

Teaching deep neural networks to localize sources in super-resolution microscopy by combining simulation-based learning and unsupervised learning

Artur Speiser ^{*1,2,3}, Srinivas C. Turaga ^{†4,**}, Jakob H. Macke ^{‡1,2,**}

¹Computational Neuroengineering, Department of Electrical and Computer Engineering, Technical University of Munich, Munich, Germany

²research center caesar, an associate of the Max Planck Society, Bonn, Germany

³International Max Planck Research School ‘Brain and Behavior’, Bonn/Florida

⁴HHMI Janelia Research Campus, Ashburn VA, USA

**equal contribution

Abstract

Single-molecule localization microscopy constructs super-resolution images by the sequential imaging and computational localization of sparsely activated fluorophores. Accurate and efficient fluorophore localization algorithms are key to the success of this computational microscopy method. We present a novel localization algorithm based on deep learning which significantly improves upon the state of the art. Our contributions are a novel network architecture for simultaneous detection and localization, and a new training algorithm which enables this deep network to solve the Bayesian inverse problem of detecting and localizing single molecules. Our network architecture uses temporal context from multiple sequentially imaged frames to detect and localize molecules. Our training algorithm combines simulation-based supervised learning with autoencoder-based unsupervised learning to make it more robust against mismatch in the generative model. We demonstrate the performance of our method on datasets imaged using a variety of point spread functions and fluorophore densities. While existing localization algorithms can achieve optimal localization accuracy in data with low fluorophore density, they are confounded by high densities. Our method significantly outperforms the state of the art at high densities and thus, enables faster imaging than previous approaches. Our work also more generally shows how to train deep networks to solve challenging Bayesian inverse problems in biology and physics.

Introduction

Super-resolution microscopy techniques such as stochastic optical reconstruction microscopy (STORM, [1]) and photo-activated localization microscopy (PALM, [2]) have made it possible to observe biological structures and processes that were not accessible through optical microscopy due to the limitations posed by the Abbe diffraction limit. These methods critically rely on computational methods for accurately localizing point-spread functions (PSF) in low resolution images of sparsely activated fluorophores [3]. State-of-the-art localization algorithms typically operate in two steps: first, single fluorophore candidates are detected and

*artur.speiser@tum.de

†turagas@janelia.hhmi.org

‡macke@tum.de

extracted from the images, and second, fluorophores are localized by fitting a high resolution "generative" model of the PSF to the image. To deal with overlapping fluorophores, peaks are either rejected based on a statistical test for the presence of multiple fluorophores (single emitter fitting [4, 5, 6]), or emitters are added throughout the fitting procedure until a predetermined threshold for the goodness of fit is met (multi-emitter fitting [7, 8, 9]). More recently deep learning approaches have been used to perform the localization step [10, 11].

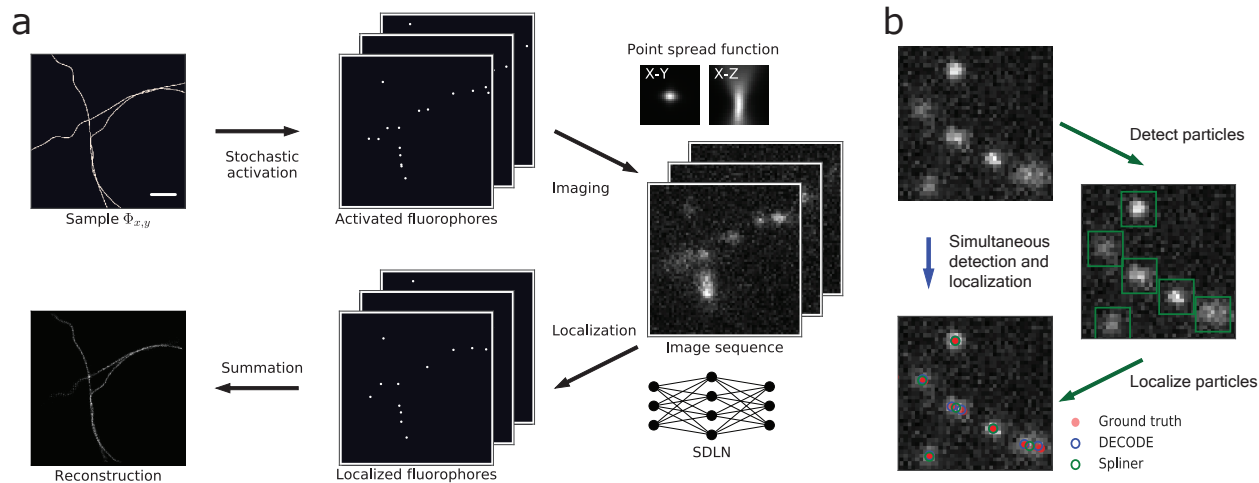


Figure 1: **Source reconstruction for Single-Molecule Localization Microscopy (SMLM):** (a) Fluorophores are stochastically activated and recorded using fluorescence microscopy. A localization algorithm is used to infer the underlying sources from noisy and blurry imaging measurements, and summing over the inferred sources yields an estimate of the underlying structure. (b) Classical image-processing algorithms for SMLM source localization (such as Spliner [7]) are based on a two-step approach (detect/localize), whereas our approach (DECODE) uses a neural network for simultaneous detection and localization.

This general approach can be highly effective under favourable conditions of high SNR and low fluorophore density [12]. However, even multi-emitter approaches produce sub-par results in datasets with high fluorophore densities. As was noted in [13], they perform even worse than single-emitter algorithms for 3D data. These limitations imply that current fitting-based approaches to SMLM can not be applied to experiments with high emitter densities, which would be critical for the investigation of living or moving structures.

Furthermore, most previous algorithms base their predictions on a single observed image. Thus, they ignore potentially useful information in the sequence of imaging frames which can enable detecting and separating fluorophores in crowded high density data by taking into account temporal dynamics. Nevertheless, attempts at using information from multiple images during inference are rare [14, 15], and have not yet yielded state-of-the-art performance.

Deep learning methods have revolutionized computer vision, and biological image analysis is no different [16, 17, 18, 19]. However, many of these advances are the result of the supervised training of deep neural networks using large training datasets of pairs of example input images and desired output predictions. The two first applications of deep learning to SMLM, Deep-STORM [20] and DeepLoco [21], took a similar approach. These two methods simulated synthetic SMLM data and trained deep networks to localize single molecules, an approach we call "simulator learning". However, such an approach should be used with caution, since imperfect simulations can significantly impair the performance of the network. Simultaneously

with our work, [22] presented a deep-learning approach using simulator learning for optimizing point-spread functions and 3D source reconstruction from high-density imaging data.

We here present a novel deep learning based approach called DECODE for fast and efficient single-molecule localization which achieves state of the art performance. First, we show how combining simulator learning with a variational ‘‘autoencoder learning’’ method (AEL) [23, 24] for Bayesian model inversion can increase the robustness of the deep network. Second, we develop a network architecture which simultaneously detects and localizes single molecules. This single DECODE network is trained to produce accurate predictions at both low and high densities, alleviating the need for analysis methods which deal with these ‘single emitter’ and ‘multi emitter’ cases separately. Third, our DECODE network uses local temporal context contained within the temporally neighboring frames to detect and localize emitters in any given frame.

