Thèse n° 10 044

Practical computational imaging by use of spatiotemporal light modulation: from simulations to applications in biological microscopy

Présentée le 5 octobre 2023

Faculté des sciences et techniques de l'ingénieur
Laboratoire de l'IDIAP
Programme doctoral en génie électrique

pour l'obtention du grade de Docteur ès Sciences

par

François Thierry M MARELLI

Acceptée sur proposition du jury

Prof. M. Unser, président du jury
Prof. F. Fleuret, Dr M. S. D. Liebling, directeurs de thèse
Dr W. Supatto, rapporteur
Dr N. Ducros, rapporteur
Prof. D. Psaltis, rapporteur

École
polytechnique
fédérale
de Lausanne

2023

Research is what I’m doing when
I don’t know what I’m doing.

— Wernher von Braun

To my parents and sister.

Acknowledgements

I would like to thank my supervisor Dr. Michael Liebling for his support and guidance throughout my thesis. His encouragements helped me stay motivated even when facing the many difficulties introduced by a global pandemic. I particularly appreciated his advice in both written and spoken scientific communication, which pushed me to constantly polish my works and enabled me to improve my overall methodology. I also want to thank Prof. Hervé Bourlard for welcoming me when I started reaserch at Idiap Research Institute, and for developing this research institute into the wonderful work environment it is now. For this I also want to thank the whole administration and system teams of Idiap. I thank Prof. Thierry Dutoit for inspiring me and putting me on the path of research.

I thank the Swiss National Science Foundation for funding my research through the grants 200020_179217 and 206021_164022. I also thank the members of my thesis committee, Prof. Michael Unser, Prof. François Fleuret, Prof. Demetri Psaltis, Dr. Nicolas Ducros, and Dr. Willy Supatto for taking the time to read my thesis and for providing their insightful comments.

Moving to Martigny had me live far (for a Belgian) away from home for the first time, and I am very grateful to all the members of the Martigny Municipal Band for welcoming me into their ranks and facilitating my integration in the city. Practicing my passion for music in the company of similarly passionate people in Martigny was a wonderful experience that I will never forget. I also thank the musicians for introducing me to the culture of Valais and letting me live it through the many events we attended together.

During my stay at Idiap, I met and became friends with many incredible people that became my family away from home. I am extremely grateful for the countless and yet precious moments we shared together (be it cook and game evenings, dinners, hikes, kicker matches, skiing sessions or simply coffee-apple-tea breaks). Although there would be too many names to cite here for concision's sake, I must list a few lest I be tormented endlessly: Apoorv, Sargam, Bastian, Sophie, Banri, Hermanito, Julian, Fabio, Laurent, Neha, Pablo, Zohreh, Amir, Andrei, Anshul, Arnaud, Eklavya, Louise, Tilak, Christian, Olivia, and many more amazing people who made these years in Switzerland unforgettable.

I want to thank my long-time friends from Belgium, who kept in touch regardless of the distance separating us. Thanks to the unyielding support of the Team Chaussettes, I never felt alone even at times when I was confined to my single room, and I could release stress in the most crucial moments. I also thank Véro and Alexandre, who were working on their own theses at the same time as me. It has been very helpful for me to be able to discuss and share my interrogations with friends going through similar experiences. I wish to thank Florian, who brought a bit of home with him when he joined Switzerland too, and was instrumental to maintaining my physical (and mental) condition by dragging me out of my lair for memorable hikes on the mountain peaks. I wonder how he was able to put up with my company for so long without losing his sanity, but I am extremely grateful for it.

Last but not least, I wish to thank my whole family, and especially my parents and sister. In many different ways, they have encouraged me to nourish my curiosity, and helped to shape the person I am now. It is thanks to their relentless support and encouragements throughout my entire life that I can write these lines today, and I am wholeheartedly grateful to them.

Martigny, June 5, 2023 François Marelli

Abstract

Optical microscopy is an essential tool for biologists, who are often faced with the need to overcome the spatial and temporal resolution limitations of their devices to capture finer details. As upgrading imaging hardware is expensive, computational methods offer cost-efficient ways to enhance the capabilities of the devices by combining acquisition techniques and post-processing. Despite the freedom offered by mechanized stages and variable intensity lights in modern microscopes, existing custom control solutions do not offer sufficient flexibility to explore various subframe illumination and motion patterns. Moreover, the high computational cost of 3D simulations makes it difficult to characterize imaging setups at scale.

The contributions in this thesis address the above issue by introducing illumination control hardware that allows exploiting the optomechanical components commonly found in most imaging platforms to increase their spatial and temporal resolutions through computational imaging approaches. A complete simulation environment accelerated by parallel computing complements the hardware setup. I illustrate the strengths of this framework by imaging biological samples in three applications: focal sheet scanning optical projection tomography, multichannel low-speed cardiac imaging, and optomechanical modulation tomography.

As my first contribution, I extend the OpenSPIM microscope, an open light sheet device, to enable sub-frame control of illumination modulation in conjunction with the camera and stage. This allows implementing complex acquisition procedures that require a precisely timed control of the microscope's optomechanical components. Then, I show that illumination shaping using focal plane scanning increases the resolution of optical projection tomography images. This result is supported by my simulation framework, which allows testing a range of optical settings at scale thanks to parallel programming. Next, I take advantage of the subframe illumination capabilities of my system and introduce a method to reconstruct multichannel videos with virtually increased frame rate from single-channel cardiac fluorescence imaging using a low frame rate camera. The proposed method uses a paired acquisition approach that alternates between illumination modalities to allow reconstructing videos with improved temporal resolution in post-processing. Finally, I investigate a compressed sensing method that combines the spatial and temporal aspects of illumination shaping to reconstruct high-resolution volumes from few images. I formulate an efficient 1+2D regularization function that offers high reconstruction fidelity while being fast to compute using parallel computing.

In conclusion, this thesis provides hardware and software tools to implement spatiotemporal light modulation methods with good cost efficiency. It shows that the combination of long camera integration times with active illumination modulation or sample movement can be exploited to enhance spatial and temporal resolutions in various biological imaging applications, through the optical computation of continuous integrals.

Keywords: optical microscopy, computational imaging, active illumination.

Résumé

La microscopie optique est un outil essentiel pour les biologistes, qui sont souvent confrontés à la nécessité de surmonter les limites de résolution spatiale et temporelle de leur équipement pour capturer des détails de plus en plus fins. Étant donné le coût élevé d'une mise à jour du matériel d'imagerie, les méthodes computationnelles offrent une alternative économique pour dépasser les capacités des appareils en combinant techniques d'acquisition et de post-traitement. Malgré la liberté offerte par les éléments mécanisés et l'illumination contrôlable en intensité des microscopes modernes, les solutions de contrôle existantes n'offrent pas suffisamment de flexibilité que pour explorer diverses combinaisons de mouvements et d'illumination. De plus, il est difficile de tester les méthodes d'imagerie à l'échelle en raison du temps de calcul élevé des simulations 3D.

Les contributions dans cette thèse abordent ce problème en introduisant un outil de contrôle de l'illumination qui exploite les composants optomécaniques présents dans la plupart des microscopes afin d'améliorer leurs résolutions spatiale et temporelle via des approches d'imagerie computationnelle. Un environnement de simulation accéléré par parallélisme complémente ce dispositif matériel. J'illustre les atouts de cet environnement en démontrant son utilité pour trois applications d'imagerie biologique : la tomographie par projection optique, l'imagerie cardiaque basse vitesse et la tomographie par modulation optomécanique.

Ma première contribution est de modifier l'OpenSPIM, un microscope à feuillet de lumière, pour permettre la modulation rapide de l'illumination en synchronisation avec la caméra et un bras mécanisé. Cela permet d'implémenter des techniques d'acquisition complexes qui demandent un contrôle temporel précis des composants optomécaniques du microscope. Ensuite, je montre que scanner le plan focal pendant l'acquisition améliore la résolution de la tomographie par projection optique. Ce résultat est confirmé par mon environnement de simulation, qui permet de tester une large gamme de techniques à échelle réelle grâce au parallélisme. Après, j'exploite les capacités de mon dispositif à moduler l'illumination au cours d'une acquisition pour introduire une méthode permettant de reconstruire des vidéos haute-vitesse à plusieurs canaux à partir d'imagerie par fluorescence utilisant une caméra basse-vitesse à un seul canal. Ma méthode capture les images par paires, alternant entre plusieurs modalités d'illumination pour permettre la reconstruction de vidéos à haute résolution temporelle en post-traitement. Enfin, j'étudie une méthode d'acquisition comprimée qui combine les aspects temporels et spatiaux de la modulation d'illumination pour générer des volumes en haute définition à partir de peu d'images. Je formule une régularisation 1+2D qui offre une haute fidélité de reconstruction tout en étant rapide grâce au parrallélisme.

En conclusion, cette thèse fournit des outils matériels et logiciels permettant d'implémenter des techniques de modulation spatiotemporelle de la lumière avec un rapport coût-efficacité élevé. Elle montre que la combinaison de temps d'exposition longs avec une modulation spatiotemporelle de l'illumination permet d'améliorer les résolutions spatiale et temporelle dans diverses applications d'imagerie microscopique, grâce au calcul d'intégrales optiques.

Mots-clés : microscopie optique, imagerie computationnelle, illumination active.

Contents

Introduction

Optical microscopy is an essential tool for modern biological research. Enabling the detailed study of small organisms, it is widely used to understand the development of animal embryos. One of the many applications of optical microscopy is medical research, through the imaging of animal models that share similar organ formation with humans. For example, the first stages of cardiac development of humans are shared with the zebrafish (Bakkers, 2011; Nguyen, Lu, Y. Wang, and J.-N. Chen, 2008), which is a widely studied model in cardiovascular research for that reason. Microscopic imaging of the zebrafish embryo led to a better understanding of mechanisms underlying the formation of the heart (Rohr, Bit-Avragim, and Abdelilah-Seyfried, 2006; Bussmann, Bakkers, and Schulte-Merker, 2007; D. Y. Stainier and Fishman, 1994; Supatto and Vermot, 2011), as well as the early apparition of cardiac defects and diseases and possible ways to treat them (C.-J. Huang, Tu, Hsiao, Hsieh, and Tsai, 2003; Chico, Ingham, and Crossman, 2008; Scherz, Huisken, Sahai-Hernandez, and D. Y. R. Stainier, 2008; Vermot, A. S. Forouhar, Liebling, D. Wu, Plummer, Gharib, and S. E. Fraser, 2009; Marques, Lupi, and Mercader, 2019). These advances in developmental biology drive a constant need to surpass the capacities of imaging devices, as microscopists try to capture ever finer details at ever faster speeds.

Enhancing microscopy platforms through upgrading their hardware has become increasingly expensive due to the high sophistication of the optical elements used in modern imaging devices. To alleviate this cost, computational imaging methods offer an alternative way to extend the capacities of microscopes, based on a smart combination of acquisition strategies and post-processing algorithms. They allow extracting more information from the measurements, therefore maximizing the performance of imaging platforms. Diverse needs and optical configurations have led to the development of various computational imaging techniques, spanning a broad range of applications and offering specific benefits. I briefly review here some of these methods that are used in modern research.

Taking advantage of interference patterns generated when working with coherent light, Digital Holographic Microscopy (DHM) allows reconstructing volumetric information from a single image, called hologram, by extracting its phase and amplitude information (Gabor, n.d.; Kim, 2010). Working instead with low coherence light, Optical Coherence Tomography (OCT) detects the interference between a reference beam and light reflected by the sample to reconstruct its structure at chosen depth (D. Huang, Swanson, Lin, Schuman, Stinson,

W. Chang, Hee, Flotte, Gregory, Puliafito, and Fujimoto, 1991). In multimode fiber imaging, contrast enhancement is obtained by recording a hologram and applying a digital confocal filtering method that simulates the propagation of light through a pinhole (Loterie, Farahi, I. Papadopoulos, Goy, Psaltis, and Moser, 2015).

Compressed sensing exploits prior knowledge about the information contained in the sample to reconstruct high quality images from few measurements (Donoho, 2006; Candes and Wakin, 2008). Its application in fluorescence microscopy allows reducing photobleaching (an unwanted reduction of fluorophore emission due to sustained illumination) and increase acquisition speed (Calisesi, Ghezzi, Ancora, D'Andrea, Valentini, Farina, and Bassi, 2022). Spectral imaging also attempts to maximize the quantity of information per measurement by imaging multiple fluorophores with different emission spectra simultaneously, separating them via a linear unmixing of their contribution in each imaged channel (Zimmermann, Rietdorf, Girod, Georget, and Pepperkok, 2002; Zimmermann, Rietdorf, and Pepperkok, 2003). Temporal pixel multiplexing on the other hand acquires multiple images within a single camera exposure by activating selected groups of pixels sequentially, trading spatial resolution for increased acquisition speed (Bub, Tecza, Helmes, Lee, and Kohl, 2010). With a more complex multi-spectral LED array illumination, it is possible to compress both time and colour information into a single image using multiple modulation patterns for different wavelengths inside a single frame (Ma, Yuan, Fu, and Arce, 2021).

In order to provide isotropic resolution in 3D microscopy, Optical Projection Tomography (OPT) acquires images of the sample from different orientations (obtained by rotating the object or the imaging device), combining them to recover the volumetric information (Sharpe, Ahlgren, Perry, Hill, Ross, Hecksher-Sørensen, Baldock, and D. Davidson, 2002). Using two perpendicular imaging paths instead of rotating the sample, dual-view plane illumination microscopy fuses images acquired with Single Plane Illumination Microscopy (SPIM) to achieve isotropic resolution while operating with low photodamage (Y. Wu, Wawrzusin, Senseney, Fischer, Christensen, Santella, York, Winter, Waterman, Z. Bao, et al., 2013; Huisken, Swoger, Bene, Wittbrodt, and E. H. K. Stelzer, 2004). Multicolor two-photon light-sheet microscopy further reduces the photobleaching of 3D acquisitions while providing high imaging depth, and can be used to obtain 4D movies by using a temporal registration procedure (Mahou, Vermot, Beaurepaire, and Supatto, 2014).

To overcome the resolution limit imposed by optical diffraction, PhotoActivated Localization Microscopy (PALM) and Stochastic Optical Reconstruction Microscopy (STORM) randomly activate and localize small populations of fluorophores. They then aggregate multiple measurements to reconstruct images with a very high resolution (Betzig, Patterson, Sougrat, Lindwasser, Olenych, Bonifacino, M. W. Davidson, Lippincott-Schwartz, and H. F. Hess, 2006; Rust, Bates, and Zhuang, 2006; S. T. Hess, Girirajan, and Mason, 2006). Relying instead on an LED array, Fourier ptychography microscopy illuminates the sample from various angles, allowing to combine multiple low resolution measurements to reconstruct a high resolution image, or to recover depth information (Zheng, Horstmeyer, and C. Yang, 2013; Tian, X. Li, Ramchandran, and Waller, 2014; Tian, J. Wang, and Waller, 2014).

Among computational imaging methods, some exploit the potential of illumination modulation. They draw their performance from shaping the light in both time and space, using precise patterns to design imaging processes that yield improved spatial and temporal resolutions.

Structured Illumination Microscopy (SIM) pioneered this concept, using multiple acquisitions with carefully designed illumination patterns to overcome the diffraction limit and achieve superresolution (Gustafsson, Agard, and Sedat, 1999; M. G. L. Gustafsson, 2000; Heintzmann, Jovin, and Cremer, 2002). More recently, implementations with finer spatial control over the illumination intensity helped reduce its phototoxicity (Chakrova, Canton, Danelon, Stallinga, and Rieger, 2016).

Light Sheet Fluorescence Microscopy (LSFM), in particular, has seen many developments using illumination shaping to improve its performance. For hyperspectral imaging using a spectrometer, shaping the light sheet with a spatial illumination pattern allows reconstructing images with high spectral and spatial resolutions (Crombez, Leclerc, Ray, and Ducros, 2022). Another example is the generation of a virtually uniformly thin light sheet to achieve very precise optical sectioning over a large field of view. This is achieved by shaping a Bessel light sheet through sweeping the focus point of the illumination optics and synchronously scanning a strip of active pixels in the camera sensor (Dean, Roudot, Welf, Danuser, and R. Fiolka, 2015). This can be further improved by combining it with structured illumination to improve its optical sectioning performance in light scattering samples (Frantz, Karamahmutoglu, Schaser, Kirik, and Berrocal, 2022). The full potential of using SIM with LSFM is unlocked by using an oblique plane microscopy setup, which offers isotropic resolution doubling thanks to the possibility to freely orient the light sheet in 3D (B. Chen, B.-J. Chang, Roudot, F. Zhou, Sapoznik, Marlar-Pavey, Hayes, Brown, Zeng, Lambert, Friedman, C.-L. Zhang, Burnette, Shepherd, Dean, and R. P. Fiolka, 2022).

For volumetric imaging, systems able to generate multiple simultaneous light sheets have reduced acquisition time and phototoxicity (Calisesi, Castriotta, Candeo, Pistocchi, D'Andrea, Valentini, Farina, and Bassi, 2019). A similar illumination shaping can also improve contrast in extended depth of field applications (Zunino, Garzella, Trianni, Saggau, Bianchini, Diaspro, and Duocastella, 2021). A more flexible 4D light patterning achieved with SPIM helps control light dose and reduce photodamage in 3D video acquisition (Power and Huisken, 2018).

However, the previously listed methods require the use of dedicated devices, or costly specific modifications of existing optical setups such as the addition of digital micromirror devices, piezoelectric actuators, electrically tunable lenses or acousto-optic deflectors in the illumination path. Although there has been an effort to propose a versatile platform that provides more generic spatial light modulation for SPIM devices (Garbellotto and Taylor, 2018), its implementation still requires costly hardware modifications that are not compatible with all imaging devices.

Recent works have implemented spatiotemporal light shaping by exploiting the optomechanical components commonly found in modern imaging platforms: mechanical stages and intensity controllable light sources. These methods have the advantage of requiring little to no modifications for use with existing microscopy platforms, therefore being usable for a broad range of applications and with a good cost efficiency.

For example, modulating illumination in time to generate laser pulses shorter than the exposure time of the camera to reduces photobleaching and motion blur when imaging fast moving samples (Staudt, J. Liu, Thorn, Stuurman, Liebling, and D. Y. R. Stainier, 2014). To completely negate motion blur, it is possible to use temporally coded illumination by turning the illumination on and off through time following a pseudo-random sequence. Motion blur in the obtained images can then be efficiently removed by deconvolution (Gorthi, Schaak, and Schonbrun, 2013).

Using multiple light sources of different wavelengths synchronized in time allows increasing the acquisition speed of colour cameras when imaging monochromatic objects (Jaques, Ernst, Mercader, and Liebling, 2020). A similar setup can reduce aliasing in videos by implementing generalized sampling using temporal light modulation (Jaques and Liebling, 2020).

In previous works, we showed that illumination modulation through time allows acquiring pairs of images that contain complementary information. For example, it is possible to encode motion by capturing a blurry image, and pair it with a sharp image acquired in quick succession. This allows sorting image sequences and disambiguating similar frames using the motion information (Mariani, Marelli, Jaques, Ernst, and Liebling, 2021).

Exploiting the mechanized stage allows virtually increasing the depth of field of OPT by scanning the focal plane through the sample during the acquisition of a single projection (Miao, Hayenga, Meyer, Neumann, Nelson, and Seibel, 2010). This can be coupled with resolution enhancement using PSF deblurring on the acquired scanned projections (Kikuchi, Sonobe, Sidharta, and Ohyama, 1994; K. G. Chan and Liebling, 2017). By combining mechanical scanning of a SPIM device with temporal illumination modulation, Woringer et al. (Woringer, Darzacq, Zimmer, and Mir, 2017) proposed a versatile technique to perform 3D compressed sensing.

Although these methods use common optomechanical components that require no major hardware modifications to implement, they use custom control solutions for spatiotemporal illumination which are not flexible enough to generalize to other techniques or explore new methods. Designing such controllers is time-consuming and requires technical skills in low-level programming and electronics, which makes it difficult to switch between different modulation strategies and slows down the research process towards new illumination shaping techniques. Moreover, the high dimensionality of modern microscope images makes simulating and testing such systems in 3D at scale highly computationally demanding, further slowing down the development and characterization of these light shaping methods.

In this thesis, I address this issue by designing a versatile optomechanical modulation framework that allows implementing a wide range of spatiotemporal light modulation techniques using common optomechanical components. Namely, I target the synchronization of a mechanized stage with a digital camera, and the temporal intensity modulation of multiple light sources. I provide a corresponding simulation framework that allows quick prototyping and characterization of these methods, and can efficiently run at scale thanks to parallel computing. I demonstrate its practical usability in three different applications in biological imaging: optical projection tomography, cardiac imaging and compressed sensing for 3D fluorescence microscopy. In each of these examples, I measure the benefits offered by the spatiotemporal light modulation in terms of spatial and temporal resolutions on experimental images.

Thesis outline and contributions

Throughout this thesis, I describe my hardware and software platform for spatiotemporal illumination shaping and its applications to biological imaging. Chapters 2 to 4 are based on scientific publications that illustrate its versatility by implementing illumination shaping methods of increasing complexity. First I consider the spatial component alone, then I study temporal aspect of modulation, and finally I describe the combination of both in a spatiotemporal illumination shaping technique.

In Chapter 1, I briefly describe the hardware and software design of my platform. I give an overview of its characteristics and advantages, and how it implements spatiotemporal illumination shaping within a popular open source microscopy control software.

In Chapter 2, which is based on (Marelli and Liebling, 2021), I describe an application of spatial light shaping using the mechanized stage to improve the resolution of OPT imaging. High resolution OPT is achieved by deconvolving Focal-Plane-Scanning (FPS-OPT) data, but this requires an accurate knowledge of the system's Point Spread Function (PSF). While the presence of noise and inaccuracies in the PSF model or parameters affects reconstruction quality, their effect is difficult to assess quantitatively in practice and the computational cost of naive simulations is prohibitively expensive. I present an efficient approach to carry out FPS-OPT simulations for a wide range of illumination geometries, including Focal-Sheet-Scanning OPT (FSS-OPT), a method using a lateral light-sheet illumination to perform FPS-OPT. I implement a simulation framework that can accommodate large size 3D data by dividing the forward model into elements that can be efficiently processed by GPUs. I compare the performance of FPS-OPT and FSS-OPT on simulated data. In the presence of Poisson noise, I show that FSS-OPT outperforms FPS-OPT with deconvolution even if all model parameters are accurately known. I then validate these results on experimentally acquired data.

In Chapter 3, based on (Marelli, Ernst, Mercader, and Liebling, 2023), I investigate the application of temporal illumination modulation to live cardiac imaging using fluorescence microscopy. In this application, a high acquisition frame rate is necessary to study the fast contractions of the beating heart, but the limited fluorescence intensity requires sensitive cameras that are often too slow. Moreover, the problem is even more complex when imaging distinct tissues in the same sample using different fluorophores. I present Paired Alternating AcQuisitions (PAAQ), a method to image cyclic processes in multiple channels, which requires only a single (possibly slow) camera.

I generate variable temporal illumination patterns in each frame, alternating between channel-specific illuminations (fluorescence) in odd frames and a motion-encoding brightfield pattern as a common reference in even frames. Starting from the image pairs, I find the position of each reference frame in the cardiac cycle through a combination of image-based sorting and regularized curve fitting. Thanks to these estimated reference positions, I assemble multichannel videos whose frame rate is virtually increased.

I characterize my method on synthetic and experimental images collected in zebrafish embryos, showing quantitative and visual improvements in the reconstructed videos over existing non-gated sorting-based alternatives. Using a 15 Hz camera, I showcase a reconstructed video containing 2 fluorescence channels at 100 fps.

In Chapter 4, which corresponds to (Marelli and Liebling, 2023), I implement an illumination shaping technique that exploits both spatial and temporal modulations: OptoMechanical Modulation Tomography (OMMT). This method exploits compressed sensing to reconstruct high resolution microscopy volumes from fewer measurement images compared to exhaustive section sampling in conventional light sheet microscopy. Nevertheless, the volumetric reconstruction process is computationally expensive, making it impractically slow to use on large-size images, and prone to generating visual artefacts. I propose a reconstruction approach that uses a 1+2D Total Variation (TV_{1+2}) regularization that does not generate such artefacts and is amenable to efficient implementation using parallel computing. I evaluate my method for accuracy and scaleability on simulated and experimental data. Using a high quality, but computationally expensive, Plug-and-Play (PnP) method that uses the BM4D denoiser as a benchmark, I observe that my approach offers an advantageous trade-off between speed and accuracy.

Finally, I conclude this thesis in Chapter 5, summarizing the impact of my research and offering possible directions for future works.

1 A versatile framework for spatiotemporal illumination shaping

Implementing spatiotemporal light modulation requires a precise synchronized control of the multiple optomechanical elements used for imaging. Commercially available universal function generators are very expensive and lack flexibility, leading to the use of custom-made modulation controllers in research. The open-source Arduino microcontroller is a common choice to develop such control solutions thanks to its relative ease of use, good availability and high cost efficiency (Jaques, Ernst, Mercader, and Liebling, 2020; Jaques and Liebling, 2020; Mariani, Marelli, Jaques, Ernst, and Liebling, 2021). However, the implementations provided with the previously cited works are specific to their application and hardware, and cannot be generalized to multiple imaging setups. Although they share the same basic principle of generating voltage curves to modulate the intensity of light sources, their design does not allow to freely choose the modulation function. Moreover, they do not offer any user-friendly control interface, and each update of the modulation parameters (such as exposure time or maximum intensity) requires modifying their source code, recompiling it and uploading it to the microcontroller. This mode of operation make these devices very impractical to use, and require the user to have some degree of technical skills. In addition to this, these custom controllers do not implement communication with a computer, which makes synchronization with other devices (i.e., a mechanical stage) impossible.

The solution proposed in (Woringer, Darzacq, Zimmer, and Mir, 2017), also based on Arduino, is more user-friendly as it provides a graphical user interface for controlling the modulation. It also allows the generation of any arbitrary modulation. However, its implementation is still specific to the optical setup used in their work, and cannot be easily adapted to a different imaging hardware. Moreover, the authors do not provide any instructions or schematics for reproducing their electronic setup, making it almost impossible to reimplement or adapt their controller without reverse engineering their published drivers.

The absence of a versatile and user-friendly control solution slows down the process of developing new imaging methods based on spatiotemporal illumination modulation, as any new implementation requires repeating the steps of designing a custom controller. This process is time-consuming and requires technical skills in electronics and low level programming. In order to facilitate the prototyping of new spatiotemporal light shaping methods, I propose a design for a flexible modulation device aimed at being adaptable to various imaging platforms, controlled via a user-friendly graphical interface.

1.1 Flexible hardware design

My design is based on the Arduino Nano microcontroller, which offers the same performance and advantages as the standard Arduino Uno with a smaller sized board. I take advantage of the Pulse Width Modulation (PWM) outputs of the controller to implement 4 intensity modulation outputs that can output voltages in the range of 0 to 5 V. I use the maximum available PWM frequency of 62.5 kHz to avoid flickering artefacts when imaging at high speed. For cases where analogue intensity modulation is required, 2 of these outputs are filtered using RC low-pass filters with a cut-off frequency of 10.6 kHz. My design also includes 4 digital modulation outputs that enable on/off switching at lower speed using a 5 V logic level. The Arduino offers a 4 µs resolution for the modulation outputs.

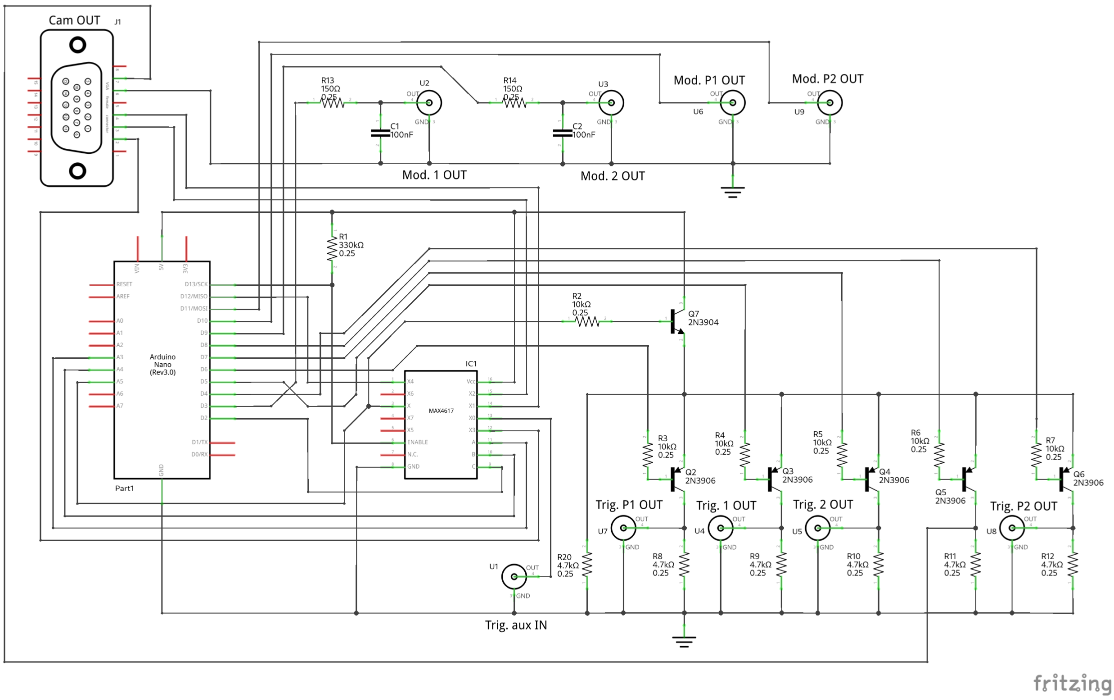

Figure 1.1: The electrical schematic of the designed hardware controller. It features input trigger selection by multiplexing, and offers 4 analogue and 4 digital modulation outputs.

