

# Applications and Challenges of Machine Learning to Enable Realistic Cellular Simulations

R. Vasan<sup>1</sup>, M. P. Rowan<sup>2</sup>, C. T. Lee<sup>1</sup>, G. R. Johnson<sup>3</sup>, P. Rangamani<sup>1</sup>, and M. J. Holst<sup>4,5,\*</sup>

November 14, 2019

<sup>1</sup>Department of Mechanical and Aerospace Engineering, University of California San Diego, La Jolla, CA 92093, USA

<sup>2</sup>Department of Bioengineering, University of California San Diego, La Jolla, CA 92093, USA

<sup>3</sup>Allen Institute of Cell Science, Seattle, WA 98109, USA

<sup>4</sup>Department of Mathematics, University of California San Diego, La Jolla, CA 92093, USA

<sup>5</sup>Department of Physics, University of California San Diego, La Jolla, CA 92093, USA

\*Corresponding Author

|Email: mholst@ucsd.edu |

## Abstract

In this perspective, we examine three key aspects of an end-to-end pipeline for realistic cellular simulations: reconstruction and segmentation of cellular structures; generation of cellular structures; and mesh generation, simulation, and data analysis. We highlight some of the relevant prior work in these distinct but overlapping areas, with a particular emphasis on current use of machine learning technologies, as well as on future opportunities.

# 1 Introduction

Machine learning (ML) approaches, including both traditional and deep learning methods, are revolutionizing biology. Owing to major advances in experimental and computational methodologies, the amount of data available for training is rapidly increasing. The timely convergence of data availability, computational capability, and new algorithms is a boon for biophysical modeling of subcellular and cellular scale processes such as biochemical signal transduction and mechanics [1]. To date, many simulations are performed using idealized geometries that allow for the use of commonly used techniques and software [2, 3, 4, 5, 6]. This is historically due to the lack of high-resolution structural data as well as the theoretical and computational challenges for simulations in realistic cellular shapes, due to the complexity of generating high-quality, high-resolution meshes for simulation, and the need to develop specialized fast numerical solvers that can be used with very large unstructured mesh representations of the physical domain.

As biophysical methods have improved, the complexity of our mathematical and computational models is steadily increasing [7, 8, 9]. A major frontier for physics-based study of cellular processes will be to simulate biological processes in realistic cell shapes derived from various structural determination modalities [10, 11]. For biological problems ranging from cognition to cancer, it has long been understood that cell shapes are often correlated with mechanism [12, 13, 14, 15, 16]. Despite such clear correlations, there remain gaps in our understanding of how cellular ultrastructure contributes to cellular processes and the feedback between cellular structure and function. Challenges such as the diffraction limit of light and difficulties in manipulation of intracellular ultrastructure constrain the potential scope of what can be achieved experiments. Much like the partnership between biophysics and molecular dynamics simulations have enabled the modeling of invisible protein motions to shed insights on experimental observations, simulations of cellular processes can also aid in the validation and generation of hypothesis currently inaccessible by experimental methods. Recently, we and others have shown that, for example, cell shape and localization of proteins can impact cell signaling [4, 5, 12, 17, 18].

The major bottleneck for the widespread use of cell scale simulations with realistic geometries is not the availability of structural data. Indeed, there exist many three-dimensional imaging modalities such as confocal microscopy, multiphoton microscopy, super-resolution fluorescence and electron tomography [19, 20]. For example, automation of modalities such as Serial Block-Face Scanning Electron Microscopy is already enabling the production of data at rapid rates. The bottleneck lies in the fact that much of the data generated from these imaging modalities need to be manually curated before it can be used for physics-based simulations. This current *status quo* of manually processing and curating these datasets for simulations is a major obstacle to our progress. In order to bridge the gap between abundance of cellular ultrastructure data generated by 3D electron microscopy (EM) techniques and simulations in these realistic geometries, innovations in machine learning (ML) methods will be necessary to reduce the time

it takes to go from structural datasets to initial models. There are already many similar efforts at the organ/tissue and connectomics length scales [21, 22, 23]. In this work, we summarize the main steps necessary to construct simulations with realistic cellular geometries (Fig. 1) and highlight where innovation in ML efforts are needed and will have significant impacts. We further discuss some of the challenges and limitations in the existing methods, setting the stage for new innovations for ML in physics-based cellular simulations.

## **2 Sources of error in imaging modalities**

Images generated by the various microscopy modalities must undergo pre-processing to correct for errors such as uneven illumination or background noise [25, 26]. The choice of suitable algorithms for error correction depends on multiple factors, some of which are listed here – the length scale of the experiment being conducted, scalability and reproducibility of the experiment, optical resolution of the microscope, sensitivity of the detector, specificity of the staining procedure, imaging mode (2D, 3D, 3D time-series), imaging modality (fluorescence, EM, ET etc.), and other imaging artifacts like electronic noise, lens astigmatism, mechanical tilting/vibration, sample temperature, and discontinuous staining [25, 26, 27, 28]. These sources of error are an important consideration for model implementation further downstream [6].

Electron tomography (ET) remains one of the most popular methods of cell imaging for modeling purposes [29, 30, 31], as it retains the highest resolution of all the 3D cell imaging techniques [26] by reconstructing a 3D object from a series of 2D images collected at different tilt angles [32]. However, images from ET also have a low signal to noise ratio (SNR) and have anisotropic resolution (for example, 1 nm resolution in x, y and 10 nm resolution in z) [25]. This is partly because biological samples can withstand only a limited dose of electron beam radiation (SNR is proportional to the square root of the electron beam current) before the specimen is damaged [33]. Other sources of error such as misalignment of projections and missing wedges from an incomplete tilt angular range can significantly affect the quality of the reconstruction. To work with data such as these, image processing steps are required for high resolution 3D reconstruction [25, 34]. Commonly used software packages for image processing such as IMOD [35] and TomoJ[36] use reconstruction algorithms such as Weighted Backprojection (WBP) and Simultaneous Iterative Reconstruction Technique (SIRT). While these have been very effective at reconstruction, sources of error can still accumulate, leading to further manual adjustment [37].

### 3 Applications of ML for the segmentation and reconstruction of cellular structures

Given a noisy 3D reconstruction, how can we segment cellular structures of interest? One approach is to employ manual segmentation tools applied to 3D tomograms such as XVOXTRACE [32, 38], and more generally, manual contouring, interpolation, and contour stacking (Fig. 1). The advantage of such methods is that the human eye performs exceptionally well at detecting objects in an image [27, 39]. Consequently, semi-manual and manual segmentation are widely adopted, favoring accuracy over efficiency. However, such methods can be extremely tedious and not always reproducible. Alternatively, numerous semi-automatic segmentation algorithms such as interpolation, watershed, thresholding, and clustering are available as plugins in software packages like IMOD [35] and ImageJ [40] (Fig. 1, classical). However, the accuracy of open platform algorithms is debatable [41] because of two main reasons – (i) Even with a ‘perfect’ ET reconstruction (no tilt misalignment, no missing wedge, no noise), the application of filtering algorithms like Gaussian blur or non-linear anisotropic diffusion (NAD) [42] can cause artefacts that lead to misclassifications, rendering the image unsuitable for downstream quantitative simulations and analysis. (ii) Segmentation workflows are often designed for a specific structure and/or imaging modality, limiting their generalizability and applicability.

