

Manuscript version: Working paper (or pre-print)

The version presented here is a Working Paper (or 'pre-print') that may be later published elsewhere.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/145930>

How to cite:

Please refer to the repository item page, detailed above, for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Adaptive Partial Scanning Transmission Electron Microscopy with Reinforcement Learning

Jeffrey M. Ede^{1,a}

¹University of Warwick, Department of Physics, Coventry, CV4 7AL, UK
aj.m.ede@warwick.ac.uk

ABSTRACT

Compressed sensing can decrease scanning transmission electron microscopy electron dose and scan time with minimal information loss. Traditionally, sparse scans used in compressed sensing sample a static set of probing locations. In contrast, we present a prototype for a contiguous sparse scan system that piecewise adapts scan paths to specimens as they are scanned. Sampling directions for scan segments are chosen by a recurrent neural network based on previously observed scan segments. The recurrent actor is trained by reinforcement learning to cooperate with a feedforward convolutional neural network that completes sparse scans. This paper presents our learning policy, experiments, and example partial scans, and discusses future research directions. Source code, pretrained models, and training data is openly accessible at <https://github.com/Jeffrey-Ede/adaptive-scans>.

Keywords: compressed sensing, deep learning, electron microscopy, reinforcement learning.

1 Introduction

Most scan systems sample signals at sequences of discrete probing locations. Examples include atomic force microscopy^{1,2}, computerized axial tomography^{3,4}, electron backscatter diffraction⁵, scanning electron microscopy⁶, scanning Raman spectroscopy⁷, scanning transmission electron microscopy⁸ (STEM) and X-ray diffraction spectroscopy⁹. In STEM, the high current density of electron probes produces radiation damage in many materials, limiting the range and types of investigations that can be performed^{10,11}. In addition, most STEM signals are oversampled¹² to ease visual inspection and decrease sub-Nyquist artefacts¹³. As a result, a variety of compressed sensing¹⁴ algorithms have been developed to enable decreased STEM probing¹⁵. In this paper, we introduce a new approach to STEM compressed sensing where a scan system learns to adapt partial scans¹⁶ to specimens by deep reinforcement learning¹⁷ (RL).

Established compressed sensing strategies include random sampling^{18–20}, uniformly spaced sampling^{19,21–23}, sampling based on a model of a sample^{24,25}, partials scans with fixed paths¹⁶, dynamic sampling to minimize entropy^{26–29} and dynamic sampling based on supervised learning³⁰. Complete signals can be extrapolated from partial scans by an infilling algorithm, estimating their fast Fourier transforms³¹ or inferred by an artificial neural network^{16,23} (ANN). In general, the best sampling strategy varies for different specimens. For example, uniformly spaced sampling is often better than spiral paths for oversampled STEM images¹⁶. However, sampling strategies designed by humans usually have limited ability to leverage an understanding of physics to optimize sampling. As proposed by our earlier work¹⁶, we have therefore developed ANNs to dynamically adapt scan paths to specimens. Expected performance of dynamic scans can always match or surpass expected performance of static scans as static scan paths are a special case of dynamic scan paths.

Exploration of STEM specimens is a finite-horizon partially observed Markov decision process^{32,33} (POMDP) with sparse losses. A partial scan can be constructed from path segments sampled at each step of the POMDP and a loss can be based on the quality of a completion generated from the partial scan. Most scan systems support custom scan paths or can be augmented with a field programmable gate array^{34,35} (FPGA) to support custom scan paths. However, there is a delay before a scan system can execute or is ready to receive a new command. Latency can be reduced by using both fewer and larger steps, and decreasing steps may also reduce distortions due to cumulative errors in probing positions³⁴ after commands are issued. Command execution can also be delayed by ANN inference. However, inference delay can be minimized by using a computationally lightweight ANN or inferring future commands while previous commands are executing.

Markov decision processes (MDPs) can be optimized by recurrent neural networks (RNNs) based on long short-term memory^{36,37} (LSTM), gated recurrent unit³⁸ (GRU), or other cells^{39–41}. LSTMs and GRUs are popular as they solve the vanishing gradient problem⁴² and have consistently high performance⁴⁰. Small RNNs are computationally inexpensive and are often applied to MDPs as they can learn to extract and remember state information to inform future decisions. To solve

dynamic graphs, an RNN can be augmented with dynamic external memory to create a differentiable neural computer⁴³ (DNC). To optimize a MDP, a discounted future loss, Q_t , at step t in a MDP with T steps can be calculated from step losses, L_t , with Bellman’s equation,

$$Q_t = \sum_{t'=t}^T \gamma^{t'-t} L_{t'}, \quad (1)$$

where $\gamma \in [0, 1)$ discounts future step losses. Equations for RL are often presented in terms of rewards, e.g. $r_t = -L_t$; however, losses are an equivalent representation that avoids complicating our equations with minus signs. Discounted future loss backpropagation through time⁴⁴ (BPTT) enables RNNs to be trained by gradient descent⁴⁵. However, losses for partial scan completions are not differentiable with respect to (w.r.t.) RNN actions, (a_1, \dots, a_T) , controlling which path segments are sampled.

Many MDPs have losses that are not differentiable w.r.t. agent actions. Examples include agents directing their vision^{46,47}, managing resources⁴⁸, and playing score-based computer games^{49,50}. Actors are usually trained with non-differentiable losses by introducing a differentiable surrogate⁵¹ or critic⁵² to predict losses that can be backpropagated to actor parameters. However, non-differentiable losses can also be backpropagated to agent parameters if actions are sampled from a differentiable probability distribution^{46,53} as training losses given by products of losses and sampling probabilities are differentiable. There are also a variety of alternatives to gradient descent, such as simulated annealing⁵⁴ and evolutionary algorithms⁵⁵, that do not require differentiable loss functions. Such alternatives can outperform gradient descent⁵⁶; however, they usually achieve similar or lower performance than gradient descent for deep ANN training.

2 Training

In this section, we outline our training environment, ANN architecture and learning policy. Our ANNs were developed in Python with TensorFlow⁵⁷. Detailed architecture and learning policy is in supplementary information. In addition, source code and pretrained models are available via GitHub⁵⁸, and training data is openly accessible^{12,59}.

2.1 Environment

To create partial scans from STEM images, an actor, μ , infers L2 normalized action vectors, $\mu(h_t)$, based on a history, $h_t = (o_1^i, a_1, \dots, o_{t-1}, a_{t-1})$, of previous actions, a , and observations, o . To encourage exploration, $\mu(h_t)$ is rotated to a_t by Ornstein-Uhlenbeck⁶⁰ (OU) exploration noise⁶¹, ε_t ,

$$a_t = \begin{bmatrix} \cos \varepsilon_t & -\sin \varepsilon_t \\ \sin \varepsilon_t & \cos \varepsilon_t \end{bmatrix} \mu(h_t) \quad (2)$$

$$\varepsilon_t = \theta(\varepsilon_{\text{avg}} - \varepsilon_{t-1}) + \sigma W \quad (3)$$

where we chose $\theta = 0.1$ to decay noise to $\varepsilon_{\text{avg}} = 0$, a scale factor, $\sigma = 0.2$, to scale a standard normal variate, W , and start noise $\varepsilon_0 = 0$. OU noise is linearly decayed to zero throughout training. Correlated OU exploration noise is recommended for continuous control tasks optimized by deep deterministic policy gradients⁴⁹ (DDPG) and recurrent deterministic policy gradients⁵⁰ (RDPG). Nevertheless, follow-up experiments with TD3⁶² and D4PG⁶³ have found that uncorrelated Gaussian noise can produce similar results.

