

---

# DLBI: deep learning guided Bayesian inference for structure reconstruction of super-resolution fluorescence microscopy

Yu Li<sup>1,†</sup>, Fan Xu<sup>2,†</sup>, Fa Zhang<sup>2</sup>, Pingyong Xu<sup>3</sup>, Mingshu Zhang<sup>3</sup>,  
Ming Fan<sup>4</sup>, Lihua Li<sup>4</sup>, Xin Gao<sup>1,\*</sup> and Renmin Han<sup>1,\*</sup>

<sup>1</sup>King Abdullah University of Science and Technology (KAUST), Computational Bioscience Research Center (CBRC), Computer, Electrical and Mathematical Sciences and Engineering (CEMSE) Division, Thuwal, 23955-6900, Saudi Arabia, <sup>2</sup>High Performance Computer Research Center, Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100190, China, <sup>3</sup>Key Laboratory of RNA Biology, Institute of Biophysics, Chinese Academy of Sciences, Beijing 100101, China and <sup>4</sup>Institute of Biomedical Engineering and Instrumentation, Hangzhou Dianzi University, Hangzhou 310018, China

\*To whom correspondence should be addressed.

<sup>†</sup>The authors wish it to be known that, in their opinion, the first two authors should be regarded as Joint First Authors.

## Abstract

**Motivation:** Super-resolution fluorescence microscopy with a resolution beyond the diffraction limit of light, has become an indispensable tool to directly visualize biological structures in living cells at a nanometer-scale resolution. Despite advances in high-density super-resolution fluorescent techniques, existing methods still have bottlenecks, including extremely long execution time, artificial thinning and thickening of structures, and lack of ability to capture latent structures.

**Results:** Here, we propose a novel deep learning guided Bayesian inference (DLBI) approach, for the time-series analysis of high-density fluorescent images. Our method combines the strength of deep learning and statistical inference, where deep learning captures the underlying distribution of the fluorophores that are consistent with the observed time-series fluorescent images by exploring local features and correlation along time-axis, and statistical inference further refines the ultrastructure extracted by deep learning and endues physical meaning to the final image. In particular, our method contains three main components. The first one is a simulator that takes a high-resolution image as the input, and simulates time-series low-resolution fluorescent images based on experimentally calibrated parameters, which provides supervised training data to the deep learning model. The second one is a multi-scale deep learning module to capture both spatial information in each input low-resolution image as well as temporal information among the time-series images. And the third one is a Bayesian inference module that takes the image from the deep learning module as the initial localization of fluorophores and removes artifacts by statistical inference. Comprehensive experimental results on both real and simulated datasets demonstrate that our method provides more accurate and realistic local patch and large-field reconstruction than the state-of-the-art method, the 3B analysis, while our method is more than two orders of magnitude faster.

**Availability and implementation:** The main program is available at <https://github.com/lykaust15/DLBI>

**Contact:** [xin.gao@kaust.edu.sa](mailto:xin.gao@kaust.edu.sa) or [renmin.han@kaust.edu.sa](mailto:renmin.han@kaust.edu.sa)

**Supplementary information:** [Supplementary data](#) are available at *Bioinformatics* online.

---

## 1 Introduction

Fluorescence microscopy with a resolution beyond the diffraction limit of light (i.e. super-resolution) has played an important role in biological sciences. The application of super-resolution fluorescence microscope techniques to living-cell imaging promises dynamic information on complex biological structures with nanometer-scale resolution.

Recent development of fluorescence microscopy takes advantages of both the development of optical theories and computational methods. Living cell stimulated emission depletion (STED) (Hein *et al.*, 2008), reversible saturable optical linear fluorescence transitions (RESOLFT) (Schwentker *et al.*, 2007) and structured illumination microscopy (SIM) (Gustafsson, 2005) mainly focus on the innovation of instruments, which requires sophisticated, expensive optical setups and specialized expertise for accurate optical alignment. The time-series analysis based on localization microscopy techniques, such as photoactivatable localization microscopy (PALM) (Hess *et al.*, 2006) and stochastic optical reconstruction microscopy (STORM) (Rust *et al.*, 2006), is mainly based on the computational methods, which build a super-resolution image from the localized positions of single molecules in a large number of images. Though compared with STED, RESOLFT and SIM, PALM and STORM do not need specialized microscopes, the localization techniques of PALM and STORM require the fluorescence emission from individual fluorophores to not overlap with each other, leading to long imaging time and increased damage to live samples (Lippincott-Schwartz and Manley, 2009). More recent methods (Holden *et al.*, 2011; Huang *et al.*, 2011; Quan *et al.*, 2011; Zhu *et al.*, 2012) alleviate the long exposure problem by developing multiple-fluorophore fitting techniques to allow relatively dense fluorescent data, but still do not solve the problem completely.

Bayesian-based time-series analysis of high-density fluorescent images (Cox *et al.*, 2012; Xu *et al.*, 2015, 2017) further pushes the limit. By using data from overlapping fluorophores as well as information from blinking and bleaching events, it extends the super-resolution imaging to the large-field imaging of living cells. Despite its potential to resolve ultrastructures and fast cellular dynamics in living cells, several bottlenecks still remain. The state-of-the-art methods, such as Bayesian analysis of the blinking and bleaching (i.e. the 3B analysis) (Cox *et al.*, 2012), are computationally expensive, and may cause artificial thinning and thickening of structures due to local sampling. Significant improvements on runtime and accuracy have been achieved by single molecule-guided Bayesian localization microscopy (SIMBA) (Xu *et al.*, 2017) with the introduction of dual-channel fluorescent imaging and single molecule-guided Bayesian inference. However, the enhanced process is severely limited by the specialized class of proteins.

Deep learning has accomplished great success in various fields, including super-resolution imaging (Kim *et al.*, 2016; Ledig *et al.*, 2016; Lim *et al.*, 2017). Among different deep learning architectures, the generative adversarial network (GAN) (Goodfellow *et al.*, 2014) achieved the state-of-the-art performance on single image super-resolution (SISR) (Ledig *et al.*, 2016). However, there are two fundamental differences between the SISR and super-resolution fluorescence microscopy. First, the input of SISR is a downsampled (i.e. low-resolution) image of a static high-resolution image and the expected output is the original image, whereas the input of super-resolution fluorescence microscopy is a time-series of low-resolution fluorescent images and the output is the high-resolution image containing estimated locations of the fluorophores (i.e. the reconstructed structure). Second, the nature of SISR ensures that there are

readily a huge amount of existing data to train deep learning models, whereas for fluorescence microscopy, there are only limited time-series datasets. Furthermore, most of these datasets do not have the ground-truth high-resolution images, which make supervised deep learning infeasible.