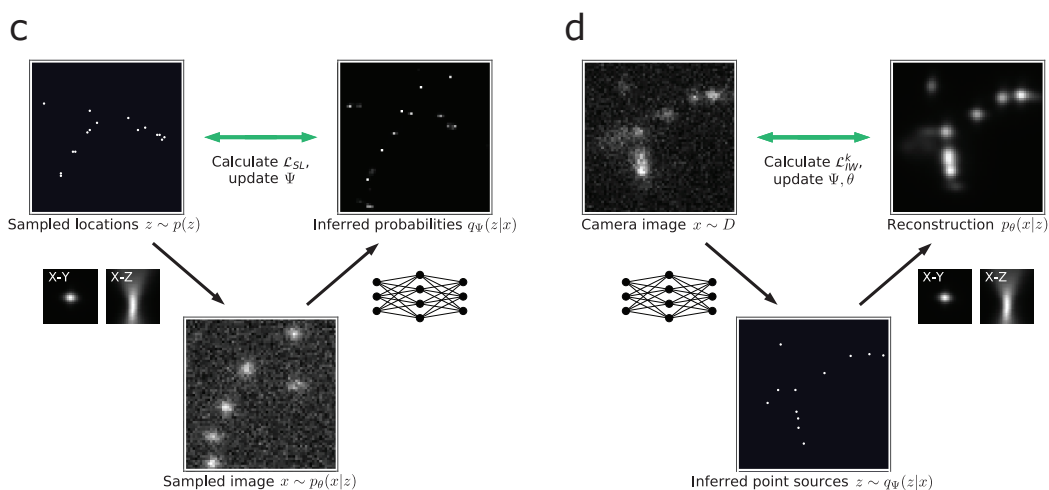


Figure 2: **Approaches for training a neural network to localize sources:** (c) Simulator learning (SL). Synthetic images are constructed by the simulated imaging of randomly located fluorophore point sources using a generative model, and a network is trained to detect and localize the fluorophores using supervised learning. (d) Auto-encoder learning (AEL). A neural network used to infer putative locations from a measured camera image, and subsequently the generative model is used to reconstruct the original camera image. Both the parameters of the generative model and of the DECODE network are optimized.

Results

Our method has three main components: First, the generative or forward model, which is a simplified model of the imaging process, and specifies how the noisy images measured by the cameras can be generated from the underlying fluorophores. This component is an essential part of all SMLM reconstruction algorithms and is described in the supplement. It includes a parametric description of the point spread function of the microscope, the image acquisition on the pixels of a camera, and the noise associated with the whole measurement process. The second component is our DECODE network for simultaneous detection and localization of fluorophores. And third, we describe how we jointly train both the generative model and the DECODE network to optimize the performance of the algorithm.

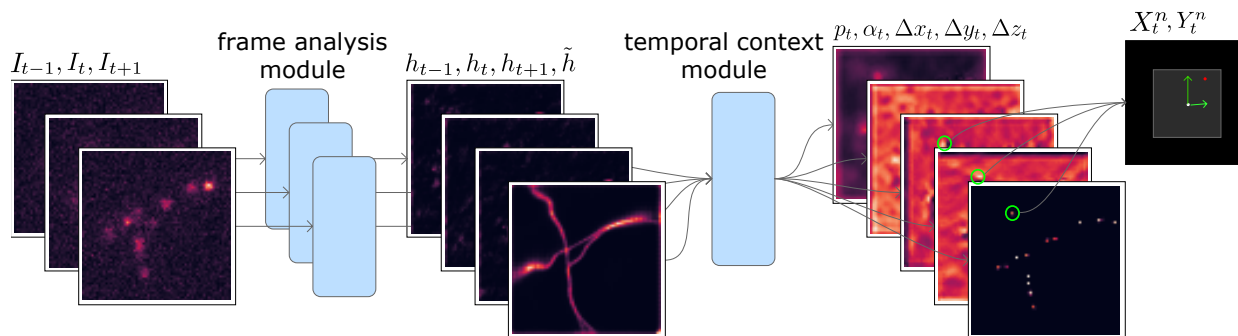


Figure 3: **DECODE network for simultaneous detection and localization of fluorophores.** Hidden features are extracted from each consecutive imaging frame by the first stage of the network by the “frame analysis module”. These frame specific features are integrated by the “temporal context module” leading to a prediction of five output maps. The first output map is a binary map of fluorophore detections. The second output map predicts the brightness of the corresponding detected fluorophore. The final three output maps are the predicted coordinates of the detected fluorophore, relative to the center of the detected pixel.

DECODE network for simultaneous detection and localization of fluorophores

Our goal is to design and train a deep learning based function approximator to simultaneously detect and localize an unknown number of fluorophores. The input to the deep network is a sequence of imaging frames containing sparsely activated fluorophores, and the outputs are locations of an unknown and changing number of fluorophores active from frame to frame.

Output representation. Previous deep learning approaches have treated imaging frames independently and have taken one of two extreme approaches, which we combine in our work. Deep-STORM [20] directly predicts super-resolved images using a U-net architecture [25]. This approach essentially consists of many independent classifiers, one for each super-resolution pixel, indicating the detection of a fluorophore in that pixel. This has two main drawbacks: First, their method is restricted to 2D localization and implicitly constrains localizations to a discrete lattice of super-resolution pixels, and it is computationally inefficient to extend this strategy to 3D detection and localization since the number of voxels predicted scales with the volume. Second, it can’t make use of the large amount of post-processing methods available for tasks like drift correction and visualization, as the output does not conform to the usual format of a list of localizations [26, 27]. In contrast, DeepLoco [21] effectively combines classification and regression by predicting a fixed sized 256×4 matrix for each imaging frame, with each row representing the presence or absence of a fluorophore, followed by the 3D vector containing the x , y and z coordinate of the molecule. This has the advantage that the computational complexity of the output scales only with the maximum number of possible active fluorophores in any frame, but not the volume of imaged field of view. However, it requires the network to learn a highly non-local and non-linear transformation from images into 3D coordinates in an undetermined ordering.

Our DECODE network architecture is a hybrid of the two extremes and is based on stacking two U-nets [25] which predict five channels for each imaged frame with the same pixel size as the imaging system (fig. 3, 1.4). The first channel indicates probability of the presence or absence of a fluorophore within that pixel in an imaging frame. The second channel predicts the brightness of the detected fluorophore, and the remaining three channels collectively describe the continuous valued x , y , z coordinate of the fluorophore detected

within a particular pixel. Our approach has the limitation that it cannot detect or localize more than one fluorophore per imaging pixel, however it is challenging to accurately detect and localize simultaneously active fluorophores less than a pixel-width apart anyway. In a post-processing step we cluster connected probabilities to obtain our final list of localizations (see 1.5). This approach allows us to produce lists of localizations in three dimensions, while preserving the spatial correlation of input and output which makes the transformation easier to learn.

Using temporal and spatial context. We introduce a new mechanism to integrate information across time, and show that improved detection and localization for a specific frame of an imaging sequence can be achieved by taking into account other frames. The signals contained in individual imaging frames are not independent: First, the activation dynamics of the fluorophores are such that a fluorophore can be active across multiple adjacent frames, inducing correlations which are local in time. Second, the fluorophore locations are not randomly distributed, but are concentrated around the biological structures which have been labeled by the fluorophores, which induces correlations which are global in time. For these reasons, we designed neural networks which use local and global temporal context to exploit local and global correlations in the images. The DECODE network architecture (fig. 3) is able to make use of both forms of context by integrating the inferred hidden states from three consecutive images, as well a running average as input. We show that local context has an especially large positive impact on performance (section).

