

A Bayesian framework for the analog reconstruction of kymographs from fluorescence microscopy data

Denis K. Samuylov, Gábor Székely, and Grégory Paul

Abstract—Kymographs are widely used to represent and analyse spatio-temporal dynamics of fluorescence markers along curvilinear biological compartments. These objects have a singular geometry, thus kymograph reconstruction is inherently an analog image processing task. However, the existing approaches are essentially digital: the kymograph photometry is sampled directly from the time-lapse images. As a result, such kymographs rely on raw image data that suffer from the degradations entailed by the image formation process and the spatio-temporal resolution of the imaging setup. In this work, we address these limitations and introduce a well-grounded Bayesian framework for the analog reconstruction of kymographs. To handle the movement of the object, we introduce an intrinsic description of kymographs using differential geometry: a kymograph is a photometry defined on a parameter space that is embedded in physical space by a time-varying map that follows the object geometry. We model the kymograph photometry as a Lévy innovation process, a flexible class of non-parametric signal priors. We account for the image formation process using the virtual microscope framework. We formulate a computationally tractable representation of the associated maximum *a posteriori* problem and solve it using a class of efficient and modular algorithms based on the alternating split Bregman. We assess the performance of our Bayesian framework on synthetic data and apply it to reconstruct the fluorescence dynamics along microtubules *in vivo* in the budding yeast *S. cerevisiae*. We demonstrate that our framework allows revealing patterns from single time-lapse data that are invisible on standard digital kymographs.

Index Terms—Quantitative fluorescence microscopy, model-based image processing, Bayesian modelling, alternating split Bregman, Lévy innovation processes, virtual microscope framework

I. INTRODUCTION

FLUORESCENCE microscopy is a powerful tool for studying the dynamics of biological processes. In many cases, these processes take place in specific compartments. As geometries, these compartments represent different levels of spatial restriction: inside a volume (*e.g.* cytoplasm), on a surface (*e.g.* membranes), or along a curvilinear object (*e.g.* axons, microtubules, actin filaments). In this paper, we focus on processes restricted to curvilinear geometries. A standard representation for characterising the spatio-temporal dynamics of such processes is a *kymograph* [1]: a two-dimensional representation of the signal along the curvilinear object (*e.g.* displayed as the ordinate) varying in time (*e.g.* displayed as the abscissa). Kymographs are custom in biology to study regulatory mechanisms coordinating spatio-temporal protein interactions (see *e.g.* [2], [3], [4], [5]). Thus, there is a high interest in developing computational tools to build and analyse kymographs [6], [7], [8], [9].

Reconstructing kymographs from a sequence of images is a challenging image processing task. It requires jointly estimating *the geometry and the photometry* of a curvilinear object. When the objects move during the time-lapse imaging, the complexity of the task increases as in addition it requires solving a *tracking problem*. Reconstructing a kymograph is an *inherently analog image processing task* because the typical diameter of such biological objects is smaller than the pixel size, and not aligned with the pixel grid geometry. To the best of our knowledge, the existing kymograph reconstruction frameworks have been mostly digital in nature: the kymograph is reconstructed by directly sampling the image data, with an optional pre- or post-processing. In addition, the geometry and the photometry are estimated sequentially.

A standard approach to reconstruct kymographs consists of mainly two steps: estimating the curvilinear object geometry and sampling the grey values along this geometry. When the image data are acquired at a single focal plane, the geometry is outlined manually [3], [10], [6], [7], [8], [9], semi-automatically [3], [11], or automatically [3] on the average or on the maximum projection of the time-lapse data. When the image data are acquired at multiple focal planes, the maximum projection is first applied along the axial direction [12], [3]. The axial coordinate is either discarded [12] or obtained from a manual annotation of the curvilinear object geometry on a cross-sectional image between the time-lapse images projected in time and the outlined two-dimensional geometry [3]. It is common to dilate the estimated geometry and orthogonally project the enclosed pixel onto the estimated curve by computing the maximum projection [13], [11], [14], [15], [6]. It makes the reconstructed kymographs robust to small errors in the estimated geometry and its variations during time-lapse imaging. It also allows observing the intensity variations within a region of interest [13]. However, the kymograph reconstruction will fail in case of high object displacement due to drifting or to the underlying object dynamics [15]. To address the problem of reconstructing kymographs, it has been suggested to track their geometry using semi- [15] or fully-automated [14] approaches. To facilitate the kymograph processing (*e.g.* estimating trajectory, direction, velocity of vesicles that moves along the curvilinear geometries), it has been suggested to denoise [6] and deconvolve the original time-lapse [3] as a pre-processing step, or to apply digital filters to the kymographs as a post-processing step [3], [11], [15]. In [13], a super-resolution microscopy technique (structured-illumination microscopy operated in an total internal reflection fluorescence mode) yields high-resolution images, and hence high-resolution kymographs.

A kymograph aims to represent the fluorescence signal along a curvilinear object in order to characterise the underlying spatio-temporal dynamics. However, the aforementioned digital techniques reconstruct the kymograph by sampling or averaging image grey values that suffer from two sources of degradation: first, *the distortions* due to light emission, propagation throughout the environment (*e.g.* cytoplasm, coverslip, microscope objective) and conversion into grey values; second, *the spatio-temporal resolution limit* entailed by the pixel grid and the acquisition duration. Therefore, these approaches implicitly assume that these degradations are negligible. This assumption is reasonable for objects that are large and immobile compared to the spatio-temporal resolution of the imaging setup. Those factors set limits on the spatio-temporal scales that can be directly resolved, decrease the quality of the kymograph reconstructions and complicate their subsequent analysis [16], [10].

To address these limitations, we establish a minimal Bayesian framework that allows formulating inverse problems for the analog reconstruction of kymographs in fluorescence microscopy. We assume that the geometry is given as the solution of another inverse problem that is not the focus of this paper. The main challenge addressed in this work is to formulate and solve an *analog inverse problem* to reconstruct the kymograph before degradation.

II. MOTIVATION: A SUPER-RESOLUTION PERSPECTIVE

Super-resolution is an inverse problem characterised by a reconstruction space that is different from the sampling space. Two archetypal super-resolution problems found in the literature are *zooming* and *source localisation*: zooming amounts to reconstructing an image on a grid finer than the sampled data; source localisation amounts to reconstructing positions off the sampling grid.

Following [17], imaged objects are characterised by a geometry and a photometry attribute. Informally, geometry describes *where* the light sources are, and photometry describes their intensity value. In zooming, the main focus is to reconstruct the photometry, whereas the geometry is either given or a nuisance parameter. In contrast, in source localisation, the main interest is to reconstruct the geometry (*i.e.* source positions), and the photometry is treated as a nuisance parameter.

In the literature, the kymograph reconstruction problem assumes a given geometry and concentrate on reconstructing the photometry. Therefore, the problem we address in this work is similar to the zooming problem, on a curve.

1) The geometry and photometry of deconvolution: In this section we revisit the convolution linear inverse problem from the perspective of geometry and photometry. The mean intensity μ generated by a distribution of n_s point sources sampled on a grid of n_p pixels is modelled as a linear equation:

$$\underset{n_p \times 1}{\mu} = \underset{n_p \times 1}{\mu_{bg}} + \underset{n_p \times n_s}{M} \underset{n_s \times 1}{\varphi}, \quad (1)$$

where μ_{bg} is an offset due to the background, and M is a matrix representing both the convolution and the sampling (integration in space and time). Each line of M corresponds to a pixel grid location, and each column a point source location.

When these two sets of locations coincide, we obtain the custom convolution matrix [18]. Another way of looking at M is that each column holds the integrated/sampled point spread function (PSF) shifted at each point source location. The vector φ hence represents the intensity of each point source, the mean expectation image resulting from the sum of the shifted, scaled, and integrated PSF kernels.

This straightforward identification reveals that the geometry is *implicitly* encoded in the column space of the matrix M , whereas the photometry is *explicitly* encoded in φ . This insight is crucial to motivate the approach we develop for solving the kymograph reconstruction problem.

2) Interplay between geometry and photometry in the conditioning of deconvolution: The sensitivity of the solutions of the linear inverse problem (1) to perturbations in $\mu - \mu_{bg}$ is quantified by the condition number of the matrix M , denoted $\kappa(M) \in [1, \infty]$, see [19], [20]. Another insightful interpretation is given by [21], [22]. These authors show that the reciprocal condition number (RCN), denoted $1/\kappa \in [0, 1]$, is a *scaled distance to the nearest ill-posed problem*. It assumes the value 0 for ill-posed problem, and the closer to 0, the more ill-conditioned. Ill-conditioning is an intrinsic property of the inverse problem that makes it hard to solve in practice with finite precision arithmetic.

In the previous section, we have shown that the column space of M is related to the localisation of the point sources. Therefore, the conditioning of the deconvolution problem (1) is related to the geometry of the underlying point sources. To quantify this relationship, we need another insight about the condition number. It quantifies the amount of collinearity in column space: if there exists at least two colinear column vectors, M is *rank deficient* and the linear inverse problem (1) is ill-posed, *i.e.* $1/\kappa(M) = 0$.

Result 1. *RCN for the linear deconvolution problem (1). The deconvolution problem is ill-posedness whenever at least two point sources occupy exactly the same position ($1/\kappa(M) = 0$). In such a case, the deconvolution problem (1) admits an infinite number of solutions.*

*The deconvolution problem is well-posed when all point sources occupy different positions. The RCN decreases with the total overlap of the PSF kernels. For compactly supported PSFs, the maximum RCN is obtained when none of the support overlap. In such a case, the RCN equals the ratio between the smallest norm and the largest among the column vectors of M , *i.e.* $\min_k \|M_{\cdot k}\|_2 / \max_k \|M_{\cdot k}\|_2$. For shift-invariant compactly-supported PSF kernels far from the image boundary, the RCN is equal to one.*

Proof. The geometry is implicitly encoded in the column space. If at least two sources occupying the same location, this translates in many identical columns in M , and hence a zero RCN.

If all the sources occupy different locations, the columns are all different and M has maximal rank. However, a small singular value of order k means that there is k column vectors that are nearly colinear. This happens when the corresponding sources are close enough to have their PSF affecting each other, *i.e.* when their supports overlap.

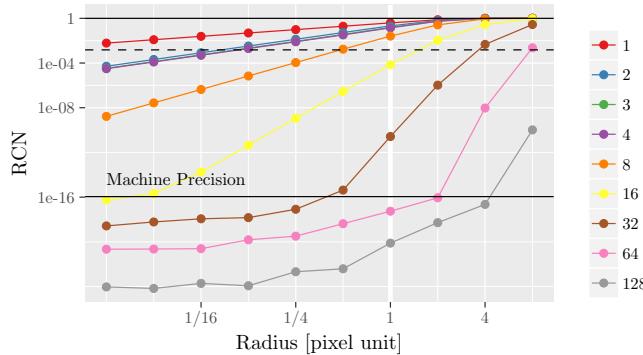


Figure 1. Influence of the geometry on the reciprocal condition number. The reciprocal condition number (RCN) is computed for different numbers of neighbors (between 1 and 128) lying on a circle of radii 2^k with k an integer in $[-6, 3]$. One source is located at the center of the imaging plane, its n neighbors are positioned around this central source according to the n -th roots of unity $(\cos 2\pi/k, \sin 2\pi/k)_{k=0:(n-1)}$, and scaled at different radii. For reference, we show the RCN of the custom convolution matrix builds on the whole pixel grid, with Dirichlet boundary conditions (dashed line), and the machine precision. The super-resolution regime starts at a radius below 1.

For a compactly supported PSF, disjoint supports translate in $\mathbf{M}^T \mathbf{M}$ being diagonal, with $(\mathbf{M}^T \mathbf{M})_{ii} = \|M_i\|_2^2$ being the eigenvalues of $\mathbf{M}^T \mathbf{M}$. The RCN of $\mathbf{M}^T \mathbf{M}$ is therefore the ratio between the smallest and the largest eigenvalues. Using

$$\kappa(\mathbf{M}) = \sqrt{\kappa(\mathbf{M}^T \mathbf{M})},$$

see for example [23], one concludes the result for compactly supported PSF. For shift-invariant and compactly supported PSFs, the columns of \mathbf{M} will have the same norm, except for sources located close to the image boundaries. For example, with Dirichlet boundary conditions, point sources close enough to the boundaries loose the PSF contributions beyond the image boundary. \square

To illustrate this theoretical result, we computed the RCN for increasing amounts of PSF support overlap. We increase the overlap by either increasing the number of point sources, or by reducing the distance between them.

For the PSF and the pixel size, we use the imaging parameters shown in Table I. To make the analysis more tractable, we consider a single imaging plane (*i.e.* a 2D problem). We assemble \mathbf{M} for different number of sources and configurations. A single source is placed at the centre of the imaging plane, and neighbouring sources are placed on a circle according to the roots of unity. Reducing the radius of the circle and increasing the number of sources lying on the circle both increase the total overlap between the PSF supports.

In Fig. 1 we observe that for the largest radius (8 pixels), a small number of sources achieve the optimal RCN of 1. From 32 sources on, their PSF start interacting laterally, and the amount of overlap increases with the number of sources: the RCN decreases accordingly. When the circle radius decreases, the sources interact even more. Already for two sources (*i.e.* for one neighbouring source, in red), the impact is moderate and the RCN decreases slowly, even at sub-pixel distances (*i.e.* radius below 1). However, for an increasing number of

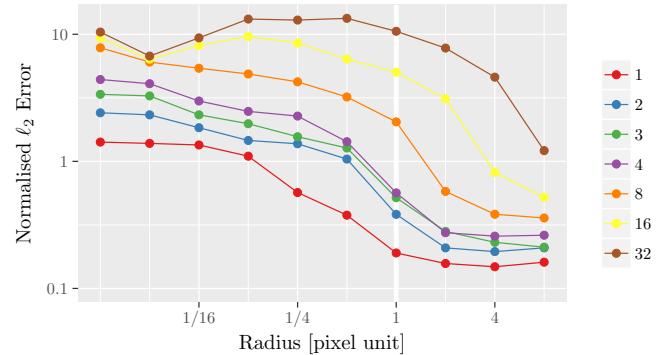


Figure 2. Geometry-photometry interplay in reconstruction accuracy. We use the same setup as in Fig. 1. The background intensity is uniform and set to $\mu_{bg} = 10$. Each source have an intensity of $\varphi_{true} = 400$. We report the median of the ℓ_2 error between the estimated intensity vector $\hat{\varphi}$ and the true vector among 100 independent Poisson noise realisation. The ℓ_2 error is normalised by φ_{true} . Due to very bad conditioning beyond 32 neighbours (see Fig. 1) L-BFGS-B fails and the results are not shown.

sources, the decrease in RCN is faster with the number of sources.

