

Unpaired Modality Translation for Pseudo Labeling of Histology Images

Arthur Boschet², Armand Collin^{1,2}, Nishka Katoch¹, and Julien Cohen-Adad^{1,2}

¹ NeuroPoly Lab, Institute of Biomedical Engineering, Polytechnique Montréal, Montréal, Québec, Canada

² Mila - Québec Artificial Intelligence Institute, Montréal, Québec, Canada

Abstract. The segmentation of histological images is critical for various biomedical applications, yet the lack of annotated data presents a significant challenge. We propose a microscopy pseudo labeling pipeline utilizing unsupervised image translation to address this issue. Our method generates pseudo labels by translating between labeled and unlabeled domains without requiring prior annotation in the target domain. We evaluate two pseudo labeling strategies across three image domains increasingly dissimilar from the labeled data, demonstrating their effectiveness. Notably, our method achieves a mean Dice score of 0.736 ± 0.005 on a SEM dataset using the tutoring path, which involves training a segmentation model on synthetic data created by translating the labeled dataset (TEM) to the target modality (SEM). This approach aims to accelerate the annotation process by providing high-quality pseudo labels as a starting point for manual refinement.

Keywords: image segmentation · image translation · histology · axon · myelin · pseudo labels

1 Introduction

Histology images play a crucial role in the study of neurodegenerative disorders, such as Parkinson’s [13] or Alzheimer’s disease [9]. High resolution histology imaging of neurological tissues can be obtained using scanning (SEM) or transmission (TEM) electron microscopy, bright field imaging (BF) or Coherent anti-Stokes Raman spectroscopy (CARS), each producing images of different resolution and contrasts, making it difficult to train a segmentation model that works across a variety of imaging technologies and acquisition parameters. Recent advances in deep learning have largely improved models to segment microscopy images [19, 8]. For instance, the U-Net architecture [11], a CNN-based model has become a standard in biomedical image segmentation due to its ability to produce precise segmentation with the help of annotated data [4]. Unfortunately, one challenge in developing effective models of biomedical segmentation is the low amount of labels, which are modality specific. Annotating histology images is a labor and time-intensive process that requires expert knowledge. A new annotation strategy is to use image translation techniques to augment existing

datasets and generate synthetic annotated images. With translations often unpaired, traditionally generative adversarial networks (GANs) [1] have been used. Diffusion models [2] have recently emerged as a robust alternative, based on explicit likelihood characterization and a gradual sampling process. We utilize SynDiff [10], an adversarial diffusion model that uses conditional diffusion for efficient modality translation.

Our main contribution is a microscopy pseudo labeling pipeline based on unsupervised image translation. The proposed weakly supervised method generates pseudo labels by image translation between the labeled and unlabeled data, and does not require any prior annotation in the target domain. This contrasts with transfer learning, which is the most prevalent approach to mitigate data scarcity in medical deep learning. Unlike [18], where image translation was used to perform pseudo labeling for object detection in microscopy images, we explore this idea for a semantic segmentation task and compare two pseudo labeling strategies on three increasingly challenging target datasets, while providing recommendations as to when to deploy this pipeline for optimal results. The first strategy trains a proxy segmentation model on synthetic data derived from modality-translated labeled images, while the second translates unlabeled images to match a labeled data distribution, enabling segmentation with a pre-trained model. Our code is available online ³.

2 Methods

We have two similar image domains: L (labeled) and U (unlabeled). The task we perform on these two domains is shared — in our case, segmentation of the axon and myelin tissues. Images X_L sampled from L are annotated with segmentation labels Y_L , whereas images X_U sampled from U are unlabeled. Typically, we have a model M_L pre-trained on $\{X_L, Y_L\}$ but it doesn't generalize well to U . Our goal is ultimately to generate pseudo labels \tilde{Y}_U to bootstrap and accelerate the annotation of images X_U .