Combining simulator learning and autoencoder learning

Deep learning is traditionally the domain of big data, black-box function approximation, and supervised learning, where we replace our ignorance of the relationship between inputs and outputs with labeled data. In contrast, the prediction problem in localization microscopy is a Bayesian inverse problem where we have detailed knowledge of the forward stochastic generative process, which describes how the measurement is the result of the noisy imaging of sparsely activated single molecule fluorophores. However, we do not have access to large amounts of labeled data in the form of ground truth localization data. One possible solution, which both Deep-STORM and DeepLoco use, is to train the network on simulated data – given a forward generative description of the imaging process, one can simulate a large synthetic dataset from the model and use it for supervised learning. We call this approach “simulator learning”.

The effectiveness of simulator learning depends critically on the availability of an accurate generative model at training time, as deviations of the true forward model from the simulated forward model can degrade performance significantly. This problem can be solved by simultaneously estimating the parameters of the true forward model, and training the detection and localization network using the real measured data, rather than a fictitious simulation. This is possible using the recently developed framework of variational autoencoders (VAEs) [23, 24]. In the VAE framework, the stochastic forward generative model and the DECODE network are stacked to form a stochastic autoencoder. This autoencoder is then used to simultaneously optimize the parameters of the deep network and the forward model, with the goal of achieving image-reconstructions which are similar to the original measurement. Formally, this can be achieved by maximizing a so-called ‘evidence-lower bound’ via stochastic gradient optimization (see 1.6 for details). VAEs have previously been used on the related problem of inferring action potentials from calcium imaging data [28]. A drawback of VAE-approaches are that gradients for network training need to be approximated using Monte Carlo sampling, which can make optimization more challenging.

We empirically compared simulator-learning (SL) and autoencoder learning (AEL) to optimize our novel

network architecture. We found that pure simulator learning performs best when the true underlying model is known, but that variational learning is more robust to model-mismatch. Finally, we show that the reweighted wake-sleep algorithm, which alternates between the two methods [29] (AEL+SL), combines the advantages of both approaches.

Quantitative evaluation on simulated datasets from the SMLM challenge show DECODE outperforming all algorithms across a variety of conditions

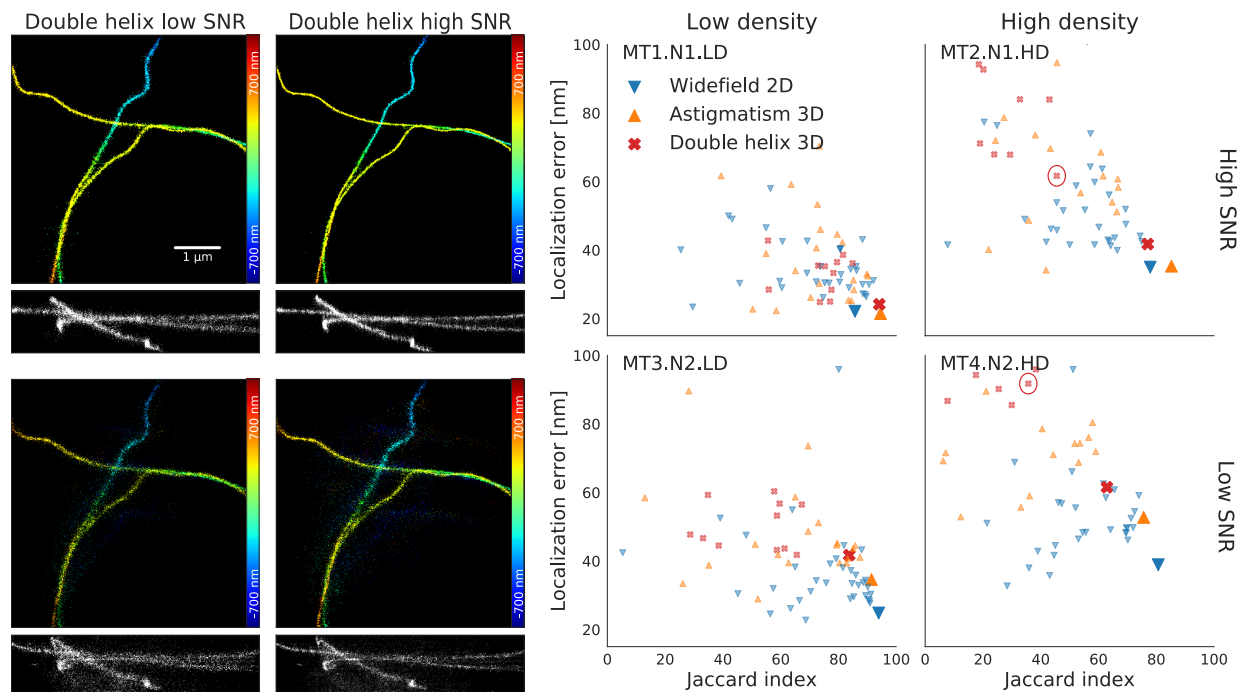


Figure 4: **Performance comparison on simulated microtubules from the SMLM2016 challenge. (a)** Reconstructions by DECODE and the Spliner algorithm on high density double helix challenge data. Upper panels show the color-coded $x - y$ view, lower panels the $x - z$ cross-section. Images are rendered as histograms with 5 nm pixel size, localizations are convolved with Gaussians with $\sigma = 15$ nm. **(b)** Performance evaluation on the twelve test datasets with low/high density, low/high SNR and different modalities using the Jaccard index (higher is better) and localization error (either lateral or volume, lower is better) as metrics. Each marker indicates one benchmarked algorithm, large solid markers indicate DECODE, red circles indicate Spliner results for the conditions in **(a)**.

Two contests have been set up to enable objective, quantitative evaluations of the plethora of available localization algorithms [30, 13]. The 2016 SMLM challenge¹ offers synthetic datasets for training and evaluation that were created to emulate various experimental conditions. We compared the reconstructions obtained with DECODE to the multi-emitter fitting approach Spliner [7] on two 3D double-helix datasets with high fluorophore densities². Visual comparison shows that DECODE achieves superior resolution and produces far fewer spurious localizations (fig. 4a).

¹<http://bigwww.epfl.ch/smlm/challenge2016/index.html?p=datasets>

²We used settings provided by the authors: <https://github.com/ZhuangLab/storm-analysis>

We compared DECODE on multiple data-sets to all other contenders, using the data, benchmark-implementations and performance metrics provided by the challenge. We analyzed the performance of DECODE on the low (N2) and high (N1) signal-to-noise (SNR) datasets, with low (LD) or high (HD) emitter density for 2D, Astigmatism (AS) and Double Helix (DH) modalities. We quantify performance using the lateral or volume localization error in nm for 2D and 3D data respectively and the Jaccard index J which measures how well an algorithm does at detecting all the fluorophores while avoiding false positives. DECODE achieves state of the art performance in all settings when compared to all other algorithms that submitted their results on the test datasets to the challenge so far (fig. 4b). The difference is particularly large in difficult conditions, i.e. for high density and low SNR. Details of the evaluation procedure are described in [13], but our algorithm was submitted after the closing deadline for the publication describing the challenge. The results, including extensive evaluations and side by side comparisons are available online³. A direct comparison with the top contenders, taken from the website is shown figure 8. This shows that DECODE outperforms other approaches across most performance metrics and applicable data-sets, and often by a substantial margin.




Figure 5: **Performance comparison on real data.** (a) Reconstruction of one section of the Tubulin-A647-3D dataset. (b) Comparison of reconstructions obtained for sub-sampled datasets analyzed with DECODE and Spliner.