An action, a_t , is the direction to move to observe a path segment, o_t , relative to the position at the end of the previous segment. Partial scans are constructed from complete histories of actions and observations, h_T . A simplified partial scan is shown in fig. 1. In our experiments, partial scans, s , are constructed from $T = 20$ straight path segments selected from 96×96 STEM images. Each segment has 20 probing positions separated by $d = 2^{1/2}$ px and positions can be outside an image. The pixels in the image nearest each probing position are sampled, so a separation of $d \geq 2^{1/2}$ prevents probing positions in a segment from sampling the same pixel. A separation of $d < 2^{1/2}$ would allow a pixel to be sampled more than once by moving diagonally, potentially incentivising orthogonal scan motion to sample more pixels.

Following our earlier work^{16,23,64}, we select subsets of pixels from STEM images to create partial scans to train ANNs for compressed sensing. Selecting a subset of pixels is easier than preparing a large, carefully partitioned and representative dataset^{65,66} containing experimental partial scan and full image pairs, and selected pixels have realistic noise characteristics as they are from experimental images. However, selecting a subset of pixels does not account for probing location errors varying with scan shape³⁴. We use a Warwick Electron Microscopy Dataset (WEMD) containing 19769 32-bit 96×96 images cropped and downsampled from full images^{12,59}. Cropped images were blurred by a symmetric 5×5 Gaussian kernel with a 2.5 px standard deviation to decrease any training loss variation due to varying noise characteristics. Finally, images, I , were linearly transformed to normalized images, I_N , with minimum and maximum values of -1 and 1 . To test performance, the 19769 images were split, without shuffling, into a training set containing 15815 images and a test set containing 3954 images.




Figure 1. Example 8×8 partial scan with $T = 5$ straight path segments. Each segment in this example has 3 probing positions separated by $d = 2^{1/2}$ px and their starts are labelled by step numbers, t . Partial scans are selected from STEM images by sampling pixels nearest probing positions, even if the probing position is nominally outside an imaging region.

2.2 Architecture

For training, our adaptive scan system consists of an actor, μ , target actor, μ' , critic, Q , target critic, Q' , and generator, G . To minimize latency, our actors and critics are computationally inexpensive deep LSTMs⁶⁷ with a depth of 2 and 256 hidden units. Our generator is a convolutional neural network^{68,69} (CNN). A recurrent actor selects actions, a_t and observes path segments, o_t , that are added to an experience replay⁷⁰, R , containing 10^5 sequences of actions and observations, $h_T = (o_1, a_1, \dots, o_T, a_T)$. Partial scans, s , are constructed from histories sampled from the replay to train a generator to complete partial scans, $I_G^i = G(s^i)$. The actor and generator cooperate to minimize generator losses, L_G , and are the only networks needed for inference.

Generator losses are not differentiable w.r.t. actor actions used to construct partial scans i.e. $\partial L_G / \partial a_t = 0$. Following RDPG⁵⁰, we therefore introduce recurrent critics to predict losses from actor actions and observations that can be backpropagated to actors for training by BPTT. Actor and critic RNNs have the same architecture, except actors have two outputs to parameterize actions whereas critics have one output to predict losses. Target networks^{49,71} use exponential moving averages of live actor and critic network parameters and are introduced to stabilize learning. For training by RDPG, live and target ANNs separately replay experiences. However, we propagate live RNN states to target RNNs at each step as a precaution against any cumulative divergence of target network behaviour from live network behaviour across multiple steps.

2.3 Learning Policy

To train actors to cooperate with a generator to complete partial scans, we developed cooperative recurrent deterministic policy gradients (CRDPG, algorithm 1). This is an extension of RDPG to an actor that cooperates with another ANN to minimize its loss. We train our networks by ADAM⁷² optimized gradient descent for $M = 10^6$ iterations with a batch size, $N = 32$. We use constant learning rates $\eta_\mu = 0.0005$ and $\eta_Q = 0.0010$ for the actor and critic, respectively. For the generator, we use an initial learning rate $\eta_G = 0.0030$ with an exponential decay factor of $0.75^{5m/M}$ at iteration m . The exponential decay envelope is multiplied by a sawtooth cyclic learning rate⁷³ with a period of $2M/9$ that oscillates between 0.2 and 1.0. Training takes two days with an Intel i7-6700 CPU and an Nvidia GTX 1080 Ti GPU.

We augment training data by a factor of eight by applying a random combination of flips and 90° rotations, mapping $s \rightarrow s'$ and $I_N \rightarrow I'_N$, similar to our earlier work^{16,23,64,74}. Our generator is trained to minimize mean squared errors (MSEs),

$$L_G = \text{MSE}(G(s'), I_N), \quad (12)$$

between scan completions, $G(s')$, and normalized target images, I_N . Generator losses decrease during training as the generator learns, and may vary due to loss spikes⁶⁴, learning rate oscillations⁷³ or other training phenomena. Normalizing losses can improve RL⁷⁵, so we divide generator losses used for critic training by their running mean,

$$L_{\text{avg}} \leftarrow \beta_L L_{\text{avg}} + \frac{1 - \beta_L}{N} \sum_i L_G, \quad (13)$$

Algorithm 1. Cooperative recurrent deterministic policy gradients (CRDPG).

Initialize actor, μ , critic, Q , and generator, G , networks with parameters ω , θ and ϕ , respectively.

Initialize target networks, μ' and Q' , with parameters $\omega' \leftarrow \omega$, $\theta' \leftarrow \theta$, respectively.

Initialize replay buffer, R .

Initialize average generator loss, L_{avg} .

for iteration $m = 1, M$ **do**

Initialize empty history, h_0 .

for step $t = 1, T$ **do**

Make observation, o_t .

$h_t \leftarrow h_{t-1}, a_t, o_t$ (append action and corresponding observation to history).

Select action, a_t , by computing $\mu(h_t)$ and applying exploration noise, ϵ_t .

end for

Store the sequence $(o_1, a_1, \dots, o_T, a_T)$ in R .

Sample a minibatch of N histories, $h_T^i = (o_1^i, a_1^i, \dots, o_T^i, a_T^i)$, from R .

Construct partial scans, s^i , from h_T^i .

Use generator to complete partial scans, $L_G^i = G(s^i)$.

Compute step losses, (L_1^i, \dots, L_T^i) , from generator losses, L_G^i , and over edge losses, E_t^i ,

$$L_t^i = E_t^i + \delta_{tT} \frac{\text{clip}(L_G^i)}{L_{\text{avg}}}, \quad (4)$$

where the Kronecker delta, δ_{tT} , is 1 if $t = T$ and 0 otherwise, and $\text{clip}(L_G^i)$ is the smaller of L_G^i and three standard deviations above its running mean.