In order to enable a maximum flexibility, I use a multiplexer to switch between different possible trigger sources for the modulation synchronization. This allows using the camera or any external device as a reference, relying on the internal clock of the Arduino, or manually triggering events from the software side. The schematics of my device are given in Fig. 1.1. Thanks to its 4 analogue and 4 digital modulation outputs, this controller is flexible enough to be implemented in any of the applications listed at the start of this chapter, and in any imaging setup using up to 4 light sources. It is however easy to extend this to any higher number of illumination devices by having multiple implementations of the modulation device connected to the same trigger source.

In order to facilitate the implementation of my controller, I propose a blueprint to build it on a standard stripboard with commonly available low-cost electronic components in Fig. 1.2. Using standard BNC connectors for the modulation outputs allows easily connecting it to most commercially available illumination devices, though it is easy to replace them with different connectors to adapt this design to specific imaging hardware. I also provide a 3D-printed case for this design to protect the electronics. Figure 1.3 shows a picture of the completed device.

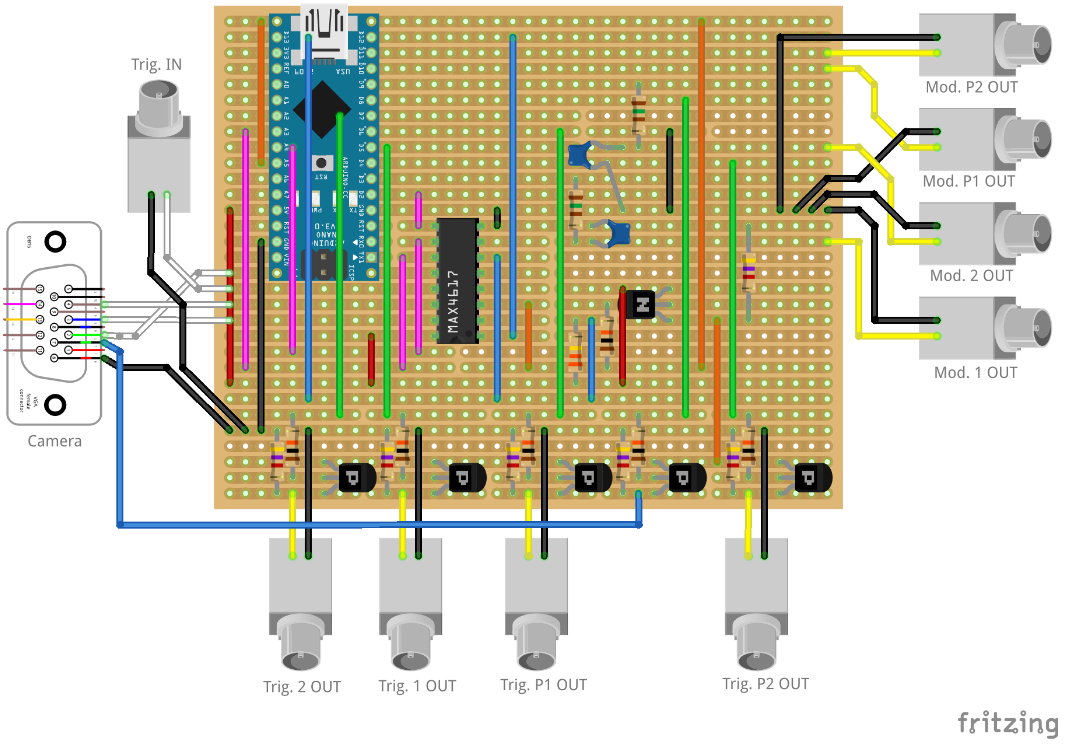

Figure 1.2: Proposed blueprint for the implementation of the modulation controller on a standard stripboard. It uses commonly available low-cost components.

1.2 User-friendly operation

To enable easy operation of my modulation controller, I provide software drivers to integrate it into the popular open-source microscopy framework μManager (Edelstein, Amodaj, Hoover, Vale, and Stuurman, 2010), which allows controlling a wide range of devices such as cameras and mechanized stages. The device is controlled using the plugin shown in Fig. 1.4. It provides a user-friendly interface with the controller to easily set the modulation parameters and

Figure 1.3: The controller in its 3D-printed protective case. Standard BNC connectors allow maximal compatibility with most imaging hardware.

visualize the output functions. To offer maximum flexibility, the controllers supports the use of modulation patterns that span multiple acquisition frames of the camera.

I also distribute a high-level scripting library that allows automating complex acquisition procedures. It is a wrapper over the low-level API of μManager to control mechanized stages and cameras via simplified function calls. This facilitates the implementation of imaging techniques that exploit non-trivial stage motion during acquisition, such as moving the focal plane during the integration time of the camera. Moreover, it allows using the modulation controller in synchronization with a mechanized stage to achieve spatiotemporal light shaping. The different high-level automation functions packaged in the scripting library include:

- standard 4D stage control (x, y, z and rotation),
- standard single image acquisition,
- high-speed sequence acquisition,
- software-triggered manual exposure acquisition,
- fast depth stack acquisition by recording a video while moving the focus stage,
- projection acquisition by sweeping the focus stage during camera integration time, used in Chapter 2 and in (Y. Liu, Dong, Pham, Marelli, and Unser, 2022) to implement FSS-OPT,
- dynamically changing the patterns of the modulation controller, used in Chapter 4 to implement OMMT.

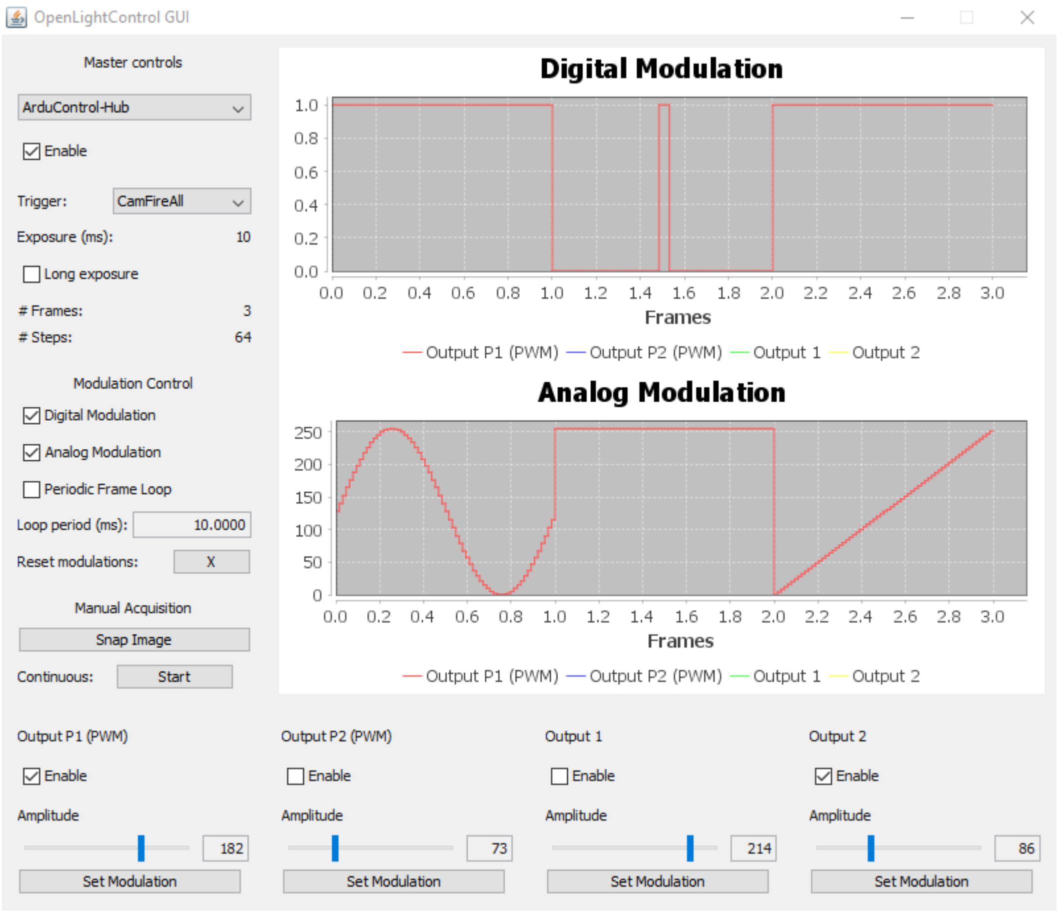

OpenLightControl GUI

Master controls

ArduControl-Hub

☑ Enable

Trigger: CamFireAll

Exposure (ms): 10

☐ Long exposure

Frames: 3

Steps: 64

Modulation Control

☑ Digital Modulation

☑ Analog Modulation

☐ Periodic Frame Loop

Loop period (ms): 10.0000

Reset modulations: X

Manual Acquisition

Continuous: Start

Output P1 (PWM): ☑ Enable, Amplitude 182

Output P2 (PWM): ☐ Enable, Amplitude 73

Output 1: ☐ Enable, Amplitude 214

Output 2: ☑ Enable, Amplitude 86

Figure 1.4: The graphical interface for controlling the device in μManager. All the modulation parameters are easily accessible for user-friendly operation. The top panel shows the generated output curves.

The schematics, drivers, and plugins described in this chapter are publicly available at the following address: https://github.com/idiap/CBI-MMTools.

1.3 Experimental imaging platform

In this thesis, I use a custom implementation of the OpenSPIM microscopy platform (Pitrone, Schindelin, Stuyvenberg, Preibisch, Weber, Eliceiri, Huisken, and Tomancak, 2013) to acquire experimental datasets. It uses two lasers (Vortran Stradus, 488 and 561 nm), a UMPLFLN 20XW semi-apochromat water dipping objective lens and a sCMOS camera (Andor Zyla 4.2) mounted on a U-TV0.5XC-3 adapter. It also includes brightfield illumination with a power LED and a Köhler illumination path. Thanks to its 3 different light sources and its 4D mechanized stage, this imaging setup is well suited for implementing spatiotemporal illumination modulation techniques based on optomechanical components. Figure 1.5 shows an overview of the setup and illustrates the different illumination paths.

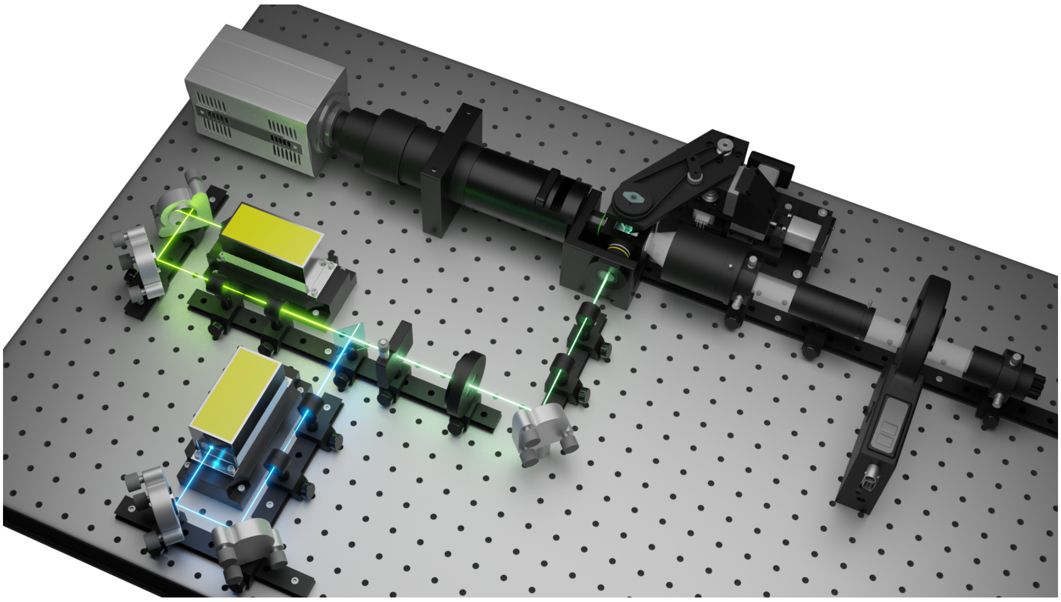

Figure 1.5: The custom OpenSPIM platform used in this thesis (3D rendering). It features a SPIM illumination with 2 lasers and a brightfield illumination using a power LED.

2 Focal plane scanning for enhanced resolution in optical projection tomography

This chapter contains the results published in F. Marelli and M. Liebling (2021). "Optics Versus Computation: Influence of Illumination and Reconstruction Model Accuracy in Focal-Plane-Scanning Optical Projection Tomography". In: *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pp. 567–570. DOI: 10.1109/ISBI48211.2021.9433834.

2.1 Introduction

Optical Projection Tomography (OPT) is a microscopy technique used for 3D imaging of mesoscopic samples, from a few micrometers to a few millimetres (Sharpe, Ahlgren, Perry, Hill, Ross, Hecksher-Sørensen, Baldock, and D. Davidson, 2002). It uses projected images taken at different angles to reconstruct volumetric information, similarly to X-ray computed tomography. To achieve approximately straight-line projections, OPT requires the sample to be smaller than the imaging depth of field. Low Numerical Aperture (NA) objectives have a large depth of field and are typically used for OPT but they have worse lateral resolution thereby limiting the resolution of OPT systems.

In order to circumvent this limitation, the depth of field of high-NA objectives can be extended by scanning the focal plane through the sample, creating high-resolution pseudoprojections (Miao, Hayenga, Meyer, Neumann, Nelson, and Seibel, 2010). A 3D image is then reconstructed from these projections using the Filtered BackProjection (FBP) (Kak, Slaney, and G. Wang, 2002). The pseudoprojections contain blurred out-of-focus information, which can be filtered out to further improve the quality of the reconstruction (Kikuchi, Sonobe, Sidharta, and Ohyama, 1994; K. G. Chan and Liebling, 2017). This filtering operation requires to know the Point Spread Function (PSF) of the system, but measuring it is very difficult in practice (McNally, Karpova, Cooper, and J. A. Conchello, 1999).

In this paper, we propose to combine Light-Sheet Fluorescence Microscopy (LSFM) (Huisken, Swoger, Bene, Wittbrodt, and E. H. K. Stelzer, 2004) with focal plane scanning in order to reduce out-of-focus blur for high-NA objectives, into Focal-Sheet-Scanning OPT (FSS-OPT).

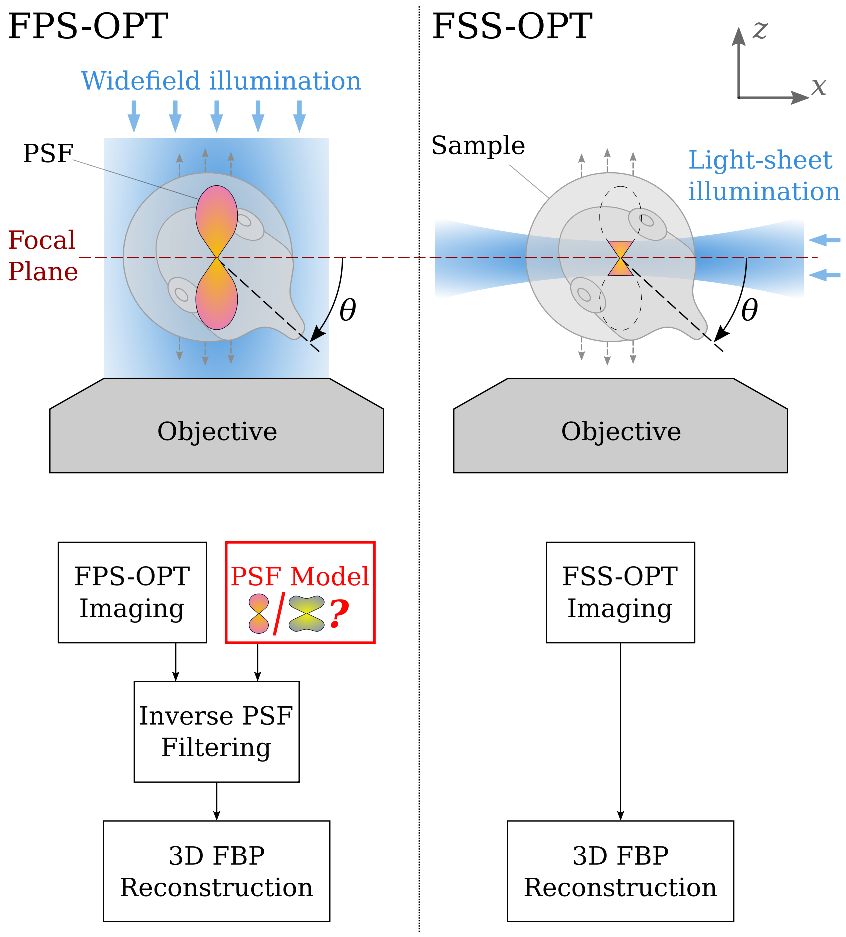

Figure 2.1: Comparison of FPS-OPT and FSS-OPT focal plane scanning OPT methods. In both methods, the sample is rotated then scanned along the z axis to get pseudoprojections. FPS-OPT allows taking into account high-NA PSFs in a traditional wide-field illumination acquisition setup; precise reconstruction requires a modified filtered back-projection and knowledge of the PSF. FSS-OPT uses an optically more involved lateral light-sheet illumination setup that truncates the PSF sufficiently for direct reconstruction with X-ray tomography-like methods.

Specifically, we present a simulation framework to compare its performance to the Focal-Plane-Scanning OPT (FPS-OPT) technique described in (K. G. Chan and Liebling, 2017), which implements a modified FBP incorporating deconvolution to take into account the PSF. Our aim is to measure the impact that potential errors on the precise PSF have during FPS-OPT's reconstruction and whether it remains competitive with the (all-optical) FSS-OPT technique.

In Section 2.2, we present the imaging processes and the reconstruction algorithms. In Section 2.3, we characterize the methods using a 3D phantom with different NAs, and describe the implementation of an efficient framework used for the simulations. In Section 2.4, we illustrate the results on OPT data from fluorescent textile fibre.

2.2 Imaging methods

Focal-plane scanning OPT methods, such as the method proposed by Miao et al. (Miao, Hayenga, Meyer, Neumann, Nelson, and Seibel, 2010), use data obtained via pseudoprojections, that is, by integrating the images obtained when scanning the focal plane through the whole sample. In practice, this is done by taking a single long exposure image while sweeping the focal plane through the object, or by acquiring a depth stack of the sample and averaging the images. This scanning procedure is common to both FPS-OPT and FSS-OPT, which we will describe shortly.

For the following, let $f(x,y,z)$ be a 3D object to be imaged, with x,y,z the horizontal, vertical and focus axes, respectively; and y is used as the rotation axis for the acquisitions.

2.2.1 FPS-OPT

We briefly recall FPS-OPT (Miao, Hayenga, Meyer, Neumann, Nelson, and Seibel, 2010; K. G. Chan and Liebling, 2017), which uses wide-field imaging to acquire the projections, as illustrated in Fig. 2.1. The entire sample is illuminated, and $p(s,y,\theta)$, the pseudoprojection acquired at an angle θ, results from a convolution between the object and the hourglass-shaped PSF of the system $h(x,y,z)$ according to:

$$p(s,y,\theta) = \mathscr{X}f \circledast \left\{ \int_{\mathbb{R}} h(x,y,z)\,\mathrm{d}z \right\} = \mathscr{X}f \circledast h_\perp(x,y), \tag{2.1}$$

where $\circledast$ is a convolution on the first two axes, and $h_\perp(x,y)$ is the projection of the PSF along z; $\mathscr{X}f(s,y,\theta)$ is the 2D projection of f along the direction θ using to the X-ray transform, with s the horizontal axis in projected space, and $\delta(\cdot)$ the Dirac delta function:

$$\mathscr{X}f(s,y,\theta) = \iint_{\mathbb{R}^2} f(x,y,z)\,\delta(z\cos(\theta) + x\sin(\theta) - s)\,\mathrm{d}x\,\mathrm{d}z. \tag{2.2}$$

Following (K. G. Chan and Liebling, 2017), the blurred projections are filtered in Fourier space using the regularized inverse of the projected PSF:

$$H_{inv}(\omega_x,\omega_y) = \frac{\mathscr{F}^{2*}\left\{h_\perp(x,y)\right\}}{\left|\mathscr{F}^2\left\{h_\perp(x,y)\right\}\right|^2 + \lambda\left|\mathscr{F}^2\left\{r(x,y)\right\}\right|^2}, \tag{2.3}$$

where $\mathscr{F}^2$ is the 2D Fourier transform, λ is a regularization weight and $r(x,y)$ is a 2D Laplacian high-pass regularization filter. The FBP (Kak, Slaney, and G. Wang, 2002) is then applied to reconstruct the volumetric information.

The filtering step requires to know the PSF of the imaging system. However, experimentally measuring it is a tedious process sensitive to noise, and theoretical models can only give an ideal approximation of the actual optics (McNally, Karpova, Cooper, and J. A. Conchello, 1999).

Therefore, the PSF used for deconvolution in practice can contain discrepancies that could degrade the quality of the final reconstruction.

2.2.2 FSS-OPT

To overcome the need for using a deconvolution, we propose to use LSFM illumination to truncate the PSF of the system when acquiring the pseudoprojections. The sample is illuminated from the side using a thin light sheet, as illustrated in Fig. 2.1. This eliminates most of the out-of-focus contribution to the images (Huisken, Swoger, Bene, Wittbrodt, and E. H. K. Stelzer, 2004), at the expense of more costly and complex optics required to generate the light sheet, as compared to the wide-field illumination for FPS-OPT.

Integrated platforms built for both OPT and LSFM imaging already exist, and most LSFM devices can be extended to integrate OPT modality without modifying the light sheet system (Bassi, Schmid, and Huisken, 2015). Such platforms could be operated such as to acquire FSS-OPT images without optical hardware modifications.

In the case of FSS-OPT, the pseudoprojection $q(s,y,\theta)$ obtained at an angle θ is given by:

$$q(s,y,\theta) = \int_{\mathbb{R}} \left[\left(R_\theta \left\{ f \right\} \cdot T_z\{g\} \right) * h \right] (s,y,z)\,\mathrm{d}z, \tag{2.4}$$

where $g(x,y,z)$ is the LSFM illumination function, $T_z\{\cdot\}$ is a transformation operator translating a function by a distance z along the focal axis, and $*$ is a 3D convolution operator.

The effective imaging PSF is locally equal to the multiplication of the original PSF h with the illumination function g, and brings much less out-of-focus information as the light sheet thickness is smaller than the PSF (Engelbrecht and E. H. Stelzer, 2006). The pseudoprojections are therefore less blurred, and the FBP (or other X-ray tomography methods) can be directly applied to the acquired images without any deconvolution. As compared to FPS-OPT, this method does not require any prior knowledge on the system's PSF. FSS-OPT shares similarities with multi-view LSFM techniques (Swoger, Verveer, Greger, Huisken, and E. H. Stelzer, 2007), but requires less heavy computations for reconstruction as no registration or deconvolution are required. Also, the data storage requirements are the same as for OPT.

2.3 Proposed Simulation Method

To evaluate both methods, we simulate the imaging processes with multiple NAs. In order to get a meaningful sample size of at least 0.1 mm with a realistic sampling resolution, the 3D array representing the object must be at least 256 × 256 × 256. A naive implementation of the imaging equations is very slow for 3D arrays of this size, and does not allow to run the simulations in a reasonable time. In the following section, we describe how we optimized the implementation for a major speed up, before presenting the simulation results.

2.3.1 Efficient framework implementation

We use spline-based algorithms from (Horbelt, Liebling, and Unser, 2002) to implement the X-ray transform and FBP. Only a 2D implementation is provided, and looping over the slices of a 3D array is very inefficient. Radon transform and FBP are highly parallel algorithms that greatly benefit from multithreaded implementations and Graphics Processing Unit (GPU) acceleration (Fasih and Hartley, 2010). Our implementation uses a GPU-accelerated version of the X-ray transform and FBP whose efficiency stems from sharing computations between layers along the rotation axis. This saves both memory and computation time, and gives a 200-fold speed improvement over the initial implementation.

The FPS-OPT simulation is based on Eq. (2.4) which, although compact in appearance, would be very inefficient if implemented as such because the convolution with h would be computed individually for every z and every θ. By explicitly expanding $T_z\{g\}$ and the convolution with h we can reorder the integrals:

$$\begin{aligned} q(s,y,\theta) &= \int_{\mathbb{R}} \iiint_{\mathbb{R}^3} R_\theta f(u,v,w) \cdot g(u,v,w-z) \\ &\qquad\qquad \cdot h(s-u, y-v, z-w)\,\mathrm{d}u\mathrm{d}v\mathrm{d}w\ \mathrm{d}z && (2.5) \\ &= \iiint_{\mathbb{R}^3} R_\theta f(u,v,w) \int_{\mathbb{R}} g(u,v,w-z) \\ &\qquad\qquad \cdot h(s-u, y-v, z-w)\,\mathrm{d}z\ \mathrm{d}u\mathrm{d}v\mathrm{d}w && (2.6) \\ &= \iint_{\mathbb{R}^2} \int_{\mathbb{R}} R_\theta f(u,v,w)\,\mathrm{d}w \int_{\mathbb{R}} g(u,v,\alpha) \\ &\qquad\qquad \cdot h(s-u, y-v, -\alpha)\,\mathrm{d}\alpha\ \mathrm{d}u\mathrm{d}v && (2.7) \end{aligned}$$

The variable substitution $\alpha = w - z$ between Eqs. (2.6) and (2.7) is valid under the hypothesis that the object has a finite extent in space. The integral over w can then be computed on a finite ensemble $D \subset \mathbb{R}$. We obtain a new formulation for the FSS-OPT pseudoprojection:

$$q(s,y,\theta) = \iint_{\mathbb{R}^2} \mathscr{X} f(u,v,\theta) \cdot h_\perp^{uv}(s,y)\,\mathrm{d}u\mathrm{d}v, \tag{2.8}$$

were we use a local projected PSF $h_\perp^{uv}(x,y)$ defined as:

$$h_\perp^{uv}(x,y) = \int_{\mathbb{R}} g(x,y,z) \cdot h(u-x, v-y, -z)\,\mathrm{d}z. \tag{2.9}$$

The expression derived above for $q(s,y,\theta)$ yields a very efficient implementation, as the optimized X-ray transform can be used to compute $\mathscr{X} f$. Moreover, the number of total operations is divided by N_z, the depth of the object array. As $h_\perp^{uv}$ does not depend on θ, the computation of the local projected PSF can be shared between all angles. This formulation is also suitable for multithreading and GPU acceleration, making its implementation very efficient.

2.3.2 Simulation results

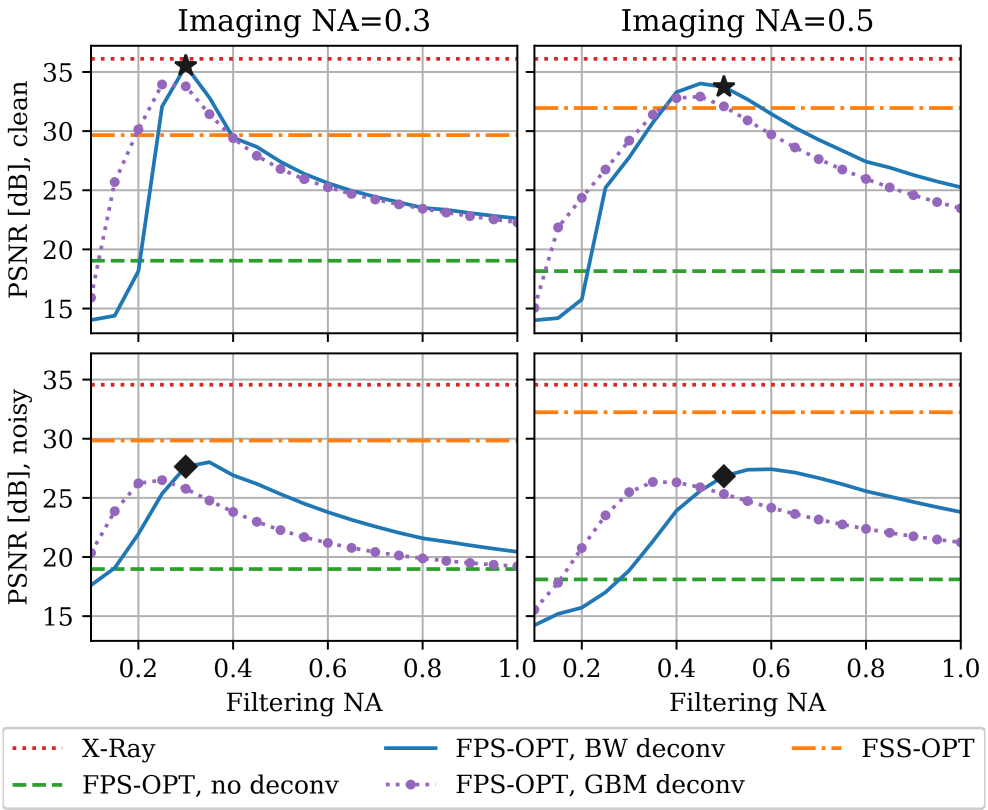

Figure 2.2: Comparison of the reconstruction PSNR with both methods. Without noise, FPS-OPT beats FSS-OPT if the error on the PSF model is small (★). With Poisson noise, FSS-OPT is better even with perfect PSF knowledge (⧫). Ideal X-ray and non-filtered FPS-OPT are shown as anchors.