Evaluation on a real 3D dataset of microtubules imaged with dSTORM

We also trained the DECODE (using AEL+SL training and local context) algorithm on a real 3D AS dataset of microtubules imaged with dSTORM. Details of the imaging procedure are provided in [4]. The experimental conditions roughly correspond to high SNR and low density settings modelled in the challenge datasets. To highlight the robustness of DECODE under more difficult conditions we generated artificial datasets with lower SNR and higher densities. To this end we subsampled the dataset by summing a number of non consecutive images. Specifically, the first image of the 6x subsampled dataset was created by summing frames 1, 301, 601, 901, 1201, the second image used frames 2, 302, 602, 902, 1202 and so on. This way local context is preserved while drift within a single image remains negligible. This procedure not only

³<http://bigwww.epfl.ch/smlm/challenge2016/leaderboard.html>

increases emitter density, but also reduces the SNR as the noise is summed while the signal per emitter stays constant. Fig. 5a show the reconstruction for a subsection of the original dataset obtained with DECODE. In fig. 5b we analyze how the quality of the reconstruction deteriorates for 6x and 12x subsampled datasets and compare to reconstructions achieved with the Spliner algorithm. Similarly to the challenge datasets we observe a sharper image and less spurious localizations for the original dataset, with the difference becoming more marked under more challenging conditions.

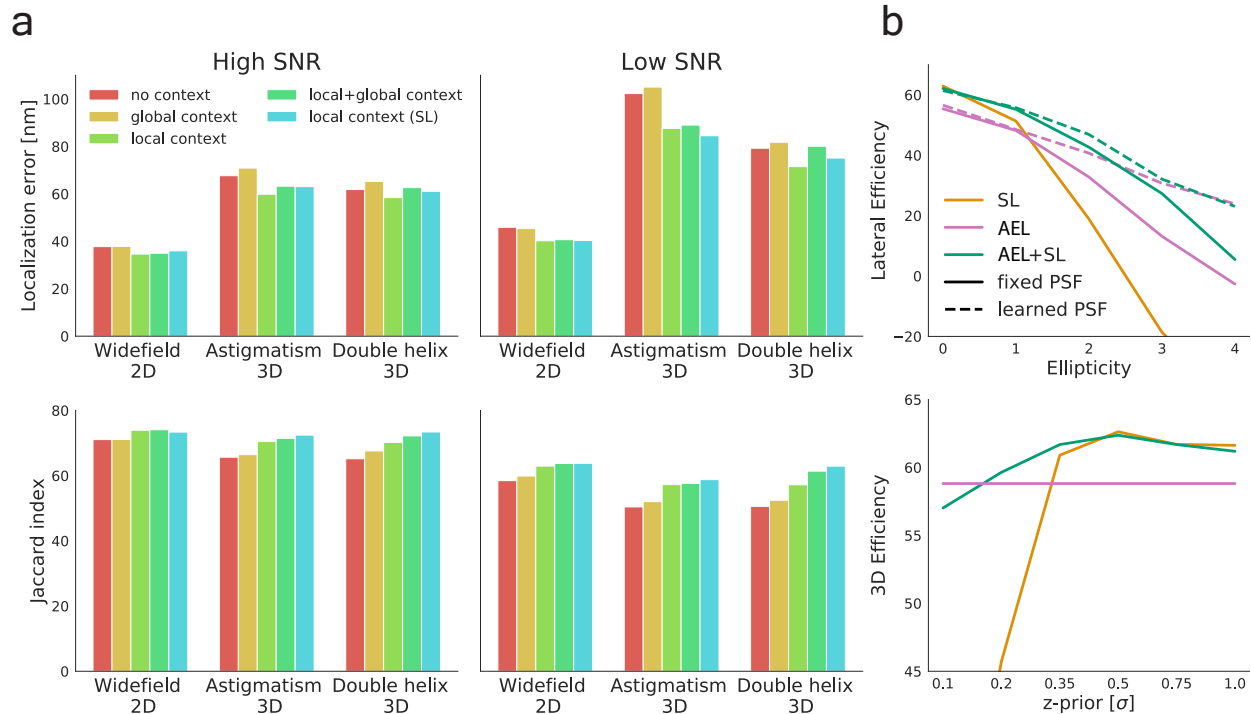


Figure 6: **DECODE performance for different settings.** (a) DECODE evaluated on the 6 high density challenge training datasets. Models were trained using AEL+SL either without context, local, global and both forms of context. Additionally we trained a model with SL and local context. (b) **Performance of different training methods for different degrees of model mismatch.** **Upper panel: PSF mismatch** Models using a circular PSF are fit to 5 datasets simulated from PSFs with varying ellipticity. PSF parameters for AEL / AEL+SL learning could be either fixed (solid line) or learned (dashed line). **Lower panel: Prior mismatch.** We created a dataset by simulating data with our astigmatism PSF, sampling axial offset values from a normal distribution with $\sigma = 0.5$. We trained models using SL / AEL+SL where the simulator used σ values between 0.1 and . We also trained a network purely unsupervised (AEL). The PSF parameters were fixed at ground truth values for all models. Performance is evaluated using the efficiency metric introduced in [13] which weights the RMSE and Jaccard values to obtain a single indicator.

Features of the algorithm and their impact on performance

Local and global context. We investigated how different features and possible settings for training DECODE models impact algorithm performance. To probe the impact of local and global context we trained DECODE models using AEL+SL learning on the 6 high density datasets of the challenge, using either no context, local context, global context or both forms of context. Local and global context have different ef-

fects on the performance (Fig. 6a). On average, the addition of local context increases the Jaccard index, but often at the cost of reduced precision: While global context can increase the chance to correctly identify dim fluorophores in regions where the likelihood of structure being present is high, it cannot provide information on the exact location as the structure has a certain extent in space. Local context, on the other hand, yields improvements in both metrics as it increases the number of collected photons available for inference when fluorophores are active over multiple frames.

Simulator and autoencoder learning. To highlight the difference between simulator and autoencoder learning we tested how the two approaches behave for different degrees of model mismatch. Given our generative model, which first samples the fluorophore localizations and their intensity and then evaluates the PSF model, mismatch can stem from two sources which we consider separately – a misspecified PSF model, or wrong assumptions about the prior distributions of the fluorophores with respect to their intensity, density, blinking behaviour and distribution along the x, y and z axis. To simulate PSF mismatch we generated datasets using elliptical Gaussians with increasing degrees of ellipticity as PSF functions and then trained DECODE models that use a symmetric Gaussian (Upper panel fig. 6, solid lines show performance of models with fixed generative parameters.) For an ellipticity of zero the models have access to the true underlying generative model. In this case SL training sets an upper bound to the achievable performance with a given network as it has access to an infinite amount of labelled data, and outperforms AE. However, pure simulator learning is more sensitive to parameter mismatch as it never ‘sees’ elliptical PSFs during training. Unsupervised learning will still try to infer the correct positions as placing the circular PSF into the middle of the elliptic one achieves the best reconstruction. Alternating between the two methods retains the advantages of both: Performance is virtually the same as simulator learning when there is no mismatch and performance degrades more gracefully when the mismatch is increased. When we add training of the generative model parameters (dashed lines), δ_{xy} learns to account for some of the PSF-mismatch, which further increases performance. To test the influence of an incorrect prior we generated data using our astigmatism model with axial offset values sampled from a normal distribution with $\sigma_{z_i} = 0.5$. We then trained SL and AEL+SL models using prior distributions with σ_{z_i} between 0.1 and 1. All other parameters are set to ground truth values. We also trained one AEL model as training uses only real data and is therefore independent of σ_{z_i} . We observed similar properties as in the PSF experiment (fig. 6b, lower panel). Simulator learning is specifically sensitive to an underestimated prior range. AEL+SL is more robust to mismatch and performs better than AEL learning which suffers from high gradient variance.