As shown in Fig. 1b, our approach is to perform image translation between domains L and U , resulting in synthetic images $X_{L \rightarrow U}$ and $X_{U \rightarrow L}$. We denote the translation directions $L \rightarrow U$ as *tutoring path* and $U \rightarrow L$ as *adaptive path*. Using the tutoring path, we obtain a synthetic dataset $\{X_{L \rightarrow U}, Y_L\}$, effectively recycling the previously available annotations. A segmentation model is trained on this synthetic dataset and subsequently applied to X_U to generate pseudo labels $\tilde{Y}_{U,t}$. We denote this process pseudo labeling by *tutorship*, because we use the annotations of data from L to guide prediction on images from U , as illustrated in Fig. 1c.

X_L and X_U are unpaired, so we use an unsupervised image translation method based on cycle-consistency. Thus, although we are interested in the tutoring path yielding $X_{L \rightarrow U}$, we also get $X_{U \rightarrow L}$ from the adaptive path at no additional cost. We apply the pre-trained model M_L on $X_{U \rightarrow L}$ to generate

³ <https://github.com/axondeepseg/AxonDeepSynth>

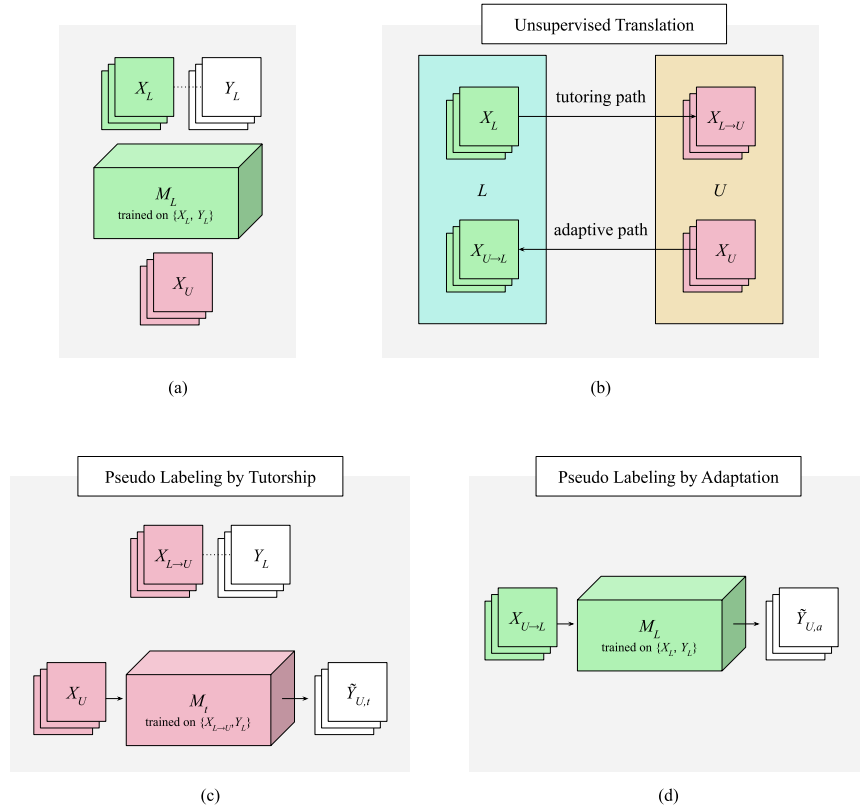


Fig. 1. Methodology overview. (a) Initial data and model (b) Unpaired translation between L and U via Adversarial Diffusion (c) Tutorship: Training a segmentation model on synthetic data $\{X_{L \rightarrow U}, Y_L\}$ for inference in U (d) Adaptation: Applying a segmentation model pre-trained on L to the translated images $X_{U \rightarrow L}$

another set of pseudo labels $\tilde{Y}_{U,a}$. We refer to this complementary process as pseudo labeling by *adaptation*, because we are transferring the style of images X_L (used to train M_L) onto images X_U , as shown in Fig. 1d.

