, Saptarshi Das, Gerald Matz, and Hans G Feichtinger. Practical estimation of rapidly varying channels for OFDM systems. *Communications, IEEE Transactions on*, 59(11) :3040–3048, 2011.
- [137] Jiang Hsieh. *Computed tomography : principles, design, artifacts, and recent advances*. SPIE Bellingham, WA, 2009.
- [138] Yu-Mei Huang, Lionel Moisan, Michael K Ng, and Tieyong Zeng. Multiplicative noise removal via a learned dictionary. *Image Processing, IEEE Transactions on*, 21(11) :4534–4543, 2012.

- [139] Jan Huisken, Jim Swoger, Filippo Del Bene, Joachim Wittbrodt, and Ernst HK Stelzer. Optical sectioning deep inside live embryos by selective plane illumination microscopy. *Science*, 305(5686) :1007–1009, 2004.
- [140] Stephane Jaffard. Wavelet methods for fast resolution of elliptic problems. *SIAM Journal on Numerical Analysis*, 29(4) :965–986, 1992.
- [141] MJ Jee, JP Blakeslee, M Sirianni, AR Martel, RL White, and HC Ford. Principal component analysis of the time-and position-dependent point-spread function of the advanced camera for surveys. *Publications of the Astronomical Society of the Pacific*, 119(862) :1403, 2007.
- [142] Raphael Jorand, Gwénaële Le Corre, Jordi Andilla, Amina Maandhui, Céline Frongia, Valérie Lobjois, Bernard Ducommun, and Corinne Lorenzo. Deep and clear optical imaging of thick inhomogeneous samples. *PloS one*, 7(4) :e35795, 2012.
- [143] Thomas Kailath. Sampling models for linear time-variant filters. Master’s thesis, Massachusetts Institute of Technology, Dept. of Electrical Engineering, 1959.
- [144] Julie Kamm and James G. Nagy. Kronecker product and SVD approximations for separable spatially variant blurs. In *SPIE’s International Symposium on Optical Science, Engineering, and Instrumentation*, pages 358–369. International Society for Optics and Photonics, 1998.
- [145] Yitzhak Katznelson. *An introduction to harmonic analysis*. Cambridge University Press, 2004.
- [146] Younguk Kim, Jongduk Baek, and Dosik Hwang. Ring artifact correction using detector line-ratios in computed tomography. *Optics Express*, 22(11) :13380–13392, 2014.
- [147] Hagai Kirshner, Daniel Sage, and Michael Unser. 3D PSF models for fluorescence microscopy in ImageJ. In *Proceedings of the Twelfth International Conference on Methods and Applications of Fluorescence Spectroscopy, Imaging and Probes (MAF’11)*, page 154, 2011.
- [148] Werner Kozek and Andreas F Molisch. Nonorthogonal pulseshapes for multicarrier communications in doubly dispersive channels. *Selected Areas in Communications, IEEE Journal on*, 16(8) :1579–1589, 1998.
- [149] Werner Kozek and Götz E. Pfander. Identification of operators with bandlimited symbols. *SIAM journal on mathematical analysis*, 37(3) :867–888, 2005.
- [150] Aleh Kryvanos, Juergen Hesser, and Gabriele Steidl. Nonlinear image restoration methods for marker extraction in 3D fluorescent microscopy. In *Electronic Imaging 2005*, pages 432–443. International Society for Optics and Photonics, 2005.