Discussion

We developed a new deep learning based method for the simultaneous detection and localization of fluorophores for single molecule localization microscopy. Our contributions are a novel deep network architecture and a method for training the network in a fully unsupervised manner. We pursued two main goals when designing our network architecture: first, to directly transform images into continuous location estimates without an intermediate ROI identification step, and second, to integrate both local and global temporal context into the inference procedure. We showed that the combination of both of these properties enables our algorithm to outperform other approaches significantly on high density data. We demonstrated the superiority and generality of our method on the benchmark SMLM challenge across a variety of point spread functions, densities, and signal-to-noise ratios.

We observed that local temporal context was highly beneficial, and increases both the number of correctly identified fluorophores and the localization precision. This result is in contrast with the observation in [15], which reported that local context offered no substantial advantage over global context. These differences in

result could be a consequence of the way local context is incorporated in their algorithm, or a consequence of the fact that their simulated data was limited to discrete lines.

To train our networks we used the ‘reweighted wake-sleep algorithm’ [29], which combines two different training approaches – simulator learning (SL) and variational auto encoder learning (AEL). We found that this combined approach worked at least as well as each of the two approaches in isolation, and sometimes substantially better. The benefits of this algorithm compared to other approaches for training deep latent variable model were recently highlighted in other domains as well [31]. For the problem of SMLM, the forward stochastic generative model is strongly constrained by the physics of the imaging process, and involves only a small number of parameters. It is therefore possible to generate simulated data which is close to the real data distribution throughout training. Therefore we believe reweighted wake-sleep to be especially advantageous to solve Bayesian inverse problems in these settings.

SMLM is a computational microscopy technique which depends critically on good algorithms for the reconstruction of a super-resolution image. Advances in reconstruction algorithms (partially fueled by advances in deep learning), can improve resolution and speed up image acquisition through improvements in reconstructions on high-density data.

Acknowledgments

This work was supported by the German Research Foundation (DFG) through SFB 1089, the German Federal Ministry of Education and Research (BMBF, project ‘ADMIMEM’, FKZ 01IS18052 A-D), and the Howard Hughes Medical Institute. We thank Jonas Ries, Lucas Müller, Daniel Sage and Christopher Obara for useful discussions, and David Greenberg for comments on the manuscript.

References

- [1] M. J. Rust, M. Bates, and X. Zhuang, “Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (storm),” *Nature methods*, vol. 3, no. 10, p. 793, 2006.
- [2] E. Betzig, G. H. Patterson, R. Sougrat, O. W. Lindwasser, S. Olenych, J. S. Bonifacino, M. W. Davidson, J. Lippincott-Schwartz, and H. F. Hess, “Imaging intracellular fluorescent proteins at nanometer resolution,” *Science*, vol. 313, no. 5793, pp. 1642–1645, 2006.
- [3] H. Deschout, F. C. Zanicchi, M. Mlodzianoski, A. Diaspro, J. Bewersdorf, S. T. Hess, and K. Braeckmans, “Precisely and accurately localizing single emitters in fluorescence microscopy,” *Nature methods*, vol. 11, no. 3, p. 253, 2014.
- [4] Y. Li, M. Mund, P. Hoess, J. Deschamps, U. Matti, B. Nijmeijer, V. J. Sabinina, J. Ellenberg, I. Schoen, and J. Ries, “Real-time 3d single-molecule localization using experimental point spread functions,” *Nature methods*, vol. 15, no. 5, p. 367, 2018.
- [5] S. Wolter, A. Löschberger, T. Holm, S. Aufmkolk, M.-C. Dabauvalle, S. Van De Linde, and M. Sauer, “rapidstorm: accurate, fast open-source software for localization microscopy,” *Nature methods*, vol. 9, no. 11, p. 1040, 2012.
- [6] P. Dedecker, S. Duwé, R. K. Neely, and J. Zhang, “Localizer: fast, accurate, open-source, and modular software package for superresolution microscopy,” *Journal of biomedical optics*, vol. 17, no. 12, p. 126008, 2012.

- [7] H. P. Babcock and X. Zhuang, “Analyzing single molecule localization microscopy data using cubic splines,” *Scientific reports*, vol. 7, no. 1, p. 552, 2017.
- [8] H. Babcock, Y. M. Sigal, and X. Zhuang, “A high-density 3d localization algorithm for stochastic optical reconstruction microscopy,” *Optical Nanoscopy*, vol. 1, no. 1, p. 6, 2012.
- [9] M. Ovesný, P. Křížek, J. Borkovec, Z. Švindrych, and G. M. Hagen, “Thunderstorm: a comprehensive imagej plug-in for palm and storm data analysis and super-resolution imaging,” *Bioinformatics*, vol. 30, no. 16, pp. 2389–2390, 2014.
- [10] T. Kim, S. Moon, and K. Xu, “Information-rich localization microscopy through machine learning,” *Nature communications*, vol. 10, no. 1, p. 1996, 2019.
- [11] P. Zelger, K. Kaser, B. Rossboth, L. Velas, G. Schütz, and A. Jesacher, “Three-dimensional localization microscopy using deep learning,” *Optics express*, vol. 26, no. 25, pp. 33166–33179, 2018.
- [12] B. Rieger and S. Stallinga, “The lateral and axial localization uncertainty in super-resolution light microscopy,” *ChemPhysChem*, vol. 15, no. 4, pp. 664–670, 2014.
- [13] D. Sage, T.-A. Pham, H. Babcock, T. Lukes, T. Pengo, R. Velmurugan, A. Herbert, A. Agarwal, S. Colabrese, A. Wheeler, *et al.*, “Super-resolution fight club: A broad assessment of 2d & 3d single-molecule localization microscopy software,” *bioRxiv*, p. 362517, 2018.
- [14] S. Cox, E. Rosten, J. Monypenny, T. Jovanovic-Talisman, D. T. Burnette, J. Lippincott-Schwartz, G. E. Jones, and R. Heintzmann, “Bayesian localization microscopy reveals nanoscale podosome dynamics,” *Nature methods*, vol. 9, no. 2, p. 195, 2012.
- [15] R. Sun, E. Archer, and L. Paninski, “Scalable variational inference for super resolution microscopy,” *bioRxiv*, p. 081703, 2016.
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- [17] C. Belthangady and L. A. Royer, “Applications, promises, and pitfalls of deep learning for fluorescence image reconstruction,” 2018.
- [18] T. Ching, D. S. Himmelstein, B. K. Beaulieu-Jones, A. A. Kalinin, B. T. Do, G. P. Way, E. Ferrero, P.-M. Agapow, M. Zietz, M. M. Hoffman, *et al.*, “Opportunities and obstacles for deep learning in biology and medicine,” *Journal of The Royal Society Interface*, vol. 15, no. 141, p. 20170387, 2018.
- [19] M. Weigert, U. Schmidt, T. Boothe, A. Müller, A. Dibrov, A. Jain, B. Wilhelm, D. Schmidt, C. Broadus, S. Culley, *et al.*, “Content-aware image restoration: pushing the limits of fluorescence microscopy,” *Nature methods*, vol. 15, no. 12, p. 1090, 2018.
- [20] E. Nehme, L. E. Weiss, T. Michaeli, and Y. Shechtman, “Deep-storm: super-resolution single-molecule microscopy by deep learning,” *Optica*, vol. 5, pp. 458–464, Apr 2018.
- [21] N. Boyd, E. Jonas, H. P. Babcock, and B. Recht, “Deeploco: Fast 3d localization microscopy using neural networks,” *bioRxiv*, 2018.
- [22] E. Nehme, D. Freedman, R. Gordon, B. Ferdman, T. Michaeli, and Y. Shechtman, “Dense three dimensional localization microscopy by deep learning,” *arXiv e-prints*, p. arXiv:1906.09957, Jun 2019.

