SeidelNet: An aberration informed deep learning model for spatially-varying deblurring

Esther Whang*[a], David McAllister*[b], Ashwin Reddy[b], Amit Kohli[b], and Laura Waller[b]

[a]The Cooper Union for the Advancement of Science and Art, New York, USA
[b]University of California, Berkeley, California, USA

ABSTRACT

Practical imaging systems form images with spatially-varying blur, making it challenging to deblur them and recover critical scene features. To address such systems, we introduce SeidelNet, a deep-learning approach for spatially-varying deblurring which learns to invert an imaging system's blurring process from a single calibration image. SeidelNet leverages the rotational symmetry present in most imaging systems by incorporating the primary Seidel aberration coefficients into the deblurring pipeline. We train and test SeidelNet on synthetically blurred images from the CARE fluorescence microscopy dataset, and find that, despite relatively few parameters, SeidelNet outperforms both analytical methods as well as a standard deblurring neural network.

Keywords: Deep Learning, Image Deblurring, Computational Imaging, Aberration Estimation, Aberration Correction, Point Spread Function, Deconvolution

1. INTRODUCTION

Whether miniature microscopes, large-scale space telescopes, or cell phone cameras, many practical imaging systems are subject to imperfections that are infeasible—and in some cases impossible—to fix with optical design and engineering. Size and cost constraints, among others, may force one to turn instead to digital solutions, namely image deblurring. It is no surprise, then, that there exists a plethora of deblurring techniques. What is more surprising is that existing literature has largely been polarized—either deblurring is done analytically by inverting a physically-informed image formation model[1–9] or it is done purely phenomenologically, primarily with deep learning.[10–12] Thus one must trade off slow, calibration-heavy, but robust methods with fast, accurate methods which require large datasets and are potentially unreliable at test time.

Only more recently do we see entries emerge which exist in between.[13–17] However these techniques are often specialized to a particular system and/or still require nontrivial calibration. In this work, we develop a general-purpose deep learning strategy called SeidelNet which incorporates the 5 primary Seidel aberration coefficients in order to perform fast, accurate, and robust deblurring with only a single calibration image per imaging system. The key idea behind our hybrid model is that most imaging systems are rotationally symmetric which means their system aberrations can be parameterized with the Seidel polynomial.[18,19] We learn the 5 primary coefficients of the Seidel polynomial from a single image of a few randomly scattered point sources and use them to obtain a system-specific, aberration-aware deblurring neural network.

To test SeidelNet we first create a dataset by applying radially-varying blur to images from the CARE dataset[10,20] using the linear revolution-invariant forward model.[21] Then, we build and train a series of different models, each of which takes both a blurry image and a corresponding set of Seidel aberration coefficients as input, and outputs a deblurred estimate of the perfect image. Overall, we show that the models which incorporate the Seidel coefficients outperform their unconditioned counterparts as well as standard analytical methods.

*These authors contributed equally
Further author information: (Send correspondence to Esther Whang)
Esther Whang: E-mail: esther.whang@cooper.edu

