PHILOSOPHICAL TRANSACTIONS A

rsta.royalsocietypublishing.org

Research

Article submitted to journal

Subject Areas:
light microscopy, machine learning

Keywords:
microscopy, deep learning, structure illumination, RCAN, SIM

Author for correspondence:
Miguel Boland, mdb119@ic.ac.uk

Improving axial resolution in SIM using deep learning

Boland, Miguel A.; Cohen, Edward A. K.; Flaxman, Seth R.; Neil, Mark A. A.

Department of Mathematics
Imperial College
South Kensington Campus
180 Queen's Gate
London SW7 2RH
United Kingdom

Structured Illumination Microscopy is a widespread methodology to image live and fixed biological structures smaller than the diffraction limits of conventional optical microscopy. Using recent advances in image up-scaling through deep learning models, we demonstrate a method to reconstruct 3D SIM image stacks with twice the axial resolution attainable through conventional SIM reconstructions. We further demonstrate our method is robust to noise & evaluate it against two point cases and axial gratings. Finally, we discuss potential adaptions of the method to further improve resolution.

THE ROYAL SOCIETY PUBLISHING

© The Author(s) Published by the Royal Society. All rights reserved.

1. Introduction

Structured Illumination Microscopy (SIM) is a super-resolution technique which enables a two-fold increase in lateral resolution (X/Y axis, perpendicular to the line of sight of the microscope) when compared to conventional fluorescence microscopy [1] [2]. SIM functions via the illumination of structured light onto a specimen to obtain spectral information that would be out of the range visible to a wide-field microscope. Repeating this procedure across a set of structured light patterns constructs a stack of raw SIM images of a specimen, the ensemble of which contains a wider spectral range of frequencies than obtained via a single exposure. Image stacks are then reconstructed into a single super-resolved image using computational algorithms. A wide range of illumination patterns can be used, including 3D patterns that allow for improvements in axial resolution (Z axis, parallel to line of sight of the microscope) [3]. This has proven to be a useful investigative tool for the analysis of live biological processes as it avoids most sample damage traditionally associated with electron microscopy [4]. Example applications include in-vivo imaging of synapses in mice [5], or analysing the distribution of proteins on chromosomes, a task previously unfeasible due to the relative scale of the protein and a confocal microscope's diffraction limit [6]. As such, the further improvement of image reconstruction methodologies for microscopy may allow the observation of previously unobserved live biological processes.

Improvements in axial resolution have historically been achieved by advances in optical setup and illumination techniques, such as 3D-SIM [7], TIRF-SIM [8], spot-scanning SR-SIM [9] or 2-photon SR-SIM [10]. As these techniques can incur significant costs and complexity, recent developments in deep learning have been adapted and optimised to address these issues from a purely computational approach. Proposed deep learning methods for SIM have successfully demonstrated reductions in aberrations and reconstruction artifacts due to movements of the specimen [11] [12] or low-light conditions [13] [14] [15], but as of yet, deep learning networks have not been applied to the problem of reconstructing SIM image stacks to a resolution beyond that attainable by standard SIM methodologies. Training a network to take as its input a set of SIM images as might be produced by a 2D-SIM microscope setup (2-beam, 3-angle) and with a target that has the resolution characteristics of a 3D-SIM microscope, would allow the network to achieve an axial resolution comparable to 3D SIM microscopes without the motion artifacts or photo-bleaching caused by high image counts [16].

As the range of imaged light frequencies is reduced in this microscope setup, the use of convolutional layers will allow us to infer additional information (as is currently done through Richardson-Lucy deconvolution [17]). Building on the work of deep learning networks applied to SIM reconstructions, a modified Residual Channel Attention Network (RCAN) [18] deep neural net is trained to reconstruct SIM image stacks at double the axial resolution than that attainable via the original SIM reconstruction methodology [2]. High resolution images are simulated as confocal images with an illumination wavelength reduced by a factor of $\sqrt{2}$ to match the lateral resolution obtained via seminal SIM reconstructions, whilst axial resolution is improved by decreasing the axial slice thickness of each image in the stack. Simulated training and test data is generated from both dense and biologically realistic 3D point clouds, at both standard resolution (SR) and high resolution (HR) as detailed in Section 3. The network is evaluated against test data to investigate the method's effectiveness at resolving structures composed of point illuminations (Section 4). As we expect the network to output spatial frequencies which are not inherent in the input data (which is already achieved through deconvolution techniques [17]), the network should be able to infer these from the widefield images using the mapping from simulated widefield images to simulated high resolution confocal images. Further analysis is performed that demonstrates robustness to low-lighting conditions (Section 4). Limitations in network generalisability and the broader potential of the work is discussed in Section 5.

2. Existing work

Computational methods for SIM image reconstruction have set benchmarks of attainable resolution, and demonstrated flexibility to imprecise or varying illumination patterns. Examples of these are blind-SIM [19], in which gradient descent is used to minimise the least squared error between the measured data and data predicted by the convolution of an estimated biological sample with an estimated illumination pattern. By alternatively fixing the estimated data and estimated illumination pattern, a high resolution image can be retrieved with limited knowledge of the illumination pattern but at the cost of significant compute time. Similarly, the use of Richardson-Lucy deconvolutions [17] to further improve the quality of the seminal SIM algorithm [2] have demonstrated significant improvements in resolution when experimentally measuring the PSF of the optical apparatus.