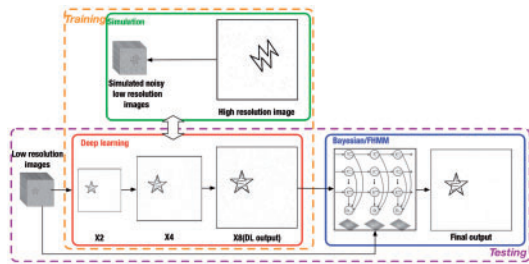
In this article, we propose a novel deep learning guided Bayesian inference (DLBI) framework, for structure reconstruction of high-resolution fluorescent microscopy. Our framework combines the strength of stochastic simulation, deep learning and statistical inference. To our knowledge, this is the first deep learning-based super-resolution fluorescent microscopy method. In particular, the stochastic simulation module simulates time-series low-resolution images from high-resolution images based on experimentally calibrated parameters of fluorophores and stochastic modeling, which provides supervised training data for deep learning models. The deep learning module takes the simulated time-series low-resolution images as inputs, captures the underlying distribution that generates the ground-truth super-resolution images by exploring local features and correlation along time-axis of the low-resolution images, and outputs a predicted high-resolution image. To achieve this goal, we develop a GAN in which a generator network and a discriminator network contest with each other. The generator network tries to learn the distribution of the high-resolution images in a multi-scale manner, whereas the discriminator network tries to discriminate the ground-truth images and the images produced by the generator network. In order to capture the deep features in the images, we further ease the degradation issue by integrating residual networks (He *et al.*, 2016) into our GAN model, where degradation means that stacking more network layers does not lead to better accuracy. The high-resolution image produced by the deep learning module is often very close to the ground-truth image. However, it can still contain some artifacts, and more importantly, lacks the physical meaning. Thus, we develop the Bayesian inference module to take the predicted high-resolution image from deep learning, run Bayesian inference from the initial locations of fluorophores in the predicted image, and predict a more accurate high-resolution image.

Comprehensive experimental results on two simulated and three real-world datasets demonstrate that DLBI provides more accurate and realistic local-patch as well as large-field reconstruction than the state-of-the-art method, the 3B analysis (Cox *et al.*, 2012). Meanwhile, our method is more than 100 times faster than the 3B analysis.

## 2 Materials and methods

As shown in Figure 1, DLBI contains three modules: (i) stochastic simulation (Section 2.1), (ii) deep neural networks (Section 2.2) and (iii) Bayesian inference (Section 2.3).

Although deep learning has proved its great superiority in various fields, it has not been used for fluorescent microscopy image analysis. One of the possible reasons is the lack of supervised training data, which means the number of time-series low-resolution image datasets is limited and even for the existing datasets, the ground-truth high-resolution images are often unknown. Here, a stochastic simulation based on the experimentally calibrated parameters is designed to solve this issue, without the need of collecting a massive amount of real fluorescent images. This empowers our deep neural networks to effectively learn the latent structures under the low-resolution, high-noise and stochastic fluorescing conditions. The primitive super-resolution images produced by deep neural networks still contain artifacts and lack physical meaning, we finally



**Fig. 1.** The overall workflow of DLBI. The three modules are shown in solid boxes: the simulation module (green), the deep learning module (red) and the Bayesian inference module (blue). The training (orange) and testing (purple) procedures are shown in dashed boxes. For training, time-series low-resolution images simulated from the simulation module are used to train the deep learning module in a multi-scale manner, from 2 $\times$ , 4 $\times$ , to 8 $\times$  resolution. For testing, given certain time-series low-resolution images, the deep learning module predicts the 2 $\times$  to 8 $\times$  super-resolution images, and the Bayesian inference module takes the predicted 8 $\times$  image (which contains artifacts) and produces the final high-resolution image

develop a Bayesian inference module based on the mechanism of fluorophore switching to produce high-confident images.

Our method combines the strength of deep learning and statistical inference, where deep learning captures the underlying distribution that generates the training super-resolution images by exploring local features and correlation along time-axis, and statistical inference removes artifacts and refines the ultrastructure extracted by deep learning, and further endues physical meaning to the final image.

## 2.1 The stochastic simulation module


The input of our simulation module is a high-resolution image that depicts the distribution of the fluorophores and the output is a time-series of low-resolution fluorescent images with different fluorescing states. We refer the readers to [Supplementary Section S1](#) for terminologies in fluorescence microscopy.

In our simulation, Laplace-filtered natural images and sketches are used as the ground-truth high-resolution images that contain the fluorophore distribution. If a gray-scale image is given, the depicted shapes are considered as the distribution of fluorophores and each pixel value on the image is considered as the density of fluorophores at the location. We then create a number of simulated fluorophores that are distributed according to the distribution and the densities. For each fluorophore, it switches according to a Markov model, i.e. among states of emitting (activated), not emitting (inactivated), and bleached. The emitting state means that the fluorophore emits photons and a spot according to the point spread function (PSF) is depicted on the canvas. All the spots of the emitting fluorophores thus result in a high-resolution fluorescent image. Applying the Markov model on the initial high-resolution image generates a time-series of high-resolution images. After adding the background to the high-resolution images, they are downsampled to low-resolution images and noise is finally added. [Figure 2](#) summarizes the stochastic simulation procedure.

Here, the success of simulation relies on three factors: (i) the principal of the linear optical system, (ii) experimentally calibrated parameters of fluorophores and (iii) stochastic modeling.

### 2.1.1 Linear optics

A fluorescence microscope is considered as a linear optical system, in which the superposition principle is valid, i.e. Image



**Fig. 2.** The workflow of stochastic simulation. Firstly, a high-resolution image is inputted as the distribution and density of fluorophores. Then, the emitting of photons is simulated based on the stochastic parameters for each time frame. A random background (DC offset) is added to each image. The images are then downsampled to low-resolution and noise is added, which results in a time-series of low-resolution images

(Obj1 + Obj2) = Image (Obj1) + Image (Obj2). The behavior of fluorophores is considered invariant to mutual interaction. Therefore, for high-density fluorescent images, the pixel density can be directly calculated from the light emitted from its surrounding fluorophores.

When a fluorophore is activated, an observable spot can be recorded by the sensor, the shape of which is called the PSF. Considering the limitation of sensor capability, the PSF of an isotropic point source is often approximated as a Gaussian function:

$$I(x, y) = I_0 \exp\left(-\frac{1}{2\sigma^2} \left((x - x_0)^2 + (y - y_0)^2\right)\right), \quad (1)$$

where  $\sigma$  is calculated from the fluorophore in the specimen that specifies the width of the PSF,  $I_0$  is the peak intensity and is proportional to the photon emission rate and the single-frame acquisition time,  $(x_0, y_0)$  is the location of the fluorophore.

While PSF describes the shape, the full width at half maximum (FWHM) describes the distinguishability. It is defined to be the half width of the maximum amplitude of PSF. If PSF is modeled as a Gaussian function, the relationship between FWHM and  $\sigma$  is given by

$$\text{FWHM} = 2\sqrt{2\ln 2} \sigma \approx 2.355 \sigma. \quad (2)$$

Considering the probability of linear optics, a high-density fluorescent image is composed by PSFs of the fluorophores.