- [151] Rasmus Kyng, Anup Rao, and Sushant Sachdeva. Fast, Provable Algorithms for Isotonic Regression in all L_p -norms. In *Advances in Neural Information Processing Systems*, pages 2701–2709, 2015.
- [152] Arthur Leclaire and Lionel Moisan. No-reference image quality assessment and blind deblurring with sharpness metrics exploiting fourier phase information. *Journal of Mathematical Imaging and Vision*, 52(1) :145–172, 2015.
- [153] Anat Levin and Boaz Nadler. Natural image denoising : Optimality and inherent bounds. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 2833–2840. IEEE, 2011.
- [154] Anat Levin, Yael Weiss, Frederic Durand, and William T Freeman. Understanding and evaluating blind deconvolution algorithms. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1964–1971. IEEE, 2009.
- [155] Jingwei Liang, Jalal Fadili, and Gabriel Peyré. Activity identification and local linear convergence of inertial forward-backward splitting. *arXiv preprint arXiv :1503.03703*, 2015.
- [156] Jacques Louis Lions and Enrico Magenes. *Non-homogeneous boundary value problems and applications*, volume 1. Springer Science & Business Media, 2012.
- [157] Cécile Louchet and Lionel Moisan. Posterior expectation of the total variation model : properties and experiments. *SIAM Journal on Imaging Sciences*, 6(4) :2640–2684, 2013.
- [158] Elie Maalouf, Bruno Colicchio, and Alain Dieterlen. Fluorescence microscopy three-dimensional depth variant point spread function interpolation using Zernike moments. *JOSA A*, 28 :1864–1870, 2011.
- [159] Aleandro S Macedo and Elvino S Sousa. Coded OFDM for broadband indoor wireless systems. In *Communications, 1997. ICC’97 Montreal, Towards the Knowledge Millennium. 1997 IEEE International Conference on*, volume 2, pages 934–938. IEEE, 1997.
- [160] Luc Le Magoarou and Rémi Gribonval. Flexible multi-layer sparse approximations of matrices and applications. *arXiv preprint arXiv :1506.07300*, 2015.
- [161] Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro. Online dictionary learning for sparse coding. In *Proceedings of the 26th annual international conference on machine learning*, pages 689–696. ACM, 2009.
- [162] François Malgouyres. A Framework for Image Deblurring Using Wavelet Packet Bases. *Applied and Computational Harmonic Analysis*, 12(3) :309–331, 2002.
- [163] Stéphane Mallat. *A wavelet tour of signal processing*. Academic press, 1999.

- [164] Stéphane Mallat. *A Wavelet Tour of Signal Processing – The Sparse Way*. Third Edition. Academic Press, 2008.
- [165] Aurore Masson, Paul Escande, Céline Frongia, Grégory Clouvel, Bernard Ducommun, and Corinne Lorenzo. High-resolution in-depth imaging of optically cleared thick samples using an adaptive spim. *Scientific reports*, 5, 2015.
- [166] Georges Matheron. *Random sets and integral geometry*. John Wiley & Sons, 1975.
- [167] FM Ngolè Mboula, J-L Starck, Samuel Ronayette, Koryo Okumura, and Jérôme Amiaux. Super-resolution method using sparse regularization for point-spread function recovery. *Astronomy & Astrophysics*, 575 :A86, 2015.
- [168] Scott R McNown and Bobby R Hunt. Approximate shift-invariance by warping shift-variant systems. In *SPIE’s 1994 International Symposium on Optics, Imaging, and Instrumentation*, pages 156–167. International Society for Optics and Photonics, 1994.
- [169] Yves Meyer. Wavelets and operators. *Analysis at Urbana*, 1 :256–365, 1989.
- [170] Yves Meyer. *Wavelets and operators*. Cambridge Univ Press, 1992.
- [171] Yves Meyer and Ronald Coifman. *Wavelets, Calderón-Zygmund and multilinear operators*. 1997.
- [172] C. A. Micchelli and M. Pontil. Kernels for multi-task learning. In *Proceedings of NIPS 2004*, 2004.
- [173] C. A. Micchelli and M. Pontil. On learning vector-valued functions. *Neural Computation*, 17(1) :177–204, 2005.
- [174] Marvin Minsky. Microscopy apparatus, 1961. US Patent 3,013,467.
- [175] David Miraut and Javier Portilla. Efficient shift-variant image restoration using deformable filtering (Part I). *EURASIP J. Adv. Sig. Proc.*, 2012 :100, 2012.
- [176] Pascal Monasse and Frederic Guichard. Fast computation of a contrast-invariant image representation. *IEEE Transactions on Image Processing*, 9(5) :860–872, 2000.
- [177] Peter G Moschopoulos. The distribution of the sum of independent gamma random variables. *Annals of the Institute of Statistical Mathematics*, 37(1) :541–544, 1985.
- [178] Beat Münch, Pavel Trtik, Federica Marone, and Marco Stampanoni. Stripe and ring artifact removal with combined wavelet Fourier filtering. *Optics express*, 17(10) :8567–8591, 2009.
- [179] James G Nagy and Dianne P O’leary. Fast iterative image restoration with a spatially varying PSF. In *Optical Science, Engineering and Instrumentation’97*, pages 388–399. International Society for Optics and Photonics, 1997.

