

Learning To Stop While Learning To Predict

Xinshi Chen¹ Hanjun Dai² Yu Li³ Xin Gao³ Le Song^{1,4}

Abstract

There is a recent surge of interest in designing deep architectures based on the update steps in traditional algorithms, or learning neural networks to improve and replace traditional algorithms. While traditional algorithms have certain stopping criteria for outputting results at different iterations, many algorithm-inspired deep models are restricted to a “fixed-depth” for all inputs. Similar to algorithms, the optimal depth of a deep architecture may be different for different input instances, either to avoid “over-thinking”, or because we want to compute less for operations converged already. In this paper, we tackle this varying depth problem using a steerable architecture, where a feed-forward deep model and a variational stopping policy are learned together to sequentially determine the optimal number of layers for each input instance. Training such architecture is very challenging. We provide a variational Bayes perspective and design a novel and effective training procedure which decomposes the task into an oracle model learning stage and an imitation stage. Experimentally, we show that the learned deep model along with the stopping policy improves the performances on a diverse set of tasks, including learning sparse recovery, few-shot meta learning, and computer vision tasks.

1. Introduction

Recently, researchers are increasingly interested in the connections between deep learning models and traditional algorithms: deep learning models are viewed as parameterized algorithms that operate on each input instance iteratively, and traditional algorithms are used as templates for designing deep learning architectures. While an important con-

¹Georgia Institute of Technology, USA ²Google Research, USA ³King Abdullah University of Science and Technology, Saudi Arabia ⁴Ant Financial, China. Correspondence to: Xinshi Chen <xinshi.chen@gatech.edu>, Le Song <lsong@cc.gatech.edu>.

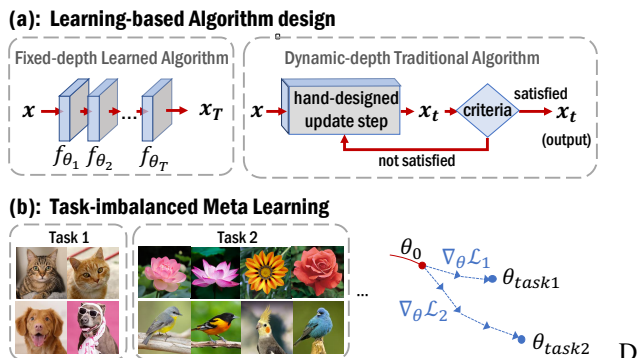


Figure 1. Motivation for learning to stop.

cept in traditional algorithms is the stopping criteria for outputting the result, which can be either a *convergence* condition or an *early stopping* rule, such stopping criteria has been more or less ignored in algorithm-inspired deep learning models. A “fixed-depth” deep model is used to operate on all problem instances (Fig. 1 (a)). Intuitively, for deep learning models, the optimal depth (or the optimal number of steps to operate on an input) can also be different for different input instances, either because we want to compute less for operations converged already, or we want to generalize better by avoiding “over-thinking”. Such motivation aligns well with both the cognitive science literature (?) and many examples below:

- In learning to optimize (Andrychowicz et al., 2016; Li & Malik, 2016), neural networks are used as the optimizer to minimize some loss function. Depending on the initialization and the objective function, an optimizer should converge in different number of steps;
- In learning to solve statistical inverse problems such as compressed sensing (Chen et al., 2018; Liu et al., 2019), inverse covariance estimation (Shrivastava et al., 2020), and image denoising (Zhang et al., 2019), deep models are learned to directly predict the recovery results. In traditional algorithms, problem-dependent early stopping rules are widely used to achieve regularization for a variance-bias trade-off. Deep learning models for solving such problems maybe also achieve a better recovery accuracy by allowing instance-specific computation steps;
- In meta learning, MAML (Finn et al., 2017) used an unrolled and parametrized algorithm to adapt a common




Figure 2. Two-component model: learning to predict (blue) while learning to stopping (green).

parameter to a new task. However, depending on the *similarity* of the new task to the old tasks, or, in a more realistic *task-imbalanced* setting where different tasks have different numbers of data points (Fig. 1 (b)), a task-specific number of adaptation steps is more favorable to avoid under or over adaptation.

To address the varying depth problem, we propose to learn a steerable architecture, where a shared feed-forward model for normal prediction and an additional stopping policy are learned together to sequentially determine the optimal number of layers for each input instance. In our framework, the model consists of (see Fig. 2)

- A **feed-forward or recurrent mapping** \mathcal{F}_θ , which transforms the input \mathbf{x} to generate a path of features (or states) $\mathbf{x}_1, \dots, \mathbf{x}_T$; and
- A **stopping policy** $\pi_\phi : (\mathbf{x}, \mathbf{x}_t) \mapsto \pi_t \in [0, 1]$, which sequentially observes the states and then determines the probability of stopping the computation of \mathcal{F}_θ at layer t .

These two components allow us to sequentially **predict** the next targeted state while at the same time determining when to **stop**. In this paper, we propose a single objective function for learning both θ and ϕ , and we interpret it from the perspective of variational Bayes, where the stopping time t is viewed as a latent variable conditioned on the input \mathbf{x} . With this interpretation, learning θ corresponds to maximizing the marginal likelihood, and learning ϕ corresponds to the inference step for the latent variable, where a variational distribution $q_\phi(t)$ is optimized to approximate the posterior. A natural algorithm for solving this problem could be the Expectation-Maximization (EM) algorithm, which can be very hard to train and inefficient.

How to learn θ and ϕ effectively and efficiently? We propose a principled and effective training procedure, where we decompose the task into an oracle model learning stage and an imitation learning stage (Fig. 3). More specifically,

- During the oracle model learning stage, we utilize a closed-form oracle stopping distribution $q^*|\theta$ which can leverage label information not available at testing time.
- In the imitation learning stage, we use a sequential policy π_ϕ to mimic the behavior of the oracle policy obtained in the first stage. The sequential policy does not have access to the label so that it can be used during testing phase.

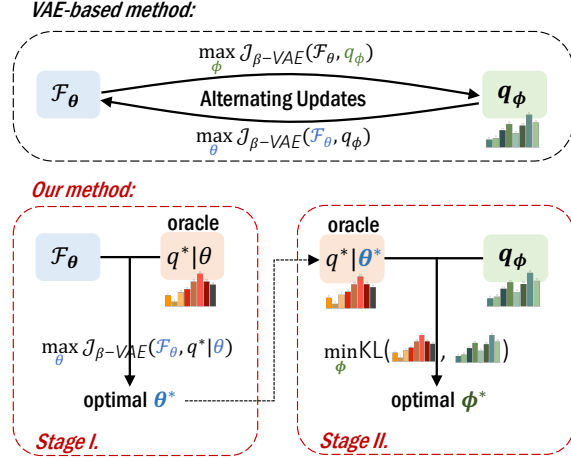


Figure 3. Two-stage training framework.