To compare the performance of FPS-OPT and FSS-OPT, we simulated the imaging processes on a 3D Shepp-Logan phantom. We used the Born & Wolf (BW) model (Born and Wolf, 2013; Kirshner, Aguet, Sage, and Unser, 2013) to generate the PSFs. The LSFM illumination was simulated to match the OpenSPIM platform (Pitrone, Schindelin, Stuyvenberg, Preibisch, Weber, Eliceiri, Huisken, and Tomancak, 2013) using Fresnel propagation in POPPY (Perrin, Soummer, Elliott, Lallo, and Sivaramakrishnan, 2012), and the beam thickness and shape were validated by comparison with (Remacha, Friedrich, Vermot, and Fahrbach, 2020). The projections were simulated for 360 angles evenly spaced between 0° and 180°. We first considered a noise-free scenario, and then added shot noise modelled by a Poisson distribution on the noise-free data rescaled to the range $[3, 10^4]$.

To quantify the impact of incorrect PSF on FPS-OPT reconstruction, we used multiple PSFs with varying NAs between 0.1 and 1 for deconvolution. We also used an approximated Gaussian Beam Model (GBM) as defined in (Trull, Horst, Palenstijn, Vliet, Leeuwen, and Kalkman, 2017) to account for errors in the model itself. We used the Peak Signal-to-Noise Ratio (PSNR)

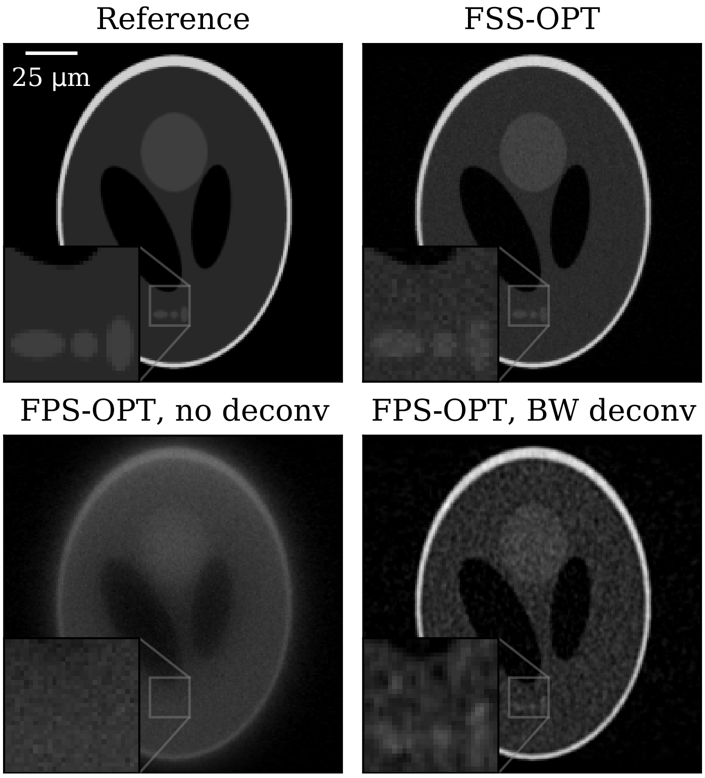

Figure 2.3: In the presence of Poisson noise, FSS-OPT improves the resolution over FPS-OPT even if the deconvolution PSF is exact. The central xy section of a 3D simulation is shown, imaging NA=0.5.

as criterion for the quality $\mathrm{PSNR} = 10\log_{10}\left(\max(f)^2/\mathrm{MSE}\right)$, where MSE is the Mean Squared Error between the reconstruction and the original object f. We used multiple values of λ in the inverse filter, and selected the reconstructions with the best PSNR.

Figure 2.2 compares the reconstruction PSNR for imaging NAs of 0.3 and 0.5, in the ideal noise-free situation (row 1) and in the presence of Poisson noise (row 2). Ideal X-ray (imaging without PSF blurring) represents the best possible reconstruction, while FPS-OPT without deconvolution is the basic scanning OPT. On noise-free images, FPS-OPT outperforms FSS-OPT if the NA of the deconvolution PSF is within 0.1 of the imaging NA, the performance being only slightly reduced when using the GBM approximation. When Poisson noise is added, FSS-OPT outperforms FPS-OPT even if the deconvolution PSF exactly matches the imaging one, yielding a much better resolution in the reconstruction as shown in Fig. 2.3.

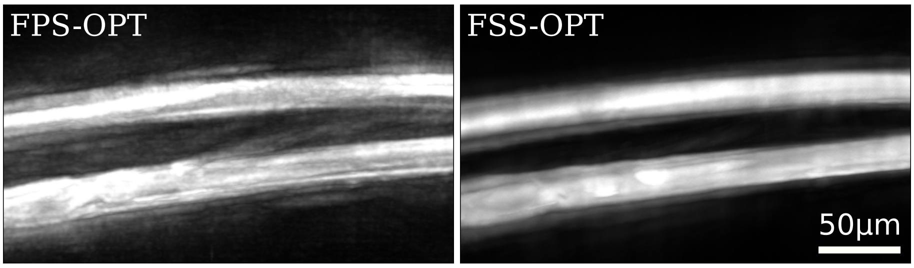

Figure 2.4: Comparison of FPS-OPT and FSS-OPT on fluorescent textile fibres. Mean intensity projection of the 3D image is shown.

2.4 Experiments

To validate the simulation results, we imaged two fluorescent textile fibres with a diameter of 25 µm using the OpenSPIM implementation described in Section 1.3. The fibres were mounted in a 1.5% low melting agarose solution, inside a fluorinated ethylene propylene tube. We acquired 180 pseudoprojections over 180° by scanning the focal plane over a depth of 300 µm. FPS-OPT was deconvolved with the BW PSF model with a 0.5 NA matching the objective. As visible in Fig. 2.4, FSS-OPT contains less out-of-focus blur than FPS-OPT, which validates the simulation results.

2.5 Conclusions

We have introduced a simulation framework for FPS-OPT that is flexible and efficient enough to accommodate various 3D illumination geometries. In particular, we proposed FSS-OPT, a scanning OPT imaging technique that reduces the out-of-focus blur using an LSFM illumination, which our simulation framework allowed us to characterize on real-scale 3D data. Our simulations have shown that FSS-OPT outperforms FPS-OPT when Poisson noise is present, improving the resolution without requiring any prior knowledge on the imaging PSF. Furthermore, we have validated the simulated results using experimentally-acquired data. Our simulation framework allowed to identify and characterize the benefits of an alternative all-optical illumination geometry (FSS-OPT).

The source code of the simulation framework is available at:
`https://github.com/idiap/cbi_toolbox`.

3 Paired illumination for multichannel virtual high frame rate

This chapter contains the works published in F. Marelli, A. Ernst, N. Mercader, and M. Liebling (2023). "PAAQ: Paired Alternating AcQuisitions for Virtual High Frame Rate Multichannel Cardiac Fluorescence Microscopy". Submitted for publication.

3.1 Introduction

In vivo fluorescence microscopy is an essential tool to study the early stages of organ development in animal embryos thanks to its intrinsically selective contrast (Lichtman and J.-A. Conchello, 2005; Sanderson, Smith, Parker, and Bootman, 2014; C. Vonesch, Aguet, J.-L. Vonesch, and Unser, 2006). When used to image the zebrafish, an excellent animal model for cardiovascular research (Bakkers, 2011; Nguyen, Lu, Y. Wang, and J.-N. Chen, 2008), it allows studying the development stages of the heart (D. Y. Stainier and Fishman, 1994; Rohr, Otten, and Abdelilah-Seyfried, 2008; Bussmann, Bakkers, and Schulte-Merker, 2007) and understanding the progression of early cardiac defects and diseases (Chico, Ingham, and Crossman, 2008; Scherz, Huisken, Sahai-Hernandez, and D. Y. R. Stainier, 2008; C.-J. Huang, Tu, Hsiao, Hsieh, and Tsai, 2003). Correctly imaging the dynamics of the heartbeat requires high-speed acquisition due to the rapid beating of the heart. This becomes a limitation when the low fluorescence intensity emitted by the labelled structures requires a longer exposure time or extremely sensitive cameras that are either too slow or prohibitively expensive. In addition to this, the use of several fluorophores to label multiple tissues usually requires that their emission be recorded sequentially in different channels.

Some acquisition platforms are able to perform simultaneous high-speed acquisition in multiple fluorescent channels. This is usually achieved by splitting the emitted light into multiple beams that are redirected to different cameras, through the use of mirrors (Mickoleit, Schmid, Weber, Fahrbach, Hombach, Reischauer, and Huisken, 2014; Gregor, Butkevich, Enderlein, and Mojiri, 2021) and prisms (Cai, Y. L. Wang, Wainner, Iftimia, Gabel, and Chung, 2019). However, these come at the cost of an increased complexity in the optics due to the addition of the splitting components. These optics and the use of multiple high-speed cameras make these systems very expensive and could make them difficult to integrate into existing imaging platforms.

As an alternative, when such parallel high-speed imaging is not available, images acquired from different channels at low speed could still be registered to their position in the heartbeat period (phase). Approaches to solve this problem fall into two categories: prospective gating and retrospective gating (reviewed in (Taylor, 2014)).

In prospective gating, the acquisition of images is triggered at precise timings that correspond to desired sampling phases in the heartbeat cycle. On big samples, this trigger can be extracted from cardiac probes that precisely measure the start and duration of a heartbeat (Jenkins, Chughtai, Basavanhally, Watanabe, and Rollins, 2007; Sablong, Rengle, Ramgolam, Saint-Jalmes, and Beuf, 2014; Brau, Wheeler, Hedlund, and G. A. Johnson, 2002). Such signals would be difficult to access in zebrafish embryos due to their small size. Instead, the trigger can be obtained by processing (in real-time) a video signal captured by a dedicated second camera (Taylor, Saunter, Love, Girkin, Henderson, and Chaudhry, 2011; Taylor, Girkin, and Love, 2012). These prospective gating methods give very good results even with slow acquisition devices, but they also require a dedicated optical setup and real-time triggering hardware and processing capabilities.

Retrospective gating methods acquire images at arbitrary phases and then attempt to estimate these phases via post-acquisition algorithms. Their results are usually not as precise as those of prospective gating, but as they do not require additional triggering hardware they can be easier to implement with existing imaging setups. For example, it is possible to sort images acquired at random times in a period based only on image distance to reconstruct a singe-channel video with a virtually increased frame rate (Q. Zhang and Pless, 2005; Mariani, K. Chan, Ernst, Mercader, and Liebling, 2019; Mariani, Ernst, Mercader, and Liebling, 2020). In addition to their limited accuracy, these methods can struggle with multichannel data. In some cases, it is possible to perform time registration between multiple channels by using a mutual information criterion (Liebling and Ranganathan, 2009; Liebling, Vermot, A. Forouhar, Gharib, M. Dickinson, and S. Fraser, 2006) or by registering them with a common reference channel (Ohn, J. Yang, S. E. Fraser, Lansford, and Liebling, 2011). Nevertheless, these channel alignment methods require that high-speed sequences be available in each channel to work.

When imaging with slower cameras, one possibility would be to first generate virtual high frame rate videos of each channel separately using the sorting methods above (Mariani, K. Chan, Ernst, Mercader, and Liebling, 2019; Mariani, Ernst, Mercader, and Liebling, 2020). These *virtual* high-speed sequences could then be synchronized using a mutual information-based algorithm designed for raw high-speed movies (Liebling, A. S. Forouhar, Gharib, S. E. Fraser, and M. E. Dickinson, 2005) to obtain the final a multichannel video. We see two main drawbacks to this potential method. Firstly, the virtual high frame rate sequences obtained with sorting are not uniformly sampled, and this sampling is different for each channel. This means that a perfect frame-to-frame pairing from one channel to another does not exist, which will negatively impact the performance of channel alignment. Secondly, high mutual information between the different channels is not guaranteed, and it can be difficult to obtain correct channel registration depending on the imaged region or the fluorophores used. These two drawbacks limit the expected performance of such a method, and could even make it unusable in some scenarios involving highly uncorrelated signals.

Instead of solving the virtual frame rate increase and channel registration problems sequentially, we propose a method that considers them jointly. In a previous work, we introduced a method to disambiguate images that appeared similar despite being captured in different phases of the heartbeat. We paired sharp images with blurred images that encode the motion at the time of acquisition (Mariani, Marelli, Jaques, Ernst, and Liebling, 2021). In this paper, we extend this technique into a general approach centred around Paired Alternating AcQuisitions (PAAQ), which result in image sequences whose frames alternate between a common reference modality and other channels (e.g., fluorescence channels). We then use this common reference to sort images from all the channels at once, achieving simultaneously a virtual high frame rate and multichannel registration. In order to address the lack of precision of naive frame sorting approaches, we also propose a phase estimation algorithm. In addition to improving phase accuracy, it also overcomes two central limitations of plain image sorting: the resulting movies can be inverted in time and their phases cannot be assigned quantitative time units to measure durations.

3.2 Problem definition and metrics

We define f_c as the intensity of the signal obtained by imaging a periodic phenomenon (the embryo heartbeat) using fluorescence microscopy, with each channel number $c = 0, \ldots, C-1$ referring to a distinct combination of illumination and emission wavelengths. The measured signal varies according to its spatial location (x, y) and its phase θ. The intensity f_c has a period of 2π, which is expressed as follows:

$$f_c(x, y, \theta) = f_c(x, y, \theta + k2\pi),\ \forall k \in \mathbb{Z}. \tag{3.1}$$

The phase itself is a function of time expressed as $\theta(t)$. In ideal time-periodic systems, it increases linearly over time t following $\theta = \omega t$, where ω is the angular frequency. In practice, such a relation is too simplistic for the cardiac cycle. Indeed, some variability stems from both the biological nature of the phenomenon and from environmental changes (e.g., temperature rising due to prolonged exposure to light leading to an increase in heart rate (Ohn and Liebling, 2011). Although the exact relationship between time and phase $\theta(t)$ is unknown, it still follows a trend that is roughly linear $\theta(t) \simeq \bar{\omega} t$, where $\bar{\omega}$ is the average angular frequency calculated over multiple periods.

We consider C sequences of N images $\mathbf{f}_c[:,:,n]$ (where $:,:$ is a shorthand notation to describe all rows and columns of the image (Golub and Van Loan, 2013), and $n = 0,\dots,N-1$) obtained by sampling the signal f_c at unknown phases $\theta_{c,n}$ with a frame rate F_{acq}, using a procedure that we detail below. Our objective is to retrieve the phases $\theta_{c,n}$, which contain all the information required to reconstruct synchronized image sequences across all channels with an increased temporal resolution, as the phases give the position of each frame in the cycle.

More specifically, we want to compute estimates $\tilde{\theta}_{c,n}$ as close as possible to the unknown real phases $\theta_{c,n}$. We define our evaluation criterion as the following error function:

$$\mathscr{E} = \min_{\Theta} \frac{1}{NC} \sum_{c=0}^{C-1} \sum_{n=0}^{N-1} \left| \theta_{c,n} \ominus \left(\tilde{\theta}_{c,n} + \Theta\right) \right|, \tag{3.2}$$

where Θ is a phase shift common to all estimations, and $\ominus$ is a phase difference operator:

$$\alpha \ominus \beta = \alpha - \beta + k2\pi \;\text{ with }\; k \in \mathbb{Z} \;\text{ s.t. }\; \alpha \ominus \beta \in (-\pi, \pi]. \tag{3.3}$$

The global phase shift Θ in the error signifies that we do not care about the absolute phases of our measurements, but only about their position relative to each other in the period. Indeed, the phase of the first frame in our signal is unknown, but this information is not necessary to derive meaningful information, such as the time elapsed between specific events in the period, or to synchronize the channels.

3.3 Methods: acquisition and reconstruction

In order to solve the previously defined problem, we introduce a two-stage method that combines PAAQ imaging and post-acquisition processing to virtually increase the frame rate. The purpose of PAAQ is to create a reference signal common to all fluorescence channels that allows synchronizing the sequences with high precision. We then use the image-processing step to refine the estimate of the phase of each frame, as well as to compute the main properties of the imaged signal. We illustrate the whole pipeline in Fig. 3.1. Sections 3.3.1 to 3.3.4 detail the different stages of our method.

3.3.1 Implementing PAAQ with active illumination

In previous work (Mariani, Marelli, Jaques, Ernst, and Liebling, 2021), we have introduced a technique that uses active illumination to acquire pairs of images in quick succession by alternating between illumination patterns. The first image in a pair encodes movement information for unequivocally sorting the sequence in order of increasing phase, while the second uses a short light flash that produces a sharp image used for analysis and display of the reconstructed sequence. We extend this method to achieve synchronization of multiple fluorescent channels by switching between imaging modalities at the same time as illumination patterns.

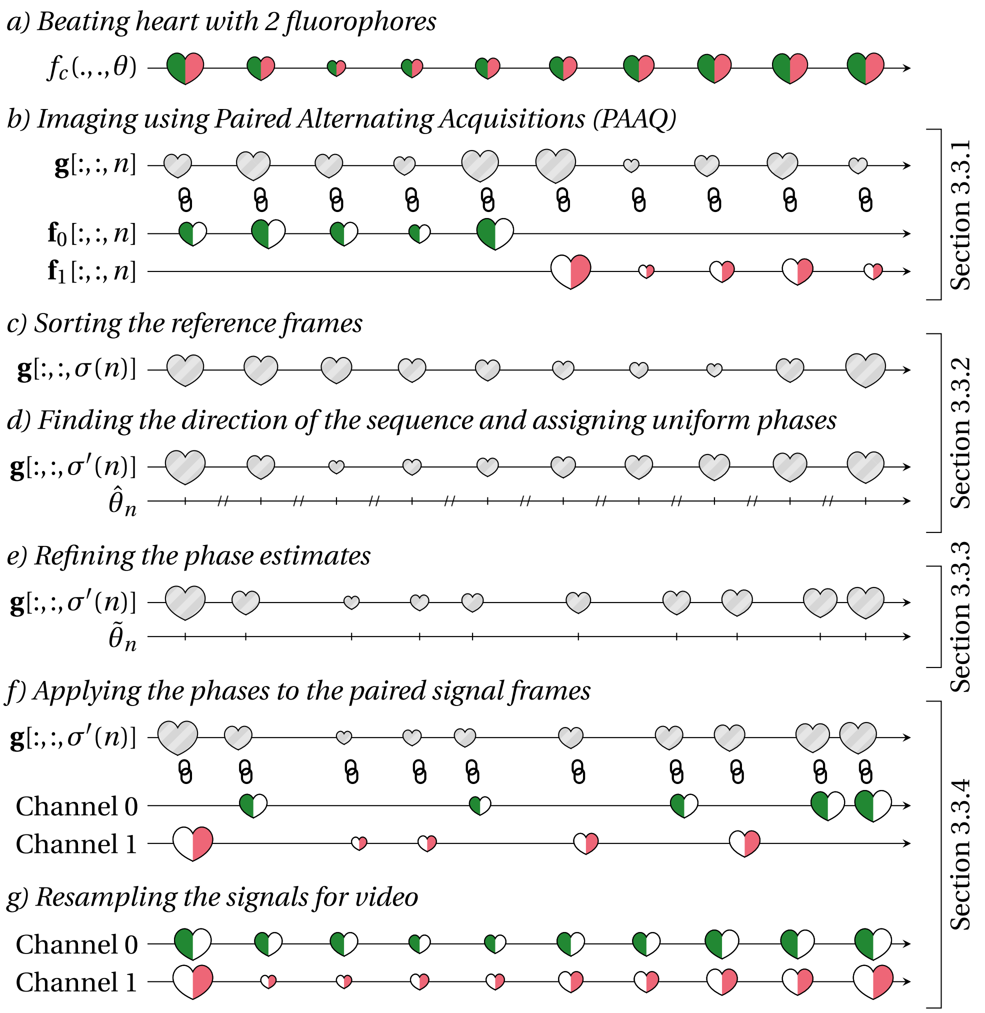

Figure 3.1: Overview of proposed PAAQ method for virtual high frame rate cardiac imaging.

We start by defining a reference signal $g(x, y, \theta)$ that is obtained by imaging the beating heart using brightfield microscopy. Any modality is suitable to acquire g, and even one of the f_c could be the reference. However, since it is a sacrificial signal that will be discarded after imaging (as the time-encoding illumination results in blur), we choose to use brightfield as it causes little to no photobleaching. The intensity of this reference varies over the same space and phase as f_c:

$$g(x, y, \theta) = g(x, y, \theta + k2\pi),\ \forall k \in \mathbb{Z}. \tag{3.4}$$

For each channel, we use PAAQ to capture a sequence of N image pairs $\big(\mathbf{g}_c[:,:,n],\mathbf{f}_c[:,:,n]\big)$, consisting of a reference image and a signal image captured in quick succession by switching from brightfield to fluorescence illumination every other frame:

$$\text{even frames:}\quad \mathbf{g}_c[l,m,n]=\int_{t_c+2n\Delta_T}^{t_c+2n\Delta_T+\Delta_E} r(t-(t_c+2n\Delta_T))\, g\big(l\Delta x, m\Delta y, \theta(t)\big)\, dt \tag{3.5}$$

$$\text{odd frames:}\quad \mathbf{f}_c[l,m,n]=\int_{t_c+(2n+1)\Delta_T}^{t_c+(2n+1)\Delta_T+\Delta_P} f_c\big(l\Delta x, m\Delta y, \theta(t)\big)\, dt, \tag{3.6}$$

where Δ_x and Δ_y are the pixel width and height, $l=0,\dots,L-1$ and $m=0,\dots,M-1$ the row and column index pairs, and $n=0,\dots,N-1$ the time frame index. Δ_T is the time interval between two consecutive frames, and t_c the arbitrary time at which the first reference image associated to channel c is acquired. Since we capture the images without interruption, Δ_T directly relates to the acquisition frame rate of the camera $F_{\text{acq}}=1/\Delta_T$. The phase corresponding to each image $\mathbf{f}_c[:,:,n]$ is $\theta_{c,n}=\theta\,(t_c+(2n+1)\Delta_T)$ and $r(t)$ is the illumination pattern used to capture images of the reference signal. It is defined over the interval $[0,\Delta_E)$ where Δ_E is the exposure time of the camera. When acquiring images of the fluorescent signals, a short light pulse of duration Δ_P illuminates the sample, with $\Delta_P\le\Delta_E$. By using a short pulse, we can ensure that the fluorescent images are sharp even if the frame rate of the camera is low. Figure 3.2 illustrates PAAQ imaging, which corresponds to Step b) in Fig. 3.1.

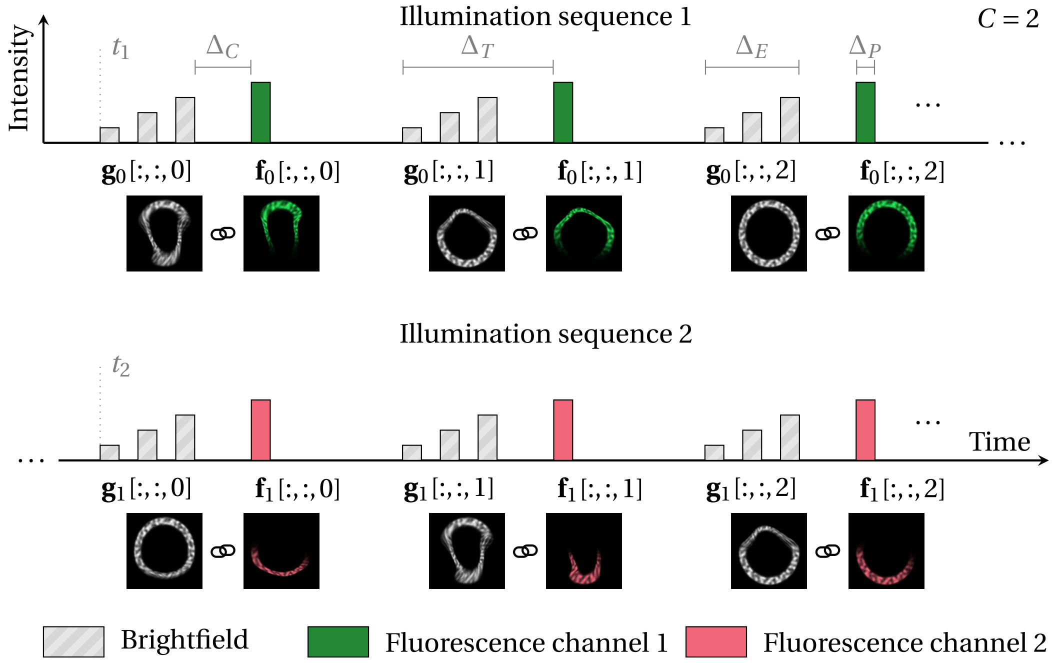

Figure 3.2: Illumination modulation for PAAQ imaging. Quickly switching between different channels and modalities allows associating fluorescence frames to a common brightfield reference.

The PAAQ method is based on the assumption that the delay between a reference image and its associated signal image is small. By acquiring the reference image at the very end of the exposure time of the camera, we make this delay entirely independent of the frame rate of the device. Indeed, this delay becomes $\Delta_C = \Delta_T - \Delta_E$ the transfer time of the camera, which is the time between two exposure periods during which the device reads the pixel information and gets ready to capture a new frame. Even on devices with low frame rate, this transfer time is typically very low (on the order of a few ms). Therefore, we shape $r(t)$ to capture most information at the end of the exposure period, minimizing the delay between images in a pair. Instead of the continuous ramp used in the original method (Mariani, Marelli, Jaques, Ernst, and Liebling, 2021) we use a series of pulses of increasing amplitude for $r(t)$, as this pattern is easier to generate while retaining the same motion-encoding properties as the ramp.

Given that the delay between a signal frame and its associated reference is very small, the phase difference between the two will also be small due to the linear trend between time and phase. Over such a low duration, the variability in the heartbeat cycle is negligible, and we can consider the phase delay between each reference and its paired signal frame to be a constant. Therefore, since a constant phase shift does not impact the global error defined in Eq. (3.2), finding the phase of each signal frame is equivalent to finding the phase of its associated reference. This is a much easier problem to solve, as all references are of the same modality, and the results obtained on reference sequences will be directly applicable to their associated signals from different channels. We are thus able to extend the PAAQ technique to generate synchronized virtual high frame rate sequences from multiple fluorescent channels at no additional computational cost.

3.3.2 Sorting and orienting the references

We combine all the reference images into a single sequence $\mathbf{g}[:,:,n]$ by concatenating the acquired images $\mathbf{g}_c$ for each channel, containing a total of $N_g = NC$ reference images. Similarly, we concatenate signal frames from all channels into a single sequence $\mathbf{f}[:,:,n]$. We can then apply a phase-sorting algorithm to the reference sequence $\mathbf{g}$ to obtain a permutation $\sigma : \{0,\dots,N_g-1\} \to \{0,\dots,N_g-1\} : \; n \to \sigma(n)$ that sorts the references by increasing phase. For this we use a travelling salesman method introduced in previous works (Mariani, K. Chan, Ernst, Mercader, and Liebling, 2019) that finds the permutation σ by minimizing the total absolute image difference between consecutive images. The obtained sorted sequence $\mathbf{g}[:,:,\sigma(n)]$ contains all the reference images reordered to reproduce a single period of the heartbeat, with a virtually increased frame rate that depends only on the total amount of images acquired N_g. This corresponds to Step c) in Fig. 3.1.

This sorting-based method assumes that the images in the reordered sequence correspond to a uniform sampling of the signal. It assigns a linear phase estimate based on the position of each frame in the sorted sequence:

$$\hat{\theta}_n = \frac{\sigma^{-1}(n)}{N_g} 2\pi. \tag{3.7}$$

The travelling salesman solution does not contain directional information. Indeed, the permutation $\sigma'(n) = N_g - \sigma(n) + 1$ corresponds to a tour with the exact same cost. However, direction is important for the reconstruction of a faithful video that follows the chronological order of events. If we assume that the sampling frequency of the reference frames (i.e., half the frame rate of the camera) is at least twice as big as the frequency of the imaged phenomenon (similarly to the Nyquist criterion), we can retrieve the direction of the sequence by computing the average phase distance between consecutively acquired frames:

$$\bar{\Delta}_{\hat{\theta}} = \frac{1}{C(N-1)} \sum_{c=0}^{C-1} \sum_{n=0}^{N-2} \left(\hat{\theta}_{c,n+1} \ominus \hat{\theta}_{c,n}\right). \tag{3.8}$$

When the sampling frequency assumption is met, the phase distance between consecutive frames must be smaller than π (half a period), and $\bar{\Delta}_{\hat{\theta}}$ must be positive. If $\bar{\Delta}_{\hat{\theta}}$ is negative, the permutation obtained using sorting is in the wrong direction, and we must use σ' instead to sort the sequence and compute the uniform phase estimate $\hat{\theta}$. This is shown as Step d) in Fig. 3.1.

3.3.3 Non-uniform phase estimation using image distance

The assumption that the sorted sequence corresponds to a uniform sampling of a single period of the signal is imprecise and leads to artefacts in the reconstructed sequence. Due in part to the ratio between the acquisition frequency and the frequency of the imaged signal, and in part to the variations in the periodicity of the phenomenon, different regions of the period will contain more or less sampling points, as shown in Fig. 3.3 b). When reconstructing a video, assuming that the sampling is uniform will generate distortions as stretching will appear in undersampled areas, and shrinking in oversampled regions, as illustrated in Fig. 3.3 c). It is therefore important to refine the phase estimate in order to reconstruct a faithful signal that correctly represents the dynamics of the imaged system.

In order to correct distortions in volumetric electron microscopy imaging caused by inhomogeneous sampling in the axial direction, Hanslovsky and al. have developed an image-base method that uses image-to-image similarity to estimate the position of slices in the volume (Hanslovsky, Bogovic, and Saalfeld, 2016). Given the similarity to our problem, we introduce a method that is inspired by this algorithm, modified as indicated below to be suitable for the specificities of our problem.