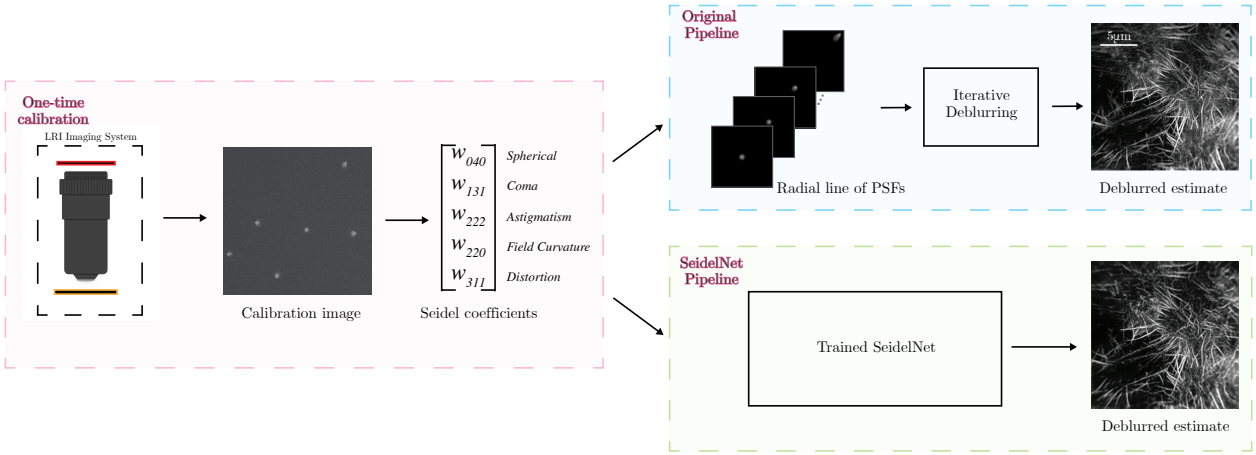

Figure 1: **Comparing the LRI Deblurring pipelines.** The original, pictured in blue, requires an additional step of generating a stack of PSFs. The SeidelNet pipeline, on the other hand, can directly utilize the Seidel Coefficients.

2. BACKGROUND

2.1 Linear revolution-invariant deblurring

The blur of an aberrated imaging system, in general, varies across the field-of-view. Systems for which this field-dependence is small may be modeled as linear, shift-invariant (LSI). Their images are accurately deblurred via deconvolution with a single point spread function (PSF).[1–3] However, systems which do not have negligible field variance but are rotationally symmetric are more appropriate modeled as linear, revolution-invariant or LRI; the PSFs vary, but only radially (see Fig. 1 Radial line of PSFs). The images from such systems require LRI deblurring which uses many PSFs—one at every radius of the FoV—and is N times slower where N is the image sidelength in pixels.[21] However, it is possible to obtain these needed PSFs by first fitting a system's Seidel aberration coefficients to a single calibration image of randomly scattered point sources, and then generating the necessary PSFs (see Fig. 1 original pipeline). Yet such a procedure is incredibly inefficient: it expands 5 numbers to N PSFs which is a total of N^3 pixels. In this paper we propose to instead directly use the Seidel aberration coefficients when performing image deblurring. But before we introduce our technique, SeidelNet, we will provide a brief background on Seidel aberration coefficients. A more comprehensive guide may be found elsewhere.[18,19]

2.2 Seidel aberration coefficients

The aberrations of a linear imaging system are often characterized by the wavefront error function w whose values correspond to the optical path length difference between the system's wavefront and an ideal spherical wavefront at the exit pupil of the system.[18,22] For rotationally symmetric systems, one can expand w as a infinite power series whose third-order coefficients are the 5 primary Seidel aberration coefficients.[18,19] Each coefficient corresponds to a well-known type of aberration: sphere, coma, astigmatism, field curvature, and distortion. Using the Seidel calibration procedure described in[21] (see Fig. 1), it is possible to obtain the Seidel coefficients from a single calibration image of a few randomly scattered point sources. Seidel coefficients provide an incredibly compact summary—just 5 numbers—of the system aberrations. It is then natural to conclude that they can enhance a deblurring network when used as co-input with the blurry measurement, which we describe as follows.

3. METHODS

We implement two strategies for incorporating Seidel coefficients into deep deblurring—Seidel U-Net and Seidel HyperU-Net. As a baseline we use a standard U-Net.[10,11,23]