### 2.1.2 Calibrated parameters of fluorophores

In most imaging systems, the characteristics of a fluorescent protein can be calibrated by experimental techniques. With all the calibrated parameters, it is not difficult to describe and simulate the fluorescent switching of a specialized protein.

The first characteristic of a fluorophore is its switching probability. A fluorophore always transfers among three states, emitting, not emitting and bleached, which can be specified by a Markov model ([Fig. 3](#)). If the fluorophore transfers from not emitting to bleached, it will not emit any photon anymore. As linear optics, each fluorophore's transitions are assumed to be independent.

The second characteristic of a fluorophore is its PSF. When a real-world fluorophore is activated, the emitted photons and its corresponding PSF will not stay unchanged over time. The stochasticity of the PSF and photon strength describes the characteristics of a fluorescent protein. To simulate the fluorescence, we should not ignore these properties. Fortunately, the related parameters can be well-calibrated. The PSF and FWHM of a fluorescent protein can be measured in low molecule density. In an instrument for PALM or




Fig. 3. The Markov model describing state transition of a fluorophore

STORM, the PSF of the microscope can be measured by acquiring image frames, fitting the fluorescent spots parameter, normalizing and then averaging the aligned single-molecule images. The distribution of FWHM can be obtained from statistical analysis. The principle of linear optics ensures that the parameters measured in single-molecule conditions are also applicable to high-density conditions.

In our simulation, a log-normal distribution (Cox *et al.*, 2012; Zhu *et al.*, 2012) is used to approximate the experimentally measured single fluorophore photon number distribution. Firstly, a table of fluorophore’s experimentally calibrated FWHM parameters is used to initialize the PSF table in our simulation, according to Equations (1) and (2). Then for each fluorophore recorded in the high-resolution image, the state of the current image frame is calculated according to the transfer table [P1, P2, P3, P4, P5] (Fig. 3) and a random PSF shape is produced if the corresponding fluorophore is at the ‘emitting’ state. This procedure is repeated for each fluorophore, which results in the final fluorescent image.

### 2.1.3 Stochastic modeling

The illumination of real-world objects is different at different time. In general, the illumination change of real-world objects can be suppressed by high-pass filtering with a large Gaussian kernel. However, this operation will sharpen the random noise and cannot remove the background or DC offset (DC offset, DC bias or DC component denotes the mean value of a signal. If the mean amplitude is zero, there is no DC offset. For most microscopy, the DC offset can be calibrated but cannot be completely removed.) To make our simulation more realistic, several stochastic factors are introduced. First, for a series of simulated fluorescent images, a background value calculated from the multiplication between a random strength factor and the average image intensity is added to the fluorescent images to simulate the DC offset. For the same time-series, the strength factor remains unchanged but the background strength changes with the image intensity. Second, the high-resolution fluorescent image is downsampled and random Gaussian noise is added to the low-resolution image. Here, the noise is also stochastic for different time-series and close to the noise strength that is measured from the real-world microscopy.

The default setting of our simulation takes a  $480 \times 480$  pixel high-resolution image as the input and simulates 200 frames of  $60 \times 60$  pixel (i.e.  $8 \times$  binned) low-resolution images.

## 2.2 The deep learning module

We build a deep residual network under the GAN framework (Goodfellow *et al.*, 2014; Ledig *et al.*, 2016) to estimate the primitive super-resolution image  $I^{SR}$  (the latent structure features) from time-series of low-resolution fluorescent images  $\mathcal{T} = \{I_k^{LR}\}_{k=1, \dots, K}$ . Instead of building just one generative model, our approach builds a pair of models, a generator model,  $G$ , which produces the estimation of the underlying structure of the training images, and a discriminator model,  $D$ , which is trained to distinguish the reconstructed super-resolution image from the ground-truth one. Figure 4 demonstrates the overview of our deep learning framework.




Fig. 4. Overview of our deep learning framework. There are two main components of the model, the generator network ( $G$ ) and the discriminator network ( $D$ ).  $G$  is used to convert the time-series noisy, low-resolution images into a noise-free, super-resolution image while  $D$  is used to distinguish the ground-truth high-resolution image from the one produced by  $G$ .  $G$  and  $D$  are designed to contest each other, which are trained simultaneously. As the training goes on, both  $G$ ’s ability to generate better super-resolution images and  $D$ ’s ability to distinguish the generated images are improved, which results in more and more similar images to the ground-truth from  $G$ . During testing, only  $G$  is used

### 2.2.1 Basic concepts

The goal of training a generator neural network is to obtain the optimized parameters,  $\theta_G$ , for the generating function,  $G$ , with the minimum difference between the output super-resolution image,  $I^{SR}$ , and ground-truth,  $I^{HR}$ :

$$\hat{\theta}_G = \arg \min_{\theta_G} \frac{1}{N} \sum_{n=1}^N I^{SR}(G(\mathcal{T}_n, \theta_G), I_n^{HR}), \quad (3)$$

where  $G(\mathcal{T}_n, \theta_G)$  is the generated super-resolution image by  $G$  for the  $n$ th training sample,  $N$  is the number of training images, and  $I^{SR}$  is a loss function that will be specified later.

For the discriminator network  $D$ ,  $D(x)$  represents the probability of the data being the real high-resolution image rather than from  $G$ . When training  $D$ , we try to maximize its ability to differentiate ground-truth from the generated image, to force  $G$  to learn better details. When training  $G$ , we try to minimize  $\log(1 - D(G(\mathcal{T}_n, \theta_G), \theta_D))$ , which is the log likelihood of  $D$  being able to tell that the image generated by  $G$  is not ground-truth. That is, we minimize the following function:

$$\begin{aligned} \min_{\theta_G} \max_{\theta_D} & \mathbb{E}_{I^{HR} \sim p_{train}(I^{HR})} [\log(D(I^{HR}, \theta_D))] \\ & + \mathbb{E}_{I^{HR} \sim p_G(\mathcal{T})} [\log(1 - D(G(\mathcal{T}, \theta_G), \theta_D))]. \end{aligned} \quad (4)$$

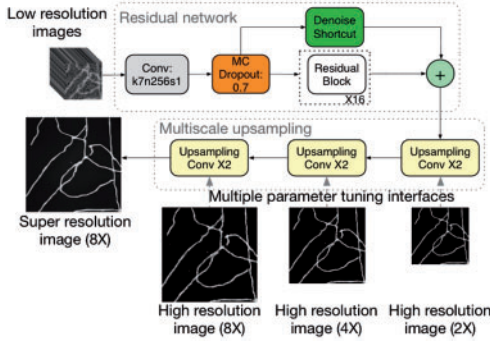
In this way, we force the generator to optimize the generative loss, which is composed of perceptual loss, content loss and adversarial loss (more details of the loss function will be introduced in Section 2.2.3).