3) Resolution/well-posedness tradeoff: Choosing a geometry to reconstruct the photometry, *i.e.* choosing \mathbf{M} in equation (1), entails a geometric resolution defined by the minimum distance between the point sources. However, for diffraction-limited imaging, the PSF shape imposes a practical resolution limit, preventing choosing arbitrarily fine geometries [24]. In Fig. 1 we give an inverse problem stability perspective by showing that refining the geometric configuration brings the deconvolution problem (1) closer to an ill-posed problem. However, in practice, the interest is to be able to reconstruct the intensity accurately, especially for the kymograph problem, where the photometry reconstruction accuracy is crucial for studying the underlying biological processes.

In Fig. 2, we simulate for each experimental design used in Fig. 1 100 independent Poisson noise realisation, where the forward problem (1) is used with $\mu_{bg} = 10$, and all point sources having the same intensity, $\varphi_{true} = 400$. For each artificial image, the maximum likelihood estimate is computed using L-BFGS-B with positivity constraints [25], and the ℓ_2 error is computed and normalised to φ_{true} . We show the median among the 100 repeats.

This analysis shows that the photometry of sub-pixel geometric configurations is a hard inverse problem, but can still be approached when only few sources are considered. This agrees with the stability perspective shown in section II-2. Therefore there is a tradeoff between the geometric resolution where one can reasonably reconstruct the photometry and the well-posedness of the deconvolution inverse problem.

4) Reconstructing kymographs sequentially is at worst ill-posed and at best limited by the pixel resolution: The previous insights allow us to motivate the challenges of the kymograph reconstruction problem. Custom digital kymograph reconstruction algorithms amount to first deconvolve/denoise the image on the pixel grid, and then to estimate the photometry at any given point along the curve by reading the reconstructed photometry at the nearest pixel. This can be formalised as a

modified version of equation (1):

$$\underset{n_p \times 1}{\boldsymbol{\mu}} = \underset{n_p \times 1}{\boldsymbol{\mu}_{\text{bg}}} + \underset{n_p \times n_p}{\boldsymbol{M}} \underset{n_p \times n_s}{\boldsymbol{S}} \underset{n_s \times 1}{\boldsymbol{\varphi}}, \quad (2)$$

where \boldsymbol{M} is now the standard digital convolution matrix, and \boldsymbol{S} is a matrix that assigns each n_s position along the curvilinear geometry to its nearest neighbour pixel. Mathematically it is a binary matrix with exactly a single one in each column. The matrix \boldsymbol{S} acts on \boldsymbol{M} by selecting a multi-set (*i.e.* repetitions are possible) of its columns. Once this equation is inverted for \boldsymbol{M} using any standard deconvolution algorithm, \boldsymbol{S} has a trivial left inverse (its transpose) acting as a selection matrix that picks up the nearest neighbour pixels.

However our previous analysis tells us that this inverse problem (1) is either ill-posed, or the finest kymograph resolution is limited to the pixel size. If one attempts to reconstruct a kymograph with at least two points that corresponds to the same nearest pixel, \boldsymbol{S} , and hence \boldsymbol{MS} will have duplicated columns, and therefore be ill-posed. The only way to be well-posed is to ensure to choose points along the curve that are assigned to unique pixels. However, this has three main limitations: the resolution cannot be finer than a pixel; the kymograph sampling depends on the embedding of the curve in physical space; the kymograph sampling is uniform only for curves aligned with the image axes.

Nevertheless, usually $n_s \ll n_p$ and hence \boldsymbol{MS} have much less columns than \boldsymbol{M} . From the previous insights, this is a very attractive aspect. Exploiting the knowledge of the underlying geometry of the kymograph offers the opportunity to reduce both the number of locations needed to reconstruct the photometry, and the amount of overlap between the associated PSF supports. Indeed, in a 1D topology, points have less neighbours than on a 3D digital grid.

In what follows we show how to overcome these limitations and benefit from exploiting the geometry underlying the kymograph using the virtual microscope framework [17].

III. FORWARD PROBLEM

A. Object model

1) *Measure-theoretic object model for incoherent imaging:* In [17], we define objects as a spatio-temporal distribution of light sources (*photometry*) restricted to a subspace of the physical space (*geometry*). In fluorescence microscopy, an object corresponds to the spatio-temporal distribution of the fluorescently labeled proteins under scrutiny within a biological compartment. Mathematically, this is captured by two ingredients: for an object indexed by l , the photometry is defined as a positive measure in space and time encoding generalised distributions of light sources, denoted $\phi_t^l(\mathbf{dy} \times dt)$, and the geometry is encoded by a time-varying piecewise-Riemannian manifold, denoted \mathcal{M}_t^l . An object entails an *object measure*, defined as the photometry measure restricted to the manifold:

$$\phi_{t,\mathcal{M}_t^l}(\mathbf{dy} \times dt) := \mathbb{1}_{\mathcal{M}_t^l}(\mathbf{y}) \phi_t^l(\mathbf{dy} \times dt), \quad (3)$$

where $\mathbb{1}_{\mathcal{M}_t^l}(\mathbf{y})$ is the indicator function assuming 1 if \mathbf{y} is on the manifold \mathcal{M}_t^l and 0 otherwise. In what follows, we

assume that the objects emit light continuously in time, which is modelled as a photometry measure proportional to the Lebesgue measure dt :

$$\phi_t^l(\mathbf{dy} \times dt) = \phi_t^l(\mathbf{dy}) dt.$$

Due to the random nature of photon emission, the fluorescence microscopy is an incoherent imaging process [26]. Therefore, the total photon flux emitted by a set of objects in the imaging volume is a linear combination of the flux emitted by each individual object:

$$\phi_t(\mathbf{dy} \times dt) := \sum_{l \in \mathcal{L}} \phi_{t,\mathcal{M}_t^l}(\mathbf{dy} \times dt). \quad (4)$$

where \mathcal{L} is the set of object labels.

2) *Object model:* In fluorescence microscopy, the objects of interest cannot be imaged in isolation. Photons emitted by objects distant from the fluorescent marker of interest also contribute to the total photon counts. Therefore, we model two kinds of objects: the signal emitted by the background (due to the auto-fluorescence of the medium and the diffuse component of the labeled proteins in the available cellular volume) and the fluorescent proteins restricted to a curvilinear compartment.

a) *Background object model:* We assume that the background signal is uniform in space but decreases in time due to photobleaching, $\phi_t^{\text{bg}} := \phi_{t,\Omega}^{\text{bg}}$, where:

$$\phi_t^{\text{bg}}(\mathbf{dy} \times dt) = \mathbb{1}_\Omega(\mathbf{y}) \varphi_t^{\text{bg}} \mathbf{dy} dt, \quad (5)$$

where $\Omega \subset \mathbb{R}^3$ denotes the imaging volume (subset of the physical space), and $\varphi_t^{\text{bg}} \in \mathbb{R}_+$ is the background intensity. The imaging volume Ω is omitted because the indicator function $\mathbb{1}_\Omega$ will assume one in practice.

b) *Curvilinear object model:* We assume that the curvilinear compartment is described by an open curve, denoted \mathcal{C}_t , and that the signal emitted by the fluorescent markers attached along the geometry has a density with respect to the spatio-temporal Lebesgue measure, denoted $\varphi_t(\mathbf{y}) \in \mathbb{R}_+$. We write the object measure as $\phi_{t,c_t} := \phi_{t,c_t}$, where:

$$\phi_{t,c_t}(\mathbf{dy} \times dt) = \mathbb{1}_{c_t}(\mathbf{y}) \varphi_t(\mathbf{y}) \mathbf{dy} dt. \quad (6)$$

3) *Kymograph-to-object mapping:* We show that the notion of parameterisation of the geometry developed in [17] is the key mathematical concept to define the notion of kymograph generically.

As introduced in the virtual microscope framework [17], we parameterise the manifold encoding the geometry and use a map to embed it into physical space. For an open curvilinear object of length Λ_t at time t , the parameter space is a one-dimensional interval, denoted $\mathcal{D}_t := [0, \Lambda_t] \subset \mathbb{R}_+$. The parameter space represents an intrinsic coordinate system attached to the curvilinear object. The embedding is a time-varying map from the parameter space \mathcal{D}_t to the curvilinear manifold \mathcal{C}_t , denoted $\sigma_t : \mathcal{D}_t \rightarrow \mathcal{C}_t \subset \Omega$. It defines the geometry of the object as $\mathcal{C}_t := \sigma_t(\mathcal{D}_t)$ and allows the intrinsic coordinate system to follow the geometry evolution in time.

In fluorescence microscopy, a kymograph represents the time evolution of the distribution of light sources along a coordinate system intrinsic to the curvilinear structure under

scrutiny. This precisely corresponds to the time evolution of the photometry measure in parameter space. Therefore, we call the set of parameter spaces the *kymograph space* or *kymospace*, denoted $\mathcal{K} := \mathcal{D}_{\mathcal{T}}$, where $\mathcal{T} \subset \mathbb{R}_+$ is the time interval during which the object is observed. The set of maps embedding the kymograph in the spatio-temporal volume, denoted $\sigma_{\mathcal{T}}$, defines a bijective mapping between a point \mathbf{y} on the curvilinear object \mathcal{C}_t in physical space and a point (t, ℓ) in kymograph space, where $\ell := \sigma_t^{-1}(\mathbf{y})$. We denote the fluorescence signal of a curvilinear object in kymograph space using the same notation as in physical space and we define it at position (t, ℓ) as $\varphi_t(\ell) := \varphi_t(\sigma_t(\ell))$. Finally, we define the *kymograph* on the kymospace $\mathcal{K}_{\mathcal{T}}$ as:

$$\mathcal{K}(\varphi) := \left\{ (t, \ell, \varphi_t(\ell)) : t \in \mathcal{T}, \ell \in \mathcal{D}_t \right\}. \quad (7)$$

4) Lévy process modelling the object photometry: In order to develop a Bayesian formulation of the kymograph reconstruction problem, we use the general framework of [27] to define a large class of priors for the photometry density φ_t .

a) Reconstruction space: We digitalise the continuous domain fluorescence signal by projecting it onto a reconstruction space. Given a reconstruction space at resolution Δ_{ℓ} , the approximated continuous-domain signal is [28]:

$$\varphi_t(\ell) = \sum_{p=0}^{n_t^b-1} \varphi_t[p] \beta\left(\frac{\ell}{\Delta_{\ell}} - p\right) = \boldsymbol{\beta}_t^T(\ell) \boldsymbol{\varphi}_t, \quad (8)$$

where β is an interpolation basis function defined on each parameter space, $n_t^b = \text{ceil}(\Lambda_t/\Delta_{\ell})$ is the number of knots, $\boldsymbol{\beta}_t(\ell) \in \mathbb{R}_+^{n_t^b}$ is the vector holding the shifted/scaled basis functions, and $\boldsymbol{\varphi}_t \in \mathbb{R}_+^{n_t^b}$ is the vector of *digital intensities* $\varphi_t[p] := \varphi_t(\ell)|_{\ell=p\Delta_{\ell}}$. Therefore, the reconstruction of the continuous fluorescence signal along a curvilinear object amounts to estimating a finite set of weights $\boldsymbol{\varphi}_t$.

b) Statistical model of the photometry: Following [27], a large class of stochastic processes is derived from the principle that a whitening linear operator L_t transforms the process φ_t into a canonical *Lévy innovation process* w_t :

$$L_t \varphi_t = w_t. \quad (9)$$

The nature of the interpolation basis function β in (8) is related to the whitening operator L_t (see [27], [28]).

The class of signals described in terms of a whitening operator and a Lévy innovation process generalises classical stochastic processes (*e.g.* Brownian diffusion, Poisson processes) and allows a better understanding of sparse stochastic processes (see [27]). Sparsity is a unifying concept of modern statistical models in signal processing (*e.g.* [29]) as well as a fundamental principle underlying modern imaging and reconstruction techniques in bio-imaging, such as super-resolution (*e.g.* [24], [30]).

Following [28], the discretisation at resolution Δ_{ℓ} of the generalised stochastic process described by (9) writes:

$$\mathbf{L}_t \boldsymbol{\varphi}_t = \mathbf{u}_t, \quad (10)$$

where $\mathbf{L}_t \in \mathbb{R}^{n_t^b \times n_t^b}$ is the matrix representation of the discrete counterpart of the operator L_t , and $\mathbf{u}_t \in \mathbb{R}^{n_t^b}$ is the discrete

innovation process. The statistical features of the discretised signal and the continuous-domain innovation process are directly related via the Lévy exponent ([28], Theorem 3).