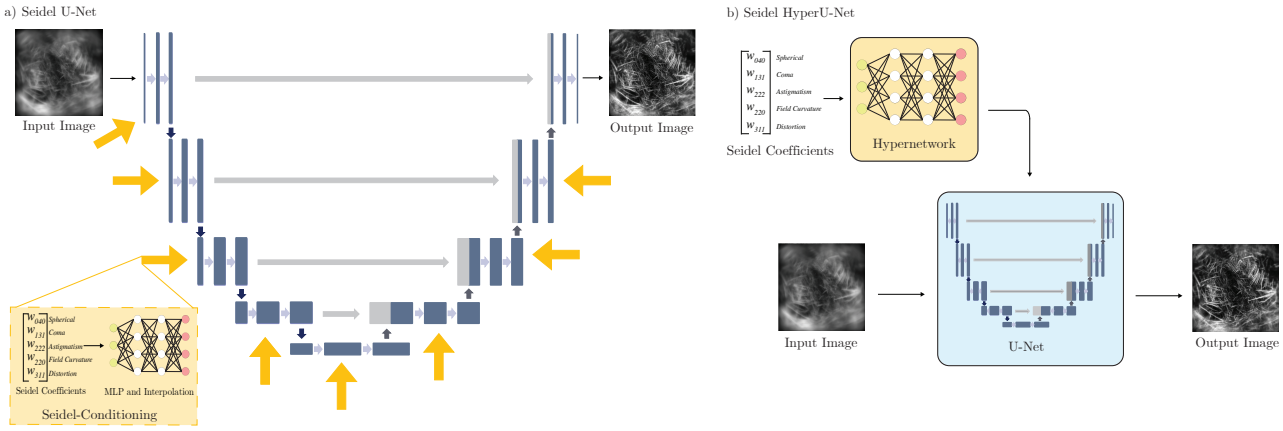

Figure 2: **SeidelNet Architectures**. a) Seidel U-Net takes in the blurry measurement and Seidel coefficients at every convolutional block. b) Seidel HyperU-Net first takes in the Seidel coefficients and produces a deblurring U-Net, which then takes in the blurry measurement.

3.0.1 Seidel U-Net

Seidel U-Net utilizes the classic U-Net architecture but takes in Seidel coefficients in addition to the blurry image as input (see Fig. 2a). To incorporate the Seidel coefficients, we jointly optimize a small three-layer multilayer perceptron (MLP) alongside the usual U-Net. This MLP lifts the Seidel coefficients into a low-dimensional image space, which is appended to the convolutional blocks of the U-Net along the channel dimension. To ensure the MLP output matches the correpsonding convolutional block dimension, we apply a combination of upscale interpolations and small convolutional layers. To compensate for these extra parameters of the MLP, we reduce the kernel size to 3x3, resulting in a model of ~ 54.4 million parameters compared to the ~ 92.7 million parameters of the baseline U-Net.

3.0.2 Seidel HyperU-Net

Seidel HyperU-Net is an interpretable and parameter-efficient alternative to Seidel U-Net. Seidel HyperU-Net builds on Hypernetworks,[24] a model that tasks a "hypernetwork" with predicting the weights of a "primary" network. In our case, the hypernetwork takes in a set of Seidel coefficients as input and maps them to the convolutional kernels of a scaled-down U-Net. This U-Net then takes in a blurry image as input and outputs a deblurred estimate (see Fig. 2b). In essence, each Seidel coefficient set essentially indexes a deblurring network specifically optimized to handle the aberrations associated with those coefficients. We implement the hypernetwork with a shared MLP and a series of learned embeddings for each layer of Seidel HyperU-Net. The output of this MLP is then multiplied with a learned matrix and reshaped to be a convolutional kernel.

3.1 Dataset

To produce a training dataset for SeidelNet, we synthetically generate blurred input images from the CARE fluorescence microscopy dataset[10,20] using the linear revolution-invariant forward model with randomly chosen sets of Seidel coefficients.[21] Thus, each training pair consists of a two-part input, the LRI-blurred image and its corresponding five primary Seidel coefficients. The unblurred ground truth image from the CARE dataset supervises the model output.

For training, we must sample Seidel coefficients which adequately cover realistic imaging systems—we find that the range $0-3$ for each coefficient is sufficient. We elect to combine two random sampling grids on $\mathbb{R}^5$. The first is a uniform distribution across Seidel coefficients between 0 and 3. The second is a nonrandom, evenly-spaced grid across the same range of coefficients with additive zero-mean Gaussian noise at each grid point.

