

## LEARNING TO SCAN: A DEEP REINFORCEMENT LEARNING APPROACH FOR PERSONALIZED SCANNING IN CT IMAGING

Ziju Shen\*

Academy for Advanced Interdisciplinary Studies  
Peking University, China

Yufei Wang\*

Computer Science Department  
Carnegie Mellon University, USA

Dufan Wu

Center for Advanced Medical Computing and Analysis  
Massachusetts General Hospital and Harvard Medical School, USA

Xu Yang

Department of Mathematics  
University of California, Santa Barbara, USA

Bin Dong†

Beijing International Center for Mathematical Research  
Center for Data Science, Institute for Artificial Intelligence  
Peking University, China

(Communicated by Hui Ji)

**ABSTRACT.** Computed Tomography (CT) takes X-ray measurements on the subjects to reconstruct tomographic images. As X-ray is radioactive, it is desirable to control the total amount of dose of X-ray for safety concerns. Therefore, we can only select a limited number of measurement angles and assign each of them limited amount of dose. Traditional methods such as compressed sensing usually randomly select the angles and equally distribute the allowed dose on them. In most CT reconstruction models, the emphasize is on designing effective image representations, while much less emphasize is on improving the scanning strategy. The simple scanning strategy of random angle selection and equal dose distribution performs well in general, but they may not be ideal for each individual subject. It is more desirable to design a personalized scanning strategy for each subject to obtain better reconstruction result. In this paper, we propose to use Reinforcement Learning (RL) to learn a personalized scanning policy to select the angles and the dose at each chosen angle for each individual subject. We first formulate the CT scanning process as an Markov Decision Process (MDP), and then use modern deep RL methods to solve it. The learned personalized scanning strategy not only leads to better reconstruction results, but also shows strong generalization to be combined with different reconstruction algorithms.

---

2020 *Mathematics Subject Classification.* Primary: 94A08, 92C55.

*Key words and phrases.* Reinforcement learning, computed tomography, CT, neural network, PPO.

\* Equal Contribution.

† Corresponding Author.

. In this paper, we focus on the angle and dose's sample process in the whole computerized tomography. Due to the detector's limitation, we have restrictions on angle and dose of X-ray. In order to get higher precision, we can select the angle and dose from the information we have got from before data collection to achieve personalized selection. After formulating the sample process to Markov Decision Process (MDP), we use Proximal Policy Optimization (PPO) algorithm from reinforcement learning combined with neural networks and successfully get higher precision in the CT reconstruction process.

**1. Introduction.** X-ray based computed tomography (CT) is a medical imaging procedure that reconstructs tomographic images by taking X-ray measurements from different angles. To obtain high-quality reconstructions, in early reconstruction algorithms such as filtered backprojection (FBP) [35] and algebra reconstruction technique (ART) [28], a number of different angles need to be measured. However, since X-ray is radioactive, the total dose of X-ray needs to be restricted in the scanning process, and thus we need to either decrease the X-ray intensity at each chosen angle, or to reduce the total number of angles taken. Decreasing X-ray intensity in each angle will result in more noisy measurements, while fewer angles will reduce the information we need for a high-quality reconstruction. This causes great challenges in designing efficient and effective reconstruction algorithms.

Compressed sensing [19] resolves the issue to a certain extent. According to the theory of compressed sensing, if an image has a sparse property after certain transformations (e.g., wavelet transform), then it can be robustly reconstructed with a reduced number of random measurements by solving an  $l_1$ -minimization problem when the measurements and the transformation satisfy the D-RIP condition [10]. We can use the alternating direction method of multipliers (ADMM) [26, 6, 8] or primal-dual hybrid gradient method (PDHG) [12, 79, 22] to solve this  $l_1$ -minimization problem to obtain a reconstructed image.

In the literature of CT image reconstruction or image restoration in general, people focus on designing effective regularizations, which includes the total variation (TV) [56], nonlocal means [7], block-matching and 3-D filtering (BM3D) [14], weighted nuclear norm minimization (WNNM) [29], wavelets and wavelet frame models [16, 40, 18], K-SVD [21], data-driven (tight) frame [9, 65], low dimensional manifold method (LDMM) [51], etc. More recently, the rapid development of machine learning, especially deep learning, has led to a paradigm shift of modeling and algorithmic design in computer vision and medical imaging [66, 68, 41, 69, 77]. Deep learning based models are able to leverage large image datasets to learn better image representations and produce better image reconstruction results than traditional methods [13, 32, 34, 78, 72, 60].

In most CT reconstruction models, the emphasize is on designing effective image representation, while much less emphasize is on improving the scanning strategy. In compressed sensing, the scanning strategy is entirely random [11, 10], i.e., the measurement angles are selected randomly and the dose are allocated uniformly across the angles. In theory, such random sampling is proven for exact recovery using a convex model for MRI. However, such result is generally untrue for CT imaging due to the coherence structure of Radon transform. In practice, uniform sampling is often adopted. However, for each individual subject, a uniform or random scanning strategy may not be ideal. It is more desirable to design a personalized scanning strategy for each subject to achieve better reconstruction results. Our

key observation is that the measurements collected in the early stage during the scanning process can be used to guide the later scanning.

Despite the potential improvement of a personalized scanning strategy for each individual subject, it is very difficult to handcraft such a strategy by a human expert. This is where machine learning can help. The personalized scanning strategy can be learned using either active learning [59] or Reinforcement Learning (RL) [38]. In this paper, we propose to use reinforcement learning to learn such a personalized scanning strategy for each subject. The reason we choose RL over active learning is that RL is non-greedy and naturally guarantees the long-term reconstruction quality. We formulate the CT scanning process as a Markov Decision Process (MDP), where the state includes currently collected measurements, the action determines the next measurement angle and the dose usage, and the reward depends on the reconstruction quality. We further use modern Deep RL algorithms to solve it. We show in the experiments that the personalized scanning policy learned by RL significantly outperforms the random scanning strategy in terms of the reconstruction quality, and can generalize to be combined with different reconstruction algorithms.

We note that current commercial CT system does not support to select angles and dose as freely as what the trained RL agent suggests since it requires fast and irregular acquisition motion of the gantry. The proposed approach serves as a proof of concept and experimental study to investigate the potential benefits of the scanning strategies suggested by the trained RL agents. Nonetheless, there are multi-beam X-ray systems available to which the proposed scanning strategy can be applied to [74, 70, 52].

**1.1. Related works.** For compressed sensing, there have been two primary categories of scanning strategies: static and dynamic. Static scanning strategy refers to the method which collects measurements in a fixed order. Low-discrepancy sampling [50] and uniformly spaced sparse sampling methods [46] are two examples of static scanning strategy. Non-uniform static scanning strategy based on the model of the subject to be scanned is proposed in [48, 67]. However, because the order of measurements is predetermined, static scanning strategy is not flexible for different subjects and may lead to poor results for some of them.

Dynamic scanning strategy refers to the methods which collect measurement adaptively based on information obtained from previous measurements. One traditional method tries to find the most suitable measurements which can minimize the entropy to decrease uncertainty of images, such as BCS [31, 58]. Similarly, other methods [4, 15] use the information gained at each additional scan to guide the selection of the next measurement. However, these methods are typically greedy methods in nature, have many hyperparameters to be properly tuned, and are slow during inference as they either need to take inverse of large matrices, or to run the reconstruction algorithm for many times when determining the best next angle. More recently, deep neural networks are used to estimate the expected reduction in distortion (ERD) in the reconstructed image when an additional measurement is selected [24, 25, 76, 75, 30, 47]. However, for the estimate of ERD to be accurate, it requires a large number of measurements in training.