### 2.2.2 Model architecture

Our network is specialized for the analysis of time-series images through: (i) 3D filters in the neural network that take all the image frames into consideration, which extracts the time dependent information naturally, (ii) two specifically designed modules in the generator residual network, i.e. Monte Carlo dropout (Gal and Ghahramani, 2015) and denoise shortcut, to cope with the stochastic switching of fluorophores and random noise and (iii) a novel incremental multi-scale architecture and parameter tuning scheme, which is designed to suppress the error accumulation in large upscaling factor neural networks.

Figure 5 illustrates the entire architecture of the generator model. The input is time-series low-resolution images. We first use a convolutional layer with the filter size as seven by seven (Here, the filter size depends on the FWHM of a PSF. Generally, a fluorescence microscope produces low-resolution images with PSF spanning





**Fig. 5.** Architecture of the generator network, which is composed of a residual network component and a multi-scale upsampling component. The low-resolution images are firstly fed to the residual network to extract information from the original 3D space, during which denoising is performed by the Monte Carlo dropout and denoise shortcut. The extracted feature maps are fed into the multi-scale upsampling component to increase the resolution gradually. We increase the resolution by a factor of two during each upsampling, resulting in three parameter tuning interfaces. Using the three interfaces, we can use all the  $2\times$ ,  $4\times$  and  $8\times$  ground-truth images to train the generator network, reducing the artifacts in the final  $8\times$  super-resolution images greatly

$3 \sim 7$  pixels. The specially designed filter size (which can balance between the computational time and physical meaning.), which is larger than the commonly used filter, to capture meaningful features of the input fluorescence microscope images. The Monte Carlo dropout layer, which drops out some pixels from the input feature maps during both training and testing, is applied to the output of the first layer to suppress noise. To further alleviate the noise issue, we use another technique, the denoise shortcut. It is similar to the identical shortcut in the residual network block. However, instead of being exactly the same as the input, we set each channel of the input feature map as the average of all the channels. The denoise shortcut is added to the output of the convolutional layer, which is after 16 residual blocks (see [Supplementary Section S2](#)), element-wise. After this feature map extraction process, we use a pixel shuffle layer combined with the convolutional layer to increase the dimensionality of the image gradually (upsampling Conv X2 in [Fig. 5](#)).

Here, we adopt a novel multi-scale tuning procedure to stabilize the  $8\times$  images. As shown in [Figure 5](#), our generator can output and thus calculate the training error of multi-scale super-resolution images, ranging from  $2\times$  to  $8\times$ , which means that our model has multiple training interfaces for back propagation. Thus during training, we use the  $2\times$ ,  $4\times$ ,  $8\times$  high-resolution ground-truth images to tune the model simultaneously to ensure that the dimensionality of the images increases smoothly and gradually without introducing too much fake detail.

For the discriminator network, we adopt the traditional convolutional neural network, which contains eight convolutional layers, one residual block and one sigmoid layer (see [Supplementary Section S2](#)). The convolutional layers increase the number of channels gradually to 2048 and then decrease it using one by one filters. Those convolutional layers are followed by a residual block, which further increases the model ability of extracting features. [Supplementary Section S2](#) provides detailed descriptions of the generator and discriminator networks.

### 2.2.3 Model training and testing

GAN is known to be difficult to train ([Salimans et al., 2016](#)). We use the following techniques to obtain stable models. For the generator model, we do not train GAN immediately after initialization.

Instead, we pretrain the model. During the pretrain process, we minimize the mean squared error between the super-resolution image and the ground-truth, i.e. with the pixel-wise MSE loss as

$$I_{MSE_\mu}^{SR} = \frac{1}{\mu^2 WH} \sum_x \sum_y^{\mu H} \left( G(T, \theta_{G_\mu}) - I_{x,y}^{HR} \right)^2, \quad (5)$$

where  $W$  is the width of the low-resolution image,  $H$  is the height of the low-resolution image, and  $\mu = 2, 4, 8$  is the upscaling factor. During pretraining, we optimize  $I_{MSE_8}^{SR}$ ,  $I_{MSE_4}^{SR}$ ,  $I_{MSE_2}^{SR}$  simultaneously, instead of optimizing the sum of them.

Only after the model has been well-pretrained do we start training the GAN. During that process, we also use VGG19 ([Simonyan and Zisserman, 2014](#)) to calculate the perceptual loss ([Johnson et al., 2016](#)) and use Adam optimizer ([Kingma and Ba, 2014](#)) with learning rate decay as the optimizer. When feeding an image to the VGG model, we resize the image to fulfill the dimensionality requirement:

$$I_{VGG_\mu}^{SR} = \sum_{i=1}^V \left( VGG(G(T, \theta_{G_\mu}))_i - VGG(I^{HR})_i \right)^2, \quad (6)$$

where  $V$  is the dimensionality of the VGG embedding output.

During final tuning, we simultaneously optimize the  $2\times$ ,  $4\times$ , and  $8\times$  upscaling by the generative loss:

$$I_{GAN_\mu}^{SR} = 0.4 * I_{MSE_\mu}^{SR} + 10^{-6} * I_{VGG_\mu}^{SR}, \quad (7)$$

and

$$I_{GAN_8}^{SR} = 0.5 * I_{MSE_8}^{SR} + 10^{-3} * I_{ADV_8}^{SR} + 10^{-6} * I_{VGG_8}^{SR}, \quad (8)$$

where  $\mu = 2, 4$  and the  $8\times$  upscaling has an additional term, the adversarial loss  $I_{ADV_8}^{SR} = \sum_{n=1}^N \log(1 - D(G(T_n, \theta_G), \theta_D))$ . For the discriminator network, we use the following loss function:

$$I_{DIS}^{SR} = \sum_{n=1}^N \log(D(G(T_n, \theta_G), \theta_D)) + \sum_{n=1}^N \log(1 - D(I_n^{HR}, \theta_D)). \quad (9)$$

During testing, for the same input time-series images, we run the model multiple times to get a series of super-resolution images. Because of the Monte Carlo dropout layer in the generator model, all of the super-resolution images are not identical. We then compute the average of these images as the final prediction, with another map showing the  $P$ -value of each pixel. We use Tensorflow combined with TensorLayer ([Dong et al., 2017](#)) to implement the deep learning module. Trained on a workstation with one Pascal Titan X, the model gets converged in around 8 h.

### 2.3 The Bayesian inference module

Our Bayesian inference module takes both the time-series low-resolution images and the primitive super-resolution image produced by the deep learning module as inputs, and generates a set of optimized fluorophore locations, which are further interpreted as a high-confident super-resolution image. Since the deep learning module has already depicted the ultrastructures in the image, we use these structures as the initialization of the fluorophore locations, re-sampling with a random punishment against artifacts. For each pixel, we re-sample the fluorophore intensity by  $\sqrt{I_{x,y}}$  and the location by  $(x, y) \pm rand(x, y)$ , where  $I_{x,y}$  is the pixel value in the image produced by deep learning,  $rand(x, y)$  is limited in  $\pm 8$ . In this way, the extremely high illumination can be suppressed and fake structures will be re-estimated.