This procedure provides us a very good initial predictive model and a stopping policy. We can either directly use these learned models, or plug them back to the variational EM framework and reiterate to further optimize both together.

Our proposed learning to stop method is a generic framework that can be applied to a diverse range of applications. To summarize, our contribution in this paper includes:

1. a variational Bayes perspective to understand the proposed model for learning both the predictive model and the stopping policy together;
2. a principled and efficient algorithm for jointly learning the predictive model and the stopping policy; and the relation of this algorithm to reinforcement learning;
3. promising experiments on various tasks including learning to solve sparse recovery problems, task-imbalanced few-shot meta learning, and computer vision tasks, where we demonstrate the effectiveness of our method in terms of both the prediction accuracy and inference efficiency.

2. Related Works

Unrolled algorithm. A line of recent works unfold and truncate iterative algorithms to design neural architectures. These algorithm-based deep models can be used to automatically learn a better algorithm from data. This idea has been demonstrated in different problems including sparse signal recovery (Gregor & LeCun, 2010; Sun et al., 2016; Borgarding et al., 2017; Metzler et al., 2017; Zhang & Ghanem, 2018; Chen et al., 2018; Liu et al., 2019), sparse inverse covariance estimation (Shrivastava et al., 2020), sequential Bayesian inference (Chen et al., 2019), parameter learning in graphical models (Domke, 2011), non-negative matrix factorization (Yakar et al., 2013), etc. Unrolled algorithm based deep module has also be used for structured prediction (Belanger et al., 2017; Ingraham et al., 2019; Chen et al.,

2020). Before the training phase, all these works need to assign a fixed number of iterations that is used for every input instance regardless of their varying difficulty level. Our proposed method is orthogonal and complementary to all these works, by taking the variety of the input instances into account via adaptive stopping time.

Meta learning. Optimization-based meta learning techniques are widely applied for solving challenging few-shot learning problems (Ravi & Larochelle, 2017; Finn et al., 2017; Li et al., 2017). Several recent advances proposed task-adaptive meta-learning models which incorporate task-specific parameters (Qiao et al., 2018; Lee & Choi, 2018; Na et al., 2020) or task-dependent metric scaling (Oreshkin et al., 2018). In parallel with these task-adaptive methods, we propose a task-specific number of adaptation steps and demonstrate the effectiveness of this simple modification under the task-imbalanced scenarios.

Other adaptive-depth deep models. In image recognition, ‘early exits’ is proposed mainly aimed at improving the computation efficiency during the inference phase (Teerapittayanon et al., 2016; Zamir et al., 2017; Huang et al., 2018), but these methods are based on specific architectures. Kaya et al. (2019) proposed to avoid “over-thinking” by early stopping. However, the same as all the other ‘early exits’ models, some heuristic policies are adopted to choose the output layer by confidence scores of internal classifiers. Also, their algorithms for training the feed-forward model \mathcal{F}_θ do not take into account the effect of the stopping policy.

Optimal stopping. In optimal control literature, optimal stopping is a problem of choosing a time to take a given action based on sequentially observed random variables in order to maximize an expected payoff (Shiryaev, 2007). When a policy for controlling the evolution of random variables (corresponds to the output of \mathcal{F}_θ) is also involved, it is called a “mixed control” problem, which is highly related to our work. Existing works in this area find the optimal controls by solving the Hamilton-Jacobi-Bellman (HJB) equation, which is theoretically grounded (Pham, 1998; Ceci & Basan, 2004; Dumitrescu et al., 2018). However, they focus on stochastic differential equation based model and the proposed algorithms suffer from the curse of dimensionality problem. Becker et al. (2019) use DL to learn the optimal stopping policy, but the learning of θ is not considered. Besides, Becker et al. (2019) use reinforcement learning (RL) to solve the problem. In Section 4, we will discuss how our variational inference formulation is related to RL.

3. Problem Formulation

In this section, we will introduce how we model the stopping policy together with the predictive deep model, define the joint optimization objective, and interpret this framework

from a variational Bayes perspective.

3.1. Steerable Model

The predictive model, \mathcal{F}_θ , is a typical T -layer deep model that generates a path of embeddings $(\mathbf{x}_1, \dots, \mathbf{x}_T)$ through:

$$\textbf{Predictive model: } \mathbf{x}_t = f_{\theta_t}(\mathbf{x}_{t-1}), \text{ for } t=1, \dots, T \quad (1)$$

where the initial \mathbf{x}_0 is determined by the input \mathbf{x} . We denote it by $\mathcal{F}_\theta = \{f_{\theta_1}, \dots, f_{\theta_T}\}$ where $\theta \in \Theta$ are the parameters. Standard supervised learning methods learn θ by optimizing an objective estimated on the final state \mathbf{x}_T . In our model, the operations in Eq. 1 can be stopped earlier, and for different input instance \mathbf{x} , the stopping time t can be different.

Our stopping policy, π_ϕ , determines whether to stop at t -th step after observing the input \mathbf{x} and its first t states $\mathbf{x}_{1:t}$ transformed by \mathcal{F}_θ . If we assume the Markov property, then π_ϕ only needs to observe the most recent state \mathbf{x}_t . In this paper, we only input \mathbf{x} and \mathbf{x}_t to π_ϕ at each step t , but it is trivial to generalize it to $\pi_\phi(\mathbf{x}, \mathbf{x}_{1:t})$. More precisely, π_ϕ is defined as a randomized policy as follows:

$$\textbf{Stopping policy: } \pi_t = \pi_\phi(\mathbf{x}, \mathbf{x}_t), \text{ for } t=1, \dots, T-1 \quad (2)$$

where $\pi_t \in [0, 1]$ is the probability of stopping. We abuse the notation π to both represent the parametrized policy and also the probability mass.

This stopping policy sequentially makes a decision whenever a new state \mathbf{x}_t is observed. Conditioned on the states observed until step t , whether to stop before t is independent on states after t . Therefore, once it decides to stop at t , the remaining computations can be saved, which is a favorable property when the inference time is a concern, or for some optimal stopping problems such as option trading where getting back to earlier states is not allowed.