3.2 Implementation

We implement our deblur models using the PyTorch framework in Python. The U-Net models are modified from an open-source implementation,[25] and the HyperU-Net model incorporates elements from an open-source Hypernetwork implementation.[26] All models are trained on the same training set and tested on the same test set.

4. RESULTS AND DISCUSSION

Below we compare the results of the aforementioned deep learning models as well as the analytical deblurring algorithms for the LSI and LRI models. The models' average mean squared errors against the ground truth are presented in Table 1 and image outputs displayed in Fig. 3. All results were performed on the same test set of images and Seidel coefficients not present in the training set.

Table 1: Quantitative results of deblur models on test dataset.

Model	Number of Parameters	MSE	PSNR (avg)	Runtime* (s)
LSI	262K	0.01022	69.307	—
LRI	134M	0.00322	75.082	63.34
U-Net	92.7M	0.00800	70.543	0.302
Seidel U-Net	54.4M	0.00626	71.305	0.234
HyperU-Net	6.9M	0.00583	71.716	0.180

4.1 Quantitative Results

As shown in Table 1 the SeidelNet models, Seidel U-Net and Seidel HyperU-Net, both outperform the baseline U-Net despite having fewer parameters. Notably, HyperU-Net has the best performance of all deep learning models while being roughly 13 times smaller than the baseline U-Net and about 8 times smaller than Seidel U-Net. This parameter efficiency also means that Seidel HyperU-Net is much faster to evaluate than the other models. Though not shown here, the HyperU-Net also has the smallest loss on the training set, which indicates that its superior performance is not due to the larger models overfitting (by virtue of their size). Indeed, it appears that access to Seidel coefficients in tandem with the unique structure of the Seidel HyperU-Net allows deblurring to be learned from data efficiently and effectively.

If we additionally consider the analytical methods, we see that LRI deblurring has the best performance of all methods; however, this is in larger part because the LRI model was used to blur the images in the dataset. Moreover, since it does not require learning, LRI deblurring is not biased to any particular training set. However, it remains prohibitively slow, suggesting that the SeidelNet models are the practical frontrunners. Finally, as expected, we see that LSI deconvolution performs the worst—unlike the other models, it only has access to information about the center PSF, and does not accurately model the spatially varying blur.

4.2 Qualitative Results

In Fig. 3 we see that the various deblur methods perform in accordance with the quantitative results. The largest difference between the methods occurs near the edges and corners of the images, where the blur becomes the worst. Amongst the learning strategies, the two SeidelNet models provide the sharpest deblur in the corners and edges, resolving small features with high contrast (see the top right section of images in the second row for example). Meanwhile, the baseline U-Net tends to create blurrier, lower contrast reconstructions near the edges and corners. For the analytical models, LRI also performs well in the corners but suffers from some artifacts due to image cropping; LSI on the other hand only deblurs the center but fails to deblur the edges due to its inadequate modeling.

*Amortized inference time on GTX 1080ti GPU. We omit LSI runtime since we did not implement this for GPU.