Compute target values, (y_1^i, \dots, y_T^i) , with target networks,

$$y_t^i = L_t^i + \gamma Q'(H_Q^i, o_{t+1}^i, a_{t+1}^i, \mu'(H_\mu^i, o_{t+1}^i, a_{t+1}^i)), \quad (5)$$

where H_Q^i and H_μ^i are states of live networks after computing $Q(h_t^i, a_t^i)$ and $\mu(h_t^i)$, respectively.

Compute critic update (using BPTT),

$$\Delta\omega = \frac{1}{NT} \sum_i^N \sum_t^T (y_t^i - Q(h_t^i, a_t^i)) \frac{\partial Q(h_t^i, a_t^i)}{\partial \omega}. \quad (6)$$

Compute actor update (using BPTT),

$$\Delta\theta = \frac{1}{NT} \sum_i^N \sum_t^T \frac{\partial Q(h_t^i, a_t^i)}{\partial \mu(h_t^i)} \frac{\partial \mu(h_t^i)}{\partial \theta}. \quad (7)$$

Compute generator update,

$$\Delta\phi = \frac{1}{N} \sum_i^N \frac{\partial L_G^i}{\partial \phi}. \quad (8)$$

Update the actor, critic and generator by gradient descent.

Update the target networks and average generator loss,

$$\omega' \leftarrow \beta_\omega \omega' + (1 - \beta_\omega) \omega, \quad (9)$$

$$\theta' \leftarrow \beta_\theta \theta' + (1 - \beta_\theta) \theta, \quad (10)$$

$$L_{\text{avg}} \leftarrow \beta_L L_{\text{avg}} + \frac{1 - \beta_L}{N} \sum_i^N (L_G^i). \quad (11)$$

end for

where we chose $\beta_L = 0.997$ and L_{avg} is updated at each training iteration.

Heuristically, an optimal policy does not go over image edges as there is no information there in our training environment. To accelerate convergence, we therefore added a small loss penalty, $E_t = 0.1$, at step t if an action results in a probing position being over an image edge. The total loss at each step is

$$L_t = E_t + \delta_{tT} \frac{\text{clip}(L_G)}{L_{\text{avg}}}, \quad (14)$$

where $\text{clip}(L_G)$ clips losses used for RL to three standard deviations above their running mean. This adaptive loss clipping is inspired by adaptive learning rate clipping⁶⁴ (ALRC) and reduces learning destabilization by high loss spikes. However, we expect that clipping normalized losses to a fixed threshold⁷¹ would achieve similar results. The Kronecker delta, δ_{tT} , in eqn. 14 is 1 if $t = T$ and 0 otherwise, so it only adds the generator loss at the final step, T .

To estimate discounted future losses, Q_t^{rl} , for RL, we use a target actor and critic,

$$Q_t^{\text{rl}} = L_t + \gamma Q'(h_{t+1}, \mu'(h_{t+1})), \quad (15)$$

where we chose $\gamma = 0.97$. Target networks stabilize learning and decrease policy oscillations^{76–78}. The critic is trained to minimize mean squared differences, L_Q , between predicted and target losses, and the actor is trained to minimize losses, L_μ , predicted by the critic,

$$L_Q = \frac{1}{2T} \sum_{t=1}^T (y_t - Q(h_t, a_t))^2, \quad (16)$$

$$L_\mu = \frac{1}{T} \sum_{t=1}^T Q(h_t, a_t). \quad (17)$$

Our target actor and critic have trainable parameters ω' and θ' , respectively, that track live parameters, ω and θ , by soft updates⁴⁹,

$$\omega'_m = \beta_\omega \omega'_{m-1} + (1 - \beta_\omega) \omega_m, \quad (18)$$

$$\theta'_m = \beta_\theta \theta'_{m-1} + (1 - \beta_\theta) \theta_m, \quad (19)$$

where we chose $\beta_\omega = \beta_\theta = 0.9997$. We also investigated hard updates⁷¹, where target networks are periodically copied from live networks; however, we found that soft updates result in faster convergence and more stable training.

3 Experiments

In this section, we present examples of adaptive partial scans and select learning curves for architecture and learning policy experiments. Examples of 1/23.04 px coverage partial scans, target outputs and generator completions are shown in fig. 2 for 96×96 crops from test set STEM images. They show both adaptive and spiral scans after flips and rotations to augment data for the generator. The first actions select a path segment from the middle of image in the direction of a corner. Actors then use the first and following observations to inform where to sample the remaining $T - 1 = 19$ path segments. Actors adapt scan paths to specimens. For example, if an image contains regular atoms, an actor might cover a large area to see if there is a region where that changes. Alternatively, if an image contains a uniform region, actors, may explore near image edges and far away from the uniform region to find region boundaries.

The main limitation of our experiments is that generators trained to complete a variety of partial scan paths generated by an actor achieves lower performance than a generate trained to complete partial scans with a fixed path. For example, fig. 3a shows that generators trained to cooperate with LSTM or GRU actors are outperformed by generators trained with fixed spiral or other scan paths shown in fig. 3b. Spiral paths outperform fixed scan paths; however, we emphasize that paths generated by actors are designed for individual training data, rather than all training data. Freezing actor training to prevents changes in actor policy does not result in clear improvements in generator performance. Consequently, we think that improvements to generator architecture or learning policy should be a starting point for further investigation. To find the best practical actor policy, we think that a generator trained for a variety of scan paths should achieve comparable performance to generators trained for single scan paths.

We investigated a variety of popular RNN architectures to minimize inference time. Learning curves in fig. 3a show that performance is similar for LSTMs and GRUs. GRUs require less computation. However, LSTM and GRU inference time is comparable and GRU training seems to be more prone to loss spikes, so LSTMs may be preferable. We also created a DNC