3.2. From Sequential Policy To Stop Time Distribution

The stopping policy π_ϕ makes sequential actions based on the observations, where $\pi_t := \pi_\phi(\mathbf{x}, \mathbf{x}_t)$ is the probability of stopping when \mathbf{x}_t is observed. These sequential actions π_1, \dots, π_{T-1} jointly determines the random time t at which the stop occurs. Induced by π_ϕ , the probability mass function of the stop time t , denoted as q_ϕ , can be computed by

Variational stop time distribution:

$$\begin{cases} q_\phi(t) = \pi_t \prod_{\tau=1}^{t-1} (1 - \pi_\tau) & \text{if } t < T, \\ q_\phi(T) = \prod_{\tau=1}^{T-1} (1 - \pi_\tau) & \text{else.} \end{cases} \quad (3)$$

In this equation, the product $\prod_{\tau=1}^{t-1} (1 - \pi_\tau)$ indicates the probability of ‘not stopped before t ’, which is the survival probability. Multiply this survival probability with π_t , we have the stop time distribution $q_\phi(t)$. For the last time step

T , the stop probability $q_\phi(T)$ simply equals to the survival probability at T , which means if the process is ‘not stopped before T ’, then it must stop at T .

Note that we only use π_ϕ in our model to sequentially determine whether to stop. However, we use the induced probability mass q_ϕ to help design the training objective and also the algorithm.

3.3. Optimization Objective

Note that the stop time t is a discrete random variable with distribution determined by $q_\phi(t)$. Given the observed label \mathbf{y} of an input \mathbf{x} , the loss of the predictive model stopped at position t can be computed as $\ell(\mathbf{y}, \mathbf{x}_t; \theta)$ where $\ell(\cdot)$ is a loss function. Taking into account all possible stopping positions, we will be interested in the loss in expectation over t ,

$$\mathcal{L}(\theta, q_\phi; \mathbf{x}, \mathbf{y}) := \mathbb{E}_{t \sim q_\phi} \ell(\mathbf{y}, \mathbf{x}_t; \theta) - \beta H(q_\phi), \quad (4)$$

where $H(q_\phi) := -\sum_t q_\phi(t) \log q_\phi(t)$ is an entropy regularization and β is the regularization coefficient. Given a data set $\mathcal{D} = \{(\mathbf{x}, \mathbf{y})\}$, the parameters of the predictive model and the stopping policy can be estimated by

$$\min_{\theta, \phi} \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \mathcal{L}(\theta, q_\phi; \mathbf{x}, \mathbf{y}). \quad (5)$$

To better interpret the model and objective, in the following, we will make a connection from the perspective of variational Bayes, and how the objective function defined in Eq. 4 is equivalent to the β -VAE objective.

3.4. Variational Bayes Perspective

In the Bayes’ framework, a probabilistic model typically consists of prior, likelihood function and posterior of the latent variable. We find the correspondence between our model and a probabilistic model as follows (also see Table 1)

- we view the adaptive stopping time t as a *latent variable* which is unobserved;
- The conditional *prior* $p(t|\mathbf{x})$ of t is a uniform distribution over all the layers in this paper. However, if one wants to reduce the computation cost and penalize the stopping decisions at deeper layers, a prior with smaller probability on deeper layers can be defined to regularize the results;
- The *likelihood* function $p_\theta(\mathbf{y}|t, \mathbf{x})$ of the observed label \mathbf{y} is controlled by θ , since \mathcal{F}_θ determines the states \mathbf{x}_t ;
- The *posterior* distribution over the stopping time t can be computed by Bayes’ rule $p_\theta(t|\mathbf{y}, \mathbf{x}) \propto p_\theta(\mathbf{y}|t, \mathbf{x})p(t|\mathbf{x})$, but it requires the observation of the label \mathbf{y} , which is infeasible during testing phase.

In this probabilistic model, we need to learn θ to better fit the observed data and learn a variational distribution q_ϕ over t that only takes \mathbf{x} and the transformed internal states as inputs to approximate the true posterior.

Table 1. Corresponds between our model and Bayes’ model.

stop time t	latent variable
label \mathbf{y}	observation
loss $\ell(\mathbf{y}, \mathbf{x}_t; \theta)$	likelihood $p_\theta(\mathbf{y} t, \mathbf{x})$
stop time distribution q_ϕ	posterior $p_\theta(t \mathbf{y}, \mathbf{x})$
regularization	prior $p(t \mathbf{x})$

More specifically, the parameters in the likelihood function and the variational posterior can be optimized using the variational autoencoder (VAE) framework (Kingma & Welling, 2013). Here we consider a generalized version called β -VAE (Higgins et al., 2017), and obtain the optimization objective for data point (\mathbf{x}, \mathbf{y})

$$\mathcal{J}_{\beta\text{-VAE}}(\theta, q_\phi; \mathbf{x}, \mathbf{y}) := \mathbb{E}_{q_\phi} \log p_\theta(\mathbf{y}|t, \mathbf{x}) - \beta \text{KL}(q_\phi(t) \| p(t|\mathbf{x})), \quad (6)$$

where $\text{KL}(\cdot \| \cdot)$ is the KL divergence. When $\beta = 1$, it becomes the original VAE objective, i.e., the evidence lower bound (ELBO). Now we are ready to present the equivalence relation between the β -VAE objective and the loss defined in Eq. 4. See Appendix A.1 for the proof.

Lemma 1. *Under assumptions: (i) the loss function ℓ in Eq. 4 is defined as the negative log-likelihood (NLL), i.e.,*

$$\ell(\mathbf{y}, \mathbf{x}_t; \theta) := -\log p_\theta(\mathbf{y}|t, \mathbf{x});$$

(ii) the prior $p(t|\mathbf{x})$ is a uniform distribution over t ;

then minimizing the loss \mathcal{L} in Eq. 4 is equivalent to maximizing the β -VAE objective $\mathcal{J}_{\beta\text{-VAE}}$ in Eq. 6.

For classification problems, the cross-entropy loss is aligned with NLL. For regression problems with mean squared error (MSE) loss, we can define the likelihood as $p_\theta(\mathbf{y}|t, \mathbf{x}) \sim \mathcal{N}(\mathbf{x}_t, I)$. Then the NLL of this Gaussian distribution is $-\log p_\theta(\mathbf{y}|t, \mathbf{x}) = \frac{1}{2} \|\mathbf{y} - \mathbf{x}_t\|_2^2 + C$, which is equivalent to MSE loss. More generally, we can always define $p_\theta(\mathbf{y}|t, \mathbf{x}) \propto \exp(-\ell(\mathbf{y}, \mathbf{x}_t; \theta))$.