2.1 Data

The 3 unlabeled datasets used in this project are described in Table 1, along with the reference annotated TEM dataset. Note the diversity in input modality (TEM/SEM/BF), species, pathology, body part, spatial resolution and median size. This diversity is illustrated in Fig. S1 in the Supplementary Material for a dataset preview. The TEM dataset represents the X_L annotated images across all experiments. We evaluate our pipeline in 3 different settings, using in turn

the TEM-MACAQUE, SEM and BF datasets to simulate the unannotated data pool X_U . In reality, these datasets are annotated so the pseudo labels can be evaluated quantitatively. Conveniently, the X_L data pool has the smallest pixel size (highest magnification). As such, X_L is resized to match the X_U pixel size in all 3 experiments without significant loss of details.

Table 1. Dataset details

dataset	TEM	TEM-MACAQUE	SEM	BF
species	mouse	macaque	rat	rat, rabbit, human
pathology	healthy	healthy	healthy	healthy, demyelination
organ	brain	brain	spinal cord	brain, spinal cord, nerves, muscles
pixel size ($\mu\text{m}/\text{px}$)	0.00236	0.009	0.1	0.1 0.211 0.226
median size	2286x3762	3040x2300	924x1278	1255x1440
nb. of images	158	16	10	39

2.2 Models

Image Translation To perform unpaired modality translation, we utilized the SynDiff framework, which leverages conditional adversarial diffusion models for translation [10]. Training on unpaired data is enabled by the use of a cycle consistency loss on a pair of non-diffusive GANs which generate initial unrefined estimates of the target translations. These GANs employ a 3-layer discriminator and a 6-block ResNet generator, both utilizing instance normalization, adapting the architecture from the pix2pix repository [7]. Subsequently, these preliminary translations are passed as conditioning inputs to the denoising diffusion GANs, which are thus trained to produce translations between two modalities in an unpaired manner. At inference time, the non-diffusive modules are ignored, and the source image is directly used as the conditioning input for the denoising diffusion GANs.

The use of denoising diffusion GANs is motivated by resolving the trade-offs between the sampling efficiency of GANs and the training stability and mode coverage of diffusion models [17]. It has been demonstrated that with infinitesimally small time-steps, the reverse diffusion process mirrors the Gaussian nature of the forward diffusion process. Thus, if we model the forward diffusion process with a Gaussian distribution, we can assume the reverse process will also follow

a Gaussian distribution [2, 14, 17]. As such, one major limitation of standard diffusion models is their requirement for a large number of timesteps with minimal variance steps for this assumption to hold, significantly increasing sampling time. This is suboptimal for tasks such as translating whole slide images. In contrast, GANs generate high-quality samples with low sampling times but often suffer from reduced sample diversity due to issues like mode collapse [12, 20].

Denosing diffusion GANs combine the advantages of both approaches—high quality and sampling speed with extensive mode coverage—by bypassing the assumption of normally distributed data at each diffusion step and using fewer than 8 timesteps. To achieve this, they employ GANs to model transitions between each denosing step [17]. It can be shown, based on the formulation by Ho et al., that the original DDPM can be equivalent to first predicting x_0 and then reapplying noise via $p_\theta(x_{t-1}|x_t) = q(x_{t-1}|x_t, x_0 = f_\theta(x, t))$. Denosing diffusion GANs parameterize f_θ using GANs instead of deterministically predicting the original image x_0 [17]. Specifically, SynDiff parameterizes f_θ using the NCSN++ generator [15], with a U-net structure featuring ResNet and Attention blocks, z-conditioning via mapping networks, and adaptive group normalization, while its time-dependent discriminator employs a convolutional ResNet architecture, sinusoidal positional embeddings, and minibatch standard deviation layers, following the implementation described by Xiao et al. [17]. For more details on the training objective, please refer to the supplementary materials.