All the above methods are not specific to CT scanning. They are greedy in nature and do not provide a strategy for dose allocation. In contrast, RL is able to generate a non-greedy policy that aims at maximizing long-term rewards which, in this paper, is the quality of the reconstructed CT image. Furthermore, the setting of RL is flexible enough to handle both angle selection and dose allocation, and

even more decision options during scanning. Therefore, in this paper we use RL to design a scanning policy that acts optimally on each individual subject. In Scanning Transmission Electron Microscopy (STEM), a recent work by [20] proposes to use RL to guide the movement of the detector and uses a generator to generate reconstructed images. However, since the image modality is drastically different from CT, the proposed MDP (especially the state, action and architecture of the policy network) is vastly different from what is proposed in this paper.

Recently, there has been a line of work that uses RL to solve combinatorial optimization problems, such as Travelling Salesman Problem [5], Vehicle Routing Problem [37], Influence Maximization [43], Autonomous Exploration [39], and they all show RL can obtain better results than traditional solvers. Given the success of these prior works, and that the angle selection in CT scanning is also a combinatorial optimization problem, it naturally motivates us to try to use RL to solve the problem and see if it brings any benefit.

## 2. Preliminaries.

**2.1. A brief review on MDP and reinforcement learning [63].** A sequential decision problem can be formulated as a Markov Decision Process (MDP). MDP is a tuple  $(\mathcal{S}, \mathcal{A}, \gamma, \mathbb{P}, r)$  that consists of the state space  $\mathcal{S}$ , the action space  $\mathcal{A}$ , the discount factor  $\gamma$ , the transition probability of the environment  $\mathbb{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  and the reward  $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{R}$ . A policy  $\pi$  in RL is a probability distribution on the action  $\mathcal{A}$  over  $\mathcal{S}$ :  $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ . Denote the interactions between the agent and the environment as a trajectory  $\tau = (s_0, a_0, r_0, \dots, s_T, a_T, r_T, \dots)$ . The return of  $\tau$  is the discounted sum of all its future rewards:

$$G(\tau) = \sum_{t=0}^{\infty} \gamma^t r_t.$$

Given a policy  $\pi$ , the value of a state  $s$  is defined as the expected return of all the trajectories when the agent starts at  $s$  and then follows  $\pi$ :

$$V^\pi(s) = E_\tau[G(\tau) | \tau(s_0) = s, \tau \sim \pi]$$

Similarly, the value of a state-action pair is defined as the expected return of all trajectories when the agent starts at  $s$ , takes action  $a$ , and then follows  $\pi$ :

$$Q^\pi(s, a) = E_\tau[G(\tau) | \tau(s_0) = s, \tau(a_0) = a, \tau \sim \pi]$$

Given an MDP, the goal of a reinforcement learning algorithm is to find a policy  $\pi$  that maximizes the discounted accumulated rewards in this MDP:

$$(1) \quad \max_{\pi} \eta(\pi) = \mathbb{E}_{s_0 \sim \rho(s)}[V^\pi(s)].$$

Many effective RL algorithms have been developed to find the optimal policy  $\pi$ . They can be generally classified into two categories: valued-based methods and policy gradient methods.

Value-based methods such as Q-learning [71, 44] uses the Bellman Equation to learn the optimal Q function, and then derive the optimal policy by acting greedily according to the optimal Q function. Formally, the Bellman Optimal Equation is:

$$Q^*(s, a) = \max_{\pi} Q^\pi(s, a)$$

$$Q^*(s_t, a_t) = r(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1} \sim P(\cdot | s_t, a_t)} \left[ \max_{a_{t+1}} Q^*(s_{t+1}, a_{t+1}) \right]$$

Methods like DQN [45] uses a deep neural network  $Q_w$  to learn the optimal  $Q^*$  by storing transitions  $\{s, a, r, s'\}$  in an off-line replay buffer and minimizing the following Bellman Error:

$$L(w) = \mathbb{E}_{s,a,r,s'}[(Q_w(s, a) - (r + \gamma \max_{a'} Q_{\tilde{w}}(s', a'))^2]$$

where  $Q_{\tilde{w}}$  is the target Q network whose parameters  $\tilde{w}$  are periodically copied from  $w$  to stabilize training. After the optimal Q function is learned, the policy is simply  $\pi(s) = \arg \max_a Q_w(s, a)$ . Due to this  $\arg \max$  operator, value-based methods are mostly suitable for discrete actions. Because our action contains continuous variables (the dose allocation), we do not use value-based methods as our RL algorithm.

Policy gradient methods [64, 62] directly optimize a parameterized policy by computing a surrogate objective. Given a policy  $\pi_\theta$  parameterized by a neural network  $\theta$ , the policy gradient theorem [64] states that:

$$\nabla_\theta \eta(\pi_\theta) = \mathbb{E}_{s,a \sim \pi_\theta} [\nabla_\theta \log \pi_\theta(a|s) \cdot Q^{\pi_\theta}(s, a)]$$

Policy gradient methods is suitable for continuous action space. One drawback of vanilla policy gradient method is that the gradient might have high variance and make the update of  $\theta$  unstable. Many more advanced methods have then been proposed to address this. In this paper, we use the Proximal Policy Optimization (PPO) algorithm [57], which updates the policy in a proximal region to avoid unstable updates. We now briefly review how it works. Given a parameterized policy  $\pi_\theta$ , its advantage function is defined as  $A^{\pi_\theta}(s, a) = Q^{\pi_\theta}(s, a) - V^{\pi_\theta}(s)$ . Given an old policy  $\pi_{\theta_{old}}$ , let  $b_\theta(s, a) = \frac{\pi_\theta(a|s)}{\pi_{\theta_{old}}(a|s)}$ , PPO optimizes  $\pi_\theta$  w.r.t. the following surrogate objective using gradient descent:

(2)

$$J^{\text{PPO}}(\theta) = \mathbb{E}_{s,a \sim \pi_{\theta_{old}}} [\min(b_\theta(s, a) A^{\pi_{\theta_{old}}}(s, a), \text{clip}(b_\theta(s, a), 1 - \varepsilon, 1 + \varepsilon) A^{\pi_{\theta_{old}}}(s, a))],$$

where  $\text{clip}(x, a, b) = \max(\min(x, b), a)$ .

**2.2. CT reconstruction.** One of the common CT systems is the cone-beam CT system. In the 2-dimensional case, it is known as the fan-beam CT, and this is the CT system that we focus on in this paper. For a given angle  $\theta$  and X-ray beamlet  $r$ , the X-ray projection operator  $A^{\theta,r}$  is defined as follows:

$$(3) \quad A^{\theta,r}[f](t) = \int_0^{L(t)} f(\mathbf{x}_\theta + \mathbf{n}l) dl$$

where  $f$  is the unknown image (X-ray attenuation coefficients) that needs to be reconstructed,  $\mathbf{x}_\theta = (x_\theta, y_\theta)$  represents the coordinate of the X-ray source which is different for different projection angle  $\theta$ ,  $\mathbf{n} = (n_x, n_y)$  is the direction vector of beamlet  $r$ ,  $t$  is the coordinate on the X-ray imager which is precisely the intersection of the beamlet  $r$  with the X-ray imager.  $L(t)$  is the length of the X-ray beamlet from the source to the location  $t$  on the imager. If  $A^{\theta,r}[f](t)$  is sampled with respect to  $t$  for each angle  $\theta$ , the resulting data projection can essentially be written as a vector  $p_\theta$ . Now, putting the vectors  $p_\theta$  together for all different angles  $\theta$ , we obtain an image denoted as  $p$  whose columns are formed by  $p_\theta$ .