In contrast, deep learning has been applied to SIM reconstruction with varying goals, such as accelerating the imaging process by reducing the required number of raw frames (i.e the number of variations in illumination patterns) and inferring the spectral information from a smaller subset of SIM frames. These demonstrate an objective of high-throughput imaging, where a known microscope configuration is coupled with deep learning methods to minimise the imaging and processing time whilst improving the obtained resolution.

This has been achieved using U-Nets [14] and CycleGAN (a variation of a Conditional Generative Adversarial Network) [15]. These diverge from our goal which is to attain a higher axial resolution whilst maintaining a $2\times$ increase in lateral resolution. Image super-resolution for generic photography has progressed considerably in the past decade, resulting in a large variety of architectures and techniques [20]. Notably, these form paradigms regarding the pre or post upsampling of inputs (relative to convolutional or other image processing), methods of image interpolation, residual learning, network recursiveness or parallel paths in addition to a large number of architecture-specific features.

Amongst these exists the Residual Channel Attention Network (RCAN) [18], which is engineered to use inter-channel features (RGB colors in the original network) in addition to sub-pixel upsampling. These are particularly suited to SIM reconstruction, as this process entails the recombination of multiple raw images (i.e channels) into an up-scaled image. An RCAN for SIM reconstruction was successfully trained using simulated SIM stacks of generic stock imagery [13], which forms the basis for our architecture.

3. Methods

(a) Deep learning

(i) Network architecture

The network architecture is based on the RCAN [18] which was previously used to implement a form of SIM reconstruction without up-scaling [13]. We modified the network to re-implement the lateral up-scaling module found in the original implementation, therefore doubling the lateral output image size. The up-scaling module was necessary to generate images comparable to our target data, which is Nyquist sampled without aliasing. The use of an up or down-scaling algorithm to adapt the target to the network output or vice-versa could cause undesirable patterns to be learned by the network, whilst using the up-scaling module between learned layers allows for some learned adaption to the underlying mapping of the low-resolution image stacks and high-resolution confocal target images.

Furthermore, the first 2D convolutional layer was replaced by a 3D convolution in order to infer information from other SIM stacks within the same processing chunk (see section below for chunk-generating procedure) prior to higher order operations (which take place in the residual blocks of the network). The number of output channels was also changed according to the chunk

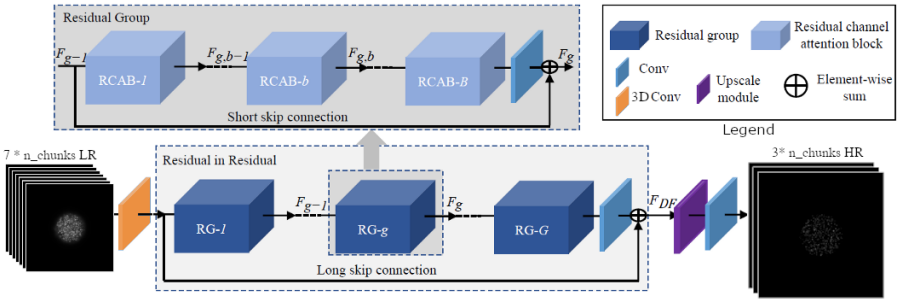


Figure 1: Modified RCAN network architecture, adapted from [18]. Note the use of a 3D convolution at the input, and the use of a single input channel and 3 output channels.

size, such that 3 consecutive gray scale (i.e. single channel) images were produced for every inputted SIM stack.

The final network configuration utilised 12 groups of 3 residual blocks, a kernel size of 3x3 and a chunk size of 3; these parameters were chosen experimentally as detailed in Supplementary Materials - 6 - Network architecture & Optimisation.

(ii) Chunk processing

Input SIM stacks were processed in chunks to allow varying depths of image stacks to be compatible with the network and improve the model's performance by inferring relations between consecutive axial images. In this scheme, a chunk size of two corresponds to a network input of 14 images (7 frames per SIM image * 2 chunk size) and a network output of 6 frames (3 frames per SIM image * 2 chunk size). Chunk size was determined experimentally by modifying the training data and input channels for the architecture described in Section 3; these were then evaluated using 24 withheld chromatin structure images and a series of images of spherical point clouds with increasing point densities. The evaluation of chunk size on network performance can be found in the supplementary materials Section A5.

The chunk size had no significant effect on the network's Mean Squared Error (MSE) (Mann-Whitney U-Test, $p>0.05$ for all pairs of chunked data), but the Structural Similarity Index (SSIM, a comparative metric of conservation of structural information between an image and reference image [21]) was higher for a chunk size of 3 (U-Test, $p<0.05$ for 6/9 comparisons). This indicates that a chunk size of 3 maximises the conservation of the structure of the images, as reflected by the mean lateral FWHM (smaller than all other chunk size's means, U-Test, $p<0.05$) and axial mean (smaller than all other chunk sizes, U-Test, $p<0.05$).

(iii) Data pre & post processing

Pairs of SIM image stacks and HR confocal images were axially sliced to produce chunked sets of images as described above. As the simulated images extended beyond the simulated underlying data, image pairs with little signal (mean pixel value below 5×10^{-7}) were discarded. The remaining set of chunked images was then normalised to a range of *[0-1]*.

