

# SELF-TUNING NETWORKS: BILEVEL OPTIMIZATION OF HYPERPARAMETERS USING STRUCTURED BEST-RESPONSE FUNCTIONS

Matthew MacKay\*, Paul Vicol\*, Jon Lorraine, David Duvenaud, Roger Grosse

{mmackay, pvicol, lorraine, duvenaud, rgrosse}@cs.toronto.edu

University of Toronto

Vector Institute

## ABSTRACT

Hyperparameter optimization can be formulated as a bilevel optimization problem, where the optimal parameters on the training set depend on the hyperparameters. We aim to adapt regularization hyperparameters for neural networks by fitting compact approximations to the *best-response function*, which maps hyperparameters to optimal weights and biases. We show how to construct scalable best-response approximations for neural networks by modeling the best-response as a single network whose hidden units are gated conditionally on the regularizer. We justify this approximation by showing the exact best-response for a shallow linear network with  $L_2$ -regularized Jacobian can be represented by a similar gating mechanism. We fit this model using a gradient-based hyperparameter optimization algorithm which alternates between approximating the best-response around the current hyperparameters and optimizing the hyperparameters using the approximate best-response function. Unlike other gradient-based approaches, we do not require differentiating the training loss with respect to the hyperparameters, allowing us to tune discrete hyperparameters, data augmentation hyperparameters, and dropout probabilities. Because the hyperparameters are adapted online, our approach discovers *hyperparameter schedules* that can outperform fixed hyperparameter values. Empirically, our approach outperforms competing hyperparameter optimization methods on large-scale deep learning problems. We call our networks, which update their own hyperparameters online during training, *Self-Tuning Networks (STNs)*.

## 1 INTRODUCTION

Regularization hyperparameters such as weight decay, data augmentation, and dropout (Srivastava et al., 2014) are crucial to the generalization of neural networks, but are difficult to tune. Popular approaches to hyperparameter optimization include grid search, random search (Bergstra & Bengio, 2012), and Bayesian optimization (Snoek et al., 2012). These approaches work well with low-dimensional hyperparameter spaces and ample computational resources; however, they pose hyperparameter optimization as a black-box optimization problem, ignoring structure which can be exploited for faster convergence, and require many training runs.

We can formulate hyperparameter optimization as a bilevel optimization problem. Let  $\mathbf{w}$  denote parameters (e.g. weights and biases) and  $\lambda$  denote hyperparameters (e.g. dropout probability). Let  $\mathcal{L}_T$  and  $\mathcal{L}_V$  be functions mapping parameters and hyperparameters to training and validation losses, respectively. We aim to solve<sup>1</sup>:

$$\lambda^* = \arg \min_{\lambda} \mathcal{L}_V(\lambda, \mathbf{w}^*) \quad \text{subject to} \quad \mathbf{w}^* = \arg \min_{\mathbf{w}} \mathcal{L}_T(\lambda, \mathbf{w}) \quad (1)$$

\*Equal contribution.

<sup>1</sup>The uniqueness of the arg min is assumed.

Substituting the *best-response function*  $\mathbf{w}^*(\boldsymbol{\lambda}) = \arg \min_{\mathbf{w}} \mathcal{L}_T(\boldsymbol{\lambda}, \mathbf{w})$  gives a single-level problem:

$$\boldsymbol{\lambda}^* = \arg \min_{\boldsymbol{\lambda}} \mathcal{L}_V(\boldsymbol{\lambda}, \mathbf{w}^*(\boldsymbol{\lambda})) \quad (2)$$

If the best-response  $\mathbf{w}^*$  is known, the validation loss can be minimized directly by gradient descent using Equation 2, offering dramatic speed-ups over black-box methods. However, as the solution to a high-dimensional optimization problem, it is difficult to compute  $\mathbf{w}^*$  even approximately.

Following Lorraine & Duvenaud (2018), we propose to approximate the best-response  $\mathbf{w}^*$  directly with a parametric function  $\hat{\mathbf{w}}_\phi$ . We jointly optimize  $\phi$  and  $\boldsymbol{\lambda}$ , first updating  $\phi$  so that  $\hat{\mathbf{w}}_\phi \approx \mathbf{w}^*$  in a neighborhood around the current hyperparameters, then updating  $\boldsymbol{\lambda}$  by using  $\hat{\mathbf{w}}_\phi$  as a proxy for  $\mathbf{w}^*$  in Eq. 2:

$$\boldsymbol{\lambda}^* \approx \arg \min_{\boldsymbol{\lambda}} \mathcal{L}_V(\boldsymbol{\lambda}, \hat{\mathbf{w}}_\phi(\boldsymbol{\lambda})) \quad (3)$$

Finding a scalable approximation  $\hat{\mathbf{w}}_\phi$  when  $\mathbf{w}$  represents the weights of a neural network is a significant challenge, as even simple implementations entail significant memory overhead. We show how to construct a compact approximation by modelling the best-response of each row in a layer’s weight matrix/bias as a rank-one affine transformation of the hyperparameters. We show that this can be interpreted as computing the activations of a base network in the usual fashion, plus a correction term dependent on the hyperparameters. We justify this approximation by showing the exact best-response for a shallow linear network with  $L_2$ -regularized Jacobian follows a similar structure. We call our proposed networks *Self-Tuning Networks* (STNs) since they update their own hyperparameters online during training.

STNs enjoy many advantages over other hyperparameter optimization methods. First, they are easy to implement by replacing existing modules in deep learning libraries with “hyper” counterparts which accept an additional vector of hyperparameters as input<sup>2</sup>. Second, because the hyperparameters are adapted online, we ensure that computational effort expended to fit  $\phi$  around previous hyperparameters is not wasted. In addition, this online adaption yields hyperparameter schedules which we find empirically to outperform fixed hyperparameter settings. Finally, the STN training algorithm does not require differentiating the training loss with respect to the hyperparameters, unlike other gradient-based approaches (Maclaurin et al., 2015; Larsen et al., 1996), allowing us to tune discrete hyperparameters, such as the number of holes to cut out of an image (DeVries & Taylor, 2017), data-augmentation hyperparameters, and discrete-noise dropout parameters. Empirically, we evaluate the performance of STNs on large-scale deep-learning problems with the Penn Treebank (Marcus et al., 1993) and CIFAR-10 datasets (Krizhevsky & Hinton, 2009), and find that they substantially outperform baseline methods.

## 2 BILEVEL OPTIMIZATION

A bilevel optimization problem consists of two sub-problems called the *upper-level* and *lower-level* problems, where the upper-level problem must be solved subject to optimality of the lower-level problem. Minimax problems are an example of bilevel programs where the upper-level objective equals the negative lower-level objective. Bilevel programs were first studied in economics to model leader/follower firm dynamics (Von Stackelberg, 2010) and have since found uses in various fields (see Colson et al. (2007) for an overview). In machine learning, many problems can be formulated as bilevel programs, including hyperparameter optimization, GAN training (Goodfellow et al., 2014), meta-learning, and neural architecture search (Zoph & Le, 2016).

Even if all objectives and constraints are linear, bilevel problems are strongly NP-hard (Hansen et al., 1992; Vicente et al., 1994). Due to the difficulty of obtaining exact solutions, most work has focused on restricted settings, considering linear, quadratic, and convex functions. In contrast, we focus on obtaining local solutions in the nonconvex, differentiable, and unconstrained setting. Let  $F, f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$  denote the upper- and lower-level objectives (e.g.,  $\mathcal{L}_V$  and  $\mathcal{L}_T$ ) and  $\boldsymbol{\lambda} \in \mathbb{R}^n, \mathbf{w} \in \mathbb{R}^m$  denote the upper- and lower-level parameters. We aim to solve:

$$\min_{\boldsymbol{\lambda} \in \mathbb{R}^n} F(\boldsymbol{\lambda}, \mathbf{w}) \quad (4a)$$

$$\text{subject to } \mathbf{w} \in \arg \min_{\mathbf{w} \in \mathbb{R}^m} f(\boldsymbol{\lambda}, \mathbf{w}) \quad (4b)$$

<sup>2</sup>We illustrate how this is done for the PyTorch library (Paszke et al., 2017) in Appendix G.

It is desirable to design a gradient-based algorithm for solving Problem 4, since using gradient information provides drastic speed-ups over black-box optimization methods (Nesterov, 2013). The simplest method is simultaneous gradient descent, which updates  $\lambda$  using  $\partial F/\partial \lambda$  and  $\mathbf{w}$  using  $\partial f/\partial \mathbf{w}$ . However, simultaneous gradient descent often gives incorrect solutions as it fails to account for the dependence of  $\mathbf{w}$  on  $\lambda$ . Consider the relatively common situation where  $F$  doesn't depend directly on  $\lambda$ , so that  $\partial F/\partial \lambda \equiv \mathbf{0}$  and hence  $\lambda$  is never updated.

## 2.1 GRADIENT DESCENT VIA THE BEST-RESPONSE FUNCTION

A more principled approach to solving Problem 4 is to use the *best-response function* (Gibbons, 1992). Assume the lower-level Problem 4b has a unique optimum  $\mathbf{w}^*(\lambda)$  for each  $\lambda$ . Substituting the best-response function  $\mathbf{w}^*$  converts Problem 4 into a single-level problem:

$$\min_{\lambda \in \mathbb{R}^n} F^*(\lambda) := F(\lambda, \mathbf{w}^*(\lambda)) \quad (5)$$

If  $\mathbf{w}^*$  is differentiable, we can minimize Eq. 5 using gradient descent on  $F^*$  with respect to  $\lambda$ . This method requires a unique optimum  $\mathbf{w}^*(\lambda)$  for Problem 4b for each  $\lambda$  and differentiability of  $\mathbf{w}^*$ . In general, these conditions are difficult to verify. We give sufficient conditions for them to hold in a neighborhood of a point  $(\lambda_0, \mathbf{w}_0)$  where  $\mathbf{w}_0$  solves Problem 4b given  $\lambda_0$ .

**Lemma 1.** (Fiacco & Ishizuka, 1990) *Let  $\mathbf{w}_0$  solve Problem 4b for  $\lambda_0$ . Suppose  $f$  is  $C^2$  in a neighborhood of  $(\lambda_0, \mathbf{w}_0)$  and the Hessian  $\partial^2 f/\partial \mathbf{w}^2(\lambda_0, \mathbf{w}_0)$  is positive definite. Then for some neighborhood  $U$  of  $\lambda_0$ , there exists a continuously differentiable function  $\mathbf{w}^* : U \rightarrow \mathbb{R}^m$  such that  $\mathbf{w}^*(\lambda)$  is the unique solution to Problem 4b for each  $\lambda \in U$  and  $\mathbf{w}^*(\lambda_0) = \mathbf{w}_0$ .*

*Proof.* See Appendix B.1. □

The gradient of  $F^*$  decomposes into two terms, which we term the *direct gradient* and the *response gradient*. The direct gradient captures the direct reliance of the upper-level objective on  $\lambda$ , while the response gradient captures how the lower-level parameter responds to changes in the upper-level parameter:

$$\frac{\partial F^*}{\partial \lambda}(\lambda_0) = \underbrace{\frac{\partial F}{\partial \lambda}(\lambda_0, \mathbf{w}^*(\lambda_0))}_{\text{Direct gradient}} + \underbrace{\frac{\partial F}{\partial \mathbf{w}}(\lambda_0, \mathbf{w}^*(\lambda_0)) \frac{\partial \mathbf{w}^*}{\partial \lambda}(\lambda_0)}_{\text{Response gradient}} \quad (6)$$

Even if  $\partial F/\partial \lambda \neq 0$  and simultaneous gradient descent is possible, including the response gradient can stabilize optimization by converting the bilevel problem into a single-level one, as noted by Metz et al. (2016) for GAN optimization. Conversion to a single-level problem ensures that the gradient vector field is conservative, avoiding pathological issues described by Mescheder et al. (2017).