We can write the CT image reconstruction problem as a linear inverse problem

$$(4) \quad p = Af + \epsilon,$$

where  $A$  is the linear operator represents the collection of discrete line integration at different projection angles and along different beamlets, and  $\epsilon$  is an additive noise. In our simulations, the matrix  $A$  is generated by Siddon's algorithm [54] As equation

(4) is a linear equation, it can be directly solved by ART or SART algorithm [28, 49]. However, as  $m$  can be far smaller than  $n$ , equation (4) has far less equations than unknowns. In order to obtain high-quality solutions, regularization-based models are often used, which typically take the form as follows:

$$(5) \quad \min_f \frac{1}{2} \|p - Af\|_2^2 + \lambda R(f),$$

where  $R(f)$  is the regularization term. Two benchmarking regularization terms are TV regularization  $R(f) = \|\nabla f\|_1$  [61] and wavelet regularization  $R(f) = \|Wf\|_1$  [17], where  $W$  is the wavelet transform. Both of these two optimization problems can be solved by ADMM or PDHG.

**2.3. Relationship between measurement noise and dose.** Noise intensity on measurements heavily relies on the X-ray dose. It is common to assume that the measurement noise follows a Gaussian distribution [73],  $\epsilon \sim \mathcal{N}(0, \sigma)$ , and

$$(6) \quad \sigma \propto \frac{1}{\sqrt{n_{\max} d \exp(-P)}},$$

where  $d$  is the X-ray dose used in a measurement,  $P$  is the average intensity of measurement, and  $n_{\max}$  is the maximum number of photons the source can generate. We can easily see that if we use more dose, the noise level becomes smaller. We note that the Poisson distribution can also be used here. However, according to [73], the Gaussian distribution approximates the noise well enough.

**2.4. Some further discussions.** As equation (4) shows, the measurements we obtain from a CT scan depends both on the angle (which determines  $A$ ), and the X-ray dose (which determines  $\epsilon$ ). Due to the limitation on X-ray dose usage, we can only select a limited number of angles and assign each of them limited amount of dose. Traditional methods simply randomly select the angles and equally distribute the allowed dose on them. Our goal is to use RL to learn a personalized policy to select the angles and the dose at each chosen angle for each individual subject.

**3. Method.** Our goal is to learn a policy that can decide the next measurement angle and its corresponding X-ray dose based on the measurements that we have already obtained in the scanning process. We now present how the scanning process can be formulated as a Markov Decision Process (MDP) and solved by reinforcement learning algorithms.

**3.1. MDP formulation of personalized scanning.** We note that the angle selection problem in CT scanning itself is a combinatorial optimization problem and is NP-Hard. However, we can view it as a sequential decision process and use RL to solve it effectively. This is similar to works that use RL to solve other combinatorial optimization problems such as Influence Maximization [43]. Specifically, the CT scanning process can be viewed as a sequential decision process, where at each time step we need to decide on the measurement angle and the corresponding X-ray dose. Given an Image  $I$  and the number of all possible angles  $N$  (e.g.,  $N = 360$  if we can choose all the integer angles from  $0^\circ$  to  $359^\circ$ ), we now elaborate how the CT scanning process on  $I$  can be formulated as an MDP:

- 1) The **state** is a sequence  $\vec{s}_t = (s_1, s_2, \dots, s_t)$ , where  $s_t = (p_t, d_t^{\text{ac}}, d_t^{\text{rest}})$ .  $p_t$  is the collected measurement at time step  $t$ .  $d_t^{\text{ac}}$  records the used dose distribution up to time step  $t$ . It is an  $N$  dimensional vector, and the value at each

entry represents the used X-ray dose at that corresponding angle.  $d_t^{\text{rest}}$  is a scalar that represents the amount of the remaining dose that we can use. Because the reconstructed image at time step  $t$  relies on all previously collected measurements, we include all of them in the state.

- 2) The **action** is  $a_t = (a_t^{\text{angle}}, a_t^{\text{dose}})$ .  $a_t^{\text{angle}}$  is the angle we choose at time step  $t$ , and it is a one-hot vector of dimension  $N$ .  $a_t^{\text{dose}} \in [0, 1]$  is the fraction of dose that we apply at the corresponding angle. If at a certain time  $t$ , the total used dose exceeds the total allowed dose, we clip the exceeding dose and terminate the MDP.
- 3) The **reward** is computed as  $r(s_t, a_t) = \text{PSNR}(I_t, I) - \text{PSNR}(I_{t-1}, I)$ , where  $I$  is the groundtruth image,  $I_t$  is the reconstructed image at time step  $t$ , and  $\text{PSNR}(\hat{I}, I)$  represents the Peak Signal to Noise Ratio (PSNR) value of the reconstructed image  $\hat{I}$ . We use the increment of PSNR to evaluate how much benefit the new chosen angle/dose brings. The reconstructed image  $I_t$  can be obtained from any reconstruction algorithm such as SART, TV regularization [56, 27], wavelet frame regularization [55, 8], and any modern deep learning based methods [3]. Note that a more refined image reconstruction algorithm leads to higher quality reconstructed image which may bring benefit for training the RL policy. However, regularization based methods normally have at least one hyperparameter that needs to be tuned when number of projections changes for optimal reconstruction. On the other hand, deep learning methods normally have generalization issue when number of projections vary drastically. Therefore, for practical concern, it is more convenient to apply SART or other iterative reconstruction methods that do not have generalization issues and do not require hyperparameter tuning. Furthermore, angle selection and dose allocation may not require very refined reconstructions, but rather a global reconstruction of anatomical structures. Therefore, we shall use SART to compute the reward.
- 4) The **transition model**  $\mathbb{P}$  represents the scanning process of CT. At time step  $t$ , given the state  $\vec{s}_t$  and action  $a_t$ , the next state  $\vec{s}_{t+1}$  is simply the concatenation of  $\vec{s}_t$  and  $s_{t+1} = (p_{t+1}, d_{t+1}^{\text{ac}}, d_{t+1}^{\text{rest}})$ . We now show how each of the three elements in  $s_{t+1}$  can be computed.
  1. The new measurement is obtained as  $p_{t+1} = p_{t+1}^{\text{true}} + \epsilon$ , where  $p_{t+1}^{\text{true}}$  is the clean projected value obtained using the chosen angle  $a_t^{\text{angle}}$ , and  $\epsilon$  is the measurement noise. The noise depends on the chosen dose  $a_t^{\text{dose}}$  as mentioned in section 2.3:  $\epsilon \sim \mathcal{N}(0, \sigma)$ ,  $\sigma \propto \frac{1}{\sqrt{n_{\max} a_t^{\text{dose}} \exp(-P)}}$ , where  $P$  is the average of  $p_{t+1}^{\text{true}}$ .
  2. The new dose distribution is obtained by adding the new decision:  $d_{t+1}^{\text{ac}} = d_t^{\text{ac}} + \mathbf{1}_{a_t^{\text{angle}}} \cdot a_t^{\text{dose}}$ , where  $\mathbf{1}_{a_t^{\text{angle}}}$  is the one-hot vector of the chosen angle  $a_t^{\text{angle}}$ .
  3. The rest amount of dose is updated by subtracting the used dose:  $d_{t+1}^{\text{rest}} = d_t^{\text{rest}} - a_t^{\text{dose}}$ . The MDP terminates once the dose is used up.

As we choose the increment in PSNR as the reward, the total sum of reward (when there is no discounting, as in our experiments) is the PSNR value of the final reconstructed image. Therefore, if we find the optimal policy to this MDP, it will also have the best reconstruction result for the image.

We note that, the final output of the RL policy is a set of angles (as well as the corresponding allocated doses). Although the angles are selected by the RL policy

in an orderly fashion, in terms of image reconstruction afterwards, such order is not essential. However, the order is essential during the angle selection stage, i.e., the inference stage of the trained RL policy. This is because during inference, we do not know the structure of the object to be imaged a priori. Thus, it is difficult to make a global personalized decision by selecting all angles at once. What the RL policy does is by making smart angle selections step by step (e.g. by first selecting a few angles around the object to acquire a coarse estimation of the global structure and refine it afterwards) so that the final collection of angles is good. A poorly trained policy may make poor sequential angle selection from the beginning which can eventually lead to an inferior collection of angles.