This VAE view allows us to design a two-step procedure to effectively learn θ and ϕ in the predictive model and stopping policy, which is presented in the next section.

4. Effective Training Algorithm

VAE-based methods perform optimization steps over θ (M step for learning) and ϕ (E step for inference) alternatively until convergence, which has two limitations in our case:

- The alternating training can be slow to converge and requires tuning the training scheduling;
- The inference step for learning q_ϕ may have the mode collapse problem, which in this case means q_ϕ only captures the time step t with highest averaged frequency.

To overcome these limitations, we design a training procedure followed by an optional fine-tuning stage using the variational lower bound in Eq. 6. More specifically,

Stage I. Find the optimal θ by maximizing the conditional marginal likelihood when the stop time distribution follows an oracle distribution q_θ^* .

Stage II. Fix the optimal θ learned in Stage I, and only learn the distribution q_ϕ to mimic the oracle by minimizing the KL divergence between q_ϕ and q_θ^* .

Stage III. (Optional) Fine-tune θ and ϕ jointly towards the joint objective in Eq. 6.

The overall algorithm steps are summarized in Algorithm 1. In the following sections, we will focus on the derivation of the first two training steps. Then we will discuss several methods to further improve the memory and computation efficiency for training.

4.1. Oracle Stop Time Distribution

We first give the definition of the oracle stop time distribution q_θ^* . For each fixed θ , we can find a closed-form solution for the optimal q_θ^* that optimizes the joint objective.

$$q_\theta^*(\cdot|\mathbf{y}, \mathbf{x}) := \arg \max_{q \in \Delta^{T-1}} \mathcal{J}_{\beta\text{-VAE}}(\theta, \mathbf{q}; \mathbf{x}, \mathbf{y})$$

Alternatively, $q_\theta^*(\cdot|\mathbf{y}, \mathbf{x}) = \arg \min_{q \in \Delta^{T-1}} \mathcal{L}(\theta, \mathbf{q}; \mathbf{x}, \mathbf{y})$. Under the mild assumptions in Lemma 1, these two optimizations lead to the same optimal oracle distribution.

Oracle stop time distribution:

$$q_\theta^*(t|\mathbf{y}, \mathbf{x}) = \frac{p_\theta(\mathbf{y}|t, \mathbf{x})^{\frac{1}{\beta}}}{\sum_{t=1}^T p_\theta(\mathbf{y}|t, \mathbf{x})^{\frac{1}{\beta}}} \quad (7)$$

$$= \frac{\exp(-\frac{1}{\beta}\ell(\mathbf{y}, \mathbf{x}_t; \theta))}{\sum_{t=1}^T \exp(-\frac{1}{\beta}\ell(\mathbf{y}, \mathbf{x}_t; \theta))} \quad (8)$$

This closed-form solution makes it clear that the oracle picks a step t according to the smallest loss or largest likelihood with an exploration coefficient β .

Remark: When $\beta = 1$, q_θ^* is the same as the posterior distribution $p_\theta(t|\mathbf{y}, \mathbf{x}) \propto p_\theta(\mathbf{y}|t, \mathbf{x})p(t|\mathbf{x})$.

Note that there are no new parameters in the oracle distribution. Instead, it depends on the parameters θ in the predictive model. Overall, the oracle q_θ^* is a function of θ , t , \mathbf{y} and \mathbf{x} that has a closed-form. Next, we will introduce how we use this oracle in the first two training stages.

4.2. Stage I. Predictive Model Learning

In Stage I, we optimize the parameters θ in the predictive model by taking into account the oracle stop distribution q_θ^* . This step corresponds to the M step for learning θ , by maximizing the marginal likelihood. The difference with

Algorithm 1 Overall Algorithm

Randomly initialized θ and ϕ .

For $itr = 1$ to $\#iterations$ **do** ▷ Stage I.

Sample a batch of data points $\mathcal{B} \sim \mathcal{D}$.

Take an optimization step to update θ towards the marginal likelihood function defined in Eq. 9.

For $itr = 1$ to $\#iterations$ **do** ▷ Stage II.

Sample a batch of data points $\mathcal{B} \sim \mathcal{D}$.

Take an optimization step to update ϕ towards the reverse KL divergence defined in Eq. 10.

For $itr = 1$ to $\#iterations$ **do** ▷ Optional Step

Sample a batch of data points $\mathcal{B} \sim \mathcal{D}$.

Update both θ and ϕ towards β -VAE objective in Eq. 6.

return θ, ϕ

the normal M step is that here q_ϕ is replaced by the oracle q_θ^* that gives the optimal stopping distribution so that the marginal likelihood is independent on ϕ . More precisely, stage I finds the optimum of:

$$\max_{\theta} \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \mathcal{J}_{\beta\text{-VAE}}(\theta, q_\theta^*; \mathbf{x}, \mathbf{y}), \quad (9)$$

where the β -VAE objective here is $\mathcal{J}_{\beta\text{-VAE}}(\theta, q_\theta^*; \mathbf{x}, \mathbf{y}) = \sum_{t=1}^T q_\theta^*(t|\mathbf{y}, \mathbf{x}) \log p_\theta(\mathbf{y}|t, \mathbf{x}) - \beta \text{KL}(q_\theta^*(t)||p(t|\mathbf{x}))$.

Remark. For experiments that require higher memory costs (e.g., MAML), we prefer to drop the entropy term, $\beta \text{KL}(q_\theta^*(t)||p(t|\mathbf{x}))$, in the objective, so that stochastic sampling can be applicable to reduce the memory cost (see more details of the efficient training algorithm in Sec. 4.5). Since we can adjust β in the oracle q^* to control the concentration level of the distribution, dropping the entropy term in the objective in stage I does not affect much the performance.

Since q_θ^* has a differentiable closed-form expression in terms of θ , \mathbf{x} , \mathbf{y} and t , the gradient can also propagate through q_θ^* , which is also different from the normal M step.

To summarize, in Stage I., we learn the predictive model parameter θ , by assuming that the stop time always follows the best stopping distribution that depends on θ . In this case, the learning of θ has already taken into account the effect of the data-specific stop time.

However, we note that the oracle q_θ^* is not in the form of sequential actions as in Eq. 2 and it requires the access to the true label \mathbf{y} , so it can not be used for testing. However, it plays an important role in obtaining a sequential policy which will be explained next.