A total of 100 chromatin structures and 68 spherical datasets were simulated to generate 168 image stacks. Of these, 10 chromatin structures and 20 spherical datasets were reserved for a test dataset, leaving 138 stacks for training & validation. The input/output images were then axially chunked into sets of 3 SIM images, generating 12 training datapoints per stack for a total of 1656 training points. As the simulated stacks contain little signal at the start and end of the image stacks, chunks with a mean pixel value (measured as a 16-bit float) below 10^{-7} are discarded, leaving a total of 1178 training points and 130 validation points.

Negative pixel values did occur in the RCAN outputs due to small normally distributed random errors in output values; as a majority of test images are zero-valued backgrounds, many of the errors cause small negative values. The output was therefore post-processed to truncate negative values, which is in effect the equivalent to using an additional *reLu* layer on the network but without causing a loss of learning gradient.

(iv) Training methodology

All deep learning models were implemented using the PyTorch library (version 1.5) [22] in Python 3.7. The starting learning rate was set to 10^{-6}, was adjusted if the validation rate plateaued during training through the use of the *ReduceLROnPlateau* method provided by PyTorch, and a further 1% decrease was applied every 5 epochs to ensure fine tuning in later epochs. Network gradients were clipped at $\pm$ 0.1 to avoid exploding gradients. Training / validation sets were drawn from all generated data except for withheld test cases, which were randomly selected prior to training. Training took approximately 8 hours on a set of 2 Nvidia GTX 2080 TI GPUs. Over-fitting does not appear to have occurred in our training scheme as demonstrated by the continual decrease in mean MSE values for the validation dataset (see Supplementary Materials A6).

(b) Data simulation

The model was trained using simulated SIM image stacks as input, and an up-sampled confocal image as ground truth. The methods used to generate these and benchmark SIM reconstructions are described below.

(i) SIM image stacks

Raw SIM images were generated from 3D point cloud data; these simulated a biologically realistic chromatin structure [23] (approximately bounded in a square of 16×16 µm) or a point cloud encompassed in a sphere (radius of 5 µm) with varying numbers of point emitters. The simulated microscope utilised a 7-frame setup based on the interference of 3 coherent beams to produce a 2D hexagonal illumination pattern on the sample; this is well illustrated in Figure 2 D-F of Ingerman et. al., 2019 [3]. The structured illumination is simulated as a thin light-sheet projected onto the specimen from the side [24] thereby avoiding the issues of reduced resolution enhancement usually found with this SIM geometry in an epi-illumination setup. The simulated microscope setup and image simulation approach is further described by Gong et al [25], and a mathematical description can be found in the supplementary material of Gong et al. The system was simulated with a camera size of 256x256 pixels of size 6.5 μm square, an objective magnification of 60, with a numerical aperture of 1.1 and an illumination/emission wavelength of 525 nm. The chosen illumination/emission wavelengths are not required to be identical and were chosen for convenience in this experiment. No averse effects are expected from the use of wavelengths differing by up to 10%. Photobleaching and other image quality degradation were not simulated as part of the experiment, though low-lighting conditions were simulated using Poisson statistics.

(ii) High-resolution confocal microscope simulation

Target output stacks were generated using a simulation of a confocal microscope on the same cloud point data. The parameters of a simulated confocal microscope were set to out-perform the simulated SIM microscope (halved pixel size, illumination wavelength is divided by $\sqrt{2}$, axial & lateral Nyquist sampling rate doubled, squared PSF function). These parameters ensure a doubling in both axial and lateral resolutions.

The number of output frames was also increased from 40 to 120 frames in order to improve the precision of measurements of axial resolution.

(iii) SIM reconstructions

While complete high resolution images can be reconstructed with a set of just 7 images, this hexagonal patterned SIM illumination reconstructs according to the same carrier spatial frequencies as are found in a 2-beam, 3-angle setup using 9-frames [1] [2]. As a consequence the processing of the 7-frame data to a super-resolved image with enhanced lateral resolution [25] was achieved in a similar way to the standard SIM reconstruction algorithm [2]. The estimation of reconstruction parameters was found to be imprecise due to the sub-sampling of image stacks in which the number of point emitters was limited. Constant reconstruction parameters matching those used to simulate the images were therefore used to reduce any source of variability in the reconstructed SIM images. Reconstructed SIM images were calculated to provide a comparative baseline for our RCAN-reconstructed images. This was implemented in Python (see Section 5), and was modified to output 120 frames to match the number of output frames in our high resolution confocal simulations and network outputs.

(c) Measuring resolution

The Full Width at Half Maximum (FWHM) metric has commonly been used to evaluate the resolution of microscopy images [7] [26]. It is defined as the width of a point spread function in a given direction at half the peak intensity; a large FWHM would indicate a wider point spread function (PSF) resulting in a less detailed image, whilst a smaller FWHM reflects a narrower PSF and a sharper image. Alternative methods such as fourier-ring correlation [27] could be viable to measure resolution; we instead elected to use FWHM as it provided a high-throughput measure of resolution for images of point clouds.

The FWHM was estimated using a method based on auto-correlation functions [28], thereby evaluating the FWHM for every point light source in an image rather than individually and manually selected data points. The auto-correlation function (ACF) of an image was calculated using the inverse Fourier transform of the power spectrum on individual 3D datasets. High densities of data-points can cause significant cross-correlation of signal from neighbouring PSFs; these are removed by modelling the image auto-correlation as the sum of two Gaussians. The first Gaussian models the auto-correlation of individual PSFs in the image, whilst the second models the cross-correlation between neighbouring randomly distributed PSFs. These are estimated using a bounded Least-Squares regression, and the FWHM is calculated from the Gaussian with the smallest standard deviation. An example of the modelling process is illustrated in supplementary materials Section A4. The Python & Scipy implementation of this methodology is available via GitHub, and an equivalent mathematical notation is available in the supplementary materials Section A1.