Annual cell segmentation challenges are evidence of the demand for automatic segmentation [43, 44], with many of its past winners responding with ML-based programs [45, 46]. Training labels for ML techniques requires a relatively small percentage (as small as 10%) of manually segmented labels, allowing for very large data sets to be processed significantly faster than previously implemented semi-automatic segmentation methods. The most successful teams utilized ML techniques such as random forest classifiers, support vector machines, or a combination of these to get segmentations comparable or often even better than their human counterparts [43, 44, 45, 46] (Fig. 1, machine learning). These techniques function by imputing several image features such as noise reduction, and texture and edge detection filters [47]. These filters are then used to train a classification algorithm in an interactive manner, achieving better classification accuracy at the cost of increased training time compared to the direct application of a filter. However, because the algorithm is interactive, it still requires manual input and both the training time and accuracy can depend on the user.

More recently, deep learning-based ML algorithms (Fig. 1, deep learning), and more specifically, convolutional neural networks (CNNs) have surged in popularity due to the success of AlexNet in the ImageNet classification challenge [48]. CNNs are complex learnable non-linear functions that do not require the imputation of data-specific features. Indeed, CNNs learn the feature mapping directly from the image. The U-Net convolutional neural network architecture [46] further generalized deep learning, winning the ISBI neuronal structure segmentation challenge in

2015 with a quicker speed and with fewer training images. It functions by using the feature mapping imputed by a CNN to map the classification vector back into a segmented image. Such is the achievement of the U-Net that its variants are now the state-of-the-art in tasks like calling genetic variation from gene-sequencing data [49], brain tumor detection [50] and segmentation of medical image datasets [51]. However, such deep learning based methods have their own challenges. They require both quality and quantity of annotated training data, significant amount of training time, graphics processing unit computing, and can generalize poorly to a different dataset.

Both the difficulty and cost of generating annotated training data increases exponentially when dealing with Volumetric (3D) images compared with 2D, which are the desired inputs for biophysical simulations. Since the U-Net is a 2D architecture [46], it cannot be applied directly to 3D images without modifications. To this end, 3D U-net used sparsely annotated 2D slices to generate volumetric segmentations of brain tumors [52]. Similarly, VoxRestNet [53] introduced residual learning using ResNet [54], a deep residual network capable of training hundreds to thousands of layers without a performance drop, to a voxelwise representation of 3D magnetic resonance (MR) images of the brain, paving the way for scalable 3D segmentation.

Excitingly, such algorithms are being made openly accessible and easy-to-use. For example, iLastik [45, 55] and Trainable Weka Segmentation [47] are both available as plugins in software packages like ImageJ. These tools provide an interactive platform for segmentation, employing supervised classification techniques like random forests as well as unsupervised clustering such as K-means [47]. Similarly, deep learning tools such as DeepCell [56] and U-Net [46, 57] are also available in various bioimage software packages. Other stand-alone tools like the Allen Cell Structure Segmenter provide a lookup table of 20 segmentation workflows that feed into an iterative deep learning model [58]. Cloud compute based segmentation plugins like CDeep3M [59] leverage Amazon Web Services (AWS) images to provide an efficient and compute-scalable tool for both 2D and 3D biomedical images.

Generating well-organized and annotated training data continues to be the major challenge for most ML segmentation methods. Crowdsourced annotation tools like Amazon's Mechanical Turk can be useful in this context, but are still limited by the difficulty of training naive users on tracing specific structural images. Alternatively, many ML algorithms leverage transfer learning approaches using pre-trained networks such as VGG-net [60, 61, 62], AlexNet [48] and GoogleNet [63]. In fact, popular semantic segmentation and clustering networks like Fully Convolutional Networks (FCN) [64] and DECAF [65] are themselves implemented using transfer learning approaches. Such transfer learning can also be used to generalize models trained on biological data to a different cell type or experimental condition, significantly reducing the time for training and accompanying computing resources required. More recently, label-free approaches employing a U-net variant have been applied to predict cellular structure from unlabeled brightfield images [66, 67]. These methods can serve as a platform for building low cost, scalable, and efficient segmentation of 3D cellular structure.

## 4 Applications of ML for the generation of synthetic cellular structures

There are two main aspects involved in the development of comprehensive biophysical models – (1) what is the process being modeled? and (2) what is the geometry in which this process is being modeled? Answers to the first question are based on experimental observations and specific biology. Answering the latter is significantly more challenging because of the difficulties in – (i) obtaining accurate segmentations, (ii) discovering new structure from experiments, and (iii) simultaneously visualizing multiple structures. The use of synthetically generated geometries, which can probe different arrangements of organelles within cells could be relevant for generating biologically relevant hypotheses.

A subset of ML models, called *generative* models, deal with the task of generating new synthetic but realistic images that match the training set distribution. For our purposes, such methods are relevant in the context of generating (i) noise-free images, (ii) images representative of a different cell type, structure, or time-point, and (iii) unlikely images that represent the most unique shapes of the structure being imaged. For example, by capturing the unlikely and likely shapes in our dataset, we could generate sequences of synthetic images that transition from one shape to the next. These synthetic images can be used in biophysical simulations to generate biologically relevant hypotheses.

In recent years, there has been rapid progress in applying deep generative models to natural images, text, and even medical images. Popular classes of deep generative models like Variational Autoencoders [68], Generative Adversarial Networks [69], and Autoregressive models such as PixelRNN [70] and PixelCNN [71] have achieved state of the art performance on popular image datasets such as MNIST [72], CIFAR [73] and ImageNet [74]. Each class of models has numerous modified implementations. For example, GANs alone include models like deep convolutional GAN (DCGAN) [75], conditional GAN (cGAN) [76], StackGAN [77], InfoGAN [78] and Wasserstein GAN [79] to name a few. Each model has its own distinct set of advantages and disadvantages. GANs can produce photo-realistic images at the cost of tricky training and no dimensionality reduction. VAEs allow for both generation and inference, but their naive implementation results in less photo-realistic generative examples. Autoregressive models obtain the best log-likelihoods at the cost of poor dimensionality reduction. Importantly, all of these models are unsupervised, implying that they are not limited by manual annotation that is otherwise a common challenge to supervised learning approaches.