### 2.3.1 Basic concepts

As shown in Figure 3, a fluorophore has three states: emitting (light), not emitting and bleached. In classic Bayesian-based time-series analysis, the switching procedure of fluorophores is modeled by Bayesian inference, i.e. given an observed region  $R$ , deciding whether there is a fluorophore ( $F$ ) or not ( $N$ ) by

$$\frac{P(F|R)}{P(N|R)} = \frac{P(R|F)P(F)}{P(R|N)P(N)}, \quad (10)$$

where  $P(F)$  and  $P(N)$  are constants which are based on experimental prior,  $P(R|F)$  is the probability of the observed data region  $R$  given the location of the fluorophore,  $P(R|N)$  is the probability of the observed data region  $R$  if there is no fluorophore, which can be calculated by integrating all the probability of observing pixels given the noise model.

For a single fluorophore, the switching procedure can be modeled by a hidden Markov model (HMM) (Rabiner, 1989), as shown in Figure 6A. However, for high-density fluorophores, each fluorophore transfers the state independently with a stable probability (Cox *et al.*, 2012) and all the fluorophores together can be modeled by a factorial hidden Markov model (FHMM) (Ghahramani and Jordan, 1996), as shown in Figure 6B, which has been used and proved in (Xu *et al.*, 2017).

### 2.3.2 Refining results with physical meaning

Although HMM and FHMM are capable of modeling the fluorophore switching process, they are localization-guided, which often ignore the global information, and are computationally expensive to learn. Thus, we initialize the fluorophores' locations by using the image generated by deep learning and use Bayesian inference to further refine the results.


We apply the FHMM model to deal with high-density fluorescence microscopy. The parameters of FHMM are estimated by the expectation-maximization algorithm:

$$\mathcal{Q}(\phi^{new}|\phi) = E\{\log P(\{F_t, D_t\}|\phi^{new})|\phi, \{D_t\}\}, \quad (11)$$

where the observation sequence has  $T$  frames,  $\{D_t\}$ ,  $t = 1, \dots, T$ . The hidden states are  $\{F_t\}$ , where each fluorophore has three possible states in the model.  $\mathcal{Q}$  is a function of the fluorophore parameters  $\phi^{new}$  given the current parameter estimation and the observation sequence  $\{D_t\}$ . The procedure iterates between a step that fixes the current parameters and computes posterior probabilities over the hidden states (the E-step), and a step that uses these probabilities to maximize the expected log likelihood of the observations as a function of the parameters (the M-step).

In the E-step, we fix the fluorophore parameters in the model and utilize the hybrid of Markov chain Monte Carlo and forward algorithm to sample the initial model. When a new fluorophore is determined, we take samples of this fluorophore using the forward filtering backward sampling algorithm (Godsill *et al.*, 2004). Thus, the sampled image sequence contains this fluorophore. In the M-step, we optimize the fluorophore parameters and find the maximum a posteriori (MAP) fluorophore positions using the conjugate gradient. Then, based on already known positions of fluorophores, the surrounding fluorophores with high probability are expended. The final super-resolution image is obtained by iterating these two steps until convergence.

The detailed method description and parameter setting are given in Supplementary Section S3.



**Fig. 6.** The Bayesian inference model used for fluorophore switching. (A) For one fluorophore, its transition between different states can be modeled by a HMM. (B) For high-density fluorophores, the observed fluorescence and their underlying transitions can be modeled by a FHMM

## 3 Experimental results

### 3.1 Training deep learning


To train our deep learning module, the stochastic simulation module was used to simulate time-series low-resolution images from 12 000 gray-scale high-resolution images. These images were downloaded from two databases: (i) 4000 natural images were downloaded from ILSVRC (Russakovsky *et al.*, 2015) and Laplace filtered and (ii) 8000 sketches were downloaded from the Sketchy Database (Sangkloy *et al.*, 2016). Note that our simulation is a generic method, which does not depend on the type of the input images. Thus any gray-scale image can be interpreted as the fluorophore distribution and used to generate the corresponding time-series low-resolution images.

To initialize all the weights of the deep learning models, we used the random normal initializer with the mean as 0 and standard deviation as 0.02. As for the Monte Carlo dropout layer, we set the keep ratio as 0.8. In terms of the Adam optimizer (Kingma and Ba, 2014), we followed the setting in (Dai *et al.*, 2017; Li *et al.*, 2018) and set the learning rate as  $1 * 10^{-4}$  and the  $\beta_{1,1}$ , which is the exponential decay rate for the first moment estimates, as 0.9. During training, we set the batch size as 8, the initialization training epoch as 2 and the GAN training epoch as 40. When performing the real GAN training, we utilized the learning rate decay technique, reducing the learning rate by half every 10 epochs.

### 3.2 Evaluation datasets

Two simulated datasets and three real-world datasets are used to evaluate the performance of the proposed method. Simulated datasets are used due to the availability of ground-truth.

The first two datasets are simulated datasets, for which the ground-truth (i.e. high-resolution images) is downloaded from the Single-Molecule Localization Microscopy (SMLM) challenge (<http://bigwww.epfl.ch/smlm/datasets/>) (Sage *et al.*, 2015). The two datasets correspond to two structures: MT0.N1.HD (MT) and Tubulin ConjAL647 (Tub). For each structure, single molecule positions were downloaded and then transformed to fluorophore densities according to Section 2.1. For simulation, the photo-convertible fluorescent protein (PCFP) mEos3.2 (Zhang *et al.*, 2012) and its associated PSF, FWHM and state transfer table were used. For the convenience of calculation, we cropped the large-field structure into four separate areas, each with  $480 \times 480$  pixels ( $1px = 20nm$ ). For each high-resolution image, 200 frames of low-resolution fluorescence images were generated, each with  $60 \times 60$  pixels.



**Fig. 7.** Visualization of the ground-truth high-resolution images, representative low-resolution input images, the reconstruction results of the 3B analysis (Cox *et al.*, 2012), and the results of our method on three representative areas of each simulated dataset: MT (columns 1–3) and Tub (columns 4–6). The four rows show the ground-truth high-resolution images, the first frames of the simulated time-series low-resolution images, the reconstruction results of the 3B analysis and the reconstruction results of DLBI, respectively

The third dataset is a real-world dataset, which was used in recent work (Xu *et al.*, 2017). The actin was labeled with mEos3.2 (For the convenience of cellular labeling and instrument setup, here all the experiments were carried out by mEos3.2.) in U2OS cells (Actin1) and taken with an exposure time of 50 ms per image frame. The actin network is highly dynamic and exhibits different subtype structures criss-crossing at various distances and angles, including stress fibers and bundles with different sizes and diameters. The dataset has 200 frames of high-density fluorescent images, each with  $249 \times 395$  pixels ( $1px = 160nm$ ) in the green channel. This is a good benchmark set that has been well tested which can compare our method with SIMBA (Xu *et al.*, 2017), a recent Bayesian approach based on dual-channel imaging and PCFPs.