## 2.2 APPROXIMATING THE BEST-RESPONSE FUNCTION

In general, the solution to Problem 4b is a set, but assuming uniqueness of a solution and differentiability of  $\mathbf{w}^*$  can yield fruitful algorithms in practice. In fact, gradient-based hyperparameter optimization methods can often be interpreted as approximating either the best-response  $\mathbf{w}^*$  or its Jacobian  $\partial \mathbf{w}^*/\partial \lambda$ , as detailed in Section 5. However, these approaches can be computationally expensive and often struggle with discrete hyperparameters and stochastic hyperparameters like dropout probabilities, since they require differentiating the training loss with respect to the hyperparameters. Promising approaches to approximate  $\mathbf{w}^*$  directly were proposed by Lorraine & Duvenaud (2018), and are detailed below.

**1. Global Approximation.** The first algorithm proposed by Lorraine & Duvenaud (2018) approximates  $\mathbf{w}^*$  as a differentiable function  $\hat{\mathbf{w}}_\phi$  with parameters  $\phi$ . If  $\mathbf{w}$  represents neural net weights, then the mapping  $\hat{\mathbf{w}}_\phi$  is a hypernetwork (Schmidhuber, 1992; Ha et al., 2016). If the distribution  $p(\lambda)$  is fixed, then gradient descent with respect to  $\phi$  minimizes:

$$\mathbb{E}_{\lambda \sim p(\lambda)} [f(\lambda, \hat{\mathbf{w}}_\phi(\lambda))] \quad (7)$$

If  $\text{support}(p)$  is broad and  $\hat{\mathbf{w}}_\phi$  is sufficiently flexible, then  $\hat{\mathbf{w}}_\phi$  can be used as a proxy for  $\mathbf{w}^*$  in Problem 5, resulting in the following objective:

$$\min_{\lambda \in \mathbb{R}^n} F(\lambda, \hat{\mathbf{w}}_\phi(\lambda)) \quad (8)$$

**2. Local Approximation.** In practice,  $\hat{\mathbf{w}}_\phi$  is usually insufficiently flexible to model  $\mathbf{w}^*$  on support( $p$ ). The second algorithm of Lorraine & Duvenaud (2018) locally approximates  $\mathbf{w}^*$  in a neighborhood around the current upper-level parameter  $\boldsymbol{\lambda}$ . They set  $p(\epsilon|\boldsymbol{\sigma})$  to a factorized Gaussian noise distribution with a fixed scale parameter  $\boldsymbol{\sigma} \in \mathbb{R}_+^n$ , and found  $\phi$  by minimizing the objective:

$$\mathbb{E}_{\epsilon \sim p(\epsilon|\boldsymbol{\sigma})} [f(\boldsymbol{\lambda} + \epsilon, \hat{\mathbf{w}}_\phi(\boldsymbol{\lambda} + \epsilon))] \quad (9)$$

Intuitively, the upper-level parameter  $\boldsymbol{\lambda}$  is perturbed by a small amount, so the lower-level parameter learns how to respond. An alternating gradient descent scheme is used, where  $\phi$  is updated to minimize equation 9 and  $\boldsymbol{\lambda}$  is updated to minimize equation 8. This approach worked for problems using  $L_2$  regularization on MNIST (LeCun et al., 1998). However, it is unclear if the approach works with different regularizers or scales to larger problems. It requires  $\hat{\mathbf{w}}_\phi$ , which is *a priori* unwieldy for high dimensional  $\mathbf{w}$ . It is also unclear how to set  $\boldsymbol{\sigma}$ , which defines the size of the neighborhood on which  $\phi$  is trained, or if the approach can be adapted to discrete and stochastic hyperparameters.

### 3 SELF-TUNING NETWORKS

In this section, we first construct a best-response approximation  $\hat{\mathbf{w}}_\phi$  that is memory efficient and scales to large neural networks. We justify this approximation through analysis of simpler situations. Then, we describe a method to automatically adjust the scale of the neighborhood  $\phi$  is trained on. Finally, we formally describe our algorithm and discuss how it easily handles discrete and stochastic hyperparameters. We call the resulting networks, which update their own hyperparameters online during training, Self-Tuning Networks (STNs).

#### 3.1 AN EFFICIENT BEST-RESPONSE APPROXIMATION FOR NEURAL NETWORKS

We propose to approximate the best-response for a given layer’s weight matrix  $\mathbf{W} \in \mathbb{R}^{D_{out} \times D_{in}}$  and bias  $\mathbf{b} \in \mathbb{R}^{D_{out}}$  as an affine transformation of the hyperparameters  $\boldsymbol{\lambda}$ <sup>3</sup>:

$$\hat{\mathbf{W}}_\phi(\boldsymbol{\lambda}) = \mathbf{W}_{\text{elem}} + (\mathbf{V}\boldsymbol{\lambda}) \odot_{\text{row}} \mathbf{W}_{\text{hyper}}, \quad \hat{\mathbf{b}}_\phi(\boldsymbol{\lambda}) = \mathbf{b}_{\text{elem}} + (\mathbf{C}\boldsymbol{\lambda}) \odot \mathbf{b}_{\text{hyper}} \quad (10)$$

Here,  $\odot$  indicates elementwise multiplication and  $\odot_{\text{row}}$  indicates row-wise rescaling. This architecture computes the usual elementary weight/bias, plus an additional weight/bias which has been scaled by a linear transformation of the hyperparameters. Alternatively, it can be interpreted as directly operating on the pre-activations of the layer, adding a correction to the usual pre-activation to account for the hyperparameters:

$$\hat{\mathbf{W}}_\phi(\boldsymbol{\lambda})\mathbf{x} + \hat{\mathbf{b}}_\phi(\boldsymbol{\lambda}) = [\mathbf{W}_{\text{elem}}\mathbf{x} + \mathbf{b}_{\text{elem}}] + [(\mathbf{V}\boldsymbol{\lambda}) \odot (\mathbf{W}_{\text{hyper}}\mathbf{x}) + (\mathbf{C}\boldsymbol{\lambda}) \odot \mathbf{b}_{\text{hyper}}] \quad (11)$$

This best-response architecture is tractable to compute and memory-efficient: it requires  $D_{out}(2D_{in} + n)$  parameters to represent  $\hat{\mathbf{W}}_\phi$  and  $D_{out}(2 + n)$  parameters to represent  $\hat{\mathbf{b}}_\phi$ , where  $n$  is the number of hyperparameters. Furthermore, it enables parallelism: since the predictions can be computed by transforming the pre-activations (Equation 11), the hyperparameters for different examples in a batch can be perturbed independently, improving sample efficiency. In practice, the approximation can be implemented by simply replacing existing modules in deep learning libraries with “hyper” counterparts which accept an additional vector of hyperparameters as input<sup>4</sup>.


#### 3.2 EXACT BEST-RESPONSE FOR TWO-LAYER LINEAR NETWORKS

Given that the best-response function is a mapping from  $\mathbb{R}^n$  to the high-dimensional weight space  $\mathbb{R}^m$ , why should we expect to be able to represent it compactly? And why in particular would equation 10 be a reasonable approximation? In this section, we exhibit a model whose best-response function can be represented exactly using a minor variant of equation 10: a linear network with Jacobian norm regularization. In particular, the best-response takes the form of a network whose hidden units are modulated conditionally on the hyperparameters.

Consider using a 2-layer linear network with weights  $\mathbf{w} = (\mathbf{Q}, \mathbf{s}) \in \mathbb{R}^{D \times D} \times \mathbb{R}^D$  to predict targets  $t \in \mathbb{R}$  from inputs  $\mathbf{x} \in \mathbb{R}^D$ :

<sup>3</sup>We describe modifications for convolutional filters in Appendix C.

<sup>4</sup>We illustrate how this is done for the PyTorch library (Paszke et al., 2017) in Appendix G.

Figure 1: Best-response architecture for an  $L_2$ -Jacobian regularized two-layer linear network.

$$\mathbf{a}(\mathbf{x}; \mathbf{w}) = \mathbf{Q}\mathbf{x}, \quad y(\mathbf{x}; \mathbf{w}) = \mathbf{s}^\top \mathbf{a}(\mathbf{x}; \mathbf{w}) \quad (12)$$

Suppose we use a squared-error loss regularized with an  $L_2$  penalty on the Jacobian  $\partial y / \partial \mathbf{x}$ , where the penalty weight  $\lambda$  lies in  $\mathbb{R}$  and is mapped using  $\exp$  to lie  $\mathbb{R}_+$ :

$$\mathcal{L}_T(\lambda, \mathbf{w}) = \sum_{(\mathbf{x}, t) \in \mathcal{D}} (y(\mathbf{x}; \mathbf{w}) - t)^2 + \frac{1}{|\mathcal{D}|} \exp(\lambda) \left\| \frac{\partial y}{\partial \mathbf{x}}(\mathbf{x}; \mathbf{w}) \right\|^2 \quad (13)$$

**Theorem 2.** Let  $\mathbf{w}_0 = (\mathbf{Q}_0, \mathbf{s}_0)$ , where  $\mathbf{Q}_0$  is the change-of-basis matrix to the principal components of the data matrix and  $\mathbf{s}_0$  solves the unregularized version of Problem 13 given  $\mathbf{Q}_0$ . Then there exist  $\mathbf{v}, \mathbf{c} \in \mathbb{R}^D$  such that the best-response function<sup>5</sup>  $\mathbf{w}^*(\lambda) = (\mathbf{Q}^*(\lambda), \mathbf{s}^*(\lambda))$  is:

$$\mathbf{Q}^*(\lambda) = \sigma(\lambda \mathbf{v} + \mathbf{c}) \odot_{\text{row}} \mathbf{Q}_0, \quad \mathbf{s}^*(\lambda) = \mathbf{s}_0,$$

where  $\sigma$  is the sigmoid function.

*Proof.* See Appendix B.2. □

Observe that  $y(\mathbf{x}; \mathbf{w}^*(\lambda))$  can be implemented as a regular network with weights  $\mathbf{w}_0 = (\mathbf{Q}_0, \mathbf{s}_0)$  with an additional sigmoidal gating of its hidden units  $\mathbf{a}(\mathbf{x}; \mathbf{w}^*(\lambda))$ :

$$\mathbf{a}(\mathbf{x}; \mathbf{w}^*(\lambda)) = \mathbf{Q}^*(\lambda) \mathbf{x} = \sigma(\lambda \mathbf{v} + \mathbf{c}) \odot_{\text{row}} (\mathbf{Q}_0 \mathbf{x}) = \sigma(\lambda \mathbf{v} + \mathbf{c}) \odot_{\text{row}} \mathbf{a}(\mathbf{x}; \mathbf{w}_0) \quad (14)$$

This architecture is shown in Figure 1. Inspired by this example, we use a similar gating of the hidden units to approximate the best-response for deep, nonlinear networks.

### 3.3 LINEAR BEST-RESPONSE APPROXIMATIONS

The sigmoidal gating architecture of the preceding section can be further simplified if one only needs to approximate the best-response function for a small range of hyperparameter values. In particular, for a narrow enough hyperparameter distribution, a smooth best-response function can be approximated by an affine function (i.e. its first-order Taylor approximation). Hence, we replace the sigmoidal gating with linear gating, in order that the weights be affine in the hyperparameters. The following theorem shows that, for quadratic lower-level objectives, using an affine approximation to the best-response function and minimizing  $\mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)} [f(\boldsymbol{\lambda} + \epsilon, \hat{\mathbf{w}}_\phi(\boldsymbol{\lambda} + \epsilon))]$  yields the correct best-response Jacobian, thus ensuring gradient descent on the approximate objective  $F(\boldsymbol{\lambda}, \hat{\mathbf{w}}_\phi(\boldsymbol{\lambda}))$  converges to a local optimum:

**Theorem 3.** Suppose  $f$  is quadratic with  $\partial^2 f / \partial \mathbf{w}^2 \succ \mathbf{0}$ ,  $p(\epsilon|\sigma)$  is Gaussian with mean  $\mathbf{0}$  and variance  $\sigma^2 \mathbf{I}$ , and  $\hat{\mathbf{w}}_\phi$  is affine. Fix  $\boldsymbol{\lambda}_0 \in \mathbb{R}^n$  and let  $\phi^* = \arg \min_\phi \mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)} [f(\boldsymbol{\lambda} + \epsilon, \hat{\mathbf{w}}_\phi(\boldsymbol{\lambda} + \epsilon))]$ . Then we have  $\hat{\mathbf{w}}_\phi / \partial \boldsymbol{\lambda}(\boldsymbol{\lambda}_0) = \partial \mathbf{w}^* / \partial \boldsymbol{\lambda}(\boldsymbol{\lambda}_0)$ .

