Thèse présentée

pour obtenir le grade de

Docteur

de l'Université de Bordeaux

École Doctorale Sciences Physiques & de l'Ingénieur

Spécialité : Mécanique

Par **Florian Crouau**

Méthodes de super-résolution et sous-échantillonnage en imagerie thermique

Sous la direction de : **Jean-Luc Battaglia**
la co-direction de : **Stéphane Chevalier**
et le co-encadrement de : **Jérémie Maire**

Soutenue le 10 Septembre 2026

Membres du jury :

Mme. Audrey Giremus	Professeure	IMS, Université de Bordeaux	Présidente
M. Thomas Pierre	Professeur	IRDL, Univ. Bretagne Sud	Rapporteur
M. Vincent Schick	Professeur	Mines Nancy, Univ. Lorraine	Rapporteur
M. Julien Lecompagnon	Docteur	B.A.M.	Examinateur
M. Jean-Luc Battaglia	Professeur	I2M, Université de Bordeaux	Directeur
M. Stéphane Chevalier	Maître de Conférences	I2M, ENSAM	Co-directeur
M. Jérémie Maire	Chargé de Recherche	I2M, Université de Bordeaux	Invité

Méthodes de super-résolution et sous-échantillonnage en imagerie thermique

Résumé : L'imagerie infrarouge permet le contrôle et l'évaluation non destructifs des matériaux, de l'échelle micrométrique à l'échelle macroscopique. Elle fournit des informations sur la surface des matériaux, comme pour l'imagerie visible, mais permet également la visualisation de la distribution de la chaleur ainsi que la caractérisation des propriétés thermiques en surface et en profondeur. La gamme de matériaux étudiés est vaste : composites pour applications aérospatiales, tissus biologiques, matériaux de construction tels que le bois ou le béton, semi-conducteurs et microbatteries. Cependant, nombre de ces applications se heurtent à des limitations physiques liées à la nature de la lumière infrarouge et aux performances actuelles des détecteurs, empêchant d'atteindre une résolution inférieure au micromètre.

Cette thèse explore différentes approches pour obtenir une super-résolution et un sous-échantillonnage. Le chapitre 2 définit un modèle général d'acquisition d'images afin de mettre en évidence les deux principales limites de résolution : la limite de diffraction et la limite d'échantillonnage. Il présente également un aperçu des approches possibles, décrites dans la littérature, qui permettent de surmonter, de contourner ou de repousser ces limites. Ensuite, deux méthodes infrarouges point par point bien connues sont utilisées : la thermographie laser Flying Spot (FS) pour le chapitre 4 et la Radiométrie Photothermique par Balayage (SPR) pour le chapitre 5. L'objectif est de sous-échantillonner un objet, c'est-à-dire de ne retenir que quelques points d'une grille classique à échantillonnage régulier, puis de reconstruire une image équivalente à échantillonnage complet. Différentes méthodes de reconstruction sont comparées, avec un intérêt particulier pour le Compressive Sensing (Echantillonnage compressé). Ce paradigme mathématique postule que la plupart des images non bruitées sont très compressibles, en ce sens qu'elles peuvent être exprimées avec beaucoup moins de données grâce à une base adaptée. Dans les deux cas, une réduction d'au moins 80% du nombre d'échantillons a été obtenue. Puis, au chapitre 3, la super-résolution est explorée du point de vue de la diffraction et de l'échantillonnage. Ces études ont montré que les applications d'imagerie infrarouge en laboratoire souffrent souvent de la limite d'échantillonnage, ce qui a conduit à l'exploration d'une méthode d'échantillonnage compressif basée sur une généralisation des Single Pixel Camera (Caméra à 1 pixel), dans laquelle chaque pixel est subdivisé grâce à des masques binaires plus petits. Une preuve de concept a été réalisée à l'aide de caméras visibles, ce qui a permis de multiplier par 16 le nombre effectif de pixels, mais de nombreuses limitations doivent encore être surmontées afin d'obtenir des résultats cohérents et reproductibles dans l'infrarouge, ce qui constitue la principale perspective de développement de ce travail.

Mots-clés : Méthodes Inverses, Transferts thermiques, Super-résolution, Instrumentation, Optique, Sous-échantillonnage

---

I2M
UMR 5295 Université de Bordeaux, 33400 Talence, France.

Sauf mention contraire, ce document est sous licence
http://creativecommons.org/licenses/by/4.0/

Super-resolution and undersampling methods for thermal imaging

Abstract: Infrared imaging allows for non-destructive testing and evaluation of materials from the micro to the macro scale. It gives information about a materials surface just like visible imaging, but also allows for the visualisation of heat distribution as well as the characterisation of thermal properties on the surface but also in depth. Thus, the panel of materials studied range from composites for aerospace applications, to biological tissue, construction materials such as wood or concrete, or semiconductors and microbatteries. However, many of these applications have been faced with physical limitations related to both the nature of infrared light and the current performance of detectors which prevents the best achievable resolution to be lower than the micrometre scale.

In this thesis, multiple approaches for achieving super-resolution and under-sampling are explored. In chapter 2, a model of image acquisition is defined in order to highlight the two main limits to resolution that are the diffraction limit and the sampling limit, with a brief overview of the possible approaches in the literature that overcome, bypass or otherwise push those limits. Then, two well-known infrared point-by-point methods are used, namely Flying Spot (FS) laser thermography for chapter 4 and Scanning Photothermal Radiometry (SPR) for chapter 5 in order to first under-sample an object, that is taking only a few points from a classical regularly sampled grid, then reconstruct an equivalent fully sampled image. Various reconstruction methods are compared with a specific focus on compressed sensing, which is a mathematical paradigm that postulates that most non-noise images are very compressible in the sense that they can be expressed with way less data by expressing it in an adapted basis. In both cases, a reduction by at least 80% of the number of samples was achieved. Then, in chapter 3, super-resolution is explored in both the diffraction and sampling sense. These studies showed that infrared imaging applications in the lab often suffer from the sampling limit first, which has led to the exploration of a compressive sensing based method based on a generalisation of the single pixel camera, in which each pixel is sub-divided thanks to smaller binary masks. A proof of work was made using visible cameras, that allowed to multiply by 16 the effective number of pixels, but multiple limitations still need to be overcome in order to achieve consistent repeatable results in infrared, making this the main perspective development of this work.

Keywords: Inverse methods, Heat transfer, Super-resolution, Under-sampling, Instrumentation, Optics

I2M
UMR 5295 Université de Bordeaux, 33400 Talence, France.

Except where otherwise noted, this work is licensed under
http://creativecommons.org/licenses/by/4.0/

Résumé substantiel en français

Ces travaux de thèse s'inscrivent dans le cadre du contrôle non destructif par imagerie infrarouge. L'objectif est de compléter de l'information manquante dans les cartographies obtenues en les reconstruisant à l'aide de méthodes inverses. Deux concepts correspondent à cet objectif, d'un côté la super-résolution et de l'autre le sous échantillonnage. Dans le premier cas l'objectif est d'améliorer la résolution de l'image, c'est à dire aller au delà des limites physiques classiques. Dans le second, l'objectif est de compléter l'image à partir de quelques points.

Cette thèse s'articule en 3 grandes parties sur 4 chapitres. Dans la première (Chapitre 2) le concept de super-résolution est défini et les principales méthodes utilisées dans la littérature sont présentées en mettant l'accent sur celles qui sont adaptables dans le contexte du contrôle non destructif thermique. La seconde (Chapitre 3) regroupe des méthodes de super-résolution utilisées au cours de cette thèse, avec 3 approches présentées qui sont le dépliement spectral, l'illumination structurée et une méthode utilisant le principe de l'échantillonnage compressé. Enfin, la dernière (Chapitre 3 et 4) utilise aussi en grande partie ce dernier principe et l'idée générale de la compressibilité des données acquises afin de sous-échantillonner des mesures faites par balayage, que ce soit en Flying Spot ou en Radiométrie Photothermique afin de reconstruire des cartographies de propriétés thermiques avec le moins d'acquisitions possibles.

chapter 2 Définitions et état de l'art

Un système d'acquisition est représentable comme une boîte noire. La plupart des systèmes imageurs peuvent être subdivisés en une partie analogique et une partie numérique, comme illustré dans Figure 2.1 et Equation 2.2.

La partie analogique correspond à l'optique, c'est à dire les lentilles et miroirs qui, de part leur dimensions limitées causent un phénomène de diffraction ayant pour effet d'élargir (flouter) une source ponctuelle, prenant la forme d'une tâche d'Airy dans le cas d'optiques circulaires. Il s'agit donc d'un problème convolutif. La distance minimale à partir de laquelle deux points source sont distinguables est donc la limite d'Abbe Equation 2.18 $\frac{\lambda}{2NA}$. Elle s'exprime en termes de distance minimale peut aussi s'écrire en fréquence spatiale de coupure, définie par la transformée de Fourier de la réponse impulsionnelle Equation 2.16. Cette limite ne peut pas être dépassée uniquement par le calcul (déconvolutions), mais demande nécessairement une approche instrumentale différente en plus des méthodes inverses afin de la repousser.

Quelques méthodes existantes à cet effet dans la littérature sont présentées: La microscopie confocale, qui permet en théorie de doubler la fréquence de coupure en ajoutant une paire de petites ouvertures en entrée et sortie du système, l'image doit ensuite être construite par balayage. Tout un panel de méthodes d'imagerie à balayage existent aussi en s'éloignant un peu plus du microscope traditionnel, ce sont les méthodes à sonde locale ou à champ proche, ces dernières sont sans lentilles et suffisamment proche de l'objet pour que la limite de diffraction ne s'applique pas. Parmi ces

techniques se trouvent notamment la microscopie à force atomique et la microscopie thermique à balayage employées au laboratoire. La microscopie à fluorescence quant a elle est une méthode très répandue en biologie et permet une amélioration de la résolution en incluant des particules fluorescentes dans le milieu à imager qui agissent comme des sources ponctuelles qui peuvent être activées ou désactivées indépendamment les unes de autres. L'idée générale est qu'avec plusieurs acquisitions, il est possible de distinguer deux sources proches en activant l'une puis l'autre au lieu des deux à la fois. Ces méthodes sont difficilement envisageable dans le contexte du contrôle non destructif car elles nécessitent l'inclusion de particules dans le matériau. Enfin, les méthodes d'illumination structurée peuvent être vues comme conceptuellement similaire, car elles consistent en une série d'acquisitions où seule une partie du matériau est éclairée à la fois, mais cette sélection est faite par une structuration de la source lumineuse et donc permet de s'affranchir du problème de l'inclusion de particules. Cette méthode est donc étudiée expérimentalement dans le chapitre suivant.

La super-résolution peut aussi être interprétée du point de vue numérique, c'est à dire en considérant un problème d'échantillonnage. Le théorème associé tel que formulé par Shannon est le suivant : Tout signal à bande limitée ayant une fréquence maximum ν_{max} peut être entièrement déterminé à partir d'échantillons régulièrement espacés d'une distance $\frac{1}{2\nu_{max}}$. Ce théorème est une conséquence directe du phénomène de repliement spectral causé par l'échantillonnage et facilement exprimable dans l'espace de Fourier, voir Figure 2.11 et 2.12. Ce théorème, usuellement énoncé pour les signaux 1D reste vrai en images. Cette discrétisation des images nécessite de représenter le problème mathématiquement par des produits matriciels, avec $\mathbf{H}$ la fonction de transfert entre la mesure et l'objet décrite dans les équations 2.37 à 2.41.

Cette limite d'échantillonnage peut être repoussée lorsque ses hypothèses associées ne sont plus vérifiées. Il est par exemple envisageable de ne plus travailler avec des échantillonnages réguliers, auquel cas la limite de résolution ne sera pas la même sur toute l'image, et le problème direct devient mal posé, avec plus de points à reconstruire que de mesures. Des méthodes d'interpolations irrégulières existent dans ces cas là, notamment le Krigeage et l'interpolation à fonctions de base radiale. Une autre méthode est celle de l'échantillonnage compressé, où au lieu de considérer des signaux quelconques à bande limitée seuls des signaux fortement compressibles sont considérés, c'est à dire des signaux pouvant s'exprimer dans une base donnée (Fourier ou autre) en uniquement quelques modes. Cette propriété plus restrictive permet de partiellement s'affranchir de la limite de Shannon. Dans ce paradigme, le problème est posé de la forme $y = \mathbf{H}\mathbf{\Psi}s$ où $\mathbf{\Psi}$ est un changement de base, s un signal dit éparse, c'est à dire composé essentiellement de 0, et $\mathbf{\Psi}s = x$. L'objectif est alors de minimiser le nombre de modes dans s, requérant des minimisations en norme ℓ_0 ou ℓ_1 faites dans la suite de cette thèse à partir des algorithmes 1 (FISTA), 2 (OWL-QN) et 3 (TVAL3).

chapter 3 Super-résolution en Thermographie IR

Trois méthodes de super-résolution sont abordées dans cette section en commençant par le dépliement spectral. Cette méthode a permis une réduction de la taille effective des pixels de $15\,\mu\mathrm{m}$ à $10.4\,\mu\mathrm{m}$ selon un axe pour des mesures de thermotransmittance autour d'un fil d'or chauffé. Le principe de la méthode consiste à imager le fil à un certain angle par rapport à la grille de pixel. La transformée de Fourier de cette image contient de l'information le long d'une ligne perpendiculaire au fil d'or. Le repliement spectral

causé par le non respect du théorème de Shannon s'exprime sous la forme de lignes non pas superposées au signal principal, mais parallèle, il est donc possible d'extraire les signaux repliés pour les replacer à leur position originelle, voir Figure 3.3. Cette méthode peut être utilisée sur des images avec des formes géométriques simples (circuits imprimés), mais dès lors que le repliement spectral se fait sur une région de l'espace de Fourier contenant déjà de l'information il n'est plus possible d'inverser l'effet, du moins avec une seule mesure.

L'illumination structurée permet une amélioration de la résolution au sens d'Abbe grâce à l'envoi d'une série d'éclairages non-homogène, qui peuvent être structurés (sinus), aléatoires, connus ou non. Ici, la variante testée expérimentalement tel qu'illustré en Figure 3.5 est aléatoire et connue. A partir d'une série de mesures, une image haute résolution est reconstruite avec l'algorithme 4 en minimisant la fonction coût par descente de gradient pour chaque mesure. Cet algorithme est d'abord testé sur un cas numérique Figure 3.6, où un gain de résolution est observé sur les hautes fréquences d'une étoile de Siemens. De plus, Figure 3.7 montre que, en théorie, si le motif est connu avec une meilleure précision que la résolution de la caméra, il est possible de reconstruire des plus hautes fréquences que dans le cas précédent. Cependant, l'application expérimentale n'a pas été concluante en Figure 3.8 pour deux raisons, la première étant l'incertitude sur la connaissance et la mesure de la fonction de transfert, et la seconde plus fondamentale et cause de la première, étant qu'en imagerie infrarouge la limite de résolution principale est la limite d'échantillonnage.

C'est pourquoi ces travaux se sont ensuite tournés vers l'échantillonnage compressé avec une caméra et un filtre spatial intermédiaire. Le principe expérimental est similaire au cas précédent, sauf que la sortie du signal est structurée (juste avant la caméra) à la place de la source Figure 3.9. Si le problème reste globalement de la même forme $y = \mathbf{H}x$, il ne peut plus être décrit comme linéaire invariant car la réponse impulsionnelle peut varier selon sa position dans une grille de sous-pixels par rapport à la grille de pixels caméra, en particulier lorsque ces deux grilles ne sont pas parfaitement alignées, ce qui est généralement le cas expérimentalement. Il faut donc plusieurs mesures pour reconstruire une image, mais aussi caractériser le système en mesurant chaque réponse impulsionnelle en chaque position. Ce dernier point étant expérimentalement illusoire pour des images avec un nombre conséquent de pixels, l'échantillonnage compressé est utilisé pour minimiser le nombre de mesures à la fois pour l'image et la caractérisation du système. La méthode est testée numériquement et expérimentalement sous plusieurs variantes. Pour le cas numérique, le filtre spatial intermédiaire est composé de masques dit de Hadamard, illustrés en Figure 3.11, qui a permis une une amélioration de la résolution en Figure 3.13e, même lorsque le nombre de mesures utilisée est inférieur au gain en résolution. Pour le cas expérimental, des motifs aléatoires ont été utilisés afin d'être moins semsible aux problèmes d'alignement. Deux méthodes de création de ces motifs sont utilisés, un système de micro miroirs et des masques imprimés par photolithographie, voir Figure 3.14. Des preuves de concept en imagerie visible ont permis des amélioration des résolutions en Figure 3.16 et 3.17, même si l'application en infrarouge n'a pas encore été concluante notamment à cause du niveau de bruit plus élevé affectant à la fois la reconstruction d'image et des réponses impulsionnelles.

chapter 4 Flying Spot sous-échantillonné

La thermographie laser à balayage dit « Flying Spot » regroupe plusieurs variantes permettant la caractérisation de matériaux, où une caméra enregistre la réponse d'un matériau à un laser le chauffant ponctuellement. Ce point peut ensuite être déplacé relativement au matériau afin de le caractériser, un exemple de montage est schématisé en Figure 4.1. Ici, ce montage est utilisé pour créer une série d'impulsions à différentes positions du matériau, dont l'évolution de la réponse au cours du temps permet de déduire l'éffusivité locale. Ces positions sont choisies aléatoirement parmi une grille de points qui correspondrait à l'échantillonnage classique au sens de Shannon. Le principe est décrit en Figure 4.3. La cartographie des éffusivités est reconstruite à partir des mesures sous-échantillonnées grâce à des algorithmes d'échantillonnage compressés (TVAL, OWLQN et FISTA) qui sont comparés expérimentalement selon leur capacité à reconstruire fidèlement l'image avec le minimum d'acquisitions possibles. Figure 4.6 montrent qu'il est possible de n'utiliser que 20% des points de la grille régulière pour obtenir des images où la structure du matériau est préservée. Cependant les trois algorithmes ne sont pas aussi performants, en particulier l'algorithme TVAL3 qui a produit une image plus dégradée pour les taux d'échantillonnage les plus faibles, ce qui s'explique à la fois par le fait que la minimisation éparse est faite sur le gradient, alors que les deux autres algorithmes utilisent la transformée en cosinus, mais aussi par le fait que la méthode demande l'optimisation de 2 paramètres de régularisation, rendant TVAL plus difficile d'utilisation.

chapter 5 Photothermie radiométrique à balayage sous-échantillonnée

La photothermie radiométrique à balayage (SPR) est une technique de caractérisation basée sur la photothermie radiométrique modulée (MPtR). Elle consiste en la mesure de l'émittance d'un matériau en réponse à un laser le chauffant périodiquement, la fréquence choisie modifiant la surface et la profondeur chauffée et donc sondée, voir Figure 5.1. Un lock-in permet ensuite d'extraire le signal correspondant à la faible variation de signal induite. La SPR consiste à répéter cette mesure d'amplitude et phase afin de cartographier un échantillon, le montage expérimental est schématisé en Figure 5.2. L'acquisition de ces images peut être coûteuse en temps, en particulier si plusieurs fréquences de modulation sont étudiées, d'où l'intérêt de les sous-échantillonner. L'algorithme OWL-QN est comparé à l'interpolation à base radiale (RBF) afin de comparer à la fois la robustesse au bruit (Figure 5.8) et aux taux de sous-échantillonnage (Figure 5.6) des deux méthodes. Le bruit empêche la reconstruction d'images assez rapidement quelque soit la méthode, cependant la RBF permet un sous-échantillonnage plus agressif, avec sur le matériau étudié 5% des mesures suffisantes pour retrouver une image proche de la référence, contre 15% pour OWL-QN. Les deux méthodes réduisent ainsi des temps d'acquisition de l'ordre de l'heure vers la dizaine de minutes. Enfin, la méthode SPR permet la cartographie 3D en modulant la fréquence laser, donc un processus d'acquisition est proposé en Figure 5.4 où une première image est acquise avec un sous-échantillonnage aléatoire, reconstruite, puis les autres fréquences supérieures sont choisies préférentiellement là où les variations sont les plus fortes, ce qui a permis expérimentalement d'obtenir des images plus contrastées en Figure 5.9.

Acknowledgements

I would like to thank all people who accompanied me for those last three years, and directly or indirectly helped with this thesis.

Starting with all the members of the jury for accepting to contribute in the revewing process. I am thanking both Thomas Pierre and Vincent Schick for reviewing this thesis, as well as Audrey Giremus for presiding this jury. I am thanking as well Julien Lecompagnon for accepting to be an examiner, as his and his colleagues work has been a source of inspiration for this thesis.

I would like to thank as well my supervisors. Starting by Jérémie Maire, who has been very present and involved in the whole process. Sadly for administrative reasons this doesn't show up on the front page as both the terms *co-encadrement* and *Invité* don't reflect his involvement. Thanks to Stéphane Chevalier as well for his regular support as well as advices on both science and communication. I am also thanking Jean-Luc Battaglia, who, despite the fact that I am sure he would say that it wasn't much, had extremely useful insights and discussions which were always encouraging.

I would like to thank as well all the rest of the ICT research team, namely: Jean-Christophe Batsale, Marie-Marthe Groz, Andrzej Kusiak, Paule Lapeyre. A special thanks to Alain Sommier as well who has saved the day a couple of times in my experiments, research wouldn't be anywhere near where it is at without good research engineers such as him (And the sometimes dubious jokes are a bonus).

I thank all the others PhD students and postdocs in ICT, as we are the other half of the team: Kilian, Théo, Célia, Mariana, Jordan, Ida, Alejandro, Thomas, Hugo, Malo, Marine, Kevin, Coline, Caroline and obviously Alisa as well as the many interns whom I've met during those three years. This thanks extends as well the many PhD students of the lab, and especially of the TREFLE department with whom we had regular meetings, which were welcome breaks from the usual routine.

Thanks as well to the people I have met as a representative for the Doctoral School SPI, the other PhD students representatives as well as the directors Sakir Amiroudine and Nathalie Malbert and administrators Mariann Touzac and Brigitte Bordes. Thanks to Jean-François Giovannelli and Timothée Lévi for having participated in my yearly *Comités de Suivi*. Thanks to all the other research support staff, from administrative to IT, without whom no lab would be operational.

This might be a weird one, but I wouldn't be here without extensive support from the State which paid for all my education from the earliest age to today, including scholarships. So, thanks to the French Republic which, despite the successive governments, brought me there and I hope that the ladder I managed to climb will be improved, not broken after me. On a similar note, I would like to thank Sylvie who, 11 years ago brought me to Toulouse in order to register to the University, it might sound small but it probably changed my life, and the first mention of me doing a PhD was on that day.

A final thanks to my family and friends for their emotional support, and for bearing with me both online and offline.

Achievements and Contributions

The main content of this thesis contains what I have done during those three years as well as its required context to understand it. In order to make sure to clarify which is which, in this short section I will list summarily my work in order to distinguish it from what was already developed by previous researchers.

chapter 2

This is the bibliography chapter. The goal here was to merge the important information on the subject of super-resolution that I have gathered over the years. Obviously nothing is new here from the point of view of global knowledge, but as this subject is new within the lab, I think it was necessary to clarify as much as possible the problem for the future students (and researchers) in this field within the lab especially, and to give a brief overview of the many existing approaches. If "a month in the laboratory can often save an hour in the library", as Frank Westheimer is often quoted to have said, then many theses worth of time in the lab were spared with this chapter.

chapter 3

In section 3.1, the setup was already developed during the thesis of Coline Bourges. I did the measurements for the transfer function using an USAF 1951 target on which I measured contrasts. I then did the measurement on the heated gold wire, and processed it using the demodulation technique described in Coline's work. Then I did the processing and filtering in the spatial Fourier transform domain. I did as well some deconvolutions on this image. All the video files generated by the FLIR softwares are read on MatLab using it's associated SDK and programs developed within the ICT.

In section 3.2, I programmed the Algorithm 4, that is based on the formulas described by Yeh et al. I did the entirety of the numerical tests, generating the object, the random speckles, the transfer functions and the measurements in order to use the algorithm. For the experimental part, I built the setup using a gold diffuser that was already present on the lab, and I made the successive measurements by slightly rotating and translating it every time. I also tested early on other approaches, notably from Murray et al.'s article, but my implementation of the code was most likely off which lead to inconclusive results at the time.

In section 3.3, I made the two setups from Figure 3.14. Jérémie and I procured the DMD from a Texas Instrument distributor, I built a support with 3D printing to hold it firmly in position and at $45°$ angle, and did the optical alignment of the rest of the setup, both with an IR detector and a visible camera. For the photolithography masks, I designed the patterns at 20 and 40 μm and made them into gds files, that were then sent to be printed by a third party company. I also made the program controlling the motors to move the masks by a given distance every few seconds. On the computation side, I built all the necessary data for the direct problem, namely the Hadamard and random

patterns, the ground truth, the transfer function and the measurements. I made two programs based on the TVAL algorithm described by Chengbo Li (Its implementation was found in Liu Yang's GitHub page, see the corresponding Chapter) in order to calculate the transfer matrix and the image following the required structure and reshape and I chose the optimum regularisation parameters. The computation of the experimental measurement was done the same way, with the random patterns either generated for the DMD or taken from the reference image of the mask. The alignment issue made this step difficult if no high-resolution reference is known. I made the measurements with visible light using only 1 channel of a color camera, and processed it accordingly.

chapter 4

The Flying Spot setup was already in place in the lab, thanks to the work of Ludovic Gavérina as well as Alain Sommier and Christophe Pradère. I had to learn how the LabView program worked and made some modifications to in particular in order to read an arbitrary set of positions and laser intensities. These sets were created thanks to a Python script that I built in order to generate the random subset of samples, which then allowed me to make the measurement presented in the chapter. Both script and LabView modifications then went on to be used for other research in the lab for which I also lent a hand. Some troubleshooting was also required in order to work with the Alphanov laser used in the experiment instead of the more powerful but wider one usually in place, as well as a quick characterisation of the beam (size, intensity).

On the algorithmic reconstruction side, the main functions for FISTA, OWL-QN and TVAL3 were already developed by their corresponding authors, and I did all the preprocessing in order to use them, including reshapes, construction of the transfer matrix, and choice of regularisation parameters. Installation of the package for the FISTA algorithm was straightforward, but OWL-QN required some tweaks due to the fact that it calls a C program (For the Hessian calculation using LBFGS) and doesn't behave well on all machines (I made it work with a machine using Debian 13). Also, I did some processing using Matlab to find the time and position of each pulse as well as to fit curves to $1/\sqrt{t}$. I also did reimplement the FISTA algorithm from scratch, but this was set aside due to the fact that the Pylops package version was better performing.

chapter 5

The Scanning Photothermal Radiometry setup was developed in the context of Alejandro Mateos-Canseco's thesis by him, Andrzej Kusiak and Jean-Luc Battaglia. Alejandro and I added a feature to his LabView program in order to read files in the same overall format as for the Flying Spot case. Files were then created using a similar Python script as above for fully random sampling. I then made a script that takes a first reconstructed image, generate a weighted random subset of pixels that are then turned into a LabView readable file. I was also taught by Alejandro how the software and the setup works which allowed me to make measurements at different scales and on different samples, before choosing to work with and present the one included in the chapter. The rest of the programming side is similar to the Flying Spot chapter, with the addition of the RBF calculations requiring the values measured and their coordinates.

Articles:

- [1] (published, first author) see chapter 4,
- [2] (pre-print, first author) see chapter 5.

International conference:

- QIRT 2026: (speaker) for the work in the last article above [2].

National conferences:

- SFT 2024: (poster and article) on spectral unfolding, see section 3.1,
- SFT 2025: (poster) on Compressive sensing Super Resolution see section 3.3.

... And on a more local level:

- Doctoral School Day 2025: (poster), same subject as for SFT 2025,
- Presentation for the TREFLE research department Day 2026, on preliminary work for what became [2], Still available here.

Other responsabilities:

- Teaching: Travaux Pratiques de Thermique (Tutorials in Thermal Engineering) for 1st year engineering[1] students at ENSAM - 50ish h (2024-2026)
- PhD student representative for the Doctoral School for Physics and Engineering (SPI), elected twice (2024-2026), including once additionally as representative for the University of Bordeaux's Graduate Research School (Collège des Écoles Doctorales). This included assisting to meetings, in which among other things we proposed changes for the *Comité de Suivi Individuel*. As well as organising two Doctoral Schools Days (2025 and 2026) conferences.

[1] Confusingly, that's equivalent to 3rd year of Bachelor with a mix of pride and stubbornness leading to them not yet accepting the Bologna Process

Contents

Abstract iii

Acknowledgements ix

Achievements and Contributions x

List of Figures xv

Acronyms xvii

Symbols xviii

1 Introduction 1
- 1.1 General context of the lab . . . 2
- 1.2 Super-resolution and under-sampling . . . 2
- 1.3 Infrared imaging . . . 3
- 1.4 Non-destructive testing . . . 4
- 1.5 Objectives and Summary . . . 4

2 Super resolution: Definitions and techniques 6
- 2.1 General instrumental model and direct problem . . . 7
- 2.2 Super resolution in the analog sense . . . 8
  - 2.2.1 Resolution limit and definition of the convolutive problem . . . 8
  - 2.2.2 Deconvolution . . . 15
  - 2.2.3 Scanning methods . . . 16
    - Confocal Microscopy . . . 16
    - Scanning Probe Microscopy . . . 17
  - 2.2.4 Super resolution in Fluorescence Microscopy . . . 18
  - 2.2.5 Fourier Ptychography . . . 20
  - 2.2.6 Structured Illumination Microscopy . . . 21
- 2.3 Super resolution in the digital sense . . . 24
  - 2.3.1 Definition of the under sampled problem . . . 24
  - 2.3.2 Interpolations . . . 28
    - Radial Basis Functions . . . 30
    - Kriging . . . 31
  - 2.3.3 Compressive sensing . . . 32
    - Hypotheses . . . 32
    - How to promote sparsity . . . 33
    - Applications . . . 35
    - Algorithms used in this thesis . . . 36

2.3.4 Specificities of thermal imaging in super resolution 40
2.4 Conclusion 42

3 Sampling and Diffraction limits in IR Thermography: First attempts 43
3.1 Spectral unfolding 44
3.1.1 Setup characterisation 45
3.1.2 Spectral unfolding 45
3.2 Structured illumination (Result/Limitations) 48
3.2.1 Method 48
3.2.2 Numerical results 49
3.2.3 Experimental limitations 51
3.2.4 Perspectives 52
3.3 Compressive sensing with cameras and an intermediary spatial filter . . 54
3.3.1 Principles 54
3.3.2 Measurement of the transfer matrix C 56
3.3.3 Hadamard sampling 57
Numerical results 59
3.3.4 Random sampling 60
3.3.5 Experimental results using DMD 61
3.3.6 Experimental results using photolithography masks 64
3.3.7 Conclusion and Prespectives 66

4 Under sampled Flying Spot 68
4.1 Introduction 69
4.2 Flying spot: operating principle 69
4.3 Undersampled Flying Spot Principle 71
4.4 Algorithms used 74
4.5 Experimental results 75
4.5.1 Thermal measurements and regular sampling 75
4.5.2 Comparison of sampling ratios 78
4.6 Conclusion and Perspectives 81

5 Under sampled Photothermal Radiometry 82
5.1 Method and operating principle 83
5.2 Algorithms used 88
5.3 Comparison on sampling ratios 89
5.4 Comparison on noise levels 91
5.5 Multiple frequencies scan 91
5.6 Conclusion and Perspectives 93

6 Conclusion and perspectives 95
6.1 Conclusion 96
6.2 Perspectives 96

List of Figures

1.1 Introduction: Time gain . . . 5

2.1 $x = Hy$ . . . 7
2.2 Simple imager setup . . . 9
2.3 Typical PSF and OTF . . . 13
2.4 Rayleigh, Abbe and Sparrow limits . . . 14
2.5 Confocal microscope: OTF . . . 17
2.6 Confocal microscope: Setup . . . 18
2.7 SPM: Examples of Scanning probes . . . 19
2.8 Fourier Ptychography: Setup . . . 20
2.9 SIM: Principle . . . 21
2.10 RIM: Principle . . . 22
2.11 Aliasing Principle . . . 25
2.12 FT, DTFT, DFT . . . 26
2.13 RBF Principle . . . 31
2.14 CS Principle . . . 33
2.15 CS: ℓ_p distances . . . 34
2.16 CS: Single Pixel Camera Setup . . . 35

3.1 Thermotransmittance: Setup . . . 44
3.2 Thermotransmittance: Acquisition . . . 46
3.3 Thermotransmittance: frequency filters . . . 47
3.4 Thermotransmittance: Result . . . 47
3.5 SIM: Setup . . . 48
3.6 SIM: Numerical result . . . 50
3.7 SIM: Numerical result 2 . . . 51
3.8 SIM: Experimental attempt . . . 52
3.9 CS Camera: Setup . . . 55
3.10 CS Camera: Transfer Matrix construction . . . 56
3.11 CS Camera: Hadamard masks . . . 58
3.12 CS Camera: Example of high-res to low-res transfer matrix . . . 60
3.13 CS Camera: Numerical results (Hadamard patterns) . . . 61
3.14 CS Camera: Setups . . . 62
3.15 CS Camera: Datasheet DMD . . . 63
3.16 CS Camera: Experimental results (Random patterns on DMD) . . . 64
3.17 CS Camera: Experimental results (Random patterns on masks) . . . 65
3.18 CS Camera: Experimental results (Transfer matrix - Masks) . . . 66

4.1 Flying Spot: Setup . . . 69
4.2 Flying Spot: Example of measurement with parabolas . . . 72

4.3 Flying Spot Undersampling scheme 73
4.4 Flying Spot: Example of acquisition and pre-processing 76
4.5 Flying Spot: Full effusivity map 78
4.6 Flying Spot: Comparison of reconstruction algorithms 79
4.7 Flying Spot comparison: Profiles and RMSE 80

5.1 SPR: Frequency and depth 84
5.2 SPR: Setup 85
5.3 SPR: Regular measurement 86
5.4 SPR: Weighted Random Principle 87
5.5 SPR: Sample 88
5.6 SPR: Comparison (Sampling Ratios) 90
5.7 SPR: Equivalent regular sampling 91
5.8 SPR: Comparison (Noise levels) 92
5.9 SPR: Weighted random 93
5.10 SPR: Sample 94

6.1 Laminography: CS Test 97

Acronyms

AFM Atomic Force Microscopy 17, 35
AOM Acousto-Optic Modulator 85
BCCB Block Circulant with Circulant Blocks 28, 57
BLUE Best Linear Unbiased Estimator
CCD Charge Coupled Device 46
CS Compressive Sensing 32, 33, 41
CTF Contrast Transfer Function 45
DAQ Data Acquisition Card
DCT Discrete Cosine Transform 74
DFT Discrete Fourier Transform 25
DL Digital Level
DMD Digital Micromirrors Device 35, 55, 56, 60
DTFT Discrete-Time Fourier Transform 25
ESF Edge Spread Function
FDTR Fourier Domain Thermoreflectance 41, 84
FFT Fast Fourier Transform algorithm 25
FISTA Fast Iterative Schrikage Threshold Algorithm 36, 74
FS Flying Spot 69,
FT Fourier Transform 9, 21
FWHM Full-Width at Half Maximum 75
GPU Graphical Processing Unit
HDMI High-Definition Multimedia Interface
I2M Institut d'Ingénierie et de Mécanique 2
ICT Imagerie et Caractérisation Thermique 2, 44
iFFT Inverse FFT 26
IR InfraRed
ISO International Standardisation Organisation 47
ISTA Iterative Schrikage Threshold Algorithm 36
L-BFGS Limited-memory Broyden Fletcher Goldfarb Shanno Algorithm 37
LTI Linear Time-Invariant 7, 9
MPtR Modulated Photothermal Radiometry 83
MRI Magnetic Resonance Imaging 12, 35
MTF Modulation Transfer Function 11, 21, 45
NA Numerical Aperture 12, 25, 48
NDT-E Non-Destructive Testing and Evaluation 2, 4
NRMSE Normalised Root-Mean Squared Error 78
OTF Optical Transfer Function 11, 13, 20, 21
OWL-QN Orthant-Wise Limited-memory Quasi-Newton Algorithm 37, 74
PALM Photo Activated Localisation Microscopy 19, 20, 42
PSF Point Spread Function 11, 19, 21
QIRT Quantitative IR Thermogaphy
RBF Radial Basis Functions 30
RIM Random Illumination Microscopy 22
RLD Richardson Lucy Deconvolution
RMSE Root-Mean Squared Error
ROI Region Of Interest 66
SFT Société Française de Thermique 43
SIM Structured Illumination Microscopy 21, 22, 36, 47, 48
SLM Spatial Light Modulator 48
SNOM Scanning Near Field Optical Microscopy 17
SNR Signal to Noise Ratio
SPC Single Pixel Camera 35, 40, 54
SPM Scanning Probe Microscopy 17
SPR Scanning Photothermal Radiometry 83
SSIM Structural Similarity Index Measure
STEDM Stimulated Emission Depletion Microscopy 19, 35, 42
SThM Scanning Thermal Microscopy 17
STM Scanning Tunneling Microscopy 17
STORM Stochastic Optical Recons Microscopy 19, 20, 23, 42
SVD Singular Value Decomposition 32
TV Total Variation
TVAL3 Total Variation ALternating Lagrangian Algorithm 38, 55, 74
USAF US Air Force 1951 Calibration Target 45

Symbols

Symbol	Description	Unit	Page
$\mathbf{a}$	Diffusivity tensor, components a_{ij}	m^2/s	70
$\mathbf{\Lambda}$	Conductivity tensor, components Λ_{ij}		69
$\mathbf{\Psi}$	Transform basis		74
(ν_x, ν_y)	spatial frequencies	m^{-1}	
E	Thermal effusivity		70
I	Irradiance of light signal	W/m^2	11
M	number of samples of y, $M \leq N$		29
M_e	Emittance		83
M_z	Lens Magnification		10
N	number of samples of x		29
P	Pupil function (opening)		10
Q	Energy deposited (laser)	J	70
U	Amplitude of light signal		10
c_p	Specific heat capacity		69
d	sampling step, distance	m	25
f	Frequency (temporal)	Hz	44
f/n°	f-number		12, 54
h	impulse response		8
x	Input		8
y	Output		8
z	distance along optical axis		
δ	Dirac Delta		
δ_p	Penetration depth	m	83
ϵ	noise		8
μ	Regularisation parameter		
ρ	density		69
ε	emissivity coefficient		83
ϱ	source illumination		8
φ	phase		
$\otimes$	convolution product		8
Δd	Distance between points		8
$\mathbf{C}$	filter-to-camera transfer matrix		55
$\mathbf{F}$	DFT Matrix		27
G	Green function		41
$\check{h}$	impulse resp. in amplitude		10, 21
$\tilde{h}$	FT of h		11
$\mathbf{H}$	Transfer Matrix, size $N \times M$		15, 29, 54
H	Transfer function		7
$\mathbf{H_{2^k}}$	Hadamard matrix, size 2^k		58
$\mathbf{H_f}$	Hessian matrix of $f(x)$		37
ℓ_0	Pseudo-norm, counts non-zeros in a vector		33, 34

Symbol	Description	Unit	Page
ℓ_1	Taxicab norm		34
ℓ_2	Euclidian norm		15, 34
λ	wavelength	m	
$\mathcal{F}$	Fourier Transform operator		9
$\hat{}$	Reconstruction estimate		
$\tilde{}$	Fourier transform		
$\vec{}$	2D coordinates		
(u, v)	positions on the plane (image)	m	
r_d	radius	m	12
r_i	i^{th} random pattern vector		54
s	sparse vector		33
III	III_d Dirac broom of step d		8, 24
w_i	weight (RBF Interpolation)		

Chapter 1

Introduction

Contents

1.1 General context of the lab 2
1.2 Super-resolution and under-sampling 2
1.3 Infrared imaging 3
1.4 Non-destructive testing 4
1.5 Objectives and Summary 4

1.1 General context of the lab

The work described by this thesis was done in the Engineering and Mechanics Institute *Institut d'Ingénierie et de Mécanique (I2M)* laboratory within the Thermal Imaging Fields and Characterization group *Imagerie et Caractérisation Thermique (ICT)*.The core activities are Non-Destructive Testing and Evaluation (NDT-E) of materials and characterisation of thermal properties from tens of nanometers to the meter scale usually using infrared thermography and inverse methods. In this context, the need for precise detection of small defects while minimising the costs in time and data often arise, and this applies to many imaging systems.

Image acquisition (and display) can be made in two main ways. The first one consists in the simultaneous acquisition of the image plane, this corresponds to both film and digital photography[1], where every position of the picture was exposed to light at the same moment. This requires a series of detectors distributed all along the image plane. The other way is scanning, that is to use a single detector and to take a successive series of measurement at all points within the plane fast enough relative to any potential movement in the scene. Despite being less common in everyday life nowadays, it is still widely used, for instance a lot of satellite imaging applications behave that way where a single detector gathers quickly a series of data point along a line thanks to a moving mirror, then the whole satellite moves relatively slowly a bit further along where the next line is measured, etc. Hybrid cases of line scan also exists, where one axis is taken at once, and the other in succession.[2] Both ways have their pros and cons, with the simultaneousness being desirable but requires compromises on noise levels or cost principally that might not always be worth it.

In imaging, the precision in terms of distances, ie. the resolution, is limited both by the nature of the light and the acquisition system. Indeed, there are two main limitations, an optical one called the diffraction limit that states that the minimum achievable size is roughly half the wavelength, as well as a limit depending on the size of the detector, or more accurately the distance between two samples, called the sampling limit.

Considering that, in infrared imaging: 1. The camera setups are such that the effective distance between two pixels is in the tens of micrometres and 2. The order of magnitude of the wavelength can not be meaningfully changed since the objective is to measure thermal-induced infrared signals,

A question then arises: Is it possible to be more precise than those limits? Answering this is called super-resolution.

1.2 Super-resolution and under-sampling

The term of super-resolution is ambiguous and can be used in different contexts that have the same general objective, that is to "see better", but are fundamentally different in how they are approached due to the physical limitations encountered being different:

[1] Lets omit rolling shutters here

[2] The two methods can also be illustrated on the display side by comparing CRT TV images that are made point by point to LCD screens where each pixel is displayed simultaneously, but the comparison is probably already no longer relevant for the younger generations

A diffraction limited problem cannot be treated the same way as a sampling problem, and yet all is encompassed in the literature by this same word. The super-resolution as going beyond the sampling limit is often related to undersampling. To undersample is to take a small subset of data point that suffice to get a good idea of the whole signal or image. This is, if not exactly identical, at least very close to wanting to go beyond the sampling limit which consists in estimating the better resolved image from a (set of) less precise one, and thus using a subset of data. These can be viewed as two sides of the same coin, with the first side of super-resolution putting the emphasis more on prediction of the remaining points, while in under-sampling the focus is made more on compression of the measured data.

The question of super-resolution and under-sampling arises in thermal imaging for multiple reasons beside mere curiosity. First the best achievable resolution scale, of slightly less than the micrometer can be insufficient for some applications. For instance in micro to nanophysics, most notably in semi-conductors but also in lab-on-chip (microfluidics) applications, thermal imaging could be desirable in order to localise hot spots yet conventional method struggle at this due to the aforementioned limitations. Second, the methods that are the most relevant at the smallest scale are more often than not point by point methods, in which the limiting factor is time. As each point is taken successively, maintaining a reasonably wide field of view with those very small movements requires a long acquisition time, thus limiting the amount of samples that can be tested. Third and related to the previous point is about the amount of data: As the number of pixel increases, and as those pixels turn into voxels with the apparition of 3D imaging as well as videos, the sizes of acquisition files can very quickly enter in the gigabyte range and beyond, which can become an issue both for storing and for processing this data, therefore finding the lowest amount of samples that contain the useful information is is an approach to alleviate the data issue.

1.3 Infrared imaging

Infrared thermography is a contactless method for measurement, and ultimately imaging, of temperatures that uses infrared radiation. It is based on a couple of fundamental physical laws, notably the Stefan-Boltzmann law which states that any matter that possesses a temperature above zero Kelvin emits photons, and the total radiated power is proportional to the temperature to the fourth power.

This radiation is distributed along all directions, but also in a wide spectrum of wavelengths that are described by Plank's law. This law shows that for close to ambient temperature, the wavelengths involved are in the range of $3\,\mu\mathrm{m}$ and above, which is well into the mid to long wave infrared, with only very high temperatures ($\sim 500\,\mathrm{K}$) causing a shift toward visible light.

There are two important caveats that need to be underlined when attempting to measure a temperature from thermal radiation: First, the Stefan-Boltzmann law is also dependant on the emissivity of the radiating body, this value between 0 and 1 indicates how close the material is from a black body, i.e. an idealised object that absorbs all radiation and emits the maximum amount of energy for its temperature. In practice this value is not known, but some black paints can help with bringing it close to 1. Second, the radiation is emitted in all directions in a similar manner to diffuse reflection, thus

the angle at which the material is observed matters, and being as close as possible to perpendicular to the surface is often necessary to give quantitative values.

1.4 Non-destructive testing

NDT-E encompasses a wide panel of methods with the same objective to detect defects within a material without destroying it. These tests are one of the pillars of both industrial production and medicine. In order to do so, a wide array of signals emitted or re-emitted by the sample can be measured, going from simple visible light in visual testing, to the rest of the light spectrum, including notably X-ray, as well as acoustic waves and a few others. Among those, thermographic testing proved itself useful in the context of NDT especially in the last few decades with the rise of infrared cameras. These methods consists in measuring temperatures and their evolution, as those depend on the type of material. From this, information about the structure of both the surface and the depth of the material can be gathered. Thermography can be separated into two main groups. On one hand, passive thermography consists in measuring an object without bringing any form of outside energy on the studied object. This monitoring is widely used in the construction industry, most notably for the thermal insulation of buildings, but also for verifications of the electricity and plumbing network. On the other hand, active thermography consists in sending a heat source on the sample that can be punctual, periodic or constant and monitoring the response of the material to it. The most common contactless infrered thermography method is the flash method, and consists in sending energy on the sample thanks to powerful lamps illuminating the sample. Many variants of this general concept exist, the source can be placed on the same side as the measurement (front facing flash), on the opposite side (rear facing flash), the source can be a pulse or a step, as well as a sine wave (lock-in thermography), the source can be a lamp lighting the whole sample, or a laser that acts locally and then can be moved (flying spot), etc.

1.5 Objectives and Summary

Upon starting the work on this thesis, the two main objectives were:

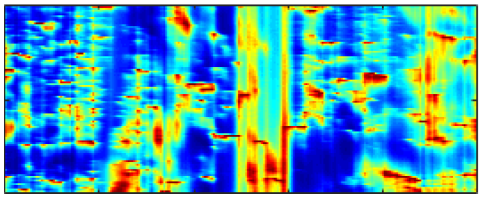
- To investigate already existing super-resolution methods,
- To implement, test and validate the methods that can be applied in infrared thermography.

As work on this thesis progressed, and while experimenting with Compressive Sensing in particular, the idea of undersampling images measured via the scanning methods employed in the lab grew, thus adding this third related objective:

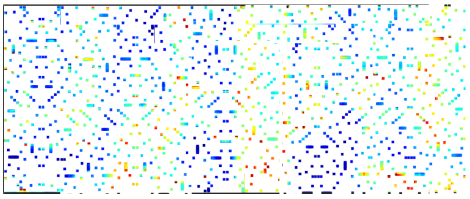
- To reduce acquisition times of point-by-point methods by implementing undersampling techniques.