**3.2. Policy network architecture.** Because we include all the previous measurements in the state, the dimension of the state vector increases as we take more measurements. To handle the varying dimensionality of the state vector, we represent the policy network as a Recurrent Neural Network (RNN), so all the information from the past measurements can be encoded in the hidden state of the RNN. Specifically, we use the Gated Recurrent Unit (GRU). Besides, the policy network needs to output two different actions: the discrete action for choosing the angle  $a_t^{\text{angle}}$ , and the continuous action for choosing the dose  $a_t^{\text{dose}}$ . To handle this, we design a special architecture for the policy network, as shown in Figure 1. We use separate Multi-Layer Perceptron (MLP) after the RNN hidden states for these two actions.  $a_t^{\text{angle}}$  is sampled from a probability vector of length  $N$ , where the value at each entry represents the probability of choosing that angle. We use softmax after the final linear layer to obtain the probability vector. We also introduce a mask to remove the previously chosen angles, which is obtained from  $d_t^{\text{ac}}$  that records the used dose distribution on all angles up to time  $t$ . For the dose usage  $a_t^{\text{dose}}$ , we assume  $a_t^{\text{dose}} \sim \mathcal{N}(\mu^{\text{dose}}, \sigma^{\text{dose}})$ , with the mean and std both learned by a MLP. It is natural to determine the amount of dose after the angle is chosen, i.e.,  $\pi_\theta(a|\vec{s}) = \pi_\theta(a^{\text{dose}}, a^{\text{angle}}|\vec{s}) = \pi_\theta(a^{\text{angle}}|\vec{s})\pi_\theta(a^{\text{dose}}|\vec{s}, a^{\text{angle}})$ , so we concatenate the one-hot vector of the chosen angle as part of the input for the dose MLP.

#### 4. Experiments.

**4.1. Experiment setup.** We train the RL policy on 250 CT images of size  $512 \times 512$  from the AAPM dataset of the “2016 NIH-AAPM-Mayo Clinic Low Dose CT Grand Challenge” [42]. The 250 training images are slices of the 3D CT image of one patient. The ground-truth images are obtained by using the commercial reconstruction algorithm with normal dose. We use ASTRA Toolbox [1, 2] to generate the sinogram data. During training, we use SART as the reconstruction algorithm for computing the reward. The possible angles are all integers in  $[0^\circ, 360^\circ]$ . We use Adam [36] to optimize both the policy network and the value network, with a learning rate of 0.0004, and  $\beta_1, \beta_2 = (0.5, 0.999)$ . More detailed hyperparameters for PPO and network architecture can be found in the code which will be released upon acceptance of this paper.

After training, we test the learned RL policy on 350 slices of CT images of another patient from the AAPM dataset. We compare the following three scanning strategies: (1) **UF-AEC**, which selects angles uniformly and distributes the dose by automatic exposure control (AEC)[23, 33]; (2) **DS-ED**, which selects angle by a dynamic sample strategy based on entropy from [53], while distributes the dose uniformly; (3) **RL-AD**: which uses the learned personalized policy for both angle

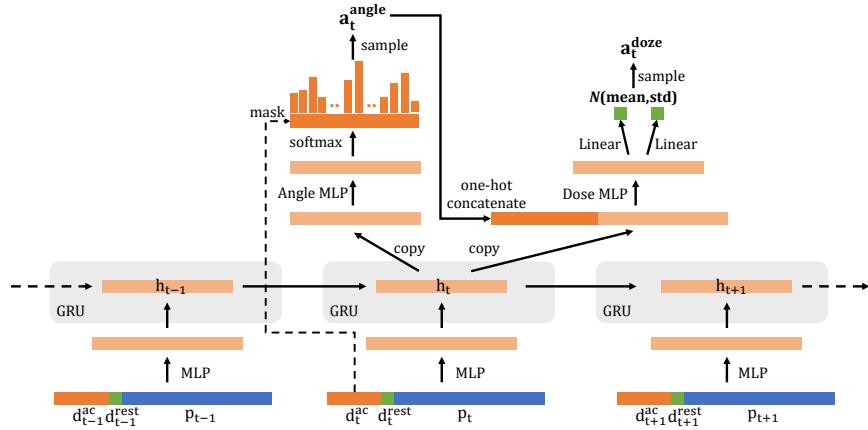


FIGURE 1. Policy network architecture. Each MLP contains two hidden layers with 512 neurons. We use a multi-layer GRU which contains 3 recurrent layers and each layer has 256 neurons. The Angle MLP has one hidden layer of 512 neurons, and the Dose MLP has 2 hidden layers with 512 neurons.

selection and dose allocation at each chosen angle. During testing, we use four different reconstruction algorithms: SART, TV regularization (TV) [56, 27], wavelet frame (WF) regularization [55, 8], and the recently proposed deep learning method PD-net [3]. Note that during the angle selection stage, i.e. the inference stage of the trained RL policy, we do not need to conduct image reconstruction. The decision on the next angle and the associated dose is determined from the sinogram formed by the angles and doses that have already been selected by the RL policy. We run the reconstruction algorithm after all angles are selected by the RL policy. In other words, the angles/dose selection and image reconstruction are two independent stages. The evaluation metric is PSNR and the structure similarity metric (SSIM) of the reconstructed images. In the testing phase, we add an additional Poisson noise to the sinogram. Recall that the original noise in the sinogram is Gaussian as described in Section 2.3. Note that we add additional noise to test the robustness and generalization ability of our trained RL policy, and we choose Poisson noise since it is widely accepted in CT imaging that the measurement noise is Poisson. The incident photon intensities of the additional Poisson noise are  $10^7$ ,  $10^6$  and  $10^5$ . We denote these three levels of noise as Noise 1, Noise 2 and Noise 3.

A difficulty in conducting a fair comparison of UF-AEC and DS-ED with RL-AD is that the number of selected angles of RL-AD is personalized and hence different for different subjects (see Figure 4). In our experiments below, we choose the number of measurement angles for RD-AEC and DS-ED to be 60, which is the mean number of measurement angles selected by RL-AD over all the 350 test images. Thus, the dose on each measurement angle of UF-AEC and DS-ED is 1/60. We also note that the deep reconstruction model PD-net is trained from scratch on the 250 images in the training set using 60 angles and cone-beams geometry.

**4.2. Results.** Table 1 presents the mean and standard deviation of the PSNR and SSIM values of the reconstructed images of all compared scanning strategies and reconstruction algorithms. As one can see that the proposed scanning strategy