- [180] James G. Nagy and Dianne P. O’Leary. Restoring images degraded by spatially variant blur. *SIAM Journal on Scientific Computing*, 19(4) :1063–1082, 1998.
- [181] Laurent Najman and Michel Couprie. Building the component tree in quasi-linear time. *Image Processing, IEEE Transactions on*, 15(11) :3531–3539, 2006.
- [182] Ramesh Neelamani, Hyeokho Choi, and Richard Baraniuk. Forward : Fourier-wavelet regularized deconvolution for ill-conditioned systems. *Signal Processing, IEEE Transactions on*, 52(2) :418–433, 2004.
- [183] Yurii Nesterov. A method of solving a convex programming problem with convergence rate $o(1/k^2)$. In *Soviet Mathematics Doklady*, volume 27, pages 372–376, 1983.
- [184] Yurii Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1) :125–161, 2013.
- [185] Yurii Nesterov. *Introductory lectures on convex optimization : A basic course*, volume 87. Springer Science & Business Media, 2013.
- [186] Yurii Nesterov, Arkadii Nemirovskii, and Yinyu Ye. *Interior-point polynomial algorithms in convex programming*, volume 13. SIAM, 1994.
- [187] Michael K Ng, Pierre Weiss, and Xiaoming Yuan. Solving constrained total-variation image restoration and reconstruction problems via alternating direction methods. *SIAM journal on Scientific Computing*, 32(5) :2710–2736, 2010.
- [188] Ivan Oseledets and Eugene Tyrtyshnikov. TT-cross approximation for multidimensional arrays. *Linear Algebra and its Applications*, 432(1) :70–88, 2010.
- [189] VY Panin, GL Zeng, and GT Gullberg. Total variation regulated EM algorithm. *Nuclear Science, IEEE Transactions on*, 46(6) :2202–2210, 1999.
- [190] Panos M Pardalos and Guoliang Xue. Algorithms for a class of isotonic regression problems. *Algorithmica*, 23(3) :211–222, 1999.
- [191] Götz E Pfander. Sampling of operators. *Journal of Fourier Analysis and Applications*, 19(3) :612–650, 2013.
- [192] Götz E Pfander, Holger Rauhut, and Jared Tanner. Identification of matrices having a sparse representation. *Signal Processing, IEEE Transactions on*, 56(11) :5376–5388, 2008.
- [193] Allan Pinkus. *N-widths in Approximation Theory*, volume 7. Springer Science & Business Media, 2012.
- [194] Aleksandra Pizurica, Wilfried Philips, Ignace Lemahieu, and Marc Acheroy. A versatile wavelet domain noise filtration technique for medical imaging. *Medical Imaging, IEEE Transactions on*, 22(3) :323–331, 2003.

- [195] Chrysanthé Preza and José-Angel Conchello. Depth-variant maximum-likelihood restoration for three-dimensional fluorescence microscopy. *JOSA A*, 21 :1593–1601, 2004.
- [196] Nelly Pustelnik, Amel Benazza-Benhayia, Yuling Zheng, and Jean-Christophe Pesquet. Wavelet-based Image Deconvolution and Reconstruction. Jun 2015.
- [197] Ralph Tyrell Rockafellar. *Convex analysis*. Princeton university press, 2015.
- [198] June Chul Roh and Bhaskar D. Rao. An efficient feedback method for MIMO systems with slowly time-varying channels. In *Wireless Communications and Networking Conference, 2004. WCNC. 2004 IEEE*, volume 2, pages 760–764. IEEE, 2004.
- [199] Vladimir Rokhlin. Rapid solution of integral equations of classical potential theory. *Journal of Computational Physics*, 60(2) :187–207, 1985.
- [200] Leonid Rudin, Pierre-Louis Lions, and Stanley Osher. Multiplicative denoising and deblurring : theory and algorithms. In *Geometric Level Set Methods in Imaging, Vision, and Graphics*, pages 103–119. Springer, 2003.
- [201] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D : Nonlinear Phenomena*, 60 :259–268, 1992.
- [202] Yousef Saad. *Iterative methods for sparse linear systems*. Siam, 2003.
- [203] Alexander A. Sawchuk. Space-variant image motion degradation and restoration. *Proceedings of the IEEE*, 60(7) :854–861, 1972.
- [204] Alexander A Sawchuk. Space-variant image restoration by coordinate transformations. *JOSA*, 64(2) :138–144, 1974.
- [205] Johannes Schindelin, Ignacio Arganda-Carreras, Erwin Frise, Verena Kaynig, Mark Longair, Tobias Pietzsch, Stephan Preibisch, Curtis Rueden, Stephan Saalfeld, Benjamin Schmid, et al. Fiji : an open-source platform for biological-image analysis. *Nature methods*, 9(7) :676–682, 2012.
- [206] Isaac Jacob Schoenberg. *Cardinal spline interpolation*, volume 12. SIAM, 1973.
- [207] J Schur. Bemerkungen zur theorie der beschränkten bilinearformen mit unendlich vielen veränderlichen. *Journal für die reine und Angewandte Mathematik*, 140 :1–28, 1911.
- [208] Laurent Schwartz. Théorie des noyaux. In *Proceedings of the International Congress of Mathematicians*, volume 1, pages 220–230, 1950.
- [209] Jean Serra. *Image analysis and mathematical morphology*. Academic press, 1982.