*Proof.* See Appendix B.3. □




Figure 2: **The effect of the sampled neighborhood.** **Left:** If the sampled neighborhood is too small (e.g., a point mass) the approximation learned will only match the exact best-response at the current hyperparameter, with no guarantee that its gradient matches that of the best-response. **Middle:** If the sampled neighborhood is not too small or too wide, the gradient of the approximation will match that of the best-response. **Right:** If the sampled neighborhood is too wide, the approximation will be insufficiently flexible to model the best-response, and again the gradients will not match.

### 3.4 ADAPTING THE HYPERPARAMETER DISTRIBUTION

The entries of  $\sigma$  control the scale of the hyperparameter distribution on which  $\phi$  is trained. If the entries are too large, then  $\hat{\mathbf{w}}_\phi$  will not be flexible enough to capture the best-response over the samples. However, the entries must remain large enough to force  $\hat{\mathbf{w}}_\phi$  to capture the shape locally around the current hyperparameter values. We illustrate this in Figure 2. As the smoothness of the loss landscape changes during training, it may be beneficial to vary  $\sigma$ .

To address these issues, we propose adjusting  $\sigma$  during training based on the sensitivity of the upper-level objective to the sampled hyperparameters. We include an entropy term weighted by  $\tau \in \mathbb{R}_+$  which acts to enlarge the entries of  $\sigma$ . The resulting objective is:

$$\mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)}[F(\boldsymbol{\lambda} + \epsilon, \hat{\mathbf{w}}_\phi(\boldsymbol{\lambda} + \epsilon))] - \tau \mathbb{H}[p(\epsilon|\sigma)] \quad (15)$$

This is similar to a variational inference objective, where the first term is analogous to the negative log-likelihood, but  $\tau \neq 1$ . As  $\tau$  ranges from 0 to 1, our objective interpolates between variational optimization (Staines & Barber, 2012) and variational inference, as noted by Khan et al. (2018). Similar objectives have been used in the variational inference literature for better training (Blundell et al., 2015) and representation learning (Higgins et al., 2017).

Minimizing the first term on its own eventually moves all probability mass towards an optimum  $\boldsymbol{\lambda}^*$ , resulting in  $\sigma = \mathbf{0}$  if  $\boldsymbol{\lambda}^*$  is an isolated local minimum. This compels  $\sigma$  to balance between shrinking to decrease the first term while remaining sufficiently large to avoid a heavy entropy penalty. When benchmarking our algorithm’s performance, we evaluate  $F(\boldsymbol{\lambda}, \hat{\mathbf{w}}_\phi(\boldsymbol{\lambda}))$  at the deterministic current hyperparameter  $\boldsymbol{\lambda}_0$ . (This is a common practice when using stochastic operations during training, such as batch normalization or dropout.)

### 3.5 TRAINING ALGORITHM

We now describe the complete STN training algorithm and discuss how it can tune hyperparameters that other gradient-based algorithms cannot, such as discrete or stochastic hyperparameters. We use an unconstrained parametrization  $\boldsymbol{\lambda} \in \mathbb{R}^n$  of the hyperparameters. Let  $r$  denote the element-wise function which maps  $\boldsymbol{\lambda}$  to the appropriate constrained space, which will involve a non-differentiable discretization for discrete hyperparameters.

Let  $\mathcal{L}_T$  and  $\mathcal{L}_V$  denote training and validation losses which are (possibly stochastic, e.g., if using dropout) functions of the hyperparameters and parameters. Define functions  $f, F$  by  $f(\boldsymbol{\lambda}, \mathbf{w}) = \mathcal{L}_T(r(\boldsymbol{\lambda}), \mathbf{w})$  and  $F(\boldsymbol{\lambda}, \mathbf{w}) = \mathcal{L}_V(r(\boldsymbol{\lambda}), \mathbf{w})$ . STNs are trained by a gradient descent scheme which alternates between updating  $\phi$  for  $T_{train}$  steps to minimize  $\mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)}[f(\boldsymbol{\lambda} + \epsilon, \hat{\mathbf{w}}_\phi(\boldsymbol{\lambda} + \epsilon))]$  (Eq. 9) and updating  $\boldsymbol{\lambda}$  and  $\sigma$  for  $T_{valid}$  steps to minimize  $\mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)}[F(\boldsymbol{\lambda} + \epsilon, \hat{\mathbf{w}}_\phi(\boldsymbol{\lambda} + \epsilon))] - \tau \mathbb{H}[p(\epsilon|\sigma)]$  (Eq. 15). We give our complete algorithm as Algorithm 1 and show how it can be implemented in code in Appendix G. The possible non-differentiability of  $r$  due to discrete hyperparameters poses no problem. To estimate the derivative of  $\mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)}[f(\boldsymbol{\lambda} + \epsilon, \hat{\mathbf{w}}_\phi(\boldsymbol{\lambda} + \epsilon))]$  with respect to  $\phi$ , we can use the reparametrization trick and compute  $\partial f / \partial \mathbf{w}$  and  $\partial \hat{\mathbf{w}}_\phi / \partial \phi$ , neither of whose computation paths involve the discretization  $r$ . To differentiate  $\mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)}[F(\boldsymbol{\lambda} + \epsilon, \hat{\mathbf{w}}_\phi(\boldsymbol{\lambda} + \epsilon))] - \tau \mathbb{H}[p(\epsilon|\sigma)]$  with respect to a discrete hyperparameter  $\lambda_i$ , there are two cases we must consider:

<sup>5</sup>This is an abuse of notation since there is not a unique solution to Problem 13 for each  $\lambda$  in general.



**Algorithm 1** STN Training Algorithm

---

**Initialize:** Best-response approximation parameters  $\phi$ , hyperparameters  $\lambda$ , learning rates  $\{\alpha_i\}_{i=1}^3$   
**while** not converged **do**  
  **for**  $t = 1, \dots, T_{train}$  **do**  
     $\epsilon \sim p(\epsilon|\sigma)$   
     $\phi \leftarrow \phi - \alpha_1 \frac{\partial}{\partial \phi} f(\lambda + \epsilon, \hat{\mathbf{w}}_\phi(\lambda + \epsilon))$   
  **for**  $t = 1, \dots, T_{valid}$  **do**  
     $\epsilon \sim p(\epsilon|\sigma)$   
     $\lambda \leftarrow \lambda - \alpha_2 \frac{\partial}{\partial \lambda} (F(\lambda + \epsilon, \hat{\mathbf{w}}_\phi(\lambda + \epsilon)) - \tau \mathbb{H}[p(\epsilon|\sigma)])$   
     $\sigma \leftarrow \sigma - \alpha_3 \frac{\partial}{\partial \sigma} (F(\lambda + \epsilon, \hat{\mathbf{w}}_\phi(\lambda + \epsilon)) - \tau \mathbb{H}[p(\epsilon|\sigma)])$

---

**Case 1:** For most regularization schemes,  $\mathcal{L}_V$  and hence  $F$  does not depend on  $\lambda_i$  directly and thus the only gradient is through  $\hat{\mathbf{w}}_\phi$ . Thus, the reparametrization gradient can be used.

**Case 2:** If  $\mathcal{L}_V$  relies explicitly on  $\lambda_i$ , then we can use the REINFORCE gradient estimator (Williams, 1992) to estimate the derivative of the expectation with respect to  $\lambda_i$ . The number of hidden units in a layer is an example of a hyperparameter that requires this

approach since it directly affects the validation loss. We do not show this in Algorithm 1, since we do not tune any hyperparameters which fall into this case.

## 4 EXPERIMENTS

We applied our method to convolutional networks and LSTMs (Hochreiter & Schmidhuber, 1997), yielding self-tuning CNNs (ST-CNNs) and self-tuning LSTMs (ST-LSTMs). We first investigated the behavior of STNs in a simple setting where we tuned a single hyperparameter, and found that STNs discovered *hyperparameter schedules* that outperformed fixed hyperparameter values. Next, we compared the performance of STNs to commonly-used hyperparameter optimization methods on the CIFAR-10 (Krizhevsky & Hinton, 2009) and PTB (Marcus et al., 1993) datasets.

### 4.1 HYPERPARAMETER SCHEDULES

Due to the joint optimization of the hypernetwork weights and hyperparameters, STNs do not use a single, fixed hyperparameter during training. Instead, STNs discover schedules for adapting the hyperparameters online, which can outperform *any* fixed hyperparameter. We examined this behavior in detail on the PTB corpus (Marcus et al., 1993) using an ST-LSTM to tune the output dropout rate applied to the hidden units.

The schedule discovered by an ST-LSTM for output dropout, shown in Figure 3, outperforms the best, fixed output dropout rate (0.68) found by a fine-grained grid search, achieving 82.58 vs 85.83 validation perplexity. We claim that this is a consequence of the schedule, and not of regularizing effects from sampling hyperparameters or the limited capacity of  $\hat{\mathbf{w}}_\phi$ .

To rule out the possibility that the improved performance is due to stochasticity introduced by sampling hyperparameters during STN training, we trained a standard LSTM while perturbing its dropout rate around the best value found by grid search. We used (1) random Gaussian perturbations, and (2) sinusoid perturbations for a cyclic regularization schedule. STNs outperformed both perturbation methods (Table 1), showing that the improvement is not merely due to hyperparameter stochasticity. Details and plots of each perturbation method are provided in Appendix F.

Method	Val	Test
$p = 0.68$ , Fixed	85.83	83.19
$p = 0.68$ w/ Gaussian Noise	85.87	82.29
$p = 0.68$ w/ Sinusoid Noise	85.29	82.15
$p = 0.78$ (Final STN Value)	89.65	86.90
<b>STN</b>	<b>82.58</b>	<b>79.02</b>
LSTM w/ STN Schedule	82.87	79.93

Table 1: Comparing an LSTM trained with fixed and perturbed output dropouts, an STN, and LSTM trained with the STN schedule.




Figure 3: Dropout schedules found by the ST-LSTM for different initial dropout rates.

Method	PTB		CIFAR-10	
	Val Perplexity	Test Perplexity	Val Loss	Test Loss
<b>Grid Search</b>	97.32	94.58	0.794	0.809
<b>Random Search</b>	84.81	81.46	0.921	0.752
<b>Bayesian Optimization</b>	72.13	69.29	0.636	0.651
<b>STN</b>	<b>70.30</b>	<b>67.68</b>	<b>0.575</b>	<b>0.576</b>

Table 2: Final validation and test performance of each method on the PTB word-level language modeling task, and the CIFAR-10 image-classification task.

To determine whether the limited capacity of  $\hat{\mathbf{w}}_\phi$  acts as a regularizer, we trained a standard LSTM from scratch using the schedule for output dropout discovered by the ST-LSTM. Using this schedule, the standard LSTM performed nearly as well as the STN, providing evidence that the schedule itself (rather than some other aspect of the STN) was responsible for the improvement over a fixed dropout rate. To further demonstrate the importance of the hyperparameter schedule, we also trained a standard LSTM from scratch using the final dropout value found by the STN (0.78), and found that it did not perform as well as when following the schedule. The final validation and test perplexities of each variant are shown in Table 1.

Next, we show in Figure 3 that the STN discovers the same schedule regardless of the initial hyperparameter values. Because hyperparameters adapt over a shorter timescale than the weights, we find that at any given point in training, the hyperparameter adaptation has already equilibrated. As shown empirically in Appendix F, low regularization is best early in training, while higher regularization is better later on. We found that the STN schedule implements a curriculum by using a low dropout rate early in training, aiding optimization, and then gradually increasing the dropout rate, leading to better generalization.