Two other real-world datasets labeled with mEos3.2 were also used. One is an actin cytoskeleton network (Actin2), which is labeled and taken under a similar exposure condition with Actin1, but is completely new and has not been used by previous works. The other one is an Endoplasmic reticulum structure (ER), which has a more complex structure. It is a type of organelle that forms an interconnected network of flattened, membrane-enclosed sacs or tubes known as cisternae, which exhibits different circular-structures and connections at different scales. For the ER dataset, the exposure time is 6.7 ms per frame. The resolution of each image in Actin2 is  $263 \times 337$  pixels ( $1px = 160nm$ ) and that in ER is  $256 \times 170$  pixels ( $1px = 100nm$ ). Both datasets have 200 frames of high-density fluorescent images and the same photographing parameters as Actin1. These datasets were used to demonstrate the power of our method in diverse ultrastructures. The detailed procedure for collecting the real-world datasets is given in Supplementary Section S4.

Since the 3B analysis (Cox *et al.*, 2012) is one of the most widely used high-density fluorescent super-resolution techniques, which can deal with high temporal and spatial resolutions (Cox *et al.*, 2012; Lidke, 2012), it was chosen to compare with our method.

### 3.3 Performance on simulated datasets

#### 3.3.1 Visual performance

Figure 7 shows the visualization of the ground-truth high-resolution images, representative low-resolution input images, the reconstruction results of the 3B analysis, and the results of our method on the simulated datasets. Due to the space limitation, we illustrate three representative areas of each dataset and leave the fourth in Supplementary Section S5.

As shown in Figure 7, the ground-truth images have very clear structures while the low-resolution image frames are very blurry and noisy ( $8\times$  downsampled). To reconstruct the ultrastructures, we ran the 3B analysis with 240 iterations and ran our Bayesian inference module after the deep learning module with 60 iterations. In each iteration, the Bayesian inference module of our method searches four neighbor points for each fluorophore, whereas the 3B analysis takes isolated estimation strategy. Thus the difference in iteration numbers is comparable. Due to the high computational expense of the 3B analysis, each  $60 \times 60$  image was subdivided into nine overlapped subareas for multi-core process, whereas for our method, the entire image was processed by a single CPU core.

It is clear that the reconstructions of our method are very similar to the ground-truth in terms of smoothness, continuity, and thickness. On the other hand, the reconstructions of the 3B analysis consist of a number of interrupted short lines and points with thin structures. In general, two conclusions can be drawn from the visual inspection.

First, DLBI discovered much more natural structures than the 3B analysis. For example, in the bottom part of Figure 7B, there are two lines overlapping with each other and a bifurcation at the tail. Due to the very low resolution in the input time-series images (e.g. Fig. 7H), neither DLBI nor the 3B analysis was able to recover the overlapping structure. However, DLBI reconstructed the proper thickness of that structure (Fig. 7T), whereas the 3B analysis only recovered a very thin line structure (Fig. 7N). Moreover, the



bifurcation structure was reconstructed naturally by DLBI. Similar conclusions can be drawn on the more complex structures in the Tub dataset (columns 4–6 in Fig. 7).

Second, DLBI discovered much more latent structures than the 3B analysis. The Tub dataset consists of a lot of lines (tubulins) with diverse curvature degrees (Fig. 7D–F). The reconstructions of the 3B analysis successfully revealed most of the tubulin structures but left the crossing parts interrupted (Fig. 7P–R). As a comparison, the reconstruction results of DLBI recovered both the line-like tubulin structures and most of the crossing parts accurately (Fig. 7V–X).

### 3.3.2 Quantitative performance

For single-molecule super-resolution fluorescence microscopy, the quantitative performance has been measured by assessing the localization accuracy of single-emitters in each frame (Huang *et al.*, 2011; Ram *et al.*, 2006; Small, 2009). For high-density super-resolution fluorescence microscopy, the entire time-series are analyzed and the production is the probability map of the locations of fluorophores.

Since the ground-truth is known for the simulated datasets, here we use peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) to measure the reconstruction performance, both of which are widely-used criteria for image reconstruction in computer vision. The performance of the 3B analysis and DLBI on the two simulated datasets are given in Table 1. Here, we denote the four areas (Section 3.2) of each dataset as ‘01’, ‘02’, ‘03’ and ‘04’, respectively. It can be seen that DLBI clearly outperforms the 3B analysis in terms of both PSNR and SSIM on all the areas of the two datasets.

### 3.4 Performance on real datasets

Figure 8 shows the first frame of the time-series fluorescent images for each of the three real-world datasets. Here, we evaluate the performance of our method for both local-patch reconstruction (areas selected by green rectangles) and large-field reconstruction (areas selected by yellow rectangles).

#### 3.4.1 Local-patch reconstruction

Figure 9 shows the first frames of the low-resolution images of the three local-patches, and the reconstruction results of the 3B analysis and DLBI. The regions of interests of the selected patches are (60, 60, 60, 60), (120, 120, 60, 60) and (60, 60, 60, 60) for the three datasets, respectively (Region of interest is usually denoted as  $(X, Y, W, H)$ , where  $(X, Y)$  are the coordinates of the top left point of the rectangle,  $W$  is the width of the rectangle, and  $H$  is the height of the rectangle.). The temporal resolutions for the two actin datasets and the ER dataset were 10 s and 1.34 s, respectively, according to the exposure time of the image frames.

It can be seen that the reconstruction results of the 3B analysis capture the main structures in the fluorescent images, but mainly

consist of isolated high-illuminating spots, with details being interrupted (Fig. 9D–F). In contrast, the results of DLBI recover most of the latent ultrastructures, and the reconstructed structures have well-estimated fluorophore distribution and continuous depiction (Fig. 9G–I).