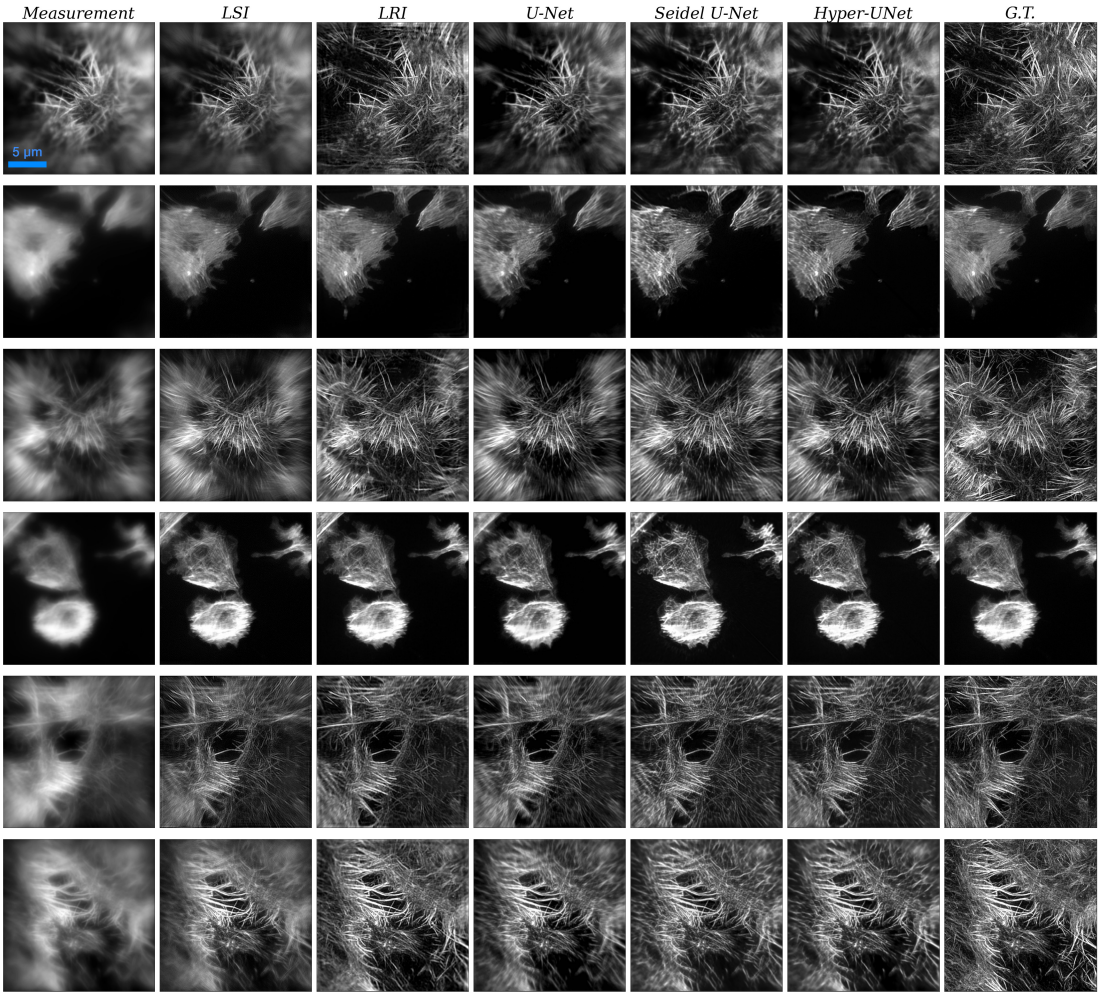

Figure 3: **Comparison of model performance on the test set** Models are stratified by column. Quantitative performance metrics for the same test set are available in Table 1.

4.3 Discussion

SeidelNet offers a path toward a new regime of general purpose image deblurring which is as easy to calibrate as standard deconvolution and shares the ever-impressive performance of deep learning models. It is our hope that SeidelNet and the research which builds upon it will form a new standard toolkit for all users of imaging systems, from microscopes to space telescopes. We intend to further test and improve SeidelNet, particularly with real images taken with a wide variety of imaging systems.

ACKNOWLEDGMENTS

This research was funded in part by AFSOR FA9550-21-S-0001 and the National Science Foundation SUPERB program.

REFERENCES

[1] Wiener, N. et al., [*Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications*], vol. 8, MIT press Cambridge, MA: (1964).