#### 4.2 LANGUAGE MODELING

We evaluated an ST-LSTM on the PTB corpus (Marcus et al., 1993), which is widely used as a benchmark for RNN regularization due to its small size (Gal & Ghahramani, 2016; Merity et al., 2018; Wen et al., 2018). We used a 2-layer LSTM with 650 hidden units per layer and 650-dimensional word embeddings. We tuned 7 hyperparameters: variational dropout rates for the input, hidden state, and output; embedding dropout (that sets rows of the embedding matrix to  $\mathbf{0}$ ); Drop-Connect (Wan et al., 2013) on the hidden-to-hidden weight matrix; and coefficients  $\alpha$  and  $\beta$  that control the strength of activation regularization and temporal activation regularization, respectively. For LSTM tuning, we obtained the best results when using a fixed perturbation scale of 1 for the hyperparameters. Additional details about the experimental setup and the role of these hyperparameters can be found in Appendix D.




Figure 4: **(a)** A comparison of the best validation perplexity achieved on PTB over time, by grid search, random search, Bayesian optimization, and STNs. STNs achieve better (lower) validation perplexity in less time than the other methods. **(b)** The hyperparameter schedule found by the STN for each type of dropout. **(c)** The hyperparameter schedule found by the STN for the coefficients of activation regularization and temporal activation regularization.

We compared STNs to grid search, random search, and Bayesian optimization.<sup>6</sup> Figure 4a shows the best validation perplexity achieved by each method over time. STNs outperform other meth-

<sup>6</sup>For grid search and random search we used the Ray Tune libraries (<https://github.com/ray-project/ray/tree/master/python/ray/tune>). For Bayesian optimization, we used Spearmint (<https://github.com/HIPS/Spearmint>).






Figure 6: The hyperparameter schedule prescribed by the STN while training for image classification. The dropouts are indexed by the convolutional layer they are applied to. FC dropout is for the fully-connected layers.

ods, achieving lower validation perplexity more quickly. The final validation and test perplexities achieved by each method are shown in Table 2. We show the schedules the STN finds for each hyperparameter in Figures 4b and 4c; we observe that they are nontrivial, with some forms of dropout used to a greater extent at the start of training (including input and hidden dropout), some used throughout training (output dropout), and some that are increased over the course of training (embedding and weight dropout).

### 4.3 IMAGE CLASSIFICATION

We evaluated ST-CNNs on the CIFAR-10 (Krizhevsky & Hinton, 2009) dataset, where it is easy to overfit with high-capacity networks. We used the AlexNet architecture (Krizhevsky et al., 2012), and tuned: (1) continuous hyperparameters controlling per-layer activation dropout, input dropout, and scaling noise applied to the input, (2) discrete data augmentation hyperparameters controlling the length and number of cut-out holes (DeVries & Taylor, 2017), and (3) continuous data augmentation hyperparameters controlling the amount of noise to apply to the hue, saturation, brightness, and contrast of an image. In total, we considered 15 hyperparameters.

We compared STNs to grid search, random search, and Bayesian optimization. Figure 5 shows the lowest validation loss achieved by each method over time, and Table 2 shows the final validation and test losses for each method. Details of the experimental setup are provided in Appendix E. Again, STNs find better hyperparameter configurations in less time than other methods. The hyperparameter schedules found by the STN are shown in Figure 6.




Figure 5: A comparison of the best validation loss achieved on CIFAR-10 over time, by grid search, random search, Bayesian optimization, and STNs. STNs outperform other methods for many computational budgets.

## 5 RELATED WORK

**Bilevel Optimization.** Colson et al. (2007) provide an overview of bilevel problems, and a comprehensive textbook was written by Bard (2013). When the objectives/constraints are restricted to be linear, quadratic, or convex, a common approach replaces the lower-level problem with its KKT conditions added as constraints for the upper-level problem (Hansen et al., 1992; Vicente et al., 1994). In the unrestricted setting, our work loosely resembles trust-region methods (Colson et al., 2005), which repeatedly approximate the problem locally using a simpler bilevel program. In closely related work, Sinha et al. (2013) used evolutionary techniques to estimate the best-response function iteratively.

**Hypernetworks.** First considered by Schmidhuber (1993; 1992), hypernetworks are functions mapping to the weights of a neural net. Predicting weights in CNNs has been developed in various forms (Denil et al., 2013; Yang et al., 2015). Ha et al. (2016) used hypernetworks to generate weights for modern CNNs and RNNs. Brock et al. (2017) used hypernetworks to globally approximate a best-response for architecture search. Because the architecture is not optimized during training, they require a large hypernetwork, unlike ours which locally approximates the best-response.

**Gradient-Based Hyperparameter Optimization.** There are two main approaches. The first approach approximates  $\mathbf{w}^*(\lambda_0)$  using  $\mathbf{w}_T(\lambda_0, \mathbf{w}_0)$ , the value of  $\mathbf{w}$  after  $T$  steps of gradient descent on  $f$  with respect to  $\mathbf{w}$  starting at  $(\lambda_0, \mathbf{w}_0)$ . The descent steps are differentiated through to approximate  $\partial \mathbf{w}^* / \partial \lambda(\lambda_0) \approx \partial \mathbf{w}_T / \partial \lambda(\lambda_0, \mathbf{w}_0)$ . This approach was proposed by Domke (2012) and used by Maclaurin et al. (2015), Luketina et al. (2016) and Franceschi et al. (2018). The second approach uses the Implicit Function Theorem to derive  $\partial \mathbf{w}^* / \partial \lambda(\lambda_0)$  under certain conditions. This was first developed for hyperparameter optimization in neural networks (Larsen et al., 1996) and developed further by Pedregosa (2016). Similar approaches have been used for hyperparameter optimization in log-linear models (Foo et al., 2008), kernel selection (Chapelle et al., 2002; Seeger, 2007), and image reconstruction (Kunisch & Pock, 2013; Calatroni et al., 2015). Both approaches struggle with certain hyperparameters, since they differentiate gradient descent or the training loss with respect to the hyperparameters. In addition, differentiating gradient descent becomes prohibitively expensive as the number of descent steps increases, while implicitly deriving  $\partial \mathbf{w}^* / \partial \lambda$  requires using Hessian-vector products with conjugate gradient solvers to avoid directly computing the Hessian.

**Model-Based Hyperparameter Optimization.** A common model-based approach is Bayesian optimization, which models  $p(r|\lambda, \mathcal{D})$ , the conditional probability of the performance on some metric  $r$  given hyperparameters  $\lambda$  and a dataset  $\mathcal{D} = \{(\lambda_i, r_i)\}$ . We can model  $p(r|\lambda, \mathcal{D})$  with various methods (Hutter et al., 2011; Bergstra et al., 2011; Snoek et al., 2012; 2015).  $\mathcal{D}$  is constructed iteratively, where the next  $\lambda$  to train on is chosen by maximizing an acquisition function  $C(\lambda; p(r|\lambda, \mathcal{D}))$  which balances exploration and exploitation. Training each model to completion can be avoided if assumptions are made on learning curve behavior (Swersky et al., 2014; Klein et al., 2017). These approaches require building inductive biases into  $p(r|\lambda, \mathcal{D})$  which may not hold in practice, do not take advantage of the network structure when used for hyperparameter optimization, and do not scale well with the number of hyperparameters. However, these approaches have consistency guarantees in the limit, unlike ours.

**Model-Free Hyperparameter Optimization.** Model-free approaches include grid search and random search. Bergstra & Bengio (2012) advocated using random search over grid search. Successive Halving (Jamieson & Talwalkar, 2016) and Hyperband (Li et al., 2017) extend random search by adaptively allocating resources to promising configurations using multi-armed bandit techniques. These methods ignore structure in the problem, unlike ours which uses rich gradient information. However, it is trivial to parallelize model-free methods over computing resources and they tend to perform well in practice.

**Hyperparameter Scheduling.** Population Based Training (PBT) (Jaderberg et al., 2017) considers schedules for hyperparameters. In PBT, a population of networks is trained in parallel. The performance of each network is evaluated periodically, and the weights of under-performing networks are replaced by the weights of better-performing ones; the hyperparameters of the better network are also copied and randomly perturbed for training the new network clone. In this way, a single model can experience different hyperparameter settings over the course of training, implementing a schedule. STNs replace the population of networks by a single best-response approximation and use gradients to tune hyperparameters during a single training run.

## 6 CONCLUSION

We introduced Self-Tuning Networks (STNs), which efficiently approximate the best-response of parameters to hyperparameters by scaling and shifting their hidden units. This allowed us to use gradient-based optimization to tune various regularization hyperparameters, including discrete hyperparameters. We showed that STNs discover hyperparameter schedules that can outperform fixed hyperparameters. We validated the approach on large-scale problems and showed that STNs achieve better generalization performance than competing approaches, in less time. We believe STNs offer a compelling path towards large-scale, automated hyperparameter tuning for neural networks.

### ACKNOWLEDGMENTS

We thank Matt Johnson for helpful discussions and advice. MM is supported by an NSERC CGS-M award, and PV is supported by an NSERC PGS-D award. RG acknowledges support from the CIFAR Canadian AI Chairs program.

## REFERENCES

- Eugene L Allgower and Kurt Georg. *Numerical Continuation Methods: An Introduction*, volume 13. Springer Science & Business Media, 2012.
- Jonathan F Bard. *Practical Bilevel Optimization: Algorithms and Applications*, volume 30. Springer Science & Business Media, 2013.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *International Conference on Machine Learning*, pp. 41–48. ACM, 2009.
- James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13:281–305, 2012.
- James S. Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems*, pp. 2546–2554. 2011.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424*, 2015.
- Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. SMASH: One-shot model architecture search through hypernetworks. *arXiv preprint arXiv:1708.05344*, 2017.
- Luca Calatroni, Cao Chung, Juan Carlos De Los Reyes, Carola-Bibiane Schönlieb, and Tuomo Valkonen. Bilevel approaches for learning of variational imaging models. *arXiv preprint arXiv:1505.02120*, 2015.
- Olivier Chapelle, Vladimir Vapnik, Olivier Bousquet, and Sayan Mukherjee. Choosing multiple parameters for Support Vector Machines. *Machine Learning*, 46(1-3):131–159, 2002.
- Benoît Colson, Patrice Marcotte, and Gilles Savard. A trust-region method for nonlinear bilevel programming: Algorithm and computational experience. *Computational Optimization and Applications*, 30(3):211–227, 2005.
- Benoît Colson, Patrice Marcotte, and Gilles Savard. An overview of bilevel optimization. *Annals of Operations Research*, 153(1):235–256, 2007.
- Misha Denil, Babak Shakibi, Laurent Dinh, Nando De Freitas, et al. Predicting parameters in deep learning. In *Advances in Neural Information Processing Systems*, pp. 2148–2156, 2013.
- Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- Justin Domke. Generic methods for optimization-based modeling. In *Proceedings of Machine Learning Research*, pp. 318–326, 2012.
- John Duchi. Properties of the trace and matrix derivatives, 2007. URL [https://web.stanford.edu/~jduchi/projects/matrix\\_prop.pdf](https://web.stanford.edu/~jduchi/projects/matrix_prop.pdf).
- Anthony V Fiacco and Yo Ishizuka. Sensitivity and stability analysis for nonlinear programming. *Annals of Operations Research*, 27(1):215–235, 1990.
- Chuan-sheng Foo, Chuong B Do, and Andrew Y Ng. Efficient multiple hyperparameter learning for log-linear models. In *Advances in Neural Information Processing Systems*, pp. 377–384, 2008.
- Luca Franceschi, Paolo Frasconi, Saverio Salzo, and Massimiliano Pontil. Bilevel programming for hyperparameter optimization and meta-learning. *arXiv preprint arXiv:1806.04910*, 2018.
- Yarin Gal and Zoubin Ghahramani. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in Neural Information Processing Systems*, pp. 1027–1035, 2016.
- Robert Gibbons. *A Primer in Game Theory*. Harvester Wheatsheaf, 1992.

- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2672–2680, 2014.
- David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
- Pierre Hansen, Brigitte Jaumard, and Gilles Savard. New branch-and-bound rules for linear bilevel programming. *SIAM Journal on Scientific and Statistical Computing*, 13(5):1194–1217, 1992.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer Series in Statistics. Springer, New York, NY, USA, 2001.
- Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. Beta-VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8): 1735–1780, 1997.
- Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *International Conference on Learning and Intelligent Optimization*, pp. 507–523. Springer, 2011.
- Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, et al. Population-based training of neural networks. *arXiv preprint arXiv:1711.09846*, 2017.
- Kevin Jamieson and Ameet Talwalkar. Non-stochastic best arm identification and hyperparameter optimization. In *International Conference on Artificial Intelligence and Statistics*, 2016.
- Mohammad Emtiyaz Khan, Didrik Nielsen, Voot Tangkaratt, Wu Lin, Yarin Gal, and Akash Srivastava. Fast and scalable Bayesian deep learning by weight-perturbation in Adam. *arXiv preprint arXiv:1806.04854*, 2018.
- Aaron Klein, Stefan Falkner, Jost Tobias Springenberg, and Frank Hutter. Learning curve prediction with Bayesian neural networks. In *International Conference on Learning Representations*, 2017.
- Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. In *Technical Report, University of Toronto*, 2009.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pp. 1097–1105, 2012.
- Karl Kunisch and Thomas Pock. A bilevel optimization approach for parameter learning in variational models. *SIAM Journal on Imaging Sciences*, 6(2):938–983, 2013.
- Jan Larsen, Lars Kai Hansen, Claus Svarer, and M Ohlsson. Design and regularization of neural networks: The optimal use of a validation set. In *IEEE Workshop on Neural Networks for Signal Processing*, pp. 62–71. IEEE, 1996.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: Bandit-based configuration evaluation for hyperparameter optimization. *Journal of Machine Learning Research*, 18(1):6765–6816, 2017.
- Jonathan Lorraine and David Duvenaud. Stochastic hyperparameter optimization through hypernetworks. *arXiv preprint arXiv:1802.09419*, 2018.
- Jelena Luketina, Mathias Berglund, Klaus Greff, and Tapani Raiko. Scalable gradient-based tuning of continuous regularization hyperparameters. In *International Conference on Machine Learning*, pp. 2952–2960, 2016.

- Dougal Maclaurin, David Duvenaud, and Ryan Adams. Gradient-based hyperparameter optimization through reversible learning. In *International Conference on Machine Learning*, pp. 2113–2122, 2015.
- Mitchell P Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2):313–330, 1993.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. Regularizing and optimizing LSTM language models. *International Conference on Learning Representations*, 2018.
- Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. The numerics of GANs. In *Advances in Neural Information Processing Systems*, pp. 1825–1835, 2017.
- Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163*, 2016.
- Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*, volume 87. Springer Science & Business Media, 2013.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *Advances in Neural Information Processing Workshop*, 2017.
- Fabian Pedregosa. Hyperparameter optimization with approximate gradient. In *International Conference on Machine Learning*, pp. 737–746, 2016.
- Jürgen Schmidhuber. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation*, 4(1):131–139, 1992.
- Jürgen Schmidhuber. A ‘self-referential’ weight matrix. In *International Conference on Artificial Neural Networks*, pp. 446–450. Springer, 1993.
- Matthias Seeger. Cross-validation optimization for large scale hierarchical classification kernel methods. In *Advances in Neural Information Processing Systems*, pp. 1233–1240, 2007.
- Ankur Sinha, Pekka Malo, and Kalyanmoy Deb. Efficient evolutionary algorithm for single-objective bilevel optimization. *arXiv preprint arXiv:1303.3901*, 2013.
- Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, pp. 2951–2959, 2012.
- Jasper Snoek, Oren Rippel, Kevin Swersky, Ryan Kiros, Nadathur Satish, Narayanan Sundaram, Mostofa Patwary, M Prabhat, and Ryan Adams. Scalable Bayesian optimization using deep neural networks. In *International Conference on Machine Learning*, pp. 2171–2180, 2015.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Joe Staines and David Barber. Variational optimization. *arXiv preprint arXiv:1212.4507*, 2012.
- Kevin Swersky, Jasper Snoek, and Ryan Prescott Adams. Freeze-thaw Bayesian optimization. *arXiv preprint arXiv:1406.3896*, 2014.
- Luis Vicente, Gilles Savard, and Joaquim Júdice. Descent approaches for quadratic bilevel programming. *Journal of Optimization Theory and Applications*, 81(2):379–399, 1994.
- Heinrich Von Stackelberg. *Market Structure and Equilibrium*. Springer Science & Business Media, 2010.
- Li Wan, Matthew Zeiler, Sixin Zhang, Yann Le Cun, and Rob Fergus. Regularization of neural networks using Dropconnect. In *International Conference on Machine Learning*, pp. 1058–1066, 2013.

Yeming Wen, Paul Vicol, Jimmy Ba, Dustin Tran, and Roger Grosse. Flipout: Efficient pseudo-independent weight perturbations on mini-batches. *International Conference on Learning Representations*, 2018.

Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256, 1992.

Zichao Yang, Marcin Moczulski, Misha Denil, Nando de Freitas, Alex Smola, Le Song, and Ziyu Wang. Deep fried convnets. In *International Conference on Computer Vision*, pp. 1476–1483, 2015.

Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*, 2014.

Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016.



## A TABLE OF NOTATION

Table 3: Table of Notation

$\lambda, \mathbf{w}$	Hyperparameters and parameters
$\lambda_0, \mathbf{w}_0$	Current, fixed hyperparameters and parameters
$n, m$	Hyperparameter and elementary parameter dimension
$f(\lambda, \mathbf{w}), F(\lambda, \mathbf{w})$	Lower-level & upper-level objective
$r$	Function mapping unconstrained hyperparameters to the appropriate restricted space
$\mathcal{L}_T(\lambda, \mathbf{w}), \mathcal{L}_V(\lambda, \mathbf{w})$	Training loss & validation loss - $(\mathcal{L}_T(r(\lambda), \mathbf{w}), \mathcal{L}_V(r(\lambda), \mathbf{w})) = (f(\lambda, \mathbf{w}), F(\lambda, \mathbf{w}))$
$\mathbf{w}^*(\lambda)$	Best-response of the parameters to the hyperparameters
$F^*(\lambda)$	Single-level objective from best-response, equals $F(\lambda, \mathbf{w}^*(\lambda))$
$\lambda^*$	Optimal hyperparameters
$\hat{\mathbf{w}}_\phi(\lambda)$	Parametric approximation to the best-response function
$\phi$	Approximate best-response parameters
$\sigma$	Scale of the hyperparameter noise distribution
$\sigma$	The sigmoid function
$\epsilon$	Sampled perturbation noise, to be added to hyperparameters
$p(\epsilon \sigma), p(\lambda \sigma)$	The noise distribution and induced hyperparameter distribution
$\alpha$	A learning rate
$T_{train}, T_{valid}$	Number of training steps on the training and validation data
$\mathbf{x}, t$	An input datapoint and its associated target
$\mathcal{D}$	A data set consisting of tuples of inputs and targets
$D$	The dimensionality of input data
$y(\mathbf{x}, \mathbf{w})$	Prediction function for input data $\mathbf{x}$ and elementary parameters $\mathbf{w}$
$\odot_{\text{row}}$	Row-wise rescaling - <i>not elementwise multiplication</i>
$\mathbf{Q}, \mathbf{s}$	First and second layer weights of the linear network in Problem 13
$\mathbf{Q}_0, \mathbf{s}_0$	The basis change matrix and solution to the unregularized Problem 13
$\mathbf{Q}^*(\lambda), \mathbf{s}^*(\lambda)$	The best response weights of the linear network in Problem 13
$\mathbf{a}(\mathbf{x}; \mathbf{w})$	Activations of hidden units in the linear network of Problem 13
$\mathbf{W}, \mathbf{b}$	A layer's weight matrix and bias
$D_{out}, D_{in}$	A layer's input dimensionality and output dimensionality
$\frac{\partial \mathcal{L}_T(\lambda, \mathbf{w})}{\partial \lambda}$	The (validation loss) direct (hyperparameter) gradient
$\frac{\partial \mathbf{w}^*(\lambda)}{\partial \lambda}$	The (elementary parameter) response gradient
$\frac{\partial \mathcal{L}_T(\lambda, \mathbf{w}^*(\lambda))}{\partial \mathbf{w}^*(\lambda)} \frac{\partial \mathbf{w}^*(\lambda)}{\partial \lambda}$	The (validation loss) response gradient
$\frac{d \mathcal{L}_T(\lambda, \mathbf{w})}{d \lambda}$	The hyperparameter gradient: a sum of the validation losses direct and response gradients

## B PROOFS

### B.1 LEMMA 1

Because  $\mathbf{w}_0$  solves Problem 4b given  $\lambda_0$ , by the first-order optimality condition we must have:

$$\frac{\partial f}{\partial \mathbf{w}}(\lambda_0, \mathbf{w}_0) = 0 \quad (16)$$

The Jacobian of  $\partial f / \partial \mathbf{w}$  decomposes as a block matrix with sub-blocks given by:

$$\left[ \begin{array}{c|c} \frac{\partial^2 f}{\partial \lambda \partial \mathbf{w}} & \frac{\partial^2 f}{\partial \mathbf{w}^2} \end{array} \right] \quad (17)$$

We know that  $f$  is  $\mathcal{C}^2$  in some neighborhood of  $(\lambda_0, \mathbf{w}_0)$ , so  $\partial f / \partial \mathbf{w}$  is continuously differentiable in this neighborhood. By assumption, the Hessian  $\partial^2 f / \partial \mathbf{w}^2$  is positive definite and hence invertible at  $(\lambda_0, \mathbf{w}_0)$ . By the Implicit Function Theorem, there exists a neighborhood  $V$  of  $\lambda_0$  and a unique continuously differentiable function  $\mathbf{w}^* : V \rightarrow \mathbb{R}^m$  such that  $\partial f / \partial \mathbf{w}(\lambda, \mathbf{w}^*(\lambda)) = 0$  for  $\lambda \in V$  and  $\mathbf{w}^*(\lambda_0) = \mathbf{w}_0$ .

Furthermore, by continuity we know that there is a neighborhood  $W_1 \times W_2$  of  $(\lambda_0, \mathbf{w}_0)$  such that  $\partial^2 f / \partial \mathbf{w}^2$  is positive definite on this neighborhood. Setting  $U = V \cap W_1 \cap (\mathbf{w}^*)^{-1}(W_2)$ , we can conclude that  $\partial^2 f / \partial \mathbf{w}^2(\lambda, \mathbf{w}^*(\lambda)) \succ 0$  for all  $\lambda \in U$ . Combining this with  $\partial f / \partial \mathbf{w}(\lambda, \mathbf{w}^*(\lambda)) = 0$  and using second-order sufficient optimality conditions, we conclude that  $\mathbf{w}^*(\lambda)$  is the unique solution to Problem 4b for all  $\lambda \in U$ .