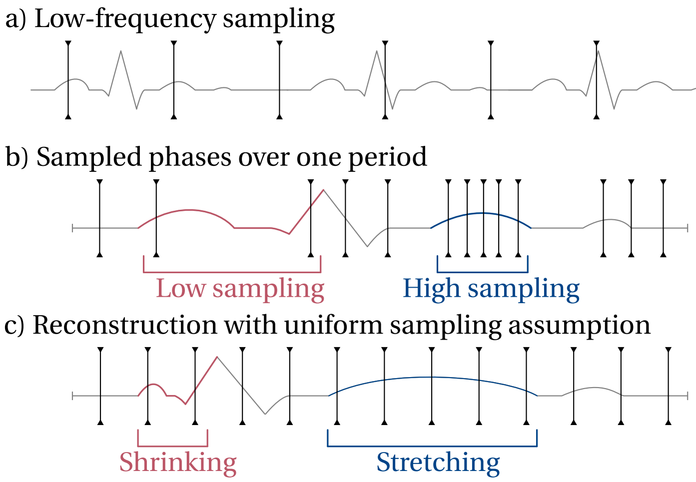

Figure 3.3: Reconstruction artefacts introduced by the uniform sampling assumption. a) Fast repeating process sampled with a low frequency. b) Rearranged over a single period, the sampled points are not uniformly spaced. c) If assuming a uniform phase sampling the reconstructed signal is deformed.

The method is based on the assumption that the pixel-based distance between images in the signal increases with their phase difference. Indeed, images that are temporally close to one another in the signal should be similar, while images further away are more likely to look different. This assumption echoes the hypothesis on which the sorting-based virtual high frame rate technique is built (Mariani, K. Chan, Ernst, Mercader, and Liebling, 2019).

For each acquired image at index n in our sequence, we define a curve that represents the evolution of image-to-image distance with respect to phase distance:

$$\mathscr{D}_n(u) = \mathscr{D}_n(u + k2\pi), \ \forall k \in \mathbb{Z}. \tag{3.9}$$

The periodicity of $\mathscr{D}_n(u)$ stems directly from the periodicity of the imaged signal. We define this function to be 0 at $u = 0$, then monotonically increasing up to a maximum, and monotonically decreasing back to 0 at $u = 2\pi$. This monotonicity is a direct translation of the assumption that images will look less similar as the phase interval separating them increases. The method introduced for volumetric imaging (Hanslovsky, Bogovic, and Saalfeld, 2016) uses a monotonically decreasing similarity function over the whole domain, but this cannot represent the periodicity of the signal. The use of a distance metric instead of a similarity measure only changes the function from decreasing to increasing, which does not affect the method. More specifically, we use the Minkowski distance of order 1 as the distance metric d between two images $\mathbf{a}$ and $\mathbf{b}$ of size $L \times M$:

$$d(\mathbf{a}, \mathbf{b}) = \sum_{l=0}^{L-1} \sum_{m=0}^{M-1} |\mathbf{a}[l,m] - \mathbf{b}[l,m]| \, . \tag{3.10}$$

The second assumption of the method is that distance curves $\mathscr{D}_n(u)$ corresponding to points arbitrarily close in phase are near identical. This is referred to as local constancy of shape by the authors of the volumetric imaging method (Hanslovsky, Bogovic, and Saalfeld, 2016). This assumption seems reasonable for our problem as the motion of the heart is continuous, and over any arbitrarily short timespan its speed is near constant.

Ideally, the distance between any two images at indexes n and i in our sequence $d(\mathbf{g}[:,:,n], \mathbf{g}[:,:,i])$ should match the value of the distance curve evaluated at the corresponding phase distance $\mathscr{D}_n(\tilde{\theta}_i - \tilde{\theta}_n)$. If we knew $\mathscr{D}_n$, we could use it to find the phase of each frame in our sequence. Therefore, our method is based on two steps: first we find an estimate for the distance curve, and then we use it to compute new phase estimates, as illustrated in Fig. 3.4. We iterate over these two steps until the phase estimation converges which corresponds to Step e) in Fig. 3.1. The following subsections detail each step of the algorithm.

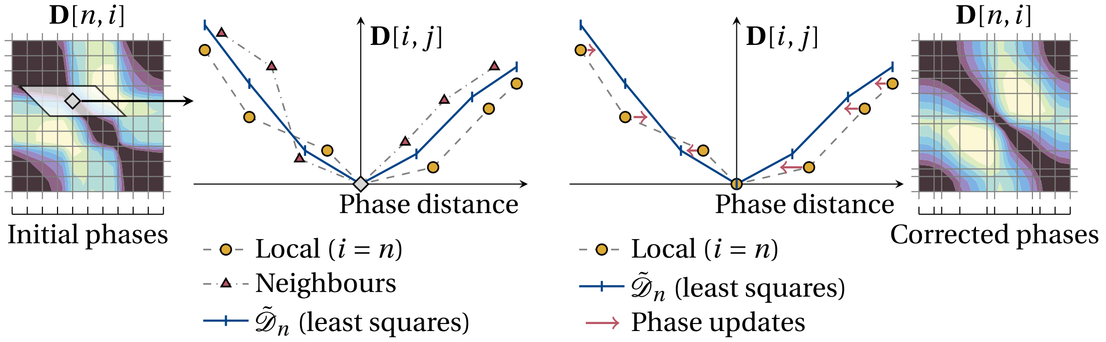
a) Estimating the distance curve locally

b) Fitting the phases to the curve

Figure 3.4: Image-based iterative phase estimation algorithm. a) We compute a local average distance curve $\tilde{\mathscr{D}}_n$ via least-squares fitting on the measurements $\mathbf{D}[:,:]$ b) We update the phase estimates by minimizing the distance between local measurements and the average curve. We iterate over these steps until convergence. The distance matrix $\mathbf{D}[:,:]$ gets smoother as the phase estimates improve.

Estimating the distance curve locally

We approximate the distance curve $\mathscr{D}_n$ by defining a piecewise-linear function $\tilde{\mathscr{D}}_n$ characterized by values $\tilde{\mathbf{d}}_n[l]$, $l = 0, \dots, L-2$ which we will compute below, defining L intervals over $u \in [0, 2\pi)$:

$$\tilde{\mathscr{D}}_n(u) = \begin{cases} L\frac{u}{2\pi}\,\tilde{\mathbf{d}}_n[0] & \text{if } \frac{u}{2\pi} < \frac{1}{L} \\ \left(L\frac{u}{2\pi} - l\right)\tilde{\mathbf{d}}_n[l] + \left(l+1-L\frac{u}{2\pi}\right)\tilde{\mathbf{d}}_n[l-1] & \text{if } \frac{l}{L} \le \frac{u}{2\pi} < \frac{l+1}{L},\ l \in \{1,\dots,L-2\} \\ \left(L - L\frac{u}{2\pi}\right)\tilde{\mathbf{d}}_n[L-2] & \text{if } \frac{L-1}{L} \le \frac{u}{2\pi} < 1, \end{cases} \tag{3.11}$$

$$\tilde{\mathscr{D}}_n(u) = \tilde{\mathscr{D}}_n(u + k2\pi),\ \forall k \in \mathbb{Z}. \tag{3.12}$$

We impose constraints on the values $\tilde{\mathbf{d}}_n[:]$ to fulfil the monotonicity assumption on $\mathscr{D}_n$ (stemming from the hypothesis that images farther away in phase look less similar):

$$\tilde{\mathbf{d}}_n[0] \geq 0 \tag{3.13}$$

$$\tilde{\mathbf{d}}_n[l] \geq \tilde{\mathbf{d}}_n[l-1] \qquad \forall l \in \{1,\ldots,L_n^{\max}\} \tag{3.14}$$

$$\tilde{\mathbf{d}}_n[l] \leq \tilde{\mathbf{d}}_n[l-1] \qquad \forall l \in \{L_n^{\max}+1,\ldots,L-2\} \tag{3.15}$$

$$\tilde{\mathbf{d}}_n[L-2] \geq 0, \tag{3.16}$$

where $L_n^{\max}$ is the position of the maximum of the function $\tilde{\mathscr{D}}_n$, which is unknown a priori.

Taking advantage of the local constancy of shape of $\mathscr{D}_n$, we can consider all distance measurements in a local neighbourhood as noisy samples of the distance curve. This allows us to find the values $\tilde{\mathbf{d}}_n[:]$ by solving a weighted linear least squares fitting problem that minimizes:

$$\mathscr{L}_{\text{dist}}(\tilde{\mathbf{d}}_n) = \sum_{i=0}^{N_g-1} \sum_{j=0}^{N_g-1} \left(\mathbf{D}[i,j] - \tilde{\mathscr{D}}_n(\tilde{\theta}_j \ominus \tilde{\theta}_i)\right)^2 w_\sigma\left(\tilde{\theta}_i \ominus \tilde{\theta}_n\right), \tag{3.17}$$

where $\mathbf{D}[i,j] = d(\mathbf{g}[:,:,i], \mathbf{g}[:,:,j])$ is the matrix containing all pairwise image distances between reference frames, and $w_\sigma(u)$ is a Gaussian windowing function of given standard deviation σ that gives more importance to measurements in a close neighbourhood. This step is illustrated in Fig. 3.4a).

If the phase estimates $\tilde{\theta}_n$ are close enough to the real phases, then the $\tilde{\mathscr{D}}_n$ will be good approximations of the distance curves $\mathscr{D}_n$. When solving this step for the first iteration, we initialize $\tilde{\theta}_n$ with the uniform sampling approximation $\hat{\theta}_n$ defined in Eq. (3.7) which gives a mean error of less than 10% of the period if the sorted sequence contains at least 10 images (Mariani, Ernst, Mercader, and Liebling, 2020).

Taking into account the monotonicity constraints in Eqs. (3.13) and (3.16), we use an efficient quadratic cone programming algorithm (O'Donoghue, 2021) to minimize Eq. (3.17). However, writing the constraints requires knowing $L_n^{\max}$, which we find by solving the problem without monotonicity constraints beforehand.

Unlike the volumetric imaging method (Hanslovsky, Bogovic, and Saalfeld, 2016), we do not resample the distance matrix $\mathbf{D}[:,:]$ to compute the loss. This gives a solution that corresponds more closely to the measures, but makes the least squares problem more complex to write in matrix form. The use of a constrained solver allows us to guarantee the monotonicity of the estimate, thus removing the need of an additional truncating step to ensure the validity of the solution. We choose to use all frames of the sequence to compute each curve estimate instead of using an additional windowing that would exclude distant points from the computations. This allows us to obtain a distance curve defined over the whole period, which we will use in the following step to update the phase estimates.

The linear least squares problem is expressed as follows:

$$\left(\mathbf{X}_n^T \mathbf{W}_n \mathbf{X}_n\right) \tilde{\mathbf{d}}_n = \mathbf{X}_n^T \mathbf{W}_n \mathbf{y}_n, \tag{3.18}$$

where the matrices are defined as:

$$\mathbf{y}_n[N_g i + j] = \mathbf{D}[i, j] \tag{3.19}$$

$$\mathbf{W}_n[N_g i + j, N_g i + j] = w_\sigma\left(\tilde{\theta}_i \ominus \tilde{\theta}_n\right) \tag{3.20}$$

$$\mathbf{X}_n[N_g i + j, l] = \begin{cases} l + 2 - \Delta_{ij} & \text{if } l + 1 \leq \Delta_{ij} < l + 2 \\ \Delta_{ij} - l & \text{if } l \leq \Delta_{ij} < l + 1 \\ 0 & \text{otherwise}\,, \end{cases} \tag{3.21}$$

with $i = 0, \ldots, N_g$, $j = 0, \ldots, N_g$ and $l = 0, \ldots, L - 2$, and

$$\Delta_{ij} = L \frac{\tilde{\theta}_j - \tilde{\theta}_i}{2\pi} + kL, \; k \in \mathbb{Z} \;\text{ s.t. } \Delta_{ij} \in [0, L). \tag{3.22}$$

The quadratic programming problem is then:

$$\mathscr{L}_{\text{dist}}(\tilde{\mathbf{d}}_n) = \frac{1}{2} \tilde{\mathbf{d}}_n \mathbf{Q}_n \tilde{\mathbf{d}}_n + \mathbf{c}_n^T \tilde{\mathbf{d}}_n \qquad \text{with } \mathbf{A}_n \tilde{\mathbf{d}}_n \leq \mathbf{b}_n \tag{3.23}$$

$$\mathbf{Q}_n = 2\mathbf{X}_n^T \mathbf{W}_n \mathbf{X}_n \qquad \mathbf{c}_n = -2\mathbf{y}_n^T \mathbf{W}_n \mathbf{X}_n \qquad \mathbf{b}_n = \mathbf{0}, \tag{3.24}$$

with the monotonicity constraint matrix:

$$\mathbf{A}_n = \left[\begin{array}{cccccccc} -1 & & & & & & & \\ 1 & -1 & & & & & & \\ & & \ddots & & & & & \\ & & 1 & -1 & & & & \\ \hdashline & & & -1 & 1 & & & \\ & & & & & \ddots & & \\ & & & & & & -1 & 1 \\ & & & & & & & -1 \end{array}\right] \tag{3.25}$$

of size $(L \times L - 1)$, where the change of sign on the diagonal (highlighted) happens after line $L_n^{\max}$.

Fitting the phases to the curve

Using the results of the previous step, we can now update the phases $\tilde{\theta}_n$ in order to make the measurements $\mathbf{D}[:,:]$ match the estimated distance curve $\tilde{\mathscr{D}}_n$ as closely as possible. We write this as a weighted least squares cost function:

$$\mathscr{C}\left(\tilde{\theta}\right) = \frac{\sum_{n=0}^{N_g-1}\sum_{i=0}^{N_g-1}\left(\mathbf{D}[n,i] - \tilde{\mathscr{D}}_n\left(\tilde{\theta}_i \ominus \tilde{\theta}_n\right)\right)^2 w_\sigma\left(\tilde{\theta}_i \ominus \tilde{\theta}_n\right)}{\sum_{n=0}^{N_g-1}\sum_{i=0}^{N_g-1} w_\sigma\left(\tilde{\theta}_i \ominus \tilde{\theta}_n\right)}. \tag{3.26}$$

We use the same Gaussian windowing function w_σ as in Eq. (3.17) to give less importance to samples far from the reference point of the distance curve. This addresses the fact that these samples are more noisy and less reliable.

Unlike the volumetric imaging method (Hanslovsky, Bogovic, and Saalfeld, 2016), we use the distance metrics rather than the phases to compute this cost. This circumvents the need to use the inverse of $\tilde{\mathscr{D}}_n$, which is not clearly defined as the distance curve is not a bijective function.

We introduce a regularization term derived from the fact that we acquire images consecutively for each channel, at a constant sampling rate. We consider the neighbourhood containing any three reference frames acquired successively. Over this region, we approximate the evolution of the phase over time using a partial Taylor sum of degree 1. As the time delay between the acquisition of consecutive reference frames is constant, according to this local model the phase step between these frames should be nearly constant. We express this as the following regularization cost:

$$\mathscr{R}\left(\tilde{\theta}\right) = \frac{1}{C(N-2)}\sum_{c=0}^{C-1}\sum_{n=1}^{N-2}\left(\left(\tilde{\theta}_{c,n+1} \ominus \tilde{\theta}_{c,n}\right) - \left(\tilde{\theta}_{c,n} \ominus \tilde{\theta}_{c,n-1}\right)\right)^2. \tag{3.27}$$

We do not use the scaling factor introduced in the volumetric imaging method (Hanslovsky, Bogovic, and Saalfeld, 2016), as its purpose is to correct for varying degrees of noise between images, and localized artefacts on some frames. Given the imaging method we use, the noise level is comparable in all images, and using this scaling factor would only add more complexity to the algorithm with no expected major performance improvement.

We combine the two cost functions using a regularization strength λ:

$$\mathscr{L}_{\text{phase}}\left(\tilde{\theta}_n\right) = (1-\lambda)\mathscr{C}\left(\tilde{\theta}_n\right) + \lambda\mathscr{R}\left(\tilde{\theta}_n\right). \tag{3.28}$$

We minimize this objective function using a gradient descent approach. More specifically, we use the Adam adaptive learning rate optimizer (Kingma and Ba, 2015) to find an optimal estimate for $\tilde{\theta}_n$.

Since we update the phase estimates in this step, the distance curve $\tilde{\mathscr{D}}_n$ is no longer valid and we must recompute it with the new phases. We iteratively minimize Eqs. (3.17) and (3.28) until convergence of the phase estimates. We describe a reliable way to choose the values of the different hyperparameters for minimization in Section 3.3.6.

3.3.4 Reconstructing signal sequences

As established in Section 3.3.1, we can directly assign the reference phase estimates $\tilde{\theta}_n$ to their associated signal frames, as shown in Fig. 3.1 f), under the assumption that the phase delay between reference and signal images is constant. Knowing the phase of each signal frame allows us to reconstruct virtual high frame rate sequences for each channel by sorting the images in increasing phase order. It also allows us to synchronize the different channels, as the phases give us the position of each frame in any channel with respect to all the frames in all the channels, as shown in Step f) of Fig. 3.1.

Simply sorting the frames in each channel according to their phase gives non-uniformly sampled sequences that are not suitable for being viewed as videos. In order to reconstruct high frame rate synchronized videos of each channel, we resample the sorted sequences at an arbitrary fixed rate. Each video frame contains the signal frame of the corresponding channel with the closest possible phase, as shown in Step g) of Fig. 3.1. In doing so, some images may be repeated in undersampled regions, and some may be dropped in oversampled regions depending on the chosen video frame rate. However, the reconstructed videos will be more faithful to the imaged signals, with fewer of the distortion artefacts described in Fig. 3.3. We stress that the resampled video is meant mostly for displaying, and that using the phases directly is more precise to measure the time intervals between events in the heartbeat.

We can use the phase estimates to compute the average frequency of the imaged signal over any region of the sequence. Indeed, if we plot the phases in acquisition order with respect to time, they follow the linear trend of $\theta(t)$ and the average angular frequency of the signal $\bar{\omega}$ at any time corresponds its slope. As the estimated phases are contained within a single period, the phase-time plot has the shape of a sawtooth wave rather than a line. This is due to the fact that the phases drop back to 0 at the start of every new period. We must therefore "unfold" the phases before computing the slope as follows:

$$\tilde{\theta}'_{c,n} = \tilde{\theta}_{c,n} + \phi_{c,n} \tag{3.29}$$

$$\phi_{c,1} = 0 \qquad \forall c \in \{0,\ldots,C-1\} \tag{3.30}$$

$$\phi_{c,n} = \begin{cases} \phi_{c,n-1} + 2\pi & \text{if } \tilde{\theta}_{c,n} - \tilde{\theta}_{c,n-1} < -\frac{\pi}{2} \\ \phi_{c,n-1} & \text{otherwise} \end{cases} \qquad \forall n \in \{1,\ldots,N-1\}, \tag{3.31}$$

where $\tilde{\theta}'_{c,n}$ are the unfolded phases for each channel, and $\phi_{c,n}$ represents how many periods have elapsed before its associated frame. It increments of a full period when the difference between two consecutive frames is lower than a threshold of $-\frac{\pi}{2}$, which prevents counting too many periods because of small errors in the phase estimates. We can then compute the average angular frequency over a chosen region of the signal by fitting a linear model on the points $(t_{c,n}, \tilde{\theta}'_{c,n})$ within that interval using linear least squares. The slope of the model directly gives $\bar{\omega}$ over the analysed region.

The frequency information, which was missing in previous sorting-based virtual high frame rate methods (Q. Zhang and Pless, 2005; Mariani, Marelli, Jaques, Ernst, and Liebling, 2021; Mariani, K. Chan, Ernst, Mercader, and Liebling, 2019), allows to study the heart rate of the sample. It also allows converting phases into time units, which is better suited for analysing the dynamics of the heartbeat.

3.3.5 Theoretical analysis of the error caused by heart rate variability

In practice, as the heart rate varies, so does the phase delay between paired frames. This introduces errors in the reconstruction step described in the previous section, as it is based on the assumption that these phase variations are negligible. In order to find in which range this approximation is valid, we want to quantify the impact of the cardiac frequency variability on the performance of our method.

In order to measure only the error caused by heart rate variability, we imagine that our phase estimation algorithm finds the exact phases of each reference frame: $\tilde{\theta}_{c,n} = \theta^g_{c,n}$. We can write the phase delay between a reference frame and its associated signal frame as the product of the time delay between the images Δ_C and the instantaneous frequency of the heartbeat during that delay $\omega_{c,n}$. More specifically, this instantaneous frequency is the sum of the average frequency over the whole imaging period $\bar{\omega}$ and a deviation ϵ that can be represented as samples of a zero-mean statistical distribution: $\omega_{c,n} = \bar{\omega} + \epsilon_{c,n}$. We introduce this in Eq. (3.2) to compute the error $\mathscr{E}_\epsilon$ that results from the variability in the heart rate:

$$\mathscr{E}_\epsilon = \min_{\Theta} \frac{1}{NC} \sum_{c=0}^{C-1} \sum_{n=0}^{N-1} \left| \theta^g_{c,n} + (\bar{\omega} + \epsilon_{c,n})\Delta_C \ominus \left(\theta^g_{c,n} + \Theta\right) \right|. \tag{3.32}$$

From the definition of the phase operator in Eq. (3.3), it stems that $|\alpha \ominus \beta| \leq |\alpha - \beta|$, and we can find an upper bound for our error:

$$\begin{aligned} \mathscr{E}_\epsilon &\leq \min_{\Theta} \frac{1}{NC} \sum_{c=0}^{C-1} \sum_{n=0}^{N-1} \left| \theta^g_{c,n} + (\bar{\omega} + \epsilon_{c,n})\Delta_C - \left(\theta^g_{c,n} + \Theta\right) \right| && (3.33) \\ &\leq \min_{\Theta} \frac{1}{NC} \sum_{c=0}^{C-1} \sum_{n=0}^{N-1} \left| \bar{\omega}\Delta_C - \Theta + \epsilon_{c,n}\Delta_C \right|. && (3.34) \end{aligned}$$

Given that $\epsilon_{c,n}$ is sampled from a distribution with a zero mean, the right term is minimized when $\Theta = \bar{\omega}\Delta_C$, leaving the final expression of the upper bound:

$$\mathscr{E}_\epsilon \leq \frac{1}{NC} \sum_{c=0}^{C-1} \sum_{n=0}^{N-1} \left| \epsilon_{c,n}\Delta_C \right| = \frac{\Delta_C}{NC} \sum_{c=0}^{C-1} \sum_{n=0}^{N-1} \left| \epsilon_{c,n} \right|. \tag{3.35}$$

Thus, we can compute the upper bound of the error introduced by heart rate variability if we know the statistical distribution of the cardiac frequency during the acquisition. For practical applications, it is easier to describe it by using $\omega_{\min}$ and $\omega_{\max}$, the minimal and maximal values (respectively) that this frequency can take over the imaging process.

In the worst possible case where the heart rate is maximal during half of the acquisition and jumps without transition to its minimal value for the other half of the acquisition, the bound in Eq. (3.35) becomes:

$$\mathscr{E}_\epsilon \leq \Delta_C \frac{\omega_{\max} - \omega_{\min}}{2}. \tag{3.36}$$

If the cardiac frequency is uniformly distributed over its range during the acquisition (e.g., if the heart linearly accelerates at a constant rate), the bound gets smaller:

$$\mathscr{E}_\epsilon \leq \Delta_C \frac{\omega_{\max} - \omega_{\min}}{4}. \tag{3.37}$$

And if its distribution is a normal (e.g., if the heart rate oscillates around its average value), using the three-sigma rule to set the interval, the bound tends towards an even lower value:

$$\mathscr{E}_\epsilon \leq \Delta_C \sqrt{\frac{2}{\pi}} \frac{\omega_{\max} - \omega_{\min}}{6}. \tag{3.38}$$

In the case of a healthy heart that smoothly accelerates or decelerates, and may settle around a given cardiac frequency, the distribution lies between the uniform and the normal.

For example, we compute the upper bound in a realistic scenario with a delay of $\Delta_C = 10\,\text{ms}$ between reference and signal images, and an average heart rate of 2.5 beats per second, considering the cardiac frequency distribution to be at worst uniform. If the heart rate stays between 0.5 and 4.5 beats per second, the error due to cardiac frequency variability will be lower than 1% of the period: $\mathscr{E}_\epsilon \leq 10^{-2} \frac{4.5-0.5}{4} 2\pi$.

We consider a different scenario with all parameters equal except the interval $\Delta_C = 0.033\,\text{s}$. This corresponds for example to using a camera with a frame rate of 30 fps in a setup where triggering the reference acquisition at the end of the exposure time is not possible. To guarantee the same 1% upper bound as above in this scenario, the heart would have to stay between 1.9 and 3.1 beats per second: $\mathscr{E}_\epsilon \leq 0.033 \frac{3.1-1.9}{4} 2\pi$.

We can use these bounds to predict what performance to expect, based on rough estimates of the range of heart rates that are expected or measured. This error induced by frequency variability will then be combined with the precision of the phase estimation algorithm when measuring the overall error of our method.

3.3.6 Reliable computation of the hyperparameter values

The performance of our phase estimation algorithm depends on the values of multiple hyperparameters. In order to be able to apply our method in practice, we propose a reliable way to compute optimal values for these parameters based on the acquired images. This guarantees the reproducibility of our results, and allows to apply our method in new scenarios without requiring manual hyperparameter tuning.

The regularization strength λ is the parameter with the most impact on performance. We find its optimal value using an L-curve (Hansen, 1992). This consists in plotting the values of the residual cost from Eq. (3.26) against the final regularization cost from Eq. (3.27) for multiple values of λ. The curve resembles the shape of an L, as seen in Fig. 3.6c). To compute the optimal λ, we first apply min-max normalization to both costs then multiply the regularization cost by a factor 10. This is because in our case, the regularization corresponds to a strong prior about the system, and must play a bigger role in selecting λ. We then choose the point in the rescaled L-curve that is the closest to the origin as the optimal λ.

The width of the Gaussian window σ used in Eqs. (3.17) and (3.26) also influences performance. Indeed, a window too large will decrease precision due to too much averaging, while a window too narrow will be too sensitive to noise due to not including enough measurements. To find good value for σ, we plot the residual error obtained when minimizing Eq. (3.17) with the initial phase estimates, using different window widths up to 2π. We select the last local minimum in that curve as the best value for σ. This represents the best fitting performance, ignoring the minima obtained with very small σ that correspond to overfitting the noise in the data.

The number of points L used to represent the distance curve has a simpler impact on performance. The precision of the method tends to increase with L, as the curve approximation gets finer, but the convergence time of the method increases as well due to the added complexity. Generally speaking, we have found that our algorithm fails to converge when L is too big (typically bigger than $2N_g$), and that using more than $L = 200$ does not seem to improve performance further.

To enhance the performance of phase estimation, we can restart our iterative algorithm after convergence, using the obtained phase estimates as the initial guess for the new run. We call this additional step a precision pass. When restarting the minimization, we halve the width of the Gaussian window σ and double L to get finer phase estimates. Even with the higher L and smaller σ, the minimization converges as it starts from a better initial guess. This allows getting better precision using bigger values of L that would have caused convergence issues if used directly on the initial uniform phase estimates.

The update step size of the gradient descent optimizer does not impact the precision much (thanks to the adaptive scaling of the Adam optimizer), but choosing a good value can speed convergence up. We have found that a default value of 10^{-3} works well in all scenarios.

3.4 Methods: fluorescence cardiac imaging simulation

Previous works used simplified simulation models of the heart for performance characterization, consisting of a circle (or cylinder) textured with a sinusoidal pattern, with a radius varying over time to represent beating (Mariani, Marelli, Jaques, Ernst, and Liebling, 2021; Liebling, A. S. Forouhar, Gharib, S. E. Fraser, and M. E. Dickinson, 2005; Mariani, K. Chan, Ernst, Mercader, and Liebling, 2019; Mariani, Ernst, Mercader, and Liebling, 2020). These models lack properties such as irregular texture, asymmetric contraction or small variations in the periodicity that we find in experimental cardiac images. In order to test the different methods and characterize their performance as reliably as possible, we designed a simulation tool that reproduces the relevant properties of fluorescent beating heart images. We aim to synthesize a smooth organic-like textured object that deforms following a nonlinear contraction wave that propagates periodically through time and space asymmetrically. We also want to generate multiple channels that contain variable amounts of mutual information, and to emulate the small variations in the periodicity of the heartbeat.

3.4.1 Simulating a beating heart section

We define a function $h(x, y)$ over a two-dimensional space to simulate a heart section:

$$h(x, y) = e(x, y)\, \sigma(x, y), \tag{3.39}$$

where $e(x, y)$ represents an ellipse with a thick outline and $\sigma(x, y)$ is a two-dimensional simplex noise (Perlin, 1985; Perlin, 2002) used for texturing. Simplex noise is a gradient noise function that generates visually isotropic and continuous textures, and is commonly used in computer-generated graphics to procedurally create natural looking images. Its properties make it suitable for creating synthetic textures with realistic features in biomedical imaging (Barufaldi, Abbey, Lago, Vent, Acciavatti, P. R. Bakic, and Maidment, 2021; Dustler, Förnvik, and Lång, 2018; Dustler, P. Bakic, Petersson, Timberg, Tingberg, and Zackrisson, 2015; Abdolhoseini, Kluge, Walker, and S. J. Johnson, 2019), and we use it here to generate an organic-like texture for the heart walls.

For later use in our simulations, we introduce a periodic continuously smooth pulse function obtained by composing a sigmoid and a sine function:

$$\operatorname{sigsin}(\theta) = \frac{\operatorname{sig}\left((\sin(\theta) - b)\, s\right)}{\operatorname{sig}\left((1 - b)\, s\right)} \tag{3.40}$$

$$\operatorname{sig}(x) = \frac{1}{1 + \exp(-x)}, \tag{3.41}$$

where b is the sigsin bias and s its slope. b controls the width of the pulse, with higher bias generating shorter pulses, while s affects the overall slope of the rise and fall of the function, with higher slope creating a fast rising pulse.