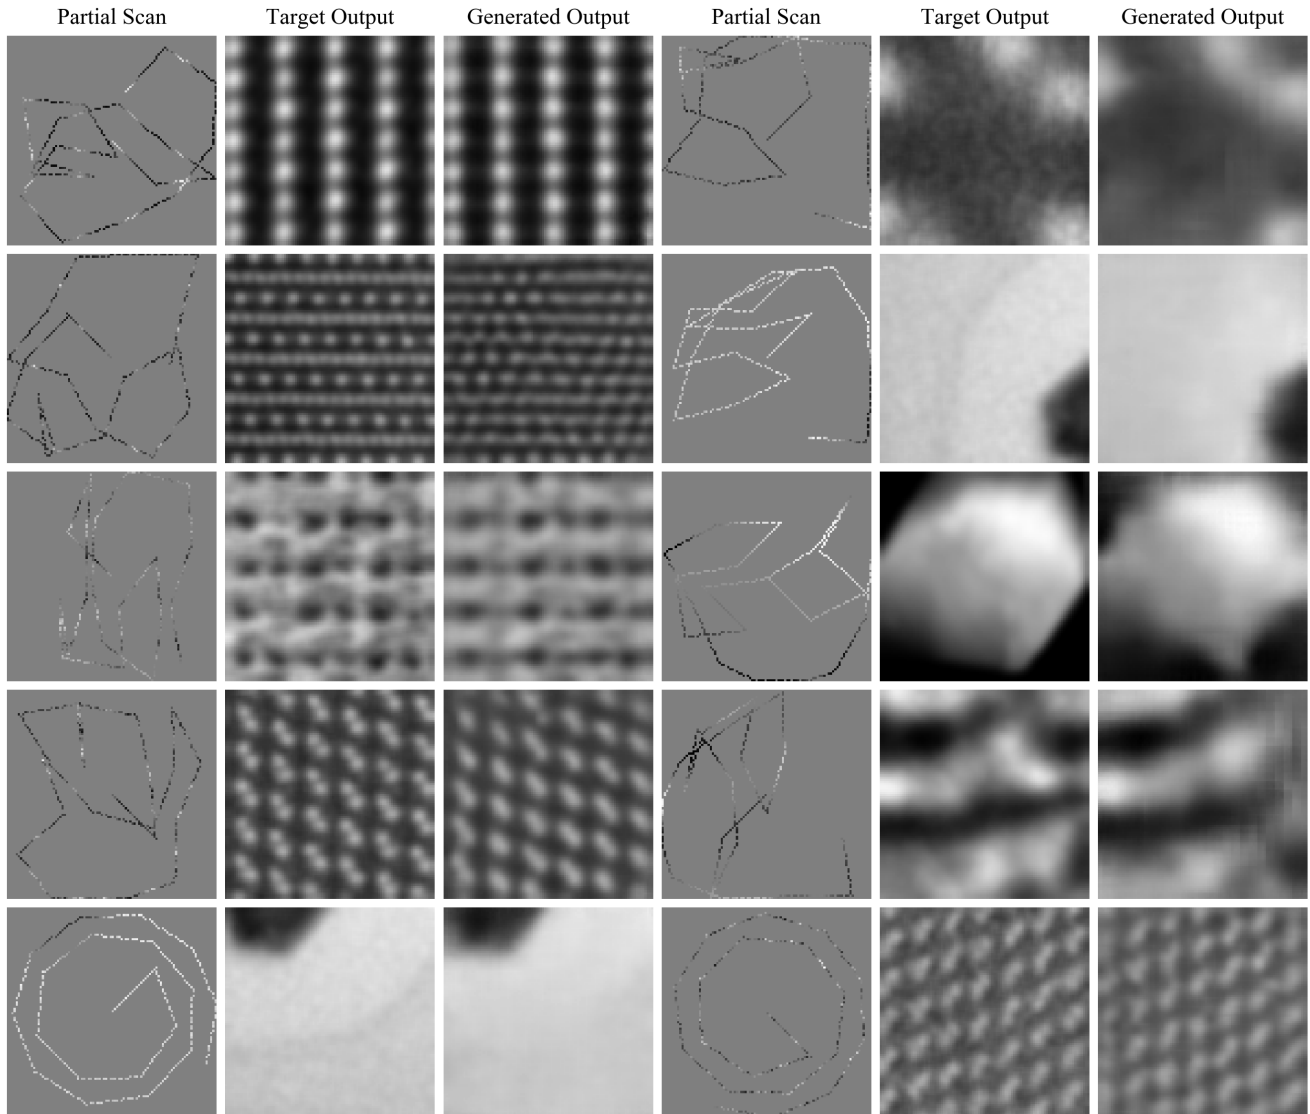


Figure 2. Test set 1/23.04 px coverage partial scans, target outputs and generated partial scan completions for 96×96 crops from STEM images. The top four rows show adaptive scans, and the bottom row shows spiral scans. Input partial scans are noisy, whereas target outputs are blurred.

by augmenting a deep LSTM with dynamic external memory. However, fig. 3c shows that LSTM and DNC performance is similar, and inference time and computational requirements are much higher for our DNC. We tried to reduce computation and accelerate convergence by applying projection layers to LSTM hidden states⁷⁹. However, we found that performance decreased with decreasing projection layer size.

Experienced replay buffers for RL often have heuristic sizes, such as 10^6 examples. However, RL can be sensitive to replay buffer size⁷⁰. Indeed, learning curves in fig. 3d show that increasing buffer size improves learning stability and decreases test set errors. Increasing buffer size usually improves learning stability and decreases forgetting by exposing actors and critics to a higher variety of past policies. However, we expect that convergence would be slowed if the buffer became too large as increasing buffer size increases expected time before experiences with new policies are replayed. We also found that increasing buffer sized decreased the size of small loss oscillations⁷⁶⁻⁷⁸, which have a period near 2000 iterations. However, the size of loss oscillations does not appear to affect performance.

We found that initial convergence is usually delayed if a large portion of initial actions go outside the imaging region. This would often delay convergence by about 10^4 iterations before OU noise led to the discovery of better exploration strategies away from image edges. Although 10^4 iterations is only 1% of our 10^6 iteration learning policy, it often impaired development