4.3. Stage II. Imitation With Sequential Policy

In Stage II, we learn the sequential policy π_ϕ that can best mimic the oracle distribution q_θ^* , where θ is fixed to be the optimal θ learned in Stage I. The way of doing so is

to minimize the divergence between the oracle q_θ^* and the variational stop time distribution q_ϕ induced by π_ϕ (Eq. 3). There are various variational divergence minimization approaches that we can use (Nowozin et al., 2016). For example, a widely used objective for variational inference is the **reverse KL divergence**:

$$\text{KL}(q_\phi \| q_\theta^*) = \sum_{t=1}^T -q_\phi(t) \log q_\theta^*(t | \mathbf{y}, \mathbf{x}) - H(q_\phi).$$

Remark. We write $q_\phi(t)$ instead of $q_\phi(t | \mathbf{x}_{1:T}, \mathbf{x})$ for notation simplicity, but q_ϕ is dependent on \mathbf{x} and $\mathbf{x}_{1:T}$ (Eq. 3).

If we rewrite q_ϕ using π_1, \dots, π_{T-1} as defined in Eq. 3, we can find that minimizing the reverse KL is equivalent to finding the optimal policy π_ϕ in a reinforcement learning (RL) environment, where the state is \mathbf{x}_t , action $a_t \sim \pi_t := \pi_\phi(\mathbf{x}, \mathbf{x}_t)$ is a stop/continue decision, the state transition is determined by θ and a_t , and the reward is defined as

$$r(\mathbf{x}_t, a_t; \mathbf{y}) := \begin{cases} -\beta \ell(\mathbf{y}, \mathbf{x}_t; \theta) & \text{if } a_t = 0 \text{ (i.e. stop)} \\ 0 & \text{if } a_t = 1 \text{ (i.e. continue)} \end{cases}$$

where $\ell(\mathbf{y}, \mathbf{x}_t; \theta) = -\log p_\theta(\mathbf{y} | t, \mathbf{x})$. More details and also the derivation are given in Appendix A.2 to show that minimizing $\text{KL}(q_\phi \| q_\theta^*)$ is equivalent to solving the following **maximum-entropy RL**:

$$\max_{\phi} \mathbb{E}_{\pi_\phi} \sum_{t=1}^T [r(\mathbf{x}_t, a_t; \mathbf{y}) + H(\pi_t)].$$

In some related literature, optimal stopping problem is often formulated as an RL problem (Becker et al., 2019). Above we bridge the connection between our variational inference formulation and the RL-based optimal stopping literature.

Although reverse KL divergence is a widely used objective, it suffers from the mode collapse issue, which in our case may lead to a distribution q_ϕ that captures only a common stopping time t for all \mathbf{x} that on average performs the best, instead of a more spread-out stopping time. Therefore, we consider the **forward KL divergence**:

$$\text{KL}(q_\theta^* \| q_\phi) = - \sum_{t=1}^T q_\theta^*(t | \mathbf{y}, \mathbf{x}) \log q_\phi(t) - H(q_\theta^*), \quad (10)$$

which is equivalent to the **cross-entropy** loss, since the term $H(q_\theta^*)$ can be ignored as θ is fixed in this step. Experimentally, we find forward KL leads to a better performance.

4.4. The Optional Fine Tuning Stage

It is easy to see that our two-stage training procedure also has an EM flavor. However, with the oracle q_θ^* incorporated, the training of θ has already taken into account the effect of the optimal stopping distribution. Therefore, we can save a lot of alternation steps. After the two-stage training, we can fine-tune θ and ϕ jointly towards the β -VAE objective. Experimentally, we find this additional stage does not improve much the performance trained after the first two stages.

4.5. Implementation Details For Efficient Training

Since both objectives in oracle learning stage (Eq. 9) and imitation stage (Eq. 10) involve the summation over T layers, the computation and memory costs during training are higher than standard learning methods. The memory issue is especially important in meta learning. In the following, we introduce several ways of improving the training efficiency.

Fewer output channels. Instead of allowing the model to output \mathbf{x}_t at any layer, we can choose a smaller number of output channels that are evenly placed along with the layers.

Stochastic sampling in Step I. A Monte Carlo method can be used to approximate the expectation over q_θ^* in Step I. More precisely, for each (\mathbf{x}, \mathbf{y}) we can randomly sample a layer $t_s \sim q_\theta^*(t | \mathbf{y}, \mathbf{x})$ from the oracle, and only compute $\log p_\theta(\mathbf{y} | t_s, \mathbf{x})$ at t_s , instead of summing over all $t \in [T]$. Note that, in this case, the gradient will not back-propagate through $q_\theta^*(t | \mathbf{y}, \mathbf{x})$. As explained earlier in Sec 4.2, the entropy term can be dropped to reduce the memory cost.

MAP estimate in Step II. Instead of approximating the distribution q_θ^* , we can approximate the maximum a posteriori (MAP) estimate $\hat{t}(\mathbf{x}, \mathbf{y}) = \arg \max_{t \in [T]} q_\theta^*(t | \mathbf{y}, \mathbf{x})$ so that the objective for each sample is $-\log q_\theta(\hat{t}(\mathbf{x}, \mathbf{y}))$, which does not involve the summation over t . Except for efficiency, we also find this MAP estimate can lead to a higher accuracy, by encouraging the learning of q_ϕ to focus more on the sample-wise best layer.

5. Experiments

We conduct experiments on (i) learning-based algorithm for sparse recovery, (ii) few-shot meta learning, and (iii) image denoising. The comparison is in an ablation study fashion to better examine whether the stopping policy can improve the performances given the same architecture for the predictive model, and whether our training algorithm is more effective compared to the alternating EM algorithm. In the end, we also discuss our exploration of the image recognition task. Pytorch implementation of the experiments is released at <https://github.com/xinshi-chen/l2stop>.