To evaluate SynDiff checkpoints, we used SSIM [16, 3], PSNR [3], and $L1$ -loss. These metrics, requiring ground truths, were computed after a full $L \rightarrow U \rightarrow L$ translation cycle, comparing the reconstruction $X_{L \rightarrow U \rightarrow L}$ with the original labeled image X_L .

Image Segmentation The nnU-Net framework has been selected for our study due to its demonstrated robustness and superior performance in medical image segmentation tasks. Its automated configuration and adaptability make it particularly effective for small and diverse medical datasets. nnU-Net’s design principles ensure rigorous validation and optimal architectural choices, consistently delivering state-of-the-art results [5, 6]. nnU-Net employs a 5-fold cross-validation approach, training five different models on five different folds. The best checkpoint from each fold is then used to compute the standard deviation of the segmentation scores.

2.3 Experiments

The following 3 experiments are increasing challenging, as the domains L and U are progressively more dissimilar. In each case, we evaluate pseudo labeling strategies using the tutoring and adaptive paths. Both methods are illustrated in Fig. 2.

TEM \leftrightarrow TEM-MACAQUE In this experiment, both X_L and X_U were acquired using transmission electron microscopy at similar magnifications. The

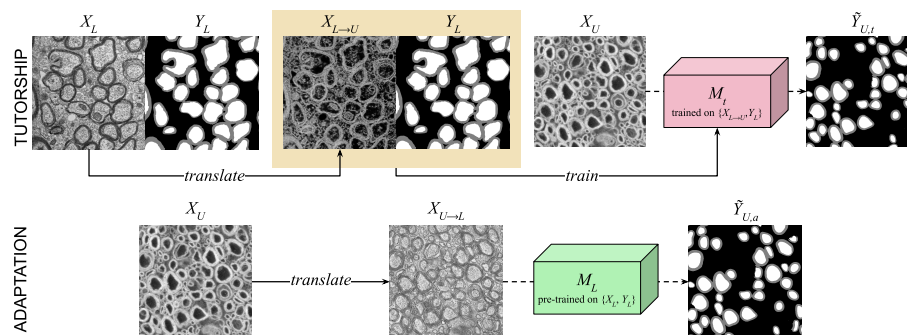


Fig. 2. Overview of the two pseudo labeling strategies using TEM as the labeled dataset and SEM as the unlabeled data. (Top) The Tutorship strategy involves generating a synthetic dataset by translating labeled TEM images and training a proxy segmentation model on this synthetic data. (Bottom) The Adaptation strategy converts unlabeled SEM images to match the TEM data distribution, enabling segmentation with a pre-trained model.

main difference between these datasets is the provenance of the samples in terms of species, and a slightly different intensity profile (most likely due to sample preparation and acquisition parameters). This experiment acts as a baseline and shows how our pseudo labeling method performs on a small domain shift.

TEM \leftrightarrow SEM Here, X_L and X_U were acquired with different modalities. TEM and SEM have opposite contrasts (i.e. myelin is dark in TEM, white in SEM. See Supplemental Materials, Fig. S1). The species, body part and pixel size are also different, resulting in an intermediate domain shift.

TEM \leftrightarrow BF This final experiment investigates a larger domain shift between X_L and X_U . The BF dataset exhibits a wide diversity across 4 organs, including 3 different species, pathology and pixel size.

3 Results and Discussion

Learning curves featuring the peak signal-to-noise ratio (PSNR), the structural similarity index measure (SSIM), and the $L1$ -loss of the image translation models after an $L \rightarrow U \rightarrow L$ translation cycle are reported in Fig. 3. These curves, along with visual assessment of the generation process, guided the selection of the best checkpoints. Specifically, checkpoints for TEM \leftrightarrow TEM-MACAQUE, TEM \leftrightarrow SEM, and TEM \leftrightarrow BF were selected at epochs 100, 50, and 40, respectively.