We further assessed the reconstruction quality of the 3B analysis and DLBI by SQUIRREL (super-resolution quantitative image rating and reporting of error locations) (Culley *et al.*, 2018). SQUIRREL compares the diffraction-limited image (the reference image) and the reconstructed equivalents to generate a quantitative map, in which two scores are calculated: the resolution-scaled Pearson coefficient (RSP) and the resolution-scaled error (RSE). The higher RSP and lower RSE values, the higher the image quality is. Table 2 shows the RSP and RSE scores for the 3B analysis and DLBI. It is clear




Fig. 8. The first frame of the time-series fluorescent images for each of the three real-world datasets: (A) Actin1 (Xu *et al.*, 2017), (B) Actin2 and (C) ER




Fig. 9. Reconstructions of the local patches of the three real datasets. First column: the Actin1 dataset (the green box in Fig. 8A). Second column: the Actin2 dataset (the green box in Fig. 8B). Third column: the ER dataset (the green box in Fig. 8C). The first row shows the first frames of the time-series low-resolution images. The second row shows the reconstructions of the 3B analysis. The third row shows the reconstructions of DLBI. The reconstructed images are  $480 \times 480$  pixels and the local-patch images are  $60 \times 60$  pixels

Table 1. Performance comparison between the 3B analysis and DLBI on the four areas of the two simulated datasets in terms of PSNR and SSIM

Datasets	MT0.N1.HD				Tubulin ConjAL647				
	01	02	03	04	01	02	03	04	
PSNR (dB)	3B	17.99	17.62	17.84	17.89	13.42	15.49	15.00	13.21
	DLBI	<b>18.59</b>	<b>19.16</b>	<b>18.51</b>	<b>20.42</b>	<b>18.72</b>	<b>19.17</b>	<b>18.72</b>	<b>16.63</b>
SSIM	3B	0.89	0.89	0.90	0.90	0.74	0.81	0.75	0.69
	DLBI	<b>0.92</b>	<b>0.92</b>	<b>0.93</b>	<b>0.94</b>	<b>0.82</b>	<b>0.85</b>	<b>0.80</b>	<b>0.76</b>

The best performance is shown in bold.



that DLBI significantly outperforms the 3B analysis. More detailed comparisons are given in [Supplementary Section S5.3](#).

### 3.4.2 From deep learning to Bayesian inference

Our method combines the strength of deep learning and statistical inference, where deep learning captures both local features in the images and the time-course correlation, and statistical inference removes artifacts from deep learning and enhances physical meaning to the final results. Conceptually, this is equivalent to using the power of deep learning to automatically and systematically explore and extract spatial and temporal features, and taking advantages of the explicit and rigorous mathematical foundation of probabilistic graphical models. Here, we investigate the effectiveness of this combination.


[Figure 10](#) demonstrates the outputs of the deep learning module and the Bayesian inference module. It can be seen that the super-resolution images outputted from the deep learning module are very close to the final images from the Bayesian inference module, except for some artifacts and false structures. This is due to two reasons: (i) the abundance of training data provided by our simulation module, which are simulated under the real experimentally-calibrated parameters, enable deep learning to effectively learn spatial and temporal features and (ii) the high diversity of biological structures is still a challenge, which causes the artifacts and false structures to be learned by deep learning.

After the deep learning module generates the super-resolution image, the Bayesian inference module uses both the original time-series low-resolution images and the deep learning image to statistically infer a ‘false/true’ determination on each fluorophore location and produce the final image. In particular, the false structures are not directly rejected but used as seeds to search for true structures. Therefore, as shown in [Figure 10](#), although the deep learning module outputted

**Table 2.** Performance comparison between the 3B analysis and DLBI on the real datasets in terms of RSP and RSE with SQUIRREL

Dataset	Actin1-patch		Actin2-patch		ER-patch	
	RSP	RSE	RSP	RSE	RSP	RSE
3B	0.583	3915.096	0.770	2196.068	0.827	4077.037
DLBI	<b>0.721</b>	<b>3326.007</b>	<b>0.878</b>	<b>1648.919</b>	<b>0.916</b>	<b>2904.707</b>

The higher RSP and lower RSE values, the higher the image quality is. The best performance is shown in bold.



**Fig. 10.** Reconstructions of the three local patches of the three real datasets (the first column) by the deep learning module (the second column) and by deep learning guided Bayesian inference (the third column)

some unnatural structures for the Actin2 and ER datasets, these structures were further corrected by the Bayesian inference module.


### 3.4.3 Runtime analysis

After being trained, running the deep learning model is very computationally inexpensive. Furthermore, the results of deep learning provide a close-to-optimal initialization for Bayesian inference, which also significantly reduces trial-and-error and leads to faster convergence. [Figure 11](#) shows the runtime comparison of the deep learning module, the entire DLBI pipeline, and the 3B analysis on the nine reconstruction tasks (i.e. the six areas of the simulated datasets shown in [Fig. 7](#) and the three local patches of the real datasets shown in [Fig. 9](#)). It can be seen that the runtime for the deep learning module ranges between 1–3 min and that of DLBI ranges between 30–40 min. In contrast, the runtime for the 3B analysis is around 75 h, which is more than 110 times higher than that for DLBI. Our results have demonstrated that the super-resolution images from the deep learning module alone is a good estimation to the ground-truth. Therefore, for users who value time and can compromise accuracy, the results from the deep learning module provide a good tradeoff, and thus a good estimation of the ground-truth.


### 3.4.4 Large-field reconstruction

To analyze a dataset with 200 frames, each with about  $200 \times 300$  pixels, it takes our method about 7–10 h on a single CPU core. Therefore, our method is able to achieve large-field reconstruction on the yellow areas shown in [Figure 8](#). [Figure 12](#) shows the large-field reconstruction images of the three real datasets. For the Actin1 dataset, the selected area is  $200 \times 300$  pixels and the reconstructed super-resolution image is  $1600 \times 2400$  pixels. For the Actin2 dataset, the selected area is  $250 \times 240$  pixels and the reconstructed image is  $2000 \times 1920$  pixels. And for the ER dataset, the selected area is  $200 \times 150$  pixels and the reconstructed image is  $1600 \times 1200$  pixels.

As shown in [Figure 12A and B](#), the actin networks in the two datasets have been successfully recovered by DLBI. The thinning and thickening trends of the cytoskeleton have been clearly depicted, as well as the small latent structures, including actin filaments, actin bundles and ruffles. For the endoplasmic reticulum structure ([Fig. 12C](#)), the circular-structures and connections of the cytoskeleton have also been accurately reconstructed.



**Fig. 11.** Runtime comparison of the deep learning module (DNN), the entire DLBI pipeline (DLBI) and the 3B analysis (3B) on the nine reconstruction tasks (i.e. the six areas of the simulated datasets shown in [Fig. 7](#) and the three local patches of the real datasets shown in [Fig. 9](#)). The runtime was measured on a Fedora 25 system with 128 Gb memory and E5-2667v4 (3.2 GHz) CPU



**Fig. 12.** Large-field reconstructions of the three real datasets: (A) Actin1, (B) Actin2 and (C) ER. (D) The single-molecule reconstruction of the Actin1 dataset by PALM based on 20 000 frames. (E) Overlap of the reconstruction images by DLBI (in red) and PALM (in green)

For the Actin1 dataset, the single-molecule reconstruction of the red channel is available (Fig. 12D). This reconstruction was produced by PALM (Hess *et al.*, 2006) using 20 000 frames, whereas the reconstruction image of DLBI (Fig. 12A) used only 200 frames. We further overlaid the image produced by DLBI with that of PALM to check how well they overlap (Fig. 12E). It is clear that the main structures of the two images almost perfectly agree with each other. In addition, our method was able to recover the latent structure on the top-left part, which was not photographed by PALM due to out of range of views in dual-channel photographing. If we carefully check the original low-resolution fluorescent images, we could find that this predicted structure indeed exists, which is consistent with our reconstruction.