To simplify the inverse problem algorithm, we restrict ourselves to first-order whitening operators. In that case, the signal φ_t is a Lévy process, and the interpolation basis function is a B-spline of degree zero [28], *i.e.* $\beta(x)$ assumes 1 if $x \in [0, 1]$ and 0 otherwise. Since the basis functions are non-overlapping, the increments of the innovation process are independent and identically distributed (see [31], [28]): $\mathbb{P}(\mathbf{u}_t) = \prod_{p=0}^{n_t^b-1} \mathbb{P}_U(\mathbf{u}_t[p])$. Introducing the negative log-transformed probability $\varpi_u := -\log \mathbb{P}_U$ to rewrite the latter equation in potential form:

$$\varpi(\mathbf{u}_t) := -\log \mathbb{P}(\mathbf{u}_t) = \sum_{p=0}^{n_t^b-1} \varpi_u(\mathbf{u}_t[p]). \quad (11)$$

B. Image formation model

In fluorescence microscopy, the image formation process consists of two steps. First, the signal emitted by fluorescently labeled objects is distorted due to the random nature of light emission, the propagation through the environment (*e.g.* cytoplasm, coverslip, immersion layer, optics in the microscope objective), and the sampling at the pixel grid. Second, the signal is corrupted by the measurement noise (*e.g.* spurious charge, amplification noise, readout noise) and the quantisation at the camera detector incurred by its encoding into grey values.

1) Object-to-pixel mapping:

a) General object-to-pixel mapping: We assume that the microscope is operated in a regime where Poisson shot noise (due to the random nature of photon emission) dominates the photon counting statistics [32]. Therefore, we assume that the number of photons collected at pixel $\mathcal{P}_j \subset \Omega$ at time t_k during t_a (*i.e.* during the interval $\mathcal{T}_{t_k, t_a} := [t_k, t_k + t_a]$) is a random variable following a Poisson distribution parameterised by the *total expected photon count* μ_{kj} , *i.e.* $N_{kj}^{\text{photon}} \sim \text{Poisson}(\mu_{kj})$. We also assume that the family of random variables $\{N_{kj}^{\text{photon}}\}_{kj}$ is mutually independent, but not identically distributed. Indeed, the optical distortions correlate the photon statistics in space. This is captured by the point spread function (PSF) kernel, denoted κ . We assume that the PSF is shift-invariant, hence acting as a convolution operator on the object measure. The resulting measure characterises the expected photon flux in space and time: $\Phi(\phi_t) := \kappa * \phi_t$, where the convolution is understood in the sense of measures, (see [17]). The expected number of photons collected on the pixel surface \mathcal{P}_j during the acquisition interval \mathcal{T}_{t_k, t_a} corresponds to the integral of the convolution measure:

$$\mu_{kj}(\phi_t) := \int_{\mathcal{P}_j \times \mathcal{T}_{t_k, t_a}} (\kappa * \phi_t)(d\mathbf{x} \times dt).$$

The latter operator is linear in the object measure, thus the superposition principle for the object measures (4) induces the linearity of the expected number of photons:

$$\mu_{kj} \left(\sum_{l \in \mathcal{L}} \phi_{t, \mathcal{M}_t^l} \right) = \sum_{l \in \mathcal{L}} \mu_{kj}(\phi_{t, \mathcal{M}_t^l}). \quad (12)$$

Therefore, to compute the total expected photon count, we can consider each object separately. In the following, we denote the contribution to the expected photon flux coming from the l -th object by $\Phi_t^l := \Phi(\phi_{\mathcal{M}_t^l})$.

b) *Background and curvilinear objects:* Following [32], we normalise the PSF kernel to a probability. Therefore, the contribution from the background object described by (5) to the total expected photon flux writes as:

$$\Phi_t^{\text{bg}}(\mathbf{d}\mathbf{x} \times dt) = \varphi_t^{\text{bg}} d\mathbf{x} dt =: \Phi^{\text{bg}}(t) d\mathbf{x} dt. \quad (13)$$

For the curvilinear object (6), the photon flux measure has a density with respect to the spatio-temporal Lebesgue measure: $\Phi_t^c(\mathbf{d}\mathbf{x} \times dt) := \Phi^c(\mathbf{x}, t) d\mathbf{x} dt$. This density involves an integral on the curve \mathcal{C}_t :

$$\Phi^c(\mathbf{x}, t) = \int_{\mathcal{C}_t} \varphi_t(\mathbf{y}) \kappa(\mathbf{x} - \mathbf{y}) d\mathbf{y}.$$

The latter formula requires potentially integrating over the three dimensional imaging volume. However, using the bijective parameterisation introduced in Section III-A3, we pull back the integration into the parameter space, thus reducing the dimensionality of the integral:

$$\Phi^c(\mathbf{x}, t) = \int_{\mathcal{D}_t} \varphi_t(\ell) \kappa(\mathbf{x} - \boldsymbol{\sigma}_t(\ell)) g_t(\ell) d\ell, \quad (14)$$

where g_t is the Riemannian metric induced by $\boldsymbol{\sigma}_t$ and defined as the Euclidean norm of the derivative along the curve of the parameterisation: $g_t := |d/d\ell \boldsymbol{\sigma}_t|_2$. The total expected photon flux sampled by the pixel array is a spatio-temporal measure with density: $\Phi(\mathbf{x}, t) := \Phi^{\text{bg}}(t) + \Phi^c(\mathbf{x}, t)$.

c) *Discretised object-to-pixel mapping:* In order to discretise the expected photon flux, two levels of integration need to be approximated: the object-level integration and the sampling integration on the pixel array. The former amounts to the virtual source approximation introduced in [17]. It involves approximating an integral in parameter space and the integration during the acquisition interval (see [17]). The convolution integral for a curvilinear objects requires a one dimensional quadrature:

$$\Phi(\mathbf{x}, t) \approx \varphi_t^{\text{bg}} + \sum_{s \in \mathcal{VS}_t} w_{t,s}(\ell_s) \varphi_t(\ell_s) \kappa(\mathbf{x} - \boldsymbol{\sigma}_t(\ell_s)), \quad (15)$$

where \mathcal{VS}_t is the set of n_t^v virtual point source indices, and the weight function is the product of the Riemannian metric accounting for the geometry and the quadrature weight, denoted $w_{t,s}^q$: $w_{t,s}^q := g_t(\ell_s) w_{t,s}$. The *virtual source approximation* amounts to approximating the integral of the expected photon flux Φ during the acquisition time interval. We choose a *right continuous with left limits* piecewise constant approximation in time for the quadrature. This means that any function in time is approximated by its value at the beginning of each integration interval:

$$\Phi(\mathbf{d}\mathbf{x} \times \mathcal{T}_{t_k, t_a}) = \int_{t_k}^{t_k + t_a} \Phi(\mathbf{x}, t) dt d\mathbf{x} \approx t_a \Phi(\mathbf{x}, t_k) d\mathbf{x}.$$

To approximate the expected photon count at pixel \mathcal{P}_j , we use a midpoint rule to integrate the previous equation:

$$\mu_{kj} \approx |\mathcal{P}| t_a \Phi(\mathbf{x}_j, t_k) =: \bar{\Phi}(\mathbf{x}_j, t_k), \quad (16)$$

where \mathbf{x}_j is the centre of j -th pixel, and the bar over a quantity q denotes integration in space and time, i.e. $\bar{q} := |\mathcal{P}| t_a q$. Using (15), we obtain the following approximation of the expected photon flux:

$$\bar{\Phi}(\mathbf{x}_j, t_k) = \bar{\varphi}_k^{\text{bg}} + \sum_{s \in \mathcal{VS}_k} \bar{w}_{k,s}(\ell_s) \varphi_k(\ell_s) \kappa(\mathbf{x}_j - \boldsymbol{\sigma}_k(\ell_s)).$$

To obtain the final *digital approximation* of the expected photon count, denoted $\mu_k \in \mathbb{R}_+^{n^p}$, as a function of the fluorescence signal along a curvilinear object, we stack (16) into a vector storing the expected photon count of the n^p pixels registered at time t_k :

$$\mu_k(\varphi_k) = \bar{\varphi}_k^{\text{bg}} \mathbf{1}_{n^p} + \sum_{s \in \mathcal{VS}_k} \bar{w}_{k,s}(\ell_s) \varphi_k(\ell_s) \boldsymbol{\kappa}_k(\ell_s),$$

where $\mathbf{1}_{n^p} \in \{1\}^{n^p}$ is the vector of ones, $\boldsymbol{\kappa}_k(\ell_s) \in \mathbb{R}_+^{n^p}$ is the vector with j -th element defined as $\kappa(\mathbf{x}_j - \boldsymbol{\sigma}_k(\ell_s))$. Inserting (8) and reordering the terms, we obtain:

$$\mu_k(\varphi_k) = \bar{\varphi}_k^{\text{bg}} \mathbf{1}_{n^p} + \sum_{s \in \mathcal{VS}_k} \boldsymbol{\kappa}_k(\ell_s) \bar{w}_{k,s}(\ell_s) \boldsymbol{\beta}_k^T(\ell_s) \varphi_k.$$

Introducing the convolution matrix:

$$\mathbf{K}_k := [\boldsymbol{\kappa}_k(\ell_1) \cdots \boldsymbol{\kappa}_k(\ell_s) \cdots \boldsymbol{\kappa}_k(\ell_{n_k^v})] \in \mathbb{R}_+^{n^p \times n_k^v},$$

the basis matrix:

$$\mathbf{B}_k := [\boldsymbol{\beta}_k(\ell_1) \cdots \boldsymbol{\beta}_k(\ell_s) \cdots \boldsymbol{\beta}_k(\ell_{n_k^v})] \in \mathbb{R}_+^{n_k^b \times n_k^v},$$

and stacking the integrated weights into a vector, denoted $\bar{\mathbf{w}}_t$, we finally write the digital expected photon count vector at time t as:

$$\mu_k(\varphi_k) = \bar{\varphi}_k^{\text{bg}} \mathbf{1}_{n^p} + \mathbf{K}_k \text{ diag}(\bar{\mathbf{w}}_k) \mathbf{B}_k^T \varphi_k. \quad (17)$$

The linear part of the latter equation represents a mapping between the intensity vector along the curvilinear object in kymospace and the expected photon count vector parameterising the Poissonian photon statistics: $\mathbf{M}_k := \mathbf{K}_k \text{ diag}(\bar{\mathbf{w}}_k) \mathbf{B}_k^T$.

2) *Pixel-to-image mapping:* The pixel-to-image mapping models the conversion of photons hitting the camera into grey values. This mapping, denoted ν , is camera-specific. In this work, we assume an electron multiplication charge-coupled device (EM CCD). For the sake of simplicity, we model this mapping deterministically:

$$N_{kj}^{\text{grey}} = q_\lambda M f^{-1} N_{kj}^{\text{photon}} + b =: \nu(N_{kj}^{\text{photon}}), \quad (18)$$

where q_λ is the quantum efficiency, M is the multiplication gain, f is the analog-to-digital proportionality factor, and b is the camera bias.

IV. INVERSE PROBLEM

In this work, we focus on estimating the distribution of light sources in kymospace. Therefore, we assume that the dynamics of the geometry of a curvilinear object and of the background intensity are estimated beforehand.

We formulate the inverse problem to reconstruct the kymograph along a curvilinear object from a sequence of n^f images. At each time point $t_k := (k-1) t_a$ for $k \in \mathcal{F} := \{1, \dots, n^f\}$,

the signal is collected at n^s focal planes resulting in a stack of images, denoted $\mathbf{Z}_k \in \mathbb{Z}^{n^s \times n^h \times n^w}$, and $n^h \times n^w$ is the number of pixels in camera array. We make the standard assumption of neglecting the objects dynamics during the acquisition of a stack, thus the number of pixels in (17) is $n^p := n^s n^h n^w$. The kymograph space and its embedding are defined by the estimated geometry of the curvilinear object: $\mathcal{K}_{\mathcal{F}} = \{\widehat{\mathcal{D}}_k\}_{k \in \mathcal{F}}$ and $\boldsymbol{\sigma}_{\mathcal{F}} = \{\widehat{\boldsymbol{\sigma}}_k\}_{k \in \mathcal{F}}$, respectively.

A. Nearest neighbour (NN) kymograph

A straightforward and widespread kymograph estimate is to sample the grey values along the estimated curvilinear geometry. Each point (t_k, ℓ_s) in kymospace is embedded in physical space by $\widehat{\boldsymbol{\sigma}}_k$, and then the nearest pixel centre index is selected: $j_{ks} := \arg \min_{j \in \mathcal{J}} \|\widehat{\boldsymbol{\sigma}}_k(\ell_s) - \mathbf{x}_j\|_2$, where $\mathcal{J} := \{1, \dots, n^p\}$ is the set of all pixels in the image stack. The *nearest neighbour kymograph* assigns the signal intensity to the grey value of the selected pixel:

$$\widehat{\mathcal{K}}^{\text{NN}} := \left\{ (t_k, \ell_s, \mathbf{Z}_k[j_{ks}]) : t_k \in \mathcal{F}, \ell_s \in \widehat{\mathcal{D}}_k \right\}.$$

The nearest neighbour estimate suffers from all the distortions and degradation affecting the image space, *e.g.* sampling resolution limit, blurring, photobleaching, measurement noise.

To overcome this issue, we use the forward problem model developed in Section III to derive a maximum *a posteriori* (MAP) formulation of the kymograph estimation problem.

B. Sub-pixel resolution kymograph: a MAP formulation

We aim at reconstructing the kymograph that represents the photometry dynamics along an estimated curvilinear geometry:

$$\widehat{\mathcal{K}}^{\text{MAP}} := \left\{ (t_k, \ell_s, \widehat{\boldsymbol{\varphi}}_{k,s}^{\text{MAP}}) : k \in \mathcal{F}, \ell_s \in \widehat{\mathcal{D}}_k \right\},$$

where from (8): $\widehat{\boldsymbol{\varphi}}_{k,s}^{\text{MAP}} := \boldsymbol{\beta}_k^T(\ell_s) \widehat{\boldsymbol{\varphi}}_k^{\text{MAP}}$. The vector of digital intensities $\widehat{\boldsymbol{\varphi}}_k^{\text{MAP}} \in \mathbb{R}_{+}^{n_k^b}$ is estimated as the solution of the MAP problem associated to the forward problem described in Section III.