- [23] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [24] D. J. Rezende, S. Mohamed, and D. Wierstra, “Stochastic backpropagation and approximate inference in deep generative models,” *arXiv preprint arXiv:1401.4082*, 2014.
- [25] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” pp. 234–241, 2015.
- [26] M. El Beheiry and M. Dahan, “Visp: representing single-particle localizations in three dimensions,” *Nature methods*, vol. 10, no. 8, p. 689, 2013.
- [27] T. Pengo, S. J. Holden, and S. Manley, “Palmsiever: a tool to turn raw data into results for single-molecule localization microscopy,” *Bioinformatics*, vol. 31, no. 5, pp. 797–798, 2014.
- [28] A. Speiser, J. Yan, E. W. Archer, L. Buesing, S. C. Turaga, and J. H. Macke, “Fast amortized inference of neural activity from calcium imaging data with variational autoencoders,” *Advances in Neural Information Processing Systems 30*, pp. 4024–4034, 2017.
- [29] J. Bornschein and Y. Bengio, “Reweighted wake-sleep,” *CoRR*, vol. abs/1406.2751, 2014.
- [30] A. Small and S. Stahlheber, “Fluorophore localization algorithms for super-resolution microscopy,” *Nature methods*, vol. 11, no. 3, p. 267, 2014.
- [31] T. A. Le, A. R. Kosiosek, N. Siddharth, Y. W. Teh, and F. Wood, “Revisiting reweighted wake-sleep,” *arXiv preprint arXiv:1805.10469*, 2018.
- [32] B. Huang, W. Wang, M. Bates, and X. Zhuang, “Three-dimensional super-resolution imaging by stochastic optical reconstruction microscopy,” *Science*, vol. 319, no. 5864, pp. 810–813, 2008.
- [33] S. R. P. Pavani, M. A. Thompson, J. S. Biteen, S. J. Lord, N. Liu, R. J. Twieg, R. Piestun, and W. Moerner, “Three-dimensional, single-molecule fluorescence imaging beyond the diffraction limit by using a double-helix point spread function,” *Proceedings of the National Academy of Sciences*, vol. 106, no. 9, pp. 2995–2999, 2009.
- [34] A. Aitken, C. Ledig, L. Theis, J. Caballero, Z. Wang, and W. Shi, “Checkerboard artifact free sub-pixel convolution: A note on sub-pixel convolution, resize convolution and convolution resize,” *arXiv preprint arXiv:1707.02937*, 2017.
- [35] S. Stallinga and B. Rieger, “Accuracy of the gaussian point spread function model in 2d localization microscopy,” *Optics express*, vol. 18, no. 24, pp. 24461–24476, 2010.
- [36] Y. Burda, R. Grosse, and R. Salakhutdinov, “Importance weighted autoencoders,” *arXiv preprint arXiv:1509.00519*, 2015.
- [37] A. Mnih and D. J. Rezende, “Variational inference for monte carlo objectives,” in *Proceedings of the 33rd International Conference on Machine Learning*, 2016.

Supplementary Information

1 Methods

1.1 Generative model for data generation and AE learning

We assume space to be discretized with the resolution of the EMCCD camera. When recording images of a fluorescent specimen Φ_{xy} , the probability for a fluorophore to be activated at pixel x, y and at time t is given by

$$S_{t,x,y} \sim \text{Bernoulli}(p_{act}\Phi_{x,y} + \gamma S_{t-1,x,y}), \quad (1)$$

$$(2)$$

where p_{act} is the activation rate of all fluorophores within this pixel. Fluorophores are frequently active across multiple frames. We model this behaviour through the survival probability γ . The distribution for an image of a specimen of n activated fluorophores S_n is

$$I_t \sim \text{Gamma}((\sum_n^i A_n S_n + \beta)/\eta, \eta). \quad (3)$$

Here, β is the baseline activity and $A_n = \alpha_n L \cdot \text{PSF}(Z_n)$ is a matrix implementing convolution with a normalized PSF, down sampling and scaling by the number of photons emitted by a fluorophore during one frame times the camera gain which we express as the maximum number of photons emitted L times the fraction of the timeframe the photon is active. The light emitted from a source follows a Poisson distribution (shot noise). For an EMCCD camera these counts are convolved with a gamma distribution that models the EM gain and with a Gaussian distribution that accounts for read-out noise. This distribution can not be expressed analytically. This is irrelevant when doing simulator learning, however AE learning requires a differentiable expression. Therefore, we approximate the noise model using a single gamma distribution with a scale parameter η that can be estimated from the background distribution of the images.

1.2 PSF models

We focused on 2D and 3D datasets that were obtained with the astigmatism (AS, [32]) and the double helix (DH, [33]) technique. AS data is created with a cylindrical lens to distort the PSF into an elliptical shape depending on the axial position of a fluorophore. For DH data a dichroic beam splitter is used to split the PSF into two components that rotate around the actual fluorophore position. The PSF models that we use for these modalities are:

$$PSF_{2D} = \sum_n^2 (\omega_n \mathcal{N}^2(\{X_n, Y_n\}, (1 + \text{softplus}(Z_n)) \cdot \sigma_n)) + \delta_{xy}(Z_n) \quad (4)$$

$$PSF_{AS} = \mathcal{N}^2(\{X_n, Y_n\}, \{\sigma_x(Z_n), \sigma_y(Z_n)\}) + \delta_{xy}(Z_n) \quad (5)$$

$$\sigma_{x,y}(z) = \sigma_0 \sqrt{1 + \frac{z - c_{x,y}}{d}}^2 \quad (6)$$

$$PSF_{DH} = \mathcal{N}^2(\{X_n - x(Z_n), Y_n - y(Z_n)\}, \sigma_n) + \mathcal{N}^2(\{X_n + x(Z_n), Y_n + y(Z_n)\}, \sigma_n) + \delta_{xy}(Z_n) \quad (7)$$

$$x(z) = r \cdot \cos(\beta \cdot z) \quad y(z) = r \cdot \sin(\beta \cdot z) \quad (8)$$

X_n, Y_n, Z_n indicate the position of the n_{th} fluorophore within the image. We use a weighted sum of two two-dimensional circular Gaussians for the 2D case. The more their width is increased, the further out of focus the particle lies. For AS data, we use an elliptical PSF with diagonal variance. We adopted the expression for the dependency of σ_x, σ_y on z from [8]. The DH helix data is modelled as two circular Gaussians that rotate around the fluorophore position in a radius r when changing Z . δ_{xy} are pixel maps that allow the model to approximate details of the PSF that can't be easily captured by a parametric model (for example diffraction rings). For each model we determine the maximum axial range ahead of training and then learn one map per 100 nm. During training the learned maps are indexed by the inferred offset variable. Trilinear interpolation is used to smooth the transition between two maps and allow for sub-pixel lateral shifts.