5.1. Learning To Optimize: Sparse Recovery

We consider a sparse recovery task which aims at recovering $\mathbf{x}^* \in \mathbb{R}^n$ from its noisy linear measurements $\mathbf{b} = A\mathbf{x}^* + \epsilon$, where $A \in \mathbb{R}^{m \times n}$, $\epsilon \in \mathbb{R}^m$ is Gaussian white noise, and $m \ll n$. A popular approach is to model the problem as the LASSO formulation $\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{b} - A\mathbf{x}\|_2^2 + \rho \|\mathbf{x}\|_1$ and solves it using iterative methods such as the ISTA (Blumensath & Davies, 2008) and FISTA (Beck & Teboulle, 2009) algorithms. We choose the most popular model named Learned ISTA (LISTA) as the baseline and also as our predictive model. LISTA is a T -layer network with update steps:

$$\mathbf{x}_t = \eta_{\lambda_t} (W_t^1 \mathbf{b} + W_t^2 \mathbf{x}_{t-1}), \quad t = 1, \dots, T, \quad (11)$$

where $\theta = \{(\lambda_t, W_t^1, W_t^2)\}_{t=1}^T$ are learnable parameters.

Experiment setting. We follow [Chen et al. \(2018\)](#) to generate the samples. The signal-to-noise ratio (SNR) for each sample is uniformly sampled from 20, 30, and 40. The training loss for LISTA is $\sum_{t=1}^T \gamma^{T-t} \|\mathbf{x}_t - \mathbf{x}^*\|_2^2$ where $\gamma \leq 1$. It is commonly used for algorithm-based deep learning, so that there is a supervision signal for every layer. For ISTA and FISTA, we use the training set to tune the hyperparameters by grid search. See [Appendix B.1](#) for more details.

Table 2. Recovery performances of different algorithms/models.

SNR	mixed	20	30	40
FISTA ($T = 100$)	-18.96	-16.75	-20.46	-20.97
ISTA ($T = 100$)	-14.66	-13.99	-14.99	-15.07
ISTA ($T = 20$)	-9.17	-9.12	-9.24	-9.16
FISTA ($T = 20$)	-11.12	-10.98	-11.19	-11.19
LISTA ($T = 20$)	-17.58	-16.52	-18.16	-18.29
LISTA-stop ($T \leq 20$)	-22.41	-20.29	-23.90	-24.21

Recovery performance. (Table 2) We report the NMSE (in dB) results for each model/algorithm evaluated on 1000 fixed test samples per SNR level. It is revealed in Table 2 that learning-based methods have better recovery performances, especially for the more difficult tasks (i.e. when SNR is 20). Compared to LISTA, our proposed adaptive-stopping method (LISTA-stop) significantly improve recovery performance. Also, LISTA-stop with ≤ 20 iterations performs better than ISTA and FISTA with 100 iterations, which indicates a better convergence.

Stopping distribution. The stop time distribution $q_\phi(t)$ induced by π_ϕ can be computed via Eq. 3. We report in Fig. 4 the stopping distribution averaged over the test samples, from which we can see that with a high probability LISTA-stop terminates the process before arriving at 20-th iteration.




Figure 4. Left: Stop time distribution $\frac{1}{|\mathcal{D}_{test}|} \sum_{\mathbf{x} \in \mathcal{D}_{test}} q_\phi(t|\mathbf{x})$ averaged over the test set. Right: Convergence of different algorithms. For LISTA-stop, the NMSE weighted by the stopping distribution q_ϕ is plotted. In the first 13 iterations $q_\phi(t) = 0$, so no red dots are plotted.

Convergence comparison. Fig. 4 shows the change of

NMSE as the number of iterations increases. Since LISTA-stop outputs the results at different iteration steps, it is not meaningful to draw a unified convergence curve. Therefore, we plot the NMSE weighted by the stopping distribution q_ϕ , i.e., $10 \log_{10} \left(\frac{\sum_{i=1}^N q_\phi(t|i) \|\mathbf{x}_t - \mathbf{x}^{*,i}\|_2^2}{\sum_{i=1}^N q_\phi(t|i)} \right) / \left(\frac{\sum_{i=1}^N \|\mathbf{x}^{*,i}\|_2^2}{N} \right)$, using the red dots. We observe that for LISTA-stop the expected NMSE increases as the number of iterations increase, this might indicate that the later stopped problems are more difficult to solve. Besides, at 15th iteration, the NMSE in Fig. 4 (b) is the smallest, while the averaged stop probability mass $q_\phi(15)$ in Fig. 4 (a) is the highest.

Table 3. Different algorithms for training LISTA-stop.

SNR	mixed	20	30	40
AEVB algorithm	-21.92	-19.92	-23.27	-23.58
Stage I. + II.	-22.41	-20.29	-23.90	-24.21
Stage I.+II.+III.	-22.78	-20.59	-24.29	-24.73

Ablation study on training algorithms. To show the effectiveness of our two-stage training, in Table 3, we compare the results with the auto-encoding variational Bayes (AEVB) algorithm ([Le et al., 2018](#)) that jointly optimizes \mathcal{F}_θ and q_ϕ . We observe that the distribution q_ϕ in AEVB gradually becomes concentrated on one layer and does not get rid of this local minimum, making its final result not as good as the results of our two-stage training. Moreover, it is revealed that Stage III does not improve much of the performance of the two-stage training, which also in turn shows the effectiveness of the oracle-based two-stage training.

5.2. Task-imbalanced Meta Learning

In this section, we perform meta learning experiments in the few-shot learning domain ([Ravi & Larochelle, 2017](#)).

Experiment setting. We follow the setting in MAML ([Finn et al., 2017](#)) for the few-shot learning tasks. Each task is an N-way classification that contains meta-{train, valid, test} sets. On top of it, the macro dataset with multiple tasks is split into train, valid and test sets. We consider the more realistic *task-imbalanced* setting proposed by [Na et al. \(2020\)](#). Unlike the standard setting where the meta-train of each task contains k -shots for each class, here we vary the number of observation to perform k_1 - k_2 -shot learning where $k_1 < k_2$ are the minimum/maximum number of observations per class, respectively. Build on top of MAML, we denote our variant as MAML-stop which learns how many adaptation gradient descent steps are needed for each task. Intuitively, the tasks with less training data would prefer fewer steps of gradient-update to prevent overfitting. As we mainly focus on the effect of learning to stop, the neural architecture and other hyperparameters are largely the same as MAML. Please refer to [Appendix B.2](#) for more details.

Table 4. Few-shot classification in vanilla meta learning setting (Finn et al., 2017) where all tasks have the same number of data points.