This thesis is structured the following way in order to answer to these objectives:

- chapter 2 intends to answer the first objective, by first clarifying what is super-resolution, what causes the resolution limits, and what are the main methods that are used in the literature to overcome them. These answers are in two part, one for the diffraction problem, also called in this thesis the analog case, and one for the digital sampling problem.
- chapter 3 is associated to the second objective. A few methods of super-resolution were deemed appropriate in the context of small-scale infrared thermography and thus were explored, notably Structured Illumination Microscopy and a version of Single Pixel Cameras using small intermediary filters.
- chapter 4 and 5 are applications of undersampling techniques, especially using Compressive Sensing in which performance of a few algorithms are tested in order to reduce the number of required samples down to between 10% and 20%. In chapter 4 the Flying Spot technique is employed, while in chapter 5 measurement are made by Scanning Photothermal Radiometry.



(A) Full image, 200×500 points, taking 3h [3]



(B) Subsampled image principle, 10% of the points, taking 18min

FIGURE 1.1: Illustration of the time gain and data reduction with undersampling. The idea is to greatly reduce the acquisition time by reducing the samples, at the cost of a comparatively very small increase of processing (less than a minute).

Chapter 2

Super resolution: Definitions and techniques

This chapter's goal is to give an overview of the main methods that have been developed over the years in various fields to achieve super resolution. It is separated in two main axes: analog and digital, with an emphasis on irregular sampling techniques such as Compressive Sensing as they are at the core of this thesis work.

Contents

2.1 General instrumental model and direct problem . . . 7
2.2 Super resolution in the analog sense . . . 8
2.2.1 Resolution limit and definition of the convolutive problem . . . 8
2.2.2 Deconvolution . . . 15
2.2.3 Scanning methods . . . 16
Confocal Microscopy . . . 16
Scanning Probe Microscopy . . . 17
2.2.4 Super resolution in Fluorescence Microscopy . . . 18
2.2.5 Fourier Ptychography . . . 20
2.2.6 Structured Illumination Microscopy . . . 21
2.3 Super resolution in the digital sense . . . 24
2.3.1 Definition of the under sampled problem . . . 24
2.3.2 Interpolations . . . 28
Radial Basis Functions . . . 30
Kriging . . . 31
2.3.3 Compressive sensing . . . 32
Hypotheses . . . 32
How to promote sparsity . . . 33
Applications . . . 35
Algorithms used in this thesis . . . 36
2.3.4 Specificities of thermal imaging in super resolution . . . 40
2.4 Conclusion . . . 42

2.1 General instrumental model and direct problem

In order to lift the ambiguity regarding the concept of super-resolution, some context is needed on the general instrumental model and the two main categories of problems it implies.

Most measurements are made by an acquisition system that can accurately be described as a linear system, sometimes called linear black box. In this paradigm, a measurement produces an output, the measure y, from an input, often called ground truth x. They are linked by a linear relationship between the two meaning that if there are two pairs of input and output x_1, y_1 and x_2, y_2 then for an input $x_1 + x_2$ the corresponding output is $y_1 + y_2$. Often, the system is said to be Linear Time-Invariant (LTI), where a second condition of time-invariance is added, or in the case of images space-invariance. It states that if the input $x(t)$ is shifted to $x(t + T)$, then so is the output $y(t)$ to $y(t + T)$. A common schematic of such systems is illustrated in Figure 2.1 where the acquisition is described by an operator H leading to an equation of the form:

$$y(t) = H(x(t)) \tag{2.1}$$

Figure 2.1 is called the direct problem, and the objective in super resolution is to solve the associated inverse problem, that is to find x knowing y such that Figure 2.1 is true. Doing so is to revert the effects of the acquisition system and finding the ground truth. In the case of imaging, these effects can be described by separating the acquisition in two distinct parts. The first part will be called the analog system and the second one will be the digital one. The first part corresponds to the imprecision caused by the optical setup, while the second one is associated to the spatial digitisation[1] of the image as seen in Figure 2.1.

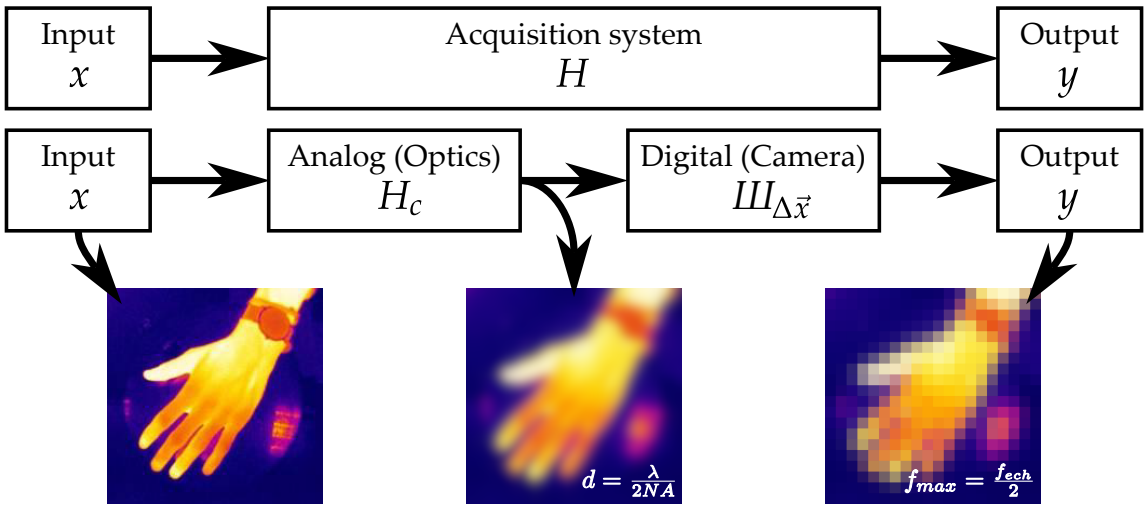


FIGURE 2.1: Schematic of a linear acquisition system (top), separated into continuous and digital parts (bottom)

[1] An image digitisation happens on three axes, the two spatial ones where a step corresponds to a pixel as well of the intensity of the signal where each pixel can take any integer in for instance $[0, 255]$

In the literature, the term of super-resolution can be used interchangeably to solve either of these problems even if there are differences of causes but also in terms of dimensions. Indeed, in the "analog"[2] case the size of the input x and the output y is the same, while in the digital one x is of higher dimensions than the measurement y. The first case could also be described as a convolutive problem in which the objective is to de-blur the image, while the second is a sampling one where the objective is to increase the number of points. Therefore both terms of deconvolution and undersampling (Or at the very least re-sampling) respectively are more precise terms encompassed by super-resolution. The reason they are placed under the same term is that the end objective is the same: To be able to distinguish two points that were too close together in the original measurement.

The linear model of image acquisition can be developed in order to write the aforementioned dichotomy as the following formula, at a given wavelength:

$$y = \text{Ш}_{\Delta\vec{d}}(((x \cdot \varrho) \otimes h_c) \otimes h_{px}) + \epsilon \tag{2.2}$$

In this equation:

- x represents the input (ground truth),
- y represents the output (measurement),
- ϱ in active imaging problems corresponds to the light incoming on the sample,
- $\otimes$ is the convolution product,
- h_c is the convolution kernel associated to the analog problem,
- h_{px} is the convolution kernel corresponding to the size and shape of a single pixel,
- $\text{Ш}_{\Delta\vec{d}}$ is a finite 2 dimensional Dirac comb, where dirac distributions are regularly spaced by distances Δd_x and Δd_y,
- ϵ corresponds to the additive measurement noise.

The goal of the next two sections is to explain what does this equation means, where does it come from and what are the existing methods to solve the inverse problem are. To do so, it is separated into two extreme cases: In section 2.2 the sampling is small enough that the image can be considered spatially continuous, while in section 2.3 the sampling is the only issue.

2.2 Super resolution in the analog sense

2.2.1 Resolution limit and definition of the convolutive problem

Before starting the demonstration of the resolution limit, a couple of mathematical definitions are required. First, the convolution product between 2 functions, noted $\otimes$ in this document[3], is defined as:

$$(f \otimes g)(t) \triangleq \int_{-\infty}^{\infty} f(\tau)g(t-\tau)d\tau \tag{2.3}$$

[2]Sufficiently sampled would be a more correct term here, since in order to be written as vector and matrices that a computer can process, it has to be digital

[3]$*$ is sometimes used in the literature, and $\otimes$ can represent a tensor product

Or, similarly in the discrete case as:

$$(f \otimes g)(n) \triangleq \sum_{-\infty}^{m=\infty} f(m)g(n-m) \tag{2.4}$$

In signal processing, this product is used to represent the effect on the output y of a LTI system acting as a filter on a given input signal x. This filter, usually denoted h, is the impulse response. The product is also known as a weighted moving average which gives a good intuition of how it behaves in practice: Each point of y is calculated as an average of the neighbouring points of x with weight and distance defined in h.

One of the principal properties of this convolution is its behaviour regarding the Fourier and Laplace transforms, where it turns into a pointwise product.

$$\mathcal{F}(f \otimes g) = \mathcal{F}(f) \cdot \mathcal{F}(g) \tag{2.5}$$

with $\mathcal{F}$ symbolising the Fourier Transform (FT) such that:

$$\mathcal{F}(f(x))(v) \triangleq \int_{-\infty}^{\infty} f(x)e^{-i2\pi xv}dx \tag{2.6}$$

This property means that it is more practical and simpler to work and compute in the reciprocal space.

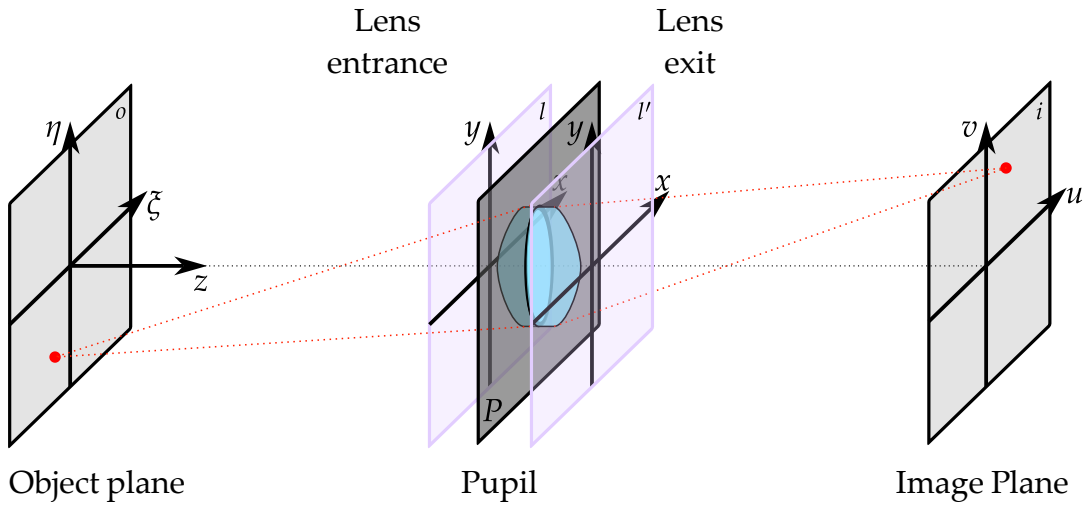


FIGURE 2.2: Schematic of a simple optics setup, with a lens of limited size imaging an object plane onto a screen. The output image (subscript $_o$, coordinates (u, v)) is to be calculated from the input object (subscript $_i$, coordinates (ξ, η)), with evaluations of the amplitude at intermediary steps $_l$ and $_{l'}$ right before and right after the thin lens (Coordinates (x_l, y_l)).

Now, in order to demonstrate the optical resolution limit and its associated convolutive equation, it is possible to describe the imaging system in wave optics, as shown in Introduction to Fourier Optics [4]. Let's start by describing the simplest imaging system comprised of a single lens, illustrated in Figure 2.2, which images an object on a

screen at a given wavelength. The image (output) on the screen is described by a complex amplitude field $U_i(u,v)$ while the object (input) is given by $U_o(\xi,\eta)$, also called the ground truth. These can be linked together by the impulse response in amplitude $\breve{h}(u,v,\xi,\eta)$ of the lens which describes how a single point on the object (Coordinates (ξ,η)) behaves on the image ((u,v)) by summing all of the possible responses, ie. via the following superposition integral:

$$U_i(u,v) = \iint_{-\infty}^{\infty} \breve{h}(u,v,\xi,\eta) U_o(\xi,\eta) d\xi d\eta \tag{2.7}$$

The impulse response can be calculated by considering a single point source as the object. First the light will propagate from (ξ,η) as a spherical wave before interacting with the lens, thus, the amplitude $U_l(x_l,y_l)$ can be written as the solution to the Helmoltz equation for small angles:

$$U_l(x_l,y_l) = \frac{1}{i\lambda z_1} e^{i\frac{k}{2z_1}((x_l-\xi)^2+(y_l-\eta)^2)} \tag{2.8}$$

It then goes through the lens, which implies a phase delay that can be written as a multiplicative factor depending on the path the wave took within and outside the lens for small angles. Therefore it depends on the lens properties such as the refractive index and geometry (Concave or convex curves) which are directly linked to its focal length f. In order to represent the finite dimension of the lens, a function $P(x_l,y_l)$ is used, called pupil function. For a perfect lens without aberrations, it is usually defined as 1 if the amplitude is within the lens, and 0 otherwise. With a thin lens approximation, the distance between the planes l and l' is negligible their coordinates are both (x_l,y_l) and the field on the other side of the lens can be written as a product of the input, the pupil and the phase shift:

$$U_{l'}(x_l,y_l) = U_l(x_l,y_l) P(x_l,y_l) e^{-i\frac{k}{2f}(x_l^2+y_l^2)} \tag{2.9}$$

The field $U_{l'}$ is then diffracted in order to give the impulse response.

$$\breve{h}(u,v,\xi,\eta) = \frac{1}{i\lambda z_2} \int\int U_{l'}(x_l,y_l) e^{i\frac{k}{2z_2}((u-x_l)^2+(v-y_l)^2)} dxdy \tag{2.10}$$

This general equation of the impulse response involves quadratic terms, but considering that U_i and U_o are on the image and object "plane" respectively, and that these planes are in fact spherical caps, ie. part of the surface of spheres of radii z_1 and z_2 [4], $\breve{h}$ can be simplified as a space invariant result:

$$\breve{h}(u-M_z\xi, v-M_z\eta) = \frac{1}{\lambda^2 z_1 z_2} \iint P(x_l,y_l) e^{-i\frac{k}{z_2}((u-M_z\eta\xi)x+(v-M\eta)y_l)} dx_l dy_l \tag{2.11}$$

With $M_z = -\frac{z_2}{z_1}$ the magnification of the lens. This equation is equivalent to Fraunhofer diffraction of the lens aperture defined by $P(x_l,y_l)$. It is therefore proportional to a

[4] A more comprehensive list of possible hypotheses can be found in Introduction to Fourier Optics [4, p.170]

Fourier transform of said aperture, ie. $\mathcal{F}(P) \propto \breve{h}$.

Equation 2.7 then becomes a convolution product $\breve{h} \otimes U_o$:

$$U_i(u,v) = \iint_{-\infty}^{\infty} \breve{h}(u - M_z\xi, v - M_z\eta) U_o(M_z\xi, M_z\eta) d\xi d\eta \tag{2.12}$$

Although the system modelled here is one of the simplest possible imager, this conclusion that a diffraction-limited system can be viewed as a convolution in amplitude between the ground truth and a Fourier transform of a pupil-equivalent is more general. It applies to most linear black box optical systems that produce a real image. It can also be understood as a blurring added to the ideal image that is described by geometrical optics.

These equations, up to Equation 2.12 are on complex amplitude fields. However what is effectively measured is an intensity. The link between the two depends on light coherence, as coherent light systems are linear in amplitude while incoherent ones are linear in intensity. Here the incoherent case will be studied, as images here are made up of a series of secondary sources that are not coherent. The intensity is given by a time average of the squared norm of the amplitude $I(u,v) = \langle |U(u,v,t)|^2 \rangle$, this time corresponds to the integration time of the detector.

$$I_i(u,v) = \iint_{-\infty}^{\infty} \iint_{-\infty}^{\infty} \breve{h}(u - M_z\xi_1, v - M_z\eta_1) \breve{h}^*(u - M_z\xi_2, v - M_z\eta_2) \cdot \langle U_o(M_z\xi_1, M_z\eta_1, t) U_o^*(M_z\xi_2, M_z\eta_2, t) \rangle d\xi d\eta$$

The impulse response is time independent. The amplitudes are considered statistically independent from each other meaning that the average product of amplitudes can only be non-zero when those amplitudes are the same, ie. $\langle U_o(M_z\xi_1, M_z\eta_1, t) U_o^*(M_z\xi_2, M_z\eta_2, t) \rangle = I_0(M_z\xi, M_z\eta)$ only if $(\xi_1, \eta_1) = (\xi_2, \eta_2)$, therefore the measured intensity becomes:

$$\begin{aligned} I_i(u,v) &= \iint_{-\infty}^{\infty} |\breve{h}(u - M_z\xi, v - M_z\eta)|^2 I_o(M_z\xi, M_z\eta) d\xi d\eta \\ &= h \otimes I_o \end{aligned} \tag{2.13}$$

The intensity on the image remains a convolution between an impulse response and the intensity of the object. The intensity impulse response being $h \triangleq |\breve{h}|^2$ is also called Point Spread Function (PSF) in optics, or sometimes blur convolution kernel in image processing. Its Fourier transform $\tilde{h}$ is called the Optical Transfer Function (OTF), in practice the phase part is often discarded, in order to look only at the modulus $|\tilde{h}|$ called Modulation Transfer Function (MTF). According to the convolution theorem, it can also be written as an autocorrelation:

$$\tilde{h} \triangleq \mathcal{F}(h) = \mathcal{F}(|\breve{h}|^2) = \mathcal{F}(\breve{h}) \otimes \mathcal{F}(\breve{h}^*) \tag{2.14}$$

Since according to Equation 2.11 $\breve{h}$ is itself a Fourier transform of the pupil P, the OTF is therefore the autocorrelation of $P(-x_l, -y_l)$, with a normalisation factor. Since the OTF is a spatial FT of the PSF, it is expressed in the reciprocal space, meaning that its axes

are spatial frequencies and expressed in reciprocal distances (m^{-1}) usually noted ν or k in reference to wavenumbers[5], and with the reciprocal space sometimes called k-space, especially in Magnetic Resonance Imaging (MRI) and crystallography.

In practice, the pupil is most of the time circular as it corresponds to the shape of the lens. It can be defined as a function of the radial distance r_d such as:

$$P_c(r_d) = \begin{cases} 1 & \text{if } r_d < R_d \\ 0 & \text{else} \end{cases} \tag{2.15}$$

The $P \otimes P$ OTF corresponding to the pupil can be calculated analytically [4, p.202] to obtain:

$$\tilde{h}_c(\nu) = \begin{cases} \frac{\pi}{2}(\arccos(\frac{r_d}{\nu_0}) - \frac{r_d}{\nu_0}\sqrt{1-(\frac{\nu}{\nu_0})^2}) & \text{if } \nu < \nu_0 \\ 0 & \text{else} \end{cases} \tag{2.16}$$

Where $\nu_0 = \frac{2R_d}{\lambda z_2}$ is the theoretical cutoff frequency.

The impulse response can be calculated using the Fourier-Bessel (aka. Hankel) transform due to the geometry of the problem. The convolution theorem also states that $\mathcal{B}(P \otimes P) = \mathcal{B}(P)^2$, which leads to:

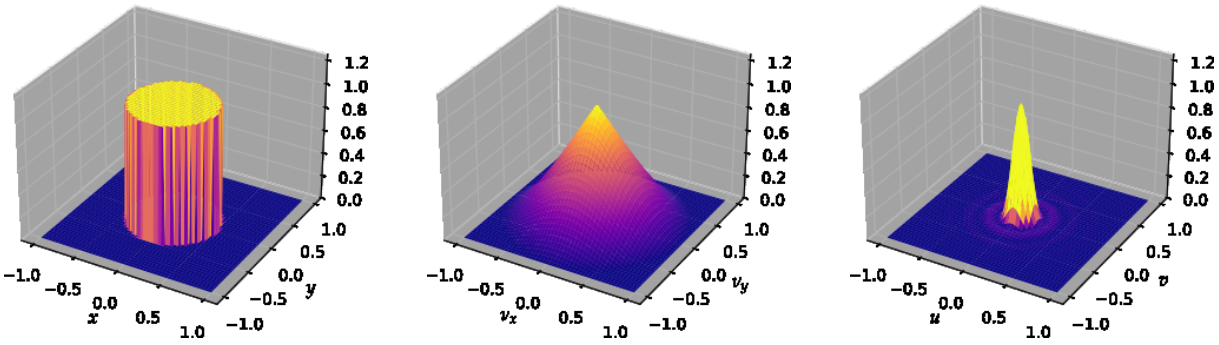
$$I_i(r) = I_0 \frac{2J_1(\frac{kR_d}{z_2} r_d)^2}{\frac{kR_d}{z_2} r_d} \tag{2.17}$$

With J_1 the Bessel function of order 1. This formula corresponds to the well known Airy spot, first noticed experimentally at least 2 centuries ago and its minima were mathematically described in 1835 by G.B. Airy with the famous formula $\sin\theta = 1.22\frac{\lambda}{R_d}$, or expressed in terms of radius $r_{dA} = 1.22\frac{z_2\lambda}{R_d}$. These equations can be viewed as the 2D circular counterpart to the relations between a 1D boxcar function, triangle and cardinal sine function, for P, $\tilde{h}$, and h respectively as illustrated in Figure 2.3.

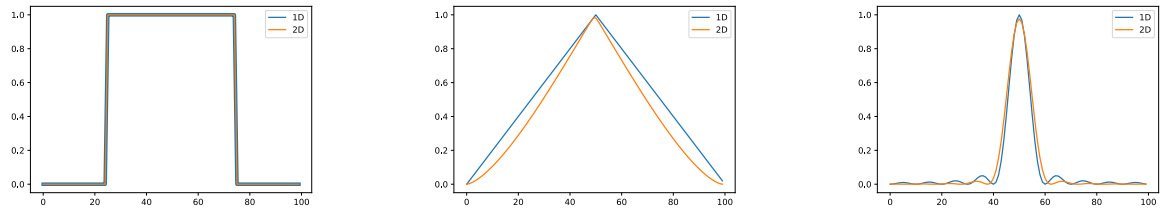
Multiple ways of defining the resolution limit do exist as there is a degree of arbitrariness to it, however the most commonly used limits are the Abbe diffraction limit and the Rayleigh criterion. This last criterion states that the minimum distance at which 2 points of same intensity can be distinguished from each other is of 1 Airy disk radius. It is usually written under the additional assumption that the distance to the lens is close to its focal length in terms of either f-number ($f/n°$) or Numerical Aperture (NA), with $NA \simeq \frac{n}{2f/n°}$ such that $d_{min} = 1.22\frac{\lambda}{2NA}$. Although this definition is experimentally useful, the limit usually remembered is Abbe's diffraction limit which was discovered by considering the object as a diffraction grating with only a small subset of all its diffracted modes going effectively through the lens, leading to a loss of high frequency information. This limit is defined as:

$$d_{min} = \frac{\lambda}{2NA} \tag{2.18}$$

[5]To avoid confusions here: These simply represent spatial frequencies no matter their origin, and aren't linked with the wavelength of the light



(A) 2D Pupil (B) 2D OTF (C) 2D PSF



(D) 1D Boxcar Function (E) 1D Transfer Function (F) 1D Impulse Response

FIGURE 2.3: Typical normalised Transfer Functions and Impulse Responses for a 2D circular lens and a 1D signal. (b) is the normalised OTF of cutoff frequency $\nu_{max} = 1$ obtained by convolution of (a), and (c) is the shape of the inverse Fourier Transform of (b) (Not to scale for lisibility). In (d)-(f) the profiles of (a)-(c) are shown in orange, with a comparison with the similar 1D case of a boxcar function (d), its autocorrelation forming a triangle (e), and a sinus cardinal (f) all shown in blue.

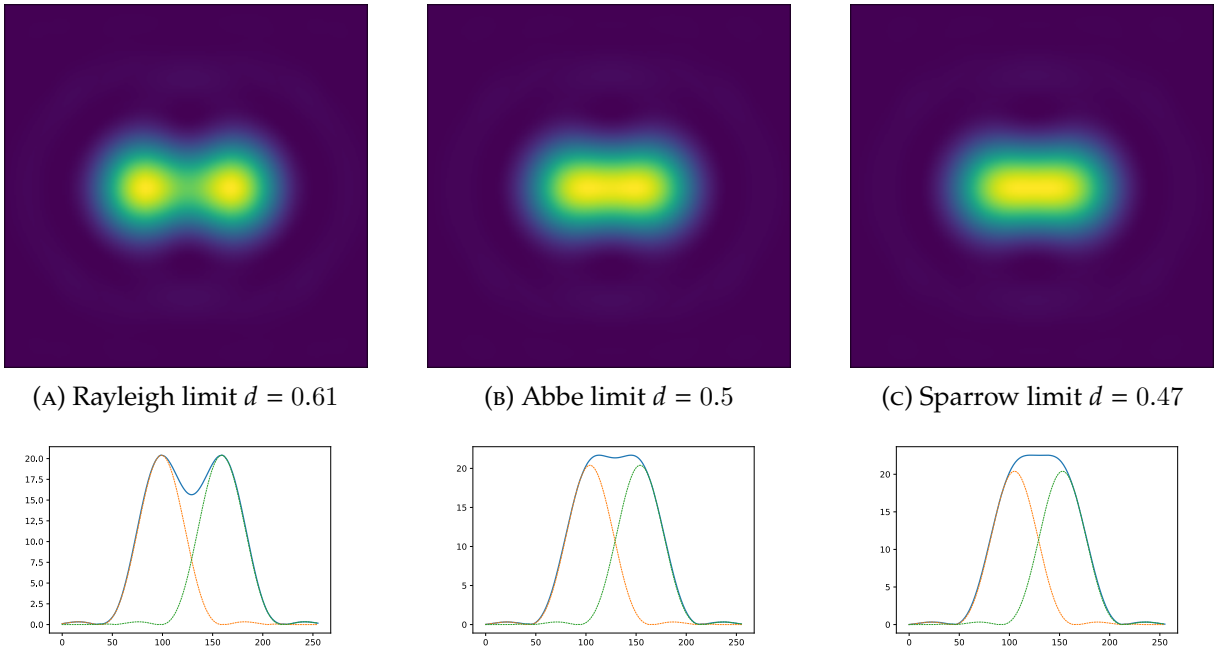
Which is also the limit given by the cutoff frequency given in Equation 2.16 $d_{min} = \frac{1}{\nu_{max}} = \frac{\lambda z_2}{2R_d}$.

The aforementioned arbitrariness stems from the fact that although the cutoff frequency is clearly defined, in practice it is a best case scenario as this minimum distinguishable distance is clearest for 2 spots of same intensities, while if one of the spot is less bright than the other, it will be harder to separate. In other words, it means that low contrast details will be lost sooner than high contrast ones, which also means that separating 2 spots from each other is often more difficult than detecting one. Example of the commonly used limits for sources of same intensities is shown in Figure 2.4.

In addition, any manufactured lens or mirror will cause distortions called aberrations more or less noticeable depending on its process of fabrication, often leading to asymmetries. These aberrations can also be described by their PSF and their OTF, and decomposed in Zernike polynomials. Sometimes, the PSF is modified intentionally in order to filter out unwanted effects, such as in apodisation for instance. For any given frequency, Equation 2.16 acts as a maximum achievable value.

This section shows that in the continuous case or when the sampling is small, that is $\Delta\vec{d} \ll \frac{\lambda}{2NA}$, Equation 2.2 becomes equivalent to the relation between the object and the image is given by Equation 2.13:

$$y = (x \cdot \varrho) \otimes h_c + \epsilon \tag{2.19}$$



(A) Rayleigh limit $d = 0.61$ (B) Abbe limit $d = 0.5$ (C) Sparrow limit $d = 0.47$

FIGURE 2.4: Comparison of commonly used diffraction limits in 2D and their profile at maximum. The Rayleigh limit is defined such that the two points are separated by the radius of the Airy Disk, the Abbe limit is defined at $d = \frac{\lambda}{2NA}$, and the Sparrow limit is the smallest distance between 2 identical sources at which a minimum can be observed between the 2 maxima.

with $(x \cdot \varrho)$ the product between the true object and the illumination equal to the object projected I_o from Equation 2.13.

It shows as well, in Equation 2.18 that to improve resolution is to either decrease λ or increase NA. Working at lower wavelengths is a common solution and is one of the reason why the highest resolution images are acquired thanks to electron microscopy, it also explains the push toward deep UV in semiconductors manufacture. NA is more difficult to change as it depends on the quality of manufacture of the lens, its value is usually in the $[0, 1]$ range. Since $NA = n \sin(\theta) < n$, it is possible to modify n, and NA can be pushed to around 1.5 in the best case scenarios by modifying the refraction index of the medium, that is by replacing the air by some specialised immersion oil [5]. These oils are made such that their refractive index is close to the one of glass in order to avoid refractions. As a general rule, increasing NA also decreases the field of view.

In this thesis, the objective is to measure thermal-induced infrared signals, the order of magnitude of λ can not be meaningfully changed, and the lenses available at those wavelengths have a NA of at most 1 in theory with no known possibility of changing the refraction index. In practice, the objectives with the highest NA are Cassegrain objectives with $NA < 0.8$ [6], and these types of objectives are likely not well suited to any sort of immersion oil if those existed. Therefore, super resolution techniques are studied here.

2.2.2 Deconvolution

Deconvolutions are methods used, as the name implies, to revert the effect of convolution and are therefore inverse methods on Equation 2.19. Although these methods can not be strictly called super resolution, they still are useful to enhance the resolution and to de-blur up to the resolution limit Equation 2.18. Due to the numerical nature of these methods, the problem needs to be written as a matrix product and therefore discretised. In this section it is assumed that the signal is well-sampled such that each step is small compared to the diffraction limit $\Delta\vec{d} \ll \frac{\lambda}{2NA}$. Equation 2.19 is equivalent to:

$$y = \mathbf{H}x + \epsilon \tag{2.20}$$

Where x, y, ϵ are discretised and vectorised versions of their continuous counterparts. $\mathbf{H}$ represents the convolution product with the impulse response and is a block circulant matrix where each line is a vectorised impulse response h. $\mathbf{H}$ is therefore of high dimensions $(n_u n_v) \times (n_u n_v)$. For more detail on the digitisation of the images and matrix, see Figure 2.3.1. Usually, $\mathbf{H}$ is not invertible therefore even if the noise is negligible, it is an ill-posed problem in the Hadamard sense. A wide array of methods exist in order to compute this problem.

The first and most straightforward method is least squares, which attempts to find an input image such that the ℓ_2-norm distance between the measurements and the model output, ie. what the input image would have given as an output considering that the system is known. Its solution is stated as follows:

$$\hat{x}_{LS} = \min_x ||y - Hx||^2 \tag{2.21}$$

It is also possible to compute this minimisation via Fourier transforms:

$$\mathcal{F}(y) = \tilde{h} \cdot \mathcal{F}(x) \tag{2.22}$$

This solution is often described as the "naïve" method, as one would like to simply write the solution of the inverse problem as a point by point division:

$$\hat{x} = \mathcal{F}^{-1}\frac{\tilde{y}}{\tilde{h}} \tag{2.23}$$

However this cannot work as first of all $\tilde{h}$ has to be non-zero on each point[6], and when this value is small (At high frequencies), small variations of y are strongly amplified which, although is the point, means that the noise will be amplified as well, making it impractical.

Effectively applied deconvolution methods are attempts at overcoming this limitation. To do so some prior information about the image have to be brought up in order to create a compromise between the least square result and the general expected allure. This can be done via regularisation, which often consists in adding a penalty constraint to the least square equation:

[6] Even in the best case, Equation 2.16 states that it is 0 after the cutoff frequency, and since computations are made on a bounding square, there always are zeros

$$\hat{x}_{RLS} = \min_x(||y - Hx||^2 + \mu R(x)) \tag{2.24}$$

The regularising function $R(x)$ can be defined in a myriad ways depending on the prior information. One of the most common type of regulariser is the quadratic penalty, called a Tikhonov or ℓ_2 regularisation, which can be defined as $||\mathbf{D}x||^2$ with $\mathbf{D}$ being at least block circulant similarly to $\mathbf{H}$, with d its defining line vector and $\tilde{d}$ its Fourier transform. This can be expressed in the Fourier space as a simple equation [7]:

$$\mathcal{F}(\hat{x}) = \frac{\tilde{h}^*}{|\tilde{h}|^2 + \mu|\tilde{d}|^2}\tilde{y} \tag{2.25}$$

This process can be identified to the Wiener deconvolution when $\tilde{d}$ is defined as the Noise to Signal Ratio (Or at least its best available estimation) ie. the ratio between the power spectral density of the noise ϵ and the power spectral density of the signal x. However since the estimation of these power densities is not always easy, $\tilde{d}$ can be defined in various other ways such as a high-pass filter. Some methods are based on this principle and also estimate these parameters [8], as implemented for instance in Scikit-Image.

Other methods stem from Equation 2.24, but usually don't have an analytical solution, and are instead calculated iteratively. This is the case for the Richardson-Lucy deconvolution [9] which is one of the most used deconvolution algorithm. It works under the assumption of Poisson noise, which makes it especially relevant when working in systems where the shot noise is important (such as low light conditions), and solves iteratively an equation of the form:

$$\hat{x}^{(t+1)} = \hat{x}^{(t)} \cdot \left(\frac{y}{\hat{x}^{(t)} \otimes h} \otimes h^*\right) \tag{2.26}$$

2.2.3 Scanning methods

As shown in subsection 2.2.1 the resolution of any lens-based imaging system is limited by its aperture and therefore its transfer function. One category of solutions to this problem are based on scanning techniques: Instead of capturing an entire image at once, a probe travels along the material under study going point by point. It can be advantageous to do so in terms of resolution because these methods can either get rid of lenses entirely bypassing the resolution limit in the case of Scanning Probe Microscopy techniques or by using techniques such as confocal microscopy that exchange field of view and depth of field for resolution. In this subsection the main relevant methods are briefly described but for a more exhaustive overview of scanning microscopy see [10].

Confocal Microscopy

A confocal microscope uses small pinholes to filter out light that is outside of the image plane, increasing the final image resolution. Besides those pinholes, it uses the same components as a conventional microscope, making it a very widespread method.

By placing a small pinhole at the source, the sample is illuminated locally by a spot that is described by the point spread function, then a pinhole at the detector will collect light along the point spread function. Therefore the illumination on the detector is defined by a product of impulse responses that are identical, which can be expressed in the Fourier space as a convolution with itself ie. an autocorrelation. This implies that the theoretical cutoff frequency is twice that of the same system in widefield (ie. non-confocal) mode. In practice the effective maximum frequency is lower as the gain of the OTF at high frequencies is very low, the design of the pinholes is never exactly a Dirac point source but roughly the size of the first Airy disk [11], and the pinholes limit the amount of light that goes on the sample and on the detector.

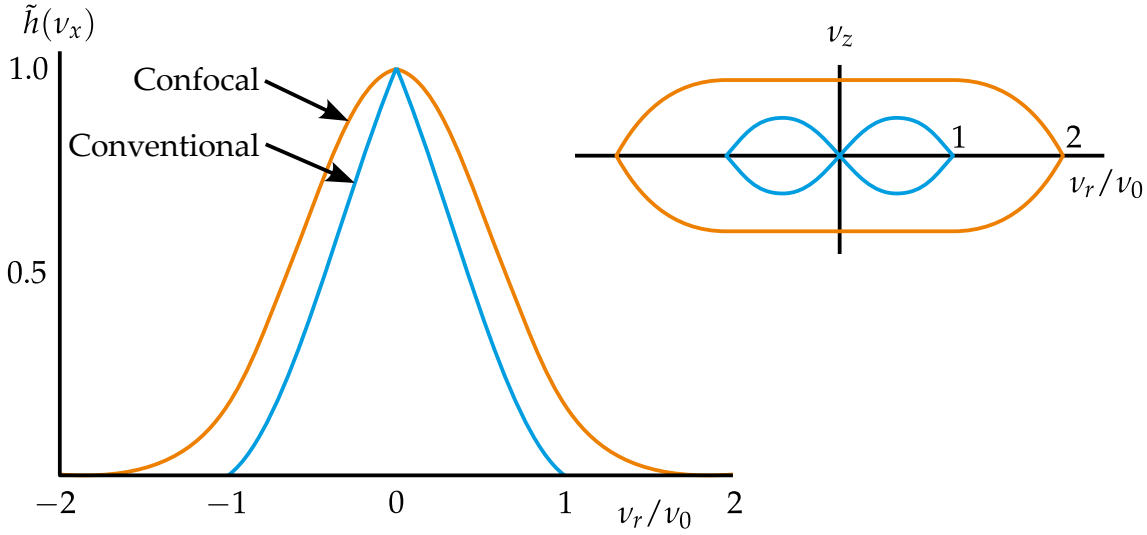
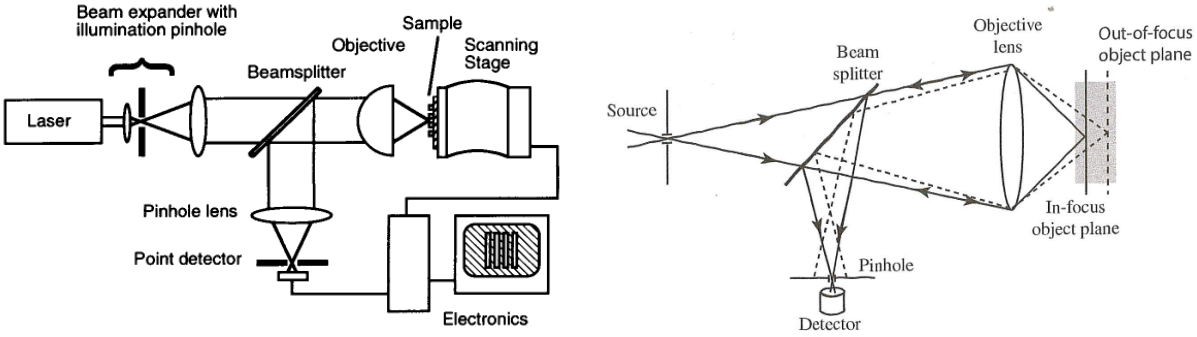


FIGURE 2.5: Profile of OTF for equivalent classical and confocal microscopes [4], with a schamatic of its shape in depth (top right) [12, p.298]

Although there is an increase in lateral resolution as shown in Figure 2.5, confocal microscopes are used for their property regarding their depth resolution. As the detector pinhole also physically filter out light that is out of focus, it reduces the depth of field of a single measurement, as shown in Figure 2.6b. This makes it possible to also scan the sample in depth with a good resolution along this axis as well, especially for transparent media but also for scanning rough and uneven surfaces.

Scanning Probe Microscopy

Scanning Probe Microscopy (SPM) encompasses a wide variety of microscopy techniques and their variants such as Atomic Force Microscopy (AFM), Scanning Tunneling Microscopy (STM), Scanning Near Field Optical Microscopy (SNOM), Scanning Thermal Microscopy (SThM) and many others. These systems are lensless and can measure various physical properties, not necessarily light, their resolution can't be limited by diffraction both due to the lack of lens. The comparison can be made more precisely for methods such as SNOM that measure light: Because the measurement is made in the near-field, the approximations that led to the Fraunhoffer diffraction result are not applicable. A description using Fresnel diffraction shows that, when the distance to the



(A) Instrumentation (B) Geometric blocking of light

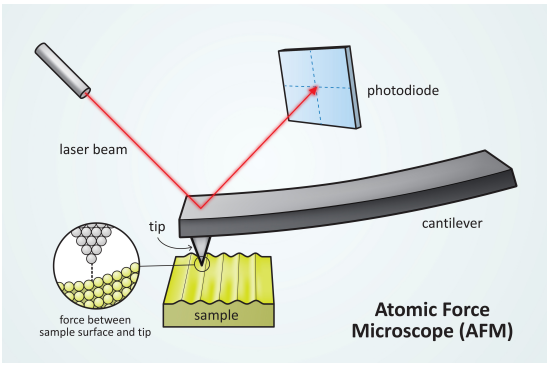
FIGURE 2.6: General setup of a confocal scanning laser microscope (a) [10]. Out-of-focus light is physically blocked by the pinholes, as illustrated in (b) [4]

probe is similar or smaller than its aperture, the size of the aperture and of the resulting diffraction pattern are similar. SNOM is therefore, like the other SPM methods limited by the size of the probe or the light source, and, as will be detailed in the following chapters, especially 5, the minimum step size between successive measurements. The type of probe used vary depending on the property measured, with the two main ones being thin optic fibre and cantilever probes. An example of both is shown in Figure 2.7b in the context of SThM. SPM encompasses active and passive methods, with and without contact. In passive methods such as AFM, no input source is required on the sample, even if the method usually involves a laser it is only used to measure the displacement of the lever and doesn't interact with the sample as shown in Figure 2.7a. In active methods, the light source and detectors can be in opposite configurations as illustrated in Figure 2.7c, as long as one of them is thin enough.

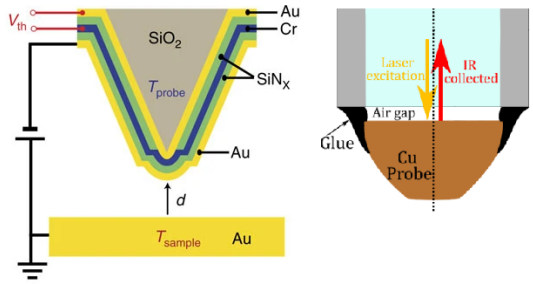
2.2.4 Super resolution in Fluorescence Microscopy

Fluorescence microscopy is a widely adopted technique in life sciences, the principle consists in labelling molecules by fixing a known fluorophore to it (ie. A particle that can emit fluorescent light). This fluorophore will, when exited at a certain wavelength, re-emit a photon of lower energy (ie. longer wavelength) within a few nanoseconds in a process called fluorescence. Although it can be viewed at first as a contrast-enhancing technique[7], fluorescence microscopy has been the basis of many developments related to super-resolution, which could seem counter-intuitive as it implies an upward shift in wavelength. However, in practice this effect matters less than the fact that these images are sometimes made on tagged particles that are spatially spread out from each other, and where each fluorophore is considered a punctual source. This simplifies greatly the problem as the main issue is rarely to distinguish two spots from each other, but instead simply locate them by their PSF assuming no overlap. The resolution limit has been shown in [15] to be approximately $\Delta x = \frac{\lambda}{\sqrt{N_\gamma} 2NA}$ with N_γ the number of photons received on the image from the fluorophore, in other words, if enough photons are

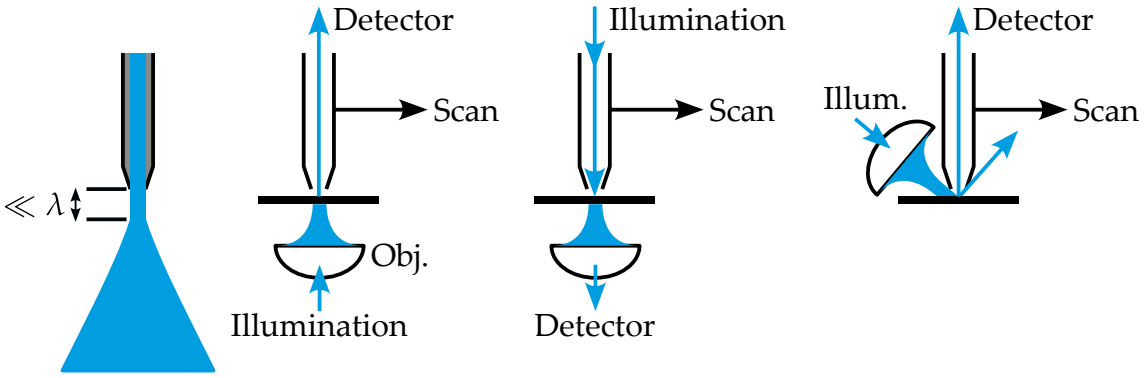
[7] There are some conceptual similitudes between this and Dye Penetrant inspection in industrial NDT



(A) AFM setup



(B) SThM probes



(C) Main Scanning modes

FIGURE 2.7: (a) AFM Principle [13]. (b) Probes used in SThM : cantilever tip [14] (left) and fibre probe [3] (right). (c) Usual scanning modes, with different positions of source and detector. The source can be wide field in transmission or reflection, and the detector gathers the signal going through the thin probe, or vice-versa [10], as long as the probe is close enough ($\ll \lambda$ in SNOM)

measured, their average position (ie. the centre of the PSF) is a precise position of the source.

Resolution in fluorescence microscopy can be enhanced via a few methods, such as the aforementioned confocal microscopy, techniques such as Stimulated Emission Depletion Microscopy (STEDM) [16], and in wide field Photo Activated Localisation Microscopy (PALM) [17] and Stochastic Optical Recons Microscopy (STORM) [18]. These last three methods developed in the late 1980s to 1990s were awarded the 2014 Nobel prize in Chemistry. The overarching principle is to engineer a setup which can activate only part of the fluorophores such that the problem goes from distinguishing adjacent sources to locating a single source with no currently emitting neighbour.

STEDM is a scanning method where 2 signals are sent toward the sample. The first one is the fluorescence excitation that all local fluorophores will absorb. The second one is at a wavelength such that it can induce a des-excitation via a process of stimulated emission, while still being different from the fluorescence wavelength. This second beam is carefully engineered such that it is a torus shape, usually by creating an optical vortex, which is a phenomenon of interferences within a beam of circular polarisation. This implies that the sample can only re-emit a fluorescence signal within the inside

of the torus shape, since the inner part of the torus can be made very small due to it being made via interferometry with itself thanks to a phase plate (vortex). This method doesn't require more processing than in the confocal case.

PALM and STORM are widefield methods based on the use of photo-activable or switchable fluorophores. First, some of the particles are activated by applying a weak light at the wavelength required for photo-activation, which activates at random a small subset of fluorophores. Then the sample is illuminated so that the activated particles can fluoresce while the others remain inactive, an image of the response at this step is taken, which should be made of only a few clearly separated points. Thirdly the currently activated fluorophores are turned off, via photobleaching, which is a fading effect that happens irreversibly when exposed to sufficient illumination, or other similar means of de-activation. These three steps are repeated until most of the particles have been measured. The resulting set of data is therefore a series of images each comprising of a small subset of particles. They therefore can first be accurately localised separately on each images, then the sum of the deduced locations can be plotted in which 2 nearby particles can be distinguished as long as they were activated on different images.

These methods require the addition of those carefully chosen particles. The limitations are therefore that they are somewhat invasive and not contactless, and that the fluorophore (Or similar photo-switchable particle) that is measured as a "proxy" needs to be such that it somehow attaches to the object of interest while still behaving as a point source. These methods are therefore mostly used in life science and adjacent fields.

2.2.5 Fourier Ptychography

Fourier ptychography [19] is a technique where the goal is to shift frequencies. This shift is achieved by illuminating with coherent light the object with a source outside the main optical axis.