1.3 Fitting of 3D PSFs from bead stacks

When performing 3D inference, it is important to calibrate the PSF model on data with known axial offset as the exact relationship between its shape and the position cannot be estimated from unlabeled data. We estimated the AS and DH PSF's using calibration bead stacks which are images of single fluorophores at different offsets with very high signal to noise ratios. We first obtain a rough estimate of the bead locations using a basic peak-finding routine. We then maximize the likelihood $p_\theta(x|z)$ by performing stochastic gradient descent on the exact x - and y - coordinates of each bead (which are constant across images), the shape parameters of the PSF model (4) and the pixel maps δ_{xy} . This simple method achieves localization errors of less than 0.3 nm on the challenge calibration stacks where the ground truth locations are available. During training of the DECODE network we keep the PSF model fixed. On the other hand, for 2D datasets the PSF model can be learned together with the network parameters in a completely unsupervised way. Example fits for the different modalities are shown in fig. 7. We emphasize that neither the training algorithm, nor the network architecture, depends on the specifics of the generative model or the PSF model, and both could well be combined with more flexible functional forms of PSFs.

1.4 Simultaneous detection and localization network

Architecture Our recognition model consists of two stacked U-net architectures ([25]). The first one extracts features from each single image, while the second takes hidden states from three consecutive frames (local context) and the average of the hidden state of all images (global context) $\tilde{h} = \sum_T h_t$ as input.

$$h_t = f_1(I_t) \tag{9}$$

$$p, \alpha, \Delta x, \Delta y, \Delta z = f_2(h_t, h_{t-1}, h_{t+1}, \tilde{h}) \tag{10}$$

The DECODE network predicts five image channels of the same dimensionality as the imaging frame. This prevents the correct identification of multiple fluorophores within one pixel, however this occurs rarely and for typical experimental conditions fluorophores that are that close together are not separable. p is the probability for an active fluorophore in the respective pixel. To obtain continuous localizations that are not limited by the imaging resolution we additionally predict offset variables $\Delta x, \Delta y$ that determine the exact position of an active fluorophore within a pixel. Δz is the axial shift from the imaging plane and α the fraction of imaging time the fluorophore was active.

Our U-nets are built from downsampling blocks consisting of two convolutional layers and a strided convolution, and upsampling blocks, also with two convolutional and a deconvolution layer. The first U-net, which extracts hidden features from each input image, consists of two down- and upsampling blocks. The

second one, which combines the features, has one down- and upsampling block. The regular convolutional layers use 3x3 filters and exponential linear units. A final layer with linear activations is used for the continuous outputs. For the upsampling layers we found it to be important to use the ICNR initialization to avoid checkerboard patterns in the outputs [34]. Another layer with a sigmoid nonlinearity outputs the probabilities for the discrete activations. It is important to use a high negative initial bias (around -6), otherwise the number of inferred fluorophores will be very high at the start of the training process which might cause memory issues as the generative model has to compute a PSF for each sampled fluorophore.

Training Training requires accurate estimates of the baseline background activity β and the noise scale parameter η . We first identify the 10% of pixels in the image sequence with the lowest mean activity assuming that these regions don't contain any fluorescent sample. Then we fit a gamma distribution to the histogram of all intensity values of those pixels. When working with real data background activity cannot be assumed to be homogeneous. We therefore use a sliding window approach to obtain time and space dependent background variables.

Training is performed on 32×32 regions randomly selected from the images at each iteration. This speeds up training, increases the efficiency of the importance weights for AE training and acts as a form of data augmentation. If the network is trained to make use of global context, we use a running average of the hidden states collected over the last 100 training batches. At test time we perform two passes over the dataset: the first one to collect the average hidden state $\tilde{h} = \sum_T h_t$ and the second one to obtain the inference results.

When training with local context we employ different strategies for SL and AE training steps. For simulator learning, when sampling data we align the variables $\Delta x, \Delta y, \Delta z$ to be identical when a fluorophore is active in consecutive frames. For AE training, for each set of variables $S_t, \Delta x_t, \Delta y_t, \Delta z_t$ which are inferred from the images I_{t-1}, I_t, I_{t+1} we also infer variables for $t + 1$ and $t - 1$ to provide context.

We use these variables to calculate an error term that is the sum of square errors between the offset variables at each pixel with consecutive activations:

$$\delta_{xyz} = \sum_{x,y} S_t \cdot S_{t-1} \{(\Delta x_t - \Delta x_{t-1})^2 + (\Delta y_t - \Delta y_{t-1})^2 + (\Delta z_t - \Delta z_{t-1})^2\} \quad (11)$$

$$+ S_t \cdot S_{t+1} \{(\Delta x_t - \Delta x_{t+1})^2 + (\Delta y_t - \Delta y_{t+1})^2 + (\Delta z_t - \Delta z_{t+1})^2\} \quad (12)$$

This term is subtracted from the ELBO during training.

1.5 Obtaining localizations at test time

Our localization network outputs probabilities of a fluorophore being located at a specific pixel. While we use samples from this distribution during training, is unsuitable to obtain localizations at test time. Given that the probabilities are factorized we would obtain varying number of fluorophores when sampling from probabilities that are spread out over multiple pixels. While we could use our generative model to calculate likelihoods for different samples and perform importance weighted resampling to get more realistic samples, this would drastically increase the computation time. To get deterministic, fast and precise pseudo samples we instead use a variant of non-maximum suppression to obtain final localizations. To obtain a binary mask of fluorophore candidates for a given frame we identify probability peaks, i.e. pixels with values that are above 0.3 and higher than all values in a surrounding 3x3 patch. We then add the probability mass from the directly adjacent 4 pixels to the values at the candidate positions by convolving the probability map

with a cross shaped filter and applying the mask. All candidates with added probability values above 0.7 are counted towards the localizations. For each inferred discrete position the offset variables $\Delta x, \Delta y, \Delta z$ from the 5 contributing pixels are weighted by the corresponding probabilities and averaged to get better estimates. The algorithm can be expressed purely in the form of pooling and convolution operations and therefore runs efficiently on a GPU.

For difficult imaging conditions, i.e. high densities, low SNR values the lateral offset variables can be biased towards small absolute values. The reason for this is that during SL training the true values of $\Delta x, \Delta y$ are sampled from a uniform distribution with values in $[-0.5, 0.5]$. When the uncertainty is high the network is biased to predict values closer to zero because ground truth values are never outside of these borders. This effect can produce artifacts in the reconstructed image as localizations are concentrated at the pixel centers. This effect is specifically present for fluorophores far away from the imaging plane. To counteract this we therefore divide all localizations into equally sized bins according to their axial offset. Then we calculate an empirical CDF \hat{F}_x, \hat{F}_y from the histograms of the Δx and Δy variables in each bin. The variables $\hat{x}_{os}, \hat{y}_{os} = \hat{F}_x(\Delta x) - 0.5, \hat{F}_y(\Delta y) - 0.5$ have a uniform distribution as desired. This transformation effectively removes image artifacts while having minimal impact on the performance metrics.