In cell biology, much of the work in building generative models of cellular structures has been associated with the open source CellOrganizer [80, 81, 82, 83, 84, 85, 86], which uses a Gaussian Mixture Model given reference frames like the cell and nuclear shape in order to predict organelle shape distribution. These models also have the option to be parametric (parameters such as number of objects), which reduces the complexity of the learning task, the training time and GPU computing resources required, while also allowing for exploration and analysis of the parameters

and their effect on the spatial organization of cells. Aside from CellOrganizer, other recent efforts have begun to leverage deep generative models in cell biology. We now have models that can predict structure localization given cell and nuclear shape [87], extract functional relationships between fluorescently tagged proteins structures in cell images [88], learn cell features from cell morphological profiling experiments [89], and interpret gene expression levels from single-cell RNA sequencing data [90, 91].

The challenge going forward will be how best to use generative modeling given the data in hand. This will depend on the question we want to ask of the data. For example, if we are modeling processes associated with cell and nuclear shape, spherical harmonics based generative models might be more appropriate than deep learning based methods [92]. If we are interested in inpainting a missing wedge from a tomogram using a generative model, then GANs might be more appropriate [93]. Generated images can also be used as a source of training data for segmentation and classification tasks [94]. Taken together, these choices will help develop efficient end-to-end pipelines for segmentation and shape generation, and provide a platform for running biophysical simulations. Already, CellOrganizer can export spatial instances to cell simulation engines such as MCell [95] and VirtualCell [96], allowing us to simulate chemical reactions in different spatial compartments. Similar pipelines for deep generative models will need to be implemented in order to fully realize their downstream interpretations.

## 5 Applications of ML for meshing, simulation, and data analysis

ML is commonly applied to mesh segmentation and classification; examples include PointNet [99] (segments and classifies a point cloud), and MeshCNN [100] (segments and classifies edges in a mesh). However, although the term *machine learning* was not traditionally used to describe meshing techniques, in fact algorithms for mesh generation (cf. [101]), mesh improvement (such as mesh smoothing [102]), and mesh refinement [103, 104, 105, 106] all fundamentally involve local (cf. [107]) and/or global (cf. [108]) optimization of an *objective function* (see Fig. 2). Mesh point locations, and/or edge/face connectivity decisions are viewed as parameters that are determined (or *learned*) as part of an iterative algorithm that extremizes a local or global objective function (usually involving constraints as well) in an effort to generate, improve, or refine a given mesh. In addition, *adaptive numerical methods* for simulation of physical systems involving the solution of ordinary (ODE) and partial (PDE) differential equations are again an early example of the application of ML techniques in computational science, long before the terminology was widely used. A classic reference from the 1970's in the context of adaptive finite element methods is [109, 110]; all modern approaches to adaptive numerical methods for ODE and PDE systems continue to follow the same general framework outlined in that work: (i) Solve the ODE/PDE on the current mesh; (ii) Estimate the error using a posteriori indicators; (iii) Refine the mesh using provably non-degenerate local refinement with closure; (iv) Go back to step (i) and repeat the iteration until a target quality measure is obtained (a standard

approach is to approximately minimize a global error function, through the use of local error estimates). These types of adaptive algorithms are effectively *machine learning* the best possible choice (highest accuracy with least cost) of mesh and corresponding numerical discretization for the target ODE/PDE system. Recent work in the area is now moving toward a more explicit and sophisticated use of modern ML techniques (cf. [111, 112]). ML can further assist in simulation and data analysis further downstream. Specifically, it can accelerate (i) parameter estimation, (ii) uncertainty quantification, and (iii) dimensionality reduction, three of the most common post-processing tasks from a biophysical simulation. Incorporating specific biophysical model information like stress-strain relationships [113] or statistical molecular dynamic states [114] into ML algorithms can also reduce the computational time for numerical solvers. Finally, more standard ML approaches like clustering and dimensionality reduction can assist in both visualization and interpretation of simulation results.

## 6 Perspectives and Future Directions

In this perspective, we have discussed three key aspects of a pipeline for realistic cellular simulations: (i) Reconstruction and segmentation of cellular structure; (ii) Generation of cellular structure; and (iii) Mesh generation, refinement and simulation. While these were discussed separately, neural networks like Pixel2Mesh demonstrate the feasibility of end-to-end pipelines from a single black box [115]. Of course, black boxes are not interpretable, and recent ML frameworks like SAUCIE have begun to use regularizations to enforce mechanistic interpretability in the hidden layers of an autoencoder neural network [116]. We anticipate that future endeavours will implement a fully interpretable and end-to-end pipeline for biophysical simulations.

### Conflict of Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

### Author Contributions

All co-authors contributed to the writing of the manuscript, and also provided area-specific expertise.

*R. Vasan*: Provided expertise on generation and simulation.

*M. Rowan*: Provided expertise on imaging, segmentation, and reconstruction.



*C.T. Lee:* Provided expertise on reconstruction, meshing and simulation.

*G.R. Johnson:* Provided expertise on imaging, segmentation, reconstruction and generation.

*P. Rangamani:* Provided expertise on imaging, segmentation, reconstruction, and simulation.

*M. Holst:* Provided expertise on meshing, simulation, and analysis.

## **Funding**

R.V., M.R., C.T.L, and P.R. are supported in part by an AFOSR MURI award FA9550-18-1-0051 and ONR fund ONR N00014-17-1-2628. C.T.L. also acknowledges support from a Hartwell Foundation Postdoctoral Fellowship. M.J.H. was supported in part by NSF Awards 1934411 and 1630366.

## **Acknowledgments**

We would like to thank Prof. Pietro De Camilli and coworkers for sharing their datasets from Wu et al. [24]. We also thank Dr. Matthias Haberl, Mr. Evan Campbell, Profs. Brenda Bloodgood and Mark Ellisman for helpful discussion and suggestions. GJ thanks the Allen Institute for Cell Science founder, Paul G. Allen, for his vision, encouragement and support.

## **References**

- [1] Christopher T. Lee and Rommie E. Amaro. Exascale Computing: A New Dawn for Computational Biology. Computing in Science & Engineering, 20(5):18–25, September 2018.
- [2] Haleh Alimohamadi, Ritvik Vasani, Julian Edwin Hassinger, Jeanne C Stachowiak, and Padmini Rangamani. The role of traction in membrane curvature generation. Molecular biology of the cell, 29(16):2024–2035, 2018.
- [3] Ritvik Vasani, Matthew Akamatsu, Johannes Schöneberg, and Padmini Rangamani. Intracellular membrane trafficking: Modeling local movements in cells. In Cell Movement, pages 259–301. Springer, 2018.
- [4] Miriam Bell, Tom Bartol, Terrence Sejnowski, and Padmini Rangamani. Dendritic spine geometry and spine apparatus organization govern the spatiotemporal dynamics of calcium. The Journal of general physiology, 151(9):2221, 2019.