4. Results

(a) Doubling axial resolution using RCAN networks

The network's ability to double the axial resolution was tested using 24 unseen simulated chromatin structures; this is visible in Figure 2a. The mean axial FWHM for RCAN reconstructions (mean=436.8nm, std=21.8) was found to be significantly smaller than those produced using SIM (mean=669.67nm, std=207.2) (T-Test $t(10)=-3.5$, $p<0.05$). We further found the model to be robust to increasing point density when compared to standard SIM reconstructions (see Figure 2b). This is achieved without significantly affecting the lateral resolution (X & Y axis, perpendicular to line of sight of the microscope), as seen in Figure 2c, and maintaining a strong conservation of structure as seen in Figure 2d. This is illustrated by the orthogonal Y/Z view in the supplementary material Section A2 - Figure 2, where the high structural fidelity and low background noise of the

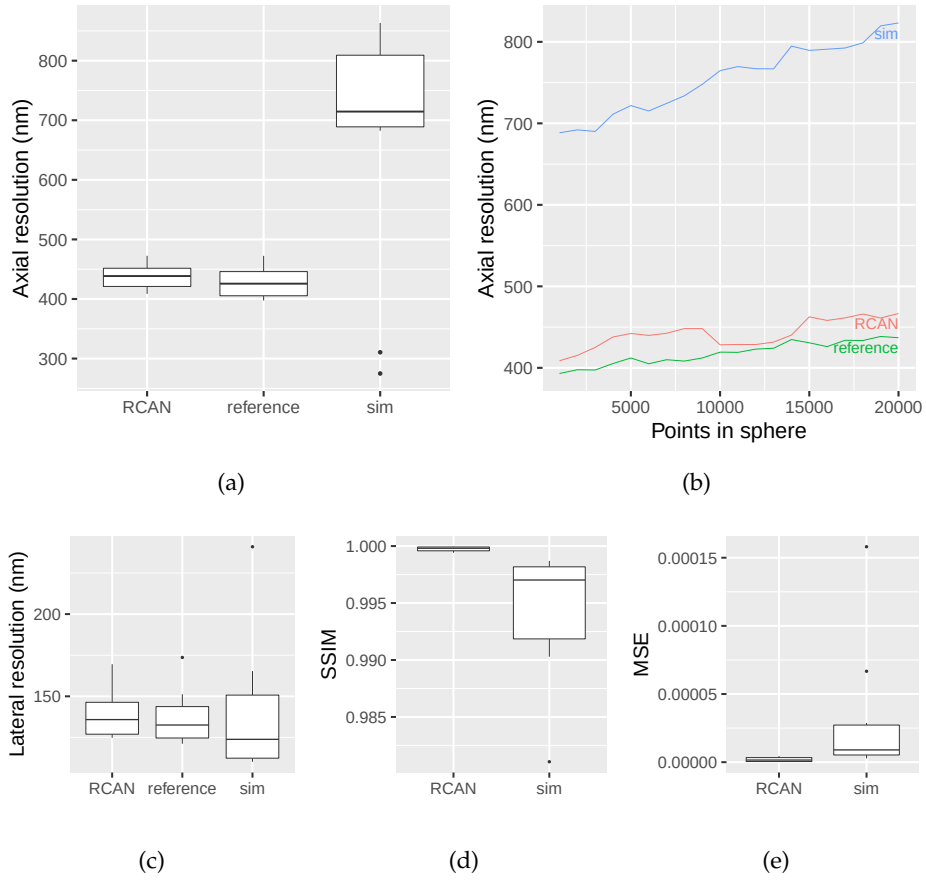


Figure 2: The proposed deep learning model (see Section 3) demonstrates a doubled axial resolution compared to a standard SIM reconstruction when evaluated against 24 simulated chromatin structures (a) (whiskers represent percentiles, boxes represent IQR). An equivalent improvement in resolution is observed when reconstructing images of point clouds of increasing densities (b). Lateral resolution is maintained when compared to SIM reconstructions (c). Finally the RCAN (mean=0.999, std=0.0002) appears to better preserve the structure of the underlying data compared to SIM (mean=0.994, std=0.006) as measured by SSIM [21] (T-test t(10)=3.06, p<0.05), although MSE was not significantly reduced (RCAN (mean=1.96e-06, std=1.69e-06), SIM (mean=3.11e-05, std=4.86e-05), T-test t(10)=-1.90, p=0.08).

RCAN output contrasts with the SIM reconstruction.