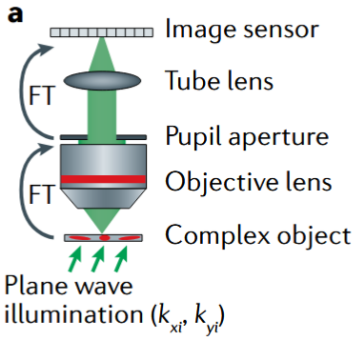


FIGURE 2.8: Fourier Ptychography principle [20]

More specifically, Fourier Ptychography uses array of point sources. These sources are then illuminated sequentially and the corresponding system response is measured. In terms of frequencies, each image remains still diffraction limited, with the same OTF caused by the experimental setup. However, the fact that the source is off-axis leads to a shift in frequencies on the same axis, meaning that it contains information of higher

frequencies. This corresponds mathematically to:

$$\begin{aligned} I_c(u,v) &= |U(u,v) \cdot e^{i\nu_{u_i}u} e^{i\nu_{v_i}v} \otimes \breve{h}(u,v)|^2 \\ &= |\mathcal{F}^{-1}(\hat{U}(\nu_x - \nu_{u_i}, \nu_v - \nu_{v_i}) \cdot P(\nu_u, \nu_v))|^2 \end{aligned}$$

with U the complex object, interacting with the tilted plane wave described by $e^{i\vec{\nu}_i\vec{u}}$ and $\vec{\nu}_i$ the position of the off-axis source, $\breve{h}$ the coherent PSF, and P its fourier transform and the pupil. The objective is to extract these and then sift them back to add them to the Fourier transform of the image in-axis. The effective resolution of the complex amplitude is $d = \frac{\lambda}{NA_{obj}+NA_{illum}}$ [20]. This is to compare with $d = \frac{\lambda}{NA_{obj}}$ the limit for coherent light that is twice as big as the limit for an incoherent case. A similar frequency-shift concept has been used in the context of pump-probe quantitative phase imaging [21].

2.2.6 Structured Illumination Microscopy

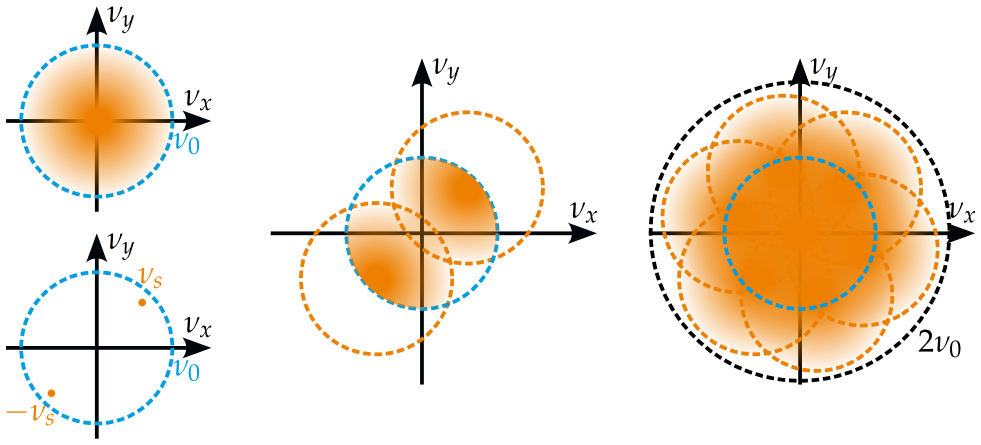


FIGURE 2.9: Structured Illumination Microscopy (SIM) principle. Fourier transform of an object (top left) and of a sine wave at a given direction (bottom left). The blue circle represent the diffraction limit. The FT of the product of the two shifts the frequencies such that some frequencies that were previously after the cutoff are brought back into the MTF, overlapped with the other components of the signal. If this information is extracted, it can be used to know the missing crescent. After repeating the process for multiple angles (right) it is possible to reconstruct an image with up to 2× the frequency limit. [22]

SIM encompasses a few techniques starting with a method [22] based on frequencies shifts similarly than in Fourier Ptychography, with the main difference being that incoherent light can be used. This frequency shift is induced by the use of a light source structured such that it contains known high-frequency fringes. The superposition of these fringes with the object causes a frequency shift of the entire spectrum such that previously filtered frequencies end up within the OTF of the setup. This effect is known

as moiré patterns, and is also very similar to aliasing effects that happen with digitised signals as will be detailed in the next section.

As represented in Equation 2.19, the light to be imaged is the product between the source ϱ and the object x. The Fourier transform of y is:

$$\tilde{y} = (\tilde{x} \otimes \tilde{\varrho}) \cdot \tilde{h}_c \tag{2.27}$$

If the illumination ϱ is chosen to be a sine wave along an axis u, ie.:

$$\varrho(u) \propto 1 + \cos(\nu_s u + \varphi) \tag{2.28}$$

Its Fourier transform in spatial frequencies ν is:

$$\tilde{\varrho}(\nu) = \delta(\nu) + \frac{1}{2}\delta(\nu - \nu_s)e^{-i\varphi} + \frac{1}{2}\delta(\nu + \nu_s)e^{i\varphi}) \tag{2.29}$$

$$\tilde{y}(\nu) = (\tilde{x}(\nu) + \frac{1}{2}\tilde{x}(\nu - \nu_s)e^{-i\varphi} + \frac{1}{2}\tilde{x}(\nu + \nu_s)e^{i\varphi}) \cdot \tilde{h}_c(\nu) \tag{2.30}$$

By repeating the measure with 3 phases φ, it is possible to extract the spectrum $\tilde{x}(\nu \pm \nu_s) \cdot \tilde{h}_c(\nu)$. Then these spectra can be added at their position $\pm\nu_s$ to the original, non-frequency shifted one. This can double the effective frequency range of the image along one axis, with a possibility to go even beyond if the illumination pattern can be made smaller in cases where it doesn't go through the same optics as the image. The same can finally be repeated along different axes in order to increase the resolution on all directions. The principle is schematically illustrated with representations of the spectral responses in Figure 2.9

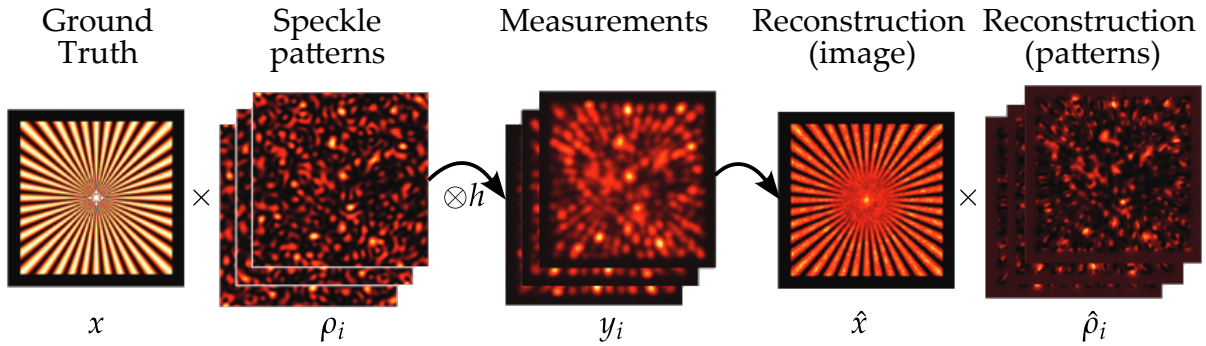


FIGURE 2.10: Random Structured Illumination principle. A series of measurement is taken of the point-by-point product of the object and a series of random unknown patterns, both are then reconstructed using a minimisation algorithm, and by assuming that the illumination is on average homogeneous [23]

Variants have been developed with illuminations structured differently, including with random illuminations, first proposed in [23] named blind-SIM, with further improvements then called Random Illumination Microscopy (RIM) [24, 25]. The main steps of the process are illustrated in Figure 2.10. Instead of sending a sine wave along

an axis, a speckle is used, made of many small light grains. These patterns contain a very high frequency information that is layered on top of the object. The resolution limit has been shown experimentally to be $d_{bSIM} \simeq \frac{\lambda}{2NA_{sys}+2NA_{illum}}$ in [26], although the explanation is way more complex than in the sine structures case [27]. Due to these differences with the more classical approach of SIM, some have called similar principles Incoherent Fourier Ptychography [28], while others link it with STORM since the structured light has a very similar effect to the repeated random excitation of fluorophores [29]. The fundamental reason for the super-resolution can be also understood in a similar way as in the STORM case, as it is easier to precisely localise a point rather than distinguish two close ones. This random structured illumination approach is studied in this thesis in section 3.2.

2.3 Super resolution in the digital sense

2.3.1 Definition of the under sampled problem

Mathematical definition

In this section, the assumption will be the opposite of the analog case. The acquired image is digitised such that each pixel (or sample) is bigger than the size of the point spread function $\Delta\vec{d} \gg \frac{\lambda}{2NA}$. In the idealised case the PSF remains a Dirac, ie. the optics are considered ideal such that the image can be described entirely through geometrical optics. Equation 2.2 then becomes:

$$y = \text{Ш}_{\Delta\vec{d}}((x \cdot \varrho) \otimes h_{px}) + \epsilon \tag{2.31}$$

The problem is dominated by the Dirac comb Ш (Sometimes called Dirac brush or bed of nails in 2D images) representing the sampling. It is defined as a periodic series of evenly spaced Dirac distributions:

$$\text{Ш}_{\Delta\vec{d}}(u, v) \triangleq \sum_l \delta(u - l\Delta d_u) \sum_m \delta(v - m\Delta d_v) \tag{2.32}$$

Δd_u and Δd_v correspond to the space between each sample along each axis. For most images taken by cameras, it is the distance between the center of 2 successive pixels, called pixel pitch, and $\Delta d_u = \Delta d_v$. l and m are indices corresponding to the pixel number $l \in [0, N_u]$. In this definition, the Dirac comb is therefore finite. In the literature l and m is usually defined on $[-\infty; \infty]$, in which case a multiplication with a boxcar function is needed in order to represent a real measurement.

Although uncommon, some instruments can be structured in other ways, such as with a honeycomb pattern. In those cases the function need to be modified to take into account an offset between successive lines, but there remains a periodic structure of evenly spaced samples which doesn't fundamentally change the principle [30].

The sampling for images can't be described only by an infinitely small point sample, as a surface is required to detect a sufficient enough amount of photons. In order to model this effect, a convolution product is made between the object and the impulse response of the pixel h_{px}, assuming all pixels are identical. Ideally, h_{px} is the shape of the pixel, thus acting as a 2D equivalent to a zero-order hold. In practice it can take into account the presence of wires, gates, microlens, etc. creating blind spots or slightly non-homogeneous responses depending on where the photon landed on the pixel. In systems where the sampling is made not via the detector but instead via the movement of the source (or the movement of both), h_{px} can represent the light spot on the sample if it is smaller than the detector, as it would represent the effectively used part of the sensor.

Equation 2.32 is written with the omission of the variables, the object x is a continuous function $x(u, v)$, while the image y is a discrete function $y(l, m)$. When shown on a screen of different resolution, $y(l, m)$ is represented as piecewise continuous (constant) function via a nearest-neighbour interpolation that can be also interpreted by a convolution with a boxcar function of size Δd.

Aliasing and the Sampling theorem

The Dirac comb function has two main interesting properties, the first one is the one that is used in Equation 2.32, where it can represent a regular sampling via a multiplication, as each translated Dirac delta extracts the value at its position. The second one is that it can model a periodisation via a convolution, due to the fact that a convolution with a translated Dirac delta gives a delayed function of the same amount. Lastly, the Fourier transform of a (infinite) Dirac comb remains a Dirac comb[8] such that $\mathcal{F}(\text{Ш}_d) = \frac{1}{d}\text{Ш}_{\frac{1}{d}}$ where d is the sampling step, and $\frac{1}{d}$ the sampling rate.

The simplest model for a sampled signal is such that:

$$y = \text{Ш}_d x \tag{2.33}$$

It can be expressed in the Fourier space via the convolution:

$$\mathcal{F}(y) = \text{Ш}_{\frac{1}{d}} \otimes \mathcal{F}(x) \tag{2.34}$$

This formula represents the Discrete-Time Fourier Transform (DTFT)[9]. The inverse of $\mathcal{F}(y)$ can be deduced from only one period, with 2 different cases: Either $\mathcal{F}(x)$ is contained within one period, in which case the inverse of $\mathcal{F}(y)$ is x, or it is bigger than one period, in which case the aliasing phenomenon arises, where the spectrum shifted to the adjacent periods overlap with the current one. In digital images this aliasing causes artefacts such as moiré patterns.



(A) DTFT (B) Aliased DTFT

FIGURE 2.11: Aliasing Principle

This is the underlying phenomenon behind the sampling theorem, which was described by various authors during the first half of the 20th century such as Shannon and Kotelnikov. The version written by Shannon [31] states that any band-limited function with a maximum frequency ν_{max} can be determined (ie. all its frequencies recovered) from samples spaced by $\frac{1}{2\nu_{max}}$.

As in images the sample rate $\nu_s = \frac{1}{d}$ tends to be fixed, it means that only signals with a maximum frequency $\nu_{max} < \frac{\nu_s}{2}$ can be reconstructed. As seen in the previous section, images are limited in spatial frequencies by the OTF of the setup, therefore ν_{max} is the cutoff frequency ν_0. A way to avoid aliasing in images is then to change the setup, usually by closing an aperture, such that its effective NA is lowered to the point where the sampling theorem is respected.

When computing the Fourier Transform digitally, the Discrete Fourier Transform (DFT) is calculated usually thanks to a Fast Fourier Transform algorithm (FFT) algorithm. The DFT is a discretised version of the DTFT. This discretisation in the Fourier

[8] Proof using Fourier series

[9] Again, here, Discrete-Space would be more accurate

space induces a periodisation in the real space, the smallest step size in the Fourier space is limited by the size of the entire signal (before it was periodised). This space periodisation can cause discontinuities, in which case applying windows to the signal (and thus making both edges equal) can alleviate artefacts. Finally, the FFT algorithm will compute a single period of the DFT, and the Inverse FFT (iFFT) will compute a single period of the signal. The FFT is such that the amount of samples is the same as the amount of samples in the signal N, although due to the symmetry sometimes only half of the spectrum is represented. Also due to the fact that the period is arbitrarily chosen, the FFT can be represented on different intervals, for images it is usually $[-\frac{\nu_s}{2}, \frac{\nu_s}{2}]$ but $[0, \nu_s]$ is equivalent. To sum up, the smallest step in the FFT is $\frac{1}{dN}$ and defined by the size of the entire signal, and its extrema frequencies $\pm\frac{1}{2d}$ are defined by the step size of the image.

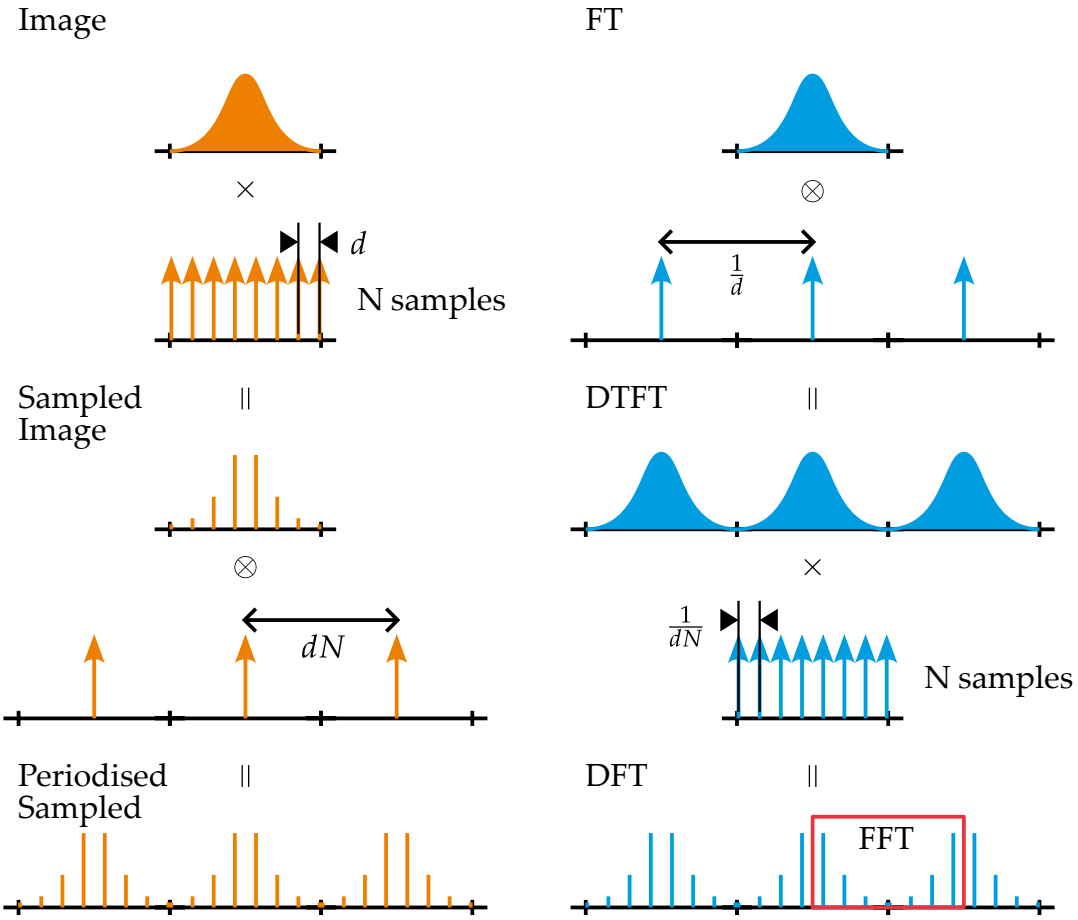


FIGURE 2.12: Steps for the digitalisation of an image and its Fourier transform

In the context of digital signals, super-resolution is defined as the act of "overcoming" the sampling theorem limit which can either be viewed as reconstructing signals with frequencies higher than $\frac{\nu_s}{2}$ or equivalently as finding sub-pixel information. As it will become clearer in the following sections, this idea needs to be nuanced a bit, as the super resolution methods are not really about "going beyond" the limit but are rather

about moving this limit to higher frequencies. Super-resolution in that context can also be viewed as equivalent to (willingly) under-sample, as under-sampling is about reducing the number of samples such that the sampling theorem is not verified.

Mathematical tools

Computations with digitised images require a few considerations that are quickly described here.

First, the black box problem from Equation 2.1 can be written as a matrix product:

$$y = \mathbf{H}x \tag{2.35}$$

in which the images y and x have to be written as vectors (Otherwise the dimensions wouldn't match) of size $N = N_u N_v$. This is done by first sampling the images along both axes, then by taking each line (or column) and concatenating them in a single line vector. A reshape can be done in either by reading line first or columns first, and are often called C and Fortran order, either way work as long as it is done consistently throughout. $\mathbf{H}$ is the operator linking the two, a matrix of the same dimension as y in columns, and same as x in lines.

Assuming that the operator is a convolution, the matrix $\mathbf{H}$ represents the convolution product with the impulse response. Each line contains the impulse response h, that is a vectorised image of the same dimensions as x. For instance, the first line of $\mathbf{H}$ contains the response image to the very first point on the image x, the second line contains the response for the second point, etc. By construction, $\mathbf{H}$ contains only h shifted by one column every time[10], ie.:

$$\mathbf{H} = \begin{pmatrix} h_1 & h_2 & h_3 & \cdots & h_N & 0 & 0 & \cdots \\ 0 & h_1 & h_2 & \cdots & h_{N-1} & h_N & 0 & \cdots \\ 0 & 0 & h_1 & \cdots & h_{N-2} & h_{N-1} & h_N & \cdots \\ \vdots & \vdots & \vdots & \cdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tag{2.36}$$

This is a Toeplitz matrix, that necessarily arises with the shift-invariance. This matrix just like the convolution it represents can be defined in a few ways depending on the edge conditions. Here, what is almost always done is to consider periodic objects, thus turning the matrix into a circulant one:

$$\mathbf{H} = \begin{pmatrix} h_1 & h_2 & h_3 & \cdots & h_N & h_{N-1} & h_{N-2} & \cdots & h_2 \\ h_2 & h_1 & h_2 & \cdots & h_{N-1} & h_N & h_{N-1} & \cdots & h_3 \\ h_3 & h_2 & h_1 & \cdots & h_{N-2} & h_{N-1} & h_N & \cdots & h_4 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \end{pmatrix} \tag{2.37}$$

The periodic hypothesis is made due to this matrix being diagonalisable in the Fourier basis:

$$\mathbf{H} = \mathbf{F}^\dagger \operatorname{diag}(\tilde{h})\mathbf{F} \tag{2.38}$$

with diag an operator that takes a vector and puts it in the diagonal of a matrix, and $\mathbf{F}$ the Discrete Fourier Transform matrix.

[10] And assuming that the impulse response is of limited size

The DFT matrix for a 1D signal[11] of size N_u is defined as the following, with $\omega = e^{-\frac{i2\pi}{N_u}}$:

$$\mathbf{F} = \frac{1}{\sqrt{N_u}} \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega & \omega^2 & \cdots & \omega^{N_u-1} \\ 1 & \omega^2 & \omega^4 & \cdots & \omega^{2(N_u-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{N_u-1} & \omega^{2(N_u-1)} & \cdots & \omega^{(N_u-1)(N_u-1)} \end{pmatrix} \tag{2.39}$$

Due to the diagonalisation, $y = \mathbf{H}x$ becomes $y = \mathbf{F}^\dagger \operatorname{diag}(\tilde{h})\mathbf{F}x$, which can be expressed in the Fourier space by multiplying by $\mathbf{F}$, leading to $\mathbf{F}y = \operatorname{diag}(\tilde{h})\mathbf{F}x$. $\mathbf{F}x$ and $\mathbf{F}y$ are the Fourier transform of x and y, and can also be calculated via the FFT, bypassing the use of the DFT matrix in calculations and thus turning a computationally very heavy matrix product into a simpler one, requiring a FFT, an iFFT and a vector product between $\tilde{h}$ and $\tilde{x}$. This is the reason why the usual advice in the literature is to not compute the transfer matrix directly [7].

When working with images of size $N = N_u \times N_v$, the matrix $\mathbf{H}$ is not only circulant but even Block Circulant with Circulant Blocks (BCCB) matrix, which are matrices of the following form [32]:

$$\mathbf{H} = \begin{pmatrix} \mathbf{A_1} & \mathbf{A_2} & \cdots & \mathbf{A_{N_u}} \\ \mathbf{A_{N_u}} & \mathbf{A_1} & \cdots & \mathbf{A_{N_u-1}} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A_2} & \mathbf{A_3} & \cdots & \mathbf{A_1} \end{pmatrix} \tag{2.40}$$

with each $\mathbf{A_i}$ circulant:

$$\mathbf{A_i} = \begin{pmatrix} h_{i,1} & h_{i,2} & \cdots & h_{i,N_v} \\ h_{i,N_v} & h_{i,1} & \cdots & h_{i,N_v-1} \\ \vdots & \vdots & \ddots & \vdots \\ h_{i,2} & h_{i,3} & \cdots & h_{i,1} \end{pmatrix} \tag{2.41}$$

This form of matrix naturally arises when attempting to express a multidimensional convolution product with periodic edge conditions, and the structure being a direct consequence of the reshape of the data into vectors. Each dimension adds an additional block circulant layer, for instance a 3D convolution is a block circulant matrix containing block-circulant sub-matrices themselves circulant. h is the impulse response image of size $N_u \times N_v$ in the example, thus $\mathbf{A_i}$ contains the impulse response of the i^{th} line. Each line of $\mathbf{H}$ still contains the vectorised impulse response as in 2.37

2.3.2 Interpolations

Although not computationally efficient, a straightforward way to turn Equation 2.33 into a matricial product is to write the following:

$$y = \mathbf{H}x \tag{2.42}$$

[11] For the 2D case, $\mathbf{F_{2D}} = \operatorname{kron}(\mathbf{F_u}, \mathbf{F_v})$, see 3.12

With x being a discretised but high resolution (N points) version of the signal, and y the low resolution measurements ($M < N$ points) [33]. $\mathbf{H}$ is then a $N \times M$ matrix representing the sampling scheme. This rectangular matrix can be constructed from a $N \times N$ identity matrix by extracting every $\frac{M}{N}$ lines. This matrix is not invertible, thus finding x from the measurements y is an ill-posed problem as there exist an infinity of x such that they go through the points of y. Therefore a least squares approach is well suited to attempt to solve this problem. If an inverse for $\mathbf{H}^T\mathbf{H}$ exists, x is such that:

$$x = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T y \tag{2.43}$$

With $\mathbf{H}$ defined as above $\mathbf{H}^T H$ is the identity matrix, therefore the simplest way to approximate x is to use the transpose matrix of $\mathbf{H}^T$, and will result in values outside of the sampling positions being set to 0. The blurring caused by the pixel can easily be modelled by constructing $\mathbf{H}$ from the Toeplitz matrix of h_{px} instead of the identity matrix. $\mathbf{H}^T\mathbf{H}$ remains the identity matrix up to a scaling factor in this case, as long as there is no overlap between pixels (ie. the pixel pitch is always bigger or equal to the pixel). With this added information to $\mathbf{H}$ estimating x from its transpose will lead to a part continuous function similar to a nearest-neighbour interpolation, also called zero-order hold for signals. The first-order hold, or linear interpolation can be defined of a similar manner, by defining h as a normalised triangle function calculated from the auto-correlation of the boxcar function of size equal to the sampling step (aka. the ideal pixel), with the notable difference that $\mathbf{H}^T\mathbf{H}$ is (merely) symmetric. The generalisation to higher sizes leads to successively wider cardinal B-splines [34, 35]. The overarching idea of this first type of methods is to calculate the value at any point thanks to a distance-weighted average of the neighbouring samples, and thus behaves like a convolutive problem.

Other methods rely on various basis change, such as in polynomial interpolation [36]. The goal in polynomial interpolation is to re-write the problem such that the signal can be expressed as a polynomial of degree equal to $N-1$ with coefficients to be determined with the measurements and their positions. $P(u) = c_0 + c_1 u + \cdots + c_{N-1}u^{N-1}$ verified for each point (u_i, y_i).

$$y = \begin{pmatrix} 1 & u_0 & \cdots & u_0^{N-1} \\ 1 & u_1 & \cdots & u_1^{N-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & u_{N-1} & \cdots & u_{N-1}^{N-1} \end{pmatrix} \begin{pmatrix} c_0 \\ c_1 \\ \vdots \\ c_{N-1} \end{pmatrix} \tag{2.44}$$

Or, simply written $y = \mathbf{V}c$, with $\mathbf{V}$ a Vandermonde matrix, c the vector of coefficients to determine, and y the measurements. $\mathbf{V}^{-1}$ can be calculated thanks to Lagrange basis polynomials.

This method trend to only work with relatively small amounts of points and low degree polynomials. Thus other related methods have been developed, such as spline interpolation, which represents the reconstruction as piecewise polynomial $p_k(t) = a_k + b_k(t - t_{k-1}) + \cdots + c_k(t - t_{k-1})^{n-1}$, $t \in [t_{k-1}; t_k]$ where t is an intrinsic parameter following the curve, and with conditions of continuity at the nodes.

Spline interpolation bridges the gap between the two methods of interpolation as they have a similar idea of processing locally the signal, while still relying on a Lagrange basis [37, 38]. For a degree 1 polynomial the result is identical to the linear interpolation,

and equivalences can be found in higher dimensions between splines interpolation and the aforementioned convolutive methods, as illustrated in the Wikipedia page up to the bicubic case [39].

In all cases, the user can choose a level of smoothness, ie. an underlying model, either by chosing the degree of the polynomial or by choosing the order and shape of the interpolant function h. This choice should be viewed as additional information added about the measurement setup and/or the measure itself. Indeed, nothing stops the user from always using a polynomial interpolation of degree $N-1$, however it would lack physical meaning, and exhibit issues related to overfitting such as the Runge's phenomenon. The only model that can be considered almost always true in images is the zero-order hold as it is based on the sensor shape.

Some of these methods can be extended to multidimensional data. These include nearest-neighbour and multilinear since their effects are separable, but also other spline interpolations, with bicubic interpolation being the highest degree commonly used.

Up to this point, only regular sampling was considered, however, in this thesis a particular interest was given to undersampling techniques requiring an irregular sampling in order to work outside the hypotheses of the sampling theorem and thus use less measurements. This added condition of being able to interpolate unstructured data complexifies the problem and limits the type of reconstruction algorithms available. Two methods will be presented here: Kriging and Radial Basis Functions interpolation.

Radial Basis Functions

Radial Basis Functions (RBF) interpolation is one of the main meshfree method working under the conditions of unstructured and multidimensional data. Its objective is to find x the ground truth from a set of measurements y at known positions u (i.e. $y_i = x(u_i)$) by solving a problem of the form:

$$y = \begin{pmatrix} A_{11} & A_{12} & \cdots & A_{1N} \\ A_{21} & A_{22} & \cdots & A_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ A_{N1} & A_{N2} & \cdots & A_{NN} \end{pmatrix} \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_N \end{pmatrix} \tag{2.45}$$

With y the measurement vector, w the vector of weights to be determined, and the matrix $\mathbf{A}$ containing elements A_{ij} that are radial basis function chosen appropriately. These RBF are a class of function that depend only on the distance between the points i and j of positions u_i and u_j, such that $A_{ij} = A(||u_i - u_j||_2)$. Most commonly used radials functions are for instance $A(r) = e^{-(\alpha r)^2}$ (Gaussian), $A(r) = r^k \ln(r)$ ($k = 2$ Thin plate spline), etc.

The weights w can be calculated by inverting $\mathbf{A}$, $w = \mathbf{A}^{-1}y$, as $\mathbf{A}$ can be shown to be non-singular as long as the measurement points are distinct [40]. It is then possible to calculate $x(u) = \sum_i^N w_i A(||u - u_i||)$ at any position u.

This method is a use-case of RBF neural network [41], as the system can be described as an input layer linked to an hidden layer containing the RBFs via the parameters u_i itself linked to the output layer via the weights w. It can also be applied to other problems, notably approximation instead of interpolation, where the positions of the centres of the RBFs have to be learned and are not necessarily fixed.

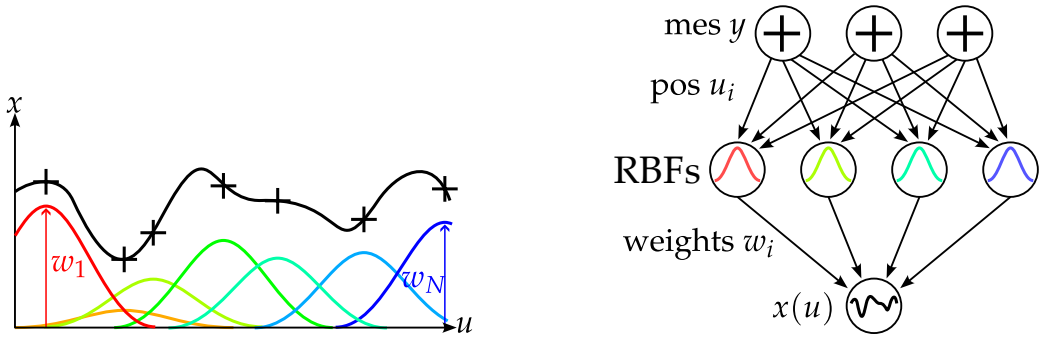


FIGURE 2.13: RBF schematic principle and neural network

The RBF interpolation can be regularised by adding a parameter λ such that $y = (\mathbf{A}+\lambda\mathbf{I})w$ [42, p.167]. This reconstruction relaxes the conditions by allowing the result to not strictly pass through the data points, it is therefore not an interpolation but instead an approximation. The resulting image is similar to the non regularised one smoothed via a convolution product. This RBF approximation has additional drawbacks by adding a bias and requiring the correct choice of regularisation parameter λ.

Kriging

Another popular interpolation method is Kriging, also called also called Gaussian process regression and Wiener Kolmogorov prediction. Kriging is also based on the same idea that any point can be deduced from the measurements and their distance to the point, via a weighted average [43].

Kriging seeks to solve a minimisation problem based on the variances. At a given position u_l, the estimation is $\hat{x}_l = \sum_i \omega_{li} y_i$, where ω are weights calculated to minimise $\sigma^2 = \sum_l (\hat{x}_l - x_l)^2$. This minimisation is equivalent to solve this system:

$$\begin{pmatrix} \gamma_{l1} \\ \gamma_{l2} \\ \vdots \\ \gamma_{lN} \end{pmatrix} = \begin{pmatrix} \gamma_{11} & \gamma_{12} & \cdots & \gamma_{1N} \\ \gamma_{21} & \gamma_{22} & \cdots & \gamma_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{N1} & \gamma_{N2} & \cdots & \gamma_{NN} \end{pmatrix} \begin{pmatrix} w_{j1} \\ w_{j2} \\ \vdots \\ w_{jN} \end{pmatrix} \tag{2.46}$$

Also written $g_l = \mathbf{\Gamma}\omega_l$. γ is a variogram function, the matrix is the variogram of the distances between the measurements u_{y_i} and u_{y_j}, and g_l is the variogram of the distance between the measurements and the point to be estimated at position u_l.

The variogram represents the dependence between points in function of their distance and can be defined as the variance of the difference between values at 2 locations. For a stationary and isotropic process, it only depends on the distance between the two points. $\gamma(||u_b - u_a||) = \frac{1}{2}\langle(y_a - y_b)^2\rangle$. This function is estimated from the known measurements by fitting a model on the experimental variogram.

This is called Simple Kriging in the litterature. Although it is the Best Linear Unbiased Estimator, it is in practice modified with additional parameters, such as in Ordinary Kriging to take into account a lack of knowledge of the global mean for instance.

Kriging and RBF interpolation were shown to be equivalent as long as the radial basis function is chosen to be identical to the variogram function [44]. As they are in practice chosen semi-arbitrarily and knowing that RBF is less costly in terms of compute time, RBF interpolation will be preferred in this thesis.

2.3.3 Compressive sensing

Hypotheses

As seen in the previous section, the sampling theorem states that there is a maximum frequency $f_{max} < \frac{f_s}{2}$ that signals must not exceed in order to correctly retrieve their spectral components. This limit is a direct consequence of the uniform sampling at frequency f_s. Thus one of the key to "avoid" (or at least postpone) this limit is to use irregular sampling schemes. However, simply removing samples removes information and can't bring better results, that's why a new paradigm called Compressive Sensing (CS) has been developed over the last 20-30 years [45–47].

CS is based on the general idea that most "natural" signals are very compressible [48, 49]. A compressible signal is such that it is possible to reduce the amount of data that needs to be stored without meaningfully altering its content. To do so, the signal is written in another basis such that it can be expressed as mostly zeros, with the exception of a few modes. This basis is signal-dependant, however, some transforms like Fourier and wavelets are applicable to most "natural" signals and thus are called universal bases. This general idea is behind most data compression techniques such as JPEG as it takes advantage first of the sparsity of images, but also of the efficiency and omnipresence of FFT related transforms. These bases are convenient due to their universality, however they can't compress the signal as efficiently as bases that are specific to the signal, sometimes called tailored bases, with the Singular Value Decomposition (SVD) derived basis being the most common example. If the problem is viewed in a more cryptographic way, the bases act as encryption/decryption keys, so a signal can be compressed using the FFT, sent, then decompressed using the iFFT since this basis is known, universal and already implemented, while if one uses the same idea with a SVD compression, it is not possible to decompress the data without prior knowledge of the basis vectors that are specific to the given signal.

CS is therefore a study of "natural" images only. There is no formal definition for this concept, besides may be equating it to compressible, as the term can be understood as being in opposition to pure noise. Natural images are a (very small) subset of all possible images, indeed if one generates all possible combination of pixel values for a given image size, most of the generated images will be made of pure noise as each pixel are uncorrelated, but a few generated images will represent "something". This is equivalent to Borel's savant monkey thought experiment.

While the sampling theorem will work for any signal whatever their spectral behaviour before a given frequency f_{max}, CS will only work if its spectrum is not dense, ie. sparse, if the Fourier basis is chosen.

Starting from the same under-determined under-sampling problem as seen in the interpolation section:

$$y = \mathbf{H}x \tag{2.47}$$

With x being a discretised but high resolution (N points) version of the signal and y low resolution measurements ($M < N$ points). H is then a $N \times M$ matrix representing the sampling scheme. As CS assumes the sparsity hypothesis, x can be written as:

$$x = \mathbf{\Psi} s \tag{2.48}$$

With $\mathbf{\Psi}$ the transform basis chosen ($N \times N$ matrix) and s the signal expressed in said basis containing only a small amount of non-zero points. Thus, the (still ill-posed) problem is:

$$y = \mathbf{H\Psi} s \tag{2.49}$$

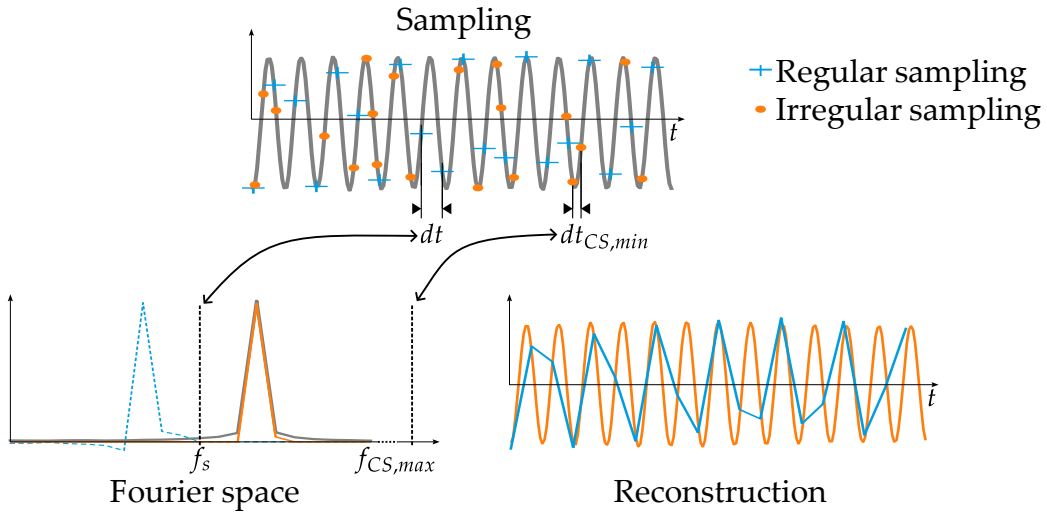


FIGURE 2.14: Principle of CS applied to a sinus signal and comparison with the classical case. (top) a sinus is sampled with a regular step and a random one, the signal is reconstructed using the classical method for regular samples and CS using the Fourier basis (bottom right), with the Fourier transform of all three signals (bottom left). As the regular sampling is too low for the signal, it aliases making the sinus frequency off, however with CS and the random sampling but the same total number of sample, the signal is reconstructed, and the theoretical, best case scenario maximum frequency limit is shifted to one inversely proportional to the minimum step size achieved between 2 measurements

How to promote sparsity

The objective is to find the most compressed form possible of the signal such that it solves Equation 2.49. Therefore, the "best" s to find is such that the signal is the sparsest, and to do so is to count the number of non-zeros in s. This process led to the introduction to the "ℓ_0-norm". A ℓ_p-norm is defined for $p > 0$ as:

$$||x||_p = \left(\sum_i |x_i|^p \right)^{\frac{1}{p}} \tag{2.50}$$

If 0^0 is defined as $0^0 \triangleq 0$, a similar sum would count the non-zero although it doesn't behaves homogeneously:

$$||x||_0 = \sum_i |x_i|^0 \tag{2.51}$$

Finding the sparsest s that corresponds to Equation 2.49 is the constrained problem:

$$\hat{s} = \min_s ||s||_0 \text{ such that } y = \mathbf{H}\mathbf{\Psi}s \tag{2.52}$$

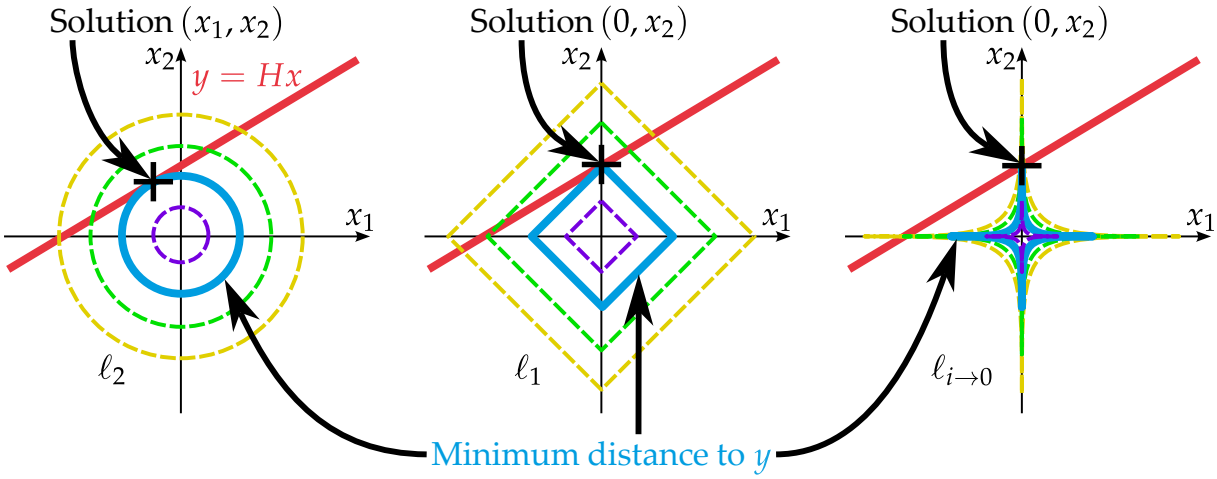


FIGURE 2.15: Comparison of ℓ_2, ℓ_1 and ℓ_0 distances

This minimisation is very difficult to compute efficiently. Thus a commonly used idea is to solve the closest convex norm minimisation. This closest norm is ℓ_1, the Manhattan norm, see Figure 2.15. This change can be made on two conditions: First, that $\mathbf{H}$ and $\mathbf{\Psi}$ are incoherent meaning that the rows of $\mathbf{H}$ are not correlated with the columns of $\mathbf{\Psi}$. Second that the number of measurement is sufficient, for a sparse signal made of K non-zero modes, $m > k_1 K \log(\frac{n}{K})$, with k_1 a factor representing the incoherence between $\mathbf{H}$ and $\mathbf{\Psi}$. If these conditions are met, Equation 2.52 becomes:

$$\hat{s} = \min_s ||s||_1 \text{ such that } y = \mathbf{H}\mathbf{\Psi}s \tag{2.53}$$

The presence of noise complicates the problem, as the sparsity hypothesis doesn't work on noise, thus related variants of Equation 2.53 can be more robust. These variants are either in the form of a constrained problem, or in a convex optimisation problem directly derived from Equation 2.24, where the regularising function is defined as the ℓ_1-norm:

$$\hat{s} = \min_s ||\mathbf{H}\mathbf{\Psi}s - y||_2 + \mu||s||_1 \tag{2.54}$$

Applications

CS has had many applications over the last 20 years in image and signal processing, with both objectives of either reducing the number of acquisitions or compressing the data for storage and transmission. [50] gives an overview of all the fields in which it has been applied. Outside of imaging, CS has been used with a wide variety of signals, from electrocardiograms to the study of DNA and genomes in biology, in communications for maximising data transfer in multiple-input multiple-output channels, in acoustics for sound localisation, monitoring and reconstruction. It has also been used in the context of pattern and object recognitions in images.

Focusing on imaging in particular, a few major points have garnered a lot of attention, starting with Single Pixel Camera (SPC). SPC were developed by researchers at Rice University [51], soon after CS was first described. The principle is described in Figure 2.16. The objective of the setup is to turn a single pixel, ie. a photodiode into a camera by structuring the light received by it in a known manner. To do so, first the light from a given object is gathered thanks to a lens in order to image it on a known binary pattern, this ensemble is then focused on a photodiode that gather the average intensity associated to the given pattern, then this process is repeated for a different binary pattern. By knowing enough sets of patterns and their associated average intensity, it is then possible to reconstruct the image thanks to CS. From an instrumental point of view, the main difference from any other imaging setup is the presence of a set of known random patterns. Multiple methods do exist in order to create them, with the most common one, as shown in Figure 2.16, being the Digital Micromirrors Device (DMD). A generalisation of this method is shown in chapter 3.

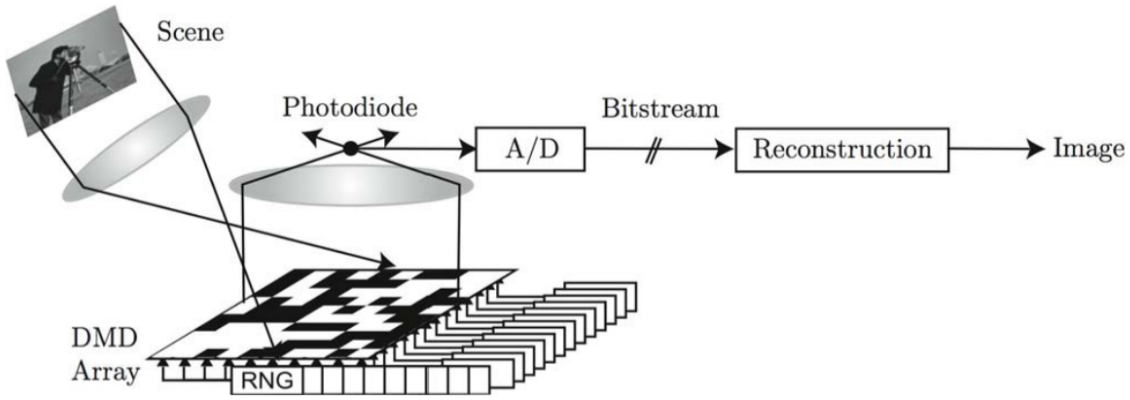


FIGURE 2.16: Single Pixel Camera operating principle as described in [52]

CS is also well suited to reduce the amount of samples required in point-by-point or line-by-line sweeping techniques that exist in many contexts, from radar [53] to STEDM [54] and AFM [55]. This is further developed in thermography for the Flying Spot technique and Scanning Photothermal Radiometry in the following chapters.

CS has also been used extensively for processing 3D data, sometimes called Snapshot Compressive Imaging [56], especially in the context of tomography. Sparse acquisitions with MRI [57] were among the first applications, the idea being to reduce the number of acquisitions that are made in the k-space (Spatial frequencies space) in

MRI. Similarly in electron tomography for instance [58] it can be used to reduce the number of acquisitions that are made in the Radon space, by skipping the projection of some angles both to make a faster acquisition but also to overcome instrumental issues when some angles are inaccessible. In tomographic reconstructions, Total Variation has imposed itself as the main sparsity inducing transform, as the objectives are to detect edges between different materials while removing line artefacts. Reconstructions of videos also has been achieved [59] by separating the problem into small blocks.

As a last note, in the literature CS based techniques and SIM are often presented separately, however the difference between them can sometimes be arbitrary. CS imaging is almost always in what can be described as structured detection [60], in the sense that the modification of the standard setup to structure the light is done at the output, as close as possible to the detector. This is both to minimise the amount of blurring of the patterns and to not have to modify anything upstream thus keeping a similar behaviour as a camera. On the other hand, SIM is always structured illumination, ie. at the input, however, it is possible to apply the CS methods there as well [61].