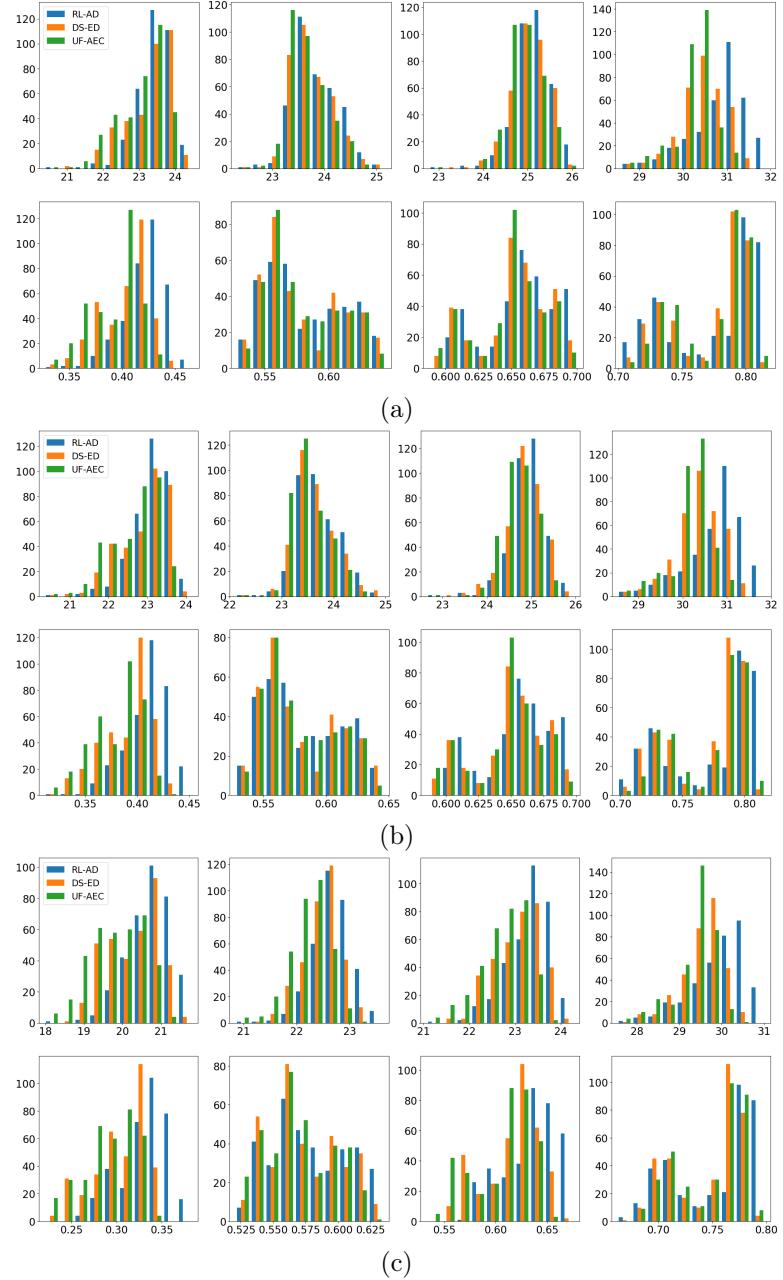


FIGURE 2. Histograms of PSNR and SSIM from all the 350 test images as shown in Table 1. Figures in (a), (b) and (c) correspond to the three different noise levels. For each (a), (b) and (c), the first row is PSNR and the second is SSIM. Figures from left to right are results from reconstruction methods SART, TV, WF and PD-net respectively. Every sub-figure contains histograms of three scanning strategies, i.e., RL-AD, DS-ED and UF-AEC.

Reconstruction Method		RL-AD	DS-ED	UF-AEC
Noise 1				
SART	PSNR	<b>23.48(0.47)</b>	23.30(0.64)	23.01(0.64)
	SSIM	<b>0.424(0.020)</b>	0.403(0.023)	0.391(0.022)
TV	PSNR	<b>23.85(0.42)</b>	23.75(0.41)	23.63(0.38)
	SSIM	<b>0.582(0.030)</b>	0.579(0.030)	0.578(0.028)
WF	PSNR	<b>25.14(0.40)</b>	25.05(0.42)	24.91(0.39)
	SSIM	<b>0.659(0.027)</b>	0.652(0.027)	0.649(0.026)
PD-net	PSNR	<b>30.87(0.64)</b>	30.44(0.51)	30.23(0.46)
	SSIM	<b>0.776(0.036)</b>	0.771(0.029)	0.773(0.028)
Noise 2				
SART	PSNR	<b>23.15(0.48)</b>	22.91(0.53)	22.60(0.64)
	SSIM	<b>0.413(0.020)</b>	0.390(0.024)	0.378(0.024)
TV	PSNR	<b>23.74(0.40)</b>	23.50(0.36)	23.27(0.40)
	SSIM	<b>0.580(0.030)</b>	0.576(0.030)	0.573(0.028)
WF	PSNR	<b>24.98(0.29)</b>	24.84(0.41)	24.68(0.39)
	SSIM	<b>0.657(0.027)</b>	0.649(0.026)	0.646(0.026)
PD-net	PSNR	<b>30.78(0.64)</b>	30.35(0.51)	30.15(0.77)
	SSIM	<b>0.774(0.037)</b>	0.769(0.030)	0.771(0.029)
Noise 3				
SART	PSNR	<b>20.71(0.55)</b>	20.26(0.72)	19.83(0.66)
	SSIM	<b>0.334(0.026)</b>	0.304(0.030)	0.291(0.029)
TV	PSNR	<b>21.73(0.57)</b>	21.43(0.48)	21.08(0.47)
	SSIM	<b>0.568(0.027)</b>	0.555(0.026)	0.545(0.026)
WF	PSNR	<b>23.35(0.48)</b>	23.05(0.51)	22.72(0.55)
	SSIM	<b>0.636(0.0326)</b>	0.616(0.027)	0.605(0.028)
PD-net	PSNR	<b>29.97(0.66)</b>	29.56(0.51)	29.36(0.47)
	SSIM	<b>0.753(0.038)</b>	0.746(0.032)	0.747(0.031)
Inference Time (s)		0.46(0.02)	0.21(0.008)	0.20(0.001)

TABLE 1. This table presents comparisons of different scanning strategies (1-3rd column for RL-AD, DS-ED and UF-AEC respectively) combined with different image reconstruction methods (1-4th row for SART, TV, WF and PD-net respectively). Last row presents the inference times of angle selection (in seconds) of the three compared scanning strategies. The mean (std) of the PSNR and SSIM of the reconstructed images and the inference times are computed among all 350 testing CT images. The best results among the compared algorithms are shown in bold numbers.

RL-AD significantly outperforms dynamic sampling (DS-ED) and uniform sampling (UF-AEC), while DS-ED outperforms UF-AEC.

We also note that the RL-policy is trained only using the SART for computing the reward function, whereas the learned policy can generalize well to three other reconstruction algorithms, i.e., the TV regularization, the wavelet frame regularization and the deep learning model PD-net, where it still brings a notable improvement upon the dynamic sampling and uniform scanning baseline in reconstruction quality. We further note that during training, only Gaussian noise was included following

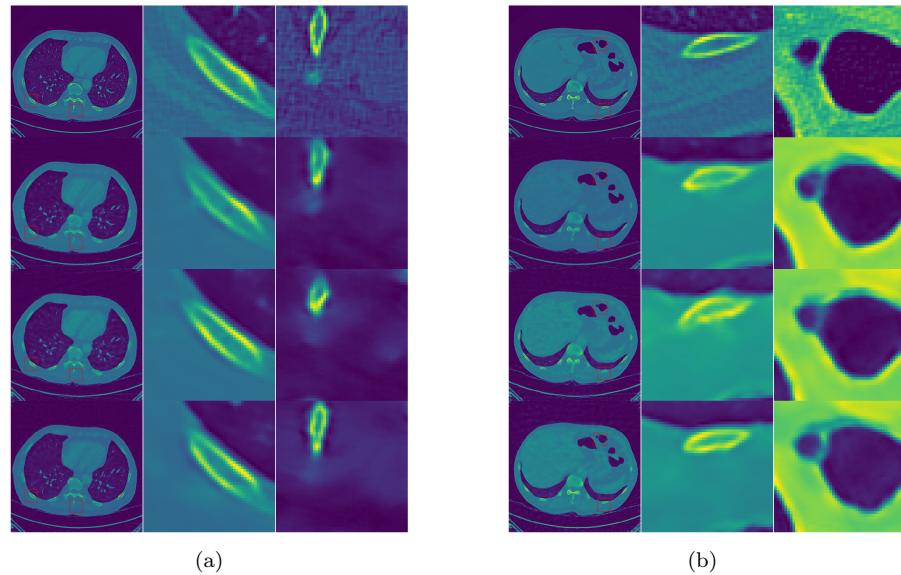


FIGURE 3. Two examples of the reconstructed images. The top row contains the ground truth images and their zoom-in views. The second through the fourth row contain results from UF-AEC, DS-ED and RL-AD respectively, and combined with PD-net's reconstruction. Note that RL-AD selects 65 measurement angles for the subject in (a) and 54 measurement angles for the subject in (b).

the formula given in Section 2.3. The results in Table 1 also shows that the trained RL-policy is also transferable to different noise levels.