- [5] Donya Ohadi and Padmini Rangamani. Geometric control of frequency modulation of camp oscillations due to  $ca^{2+}$ -bursts in dendritic spines. Biophysical journal, 2019.
- [6] Ritvik Vasan, Mary M Maleckar, Charles David Williams, and Padmini Rangamani. Dlite uses cell-cell interface movement to better infer cell-cell tensions. Biophysical journal, 2019.
- [7] Ritvik Vasan, Shiva Rudraraju, Matthew Akamatsu, Krishna Garikipati, and Padmini Rangamani. A mechanical model reveals that non-axisymmetric buckling lowers the energy barrier associated with membrane neck constriction. arXiv preprint arXiv:1906.06443, 2019.
- [8] Shiva Rudraraju, Anton Van der Ven, and Krishna Garikipati. Mechanochemical spinodal decomposition: a phenomenological theory of phase transformations in multi-component, crystalline solids. npj Computational Materials, 2:16012, 2016.
- [9] L Angela Mihai, Silvia Budday, Gerhard A Holzapfel, Ellen Kuhl, and Alain Goriely. A family of hyperelastic models for human brain tissue. Journal of the Mechanics and Physics of Solids, 106:60–79, 2017.
- [10] Robert F Murphy. Building cell models and simulations from microscope images. Methods, 96:33–39, 2016.
- [11] Christopher T Lee, Justin G Laughlin, John B Moody, Rommie E Amaro, J Andrew McCammon, Michael J Holst, and Padmini Rangamani. An open source mesh generation platform for biophysical modeling using realistic cellular geometries. arXiv preprint arXiv:1909.04781, 2019.
- [12] Padmini Rangamani, Azi Lipshtat, Evren U Azeloglu, Rhodora Cristina Calizo, Mufeng Hu, Saba Ghassemi, James Hone, Suzanne Scarlata, Susana R Neves, and Ravi Iyengar. Decoding information in cell shape. Cell, 154(6):1356–1369, 2013.
- [13] HJ Deuling and W Helfrich. Red blood cell shapes as explained on the basis of curvature elasticity. Biophysical journal, 16(8):861–868, 1976.
- [14] Thomas M Bartol Jr, Cailey Bromer, Justin Kinney, Michael A Chirillo, Jennifer N Bourne, Kristen M Harris, and Terrence J Sejnowski. Nanoconnectomic upper bound on the variability of synaptic plasticity. Elife, 4:e10778, 2015.
- [15] Raphael Ritz and Terrence J Sejnowski. Synchronous oscillatory activity in sensory systems: new vistas on mechanisms. Current opinion in neurobiology, 7(4):536–546, 1997.
- [16] Kristen M Harris and SB Kater. Dendritic spines: cellular specializations imparting both stability and flexibility to synaptic function. Annual review of neuroscience, 17(1):341–371, 1994.

- [17] Andrea Cugno, Thomas M Bartol, Terrence J Sejnowski, Ravi Iyengar, and Padmini Rangamani. Geometric principles of second messenger dynamics in dendritic spines. Scientific reports, 9(1):1–18, 2019.
- [18] Donya Ohadi, Danielle L Schmitt, Barbara Calabrese, Shelley Halpain, Jin Zhang, and Padmini Rangamani. Computational modeling reveals frequency modulation of calcium-camp/pka pathway in dendritic spines. Biophysical journal, 2019.
- [19] Fang Huang, George Sirinakis, Edward S Allgeyer, Lena K Schroeder, Whitney C Duim, Emil B Kromann, Thomy Phan, Felix E Rivera-Molina, Jordan R Myers, Irnov Irnov, et al. Ultra-high resolution 3d imaging of whole cells. Cell, 166(4):1028–1040, 2016.
- [20] Benedikt W Graf and Stephen A Boppart. Imaging and analysis of three-dimensional cell culture models. In Live cell imaging, pages 211–227. Springer, 2010.
- [21] Jeff W Lichtman, Hanspeter Pfister, and Nir Shavit. The big data challenges of connectomics. Nature neuroscience, 17(11):1448, 2014.
- [22] Gabriel Maher, Nathan Wilson, and Alison Marsden. Accelerating cardiovascular model building with convolutional neural networks. Medical & biological engineering & computing, 57(10):2319–2335, 2019.
- [23] Michał Januszewski, Jörgen Kornfeld, Peter H Li, Art Pope, Tim Blakely, Larry Lindsey, Jeremy Maitin-Shepard, Mike Tyka, Winfried Denk, and Viren Jain. High-precision automated reconstruction of neurons with flood-filling networks. Nature methods, 15(8):605, 2018.
- [24] Yumei Wu, Christina Whiteus, C. Shan Xu, Kenneth J. Hayworth, Richard J. Weinberg, Harald F. Hess, and Pietro De Camilli. Contacts Between the Endoplasmic Reticulum and Other Membranes in Neurons. Proc. Natl. Acad. Sci. USA, 114(24):E4859–E4867, jun 2017.
- [25] Wim van Aarle, Willem Jan Palenstijn, Jan De Beenhouwer, Thomas Altantzis, Sara Bals, K Joost Batenburg, and Jan Sijbers. The astra toolbox: A platform for advanced algorithm development in electron tomography. Ultramicroscopy, 157:35–47, 2015.
- [26] Diane S Lidke and Keith A Lidke. Advances in high-resolution imaging–techniques for three-dimensional imaging of cellular structures. J Cell Sci, 125(11):2571–2580, 2012.
- [27] Erick Moen, Dylan Bannon, Takamasa Kudo, William Graf, Markus Covert, and David Van Valen. Deep learning for cellular image analysis. Nature methods, page 1, 2019.
- [28] AJ Koster, H Chen, JW Sedat, and DA Agard. Automated microscopy for electron tomography. Ultramicroscopy, 46(1-4):207–227, 1992.