1.6 Gradient estimation

To develop our methods for training DNNs to perform fluorophore localization we frame the problem as an instance of a generative model with discrete latent variables. Our images $x = x_1, \dots, x_N$ are samples from the true and unknown marginal distribution $p(x)$. We know that the images are dependent on the currently activated fluorophores. This relationship is captured by the generative model $p_\theta(x, z) = p_\theta(x|z)p(z)$. Here the prior gives the distribution of fluorophore activations, and the likelihood describes how convolving the latent positions with a PSF function yields the recorded images. Our main goal is to learn an inference network $q_\psi(z|x)$ which takes images as input and transforms them into conditional distributions over the latent variables which approximates the posterior distribution $q_\psi(z|x) \sim p(z|x)$. Given a sufficient amount of labeled training data $\{(x_1, z_1), \dots, (x_N, z_N)\}$ we could train our network in a supervised fashion, however for SMLM such datasets are not available. Therefore we look at two alternative methods to train $q_\psi(z|x)$.

Simulator learning The physical process that our generative models describes is well understood [35] and can be approximated using models with a small number of parameters θ . Therefore, even with a minimal amount of labeled data we can obtain good estimates of $p_\theta(x, z)$ which allows us to sample realistic data. We can use this data as a surrogate for real labeled data to train our network. For this we take samples from our generative model $z \sim p(z), x \sim p_\theta(x|z)$ and maximize the loglikelihood $\log q_\psi(z|x)$. This procedure amounts to minimizing the Kullback-Leibler divergence (D_{KL}) between the posterior of the generative model and the recognition network, averaged over the (simulated) data distribution:

$$\mathbb{E}_{p_\theta(x)}[-D_{KL}(p_\theta(z|x)||q_\psi(z|x))] = \mathbb{E}_{p_\theta(x,z)}[\log q_\psi(z|x)] + \text{const.} \quad (13)$$

Optimizing a lower bound on $p(x)$ Using Jensen’s inequality we can derive a lower bound (ELBO) on the marginal likelihood $p(x)$:

$$\log p(x) = \log \mathbb{E}_q\left[\frac{p_\theta(x, z)}{q_\psi(z|x)}\right] \geq \mathbb{E}_q\left[\log \frac{p_\theta(x, z)}{q_\psi(z|x)}\right] = \mathcal{L}(x) \quad (14)$$

by maximizing this ELBO with respect to θ we minimize the reverse D_{KL} averaged over the true data distribution

$$\mathbb{E}_{p(x)}[D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z}|\mathbf{x}))] \quad (15)$$

Unlike simulator learning maximization with respect to ψ also allows us to learn the parameters of the generative model.

If we instead use an importance weighted average over k samples from our recognition model to estimate $p(x)$, and again apply Jensen’s inequality we obtain a tighter lower bound (which is identical to the ELBO for $k = 1$):

$$\mathcal{L}_{IW}^k(x) = \mathbb{E}_{q(z_{1:K}|x)} \left[\frac{1}{K} \sum_{k=1}^K \log \left[\underbrace{\frac{p_\theta(x, z_k)}{q_\phi(z_k|x)}}_{\omega_k(x, z_k)} \right] \right] \quad (16)$$

This objective is the basis for both the importance weighted autoencoder (IWAE) [36] and the reweighted wake-sleep algorithm (RWS) [29].

Updating θ For a given value of ψ , unbiased gradients for θ can be obtained by sampling $z_1, \dots, z_k \sim q_\phi$ and calculating the gradients:

$$\nabla_\theta \mathcal{L}_{IW}^k(x \sim \mathcal{D}) = \nabla_\theta \log \left(\frac{1}{K} \sum_{k=1}^K \omega_k \right) = \sum_{k=1}^K \tilde{\omega}_k \nabla_\theta \log p_\theta(z_k|x) \quad (17)$$

$$\tilde{\omega}_k = \frac{\omega_k(x, z_j)}{\sum_{k'=1}^K \omega(x, z_{k'})} \quad (18)$$

Updating ψ Obtaining gradients for ψ is more involved, especially in the case of discrete latents when the reparametrization trick cannot be applied.

The RWS algorithm includes two procedures to obtain gradients for ψ . The sleep phase update matches simulator learning 1.6 which minimizes the D_{KL} between p and q over data that is generated from the generative model $p_\theta(x|z)p(z)$.

The wake phase update optimizes the same D_{KL} , but over the true data distribution $p(x)$ (i.e. using samples from the data).

$$\nabla_\psi D_{KL}(p_\theta(z|x)||q_\psi(z|x)) \simeq \sum_{k=1}^K \tilde{\omega}_k \nabla_\psi \log q_\theta(z_k|x) \quad (19)$$

We also experimented with the VIMCO algorithm [37] that uses REINFORCE gradients and a per sample baseline to reduce variance. However RWS proved to perform slightly better.

1.7 Learning algorithms

Simulator learning works well in practice when good generative models are available as is the case for SMLM imaging. However, as noted in [21], this approach is sensitive to any mismatch between the training and test distribution because the network never "sees" any real data during training. On the other hand, a pure variational autoencoder approach, as used in [28], is optimized over the true data distribution. It is therefore more robust to model mismatch. On the negative side, the training procedure is less straightforward than simulator learning, especially when using discrete latents: more hyper parameter tuning is necessary for the algorithm to converge, and the final results might be worse due to higher variance of the gradient estimates. Furthermore, depending on the amount of available data, overfitting might be a concern.

Given these observations, it is natural to ask if the two approaches can be combined to gain an advantage. RWS includes simulator learning in the form of the sleep phase ψ update.

We therefore compare the following algorithms:

Simulator learning (SL) Pure simulator learning using 13

Variational Learning (AE) Maximizing $\nabla_{\psi} \mathcal{L}_{TW}^k(x)$ using 17 and 19.

Combined learning (AE+SL) The full RWS algorithm using 17, 19 and the sleep updates 13.

2 Supplementary Figures






Figure 7: **PSF model fits** PSFs were fitted with the the method described in section 1.3. Contour plots show the underlying parametric model and highlight the contribution from the pixel maps δ_{xy} .

Efficiency 2D/3D



Jaccard



RMSE Lateral/Volumetric in nm

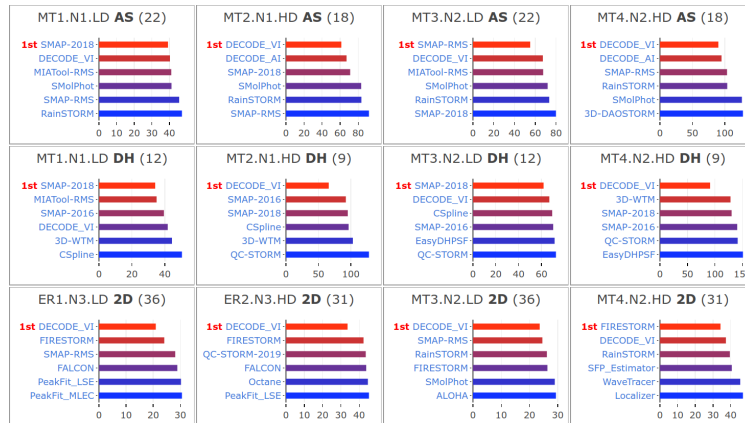


Figure 8: Performance metrics on challenge data Top 6 submissions to the SMLM Challenge 2016 for different modalities and imaging conditions when ordering them by the 2D/3D efficiency metric [13], Jaccard index and RMSE. DECODE(AI) corresponds to amortized inference (AE training using the training set, evaluation on the test set), DECODE(VI) to variational inference (training and evaluation on the test set). DECODE leads on all datasets on the efficiency metric which quantifies the payoff between Jaccard and RMSE.