As illustrated in Fig. 4, our pseudo labeling strategy is effective, especially using the tutoring path. The Dice score, sensitivity, and specificity computed on the pseudo labels, along with the standard deviations of each metric across all

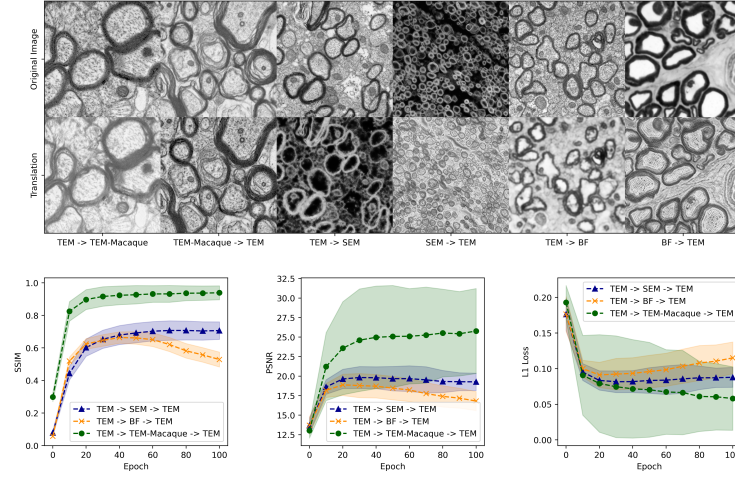


Fig. 3. (Top) Image translation examples. (Bottom) Unsupervised translation reconstruction losses: Structural Similarity Measure (SSIM), Peak Signal-to-Noise Ratio (PSNR) and $L1$. The error bars represent the standard deviations among the validation images.

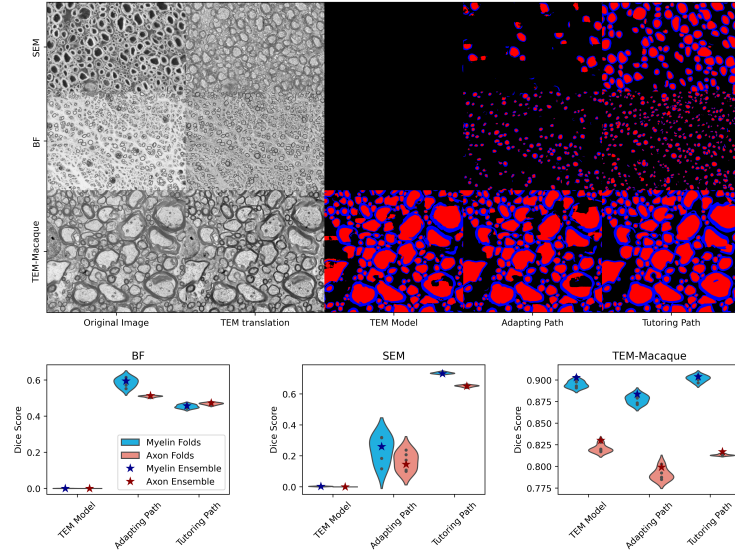


Fig. 4. Pseudo labeling results. (Top) Original and translated image alongside the pre-trained model prediction and pseudo labels. (Bottom) Dice scores under 5-fold cross-validation.

five folds of the nnU-Net segmentations, are reported in Table 2. For the TEM \leftrightarrow TEM-MACAQUE experiment, using the pre-trained model provides a similar performance to our pseudo labeling strategies. This highlights the fact that when X_L and X_U are sufficiently similar, using a model pre-trained on X_L is more than enough to produce good pseudo labels. However, this is not the case for the two remaining experiments, where the pre-trained TEM model performance drops to zero, as expected due to larger domain shifts. Both the TEM \leftrightarrow SEM and TEM \leftrightarrow BF experiments demonstrate the benefit of our pseudo labeling strategies. In the TEM \leftrightarrow SEM case, the tutoring path provides a better prediction than the adapting path, the latter having a much lower sensitivity. Using the tutorship method, the mean Dice score averages 0.736 ± 0.005 for axons and 0.652 ± 0.005 for myelin, a remarkable starting point in the context of manual annotation. In the TEM \leftrightarrow BF scenario, the adapting path marginally surpasses the tutoring path, achieving a mean Dice score of 0.586 ± 0.025 for axons and 0.510 ± 0.004 for myelin. While further research is required to elucidate the conditions under which the tutoring path outperforms the adapting path and vice versa, it is plausible that the performance is contingent on the efficacy of the translation model in either the adapting $U \rightarrow L$ or tutoring $L \rightarrow U$ direction.