Algorithms used in this thesis

A few algorithms were used in this thesis to solve ℓ_1-norm problems as in the CS and SIM cases. The objective here is to show the general philosophy of each algorithm and to give a brief explanation of the steps taken. As for most algorithms solving ill-posed problems, some choices are made that can be (and have been) argued for and against, such as for instance the step size between iterations, and won't be thoroughly explored here for clarity.

Fast Iterative Schrikage Threshold Algorithm (FISTA) [62] is an accelerated variant of Iterative Schrikage Threshold Algorithm (ISTA). It seeks a solution for:

$$\min_x ||\mathbf{H}x - y||^2 + \mu||x||_1 \tag{2.55}$$

This is achieved by solving the proximal subproblem [63], which in this case is iterating a form of gradient descent and thresholding it. more precisely with the following formula:

$$x_{k+1} = \mathcal{T}_{\mu t}(x_k - 2t\mathbf{H}^T(\mathbf{H}x_k - y)) \tag{2.56}$$

with t the stepsize defined as the inverse of the square of the maximum singular value of $\mathbf{H}$, and the shrinkage operator (Sometimes also called soft-threshold) defined for each point as:

$$\mathcal{T}_\alpha(w)_i = \text{sgn}(w_i)(|w_i| - \alpha)_+ = \begin{cases} w_i - \alpha & \text{if } w_i > \alpha \\ 0 & \text{if } -\alpha < w_i < \alpha \\ w_i + \alpha & \text{if } w_i < -\alpha \end{cases} \tag{2.57}$$

FISTA takes this same principle and optimises the number of iterations required in order to solve this problem, by using a Nesterov acceleration step in which the iterative shrinkage isn't made directly on x, but rather an auxiliary variable b that is a combination of the previous two points.

Algorithm 1 FISTA

Require: x_0, $y_1 = x_0$, $t_1 = 1$, $\alpha > 0$ ▷ α constant stepsize

1: **for** $k \in [1, N]$ **do**

2: $\quad x_k \leftarrow \mathcal{T}_{\mu\alpha}(b_k - 2\alpha \mathbf{H}^T(\mathbf{H}b_k - y))$

3: $\quad t_{k+1} \leftarrow \frac{1+\sqrt{1+4t_k^2}}{2}$

4: $\quad b_{k+1} \leftarrow x_k + \frac{t_k - 1}{t_{k+1}}(x_k - x_{k-1})$

Orthant-Wise Limited-memory Quasi-Newton Algorithm (OWL-QN) [64] minimises functions of the form:

$$f(x) = g(x) + C||x||_1 \tag{2.58}$$

for some constant C. It is built from Limited-memory Broyden Fletcher Goldfarb Shanno Algorithm (L-BFGS) which has for main objective to solve a quasi-Newton method of the form $x_{k+1} = x_k - \alpha_k \mathbf{H_f}_k \nabla f(x_k)$ by proposing an efficient calculation of the Hessian matrix $\mathbf{H_f}$ that is of the form [65]:

$$\mathbf{H_f}_{k+1} = (\mathbf{I} - \varsigma_k s_k q_k^T)\mathbf{H_f}_k(\mathbf{I} - \varsigma_k q_k s_k^T) + \varsigma_k s_k s_k^T \tag{2.59}$$

with $s_k = x_{k+1} - x_k$, $q_k = \nabla f(x_{k+1}) - \nabla f(x_k)$, $\varsigma_k = \frac{1}{q_k^T s_k}$. In practice LBFGS calculates $\mathbf{H_f}_k \nabla f(x_k)$ by only storing a few of the last vectors $\{s_i, q_i\}_{i=k-m}^{k-1}$ with m an allocated memory budget and by assuming that $\mathbf{H_f}_0$ is diagonal. The iterative process only requires products of vectors and one diagonal matrix and can be described as such, with $\mathbf{V}_n = (\mathbf{I} - \varsigma_n s_n q_n^T)$:

$$\begin{aligned}
\mathbf{V}_k \nabla f & \qquad \text{First for loop} \\
\mathbf{V}_{k-1}\mathbf{V}_k \nabla f & \\
\mathbf{V}_{k-m} \cdots \mathbf{V}_{k-1}\mathbf{V}_k \nabla f & \\
\mathbf{H_f}_0 \mathbf{V}_{k-m} \cdots \mathbf{V}_{k-1}\mathbf{V}_k \nabla f & \qquad \text{Mult by diag} \\
\mathbf{V}_{k-m}^T \mathbf{H_f}_0 \mathbf{V}_{k-m} \cdots \mathbf{V}_{k-1}\mathbf{V}_k \nabla f & \qquad \text{2nd loop} \\
\mathbf{V}_k^T \cdots \mathbf{V}_{k-m}^T \mathbf{H_f}_0 \mathbf{V}_{k-m} \cdots \mathbf{V}_{k-1}\mathbf{V}_k \nabla f &
\end{aligned}$$

The rest of the OWL-QN algorithm deviates a little further from the standard LBFGS procedure. The general idea is that when the ℓ_1 norm is restricted to a given orthant, it is linear and thus differentiable. An orthant being a subset of space in which the coordinates never change sign. For instance in $\mathbb{R}^2$, the orthants are 4 quadrants separated by the axes. The idea is then to find an optimum while staying within said linearly behaving orthant.

A pseudo gradient function is defined in order to figure out which orthant to explore:

$$\diamond_i f(x) = \begin{cases} \partial_i^- f(x) & \text{if } \partial_i^- f(x) > 0 \\ \partial_i^+ f(x) & \text{if } \partial_i^+ f(x) < 0 \\ 0 & \text{else} \end{cases} \tag{2.60}$$

with $\diamond$ a symbol representing the pseudo-gradient by analogy with ∇, $\partial_i^{\pm} f(x)$ directional (left or right) partial derivatives such that

$$\partial_i^{\pm} f(x) = \frac{\partial g(x)}{\partial x_i} + \begin{cases} \operatorname{sgn}(x_i)C & \text{if } x_i \neq 0 \\ \pm C & \text{if } x_i = 0 \end{cases}$$

The calculation of this pseudo-gradient allows to check that the current iteration is in the orthant of interest and that the gradient leads to it, corresponding to Equation 2.62. This verification can be made by defining $\pi(a, b)$ a function that checks if a and b are of equal sign and is the projection of a within the orthant of b:

$$\pi_i(a, b) = \begin{cases} a_i & \text{if } \operatorname{sgn}(a_i) = \operatorname{sgn}(b_i) \\ 0 & \text{else} \end{cases} \tag{2.61}$$

$$\xi_i^k = \begin{cases} \operatorname{sgn}(x_i^k) & \text{if } x_i^k \neq 0 \\ \operatorname{sgn}(-\diamond_i f(x^k)) & \text{if } x_i^k = 0 \end{cases} \tag{2.62}$$

Algorithm 2 OWL-QN

Require: x^0, $S \leftarrow []$, $Q \leftarrow []$, m, N

for $k \in [1, N]$ **do**
  $v^k \leftarrow -\diamond_i f(x)$ from Equation 2.60
  for $i \in [k-1, m]$ **do**
    $\gamma_i \leftarrow \varsigma_i s_i^T v^k$
    $v^k \leftarrow v^k - \gamma_i q_i$ ▷ $v^k - \varsigma_i s_i^T v^k q_i = (I - \varsigma_i q_i s_i^T) v^k$
  $d^k \leftarrow \mathbf{H_f}_k^0 v^k$, with $\mathbf{H_f}_k^0 = \frac{s_{k-1}^T q_{k-1}}{q_{k-1}^T q_{k-1}} I$
  for $i \in [k-m, k-1]$ **do**
    $\beta \leftarrow \varsigma_i q_i^T d^k$
    $d^k \leftarrow d^k - s_i(\gamma_i - \beta)$ ▷ $= (I - \varsigma_i s_i q_i^T) d^k + \varsigma_i \gamma_i$
  $p^k \leftarrow \pi(d^k; v^k)$
  $x^{k+1} \leftarrow \pi(x^k + \alpha p^k; \xi^k)$ ▷ α *stepsize,* ξ^k *implicit*
  $s^k \leftarrow x^{k+1} - x^k$ and discard s^{k-m} ▷ s^k k^{th} *parameter of* S
  $q^k \leftarrow \nabla g(x^{k+1}) - \nabla g(x^k)$ and discard q^{k-m} ▷ q^k k^{th} *parameter of* Q

The algorithm then searches like a quasi-Newton method, except only within that orthant, x is iterated as the following line search:

$$x^{k+1} = \pi(x^k + \alpha p^k; \xi^k) \tag{2.63}$$

with $p^k = \pi(\mathbf{H_f}_k v^k; v^k)$ and $v^k = -\diamond f(x^k)$.

Total Variation ALternating Lagrangian Algorithm (TVAL3) [66] (Total Variation

based on Augmented Lagrangian and Alternating directions Algorithm) is an algorithm that seeks to solve a Total Variation Regularisation problem of the following form:

$$\min_x \sum_i ||D_i x|| \text{ s.t. } \mathbf{H}x = y \tag{2.64}$$

with $D_i x$ the gradient of x at the i^{th} pixel. The algorithm relies on the augmented Lagrangian method, which allows to transform a constrained problem into non-constrained ones by adding a quadratic penalty. Without the quadratic penalty, a Lagrange multipliers method is related to Tikhonov regularisation due to the fact that it can also be written as a constrained problem.

The augmented Lagrangian method states an equivalent problem to solve:

$$\min_{w_i,x} \mathcal{L}_A = \min_{w_i,x} \sum_i (||w_i|| - \nu_i^T(D_i x - w_i) + \frac{\beta_i}{2}||D_i x - w_i||_2^2) - \lambda^T(\mathbf{H}x - y) + \frac{\mu}{2}||\mathbf{H}x - y||_2^2 \tag{2.65}$$

with β_i, μ penalty constraints, and ν_i, λ the Lagrange multipliers. This allows for the separation into two sub-problems, a minimisation on x and one on w_i.

For the sub-problem on w_i:

$$w_{i,k+1} \text{ s.t. } \min_{w_i} \mathcal{L}_A(w_i, x_k) \tag{2.66}$$

with $\min_{w_i} \mathcal{L}_A(w_i, x_k)$ equivalent to:

$$\min_{w_i} \sum_i (||w_i|| - \nu_i^T(D_i x_k - w_i) + \frac{\beta_i}{2}||D_i x_k - w_i||_2^2) \tag{2.67}$$

It can be proved [66, p.23] that there is an analytical solution for this problem, also called shrinkage-like which is for the ℓ_1-norm:

$$w_{i,k+1} = \max\left\{||D_i x_k - \frac{\nu_i}{\beta_i}||_j - \frac{1}{\beta_i}; 0\right\} \text{sgn}(D_i x_k - \frac{\nu_i}{\beta_i}) \tag{2.68}$$

Then, once $w_{i,k+1}$ are known, the sub-problem x can be solved:

$$x_{k+1} \text{ s.t. } \min_x \mathcal{L}_A(w_{i,k+1}, x) \tag{2.69}$$

where $\mathcal{L}_A(w_{i,k+1}, x)$ is a quadratic function with regards to x. Its gradient is:

$$d_k(x) = \sum_i (\beta_i D_i^T(-D_i x - w_{i,k+1}) - D_i^T \nu_i) + \mu \mathbf{H}^T(\mathbf{H}x - y) - \mathbf{H}^T\lambda \tag{2.70}$$

It could be solved in many ways, however for the sake of efficiency, a single-step gradient descent is used at every iteration k. This is done under the assumption that since there are two sub problems iterated alternatively, there should be no need for the accuracy given by either a Moore-Penrose pseudo inverse or a more fleshed out steepest

gradient method. x_{k+1} is thus calculated at every step:

$$x_{k+1} = x_k - \alpha_k d_k(x_k) \tag{2.71}$$

with α_k defined from the previous iterations $\alpha_k = \frac{(x_k - x_{k-1})^T (x_k - x_{k-1})}{(x_k - x_{k-1})^T (d_k(x_k) - d_k(x_{k-1}))}$, and then validated by a line search algorithm checking that $\mathcal{L}_A(w_{i,k+1}, x_k - \alpha_k d_k(x_k))$ is small enough.

Algorithm 3 TVAL3

Require: $v_i^0, \beta_i^0, \lambda^0, \mu^0$ and init points $w_i^0 \; \forall \; i = 1 \dots N$, x^0, $d_k(x_0)$

while $|x^{k+1} - x^k|_2 > tol$ **do**
  $w_{i,0}^{k+1} \leftarrow w_i^k \; ; x_0^{k+1} \leftarrow x^k$
  while $|x_{l+1} - x_l|_2 > tol$ **do**
    $w_{i,l+1} \leftarrow$ Equation 2.68
    $d_l(x_l) \leftarrow$ Equation 2.70
    $\alpha_l \leftarrow \frac{(x_l - x_{l-1})^T (x_l - x_{l-1})}{(x_l - x_{l-1})^T (d_l(x_l) - d_l(x_{l-1}))}$
    $x_{l+1} \leftarrow x_l - \alpha_l d_l(x_l)$
  $v_i^{k+1} \leftarrow v_i^k - \beta_i (D_i x - w_i) \; ; \lambda^{k+1} = \lambda^k - \mu(\mathbf{H}x - y)$
  $\beta_i^{k+1} \geq \beta_i^k \; ; \mu_i^{k+1} \geq \mu_i^k$

2.3.4 Specificities of thermal imaging in super resolution

Super resolution in infrared imaging has been studied mostly through the lens of reducing costs of IR cameras. From the start of SPC, its implementation in IR was a goal [51, 66], and continued with the use of cameras instead of monodetectors [67, 68]. This technique allowed for imaging in a very wide spectral range, including with THz wavelengths [69].

Often, these methods are used in the IR just like in visible imaging without taking into account the specificities of measurements at these wavelength especially regarding the thermal diffusion. This can be seen in a recent review on the subject of IR super-resolution [70], the field is dominated by a more generalist computer vision approach. This thermal diffusion makes the commonly made assumptions in image processing no longer true. These assumptions usually are that there is a clear separation between high and low frequencies components, where the high frequencies are sharp edges and the low frequencies are smooth areas and relatively homogeneous textures, which for instance made de-blurring and other reconstructions based on Total Variation minimisation (That is minimisation of the sum of the gradient) extremely prevalent. The thermal diffusion lead to these high-frequencies component to be shifted down, blurring the edges.

This raises a very fundamental question on the objectives of the measurement. Two main schools of thought emerges from this. The first one wants to map the thermal radiation or other properties depending from it, usually the temperature and with the emissivity being the main obstacle. The second one isn't interested as much in the radiation itself but instead on what caused it, the thermal sources. This second approach adds more obstacles on top of the first ones by having to revert the effect of thermal

diffusion, although since the diffusion is the dominant effect, it is sometimes possible to neglect the diffraction problem.

In order to describe this diffusion problem, it is possible to define a thermal impulse response, sometimes called Thermal PSF [71, p. 29]. This total impulse response is then:

$$h_{heat}(x,y,t) = G_{heat}(x,y,t) \otimes_{(x,y)} h(x,y) \otimes_t h_t(t) \tag{2.72}$$

where $\otimes_{(x,y)}$ is the spatial convolution, $h(x,y)$ the previously defined optical PSF, $\otimes_t$ is the temporal convolution, $h_t(t)$ is the temporal impulse response of the illumination which is separable from its spatial component (ie. the optical PSF), and most importantly $G_{heat}(x,y,t)$ is the Green function associated with the heat conduction equation. $h_t(t)$ can be omitted for most applications since a light pulse is instantaneous, but might have to be considered when measuring it since in practice the light pulse is made by a source requiring a voltage that itself isn't instantaneous.

The Green function is a generalisation of the impulse response, and gives the temperature at any point and time by a convolution between it and the boundary conditions. It is the solution of the heat equation for pulse in both space and time, for any shift $\vec{s}$ in space and τ in time:

$$\frac{1}{\alpha}\frac{\partial G(\vec{x}-\vec{s},t-\tau)}{\partial t}\frac{\partial^2 G(\vec{x}-\vec{s},t-\tau)}{\partial \vec{x}^2} = \delta(\vec{x}-\vec{s})\delta(t-\tau) \tag{2.73}$$

And is of the form:

$$G_{heat}(\vec{x},t) = \frac{1}{(4\pi\alpha t)^{\frac{3}{2}}}\exp\left(\frac{||\vec{x}||^2}{4\alpha t}\right) \tag{2.74}$$

The main issue with this function is that it is time dependant, thus adding an additional dimension. Yet, this issue is also its main benefit as this time is associated to a diffusion length $\delta_p = \sqrt{\alpha t}$ for a semi infinite homogeneous medium of thermal diffusivity α. Although it could be theoretically possible to directly solve the 3D inverse problem, it is in practice extremely prohibitive in therms of computing cost, the preferred option is to evaluate this impulse response at a given depth, and it's associated time delay, which allows for its localisation on the 2D plane by a less costly inversion. This general approach of reconstruction of in depth sources led to interesting results [72, 73], notably with methods using virtual waves [74, 75].

For point-by-point methods, in [76], super-resolution was achieved in the context of Flying Spot, which is a thermal scanning method that is described in more detail in the next chapter. This study was done with both continuous and pulsed mode. Fundamentally, the gain in resolution stems from the fact that it is possible to move the laser spots with better precision than the pixel size, or for the continuous case at speeds that are slower than the acquisition frequency of the camera. This is fundamentally a form of geometrical super-resolution. Morelet transforms are used in order to filter out contributions of the thermal responses in order to get the local properties of interest in SR while filtering out the global properties of the material as well as the signal related directly to the laser pulse. In [77], CS was used with Fourier Domain Thermoreflectance (FDTR) measurements in a similar method to what is described more detail in chapter 4 and 5 in order to create cartographies of thermal conductances and conductivities using around 20% of the samples relative to a regular sampling.

2.4 Conclusion

In this chapter, the causes of the resolution limit were explained and are separable in two groups that are usually studied separately in the litterature.

On one hand is the diffraction limit caused by the optics. This limit is $d = \frac{\lambda}{2NA}$ in the best case scenario of two high-contrast points. In visible light microscopy, especially in the context of biological applications this problem has led to a few groundbreaking methods, both in scanning and widefield, notably STEDM, STORM and PALM.

On the other hand the sampling limit is caused by the digitalisation of the signal, is better understood in terms of frequency as it states that there is a maximum spatial frequency that has to be respected in order to correctly recover a signal. This limit is rarely a problem in visible microscopy due to the good quality of cameras at those wavelengths. Instead this problem was studied in fields such as geostatistics for instance where reconstructions from irregularly undersampled images led to interpolations methods such as Kriging. In image processing and computer science, this problem is seen under the lens of data compression, that has led to the development of Compressive Sensing that allows for the reconstruction of images by compressing the data in a few modes.

The field of super-resolution is wider than what is described here. Other methods beside those presented exist [78], most notably with machine and dictionary learning, see for instance [79, 80] among many others, that weren't considered here due to the fact that they require a big training dataset. When these do exist, these methods are indeed well proven, however here a set that is relevant not just for infrared images but for thermal measurements would be required, which is a considerable endeavour, since image processing involves massive training datasets, preferably unsupervised. Additionally, other fields are interested by some form of super-resolution, notably for semi-conductors, where Inverse Lithography Technology is very much confronted to the same diffraction problem, although in this context the problem is taken upside down, that is the source has to be calculated such that the printed pattern is optimum[12].

An important point to note is that the super resolution methods are all a result of two practical compromises: The first is between spatial and temporal resolution as the gain of resolution is often at the expense of more measurements taking more time. It is therefore difficult to study images outside their steady state, either because the measurement are made point by point, or because the reconstruction is the result of a series of images. The second compromise is between acquisition time and both computing time and a priori information. What is gained in compressing the acquisition time usually requires a combination of knowing the impulse response or a training dataset, and a long processing time, but both of these steps can be made offline respectively before and after the measurement.

[12] Meaning that a graphics card calculates the optimum way to produce a itself

Chapter 3
Sampling and Diffraction limits in IR Thermography: First attempts

In this chapter, the problem of super-resolution is explored with three approaches in the three main sections that follow both chronological order and levels of complexity. In the first section that was presented in the Société Française de Thermique (SFT) 2024 national conference, a technique of de-aliasing of relatively simple images is explored. In the second section, structured illumination is used in order to find if the method can be relevant in our context. In the last section, compressive sensing is used in order to increase the effective amount of pixels in the images. These last two sections are still work in progress and weren't published.

Contents

3.1 Spectral unfolding . . . 44
3.1.1 Setup characterisation . . . 45
3.1.2 Spectral unfolding . . . 45
3.2 Structured illumination (Result/Limitations) . . . 48
3.2.1 Method . . . 48
3.2.2 Numerical results . . . 49
3.2.3 Experimental limitations . . . 51
3.2.4 Perspectives . . . 52
3.3 Compressive sensing with cameras and an intermediary spatial filter . . . 54
3.3.1 Principles . . . 54
3.3.2 Measurement of the transfer matrix C . . . 56
3.3.3 Hadamard sampling . . . 57
Numerical results . . . 59
3.3.4 Random sampling . . . 60
3.3.5 Experimental results using DMD . . . 61
3.3.6 Experimental results using photolithography masks . . . 64
3.3.7 Conclusion and Prespectives . . . 66

3.1 Spectral unfolding

The first experimental case of super-resolution that was studied during the thesis was an unfolding the spectral signal of a structurally simple image. The objective was to characterise a heating wire line within a 2D image acquired with an existing thermo-transmittance setup.

Thermotransmittance in semi-transparent materials is a field of research within the ICT team [81–83]. The experimental setup is shown in Figure 3.1. Part of the IR light source is transmitted through the sample, it is detected by an IR camera with a $\times 2$ objective, and effective pixel size of $15\,\mu\mathrm{m}$. The measurement contains the transmitted signal and its spontaneous emission. The transmitted signal is extracted thanks to a synchronisation between the camera and an optical chopper placed at the source. The frequency of the chopper f_C is such that it is twice as slow as the image acquisition. It is then possible to calculate the difference between the image with the signal and the one with the spontaneous emission only if variations of the signal happens at a comparatively low frequency compared to the chopper. The sample contains a gold wire of nominal width $30\,\mu\mathrm{m}$ in which a periodic current is applied at frequency f_I, this current induces thermal variations on the sample at frequency $f_T = 2f_I$. This in turn causes a variation in transmittance.

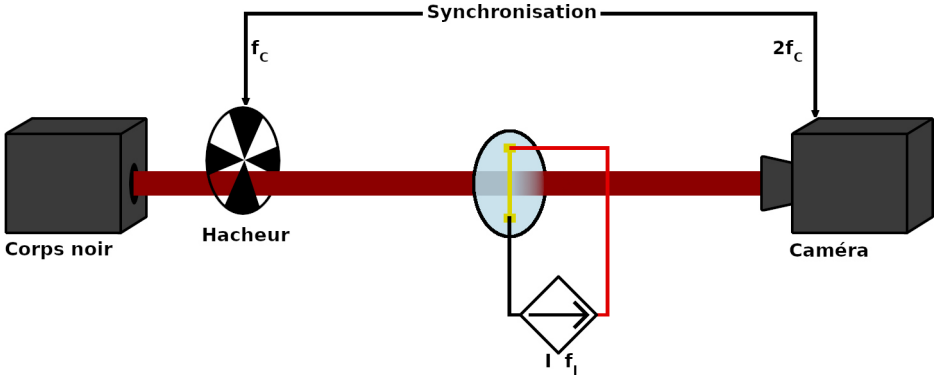


FIGURE 3.1: Setup

$$\frac{\Delta\Gamma}{\Gamma_0} = \kappa\Delta T \tag{3.1}$$

With Γ_0 the transmittance at ambient temperature and κ is the thermotransmittance coefficient of the material at the measurement wavelength ($\lambda = 3.3\,\mu\mathrm{m}$, $\kappa(\lambda) = 5.2 \times 10^{-4}\,\mathrm{K}^{-1}$) for thin materials [84].

An entire dataset contains the aforementioned pair of images extracted regularly from the overall signal ("video" feed) at frequency f_T, then a temporal Fourier transform is made on each pixel in order to extract the image at f_T in order to remove as much noise as possible.

In previous works in the lab, it was shown that the measured phenomenon is well described by a theoretical model with the exception of the close proximity to the thermal source. The objective was thus to attempt to describe the behaviour of the signal near and on the wire.

3.1.1 Setup characterisation

The first requirement was to characterise the setup. To do so the MTF was measured. The simplest method consisting of directly measuring the response to a spatial Dirac is difficult to implement accurately because it would require a strong enough localised source. Instead, the MTF was drawn from the frequency responses of known periodic signals. In order to do so, a US Air Force 1951 Calibration Target (USAF) was used and placed instead of the sample [85]. It contains series of short square wave signals each made of 3 metallic rectangles along both axes of increasing frequencies. Each frequency in USAF 51 is defined by a group number g on top, and a number within that group n on the side such that $\nu_{USAF} = 2^{g+(n-1)/6}$. For each ν_{USAF} along a direction the visibility of the signal was measured in order to obtain the response function to square waves, sometimes called Contrast Transfer Function (CTF) $CTF(\nu) = \frac{I_{max}(\nu)-I_{min}(\nu)}{I_{max}(\nu)+I_{min}(\nu)}$. The 1D MTF can be deduced from the CTF according to Coltman's formula [86]:

$$MTF(\nu) = \frac{\pi}{4}\left(CTF(\nu) + \frac{CTF(3\nu)}{3} - \frac{CTF(5\nu)}{5} + \cdots\right) \tag{3.2}$$

Assuming no aberrations for the objective, the 2D MTF is then rotationally symmetric. Otherwise the measurement would need to be repeated along multiple axes. The PSF is then calculated via a 2D iFFT. However this method is only accurate if the sampling limit is respected, which isn't the case here. In this section, it is shown that this issue then makes it impossible to accurately use deconvolutions.

3.1.2 Spectral unfolding

The image of the object's transmittance at f_T is shown in Figure 3.2. The variation of transmittance near the maximum exhibits two lines of different intensities and width. The goal here is to find what corresponds to the wire in order to know its exact position.

A Fourier transform analysis of this image in Figure 3.3 exhibits clearly some aliasing. Since the object has an axis of symmetry along the wire, its Fourier Transform signal is contained within a line represented perpendicularly in the reciprocal space. Due to the fact that the symmetry axis is not aligned with the pixel grid, the aliasing present in the FFT appears as lines parallel to the main part of the spectrum. Thus it is possible to numerically replace, ie. unfold the signal, as done in Figure 3.3. This is done in two steps, in Figure 3.3b, a binary mask is created filtering out the low frequencies as well as the frequencies outside the line going through 0 and perpendicular to the wire except those in the two parallel lines to the main one that are corresponding to the aliasing, then in Figure 3.3c these two lines are shifted laterally to the other side of the image, where by construction they end up in the continuity of the main line. The 256×320 spectrum becomes 256×460 without changing the frequency step $\Delta\nu = \frac{1}{\Delta u N} = 209\,\mathrm{m}^{-1}$. This induces that after iFFT, the spatial resolution along the horizontal axis is $\Delta u = 10.4\,\mu\mathrm{m}$ instead of $15\,\mu\mathrm{m}$. The reconstructed "pixel" is thus rectangular.

The filtered reconstruction in Figure 3.4 shows the high-frequencies components of the original image. These high frequencies correspond to the wire itself as the change of material from the semi-transparent borofloat to the gold wire is a sharp, discontinuous

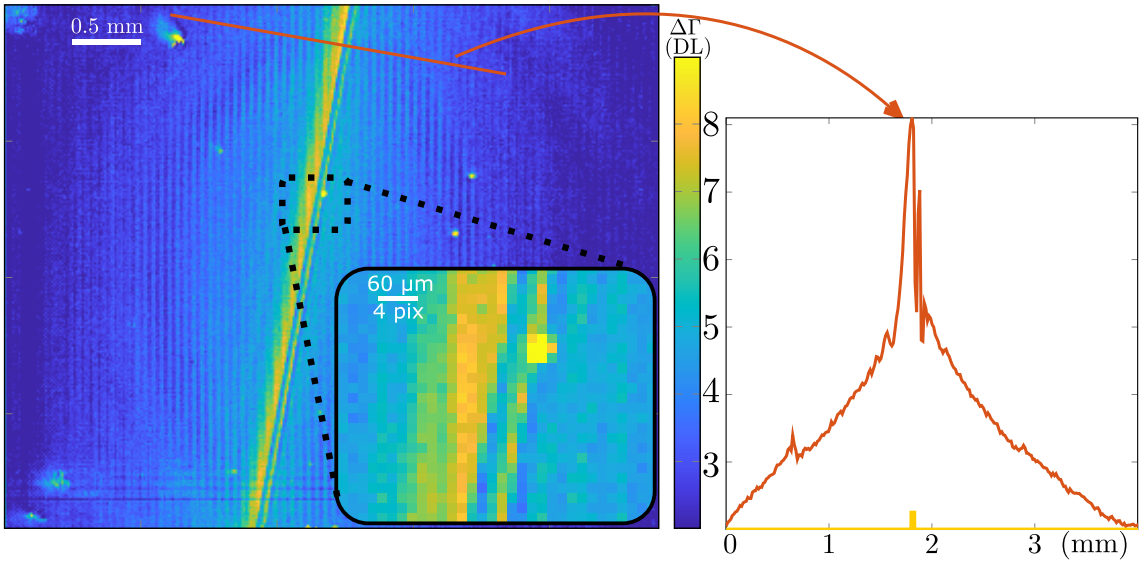
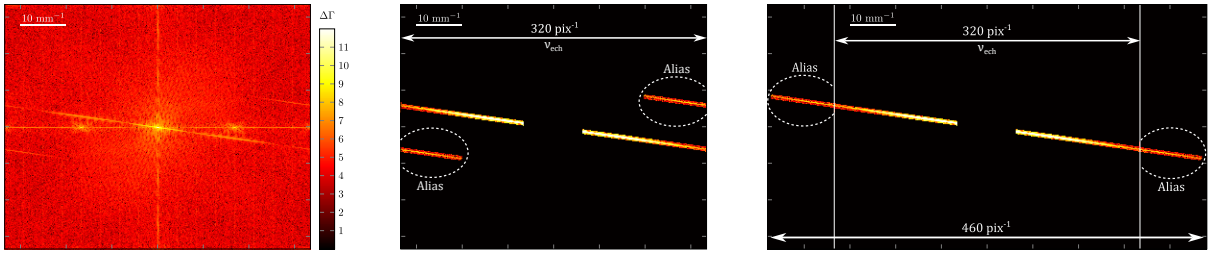


FIGURE 3.2: (left) Image acquired via thermotransmittance of the sample consisting of a thin $30\,\mu m$ gold wire emitting a periodic thermal signal demodulated at the temporal frequency of interest. (right) Profile drawn perpendicular to the gold wire.

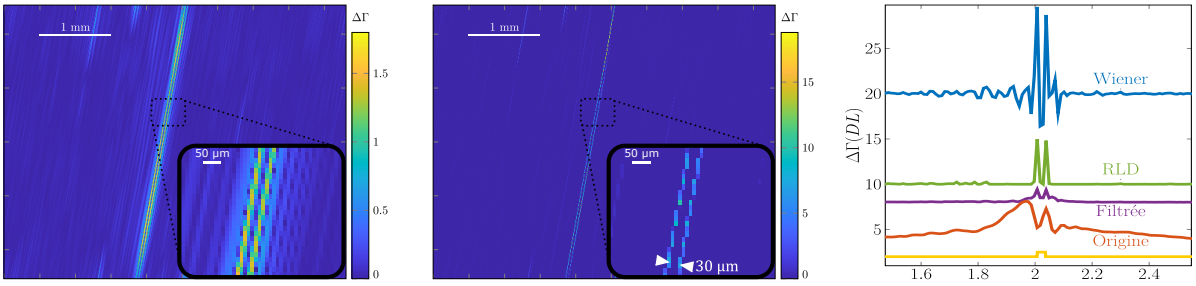
step in a thermal signal that is otherwise low frequency. By tracing the profile perpendicularly to the wire, the signal is made of two peaks, separated by roughly $30\,\mu m$ which corresponds to the known size of the wire.

This study helped to better understand some issues related to the thermotransmittance measurements made previously in the lab. First, it showed that the wire shows up in the measurement as its two edges separated by a slump of $30\,\mu m$. Since the pixel is of $15\,\mu m$, this effect is averaged out, and values at these positions cannot be used since the gold has a completely different κ. In these measurements, the maximum is thus slightly offset from the wire (The wide peak in Figure 3.2 instead of the thin one), which implies some slight misalignment. Secondly, the FFT exhibits and allows to filter a noise that is present in the acquired signal at a horizontal frequency that is exactly a quarter of the spatial frequency, as can be seen in Figure 3.3a, as well as in the form of vertical lines in Figure 3.2. The source of this defect is still uncertain but it appears after the demodulation step (FFT in time for each pixel) thus a possibility is that it is an artefact linked to the Charge Coupled Device (CCD) matrix and its operating frequency, with each column being read at a given speed such that the demodulation extract one of its harmonics. Finally, the reconstruction could be improved with better choices of filters in both spaces. In the reciprocal space, the filter chosen is a binary mask, ie. a boxcar function, which leads inevitably to some artefacts after the iFFT creating sinus cardinal like secondary lobes, however a smoother filter would inevitably include more noise. In the real space, the FFT was made on the image itself. A better scenario for its computation would be to make the image periodic, as discussed and illustrated in Figure 2.12. This is done usually by applying a window (such as Hann or Hamming) which would lead to a more accurate FFT by attenuating the two centred horizontal and vertical lines in Figure 3.3 in exchange of the removal of some of the information from the original image.



(A) FFT (B) Filtered FFT (C) Unaliased FFT

FIGURE 3.3: (a) Image in reciprocal space of the logarithm of the Fourier transform calculated via FFT, including aliasing, noise at $\nu_x/4$, and artefacts related to the non-periodicity of the input. (b) Image of the filtered high frequency signal associated to the wire. (c) Image of the de-aliased result.



(A) iFFT of Figure 3.3c (B) Richardson deconvolution (C) Transmittance variation ΔΓ

FIGURE 3.4: Result

This method is basic but has the advantages of being simple computationally and can work on some manufactured material such as printed circuit boards due to their geometries. It can viewed as conceptually similar as SIM, since the camera sampling acts as a frequency shift creating a moiré pattern, and the limitation stems from the fact that this shift is always present therefore it is in the general case impossible to extract the shifted frequencies. This method can also be compared to the Slanted Edge MTF measurement method as described in the International Standardisation Organisation (ISO) 12233 standard [87, 88], like in it, the slight incline relative to the pixel grid of a straight line allows the measurement at a sub-pixel precision. The difference resides mainly in the fact that in the Slanted edge method, the Edge Spread Function is first measured at a sub-pixel level via geometric considerations and then Fourier transformed in order to get the MTF, while here the Fourier transform on the inclined line is done first and then the unfolding causes the sub-pixel information.

3.2 Structured illumination (Result/Limitations)

Structured Illumination (SIM) is one of the main super-resolution technique in the literature, and was started with Gustafsson's work [22]. The interest for this method is that it is conceptually similar to the spectral unfolding of the previous section, as the idea is to create moiré patterns, inducing a frequency shift that allows the acquisition of higher frequencies and then to replace this high frequency in its original position outside the standard OTF. However, it relies on the ability to precisely structure the incoming light in a sine wave, which usually requires the use of Spatial Light Modulator (SLM). SLMs are made of liquid crystal array[1] that are often opaque in the mid and far-infrared, and quite expensive in the near infrared (15k to 20k€). Thus, variants were proposed in the literature using random structures [23, 89]. These techniques can greatly simplify the instrumentation as shown by Yeh *et al.* [26, 90] where the light is structured by a simple adhesive tape that is shifted around in order to generate a series of random patterns.

3.2.1 Method

To test the feasibility of these techniques with infrared imaging, a simplified version of [26] was implemented numerically and experimentally. The setup implemented consists of a laser illuminating a diffusor. This diffusor re-emits a speckle, ie. a random pattern that is infinity focused by a lens before going through a test sample. The outgoing result is measured by an infrared camera. The diffusor can be rotated, and the sample can be moved on and off the light path.

The objective here is to create a series of images containing the superposition of the object sample and different random patterns. In order to keep the test simple, these random patterns are known by measuring them on their own without the sample. This limitation is the main difference from the article of Yeh *et al.*, and should limit the resolution gain below $2\times$ as the speckle is measured with the same setup and thus the same NA as the sample.

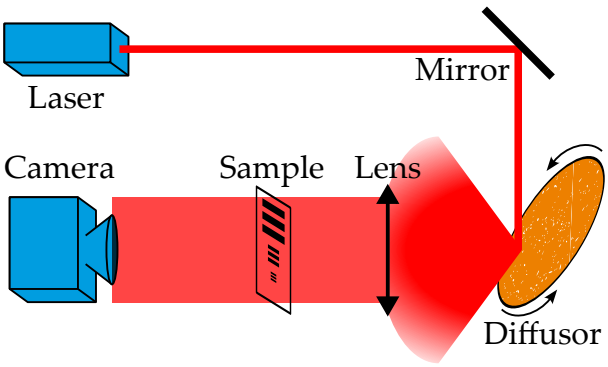


FIGURE 3.5: Setup

The problem can be described mathematically as:

[1] With a pixel pitch of usually $8\,\mu m$

$$y_l = h \otimes (\varrho_l x) + \epsilon_l \tag{3.3}$$

With y_l and ϱ_l the l^{th} measurement and pattern respectively, x the ground truth and h the PSF, all of the same dimensions pixel-wise. This can then be re-written as a matrix product after vectorisation of the images y_l, ϱ_l, x as:

$$y_l = \mathbf{H}\,\mathrm{diag}(\varrho_l)x + \epsilon_l \tag{3.4}$$

The inverse problem to solve is to find x knowing l measurements and their patterns as well as the point spread function of the setup. To do so, a cost function is defined as a cost vector squared:

$$f_l(x, y_l, \varrho_l) = \vec{f_l}^T \vec{f_l} = ||y_l - \mathbf{H}\,\mathrm{diag}(\varrho_l)x||_2^2 \tag{3.5}$$

Via a gradient descent, the function to minimise is:

$$\min_x f(x) = \sum_{l=1}^{N_m} f_l(x, y_l, \varrho_l) \tag{3.6}$$

with the gradient of the cost function being:

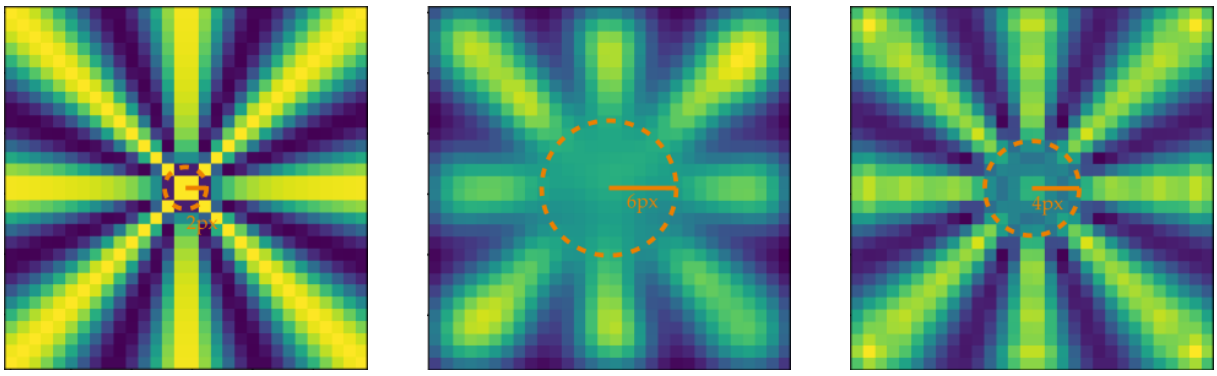
$$\left(\frac{\partial f_l}{\partial x}\right)^T = \frac{\partial \vec{f_l}}{\partial f_l}\frac{\partial \vec{f_l}}{\partial x} = -2\,\mathrm{diag}(\varrho_l)H^T\vec{f_l} \tag{3.7}$$

Algorithm 4 SIM Image reconstruction

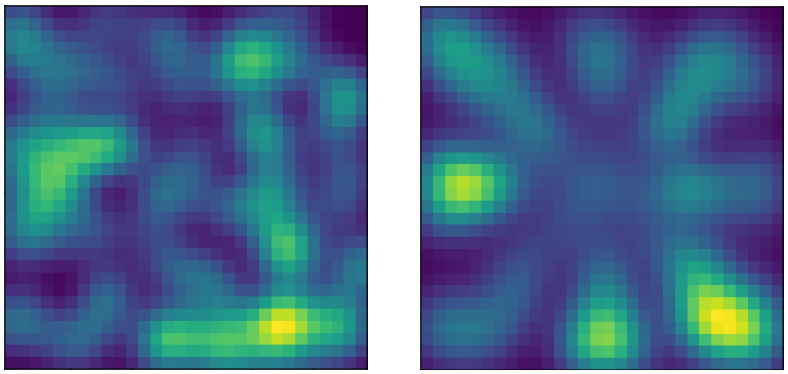
Require: $y_l, \varrho_l, h, l = 1 \dots N_m$
1: $x_{1,0} \leftarrow \sum_l y_l / N_m$
2: **for** $k = 1 : K$ **do**
3: **for** $l = 1 : N_m$ **do**
4: $x_{k,l} \leftarrow x_{k,l-1} - \nabla_x f_l(x_{k,l-1}, \varrho_l, y_l)/\max(\varrho_l)^2$
5: $t_{k+1} \leftarrow \frac{1+\sqrt{1+4t_k^4}}{2}$
6: $x_{k+1,0} \leftarrow x_{k,N_m} + \frac{t_k - 1}{t_{k+1}}(x_{k,N_m} - x_{k-1,N_m})$
7: **return** $x_{K,0}$

3.2.2 Numerical results

First, the method was implemented and tested numerically. The numerical model is made on a 32×32 image. The ground truth object is defined similarly to a Siemens star as a radial cosine, written in the polar coordinates (r, ϕ) as $x(r, \phi) \propto 1 + \cos(k\phi)$, illustrated in Figure 3.6a. This object contains low frequencies close to the edge and high frequencies near the center, it is thus often used for calibration purposes as the maximum frequency will correspond to the smallest radius where lines are distinguishable. Speckles are made by drawing each pixel value following a Poisson law with $\lambda = 0.1$ such that the image contains a few sources and then convolved with an impulse response, that is in Figure 3.6d the same PSF as the rest of the setup. Measurements



(A) Ground Truth (B) Average image (C) Reconstruction



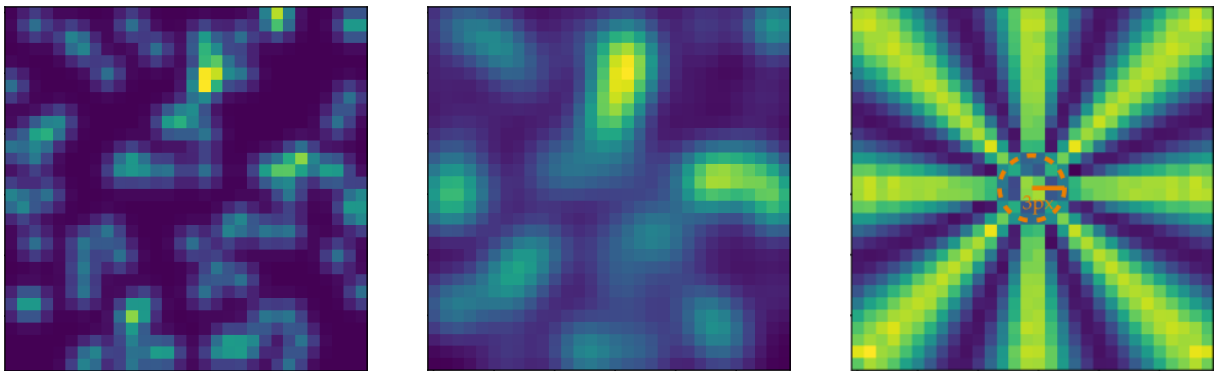
(D) Measured Speckle (E) Measurement

FIGURE 3.6: SIM Numerical result. (a) is the numerical Siemens star spatial frequencies are high near the center, and low on the edges. (b) is the average of measurement as well as the image that is obtained classically, blurred by the optics such that the high frequencies within a circle of 6 pixel in radius from the center are lost. (c) is the image calculated from a series of measurement and their associated speckle. The radius of blurred signal is of 4 pixels. On the bottom half are one example of a pair of speckle (d) and their associated measurement (e) that were used for the reconstruction.

are then simulated as the direct problem from Equation 3.3 with Figure 3.6e being one example of them. The PSF needed for those last two steps is calculated thanks to the theorical OTF equation Equation 2.16 and its Fourier Transform, at a cutoff frequency $\nu_0 = 0.25$ relative to the maximum sampling limit.

The number of simulated measurements is set to $N_m = 75$, and the number of iterations for the algorithm is set to 50. This leads to the reconstruction image in Figure 3.6c. For comparison the average of all measurements is calculated in Figure 3.6b and is equivalent to a measurement that would have been made under homogeneous light. In the classical case, the star is blurred beyond recognition for a radius of roughly 6 pixels around the center, while in the reconstruction, this radius is almost twice as small.

This algorithm can also work if the speckle pattern is known at a better resolution than the object itself. For instance, if in a similar scenario as the previous case, each



(A) Measured Speckle (B) Measurement (C) Reconstruction

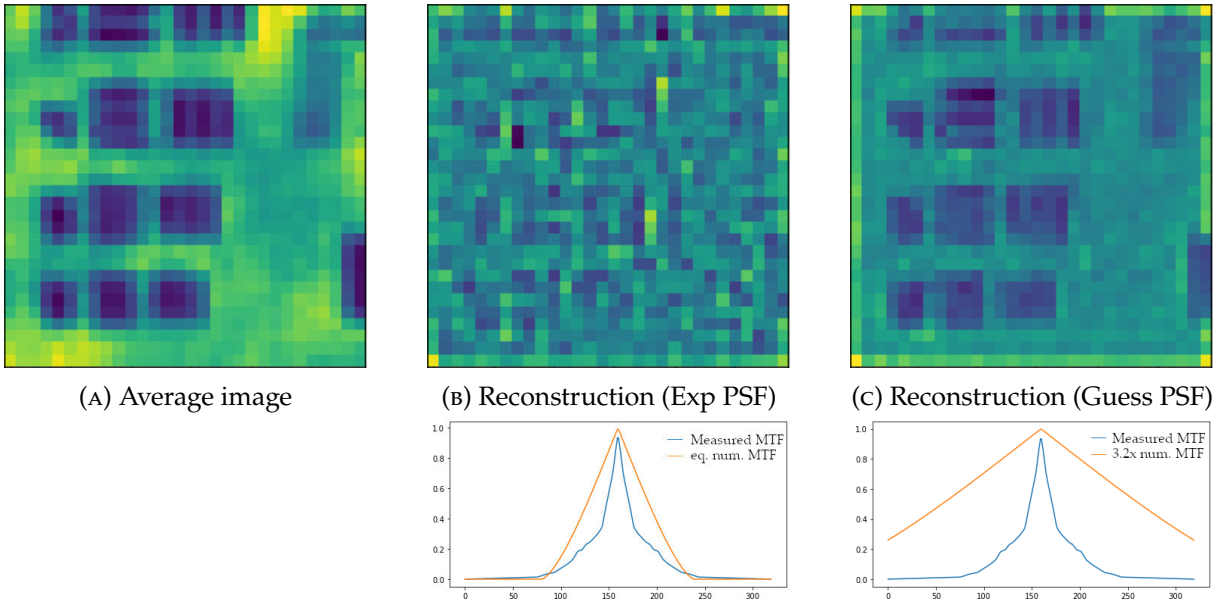
FIGURE 3.7: SIM Numerical result with better knowledge of the patterns than in Figure 3.6, as the speckles spot sizes get smaller and tend towards the size of the pixel, the reconstruction improves to the point where the resolution is the same as the ground truth (And thus sampling limited)

dot from the random pattern is now blurred by a 3×3 kernel in Figure 3.7a instead of the wider PSF, while keeping the other parameters equal including the measurement PSF, the reconstruction will tend towards the ground truth, as seen in Figure 3.7c. The method is thus numerically validated for known speckles.