(b) Robustness in low-light imaging conditions

An additional dataset was generated using 4 other simulated chromatin structures tainted with Poisson noise. The signal-to-noise ratio was decreased by reducing the expected signal photons per emitter over 7 SIM frames, thereby simulating low-lighting conditions. The RCAN model demonstrated robustness against low-lighting conditions by maintaining a consistent axial resolution (Figure 3a) MSE (Figure 3c) and SSIM [21] (Figure 3d), in contrast to the reconstructed SIM images which diverge from the reference HR images in SSIM and MSE in worsening lighting conditions. The notable decrease in axial resolution for SIM reconstructions at lower lighting

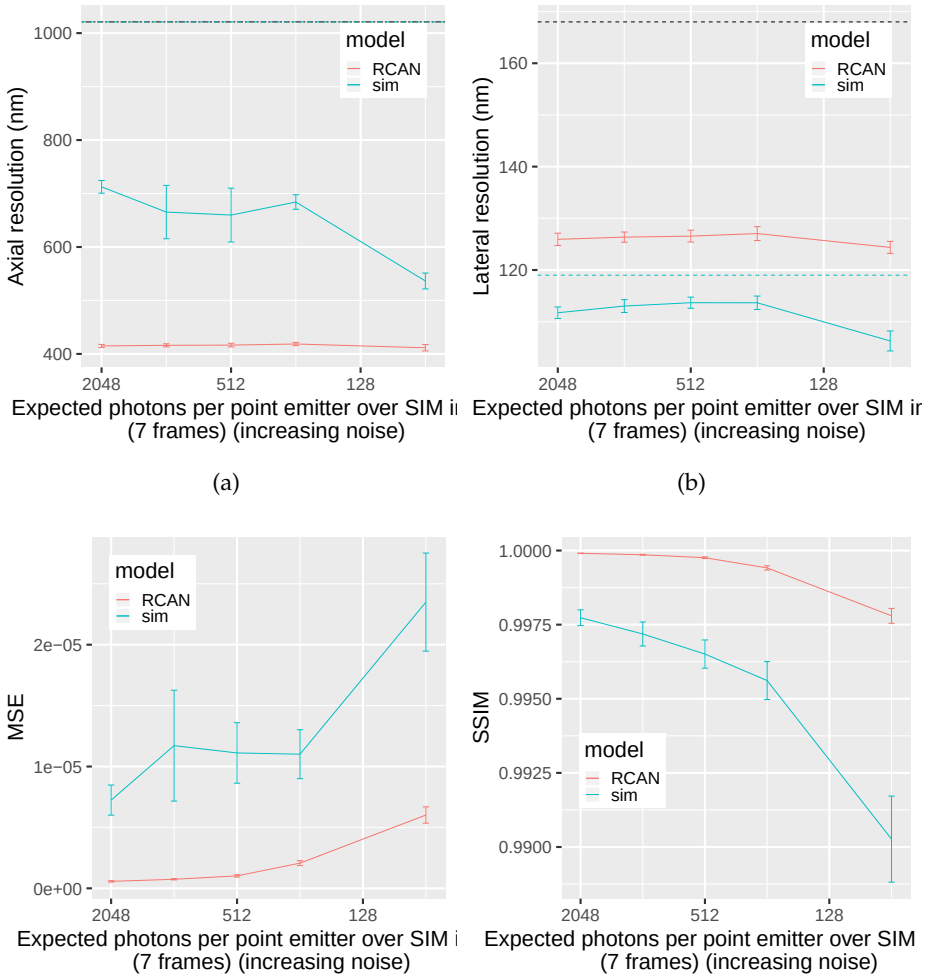


Figure 3: The effect of Poisson noise on quality of image reconstruction, run on 5 different simulated chromatin structures with 9 levels of Poisson noise. Error bars illustrate the standard error. Dotted lines in Figure 3a and Figure 3b illustrate the theoretical resolution for the wide-field images (black) and SIM reconstructions (blue). The RCAN network provides more stable axial resolutions (see figures a and b) without a major loss in image quality compared to SIM reconstructions, as shown by the divergence in MSE and SSIM (see figures c and d). The drop in resolution as noise levels increase represents a shortfall of our method for resolution estimation, as reconstructed point artifacts due to noise are not filtered when estimating the resolution. We therefore consider the measured resolution as a comparative measure of reconstruction quality rather than an attainable resolution in experimentally noisy conditions.

levels is due to a breakdown in the quality of the reconstruction, as reflected by the large increase in MSE and decrease in SSIM, whilst these are more consistent for RCAN reconstructions.

(c) Evaluation against third party data

The training data examined above differs from the biological structures observed in other SIM experiments, such as actin filaments [14], membranes [13] or other biological structures which

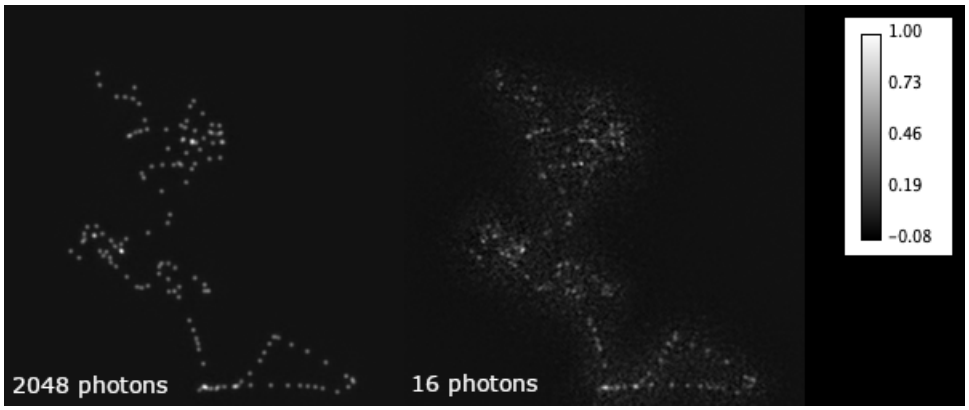


Figure 4: An example of Poisson noise corrupting a sample in low-lighting conditions. The illustrated sample is chromatin structure #10, displayed as the axially averaged pixel value and normalised for better visibility. The left image represents a relatively well illuminated image at approximately 2048 photons per emitter across 7 SIM frames, whilst the right image illustrates an excessively attenuated lighting condition of 16 photons per emitter across 7 SIM frames. As this strength of illumination is in the order of the readout noise for CCDs, the 16 photon per emitter example serves to illustrate the form of noise implemented and was not used to evaluate the relative performance of methods in this study.