Table 2. Pseudo Label Dice Score, Specificity and Sensitivity. TEM-M stands for the TEM-MACAQUE dataset. Applying the pre-trained TEM model directly on the SEM/BF datasets produced empty masks, so we omit these cases.

Class	X_U	Method	Dice Score	Specificity	Sensitivity
Axon	TEM-M	TEM model	0.895 ± 0.004	0.935 ± 0.007	0.905 ± 0.015
		Adaptation	0.878 ± 0.005	0.936 ± 0.006	0.872 ± 0.017
		Tutorship	0.902 ± 0.003	0.932 ± 0.001	0.922 ± 0.006
	SEM	Adaptation	0.238 ± 0.088	0.984 ± 0.010	0.144 ± 0.061
		Tutorship	0.736 ± 0.005	0.914 ± 0.002	0.705 ± 0.009
	BF	Adaptation	0.586 ± 0.025	0.949 ± 0.008	0.889 ± 0.022
Tutorship		0.454 ± 0.010	0.915 ± 0.003	0.942 ± 0.003	
Myelin	TEM-M	TEM model	0.820 ± 0.004	0.945 ± 0.009	0.811 ± 0.021
		Adaptation	0.791 ± 0.007	0.938 ± 0.005	0.777 ± 0.022
		Tutorship	0.813 ± 0.001	0.934 ± 0.000	0.822 ± 0.001
	SEM	Adaptation	0.165 ± 0.060	0.995 ± 0.003	0.092 ± 0.036
		Tutorship	0.652 ± 0.005	0.962 ± 0.002	0.512 ± 0.008
	BF	Adaptation	0.510 ± 0.004	0.949 ± 0.003	0.635 ± 0.020
Tutorship		0.469 ± 0.005	0.947 ± 0.002	0.566 ± 0.012	

4 Conclusion

Our results provide a better understanding of the optimal scenarios where leveraging image translation is advantageous for pseudo-labeling. When the target unlabeled dataset is sufficiently similar to the available labeled data, our pseudo

labeling strategies do not provide a significant advantage over applying a pre-trained model, as demonstrated in the TEM \leftrightarrow TEM-MACAQUE experiment. However, in both the TEM \leftrightarrow SEM and TEM \leftrightarrow BF scenarios, the target dataset is dissimilar enough that applying a pre-trained model is ineffective. In these cases, our pseudo labeling strategies are effective and provide useful initial masks to be manually corrected instead of annotated from scratch, saving precious time in the process. Realistically, a Dice score superior to 0.5 goes a long way, reducing the annotation time by 25%-50% (although this should be further quantified in future works). We recommend first using the adaptive path which does not require training an intermediate segmentation model. The tutoring path can provide a complementary set of pseudo labels, at the additional cost of training a proxy segmentation model. We hope this work inspires other groups to recycle available datasets when applicable to acquire more labeled data.