#### 4 Discussion and conclusion

In this article, we proposed a deep learning guided Bayesian inference method for structure reconstruction of super-resolution fluorescence microscopy. Our method combines the strength of deep learning and statistical inference. We further overcame the high data requirement bottleneck of deep learning by a novel stochastic simulation module based on the experimentally-calibrated parameters and problem-specific physical models. It should be noted that although our simulation provides close-to-realistic data to train the deep learning module, it still contains bias and unrealistic parts, which will be learned by the deep learning module. Bayesian inference, on the other hand, can

correct and refine the ultrastructure learned by deep learning, and thus enhance the physical meaning of the final super-resolution image.

We have comprehensively evaluated the quality of the reconstructed super-resolution images. The future work includes evaluating how well the framework can rediscover the parameters used in the simulation module.

#### Funding

This work was supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research (OSR) under Awards No. FCC/1/1976-04, URF/1/2602-01, URF/1/3007-01, URF/1/3412-01 and URF/1/3450-01, the National Key Research and Development Program of China [2017YFA0504702], the National natural Science Foundation of China [U1611263, U1611261, 61232001, 61472397, 61502455, 61672493].

*Conflict of Interest:* none declared.

#### References

- Schwentker, M.A. *et al.* (2007) Wide-field subdiffraction RESOLFT microscopy using fluorescent protein photoswitching. *Microsc. Res. Tech.*, **70**, 269–280.
- Cox, S. *et al.* (2012) Bayesian localization microscopy reveals nanoscale podosome dynamics. *Nat. Methods*, **9**, 195–200.

- Culley, S. et al. (2018) Quantitative mapping and minimization of super-resolution optical imaging artifacts. *Nat. Methods*, **15**, 263–266.
- Dai, H. et al. (2017) Sequence2vec: a novel embedding approach for modeling transcription factor binding affinity landscape. *Bioinformatics*, **33**, 3575–3583.
- Dong, H. et al. (2017). Tensorlayer: a versatile library for efficient deep learning development. In: *Proceedings of the 2017 ACM on Multimedia Conference*, pp. 1201–1204. ACM.
- Gal, Y. and Ghahramani, Z. (2015). Dropout as a Bayesian approximation: representing model uncertainty in deep learning. *arXiv:1506.02142*.
- Ghahramani, Z. and Jordan, M.I. (1996). Factorial hidden markov models. In: *Advances in Neural Information Processing Systems*, pp. 472–478.
- Godsill, S.J. et al. (2004). Monte carlo smoothing for nonlinear time series. *J. Am. Stat. Assoc.*, **99**, 156–168.
- Goodfellow, I. et al. (2014). Generative adversarial nets. In: *Advances in neural information processing systems*, pp. 2672–2680.
- Gustafsson, M.G. (2005) Nonlinear structured-illumination microscopy: wide-field fluorescence imaging with theoretically unlimited resolution. *Proc. Natl. Acad. Sci. USA*, **102**, 13081–13086.
- He, K.M. et al. (2016). Deep residual learning for image recognition. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778.
- Hein, B. et al. (2008) Stimulated emission depletion (STED) nanoscopy of a fluorescent protein-labeled organelle inside a living cell. *Proc. Natl. Acad. Sci. USA*, **105**, 14271–14276.
- Hess, S.T. et al. (2006) Ultra-high resolution imaging by fluorescence photoactivation localization microscopy. *Biophys. J.*, **91**, 4258–4272.
- Holden, S.J. et al. (2011) Daostorm: an algorithm for high-density super-resolution microscopy. *Nat. Methods*, **8**, 279–280.
- Huang, F. et al. (2011) Simultaneous multiple-emitter fitting for single molecule super-resolution imaging. *Biomed. Opt. Express*, **2**, 1377–1393.
- Johnson, J. et al. (2016). Perceptual losses for real-time style transfer and super-resolution. In: *European Conference on Computer Vision*, pp. 694–711. Springer.
- Kim, J. et al. (2016). Accurate image super-resolution using very deep convolutional networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1646–1654.
- Kingma, D. and Ba, J. (2014). Adam: a method for stochastic optimization. *arXiv preprint arXiv: 1412.6980*.
- Ledig, C. et al. (2016). Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv: 1609.04802*.
- Li, Y. et al. (2018) DEEPre: sequence-based enzyme EC number prediction by deep learning. *Bioinformatics*, **34**, 760–769.
- Lidke, K.A. (2012) Super resolution for common probes and common microscopes. *Nat. methods*, **9**, 139.
- Lim, B. et al. (2017). Enhanced deep residual networks for single image super-resolution. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, Vol. 1.
- Lippincott-Schwartz, J. and Manley, S. (2009) Putting super-resolution fluorescence microscopy to work. *Nat. Methods*, **6**, 21–23.
- Quan, T. et al. (2011) High-density localization of active molecules using structured sparse model and bayesian information criterion. *Opt. Express*, **19**, 16963–16974.
- Rabiner, L.R. (1989) A tutorial on hidden markov models and selected applications in speech recognition. *Proc. IEEE*, **77**, 257–286.
- Ram, S. et al. (2006) Beyond rayleigh's criterion: a resolution measure with application to single-molecule microscopy. *Proc. Natl. Acad. Sci. USA*, **103**, 4457–4462.
- Russakovsky, O. et al. (2015) ImageNet large scale visual recognition challenge. *Int. J. Comput. Vision*, **115**, 211–252.
- Rust, M.J. et al. (2006) Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM). *Nat. Methods*, **3**, 793–796.
- Sage, D. et al. (2015) Quantitative evaluation of software packages for single-molecule localization microscopy. *Nat. Methods*, **12**, 717–724.
- Salimans, T. et al. (2016). Improved techniques for training GANs. *CoRR*, abs/1606.03498.
- Sangkloy, P. et al. (2016) The sketchy database: learning to retrieve badly drawn bunnies. *ACM Trans. Graph*, **35**, 1.
- Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv: 1409.1556*.
- Small, A.R. (2009) Theoretical limits on errors and acquisition rates in localizing switchable fluorophores. *Biophys. J.*, **96**, L16–L18.
- Xu, F. et al. (2015) Bayesian localization microscopy based on intensity distribution of fluorophores. *Protein Cell*, **6**, 211–220.
- Xu, F. et al. (2017) Live cell single molecule-guided bayesian localization super resolution microscopy. *Cell Res.*, **27**, 713.
- Zhang, M. et al. (2012) Rational design of true monomeric and bright photoactivatable fluorescent proteins. *Nat. Methods*, **9**, 727–729.
- Zhu, L. et al. (2012) Faster STORM using compressed sensing. *Nat. Methods*, **9**, 721–723.