3.2.3 Experimental limitations

The same algorithm was then tested with experimental measurements with a setup as described in Figure 3.5, with an infrared source and camera. 50 pairs of measures were taken by alternating between the object being illuminated by a given speckle and the speckle alone measured by moving the object away, and then by repeating this process after a slight shift of the diffuser such that the speckle pattern sent on the sample is a new and independent random structure. With the same setup, the PSF was measured using the USAF 1951 target as previously described. Attempts at reconstruction are shown in Figure 3.8. As seen in Figure 3.8b, the method clearly fails and doesn't converge toward an acceptable result even when the number of iterations is increased. As mentionned previously, the measurement of the PSF seemed flawed, thus tests were made using the same measurement dataset but using a PSF calculated numerically. By using a PSF of equivalent cutoff frequency than the one observed experimentally, the results were roughly the same, however, when the cutoff frequency was dramatically increased (By a factor of roughly 3), the image from Figure 3.8c was calculated.

Although this image doesn't really improve the resolution, at least not significantly, some useful conclusions can be drawn. First, and as already hinted in the previous section, the measurement of the PSF is most likely wrong, and the reason for it is that the images here are limited by the sampling first. This sampling limitation affects the measurement of the PSF (Or more precisely the CTF from Equation 3.2) as for high frequencies, the three lines from the target end up represented by less than 5 pixels, which leads to a form of aliasing where the 3 lines are blurred together and thus undistinguishable, making the measured CTF null way before attaining the cutoff frequency



(A) Average image (B) Reconstruction (Exp PSF) (C) Reconstruction (Guess PSF)

(D) MTF Profile (E) Guessed MTF Profile

FIGURE 3.8: SIM Experimental attempt. (a) is the classically measured image. (b) is the result when using the method with the MTF that was measured. The measured MTF is shown in blue in (d) and (e). It is compared to the theorical MTF of same cutoff frequency in orange for (d), and a MTF of 3.2× the cutoff frequency, which ends up beyond the sampling limit frequency in (e). This MTF is used to reconstruct (c).

of the optical system. This point might explain why a thrice as wide numerical MTF of profile shown in Figure 3.8e led to better results, as this MTF clearly corresponds to a case where the cutoff frequency is beyond the sampling limit represented here by the edge of the graph.

3.2.4 Perspectives

The main conclusion from this experiment is that it is not possible to consider a purely diffraction limited problem, and the sampling limit has to be addressed.

This conclusion has led us to instead focus on super-resolution in the digital sense and under-sampling and in particular on compressive sensing. However, and in hindsight, this path could still be further explored, due to limitations of the tests made here. First, the issues regarding the PSF could be partly resolved by doing a joint estimation of the object and the PSF as was done in [26, 91]. This by itself would not suffice to address the sampling limit part of the issue, further study might be required. A simple interpolation step could not by itself solve the issue, but coupling this with a better knowledge of the structuring pattern could be an interesting path. Two main ways of having a better knowledge, ie. at a better resolution (both pixel-wise and diffraction wise), of the pattern are possible. The first would be to measure it at a better scale experimentally. This seems unlikely for speckles at least as it depend on wavelength and measuring it with for instance visible cameras in parallel of using it in infrared will lead to errors, therefore other structuring devices should be explored. Then an interpolation

on the IR image to match the visible resolution in terms of pixel pitch, followed by the reconstruction could work. The second possibility would be to further develop the joint minimisation of both the patterns and the image proposed by [26], as this was not explored in the study here. It allows for a better effective knowledge of the pattern beyond a mere measurement with the same setup, but in its current form doesn't address the sampling limit.

Some points regarding the use of the speckle also remain open. First it might not be the best way to create a random pattern, as it is fundamentally an interference pattern thus one could argue that the pattern that forms in the image plane is not the pattern that forms on the sample, or at the very least that it is highly dependent on the exact position of the object, detector and imaging lens. Second, using a sparser random illumination, that is a light containing only a few sources could bring reconstruction benefits in exchange for more acquisitions. Indeed, if the sources are few and far between, it is possible to work under the assumption that they are all unlikely to overlap and thus easy to distinguish from each other by deconvolution, leading to very precise knowledge of the position of the sources. However, this process might not be very beneficial with some thermal imaging setups, as the diffusion will quickly become the limiting factor.

Another possibility would be to use the thermal diffusion as a low pass filter. Indeed diffusion acts as a blurring effect on the source, meaning that even if the source is smaller than the pixel, after waiting long enough its associated response will be theoretically smooth enough to be non-aliased providing that the source was strong enough to begin with. Such a "diffusion limited" system would have to be treated differently from a diffraction limited one as the system would become time-dependant and thus the model would require the use of a more general Green function that usually leads to very challenging computing problems, as discussed in subsection 2.3.4.

3.3 Compressive sensing with cameras and an intermediary spatial filter

As seen in the previous section, infrared measurement at the microscale with cameras often requires to address the sampling problem. This is due to the fact that, outside Cassegrain-type objectives, achievable magnification in the mid-infrared is of the order of $1\times$ as the one used in the lab, to $3\times$ at least for FLIR SC7000 series camera [92]. With a f-number (noted $f/n°$) of $f/2$, the NA is roughly $NA = \frac{1}{2fn°} = 0.25$, thus the Abbe limit for $\lambda = 2\,\mu\text{m}$ is $d_{min} = 4\,\mu\text{m}$, which is lower than the pixel pitch of $15\,\mu\text{m}$. Cassegrain objectives can be used to reverse this situation, but both limits often end up close from each other, with at best NAs of 0.78 [6], while bringing other limitations most notably an obscuration of part of the image.

In this section, a method for super-resolution is explored in which an intermediary filter between the camera and the object is added, such that the resolution is limited by the size of the spots within the filter instead of the pixel camera. Two types of filter structures are used, first Hadamard, then random. These are tested first with a numerical case and illustrated with Hadamard filters, then experimentally and illustrated with random filters.

3.3.1 Principles

Thus, CS techniques were explored in order to address the sampling problem. The instrumentation itself is very similar to structured illumination, with the main difference being that in general the light is structured after going through the object instead of before, and this difference is mostly of convenience.

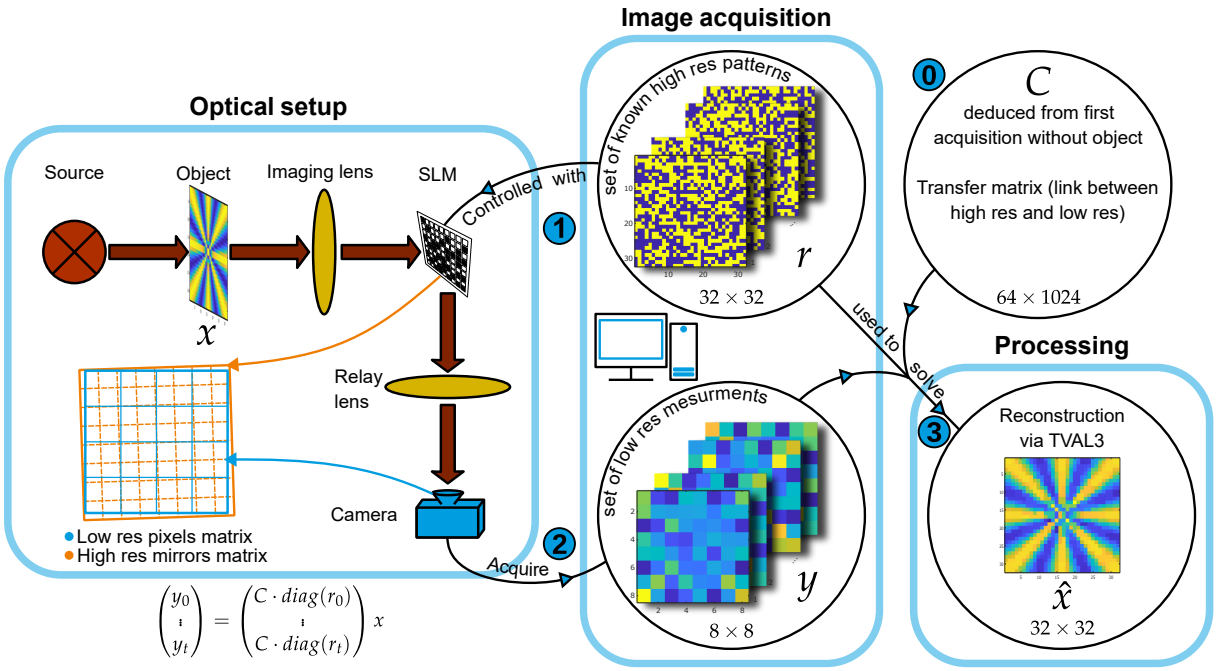
CS imaging techniques were developed first as SPC that turn a single pixel into a camera by structuring the light received by it in a known manner. The instrumental setup consists in the object, a monodetector, and an in-between filter that is usually binary. A set of lenses are placed such that the object is imaged on the filter, and then light from the filter is focused on the monodetecors, as shown in Figure 2.16.

This SPC principle was further developed in order to enhance the resolution of cameras, as described on [93] and further improvements from [68]. The overarching idea is to work with a pair of high-resolution binary patterns and low-resolution camera such that each pixel of the camera acts similarly to the photodiode of the SPC case, with the added complexity of an almost unavoidable missalignment between each pattern and the pixel, plus an increased computational complexity.

The reconstruction requires to solve an inverse problem, once again of the general form $y = \mathbf{H}x$:

$$\begin{pmatrix} y_0 \\ y_1 \\ \vdots \\ y_t \end{pmatrix} = \begin{pmatrix} \mathbf{C} \cdot \text{diag}(r_0) \\ \mathbf{C} \cdot \text{diag}(r_1) \\ \vdots \\ \mathbf{C} \cdot \text{diag}(r_t) \end{pmatrix} x \tag{3.8}$$

With y_i the i^{th} measurement, r_i the corresponding mask, x the ground truth. All of



$$\begin{pmatrix} y_0 \\ \vdots \\ y_t \end{pmatrix} = \begin{pmatrix} C \cdot diag(r_0) \\ \vdots \\ C \cdot diag(r_t) \end{pmatrix} x$$

FIGURE 3.9: Experimental setup and general operating principle of the CS camera. An object is imaged on a random binary mask generated thanks to a DMD. The product of the image and mask is then imaged on the camera. The measure y_i is repeated with multiple different masks r_i. From these set of pairs, an image $\hat{x}$ is reconstructed such that the pixel size corresponds to the size of a point on the mask that is smaller than the camera pixel.

these images are vectorised, that is reshaped to a single column. Finally $\mathbf{C}$ is the transfer matrix between the mask and the camera. Throughout this section, the TVAL3 algorithm is used to solve this problem.

The operating principle is illustrated in Figure 3.9. First, the object is imaged into the high-resolution binary mask, the mask then filters out part of the image as each point either mask or transmits the signal. The product between the mask and the object is then imaged onto the low-resolution camera. Points of the binary masks are smaller than pixels on the camera. This gives a pair of a pattern r_i and its associated response y_i. This procedure is then repeated with a new pattern. Then from all the pattern/image pairs the high-resolution image can be reconstructed, with the new "pixel" size being the size of the pattern point, and not the actual camera pixel. This last step can only be made if the transfer matrix is known beforehand.

In order to better understand the mathematical problem, a more thorough explanation of each parameter is required: Each y_i is a vector containing all pixels of the image reshaped in a single line. For a given image Y_i of dimensions $P \times Q$ given by the size of the camera array, y_i is its associated column vector of dimensions $(PQ, 1)$. y is the column vector containing all vector images y_i in order $y = (y_0, y_1, \cdots, y_t)^T$ and is thus of dimensions $(PQt, 1)$, with t the amount of measurements made and $t \leq \frac{P}{M}$, assuming that the sampling ratio is equal along both axes, ie. $\frac{P}{M} = \frac{Q}{N}$. Vectors r_i are similarly to y_i column vectors constructed from the random binary images R_i of dimensions $M \times N$

such that r_i is of dimensions $(MN, 1)$, and with $M > P$. The ground truth x here is of the same dimensions $(MN, 1)$ as r_i as our objective is to reconstruct the full image on the DMD. $\mathbf{C}$ is a transfer matrix such that $y_i = \mathbf{C} \cdot r_i x$ and thus is of dimensions (PQ, MN), and $\mathbf{H}$ is (PQt, MN).

$\mathbf{C}$ can be viewed as an impulse response, where each line is a $(1, MN)$ vector that contains all the pattern points (or mirrors[2]) that contribute to the lightning of a given pixel. When read in columns, each $(PQ, 1)$ correspond to the image on the sensor associated to a single activated mirror. This is illustrated in Figure 3.10.

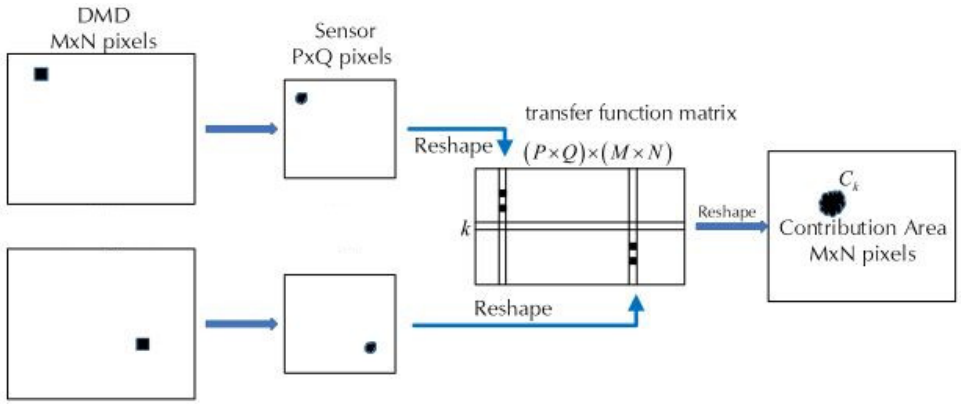


FIGURE 3.10: Transfer Matrix construction [68]

3.3.2 Measurement of the transfer matrix C

The transfer matrix $\mathbf{C}$ is the main obstacle to the resolution of the inverse problem. This matrix quickly becomes of very high dimensions, as it depends directly on the number of pixels on the reconstruction and thus very difficult to compute or construct. The most straightforward way to measure it would be to simply register every impulse response, that is to turn on a single mirror given an homogeneous object, measure its response on the camera and associate it to a column of $\mathbf{C}$. The main issue that quickly arise is that it becomes very time consuming as it has to be repeated for every single mirror (ie. NM times) leading for instance to 1024 measurements even for a small 32×32 image. Additionally it is likely to be extremely noisy especially when the effective size of the mirror is significantly smaller than the pixel, as a very small amount of intensity will be measured by the pixel, leading to high noise.

Despite this massive disadvantage, the matrix $\mathbf{C}$ has an interesting property that can be leveraged in order to simplify the problem, along with simplifying hypotheses that can be applied. Indeed, this matrix has to be extremely sparse since any contribution area of any pixel will only include a few neighbouring mirrors with the rest not contributing and thus containing only noise if measured. Or, viewed from the other axis, each mirror will contribute meaningfully to at most 4 pixels in theory.

From the first remark on the sparsity of $\mathbf{C}$, [68] proposed to use CS for its measurement. It consists in measuring the response of a series of random binary masks given

[2]The terms pattern points and mirrors will be used interchangeably from now on.

on the modulator under homogeneous light in order to deduce each line of $\mathbf{C}$. More precisely, the objective is to solve $f_k = \mathbf{\Gamma} C_k$ for every line k of $\mathbf{C}$:

$$\begin{pmatrix} (f_0)_k \\ (f_1)_k \\ \vdots \\ (f_m)_k \end{pmatrix} = \begin{pmatrix} \gamma_{0,1} & \cdots & \gamma_{0,NM} \\ \gamma_{1,1} & \cdots & \gamma_{1,NM} \\ \vdots & & \\ \gamma_{m,1} & \cdots & \gamma_{m,NM} \end{pmatrix} \begin{pmatrix} (C_1)_k \\ (C_2)_k \\ \vdots \\ (C_{NM})_k \end{pmatrix} \tag{3.9}$$

with $k \in [1, PQ]$ the pixel number, m the number of random measurements and patterns, $(f_m)_k$ the measurement in pixel k for the m^{th} pattern, $\mathbf{\Gamma}$ a matrix containing m random patterns γ_m, vectorised in line. m is at most NM, but the fact that $\mathbf{C}$ is extremely sparse m can be orders of magnitudes below this value.

The TVAL algorithm is used here and repeated for every pixel. TVAL should be well suited to solve this problem, as the expected image is made of a few well defined discontinuities. This method is extremely costly to compute, but only has to be done once as a calibration step. When tested here only small images are used, but in a more general case overcoming that problem requires at the very least to subdivide the problem in smaller regions reconstructed in parallel.

It would also be beneficial to add some prior knowledge as not only the matrix is sparse but it is so in a mostly predictable way, thus it would be also possible to give a first approximation of where the relevant mirrors are for each pixel. In fact, $\mathbf{C}$ has a very specific structure, as can be seen with numerical examples of it shown in Figure 3.12, that could be used to further optimise the problem. It is close to a BCCB matrix with some notable differences due to the fact that the problem considered here is not space invariant and thus not a convolution, explaining why the whole matrix is required, instead of computing FFTs (see Figure 2.3.1). The underlying issue stems from this fact that it is not possible to approximate the problem to a space invariant one unless a perfect alignment between mirrors and pixels is achieved with the number of mirrors per pixels being an integer value and assuming no optical aberration. This last scenario is the one illustrated in the numerical case with Figure 3.12a but is experimentally unlikely to be achieved, and its matrix could be viewed as a rectangular matrix built from lines of the circulant case taken at a regular interval. To my knowledge, I am not aware of efficient computation involving a matrix of the form of $\mathbf{C}$ but it seems that the construction of and computation with it could be greatly optimised if not by managing to turn the matrix into circulant case(s), at least by using the redundancies within $\mathbf{C}$.

3.3.3 Hadamard sampling

Due to the binary nature of the DMD and masks, the Hadamard transform proved itself to be a useful sampling basis in the context of single pixel cameras [94–96]. The main steps of the process applied to camera super resolution is shown in Figure 3.11. First a Hadamard Matrix is constructed. It can be defined recursively via the Kronecker product[3]:

$$\mathbf{H_{2^k}} = \text{kron}(\mathbf{H_2}, \mathbf{H_{2^{k-1}}}) \tag{3.10}$$

[3]The Kronecker product is often noted $\otimes$ in the literature, here, to avoid confusion with the convolution, kron is used.

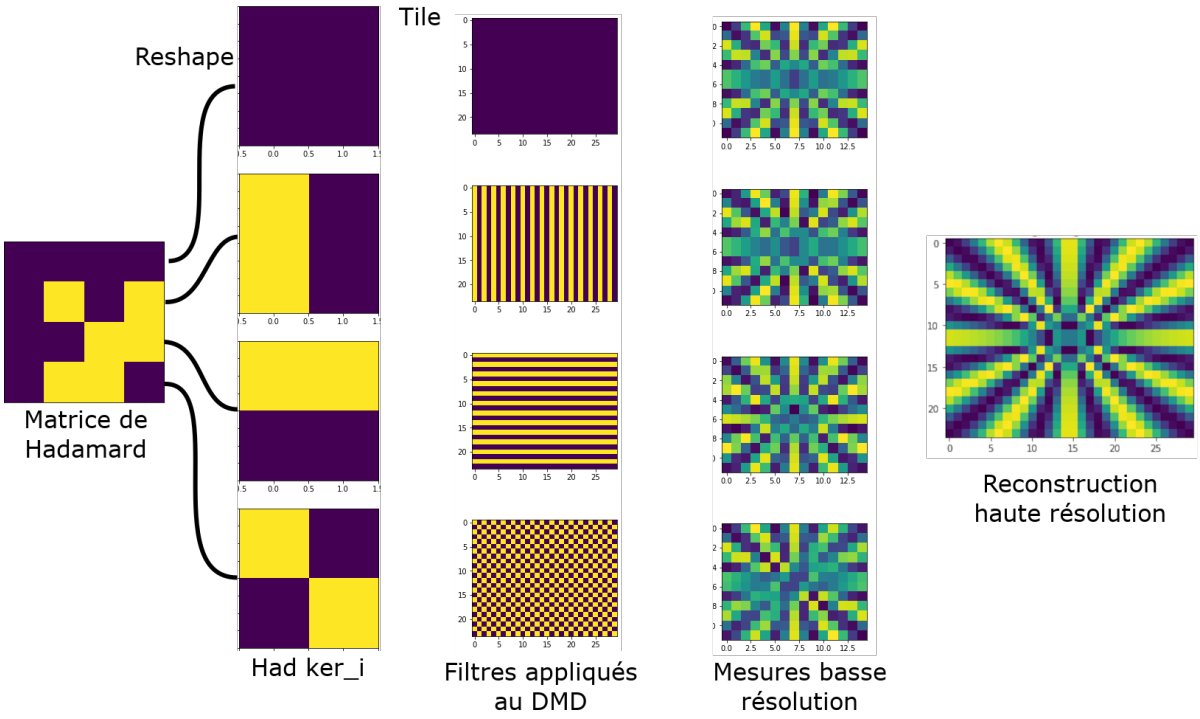


FIGURE 3.11: Hadamard masks construction. First, a Hadamard matrix is generated using its definition and the Kronecker product (left). The size of its width corresponds to the number of mirrors per pixel. Then, each line is taken and reshaped in a squares, called Hadamard kernels (2nd line). Those are the modes to which a single pixel will be exposed to. Then, those kernels are repeated for as many pixels that are on the camera by tiling, creating a series of modes for the entire image that are used as filters on the DMD (3rd line). Measures associated with each mode is made with the lower resolution camera (4th line) in order to reconstruct the high-resolution image.

$\mathbf{H_{2^k}}$ is a matrix of dimensions 2^k with

$$\mathbf{H_2} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \tag{3.11}$$

$$\text{kron}(\mathbf{P}, \mathbf{Q}) = \begin{pmatrix} p_{11}\mathbf{Q} & \cdots & p_{1n}\mathbf{Q} \\ \vdots & \ddots & \vdots \\ p_{m1}\mathbf{Q} & \cdots & p_{mn}\mathbf{Q} \end{pmatrix} \tag{3.12}$$

This matrix is then used to define all the modes that a pixel contains, by taking each line of the matrix and reshaping it into a square. These are then repeated for as many pixels that are on the camera and then applied as masks in order to take a set of measurement associated to them. Usually, the gain in resolution is either 4 or 16 times the amount of pixels of the camera. As the objective of CS is to take less measures than the gain in resolution, re-ordering the Hadamard matrix can be beneficial to sort the modes by frequencies or angles, total gradient, in dyadic order, etc. in order to filter out the modes that are expected to be the least useful.

Hadamard sampling makes sense here due to the fact that TVAL compresses the signal via its gradient, thus the restricted isometry property is respected. However this should not be done when the CS basis is for instance the Fourier transform as the product between it and the Hadamard transform are far from orthonormal. In effect, this approach turns the typical logic of CS upside down, with the roles of the measurement matrix and the sparsifying matrix inverted.

Numerical results

The method was tested numerically first using Hadamard masks, with 32×32 pixels images and masks of full size 128×128. Each pixel corresponds to a 4×4 square of mirrors exactly aligned to the pixel grid. There are a total of 16 possible Hadamard modes per pixel and thus masks. $\mathbf{C}$ is then of dimensions 1024×16384.

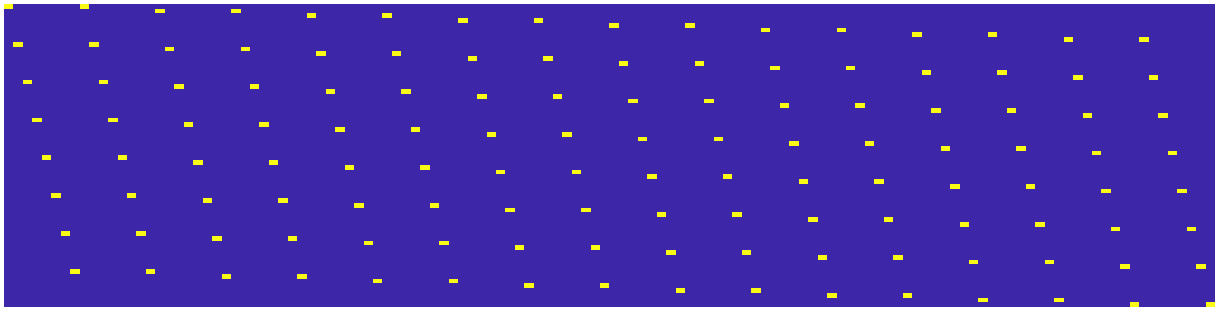
A numerical dataset is first made by calculating the direct problem from Equation 3.8 plus some gaussian noise with x a 128×128 Siemens star. $\mathbf{C}$ is made with each line being a reshape of a 128×128 image containing a single 4×4 square of equal value, the rest being 0. The difference between 2 successive lines being that the square is moved by 4 pixels horizontally or vertically. An illustration of the resulting structure is given in Figure 3.12a. r_i are the 16 images 128×128 of Hadamard masks made from tiled Hadamard matrices as described in Figure 3.11, and y_i the corresponding low res measurements. As for the inverse problem $\mathbf{C}$ is assumed to be unknown, another set of measurement is generated following $f = \mathbf{\Gamma}\mathbf{C}^T + \epsilon$ with $\mathbf{\Gamma}$ containing 100 random known patterns (100×16384), f the associated responses (100×1024), and an additive Gaussian noise of 10% of the average value of f.

With this dataset, the numerical experiment consists in reconstructing the images. This is done in two steps, first with the reconstruction of the transfer matrix from $\mathbf{\Gamma}$ and f, then the reconstruction of the image from r_i and y_i.

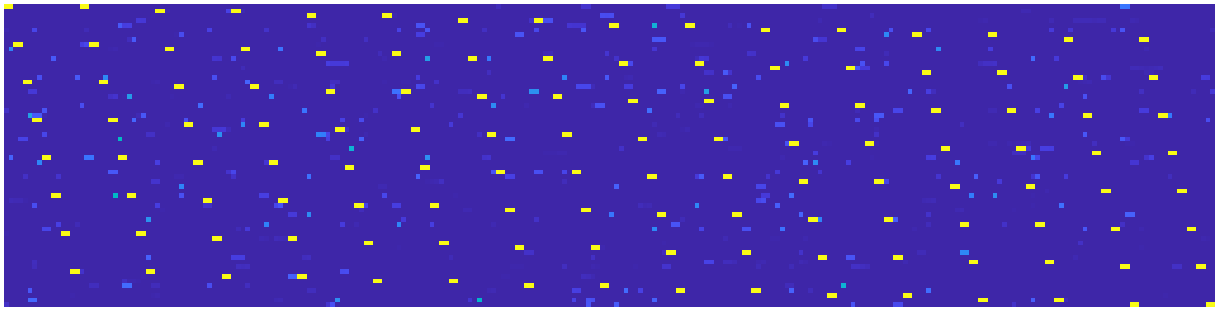
From the set $\mathbf{\Gamma}$ and f, $\mathbf{C}$ is reconstructed by using the TVAL algorithm on every Equation 3.9, resulting in a matrix similar to the one illustrated in Figure 3.12b. The reconstruction is close to Figure 3.12a although, due to the noise added on the measurements the resulting matrix is less sparse than expected, with some values being non null.

Then, from $\mathbf{C}$, r_i and y_i the image can be calculated. Results are shown in Figure 3.13, with Figure 3.13c using all 16 Hadamard modes and without additional noise. A noise level of less than 10% still brings relatively good results, with the structure still recognisable even at high frequencies in Figure 3.13e. It is also possible to not use all of the modes, that is to use less images than the pixel ratio gained, with for instance Figure 3.13d showing an image made from 12 out of the 16 masks. Some frequencies on the image are lost, especially near the center while still maintaining a better resolution than the inputs. It should be noted as well that the effective resolution limit depends on the angle, which is due to the Hadamard masks as each mode represents a frequency at a certain orientation. This numerical experiment validates the approach for the calculation of the transfer matrix from a series of random pattern as well as the reconstruction of the image from the Hadamard patterns.

The decomposition in Hadamard matrices has good reasons to be used in CS, since it acts as a binary Fourier transform by being an universal basis while also being easily constructed experimentally as a set of on/off points. However it becomes somewhat



(A) Reference (Ground truth)



(B) Reconstruction

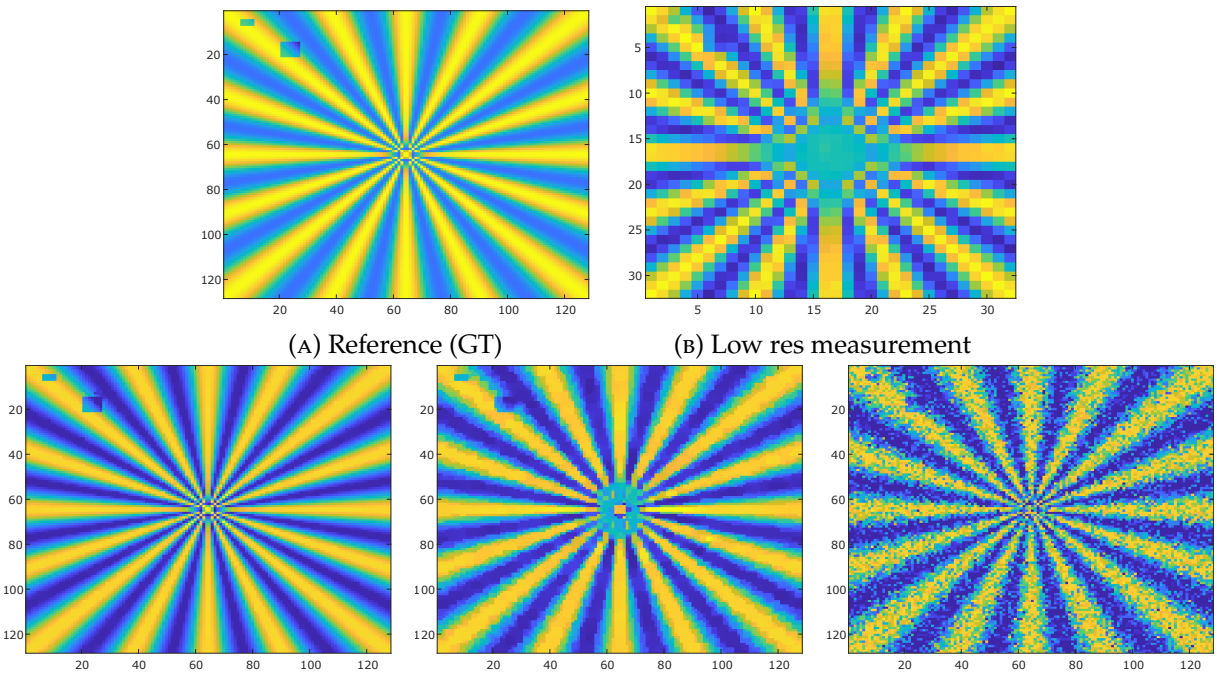
FIGURE 3.12: The matrix **C** (a) and its reconstruction (b), in a numerical case where each camera pixel corresponds exactly to a square of mirrors with no overlap. For readability it is illustrated on a reduced 64×256 matrix. Each line is thus an image of a 2×2 square within a 16×16 image reshaped into a vector, leading to 2 distinct points of 2 pixels separated by 16 rows. A line is thus identical to the previous with an offset of 2×16 up until every 8^{th} line where the whole block is repeated with an offset of 2 pixels. The reconstruction of **C** (b) contains the same pixels with some slight differences in values plus a few pixels outside those that weren't set exactly to zero

less relevant here than in the single pixel camera case for the reason that in order to associate a mode to each pixel a perfect alignment is required, which is extremely impractical. The method might still work with missalignments but it loses some of its meaning. Additionally, Hadamard modes restrict the mirrors to pixel ratios to $(2^n \times 2^n) : 1$ if the resolution for both axes has to be kept the same. At $(2 \times 2) : 1$, CS doesn't bring much as removing a single mode removes a lot of information, while not being much of a time gain over doing 4 measurements with half pixels movement for instance, while at $(32 \times 32) : 1$ and above the computation time becomes a problem.

3.3.4 Random sampling

Due to the limitations of the Hadamard based method, another approach is presented in this section, this time using fully random patterns. This method is tested in this subsection experimentally with two ways of generating the patterns: The first one using DMD, the second one using printed masks, with the setups illustrated in Figure 3.14.

Using random binary patterns is an alternative that exists in the literature [67] which allows for any ratio, is behaving the same way whether the alignment is good or not, is already used for the characterisation of the transfer matrix, and is simpler experimentally especially when using printed masks. A set of patterns relies on the idea that with



(A) Reference (GT) (B) Low res measurement

(C) 16/16 Reconstruction (D) 12 / 16 Reconstruction (E) Noisy reconstruction

FIGURE 3.13: Numerical results with Hadamard patterns. (a) is the ground truth and (b) is a classical measurement limited by the camera pixel size. (c) is the reconstruction using all the Hadamard modes, (d) uses 75% of the modes, and (e) the image made from a set of independently noisy measurement, with a noise level of 8%

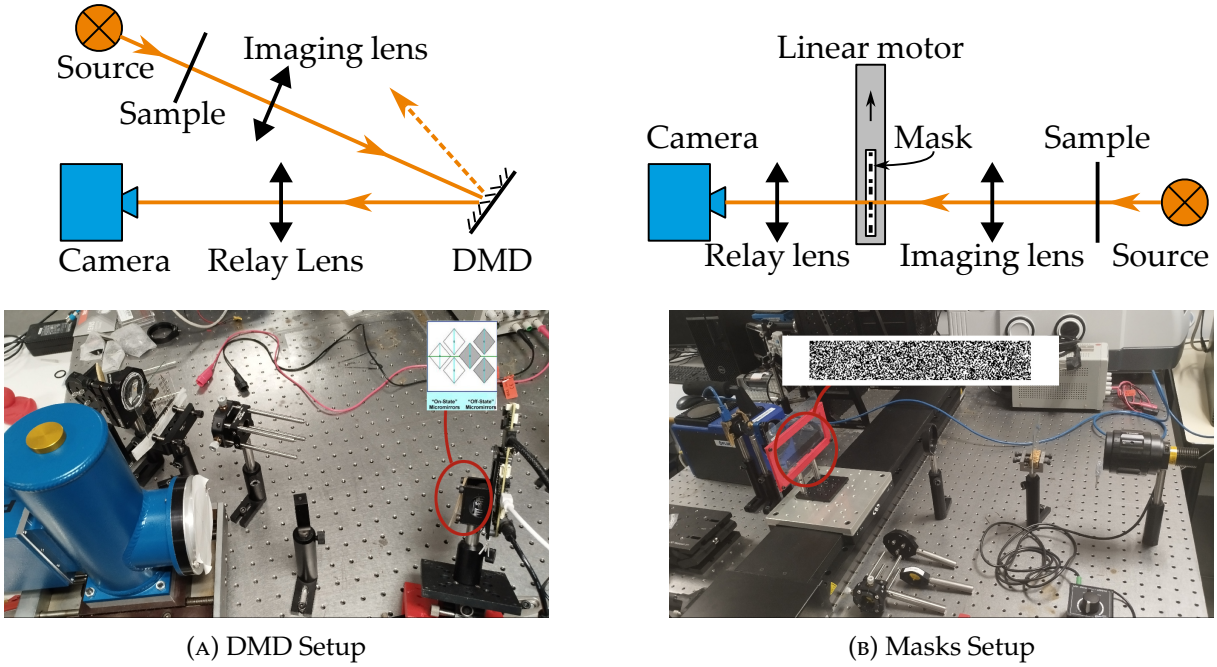
enough measurements, all mirrors can be isolated. The maximum number of modes is thus not clearly defined. It is very unlikely that the compression ratio will be as good as in the Hadamard case due to the coherence between the sampling and basis change being generally higher [48], as hinted at by results in [97] when comparing Ghost Imaging and Hadamard Transform Imaging.

For all the following experiments, random set of images are generated by a Python script in which each pixel receives either a 0 or 1 with 50% chance independently from both the neighbouring pixels and the previous images. This means that on average, each pixel receives 50% of the intensity relative to the regular measurement, just like for the Hadamard modes.

3.3.5 Experimental results using DMD

In this subsection a DMD was used to create the patterns. It consists in a grid of mirrors of tens of micrometers in size that can be programmed to reflect light in two possible directions.

More specifically here, a DLP 4500 NIR [98] is used. It was chosen because of the possibility to work with near infrared light (Up to $2.5\,\mu\mathrm{m}$). This single criterion already greatly limits the options to three models from Texas Instruments, which is the only producer of commercially available DMD. The other two alternatives being either more expensive or too low resolution, although they each have their advantages and drawbacks for this application. Measurements beyond this wavelength are complicated and



(A) DMD Setup

(B) Masks Setup

FIGURE 3.14: Experimental setups used in CS super-resolution with light structured with a DMD (a) or a mask (b)

involve either sending a way stronger illumination, or removing and replacing the glass pane protecting the mirrors, as was seemingly done in [68], which requires a clean room and might also lead to other issues regarding the maximum illumination power that can be sent on a mirror, as the datasheet hints at a thermal limitation.

DLP 4500 NIR is an array of 912×1140 pixels aligned in diamond and with a pitch of $7.6\,\mu m$. Due to the diamond structure, sometimes the distance of $10.8\,\mu m$ in diagonal is given, but this is different from the pitch. Its associated hardware and software allows for communication either via USB in order to stock in memory a small set of images, or via HDMI in which case it can be used just like a $60\,Hz$ monitor. The diamond structure leads to a grid numbering that is not straightforward and is illustrated in Figure 3.15. This means that for instance the mirrors $(0,1)$ and $(1,0)$ are directly adjacent, $(0,0)$ and $(0,1)$ are only diagonally adjacent and $(0,0)$ and $(1,1)$ are not touching (And in fact is just as far as $(3,0)$). In the tests presented here it was decided to attempt to align the mirrors with the pixel sensor, and thus the DMD is inclined at $45°$ to the side. The Python generated images of the patterns given to the DMD are reorganised to correspond to those coordinates, turned into binary black and white video files and then sent to the DMD via HDMI.

DMDs are also acting as diffraction grating, this can be a limitation on wavelengths. For instance if the first-order mode of the grating end up overlapped with the main reflection (zero-order mode) on the sensor. It also means that some energy can be lost in the higher orders. If the behaviour of the DMD is simplified to a standard blazed grating of line spacing $10.8\,\mu m$[4], the angle of the m^{th}-order mode is given by the equation $\theta_m = \arcsin(\frac{m\lambda}{d} - sin(\theta_i))$, with d the line spacing, and θ_i the incident angle. Here, $\theta_i = 12°$

[4]In practice the diamond shape makes the structure more complicated than that, the objective here is simply to illustrate the issue and to give a rough estimation of where the problem can arise

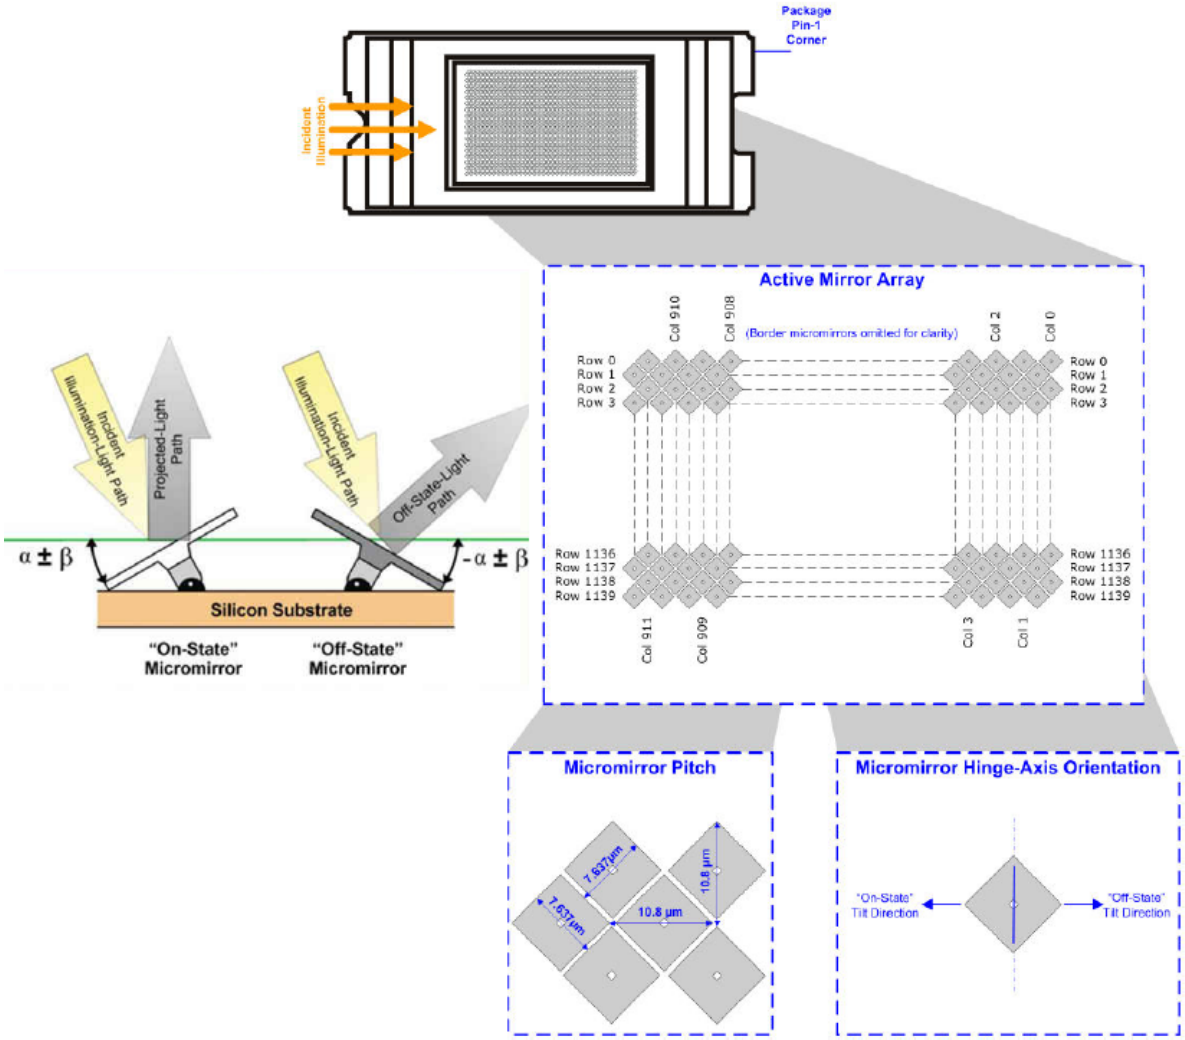
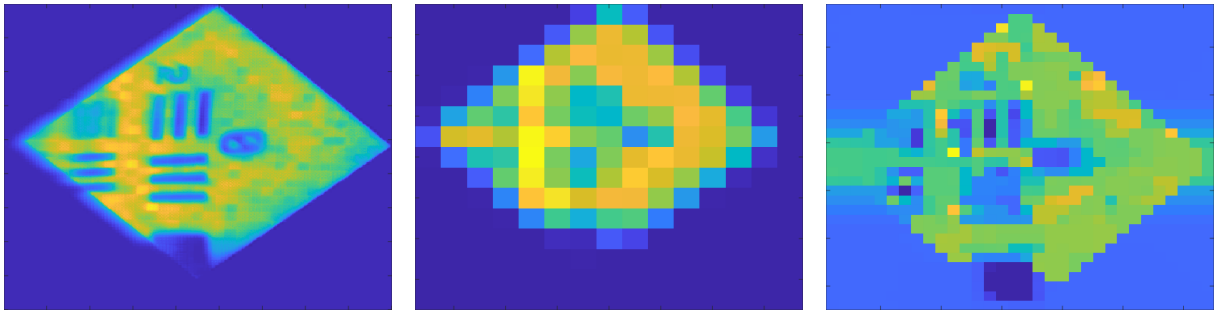


FIGURE 3.15: Digital Micromirror Device DLP4500 structure and characteristic dimensions [98]

leads to $\theta_{m=1} = -9.5°$ at the lowest visible wavelength, compared to the reflection at $-12°$ which is enough clearance for a setup where the distance between the DMD and the sensor is of more than $50\,\mathrm{cm}$.

The system was tested with the setup photographed in Figure 3.14a. As shown in the picture an attempt was made with a 64×64 infrared detector (in blue) that was inconclusive. A reason for this seems to be that the DMD comes with a window to protect the mirrors that filters out most of the signal for wavelengths above $2.5\,\mu\mathrm{m}$. A proof of concept was made with visible light instead. This was done using a 450×450 camera on which pixels were merged such that the image contains 16×16 pixels. Similarly mirrors were regrouped by 16 acting as one and the reconstruction effectively 32×32, with the edges being empty due to the inclination. This was achieved by generating a 16×32 video upscaled to 256×512 sent on the DMD. The lenses imaging on the DMD and on the camera were placed such that the magnification is as close as possible to 1 for both. One issue regarding the optical setup is visible in Figure 3.16a where the left edge



(A) Ground truth (B) Average of measurements (C) Reconstruction

FIGURE 3.16: Experimental results on visible light using a DMD and a visible camera in which pixels were artificially grouped together. (a) is the ground truth of the image, taken before any pixel grouping. (b) is the average image of measurements with pixel grouping, that is equivalent to the low resolution image that would have been classically obtained. (c) is the result after CS reconstruction with 16× the amount of pixels compared to (b).

is blurrier due to the fact that it is out of focus. This is due to the fact that the DMD forces the setup into an angle, which requires the input lens to be off axis according to the Scheimpflug principle in order to alleviate that blur in exchange for a distortion of the distances in the output relative to the input. Another notable drawback is on the amount of measurements required, for instance the image Figure 3.16c required 40 measurements, with less than 30 being low enough that adding or removing one measurement would significantly impact the result. This shows that the use of random patterns is not very well suited for low gains in resolution, as it requires enough samples to draw every pixels in enough distinct configurations. Also, this is to be added to the number of measurements required for the calculation of $\mathbf{C}$ that has to be done once for a given setup, which was 200 in that case. Despite those issues, it was possible to double the resolution on this test case, with frequencies and edges that weren't visible in the average direct measurement Figure 3.16b became visible in 3.16c.