Acknowledgments. We would like to thank Simeon Christian Daeschler, Marie-Hélène Bourget, Tessa Gordon, Gregory Howard Borschel, Charles R. Reiter, Geetanjali Bendale and Osvaldo Delbono for the BF dataset.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

Bibliography

- [1] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *Advances in neural information processing systems* **27** (2014)
- [2] Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
- [3] Hore, A., Ziou, D.: Image quality metrics: Psnr vs. ssim. In: 2010 20th international conference on pattern recognition. pp. 2366–2369. IEEE (2010)
- [4] Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (Feb 2021). <https://doi.org/10.1038/s41592-020-01008-z>
- [5] Isensee, F., Petersen, J., Klein, A., Zimmerer, D., Jaeger, P.F., Kohl, S., Wasserthal, J., Koehler, G., Norajitra, T., Wirkert, S., et al.: nnu-net: Self-adapting framework for u-net-based medical image segmentation. *arXiv preprint arXiv:1809.10486* (2018)
- [6] Isensee, F., Wald, T., Ulrich, C., Baumgartner, M., Roy, S., Maier-Hein, K., Jaeger, P.F.: nnu-net revisited: A call for rigorous validation in 3d medical image segmentation. *arXiv preprint arXiv:2404.09556* (2024)
- [7] Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1125–1134 (2017)
- [8] Litjens, G., et al.: A survey on deep learning in medical image analysis. *Medical Image Analysis* **42**, 60–88 (Dec 2017). <https://doi.org/10.1016/j.media.2017.07.005>
- [9] McKenzie, A.T., Marx, G.A., Koenigsberg, D., Sawyer, M., Iida, M.A., Walker, J.M., Richardson, T.E., Campanella, G., Attems, J., McKee, A.C., et al.: Interpretable deep learning of myelin histopathology in age-related cognitive impairment. *Acta Neuropathologica Communications* **10**(1), 131 (2022)
- [10] Özbey, M., Dalmaz, O., Dar, S.U., Bedel, H.A., Öztürk, Ş., Güngör, A., Çukur, T.: Unsupervised medical image translation with adversarial diffusion models. *IEEE Transactions on Medical Imaging* (2023)
- [11] Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation (May 2015). <https://doi.org/10.48550/arXiv.1505.04597>
- [12] Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. *Advances in neural information processing systems* **29** (2016)
- [13] Signaevsky, M., Marami, B., Prastawa, M., Tabish, N., Iida, M.A., Zhang, X.F., Sawyer, M., Duran, I., Koenigsberg, D.G., Bryce, C.H., Chahine, L.M., Mollenhauer, B., Mosovsky, S., Riley, L., Dave, K.D., Eberling, J.,

- Coffey, C.S., Adler, C.H., Serrano, G.E., White, C.L., Koll, J., Fernandez, G., Zeineh, J., Cordon-Cardo, C., Beach, T.G., Crary, J.F.: Antemortem detection of parkinson's disease pathology in peripheral biopsies using artificial intelligence. *Acta Neuropathologica Communications* **10**(1) (Feb 2022). <https://doi.org/10.1186/s40478-022-01318-7>, <http://dx.doi.org/10.1186/s40478-022-01318-7>
- [14] Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *International conference on machine learning*. pp. 2256–2265. PMLR (2015)
- [15] Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456* (2020)
- [16] Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
- [17] Xiao, Z., Kreis, K., Vahdat, A.: Tackling the generative learning trilemma with denoising diffusion gans. *arXiv preprint arXiv:2112.07804* (2021)
- [18] Xing, F., Bennett, T., Ghosh, D.: Adversarial Domain Adaptation and Pseudo-Labeling for Cross-Modality Microscopy Image Quantification, p. 740–749. *Springer International Publishing* (2019). https://doi.org/10.1007/978-3-030-32239-7_82, http://dx.doi.org/10.1007/978-3-030-32239-7_82
- [19] Zaimi, A., Wabartha, M., Herman, V., Antonsanti, P.L., Perone, C.S., Cohen-Adad, J.: AxonDeepSeg: automatic axon and myelin segmentation from microscopy data using convolutional neural networks. *Scientific Reports* **8**(1), 3816 (Feb 2018). <https://doi.org/10.1038/s41598-018-22181-4>
- [20] Zhao, S., Ren, H., Yuan, A., Song, J., Goodman, N., Ermon, S.: Bias and generalization in deep generative models: An empirical study. *Advances in Neural Information Processing Systems* **31** (2018)