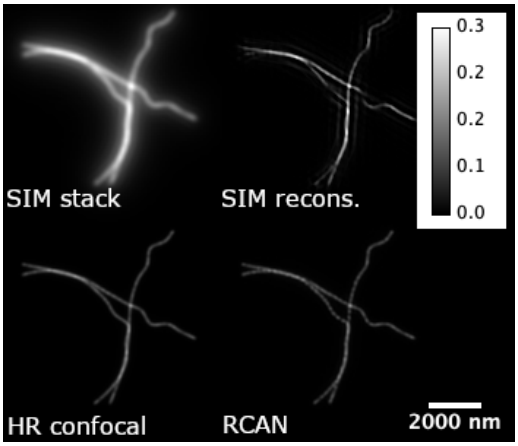


Figure 5: Axial mean of the MT1.N1.HD SMLM dataset [29]; simulated SIM stack (top-left), SIM-reconstructed stack (top-right), high-res confocal image (bottom-left) and RCAN output (bottom-right). Note the high structural fidelity and sharpness of the RCAN reconstruction, without the presence of faint side-lobe artifacts seen in the SIM reconstruction. These are visible as duplicate chromatin microtubules shifted on either side of the well illuminated structure, and are more apparent in the full video of the image stacks (available in the supplementary material). All images are rescaled from 8-bit pixel values to a range of [0,1] unless they were already outputted in that format, and no normalisation was done.

form wider continuous illuminations. To evaluate our network's performance on data containing these structures, cloud point data for Single Molecule Localisation Microscopy (SMLM) was taken from a SMLM competition's training dataset [29] (training sets MT1.N1.LD (≈ 8 000 points, low-density) and MT1.N1.HD (≈13 000 points, high-density)). These contain shallow and



Figure 6: X/Z view of the MT1.N1.HD SMLM dataset [29], averaged over Y.

narrow filament structures with high point densities, allowing us to evaluate the generalisability of our method to biological structures observed in other studies. SIM image stacks and high-resolution confocal images were simulated from this data, and processed by the RCAN network. As visualised in Figure 5, the RCAN (bottom-right) was able to produce the expected microtubule without the side-lobe artifacts seen in the SIM image; these are visible as faint duplicate microtubules on either side of a well illuminated strand, and are more distinguishable to non-expert viewers through the video provided in the supplementary material. The difference in axial resolution (RCAN mean: 331.4nm, SIM mean: 605.1nm, reference HR confocal mean: 372.9nm) reflected the previous test on chromatin data, as did the similarity in lateral resolution (RCAN: 170.8nm, SIM: 165.7nm, reference: 172.8nm), mse (RCAN: 1.6e-05, SIM: 4.4e-05) and SSIM (RCAN: 0.998, SIM: 0.996). A notable side-effect of the RCAN is the reconstruction of dense signal into a series of discrete point sources; this is visible in Figure 5 when comparing the RCAN reconstruction to the HR confocal, and in Supplementary materials A7 - Figure 10b, where a uniformly dense grating is reconstructed inconsistently. This effect is likely due to the use of sparse point clouds as training data, and could be remedied in future work by simulating artificial specimens containing uniformly dense light sources.

(d) Execution time

The execution time of our RCAN model was measured by computing SIM reconstructions of large image stacks. Reconstruction of 24 image stacks (280 frames each, i.e 40 SIM images with 7 frames per image) was repeated 10 times, allowing us to mitigate any file system interference due to shared use of the compute server. PyTorch's *DataParallel* module enables the distribution of image chunks from the same image stack to an array GPU computation units, thereby further accelerating the reconstruction of a 3D image stack. A mean execution time of 0.88 seconds per image stack (24ms per 7-frame SIM reconstruction) was demonstrated using a shared GPU cluster of two Nvidia GeForce RTX 2080 Ti. Further GPU parallelisation may be worthwhile for larger image stacks though I/O overheads allow no further performance gains for our test data.

5. Discussion

As demonstrated by the results in Section 4, the RCAN architecture described in Section 3 is capable of producing SIM reconstructions with a higher axial resolution than current 2D SIM reconstructions without a loss of lateral resolution or structural fidelity. This was demonstrated to be robust in low-light levels, which is typically a complexity of SIM and had been previously addressed by other deep-learning solutions [14].

The use of two simulated structure types in our training data (biological structures and sphere-bounded point-clouds) demonstrates some capacity for the network to generalise to varying structures in observed specimens; this was reflected by the performance of the network on simulated SMLM data in Figure 5. The network was found to approximate reconstructions of

continuous objects (microtubules) by reproducing the point-source patterns found in our training data. The task of generalising to other sources of data was addressed in varying approaches, such as training with real-world images [13]. This mitigates the likelihood of over-fitting a model to a distinct biological structure, as they are not present across all training data. In contrast, a U-Net approach found the need to re-train models to mitigate the occurrence of artifacts in reconstructed images [14]. Therefore it may be necessary to train on a wider variety of real or synthetic datasets to optimise the reconstruction of other biological structures using this method.