Following [28], the MAP associated to the Lévy innovation process U and the whitening operator $\mathbf{L}_k \in \mathbb{R}^{n_k^b \times n_k^b}$ writes as the following minimisation problem:

$$\widehat{\boldsymbol{\varphi}}_k^{\text{MAP}} = \arg \min_{\boldsymbol{\varphi} \in \mathbb{R}_{+}^{n_k^b}} \text{nll}(\boldsymbol{\varphi} | \mathbf{Z}_k) + \varpi_U(\mathbf{L}_k \boldsymbol{\varphi}), \quad (19)$$

where $\text{nll}(\boldsymbol{\varphi} | \mathbf{Z}_k) = -\log \mathbb{P}(\mathbf{Z}_k | \boldsymbol{\varphi})$ is the negative log-likelihood function, and ϖ_U is the potential function encoding the prior assumption about the fluorescence signal described in (11).

The negative log-likelihood derives from the Poissonian assumption and the linear and deterministic pixel-to-image mapping (see Section III-B1 and Section III-B2). Indeed, for a given image \mathbf{Z}_k , the invertible pixel-to-image mapping ν , allows converting grey values into photon counts that we

stack into a vector $\mathbf{n}_k := \text{vec}(\nu^{-1}(\mathbf{Z}_k))$. The negative log-likelihood writes as $\text{nll}(\boldsymbol{\varphi} | \mathbf{Z}_k) = \text{nll}(\boldsymbol{\varphi} | \mathbf{n}_k)$, where:

$$\begin{aligned} \text{nll}(\boldsymbol{\varphi} | \mathbf{n}_k) &= \left\langle \mathbf{1}_{n^p}, \mathbf{n}_k \log \frac{\mathbf{n}_k}{\boldsymbol{\mu}_k(\boldsymbol{\varphi})} + \boldsymbol{\mu}_k(\boldsymbol{\varphi}) - \mathbf{n}_k \right\rangle \\ &=: n^p \text{nll}_p\left(\overline{\boldsymbol{\varphi}}_k^{\text{bg}} \mathbf{1}_{n^p} + \mathbf{M}_k \boldsymbol{\varphi}\right). \end{aligned} \quad (20)$$

In the latter equation we define the normalised likelihood nll_p , where the number of pixels n^p is used as a normalisation to make nll values comparable between different data. Similarly, we normalise the potential function to account for a varying object length, *i.e.* $\varpi_U =: n_k^b \varpi_{k,U}$.

In order to constrain the fluorescence signal to positive values, we introduce the indicator function of a set $\mathcal{S}_k := \mathbb{R}_{+}^{n_k^b}$, denoted $\iota_{\mathcal{S}_k}(\boldsymbol{\varphi})$, as the function assuming 0 if $\boldsymbol{\varphi}$ belongs to \mathcal{S}_k , and $+\infty$ otherwise.

The final MAP optimisation problem writes in normalised form:

$$\begin{aligned} \widehat{\boldsymbol{\varphi}}_k^{\text{MAP}} &:= \arg \min_{\boldsymbol{\varphi}} \text{nll}_p\left(\overline{\boldsymbol{\varphi}}_k^{\text{bg}} \mathbf{1}_{n^p} + \mathbf{M}_k \boldsymbol{\varphi}\right) + \\ &\quad \varpi_{k,U}(\mathbf{L}_k \boldsymbol{\varphi}) + \iota_{\mathcal{S}_k}(\boldsymbol{\varphi}) \quad (21) \\ &=: \arg \min_{\boldsymbol{\varphi}} \mathcal{E}_k^{\text{MAP}}\left(\overline{\boldsymbol{\varphi}}_k^{\text{bg}} \mathbf{1}_{\text{bg}} + \mathbf{O}_k \boldsymbol{\varphi}\right), \end{aligned}$$

where $\mathbf{1}_{\text{bg}}$ is the vector selecting the background components, implemented by stacking the following vector of ones $\mathbf{1}_{n^p}$ with the vector of zeros $\mathbf{0}_{2n_k^b}$. The latter equation highlights the nature of this optimisation problem as a sum of three functionals in $\boldsymbol{\varphi}$ coupled through their arguments by the following operator: $\mathbf{O}_k := [\mathbf{M}_k^T \mathbf{L}_k^T \mathbf{I}_k]^T$, where $\mathbf{I}_k \in \mathbb{R}^{n_k^b \times n_k^b}$ is the identity matrix. The convexity of optimisation problem (21) depends only the convexity of the innovation potential $\varpi_{k,U}$ (see Table I in [28]).

C. Fully-split formulation of the algorithm

In order to exploit the additive structure of the optimisation problem (21), we use an operator splitting strategy based on the alternating split Bregman (ASB) algorithm (*e.g.* see [33] when ϖ_U is the ℓ_1 norm). We use a fully-decoupled strategy as introduced in [34] and advocated in [35]. The strategy involves three steps. The first step is to write (21) as a sum of two functionals:

$$\widehat{\boldsymbol{\varphi}}_k^{\text{MAP}} = \arg \min_{\boldsymbol{\varphi}} \left\langle \mathbf{0}_{n_k^b}, \boldsymbol{\varphi} \right\rangle + \mathcal{E}_k^{\text{MAP}}(\overline{\boldsymbol{\varphi}}_k^{\text{bg}} \mathbf{1}_{\text{bg}} + \mathbf{O}_k \boldsymbol{\varphi}).$$

The second step introduces a set of decoupling variables:

$$\mathbf{w} = [\mathbf{w}_1^T \mathbf{w}_2^T \mathbf{w}_3^T]^T := \overline{\boldsymbol{\varphi}}_k^{\text{bg}} \mathbf{1}_{\text{bg}} + \mathbf{O}_k \boldsymbol{\varphi}, \quad (27)$$

and a Bregman proximal point algorithm to enforce the constraint:

$$\begin{aligned} (\mathbf{w}^{i+1}, \boldsymbol{\varphi}^{i+1}) &= \arg \min_{\mathbf{w}, \boldsymbol{\varphi}} \left\langle \mathbf{0}_{n_k^b}, \boldsymbol{\varphi} \right\rangle + \\ &\quad \mathcal{E}_k^{\text{MAP}}(\mathbf{w}) + \frac{1}{2\gamma} \left\| \mathbf{b}^i + \overline{\boldsymbol{\varphi}}_k^{\text{bg}} \mathbf{1}_{\text{bg}} + \mathbf{O}_k \boldsymbol{\varphi} - \mathbf{w} \right\|_2^2 \\ \mathbf{b}^{i+1} &= \mathbf{b}^i + \overline{\boldsymbol{\varphi}}_k^{\text{bg}} \mathbf{1}_{\text{bg}} + \mathbf{O}_k \boldsymbol{\varphi}^{i+1} - \mathbf{w}^{i+1}. \end{aligned}$$

The third step amounts to solving the latter optimisation problem with an alternating minimisation strategy, resulting in Algorithm 1.

Algorithm 1: Fully-split ASB for solving (21)

Input : $Z_k, C_k, \bar{\varphi}_k^{\text{bg}}, \gamma$

Output: $\hat{\varphi}_k^{\text{MAP}} := w_3^\infty$

$$w_1^0 = n_k, \quad w_2^0 = w_3^0 = \mathbf{0}_{n_k^b}, \quad b^0 = \mathbf{0}_{n_p + 2n_k^b}$$

while NOT CONVERGED **do**

Least squares sub-problem:

$$\varphi^{i+1} = \arg \min_{\varphi} \left\| b^i + \bar{\varphi}_k^{\text{bg}} \mathbf{1}_{\text{bg}} + O_k \varphi - w^i \right\|_2^2 \quad (22)$$

n_{ll_p} sub-problem:

$$w_1^{i+1} = \text{Prox}_{\gamma n_{\text{ll}_p}}(b_1^i + \mu_k(\varphi^{i+1})) \quad (23)$$

Innovation potential ($\varpi_{k,U}$) sub-problem:

$$w_2^{i+1} = \text{Prox}_{\gamma \varpi_{k,U}}(b_2^i + L_k \varphi^{i+1}) \quad (24)$$

Positivity constraint sub-problem:

$$w_3^{i+1} = \text{Proj}_{S_k}(b_3^i + \varphi^{i+1}) \quad (25)$$

Bregman update (dual gradient ascent):

$$b^{i+1} = b^i + \bar{\varphi}_k^{\text{bg}} \mathbf{1}_{\text{bg}} + O_k \varphi^{i+1} - w_3^{i+1} \quad (26)$$

end

To ensure that Algorithm 1 yields a positive estimate for the fluorescence intensity $\hat{\varphi}_k^{\text{MAP}}$, we output the result of the projection step at convergence, denoted w_3^∞ .

This strategy results in a least-squares problem, a proximal evaluation problem, and a Bregman linear update. The proximal operator of a function g is defined as:

$$\text{Prox}_g(x) := \arg \min_y g(y) + \frac{1}{2} \|x - y\|_2^2.$$

This generalises the notion of projection, as for the indicator of a set, it corresponds to a projection on this set: $\text{Prox}_{\ell_S}(x) = \text{Proj}_S(x)$.

1) *Least squares sub-problem* (22): Solving the least-squares sub-problem amounts to solving the following normal equations:

$$O_k^T O_k \varphi^{i+1} = O_k^T \left(b^i - \left(b^i + \bar{\varphi}_k^{\text{bg}} \mathbf{1}_{\text{bg}} \right) \right).$$

In digital image processing, the geometry of the pixel grid allows solving efficiently the normal equations using a spectral method (e.g. applying discrete cosine transform for Neumann boundary conditions, see e.g. [34]). However, the virtual sources are not aligned with the pixel grid, preventing the use of spectral methods. Nevertheless, following [17], it is possible to efficiently compute the operator M_k using the improved fast Gaussian transform (IFGT, see [36]). For computational efficiency, the matrix inverse:

$$(O_k^T O_k)^{-1} = (M_k^T M_k + L_k^T L_k + I_k)^{-1} \in \mathbb{R}^{n_k^b \times n_k^b},$$

is pre-computed outside the main iteration loop.

2) *Likelihood sub-problem* (23): The solution of the second sub-problem requires solving the following quadratic equation for each pixel, independently (see e.g. [34], [35]):

$$(w_{1j}^{i+1})^2 + w_{1j}^{i+1} \left(\frac{\gamma}{n_p} - (b_{1j}^i + \mu_{kj}(\varphi^{i+1})) \right) - \frac{\gamma}{n_p} n_{kj} = 0.$$

This quadratic equation has two roots (positive discriminant) and always admits a positive one because the product of its roots is negative (i.e. $-\frac{\gamma}{n_p} n_{kj} < 0$). The admissible solution w_{1j}^{i+1} corresponds to the positive root because this dummy variable is related to the expected photon count (17) via the constraint (27).

3) *Innovation potential sub-problem* (24): The solution depends on the innovation process prior. For example, if the innovation process is driven by the Laplace distribution, $\varpi_U(w_2) = \eta \|w_2\|_1$, and the proximal operator is a component-wise soft thresholding then:

$$w_{2p}^{i+1} = \text{shrink} \left(b_{2p}^i + (L_k \varphi^{i+1})[p], \frac{\gamma \eta}{n_k^b} \right),$$

where $\text{shrink}(w, c) := \max(|w| - c, 0) \text{sign}(w)$. If the innovation process is driven by the Gaussian distribution, $\varpi_U(w_2) = \eta \|w_2\|_2$, and the shrinkage operator is given by $\text{shrink}(w, c) := w/(1+2c)$, see [37]. Both potentials introduce a regularisation parameter, denoted η , that controls the tradeoff between the data fidelity term n_{ll_p} and the belief in the prior about the underlying innovation process $\varpi_{k,U}$.

4) *Positivity constrain sub-problem* (25): This sub-problem decouples and is computed as the projection onto \mathbb{R}_+ :

$$w_3^{i+1} = \max(\mathbf{0}, b_3^i + \varphi^{i+1}),$$

where max is understood component-wise.

V. EXPERIMENTS

In this section, we apply our framework to reconstruct the spatio-temporal distribution of light sources along curvilinear biological structures such as microtubules (MTs). We start with discussing the difficulties associated to estimating kymographs in a challenging biological problem: studying microtubule dynamics *in vivo*. Then, we use the virtual microscope framework introduced in [17] to compute synthetic image data for different geometry and photometry scenarios. We use these data to demonstrate the capabilities of our MAP framework and compare it with the solution of two other inverse problems: the nearest neighbour signal reconstruction and the fitting of a known parametric signal model. Finally, we apply our framework to a real-world data set consisting of time-lapse images of pre-anaphase microtubule dynamics in a model system: the budding yeast *S. cerevisiae*. We show that our framework allows reconstructing details that are not accessible to conventional digital techniques. Given the estimated sub-pixel resolution kymograph as ground truth objects, we use the virtual microscope framework to imitate realistically the real data and investigate the robustness of our framework to noise, sampling and regularisation.

A. MT dynamics *in vivo*: the need for better kymographs

Microtubules are highly dynamic polar filaments that are crucial for the cell viability [38]. Their main feature is a stochastic transition between growth and shrinkage that results in tightly regulated outcomes: microtubules are essential for the intra-cellular organisation, transport, and cell division requiring microtubules to be correctly positioned in space

and time. It has been revealed that the tight regulation of microtubules is achieved by numerous microtubule associated proteins, including plus-end-tracking proteins that accumulate at the microtubule end exposed to the cytoplasm [39]. However, many regulatory mechanisms of microtubule dynamics remain largely unknown, making this topic a fundamental problem in cell biology [12], [40].

Kymographs are considered to be a *de facto* standard data representation to study microtubule dynamics *in vivo* and *in vitro*. They are used to infer parameters quantifying the dynamics at the microtubule ends [12], [13], [41], [7]. Moreover, they are used to reconstruct trajectories of motor proteins, cargos and vesicles moving along the microtubule lattice [3], [42], [10], estimate parameters quantifying their dynamics [7], [8], [9], and study the underlying regulatory mechanisms [3], [2]. Indeed, the analysis of kymographs provides a wealth of quantitative information such as the orientation, the velocity, and the pausing times of the transported particles [4], the co-localization of motor proteins and microtubule associated proteins [2], [5].