In order to simulate the propagation of heart contraction through the section, we create a continuous space transformation that preserves topology. This transformation varies on x and θ in order to model the propagation of the contraction along one spatial dimension in time. It consists locally of a small rotation followed by a scaling, representing the twisting and shrinking of the heart walls that happen during contraction. This transformation can be expressed in matrix form:

$$\mathbf{T}(x,\theta) = \mathbf{R}(x,\theta) \cdot \mathbf{S}(x,\theta) \tag{3.42}$$

where the rotation and scaling matrices $\mathbf{R}(x,\theta)$ and $\mathbf{S}(x,\theta)$ are defined as follows:

$$\mathbf{R}(x,\theta) = \begin{bmatrix} \cos(\phi(x,\theta)) & -\sin\left(\phi(x,\theta)\right) \\ \sin\left(\phi(x,\theta)\right) & \cos\left(\phi(x,\theta)\right) \end{bmatrix} \qquad \mathbf{S}(x,\theta) = \begin{bmatrix} 1+S_x(x,\theta) & 0 \\ 0 & 1+S_y(x,\theta) \end{bmatrix}. \tag{3.43}$$

The rotation angle $\phi(x,\theta)$ and shrinking factors $S_x(x,\theta)$ and $S_y(x,\theta)$ are governed by a sigsin pulse:

$$\xi(x,\theta) = A \operatorname{sigsin}\left(\theta + \frac{x}{\lambda} + \theta_0\right), \tag{3.44}$$

where A is the amplitude of the pulse, θ_0 is its initial phase shift, and λ is the spatial wavelength of the contraction pulse. The parameters A, θ_0, b and s are set independently for the functions ϕ, S_x and S_y, allowing to adjust the model to obtain a more realistic contraction.

Using the defined transform, we then compute the simulated reference signal as:

$$g^{\text{sim}}(x,y,\theta) = h\left(\mathbf{T}(x,\theta) \cdot \begin{bmatrix} x \\ y \end{bmatrix}\right). \tag{3.45}$$

In order to generate multiple fluorescent channels, we multiply this reference signal with two-dimensional masking functions $m_c(x,y)$:

$$f_c^{\text{sim}}(x,y,\theta) = h\left(\mathbf{T}(x,\theta) \cdot \begin{bmatrix} x \\ y \end{bmatrix}\right) m_c\left(\mathbf{T}(x,\theta) \cdot \begin{bmatrix} x \\ y \end{bmatrix}\right). \tag{3.46}$$

This method to simulate multiple fluorescent channels is not biologically plausible, but serves only as a means to easily control the amount of overlap between channels through the shape of the masks. This is sufficient to characterize the performance of our algorithms, and a biologically more accurate model is unnecessary in this scope. Simulated heart sections in different positions are visible in Fig. 3.2.

3.4.2 Modelling the heart rate variability

We model the variability of the heartbeat as the combination of two contributions: a random jitter that affects the phase locally, and a random acceleration that affects the phenomenon over a longer time span. The first is motivated by the imperfect nature of a biological process that creates local variations within a cycle, while the latter represents a slow variation of the average angular frequency over time. Using this model, we express the simulated phase at time t as:

$$\theta^{sim}(t) = \theta_0^{sim} + \int_0^t \omega_0^{sim} + \sigma_\omega W_\omega(t)\, dt + \sigma_\theta W_\theta(t), \tag{3.47}$$

with θ_0^{sim} the starting phase and ω_0^{sim} the initial angular frequency. We use Wiener stochastic processes to model the variability in phase $W_\theta(t)$ and in frequency $W_\omega(t)$ of the system, defined by:

$$W(0) = 0 \qquad W(t+u) - W(t) \sim \mathcal{N}(0,u),\ \forall t, u > 0, \tag{3.48}$$

with $\mathcal{N}(0,u)$ is a normal distribution of mean 0 and variance u. Their amplitudes σ_θ and σ_ω allow to control the overall amount of randomness in the simulation, as well as the strength ratio between the two sources of stochasticity. We chose Wiener stochastic processes because of their property to generate continuous paths, which is essential for modelling a heartbeat cycle as the motion of the heart is itself continuous. We illustrate the evolution of phase variability in multiple scenarios using different strengths for phase and frequency deviations in Fig. 3.5.

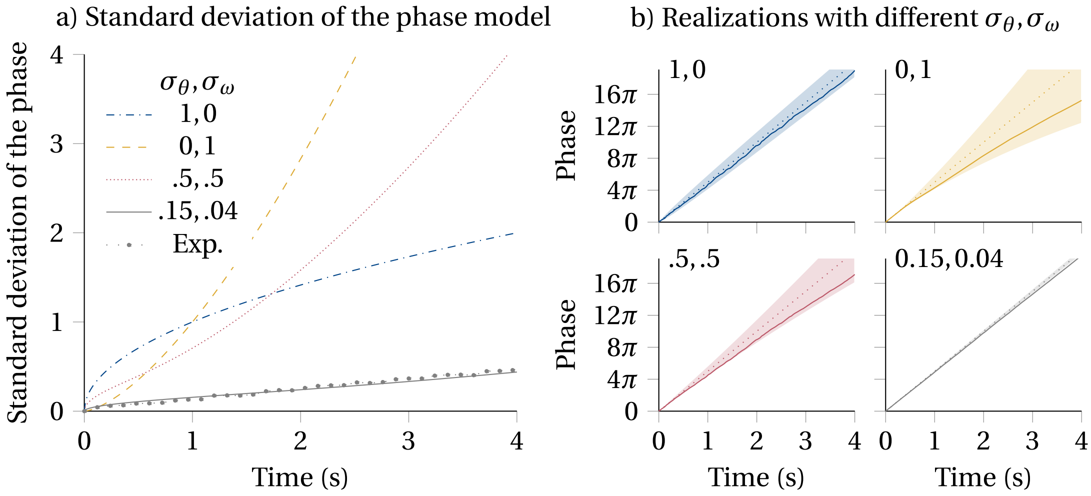

Figure 3.5: Stochastic phase model for simulating heart rate variability. Our model uses two parameters to represent variability: a phase deviation σ_θ and a frequency deviation σ_ω. Plots illustrate the model for a simulated 2.5 beats per second signal. a) Varying the ratio of these contributions changes how fast the phase uncertainty increases, with σ_ω generating a rapidly increasing variance. Measurements on experimental data match the simulation model. b) Realizations of our model in different scenarios show that σ_θ generates local noise, while σ_ω contributes to bigger smoother variations. The bottom-right panel illustrates values matching experimental measurements.

Using the phase $\theta^{sim}(t)$, we sample the simulated signals g^{sim} and f_c^{sim} as defined in Eqs. (3.5) and (3.6) to obtain simulated image pairs $\left(\mathbf{g}_c^{\text{sim}}[:,:,n], \mathbf{f}_c^{\text{sim}}[:,:,n]\right)$. Thanks to the biologically inspired simulation models, these images share the relevant properties allowing us to use them to characterize the methods exposed in this work.

3.5 Experiments and results

We characterized our method on both synthetic and real images. The following sections detail how we conducted these experiments, and discuss the obtained results. For fast computations, we implemented our gradient descent algorithm using JAX (Bradbury, Frostig, Hawkins, M. J. Johnson, Leary, Maclaurin, Necula, Paszke, VanderPlas, Wanderman-Milne, and Q. Zhang, 2018), a high-performance numerical computing framework accelerated with parallel computing. When running experiments, we noticed that keeping the state of the Adam optimizer from one iteration of the phase estimation to the next instead of resetting its parameters greatly speeds up convergence.

3.5.1 Validation of the simulation model

To confirm the plausibility of our phase simulation model, we measured the heart rate variability in a high-speed dataset that was acquired for previous works (Mariani, K. Chan, Ernst, Mercader, and Liebling, 2019). It consists in a 55 hpf wild type zebrafish embryo imaged on a Leica DMR microscope equipped with a FASTCAM SA3 camera using brightfield at 1000 fps. We randomly isolated a subsequence of 4 seconds in the data, and identified its first period using image similarity. We assigned linear phases from 0 to 2π to the frames of this first period, and then estimated the phase of all the other images in the subsequence as the phase of the most similar frame in the reference period.

We used these estimations to measure the phase deviation with respect to a perfectly periodic repetition of the first period. We repeated this measurement many times with different starting frames to build a relevant statistical representation of the experimental data. As shown in Fig. 3.5a), the measurements closely match our model when using the values $\sigma_\theta = 0.15$ and $\sigma_\omega = 0.04$. This corresponds to a relatively stable heart rate, as illustrated in the bottom-right panel of Fig. 3.5b).

3.5.2 Characterization on synthetic data

We simulated the imaging of a heart with an average of 2.44 beats per second in two fluorescent channels ($C = 2$) using a camera with a frame rate of 20 frames per second (fps). We generated images of size 256×256, and added Poisson noise to each pixel before quantizing their values to 12 bits. We used $\sigma_\omega = 0.04$ and $\sigma_\theta = 0.15$ to model the variability of the phase with comparable dynamics to what we measured on experimental data. We used a camera transfer time of 5 ms as the gap between reference and signal frames for PAAQ. We simulated imaging sequences of variable length, repeating each experiment with 10 different random seeds to generate confidence intervals.

We used $L = 50$ for the phase estimation, and set the other parameters using the strategy described in Section 3.3.6. Before applying the method we downsampled the images by a factor 2 with spatial averaging, to reduce the computation time and noise.

As a baseline for performance comparison, we considered the ungated imaging of the fluorescent channels at low frame rate without using PAAQ, in addition to the acquisition of the brightfield images without illumination patterns to serve as registration reference. We sorted each channel separately for virtual high frame rate, and synchronized the obtained sequences using a temporal registration procedure based on maximizing the mutual information between the fluorescent channels and the brightfield sequence. This corresponds to a sequential approach to solving virtual high frame rate and multichannel registration, as discussed in Section 3.1. As the direction of the sequences is not given by the sorting, we applied the registration to all possible combinations of orientations for the channels, and kept only the solution with the maximum mutual information. We used a simple temporal upsampling of the sequences to allow for sub-frame precision registration.

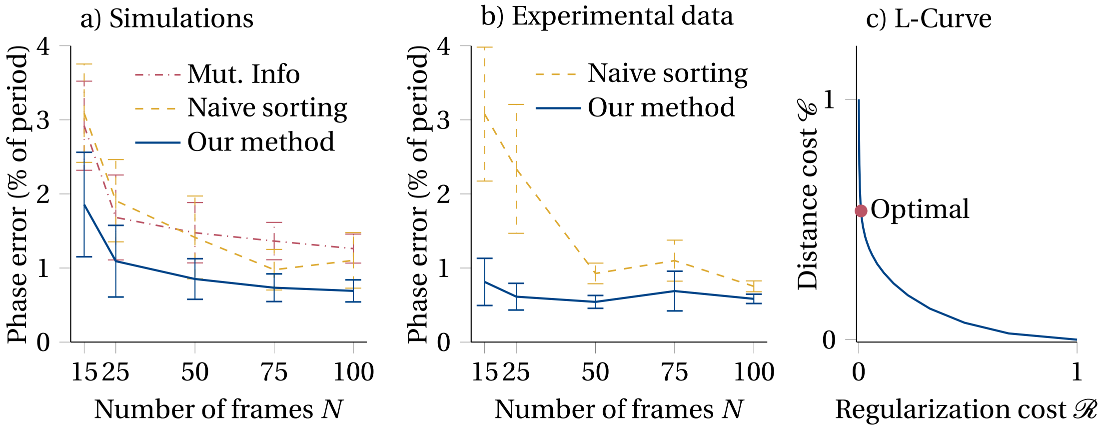

Figure 3.6: Quantitative results on synthetic and experimental data. a) On synthetic data, our method yields a relative precision improvement of 30% over competing methods, independently of the number of frames acquired. b) On experimental data, our method performs similarly, with a relative precision improvement of up to 50% (for $N \leq 25$). c) We choose the regularization strength for phase estimation using an L-curve, emphasizing the regularization cost due to strong priors.

Figure 3.6 a) compares how the phase estimation performance of the methods varies with the number of frames acquired. The naive sorting method refers to the solution given by using PAAQ imaging and estimating the phases with the uniform sampling assumption (Steps b-c in Fig. 3.1, i.e., no phase correction algorithm). The mutual information-based method and the naive sorting obtain a similar performance over the whole range of experiments, which is expected as they are both limited by the uniform sampling assumption for phase estimation. Our method consistently improves on this performance, giving a relative 30% reduction of the phase error overall. Using Eq. (3.35), we computed that the error caused by heart rate variability is bounded to a maximum of 0.13% of one period on our simulated data. Therefore, the plotted phase error mostly characterizes the precision of our phase estimates, showing the benefit of using a non-uniform phase estimation algorithm.

We computed the initial average frequency of the simulated signal over the 2 first acquired periods using the method detailed in Section 3.3.4. The obtained relative error is directly tied to the error on phase estimation, dropping from 2% to below 1% when $N \geq 50$. We repeated these experiments with increasing Poisson noise strength, with no noticeable performance impact on any of the methods.

3.5.3 Validation on high-speed experimental data

In order to confirm the applicability of our method on real data, we used the same high-speed dataset as in Section 3.4.2. To simulate slow PAAQ imaging, we randomly picked a frame as the start, then sampled the dataset every 100 frames to obtain the reference sequence, corresponding to a simulated camera frame rate of 20 fps. To emulate the paired signal frames, we sampled the data similarly starting 5 frames after the initial reference frame, corresponding to a simulated camera transfer time of 5 ms. We repeated this process with a different random starting frame to simulate the acquisition of a second channel. In this experiment, the signal sequences are all acquired using brightfield, but it does not affect the conclusions that we can draw on the precision of our phase estimation, as we use only the references for computation.

In order to generate the reference phases, we isolated a full period in the high-speed data starting from the first reference frame of the first channel, using image similarity to automatically find the start of the next period. We assigned linear phases from 0 to 2π to the frames in this initial period. We then estimated the ground-truth phase of all signal frames in all channels as the phase of the most similar frame (by image distance) in the reference period.

We simulated the slow imaging with varying sequence lengths, repeating each experiment 10 times with different random seeds to compute confidence intervals. For the phase estimation, we used $L = 50$ and set the other hyperparameters according to Section 3.3.6, running 2 additional precision passes for best performance.

Figure 3.6 b) compares the phase estimation error obtained with naive sorting (see Section 3.5.2) and our method. Due to the absence of separate fluorescent channels in the data, we did not evaluate the mutual information approach. The performance of naive sorting is consistent with what we measured on the simulated dataset. Our method achieves even better performance here, obtaining a phase error below 1% of the period for all sequence lengths. This represents a significant relative error reduction of more than 50% for sequences with few acquired frames ($N \leq 25$).

Figure 3.6 c) illustrates the L-curve obtained for an experiment with $N = 50$. The highlighted point corresponds to the selected value of regularization strength λ used for the phase estimation. It lies close to the vertical axis, as the regularization corresponds to a strong prior on the data.

3.5.4 Results on low-frame rate experimental data

We performed experiments with zebrafish (Danio rerio) embryos from animals held at Institute of Anatomy (National Licence Number 35) from the University of Bern. Adult fish needed for breeding were raised and maintained at maximal 5 fish/L with the following environmental conditions: 27.5 to 28 °C, with 14 h of light and 10 h of dark, 650 to 700 µs/cm, pH 7.5 and 10% of water exchange daily. Adult Tg(*fli1a:GFP*)^{y1Tg} (Lawson and Weinstein, 2002) and Tg(*myl7:mRFP*)^{ko08Tg} (Rohr, Bit-Avragim, and Abdelilah-Seyfried, 2006) double transgenic zebrafish were crossed to obtain homozygous embryos. Eggs were collected within 30 min and kept in E3 medium (5 mM NaCl, 0.17 mM KCl, 0.33 mM CaCl2, 0.33 mM MgSO4) with Methylene Blue (10^{-5}%). 30 min after egg collection, those that did not transition to two-cell stage were removed. The next day, at 24 h after removing the screens, chorions were removed by incubating 2 mg/mL Pronase in E3 medium (about 3 min) until gentle shaking of the Petri dish freed the larvae from the chorion. Subsequently, embryos of the same developmental stage were selected and the most frequent stage was estimated (C. B. Kimmel, Ballard, S. R. Kimmel, Ullmann, and Schilling, 1995). Before 24 hpf, we added 0.003% 1-phenyl-2-thiourea (PTU, Sigma-Aldrich) to avoid pigmentation.

We anaesthetized the embryos with Tricaine at 0.08 mg/mL, pH 7 and embedded the embryos with the anterior side (head) up in a fluorinated ethylene propylene tube in 1% low melting agarose. We imaged the sample on the custom OpenSPIM platform described in Section 1.3.

We used the modulation device introduced in Chapter 1 to generate the PAAQ illumination patterns depicted in Fig. 3.2. We acquired images with a frame rate of 14.42 Hz, corresponding to a 60 ms exposure time Δ_E and a camera transfer time of 10 ms. We set a pulse width Δ_P of 10 ms for the signal illumination, and a linear ramp containing 3 pulses of 10 ms separated by 15 ms gaps for the reference pattern. We used $L = 50$ for the phase estimation, and used Section 3.3.6 to set the other hyperparameters, running one additional precision pass for performance.

Our first experiment compares the performance of our algorithm to the mutual information-based method when processing 2 fluorescent channels with varying number of images. We extract a region of size 512 × 512 centred on the heart and downsample the image by a factor 2 before using the algorithms. Figure 3.7 a) shows a frame in each of the reconstructed videos for $N = 100$ and $N = 20$, corresponding to a single estimated phase in the signal. With a high number of frames, the two methods give visually similar results, however with fewer frames the mutual information video contains aberrations in the synchronization of the channels. This is highlighted at the cross mark, where the myocardium (red) appears to intersect with the

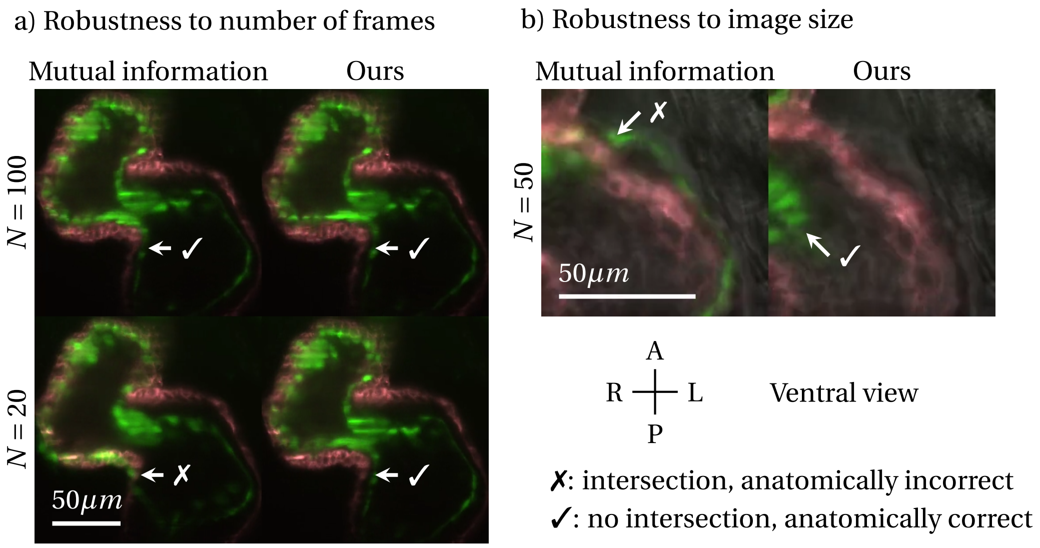

Figure 3.7: Comparison of the reconstructed videos using our method and a mutual information approach, 4 dpf transgenic Fli1V/Myl7mR zebrafish heartbeat. Cross marks highlight anatomically implausible spots where myocardium (red) intersects with endocardium (green). a) With high number of frames both methods give similar results, but mutual information gives aberrations with $N = 20$, while our method is not impacted. b) When working on small images, mutual information gives incorrect channel synchronization between red and green, while our method generates a plausible solution.

endocardium (green), which is anatomically implausible. With the same number of frames, our method does not generate this artefact, and stays consistent with the solutions computed with more frames. This illustrates the better robustness of our algorithm with respect to the amount of images available. Moreover, the solution obtained using mutual information contains anatomical aberrations even when using a high number of frames. In the same situation, our method generates a video that does not contain visible anatomical aberrations, benefitting from the better phase estimates.

Our second experiment compares the performance of the two methods when the amount of information available is limited. We achieve this by restricting the field of view to only a smaller part of the heart in a region of interest of size 256×256. We also downsample by a factor 2 before applying both methods. For this experiment, we acquired $N = 50$ images in $C = 3$ channels, treating the brightfield modality as an additional signal channel. As shown in Fig. 3.7 b), the mutual information reconstruction fails to correctly synchronize the red and green channels, even though the synchronization between red fluorescence and brightfield is correct. In the highlighted area, the myocardium (red) and endocardium (green) visibly intersect in an impossible way in the solution given by mutual information, while in the corresponding frame obtained with our method the channels are correctly synchronized. In the smaller image, the amount of mutual information between channels is lower, partly because there are not enough relevant data available to build a meaningful statistical representation of the sequences. Our method does not require mutual information between the channels to give a correct solution, making it suitable to work on data with limited size but also in cases where the different channels contain entirely uncorrelated signals.

3.6 Discussion

Our method is able to generate multichannel videos with a virtual high frame rate with high temporal precision. As derived in Section 3.3.5, the channel synchronization error caused by heart rate variability is bounded and depends only on the transfer time of the imaging device. Therefore, the overall performance of our method mainly depends on the total number of images acquired, which can be chosen as high as needed to achieve a desired temporal resolution.

Prospective optical gating methods (Taylor, Saunter, Love, Girkin, Henderson, and Chaudhry, 2011; Taylor, Girkin, and Love, 2012) provide a high accuracy, with a better control on when the cardiac cycle is sampled. They can also be used for 3D imaging, by using the same trigger to synchronize videos at different depths. However, our technique has lower hardware requirements, making it more cost-effective and easier to integrate into existing imaging platforms. Our method is currently limited to 2D imaging, although it might be generalized by implementing depth sweeping during imaging, combined with a method that removes the induced scan aberrations (Mariani, Ernst, Mercader, and Liebling, 2020). Another possibility would be to use an imaging modality with a large depth of field for the reference frames, allowing to accurately synchronize neighbouring sections of the volume through a temporal registration step.

Methods that do not use PAAQ imaging (such as the mutual information synchronization used for comparison) have no hardware requirements at all. Nevertheless, we have shown that our technique is more robust to the size of the data available, both in terms of amount of images and their size. At the cost of a more complex acquisition procedure requiring illumination switching, it guarantees a better temporal resolution. Therefore our method offers a compromise, giving high accuracy while still keeping the hardware cost and complexity low.

Although we only described the use of fluorescence imaging paired with a brightfield reference, PAAQ could exploit other combination of imaging modalities, provided that it is possible to switch quickly enough between them. The reference signal could be any modality, but one that induces minimal phototoxicity, like brithgfield, is clearly preferable.

3.7 Conclusion

We have introduced a method to reconstruct virtual high frame rate multichannel videos from slow single-channel acquisition of periodic processes. We have validated our method on experimental data, and shown that it does not contain the anatomical aberrations visible when using a state-of-the-art technique based on mutual information.

PAAQ is based on alternating acquisitions between a common reference modality and fluorescence channels, allowing to reconstruct synchronized videos of multiple modalities without requiring mutual information between channels. We have shown that the synchronization error caused by variability in the signal periodicity is small, even with big frequency variations if the transfer time of the imaging device is small enough.

We have proposed an image-only phase estimation algorithm that improves on existing sorting-based virtual high frame rate methods thanks to a non-uniform sampling assumption. We have shown on both synthetic and experimental data that our algorithm brings an overall relative reduction of the phase error by 30% compared to that of techniques that use a uniform sampling assumption. For synthetic experiments, we designed a simulation framework that reproduces the most relevant properties of cardiac fluorescence microscopy, and can model the heart rate variability with statistical properties matching experimental observations. We have proposed a reproducible way to compute the optimal hyperparameters for our method that guarantees good performance without manual tuning.

We released the schematics and drivers for a hardware controller that can modulate light sources and synchronize active illumination with imaging devices. Similarly, we released the simulation framework and the implementation of our algorithms.

Although we developed our method for the specific case of cardiac microscopy, our algorithm is generic and we foresee that it could be applied in other scenarios involving the imaging of a cyclic phenomenon with a slow device.

4 Spatiotemporal light modulation for efficient 3D compressed sensing

This chapter contains the results published in F. Marelli and M. Liebling (2023). “Efficient compressed sensing reconstruction for 3D fluorescence microscopy using OptoMechanical Modulation Tomography (OMMT) with a 1+2D regularization”. Submitted for publication.

4.1 Introduction

Compressed sensing techniques exploit prior knowledge of the imaged object to reconstruct a high quality image (Donoho, 2006; Candes and Wakin, 2008) from fewer measurements than traditional sampling would require. In 3D medical imaging, reconstruction algorithms exploit the specificities of the imaging procedure to yield improved performance and faster convergence speed (Montefusco, Lazzaro, Papi, and Guerrini, 2011).

In fluorescence microscopy, compressed sensing offers the prospect of high quality imaging while also reducing light exposure, which can lead to photobleaching or induce toxicity during in vivo imaging. Furthermore it could increase acquisition speed while taking fewer images (Calisesi, Ghezzi, Ancora, D’Andrea, Valentini, Farina, and Bassi, 2022). Compressive sensing approaches have been proposed for structured illumination microscopy (Meiniel, Spinicelli, Angelini, Fragola, Loriette, Orieux, Sepulveda, and Olivo-Marin, 2017) and widefield imaging (Q. Guo, H. Chen, Y. Wang, Y. Guo, P. Liu, Zhu, Cheng, Yu, S. Yang, M. Chen, and Xie, 2017).

For Selective Plane Illumination Microscopy (SPIM) (Huisken, Swoger, Bene, Wittbrodt, and E. H. K. Stelzer, 2004), compressive sensing has been implemented by collecting axial projections after illuminating the sample with a spatial modulation pattern, either by shaping the illumination with digital micromirror devices (Calisesi, Castriotta, Candeo, Pistocchi, D’Andrea, Valentini, Farina, and Bassi, 2019) or by applying optomechanical modulation (Woringer, Darzacq, Zimmer, and Mir, 2017). The latter approach mechanically moves the sample along the axial (focal) direction while simultaneously modulating the overall light sheet intensity in time, all while keeping the camera shutter open. Measurements are repeated for multiple scans, each with a different modulation pattern. The volume is reconstructed computationally. We will use an approach similar to (Woringer, Darzacq, Zimmer, and Mir, 2017) as it is a particularly attractive way to achieve spatial light modulation. First, it can be implemented with minimal hardware modifications of simple light sheet microscopes such as, e.g., the OpenSPIM microscope (Pitrone, Schindelin, Stuyvenberg, Preibisch, Weber, Eliceiri, Huisken, and Tomancak, 2013). Second, the spatial illumination pattern is created by modulating the intensity in time, which leaves great freedom for designing a suitable modulation pattern. We will refer to this technique as OptoMechanical Modulation Tomography (OMMT).

Despite the potential benefits of compressive sensing approaches, the computational cost of reconstructing 3D volumes is particularly high. Naive extensions of 2D reconstruction schemes to 3D are impractical for the large images that can be acquired with modern microscopes. Furthermore, overly simplified 1D or 2D regularizing terms may not capture the objects' 3D nature, a prior that is key to ensure quality reconstructions.

Recently, in the field of 2D image reconstruction, the introduction of Plug-and-Play (PnP) methods has allowed replacing explicit regularization priors with denoising algorithms for image reconstruction (Venkatakrishnan, Bouman, and Wohlberg, 2013; Kamilov, Bouman, Buzzard, and Wohlberg, 2023). This led to using high-quality denoising algorithms such as Block-Matching and 3D filtering (BM3D) (Dabov, Foi, Katkovnik, and Egiazarian, 2008) to improve image reconstruction (S. H. Chan, X. Wang, and Elgendy, 2017; Pellizzari, Trahan, H. Zhou, S. Williams, S. E. Williams, Nemati, Shao, and Bouman, 2017; Y. Sun, Wohlberg, and Kamilov, 2019). More recently, the use of pretrained deep learning based denoisers has further pushed the performance of PnP methods (Ryu, J. Liu, S. Wang, X. Chen, Z. Wang, and Yin, 2019; K. Zhang, Y. Li, Zuo, L. Zhang, Van Gool, and Timofte, 2022).

Here, we propose an efficient reconstruction method for OMMT that exploits the 3D nature of the data, thereby reducing the reconstruction artefacts from which this techniques suffered (Woringer, Darzacq, Zimmer, and Mir, 2017). To that end, we introduce a hybrid 1+2D regularization function that takes into account the anisotropy of the problem, while allowing efficient 3D computations thanks to parallel computing. We also implement a high quality, but computationally expensive, Plug-and-Play (PnP) regularization function to serve as a benchmark against which we compare our proposed approach. We evaluate our method for accuracy and scaleability on simulated and experimental data.