- [29] Tomáš Mazel, Rebecca Raymond, Mary Raymond-Stintz, Stephen Jett, and Bridget S Wilson. Stochastic modeling of calcium in 3d geometry. Biophysical journal, 96(5):1691–1706, 2009.
- [30] Matt West, Nesia Zurek, Andreas Hoenger, and Gia K Voeltz. A 3d analysis of yeast er structure reveals how er domains are organized by membrane curvature. The Journal of cell biology, 193(2):333–346, 2011.
- [31] Andrew B Noske, Adam J Costin, Garry P Morgan, and Brad J Marsh. Expedited approaches to whole cell electron tomography and organelle mark-up in situ in high-pressure frozen pancreatic islets. Journal of structural biology, 161(3):298–313, 2008.
- [32] G Perkins, C Renken, ME Martone, SJ Young, M Ellisman, and T Frey. Electron tomography of neuronal mitochondria: three-dimensional structure and organization of cristae and membrane contacts. Journal of structural biology, 119(3):260–272, 1997.
- [33] Lindsay A Baker and John L Rubinstein. Radiation damage in electron cryomicroscopy. In Methods in enzymology, volume 481, pages 371–388. Elsevier, 2010.
- [34] Sébastien Phan, Daniela Boassa, Phuong Nguyen, Xiaohua Wan, Jason Lanman, Albert Lawrence, and Mark H Ellisman. 3d reconstruction of biological structures: automated procedures for alignment and reconstruction of multiple tilt series in electron tomography. Advanced structural and chemical imaging, 2(1):8, 2017.
- [35] James R Kremer, David N Mastrorade, and J Richard McIntosh. Computer visualization of three-dimensional image data using imod. Journal of structural biology, 116(1):71–76, 1996.
- [36] Cédric Messaoudi, Thomas Boudier, Carlos Oscar Sanchez Sorzano, and Sergio Marco. Tomoj: tomography software for three-dimensional reconstruction in transmission electron microscopy. BMC bioinformatics, 8(1):288, 2007.
- [37] Rowan Leary, Zineb Saghi, Paul A Midgley, and Daniel J Holland. Compressed sensing electron tomography. Ultramicroscopy, 131:70–91, 2013.
- [38] Xinghua Yin, Grahame J Kidd, Nobuhiko Ohno, Guy A Perkins, Mark H Ellisman, Chinthasagar Bastian, Sylvain Brunet, Selva Baltan, and Bruce D Trapp. Proteolipid protein-deficient myelin promotes axonal mitochondrial dysfunction via altered metabolic coupling. J Cell Biol, 215(4):531–542, 2016.
- [39] Hervé Le Borgne, Nathalie Guyader, Anne Guérin-Dugué, and Jeanny Hérault. Classification of images: Ica filters vs human perception. In Seventh International Symposium on Signal Processing and Its Applications, 2003. Proceedings., volume 2, pages 251–254. IEEE, 2003.

- [40] Michael D Abràmoff, Paulo J Magalhães, and Sunanda J Ram. Image processing with imagej. Biophotonics international, 11(7):36–42, 2004.
- [41] Tim Jerman, Franjo Pernuš, Boštjan Likar, and Žiga Špiclin. Enhancement of vascular structures in 3d and 2d angiographic images. IEEE transactions on medical imaging, 35(9):2107–2118, 2016.
- [42] Achilleas S Frangakis and Reiner Hegerl. Noise reduction in electron tomographic reconstructions using nonlinear anisotropic diffusion. Journal of structural biology, 135(3):239–250, 2001.
- [43] Ignacio Arganda-Carreras, Srinivas C Turaga, Daniel R Berger, Dan Cireşan, Alessandro Giusti, Luca M Gambardella, Jürgen Schmidhuber, Dmitry Laptev, Sarvesh Dwivedi, Joachim M Buhmann, et al. Crowdsourcing the creation of image segmentation algorithms for connectomics. Frontiers in neuroanatomy, 9:142, 2015.
- [44] Martin Maška, Vladimír Ulman, David Svoboda, Pavel Matula, Petr Matula, Cristina Ederra, Ainhoa Urbiola, Tomás España, Subramanian Venkatesan, Deepak MW Balak, et al. A benchmark for comparison of cell tracking algorithms. Bioinformatics, 30(11):1609–1617, 2014.
- [45] Christoph Sommer, Christoph Straehle, Ullrich Koethe, and Fred A Hamprecht. Ilastik: Interactive learning and segmentation toolkit. In 2011 IEEE international symposium on biomedical imaging: From nano to macro, pages 230–233. IEEE, 2011.
- [46] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In International Conference on Medical image computing and computer-assisted intervention, pages 234–241. Springer, 2015.
- [47] Ignacio Arganda-Carreras, Verena Kaynig, Curtis Rueden, Kevin W Eliceiri, Johannes Schindelin, Albert Cardona, and H Sebastian Seung. Trainable weka segmentation: a machine learning tool for microscopy pixel classification. Bioinformatics, 33(15):2424–2426, 2017.
- [48] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pages 1097–1105, 2012.
- [49] Ryan Poplin, Pi-Chuan Chang, David Alexander, Scott Schwartz, Thomas Colthurst, Alexander Ku, Dan Newburger, Jojo Dijamco, Nam Nguyen, Pegah T Afshar, et al. A universal snp and small-indel variant caller using deep neural networks. Nature biotechnology, 36(10):983, 2018.
- [50] Hao Dong, Guang Yang, Fangde Liu, Yuanhan Mo, and Yike Guo. Automatic brain tumor detection and segmentation using u-net based fully convolutional networks. In annual conference on medical image understanding and analysis, pages 506–517. Springer, 2017.

- [51] Yu Weng, Tianbao Zhou, Yujie Li, and Xiaoyu Qiu. Nas-unet: Neural architecture search for medical image segmentation. *IEEE Access*, 7:44247–44257, 2019.
- [52] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International conference on medical image computing and computer-assisted intervention*, pages 424–432. Springer, 2016.
- [53] Hao Chen, Qi Dou, Lequan Yu, Jing Qin, and Pheng-Ann Heng. Voxresnet: Deep voxelwise residual networks for brain segmentation from 3d mr images. *NeuroImage*, 170:446–455, 2018.
- [54] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- [55] Stuart Berg, Dominik Kutra, Thorben Kroeger, Christoph N Straehle, Bernhard X Kausler, Carsten Haubold, Martin Schiegg, Janez Ales, Thorsten Beier, Markus Rudy, et al. ilastik: interactive machine learning for (bio) image analysis. *Nature Methods*, pages 1–7, 2019.
- [56] David A Van Valen, Takamasa Kudo, Keara M Lane, Derek N Macklin, Nicolas T Quach, Mialy M DeFelice, Inbal Maayan, Yu Tanouchi, Euan A Ashley, and Markus W Covert. Deep learning automates the quantitative analysis of individual cells in live-cell imaging experiments. *PLoS computational biology*, 12(11):e1005177, 2016.
- [57] Thorsten Falk, Dominic Mai, Robert Bensch, Özgün Çiçek, Ahmed Abdulkadir, Yassine Marrakchi, Anton Böhm, Jan Deubner, Zoe Jäckel, Katharina Seiwald, et al. U-net: deep learning for cell counting, detection, and morphometry. *Nature methods*, 16(1):67, 2019.
- [58] Jianxu Chen, Liya Ding, Matheus P Viana, Melissa C Hendershott, Ruian Yang, Irina A Mueller, and Susanne M Rafelski. The allen cell structure segmenter: a new open source toolkit for segmenting 3d intracellular structures in fluorescence microscopy images. *bioRxiv*, page 491035, 2018.
- [59] Matthias G Haberl, Christopher Churas, Lucas Tindall, Daniela Boassa, Sébastien Phan, Eric A Bushong, Matthew Madany, Raffi Akay, Thomas J Deerinck, Steven T Peltier, et al. Cdeep3m—plug-and-play cloud-based deep learning for image segmentation. *Nature methods*, 15(9):677, 2018.
- [60] Joan Bruna, Pablo Sprechmann, and Yann LeCun. Super-resolution with deep convolutional sufficient statistics. *arXiv preprint arXiv:1511.05666*, 2015.
- [61] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.