This paradigm relies on the hypothesis that the signal displayed in the kymograph reflects the spatio-temporal dynamics of the fluorescently labeled structures. However, as we described in Section III, many distortions impair the relationship between the labeled objects and the image data.

Our Bayesian framework accounts for the distortions of the fluorescence signal explicitly in the modelling of the forward problem, and uses a general class of signal priors to cope with these imaging limitations. We start by assessing the capability of our framework on synthetic data generated by the virtual microscope framework.

B. Reconstructing light source distributions for different geometry and photometry scenarios

We use the virtual microscope framework [17] to demonstrate the capabilities of the proposed framework to reconstruct the fluorescence signal along curvilinear objects in different geometric and photometric scenarios. We generate a synthetic data set to compare three inverse problems: the nearest neighbour estimate introduced in Section IV-A, a maximum likelihood (ML) estimate of a parametric model based on the image formation model developed in Section III-B, and our MAP estimate defined in Section IV-B. In what follows, we use models and examples inspired from the microtubule dynamics literature, but the insight applies to any curvilinear objects and is therefore generic.

1) *Synthetic examples from microtubule models:* We consider two types of geometry models (GM), and two types of photometry models (PM): straight and curved objects (GM1 and GM2, respectively), where light sources are distributed smoothly (PM1, the comet shape model) or localised in specific regions (PM2, the islands model).

a) *Straight microtubule (GM1):* We model a straight microtubule using a line segment. The mapping from the parameter space to the physical space is defined as follows: $\sigma_t(\ell) := \mathbf{x}_o^{\text{ps}} + \mathbf{d}_t \ell$, where $\mathbf{x}_o^{\text{ps}} \in \mathbb{R}^3$ is the origin of the microtubule, and $\mathbf{d}_t \in \mathbb{R}^3$ is the unit vector defining its

direction. In this model, the curve is arc-length parameterised (*i.e.* $g_t = 1$), and the weighing of the virtual sources depends only on the quadrature weights.

b) *Curved microtubule (GM2):* We model a curved microtubule using a quadratic spline function defined by four control points sampled uniformly in the imaging volume.

c) *Smooth fluorescence signal (PM1, the comet shape model):* This photometry model uses an intensity distribution similar to the one estimated in [2]. In this study, it has been shown that the protein Mal3 accumulates at the growing end of the microtubule forming a comet shape. The observed fluorescence signal was modelled using a superposition of a Gaussian and an exponential function described by a set of six parameters, denoted $\Theta^{\text{pm1}} := \{\mu, \sigma, a, b, c, d\} \in \Theta^{\text{pm1}} \subset \mathbb{R}_+^6$: $\varphi_{\theta}^{\text{pm1}}(\ell)$ is defined as $\frac{a}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(\ell-\mu)^2}$, if $\ell \geq \mu - b$, and $\left(\frac{a}{\sqrt{2\pi\sigma^2}} e^{-\frac{b^2}{2\sigma^2}} - d\right) e^{\frac{1}{c}(\ell-(\mu-b))} + d$, else. In [2], this model is fitted on the average of multiple and aligned microtubule images. However, the precise estimated parameters were not reported. We set the following values to visually reproduce the shape of the photometry density: $\mu = 4$, $\sigma = 0.3$, $a = 300$, $b = 0.2$, $c = 0.6$, and $d = 20$.

d) *Localised fluorescence model (PM2, the island model):* This model assumes a piecewise constant distribution of the light sources along the curvilinear object. This choice is motivated by the fact that some microtubule associated proteins accumulate on specific locations on the microtubule and appear as blurred spots (*e.g.* [43]). We define this distribution with a set of seven parameters, denoted $\Theta^{\text{pm2}} := \{\varphi_1, \varphi_2, \varphi_3, a, b, c, d\} \in \Theta^{\text{pm2}} \subset \mathbb{R}_+^7$: $\varphi_{\theta}^{\text{pm2}}(\ell)$ is defined as φ_1 on $[0, a]$, φ_2 on $[b, c]$, φ_3 on $[d, \Lambda_m[t]]$, and 0 else. We fix the model parameters to the following values: $\varphi_1 = 600$, $\varphi_2 = 400$, $\varphi_3 = 800$, $a = 1.0$, $b = 2.5$, $c = 3.0$, $d = 4.5$.

2) *Image formation model:* Using the forward model introduced in Section III, we compute synthetic image data for all combinations of geometry and photometry models introduced above (Fig. 3). Each image stack consists of 21 image slices of size 25×42 pixels, acquired at different focal plans. We use a widespread approximation of the microscope PSF using a Gaussian probability density function [44], [26]: $\kappa^{\text{sg}}(\mathbf{x}) = C^{\text{sg}} e^{-\frac{1}{2}\mathbf{x}^T \Sigma^{-1} \mathbf{x}}$, where $C^{\text{sg}} = (8\pi^3 \sigma_{xy}^4 \sigma_z^2)^{-0.5}$ is the normalisation constant, $\Sigma := \text{diag}(\sigma_{xy}^2, \sigma_{xy}^2, \sigma_z^2) \in \mathbb{R}^{3 \times 3}$ is the covariance matrix, σ_{xy} and σ_z are the standard deviations in lateral and axial directions, respectively. The parameter values of the image formation model are calibrated to the experimental setup described in Section V-C2 and summarised in Table I.

3) *Comparison of the inverse problems:* Given the synthetic image data, we compare the fluorescent signal reconstructed along the ground truth object geometry using three inverse problems.

a) *Nearest neighbour estimate:* We apply the nearest neighbour (NN) estimate described in Section IV-A. We observe in Fig. 3 that the reconstruction is corrupted by bias, noise, and blur. The bias is apparent in the low-intensity regions in the ground truth signal, where in the NN estimate the background level remains. The effect of blurring translates in the NN reconstruction in unsharp boundaries for the piecewise-

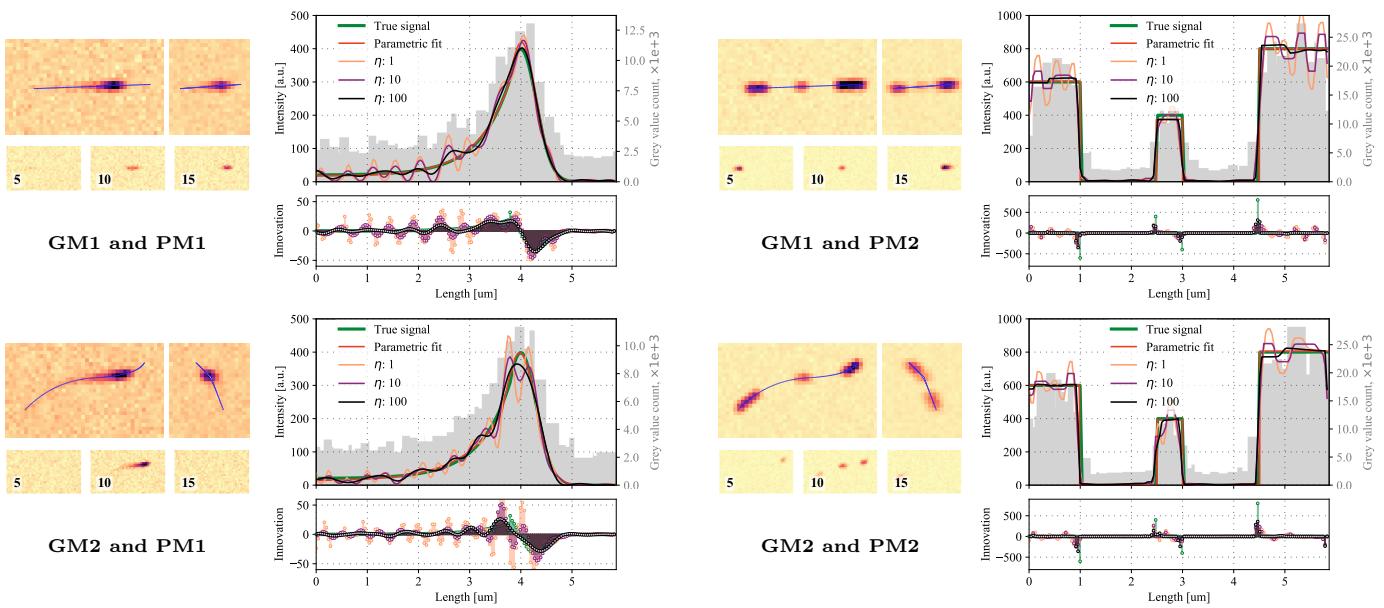


Figure 3. Reconstructing light source distributions for different geometry and photometry scenarios. We show each combination of the models in four panels. *Left, upper part:* orthogonal z - and x -mean projections of the synthetic image stack with the microtubule lattice overlaid. *Left, lower part:* three slices acquired at different focal distances. *Right, upper plot:* nearest neighbour estimate (scale on the right axis), parametric ML estimate and non-parametric MAP for three regularisation parameters (for both, scale on the left axis). *Right, lower plot:* ground truth and estimated innovation process. The resolution in reconstruction space is $0.04 \mu\text{m}$.

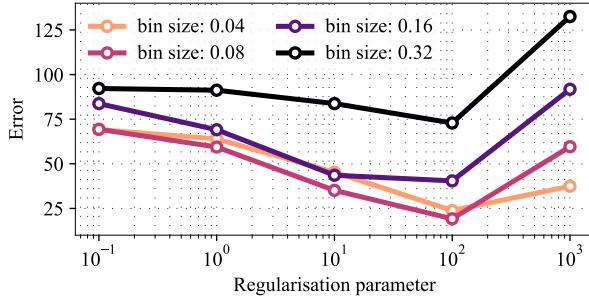


Figure 4. Influence of the bin size and the regularisation parameter on the reconstruction accuracy. The results are shown for the straight microtubule (GM1) and the localised fluorescence model (PM2). Similar trends are observed for the other combinations of geometry and photometry models.

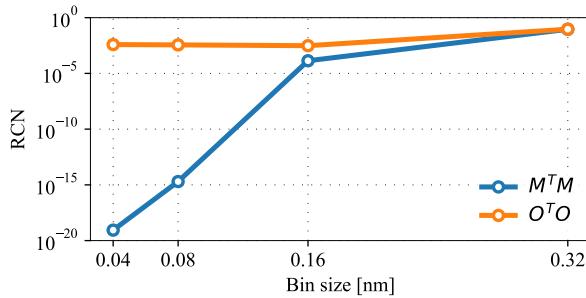


Figure 5. Influence of the bin size on the reciprocal condition number. The results are shown for the straight microtubule (GM1) and the localised fluorescence model (PM2). Similar trends are observed for the other combinations of geometry and photometry models.

constant photometry model PM2. The spatial resolution limit of the pixel size, and the Poisson noise are also prominent

Table I
PARAMETERS OF THE IMAGE FORMATION PROCESS.

Parameter	Notation	Value
Acquisition time	t_a	26.0 ms
Pixel size	—	160 nm
PSF lateral standard deviation	σ_{xy}	130 nm
PSF axial standard deviation	σ_z	255 nm
Quantum efficiency	q_λ	0.7
Multiplication gain	M	1200
ADU conversion	f	6.44
Camera bias	b	1839

in the reconstruction. Aligning and averaging multiple image stacks as in [2] would alleviate the latter issues, but not the former ones. Moreover, the NN signal is sampled directly from grey values and cannot be easily compared with the underlying ground truth intensity signal. In Fig. 3, we display a separate axis to provide the scale of the NN estimate. However, in the rest of the paper we omit it and compare only the shape of NN reconstructions with other estimates.

b) Parametric ML estimate: When the practitioner is confident in a given parametric model, one can use the forward model described in Section III to fit the model parameters by minimising the negative log-likelihood, which is equivalent to maximum likelihood. To estimate the parameters of PM1 and PM2, we solve $\hat{\theta} = \arg \min_{\theta} \text{nll}(\varphi_{\theta} | Z_k)$ using the covariance matrix adaptation evolutionary strategy (CMA-ES) [45]. We observe that the reconstructed signals are close to the underlying ground truth (Fig. 3). This improves significantly the NN reconstruction in two ways: by taking into account the image formation model (thus reconstructing the signal

in physical space), and by estimating the fluorescence signal along a microtubule from a single image stack (*i.e.* without averaging many image stacks).

We interpret the nearest neighbour estimate as a worst case scenario. On the other hand, the reconstruction obtained by fitting the parametric model corresponding to the ground truth sets a best case for the reconstruction accuracy. Nonetheless, in real-world applications, the ground truth model is not necessarily available, and one must rely on weaker models, such as non-parametric ones.

c) Non-parametric MAP estimate: The smoothed photometry model PM1 can be accommodated by penalising the squared- ℓ_2 norm of the signal transformed using a linear operator [28]. In the general framework of [27], this corresponds to choosing a Gaussian distribution to model the innovation process: $\mathbb{P}_U^{\text{pm1}}(u) \propto \exp(-\eta u^2)$ and the first-order derivative for the whitening operator.

The piecewise-constant photometry model PM2 is characterised by having sparse first-order derivatives. Therefore, we use, as for PM1, the first-order derivative as the whitening operator, but model the innovation process using a sparsity-inducing distribution. For example, it is possible to capture the sparsity in the innovation process by modelling it with a Student's or Cauchy distribution. However, in this case, the resulting optimisation problem is not convex. To preserve the convexity, we choose the Laplace distribution, that is often used as a sparsity-inducing prior: $\mathbb{P}_U^{\text{pm2}}(u) \propto \exp(-\eta|u|)$.

We set the resolution in reconstruction space to $0.04\text{ }\mu\text{m}$, fix the object geometry and the background intensity to the ground truth, and use a Gauss-Legendre quadrature with 10 points per basis function for approximating the convolution integral (15). We apply Algorithm 1 to reconstruct the fluorescence signal for different regularisation parameters. At a lower regularisation, we observe oscillations in the reconstruction. However, the overall shape (in PM1) and support (in PM2) of the signal are correctly identified. When the regularisation parameter is increased, the amplitude of the oscillations decreases and the reconstructions approach the ground truth.