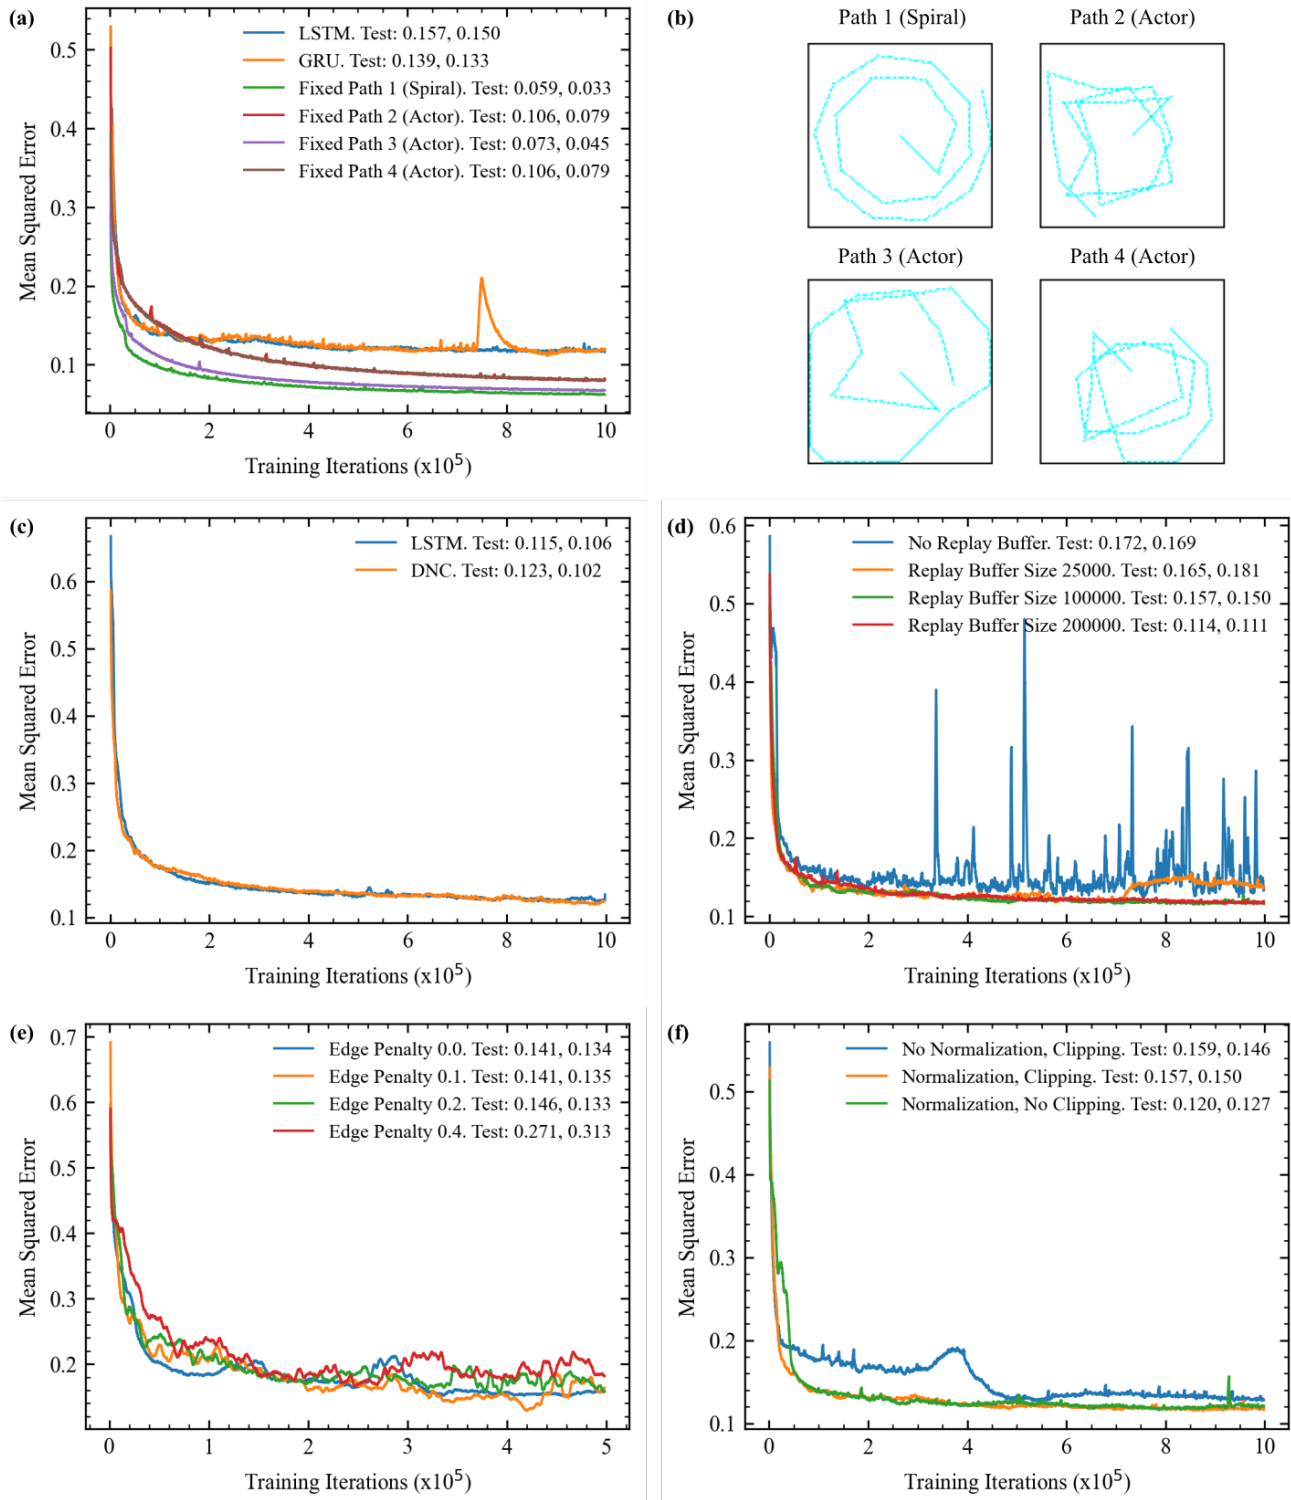


Figure 3. Learning curves for a-b) adaptive scan paths chosen by an LSTM or GRU, and fixed spiral and other fixed paths, c) adaptive paths chosen by an LSTM or DNC, d) a range of replay buffer sizes, e) a range of penalties for trying to sample at probing positions over image edges, and f) with and without normalizing or clipping generator losses used for critic training. All learning curves are 2500 iteration boxcar averaged and results in different plots are not directly comparable due to varying experiment settings. Means and standard deviations of test set errors, “Test: Mean, Std Dev”, are at the ends of labels in graph legends.

by delaying debugging or evaluation of changes to architecture and learning policy. Augmenting RL losses with subgoal-based heuristic rewards can accelerate convergence by making problems more tractable⁸⁰. Thus, we added loss penalties if actors tried to go over image edges, which accelerated initial convergence. Learning curves in fig. 3e show that over edge penalties at each step smaller than $E_t = 0.2$ have a similar effect on performance. Further, performance is lower for higher over edge penalties, $E_t \geq 0.2$. We also found that training is more stable if over edge penalties are added at individual steps, rather than propagated to past steps as part of a discounted future loss.

Our actor, critic and generator are trained together. It follows that generator losses, which our critic learns to predict, decrease throughout training as generator performance improves. However, normalizing loss sizes usually improves RL⁷⁵, so we divide by their running means in eqn. 14. Learning curves in fig. 3f show that loss normalization improves learning stability and decreases final errors. Clipping training losses can improve RL⁷¹, so we clipped generator losses used for critic training to 3 standard deviations above their running means. We found that clipping increases test set errors, possibly because most training errors are in a similar regime. Thus, we expect that clipping may be more helpful for training with sparser scans as higher uncertainty may increase likelihood of unusually high generator losses.

4 Discussion

The main limitation of our adaptive scan system is that generator errors are much higher when a generator is trained for a variety of scan paths than when it is trained for a single scan path. However, we expect that generator performance for a variety of scans could be improved to match performance for single scans by developing a larger neural network and a better learning policy. To train actors to cooperate with generators, we developed CRDPG. This is an extension of RDPG⁵⁰, and RDPG is based on DDPG⁴⁹. Alternatives to DDPG, such as TD3⁶² and D4PG⁶³, arguably achieve higher performance, so we expect that they could form the basis of a future training algorithm. Further, we expect that architecture and learning policy could be improved by AdaNet⁸¹, Ludwig⁸², or other automatic machine learning^{83–87} (AutoML) algorithms as AutoML can match or surpass the performance of human developers^{88,89}. Finally, test set losses for a variety of scans appear to be decreasing at the end of training, so we expect that performance could be improved by increasing training iterations.