More fine-grained demonstrations of the results shown in Table 1 are given in Figure 2, where we present the histograms of the compared scanning strategies and reconstructions methods. As one can see that the proposed scanning strategy by RL generally shifted the histogram towards the right and outperforms DS and UF by a significant margin.

In Figure 3, we further show two examples of the reconstructed images using the uniform sampling (UF-AEC), the dynamic sampling strategy (DS-ED) and the learned personalized policy (RL-AD), reconstructed using the deep learning model PD-net. We can see that the reconstructions using the RL policy are of higher qualities than those using random and dynamic sampling strategy, especially from the zoom-in views of the figures.

We plot the distribution of number of measurements taken by the learned personalized policy (RL-AD) in Figure 4 (a). The result demonstrates that for different subjects, the learned RL policy selects different number of angles and dose allocations. In Figure 4 (b) and (c), we take 8 images on which the learned RL policy selects 54 and 64 angles respectively and plot the distributions of the dose usage of these images. It can be seen that images using the same number of measurement angles have very similar dose allocations, and images that have more measurement angles use less dose at each angle. In Figure 5, we show 2 example images where the RL policy selects 54 and 65 measurement angles respectively. We can see that images upon which the RL policy selects more measurement angles have more structures in

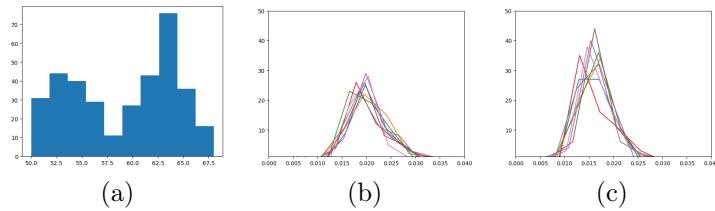


FIGURE 4. (a) Distribution of number of measurements of the learned policy (RL-AD) on all the 350 testing CT images. (b) Dose usage distribution of 8 images that use around 54 measurements. (c) Dose usage distribution of 8 images that use around 65 measurements.

the image, and thus more information/measurements need to be collected to obtain a high-quality reconstruction. In Figure 6, we present the selected angles and part of the dose allocation on the subjects shown in Figure 5.

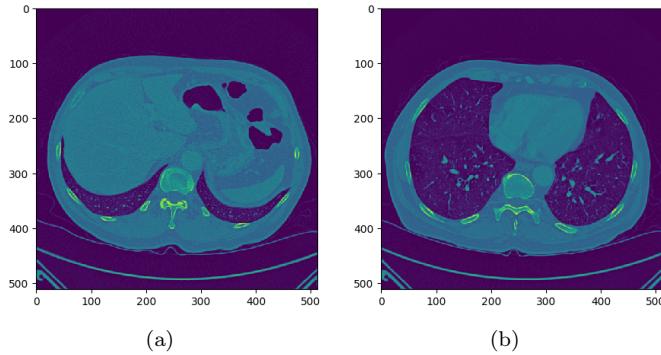


FIGURE 5. (a): an example image that takes 54 measurements. (b): an example image that takes 64 measurements. We can see that the images for which RL selects more measurement angles contains more structures.

**5. Conclusion.** In this paper, we proposed to use reinforcement learning to learn a personalized CT scanning strategy for measurement angle selection and dose allocation. We formulated the CT scanning process as a Markov Decision Process, and used the PPO algorithm to solve it. After training on 250 real 2D CT images, we validated the learned personalized scanning policy on another 350 CT images. Our validation showed that the personalized scanning policy lead to better overall reconstruction results in terms of PSNR values, and generalized well to be combined with different reconstruction algorithms. We also demonstrated that the personalized policy can indeed adjust its angle selection and dose allocations adaptively to different subjects. One drawback of the proposed method is the long training time (approximately 24 hours) even for 2D images, because RL algorithms usually need lots of simulation samples to converge, and to compute the reward in our formulated MDP requires running a reconstruction algorithm at each time step. This might prohibits the application of our method to 3D cases.

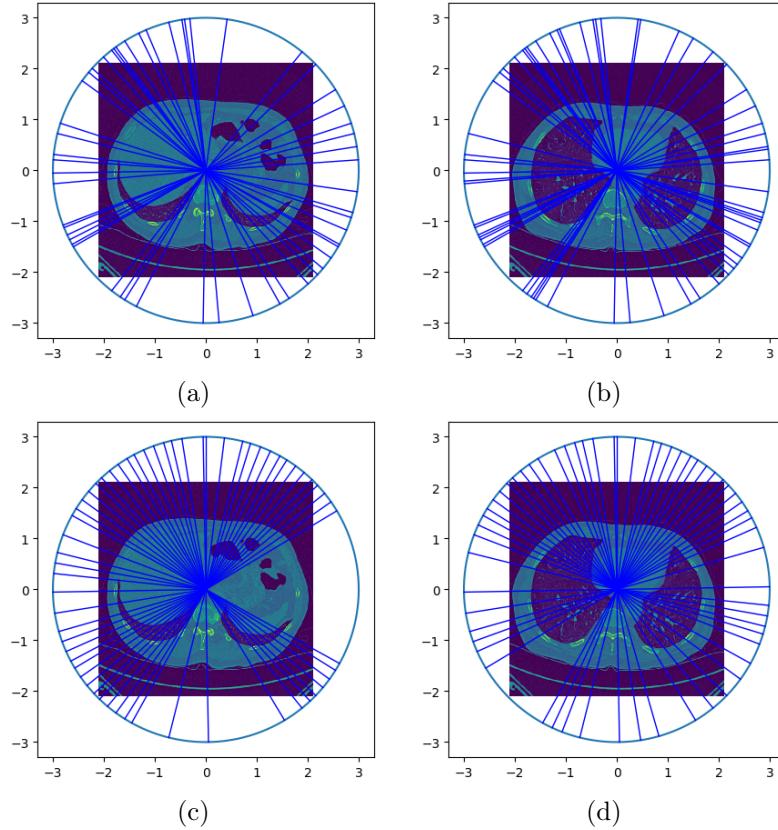


FIGURE 6. The angle selection of the CT image in Figure 5. Top row: RL-AD, bottom row: DS-ED. The lines show the selected angles.

**Acknowledgments.** Bin Dong is supported in part by National Natural Science Foundation of China (NSFC) grant No. 11831002, Beijing Natural Science Foundation (No. 180001) and Beijing Academy of Artificial Intelligence (BAAI).

#### REFERENCES

- [1] W. van Aarle, W. J. Palenstijn, J. Cant, E. Janssens, F. Bleichrodt, A. Dabovolski, J. De Beenhouwer, K. J. Batenburg and J. Sijbers, Fast and flexible x-ray tomography using the astra toolbox, *Optics Express*, **22** (2016), 25129–25147.
- [2] W. van Aarle, W. J. Palenstijn, T. Altantzis, S. Bals, K. J. Batenburg and J. Sijbers, The astra toolbox: A platform for advanced algorithm development in electron tomography, *Ultramicroscopy*, **24** (2015), 35–47.
- [3] J. Adler and O. Oktem, Learned primal-dual reconstruction, *IEEE Transactions on Medical Imaging*, **37** (2018), 1322–1332.
- [4] K. J. Batenburg, W. J. Palenstijn, P. Balázs and J. Sijbers, Dynamic angle selection in binary tomography, *Computer Vision and Image Understanding*, **117** (2013), 306–318.
- [5] I. Bello, H. Pham, Q. V Le, M. Norouzi and S. Bengio, Neural combinatorial optimization with reinforcement learning, preprint, [arXiv:1611.09940](https://arxiv.org/abs/1611.09940), 2016.
- [6] S. Boyd, N. Parikh, et al., Distributed optimization and statistical learning via the alternating direction method of multipliers, *Foundations and Trends in Machine Learning*, **3** (2011), 1–122.