- [62] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.
- [63] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 1–9, 2015.
- [64] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 3431–3440, 2015.
- [65] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In International conference on machine learning, pages 647–655, 2014.
- [66] Chawin Ounkomol, Sharmishta Seshamani, Mary M Maleckar, Forrest Collman, and Gregory R Johnson. Label-free prediction of three-dimensional fluorescence images from transmitted-light microscopy. Nature methods, 15(11):917, 2018.
- [67] Eric M Christiansen, Samuel J Yang, D Michael Ando, Ashkan Javaherian, Gaia Skibinski, Scott Lipnick, Elliot Mount, Alison O’Neil, Kevan Shah, Alicia K Lee, et al. In silico labeling: predicting fluorescent labels in unlabeled images. Cell, 173(3):792–803, 2018.
- [68] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114, 2013.
- [69] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Advances in neural information processing systems, pages 2672–2680, 2014.
- [70] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. arXiv preprint arXiv:1601.06759, 2016.
- [71] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. In Advances in neural information processing systems, pages 4790–4798, 2016.
- [72] Li Deng. The mnist database of handwritten digit images for machine learning research [best of the web]. IEEE Signal Processing Magazine, 29(6):141–142, 2012.

- [73] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [74] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009.
- [75] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434, 2015.
- [76] Grigory Antipov, Moez Baccouche, and Jean-Luc Dugelay. Face aging with conditional generative adversarial networks. In 2017 IEEE International Conference on Image Processing (ICIP), pages 2089–2093. IEEE, 2017.
- [77] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In Proceedings of the IEEE International Conference on Computer Vision, pages 5907–5915, 2017.
- [78] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In Advances in neural information processing systems, pages 2172–2180, 2016.
- [79] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. arXiv preprint arXiv:1701.07875, 2017.
- [80] Gregory R Johnson, Taraz E Buck, Devin P Sullivan, Gustavo K Rohde, and Robert F Murphy. Joint modeling of cell and nuclear shape variation. Molecular biology of the cell, 26(22):4046–4056, 2015.
- [81] Gregory R Johnson, Jieyue Li, Aabid Shariff, Gustavo K Rohde, and Robert F Murphy. Automated learning of subcellular variation among punctate protein patterns and a generative model of their relation to microtubules. PLoS computational biology, 11(12):e1004614, 2015.
- [82] Aabid Shariff, Joshua Kangas, Luis Pedro Coelho, Shannon Quinn, and Robert F Murphy. Automated image analysis for high-content screening and analysis. Journal of biomolecular screening, 15(7):726–734, 2010.
- [83] Gustavo K Rohde, Alexandre JS Ribeiro, Kris N Dahl, and Robert F Murphy. Deformation-based nuclear morphometry: Capturing nuclear shape variation in hela cells. Cytometry Part A: The Journal of the International Society for Analytical Cytology, 73(4):341–350, 2008.



- [84] Aabid Shariff, Robert F Murphy, and Gustavo K Rohde. Automated estimation of microtubule model parameters from 3-d live cell microscopy images. In 2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pages 1330–1333. IEEE, 2011.
- [85] Tao Peng and Robert F Murphy. Image-derived, three-dimensional generative models of cellular organization. Cytometry Part A, 79(5):383–391, 2011.
- [86] Ting Zhao and Robert F Murphy. Automated learning of generative models for subcellular location: building blocks for systems biology. Cytometry Part A, 71(12):978–990, 2007.
- [87] Gregory R Johnson, Rory M Donovan-Maiye, and Mary M Maleckar. Generative modeling with conditional autoencoders: Building an integrated cell. arXiv preprint arXiv:1705.00092, 2017.
- [88] Anton Osokin, Anatole Chessel, Rafael E Carazo Salas, and Federico Vaggi. Gans for biological image synthesis. In Proceedings of the IEEE International Conference on Computer Vision, pages 2233–2242, 2017.
- [89] Juan C Caicedo, Claire McQuin, Allen Goodman, Shantanu Singh, and Anne E Carpenter. Weakly supervised learning of single-cell feature embeddings. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 9309–9318, 2018.
- [90] Romain Lopez, Jeffrey Regier, Michael B Cole, Michael I Jordan, and Nir Yosef. Deep generative modeling for single-cell transcriptomics. Nature methods, 15(12):1053, 2018.
- [91] Jiarui Ding, Anne Condon, and Sohrab P Shah. Interpretable dimensionality reduction of single cell transcriptome data with deep generative models. Nature communications, 9(1):2002, 2018.
- [92] Xiongtao Ruan and Robert F Murphy. Evaluation of methods for generative modeling of cell and nuclear shape. Bioinformatics, 2018.
- [93] Guanglei Ding, Yitong Liu, Rui Zhang, and Huolin L Xin. A joint deep learning model to recover information and reduce artifacts in missing-wedge sinograms for electron tomography and beyond. Scientific reports, 9(1):1–13, 2019.
- [94] Dawei Yang and Jia Deng. Learning to generate synthetic 3d training data through hybrid gradient. arXiv preprint arXiv:1907.00267, 2019.
- [95] Joel R Stiles, Thomas M Bartol, et al. Monte carlo methods for simulating realistic synaptic microphysiology using mcell. Computational neuroscience: realistic modeling for experimentalists, pages 87–127, 2001.