[2] Richardson, W. H., "Bayesian-based iterative method of image restoration," *J. Opt. Soc. Am.* **62**(1), 55–59 (1972).
[3] Lucy, L., "An iterative technique for the rectification of observed distributions," *Astronomical Journal* **79**, 745 (1974).
[4] Schmidt, U., Rother, C., Nowozin, S., Jancsary, J., and Roth, S., "Discriminative non-blind deblurring," in [*Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*], 604–611 (2013).
[5] Xu, L., Tao, X., and Jia, J., "Inverse kernels for fast spatial deconvolution," in [*European Conference on Computer Vision*], 33–48, Springer (2014).
[6] Bahat, Y., Efrat, N., and Irani, M., "Non-uniform blind deblurring by reblurring," in [*Proceedings of the IEEE international conference on computer vision*], 3286–3294 (2017).
[7] Trussell, H. and Hunt, B., "Image restoration of space variant blurs by sectioned methods," in [*ICASSP '78. IEEE International Conference on Acoustics, Speech, and Signal Processing*], **3**, 196–198 (1978).
[8] Nagy, J. G. and O'Leary, D. P., "Restoring images degraded by spatially variant blur," *SIAM Journal on Scientific Computing* **19**(4), 1063–1082 (1998).
[9] Denis, L., Thiébaut, E., and Soulez, F., "Fast model of space-variant blurring and its application to deconvolution in astronomy," in [*2011 18th IEEE International Conference on Image Processing*], 2817–2820 (2011).
[10] Weigert, M., Schmidt, U., Boothe, T., Müller, A., Dibrov, A., Jain, A., Wilhelm, B., Schmidt, D., Broaddus, C., Culley, S., et al., "Content-aware image restoration: pushing the limits of fluorescence microscopy," *Nature methods* **15**(12), 1090–1097 (2018).
[11] Fang, Y., Zhang, H., Wong, H. S., and Zeng, T., "A robust non-blind deblurring method using deep denoiser prior," in [*Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*], 735–744 (June 2022).
[12] Zhang, K., Ren, W., Luo, W., Lai, W.-S., Stenger, B., Yang, M.-H., and Li, H., "Deep image deblurring: A survey," *International Journal of Computer Vision* **130**(9), 2103–2130 (2022).
[13] Ren, W., Zhang, J., Ma, L., Pan, J., Cao, X., Zuo, W., Liu, W., and Yang, M.-H., "Deep non-blind deconvolution via generalized low-rank approximation," *Advances in neural information processing systems* **31** (2018).
[14] Deng, Q., Wen, Z., Dong, Z., Tang, J., Chen, W., Liu, X., and Yang, Q., "Spatially variant deblur and image enhancement in a single multimode fiber imaged by deep learning," *Opt. Lett.* **47**, 5040–5043 (Oct 2022).
[15] Dong, J., Roth, S., and Schiele, B., "Deep wiener deconvolution: Wiener meets deep learning for image deblurring," *Advances in Neural Information Processing Systems* **33**, 1048–1059 (2020).
[16] Yanny, K., Monakhova, K., Shuai, R. W., and Waller, L., "Deep learning for fast spatially varying deconvolution," *Optica* **9**, 96–99 (Jan 2022).
[17] Wijesinghe, P., Corsetti, S., Chow, D. J., Sakata, S., Dunning, K. R., and Dholakia, K., "Experimentally unsupervised deconvolution for light-sheet microscopy with propagation-invariant beams," *Light: Science & Applications* **11**(1), 1–15 (2022).
[18] Born, M. and Wolf, E., [*Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*], Elsevier (2013).
[19] Voelz, D. G., [*Computational fourier optics: a MATLAB tutorial*], SPIE press Bellingham, Washington (2011).
[20] Jacquemet, G., "ZeroCostDL4Mic - CARE (2D) example training and test dataset," (Mar. 2020).
[21] Kohli, A., Angelopoulos, A., You, S., Yanny, K., and Waller, L., "Linear revolution-invariance: Modeling and deblurring spatially-varying imaging systems," (2022).
[22] Goodman, J. W., "Introduction to fourier optics. 3rd," *Roberts and Company Publishers* (2005).
[23] Ronneberger, O., Fischer, P., and Brox, T., "U-net: Convolutional networks for biomedical image segmentation," (2015).
[24] Ha, D., Dai, A., and Le, Q. V., "Hypernetworks," *arXiv preprint arXiv:1609.09106* (2016).
[25] SecretMG, "Unet-for-image-denoising (github)," (2021).
[26] Mittal, G., "Hypernetworks(github)," (2018).