After generator performance is improved, we expect the main limitation of our adaptive scan system to be distortions caused by probing position errors. Errors usually depend on scan path shape³⁴ and accumulate for each path segment. Non-linear scan distortions can be corrected by comparing pairs of orthogonal raster scans^{90,91}, and we expect this method can be extended to partial scans. However, orthogonal scanning would complicate measurement by limiting scan paths to two half scans to avoid doubling electron dose on beam-sensitive materials. Instead, we propose that a cyclic generator⁹² could be trained to correct scan distortions and provide a detailed method as supplementary information. Another limitation is that our generators do not learn to correct STEM noise⁹³. However, we expect that generators can learn to remove noise, for example, from single noisy examples⁹⁴ or by supervised learning⁷⁴.

To simplify our preliminary investigation, our scan system samples straight path segments and cannot go outside a specified imaging region. However, actors could learn to output actions with additional degrees of freedom to describe curves, multiple successive path segments, or sequences of non-contiguous probing positions. Similarly, additional restrictions could be applied to actions. For example, actions could be restricted to avoid actions that cause high probing position errors. Training environments could also be modified to allow actors to sample pixels over image edges by loading images larger than partial scan regions. In practice, actors can sample outside a scan region and being able to access extra information outside an imaging region could improve performance. However, using larger images may slow development by increasing data loading and processing times.

Not all scan systems support non-raster scan paths. However, many scan controllers can be augmented with an FPGA to enable custom scan paths^{34,35}. Recent versions of Gatan Digital Micrograph support Python⁹⁵, so our ANNs can be readily integrated into existing scan systems. Alternatively, an actor could be synthesized on the scan controlling FPGA^{96,97} to minimize inference time. There could be hundreds of path segments in a partial scan, so computationally lightweight and parallelizable actors are essential to minimize scan time. We have therefore developed actors based computationally inexpensive RNNs, which can remember state information to inform future decisions. Another approach is to update a partial scan at each step to be input to feedforward neural network (FNN), such as a CNN, to decide actions. However, we expect that FNNs are less practical than RNNs as FNNs may require additional computation to reprocess all past states at each step.

5 Conclusions

Our initial investigation demonstrates that actor RNNs can be trained by RL to direct dynamic adaption of contiguous scans to specimens for compressed sensing. We introduce CRDPG to train an RNN to cooperate with a CNN to complete STEM images from partial scans, and present learning policy, experiments, and example applications. After further development, we expect that adaptive scans will become an effective means to decrease electron beam damage and scan times by 10-100× with minimal

information loss. Static sampling strategies are a subset of possible dynamic sampling strategies, so the performance of static sampling can always be matched by or outperformed by dynamic sampling. Further, we expect that adaptive scan systems can be developed for most areas of science and technology, including for the reduction of medical radiation. To encourage further investigation, our source code, pretrained models, and training data is openly accessible.

6 Supplementary Information

Supplementary information is openly accessible at <https://doi.org/10.5281/zenodo.4304462>. Therein, we present detailed ANN architecture, additional experiments, additional examples, and a method to correct partial scan distortions.

Data Availability

The data that support the findings of this study are openly available.

Acknowledgements

Thanks go to Jasmine Clayton, Abdul Mohammed, and Jeremy Sloan for internal review. The author acknowledges funding from EPSRC grant EP/N035437/1 and EPSRC Studentship 1917382.

Competing Interests

The author declares no competing interests.

References

1. Krull, A., Hirsch, P., Rother, C., Schiffrin, A. & Krull, C. Artificial-Intelligence-Driven Scanning Probe Microscopy. *Commun. Phys.* **3**, 1–8 (2020).
2. Rugar, D. & Hansma, P. Atomic Force Microscopy. *Phys. Today* **43**, 23–30 (1990).
3. New, P. F., Scott, W. R., Schnur, J. A., Davis, K. R. & Taveras, J. M. Computerized Axial Tomography with the EMI Scanner. *Radiology* **110**, 109–123 (1974).
4. Heymsfield, S. B. *et al.* Accurate Measurement of Liver, Kidney, and Spleen Volume and Mass by Computerized Axial Tomography. *Annals Intern. Medicine* **90**, 185–187 (1979).
5. Schwartz, A. J., Kumar, M., Adams, B. L. & Field, D. P. *Electron Backscatter Diffraction in Materials Science*, vol. 2 (Springer, 2009).
6. Vernon-Parry, K. D. Scanning Electron Microscopy: An Introduction. *III-Vs Rev.* **13**, 40–44 (2000).
7. Keren, S. *et al.* Noninvasive Molecular Imaging of Small Living Subjects Using Raman Spectroscopy. *Proc. Natl. Acad. Sci.* **105**, 5844–5849 (2008).
8. Tong, Y.-X., Zhang, Q.-H. & Gu, L. Scanning Transmission Electron Microscopy: A Review of High Angle Annular Dark Field and Annular Bright Field Imaging and Applications in Lithium-Ion Batteries. *Chin. Phys. B* **27**, 066107 (2018).
9. Scarborough, N. M. *et al.* Dynamic X-Ray Diffraction Sampling for Protein Crystal Positioning. *J. Synchrotron Radiat.* **24**, 188–195 (2017).
10. Hujsak, K., Myers, B. D., Roth, E., Li, Y. & Dravid, V. P. Suppressing Electron Exposure Artifacts: An Electron Scanning Paradigm with Bayesian Machine Learning. *Microsc. Microanal.* **22**, 778–788 (2016).
11. Egerton, R. F., Li, P. & Malac, M. Radiation Damage in the TEM and SEM. *Micron* **35**, 399–409 (2004).
12. Ede, J. M. Warwick Electron Microscopy Datasets. *Mach. Learn. Sci. Technol.* **1**, 045003 (2020).
13. Amidror, I. Sub-Nyquist Artefacts and Sampling Moiré Effects. *Royal Soc. Open Sci.* **2**, 140550 (2015).
14. Binev, P. *et al.* Compressed Sensing and Electron Microscopy. In *Modeling Nanoscale Imaging in Electron Microscopy*, 73–126 (Springer, 2012).
15. Ede, J. M. Review: Deep Learning in Electron Microscopy. *arXiv preprint arXiv:2009.08328* (2020).
16. Ede, J. M. & Beanland, R. Partial Scanning Transmission Electron Microscopy with Deep Learning. *arXiv preprint arXiv:1910.10467* (2020).
17. Li, Y. Deep Reinforcement Learning: An Overview. *arXiv preprint arXiv:1701.07274* (2017).