Using a virtual microscope approach, we have shown that our framework allows reconstructing the fluorescence signal for various scenarios of object geometry and photometry.

d) Robustness of the MAP signal to regularisation and bin size: The accuracy of the reconstructed signal is influenced by the regularisation parameter (Fig. 3) and the bins size. We quantify the reconstruction error as a function of both parameters on Fig. 4. We define the error as the ℓ_1 norm of the difference between ground truth and reconstructed signals. We observe that using higher resolution of the reconstruction space results in lower reconstruction error. However, when the bin size is too small, the error starts increasing again. We also observe an V-shaped trend in the error versus regularisation parameter. Therefore, this suggests selecting an optimal regularisation, *e.g.* using the virtual microscope framework [17], given that a model of the underlying photometry model is available. Remarkably, the value of regularisation parameter resulting in the lowest reconstruction error is the same for different bin sizes.

e) Influence of the bin size on the reciprocal condition number: We use a constant number of virtual point sources (quadrature points) to integrate over each bin. If the bins are smaller, then the virtual sources are located closer to each other and have higher overlap among the PSF supports. As a result (see Result 1), the RCN of the operator $M_k^T M_k$ decreases. However, when we solve the least squares sub-problem (22), we invert the operator $O_k^T O_k = M_k^T M_k + L_k^T L_k + I_k$. Both the regulariser (contributing $L_k^T L_k$) and the fully-splitting strategy (contributing I_k) help regularising the problem. As a result, the conditioning of the least squares problem improves, and gets less sensitive to the bin size Fig. 5.

C. MT dynamics in *S. cerevisiae*

In this section we apply our framework to reconstruct sub-pixel resolution kymographs of microtubule dynamics from time-lapse images acquired *in vivo* in the budding yeast [47].

1) MT dynamics image data: We show the results for the strain Kip3 Δ , imaged in pre-anaphase. This dataset contains 743 cells. The images are characterised by the presence of a single straight cytoplasmic microtubule that is longer than the one observed in wild type yeast cells. The microtubule ends are visualised using two families of proteins labeled with the green fluorescent protein: Spc72p binding specifically to the spindle pole bodies (structures from where the microtubule minus-end originates), and Bik1 accumulating at the microtubule plus-end [47]. We observe that most of the fluorescence signal is located at the microtubule ends (Fig. 6a). However, the protein Bik1 can attach to and be transported along the microtubule lattice [47].

We assume that the distribution of light sources along the microtubule is sparse. Therefore, we model the underlying intensity signal using the first derivative as the whitening operator and assume that the increments of the innovation process follow a Laplace distribution.

2) Imaging parameters: The images are acquired by a spinning disk confocal microscope equipped with a 63X 1.4 NA objective, 493 nm solid-state laser, and EM-CCD camera (Hamamatsu ImageEM). An image stack of size $21 \times 256 \times 256$ pixels is acquired every 0.55 seconds, resulting in a sequence of 100 frames. The regions of interest containing individual cells in pre-anaphase are manually outlined and cropped out. The imaging parameters are summarised in Table I.

3) Reconstructing the background intensity: To estimate the photobleaching of the background intensity, we average the grey values of the pixels in each image stack, convert them into photon counts by applying the inverse of the pixel-to-image mapping ν given in (18), dividing the result by the spatio-temporal integration constant given by $|\mathcal{P}|t_a$ in (16). We fit an exponential decay model (characterised by three parameters: an amplitude, a decay rate and an offset). We minimise the ℓ_2 norm of the residuals using the CMA-ES algorithm [45].

4) Reconstructing the microtubule geometry: To reconstruct the sub-pixel resolution kymograph using our framework, first, we need to reconstruct the microtubule geometry for each frame. Several trackers have been designed to track microtubule ends in fluorescence microscopy image data (*e.g.*

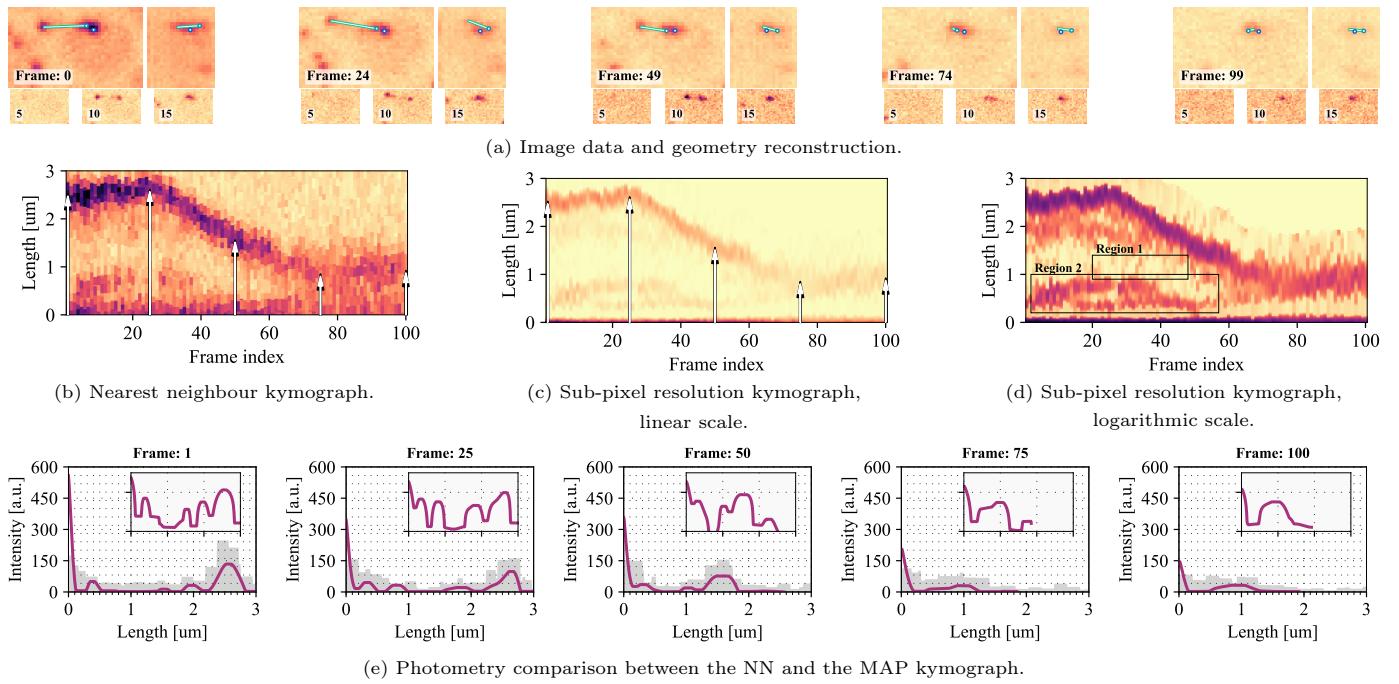


Figure 6. Sub-pixel resolution kymograph reconstruction for a single *S. cerevisiae* cell. (a) *Image data and geometry reconstruction*. Orthogonal z- and x-mean projections and image slices overlaid with the microtubule lattice and spindle pole bodies estimated by the particle filter introduced in [46]. Pixel size: 160 nm. (b) *Nearest neighbour kymograph*. Colour scale proportional to grey values. (c) *Sub-pixel resolution kymograph, linear scale*. Resolution of the reconstruction space: 8 nm. Colour scale proportional to photon counts. In (b) and (c), the microtubule geometry estimated by the particle filter is overlaid as an upward arrow. The plus-end and the minus-end are located at the head and tail respectively. (d) *Sub-pixel resolution kymograph, logarithmic scale*. Logarithmic colour scale shown for a better contrast. We observe patterns of dynamics of small clusters at a high rate (e.g. within *region 1*) and of large clusters at a slower rate (e.g. within *region 2*). (e) *Photometry comparison between the NN and the MAP kymograph*. The photometry along the arrows displayed in (b) and (c) is shown: NN estimate shown as a grey histogram; MAP estimate shown as a purple curve; Logarithmic signal shown in the inset).

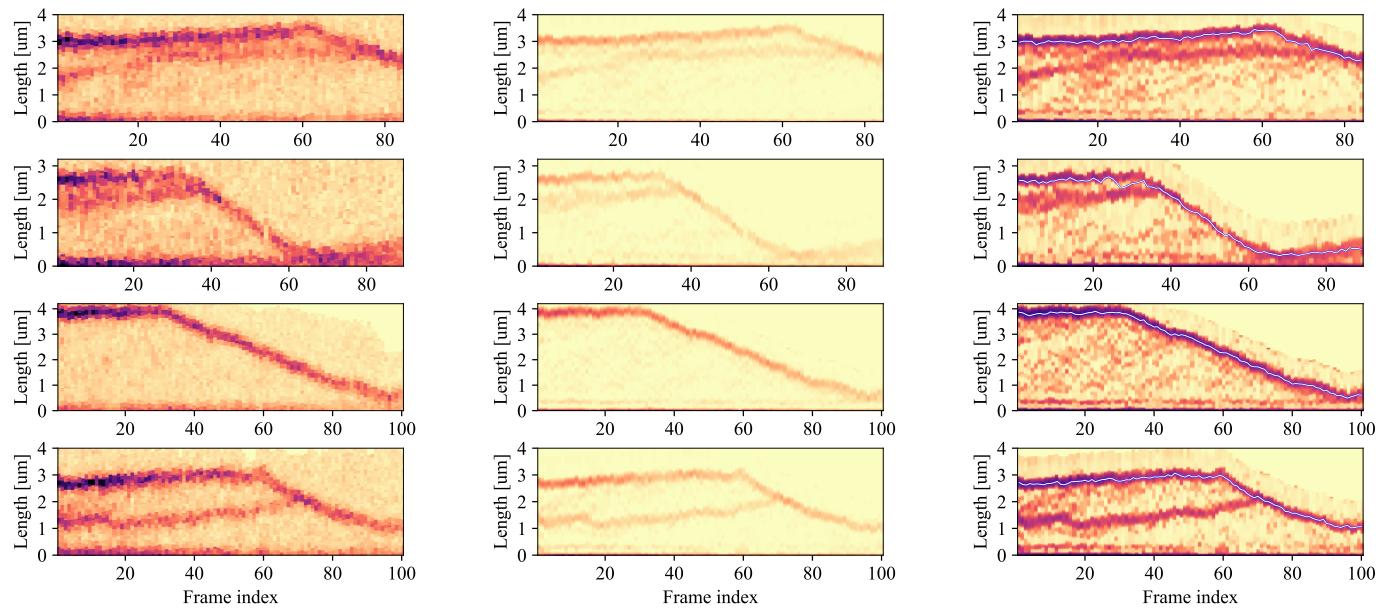


Figure 7. Diversity of the pattern of dynamics in the *S. cerevisiae* *in vivo* kymographs. Each row corresponds to a different *S. cerevisiae* cell. First column: NN kymograph, colour scale proportional to grey values. Second column: MAP kymograph, linear colour scale proportional to photon count. Third column: MAP kymograph, logarithmic colour scale for better contrast. The microtubule length estimated by the particle filter is overlaid on the kymographs as a white curve. We observe that the estimator is biased towards the maximum of the comet-shaped intensity signal at the microtubule plus-end, and it results in an underestimated microtubule length.

[48], [49], [50], [46]). In the following, we use the particle filter introduced in [46]. It is designed to track microtubules in

S. cerevisiae. It relies on the microtubule model that explicitly encodes the microtubule lattice but only accounts for the

intensity signal at its ends. We note that it is a special case of the object model presented in this paper. Modelling the fluorescence signal as a linear combination of two Dirac measures with atoms at the ends of the microtubule lattice, we recover from (6) the microtubule model from [46].

$$\phi_t^{\text{cPS}}(dy \times dt) = \mathbb{1}_{C_t}(y) \sum_{\ell \in \{0, \Lambda_t\}} \varphi_t(\ell) \delta_{\sigma_t(\ell)}(dy) dt. \quad (28)$$

The tracking results for one cell are shown in Fig. 6a. We visualise the spindle pole bodies as two points and the microtubule lattice as a line segment.

5) *Reconstructing sub-pixel resolution kymographs:* We set the reconstruction space resolution to 8 nm. This is comparable with the size of a tubulin dimer, the building block of microtubules [51]. The microtubule length estimated using the particle filter is extended up to 3 μm for the NN reconstruction and by 1 μm for signal reconstruction using our framework. The regularisation parameter is set to $\eta = 1$. Fig. 6 shows the results for a single cell, where we compare the NN reconstruction with our Bayesian framework.

The NN kymograph suffers from a coarser resolution, higher level of noise and a higher impact of blurring (Fig. 6b). The fluorescence signal at the microtubule *minus-end* (bottom of the kymograph) appears to be weaker than at the *plus-end* (top of the kymograph). The sub-pixel resolution kymograph (Fig. 6c), appears qualitatively different. The signal at the minus-end is sharply peaked, and its width is about 200 nm. It is consistent with the fact that the protein Spc72p is used as a minus-end marker. It is known that this protein is anchored in the outer plaque of the spindle pole bodies, which have a width of about 185 nm [47], [52]. The signal at the microtubule plus-end has a lower amplitude than at the minus-end, and it shows a comet-like shape similar to the pattern of the end-binding protein Mal3 observed in *in vitro* in [2]. In addition, our framework allows revealing dynamical patterns that are not observed *in vitro* by averaging multiple microtubules [2], and goes beyond the two-points photometry hypothesis used in [49], [46]. Indeed, we observe patterns along the microtubule lattice (Fig. 6d). This signal is of weaker intensity than at the microtubule ends, but still significant when compared with the patterns beyond the microtubule plus-ends. We can distinguish the dynamics of small clusters at a high rate (*e.g.* within region 1) and of large clusters at a slower rate (*e.g.* within the region 2), displaying merging and splitting events. Most of these events cannot be observed as clearly from the NN reconstruction (Fig. 6e).

