Adaptive Scanning in Ptychography through Deep Reinforcement Learning

Marcel Schloz[1], Johannes Müller[1], Thomas Pekin[1], Wouter Van den Broek[2] and Christoph Koch[3]

[1]Humboldt Universität zu Berlin, United States, [2]Humboldt Universität zu Berlin, Berlin, Berlin, United States, [3]Department of Physics, Humboldt University of Berlin, Berlin, Germany, United States

The development of new hardware and reconstruction algorithms has led to ptychography becoming a mature electron microscopy technique. The current research to further improve this technique is driven by the desire to investigate thick samples as well as to measure with a low electron dose. For the latter, researches have tried to relax the strict scanning constraint, which is required for obtaining information redundancy through consecutively overlapping probes. One approach is the optimization of the scanning pattern. A deviation from a raster grid towards a generally more beneficial pattern is supposed to improve the reconstruction quality especially at a reduced dose [1]. However, since a generally optimized scanning pattern is completely unrelated to the structure of the investigated specimen, there is obviously a limit of the achievable reconstruction quality improvement.

Here, we present an adaptive scanning algorithm that optimizes the scanning pattern based on structural information it receives from diffraction data that is acquired by the microscope. Our algorithm is based on a combination of deep learning [2] and reinforcement learning (RL) [3], or deep RL, which means that we model the described prediction process with deep neural networks and train it on previously acquired ptychographic data with RL. A recurrent neural network (RNN) is used to predict sub-sequences of scanning positions after starting from an initial sub-sequence. Input information is given by a combination X_p of processed 2D coordinates of a currently scanned sub-sequence $\mathbf{R}_p$ and processed content information z_p that is a compressed representation of the partial ptychographic reconstruction V_p. In addition, all the information from the previous prediction steps, represented by the hidden state H_p, is combined with the input information X_p and mapped by a gated recurrent unit (GRU) cell [4] to the next hidden state H_{p+1}. This hidden state is then used as a basis for the next GRU cell, but also to generate the output $\mathbf{R}_{p+1}$, which in our case corresponds to the next sub-sequence of scanning positions. Fully connected layers (FC) are used for the prediction of $\mathbf{R}_{p+1}$, as well as the processing of the location information $\mathbf{R}_p$ and content information z_p before being combined to the input information X_p. The content information z_p is generated by a pre-trained encoder part of a convolutional autoencoder. Figure 1 illustrates the RNN model in full detail.

Training of the network weights is performed by RL, as it allows to tune the weights according to a criteria that is represented by a reward r. This reward should corresponds to the quality of the ptychographic reconstruction which is desired to be high. It has been found empirically that a high reconstruction quality correlates positively with a high dynamic range in the phase. In order to evaluate by how much each scanning position contributes to the total reward, our reward is represented by a vector $\mathbf{r}$ [5]. By using the concept of a Voronoi diagram [6], where each position corresponds to a seed for one Voronoi cell (Figure 2a), we are able to assign only a part of the total phase to each position. The multi-objective function is then given by:

$$\mathbf{J}(\theta) = E_{\pi_\theta(\tau)}[\mathbf{r}(\tau)],$$

where θ corresponds to the network weights, E is the expectation operator, π is the policy and τ is the trajectory.