- [96] Leslie M Loew and James C Schaff. The virtual cell: a software environment for computational cell biology. TRENDS in Biotechnology, 19(10):401–406, 2001.
- [97] C. Lee, J. Laughlin, N. Angliviél de La Beaumelle, R. Amaro, J.A. McCammon, R. Ramamoorthi, M. Holst, and P. Rangamani. GAMer 2: A system for 3D mesh processing of cellular electron micrographs. Submitted for publication. Available as [arXiv:1901.11008 \[q-bio.QM\]](https://arxiv.org/abs/1901.11008), 2019.
- [98] C. Lee, J. Laughlin, J. Moody, R. Amaro, J.A. McCammon, M. Holst, and P. Rangamani. An open source mesh generation platform for biophysical modeling using realistic cellular geometries. Submitted for publication. Available as [arXiv:1909.04781 \[physics.comp-ph\]](https://arxiv.org/abs/1909.04781), 2019.
- [99] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 652–660, 2017.
- [100] Rana Hanocka, Amir Hertz, Noa Fish, Raja Giryes, Shachar Fleishman, and Daniel Cohen-Or. Meshcnn: a network with an edge. ACM Transactions on Graphics (TOG), 38(4):90, 2019.
- [101] C. Lee, J. Moody, R. Amaro, J. McCammon, and M. Holst. The implementation of the colored abstract simplicial complex and its application to mesh generation. ACM Trans. Math. Software, 45(3):28:1–28:20, August 2019. Available as [arXiv:1807.01417 \[math.NA\]](https://arxiv.org/abs/1807.01417).
- [102] R. E. Bank and R. K. Smith. Mesh smoothing using a posteriori error estimates. SIAM J. Numer. Anal., 34:979–997, 1997.
- [103] A. Liu and B. Joe. Quality local refinement of tetrahedral meshes based on bisection. SIAM J. Sci. Statist. Comput., 16(6):1269–1291, 1995.
- [104] J.M. Maubach. Local bisection refinement for N-simplicial grids generated by relection. SIAM J. Sci. Statist. Comput., 16(1):210–277, 1995.
- [105] J. Bey. Tetrahedral grid refinement. Computing, 55(4):355–378, 1995.
- [106] D.N. Arnold, A. Mukherjee, and L. Pouly. Locally adapted tetrahedral meshes using bisection. SIAM J. Sci. Statist. Comput., 22(2):431–448, 1997.
- [107] Z. Gao, Z. Yu, and M. Holst. Feature-preserving surface mesh smoothing via suboptimal Delaunay triangulation. Graphical Models, 75(1):23–38, 2013.
- [108] L. Chen and M. Holst. Efficient mesh optimization schemes based on optimal Delaunay triangulations. Comp. Meth. in Appl. Mech. Engr., 200(9–12):967–984, 2011.

- [109] I. Babuška and W.C. Rheinboldt. A posteriori error estimates for the finite element method. Internat. J. Numer. Methods Engrg., 12:1597–1615, 1978.
- [110] I. Babuška and W.C. Rheinboldt. Error estimates for adaptive finite element computations. SIAM J. Numer. Anal., 15:736–754, 1978.
- [111] Felix Fritzen, Mauricio Fernández, and Fredrik Larsson. On-the-fly adaptivity for nonlinear twoscale simulations using artificial neural networks and reduced order modeling. Frontiers in Materials, 6:75, 2019.
- [112] Larry M. Manevitz, Akram Bitar, and Dan Givoli. Neural network time series forecasting of finite-element mesh adaptation. Neurocomputing, 63:447–463, 2005.
- [113] Andrea Mendizabal, Pablo Márquez-Neila, and Stéphane Cotin. Simulation of hyperelastic materials in real-time using deep learning. arXiv preprint arXiv:1904.06197, 2019.
- [114] Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. Science, 365(6457):eaaw1147, 2019.
- [115] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In Proceedings of the European Conference on Computer Vision (ECCV), pages 52–67, 2018.
- [116] Matthew Amodio, David Van Dijk, Krishnan Srinivasan, William S Chen, Hussein Mohsen, Kevin R Moon, Allison Campbell, Yujiao Zhao, Xiaomei Wang, Manjunatha Venkataswamy, et al. Exploring single-cell data with deep multitasking neural networks. BioRxiv, page 237065, 2019.