18. Hwang, S., Han, C. W., Venkatakrishnan, S. V., Bouman, C. A. & Ortolan, V. Towards the Low-Dose Characterization of Beam Sensitive Nanostructures via Implementation of Sparse Image Acquisition in Scanning Transmission Electron Microscopy. *Meas. Sci. Technol.* **28**, 045402 (2017).
19. Hujsak, K., Myers, B. D., Roth, E., Li, Y. & Dravid, V. P. Suppressing Electron Exposure Artifacts: An Electron Scanning Paradigm with Bayesian Machine Learning. *Microsc. Microanal.* **22**, 778–788 (2016).
20. Anderson, H. S., Ilic-Helms, J., Rohrer, B., Wheeler, J. & Larson, K. Sparse Imaging for Fast Electron Microscopy. In *Computational Imaging XI*, vol. 8657, 86570C (International Society for Optics and Photonics, 2013).
21. Fang, L. *et al.* Deep Learning-Based Point-Scanning Super-Resolution Imaging. *bioRxiv* 740548 (2019).
22. de Haan, K., Ballard, Z. S., Rivenson, Y., Wu, Y. & Ozcan, A. Resolution Enhancement in Scanning Electron Microscopy Using Deep Learning. *Sci. Reports* **9**, 1–7 (2019).
23. Ede, J. M. Deep Learning Supersampled Scanning Transmission Electron Microscopy. *arXiv preprint arXiv:1910.10467* (2019).
24. Mueller, K. Selection of Optimal Views for Computed Tomography Reconstruction (2011). US Patent App. 12/842,274.
25. Wang, Z. & Arce, G. R. Variable Density Compressed Image Sampling. *IEEE Transactions on Image Process.* **19**, 264–270 (2009).
26. Ji, S., Xue, Y. & Carin, L. Bayesian Compressive Sensing. *IEEE Transactions on Signal Process.* **56**, 2346–2356 (2008).
27. Seeger, M. W. & Nickisch, H. Compressed Sensing and Bayesian Experimental Design. In *Proceedings of the 25th International Conference on Machine Learning*, 912–919 (2008).
28. Braun, G., Pokutta, S. & Xie, Y. Info-Greedy Sequential Adaptive Compressed Sensing. *IEEE J. Sel. Top. Signal Process.* **9**, 601–611 (2015).
29. Carson, W. R., Chen, M., Rodrigues, M. R., Calderbank, R. & Carin, L. Communications-Inspired Projection Design with Application to Compressive Sensing. *SIAM J. on Imaging Sci.* **5**, 1185–1212 (2012).
30. Godaliyadda, G. D. P. *et al.* A Framework for Dynamic Image Sampling Based on Supervised Learning. *IEEE Transactions on Comput. Imaging* **4**, 1–16 (2017).
31. Ermeydan, E. S. & Cankaya, I. Sparse Fast Fourier Transform for Exactly Sparse Signals and Signals with Additive Gaussian Noise. *Signal, Image Video Process.* **12**, 445–452 (2018).
32. Saldi, N., Yüksel, S. & Linder, T. Asymptotic Optimality of Finite Model Approximations for Partially Observed Markov Decision Processes With Discounted Cost. *IEEE Transactions on Autom. Control.* **65**, 130–142 (2019).
33. Jaakkola, T., Singh, S. P. & Jordan, M. I. Reinforcement Learning Algorithm for Partially Observable Markov Decision Problems. In *Advances in Neural Information Processing Systems*, 345–352 (1995).
34. Sang, X. *et al.* Dynamic Scan Control in STEM: Spiral Scans. *Adv. Struct. Chem. Imaging* **2**, 6 (2017).
35. Sang, X. *et al.* Precision Controlled Atomic Resolution Scanning Transmission Electron Microscopy Using Spiral Scan Pathways. *Sci. Reports* **7**, 43585 (2017).
36. Hochreiter, S. & Schmidhuber, J. Long Short-Term Memory. *Neural Comput.* **9**, 1735–1780 (1997).
37. Olah, C. Understanding LSTM Networks. Online: <https://colah.github.io/posts/2015-08-Understanding-LSTMs> (2015).
38. Cho, K. *et al.* Learning Phrase Representations Using RNN Encoder-Decoder for Statistical Machine Translation. *arXiv preprint arXiv:1406.1078* (2014).
39. Weiss, G., Goldberg, Y. & Yahav, E. On the Practical Computational Power of Finite Precision RNNs for Language Recognition. *arXiv preprint arXiv:1805.04908* (2018).
40. Jozefowicz, R., Zaremba, W. & Sutskever, I. An Empirical Exploration of Recurrent Network Architectures. In *International Conference on Machine Learning*, 2342–2350 (2015).
41. Bayer, J., Wierstra, D., Togelius, J. & Schmidhuber, J. Evolving Memory Cell Structures for Sequence Learning. In *International Conference on Artificial Neural Networks*, 755–764 (Springer, 2009).
42. Pascanu, R., Mikolov, T. & Bengio, Y. On the Difficulty of Training Recurrent Neural Networks. In *International Conference on Machine Learning*, 1310–1318 (2013).
43. Graves, A. *et al.* Hybrid Computing Using a Neural Network with Dynamic External Memory. *Nature* **538**, 471–476 (2016).

44. Werbos, P. J. Backpropagation Through Time: What It Does and How To Do It. *Proc. IEEE* **78**, 1550–1560 (1990).
45. Ruder, S. An Overview of Gradient Descent Optimization Algorithms. *arXiv preprint arXiv:1609.04747* (2016).
46. Mnih, V., Heess, N., Graves, A. & Kavukcuoglu, K. Recurrent Models of Visual Attention. In *Advances in Neural Information Processing Systems*, 2204–2212 (2014).
47. Ba, J., Mnih, V. & Kavukcuoglu, K. Multiple Object Recognition with Visual Attention. *arXiv preprint arXiv:1412.7755* (2014).
48. Vinyals, O. *et al.* AlphaStar: Mastering the Real-Time Strategy Game StarCraft II. <https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/> (2019).
49. Lillicrap, T. P. *et al.* Continuous Control with Deep Reinforcement Learning. *arXiv preprint arXiv:1509.02971* (2015).
50. Heess, N., Hunt, J. J., Lillicrap, T. P. & Silver, D. Memory-Based Control with Recurrent Neural Networks. *arXiv preprint arXiv:1512.04455* (2015).
51. Grabocka, J., Scholz, R. & Schmidt-Thieme, L. Learning Surrogate Losses. *arXiv preprint arXiv:1905.10108* (2019).
52. Konda, V. R. & Tsitsiklis, J. N. Actor-Critic Algorithms. In *Advances in Neural Information Processing Systems*, 1008–1014 (2000).
53. Zhao, T., Hachiya, H., Niu, G. & Sugiyama, M. Analysis and Improvement of Policy Gradient Estimation. In *Advances in Neural Information Processing Systems*, 262–270 (2011).
54. Rere, L. R., Fanany, M. I. & Arymurthy, A. M. Simulated Annealing Algorithm for Deep Learning. *Procedia Comput. Sci.* **72**, 137–144 (2015).
55. Young, S. R., Rose, D. C., Karnowski, T. P., Lim, S.-H. & Patton, R. M. Optimizing Deep Learning Hyper-Parameters Through an Evolutionary Algorithm. In *Proceedings of the Workshop on Machine Learning in High-Performance Computing Environments*, 1–5 (2015).
56. Such, F. P. *et al.* Deep Neuroevolution: Genetic Algorithms are a Competitive Alternative for Training Deep Neural Networks for Reinforcement Learning. *arXiv preprint arXiv:1712.06567* (2017).
57. Abadi, M. *et al.* TensorFlow: A System for Large-Scale Machine Learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, 265–283 (2016).
58. Ede, J. M. Adaptive Partial STEM Repository. Online: <https://github.com/Jeffrey-Ede/Adaptive-Partial-STEM> (2020).
59. Ede, J. M. & Beanland, R. Electron Microscopy Datasets. Online: <https://github.com/Jeffrey-Ede/datasets/wiki> (2020).
60. Uhlenbeck, G. E. & Ornstein, L. S. On the Theory of the Brownian Motion. *Phys. Rev.* **36**, 823 (1930).
61. Plappert, M. *et al.* Parameter Space Noise for Exploration. *arXiv preprint arXiv:1706.01905* (2017).
62. Fujimoto, S., Van Hoof, H. & Meger, D. Addressing Function Approximation Error in Actor-Critic Methods. *arXiv preprint arXiv:1802.09477* (2018).
63. Barth-Maron, G. *et al.* Distributed Distributional Deterministic Policy Gradients. *arXiv preprint arXiv:1804.08617* (2018).
64. Ede, J. M. & Beanland, R. Adaptive Learning Rate Clipping Stabilizes Learning. *Mach. Learn. Sci. Technol.* **1**, 015011 (2020).
65. Raschka, S. Model Evaluation, Model Selection, and Algorithm Selection in Machine Learning. *arXiv preprint arXiv:1811.12808* (2018).
66. Roh, Y., Heo, G. & Whang, S. E. A Survey on Data Collection for Machine Learning: A Big Data-AI Integration Perspective. *IEEE Transactions on Knowl. Data Eng.* (2019).
67. Zaremba, W., Sutskever, I. & Vinyals, O. Recurrent Neural Network Regularization. *arXiv preprint arXiv:1409.2329* (2014).
68. McCann, M. T., Jin, K. H. & Unser, M. Convolutional Neural Networks for Inverse Problems in Imaging: A Review. *IEEE Signal Process. Mag.* **34**, 85–95 (2017).
69. Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*, 1097–1105 (2012).
70. Zhang, S. & Sutton, R. S. A Deeper Look at Experience Replay. *arXiv preprint arXiv:1712.01275* (2017).
71. Mnih, V. *et al.* Human-Level Control Through Deep Reinforcement Learning. *Nature* **518**, 529–533 (2015).