For the experimental investigation, we acquired multiple ptychographic datasets of a monolayer molybdenum disulfide (MoS_2) specimen with a NION Ultra-HERMES microscope at 60kV acceleration voltage and a Dectris ELA direct electron detector mounted at the EELS camera port. Distortions induced by the EEL spectrometer were corrected with an in-house developed software. From these datasets we create 175 smaller datasets, each with a size corresponding to 10,000 diffraction patterns of dimension 64x64 pixels. The algorithm was then trained on the smaller datasets with the restriction to predict optimal scanning sequences of 250 out of the 10,000 possible positions, corresponding to a dose reduction by a factor of 40. Each sub-sequence contains 50 positions, where the first sub-sequence follows a quasi-random Halton sequence [7]. The partial ptychographic reconstruction to obtain Vp was performed with an optimized version of ROP [8]. The reconstruction obtained by the predicted scanning sequence resolves the structure of the material (Figure 2b), which is not the case for a purely random sequence with the same number of positions (Figure 2c). The resolution seems to be also improved with respect to the full data reconstruction (Figure 2d) at the cost of a non-homogeneous quality improvement.
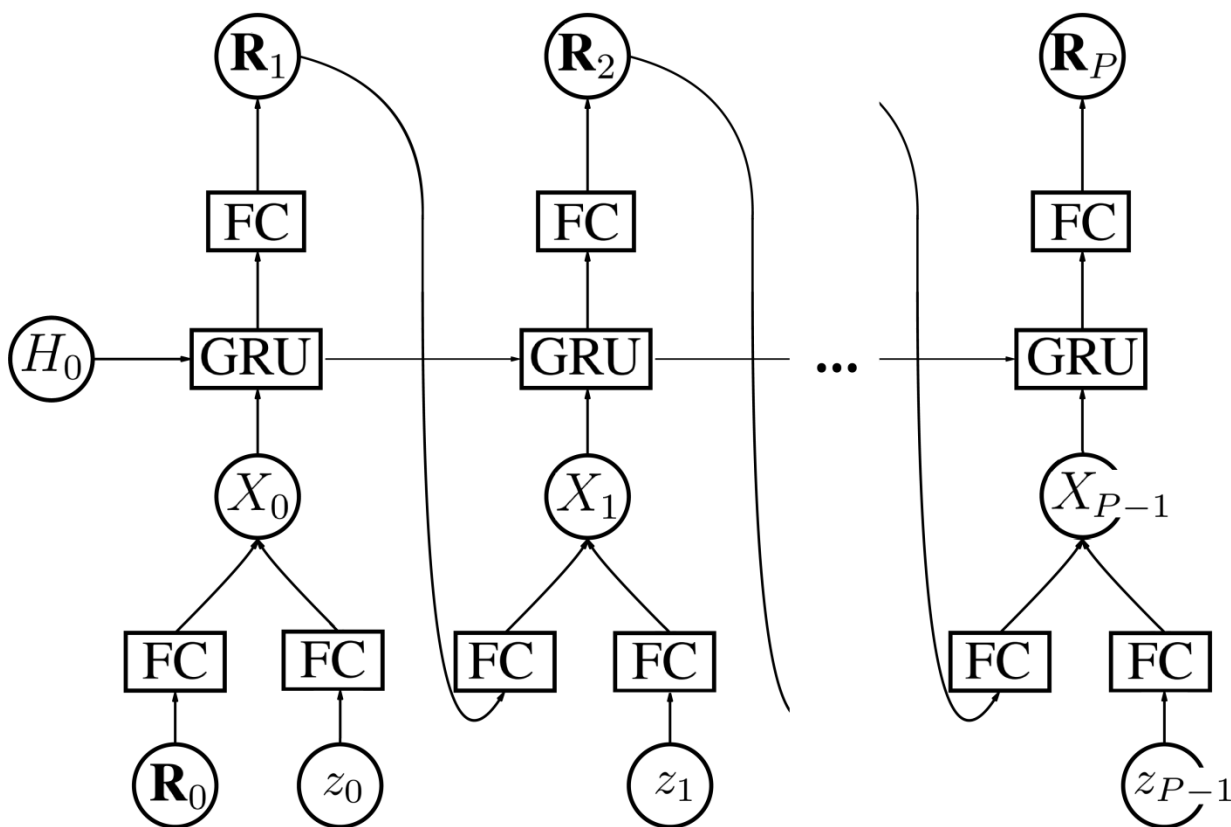
Figure 1. Schematic of the forward propagation process of the network model. The recurrent neural network consists of GRU units that use the hidden state Hp from the previous time step and the hybrid input information Xp to create a new hidden state Hp+1. The hybrid input is the concatenation of the pre-