4.2 Method

We briefly recall the method by Woringer et al. (Woringer, Darzacq, Zimmer, and Mir, 2017), who combine temporal illumination modulation combined with focal plane scanning to generate patterned illumination along the depth axis and implement 3D compressed sensing. In this approach, the mechanized focus stage moves at a constant speed v during a single camera exposure time Δ_E, enabling the acquisition of projections by optical integration. Simultaneously, the intensity of the light sheet illumination varies according to a temporal

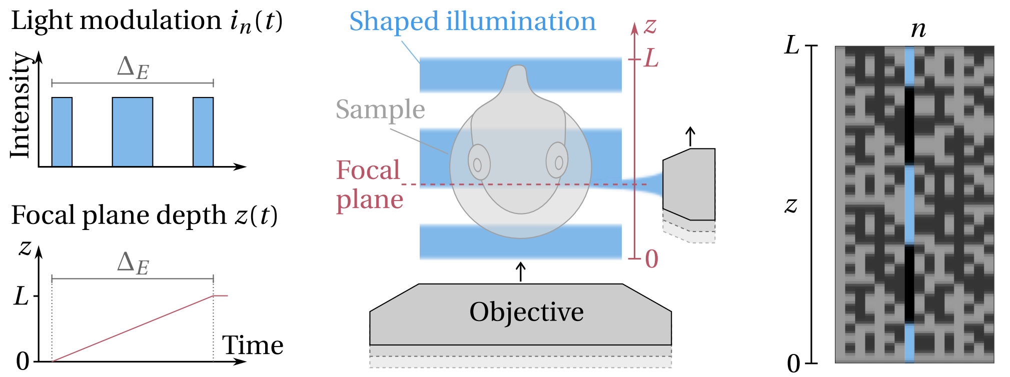

a) Optomechanical modulation

b) Resulting shaped illumination

c) Measurement matrix $\mathbf{G}'$

Figure 4.1: Overview of OMMT imaging. (a) Optomechanical modulation combines focus sweeping synchronized with temporal light modulation. (b) The modulation creates a shaped illumination pattern in the depth axis. (c) This is repeated to acquire N projections with different modulation functions, which are stacked into the measurement matrix $\mathbf{G}'$.

modulation function, which creates spatial light patterns along the depth, as illustrated in Fig. 4.1 (a-b).

4.2.1 Forward model

Let $f(x,y,z)$ be a three-dimensional object, with x, y, z the horizontal, vertical, and focus axes respectively. We write the Point Spread Function (PSF) of the system as $h(x,y,z)$, and the position in depth of the focal plane over time t as $z(t) = vt$. We acquire N projections $p_n(x,y)$ ($n = 1,\dots,N$) by modulating the light intensity over time using a function $i_n(t)$ defined over the interval $[0,\Delta_E]$, resulting in the following imaging model:

$$p_n(x,y) = \int_0^{\Delta_E} \left[f * h\right](x,y,vt)\ i_n(t)\ \mathrm{d}t\ +\ \epsilon_n(x,y), \tag{4.1}$$

where $*$ is a 3D convolution operator and $\epsilon_n(x,y)$ is an additive measurement noise.

The optomechanical modulation is particularly flexible, as the illumination functions $i_n(t)$ can be arbitrarily chosen, with the only constraint that all their values are positive. We choose to use an incomplete Hadamard basis for our experiments, as it can easily be generated by switching the illumination on and off over time, and has been used successfully in computational imaging and compressed sensing applications (Calisesi, Ghezzi, Ancora, D'Andrea, Valentini, Farina, and Bassi, 2022; Duarte, Davenport, Takhar, Laska, T. Sun, Kelly, and Baraniuk, 2008; Soldevila, Clemente, Tajahuerce, Uribe-Patarroyo, Andrés, and Lancis, 2016; Salvador-Balaguer, Latorre-Carmona, Chabert, Pla, Lancis, and Tajahuerce, 2018; Calisesi, Castriotta, Candeo, Pistocchi, D'Andrea, Valentini, Farina, and Bassi, 2019; Zunino, Garzella, Trianni, Saggau, Bianchini, Diaspro, and Duocastella, 2021; Crombez, Leclerc, Ray, and Ducros, 2022).

We start by constructing a Hadamard matrix $\mathbf{H}_M \in \mathbb{R}^{M\times M}$ of order M using Sylvester's construction, setting its negative values to 0 to conform to the positivity constraint. We then build an incomplete Hadamard matrix $\mathbf{H}_N \in \mathbb{R}^{N\times M}$ of size $N \times M$ by selecting the first row (full of ones) of $\mathbf{H}_M$ and stacking it with $N-1$ other rows of $\mathbf{H}_M$ that are randomly sampled without replacement with uniform probability. By using $N \leq M$, we can reduce the number of measurements while keeping the high frequency sampling offered by a higher order Hadamard matrix. Finally, we define our illumination modulation functions as:

$$i_n(t) = H_N[n, j] \qquad\qquad \text{s.t.}\;\; j-1 \leq \frac{Mt}{\Delta_E} < j, \tag{4.2}$$

with $n \in \{1, \dots, N\}$, and $j \in \{1, \dots, M\}$.

By applying a variable substitution $z = vt$, we can express Eq. (4.1) as an integral over depth rather than time:

$$p_n(x, y) = \int_0^L \left[f * h\right](x, y, z)\; g_n(z)\; \mathrm{d}z \;+\; \epsilon_n(x, y), \tag{4.3}$$

where $L = v\Delta_E$ is the scanning depth, and $g_n(z) = \frac{1}{v} i_n(\frac{z}{v})$, defined over $[0, L]$, is the spatial illumination pattern created by the temporal modulation i_n combined with focus sweeping, illustrated in Fig. 4.1 (b).

Reconstructing the sample f from the imaging model in Eq. (4.3) is a three-dimensional problem whose computational complexity increases with the size of the imaged volume. It makes this model unusable to reconstruct volumetric data in reasonable time when working with the typical image sizes acquired by modern microscopes.

To address this issue, Woringer et al. (Woringer, Darzacq, Zimmer, and Mir, 2017) approach the problem as a stack of 2D images, which they split into smaller regions that they reconstruct independently in a parallel computing cluster. For these 2D reconstructions, they first perform empirical measurements of the 3D PSF that they crop to a single xz plane. Because this approach solves 2D reconstructions independently, it cannot exploit the 3D nature of the sample to improve its reconstruction consistency, leading to the adjacent sections showing visual incoherences because they are computed separately. Moreover, the computational cost is still high and the method requires access to a computational cluster to run in reasonable time, as reconstructing a 512 × 512 × 100 volume on a single machine takes longer than a full day to complete.

In order to reduce the computational complexity of this problem, we simplify Eq. (4.3) by ignoring the lateral extent of the PSF and modelling it with a one-dimensional function $h'(z) = h(0, 0, z)$. This is motivated by the fact that light sheet devices have a thin truncated PSF, which vanishes much faster in the lateral directions than in the axial one (Olarte, Andilla, Gualda, and Loza-Alvarez, 2018). Using this simplified PSF, we write an approximated projection p'_n:

$$p'_n(x,y) = \int_0^L \left(\int_{-\infty}^{\infty} f(x,y,u)\ h'(u-z)\ \mathrm{d}u \right) g_n(z)\ \mathrm{d}z\ +\ \epsilon_n(x,y) \tag{4.4}$$

$$= \int_{-\infty}^{\infty} f(x,y,u) \left(\int_0^L h'(u-z)\ g_n(z)\ \mathrm{d}z \right) \mathrm{d}u\ +\ \epsilon_n(x,y) \tag{4.5}$$

$$= \int_{-\infty}^{\infty} f(x,y,z)\ g'_n(z)\ \mathrm{d}z\ +\ \epsilon_n(x,y), \tag{4.6}$$

where $g'_n(z) = g_n(z) \circledast h'(-z)$ is the modulation function blurred by the simplified PSF by the 1D convolution $\circledast$. At each location (x, y) in a projection, the imaging model reduces to a one-dimensional expression that is much simpler and faster to compute than Eq. (4.3).

When discretized, we can write Eq. (4.6) as the following linear expression:

$$\mathbf{P}[j,k,:] = \mathbf{G}'\mathbf{F}[j,k,:]\ +\ \mathbf{E}[j,k,:], \tag{4.7}$$

where $\mathbf{F} \in \mathbb{R}^{W\times H\times D}$ is a discrete sampling of f with width W, height H, and depth D, the latter corresponding to the focus axis z. $\mathbf{G}' \in \mathbb{R}^{N\times D}$ is the matrix operator corresponding to the multiplication with the modulation function g'_n and the axial integration, and constitutes the compressed sensing measurement matrix. The values in $\mathbf{P} \in \mathbb{R}^{W\times H\times N}$ correspond to our discrete measurements, and $\mathbf{E} \in \mathbb{R}^{W\times H\times N}$ is the measurement noise. We use : as a shorthand notation to describe all elements in the corresponding dimension of the array (Golub and Van Loan, 2013). For each pair of coordinates $j = 1,\ldots,W$ and $k = 1,\ldots,H$, the model is a simple matrix-vector multiplication, which can be computed very efficiently. It can therefore realistically be applied on big experimental data to compute volumetric reconstructions in reasonable time.

To derive the measurement matrix, we estimate $h'(z)$ based on a simulated PSF. Since the 1D model is an approximation, we do not require accurate experimental measurements of the 3D PSF which can be tedious to obtain. We compute the rows of $\mathbf{G}'$ by discretizing the blurred light patterns g'_n according to the axial resolution of $\mathbf{F}$, as illustrated in Fig. 4.1 (c).

4.2.2 Inverse problem

Reconstructing $\mathbf{F}$ from the measurements requires solving the linear inverse problem corresponding to Eq. (4.7). Since the reconstructed axial resolution is higher than the number of acquired projections $D > N$, this inverse problem is ill-posed and requires additional constraints to be solved uniquely. A common approach in compressed sensing is to formulate the reconstruction as a minimization problem (Calisesi, Ghezzi, Ancora, D'Andrea, Valentini, Farina, and Bassi, 2022; Kamilov, I. N. Papadopoulos, Shoreh, Goy, C. Vonesch, Unser, and

Psaltis, 2016; Woringer, Darzacq, Zimmer, and Mir, 2017):

$$\hat{\mathbf{F}} = \underset{\mathbf{F}\in\mathbb{R}^{W\times H\times D}}{\arg\min} \left(\sum_{j=1}^{W} \sum_{k=1}^{H} \mathscr{C}\left(\mathbf{P}[j,k,:], \mathbf{G}'\mathbf{F}[j,k,:]\right) + \lambda\, \mathscr{R}(\mathbf{F}) \right), \tag{4.8}$$

where $\mathscr{C} : \mathbb{R}^N \times \mathbb{R}^N \to \mathbb{R}_+$ is a loss functional that ensures consistency of the solution with the measurements, $\mathscr{R} : \mathbb{R}^{W\times H\times D} \to \mathbb{R}_+$ is a regularization term that imposes prior constraints and the parameter $\lambda \in \mathbb{R}_+$ allows adjusting the strength of the regularization. More specifically, we use a squared ℓ_2 loss for data consistency:

$$\mathscr{C}\left(\mathbf{P}[j,k,:], \mathbf{G}'\hat{\mathbf{F}}[j,k,:]\right) = ||\mathbf{P}[j,k,:] - \mathbf{G}'\hat{\mathbf{F}}[j,k,:]||_2^2. \tag{4.9}$$

We solve Eq. (4.8) using an iterative optimization algorithm, the alternating direction method of multipliers (ADMM (Glowinski and Marroco, 1975; Gabay and Mercier, 1976; Boyd, Parikh, Chu, Peleato, and Eckstein, 2011)), which alternates between minimizing $\mathscr{C}$ and $\mathscr{R}$. Choosing a good regularization $\mathscr{R}$ is crucial as it impacts the quality of the reconstruction but also the computational complexity of the minimization. We consider three different regularization candidates that we compare below.

ℓ_1 sparsity constraint

In the original method, Woringer et al. (Woringer, Darzacq, Zimmer, and Mir, 2017) use a spatial sparsity constraint expressed as:

$$\mathscr{R}(\mathbf{F}) = ||\mathbf{F}||_1. \tag{4.10}$$

This regularization favours reconstructions that are sparse (Elad and Bruckstein, 2002), and is common in compressed sensing applications (Calisesi, Ghezzi, Ancora, D'Andrea, Valentini, Farina, and Bassi, 2022; Candes and Wakin, 2008). It has a low computational cost, but it does not enforce any continuity between adjacent values in the volume and therefore cannot take advantage of the additional information provided by neighbouring voxels in the reconstruction.

TV_{1+2} sparsity constraint

Another common regularization is Total Variation (TV), which aims at generating smoother solutions while preserving sharp edges (Rudin, Osher, and Fatemi, 1992; Chambolle, 2004). We introduce a 1+2D formulation that penalizes TV separately along the z axis, which we call TV_{1+2}:

$$\mathscr{R}(\mathbf{F}) = \rho\, \mathrm{TV}_{1\mathrm{D}}(\mathbf{F}) \;+\; \mathrm{TV}_{2\mathrm{D}}(\mathbf{F}) \tag{4.11}$$

where $\mathrm{TV}_{\mathrm{1D}}$ is the 1D TV along the depth axis and $\mathrm{TV}_{\mathrm{2D}}$ is the isotropic 2D TV on xy sections:

$$\mathrm{TV}_{\mathrm{1D}}(\mathbf{F}) = \sum_{i=1}^{W} \sum_{j=1}^{H} \sum_{k=1}^{D} \left| \mathbf{F}[i,j,k+1] - \mathbf{F}[i,j,k] \right| \tag{4.12}$$

$$\mathrm{TV}_{\mathrm{2D}}(\mathbf{F}) = \sum_{i=1}^{W} \sum_{j=1}^{H} \sum_{k=1}^{D} \sqrt{\left| \mathbf{F}[i+1,j,k] - \mathbf{F}[i,j,k] \right|^2 + \left| \mathbf{F}[i,j+1,k] - \mathbf{F}[i,j,k] \right|^2}. \tag{4.13}$$

The factor ρ allows us to tune the regularization strength independently in the compression axis. We use that additional degree of freedom to factor in the anisotropic properties of the reconstructed volume, which stem from the compressed imaging and the shape of the PSF.

This regularization enforces visual coherence across space in the solution, thus benefitting from the additional information contained in neighbouring voxels. Although this comes at the cost of a higher computational complexity, we can massively speed it up using parallel computing. Indeed, all the minimization steps are simple operations that can be computed independently: both the data consistency term and $\mathrm{TV}_{\mathrm{1D}}$ are 1D operations, and $\mathrm{TV}_{\mathrm{2D}}$ is a simple 2D computation, as illustrated on Fig. 4.2. Moreover, the ADMM algorithm allows computing these terms simultaneously at each iteration, making the minimization fall into the 'embarrassingly parallel' category. We can therefore expect fast computations on machines supporting multithreaded execution, or with Graphical Processing Unit (GPU) acceleration.

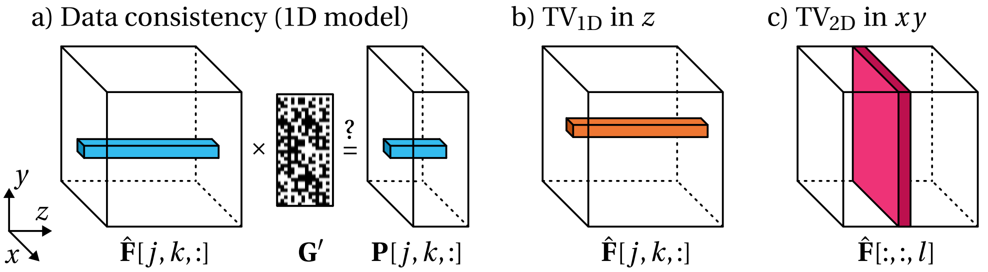

Figure 4.2: Regularization with TV_{1+2} implements 3D prior constraints while keeping efficient computations using a parallel implementation. (a) Data consistency relies on the 1D forward model. (b) $\mathrm{TV}_{\mathrm{1D}}$ regularization along the compression axis is tuned separately to account for anisotropy in the volume. (c) Isotropic $\mathrm{TV}_{\mathrm{2D}}$ on xy sections allows information sharing between the many 1D problems.

PnP BM4D constraint

ADMM does not compute the gradient of $\mathscr{R}$ but can instead rely on its proximal function to carry out the minimization. This makes it suitable for implementing PnP methods (Venkatakrishnan, Bouman, and Wohlberg, 2013) by using a denoiser as the proximal function to an implicit regularization prior. This is not possible with the SPIRAL solver (Harmany, Marcia, and Willett, 2009) used by Woringer et al. (Woringer, Darzacq, Zimmer, and Mir, 2017), as it requires an explicit regularization function.

Although deep learning denoisers provide state-of-the-art performance for PnP methods in 2D, publicly available pretrained universal 3D denoisers based on deep learning are lacking, and the massive computational cost of building a volumetric dataset and training our own model prevents us from using one in this work. Instead, we use the BM4D denoiser (Maggioni, Katkovnik, Egiazarian, and Foi, 2013), an extension of the well-known BM3D algorithm to volumetric data, as a PnP prior for our problem. By enforcing a 3D spatial coherence in the reconstruction, this prior also benefits from information sharing between adjacent voxels. The regularization strength λ in Eq. (4.8) is replaced by a new parameter σ specific to BM4D, which is the strength of the denoising applied at each iteration of the ADMM optimization.

4.2.3 Reliable choice of the hyperparameters

The performance of the reconstruction algorithm is impacted by the selected hyperparameter values for the optimization problem (depending on the regularization, these can be λ, ρ or σ). We have found the minimization of the data consistency cost $\mathscr{C}$ to be a good proxy for the performance of the reconstruction. Although the minimum cost does not always correspond to the best reconstruction accuracy, we have measured that it consistently gives a very similar performance to the actual best on our experiments.

Choosing hyperparameters that minimize the final data consistency cost is therefore a reliable way to guarantee good performance using a reproducible method. This does not require knowing the ground truth, and can be applied regardless of the regularization constraints used. That is particularly helpful in the case of the BM4D PnP prior, as the lack of an explicit regularization function prevents from using common hyperparameter selection methods such as the L-curve (Hansen, 1992).

4.3 Experiments and results

In order to take advantage of the highly parallel nature of our algorithm, we implemented the reconstruction using the scientific imaging library SCICO (Balke, Davis, Garcia-Cardona, Majee, McCann, Pfister, and Wohlberg, 2022). This framework provides an efficient implementation of ADMM based on JAX (Bradbury, Frostig, Hawkins, M. J. Johnson, Leary, Maclaurin, Necula, Paszke, VanderPlas, Wanderman-Milne, and Q. Zhang, 2018; Frostig, M. J. Johnson, and Leary, 2018), a high-performance Python library that uses just-in-time compilation to provide parallel computing and hardware acceleration (such as GPU). We used a Python implementation of an updated version of BM4D provided by its authors (Mäkinen, Azzari, and Foi, 2020; Mäkinen, Marchesini, and Foi, 2022).

We first characterize our method on simulated data, then validate it on experimental images.

4.3.1 Characterization on simulated data

We used the simulation framework previously introduced in Chapter 2 (Marelli and Liebling, 2021) to generate synthetic volumes of size 128 × 128 × 128 with physical properties matching those of data acquired with the OpenSPIM platform (Pitrone, Schindelin, Stuyvenberg, Preibisch, Weber, Eliceiri, Huisken, and Tomancak, 2013). We simulated the imaging 3D PSF using an implementation of the Born & Wolf model (Born and Wolf, 2013; Kirshner, Aguet, Sage, and Unser, 2013). After the imaging simulation, we corrupted the data with shot noise generated using a Poisson distribution. We set the strength of the noise by rescaling the data to a chosen maximum photon count before sampling the Poisson distribution, with higher photon counts corresponding to weaker overall noise. Finally, we quantized the data to a 12 bit representation to emulate acquisition with a digital camera.

In practice, the PSF simulation model used for reconstruction in our method does not exactly match the actual imaging PSF. In order to account for this discrepancy in our experiments, we used a different simplified Gaussian beam model (Trull, Horst, Palenstijn, Vliet, Leeuwen, and Kalkman, 2017) to compute the 1D PSF when building the measurement matrix.

To select the hyperparameters, we ran loose grid searches (in logarithmic space for λ and ρ, and in linear space for σ) and selected the solution with minimal consistency cost as described in Section 4.2.3. We use $N = 16$ and $M = 32$ to build our measurement matrices. We measured that this ratio of 2 for undersampling of the Hadamard matrix yields the best performance for a fixed number of projections.

In order to evaluate how the use of different regularizations can affect the artefacts present in the reconstruction, we simulated an object with a moon-shaped cross-section and imaged it with low noise (10^4 photons). The 3D reconstructions obtained using our method with each regularization prior are shown in Fig. 4.3, alongside the central xy and xz sections.

Reconstructions performed with the ℓ_1 regularization contain visual noise. While the xy section appears good visually, the quality of the xz section is visibly worse, with the strong noise leading to a very uneven intensity in the object and fuzzy edges. Similar artefacts also appeared in the results obtained by Woringer et al. (Woringer, Darzacq, Zimmer, and Mir, 2017). It is also visible on the 3D views that the ℓ_1 reconstruction contains non-zero values in empty areas around the object.

Both our proposed TV_{1+2} and BM4D priors give a much cleaner result and better defined edges, with an even intensity in the xy section and only small variations in the xz section. They also contain only few non-zero values outside the object. This reduced number of artefacts shows the benefit of using a 3D-aware regularization prior that better enforces spatial continuity. The information sharing between neighbouring voxels leads to a visually more coherent volumetric reconstruction and reduced noise.

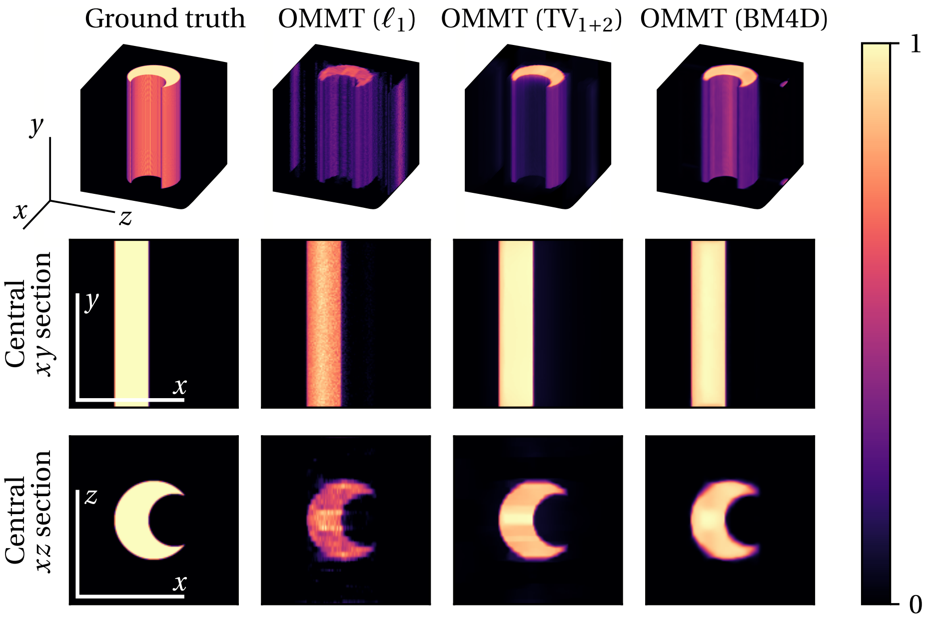

Figure 4.3: Reconstruction of a simulated sample using the different regularization priors. The ℓ_1 result contains more noise and has fuzzy edges in the xz section, while TV_{1+2} and BM4D give a much smoother result with better defined edges. 3D views use attenuated mean intensity projection. Scale bar axes are 50 µm.

Comparison with plane-by-plane acquisition for sparse object detection

When introducing compressed sensing based on optomechanical modulation, Woringer et al. (Woringer, Darzacq, Zimmer, and Mir, 2017) showed that one of its main benefits is the reduction of phototoxicity during acquisition. They demonstrated how the method outperforms plane-by-plane acquisition when matching the light dose by reducing the exposure of the SPIM acquisition, which increases the level of noise in the stack. Here we want to measure if these advantages are maintained when comparing to a plane-by-plane acquisition with fewer imaged sections, which gives a comparable level of phototoxicity without adding noise. To that end, we consider the scenario where we image a number of uniformly spaced sections equal to the number of projections $N = 16$ acquired with OMMT. By using the same average light intensity per frame and exposure time for the plane-by-plane as for OMMT, the sample receives the same light dose and the noise level in images is comparable. We then resample the SPIM stack along the z axis with cubic interpolation to achieve the same final spatial sampling as the OMMT reconstructions.

We simulate a sample composed of thin vertical cylinders uniformly spaced along the diagonal of the volume. We choose spacing along the depth axis to not be a multiple of the distance between acquired sections with the plane-by-plane method, and position the central cylinder to be exactly aligned with one of the imaged planes. That way, the cylinders are gradually less aligned with the plane-by-plane sampling as they get further away from the centre.

Figure 4.4 shows the SPIM axial sampling and the obtained volumes with both methods. While the central cylinders are correctly visible in the plane-by-plane image, misaligned objects highlighted in the figure are barely visible when they fall exactly between two sampled sections. All cylinders are clearly visible in the OMMT reconstructions regardless of their position, thanks to the fact that all regions of the sample are illuminated at least once during acquisition due to the choice of illumination functions. This makes OMMT more reliable than plane-by-plane acquisition to detect spatially sparse objects, at equal phototoxicity and noise levels.

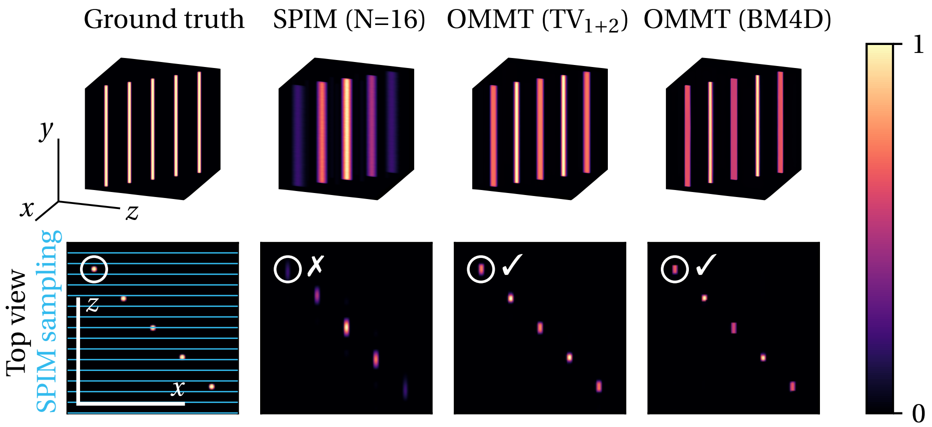

Figure 4.4: Comparison of undersampled plane-by-plane SPIM and OMMT with equal phototoxicity. SPIM fails to detect objects that are misaligned with its axial sampling, while OMMT reconstructions contain all the elements. All views use mean intensity projection, and axes are 50 µm scale bars.

Quantitative performance comparison and impact of noise

In order to quantify the performance of our algorithm, we simulated a sample containing multiple thin objects crossing the volume at various angles and with different cross-sections and intensities. To measure the impact of noise on the reconstruction, we generated images for a virtually noise-free scenario with a maximum photon count per pixel of 10^4, and a noisy scenario with a maximum photon count of 400.

Our comparison criterion is the Peak Signal-to-Noise Ratio (PSNR), defined as:

$$\mathrm{PSNR} = 10\log_{10}\left(\max(\mathbf{F})^2/\mathrm{MSE}\right), \tag{4.14}$$

where MSE is the Mean Squared Error between our reconstructions and the original sample **F**. We ran the reconstructions on a multithreaded machine (AMD EPYC 7742 2.25 GHz, limited to 8 cores), and measured the total running time of the ADMM minimization to compare the computational complexity of the different regularization priors.

We repeated each experiment 5 times with different random seeds for the noise and the choice of the illumination functions, and averaged the obtained metrics. The measured deviation between repetitions was negligible, guaranteeing good confidence in the presented numbers.

Table 4.1 summarizes the results obtained with the different regularization priors for both clean and noisy scenarios. As a reference, we also simulated a fully sampled SPIM stack with 128 sections and an undersampled SPIM stack with 16 sections corresponding to equivalent phototoxicity, as described in Section 4.3.1.

	Max photons: 10^4		*Max photons:* 400	
	PSNR	Time	PSNR	Time
SPIM (N=128)	36.08 dB	/	34.63 dB	/
SPIM (N=16)	28.64 dB	/	28.49 dB	/
OMMT (ℓ_1)	32.77 dB	**4 s**	32.22 dB	**4 s**
OMMT (TV_{1+2})	34.22 dB	26 s	34.05 dB	21 s
OMMT (BM4D)	**35.26 dB**	72 min	**35.21 dB**	68 min

Table 4.1: Performance characterization on simulated sample for different regularization priors and with varying noise strengths, with full and undersampled SPIM stacks as reference. BM4D performs best, but is significantly slower than using other priors. OMMT reconstruction consistently outperforms undersampled SPIM, and beats the full SPIM stack in the noisy scenario with BM4D prior.

The results clearly show that the PnP BM4D prior consistently outperforms all other regularizations in both scenarios. However, it is very slow to compute, with the ℓ_1 prior being 1000× faster to reach convergence. Our proposed TV_{1+2} regularization offers a trade-off between speed and accuracy, being only 5× slower than ℓ_1, while improving over its performance by 1.5 dB on average.

OMMT outperforms the undersampled SPIM regardless of the photon count. Adding noise does not significantly impact its reconstruction performance, and with the BM4D prior OMMT beats the full SPIM stack in the noisy scenario. This confirms the observations of Woringer et al. (Woringer, Darzacq, Zimmer, and Mir, 2017) that compressed sensing methods match the performance of classically sampling noisy images as the noise level increases.

We show the reconstructions obtained with TV_{1+2} and BM4D priors for the noisy scenario in Fig. 4.5 alongside the undersampled SPIM volume. OMMT reconstructions are visibly smoother in regions where the SPIM stack appears to have gaps in the objects. Using the BM4D PnP prior yields a smoother reconstruction but the visual improvement over TV_{1+2} is not significant.