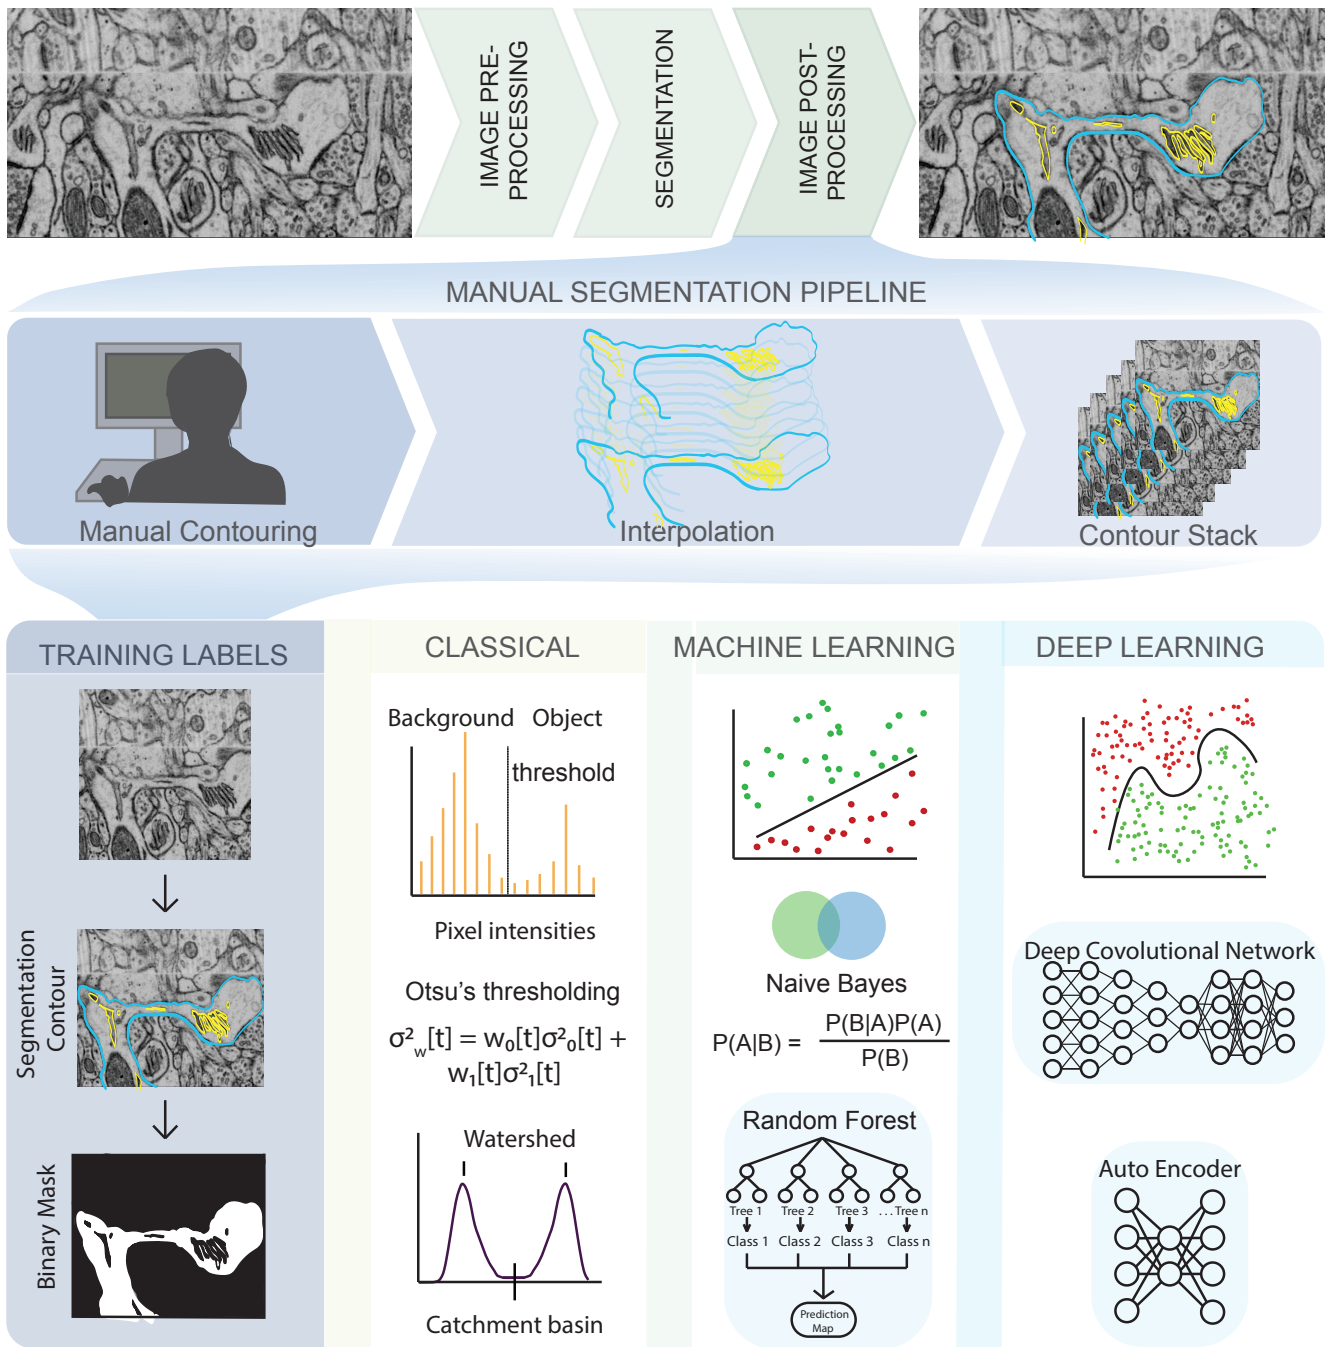


Figure 1: An illustration of the complex pipeline needed to go from imaging data to a segmented mesh, with various opportunities for emerging techniques in machine learning shown throughout the pipeline. Top row: EM images obtained from [24] of dendritic spines from mouse brain tissue. Middle row: Manual tracing or contouring, interpolation, and stacking of contours is extremely time consuming, prone to error, and relies of human judgement. Bottom row: On the other hand, development of training labels and different learning techniques can reduce both time and error, bridging the gap between biological data and simulations. The list of techniques described is representative only, and not exhaustive.

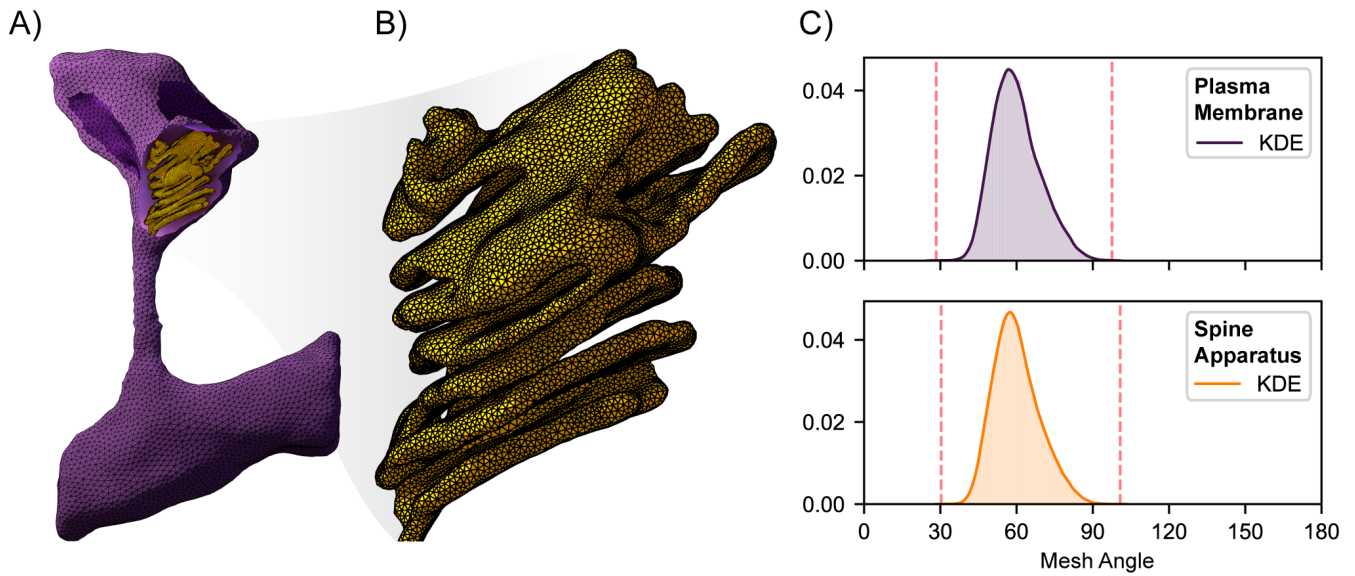


Figure 2: An illustration of complexity, size, quality, and local resolution of meshes typically needed for realistic simulation of biophysical systems. Meshes are generated using *GAMer 2* [97, 98]. A) Example surface mesh of a dendritic spine with geometry informed by electron micrographs from [24]. The plasma membrane is shown in purple with the post synaptic density rendered in dark purple. The spine apparatus, a specialized form of the endoplasmic reticulum is shown in yellow. B) A zoomed in view of the spine apparatus. Note that the mesh density is much higher in order to represent the fine structural details. C) Binned histogram distributions of mesh angles for both the plasma membrane and spine apparatus. The colored smooth lines are the result of a kernel density estimate. Dotted red lines correspond to the minimum and maximum angle values in each mesh. Both meshes are high quality with few high aspect ratio triangles (i.e., those deviating most from equilateral).