### B.2 LEMMA 2

This discussion mostly follows from Hastie et al. (2001). We let  $\mathbf{X} \in \mathbb{R}^{N \times D}$  denote the data matrix where  $N$  is the number of training examples and  $D$  is the dimensionality of the data. We let  $\mathbf{t} \in \mathbb{R}^N$  denote the associated targets. We can write the SVD decomposition of  $\mathbf{X}$  as:

$$\mathbf{X} = \mathbf{U} \mathbf{D} \mathbf{V}^\top \quad (18)$$

where  $\mathbf{U}$  and  $\mathbf{V}$  are  $N \times D$  and  $D \times D$  orthogonal matrices and  $\mathbf{D}$  is a diagonal matrix with entries  $d_1 \geq d_2 \geq \dots \geq d_D > 0$ . We next simplify the function  $y(\mathbf{x}; \mathbf{w})$  by setting  $\mathbf{u} = \mathbf{s}^\top \mathbf{Q}$ , so that  $y(\mathbf{x}; \mathbf{w}) = \mathbf{s}^\top \mathbf{Q} \mathbf{x} = \mathbf{u}^\top \mathbf{x}$ . We see that the Jacobian  $\partial y / \partial \mathbf{x} \equiv \mathbf{u}$  is constant, and Problem 13 simplifies to standard  $L_2$ -regularized least-squares linear regression with the following loss function:

$$\sum_{(\mathbf{x}, t) \in \mathcal{D}} (\mathbf{u}^\top \mathbf{x} - t)^2 + \frac{1}{|\mathcal{D}|} \exp(\lambda) \|\mathbf{u}\|^2 \quad (19)$$

It is well-known (see Hastie et al. (2001), Chapter 3) that the optimal solution  $\mathbf{u}^*(\lambda)$  minimizing Equation 19 is given by:

$$\mathbf{u}^*(\lambda) = (\mathbf{X}^\top \mathbf{X} + \exp(\lambda) \mathbf{I})^{-1} \mathbf{X}^\top \mathbf{t} = \mathbf{V} (\mathbf{D}^2 + \exp(\lambda) \mathbf{I})^{-1} \mathbf{D} \mathbf{U}^\top \mathbf{t} \quad (20)$$

Furthermore, the optimal solution  $\mathbf{u}^*$  to the unregularized version of Problem 19 is given by:

$$\mathbf{u}^* = \mathbf{V} \mathbf{D}^{-1} \mathbf{U}^\top \mathbf{t} \quad (21)$$

Recall that we defined  $\mathbf{Q}_0 = \mathbf{V}^\top$ , i.e., the change-of-basis matrix from the standard basis to the principal components of the data matrix, and we defined  $\mathbf{s}_0$  to solve the unregularized regression problem given  $\mathbf{Q}_0$ . Thus, we require that  $\mathbf{Q}_0^\top \mathbf{s}_0 = \mathbf{u}^*$  which implies  $\mathbf{s}_0 = \mathbf{D}^{-1} \mathbf{U}^\top \mathbf{t}$ .

There are not unique solutions to Problem 13, so we take any functions  $\mathbf{Q}(\lambda), \mathbf{s}(\lambda)$  which satisfy  $\mathbf{Q}(\lambda)^\top \mathbf{s}(\lambda) = \mathbf{v}^*(\lambda)$  as “best-response functions”. We will show that our chosen functions  $\mathbf{Q}^*(\lambda) = \sigma(\lambda \mathbf{v} + \mathbf{c}) \odot_{\text{row}} \mathbf{Q}_0$  and  $\mathbf{s}^*(\lambda) = \mathbf{s}_0$ , where  $\mathbf{v} = -\mathbf{1}$  and  $c_i = 2 \log(d_i)$  for  $i = 1, \dots, D$ , meet this criteria. We start by noticing that for any  $d \in \mathbb{R}_+$ , we have:

$$\sigma(-\lambda + 2 \log(d)) = \frac{1}{1 + \exp(\lambda - 2 \log(d))} = \frac{1}{1 + d^{-2} \exp(\lambda)} = \frac{d^2}{d^2 + \exp(\lambda)} \quad (22)$$

It follows that:

$$\mathbf{Q}^*(\lambda)^\top \mathbf{s}^*(\lambda) = [\sigma(\lambda \mathbf{v} + \mathbf{c}) \odot_{\text{row}} \mathbf{Q}_0]^\top \mathbf{s}_0 \quad (23)$$

$$= \left[ \text{diag} \begin{pmatrix} \sigma(-\lambda + 2 \log(d_1)) \\ \vdots \\ \sigma(-\lambda + 2 \log(d_D)) \end{pmatrix} \mathbf{Q}_0 \right]^\top \mathbf{s}_0 \quad (24)$$

$$= \mathbf{Q}_0^\top \left[ \text{diag} \begin{pmatrix} \frac{d_1^2}{d_1^2 + \exp(\lambda)} \\ \vdots \\ \frac{d_D^2}{d_D^2 + \exp(\lambda)} \end{pmatrix} \right] \mathbf{s}_0 \quad (25)$$

$$= \mathbf{V} \left[ \text{diag} \begin{pmatrix} \frac{d_1^2}{d_1^2 + \exp(\lambda)} \\ \vdots \\ \frac{d_D^2}{d_D^2 + \exp(\lambda)} \end{pmatrix} \right] \mathbf{D}^{-1} \mathbf{U}^\top \mathbf{t} \quad (26)$$

$$= \mathbf{V} \left[ \text{diag} \begin{pmatrix} \frac{d_1}{d_1^2 + \exp(\lambda)} \\ \vdots \\ \frac{d_D}{d_D^2 + \exp(\lambda)} \end{pmatrix} \right] \mathbf{U}^\top \mathbf{t} \quad (27)$$

$$= \mathbf{V} [(\mathbf{D}^2 + \exp(\lambda) \mathbf{I})^{-1} \mathbf{D}] \mathbf{U}^\top \mathbf{t} \quad (28)$$

$$= \mathbf{v}^*(\lambda) \quad (29)$$

### B.3 THEOREM 3

By assumption  $f$  is quadratic, so there exist  $\mathbf{A} \in \mathbb{R}^{n \times n}$ ,  $\mathbf{B} \in \mathbb{R}^{n \times m}$ ,  $\mathbf{C} \in \mathbb{R}^{m \times m}$  and  $\mathbf{d} \in \mathbb{R}^n$ ,  $\mathbf{e} \in \mathbb{R}^m$  such that:

$$f(\boldsymbol{\lambda}, \mathbf{w}) = \frac{1}{2} \begin{pmatrix} \boldsymbol{\lambda}^\top & \mathbf{w}^\top \end{pmatrix} \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{C} \end{pmatrix} \begin{pmatrix} \boldsymbol{\lambda} \\ \mathbf{w} \end{pmatrix} + \mathbf{d}^\top \boldsymbol{\lambda} + \mathbf{e}^\top \mathbf{w} \quad (30)$$

One can easily compute that:

$$\frac{\partial f}{\partial \mathbf{w}}(\boldsymbol{\lambda}, \mathbf{w}) = \mathbf{B}^\top \boldsymbol{\lambda} + \mathbf{C} \mathbf{w} + \mathbf{e} \quad (31)$$

$$\frac{\partial^2 f}{\partial \mathbf{w}^2}(\boldsymbol{\lambda}, \mathbf{w}) = \mathbf{C} \quad (32)$$

Since we assume  $\partial^2 f / \partial \mathbf{w}^2} \succ \mathbf{0}$ , we must have  $\mathbf{C} \succ \mathbf{0}$ . Setting the derivative equal to  $\mathbf{0}$  and using second-order sufficient conditions, we have:

$$\mathbf{w}^*(\boldsymbol{\lambda}) = -\mathbf{C}^{-1}(\mathbf{e} + \mathbf{B}^\top \boldsymbol{\lambda}) \quad (33)$$

Hence, we find:

$$\frac{\partial \mathbf{w}^*}{\partial \boldsymbol{\lambda}}(\boldsymbol{\lambda}) = -\mathbf{C}^{-1} \mathbf{B}^\top \quad (34)$$

We let  $\hat{\mathbf{w}}_\phi(\boldsymbol{\lambda}) = \mathbf{U} \boldsymbol{\lambda} + \mathbf{b}$ , and define  $\hat{f}$  to be the function given by:

$$\hat{f}(\boldsymbol{\lambda}, \mathbf{U}, \mathbf{b}, \sigma) = \mathbb{E}_{\boldsymbol{\epsilon} \sim p(\boldsymbol{\epsilon}|\sigma)} [f(\boldsymbol{\lambda} + \boldsymbol{\epsilon}, \mathbf{U}(\boldsymbol{\lambda} + \boldsymbol{\epsilon}) + \mathbf{b})] \quad (35)$$

Substituting and simplifying:

$$\begin{aligned} \hat{f}(\boldsymbol{\lambda}_0, \mathbf{U}, \mathbf{b}, \sigma) &= \mathbb{E}_{\boldsymbol{\epsilon} \sim p(\boldsymbol{\epsilon}|\sigma)} \left[ \frac{1}{2} (\boldsymbol{\lambda}_0 + \boldsymbol{\epsilon})^\top \mathbf{A} (\boldsymbol{\lambda}_0 + \boldsymbol{\epsilon}) + (\boldsymbol{\lambda}_0 + \boldsymbol{\epsilon})^\top \mathbf{B} (\mathbf{U}(\boldsymbol{\lambda}_0 + \boldsymbol{\epsilon}) + \mathbf{b}) \right. \\ &\quad \left. + \frac{1}{2} (\mathbf{U}(\boldsymbol{\lambda}_0 + \boldsymbol{\epsilon}) + \mathbf{b})^\top \mathbf{C} (\mathbf{U}(\boldsymbol{\lambda}_0 + \boldsymbol{\epsilon}) + \mathbf{b}) \right. \\ &\quad \left. + \mathbf{d}^\top (\boldsymbol{\lambda}_0 + \boldsymbol{\epsilon}) + \mathbf{e}^\top (\mathbf{U}(\boldsymbol{\lambda}_0 + \boldsymbol{\epsilon}) + \mathbf{b}) \right] \quad (36) \end{aligned}$$

Expanding, we find that equation 36 is equal to:

$$\mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)} \left[ \textcircled{1} + \textcircled{2} + \textcircled{3} + \textcircled{4} \right] \quad (37)$$

where we have:

$$\textcircled{1} = \frac{1}{2} (\lambda_0^\top \mathbf{A} \lambda_0 + 2\epsilon^\top \mathbf{A} \lambda_0 + \epsilon^\top \mathbf{A} \epsilon) \quad (38)$$

$$\textcircled{2} = \lambda_0^\top \mathbf{B} \mathbf{U} \lambda_0 + \lambda_0^\top \mathbf{B} \mathbf{U} \epsilon + \lambda_0^\top \mathbf{B} \mathbf{b} + \epsilon^\top \mathbf{B} \mathbf{U} \lambda_0 + \epsilon^\top \mathbf{B} \mathbf{U} \epsilon + \epsilon^\top \mathbf{B} \mathbf{b} \quad (39)$$

$$\begin{aligned} \textcircled{3} = \frac{1}{2} (\lambda_0^\top \mathbf{U}^\top \mathbf{C} \mathbf{U} \lambda_0 + \lambda_0 \mathbf{U}^\top \mathbf{C} \mathbf{U} \epsilon + \lambda_0 \mathbf{U}^\top \mathbf{C} \mathbf{b} + \epsilon^\top \mathbf{U}^\top \mathbf{C} \mathbf{U} \lambda_0 \\ + \epsilon^\top \mathbf{U}^\top \mathbf{C} \mathbf{U} \epsilon + \epsilon^\top \mathbf{U}^\top \mathbf{C} \mathbf{b} + \mathbf{b}^\top \mathbf{C} \mathbf{U} \lambda_0 + \mathbf{b}^\top \mathbf{C} \mathbf{U} \epsilon + \mathbf{b}^\top \mathbf{C} \mathbf{b}) \end{aligned} \quad (40)$$

$$\textcircled{4} = \mathbf{d}^\top \lambda_0 + \mathbf{d}^\top \epsilon + \mathbf{e}^\top \mathbf{U} \lambda_0 + \mathbf{e}^\top \mathbf{U} \epsilon + \mathbf{e}^\top \mathbf{b} \quad (41)$$

We can simplify these expressions considerably by using linearity of expectation and that  $\epsilon \sim p(\epsilon|\sigma)$  has mean  $\mathbf{0}$ :

$$\mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)} \left[ \textcircled{1} \right] = \frac{1}{2} \lambda_0^\top \mathbf{A} \lambda_0 \quad (42)$$

$$\mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)} \left[ \textcircled{2} \right] = \lambda_0^\top \mathbf{B} \mathbf{U} \lambda_0 + \lambda_0^\top \mathbf{B} \mathbf{b} + \mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)} \left[ \epsilon^\top \mathbf{B} \mathbf{U} \epsilon \right] \quad (43)$$

$$\begin{aligned} \mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)} \left[ \textcircled{3} \right] = \frac{1}{2} (\lambda_0^\top \mathbf{U}^\top \mathbf{C} \mathbf{U} \lambda_0 + \lambda_0 \mathbf{U}^\top \mathbf{C} \mathbf{b} + \\ \mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)} \left[ \epsilon^\top \mathbf{U}^\top \mathbf{C} \mathbf{U} \epsilon \right] + \mathbf{b}^\top \mathbf{C} \mathbf{U} \lambda_0 + \mathbf{b}^\top \mathbf{C} \mathbf{b}) \end{aligned} \quad (44)$$

$$\mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)} \left[ \textcircled{4} \right] = \mathbf{d}^\top \lambda_0 + \mathbf{e}^\top \mathbf{U} \lambda_0 + \mathbf{e}^\top \mathbf{b} \quad (45)$$