	Omniglot 5-way		Omniglot 20-way		MiniImagenet 5-way	
	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
MAML	98.7 ± 0.4%	99.1 ± 0.1%	95.8 ± 0.3%	98.9 ± 0.2%	48.70 ± 1.84%	63.11 ± 0.92%
MAML-stop	99.62 ± 0.22%	99.68 ± 0.12%	96.05 ± 0.35%	98.94 ± 0.10 %	49.56 ± 0.82%	63.41 ± 0.80%

Dataset. We use the benchmark datasets Omniglot (Lake et al., 2011) and MiniImagenet (Ravi & Larochelle, 2017). Omniglot consists of 20 instances of 1623 characters from 50 different alphabets, while MiniImagenet involves 64 training classes, 12 validation classes, and 24 test classes. We use exactly the same data split as Finn et al. (2017). To construct the imbalanced tasks, we perform 20-way 1-5 shot classification on Omniglot and 5-way 1-10 shot classification on MiniImagenet. The number of observations per class in each meta-test set is 1 and 5 for Omniglot and MiniImagenet, respectively. For evaluation, we construct 600 tasks from the held-out test set for each setting.

Table 5. Task-imbalanced few-shot image classification.

	Omniglot	MiniImagenet
	20-way, 1-5 shot	5-way, 1-10 shot
MAML	97.96 ± 0.3%	57.20 ± 1.1%
MAML-stop	98.45 ± 0.2%	60.67 ± 1.0%

Results. Table 5 summarizes the accuracy and the 95% confidence interval on the held-out tasks for each dataset. The maximum number of adaptation gradient descent steps is 10 for both MAML and MAML-stop. We can see the optimal stopping variant of MAML outperforms the vanilla MAML consistently. For a more difficult task on MiniImagenet where the imbalance issue is more severe, the accuracy improvement is 3.5%. For completeness, we include the performance on vanilla meta learning setting where all tasks have the same number of observations in Table 4. MAML-stop still achieves comparable or better performance.

5.3. Image Denoising

In this section, we perform the image denoising experiments. More implementation details are provided in Appendix B.3.

Dataset. The models are trained on BSD500 (400 images) (Arbelaez et al., 2010), validated on BSD12, and tested on BSD68 (Martin et al., 2001). We follow the standard setting in (Zhang et al., 2019; Lefkimmatis, 2018; Zhang et al., 2017) to add Gaussian noise to the images with a random noise level $\sigma \leq 55$ during training and validation phases.

Experiment setting. We compare with two DL models, DnCNN (Zhang et al., 2017) and UNLNet₅ (Lefkimmatis, 2018), and two traditional methods, BM3D (Dabov et al., 2007) and WNNM (Gu et al., 2014). Since DnCNN is one

of the most widely-used models for image denoising, we use it as our predictive model. All deep models including ours are considered in the *blind* Gaussian denoising setting, which means the noise-level is not given to the model, while BM3D and WNNM require the noise-level to be known.

Table 6. PSNA performance comparison. The sign * indicates that noise levels 65 and 75 do not appear in the training set.

σ	DnCNN-stop	DnCNN	UNLNet ₅	BM3D	WNNM
35	27.61	27.60	27.50	26.81	27.36
45	26.59	26.56	26.48	25.97	26.31
55	25.79	25.71	25.64	25.21	25.50
*65	23.56	22.19	-	24.60	24.92
*75	18.62	17.90	-	24.08	24.39

Results. The performance is evaluated by the mean peak signal-to-noise ratio (PSNR). Table 6 shows that DnCNN-stop performs better than the original DnCNN. Especially, for images with noise levels 65 and 75 which are **unseen** during training phase, DnCNN-stop generalizes significantly better than DnCNN alone. Since there is no released code for UNLNet₅, its performances are copied from the paper (Lefkimmatis, 2018), where results are not reported for $\sigma = 65$ and 75. For traditional methods BM3D and WNNM, the test is in the noise-specific setting. That is, the noise level is given to both BM3D and WNNM, so the comparison is not completely fair to learning based methods in blind denoising setting.




Figure 5. Denoising results of an image with noise level 65. (See Appendix B.3.2 for more visualization results.)

5.4. Image Recognition

We explore the potential of our idea for improving the recognition performances on Tiny-ImageNet, using VGG16 (Simonyan & Zisserman, 2014) as the predictive model. With 14 internal classifiers, after Stage I training, if the oracle q_θ^* is used to determine the stop time t , the accuracy of VGG16 can be improved to 83.26%. Similar observation is provided in SDN (Kaya et al., 2019), but their loss $\sum_t w_t \ell_t$ depends on very careful hand-tuning on the weight w_t for each layer, while we directly take an expectation using the oracle, which is more principled and leads to higher accuracy (Table 7). However, it reveals to be very hard to mimic the behavior of the oracle q_θ^* by π_ϕ in Stage II, either due to the need of a better parametrization for π_ϕ or more sophisticated reasons. Our learned π_ϕ leads to similar accuracy as the heuristic policy in SDN, which becomes the bottleneck in our exploration. However, based on the large performance gap between the oracle and the original VGG16, our result still provides a potential direction for breaking the performance bottleneck of DL on image recognition.

Table 7. Image recognition with oracle stop distribution.

VGG16	SDN training	Our Stage I. training
58.60%	77.78% (best layer)	83.26% (best layer)

6. Conclusion

In this paper, we introduce a generic framework for modelling and training a deep learning model with input-specific depth, which is determined by a stopping policy π_ϕ . Extensive experiments are conducted to demonstrate the effectiveness of both the model and the training algorithm, on a wide range of applications. In the future, it will be interesting to see whether other aspects of algorithms can be incorporated into deep learning models either to improve the performance or for better theoretical understandings.

Acknowledgement

We would like to thank anonymous reviewers for providing constructive feedbacks. This work is supported in part by NSF grants CDS&E-1900017 D3SC, CCF-1836936 FMITF, IIS-1841351, CAREER IIS-1350983 to L.S. and grants from King Abdullah University of Science and Technology, under award numbers BAS/1/1624-01, FCC/1/1976-18-01, FCC/1/1976-23-01, FCC/1/1976-25-01, FCC/1/1976-26-01, REI/1/0018-01-01, and URF/1/4098-01-01.

References

Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M. W., Pfau, D., Schaul, T., Shillingford, B., and De Freitas, N. Learning to learn by gradient descent by gradient descent.

In *Advances in Neural Information Processing Systems*, pp. 3981–3989, 2016.

Arbelaez, P., Maire, M., Fowlkes, C., and Malik, J. Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 33(5):898–916, 2010.

Beck, A. and Teboulle, M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.

Becker, S., Cheridito, P., and Jentzen, A. Deep optimal stopping. *Journal of Machine Learning Research*, 20(74):1–25, 2019.

Belanger, D., Yang, B., and McCallum, A. End-to-end learning for structured prediction energy networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 429–439. JMLR. org, 2017.