- [210] Jianing Shi and Stanley Osher. A nonlinear inverse scale space method for a convex multiplicative noise model. *SIAM Journal on Imaging Sciences*, 1(3) :294–321, 2008.
- [211] Eero P Simoncelli. Modeling the joint statistics of images in the wavelet domain. In *SPIE’s International Symposium on Optical Science, Engineering, and Instrumentation*, pages 188–195. International Society for Optics and Photonics, 1999.
- [212] Jonathan Simpkins and Robert L Stevenson. Parameterized modeling of spatially varying optical blur. *Journal of Electronic Imaging*, 23(1) :013005–013005, 2014.
- [213] Daniel A Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 81–90. ACM, 2004.
- [214] Jean-Luc Starck, Mai K Nguyen, and Fionn Murtagh. Wavelets and curvelets for image deconvolution : a combined approach. *Signal Processing*, 83(10) :2279–2283, 2003.
- [215] Jean-Luc Starck, E Pantin, and F Murtagh. Deconvolution in astronomy : A review. *Publications of the Astronomical Society of the Pacific*, 114(800) :1051–1069, 2002.
- [216] Gabriele Steidl and Tanja Teuber. Removing multiplicative noise by Douglas-Rachford splitting methods. *Journal of Mathematical Imaging and Vision*, 36(2) :168–184, 2010.
- [217] Elias M Stein and Timothy S Murphy. *Harmonic analysis : real-variable methods, orthogonality, and oscillatory integrals*, volume 3. Princeton University Press, 1993.
- [218] Thomas G. Stockham. High-speed convolution and correlation. In *Proceedings of the April 26-28, 1966, Spring joint computer conference*, pages 229–233. ACM, 1966.
- [219] Thomas Strohmer and Scott Beaver. Optimal OFDM design for time-frequency dispersive channels. *Communications, IEEE Transactions on*, 51(7) :1111–1122, 2003.
- [220] Gordon L Stüber. *Principles of mobile communication*. Springer Science & Business Media, 2011.
- [221] Frédéric Sur and Michel Grédiac. Automated removal of quasiperiodic noise using frequency domain statistics. *Journal of Electronic Imaging*, 24(1) :013003–013003, 2015.
- [222] WW Symes. The seismic reflection inverse problem. *Inverse problems*, 25(12) :123008, 2009.

- [223] Antonio Taberner, Javier Portilla, and Rafael Navarro. Duality of log-polar image representations in the space and spatial-frequency domains. *Signal Processing, IEEE Transactions on*, 47 :2469–2479, 1999.
- [224] Shaozhe Tao, Daniel Boley, and Shuzhong Zhang. Local linear convergence of ISTA and FISTA on the lasso problem. *arXiv preprint arXiv :1501.02888*, 2015.
- [225] Michael E Taylor. *Pseudodifferential operators*. Springer, 1996.
- [226] Maja Temerinac-Ott, Olaf Ronneberger, Peter Ochs, Wolfgang Driever, Thomas Brox, and Hans Burkhardt. Multiview deblurring for 3-d images from light-sheet-based fluorescence microscopy. *Image Processing, IEEE Transactions on*, 21(4) :1863–1873, 2012.
- [227] François Trèves. *Introduction to Pseudodifferential and Fourier Integral Operators Volume 2 : Fourier Integral Operators*, volume 2. Springer Science & Business Media, 1980.
- [228] Joel Trussell and Sergei Fogel. Identification and restoration of spatially variant motion blurs in sequential images. *Image Processing, IEEE Transactions on*, 1(1) :123–126, 1992.
- [229] Fuan Tsai and Walter W Chen. Striping noise detection and correction of remote sensing images. *Geoscience and Remote Sensing, IEEE Transactions on*, 46(12) :4122–4131, 2008.
- [230] Eugene Tyrtyshnikov. Incomplete cross approximation in the mosaic-skeleton method. *Computing*, 64(4) :367–380, 2000.
- [231] Michael Unser, Akram Aldroubi, and Murray Eden. B-spline signal processing. Part I - Theory. *Signal Processing, IEEE Transactions on*, 41(2) :821–833, 1993.
- [232] Florencio I Utreras. Cross-validation techniques for smoothing spline functions in one or two dimensions. In *Smoothing techniques for curve estimation*, pages 196–232. Springer, 1979.
- [233] Florencio I Utreras. Convergence rates for multivariate smoothing spline functions. *Journal of approximation theory*, 52(1) :1–27, 1988.
- [234] AH Voie, DH Burns, and FA Spelman. Orthogonal-plane fluorescence optical sectioning : Three-dimensional imaging of macroscopic biological specimens. *Journal of microscopy*, 170(3) :229–236, 1993.
- [235] Cédric Vonesch and Michael Unser. A fast thresholded Landweber algorithm for wavelet-regularized multidimensional deconvolution. *Image Processing, IEEE Transactions on*, 17(4) :539–549, 2008.