- [7] A. Buades, B. Coll and J. M. Morel, [A review of image denoising algorithms, with a new one](#), in *Multiscale Model. Simul.*, **2** (2005), 490–530.
- [8] J.-F. Cai, S. Osher and Z. Shen, [Split bregman methods and frame based image restoration](#), *Multiscale Modeling and Simulation*, **8** (2009), 337–369.
- [9] J.-F. Cai, H. Ji, Z. Shen and G. B. Ye, [Data-driven tight frame construction and image denoising](#), *Applied and Computational Harmonic Analysis*, **37** (2014), 89–105.
- [10] E. J. Candes, Y. C. Eldar, et al., *Compressed Sensing With Coherent and Redundant Dictionaries*, 2010.
- [11] E. J. Candes, J. Romberg and T. Tao, [Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information](#), *IEEE Transactions on Information Theory*, **52** (2006), 489–509.
- [12] A. Chambolle and T. Pock, [A first-order primal-dual algorithm for convex problems with applications to imaging](#), *Journal of Mathematical Imaging and Vision*, **40** (2011), 120–145.
- [13] H. Chen, Y. Zhang, M. K. Kalra, F. Lin, Y. Chen, P. Liao, J. Zhou and G. Wang, [Low-dose CT with a residual encoder-decoder convolutional neural network](#), *IEEE Transactions on Medical Imaging*, **36** (2017), 2524–2535.
- [14] K. Dabov, A. Foi, V. Katkovnik and K. Egiazarian, [Image denoising by sparse 3-d transform-domain collaborative filtering](#), *IEEE Transactions on Image Processing*, **16** (2007), 2080–2095.
- [15] A. Dabovolski, K. J. Batenburg and J. Sijbers, [Dynamic angle selection in x-ray computed tomography](#), *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, **324** (2014), 17–24.
- [16] I. Daubechies, *Ten Lectures on Wavelets*, CBMS-NSF Regional Conference Series in Applied Mathematics, 61. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1992.
- [17] B. Dong, J. Li and Z. Shen, [X-ray CT image reconstruction via wavelet frame based regularization and radon domain inpainting](#), *Journal of Scientific Computing*, **54** (2013), 333–349.
- [18] B. Dong, Z. Shen, et al., [Mra based wavelet frames and applications](#), *IAS Lecture Notes Series, Summer Program on “The Mathematics of Image Processing”, Park City Mathematics Institute*, **19** (2010), 9–158.
- [19] D. Donoho, [Compressed sensing](#), *IEEE Transactions on Information Theory*, **52** (2006), 1289–1306.
- [20] J. M. Ede, [Adaptive partial scanning transmission electron microscopy with reinforcement learning](#), preprint, [arXiv:2004.02786](https://arxiv.org/abs/2004.02786).
- [21] M. Elad and M. Aharon, [Image denoising via sparse and redundant representations over learned dictionaries](#), *IEEE Transactions on Image processing*, **15** (2006), 3736–3745.
- [22] E. Esser, X. Zhang, et al., [A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science](#), *SIAM Journal on Imaging Sciences*, **3** (2010), 1015–1046.
- [23] M. Gies, W. A. Kalender, H. Wolf and C. Suess, [Dose reduction in CT by anatomically adapted tube current modulation. i. Simulation studies](#), *Medical Physics*, **26** (1999), 2235–2247.
- [24] G. D. Godaliyadda, M. A. Uchic, D. H. Ye, M. A. Goeber, G. T. Buzzard and C. A. Bouman, [A supervised learning approach for dynamic sampling](#), *S&T Imaging. International Society for Optics and Photonics*, 2016.
- [25] G. M. D. P. Godaliyadda, D. H. Ye, M. D. Uchic, M. A. Goeber, G. T. Buzzard and C. A. Bouman, [A framework for dynamic image sampling based on supervised learning \(slads\)](#), *IEEE Trans. Comput. Imaging*, **4** (2018), 1–16.
- [26] T. Goldstein and S. Osher, [The split bregman method for  \$l\_1\$ -regularized problems](#), *SIAM Journal Imaging Sciences*, **2** (2009), 323–343.
- [27] T. Goldstein and S. Osher, [The split bregman method for  \$l\_1\$ -regularized problems](#), *SIAM Journal on Imaging Sciences*, **2** (2009), 323–343.
- [28] R. Gordon, R. Benderab and G. T. Herman, *Algebraic Reconstruction Techniques (ART) for Three-Dimensional Electron Microscopy and X-ray Photography*, Journal of Theoretical Biology, 1970.
- [29] S. Gu, L. Zhang, W. Zuo and X. Feng, [Weighted nuclear norm minimization with application to image denoising](#), in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (2014), 2862–2869.
- [30] A. Halimi, P. Ciuciu, A. Mccarthy, S. McLaughlin and G. Buller, [Fast adaptive scene sampling for single-photon 3d lidar images](#), *IEEE CAMSAP 2019 - International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, 2019.