3.3.6 Experimental results using photolithography masks

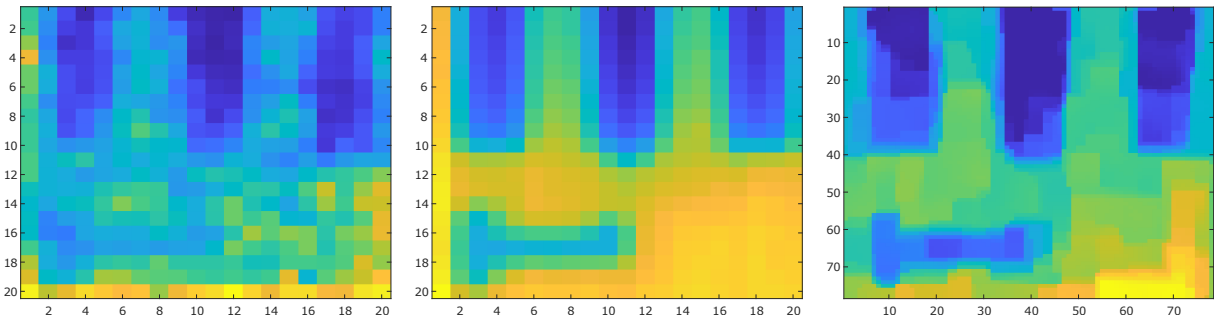
Due to the limitations associated with the DMD, another method was explored involving printed masks made of opaque points on a transparent medium. This approach is advantageous for multiple reasons. First, it can work for a wide variety of wavelengths as the only limitation is the support that can be chosen to be transparent to almost any wavelength, as was done for instance in Single Pixel Imaging in THz by working with masks made of a punctured metal sheet [69]. Additionally it can be used in transmission, that is without any angle, bypassing most distortion issues.

In this subsection, masks were designed by first generating a long line (768×5760) of random binary points of size $20\,\mu\mathrm{m}$ each, and then printed on photo-lithography masks that are transparent to IR light. This print is framed and then put on a linear motor in order to be able to select any subset from the line to be used.

Other tests were made using Hadamard basis patterns as defined in the previous section, however random masks are more relevant for this variant. Random patterns can be used for any magnification, while a Hadamard one only works for a given points

per pixel ratio which adds to the aforementioned alignment issue. Additionally, while each mode of the Hadamard pattern have to be printed on a potentially very long line, random patterns require way less space. Indeed, in order to go from a random pattern to a completely different one per pixel, only a movement of one point minimum is required, which both shorten the size of the mask, but also the time to transition between masks. A second alignment-related issue that come up with this method is that first an initial reference position has to be precisely known, in order to associate the right binary pattern to the measurement and second the mask has to be perfectly horizontal, otherwise after a certain distance the pattern will be off.

The setup is photographed in Figure 3.14b, using the same camera and lenses as in the DMD case except in transmission, and with a linear motor that displaces the sample by $100\,\mu m$ per measurement. The motor is controlled thanks to a LabView program that was created in order to regularly move the sample by that distance in sync with the camera. Pixels were merged as previously in order to work with a 20×20 image corresponding to masks of size 78×78. An example of measurement, the equivalent direct measurement and the reconstruction are shown in Figure 3.17. The result in 3.17c clearly shows sharper definition of the discontinuities on the edge of the lines and number, however the smoother variations are distorted and are forming a blocky picture. This is sometimes called a staircasing effect and is a known phenomenon when using Total Variation regularisation [99]. In order to achieve this reconstruction, the transfer matrix had to be calculated first from the measurements of 200 random patterns, part of it is shown in Figure 3.18. The overall shape is as expected, where each line corresponds after reshape to a square of neighbouring points associated with a pixel, however it is still far from perfect as it contains a lot of noise still in regions that should have been set to 0. Additional post-processing cleaning steps were attempted, by for instance using morphological operators but have been inconclusive up to that point as they tend to add some computational time without leading to meaningful improvements to the final output.



(A) A measurement (B) Average of measurements (C) Reconstruction

FIGURE 3.17: Experimental results on visible light using random patterns printed on photolithography masks (a) is one of the measurements used with a superposition of the object and mask. (b) is the average image and low resolution image . (c) is the reconstruction result with almost $4\times$ the amount of pixels

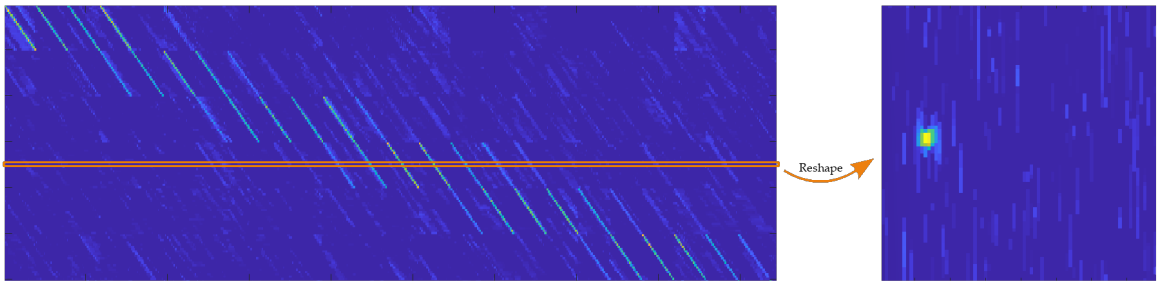


FIGURE 3.18: Experimental result for the transfer matrix $\mathbf{C}$ used in order to calculate Figure 3.17c. A 120×1900 subset of $\mathbf{C}$ (Full size 400×6084) is shown (left), with a 78×78 reshape of the 70^{th} line (right)

3.3.7 Conclusion and Prespectives

In this section super-resolution was explored by effectively increasing the amount of pixels in an image thanks to spatial filters. Two types of filters structure are proposed, both binary in order to be implementable experimentally: Hadamard masks and random masks. The first case was shown to work numerically when the masks and pixels are well aligned at least, and the number of pixels was increased by a factor of 16, even when using less than 16 masks. The second case was tested experimentally with visible light after noticing that the first is difficult to implement well due to alignment issues. Two experimental variants were tested: one using DMD which have the benefit of staying put but are limited in wavelengths, and the other using masks made of opaque points on top of a transparent background. With the DMD, it was possible to double the number of pixels, while with the masks the image went from 20×20 to 78×78. This early conclusion show that it is possible to increase resolution that way. However multiple obstacles remain in order to make the method more practical, especially with the type of infrared signals studied in the lab.

The first issue is simply related to the relatively low signals involved. As each pixels is divided in sub-pixels, so is the amount of energy to be detected. The choice of patterns somewhat alleviate this effect by making so that half the points are on, but either it is half on average with random pattern, or always half but requires precise alignment for the Hadamard patterns. This issue happens twice in the process as the noise is an issue both when making an image and when measuring the impulse responses, thus compounding the issue. As the construction of $\mathbf{C}$ itself is made by CS this leads to non-linear effects, as noise does not simply add an error on the intensity of the impulse response on a given pixel, but also can link totally unrelated mirrors and pixels.

The arguably bigger issue is in the computation times involved and the memory limitation. The method attempts at generalising Single Pixel Cameras but necessarily has to take into account a form of cross-talk as soon as the alignment between the masks and pixels isn't perfect. Which it can't be unless there is an auxiliary way to have an image of the masks at a good resolution (natively). This effect can only be addressed by computing a region simultaneously, leading to the memory limitation. To alleviate it, it is possible to work on small Region Of Interest (ROI) but this will inevitably lead to defects at the boundaries that would need to be treated in an additional step. In the literature, it is often said that CS for image reconstruction is computationally prohibitive,

for the same reason it is said to never compute a transfer function or a DFT matrix. Further research is needed in order to bring improvements focused on the computational side by better taking advantage of the a priori information on the structure of the matrices, in a conceptually similar manner to what the Cooley Tukey algorithm does for instance. Alternatively, with improvement on the hardware side, it is possible to, as is often said, "throw more compute at the problem", in order to have real-time, high resolution images. In this thesis, only personal computers were used, as the objective was to develop an effectively implementable solution on a setup, and with the long term goal of miniaturising it. Another possible different approach which can be considered would be to move the calculation process away from the setup and towards computer clusters, ideally with GPUs and especially if real-time imaging isn't required.

On a related note, the choice of algorithm should be further studied. As previously said, the TVAL algorithm seems to make sense for the construction of the transfer matrix, but is less relevant for the construction of images, especially in the context of IR imaging, as the distortion of gradients due to the staircasing effect is a common issue. A new algorithm should be developed/modified that takes this point into account, where a balance needs to be found between preserving smoother gradients and the sharper discontinuities which are what super-resolution methods intend to find to begin with.

Chapter 4
Under sampled Flying Spot

In this chapter, compressive sensing is used on data gathered via flying spot thermography. A description of the flying spot technique is made, with a specific interest for the pulsed mode and the parabolas method, of which previous results are shown for reference. Then the undersampling method is presented along with the general principles of the three CS algorithms that are used and compared in order to find the most relevant one for this application. This undersampling is applied experimentally for a layered sample using a simple thermal model. This work has led to a publication [2].

Contents

4.1 Introduction . . . 69
4.2 Flying spot: operating principle . . . 69
4.3 Undersampled Flying Spot Principle . . . 71
4.4 Algorithms used . . . 74
4.5 Experimental results . . . 75
4.5.1 Thermal measurements and regular sampling . . . 75
4.5.2 Comparison of sampling ratios . . . 78
4.6 Conclusion and Perspectives . . . 81

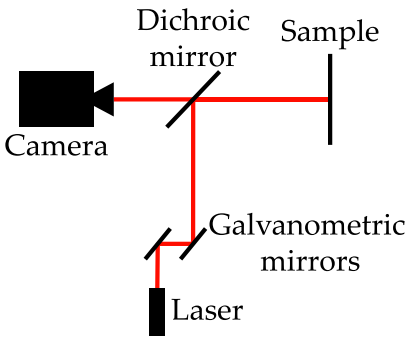


FIGURE 4.1: Flying spot setup

4.1 Introduction

Flying Spot (FS) is a group of contactless thermal technique [100] that is used to measure thermal properties, usually thermal diffusivities. It consists in sweeping a laser beam on a sample thanks to a pair of galvanometric mirrors. A camera then measures the induced field of thermal radiation linked to temperature via Stefan-Boltzmann's law. This measure can be used to detect cracks on a material as they will appear as temperature gradients. Methods have been developed over the years to extract as much accurate information as possible from the setup and were able to benefit from massive improvements in infrared cameras over the past 30 years. It is for instance possible to distinguish surface rugosities from more in depth cracks [101], to measure diffusivities along multiple axes in anisotropic materials [100, 102–104], to scan using laser lines in order to rapidly and accurately detect cracks [105], and to detect defects beyond the pixel size [76].

4.2 Flying spot: operating principle

In this chapter, a FS setup is used in a pulse mode configuration, the sample is scanned by a series of successive pulses going through the sample. This can be modelled by a Dirac thermal pulse in both space and time on a semi-infinite and orthotropic medium [103, p.112]. Starting from the heat equation:

$$\nabla \cdot (\mathbf{\Lambda} \vec{\nabla} T) = \rho c_p \frac{\partial T}{\partial t} \tag{4.1}$$

with ρ the density, c_p the specific heat capacity, and $\mathbf{\Lambda}$ the conductivity. As the medium is orthotropic $\mathbf{\Lambda}$ is a diagonal matrix such that $\mathbf{\Lambda} = \mathrm{diag}(\Lambda_{xx}, \Lambda_{yy}, \Lambda_{zz})$. The problem to solve with its boundary conditions is then defined by the set of conditions:

$$\frac{\partial T}{\partial t} = \sum_{i=x}^{z} a_{ii} \frac{\partial^2 T}{\partial i^2}$$
$$T(\vec{x}, t = 0) = 0$$
$$\Lambda_{xx} \left.\frac{\partial T}{\partial x}\right|_{x=x_{min},x_{max}} = \Lambda_{yy} \left.\frac{\partial T}{\partial y}\right|_{y=y_{min},y_{max}} = \Lambda_{zz} \left.\frac{\partial T}{\partial z}\right|_{z=z_{max}} = 0$$
$$\Lambda_{zz} \left.\frac{\partial T}{\partial z}\right|_{z=0} = Q\delta(\vec{x}, t)$$

This problem can then be reformulated in a Laplace transform space for time, and Fourier for the x and y dimensions, which leads to this general differential equation:

$$\theta(\nu_x, \nu_y, z, p) = \frac{Q}{\frac{\Lambda_{zz}}{\sqrt{a_{zz}}}\sqrt{a_{xx}\nu_x^2 + a_{yy}\nu_y^2 + p}} \exp\left(-\frac{z}{\sqrt{a_{zz}}}\sqrt{a_{xx}\nu_x^2 + a_{yy}\nu_y^2 + p}\right) \tag{4.2}$$

After inverse transforms [103, p.115], this becomes:

$$T(\vec{x}, t) = \frac{Q}{\rho c_p} \prod_{i=1}^{3} \frac{e^{\frac{-x_i^2}{4a_{ii}t}}}{\sqrt{\pi a_{ii} t}} \tag{4.3}$$

With Q the energy deposited by the laser, ρ the density of the material, c_p the specific heat capacity, $\mathbf{a}$ the tensor matrix of thermal diffusivities that is diagonal $\mathbf{a} = \mathrm{diag}(a_{xx}, a_{yy}, a_{zz})$ and such that $a_{ii} = \frac{\Lambda_{ii}}{\rho c_p}$, Λ_{ii} being the thermal conductivity of the material along the given i axis. The effusivity is also related to the diffusivity a_{ii} such that $E_i = \rho c_p \sqrt{a_{ii}}$.

The parabolas method is often used in FS in order to determine effusivities on the plane for both axes, especially in order to characterise anisotropic materials such as fibre composites. It is based on the separability of Equation 4.3, and by taking the marginal averages at depth $z = 0$:

$$\bar{T}_x = \int_{y_{min}}^{y_{max}} T(x, y, t)dy = \frac{Q}{E_z\sqrt{\pi t}} \frac{e^{\frac{-(x-x_0)^2}{4a_{xx}t}}}{\sqrt{\pi a_{xx} t}} \tag{4.4}$$

The same equation can be deduced for the y axis. The logarithms of these equations are of the form:

$$\log(\bar{T}_x) = \log\left(\frac{Q}{E_z \pi \sqrt{a_{xx}}}\right) - \log(t) - \frac{(x - x_0)^2}{4a_{xx}t} \tag{4.5}$$

$\log(\bar{T}_x)$ is a parabola that can be written as the following polynomial:

$$\log(\bar{T}_x) = \beta_0(t) + \beta_1(t)x + \beta_2(t)x^2 \tag{4.6}$$

with coefficient $\beta_2 = \frac{-1}{4a_{xx}t}$ depending only on the diffusivity. A polynomial regression is then used to deduce the diffusivities on each axis.

This logarithmic parabola method has been shown by Gaverina et al [104] to be able to measure diffusivities experimentally on a set of anisotropic materials. Figure 4.2 shows one of their results on a set of 4 samples comprised of a carbon woven in epoxy composite, plexiglass, pine wood, and unidirectional carbon in epoxy. An example of the curves obtained in x and y for the 4 materials are shown in Figure 4.2a and 4.2d. The parameters from these fits are shown to evolve linearly in time in Figure 4.2g. Measurements show that the first two materials (Woven carbon and plexiglass) are isotropic while the other two aren't, additionally the value of diffusivity for the plexiglass is within 2% of the nominal one. In order to double the resolution, the 6×7 grid (Figure 4.2b and 4.2b) was then repeated 4 times, as shown in Figure 4.2c and 4.2f. Each grid can be measured relatively rapidly due to the fact that the distance between each point is wide enough that it is not required to wait for thermal relaxation of the material.

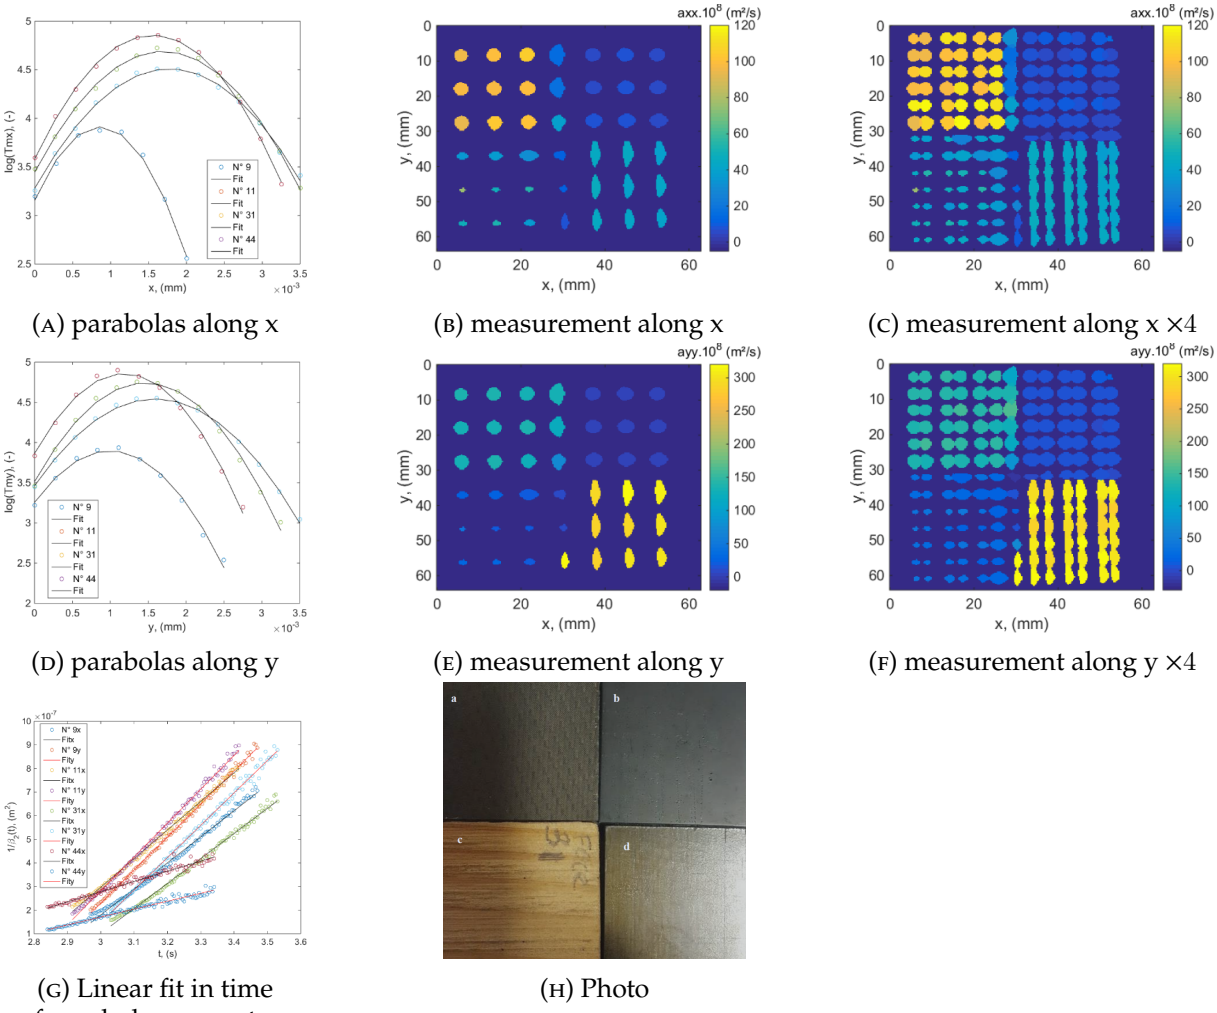
In this work, the effusivity E_z is considered. It can also be separated by taking the average temperature of the recorded image, written as:

$$\bar{T}_z(t) = \int_{x_{min}}^{x_{max}} \int_{y_{min}}^{y_{max}} T(x, y, z = 0, t) dx dy = \frac{Q}{E_z \sqrt{\pi t}} \tag{4.7}$$

The boundaries of the integral in Equation 4.7 correspond to the length and width of the region considered as long as they are both significantly bigger than the spot. In the case of the measurement within this section, it corresponds to pixels within which the signal is averaged, and should be chosen such that all the energy deposited by the laser pulse is known and remains within it during the signal acquisition. Then the temperature depends only on $\frac{1}{\sqrt{t}}$ with a slope coefficient of $\frac{Q}{E_z\sqrt{\pi}}$. Thus a linear regression directly provides the thermal effusivity E and repeating the measurement at different positions of the laser spot on the sample results in a thermal effusivity map. The choice was made to measure the material average effusivity along its depth locally within a small square. The first reason for this choice is to keep the thermal model simple. The second is that this model still applies to a monodetector instead of a camera thereby potentially reducing massively the amount of acquired data. This is however at the expense of some additional information can be gathered thanks to the camera, such as local diffusivities by using the parabolas method. The undersampling presented here can still be applied on the parabolas method as a last step, once the diffusivities along the planes are measured for some randomly selected points.

4.3 Undersampled Flying Spot Principle

In this section, FS is used to map the effusivities E_z only from a random subset of samples instead of the regular sampling that is presented in the previous paragraph. The measurement steps and image reconstruction are described in Figure 4.3 and can be separated in three parts. Starting from the undersampled acquisition, in which a few random points are taken from the full grid. The full grid chosen will correspond to the image reconstruction and the minimum possible distance between 2 samples is thus fixed by it. It is then followed by a linear regression on each temporal response in order



(A) parabolas along x

(B) measurement along x

(C) measurement along x ×4

(D) parabolas along y

(E) measurement along y

(F) measurement along y ×4

(G) Linear fit in time of parabola parameters

(H) Photo

FIGURE 4.2: Example of measurements using the parabola method, from [104], for the x (first line) and y (second line) axis. In the first column are presented $\log(\bar{T}_i)$ for a given point in time and for the 4 materials, and (g) the evolution with time of the parameter $1/\beta_2$ deduced from the parabola fit, thus giving slopes a_{ii}. In the second column a_{ii} are given for a single pass on a 6×7 grid, and repeated 4 times in the last column

to find the local effusivity E_z. Finally, the compressed sensing reconstruction turns the undersampled effusivity map into the equivalent regularly sampled one.

In order to meet the requirements of the thermal model, an additional processing step is made by locating the position of the pulse on the sample then averaging the values of a few (here 21×21) pixels around this position. This is done in order to fit to the thermal model while still being measurable. There is thus a compromise between requiring a surface big enough that the pulse can be considered as a Dirac within an infinite medium along the plane, and on the other hand not drowning the signal in noise by taking a surface so large that the energy deposited by the pulse is too weak compared to the noise on each pixel that doesn't significantly contribute to the response.

This local averaging step leads to overlaps between successive measurements. For instance, in the measurement illustrated here, the step size is $750\,\mu\text{m}$ while the averaged region is roughly 5 times bigger. The typical answer to that issue would be to add a

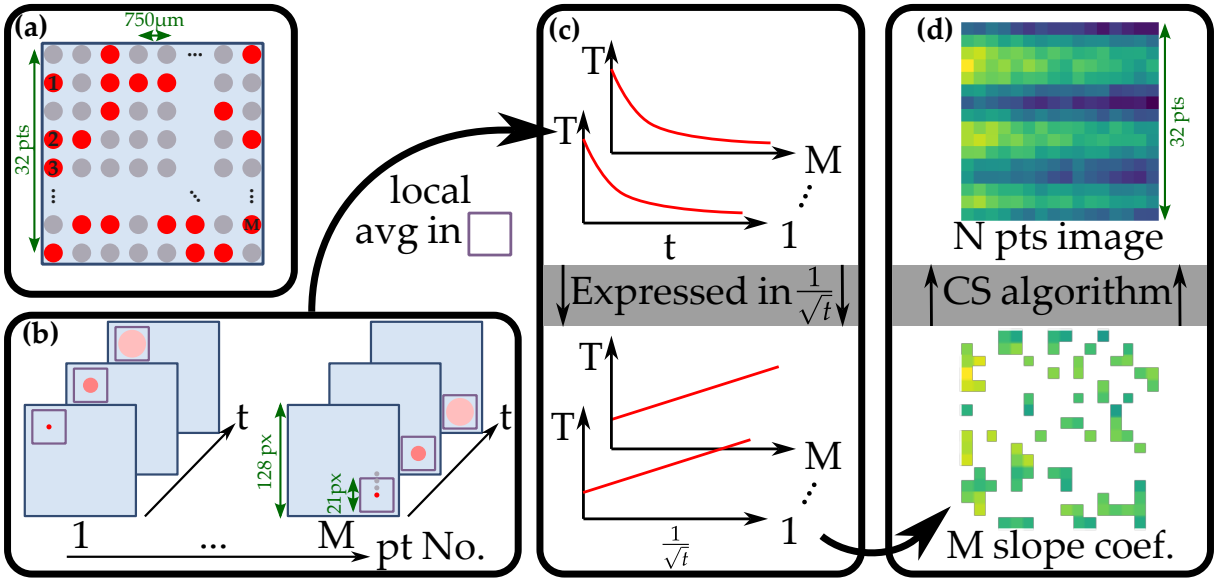


FIGURE 4.3: Overview of the method's principle. The sample is scanned by a series of M points (in red) selected randomly from a N points regular grid (in grey) (a). Each acquisition consists of a short $50\,\mathrm{ms}$ pulse followed by a $2.5\,\mathrm{s}$ recording of the material response to it (b). The local spatial average is then taken in a square centred on the pulse, such that the equation 4.7 applies in order to get M graphs of temperature versus time (c). For a region without defects within its depth these curves evolve linearly with $\frac{1}{\sqrt{t}}$ thus a slope can be deduced by linear regression for each point. From these M effusivities, a $N > M$ points map is reconstructed thanks to a CS algorithm (d). The dimensions associated to the specific measurements shown in this article are highlighted in (a) and (b) in green: A minimum step size between 2 pulses of $750\,\mu\mathrm{m}$ for $\sqrt{N} = 32$ pulses per axis in (a), and in (b) the size of the camera image 128×160 pixels and the size of the spatial averaging window 21×21 pixels with also the 3 previous closest measurement in grey showing that the window is bigger than the step size.

deconvolution step, however the associated convolution kernel is not clearly defined. For instance, it would be wrong to consider that for a measurement, all positions within the averaged square contributed equally to the measure of effusivity, as this square can be defined arbitrarily big. On the other hand, attempting to define it from the shape and size of the diffusion spot leads to the issue that it depends on the time after the pulse, and E_z is calculated from the temporal evolution of the signal, not a single point in time, as well as the structure of the material. A possible compromise would be to consider that the effusivity measured is for all the surface defined by the diffusion spot associated with the last point in time in which the effusivity is linear, but further study of the problem might be needed.

4.4 Algorithms used

In this section three CS algorithms were compared by applying them to the same dataset in the context of IR thermography for NDT: TVAL3 [66], OWL-QN [64] and a FISTA-based method [62]. The algorithms for the three methods were described previously in Algorithm 3, 2 and 1 respectively.

TVAL3 is an algorithm that reformulates a constrained problem similar to Equation 2.53 into a non constrained one via the augmented Lagrangian function. However, the minimisation is made on the total variation, which is the sum of the derivatives of the image. The problem can be written as:

$$\min_{\hat{x}} \sum^{i} ||\nabla_i \hat{x}||_{\ell_1} \text{ so that } \mathbf{H}\hat{x} = y \tag{4.8}$$

The underlying hypothesis is therefore that the gradient of the image is sparse, which is a very common hypothesis for a lot of everyday images as they mostly contain sharp discontinuities such as edges between objects. The promoted reconstructions will exhibit sharp edges (High frequencies) at the expense of smoother gradients, which is often the desired effect, for instance in object recognition, or in medical applications. It is also the case here in NDT as the type of defects are usually sharp changes of thermal properties in the transition between 2 materials, or in the presence of a crack. However it is not well suited for smoother signals and images which would be an issue if a thermal signal was directly used.

The other two algorithms are based on a more typical basis change. More specifically the Discrete Cosine Transform (DCT) space was used both for the sake of simplicity and due to the fact that it is a universal basis and therefore very likely to be sparse for a natural signal. Other strategies could be applied such as wavelets. The inverse problem is described by Equation 2.53 with $\mathbf{H}$ representing the undersampling and $\mathbf{\Psi}$ the basis change, in particular here the 2D DCT. The result $\hat{s}$ is therefore expressed in that new space and the reconstruction result is its inverse transform $\hat{x} = \mathbf{\Psi}\hat{s}$. Hence the problem can be reformulated as:

$$\min_{\hat{s}} ||\hat{s}||_{\ell_1} \text{ so that } \mathbf{H}\mathbf{\Psi}\hat{s} = y \tag{4.9}$$

With $\mathbf{H}$ that can be built as a $M \times N$ matrix from a diagonal matrix $\mathbf{H}_{full}$, where only the M lines corresponding to the measured points are used. This matrix can be the identity matrix if all the measurements are supposed to be space-invariant. $\mathbf{H}_{full}$ is strictly diagonal since in the measurement scheme the measure is made at the position of the pulse only.

Both Total Variation and Cosine Transform methods can be argued to be relevant in thermal NDT. Differences between these algorithms in their robustness to experimental noise at low sampling levels have been mentioned in numerical cases [106] but the most appropriate method is highly dependant on the setup as well as the tested object structure.

The FISTA-based method is the simplest to implement thanks to the PyLops package [107]. While the OWL-QN-based method is implemented on Python thanks to a pylbfgs wrapper [108]. The TVAL3 algorithm and the pre-processing were written using Matlab code. The main underlying algorithms are available on their respective repositories [109–111].

The regularisation parameters are defined in the literature and in programs as ϵ for FISTA and C for OWL-QN. Both correspond to a parameter on a ℓ_1-norm as they minimise a function of the form $f(x) = ||Ax - b||_2 + \epsilon||x||_1$ or more generally for OWL-QN $f(x) = g(x) + C||x||_1$ meaning that they depend on the image sparsity. The pair (μ, β) are the parameters used in the TVAL3 method that stem from the augmented Lagrangian subproblem related to noise and sparsity respectively [66, p.20].

4.5 Experimental results

4.5.1 Thermal measurements and regular sampling

The Flying Spot setup used here is as defined in Figure 4.1, with a laser (AlphaNov PDM $980\,\mathrm{nm}$) that is used to create spatio-temporal Dirac heating pulses of $50\,\mathrm{ms}$, $200\,\mathrm{mW}$ depositing $Q = 10\,\mathrm{mJ}$ and with a Full-Width at Half Maximum (FWHM) of around $1.80\,\mathrm{mm}$ effectively on the sample. The heated sample emission is then recorded for $2.5\,\mathrm{s}$ by the IR camera (FLIR SC7000 InSb) at a frame rate of $20\,\mathrm{Hz}$. The spot is then displaced thanks to a set of galvanometric mirrors (GVS112, Thorlabs) that lead to an effective movement step of the spot of less than $100\,\mu\mathrm{m}$. This is applied to the sample of interest which consists in an heterogeneous pair of materials, alternating between $3\,\mathrm{mm}$ of PVC, and $1\,\mathrm{mm}$ of a unidirectional carbon composite. Each layer is homogeneous in the through-plane direction. A photograph displayed in Figure 4.5b, shows a front view of the sample. It was chosen to illustrate the capabilities of the method due to the fact that it contains sharp geometrical features with a relatively high contrast in effusivities.

The effective power at the position of the sample and FWHM were measured for this study thanks to an IR photodetector. First the power was measured by placing the photodetector instead of the sample in the beam path, then the input voltage of the laser was progressively increased, in order to check that the power read by the detector behaves linearly and the maximum power achievable, which were measured at $180\,\mathrm{mW}$ and $205\,\mathrm{mW}$ respectively. Then the beam width was measured using a knife edge technique, which is a quick way of characterising a laser beam. It consists in progressively covering part of the beam with a straight edge in front of a detector and measuring the

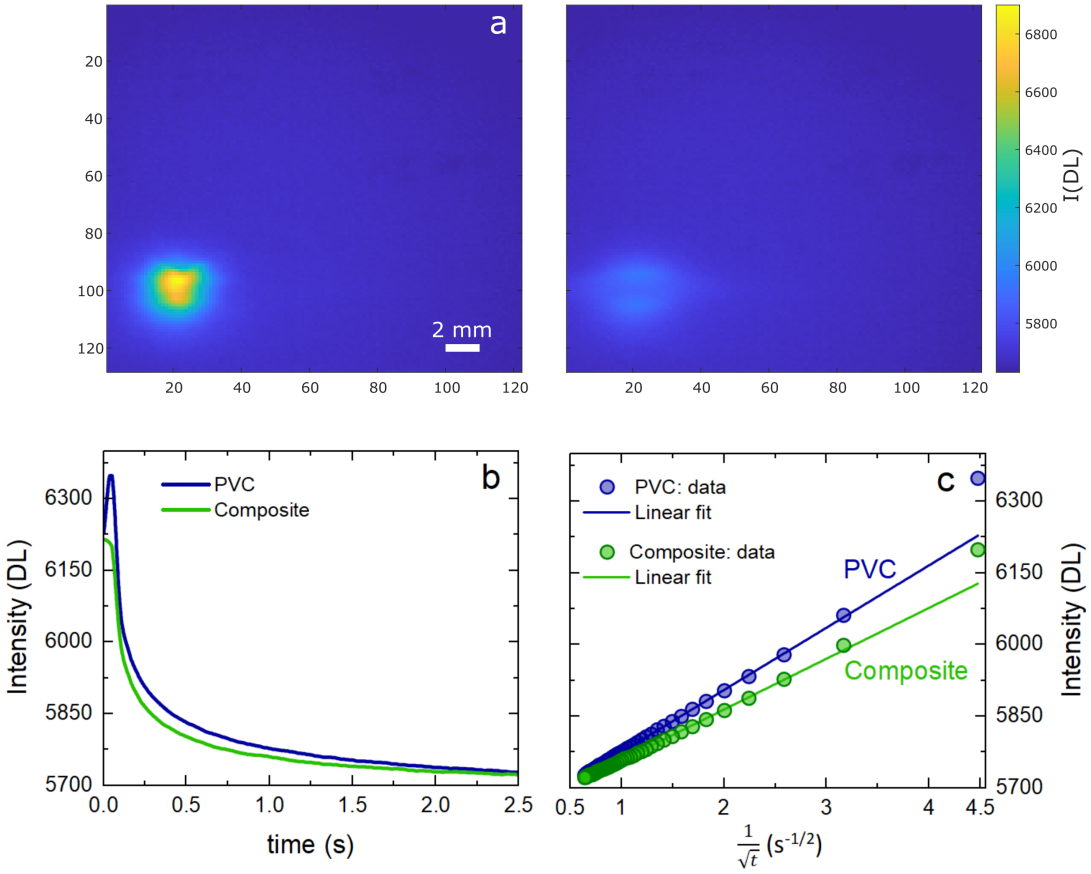


FIGURE 4.4: Examples of acquired data and pre-processing steps. Images acquired by the camera before any processing for a point at $t = 0s$ and $t = 0.5s$ are shown in (a). The evolution over time of the average around the laser spot position is shown in (b) for two different positions. It is then expressed in $\frac{1}{\sqrt{t}}$ in (c).

loss of power relative to the position of the edge. This can be done on all axes, as only the beam width in required only the vertical axis is tested here. The edge is placed at the position of the sample, and the photodetector within a centimetre behind. As the beam is Gaussian, a few characteristic distances from the peak were measured, namely at ± 1, $\sqrt{2 \ln 2}$ (FWHM) and 2σ, as they are widely known to contain 68%, 76% and 95% of the signal thus leading to the measured value of $1.80\,\mathrm{mm}$. It would also have been possible to trace the entire error function if a more precise value is needed.

The galvanometric mirrors and the laser power are controlled via DAQs and a LabView program. The laser receives input voltages up to $500\,\mathrm{mV}$ for the aforementioned linear increase in power, and still increases up to $1\,\mathrm{V}$, while each mirror takes $-10\,\mathrm{V}$ to $10\,\mathrm{V}$ to go from $-20°$ to $20°$. In order to send a series of pulses of arbitrary positions, the LabView program was modified in order to read a file containing a series of values in volt corresponding to x, y coordinates and an associated given laser power, plus a time in milliseconds. As the current setup doesn't include a physical laser shutter, the commands consisted in alternating with short pulses at given positions followed by a long interval where the voltage associated to laser power was set to 0 and the spot was

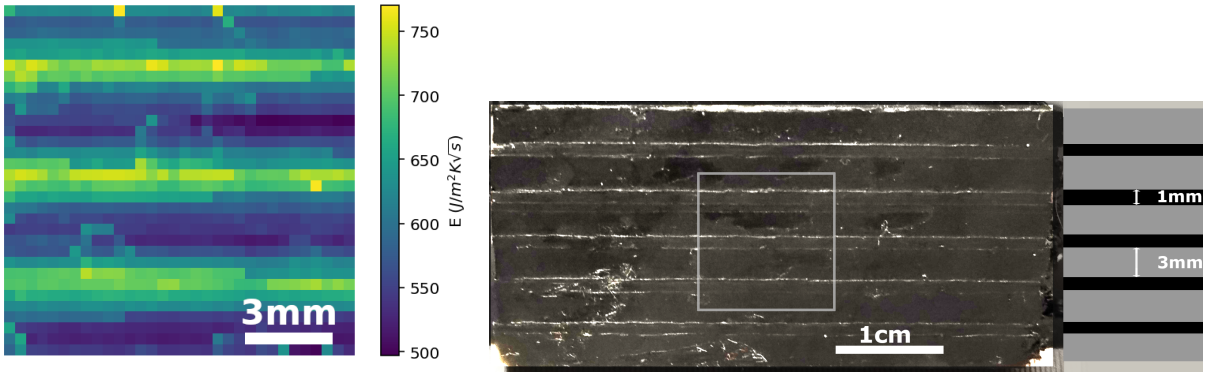
moved as far as possible from the ROI in order to avoid any leakage. These files are generated using a Python script in order to chose random samples on a 2D grid that are then sequentially ordered and turned into the LabView readable table. This sequential order makes it so successive pulses can be close together, it can be beneficial to reorder them differently such that the last heated spot doesn't interfere with the current one, making the method slightly faster.

To evaluate the performance of the compressive sensing algorithms, the sample under study was scanned on a full regular 32×32 square grid first. Different subsets were then used, corresponding to levels of compression ranging from 10 to 90%. For each subset the three compressive sensing algorithms were then used to reconstruct the images. The effective step size on the sample is $750\,\mu\mathrm{m}$. With these parameters, a full 32×32 matrix takes $43\,\mathrm{min}$ to be recorded and the corresponding full video (160×128 pixels used on the camera sensor) requires a few gigabytes of storage space.

Figure 4.4 shows an example of raw measurements followed by the temporal plot and linear regression steps that correspond to the steps described in Figure 4.3. Figure 4.4a shows two acquisitions made by the IR camera at different time delays with respect to the laser pulse, specifically the first (left, $t = 0\,\mathrm{s}$) and 10th (right, $t = 0.5\,\mathrm{s}$) frames. The average within a 21×21 camera pixel square around the maximum of the signal is taken for each point. The position of this maximum is updated at every point.The temporal evolution of this average value is illustrated in Figure 4.4b for two different positions of the laser spot on the sample. In this example the two positions were chosen to be in the middle of each of the two materials comprising the sample, first near the center of one of the PVC layers (blue) and then near the center of one of the composite layers (red). For the PVC, the highest detected signal is stronger than for the composite, and the decaying speed is different between the curves. The same signals are then plotted as a function of $\frac{1}{\sqrt{t}}$ in Figure 4.4c in order to perform a linear regression, with the two lines corresponding to the best fit. Here, the difference in decaying speed is clearly highlighted by the difference in the slope of the two curves.

As it can be seen on Figure Figure 4.4 the data are acquired in digital levels (DL), therefore a calibration step was performed in order to quantitatively convert the measurements into thermal effusivities. The slope coefficient measured in Figure 4.4c is such that it is proportional to the one described in Equation 4.7. To obtain this proportionality factor at a given emissivity, a measurement was made on a homogeneous PVC reference of known effusivity ($E_{PVC} = 498\,\mathrm{J\,m^{-2}\,K^{-1}\,s^{-\frac{1}{2}}}$) and with the same laser pulse ($Q = 10\,\mathrm{mJ}$).

Effusivity maps are shown in Figure 4.5a and 4.6 with the variations corresponding to layers made of PVC and composite. Effusivities are in the 490 to $760\,\mathrm{J\,m^{-2}\,K^{-1}\,s^{-\frac{1}{2}}}$ range. From Figure 4.5a, taking an average of the effusivity value on the horizontal line of points in the center of the PVC layers yields a value of 533 , which is within 7% of the value it was calibrated on. Similarly, the average along the horizontal line of points corresponding to the middle of the carbon composite layer is 734 The thermal effusivities measured here are in the through-plane direction, i.e. along the normal to the sample surface. In the sample under test, this direction is perpendicular to that of the carbon fibres and corresponds to the direction in which the thermal effusivity is the lowest, such that the contrast between PVC and the carbon composite thermal effusivities is



(A) Effusivity map (B) Sample

FIGURE 4.5: Effusivity map from a full 32×32 measurement (a) on the sample photographed in (b)

the smallest. The thermal effusivity estimated here for the composite is likely underestimated, as the width of the zone over which the signal is averaged (21 pixels) is larger than the composite layer. Therefore, the resulting value is also influenced by the thermal effusivity of PVC. To circumvent this problem, experiments should be performed with a smaller laser spot, which was not achievable here.

4.5.2 Comparison of sampling ratios

In order to apply and compare the CS techniques, the full image from Figure 4.5a acts as ground truth. It is then undersampled numerically by truncating the first M values from a random permutation of a sequence of integers between 1 and N, with $N = 1024$ in the present case, and M being a percentage of N. This numerical undersampling is strictly identical in terms of data to an experimental one. For each percentage, the same random subset is used for all algorithms in order to more reliably compare them. However, the randomly selected points are independent between each percentage, meaning that at 80%, some of the points omitted can be selected in 20%. Each algorithm requires some regularisation parameters as defined above, which are given between brackets in Figure 4.6.

Figure 4.6 shows the reconstruction results of the sample under study with various levels of undersampling. The main structure is relatively well reconstructed even with as few as 20% of the sampling points. As the undersampling becomes sparser, finer details however are lost and the image becomes noisier. Out of the three algorithms tested here, TVAL3 performs meaningfully worse than the others as can be seen on Figure 4.6h and 4.7. In particular, as shown in Figure 4.7b the Normalised Root-Mean Squared Error (NRMSE) is nearly 50% higher than the OWL-QN algorithms for a 20% sampling ratio. Specifically, adjacent points values trend to be grouped together as highlighted in the profile on Figure 4.7a which is a known behaviour for total variation methods. The relative underperformance of the TVAL algorithm may therefore mean that the gradient of the image is not as sparse as its cosine transform: Even though the sample itself contains a few sharp discontinuities, the effusivity is measured on a large square ($3.8\,\mathrm{mm}$) relative to the step size ($750\,\mu\mathrm{m}$), which creates a convolutive effect that

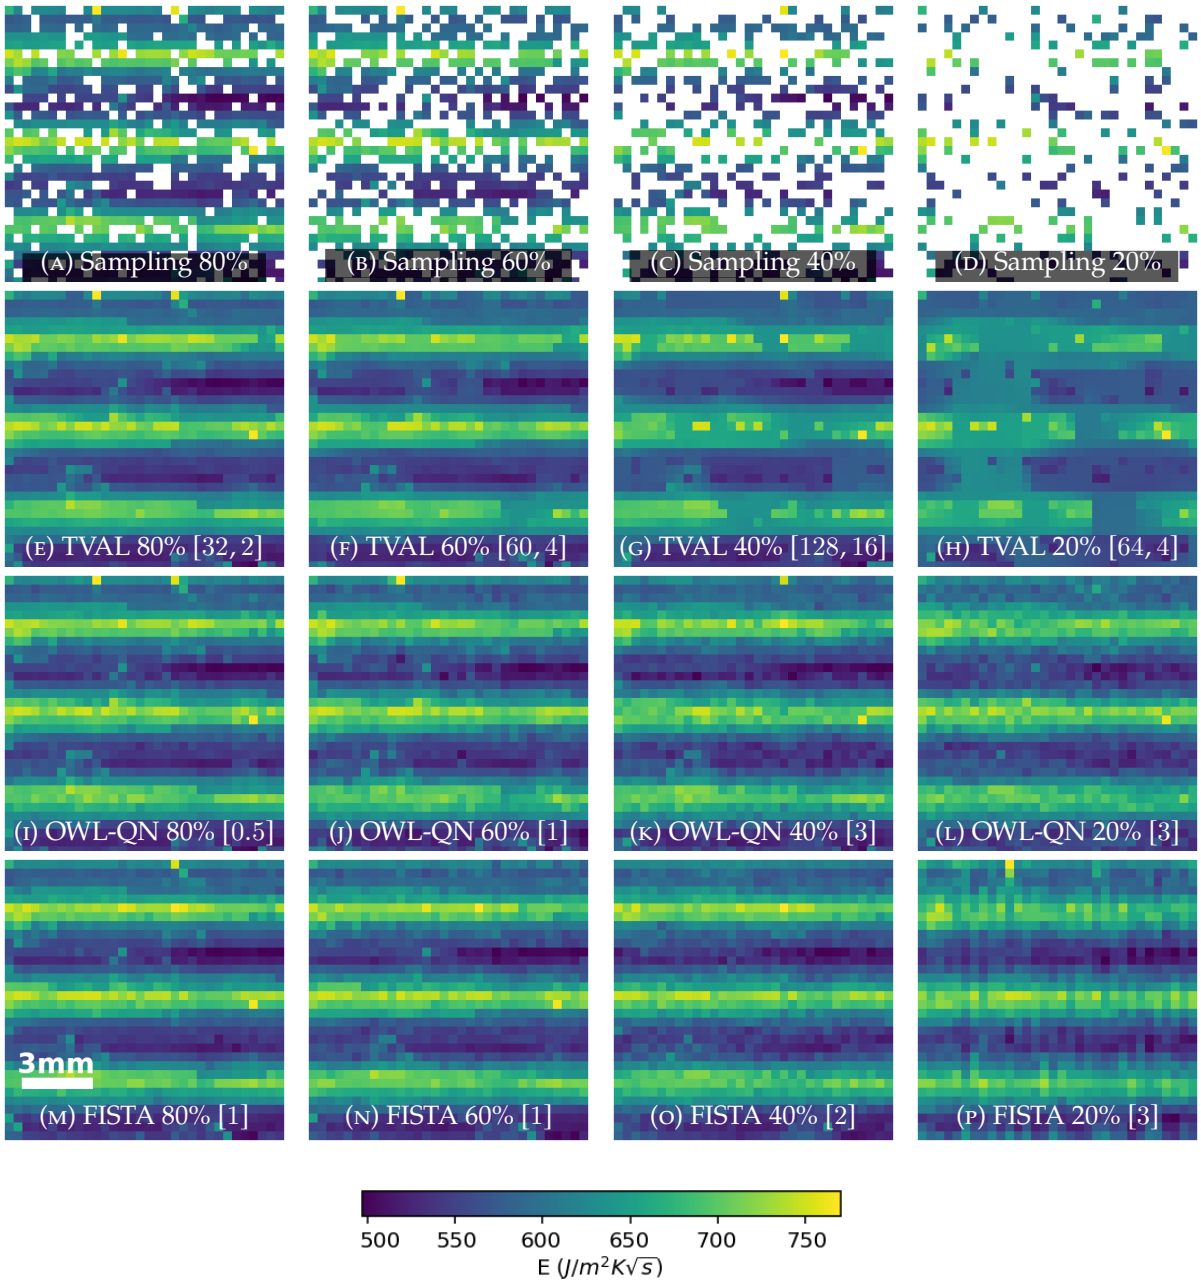


Figure 4.6: Results of the different reconstruction results for a 32×32. (a) - (d) Measurements using 80%, 60%, 40% and 20% of the samples selected randomly. The following lines correspond to the different reconstruction methods, (e) - (h) TVAL3, (i) - (l) OWL-QN, and (m) - (p) FISTA, with each column having its associated sampling ratio. The values in brackets correspond to the various regularisation parameters used: μ and β for TVAL3, C for OWL-QN, and ϵ for the FISTA algorithm, respectively.