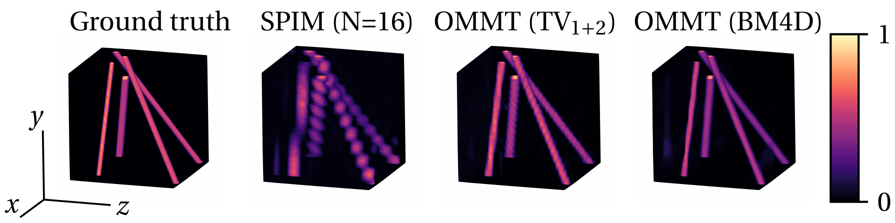

Figure 4.5: Volumetric reconstructions on simulated objects in a noisy scenario compared to an undersampled SPIM with equal phototoxicity. OMMT reconstructions are smoother, with a minor improvement when using BM4D over TV_{1+2}. Attenuated mean intensity projection is shown, and axes are 50 µm scale bars.

4.3.2 Analysis of the computational cost

To efficiently explore many scenarios and methods efficiently, we performed our simulations on undersized data. In order to apply our method in practice, it must be able to compute a solution on real-scale images with a size of at least 512 × 512 in reasonable time. To assess the scaleability of the different regularization priors, we measured the time to run a single iteration of the ADMM optimization loop for increasing data sizes. This gives us a computational efficiency measurement that is not influenced by the number of iterations to converge. The obtained single iteration run time is sufficient to estimate the relative total run time of the algorithm with various regularizations, as we observed that the different priors lead to similar total number of iterations to convergence. Any arbitrary data can be used for this speed assessment as we run a fixed number of iterations that do not need to converge. We used random uniform input data and averaged the time over multiple runs each with multiple iterations, after running 2 warm-up iterations.

We measured the performance of our implementation on a multithreaded machine (AMD EPYC 7742 2.25 GHz, limited to 8 cores). In the case of ℓ_1 and TV_{1+2}, we also measured the performance on a machine with GPU (Nvidia GeForce RTX 3090). We did not measure it for BM4D as its implementation does not support GPU acceleration.

Figure 4.6 shows our measurements for output volumes with equal depth, width, and height. We used a number of projections N equal to half the depth of the reconstructed volume. As observed before, the BM4D PnP prior is much slower than the other regularizers. Moreover, it becomes unusable for big volumes as its computation time rapidly increases with the size of the data. Even if ℓ_1 is consistently the fastest, TV_{1+2} is only slightly slower for bigger images. Their computation time grows less quickly with data size, making them realistically useable to reconstruct big volumes. Moreover, using GPU acceleration brings a massive speed-up of over 30× to the algorithm, which shows that our method is suitable for reconstructing real-scale experimental data very efficiently without requiring access to a computational cluster.

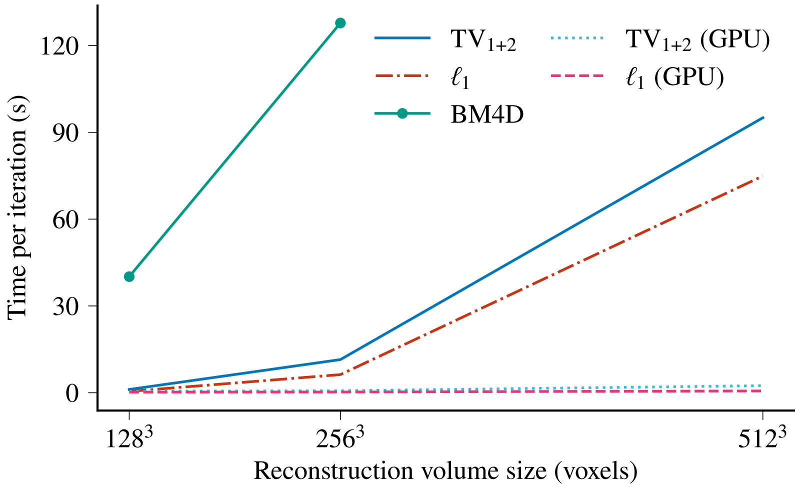

Figure 4.6: Average time per iteration of our ADMM optimization loop for different regularization priors. BM4D is not realistically useable to reconstruct big images, but both TV_{1+2} and ℓ_1 are fast enough to be applied to real size experimental data. This is particularly true with the 30× speed-up when running with GPU acceleration.

4.3.3 Validation on experimental data

To confirm the real-case applicability of our method, we mounted two fluorescent textile fibres with a diameter of 25 µm in 2 % low melting agarose solution inside a fluorinated ethylene propylene tube. We imaged the sample using the implementation of the OpenSPIM platform described in Section 1.3. We used the custom modulation controller based on Arduino that we introduced in Chapter 1 to control the intensity of the laser directly at emission, which is a cost-efficient alternative to the more expensive acousto-optic tunable filter used in the original method (Woringer, Darzacq, Zimmer, and Mir, 2017).

We acquired $N = 16$ projections with $M = 32$, sweeping the focal plane over a total depth of 600 µm and reconstructed a volume of size $640 \times 640 \times 128$ using the TV_{1+2} prior. As a reference, we acquired a full SPIM stack with 128 sections. We also imaged an undersampled SPIM containing 16 sections, which yields equivalent phototoxicity to OMMT.

Figure 4.7 shows the obtained volumes. For an equal light exposure, OMMT does not feature the jagged artefacts visible in the undersampled SPIM and does not contain gaps in the reconstructed fibres. This confirms our results obtained on simulations. The solution matches the geometry of the reference volume even with a very low number of acquired images.

4.4 Discussion

Our reconstruction method is able to use regularization priors that enforce a 3D coherence in the solution while being efficient to compute given its parallel implementation. Since our proposed TV_{1+2} formulation is specifically tailored to the geometry of the problem, we

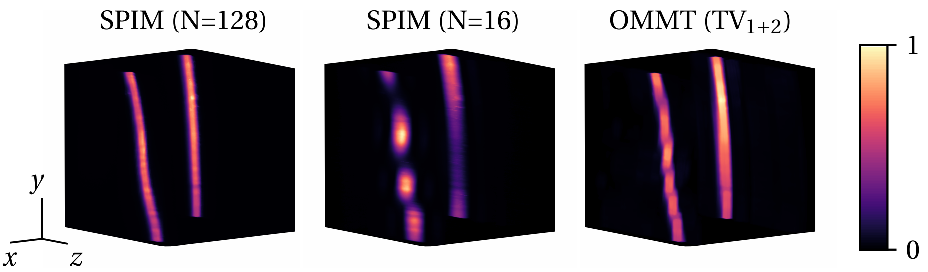

Figure 4.7: OMMT reconstruction of fluorescent textile fibres and comparison to undersampled SPIM with equal light dose. The reconstructed image with TV$_{1+2}$ prior is smoother, and looks similar to the reference SPIM volume. Views use mean intensity projection, and axes are 100 µm scale bars.

have been able to improve over the existing method with ℓ_1 prior and reduce reconstruction artefacts while having a minor additional computational cost. This better performance likely stems from the fact that the method takes into account the volumetric nature of the problem where neighbouring voxels are considered along all dimension to improve the overall accuracy. Using the same acquisition procedure as in (Woringer, Darzacq, Zimmer, and Mir, 2017)), our updated method retains the advantages of the original technique over plane-by-plane imaging (lower phototoxicity, faster acquisition) while the efficient implementation of our reconstruction method eliminates the need for a computational cluster to obtain solutions in reasonable time.

Using a universal 3D denoiser as PnP prior for reconstruction outperforms the TV$_{1+2}$ regularization but is not realistically applicable in practice due to its slower computation (up to a factor 1000) on bigger data. However, this result seems promising when considering the use of a 3D deep learning-based denoiser instead of BM4D, as neural networks are usually running efficiently with GPU acceleration. This would alleviate the speed problem by shifting part of the computational load from reconstruction time to the training of the model. Moreover, recent works in image reconstruction are trying to reduce the memory and computational requirements of PnP methods by using subsets of the data at each iteration (Y. Sun, Wohlberg, and Kamilov, 2019; Y. Sun, Z. Wu, Xu, Wohlberg, and Kamilov, 2021). These methods could further improve computation speed if they can be adapted to our 3D problem.

4.5 Conclusion

We have implemented an ADMM algorithm for OMMT reconstruction that uses parallel computing and shown that it can efficiently solve the inverse problem on a single machine in reasonable time. Our efficient 1+2D formulation of the TV regularization prior allows introducing volumetric smoothness constraints while taking into account the specific geometry of the problem while keeping the computational overhead in check.

We used simulated data to characterize the performance of our method, showing that using fully 3D-aware priors reduces visual artefacts over ℓ_1 and that a PnP method with BM4D denoising gives the best accuracy. We further established the advantages of the OMMT compressed sensing acquisition scheme over plane-by-plane imaging by demonstrating its superior ability to detect sparse objects at comparable phototoxicity.

We have measured the scaleability of our implementation on big volumes, proving its applicability to real-scale images. We validated these results by successfully applying our method to an experimental dataset.

Although we showed that using a PnP prior for reconstruction is too computationally demanding with the BM4D denoiser, we expect that efficient implementations of 3D deep learning denoising algorithms could further improve the reconstruction accuracy with a lower computational cost.

5 Conclusions

In this thesis, I have developed a versatile hardware and software framework that allows simulating and implementing computational imaging methods based on spatiotemporal illumination modulation with common optomechanical components. I have shown through imaging applications, including applications in biology, that my framework is adaptable to different optical setups and imaging modalities and can be exploited to obtain enhanced spatial and temporal resolutions with a good cost efficiency. My accelerated simulation framework enables the exploration of new light shaping methods with reliable performance quantification. I have acquired experimental datasets using my platform to validate my results.

In Chapter 1, I designed a flexible device that allows synchronizing a digital camera with multiple voltage-controlled equipments, such as light sources. I developed hardware drivers to integrate that controller into the popular open microscopy software μManager. I implemented a user-friendly interface and a high-level scripting library that allows controlling the optomechanical components of the imaging platform, without requiring low-level technical skills. The ease of use, adaptability, and good cost efficiency of my modulation framework allows microscopists to integrate it into existing imaging setups, providing spatiotemporal light modulation capacities with minimal effort.

I exploited the spatial light shaping capacities of this framework in Chapter 2. I introduced FSS-OPT, an imaging technique that uses a SPIM illumination to perform focal plane scanning OPT. I showed that FSS-OPT improves the spatial resolution of OPT without using post-processing. I implemented an accelerated real-scale simulation framework to measure how an alternative high-resolution OPT technique based on deconvolution is impacted by errors in the optical model, and used it to quantify the robustness to noise of both approaches. I showed that FSS-OPT outperforms the alternative method in a realistic scenario, which I confirmed on experimental data. In (Y. Liu, Dong, Pham, Marelli, and Unser, 2022), I used FSS-OPT to acquire a dataset designed to measure the artefacts caused by the incorrect calibration of the rotation axis in OPT imaging.

In Chapter 3, I used the subframe temporal modulation enabled by my control device to introduce PAAQ for *in vivo* cardiac fluorescence microscopy, an approach that captures pairs of images from alternating modalities to virtually increase the frame rate and register sequences imaged from different fluorescent channels. I implemented a simulation model of the beating heart that includes heart rate variability to characterize the performance of this technique. I developed a phase estimation algorithm based on a non-uniform sampling assumption, and measured on the simulations that it improves the temporal resolution of the reconstruction over state-of-the-art competing methods. I acquired an experimental dataset on a zebrafish embryo using PAAQ to show that the introduced non-uniform sampling assumption reduces the amount of artefacts caused by temporal misalignment of the different fluorescent channels.

Finally, I illustrated the combined spatial and temporal illumination modulation capacities of my framework in Chapter 4 by implementing a state-of-the-art 3D compressed sensing method for SPIM imaging. I proposed an efficient 1+2D regularization TV prior that allows solving the reconstruction problem with volumetric smoothness constraints in reasonable time thanks to parallel computing. Using my simulation framework, I measured the performance and scaleability of alternative regularizations, including an implementation of a PnP method using state-of-the-art 3D denoising. I showed that the proposed 1+2D regularization offers an advantageous trade-off between computational complexity and reconstruction accuracy, and validated its real-case applicability on experimental data.

5.1 Future directions

In this thesis, I developed a versatile framework that enables the exploration of a wide range of imaging strategies based on spatiotemporal illumination modulation. This opens the way for researchers to explore innovative acquisition designs, and I foresee that the optomechanical illumination shaping could bring benefits to illumination modalities that were not investigated in this thesis.

For example, the introduced modulation framework enables the implementation of 3D compressed sensing for hyperspectral imaging. As shown in (Crombez, Leclerc, Ray, and Ducros, 2022), using a patterned SPIM illumination allows retrieving positional information from spectrometer measurements. This technique could be extended using mechanical scanning coupled with temporal modulation similarly to OMMT in order to obtain fast 3D hyperspectral imaging, of which the advantages have already been illustrated using a different optical setup based on oblique plane microscopy (K. Guo, Kalyviotis, Pantazis, and Rowlands, 2023).

Previous works show that multi-spectral imaging of periodic processes can be achieved through the synchronization of sequences acquired in different wavelength ranges by switching between multiple illumination filters (Jaques, Bapst-Wicht, Schorderet, and Liebling, 2019). The developed modulation controller allows replacing these filters by light sources with different emission spectra, switching between wavelengths through intensity modulation rather than mechanical filter selection. This would open the possibility to compress the spectral information into fewer frames by applying continuous intensity modulation instead of the on/off switching achieved by using filters, similarly to what has been achieved with LED arrays and digital micromirror devices in (Ma, Yuan, Fu, and Arce, 2021). This would require a more complex unmixing post-processing step, but could further improve the acquisition efficiency, both in speed and phototoxicity.

In Chapter 3, I explored how acquiring pairs of images containing complementary information can be exploited in multichannel cardiac fluorescence microscopy to achieve virtually higher frame rates. This method opens the way to multi-colour volumetric imaging with enhanced frame rate. Indeed, PAAQ can be used to image sequences at different depths, using a reference modality with a large depth of field. This would allow synchronizing adjacent sections of the volume using a temporal alignment method on the sorted reference sequences (Liebling, Vermot, A. Forouhar, Gharib, M. Dickinson, and S. Fraser, 2006). Thanks to the large depth of field, images acquired at a specific depths would share information with the ones taken from different sections that are not too distant, ensuring the good performance of the temporal registration algorithm. A virtually extended depth of field obtained via mechanical scanning (similarly to FSS-OPT) could also be used for imaging the references, but this would require a step of scanning aberration correction prior to the registration (Mariani, Ernst, Mercader, and Liebling, 2020).

In Chapter 4, I highlighted that using a PnP prior for the volumetric reconstruction of OMMT images outperforms all other regularizations, but that available 3D denoisers are too slow for real-case application. Given the good performance and efficient computations offered by recent deep learning 2D denoising models (K. Zhang, Zuo, Y. Chen, Meng, and L. Zhang, 2017; K. Zhang, Zuo, and L. Zhang, 2018; Z. Wang, Cun, J. Bao, W. Zhou, J. Liu, and H. Li, 2022; Zamir, Arora, S. Khan, Hayat, F. S. Khan, and M.-H. Yang, 2022), I believe that the adaptation of these architectures to perform fast 3D denoising is a promising research direction. Indeed, it might enable the use of PnP methods for the reconstruction of large OMMT volumes in reasonable time, but it would require considerable computational resources to build a suitable 3D dataset and then train the model.

5.2 Conclusion

This thesis facilitates the implementation of spatiotemporal illumination modulation techniques based on common optomechanical components in modern microscopes. By providing a versatile and user-friendly modulation controller, it allows researchers to focus on the design of innovative imaging strategies without having to go through the time-consuming steps of creating custom hardware solutions for every new method. My results illustrate the adaptability of this illumination framework in multiple imaging applications, and provide readily usable efficient algorithms to enhance the spatial and temporal resolutions of imaging platforms using spatiotemporal light shaping. I provide a real-scale 3D simulation framework that allows fast and reliable prototyping and characterization of imaging methods, as a means to encourage the development of new acquisition procedures that exploit the computational imaging framework developed in this thesis.

The developed simulation framework and reconstruction algorithms are available on the following repository: https://github.com/idiap/cbi_toolbox, and the schematics and drivers of the modulation controller at: https://github.com/idiap/CBI-MMTools.

Bibliography

Abdolhoseini, M., M. G. Kluge, F. R. Walker, and S. J. Johnson (2019). “Neuron Image Synthesizer Via Gaussian Mixture Model and Perlin Noise”. In: *IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pp. 530–533. DOI: 10.1109/ISBI.2019.8759471.

Bakkers, J. (2011). “Zebrafish as a model to study cardiac development and human cardiac disease”. In: *Cardiovascular research* 91.2, pp. 279–288. DOI: 10.1093/cvr/cvr098.

Balke, T., F. Davis, C. Garcia-Cardona, S. Majee, M. McCann, L. Pfister, and B. Wohlberg (2022). “Scientific Computational Imaging Code (SCICO)”. In: *Journal of Open Source Software* 7.78, p. 4722. DOI: 10.21105/joss.04722.

Barufaldi, B., C. K. Abbey, M. A. Lago, T. L. Vent, R. J. Acciavatti, P. R. Bakic, and A. D. A. Maidment (2021). “Computational Breast Anatomy Simulation Using Multi-Scale Perlin Noise”. In: *IEEE Transactions on Medical Imaging* 40.12, pp. 3436–3445. DOI: 10.1109/TMI.2021.3087958.

Bassi, A., B. Schmid, and J. Huisken (2015). “Optical tomography complements light sheet microscopy for in toto imaging of zebrafish development”. In: *Development* 142.5, pp. 1016–1020. DOI: 10.1242/dev.116970.

Betzig, E., G. H. Patterson, R. Sougrat, O. W. Lindwasser, S. Olenych, J. S. Bonifacino, M. W. Davidson, J. Lippincott-Schwartz, and H. F. Hess (2006). “Imaging Intracellular Fluorescent Proteins at Nanometer Resolution”. In: *Science* 313.5793, pp. 1642–1645. DOI: 10.1126/science.1127344.

Born, M. and E. Wolf (2013). *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light.* Chap. 8.5, pp. 436–445.

Boyd, S., N. Parikh, E. Chu, B. Peleato, and J. Eckstein (2011). “Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers”. In: *Foundations and Trends in Machine Learning* 3.1, pp. 1–122. DOI: 10.1561/2200000016.

Bradbury, J., R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang (2018). *JAX: composable transformations of Python+NumPy programs.* Version 0.4.2. http://github.com/google/jax.

Brau, A. C., C. T. Wheeler, L. W. Hedlund, and G. A. Johnson (2002). “Fiber-optic stethoscope: a cardiac monitoring and gating system for magnetic resonance microscopy”. In: *Magnetic Resonance in Medicine* 47.2, pp. 314–321. DOI: 10.1002/mrm.10049.

Bub, G., M. Tecza, M. Helmes, P. Lee, and P. Kohl (2010). “Temporal pixel multiplexing for simultaneous high-speed, high-resolution imaging”. In: *Nature Methods* 7.3, pp. 209–211. DOI: 10.1038/nmeth.1429.

Bussmann, J., J. Bakkers, and S. Schulte-Merker (2007). “Early Endocardial Morphogenesis Requires Scl/Tal1”. In: *PLOS Genetics* 3.8, pp. 1–13. DOI: 10.1371/journal.pgen.0030140.

Cai, H., Y. L. Wang, R. T. Wainner, N. V. Iftimia, C. V. Gabel, and S. H. Chung (2019). “Wedge prism approach for simultaneous multichannel microscopy”. In: *Scientific Reports* 9.1, p. 17795. DOI: 10.1038/s41598-019-53581-9.

Calisesi, G., M. Castriotta, A. Candeo, A. Pistocchi, C. D’Andrea, G. Valentini, A. Farina, and A. Bassi (2019). “Spatially modulated illumination allows for light sheet fluorescence microscopy with an incoherent source and compressive sensing”. In: *Biomed. Opt. Express* 10.11, pp. 5776–5788. DOI: 10.1364/BOE.10.005776.

Calisesi, G., A. Ghezzi, D. Ancora, C. D’Andrea, G. Valentini, A. Farina, and A. Bassi (2022). “Compressed sensing in fluorescence microscopy”. In: *Progress in Biophysics and Molecular Biology* 168, pp. 66–80. DOI: 10.1016/j.pbiomolbio.2021.06.004.

Candes, E. J. and M. B. Wakin (2008). “An Introduction To Compressive Sampling”. In: *IEEE Signal Processing Magazine* 25.2, pp. 21–30. DOI: 10.1109/MSP.2007.914731.

Chakrova, N., A. S. Canton, C. Danelon, S. Stallinga, and B. Rieger (2016). “Adaptive illumination reduces photobleaching in structured illumination microscopy”. In: *Biomed. Opt. Express* 7.10, pp. 4263–4274. DOI: 10.1364/BOE.7.004263.

Chambolle, A. (2004). “An algorithm for total variation minimization and applications”. In: *Journal of Mathematical Imaging and Vision* 20, pp. 89–97. DOI: 10.1023/B:JMIV.0000011325.36760.1e.

Chan, K. G. and M. Liebling (2017). “Direct inversion algorithm for focal plane scanning optical projection tomography”. In: *Biomed. Opt. Express* 8.11, pp. 5349–5358. DOI: 10.1364/BOE.8.005349.

Chan, S. H., X. Wang, and O. A. Elgendy (2017). “Plug-and-Play ADMM for Image Restoration: Fixed-Point Convergence and Applications”. In: *IEEE Transactions on Computational Imaging* 3.1, pp. 84–98. DOI: 10.1109/TCI.2016.2629286.

Chen, B., B.-J. Chang, P. Roudot, F. Zhou, E. Sapoznik, M. Marlar-Pavey, J. B. Hayes, P. T. Brown, C.-W. Zeng, T. Lambert, J. R. Friedman, C.-L. Zhang, D. T. Burnette, D. P. Shepherd, K. M. Dean, and R. P. Fiolka (2022). “Resolution doubling in light-sheet microscopy via oblique plane structured illumination”. In: *Nature Methods* 19, pp. 1419–1426. DOI: 10.1038/s41592-022-01635-8.

Chico, T. J., P. W. Ingham, and D. C. Crossman (2008). “Modeling Cardiovascular Disease in the Zebrafish”. In: *Trends in Cardiovascular Medicine* 18.4, pp. 150–155. DOI: 10.1016/j.tcm.2008.04.002.

Crombez, S., P. Leclerc, C. Ray, and N. Ducros (2022). “Computational hyperspectral light-sheet microscopy”. In: *Opt. Express* 30.4, pp. 4856–4866. DOI: 10.1364/OE.442043.

Dabov, K., A. Foi, V. Katkovnik, and K. Egiazarian (2008). “Image restoration by sparse 3D transform-domain collaborative filtering”. In: *Image Processing: Algorithms and Systems VI*. Vol. 6812, p. 681207. DOI: 10.1117/12.766355.

Dean, K. M., P. Roudot, E. S. Welf, G. Danuser, and R. Fiolka (2015). "Deconvolution-free Subcellular Imaging with Axially Swept Light Sheet Microscopy". In: *Biophysical Journal* 108.12, pp. 2807–2815. DOI: https://doi.org/10.1016/j.bpj.2015.05.013.

Donoho, D. (2006). "Compressed sensing". In: *IEEE Transactions on Information Theory* 52.4, pp. 1289–1306. DOI: 10.1109/TIT.2006.871582.

Duarte, M. F., M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk (2008). "Single-pixel imaging via compressive sampling". In: *IEEE Signal Processing Magazine* 25.2, pp. 83–91. DOI: 10.1109/MSP.2007.914730.

Dustler, M., P. Bakic, H. Petersson, P. Timberg, A. Tingberg, and S. Zackrisson (2015). "Application of the fractal Perlin noise algorithm for the generation of simulated breast tissue". In: *Medical Imaging 2015: Physics of Medical Imaging*. Vol. 9412, 94123E. DOI: 10.1117/12.2081856.

Dustler, M., H. Förnvik, and K. Lång (2018). "Binary implementation of fractal Perlin noise to simulate fibroglandular breast tissue". In: *Medical Imaging 2018: Physics of Medical Imaging*. Vol. 10573, p. 1057357. DOI: 10.1117/12.2293234.

Edelstein, A., N. Amodaj, K. Hoover, R. Vale, and N. Stuurman (2010). "Computer Control of Microscopes Using μManager". In: *Current Protocols in Molecular Biology* 92.1, pp. 14.20.1–14.20.17. DOI: 10.1002/0471142727.mb1420s92.

Elad, M. and A. Bruckstein (2002). "A generalized uncertainty principle and sparse representation in pairs of bases". In: *IEEE Transactions on Information Theory* 48.9, pp. 2558–2567. DOI: 10.1109/TIT.2002.801410.

Engelbrecht, C. J. and E. H. Stelzer (2006). "Resolution enhancement in a light-sheet-based microscope (SPIM)". In: *Opt. Lett.* 31.10, pp. 1477–1479. DOI: 10.1364/OL.31.001477.

Fasih, A. and T. Hartley (2010). "GPU-accelerated synthetic aperture radar backprojection in CUDA". In: *2010 IEEE Radar Conference*, pp. 1408–1413. DOI: 10.1109/RADAR.2010.5494395.

Frantz, D., T. Karamahmutoglu, A. J. Schaser, D. Kirik, and E. Berrocal (2022). "High contrast, isotropic, and uniform 3D-imaging of centimeter-scale scattering samples using structured illumination light-sheet microscopy with axial sweeping". In: *Biomed. Opt. Express* 13.9, pp. 4907–4925. DOI: 10.1364/BOE.464039.

Frostig, R., M. J. Johnson, and C. Leary (2018). "Compiling machine learning programs via high-level tracing". In: *Systems for Machine Learning* 4.9.

Gabay, D. and B. Mercier (1976). "A dual algorithm for the solution of nonlinear variational problems via finite element approximation". In: *Computers & Mathematics with Applications* 2.1, pp. 17–40. DOI: 10.1016/0898-1221(76)90003-1.

Gabor, D. (n.d.). "A new microscopic principle". In: *Nature* 161 (), pp. 777–778. DOI: 10.1038/161777a0.

Garbellotto, C. and J. M. Taylor (2018). "Multi-purpose SLM-light-sheet microscope". In: *Biomed. Opt. Express* 9.11, pp. 5419–5436. DOI: 10.1364/BOE.9.005419.

Glowinski, R. and A. Marroco (1975). "Sur l'approximation, par éléments finis d'ordre un, et la résolution, par pénalisation-dualité d'une classe de problèmes de Dirichlet non linéaires”. In: *Revue française d'automatique, informatique, recherche opérationnelle. Analyse numérique* 9.R2, pp. 41–76.

Golub, G. H. and C. F. Van Loan (2013). *Matrix computations.*

Gorthi, S. S., D. Schaak, and E. Schonbrun (2013). “Fluorescence imaging of flowing cells using a temporally coded excitation”. In: *Opt. Express* 21.4, pp. 5164–5170. DOI: 10.1364/OE.21.005164.

Gregor, I., E. Butkevich, J. Enderlein, and S. Mojiri (2021). “Instant three-color multiplane fluorescence microscopy”. In: *Biophysical Reports* 1.1, p. 100001. DOI: 10.1016/j.bpr.2021.100001.

Guo, K., K. Kalyviotis, P. Pantazis, and C. J. Rowlands (2023). “Hyperspectral Oblique Plane Microscopy”. In: *bioRxiv*. DOI: 10.1101/2023.03.15.532804.

Guo, Q., H. Chen, Y. Wang, Y. Guo, P. Liu, X. Zhu, Z. Cheng, Z. Yu, S. Yang, M. Chen, and S. Xie (2017). “High-Speed Compressive Microscopy of Flowing Cells Using Sinusoidal Illumination Patterns”. In: *IEEE Photonics Journal* 9.1, pp. 1–11. DOI: 10.1109/JPHOT.2016.2644869.

Gustafsson, Agard, and Sedat (1999). “I5M: 3D widefield light microscopy with better than 100nm axial resolution”. In: *Journal of Microscopy* 195.1, pp. 10–16. DOI: 10.1046/j.1365-2818.1999.00576.x.

Gustafsson, M. G. L. (2000). “Surpassing the lateral resolution limit by a factor of two using structured illumination microscopy”. In: *Journal of Microscopy* 198.2, pp. 82–87. DOI: 10.1046/j.1365-2818.2000.00710.x.

Hansen, P. C. (1992). “Analysis of Discrete Ill-Posed Problems by Means of the L-Curve”. In: *SIAM Review* 34.4, pp. 561–580. DOI: 10.1137/1034115.

Hanslovsky, P., J. A. Bogovic, and S. Saalfeld (2016). “Image-based correction of continuous and discontinuous non-planar axial distortion in serial section microscopy”. In: *Bioinformatics* 33.9, pp. 1379–1386. DOI: 10.1093/bioinformatics/btw794.

Harmany, Z. T., R. F. Marcia, and R. M. Willett (2009). “Sparse poisson intensity reconstruction algorithms”. In: *2009 IEEE/SP 15th Workshop on Statistical Signal Processing*, pp. 634–637. DOI: 10.1109/SSP.2009.5278495.

Heintzmann, R., T. M. Jovin, and C. Cremer (2002). “Saturated patterned excitation microscopy—a concept for optical resolution improvement”. In: *J. Opt. Soc. Am. A* 19.8, pp. 1599–1609. DOI: 10.1364/JOSAA.19.001599.

Hess, S. T., T. P. Girirajan, and M. D. Mason (2006). “Ultra-High Resolution Imaging by Fluorescence Photoactivation Localization Microscopy”. In: *Biophysical Journal* 91.11, pp. 4258–4272. DOI: 10.1529/biophysj.106.091116.

Horbelt, S., M. Liebling, and M. Unser (2002). “Discretization of the radon transform and of its inverse by spline convolutions”. In: *IEEE Transactions on Medical Imaging* 21.4, pp. 363–376. DOI: 10.1109/TMI.2002.1000260.

Huang, C.-J., C.-T. Tu, C.-D. Hsiao, F.-J. Hsieh, and H.-J. Tsai (2003). “Germ-line transmission of a myocardium-specific GFP transgene reveals critical regulatory elements in the cardiac myosin light chain 2 promoter of zebrafish”. In: *Developmental Dynamics* 228.1, pp. 30–40. DOI: 10.1002/dvdy.10356.