72. Kingma, D. P. & Ba, J. ADAM: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980* (2014).
73. Smith, L. N. Cyclical Learning Rates for Training Neural Networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 464–472 (IEEE, 2017).
74. Ede, J. M. & Beanland, R. Improving Electron Micrograph Signal-to-Noise with an Atrous Convolutional Encoder-Decoder. *Ultramicroscopy* **202**, 18–25 (2019).
75. van Hasselt, H. P., Guez, A., Hessel, M., Mnih, V. & Silver, D. Learning Values Across Many Orders of Magnitude. In *Advances in Neural Information Processing Systems*, 4287–4295 (2016).
76. Czarnecki, W. M. *et al.* Distilling Policy Distillation. *arXiv preprint arXiv:1902.02186* (2019).
77. Lipton, Z. C. *et al.* Combating Reinforcement Learning’s Sisyphian Curse with Intrinsic Fear. *arXiv preprint arXiv:1611.01211* (2016).
78. Wagner, P. A Reinterpretation of the Policy Oscillation Phenomenon in Approximate Policy Iteration. In *Advances in Neural Information Processing Systems*, 2573–2581 (2011).
79. Jia, Y., Wu, Z., Xu, Y., Ke, D. & Su, K. Long Short-Term Memory Projection Recurrent Neural Network Architectures for Piano’s Continuous Note Recognition. *J. Robotics* **2017** (2017).
80. Ng, A. Y., Harada, D. & Russell, S. Policy Invariance Under Reward Transformations: Theory and Application to Reward Shaping. In *International Conference on Machine Learning*, vol. 99, 278–287 (1999).
81. Weill, C. *et al.* AdaNet: A Scalable and Flexible Framework for Automatically Learning Ensembles. *arXiv preprint arXiv:1905.00080* (2019).
82. Molino, P., Dudin, Y. & Miryala, S. S. Ludwig: A Type-Based Declarative Deep Learning Toolbox. *arXiv preprint arXiv:1909.07930* (2019).
83. He, X., Zhao, K. & Chu, X. AutoML: A Survey of the State-of-the-Art. *arXiv preprint arXiv:1908.00709* (2019).
84. Malekhosseini, E., Hajabdollahi, M., Karimi, N. & Samavi, S. Modeling Neural Architecture Search Methods for Deep Networks. *arXiv preprint arXiv:1912.13183* (2019).
85. Jaafra, Y., Laurent, J. L., Deruyver, A. & Naceur, M. S. Reinforcement Learning for Neural Architecture Search: A Review. *Image Vis. Comput.* **89**, 57–66 (2019).
86. Elsken, T., Metzen, J. H. & Hutter, F. Neural Architecture Search: A Survey. *arXiv preprint arXiv:1808.05377* (2018).
87. Waring, J., Lindvall, C. & Umeton, R. Automated Machine Learning: Review of the State-of-the-Art and Opportunities for Healthcare. *Artif. Intell. Medicine* 101822 (2020).
88. Hanussek, M., Blohm, M. & Kintz, M. Can AutoML Outperform Humans? An Evaluation on Popular OpenML Datasets Using AutoML Benchmark. *arXiv preprint arXiv:2009.01564* (2020).
89. Zoph, B., Vasudevan, V., Shlens, J. & Le, Q. V. Learning Transferable Architectures for Scalable Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8697–8710 (2018).
90. Ophus, C., Ciston, J. & Nelson, C. T. Correcting Nonlinear Drift Distortion of Scanning Probe and Scanning Transmission Electron Microscopies from Image Pairs with Orthogonal Scan Directions. *Ultramicroscopy* **162**, 1–9 (2016).
91. Ning, S. *et al.* Scanning Distortion Correction in STEM Images. *Ultramicroscopy* **184**, 274–283 (2018).
92. Zhu, J.-Y., Park, T., Isola, P. & Efros, A. A. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2223–2232 (2017).
93. Seki, T., Ikuhara, Y. & Shibata, N. Theoretical Framework of Statistical Noise in Scanning Transmission Electron Microscopy. *Ultramicroscopy* **193**, 118–125 (2018).
94. Laine, S., Karras, T., Lehtinen, J. & Aila, T. High-Quality Self-Supervised Deep Image Denoising. In *Advances in Neural Information Processing Systems*, 6968–6978 (2019).
95. Miller, B. & Mick, S. Real-Time Data Processing Using Python in DigitalMicrograph. *Microsc. Microanal.* **25**, 234–235 (2019).
96. Noronha, D. H., Salehpour, B. & Wilton, S. J. LeFlow: Enabling Flexible FPGA High-Level Synthesis of TensorFlow Deep Neural Networks. In *FSP Workshop 2018; Fifth International Workshop on FPGAs for Software Programmers*, 1–8 (VDE, 2018).
97. Ruan, A., Shi, A., Qin, L., Xu, S. & Zhao, Y. A Reinforcement Learning Based Markov-Decision Process (MDP) Implementation for SRAM FPGAs. *IEEE Transactions on Circuits Syst. II: Express Briefs* (2019).