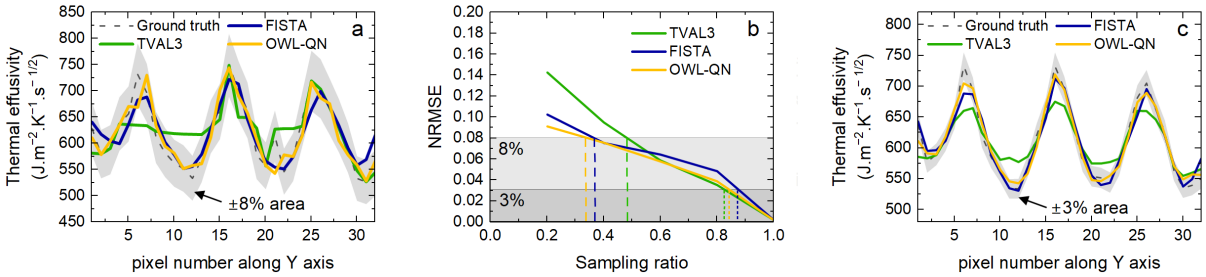


FIGURE 4.7: (a) Profiles on the column 8 of the 3 methods at 20% sampling ratio. The corresponding profile for the full sampling is shown as a dashed line, surrounded by a 8% uncertainty area. (b) Normalised RMSE error between the real image and the reconstructions at different sampling levels for all three algorithms. (c) Profiles for the fully sampled image and the three algorithms averaged over all 32 columns. The gray area represents a 3% uncertainty around the fully sampled image profile.

is due to thermal diffusion. The second issue is the difficulty to accurately select the regularisation parameters, especially when the true image (and therefore the RMSE) is not known, as they partly depend on sparsity of the image. This implies that those parameters, as well as the performance of the algorithms, may vary depending on the studied sample. Whereas OWL-QN and FISTA require one parameter, two are needed for TVAL3, which is an additional obstacle, increasing the difficulty in searching for its optimum. Here, the parameters were iteratively tested in order to plot a L-curve as one would do with an ℓ_2 regularisation (Sometimes called Pareto curve in the litterature). For TVAL3 since the L-curve is two dimensional, each parameter was treated alternatively over a few iterations. The regularisation values used for each images are given in square brackets in Figure 4.6.

The OWL-QN and FISTA methods visually perform similarly down to 40% sampling ratio as can be seen in Figure 4.6, but FISTA appears noisier at a 20%. When looking at the RMSE over a single column of reconstructed points in Figure 4.7b, the error with the FISTA algorithm is indeed higher at all sampling ratios except at 40%. Overall, the OWL-QN method is the more adapted here, giving both the most satisfactory visual results especially at low sampling ratio, together with a consistently low NRMSE. The quality of the reconstruction follows a similar trend when looking at the profile averaged over all 32 columns of the image for a sampling ratio of 20%, as shown in Figure 4.7c. Overall, in this situation, both algorithms are able to estimate the thermal effusivity of both layers of the sample within 3% of the ground truth.

In addition to the quality of the reconstruction, the calculation time is a critical factor in choosing an algorithm. In this work, the images are relatively small and the calculation time remained negligible in all cases, at most a few hundred of milliseconds. It was thus not used as a criterion here. However, for more complex systems, such as super-resolution problems where both input and output are of high dimensions [68], it should be taken into consideration.

4.6 Conclusion and Perspectives

In this work, an undersampling method based on CS applied on a FS technique is proposed. The principle consists in choosing and measuring a random subset of a regular sampling grid, followed by using CS to reconstruct the full image. Two different methods were used: the first consists in a total variation minimisation using the TVAL3 algorithm, whereas the other performs a minimisation in the cosine transform space using OWL-QN and FISTA algorithms. This measurement and reconstruction process was applied to the mapping of a sample consisting of alternating layers of PVC and carbon-based composite. A thermal effusivity map comprising 32×32 points was reconstructed with a sampling ratio as low as 20% while remaining within 10% of the ground truth, thereby enabling a five times faster mapping of our sample. In this case, the TVAL3 algorithm underperformed, whereas the OWL-QN provided the best reconstruction.

Multiple possibilities are opened to continue this work :

First, the sample was homogeneous throughout its thickness, allowing the calculation of the through plane thermal effusivity with a linear regression. The technique presented here could be expanded with the use of automatic slope identification tools for in-depth defect detection.

Also, the approach presented here could be applied to a single pixel detector for a reduction in size of the amount of pre-processed data, instead of using camera followed by an averaging around our ROI for each measurement. Two experimental modifications would be required: first a monodetector of an appropriate noise level should be placed instead of the camera so that its field of view corresponds to a ROI and second the setup should be such that only the sample moves.

Alternatively, a lot of information from the camera acquisition isn't used here. If a reduction of acquired data isn't the priority a possibility would be to use the parabolas method [103], which would allow the measurement of diffusivities along both axes in composite materials especially.

Finally, to overcome the averaging effect of the ROI, adding a deconvolution step after the reconstruction could enhance the precision, although this may amplify errors or artefacts from the previous steps. If the goal is the detection of defects without the requirement for quantitative values of thermal properties, using an arbitrary ROI such that it is smaller than the one imposed by the 1D thermal model is also an option.

Chapter 5
Under sampled Photothermal Radiometry

In this chapter, under sampling techniques are used to reduce the number of measurement on a Scanning Photothermal Radiometry setup. The general operating principles of photothermal radiometry are first presented, followed by a description of the setup and example of the resulting images previously acquired. Then the undersampling method is presented with a variant for scanning the sample at multiple depths. A couple of reconstruction techniques are then used and compared in terms of robustness to high noise and low sampling ratios. This work was presented during the QIRT 2026 conference and its publication is ongoing [1]

Contents

5.1 Method and operating principle 83
5.2 Algorithms used 88
5.3 Comparison on sampling ratios 89
5.4 Comparison on noise levels 91
5.5 Multiple frequencies scan 91
5.6 Conclusion and Perspectives 93

5.1 Method and operating principle

Scanning Photothermal Radiometry (SPR) is a novel thermal characterisation technique developed in the I2M lab [3, 112], with the objective to bridge the gap between FS and SThM. As seen on the previous chapter, FS is operating at the millimetre scale at best but is contactless while SThM is not but can achieve micrometre resolution.

Photothermal radiometry is a contactless thermal technique [113] where the infrared signal emitted from a sample in response to a photothermal excitation is measured. This excitation induces a variation of temperature T, itself related to the emittance (M_e) and is described by the Stefan-Boltzmann law, such that:

$$M_e = \varepsilon \sigma T^4 \tag{5.1}$$

Where ε the emissivity coefficient of the surface and σ the Stefan-Boltzmann constant. As the variation ΔT induced by the excitation is usually very small, the variation of emittance can be linearised as:

$$\Delta M_e = 4\varepsilon \sigma T_0^3 \Delta T \tag{5.2}$$

Further analysis of the signal can be made when the excitation is modulated, called Modulated Photothermal Radiometry (MPtR). For a periodic heat flux at frequency f, the temperature increase of the sample is written as:

$$\Delta T = \Delta T_c \Delta T_f \cos(2\pi f t + \phi) \tag{5.3}$$

Where ϕ is a phase delay depending on the material characteristics, ΔT_c the average increase and ΔT_f the amplitude of the temperature variation. Equation 5.2 thus becomes:

$$\Delta M_e = 4\varepsilon \sigma (T_0 + \Delta T_c)^3 \Delta T_f \cos(2\pi f t + \phi) \tag{5.4}$$

Measuring through a lock-in amplifier the amplitude of ΔM_e relative to the modulated input signal at frequency f gives a value proportional to ΔT_f, while the measured phase offset gives ϕ plus an eventual phase lag dependent of the detector response at higher frequencies. Both amplitude and phase are dependant on the thermal properties of the studied material and thus both can be used for testing ultimately and imaging. However they each have their own drawbacks: On one hand, the amplitude depends on ε the emissivity of the sample surface as can be seen in Equation 5.4 which makes the study of inhomogeneous materials more limited such as the composite tested here. On the other hand while the phase is independent from ε and all else being equal, its noise tends to be greater.

The measurements are representative of the entirety of the heated region of the sample. For a periodic and local heating on a semi-infinite medium it is possible to define a thermal penetration depth, illustrated in Figure 5.1, usually given at the point where the temperature is reduced back to 90% of the heat source at the surface, $\delta_p = \sqrt{\frac{\alpha}{2\pi f}}$ [114, p. 327]. This means that the probed region of the sample depends on the frequency, as f decreases the depth δ_p increases. This is also true for the radial distance, although in practice this distance is mostly dependant on the laser beam width.

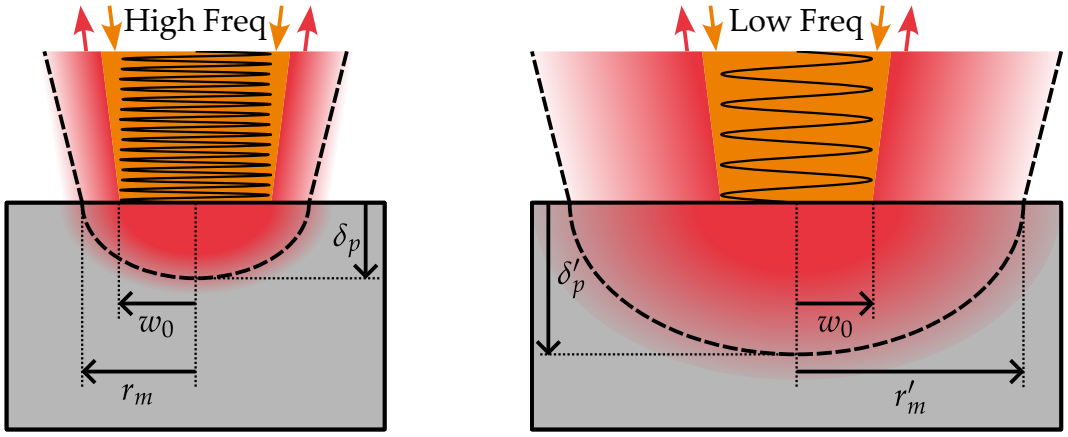


FIGURE 5.1: MPtR frequency and depth dependance. Increases in the frequency of the heat source at the surface decreases the depth δ_p in which the temperature increases, and thus the area probed. The source represented in orange is in practice a focused Gaussian beam defined by its beam waist w_0

As the name implies, Scanning Photothermal Radiometry takes this principle and extends it to imaging by scanning the sample and repeating at each point the measurement. This scanning is done by moving the sample thanks to a piezoelectric motor, thus can achieve very small sub-micrometre step sizes. Is has been used to detect defects such as cracks [115] at these scales and porosities [116], as well as for the characterisation of thin films of sub-micrometer thickness and on an area of a few tens of micrometers [117]. The effective resolution is not limited by the minimal achievable step size of the piezoelectric motor but instead a thermal PSF in its steady state that is itself dependant on the spot size of the excitation light and the frequency as well as the field of view of the detector. It thus acts as a kind of convolution product with a time-independant but very much space-dependant PSF, due to the fact that every material will have a different response [118]. The main, but indirect, limitation on the practical step size is the acquisition time as the measurement are made with a modulated signal and a lock-in, it requires a given time in order to measure the steady state, repeated over each point and frequency.

SPR shares similarities to FDTR which is an experimental setup that uses a first laser to heat the sample locally (called pump) and a second one (probe) to measure a reflectivity variation (called thermoreflectance coefficient) induced by the temperature increase of the sample. In FDTR, these variations are made to be periodic and a lock-in is used to measure the amplitude and phase lag between pump and probe [119]. The pump being frequency dependant leads to the same result as in SPR in the sense that it is possible to probe various depths since the diffusion length is inversely proportional to the frequency, although what is measured is the emmitance in one case and the thermoreflectance coefficient in the other. Due to this proximity, some overlap between the methods and philosophies can exist, for instance a similar undersampling approach to what is presented here has been done by Yang et al. [77] in FDTR while work on this thesis was ongoing.

The experiment setup is schematically described in Figure 5.2 [3, 112]. A laser at

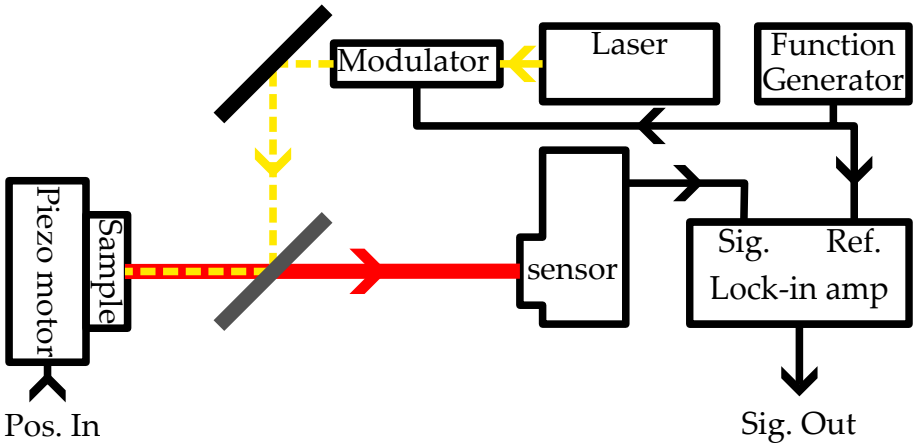


FIGURE 5.2: Simplified Scanning Photothermal Radiometry Setup

$1064\,\mathrm{nm}$ is modulated by an Acousto-Optic Modulator (AOM) at frequency f. It also is reflected toward the sample thanks to a Germanium window, and then focused with a beam waist of $r = 3\,\mu\mathrm{m}$ thanks to a parabolic mirror. The output signal at wavelengths transmitted by the Germanium window is collected by a detector with a wider field of view of $11\,\mu\mathrm{m}$. A Lock-in amplifier then extracts the amplitude A at frequency f and phase shift ϕ relative to the modulation. The measurement is then repeated in order to create an image by scanning the sample along both axes thanks to a piezoelectric translation stage. In the regular sampling case, the step size is $1.5\,\mu\mathrm{m}$ and defines the pixel size in the image. Finally each point is measured at a given set of chosen frequencies in order to retrieve information about various depths of the material as the frequency is inversely proportional to the square root of the diffusion length.

The control system consists of a computer communicating thanks to a LabView program via USB with the lock-in, the waveform generator and the motors. Communication with the lock-in (SR830 for the measurements presented here) allows to enter parameters in, such as τ the time constant, the filter and the sensitivity. The lock-in then extracts the component of the signal from a HgCdTe based monodetector (J15D12 series, Teledyne) and its associated pre-amplifier at the frequency of the waveform generator (33510B, Agilent), in order to be read in phase and amplitude by the software. This frequency can be given by the program and is itself sent on the AOM (MT250-A0.5-1064). Finally a piezo motor and a set of stepper motors are also controlled, with one or the other used depending on the scale of displacement, with the piezo motors allowing for the micrometer scale measurement presented here in the 3 dimensions, and the stepper motor for bigger surfaces at the expense of small scale precision. For regular sampling, a sampling interval has to be defined along both axes, and lines are scanned left to right one after the other from top to bottom. Within a given line, the motor doesn't stop at every point, instead, the sample is moved at a user-given speed v. This allows for faster measurement at the expense of added considerations of the optimal velocity. Indeed the continuous scan means that measurements are made along a small segment of the line, thus the choice of the lock-in time constant (and the slope

of the filter), ie. the time taken by the lock-in to reach its final value, is dependant on the velocity. The time constant also has to be related to the excitation frequency as one of the requirement of MPTR is to measure stationary signals. Typical values of settling time are $T_s = 3$ to 10τ, and v is chosen such that $v \leq \frac{\Delta x}{T_s}$. Finally, in order to reduce noise, it is also possible to sweep multiple times the same line, which can also give a standard deviation that can be used to determine the quality of the measurement, at the expense of speed.

An example of regular measurement with a silica fibre and pyrocarbon composite sample from [3] is shown in Figure 5.3 for a 15×30 image, taking roughly $10\,\mathrm{min}$ each, repeated for 5 frequencies. The fibres are especially visible in the phase plots, while the amplitude highlights the boundary resistances. As expected, the fibres appear more blurred at lower frequencies, ie. deeper within the material. Although not detrimental for the image, a slight lateral blur can be seen on some lines that show that the speed of the scan is near the limit of the system.

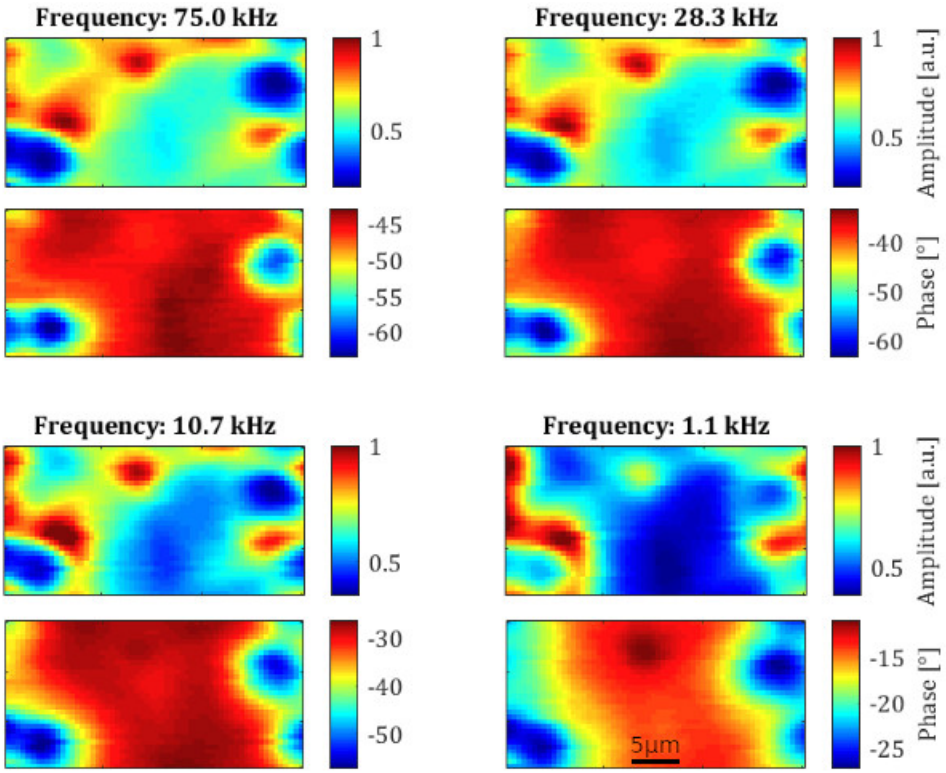


FIGURE 5.3: SPR measurement with regular sampling on PyC from [3]

For the following measurements in this section, this LabView program was further modified in order to acquire data according to a given set of positions that are sent to the motor. This set is defined thanks to a new Python script generating a random subset of point from a regular grid, and a weighted random one for the second part of the method at multiple frequencies, as described in Figure 5.4. Each point is measured at every frequency first, and then moved, unlike in the regular sampling case where it is more efficient to scan the sample for each frequency.

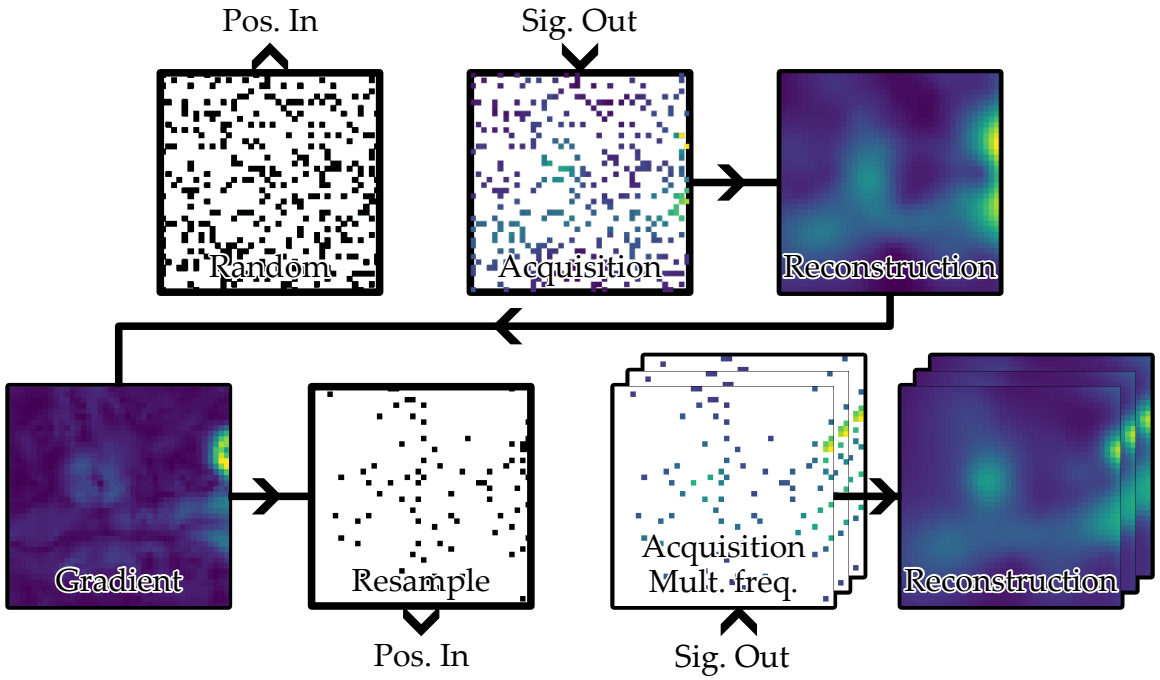
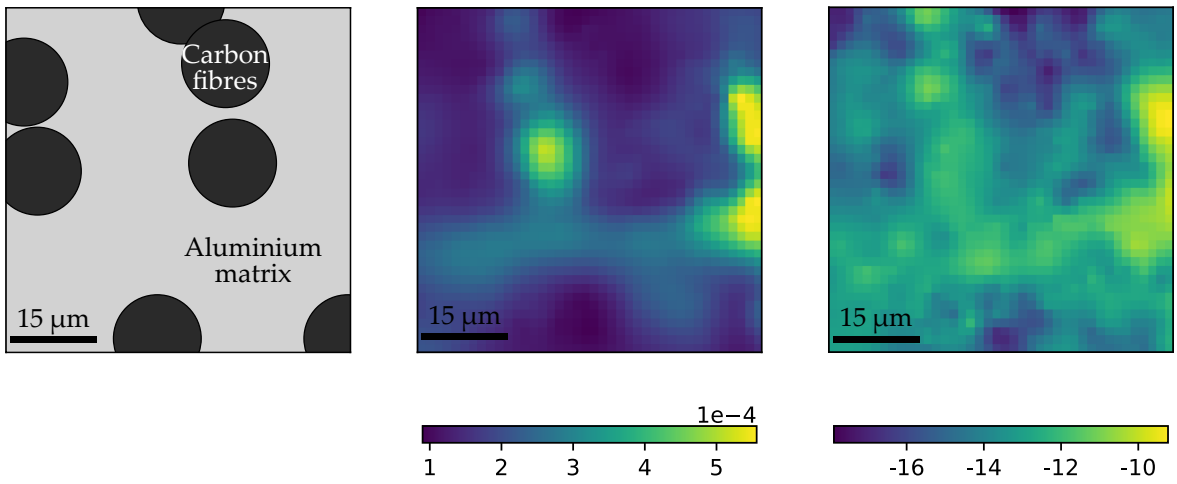


FIGURE 5.4: Sub-sampling scheme at multiple frequencies, consisting in 2 parts: A sub-sampled acquisition (top) at a given frequency, followed by a stronger sub-sampling at all the other frequencies (bottom). First a fully random series of points is created, each point acting as a set of coordinates for the piezo controller, then the associated signal at those points is measured thanks to Figure 5.2, from these acquisitions at known positions, reconstructions algorithms are used to approximate the equivalent full image. The gradient of the image is calculated and used in order to sample preferably where the gradient is high, this new sampling is then used again on the setup but this time along all the frequencies.

The two methods are compared on a given ROI of a carbon fibre in an aluminium matrix. The sample was scanned on a 40×40 grid of pixel size $1.5\,\mu\mathrm{m}$, with a lock-in frequency at $6\,\mathrm{kHz}$. The acquisition time per measured point is limited by the lock-in time constant $\tau = 30\,\mathrm{ms}$, as a settling time of 10τ is required, with additional delays caused by communication speed with the motor and data saving. Which leads to an acquisition time on the order of $1-2$ seconds per point per frequency, thus a full 40×40 image takes $30\,\mathrm{min}$ to $1\,\mathrm{h}$. A study on a wider field of view and at multiple depths can become time consuming, thus a faster acquisition is desirable. Two criterion are studied for the amplitude: Robustness to the number of samples and robustness to noise.

The regular scan in both space and frequencies can quickly become time consuming, usually taking hours, thus a weighted random scheme was developed and implemented to further optimise the acquisition time beyond the purely random sampling, as described in Figure 5.4. First, a fully random sampling is used at a first frequency, in the case presented here 25% of a 40×40 image at $f = 6\,\mathrm{kHz}$, then the full image I is reconstructed. Its gradient is calculated along both axes and used as a probability



(A) Schematic (B) Amplitude (u.a.) (C) Phase (deg)

FIGURE 5.5: Sample at 6 kHz

density. This probability proposed here is defined per pixel as:

$$P_i = \frac{G_i + \alpha \max(G)}{\sum_i (G_i + \alpha \max(G))} \tag{5.5}$$

with $G_i = |\frac{\partial I}{\partial x}| + |\frac{\partial I}{\partial y}|$ and α an arbitrary value set at $\alpha = 0.02$ here in order to still have a few samples in the homogeneous regions of the image. Then samples are drawn according to these probabilities with a number of samples inferior than in the fully random case. Measurements are then made according to this re-sampling where a range of frequencies are acquired for each point. The norm ℓ_1 is used to define G_i, as the objective is to create probabilities proportional to the gradient, an ℓ_2 norm would also work with marginal differences. The underlying assumption behind this re-sampling is that overall the material remains similar when probed at higher frequencies. In the case of the sample studied here of carbon fibre in an aluminium matrix, the fibres are along the depth. A likely spot for defects is at the interface between the fibre and the matrix, so the new sampling is made to focus on this part as the gradient at interfaces is higher than within a fibre for instance and is even greater when a defect is present.

A schematic of the ROI of the sample is drawn in Figure 5.5a, it contains carbon fibres in an aluminium matrix. The fibres are around $15\,\mu m$ in diameter and can be distinguished from the background as plateaux of amplitudes in Figure 5.5b. High variations in amplitude can correspond to defects such as detachment between fibre and matrix, or porosities within the matrix but can also be caused by the difference of thermal conductivity between the matrix and the fibre.

5.2 Algorithms used

In this work, comparisons were made between CS techniques and 2D irregular interpolations. Thus two methods were used with two associated algorithms: OWL-QN [64,

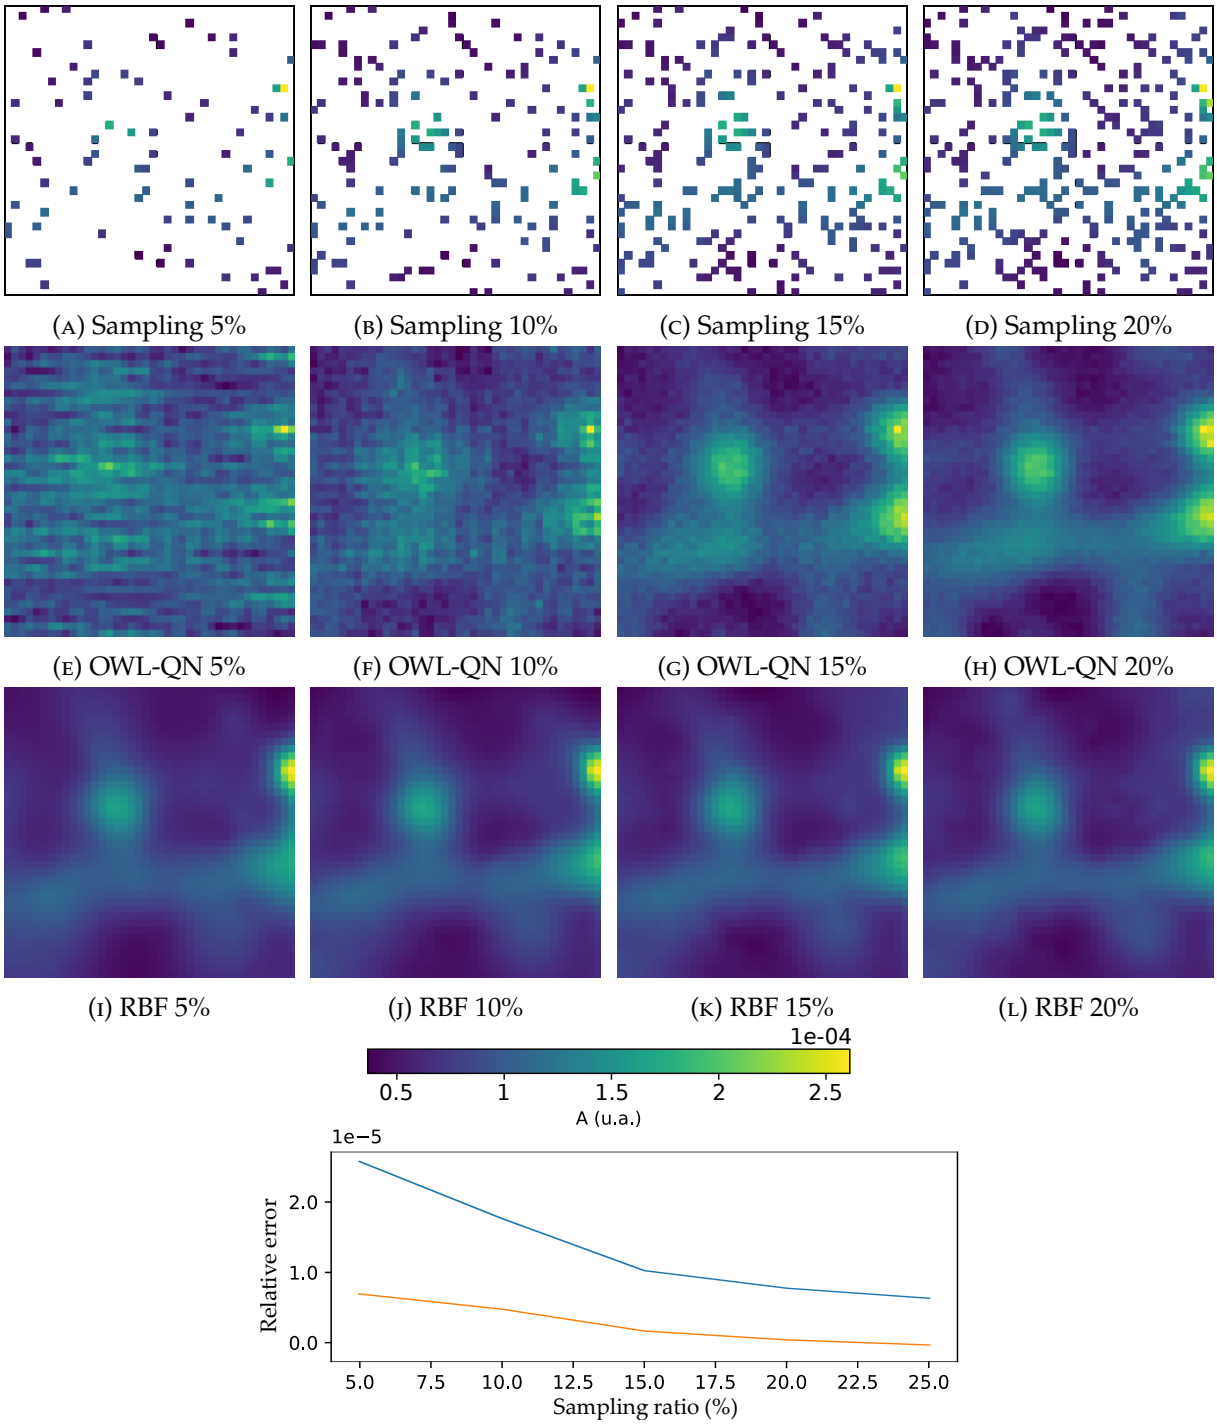
108] as it was shown in the last section to be efficient in the Flying Spot case, and RBF [42] as it is both efficient and widely accessible thanks notably to its SciPy implementation [120]. The basis chosen for CS is the DCT, as it is an universal basis and efficient to compute. For RBF, the kernel chosen is the thin plate spline $r^2 \log(r)$, this kernel has the advantage of not requiring any tweaking of the scale factor. In Figure 5.8i to 5.8l an additional comparison was made with a regularised, ie. smoothed, RBF reconstruction by adding a parameter λ_R such that $y = (\mathbf{A}+\lambda_R\mathbf{I})w$. This reconstruction relaxes the conditions by allowing the result to not strictly pass through the data points, it is therefore not an interpolation but instead an approximation. This RBF approximation has additional drawbacks by adding a bias and requiring the correct choice of regularisation parameter similarly to the OWL-QN case.

5.3 Comparison on sampling ratios

The dependence on the number of samples is studied first, with a laser power of $1\,\mathrm{W}$. The number of samples is progressively decreased from 25% (400 samples) to 5% (80 samples) with each case being a subset of the previous one. Figure 5.6 shows for each column the undersampled acquisition, followed by the reconstructions both with the CS and RBF methods. A relative error graph is provided in Figure 5.6m calculated as a root mean squared between the result and the reconstruction at 25% using RBF. While the RBF here doesn't require any regularisation parameter, OWL-QN does, and is set at $C = 0.1$.

As expected reconstruction quality decreases as the number of samples decreases. For low sampling ratios, the RBF interpolation clearly outperforms the OWL-QN reconstruction. Even a number of samples as low as 5% suffice to reconstruct the image with minimal loss of quality in RBF, while this limit is at roughly 15% for OWL-QN. Overall, this difference is general but the actual limit ratio will depend on the measured structure. Here, the image is quite smooth due to the fact that the step size is smaller than the beam waist of the laser, therefore there is a convolutive effect between successive points which is a favourable scenario for undersampling. The error graph shows that the RBF method consistently outperforms the CS one. Although this could be partly explained by the fact that the reference image used was obtained using RBF at higher sampling ratios, the evolution of both curves show that the relative error with RBF increases less quickly as the number of samples is reduced.

Equivalent regularly sampled images are given in Figure 5.7 for comparison, with Figure 5.7a containing 25% more samples than Figure 5.6a, Figure 5.7b containing 5% more samples than Figure 5.6b, and Figure 5.7c the equivalent of Figure 5.6d. 100 samples is the lowest amount of samples at which the regularly sampled image still exhibit the overall structure of the object, but any lower than that, the sampling theorem is clearly not respected as the two rightmost peaks become indistinguishable and even at Figure 5.7a they are separated by only 1 pixel which means that a different offset (ie. if the sample was moved 1 or $2\,\mu\mathrm{m}$ vertically) would have made a noticeably different result.



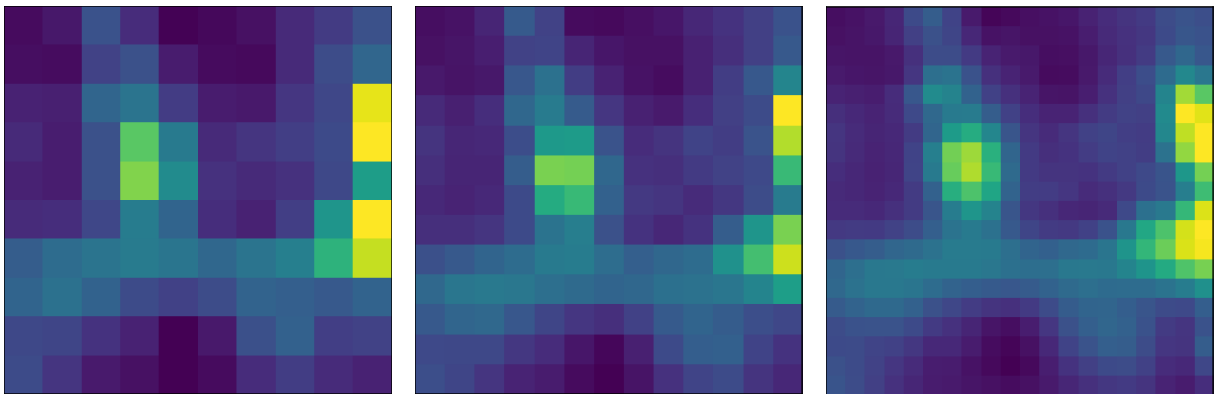
(A) Sampling 5% (B) Sampling 10% (C) Sampling 15% (D) Sampling 20%

(E) OWL-QN 5% (F) OWL-QN 10% (G) OWL-QN 15% (H) OWL-QN 20%

(I) RBF 5% (J) RBF 10% (K) RBF 15% (L) RBF 20%

(M) RMSE of the two methods

FIGURE 5.6: Reconstruction results for a 40×40 image with the two methods for different sampling ratios at a laser power of $750\,\mathrm{mW}$. The first line contains the samples used for reconstruction in the other two lines. These samples are taken randomly at a given ratio, from 5% to 20%, compared to the full image. They are also such that each increase of samples always includes the points from the previous image. For the OWL-QN algorithm, the regularisation parameter is set to $C = 0.1$.



(A) 100 samples (B) 169 samples (C) 400 samples

FIGURE 5.7: Sample at 6 kHz, with regular sampling from 100 to 400 samples. (a) is a 10×10 image of step size 6 µm, (b) 13×13 and 4.5 µm, (c) 20×20 and 2 µm.

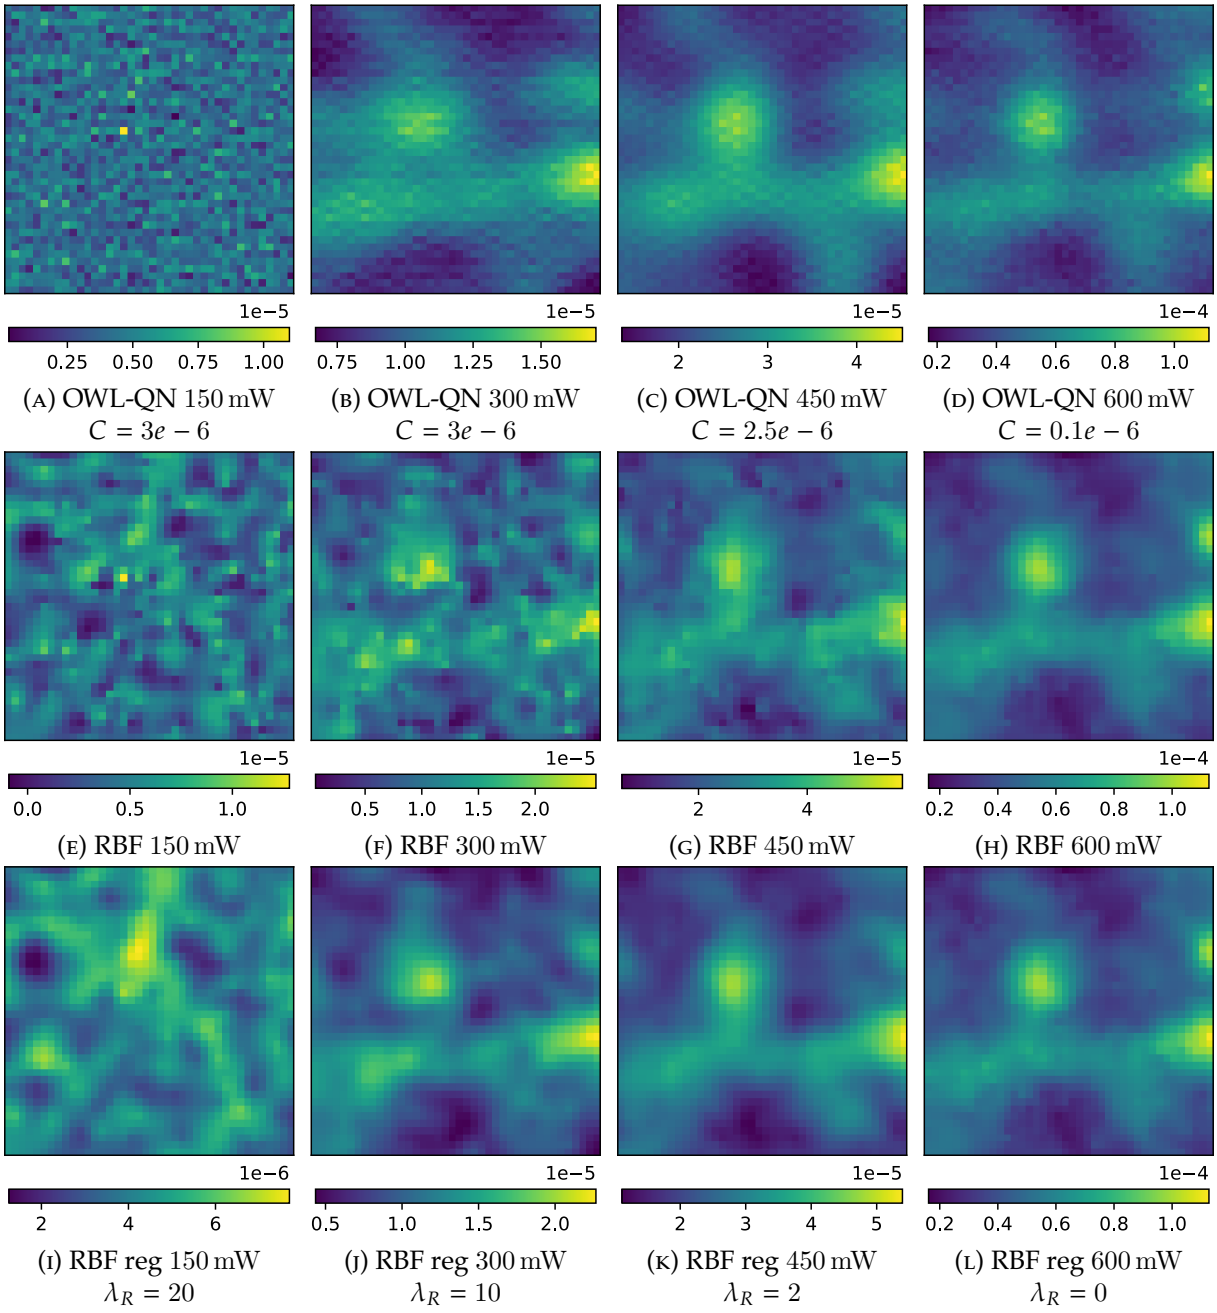
5.4 Comparison on noise levels

The robustness to noise was also studied. Successive images were made with decreasing laser input power, ranging from 1 W to 150 mW, which increases the relative noise levels accordingly. The sampling is exactly the same in every image, with a ratio of 25%. Some of the results are shown in Figure 5.8. The OWL-QN result has a better visual fidelity compared to the interpolation result especially at 300 mW and 450 mW, where it is easier to distinguish the structure, although the range of values is lowered. However, by using by adding a regularisation parameter to the RBF, the resulting approximation manages to give acceptable resulting images at 300 mW at the expense of an added bias to the values of amplitude. At the highest noise levels, ie. lowest power, both methods are unable to reconstruct the signal.

A comparison with a single graph as was done in the previous section is less straightforward, as a single reference image can't be used since the amplitudes are different at every noise level, and a normalisation of the values would mostly show differences in contrast rather than fidelity of the image. This problem was studied in the field of image processing [121] with the development of the structural similarity index measure (SSIM) that combines luminance, contrast and structure with a weight given to those three parameters. Here, it is not clear what weight should be given to the contrast in particular, as on one hand for qualitative measurements, a simple rescale of the colorbar is sufficient, as was done for instance in Figure 5.8b where the structure is visible but partly due to the fact that the color scale is narrower, but on the other hand it is impossible on that same image to give accurate quantitative values.

5.5 Multiple frequencies scan

As described in the methods section, a weighted random resampling is deduced from the previous reconstruction at 6 kHz, and then used to scan the sample at a set of frequencies in the range 1 kHz to 100 kHz. These images are compared to the fully random

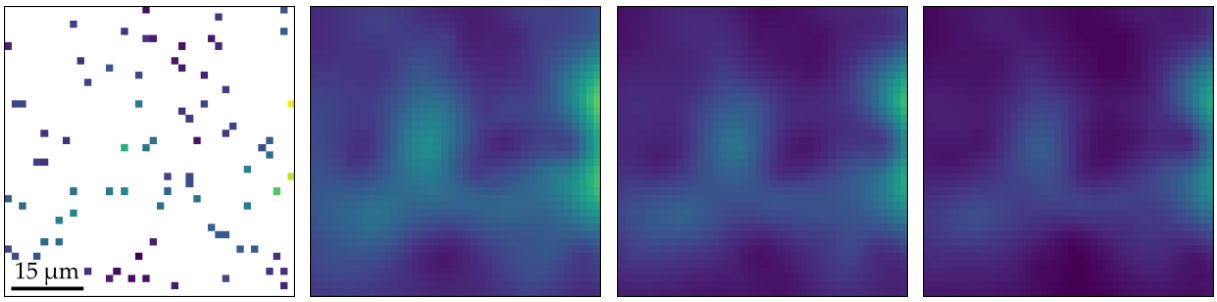


(A) OWL-QN 150 mW $C = 3e-6$ (B) OWL-QN 300 mW $C = 3e-6$ (C) OWL-QN 450 mW $C = 2.5e-6$ (D) OWL-QN 600 mW $C = 0.1e-6$

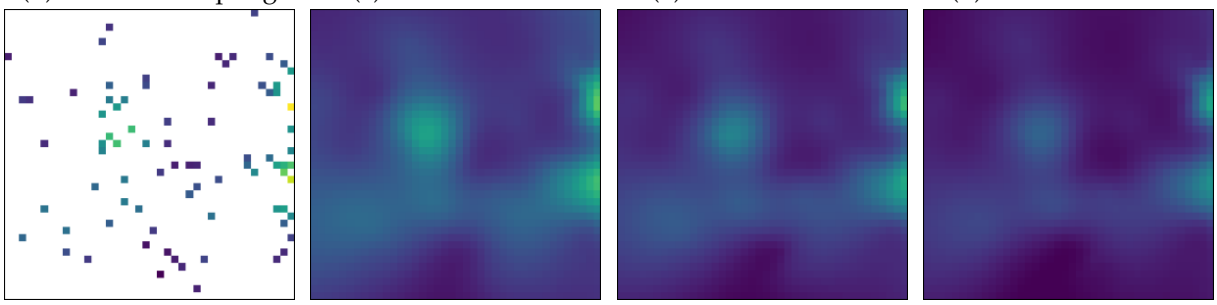
(E) RBF 150 mW (F) RBF 300 mW (G) RBF 450 mW (H) RBF 600 mW

(I) RBF reg 150 mW $\lambda_R = 20$ (J) RBF reg 300 mW $\lambda_R = 10$ (K) RBF reg 450 mW $\lambda_R = 2$ (L) RBF reg 600 mW $\lambda_R = 0$

FIGURE 5.8: Reconstruction results for a 40×40 image with the two methods for different laser powers at a sampling ratio of 25%. For OWL-QN and the regularised RBF method the associated parameters C, and λ_R are given.