In our framework, the distortions, noise and the sampling effects are accounted for explicitly in the forward model, and we do not require averaging multiple image stacks. This is a benefit as the microtubule dynamics imaged *in vivo* displays a wide spectrum of variations. Our framework allows accessing these variations, as shown in Fig. 7.

6) *Validation of the sub-pixel resolution kymograph reconstruction:* In order to assess the quality of the kymograph reconstruction for the *in vivo* yeast dataset, we use a virtual microscope approach [17] to synthesise realistic image data under various conditions. This allows alleviating the problem of lacking a ground truth reference for real data.

We use the sub-pixel reconstruction shown in Fig. 6 as the ground truth. We generate data according to the model described in Section III at different peak signal-to-noise ratio (PSNR, adapted from [26]), defined as:

$$t_a = \frac{\varphi_0^{\text{bg}} + \frac{1}{n^P} \sum_{j \in \mathcal{J}} \Phi^c(\mathbf{x}_j, 0)}{|\mathcal{P}| \max_{j \in \mathcal{J}} (\Phi^c(\mathbf{x}_j, 0))^2} 10^{0.1 \text{PSNR}},$$

where $\Phi^c(\mathbf{x}_j, 0)$ is evaluated using the virtual source representation of (14).

We use this synthetic dataset to investigate two important factors: the regularisation parameter and the spatial resolution of the reconstruction space (see Fig. 8b and Fig. 8c). We observe that at high PSNR the signal is robust to changes in both the regularisation and the resolution. However, a higher resolution of the reconstruction space yields a more accurate reconstruction of the signal at a lower PSNR.

7) *Length bias induced by a simplified photometry model:* In Fig. 7, the microtubule length estimated by the particle filter is displayed on top of the kymographs. We observe that the microtubule length is always underestimated. As discussed in Section V-C4, the tracker models the fluorescence distribution with two point sources located at the microtubule ends. However, as observed in the kymographs, the fluorescence distribution at the microtubule plus-end has a comet-like shape. Thus, the point source approximation of the photometry is biased. As a result, the particle filter estimates the microtubule plus-end at the local maximum of the comet. From the reconstructions we observe that the size of the comet changes in time, and the bias in the microtubule length using this simplified photometry model is variable. Finally, the particle filter can drift to a large cluster near the plus-end. This results in a spurious jump in the microtubule length dynamics (see Fig. 7, second row). Therefore, there is a coupling between geometry and photometry estimation: incorrect assumptions about one of affects the reconstruction of the other.

VI. DISCUSSION

In the existing literature, the degradations entailed by the image formation process are handled either before or after the kymograph reconstruction. One approach is to account for the degradations by denoising and deconvolving the original image as a pre-processing step [3], [13]. However, standard restoration algorithms are implicitly based on models that are not designed for singular objects, such as points or curvilinear objects. To achieve a sub-pixel resolution of the particles identified in kymograph space, [10] proposes to refine their position in physical space by fitting the PSF to the original image data. However, this step is performed *a posteriori* and only for individual particles. Another approach is to post-process the kymograph by applying digital filters in kymograph space [3], [11], [15]. However, this kind of restoration approach is necessarily heuristic. Indeed, the nonlinear embedding of the kymograph space into the physical space renders the physical interpretation of the signal in kymospace difficult. To achieve only denoising, it is custom to average independent image data before the kymograph reconstruction [2]. However, this is meaningful only for images of comparable objects that

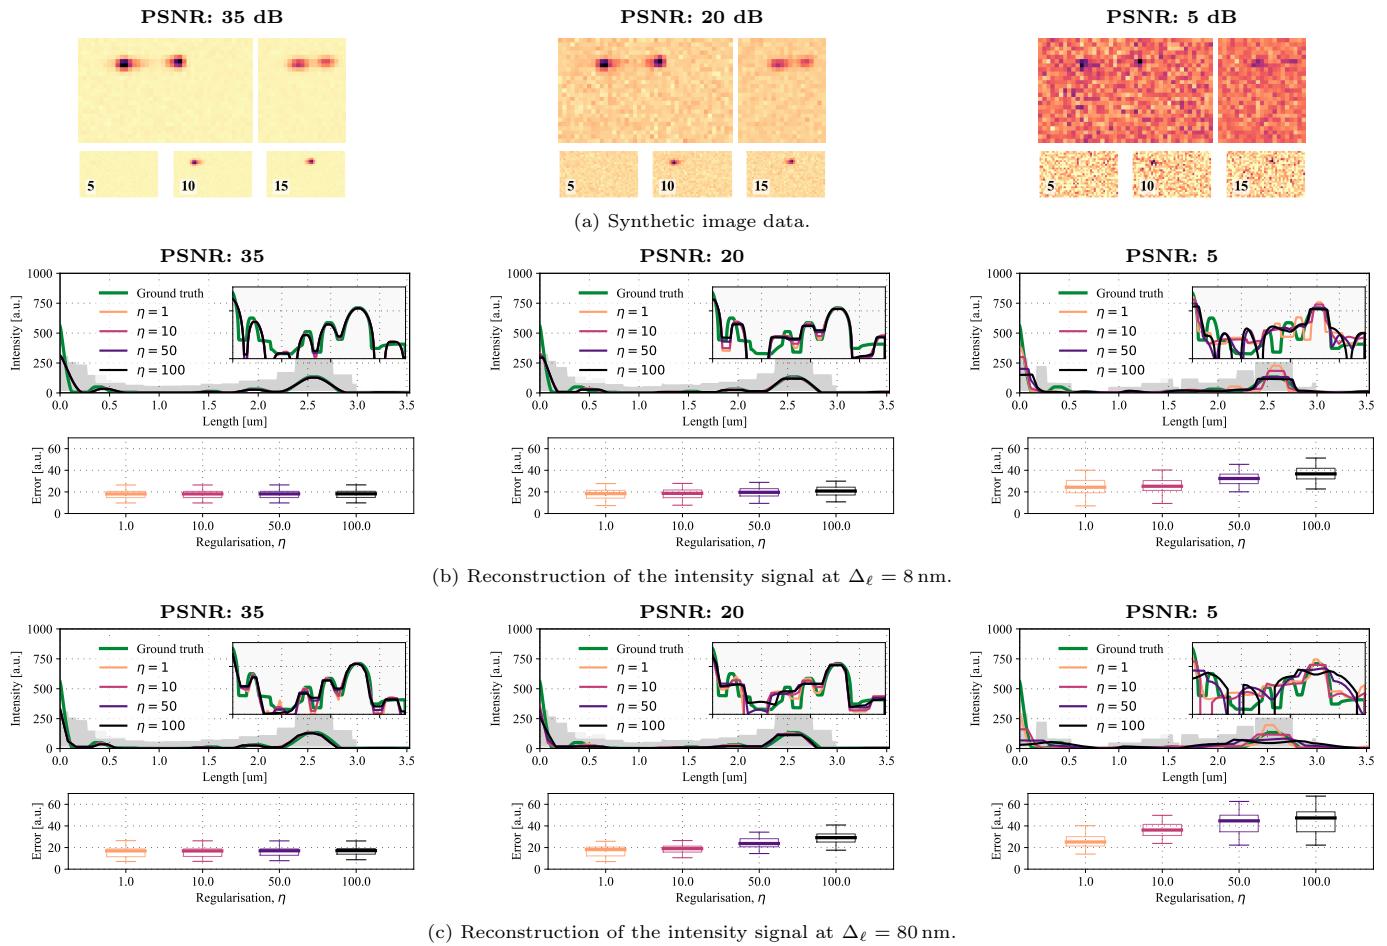


Figure 8. Robustness of the MAP kymograph to PSNR, regularisation (η), and bin size (Δ_ℓ). (a) *Image data.* Orthogonal z- and x- mean projections and image slices computed using the virtual microscope framework at different PSNR. The ground truth is the MAP estimate shown in Fig. 6 (c). *Robustness to PSNR and regularisation for two bin sizes: $\Delta_\ell = 8 \text{ nm}$ (b), and $\Delta_\ell = 80 \text{ nm}$ (c).* Each panel is organised as follows. *First row:* For each bin size, we compare the ground truth profile for the first frame with the NN estimate and the MAP estimate. The photometry reconstruction in a logarithmic scale are shown as an inset. *Second row:* box plot of the per-frame reconstruction error. The reconstruction error is the root mean squared error between the estimated and the ground truth profile.

can be registered. Beyond this fundamental limitation, such a procedure does not allow assessing the inter-object variability, and will only capture the most prominent features of the kymograph.

Our kymograph reconstruction is based on a well-grounded Bayesian framework. However, in this paper we focus on establishing a minimal version of the framework that can handle an analog reconstruction of the kymograph, at the expense of several simplifications in both the forward and inverse problem. For the forward problem, the main extension could be alleviating the over-simplified Gaussian approximation of the PSF model. We have shown that the PSF shape is crucial for reconstructing the photometry of punctual objects, and it would probably also improve the photometry reconstruction of curvilinear objects [46], [53]. In [53], we introduce a framework to model any PSF as a sparse mixture of Gaussians: it provides an accurate representation of any PSF shape, still computationally efficient to evaluate with IFGT, and compatible with the virtual microscope framework. Therefore, such an extension would require only minor modifications to

our kymograph reconstruction framework. Another direction to extend the forward problem relates to the pixel-to-image mapping. In this work, we have used a deterministic affine camera mode, but more advanced models account for the noise process in the pixel-to-image mapping [54], [55]. These camera models can be extended to handle a specific type of camera. For example, pixel-dependent noise models are crucial for the processing of images acquired using a complementary metal-oxide semiconductor (CMOS) camera [56].

For the inverse problem, a straightforward extension is to consider more general sparse stochastic processes [27]. The innovation process modelling involves a whitening operator, and a probability distribution. We restrict ourselves to first order whitening operators. This limitation allows simplifying the inverse problem to a trivial factorisation of the prior probability into its marginals. For higher-order operators, the discrete innovation process in (10) follows a Markov chain with an order depending on the whitening operator order. This entails a more involved factorisation resulting in a more complicated MAP algorithm [31]. Nonetheless, this extension

is solely based on [31] and is straightforward to integrate in the proposed kymograph reconstruction framework. We also restrict ourselves to Gaussian and Laplacian distributions. The family of sparse stochastic processes [27] contains other compatible distributions with even better sparsity-inducing capabilities. In our fully-split formulation, this would amount to changing only the innovation potential sub-problem (24). This is straightforward to do for distributions having a known proximal operator [28]. However, the resulting minimisation problem would become nonconvex, and hence more challenging in practice. This additional complexity should be justified by a given application for which the simpler alternatives are irrelevant.

Finally, we make two assumptions regarding (conditional) independence that would require significantly more efforts to alleviate. First, we treat frames independently: it allows processing each image stack separately, at the expense of simplifying the dynamical model taking place in kymograph space. Going beyond this limitation requires modelling the spatio-temporal dynamics of the fluorescence signal along a curvilinear object. A straightforward approach would be to extend the class of priors to spatio-temporal sparse stochastic processes. We believe that this exciting line of research will benefit from building spatio-temporal priors based on the expertise coming from the related biology, biochemistry and biophysics literature (*e.g.* [57]). Second, we assume the geometry to be known or already estimated. The joint reconstruction of the geometry and photometry is a problem already tackled for punctual objects [46], [53]. However, it is a more challenging task for curvilinear objects because of the more complex geometry and the feedback between the two tasks. Nevertheless, we believe that the framework that we present here could be used as a modelling/algorithmic building block of such a more advanced inverse problem.

VII. CONCLUSION

In this work, we propose a Bayesian framework of the kymograph reconstruction given the geometry of a curvilinear object. This allows a well-grounded formulation of the inverse problem that, first, involves a *realistic image formation model* accounting for optical distortions and measurement noise, second, relies on a proper reconstruction of the *fluorescence signal in physical space* modelled using a class of flexible non-parametric priors derived from *Lévy innovation processes*. Due to the singular nature of curvilinear geometries, the kymograph reconstruction problem is inherently *analog*. However, using the *virtual microscope framework*, we formulate a computationally tractable approximation that allows deriving efficient iterative algorithms based on a fully-split *alternating split Bregman* algorithm.

Using the virtual microscope framework, we assessed our Bayesian framework on synthetic and real image data. We showed that our framework allows modelling different combinations of geometry (straight/curved) and photometry (smooth/piecewise-constant) in a unified fashion, demonstrating the genericity of our approach. In addition, our framework is based on an analog reconstruction space (*i.e.* the kymograph space) that allows interpreting the kymograph directly

in terms of light source distribution dynamics in physical space, where the signal is deblurred and denoised. We applied our framework to the problem of characterising microtubule dynamics *in vivo* in the budding yeast *S. cerevisiae*. We demonstrated that the common point source approximation of the photometry is oversimplified and that it introduces a bias in the estimated microtubule length. Moreover, we show that our framework allows revealing complex patterns occurring on the microtubule lattice from single time-lapse data. These patterns are not clearly identified with a canonical approach, such as the nearest neighbour kymograph.

We expect that the framework proposed in this paper will facilitate the analysis of kymographs and enable new discoveries thanks to the increase in quality and resolution. In addition, the framework is modular and lays the ground for extensions that could better fit particular applications. It will allow formulating more specific models, widening the scope of hypotheses that can be tested in fundamental fields such as cell biology.

ACKNOWLEDGMENT

We thank A. M. Rauch and Y. Barral for sharing the original yeast data. This work has been supported by the SystemsX.ch RTD Grant #2012/192 TubeX of the Swiss National Science Foundation.