Blumensath, T. and Davies, M. E. Iterative thresholding for sparse approximations. *Journal of Fourier analysis and Applications*, 14(5-6):629–654, 2008.

Borgerding, M., Schniter, P., and Rangan, S. Amp-inspired deep networks for sparse linear inverse problems. *IEEE Transactions on Signal Processing*, 65(16):4293–4308, 2017.

Ceci, C. and Bassan, B. Mixed optimal stopping and stochastic control problems with semicontinuous final reward for diffusion processes. *Stochastics and Stochastic Reports*, 76(4):323–337, 2004.

Chen, X., Liu, J., Wang, Z., and Yin, W. Theoretical linear convergence of unfolded ista and its practical weights and thresholds. In *Advances in Neural Information Processing Systems*, pp. 9061–9071, 2018.

Chen, X., Dai, H., and Song, L. Particle flow bayes rule. In *International Conference on Machine Learning*, pp. 1022–1031, 2019.

Chen, X., Li, Y., Umarov, R., Gao, X., and Song, L. Rna secondary structure prediction by learning unrolled algorithms. *arXiv preprint arXiv:2002.05810*, 2020.

Dabov, K., Foi, A., Katkovnik, V., and Egiazarian, K. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007.

Domke, J. Parameter learning with truncated message-passing. In *CVPR 2011*, pp. 2937–2943. IEEE, 2011.

Dumitrescu, R., Reisinger, C., and Zhang, Y. Approximation schemes for mixed optimal stopping and control problems with nonlinear expectations and jumps. *arXiv preprint arXiv:1803.03794*, 2018.

- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1126–1135. JMLR. org, 2017.
- Gregor, K. and LeCun, Y. Learning fast approximations of sparse coding. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, pp. 399–406. Omnipress, 2010.
- Gu, S., Zhang, L., Zuo, W., and Feng, X. Weighted nuclear norm minimization with application to image denoising. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2862–2869, 2014.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-VAE: Learning basic visual concepts with a constrained variational framework. *ICLR*, 2(5):6, 2017.
- Huang, G., Chen, D., Li, T., Wu, F., van der Maaten, L., and Weinberger, K. Multi-scale dense networks for resource efficient image classification. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=Hk2aImxAb>.
- Ingraham, J., Riesselman, A., Sander, C., and Marks, D. Learning protein structure with a differentiable simulator. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Byg3y3C9Km>.
- Kaya, Y., Hong, S., and Dumitras, T. Shallow-deep networks: Understanding and mitigating network overthinking. In *International Conference on Machine Learning*, pp. 3301–3310, 2019.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Lake, B., Salakhutdinov, R., Gross, J., and Tenenbaum, J. One shot learning of simple visual concepts. In *Proceedings of the annual meeting of the cognitive science society*, volume 33, 2011.
- Le, T. A., Igl, M., Rainforth, T., Jin, T., and Wood, F. Auto-encoding sequential monte carlo. In *International Conference on Learning Representations*, 2018.
- Lee, Y. and Choi, S. Gradient-based meta-learning with learned layerwise metric and subspace. *arXiv preprint arXiv:1801.05558*, 2018.
- Lefkimmiatis, S. Universal denoising networks: a novel cnn architecture for image denoising. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3204–3213, 2018.
- Li, K. and Malik, J. Learning to optimize. *arXiv preprint arXiv:1606.01885*, 2016.
- Li, Z., Zhou, F., Chen, F., and Li, H. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*, 2017.
- Liu, J., Chen, X., Wang, Z., and Yin, W. ALISTA: Analytic weights are as good as learned weights in LISTA. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=B1lnzn0ctQ>.
- Martin, D., Fowlkes, C., Tal, D., and Malik, J. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pp. 416–423. IEEE, 2001.
- Metzler, C., Mousavi, A., and Baraniuk, R. Learned d-amp: Principled neural network based compressive image recovery. In *Advances in Neural Information Processing Systems*, pp. 1772–1783, 2017.
- Na, D., Lee, H. B., Lee, H., Kim, S., Park, M., Yang, E., and Hwang, S. J. Learning to balance: Bayesian meta-learning for imbalanced and out-of-distribution tasks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rkeZIJBYvr>.
- Nowozin, S., Cseke, B., and Tomioka, R. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in neural information processing systems*, pp. 271–279, 2016.
- Oreshkin, B., López, P. R., and Lacoste, A. Tadam: Task dependent adaptive metric for improved few-shot learning. In *Advances in Neural Information Processing Systems*, pp. 721–731, 2018.
- Pham, H. Optimal stopping of controlled jump diffusion processes: a viscosity solution approach. In *Journal of Mathematical Systems, Estimation and Control*. Citeseer, 1998.
- Qiao, S., Liu, C., Shen, W., and Yuille, A. L. Few-shot image recognition by predicting parameters from activations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7229–7238, 2018.
- Ravi, S. and Larochelle, H. Optimization as a model for few-shot learning. 2017.
- Shiryayev, A. N. *Optimal stopping rules*, volume 8. Springer Science & Business Media, 2007.

- Shrivastava, H., Chen, X., Chen, B., Lan, G., Aluru, S., Liu, H., and Song, L. GLAD: Learning sparse graph recovery. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=BkxpMTetPB>.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Sun, J., Li, H., Xu, Z., et al. Deep admm-net for compressive sensing mri. In *Advances in neural information processing systems*, pp. 10–18, 2016.
- Teerapittayanon, S., McDanel, B., and Kung, H.-T. Branchynet: Fast inference via early exiting from deep neural networks. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pp. 2464–2469. IEEE, 2016.
- Yakar, T. B., Litman, R., Sprechmann, P., Bronstein, A. M., and Sapiro, G. Bilevel sparse models for polyphonic music transcription. In *ISMIR*, pp. 65–70, 2013.
- Zamir, A. R., Wu, T.-L., Sun, L., Shen, W. B., Shi, B. E., Malik, J., and Savarese, S. Feedback networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1308–1317, 2017.
- Zhang, J. and Ghanem, B. Ista-net: Interpretable optimization-inspired deep network for image compressive sensing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1828–1837, 2018.
- Zhang, K., Zuo, W., Chen, Y., Meng, D., and Zhang, L. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017.
- Zhang, X., Lu, Y., Liu, J., and Dong, B. Dynamically unfolding recurrent restorer: A moving endpoint control method for image restoration. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=SJfzKiC5FX>.