- [236] Cédric Vonesch and Michael Unser. A fast multilevel algorithm for wavelet-regularized image restoration. *Image Processing, IEEE Transactions on*, 18(3) :509–523, 2009.
- [237] Grace Wahba. *Convergence rates of "thin plate" smoothing splines when the data are noisy*. Springer, 1979.
- [238] Jingyue Wang and Bradley J Lucier. Error bounds for finite-difference methods for Rudin-Osher-Fatemi image smoothing. *SIAM Journal on Numerical Analysis*, 49 :845–868, 2011.
- [239] Ruxin Wang and Dacheng Tao. Recent progress in image deblurring. *arXiv preprint arXiv :1409.6838*, 2014.
- [240] Zhou Wang, Alan Conrad Bovik, Hamid Rahim Sheikh, and Eero P Simoncelli. Image quality assessment : from error visibility to structural similarity. *Image Processing, IEEE Transactions on*, 13(4) :600–612, 2004.
- [241] Jianing Wei, Charles A Bouman, and Jan P Allebach. Fast space-varying convolution using matrix source coding with applications to camera stray light reduction. *Image Processing, IEEE Transactions on*, 23(5) :1965–1979, 2014.
- [242] Jianing Wei, Guangzhi Cao, Charles A Bouman, and Jan P Allebach. Fast space-varying convolution and its application in stray light reduction. In *IS&T/SPIE Electronic Imaging*, pages 72460B–72460B. International Society for Optics and Photonics, 2009.
- [243] Pierre Weiss, Laure Blanc-Féraud, and Gilles Aubert. Efficient schemes for total variation minimization under constraints in image processing. *SIAM journal on Scientific Computing*, 31(3) :2047–2080, 2009.
- [244] Pierre Weiss, Paul Escande, and Yiqiu Dong. Contrast invariant SNR. 2016.
- [245] Pierre Weiss, Alexandre Fournier, Laure Blanc-Féraud, and Gilles Aubert. On the illumination invariance of the level lines under directed light : Application to change detection. *SIAM Journal on Imaging Sciences*, 4(1) :448–471, 2011.
- [246] Holger Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.
- [247] Holger Wendland and Christian Rieger. Approximate interpolation with applications to selecting smoothing parameters. *Numerische Mathematik*, 101(4) :729–748, 2005.
- [248] Hermann Weyl. Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen (mit einer anwendung auf die theorie der hohlraumstrahlung). *Mathematische Annalen*, 71(4) :441–479, 1912.

- [249] Norbert Wiener. *Extrapolation, interpolation, and smoothing of stationary time series*, volume 2. MIT press Cambridge, MA, 1949.
- [250] Özdoğan Yilmaz. *Seismic data analysis*, volume 1. Society of exploration geophysicists Tulsa, 2001.
- [251] Bo Zhang, Josiane Zerubia, and Jean-Christophe Olivo-Marin. Gaussian approximations of fluorescence microscope point-spread function models. *Applied Optics*, 46(10) :1819–1829, 2007.
- [252] Xi-Le Zhao, Fan Wang, and Michael K Ng. A new convex optimization model for multiplicative noise and blur removal. *SIAM Journal on Imaging Sciences*, 7(1) :456–475, 2014.
- [253] William Y Zou and Yiyan Wu. COFDM : An overview. *Broadcasting, IEEE Transactions on*, 41(1) :1–8, 1995.