An issue which has yet to be addressed by image up-scaling networks is the limits of the up-scaling factor with regards to image fidelity. The original implementation of the RCAN architecture [18] demonstrates $4\times$ up-scaling but with a severe degradation of SSIM; this brings into question the certainty with which reconstructed images can be considered, especially if the ratio of up-scaling is further increased. The modified RCAN described here may be more effective when attempting higher ratios of up-scaling as the initial set of information contains many more frames (i.e channels) per output image, and therefore has the capacity to contain more information to be reconstructed into the HR image. A brief insight into the effect of removing one of these channels (see supplementary material A3) demonstrates the use of all channels in the reconstruction of the image, as the removal of any single image deteriorates the MSE and SSIM statistics in comparison to the reconstruction without masked layers. Ultimately, as the ratio of up-scaling goes beyond the capability of human validation, further imaging of structures via destructive methods (such as electron microscopy) could also validate final time points in reconstructed time lapse SIM images. A cross-validation could be performed using training groups formed of distinct biological structures; the sensitivity of the network to its training data would then provide better insight into potential reconstruction biases present in the network. This could further be extended by generating ensembles of trained models; this has previously been demonstrated in the use of image restoration for fluorescence microscopy [30]. In this scheme, each model would be trained on a random subset of the training data, and the output of all trained models for any given novel image reconstruction would be considered to form a statistical model for each output pixel. This could allow a degree of statistical confidence, which may be increasingly required to assert the certainty of reconstructions especially at larger ratios of image up-scaling.

Data Accessibility. All code used to generate training data, define & train deep learning models and evaluate these models is available at `https://github.com/mb1069/sim-axial-resolution`.

Funding. Funding for this project is provided by

- the European Union (under the Horizon 2020 Framework Programme: H2020 Future and Emerging Technologies (801336 - PROCHIP)).
- the Wellcome Trust's PhD program at Imperial College, London.

Competing Interests. The authors declare no competing interests.

References

1. Rainer Heintzmann and Christoph G. Cremer.
Laterally modulated excitation microscopy: improvement of resolution by using a diffraction grating.
In Irving J. Bigio, Herbert Schneckenburger, Jan Slavik, Katarina Svanberg M.D., and Pierre M. Viallet, editors, *Optical Biopsies and Microscopic Techniques III*, volume 3568, pages 185–196. International Society for Optics and Photonics, SPIE, 1999.
2. M. G. L. Gustafsson.

Surpassing the lateral resolution limit by a factor of two using structured illumination microscopy.
Journal of Microscopy, 198(2):82–87, 2000.

3. E.A. Ingerman, R.A. London, R. Heintzmann, and M.G.L. Gustafsson.
Signal, noise and resolution in linear and nonlinear structured-illumination microscopy.
Journal of Microscopy, 273(1):3–25, 2019.
4. Reto Fiolka, Lin Shao, E. Hesper Rego, Michael W. Davidson, and Mats G. L. Gustafsson.
Time-lapse two-color 3D imaging of live cells with doubled resolution using structured illumination.
Proceedings of the National Academy of Sciences, 109(14):5311–5315, 2012.
5. Raphaël Turcotte, Yajie Liang, Masashi Tanimoto, Qinrong Zhang, Ziwei Li, Minoru Koyama, Eric Betzig, and Na Ji.
Dynamic super-resolution structured illumination imaging in the living brain.
Proceedings of the National Academy of Sciences, 116(19):9586–9591, 2019.
6. Chung-Ju Rachel Wang, Peter M. Carlton, Inna N. Golubovskaya, and W. Zacheus Cande.
Interlock formation and coiling of meiotic chromosome axes during synapsis.
Genetics, 183(3):905–915, 2009.
7. Mats G. L. Gustafsson, Lin Shao, Peter M. Carlton, C. J. Rachel Wang, Inna N. Golubovskaya, W. Zacheus Cande, David A. Agard, and John W. Sedat.
Three-dimensional resolution doubling in wide-field fluorescence microscopy by structured illumination.
Biophysical journal, 94(12):4957–4970, Jun 2008.
8. L. J. Young, F. Ströhl, and C. F. Kaminski.
A Guide to Structured Illumination TIRF Microscopy at High Speed with Multiple Colors.
Journal of Visualised Experiments, 111(111), 05 2016.
9. Andrew G. York, Sapun H. Parekh, Damian Dalle Nogare, Robert S. Fischer, Kelsey Temprine, Marina Mione, Ajay B. Chitnis, Christian A. Combs, and Hari Shroff.
Resolution doubling in live, multicellular organisms via multifocal structured illumination microscopy.
Nature Methods, 9(7):749–754, Jul 2012.
10. Maria Ingaramo, Andrew G. York, Peter Wawrzusin, Oleg Milberg, Amy Hong, Roberto Weigert, Hari Shroff, and George H. Patterson.
Two-photon excitation improves multifocal structured illumination microscopy in thick scattering tissue.
Proceedings of the National Academy of Sciences, 111(14):5254–5259, 2014.
11. Ronny Förster, Kai Wicker, Walter Müller, Aurélie Jost, and Rainer Heintzmann.
Motion artefact detection in structured illumination microscopy for live cell imaging.
Opt. Express, 24(19):22121–22134, Sep 2016.
12. Ronny Förster, Walter Müller, René Richter, and Rainer Heintzmann.
Automated distinction of shearing and distortion artefacts in structured illumination microscopy.
Opt. Express, 26(16):20680–20694, Aug 2018.
13. Charles N. Christensen, Edward N. Ward, Pietro Lio, and Clemens F. Kaminski.
ML-SIM: A deep neural network for reconstruction of structured illumination microscopy images.
arXiv 2003.11064, 08 2020.
14. Luhong Jin, Bei Liu, Fenqiang Zhao, Stephen Hahn, Bowei Dong, Ruiyan Song, Timothy C. Elston, Yingke Xu, and Klaus M. Hahn.
Deep learning enables structured illumination microscopy with low light levels and enhanced speed.
Nature Communications, 11(1):1934, Apr 2020.
15. Chang Ling, Chonglei Zhang, Mingqun Wang, Fanfei Meng, Luping Du, and Xiaocong Yuan.
Fast structured illumination microscopy via deep learning.
Photonics Research, 8(8):1350–1359, Aug 2020.
16. Junchao Fan, Xiaoshuai Huang, Liuju Li, Shan Tan, and Liangyi Chen.
A protocol for structured illumination microscopy with minimal reconstruction artifacts.
Biophysics Reports, 5(2):80–90, Apr 2019.
17. Nadya Chakrova, Bernd Rieger, and Sjoerd Stallinga.