- [31] S. Ji, Y. Xue and L. Carin, *Bayesian compressive sensing*, *IEEE Transactions on Signal Processing*, **56** (2008), 2346–2356.
- [32] K. H. Jin, M. T. McCann, E. Froustey and M. Unser, *Deep convolutional neural network for inverse problems in imaging*, *IEEE Transactions on Image Processing*, **26** (2017), 4509–4522.
- [33] W. A. Kalender, H. Wolf and C. Suess, Dose reduction in CT by anatomically adapted tube current modulation. ii. Phantom measurements, *Medical Physics*, **26** (1999), 2248–2253.
- [34] E. Kang, J. Min and J. C. Ye, *A deep conversational neural network using directional wavelets for low-dose x-ray ct reconstruction*, *Medical Physics*, **44** (2017), e360–e375.
- [35] A. Katsevich, *Theoretically exact filtered backprojection-type inversion algorithm for spiral CT*, *SIAM Journal on Applied Mathematics*, **62** (2002), 2012–2026.
- [36] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980), 2014.
- [37] W. Kool, H. V. Hoof and M. Welling, Attention, learn to solve routing problems!, preprint, [arXiv:1803.08475](https://arxiv.org/abs/1803.08475), 2018.
- [38] Y. Li, Deep Reinforcement Learning: An overview, [arXiv:1701.07274](https://arxiv.org/abs/1701.07274), 2017.
- [39] L. Ly and Y.-H. R. Tsai, Autonomous exploration, reconstruction, and surveillance of 3d environments aided by deep learning, [arXiv:1809.06025](https://arxiv.org/abs/1809.06025), 2018.
- [40] S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, Inc., San Diego, CA, 1998.
- [41] M. T. McCann, K. H. Jin and M. Unser, Convolutional neural networks for inverse problems in imaging: A review, *IEEE Signal Processing Magazine*, **34** (2017), 85–95.
- [42] C. McCollough, *Tu-fg-207a-04: Overview of the low dose ct grand challenge*, *Medical Physics*, **43** (2016), 3760–3760.
- [43] A. Mittal, A. Dhawan, S. Manchanda, S. Medya, S. Ranu and A. Singh, Learning heuristics over large graphs via deep reinforcement learning, preprint, [arXiv:1903.03332](https://arxiv.org/abs/1903.03332), 2019.
- [44] V. Mnih, K. Kavukcuoglu and D. Silver, Human-level control through deep reinforcement learning, *Nature*, **518** (2015).
- [45] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al, Human-level control through deep reinforcement learning, *Nature*, **518** (2015), 529–533.
- [46] K. A. Mohan, S. V. Venkatakrishnan, E. B. Gulsoy, J. W. Gibbs, X. Xiao, M. D. Graef, P. W. Voorhees and C. A. Bouman, *Timbir: A method for time-space reconstruction from interlaced views*, *IEEE Transactions on Computational Imaging*, **1** (2015), 96–111.
- [47] E. Monier, N. Brun, T. Oberlin, X. Li, M. Tenc and N. Dobigeon, Fast reconstruction of atomic-scale stem-eels images from sparse sampling, *Ultramicroscopy*, 2020.
- [48] K. Mueller, Selection of optimal views for computed tomography reconstruction, *Patent WO*, Jan, **28** (2011).
- [49] K. Mueller, R. Yagel and J. J. Wheller, Anti-aliased three-dimensional cone-beam reconstruction of low-contrast objects with algebraic methods, *IEEE Transactions On Medical Imaging*, **6** (1999), 519–537.
- [50] R. Ohbuchi and M. Aono, *Quasi-Monte Carlo Rendering With Adaptive Sampling*, 1996.
- [51] S. Osher, Z. Shi and W. Zhu, *Low dimensional manifold model for image processing*, *SIAM Journal on Imaging Sciences*, **10** (2017), 1669–1690.
- [52] J. Park, J. Jung, A. P. Gupta, J. Soh, C. Jeong, J. Ahn, S. Cho, K. -H. Yoon, D. Kim, M. Mativenga, et al. Multi-beam x-ray source based on carbon nanotube emitters for tomosynthesis system, in *Medical Imaging 2020: Physics of Medical Imaging*, International Society for Optics and Photonics, **11312** (2020), 113122.
- [53] G. Placidi, M. Alecci and A. Sotgiu, Theory of adaptive acquisition method for image reconstruction from projections and application to epr image, *Journal of Magnetic Resonance*, (1995), 50–57.
- [54] S. RL, Fast calculation of the exact radiological path for a three-dimensional CT array, *Medical Physics*, **2** (1985), 252–5.
- [55] A. Ron and Z. Shen, *Affine systems in  $L_2(\mathbb{R}^d)$ : The analysis of the analysis operator*, *Journal of Functional Analysis*, **148** (1997), 408–447.
- [56] L. I. Rudin, S. Osher and E. Fatemi, *Nonlinear total variation based noise removal algorithms*, *Physical D: Nonlinear Phenomena*, **60** (1992), 259–268.
- [57] J. Schulman and F. Wolski, Proximal policy optimization algorithms, [arXiv:1707.06347v2](https://arxiv.org/abs/1707.06347v2), 2017.
- [58] M. W. Seeger and H. Nickisch, Compressed sensing and bayesian experimental design, in *Proceedings of the 25th International Conference on Machine Learning*, (2008), 912–919.

- [59] B. Settles, *Active Learning Literature Survey*, Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.
- [60] C. Shen, Y. Gonzalez, L. Chen, S. B. Jiang and X. Jia, Intelligent parameter tuning in optimization-based iterative CT reconstruction via deep reinforcement learning, *IEEE Transactions on Medical Imaging*, **37** (2018), 1430–1439.
- [61] E. Y. Sidky and X. Pan, Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization, *Physics in medicine and biology*, **47** (2008), 4777.
- [62] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra and M. Riedmiller, Deterministic policy gradient algorithms, In *International Conference on Machine Learning*, (2014), 387–395.
- [63] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, Cambridge, MA, 2018.
- [64] R. S. Sutton, D. A. McAllester, S. P. Singh and Y. Mansour, Policy gradient methods for reinforcement learning with function approximation, In *Advances in Neural Information Processing Systems*, (2000), 1057–1063.
- [65] C. Tai and E. Weinan, Multiscale adaptive representation of signals: I. The basic framework, *The Journal of Machine Learning Research*, **17** (2016), 4875–4912.
- [66] G. Wang, A perspective on deep imaging, *IEEE Access*, **4** (2016), 8914–8924.
- [67] Z. Wang and G. R. Arce, *Variable density compressed image sampling*, *Image Processing, IEEE Transactions*, **19** (2010), 264–270.
- [68] G. Wang, M. Kalra and C. G. Orton, Machine learning will transform radiology significantly within the next 5 years, *Medical Physics*, **44** (2017), 2041–2044.
- [69] G. Wang, J. Chu Ye, K. Mueller and J. A. Fessler, Image reconstruction is a new frontier of machine learning, *IEEE Transactions on Medical Imaging*, **37** (2018), 1289–1296.
- [70] G. Wang and H. Yu, A scheme for multisource interior tomography, *Medical physics*, **36** (2009), 3575–3581.
- [71] C. J. Watkins and P. Dayan, Q-learning, *Machine learning*, **8** (1992), 279–292.
- [72] Q. Yang, P. Yan, Y. Zhang, H. Yu, Y. Shi, X. Mou, M. K. Kalra, Y. Zhang, L. Sun and G. Wang, Low-dose CT image denoising using a generative adversarial network with wasserstein distance and perceptual loss, *IEEE Transactions on Medical Imaging*, **37** (2018) 1348–1357.
- [73] L. Yu, M. Shiung, D. Jondal and C. H. McCollough, Development and validation of a practical lower-dose-simulation tool for optimizing computed tomography scan protocols, *Journal of Computer Assisted Tomography*, **36** (20124), 477–487.
- [74] J. Zhang, G. Yang, Y. Lee, Y. Cheng, B. Gao, Q. Qiu, J. Lu and O. Zhou, A multi-beam x-ray imaging system based on carbon nanotube field emitters, in *Medical Imaging 2006: Physics of Medical Imaging*, International Society for Optics and Photonics, **6142** (2006), 614204.
- [75] S. Zhang, Z. Song, G. D. P. Godaliyadda, D. H. Ye, A. U. Chowdhury, A. Sengupta, G. T. Buzzard, C. A. Bouman and G. J. Simpson, Dynamic sparse sampling for confocal raman microscopy, *Analytical Chemistry*, **90** (2018), 4461–4469.
- [76] Y. Zhang, G. M. D. Godaliyadda, N. Ferrier, E. B. Gulsoy, C. A. Bouman and C. Phatak, *Slads-Net: Supervised Learning Approach for Dynamic Sampling Using Deep Neural Networks*, Electronic Imaging, Computational Imaging XVI, 2018.
- [77] H. -M. Zhang and B. Dong, *A review on deep learning in medical image reconstruction*, *Journal of the Operations Research Society of China*, **8** (2020) 311–340.
- [78] Z. Zhang, X. Liang, X. Dong, Y. Xie and G. Cao, A sparse-view CT reconstruction method based on combination of DenseNet and deconvolution, *IEEE Transactions on Medical Imaging*, **37** (2018), 1407–1417.
- [79] M. Zhu and T. Chan, An efficient primal-dual hybrid gradient algorithm for total variation image restoration, *UCLA CAM Report*, **34** (2008), 8–34.

Received October 2020; revised April 2021. Early access July 2021.

*E-mail address:* [zjshen@pku.edu.cn](mailto:zjshen@pku.edu.cn)

*E-mail address:* [yufeiw2@andrew.cmu.edu](mailto:yufeiw2@andrew.cmu.edu)

*E-mail address:* [dwu6@mgh.harvard.edu](mailto:dwu6@mgh.harvard.edu)

*E-mail address:* [xuyang@math.ucsb.edu](mailto:xuyang@math.ucsb.edu)

*E-mail address:* [dongbin@math.pku.edu.cn](mailto:dongbin@math.pku.edu.cn)