REFERENCES

- [1] B. Hinz, W. Alt, C. Johnen, V. Herzog, and H. W. Kaiser, "Quantifying lamella dynamics of cultured cells by SACED, a new computer-assisted motion analysis." *Experimental cell research*, vol. 251, no. 1, pp. 234–243, 1999.
- [2] P. Bieling, L. Laan, H. Schek, E. L. Munteanu, L. Sandblad, M. Dogterom, D. Brunner, and T. Surrey, "Reconstitution of a microtubule plus-end tracking system *in vitro*." *Nature*, vol. 450, no. 7172, pp. 1100–5, dec 2007.
- [3] V. Racine, M. Sachse, J. Salamero, V. Fraisier, A. Trubuil, and J.-B. Sibarita, "Visualization and quantification of vesicle trafficking on a three-dimensional cytoskeleton network in living cells." *Journal of microscopy*, vol. 225, no. Pt 3, pp. 214–28, mar 2007.
- [4] O. Welzel, D. Boening, A. Stroebel, U. Reulbach, J. Klingauf, J. Kornhuber, and T. W. Groemer, "Determination of axonal transport velocities via image cross- and autocorrelation," *European Biophysics Journal*, vol. 38, no. 7, pp. 883–889, 2009.
- [5] A. J. Roberts, B. S. Goodman, and S. L. Reck-peterson, "Reconstitution of dynein transport to the microtubule plus end by kinesin," *eLife*, vol. 3:e02641, pp. 1–16, 2014.
- [6] K. Chiba, Y. Shimada, M. Kinjo, T. Suzuki, and S. Uchida, "Simple and Direct Assembly of Kymographs from Movies Using KYMOMAKER," *Traffic*, vol. 15, no. 1, pp. 1–11, 2014.
- [7] P. Mangeol, B. Prevo, and E. J. G. Peterman, "KymographClear and KymographDirect: two tools for the automated quantitative analysis of molecular and cellular dynamics using kymographs," *Molecular Biology of the Cell*, vol. 27, no. 12, pp. 1948–1957, 2016.
- [8] A. R. Chaphalkar, K. Jain, M. S. Gangan, and C. A. Athale, "Automated multi-peak tracking kymography (AMTrak): A tool to quantify sub-cellular dynamics with sub-pixel accuracy," *PLoS ONE*, vol. 11, no. 12, pp. 1–22, 2016.
- [9] S. Neumann, R. Chassefeyre, G. E. Campbell, and S. E. Encalada, "KymoAnalyzer: a software tool for the quantitative analysis of intracellular transport in neurons," *Traffic*, vol. 18, no. 1, pp. 71–88, 2017.
- [10] K. Zhang, Y. Osakada, W. Xie, and B. Cui, "Automated image analysis for tracking cargo transport in axons," *Microscopy Research and Technique*, vol. 74, no. May 2010, pp. 605–613, 2011.
- [11] A. Mukherjee, B. Jenkins, C. Fang, R. J. Radke, G. Banker, and B. Roysam, "Automated kymograph analysis for profiling axonal transport of secretory granules," *Medical Image Analysis*, vol. 15, no. 3, pp. 354–367, 2010.

- [12] V. Mennella, G. C. Rogers, S. L. Rogers, D. W. Buster, R. D. Vale, and D. J. Sharp, "Functionally distinct kinesin-13 family members cooperate to regulate microtubule dynamics during interphase," *Nature Cell Biology*, vol. 7, no. 3, pp. 235–245, 2005.
- [13] P. Kner, B. B. Chhun, E. R. Griffis, L. Winoto, and M. G. L. Gustafsson, "Super-resolution video microscopy of live cells by structured illumination," *Nature methods*, vol. 6, no. 5, pp. 339–42, 2009.
- [14] A. J. Pereira and H. Maiato, "Improved kymography tools and its applications to mitosis," *Methods*, vol. 51, no. 2, pp. 214–219, 2010.
- [15] J. Chetta and S. B. Shah, "A novel algorithm to generate kymographs from dynamic axons for the quantitative analysis of axonal transport," *Journal of Neuroscience Methods*, vol. 199, no. 2, pp. 230–240, 2011.
- [16] L. Yuan, Y. F. Zheng, J. Zhu, L. Wang, and a. Brown, "Object tracking with particle filtering in fluorescence microscopy images: Application to the motion of neurofilaments in axons," *IEEE Transactions on Medical Imaging*, vol. 31, no. 1, pp. 117–130, 2012.
- [17] D. K. Samuylov, L. A. Widmer, G. Székely, and G. Paul, "Mapping complex spatio-temporal models to image space: The virtual microscope," *IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 707–711, 2015.
- [18] P. Hansen, J. Nagy, and D. O'Leary, *Deblurring Images*. Society for Industrial and Applied Mathematics, 2006. [Online]. Available: <https://pubs.siam.org/doi/abs/10.1137/1.9780898718874>
- [19] G. H. Golub and C. F. Van Loan, *Matrix computations*. JHU Press, 2012, vol. 3.
- [20] C. Vogel, *Computational Methods for Inverse Problems*. Society for Industrial and Applied Mathematics, 2002. [Online]. Available: <https://pubs.siam.org/doi/abs/10.1137/1.9780898717570>
- [21] S. Demko, "Condition numbers of rectangular systems and bounds for generalized inverses," *Linear Algebra and its Applications*, vol. 78, pp. 199–206, 1986. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0024379586900248>
- [22] J. W. Demmel, "On condition numbers and the distance to the nearest ill-posed problem," *Numerische Mathematik*, vol. 51, no. 3, pp. 251–289, May 1987. [Online]. Available: <https://doi.org/10.1007/BF01400115>
- [23] J. E. Gentle, *Matrix Algebra: Theory, Computations and Applications in Statistics*, 2nd ed. Springer Publishing Company, Incorporated, 2017.
- [24] B. Huang, H. Babcock, and X. Zhuang, "Primer Breaking the Diffraction Barrier: Super-Resolution Imaging of Cells," *Cell*, vol. 143, no. 7, pp. 1047–1058, 2010.
- [25] R. Byrd, P. Lu, J. Nocedal, and C. Zhu, "A limited memory algorithm for bound constrained optimization," *SIAM Journal on Scientific Computing*, vol. 16, no. 5, pp. 1190–1208, 1995. [Online]. Available: <https://doi.org/10.1137/0916069>
- [26] F. Aguet, "Super-resolution fluorescence microscopy based on physical models," Ph.D. dissertation, École Polytechnique Fédérale de Lausanne (EPFL), 2009.
- [27] M. Unser and P. D. Tafti, *An introduction to sparse stochastic processes*. Cambridge University Press, 2014.
- [28] E. Bostan, U. S. Kamilov, M. Nilchian, and M. Unser, "Sparse Stochastic Processes and Discretization of Linear Inverse Problems," *IEEE Transactions on Image Processing*, vol. 22, no. 7, pp. 2699–2710, 2013.
- [29] S. Mallat, *A wavelet tour of signal processing*. Academic press, 1999.
- [30] N. Chenouard, I. Smal, F. de Chaumont, M. Maška, I. F. Sbalzarini, Y. Gong, J. Cardinale, C. Carthel, S. Coraluppi, M. Winter, A. R. Cohen, W. J. Godinez, K. Rohr, Y. Kalaidzidis, L. Liang, J. Duncan, H. Shen, Y. Xu, K. E. G. Magnusson, J. Jaldén, H. M. Blau, P. Paul-Gilloteaux, P. Roudot, C. Kervrann, F. Waharte, J.-Y. Tinevez, S. L. Shorte, J. Willemse, K. Celler, G. P. van Wezel, H.-W. Dan, Y.-S. Tsai, C. Ortiz de Solórzano, J.-C. Olivo-Marin, and E. Meijering, "Objective comparison of particle tracking methods," *Nature methods*, vol. 11, no. 3, pp. 281–9, 2014.
- [31] A. Amini, U. S. Kamilov, E. Bostan, and M. Unser, "Bayesian estimation for continuous-time sparse stochastic processes," *IEEE Transactions on Signal Processing*, vol. 61, no. 4, pp. 907–920, 2013.
- [32] D. L. Snyder, R. L. White, and A. M. Hammoud, "Image recovery from data acquired with a charge-coupled-device camera," *J. Opt. Soc. Am. A*, vol. 10, no. 5, pp. 1014–1023, 1993.
- [33] T. Goldstein and S. Osher, "The split Bregman method for L1-regularized problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 2, pp. 323–343, 2009.
- [34] S. Setzer, G. Steidl, and T. Teuber, "Deblurring Poissonian images by split Bregman techniques," *Journal of Visual Communication and Image Representation*, vol. 21, no. 3, pp. 193–199, 2010.
- [35] G. Paul, J. Cardinale, and I. F. Sbalzarini, "Coupling image restoration and segmentation: A generalized linear model/bregman perspective," *International Journal of Computer Vision*, vol. 104, no. 1, pp. 69–93, 2013.
- [36] V. I. Morariu, B. V. Srinivasan, V. C. Raykar, R. Duraiswami, and L. S. Davis, "Automatic online tuning for fast Gaussian summation," *Transform*, pp. 1–8, 2008.
- [37] N. Parikh and S. Boyd, "Proximal Algorithms," *Foundations and Trends in Optimization*, vol. 1, no. 3, pp. 123–231, 2014.
- [38] B. Alberts, A. Johnson, J. Lewis, D. Morgan, M. Raff, K. Roberts, and P. Walter, *Molecular Biology of the Cell*. Garland Science, 2014.
- [39] A. Akhmanova and M. O. Steinmetz, "Control of microtubule organization and dynamics: two ends in the limelight," *Nature reviews. Molecular cell biology*, vol. 16, no. 12, pp. 711–26, 2015.
- [40] G. Borisy, R. Heald, J. Howard, C. Janke, A. Musacchio, and E. Nogales, "Microtubules: 50 years on from the discovery of tubulin," *Nature Reviews Molecular Cell Biology*, vol. 17, no. 5, pp. 322–328, 2016.
- [41] I. Smal, I. Grigoriev, A. Akhmanova, W. J. Niessen, and E. Meijering, "Microtubule dynamics analysis using kymographs and variable-rate particle filters," *IEEE Transactions on Image Processing*, vol. 19, no. 7, pp. 1861–76, jul 2010.
- [42] N. Chenouard and J. Buisson, "Curvelet analysis of kymograph for tracking bi-directional particles in fluorescence microscopy images," *IEEE International Conference on Image Processing (ICIP)*, pp. 3657–3660, 2010.
- [43] R. D. Vale, T. Funatsu, D. W. Pierce, L. Romberg, Y. Harada, and T. Yanagida, "Direct observation of single kinesin molecules moving along microtubules," pp. 451–453, 1996.
- [44] B. Zhang, J. Zerubia, and J. Olivo-Marin, "Gaussian approximations of fluorescence microscope point-spread function models," *Applied Optics*, vol. 46, no. 10, pp. 1819–1829, 2007.
- [45] N. Hansen, S. D. Müller, and P. Koumoustakos, "Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES)," *Evolutionary Computation*, vol. 11, no. 1, pp. 1–18, 2003.
- [46] D. K. Samuylov, G. Székely, and G. Paul, "Tracking microtubule ends is more than point tracking," *IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 808–812, 2017.
- [47] A. Rauch, E. Nazarova, and J. Vogel, *Analysis of Microtubules in Budding Yeast*. Elsevier Inc., 2010, vol. 97, no. C.
- [48] I. Smal, K. Draegestein, N. Galjart, W. Niessen, and E. Meijering, "Particle filtering for multiple object tracking in dynamic fluorescence microscopy images: application to microtubule growth analysis," *IEEE transactions on medical imaging*, vol. 27, no. 6, pp. 789–804, jun 2008.
- [49] J. Cardinale, A. Rauch, Y. Barral, G. Székely, and I. F. Sbalzarini, "Bayesian image analysis with on-line confidence estimates and its application to microtubule tracking," *IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1091–1094, 2009.
- [50] A. Matov, K. Applegate, P. Kumar, C. Thoma, W. Krek, G. Danuser, and T. Wittmann, "Analysis of microtubule dynamic instability using a plus-end growth marker," *Nature methods*, vol. 7, no. 9, pp. 761–8, sep 2010.
- [51] A. Akhmanova and M. O. Steinmetz, "Tracking the ends: a dynamic protein network controls the fate of microtubule tips," *Nature reviews. Molecular cell biology*, vol. 9, no. 4, pp. 309–322, 2008.
- [52] E. Bullitt, M. P. Rout, J. V. Kilmartin, and C. W. Akey, "The yeast spindle pole body is assembled around a central crystal of Spc42p," *Cell*, vol. 89, no. 7, pp. 1077–1086, 1997.
- [53] D. K. Samuylov, P. Prateek, G. Székely, and G. Paul, "Modeling point spread function in fluorescence microscopy with a sparse mixture of gaussians: Trade-off between accuracy and efficiency," ETH Zurich, Tech. Rep., 2017.
- [54] M. Hirsch, R. J. Wareham, M. L. Martin-Fernandez, M. P. Hobson, and D. J. Rolfe, "A Stochastic Model for Electron Multiplication Charge-Coupled Devices – From Theory to Practice," *PLoS ONE*, vol. 8, no. 1, p. e53671, 2013.
- [55] M. Konnik and J. Welsh, "High-level numerical simulations of noise in CCD and CMOS photosensors: review and tutorial," *arXiv:1412.4031*, pp. 1–21, 2014.
- [56] F. Huang, T. M. P. Hartwich, F. E. Rivera-Molina, Y. Lin, W. C. Duim, J. J. Long, P. D. Uchil, J. R. Myers, M. a. Baird, W. Mothes, M. W. Davidson, D. Toomre, and J. Bewersdorf, "Video-rate nanoscopy using sCMOS camera-specific single-molecule localization algorithms," *Nature methods*, vol. 10, no. 7, pp. 653–8, 2013.
- [57] L. Reese, A. Melbinger, and E. Frey, "Crowding of molecular motors determines microtubule depolymerization," *Biophysical Journal*, vol. 101, no. 9, pp. 2190–2200, 2011.