Deconvolution methods for structured illumination microscopy.
J. Opt. Soc. Am. A, 33(7):B12–B20, Jul 2016.

18. Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu.
Image super-resolution using very deep residual channel attention networks.
In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 294–310, Cham, 2018. Springer International Publishing.
19. Aurélie Jost, Elen Tolstik, Polina Feldmann, Kai Wicker, Anne Sentenac, and Rainer Heintzmann.
Optical sectioning and high resolution in single-slice structured illumination microscopy by thick slice blind-sim reconstruction.
PLOS ONE, 10(7):1–10, 07 2015.
20. Zhihao Wang, Jian Chen, and Steven C.H. Hoi.
Deep learning for image super-resolution: A survey.
IEEE Transactions on Pattern Analysis and Machine Intelligence, 2020.
21. Zhou Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli.
Image quality assessment: from error visibility to structural similarity.
IEEE Transactions on Image Processing, 13(4):600–612, 2004.
22. Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala.
Pytorch: An imperative style, high-performance deep learning library.
In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
23. Siyu Wang, Jinbo Xu, and Jianyang Zeng.
Inferential modeling of 3D chromatin structure.
Nucleic Acids Research, 43(8):54, 02 2015.
24. Federico Sala, Michele Castriotta, Petra Paiè, Andrea Farina, Sarah D'Annunzio, Alessio Zippo, Roberto Osellame, Francesca Bragheri, and Andrea Bassi.
High-throughput 3d imaging of single cells with light-sheet fluorescence microscopy on chip.
Biomedical Optics Express, 11(8):4397–4407, Aug 2020.
25. H Gong, W Guo, C Dunsby, and M.A.A. Neil.
GPU-accelerated, real-time reconstruction in Python of three dimensional datasets from structured illumination microscopy with hexagonal patterns.
Submitted to Philosophical Transactions of the Royal Society A, 09 2020.
26. Eva Wegel, Antonia Göhler, B. Christoffer Lagerholm, Alan Wainman, Stephan Uphoff, Rainer Kaufmann, and Ian M. Dobbie.
Imaging cellular structures in super-resolution with sim, sted and localisation microscopy: A practical comparison.
Scientific Reports, 6(1):27290, Jun 2016.
27. Sami Koho, Giorgio Tortarolo, Marco Castello, Takahiro Deguchi, Alberto Diaspro, and Giuseppe Vicidomini.
Fourier ring correlation simplifies image restoration in fluorescence microscopy.
bioRxiv, 2019.
28. Sarah L. Veatch, Benjamin B. Machta, Sarah A. Shelby, Ethan N. Chiang, David A. Holowka, and Barbara A. Baird.
Correlation functions quantify super-resolution images and estimate apparent clustering due to over-counting.
PLOS ONE, 7(2):1–13, 02 2012.
29. Daniel Sage, Thanh-An Pham, Hazen Babcock, Tomas Lukes, Thomas Pengo, Jerry Chao, Ramraj Velmurugan, Alex Herbert, Anurag Agrawal, Silvia Colabrese, Ann Wheeler, Anna Archetti, Bernd Rieger, Raimund Ober, Guy M. Hagen, Jean-Baptiste Sibarita, Jonas Ries, Ricardo Henriques, Michael Unser, and Seamus Holden.
Super-resolution fight club: assessment of 2D and 3D single-molecule localization microscopy software.
Nature Methods, 16(5):387–395, May 2019.
30. Martin Weigert, Uwe Schmidt, Tobias Boothe, Andreas Müller, Alexandr Dibrov, Akanksha

Jain, Benjamin Wilhelm, Deborah Schmidt, Coleman Broaddus, Siân Culley, Mauricio Rocha-Martins, Fabián Segovia-Miranda, Caren Norden, Ricardo Henriques, Marino Zerial, Michele Solimena, Jochen Rink, Pavel Tomancak, Loic Royer, Florian Jug, and Eugene W. Myers. Content-aware image restoration: pushing the limits of fluorescence microscopy. *Nature Methods*, 15(12):1090–1097, Dec 2018.