processed information from the coordinates of the scanned sub-sequence Rp and the corresponding compressed representation of the partial reconstruction zp. The output of the GRU cell is used to predict the coordinates of the next sub-sequence Rp+1 and is also used as the input for the next GRU cell. The process is repeated until the full length of the scanning sequence, consisting of P sub-sequences, is reached.
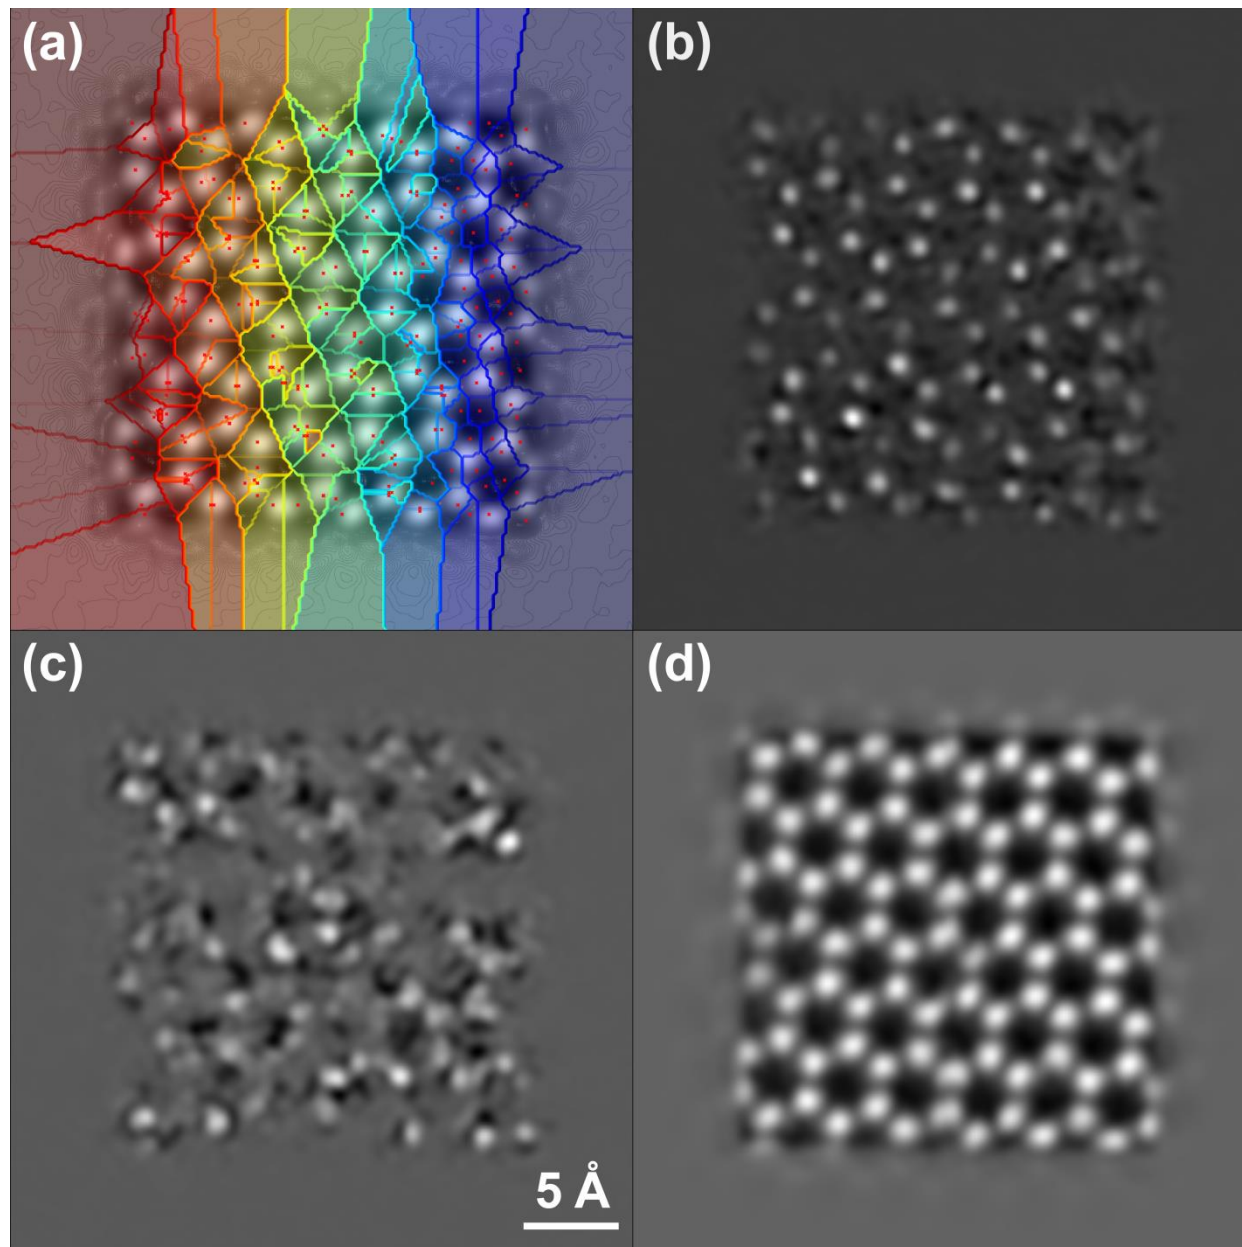

Figure 2. a) Voronoi diagram that is used to assign a unique reward to each scanning position of the predicted sequence. The scanning positions are shown as red dots, where the first 50 positions are distributed on the right side within the dark blue area. For visualization purpose, the ground truth reconstruction is included in the diagram. b) Ptychographic reconstruction from the predicted scanning sequence with a total of 250 positions. c) Reconstruction of a random sequence with a total of 250 positions. d) Reconstruction of the full data, consisting of 10,000 positions.

References
[1] Huang, Xiaojing, et al. "Optimization of overlap uniformness for ptychography." Optics Express 22.10 (2014): 12634-12644.
[2] LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. "Deep learning." nature 521.7553 (2015): 436-444.
[3] Sutton, Richard S., and Andrew G. Barto. Reinforcement learning: An introduction. MIT press, 2018.
[4] Chung, Junyoung, et al. "Empirical evaluation of gated recurrent neural networks on sequence modeling." arXiv preprint arXiv:1412.3555 (2014).
[5] Roijers, Diederik M., et al. "A survey of multi-objective sequential decision-making." Journal of Artificial Intelligence Research 48 (2013): 67-113.
[6] Okabe, Atsuyuki. "Spatial tessellations." International Encyclopedia of Geography: People, the Earth, Environment and Technology: People, the Earth, Environment and Technology (2016): 1-11.
[7] Halton, John H. "Algorithm 247: Radical-inverse quasi-random point sequence." Communications of the ACM 7.12 (1964): 701-702.
[8] Schloz, Marcel, et al. "Overcoming information reduced data and experimentally uncertain parameters in ptychography with regularized optimization." Optics Express 28.19 (2020): 28306-28323.