We can use the cyclic property of the Trace operator,  $\mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)} [\epsilon \epsilon^\top] = \sigma^2 \mathbf{I}$ , and commutability of expectation and a linear operator to simplify the expectations of  $\textcircled{2}$  and  $\textcircled{3}$ :

$$\mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)} \left[ \textcircled{2} \right] = \lambda_0^\top \mathbf{B} \mathbf{U} \lambda_0 + \lambda_0^\top \mathbf{B} \mathbf{b} + \text{Tr} [\sigma^2 \mathbf{B} \mathbf{U}] \quad (46)$$

$$\begin{aligned} \mathbb{E}_{\epsilon \sim p(\epsilon|\sigma)} \left[ \textcircled{3} \right] = \frac{1}{2} (\lambda_0^\top \mathbf{U}^\top \mathbf{C} \mathbf{U} \lambda_0 + \lambda_0 \mathbf{U}^\top \mathbf{C} \mathbf{b} + \\ \text{Tr} [\sigma^2 \mathbf{U}^\top \mathbf{C} \mathbf{U}] + \mathbf{b}^\top \mathbf{C} \mathbf{U} \lambda_0 + \mathbf{b}^\top \mathbf{C} \mathbf{b}) \end{aligned} \quad (47)$$

We can then differentiate  $\hat{f}$  by making use of various matrix-derivative equalities (Duchi, 2007) to find:

$$\frac{\partial \hat{f}}{\partial \mathbf{b}} (\lambda_0, \mathbf{U}, \mathbf{b}, \sigma) = \frac{1}{2} \mathbf{C}^\top \mathbf{U} \lambda_0 + \frac{1}{2} \mathbf{C} \mathbf{U} \lambda_0 + \mathbf{B}^\top \lambda_0 + \mathbf{e} + \mathbf{C} \mathbf{b} \quad (48)$$

$$\frac{\partial \hat{f}}{\partial \mathbf{U}} (\lambda_0, \mathbf{U}, \mathbf{b}, \sigma) = \mathbf{B}^\top \lambda_0 \lambda_0^\top + \sigma^2 \mathbf{B}^\top + \mathbf{C} \mathbf{b} \lambda_0^\top + \mathbf{e} \lambda_0^\top + \mathbf{C} \mathbf{U} \lambda_0 \lambda_0^\top + \sigma^2 \mathbf{C} \mathbf{U} \quad (49)$$

Setting the derivative  $\partial \hat{f} / \partial \mathbf{b} (\lambda_0, \mathbf{U}, \mathbf{b}, \sigma)$  equal to  $\mathbf{0}$ , we have:

$$\mathbf{b} = -\mathbf{C}^{-1} (\mathbf{C}^\top \mathbf{U} \lambda_0 + \mathbf{B}^\top \lambda_0 + \mathbf{e}) \quad (50)$$

Setting the derivative for  $\partial \hat{f} / \partial \mathbf{U} (\lambda_0, \mathbf{U}, \mathbf{b}, \sigma)$  equal to  $\mathbf{0}$ , we have:

$$\mathbf{C} \mathbf{U} (\lambda_0 \lambda_0^\top + \sigma^2 \mathbf{I}) = -\mathbf{B}^\top \lambda_0 \lambda_0^\top - \sigma^2 \mathbf{B}^\top - \mathbf{C} \mathbf{b} \lambda_0^\top - \mathbf{e} \lambda_0^\top \quad (51)$$

Substituting the expression for  $\mathbf{b}$  given by equation 50 into equation 51 and simplifying gives:

$$\mathbf{C} \mathbf{U} (\lambda_0 \lambda_0^\top + \sigma^2 \mathbf{I}) = -\sigma^2 \mathbf{B}^\top + \mathbf{C} \mathbf{U} \lambda_0 \lambda_0^\top \quad (52)$$

$$\implies \sigma^2 \mathbf{C} \mathbf{U} = -\sigma^2 \mathbf{B}^\top \quad (53)$$

$$\implies \mathbf{U} = -\mathbf{C}^{-1} \mathbf{B}^\top \quad (54)$$

This is exactly the best-response Jacobian  $\partial \mathbf{w}^* / \partial \lambda (\lambda)$  as given by Equation 34. Substituting  $\mathbf{U} = \mathbf{C}^{-1} \mathbf{B}^\top$  into the equation 50 gives:

$$\mathbf{b} = \mathbf{C}^{-1} \mathbf{B}^\top \lambda_0 - \mathbf{C}^{-1} \mathbf{B}^\top \lambda_0 - \mathbf{C}^{-1} \mathbf{e} \quad (55)$$

This is  $\mathbf{w}^* (\lambda_0) - \partial \mathbf{w}^* / \partial \lambda (\lambda_0)$ , thus the approximate best-response is exactly the first-order Taylor series of  $\mathbf{w}^*$  about  $\lambda_0$ .

#### B.4 BEST-RESPONSE GRADIENT LEMMA

**Lemma 4.** *Under the same conditions as Lemma 1 and using the same notation, for all  $\lambda \in U$ , we have that:*

$$\frac{\partial \mathbf{w}^*}{\partial \lambda}(\lambda) = - \left[ \frac{\partial^2 f}{\partial \mathbf{w}^2}(\lambda, \mathbf{w}^*(\lambda)) \right]^{-1} \frac{\partial^2 f}{\partial \lambda \partial \mathbf{w}}(\lambda, \mathbf{w}^*(\lambda)) \quad (56)$$

*Proof.* Define  $\iota^* : U \rightarrow \mathbb{R}^n \times \mathbb{R}^m$  by  $\iota^*(\lambda) = (\lambda, \mathbf{w}^*(\lambda))$ . By first-order optimality conditions, we know that:

$$\left( \frac{\partial f}{\partial \mathbf{w}} \circ \iota^* \right) (\lambda) = 0 \quad \forall \lambda \in U \quad (57)$$

Hence, for all  $\lambda \in U$ :

$$0 = \frac{\partial}{\partial \lambda} \left( \frac{\partial f}{\partial \mathbf{w}} \circ \iota^* \right) (\lambda) \quad (58)$$

$$= \frac{\partial^2 f}{\partial \mathbf{w}^2}(\iota^*(\lambda)) \frac{\partial \mathbf{w}^*}{\partial \lambda}(\lambda) + \frac{\partial^2 f}{\partial \lambda \partial \mathbf{w}}(\iota^*(\lambda)) \quad (59)$$

$$= \frac{\partial^2 f}{\partial \mathbf{w}^2}(\lambda, \mathbf{w}^*(\lambda)) \frac{\partial \mathbf{w}^*}{\partial \lambda}(\lambda) + \frac{\partial^2 f}{\partial \lambda \partial \mathbf{w}}(\lambda, \mathbf{w}^*(\lambda)) \quad (60)$$

Rearranging gives Equation 56.  $\square$

### C BEST-RESPONSE APPROXIMATIONS FOR CONVOLUTIONAL FILTERS

We let  $L$  denote the number of layers,  $C_l$  the number of channels in layer  $l$ 's feature map, and  $K_l$  the size of the kernel in layer  $l$ . We let  $\mathbf{W}^{l,c} \in \mathbb{R}^{C_{l-1} \times K_l \times K_l}$  and  $\mathbf{b}^{l,c} \in \mathbb{R}$  denote the weight and bias respectively of the  $c^{\text{th}}$  convolution kernel in layer  $l$  (so  $c \in \{1, \dots, C_l\}$ ). For  $\mathbf{u}^{l,c}, \mathbf{a}^{l,c} \in \mathbb{R}^n$ , we define best-response approximations  $\hat{\mathbf{W}}_{\phi}^{l,c}$  and  $\hat{\mathbf{b}}_{\phi}^{l,c}$  by:

$$\hat{\mathbf{W}}_{\phi}^{l,c}(\lambda) = (\lambda^{\top} \mathbf{u}^{l,c}) \odot \mathbf{W}_{\text{hyper}}^{l,c} + \mathbf{W}_{\text{elem}}^{l,c} \quad (61)$$

$$\hat{\mathbf{b}}_{\phi}^{l,c}(\lambda) = (\lambda^{\top} \mathbf{a}^{l,c}) \odot \mathbf{b}_{\text{hyper}}^{l,c} + \mathbf{b}_{\text{elem}}^{l,c} \quad (62)$$

Thus, the best-response parameters used for modeling  $\mathbf{W}^{l,c}$ ,  $\mathbf{b}^l$  are  $\{\mathbf{u}^{l,c}, \mathbf{a}^{l,c}, \mathbf{W}_{\text{hyper}}^{l,c}, \mathbf{W}_{\text{elem}}^{l,c}, \mathbf{b}_{\text{hyper}}^{l,c}, \mathbf{b}_{\text{elem}}^{l,c}\}$ . We can compute the number of parameters used as  $2n + 2(|\mathbf{W}^{l,c}| + |\mathbf{b}^{l,c}|)$ . Summing over channels  $c$ , we find the total number of parameters is  $2nC_l + 2p$ , where  $p$  is the total number of parameters in the normal CNN layer. Hence, we use twice the number of parameters in a normal CNN, plus an overhead that depends on the number of hyperparameters.

For an implementation in code, see Appendix G.

### D LANGUAGE MODELING EXPERIMENT DETAILS

Here we present additional details on the setup of our LSTM language modeling experiments on PTB, and on the role of each hyperparameter we tune.

We trained a 2-layer LSTM with 650 hidden units per layer and 650-dimensional word embeddings (similar to (Zaremba et al., 2014; Gal & Ghahramani, 2016)) on sequences of length 70 in mini-batches of size 40. To optimize the baseline LSTM, we used SGD with initial learning rate 30, which was decayed by a factor of 4 based on the non-monotonic criterion introduced by Merity et al. (2018) (i.e., whenever the validation perplexity fails to improve for 5 epochs). Following Merity et al. (2018), we used gradient clipping 0.25.

To optimize the ST-LSTM, we used the same optimization setup as for the baseline LSTM. For the hyperparameters, we used Adam with learning rate 0.01. We used an alternating training schedule in which we updated the model parameters for 2 steps on the training set and then updated the hyperparameters for 1 step on the validation set. We used one epoch of warm-up, in which we

updated the model parameters, but did not update hyperparameters. We terminated training when the learning rate dropped below 0.0003.

We tuned variational dropout (re-using the same dropout mask for each step in a sequence) on the input to the LSTM, the hidden state between the LSTM layers, and the output of the LSTM. We also tuned embedding dropout, which sets entire rows of the word embedding matrix to 0, effectively removing certain words from all sequences. We regularized the hidden-to-hidden weight matrix using DropConnect (zeroing out weights rather than activations) (Wan et al., 2013). Because DropConnect operates directly on the weights and not individually on the mini-batch elements, we cannot use independent perturbations per example; instead, we sample a single DropConnect rate per mini-batch. Finally, we used activation regularization (AR) and temporal activation regularization (TAR). AR penalizes large activations, and is defined as:

$$\alpha \|m \odot h_t\|_2 \quad (63)$$

where  $m$  is a dropout mask and  $h_t$  is the output of the LSTM at time  $t$ . TAR is a slowness regularizer, defined as:

$$\beta \|h_t - h_{t+1}\|_2 \quad (64)$$

For AR and TAR, we tuned the scaling coefficients  $\alpha$  and  $\beta$ . For the baselines, the hyperparameter ranges were:  $[0, 0.95]$  for the dropout rates, and  $[0, 4]$  for  $\alpha$  and  $\beta$ . For the ST-LSTM, all the dropout rates and the coefficients  $\alpha$  and  $\beta$  were initialized to 0.05 (except in Figure 3, where we varied the output dropout rate).