(A) Random sampling (B) Random 1 kHz (C) Random 40 kHz (D) Random 100 kHz



(E) Weighted sampling (F) Weighted 1 kHz (G) Weighted 40 kHz (H) Weighted 100 kHz

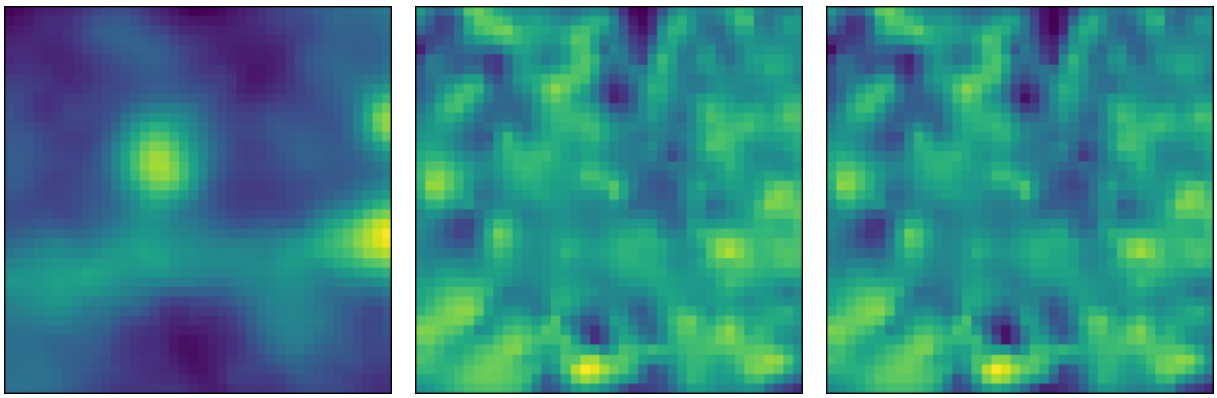
FIGURE 5.9: Reconstruction results for a 40×40 image at multiple frequencies.

sampling used previously at the same sampling ratio of 5% and a laser power of $750\,\mathrm{mW}$ in Figure 5.9. Both lead to acceptable results in the given case, although the weighted random method trends to a give slightly better contrast.

Ultimately, both the minimum sampling ratio and the choice of the coefficient α in the probability to draw samples depends on the material and the type of defects that the user expects. A low value of α means that the entire sample is supposed to be similar and therefore can lead to the non detection of defects such as the apparition of cracks at a specific given depth, especially if said depth is further than the one associated with the first frequency of the calculated gradient. Thus it is primordial in NDT to carefully chose correct hypotheses as under sampling necessarily creates a compromise between acquisition time and fidelity. The weighted random sampling can easily be modified by changing the probability function in order to better represent other scenarios. For instance in the case studied here, if the objective was to look for defects within the aluminium, the probability could be defined as the inverse of what is shown here, ie. be maximum when the gradient is minimum.

5.6 Conclusion and Perspectives

SPR is a method that allows for very precise thermal measurements and can largely be accelerated thanks to algorithmic reconstruction of images from a small subset of samples. With measurements made on a Carbon fibre Aluminum matrix composite, both RBF interpolation and CS have been used to accurately represent images with $1/6$ of the samples that would have been required for a traditional sampling scheme. Thus an equivalent reduction in time was achieved, going from roughly $1\,\mathrm{h}$ to $10\,\mathrm{min}$ for one



(A) Amplitude (B) Phase (C) Phase (direct)

FIGURE 5.10: Sample at 6 kHz

frequency in order to reconstruct a 40×40 image with an effective step size of $1.5\,\mu m$. RBF interpolation led to overall better results with lower sampling ratios (as low as $1/20$) while still performing well when confronted to noisy data. The new proposed sampling scheme drawing samples according to the image gradient in order to scan a wide array of frequencies and therefore depths opens the way to a further reduction in acquisition time.

Reconstructions were illustrated here with the amplitude signal. The phase also contains a lot of information in SPR and thus could be treated similarly. However, a couple of additional issues arise with it. First, the images trend to be noisier which, as shown in Figure 5.8, has a detrimental impact on the reconstruction results. The second issue is the possibility of presence of extremely sharp discontinuities whenever the phase change is near $\pm 180°$. To attempt to circumvent this problem, reconstructions on the real and imaginary part of the complex signal can be made and then calculated back in phase and amplitude as shown in Figure 5.10, where (a) and (b) are the reconstruction made in such manner, while (c) is the reconstruction of the phase signal directly, ie. as given by the lock-in. The two phase reconstructions are almost identical, and both show that the data is too noisy to accurately detect defects, even when the amplitude isn't.

Further improvements can be achieved by applying the undersampling not only on the spatial scanning but also on the frequency signal, as successive frequencies are strongly correlated. A more precise localisation of defects could also be made by correctly characterising the photothermal PSF, as was proposed by Seidel et al. [122] in one of the first attempt at photothermal imaging.

For the sake of speed SPR measurements done here were not repeated, which is the usual procedure. Although more costly in time, this allows for the calculation of standard deviations which gives additional information that is used for instance to aim for a given precision. This can still be done with under sampling, however it is not possible to give an uncertainty value for every pixel after reconstruction. Bayesian based methods seems to be the best path to follow, but even then since CS is based on setting values to zero in a transform space, this has a non-linear effect on noise, which might make this problem unsolvable.

Chapter 6
Conclusion and perspectives

Contents

6.1 Conclusion 96
6.2 Perspectives 96

6.1 Conclusion

The objectives were to investigate super-resolution methods in order to find which ones would make sense in the context of thermal imaging and then implement them. At their core, super-resolution and under-sampling are compromises for which there never will be a perfect and universal answer, thus many methods have been developed over the years with a wide variety of constraints. In this thesis, the focus was made on so called label-free methods in order to stay as non-intrusive as possible, and especially answers to the sampling limit with reconstructions of images from non-regular samplings, especially CS.

In chapter 2 the main limits of resolution were described by separating them in two main groups, first the diffraction limit as described by Abbe, then the sampling limit known as Shannon's limit, and eventually for thermal signals a diffusion limit, that behaves similarly to the diffraction one with the notable difference that it is usually time-dependant. An overview of the approaches already used in the literature to overcome these limits.

In chapter 3, a few approaches for super-resolution were explored. Starting with spectral filtering to extract aliases and reconstruct higher resolution images, followed by experimentations with random structured illumination. Finally, a method using a camera and masks containing points smaller than the pixels was explored to reconstruct images with a higher resolution than the camera corresponding to the size of the masks points. This method was tested using multiple variants, on the structure of the masks patterns, using random distributions and one following Hadamard modes of various frequencies, as well as using two type of experimental ways to structure the outgoing sample: One using a grid of small controllable mirrors, and the other using a series of black and transparent masks that can be changed by displacing them on a motor going side to side. Some first preliminary results were achieved using visible light, but more work is required still.

In chapter 4 and 5 undersampling and image reconstructions on point-by-point measurement was achieved using two setups, respectively FS and SPR. In 4, CS was used and three algorithms with different behaviour were compared. In these, the OWL-QN algorithm performed better and allowed to reconstruct images with 20% of the samples while still remaining close to the ground truth. In 5 the comparison was made in another context of micro-scale measurements using SPR between CS and RBF interpolation/approximation. In that context, the RBF approximation allowed for reconstruction of images even with extremely low sampling ratios, as low as 5% when the noise is low. Additionally optimisations of the sampling was proposed whenever imaging at a set of frequencies. In both SPR and FS cases, the reduction in number of samples corresponds directly to an equivalent reduction in time.

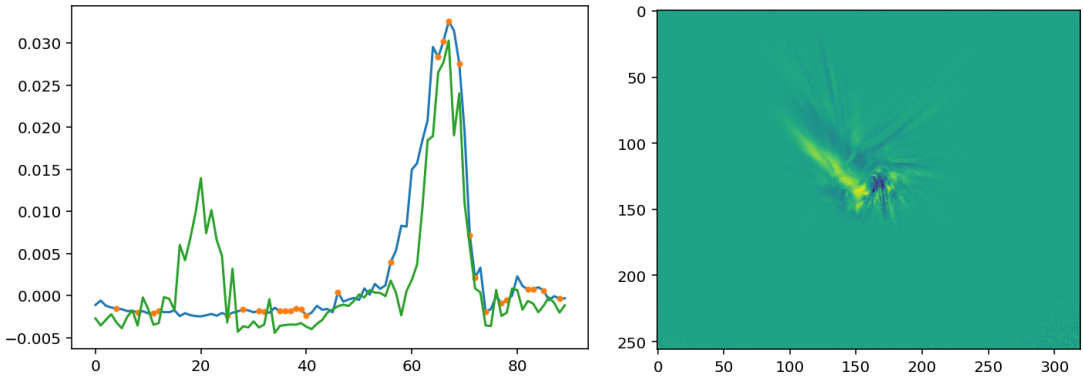
6.2 Perspectives

These are firsts steps in the lab towards super-resolution and more work can be done on a wide variety of fronts. It should be kept in mind that a lot of similitudes in the methods exists between this and other inverse methods used in the lab for instance

with reconstruction of thermal and acoustic sources, thus advances on these might (and probably should) go hand in hand.

Regarding the under-sampling methods, beside what was already mentioned as perspectives for chapter 4 and 5, it could be interesting to investigate their applications in other point-by-point methods, most notably SThM, as both the principle and the type of images is similar.

As for super-resolution, the results from chapter 3 are not yet satisfying, more work has to be made especially on the computational side in order to make super-resolution with cameras where every single pixel is subdivided viable. Another possibility related to this work would be to fully embrace the low-tech/low-cost aspect of SPC and aim for a very affordable IR camera.



(A) Graph at one pixel (B) Reconstructed image

FIGURE 6.1: Laminography of a laser beam heating a 3D sample at an angle and a naïve attempt at CS. (a) shows the values at one pixel of the camera for 90 regularly spaced angles (blue), a random subset of 50% (orange) and CS reconstruction from this (green). (b) shows a reconstructed image from one of the non-sampled spots, which clearly exhibits artefacts

Another suggestion that was made during this work was to explore both undersampling and super-resolution in the context of laminography, as this is a subject of interest within the ICT team and that some of the issues encountered are similar, especially regarding the long acquisition times and the large amount of data that has to be processed. Laminography works by taking a series of measurements at different angles in order to reconstruct 3D images from a set of slices. A quick test was made in order to reduce the number of angles that need to be acquired by undersampling them, thus taking a subset of image at random angles and then using CS on each pixels as was shown previously in this thesis. This test, in Figure 6.1, shows that this not that the method cannot be that straightforward as the minimisation struggles to work simultaneously for all pixels, leading to major artefacts. Existing work related to CS tomography and laminography exists, but usually relies on the sparsity of the gradient of the result, which is not applicable with the thermal images here.

The methods presented here and machine learning are very close, if not overlapping and are all part of information theory. Indeed, in AI the saying "compression is (artificial) intelligence" is widely used. It would thus be logical to seek to continue this work by looking at other machine learning methods. In this thesis the decision was made

to tread carefully regarding these, especially due to the many questions that arise with training. If this path is take, a few advices can be given such as starting with a specific setup in order to narrow as much as possible the training step before going to more complex cases, as well as going for unsupervised methods as the labelling of images is time consuming, especially when for thermal imaging the amount of already existing dataset might be close to non-existent.

Back in 1956, Claude Shannon wrote a short editorial called "The Bandwagon", in which he warned that information theory could be victim of its own popularity and that it cannot replace the whole scientific method of every field. His advice was simply to "keep our own house in order" by doing slow and steady research, meaning publishing less but rather focus on quality. This still very much echoes through today and should be kept in mind when going towards the very trendy (at the time of writing) field of machine learning in particular.

Bibliography

[1] F. Crouau, J. Maire, J.-L. Battaglia et al., 'Undersampled flying spot thermography using compressive sensing', *NDT & E International*, vol. 163, p. 103 780, 2026-08-01, ISSN: 0963-8695. DOI: 10.1016/j.ndteint.2026.103780 (cit. on pp. xii, 82).

[2] F. Crouau, A. Mateos-Canseco, J. Maire et al., 'Thermal characterisation by Scanning Photothermal Radiometry using a random undersampled measurement scheme', 2026-06. Accessed: 2026-06-18. [Online]. Available: https://hal.science/hal-05651735 (cit. on pp. xii, 68).

[3] A. Mateos-Canseco, 'Imaging and thermal characterization of materials by scanning photo-thermal radiometry', Thesis (cit. on pp. 5, 19, 83, 84, 86).

[4] J. W. Goodman, *Introduction to Fourier Optics*, Fourth edition. New York: W.H. Freeman, Macmillan Learning, 2017, 546 pp., ISBN: 978-1-319-11916-4 (cit. on pp. 9, 10, 12, 17, 18).

[5] 'Thorlabs · Imaging/Focusing Microscope Objectives, Oil Immersion', Accessed: 2026-06-05. [Online]. Available: https://www.thorlabs.com/imagingfocusing-microscope-objectives-oil-immersion (cit. on p. 14).

[6] K. Kadela, H. Oleś, K. Kosowska et al., 'High-resolution infrared imaging performance of new high numerical aperture Schwarzschild objectives', *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy*, vol. 331, p. 125 769, 2025-04-15, ISSN: 1386-1425. DOI: 10.1016/j.saa.2025.125769 (cit. on pp. 14, 54).

[7] J.-F. Giovannelli. 'An inverse problem course – Ressources pédagogiques du Gretsi', Accessed: 2026-06-05. [Online]. Available: https://gretsi.fr/ressources-pedagogiques/an-inverse-problem-course/ (cit. on pp. 16, 28).

[8] F. Orieux, J.-F. Giovannelli and T. Rodet, 'Bayesian estimation of regularization and point spread function parameters for Wiener-Hunt deconvolution', *Journal of the Optical Society of America. A Optics, Image Science, and Vision*, p. 1593, 2010-06. DOI: 10.1364/JOSAA.27.001593 (cit. on p. 16).

[9] W. H. Richardson, 'Bayesian-Based Iterative Method of Image Restoration*', *Journal of the Optical Society of America*, 1972. DOI: 10.1364/JOSA.62.000055 (cit. on p. 16).

[10] T. R. Corle and G. S. Kino, *Confocal Scanning Optical Microscopy and Related Imaging Systems*. San Diego: Academic Press, 1996, 335 pp., ISBN: 978-0-12-408750-7 (cit. on pp. 16, 18, 19).

[11] R. T. Borlinghaus, 'Pinhole Effect in Confocal Microscopes', 2017-04-26. Accessed: 2026-06-05. [Online]. Available: https://www.leica-microsystems.com/science-lab/life-science/pinhole-effect-in-confocal-microscopes/ (cit. on p. 17).

[12] J. Mertz, *Introduction to Optical Microscopy*, 2nd ed. Cambridge University Press, 2019-07-31, ISBN: 978-1-108-55266-0 978-1-108-42830-9. DOI: 10.1017/9781108552660 (cit. on p. 17).

[13] E. Maletz. 'Scientific Image - Atomic Force Microscope Illustration | NISE Network', Accessed: 2026-06-06. [Online]. Available: https://www.nisenet.org/catalog/scientific-image-atomic-force-microscope-illustration (cit. on p. 19).

[14] L. Cui, W. Jeong, V. Fernández-Hurtado et al., 'Study of radiative heat transfer in Ångström- and nanometre-sized gaps', *Nat Commun*, vol. 8, no. 1, p. 14 479, 2017-02-15, ISSN: 2041-1723. DOI: 10.1038/ncomms14479 (cit. on p. 19).

[15] R. J. Ober, S. Ram and E. S. Ward, 'Localization Accuracy in Single-Molecule Microscopy', *Biophysical Journal*, vol. 86, no. 2, pp. 1185–1200, 2004-02-01, ISSN: 0006-3495. DOI: 10.1016/S0006-3495(04)74193-4 (cit. on p. 18).

[16] S. W. Hell and J. Wichmann, 'Breaking the diffraction resolution limit by stimulated emission: Stimulated-emission-depletion fluorescence microscopy', *Opt. Lett.*, vol. 19, no. 11, p. 780, 1994-06-01, ISSN: 0146-9592, 1539-4794. DOI: 10.1364/OL.19.000780 (cit. on p. 19).

[17] E. Betzig, G. H. Patterson, R. Sougrat et al., 'Imaging intracellular fluorescent proteins at nanometer resolution', *Science*, vol. 313, no. 5793, pp. 1642–1645, 2006-09-15, ISSN: 00368075. DOI: 10.1126/science.1127344. PMID: 16902090 (cit. on p. 19).

[18] M. J. Rust, M. Bates and X. Zhuang, 'Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM)', *Nat Methods*, vol. 3, no. 10, pp. 793–796, 2006-10, ISSN: 1548-7105. DOI: 10.1038/nmeth929 (cit. on p. 19).

[19] G. Zheng, R. Horstmeyer and C. Yang, 'Wide-field, high-resolution Fourier ptychographic microscopy', *Nature Photon*, vol. 7, no. 9, pp. 739–745, 2013-09, ISSN: 1749-4885, 1749-4893. DOI: 10.1038/nphoton.2013.187 (cit. on p. 20).

[20] G. Zheng, C. Shen, S. Jiang et al., 'Concept, implementations and applications of Fourier ptychography', *Nature Reviews Physics*, vol. 3, no. 3, pp. 207–223, 2021-02-10, ISSN: 2522-5820. DOI: 10.1038/s42254-021-00280-y (cit. on pp. 20, 21).

[21] M. Tamamitsu, K. Toda, M. Fukushima et al., 'Mid-infrared wide-field nanoscopy', *Nature Photonics*, 2024-04-17, ISSN: 1749-4885. DOI: 10.1038/s41566-024-01423-0 (cit. on p. 21).

[22] M. G. L. Gustafsson, 'Surpassing the lateral resolution limit by a factor of two using structured illumination microscopy. SHORT COMMUNICATION', *J Microsc*, vol. 198, no. 2, pp. 82–87, 2000-05, ISSN: 0022-2720, 1365-2818. DOI: 10.1046/j.1365-2818.2000.00710.x (cit. on pp. 21, 48).

[23] E. Mudry, K. Belkebir, J. Girard et al., 'Structured illumination microscopy using unknown speckle patterns', *Nature Photon*, vol. 6, no. 5, pp. 312–315, 2012-05, ISSN: 1749-4885, 1749-4893. DOI: 10.1038/nphoton.2012.83 (cit. on pp. 22, 48).

[24] L. Mazzella, T. Mangeat, G. Giroussens et al., 'Extended-depth of field random illumination microscopy, EDF-RIM, provides super-resolved projective imaging', *Light Sci Appl*, vol. 13, no. 1, p. 285, 2024-10-10, ISSN: 2047-7538. DOI: 10.1038/s41377-024-01612-0 (cit. on p. 22).

[25] S. Labouesse, J. Idier, A. Sentenac et al., 'Random Illumination Microscopy from Variance Images', in *2020 28th European Signal Processing Conference (EUSIPCO)*, IEEE, 2021-01-24, pp. 785–789, ISBN: 978-90-827970-5-3. DOI: 10.23919/Eusipco47968.2020.9287651 (cit. on p. 22).

[26] L.-H. Yeh, S. Chowdhury and L. Waller, 'Computational structured illumination for high-content fluorescence and phase microscopy', *Biomedical Optics Express*, vol. 10, no. 4, p. 1978, 2019-04-01, ISSN: 2156-7085. DOI: 10.1364/BOE.10.001978 (cit. on pp. 23, 48, 52, 53).

[27] J. Idier, S. Labouesse, M. Allain et al., 'On the Superresolution Capacity of Imagers Using Unknown Speckle Illuminations', *IEEE Transactions on Computational Imaging*, vol. 4, no. 1, pp. 87–98, 2018-03, ISSN: 2333-9403. DOI: 10.1109/TCI.2017.2771729 (cit. on p. 23).

[28] S. Dong, P. Nanda, K. Guo et al., 'Incoherent Fourier ptychographic photography using structured light', *Photon. Res.*, vol. 3, no. 1, p. 19, 2015-02-01, ISSN: 2327-9125. DOI: 10.1364/PRJ.3.000019 (cit. on p. 23).

[29] T. Mangeat, S. Labouesse, M. Allain et al., 'Super-resolved live-cell imaging using random illumination microscopy', *Cell Reports Methods*, vol. 1, no. 1, p. 100 009, 2021-05, ISSN: 26672375. DOI: 10.1016/j.crmeth.2021.100009 (cit. on p. 23).

[30] K. Fliegel, 'Modeling and Measurement of Image Sensor Characteristics', *Radioengineering*, vol. 13, no. 4, 2004, ISSN: 1805-9600 (cit. on p. 24).

[31] C. Shannon, 'Communication in the Presence of Noise', *Proceedings of the IRE*, vol. 37, no. 1, pp. 10–21, 1949-01, ISSN: 2162-6634. DOI: 10.1109/JRPROC.1949.232969 (cit. on p. 25).

[32] N. Zhao, Q. Wei, A. Basarab et al. 'Fast Single Image Super-Resolution'. arXiv: 1510.00143 [cs], pre-published (cit. on p. 28).

[33] J. Claerbout. 'Basic Earth Imaging', Jon Claerbout's Classroom, Accessed: 2026-06-07. [Online]. Available: https://sep.stanford.edu/sep/prof/bei/toc_html/index.html (cit. on p. 29).

[34] 'B-splines and convolution » Chebfun', Accessed: 2026-06-07. [Online]. Available: https://www.chebfun.org/examples/approx/BSplineConv.html (cit. on p. 29).

[35] C. De Boor, 'On calculating with B-splines', *Journal of Approximation Theory*, vol. 6, no. 1, pp. 50–62, 1972-07-01, ISSN: 0021-9045. DOI: 10.1016/0021-9045(72)90080-9 (cit. on p. 29).

[36] '3.2: Polynomial Interpolation', Mathematics LibreTexts, Accessed: 2026-06-07. [Online]. Available: https://math.libretexts.org/Courses/Angelo_State_University/Mathematical_Computing_with_Python/3%3A_Interpolation_and_Curve_Fitting/3.2%3A_Polynomial_Interpolation (cit. on p. 29).

[37] Q. Agrapart and A. Batailly, 'Cubic and bicubic spline interpolation in Python', 2020-11-20 (cit. on p. 29).

[38] P. Walls. 'Natural Cubic Spline Interpolation — Applied Linear Algebra', Accessed: 2026-06-07. [Online]. Available: https://ubcmath.github.io/MATH307/notebooks/04_spline_interpolation.html (cit. on p. 29).

[39] *Bicubic interpolation*, in *Wikipedia*, 2026-04-25. Accessed: 2026-06-07. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Bicubic_interpolation&oldid=1350972429#Comparison_with_other_methods (cit. on p. 30).

[40] C. A. Micchelli, 'Interpolation of scattered data: Distance matrices and conditionally positive definite functions', *Constr. Approx*, vol. 2, no. 1, pp. 11–22, 1986-12-01, ISSN: 1432-0940. DOI: 10.1007/BF01893414 (cit. on p. 30).

[41] E. Viennet, 'Réseaux à fonctions de base radiales', (cit. on p. 30).

[42] G. E. Fasshauer, *Meshfree Approximation Methods with MATLAB* (Interdisciplinary Mathematical Sciences v. 6). Singapore ; Hackensack, N.J: World Scientific, 2007, 500 pp., ISBN: 978-981-270-634-8 978-981-270-633-1 (cit. on pp. 31, 89).

[43] M. Pyrcz, *Applied Geostatistics in Python: A Hands-on Guide with GeostatsPy [e-Book]*. Zenodo, 2025-04-07. DOI: 10.5281/ZENODO.15169133 (cit. on p. 31).

[44] V. S. Fazio and M. Roisenberg, 'Spatial interpolation: An analytical comparison between kriging and RBF networks', in *Proceedings of the 28th Annual ACM Symposium on Applied Computing*, Coimbra Portugal: ACM, 2013-03-18, pp. 2–7, ISBN: 978-1-4503-1656-9. DOI: 10.1145/2480362.2480364 (cit. on p. 32).

[45] D. Donoho, 'Compressed sensing', *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006-04, ISSN: 0018-9448. DOI: 10.1109/TIT.2006.871582 (cit. on p. 32).

[46] E. Candes, J. Romberg and T. Tao, 'Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information', *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006-02, ISSN: 1557-9654. DOI: 10.1109/TIT.2005.862083 (cit. on p. 32).

[47] E. J. Candes and T. Tao, 'Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies?', *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006-12, ISSN: 0018-9448. DOI: 10.1109/TIT.2006.885507 (cit. on p. 32).

[48] E. Candes and M. Wakin, 'An Introduction To Compressive Sampling', *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, 2008-03, ISSN: 1053-5888. DOI: 10.1109/MSP.2007.914731 (cit. on pp. 32, 61).

[49] S. L. Brunton and J. N. Kutz, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2019 (cit. on p. 32).

[50] M. Rani, S. B. Dhok and R. B. Deshmukh, 'A Systematic Review of Compressive Sensing: Concepts, Implementations and Applications', *IEEE Access*, vol. 6, pp. 4875–4894, 2018, ISSN: 2169-3536. DOI: 10.1109/ACCESS.2018.2793851 (cit. on p. 35).

[51] M. F. Duarte, M. A. Davenport, D. Takhar et al., 'Single-pixel imaging via compressive sampling', *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 83–91, 2008-03, ISSN: 1053-5888. DOI: 10.1109/MSP.2007.914730 (cit. on pp. 35, 40).

[52] R. G. Baraniuk, 'Compressive sensing [lecture notes]', *IEEE signal processing magazine*, vol. 24, no. 4, pp. 118–121, 2007. Accessed: 2026-05-26. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/4286571/ (cit. on p. 35).

[53] V. M. Patel, G. R. Easley, D. M. Healy et al., 'Compressed Synthetic Aperture Radar', *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 244–254, 2010-04, ISSN: 1941-0484. DOI: 10.1109/JSTSP.2009.2039181 (cit. on p. 35).

[54] L. Kovarik, A. Stevens, A. Liyu et al., 'Implementing an accurate and rapid sparse sampling approach for low-dose atomic resolution STEM imaging', *Applied Physics Letters*, vol. 109, no. 16, 2016-10-17, ISSN: 0003-6951. DOI: 10.1063/1.4965720 (cit. on p. 35).

[55] Y. Niu and G. Han, 'Fast AFM Imaging Based on Compressive Sensing Using Undersampled Raster Scan', *IEEE Transactions on Instrumentation and Measurement*, 2021. DOI: 10.1109/TIM.2020.3023215 (cit. on p. 35).

[56] X. Yuan, D. J. Brady and A. K. Katsaggelos, 'Snapshot Compressive Imaging: Theory, Algorithms, and Applications', *IEEE Signal Process. Mag.*, vol. 38, no. 2, pp. 65–88, 2021-03, ISSN: 1053-5888, 1558-0792. DOI: 10.1109/MSP.2020.3023869 (cit. on p. 35).

[57] M. Lustig, D. Donoho and J. M. Pauly, 'Sparse MRI: The application of compressed sensing for rapid MR imaging', *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007, ISSN: 1522-2594. DOI: 10.1002/mrm.21391 (cit. on p. 35).

[58] R. Leary, Z. Saghi, P. A. Midgley et al., 'Compressed sensing electron tomography', *Ultramicroscopy*, vol. 131, pp. 70–91, 2013-08, ISSN: 03043991. DOI: 10.1016/j.ultramic.2013.03.019 (cit. on p. 36).

[59] L. Zhang, J. Ke, S. Chi et al., 'High-resolution fast mid-wave infrared compressive imaging', *Optics Letters*, vol. 46, no. 10, p. 2469, 2021-05-15, ISSN: 0146-9592. DOI: 10.1364/ol.420481. PMID: 33988612 (cit. on p. 36).

[60] G. M. Gibson, S. D. Johnson and M. J. Padgett, 'Single-pixel imaging 12 years on: A review', *Opt. Express, OE*, vol. 28, no. 19, pp. 28 190–28 208, 2020-09-14, ISSN: 1094-4087. DOI: 10.1364/OE.403195 (cit. on p. 36).

[61] O. Katz, Y. Bromberg and Y. Silberberg, 'Compressive ghost imaging', *Applied Physics Letters*, vol. 95, no. 13, p. 131 110, 2009-09-28, ISSN: 0003-6951, 1077-3118. DOI: 10.1063/1.3238296. arXiv: 0905.0321 [quant-ph] (cit. on p. 36).

[62] A. Beck and M. Teboulle, 'A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems', *SIAM J. Imaging Sci.*, vol. 2, no. 1, pp. 183–202, 2009-01-01, ISSN: 1936-4954. DOI: 10.1137/080716542 (cit. on pp. 36, 74).

[63] C. W. Royer, 'Lecture notes on proximal methods and regularization', 2022. [Online]. Available: https://www.lamsade.dauphine.fr/~croyer/ensdocs/PGL/LectureNotesOML-PGL.pdf (cit. on p. 36).

[64] G. Andrew and J. Gao, 'Scalable training of L1-regularized log-linear models', *Proceedings of the 24th international conference on Machine learning*, 2007. DOI: 10.1145/1273496.1273501 (cit. on pp. 37, 74, 88).

[65] Y. Chen, 'Lecture 23: Limited-Memory BFGS (L-BFGS)', 2023. [Online]. Available: https://pages.cs.wisc.edu/~yudongchen/cs726_sp25/ (cit. on p. 37).

[66] C. Li, 'An Efficient Algorithm For Total Variation Regularization with Applications to the Single Pixel Camera and Compressive Sensing', Theses, Rice University, 2011. HDL: 1911/62229. [Online]. Available: https://hdl.handle.net/1911/62229 (cit. on pp. 38–40, 74, 75).

[67] A. Mahalanobis, R. Shilling, R. Murphy et al., 'Recent results of medium wave infrared compressive sensing', *Applied Optics*, vol. 53, no. 34, p. 8060, 2014-12-01, ISSN: 0003-6935. DOI: 10.1364/AO.53.008060 (cit. on pp. 40, 60).

[68] X.-P. Jin, A.-D. Xiong, X.-Q. Wang et al., 'Long-distance mid-wave infrared super-resolution compressive imaging', *Optics & Laser Technology*, vol. 157, p. 108 740, 2023-01-01, ISSN: 0030-3992. DOI: 10.1016/j.optlastec.2022.108740 (cit. on pp. 40, 54, 56, 62, 80).

[69] A. Vallés, J. He, S. Ohno et al., 'Broadband high-resolution terahertz single-pixel imaging', *Optics Express*, vol. 28, no. 20, p. 28 868, 2020-09-28, ISSN: 1094-4087. DOI: 10.1364/OE.404143 (cit. on pp. 40, 64).

[70] Y. Huang, T. Miyazaki, X. Liu et al. 'Infrared Image Super-Resolution: Systematic Review, and Future Trends'. arXiv: 2212.12322 [eess], pre-published (cit. on p. 40).

[71] J. Lecompagnon, 'Nondestructive defect characterization using full-frame spatially structured super resolution laser thermography', Thesis, Technischen Universität Berlin, 2023. DOI: 10.14279/depositonce-19271 (cit. on p. 41).

[72] J. Lecompagnon, P. Hirsch, C. Rupprecht et al., 'Influence of the number of measurements on detecting internal defects using photothermal super resolution reconstruction with random pixel patterns', *Quantitative InfraRed Thermography Journal*, vol. 21, no. 5, pp. 277–287, 2024-10-19, ISSN: 1768-6733, 2116-7176. DOI: 10.1080/17686733.2023.2223392 (cit. on p. 41).

[73] M. M. Groz, A. Sommier, E. Abisset-Chavanne et al., 'Thermal resistance field estimations from IR thermography using multiscale Bayesian inference', *Quantitative InfraRed Thermography Journal*, 2020. DOI: 10.1080/17686733.2020.1771529 (cit. on p. 41).

[74] S. Ahmadi, G. Thummerer, S. Breitwieser et al., 'Multidimensional Reconstruction of Internal Defects in Additively Manufactured Steel Using Photothermal Super Resolution Combined With Virtual Wave-Based Image Processing', *IEEE Trans. Ind. Inf.*, vol. 17, no. 11, pp. 7368–7378, 2021-11, ISSN: 1551-3203, 1941-0050. DOI: 10.1109/TII.2021.3054411 (cit. on p. 41).

[75] P. Burgholzer, M. Thor, J. Gruber et al., 'Three-dimensional thermographic imaging using a virtual wave concept', *Journal of Applied Physics*, vol. 121, no. 10, 2017-03-14, ISSN: 0021-8979. DOI: 10.1063/1.4978010 (cit. on p. 41).

[76] M. Groz, E. Abisset-Chavanne, J. Batsale et al., 'Active thermal super-resolution based on laser flying spot technique coupled with IR thermography and wavelet transform', *Quantitative InfraRed Thermography Journal*, pp. 1–24, 2024-01-03, ISSN: 1768-6733. DOI: 10.1080/17686733.2023.2299499 (cit. on pp. 41, 69).

[77] H. Yang, Z. Zhu, Z. Xie et al., 'Thermal-property microscopy with compressive-sensing frequency-domain thermoreflectance', *Physical Review Applied*, vol. 24, no. 1, p. 014 033, 2025-07-16, ISSN: 2331-7019. DOI: 10.1103/jqgf-zgtq (cit. on pp. 41, 84).

[78] V. N. Astratov, Y. B. Sahel, Y. C. Eldar et al., 'Roadmap on Label-Free Super-Resolution Imaging', *Laser & Photonics Reviews*, vol. 17, no. 12, p. 2 200 029, 2023-12, ISSN: 1863-8880, 1863-8899. DOI: 10.1002/lpor.202200029 (cit. on p. 42).

[79] S. Urbaś and B. Więcek, 'Application of a deep-learning neural network for image reconstruction from a single-pixel infrared camera', *Opto-Electronics Review*, pp. 148 877–148 877, 2024-01-17, ISSN: 1896-3757. DOI: 10.24425/opelre.2024.148877 (cit. on p. 42).

[80] M. Zhou, H. Chen, J. Paisley et al., 'Non-Parametric Bayesian Dictionary Learning for Sparse Image Representations', (cit. on p. 42).

[81] C. Bourges, 'Thermography in semitransparent media based on mid-infrared thermotransmittance', Theses, Université de Bordeaux, 2023-09. Accessed: 2026-05-06. [Online]. Available: https://theses.hal.science/tel-04228284 (cit. on p. 44).

[82] C. Bourges, S. Chevalier, J. Maire et al., 'Infrared thermotransmittance-based temperature field measurements in semitransparent media', *Review of Scientific Instruments*, vol. 94, no. 3, p. 034 905, 2023-03. DOI: 10.1063/5.0131422 (cit. on p. 44).

[83] C. Pradere, M. Ryu, A. Sommier et al., 'Non-contact temperature field measurement of solids by infrared multispectral thermotransmittance', *J. Appl. Phys.*, vol. 121, no. 8, p. 085 102, 2017-02-23, ISSN: 0021-8979. DOI: 10.1063/1.4976209 (cit. on p. 44).

[84] J. Letessier, A. Netter, J. Maire et al., 'Infrared Photothermal Heterodyne Imaging in Thermally Thick Medium for Thermo-optic Property Characterizations', *J. Phys. Chem. Lett.*, vol. 16, no. 46, pp. 11 987–11 995, 2025-11-20. DOI: 10.1021/acs.jpclett.5c01503 (cit. on p. 44).

[85] C. Leung and T. D. Donnelly, 'Measuring the spatial resolution of an optical system in an undergraduate optics laboratory', *American Journal of Physics*, vol. 85, no. 6, pp. 429–438, 2017-06-01, ISSN: 0002-9505, 1943-2909. DOI: 10.1119/1.4979539 (cit. on p. 45).

[86] J. W. Coltman, 'The Specification of Imaging Properties by Response to a Sine Wave Input', *J. Opt. Soc. Am., JOSA*, vol. 44, no. 6, pp. 468–471, 1954-06-01. DOI: 10.1364/JOSA.44.000468 (cit. on p. 45).

[87] 'ISO 12233:2024', ISO, Accessed: 2026-06-09. [Online]. Available: https://www.iso.org/fr/standard/88626.html (cit. on p. 47).

[88] F. Viallefont-Robinet, D. Helder, R. Fraisse et al., 'Comparison of MTF measurements using edge method: Towards reference data set', *Opt. Express, OE*, vol. 26, no. 26, pp. 33 625–33 648, 2018-12-24, ISSN: 1094-4087. DOI: 10.1364/OE.26.033625 (cit. on p. 47).

[89] L.-H. Yeh, L. Tian and L. Waller, 'Structured illumination microscopy with unknown patterns and a statistical prior', *Biomed. Opt. Express*, vol. 8, no. 2, p. 695, 2017-02-01, ISSN: 2156-7085, 2156-7085. DOI: 10.1364/BOE.8.000695 (cit. on p. 48).

[90] L.-H. Yeh, S. Chowdhury, N. A. Repina et al., 'Speckle-structured illumination for 3D phase and fluorescence computational microscopy', *Biomed. Opt. Express*, vol. 10, no. 7, p. 3635, 2019-07-01, ISSN: 2156-7085, 2156-7085. DOI: 10.1364/BOE.10.003635 (cit. on p. 48).

[91] K. Guo, Z. Zhang, S. Jiang et al., '13-fold resolution gain through turbid layer via translated unknown speckle illumination', *Biomedical Optics Express*, vol. 9, no. 1, p. 260, 2018-01-01, ISSN: 2156-7085. DOI: 10.1364/BOE.9.000260 (cit. on p. 52).

[92] FLIR, *Datasheet FLIR SC7000 Series*, 2014 (cit. on p. 54).

[93] Huaijin Chen, M. S. Asif, A. C. Sankaranarayanan et al., 'FPA-CS: Focal plane array-based compressive imaging in short-wave infrared', in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, USA: IEEE, 2015-06, pp. 2358–2366, ISBN: 978-1-4673-6964-0. DOI: 10.1109/CVPR.2015.7298849 (cit. on p. 54).

[94] W.-K. Yu, 'Super Sub-Nyquist Single-Pixel Imaging by Means of Cake-Cutting Hadamard Basis Sort', *Sensors*, vol. 19, no. 19, p. 4122, 2019-09-23, ISSN: 1424-8220. DOI: 10.3390/s19194122 (cit. on p. 57).

[95] K Szajewski, A Szajewska, S Urbaś et al., 'Imaging in NIR spectral range with a single detector using Hadamard matrix pattern- based SLM', 2022 (cit. on p. 57).

[96] P. G. Vaz, A. S. Gaudencio, L. F. Ferreira et al., 'Re-Ordering of Hadamard Matrix Using Fourier Transform and Gray-Level Co-Occurrence Matrix for Compressive Single-Pixel Imaging in Low Resolution Images', *IEEE Access*, vol. 10, pp. 46 975–46 985, 2022, ISSN: 21693536. DOI: 10.1109/ACCESS.2022.3171334 (cit. on p. 57).

[97] K. Shibuya, K. Nakae, Y. Mizutani et al., 'Comparison of reconstructed images between ghost imaging and Hadamard transform imaging', *Opt Rev*, vol. 22, no. 6, pp. 897–902, 2015-12, ISSN: 1340-6000, 1349-9432. DOI: 10.1007/s10043-015-0138-x (cit. on p. 61).

[98] 'DLP4500NIR data sheet, product information and support | TI.com', Accessed: 2026-06-10. [Online]. Available: https://www.ti.com/product/DLP4500NIR (cit. on pp. 61, 63).

[99] K. Jalalzai. 'Some remarks on the staircasing phenomenon in total-variation based image denoising'. arXiv: 1402.0091 [math.NA], pre-published (cit. on p. 65).

[100] J.-C. Krapez, 'Thermal Ellipsometry: A Tool Applied for in-Depth Resolved Characterization of Fibre Orientation in Composites', in *Review of Progress in Quantitative Nondestructive Evaluation: Volume 15A*, D. O. Thompson and D. E. Chimenti, Eds., Boston, MA: Springer US, 1996, pp. 533–540, ISBN: 978-1-4613-0383-1. DOI: 10.1007/978-1-4613-0383-1_68 (cit. on p. 69).

[101] L. Legrandjacques, J.-C Krapez, F. Lepoutre et al., 'Détection d'une fissuration débouchante à l'aide d'une caméra photothermique', 1997-01-01 (cit. on p. 69).

[102] J. Krapez, G. Gardette and D. Balageas, 'Thermal ellipsometry in steady-state or by lock-in thermography. Application for anisotropic materials characterization', in *Proceedings of the 1996 International Conference on Quantitative InfraRed Thermography*, QIRT Council, 1996. DOI: 10.21611/qirt.1996.042 (cit. on p. 69).

[103] L. Gavérina, 'Caractérisation thermique de milieux hétérogènes par excitation laser mobile et thermographie infrarouge', Theses, Université de Bordeaux, 2017-02. Accessed: 2026-05-06. [Online]. Available: https://theses.hal.science/tel-01519643 (cit. on pp. 69, 70, 81).

[104] L Gavérina, J. Batsale, A Sommier et al., 'Pulsed flying spot with the logarithmic parabolas method for the estimation of in-plane thermal diffusivity fields on heterogeneous and anisotropic materials', *Journal of Applied Physics*, vol. 121, no. 11, p. 115 105, 2017. DOI: 10.1063/1.4978919 (cit. on pp. 69, 71, 72).

[105] T. Li, D. P. Almond and D. S. Rees, 'Crack imaging by scanning pulsed laser spot thermography', *NDT & E International*, vol. 44, no. 2, pp. 216–225, 2011-03, ISSN: 09638695. DOI: 10.1016/j.ndteint.2010.08.006 (cit. on p. 69).

[106] Y. Sher, 'Review of Algorithms for Compressive Sensing of Images', 2019-08-05. arXiv: 1908.01642 (cit. on p. 75).

[107] M. Ravasi and I. Vasconcelos, 'PyLops—A linear-operator Python library for scalable algebra and optimization', *SoftwareX*, vol. 11, p. 100 361, 2020-01-01, ISSN: 2352-7110. DOI: 10.1016/j.softx.2019.100361 (cit. on p. 75).

[108] R. Taylor. 'Compressed Sensing in Python', Robert Taylor, Accessed: 2026-05-05. [Online]. Available: https://robert-taylor.me/blog/compressed-sensing-python/ (cit. on pp. 75, 88).

[109] Pylops contrib. 'PyLops', Accessed: 2026-05-06. [Online]. Available: https://pylops.readthedocs.io/en/stable/ (cit. on p. 75).

[110] R. Taylor. 'PyLBFGS', Accessed: 2026-05-06. [Online]. Available: https://github.com/humatic/pylbfgs (cit. on p. 75).

[111] L. Yang. '3DCS - TVAL3 beta2.4', GitHub, Accessed: 2026-05-06. [Online]. Available: https://github.com/liuyang12/3DCS/tree/master/packages/TVAL3_beta2.4 (cit. on p. 75).

[112] A. Mateos-Canseco, A. Kusiak and J.-L. Battaglia, 'Thermal imaging by scanning photothermal radiometry', *Review of Scientific Instruments*, vol. 94, no. 10, p. 104 902, 2023-10-01, ISSN: 0034-6748, 1089-7623. DOI: 10.1063/5.0165057 (cit. on pp. 83, 84).

[113] P.-E. Nordal and S. O. Kanstad, 'Photothermal Radiometry', *Physica Scripta*, 1979. DOI: 10.1088/0031-8949/20/5-6/020 (cit. on p. 83).

[114] F. P. Incropera, D. P. DeWitt, T. L. Bergman et al., Eds., *Principles of Heat and Mass Transfer*, 7. ed., international student version. Hoboken, NJ: Wiley, 2013, 1048 pp., ISBN: 978-0-470-50197-9 978-0-470-64615-1 (cit. on p. 83).

[115] A. Mateos-Canseco, A. Kusiak, J.-L. Battaglia et al., 'Thermal characterization of vertical interface by scanning photothermal radiometry', *Review of Scientific Instruments*, vol. 95, no. 10, p. 104 901, 2024-10-01, ISSN: 0034-6748, 1089-7623. DOI: 10.1063/5.0225690 (cit. on p. 84).

[116] A. Mateos-Canseco, A. Kusiak, J.-L. Battaglia et al., 'Investigation of the porosity of parts manufactured by electron beam melting using scanning photothermal radiometry', *Rev. Sci. Instrum.*, vol. 96, no. 6, p. 064 902, 2025-06-09, ISSN: 0034-6748. DOI: 10.1063/5.0272083 (cit. on p. 84).

[117] G. Hamaoui, E. Villarreal, H. Ban et al., 'Spatially localized measurement of isotropic and anisotropic thermophysical properties by photothermal radiometry', *Journal of Applied Physics*, 2020. DOI: 10.1063/5.0020411 (cit. on p. 84).

[118] J. Hernández-Wong, J. B. Rojas-Trigos, U. Nogal et al., 'Axisymmetric modelling of transient thermal response in solids for application to infrared photothermal radiometry technique', *Revista Mexicana de Física*, vol. 65, pp. 54–64, 1 Jan-Feb 2019, ISSN: 2683-2224. DOI: 10.31349/RevMexFis.65.54 (cit. on p. 84).

[119] D. J. Kirsch, J. Martin, R. Warzoha et al., 'An instrumentation guide to measuring thermal conductivity using frequency domain thermoreflectance (FDTR)', *Rev. Sci. Instrum.*, vol. 95, no. 10, p. 103 006, 2024-10-14, ISSN: 0034-6748. DOI: 10.1063/5.0213738 (cit. on p. 84).

[120] 'RBFInterpolator — SciPy v1.17.0 Manual', Accessed: 2026-05-06. [Online]. Available: https://docs.scipy.org/doc/scipy/reference/generated/scipy.interpolate.RBFInterpolator.html (cit. on p. 89).

[121] Z. Wang, A. Bovik, H. Sheikh et al., 'Image quality assessment: From error visibility to structural similarity', *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004-04, ISSN: 1941-0042. DOI: 10.1109/TIP.2003.819861 (cit. on p. 91).

[122] U. Seidel, K. Haupt, H. G. Walther et al., 'An attempt towards quantitative photothermal microscopy', *J. Appl. Phys.*, vol. 78, no. 3, pp. 2050–2056, 1995-08-01, ISSN: 0021-8979. DOI: 10.1063/1.360182 (cit. on p. 94).