Huang, D., E. A. Swanson, C. P. Lin, J. S. Schuman, W. G. Stinson, W. Chang, M. R. Hee, T. Flotte, K. Gregory, C. A. Puliafito, and J. G. Fujimoto (1991). "Optical Coherence Tomography". In: *Science* 254.5035, pp. 1178–1181. DOI: 10.1126/science.1957169.

Huisken, J., J. Swoger, F. D. Bene, J. Wittbrodt, and E. H. K. Stelzer (2004). "Optical Sectioning Deep Inside Live Embryos by Selective Plane Illumination Microscopy". In: *Science* 305.5686, pp. 1007–1009. DOI: 10.1126/science.1100035.

Jaques, C., L. Bapst-Wicht, D. F. Schorderet, and M. Liebling (2019). "Multi-Spectral Widefield Microscopy of the Beating Heart Through Post-Acquisition Synchronization and Unmixing". In: *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pp. 1382–1385. DOI: 10.1109/ISBI.2019.8759472.

Jaques, C., A. Ernst, N. Mercader, and M. Liebling (2020). "Temporal resolution doubling in fluorescence light-sheet microscopy via a hue-encoded shutter and regularization". In: *OSA Continuum* 3.8, pp. 2195–2209. DOI: 10.1364/OSAC.393193.

Jaques, C. and M. Liebling (2020). "Aliasing mitigation in optical microscopy of dynamic biological samples by use of temporally modulated color illumination and a standard RGB camera". In: *Journal of Biomedical Optics* 25.10, p. 106505. DOI: 10.1117/1.JBO.25.10.106505.

Jenkins, M. W., O. Q. Chughtai, A. N. Basavanhally, M. Watanabe, and A. M. Rollins (2007). "In vivo gated 4D imaging of the embryonic heart using optical coherence tomography". In: *Journal of Biomedical Optics* 12.3, p. 030505. DOI: 10.1117/1.2747208.

Kak, A. C., M. Slaney, and G. Wang (2002). "Principles of Computerized Tomographic Imaging". In: *Medical Physics* 29.1, pp. 107–107. DOI: 10.1118/1.1455742.

Kamilov, U. S., C. A. Bouman, G. T. Buzzard, and B. Wohlberg (2023). "Plug-and-Play Methods for Integrating Physical and Learned Models in Computational Imaging: Theory, algorithms, and applications". In: *IEEE Signal Processing Magazine* 40.1, pp. 85–97. DOI: 10.1109/MSP.2022.3199595.

Kamilov, U. S., I. N. Papadopoulos, M. H. Shoreh, A. Goy, C. Vonesch, M. Unser, and D. Psaltis (2016). "Optical Tomographic Image Reconstruction Based on Beam Propagation and Sparse Regularization". In: *IEEE Transactions on Computational Imaging* 2.1, pp. 59–70. DOI: 10.1109/TCI.2016.2519261.

Kikuchi, S., K. Sonobe, L. S. Sidharta, and N. Ohyama (1994). "Three-dimensional computed tomography for optical microscopes". In: *Optics Communications* 107.5, pp. 432–444. DOI: 0.1016/0030-4018(94)90361-1.

Kim, M. K. (2010). "Principles and techniques of digital holographic microscopy". In: *SPIE Reviews* 1.1, p. 018005. DOI: 10.1117/6.0000006.

Kimmel, C. B., W. W. Ballard, S. R. Kimmel, B. Ullmann, and T. F. Schilling (1995). "Stages of embryonic development of the zebrafish". In: *Developmental Dynamics* 203.3, pp. 253–310. DOI: 10.1002/aja.1002030302.

Kingma, D. P. and J. Ba (2015). "Adam: A Method for Stochastic Optimization". In: *International Conference on Learning Representations (ICLR 2015)*. DOI: 10.48550/arXiv.1412.6980.

Kirshner, H., F. Aguet, D. Sage, and M. Unser (2013). “3-D PSF fitting for fluorescence microscopy: implementation and localization application”. In: *Journal of Microscopy* 249.1, pp. 13–25. DOI: 10.1111/j.1365-2818.2012.03675.x.

Lawson, N. D. and B. M. Weinstein (2002). “In Vivo Imaging of Embryonic Vascular Development Using Transgenic Zebrafish”. In: *Developmental Biology* 248.2, pp. 307–318. DOI: 10.1006/dbio.2002.0711.

Lichtman, J. W. and J.-A. Conchello (2005). “Fluorescence microscopy”. In: *Nature Methods* 2.12, pp. 910–919. DOI: 10.1038/nmeth817.

Liebling, M., J. Vermot, A. Forouhar, M. Gharib, M. Dickinson, and S. Fraser (2006). “Nonuniform temporal alignment of slice sequences for four-dimensional imaging of cyclically deforming embryonic structures”. In: *3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro, 2006.* Pp. 1156–1159. DOI: 10.1109/ISBI.2006.1625128.

Liebling, M., A. S. Forouhar, M. Gharib, S. E. Fraser, and M. E. Dickinson (2005). “Four-dimensional cardiac imaging in living embryos via postacquisition synchronization of nongated slice sequences”. In: *Journal of biomedical optics* 10.5, p. 054001. DOI: 10.1117/1.2061567.

Liebling, M. and H. Ranganathan (2009). “Wavelet domain mutual information synchronization of multimodal cardiac microscopy image sequences”. In: *Wavelets XIII.* Vol. 7446, p. 744602. DOI: 10.1117/12.826519.

Liu, Y., J. Dong, T.-a. Pham, F. Marelli, and M. Unser (2022). “Mechanical artifacts in optical projection tomography: classification and automatic calibration”. In: *Opt. Continuum* 1.12, pp. 2577–2589. DOI: 10.1364/OPTCON.474540.

Loterie, D., S. Farahi, I. Papadopoulos, A. Goy, D. Psaltis, and C. Moser (2015). “Digital confocal microscopy through a multimode fiber”. In: *Opt. Express* 23.18, pp. 23845–23858. DOI: 10.1364/OE.23.023845.

Ma, X., X. Yuan, C. Fu, and G. R. Arce (2021). “LED-based compressive spectral-temporal imaging”. In: *Opt. Express* 29.7, pp. 10698–10715. DOI: 10.1364/OE.419888.

Maggioni, M., V. Katkovnik, K. Egiazarian, and A. Foi (2013). “Nonlocal Transform-Domain Filter for Volumetric Data Denoising and Reconstruction”. In: *IEEE Transactions on Image Processing* 22.1, pp. 119–133. DOI: 10.1109/TIP.2012.2210725.

Mahou, P., J. Vermot, E. Beaurepaire, and W. Supatto (2014). “Multicolor two-photon light-sheet microscopy”. In: *Nature Methods* 11.6, pp. 600–601. DOI: 10.1038/nmeth.2963.

Mäkinen, Y., L. Azzari, and A. Foi (2020). “Collaborative Filtering of Correlated Noise: Exact Transform-Domain Variance for Improved Shrinkage and Patch Matching”. In: *IEEE Transactions on Image Processing* 29, pp. 8339–8354. DOI: 10.1109/TIP.2020.3014721.

Mäkinen, Y., S. Marchesini, and A. Foi (2022). “Ring artifact and Poisson noise attenuation via volumetric multiscale nonlocal collaborative filtering of spatially correlated noise”. In: *Journal of Synchrotron Radiation* 29.3, pp. 829–842. DOI: 10.1107/S1600577522002739.

Marelli, F., A. Ernst, N. Mercader, and M. Liebling (2023). “PAAQ: Paired Alternating AcQuisitions for Virtual High Frame Rate Multichannel Cardiac Fluorescence Microscopy”. Submitted for publication.

Marelli, F. and M. Liebling (2021). “Optics Versus Computation: Influence of Illumination and Reconstruction Model Accuracy in Focal-Plane-Scanning Optical Projection Tomography”. In: *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pp. 567–570. DOI: 10.1109/ISBI48211.2021.9433834.

Marelli, F. and M. Liebling (2023). “Efficient compressed sensing reconstruction for 3D fluorescence microscopy using OptoMechanical Modulation Tomography (OMMT) with a 1+2D regularization”. Submitted for publication.

Mariani, O., K. Chan, A. Ernst, N. Mercader, and M. Liebling (2019). “Virtual High-Framerate Microscopy Of The Beating Heart Via Sorting Of Still Images”. In: *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pp. 312–315. DOI: 10.1109/ISBI.2019.8759595.

Mariani, O., A. Ernst, N. Mercader, and M. Liebling (2020). “Reconstruction of Image Sequences From Ungated and Scanning-Aberrated Laser Scanning Microscopy Images of the Beating Heart”. In: *IEEE Transactions on Computational Imaging* 6, pp. 385–395. DOI: 10.1109/TCI.2019.2948772.

Mariani, O., F. Marelli, C. Jaques, A. Ernst, and M. Liebling (2021). “Unequivocal Cardiac Phase Sorting From Alternating Ramp-And Pulse-Illuminated Microscopy Image Sequences”. In: *IEEE 18th International Symposium on Biomedical Imaging (ISBI 2021)*, pp. 868–872. DOI: 10.1109/ISBI48211.2021.9433858.

Marques, I. J., E. Lupi, and N. Mercader (2019). “Model systems for regeneration: zebrafish”. In: *Development* 146.18. dev167692. DOI: 10.1242/dev.167692.

McNally, J. G., T. Karpova, J. Cooper, and J. A. Conchello (1999). “Three-Dimensional Imaging by Deconvolution Microscopy”. In: *Methods* 19.3, pp. 373–385. DOI: 10.1006/meth.1999.0873.

Meiniel, W., P. Spinicelli, E. D. Angelini, A. Fragola, V. Loriette, F. Orieux, E. Sepulveda, and J.-C. Olivo-Marin (2017). “Reducing data acquisition for fast Structured Illumination Microscopy using Compressed Sensing”. In: *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, pp. 32–35. DOI: 10.1109/ISBI.2017.7950461.

Miao, Q., J. Hayenga, M. G. Meyer, T. Neumann, A. C. Nelson, and E. J. Seibel (2010). “Resolution improvement in optical projection tomography by the focal scanning method”. In: *Opt. Lett.* 35.20, pp. 3363–3365. DOI: 10.1364/OL.35.003363.

Mickoleit, M., B. Schmid, M. Weber, F. O. Fahrbach, S. Hombach, S. Reischauer, and J. Huisken (2014). “High-resolution reconstruction of the beating zebrafish heart”. In: *Nature Methods* 11.9, pp. 919–922. DOI: 10.1038/nmeth.3037.

Montefusco, L. B., D. Lazzaro, S. Papi, and C. Guerrini (2011). “A Fast Compressed Sensing Approach to 3D MR Image Reconstruction”. In: *IEEE Transactions on Medical Imaging* 30.5, pp. 1064–1075. DOI: 10.1109/TMI.2010.2068306.

Nguyen, C. T., Q. Lu, Y. Wang, and J.-N. Chen (2008). “Zebrafish as a model for cardiovascular development and disease”. In: *Drug Discovery Today: Disease Models* 5.3. Cardiovascular disease models, pp. 135–140. DOI: 10.1016/j.ddmod.2009.02.003.

O’Donoghue, B. (2021). “Operator Splitting for a Homogeneous Embedding of the Linear Complementarity Problem”. In: *SIAM Journal on Optimization* 31 (3), pp. 1999–2023. DOI: 10.1137/20M1366307.

Ohn, J. and M. Liebling (2011). “In vivo, high-throughput imaging for functional characterization of the embryonic zebrafish heart”. In: *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pp. 1549–1552. DOI: 10.1109/ISBI.2011.5872696.

Ohn, J., J. Yang, S. E. Fraser, R. Lansford, and M. Liebling (2011). “High-speed multicolor microscopy of repeating dynamic processes”. In: *genesis* 49.7, pp. 514–521. DOI: 10.1002/dvg.20774.

Olarte, O. E., J. Andilla, E. J. Gualda, and P. Loza-Alvarez (2018). “Light-sheet microscopy: a tutorial”. In: *Adv. Opt. Photon.* 10.1, pp. 111–179. DOI: 10.1364/AOP.10.000111.

Pellizzari, C. J., R. Trahan, H. Zhou, S. Williams, S. E. Williams, B. Nemati, M. Shao, and C. A. Bouman (2017). “Optically coherent image formation and denoising using a plug and play inversion framework”. In: *Appl. Opt.* 56.16, pp. 4735–4744. DOI: 10.1364/AO.56.004735.

Perlin, K. (1985). “An Image Synthesizer”. In: *SIGGRAPH Comput. Graph.* 19.3, pp. 287–296. DOI: 10.1145/325165.325247.

Perlin, K. (2002). “Improving Noise”. In: *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH ’02, pp. 681–682. DOI: 10.1145/566570.566636.

Perrin, M. D., R. Soummer, E. M. Elliott, M. D. Lallo, and A. Sivaramakrishnan (2012). “Simulating point spread functions for the James Webb Space Telescope with WebbPSF”. In: *Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave*. Vol. 8442. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 84423D, p. 84423D. DOI: 10.1117/12.925230.

Pitrone, P. G., J. Schindelin, L. Stuyvenberg, S. Preibisch, M. Weber, K. W. Eliceiri, J. Huisken, and P. Tomancak (2013). “OpenSPIM: an open-access light-sheet microscopy platform”. In: *Nature Methods* 10.7, pp. 598–599. DOI: 10.1038/nmeth.2507.

Power, R. M. and J. Huisken (2018). “Adaptable, illumination patterning light sheet microscopy”. In: *Scientific Reports* 8.1, p. 9615. DOI: 10.1038/s41598-018-28036-2.

Remacha, E., L. Friedrich, J. Vermot, and F. O. Fahrbach (2020). “How to define and optimize axial resolution in light-sheet microscopy: a simulation-based approach”. In: *Biomed. Opt. Express* 11.1, pp. 8–26. DOI: 10.1364/BOE.11.000008.

Rohr, S., N. Bit-Avragim, and S. Abdelilah-Seyfried (2006). “Heart and soul/PRKCi and nagie oko/Mpp5 regulate myocardial coherence and remodeling during cardiac morphogenesis”. In: *Development* 133.1, pp. 107–115. DOI: 10.1242/dev.02182.

Rohr, S., C. Otten, and S. Abdelilah-Seyfried (2008). “Asymmetric Involution of the Myocardial Field Drives Heart Tube Formation in Zebrafish”. In: *Circulation Research* 102.2, e12–e19. DOI: 10.1161/CIRCRESAHA.107.165241.

Rudin, L. I., S. Osher, and E. Fatemi (1992). “Nonlinear total variation based noise removal algorithms”. In: *Physica D: Nonlinear Phenomena* 60.1, pp. 259–268. DOI: 10.1016/0167-2789(92)90242-F.

Rust, M. J., M. Bates, and X. Zhuang (2006). “Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM)”. In: *Nature Methods* 3.10, pp. 793–796. DOI: 10.1038/nmeth929.

Ryu, E., J. Liu, S. Wang, X. Chen, Z. Wang, and W. Yin (2019). “Plug-and-Play Methods Provably Converge with Properly Trained Denoisers”. In: *Proceedings of the 36th International Conference on Machine Learning*. Vol. 97. Proceedings of Machine Learning Research, pp. 5546–5557.

Sablong, R., A. Rengle, A. Ramgolam, H. Saint-Jalmes, and O. Beuf (2014). “An Optical Fiber-Based Gating Device for Prospective Mouse Cardiac MRI”. In: *IEEE Transactions on Biomedical Engineering* 61.1, pp. 162–170. DOI: 10.1109/TBME.2013.2278712.

Salvador-Balaguer, E., P. Latorre-Carmona, C. Chabert, F. Pla, J. Lancis, and E. Tajahuerce (2018). “Low-cost single-pixel 3D imaging by using an LED array”. In: *Opt. Express* 26.12, pp. 15623–15631. DOI: 10.1364/OE.26.015623.

Sanderson, M. J., I. Smith, I. Parker, and M. D. Bootman (2014). “Fluorescence microscopy”. In: *Cold Spring Harbor Protocols* 2014.10, pp. 1042–1065. DOI: 10.1101/pdb.top071795.

Scherz, P. J., J. Huisken, P. Sahai-Hernandez, and D. Y. R. Stainier (2008). “High-speed imaging of developing heart valves reveals interplay of morphogenesis and function”. In: *Development* 135.6, pp. 1179–1187. DOI: 10.1242/dev.010694.

Sharpe, J., U. Ahlgren, P. Perry, B. Hill, A. Ross, J. Hecksher-Sørensen, R. Baldock, and D. Davidson (2002). “Optical Projection Tomography as a Tool for 3D Microscopy and Gene Expression Studies”. In: *Science* 296.5567, pp. 541–545. DOI: 10.1126/science.1068206.

Soldevila, F., P. Clemente, E. Tajahuerce, N. Uribe-Patarroyo, P. Andrés, and J. Lancis (2016). “Computational imaging with a balanced detector”. In: *Scientific Reports* 6.1, p. 29181. DOI: 10.1038/srep29181.

Stainier, D. Y. and M. C. Fishman (1994). “The zebrafish as a model system to study cardiovascular development”. In: *Trends in Cardiovascular Medicine* 4.5, pp. 207–212. DOI: 10.1016/1050-1738(94)90036-1.

Staudt, D. W., J. Liu, K. S. Thorn, N. Stuurman, M. Liebling, and D. Y. R. Stainier (2014). “High-resolution imaging of cardiomyocyte behavior reveals two distinct steps in ventricular trabeculation”. In: *Development* 141.3, pp. 585–593. DOI: 10.1242/dev.098632.

Sun, Y., B. Wohlberg, and U. S. Kamilov (2019). “An Online Plug-and-Play Algorithm for Regularized Image Reconstruction”. In: *IEEE Transactions on Computational Imaging* 5.3, pp. 395–408. DOI: 10.1109/TCI.2019.2893568.

Sun, Y., Z. Wu, X. Xu, B. Wohlberg, and U. S. Kamilov (2021). “Scalable Plug-and-Play ADMM With Convergence Guarantees”. In: *IEEE Transactions on Computational Imaging* 7, pp. 849–863. DOI: 10.1109/TCI.2021.3094062.

Supatto, W. and J. Vermot (2011). “Chapter two - From Cilia Hydrodynamics to Zebrafish Embryonic Development”. In: *Forces and Tension in Development*. Vol. 95. Current Topics in Developmental Biology, pp. 33–66. DOI: 10.1016/B978-0-12-385065-2.00002-5.

Swoger, J., P. Verveer, K. Greger, J. Huisken, and E. H. Stelzer (2007). “Multi-view image fusion improves resolution in three-dimensional microscopy”. In: *Opt. Express* 15.13, pp. 8029–8042. DOI: 10.1364/OE.15.008029.

Taylor, J. M. (2014). “Optically gated beating-heart imaging”. In: *Frontiers in Physiology* 5. DOI: 10.3389/fphys.2014.00481.

Taylor, J. M., J. M. Girkin, and G. D. Love (2012). “High-resolution 3D optical microscopy inside the beating zebrafish heart using prospective optical gating”. In: *Biomed. Opt. Express* 3.12, pp. 3043–3053. DOI: 10.1364/BOE.3.003043.

Taylor, J. M., C. D. Saunter, G. D. Love, J. M. Girkin, D. J. Henderson, and B. Chaudhry (2011). “Real-time optical gating for three-dimensional beating heart imaging”. In: *Journal of Biomedical Optics* 16.11, p. 116021. DOI: 10.1117/1.3652892.

Tian, L., X. Li, K. Ramchandran, and L. Waller (2014). “Multiplexed coded illumination for Fourier Ptychography with an LED array microscope”. In: *Biomed. Opt. Express* 5.7, pp. 2376–2389. DOI: 10.1364/BOE.5.002376.

Tian, L., J. Wang, and L. Waller (2014). “3D differential phase-contrast microscopy with computational illumination using an LED array”. In: *Opt. Lett.* 39.5, pp. 1326–1329. DOI: 10.1364/OL.39.001326.

Trull, A. K., J. van der Horst, W. J. Palenstijn, L. J. van Vliet, T. van Leeuwen, and J. Kalkman (2017). “Point spread function based image reconstruction in optical projection tomography”. In: *Physics in Medicine & Biology* 62.19, p. 7784. DOI: 10.1088/1361-6560/aa8945.

Venkatakrishnan, S. V., C. A. Bouman, and B. Wohlberg (2013). “Plug-and-Play priors for model based reconstruction”. In: *2013 IEEE Global Conference on Signal and Information Processing*, pp. 945–948. DOI: 10.1109/GlobalSIP.2013.6737048.

Vermot, J., A. S. Forouhar, M. Liebling, D. Wu, D. Plummer, M. Gharib, and S. E. Fraser (2009). “Reversing Blood Flows Act through klf2a to Ensure Normal Valvulogenesis in the Developing Heart”. In: *PLOS Biology* 7.11, pp. 1–14. DOI: 10.1371/journal.pbio.1000246.

Vonesch, C., F. Aguet, J.-L. Vonesch, and M. Unser (2006). “The colored revolution of bioimaging”. In: *IEEE Signal Processing Magazine* 23.3, pp. 20–31. DOI: 10.1109/MSP.2006.1628875.

Wang, Z., X. Cun, J. Bao, W. Zhou, J. Liu, and H. Li (2022). “Uformer: A General U-Shaped Transformer for Image Restoration”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17683–17693. DOI: 10.48550/arXiv.2106.03106.

Woringer, M., X. Darzacq, C. Zimmer, and M. Mir (2017). “Faster and less phototoxic 3D fluorescence microscopy using a versatile compressed sensing scheme”. In: *Opt. Express* 25.12, pp. 13668–13683. DOI: 10.1364/OE.25.013668.

Wu, Y., P. Wawrzusin, J. Senseney, R. S. Fischer, R. Christensen, A. Santella, A. G. York, P. W. Winter, C. M. Waterman, Z. Bao, et al. (2013). “Spatially isotropic four-dimensional imaging with dual-view plane illumination microscopy”. In: *Nature biotechnology* 31.11, pp. 1032–1038. DOI: 10.1038/nbt.2713.

Zamir, S. W., A. Arora, S. Khan, M. Hayat, F. S. Khan, and M.-H. Yang (2022). “Restormer: Efficient Transformer for High-Resolution Image Restoration”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5728–5739. DOI: 10.48550/arXiv.2111.09881.

Zhang, K., Y. Li, W. Zuo, L. Zhang, L. Van Gool, and R. Timofte (2022). “Plug-and-Play Image Restoration With Deep Denoiser Prior”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44.10, pp. 6360–6376. DOI: 10.1109/TPAMI.2021.3088914.

Zhang, K., W. Zuo, Y. Chen, D. Meng, and L. Zhang (2017). “Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising”. In: *IEEE Transactions on Image Processing* 26.7, pp. 3142–3155. DOI: 10.1109/TIP.2017.2662206.

Zhang, K., W. Zuo, and L. Zhang (2018). “FFDNet: Toward a Fast and Flexible Solution for CNN-Based Image Denoising”. In: *IEEE Transactions on Image Processing* 27.9, pp. 4608–4622. DOI: 10.1109/TIP.2018.2839891.

Zhang, Q. and R. Pless (2005). “Segmenting Cardiopulmonary Images Using Manifold Learning with Level Sets”. In: *Computer Vision for Biomedical Image Applications*, pp. 479–488. DOI: 10.1007/11569541_48.

Zheng, G., R. Horstmeyer, and C. Yang (2013). “Wide-field, high-resolution Fourier ptychographic microscopy”. In: *Nature Photonics* 7.9, pp. 739–745. DOI: 10.1038/nphoton.2013.187.

Zimmermann, T., J. Rietdorf, A. Girod, V. Georget, and R. Pepperkok (2002). “Spectral imaging and linear un-mixing enables improved FRET efficiency with a novel GFP2-YFP FRET pair”. In: *FEBS Letters* 531.2, pp. 245–249. DOI: 10.1016/S0014-5793(02)03508-1.

Zimmermann, T., J. Rietdorf, and R. Pepperkok (2003). “Spectral imaging and its applications in live cell microscopy”. In: *FEBS Letters* 546.1, pp. 87–92. DOI: 10.1016/S0014-5793(03)00521-0.

Zunino, A., F. Garzella, A. Trianni, P. Saggau, P. Bianchini, A. Diaspro, and M. Duocastella (2021). “Multiplane Encoded Light-Sheet Microscopy for Enhanced 3D Imaging”. In: *ACS Photonics* 8.11, pp. 3385–3393. DOI: 10.1021/acsphotonics.1c01401.

FRANÇOIS MARELLI

@ francois.marelli@idiap.ch
0000-0002-8261-238X
idiap.ch/~fmarelli
github.com/frailHand

EDUCATION AND EXPERIENCE

2018 – 2023 **Ph.D. candidate in Electrical Engineering** **EDEE, EPFL**
Practical computational imaging by use of spatiotemporal light modulation: from simulations to applications in biological microscopy. *Supervisors: Prof. M. Liebling and Prof. F. Fleuret*

2018 – 2023 **Research assistant** **Idiap Research Institute**
Computational Bioimaging Group
Developing optical microscopy imaging methods using precise illumination control to design effective reconstruction algorithms

2018 **Research intern** **Idiap Research Institute**
Speech & Audio Processing Group, supervisor: Dr. P. Garner
Integrating second order recurrent linear units in neural network systems for prosody modelling in speech synthesis

2017 **Research intern** **Idiap Research Institute**
Machine Learning Group, supervisor: Prof. F. Fleuret
Implementation and validation of DNN size reduction algorithms

2016 – 2018 **Founder member and President** **electroLAB, UMONS**
Student organization undertaking technological projects related to electronics

2016 – 2018 **M. Sc. summa cum laude** **FPMs, UMONS**
Majoring in Electrical Engineering, minor in Signal, Systems and Biomedical Engineering

2016 **Research intern** **FPMs, UMONS**
Telecom department
Investigating IoT protocols and LoRa transmission, implementing a LoRa communication testing software and prototype

2013 – 2016 **B. Sc. summa cum laude** **FPMs, UMONS**
Majoring in Electrical Engineering

PUBLICATIONS

Journal papers

2023 **Efficient Compressed Sensing Reconstruction for 3D Fluorescence Microscopy Using OptoMechanical Modulation Tomography (OMMT) with a 1+2D Regularization**
Marelli F, *Liebling M.* Submitted to Optics Express

2023 **PAAQ: Paired Alternating AcQuisitions for Virtual High Frame Rate Multichannel Cardiac Fluorescence Microscopy**
Marelli F, *Ernst A, Mercader N, Liebling M.* Submitted to Biological Imaging

2022 **Mechanical Artifacts in Optical Projection Tomography: Classification and Automatic Calibration**
Liu Y, Dong J, Pham TA, ***Marelli F****, Unser M.* Optics Continuum, 1(12):2577–89

Conference papers

2022 **HyperMixer: An MLP-based Green AI Alternative to Transformers**
Mai F, Pannatier A, Fehr F, Chen H, ***Marelli F****, Fleuret F, Henderson J.* Submitted to the 61st Annual Meeting of the Association for Computational Linguistics (ACL 2023)

2021 **Optics Versus Computation: Influence of Illumination and Reconstruction Model Accuracy in Focal-Plane-Scanning Optical Projection Tomography**
Marelli F, *Liebling M.* Presented at the IEEE 18th International Symposium on Biomedical Imaging (ISBI 2021)

2021 **Unequivocal Cardiac Phase Sorting From Alternating Ramp-And Pulse-Illuminated Microscopy Image Sequences**
Mariani O, ***Marelli F****, Jaques C, Ernst A, Liebling M.* In IEEE 18th International Symposium on Biomedical Imaging (ISBI 2021)

2021 **Perspectives and Limitations of Visible-thermal Image Pair Synthesis Via Generative Adversarial Networks**
Panchard D, ***Marelli F****, Presa ED, Wellig P, Liebling M.* In SPIE: Target and Background Signatures VII (2021)

2019 **An End-to-end Network to Synthesize Intonation Using a Generalized Command Response Model**
Marelli F*, Schnell B, Bourlard H, Dutoit T, Garner PN.* Presented at the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2019)

Poster presentations

2022 **Joint Annual Conference of the Austrian, German and Swiss Societies for Biomedical Engineering (BMT 2022)**
Post-acquisition Framerate Improvement for Multi-channel Fluorescence Cardiac Microscopy

2022 **14th IEEE EMBS-SPS International Summer School on Biomedical Imaging**
Post-acquisition Framerate Improvement for Multi-channel Fluorescence Cardiac Microscopy

2022 **BioImage Informatics 2021 (BII 2022)**
OpenLightControl: an open platform for hardware synchronization and temporal modulation of microscopic imaging devices

TEACHING

2020 – 2022 **EE311: Fundamentals of machine learning** **EPFL**
Co-writing the practical sessions, *Lecturers: F. Fleuret, M. Liebling*

2021 – 2022 **Direct calibration of a rotational stage from incomplete focal-plane scanning optical projection tomography data** **EPFL**
Supervision of N. Polvani's Master's project

OUTREACH

2022 **Intelligence Artificielle. Nos reflets dans la machine: exhibition in Musée de la main**
Designing and building an interactive demonstrator for image restoration using GANs

2019 **EPFL 50th Anniversary Open Days**
Presenting a live demo of OPT microscopy

2019 **Scientastic, the EPFL Science Festival**
Presenting a demo of the OpenSPIM platform and OPT microscopy

SKILLS

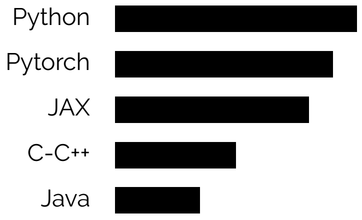

LANGUAGES

INTERESTS

- Music (multiple instruments)
- Board Games
- Aikido
- Photography
- Reading
- DIY Electronics