## E IMAGE CLASSIFICATION EXPERIMENT DETAILS

Here, we present additional details on the CNN experiments. For all results, we held out 20% of the training data for validation.

We trained the baseline CNN using SGD with initial learning rate 0.01 and momentum 0.9, on mini-batches of size 128. We decay the learning rate by 10 each time the validation loss fails to decrease for 60 epochs, and end training if the learning rate falls below  $10^{-5}$  or validation loss has not decreased for 75 epochs. For the baselines—grid search, random search, and Bayesian optimization—the search spaces for the hyperparameters were as follows: dropout rates were in the range  $[0, 0.75]$ ; contrast, saturation, and brightness each had range  $[0, 1]$ ; hue had range  $[0, 0.5]$ ; the number of cutout holes had range  $[0, 4]$ , and the length of each cutout hole had range  $[0, 24]$ .

We trained the ST-CNN’s elementary parameters using SGD with initial learning rate 0.01 and momentum of 0.9, on mini-batches of size 128 (identical to the baselines). We use the same decay schedule as the baseline model. The hyperparameters are optimized using Adam with learning rate 0.003. We alternate between training the best-response approximation and hyperparameters with the same schedule as the ST-LSTM, i.e.  $T_{train} = 2$  steps on the training step and  $T_{valid} = 1$  steps on the validation set. Similarly to the LSTM experiments, we used five epochs of warm-up for the model parameters, during which the hyperparameters are fixed. We used an entropy weight of  $\tau = 0.001$  in the entropy regularized objective (Eq. 15). The cutout length was restricted to lie in  $\{0, \dots, 24\}$  while the number of cutout holes was restricted to lie in  $\{0, \dots, 4\}$ . All dropout rates, as well as the continuous data augmentation noise parameters, are initialized to 0.05. The cutout length is initialized to 4, and the number of cutout holes is initialized to 1. Overall, we found the ST-CNN to be relatively robust to the initialization of hyperparameters, but starting with low regularization aided optimization in the first few epochs.

## F ADDITIONAL DETAILS ON HYPERPARAMETER SCHEDULES

Here, we draw connections between hyperparameter schedules and curriculum learning. Curriculum learning (Bengio et al., 2009) is an instance of a family of *continuation methods* (Allgower & Georg, 2012), which optimize non-convex functions by solving a sequence of functions that are ordered by increasing difficulty. In a continuation method, one considers a family of training criteria  $C_\lambda(\mathbf{w})$  with a parameter  $\lambda$ , where  $C_1(\mathbf{w})$  is the final objective we wish to minimize, and  $C_0(\mathbf{w})$  represents the training criterion for a simpler version of the problem. One starts by optimizing  $C_0(\mathbf{w})$  and then gradually increases  $\lambda$  from 0 to 1, while keeping  $\mathbf{w}$  at a local minimum of  $C_\lambda(\mathbf{w})$  (Bengio et al., 2009). This has been hypothesized to both aid optimization and improve generalization. In this



section, we explore how hyperparameter schedules implement a form of curriculum learning; for example, a schedule that increases dropout over time increases stochasticity, making the learning problem more difficult. We use the results of grid searches to understand the effects of different hyperparameter settings throughout training, and show that greedy hyperparameter schedules can outperform fixed hyperparameter values.

First, we performed a grid search over 20 values each of input and output dropout, and measured the validation perplexity in each epoch. Figure 7 shows the validation perplexity achieved by different combinations of input and output dropout, at various epochs during training. We see that at the start of training, the best validation loss is achieved with small values of both input and output dropout. As we train for more epochs, the best validation performance is achieved with larger dropout rates.





Figure 7: **Validation performance of a baseline LSTM given different settings of input and output dropout, at various epochs during training.** (a), (b), and (c) show the validation performance on PTB given different hyperparameter settings, at epochs 1, 10, and 25, respectively. Darker colors represent lower (better) validation perplexity.

Next, we present a simple example to show the potential benefits of greedy hyperparameter schedules. For a single hyperparameter—output dropout—we performed a fine-grained grid search and constructed a dropout schedule by using the hyperparameter values that achieve the best validation perplexity at each epoch in training. As shown in Figure 8, the schedule formed by taking the best output dropout value in each epoch yields better generalization than any of the fixed hyperparameter values from the initial grid search. In particular, by using small dropout values at the start of training, the schedule achieves a fast decrease in validation perplexity, and by using larger dropout later in training, it achieves better overall validation perplexity.



(a) Greedy schedule for output dropout, derived by taking the best hyperparameters in each epoch from a grid search.

(b) Comparison of fixed output dropout values and the dropout schedule derived from grid searches

Figure 8: **Grid search-derived schedule for output dropout.**

Figure 9 shows the perturbed values for output dropout we used to investigate whether the improved performance yielded by STNs is due to the regularization effect, and not the schedule, in Section 4.1.




Figure 9: **Comparison of output dropout schedules.** (a) Gaussian-perturbed output dropout rates around the best value found by grid search, 0.68; (b) sinusoid-perturbed output dropout rates with amplitude 0.1 and a period of 1200 mini-batches; (c) the output dropout schedule found by the ST-LSTM.

## G CODE LISTINGS

In this section, we provide PyTorch code listings for the approximate best-response layers used to construct ST-LSTMs and ST-CNNs: the `HyperLinear` and `HyperConv2D` classes. We also provide a simplified version of the optimization steps used on the training set and validation set.

Listing 1: `HyperLinear`, used as a drop-in replacement for `Linear` modules

```
class HyperLinear(nn.Module):
    def __init__(self, input_dim, output_dim, n_hparams):
        super(HyperLinear, self).__init__()

        self.input_dim = input_dim
        self.output_dim = output_dim
        self.n_hparams = n_hparams
        self.n_scalars = output_dim

        self.elem_w = nn.Parameter(torch.Tensor(output_dim, input_dim))
        self.elem_b = nn.Parameter(torch.Tensor(output_dim))

        self.hnet_w = nn.Parameter(torch.Tensor(output_dim, input_dim))
        self.hnet_b = nn.Parameter(torch.Tensor(output_dim))

        self.htensor_to_scalars = nn.Linear(self.n_hparams,
                                             self.n_scalars*2, bias=False)

        self.init_params()

    def forward(self, input, hnet_tensor):
        output = F.linear(input, self.elem_w, self.elem_b)

        if hnet_tensor is not None:
            hnet_scalars = self.htensor_to_scalars(hnet_tensor)
            hnet_wscalars = hnet_scalars[:, :self.n_scalars]
            hnet_bscalars = hnet_scalars[:, self.n_scalars:]

            hnet_out = hnet_wscalars * F.linear(input, self.hnet_w)
            hnet_out += hnet_bscalars * self.hnet_b

        output += hnet_out

    return output
```

Listing 2: `HyperConv2d`, used as a drop-in replacement for `Conv2d` modules

```
class HyperConv2d(nn.Module):
```

```

def __init__(self, in_channels, out_channels, kernel_size, padding,
             num_hparams,
             stride=1, bias=True):
    super(HyperConv2d, self).__init__()

    self.in_channels = in_channels
    self.out_channels = out_channels
    self.kernel_size = kernel_size
    self.padding = padding
    self.num_hparams = num_hparams
    self.stride = stride

    self.elem_weight = nn.Parameter(torch.Tensor(
        out_channels, in_channels, kernel_size, kernel_size))
    self.hnet_weight = nn.Parameter(torch.Tensor(
        out_channels, in_channels, kernel_size, kernel_size))
    if bias:
        self.elem_bias = nn.Parameter(torch.Tensor(out_channels))
        self.hnet_bias = nn.Parameter(torch.Tensor(out_channels))
    else:
        self.register_parameter('elem_bias', None)
        self.register_parameter('hnet_bias', None)

    self.htensor_to_scalars = nn.Linear(
        self.num_hparams, self.out_channels*2, bias=False)
    self.elem_scalar = nn.Parameter(torch.ones(1))

    self.init_params()

def forward(self, input, htensor):
    """
    Arguments:
        input (tensor): size should be (B, C, H, W)
        htensor (tensor): size should be (B, D)
    """
    output = F.conv2d(input, self.elem_weight, self.elem_bias,
                      padding=self.padding,
                      stride=self.stride)
    output *= self.elem_scalar
    if htensor is not None:
        hnet_scalars = self.htensor_to_scalars(htensor)
        hnet_wscalars = hnet_scalars[:,
            :self.out_channels].unsqueeze(2).unsqueeze(2)
        hnet_bscalars = hnet_scalars[:, self.out_channels:]

        hnet_out = F.conv2d(input, self.hnet_weight,
                             padding=self.padding,
                             stride=self.stride)
        hnet_out *= hnet_wscalars
        if self.hnet_bias is not None:
            hnet_out += (hnet_bscalars *
                         self.hnet_bias).unsqueeze(2).unsqueeze(2)
        output += hnet_out
    return output

def init_params(self):
    n = self.in_channels * self.kernel_size * self.kernel_size
    stdv = 1. / math.sqrt(n)
    self.elem_weight.data.uniform_(-stdv, stdv)
    self.hnet_weight.data.uniform_(-stdv, stdv)
    if self.elem_bias is not None:
        self.elem_bias.data.uniform_(-stdv, stdv)
        self.hnet_bias.data.uniform_(-stdv, stdv)

```

```
self.htensor_to_scalars.weight.data.normal_(std=0.01)
```

Listing 3: Stylized optimization step on the training set for updating elementary parameters

```
# Perturb hyperparameters around current value in unconstrained
# parametrization.
batch_htensor = perturb(htensor, hscale)

# Apply necessary reparametrization of hyperparameters.
hparam_tensor = hparam_transform(batch_htensor)

# Sets data augmentation hyperparameters in the data loader.
dataset.set_hparams(hparam_tensor)

# Get next batch of examples and apply any input transformation
# (e.g. input dropout) as dictated by the hyperparameters.
images, labels = next_batch(dataset)
images = apply_input_transform(images, hparam_tensor)

# Run everything through the model and do gradient descent.
pred = hyper_cnn(images, batch_htensor, hparam_tensor)
xentropy_loss = F.cross_entropy(pred, labels)
xentropy_loss.backward()
cnn_optimizer.step()
```

Listing 4: Stylized optimization step on the validation set for updating hyperparameters/noise scale

```
# Perturb hyperparameters around current value in unconstrained
# parametrization, so we can assess sensitivity of validation
# loss to the scale of the noise.
batch_htensor = perturb(htensor, hscale)

# Apply necessary reparametrization of hyperparameters.
hparam_tensor = hparam_transform(batch_htensor)

# Get next batch of examples and run through the model.
images, labels = next_batch(valid_dataset)
pred = hyper_cnn(images, batch_htensor, hparam_tensor)
xentropy_loss = F.cross_entropy(pred, labels)

# Add extra entropy weight term to loss.
entropy = compute_entropy(hscale)
loss = xentropy_loss - args.entropy_weight * entropy
loss.backward()

# Tune the hyperparameters.
hyper_optimizer.step()

# Tune the scale of the noise applied to hyperparameters.
scale_optimizer.step()
```

## H SENSITIVITY STUDIES

In this section, we present experiments that show how sensitive our STN models are to different meta-parameters.

In particular, we investigate the effect of using alternative schedules (Figure 10) for the number of optimization steps performed on the training and validation sets.

Additionally, we investigate the effect of using different initial perturbation scales for the hyperparameters, which are either fixed or tuned (Figure 11).




Figure 10: **The effect of using a different number of train/val steps.** For the CNN, we include Bayesian Optimization and the reported STN parameters for comparison. During these experiments we found schedules which achieve better final loss with CNNs.




Figure 11: **The effect of using different perturbation scales.** For the CNN, we include Bayesian Optimization and the reported STN parameters for comparison. For (a), the perturbation scales are fixed.