
The Flip Side of the Reweighted Coin: Duality of Adaptive Dropout and Regularization

Daniel LeJeune

Department of Electrical
and Computer Engineering
Rice University
Houston, TX 77005
dlejeune@rice.edu

Hamid Javadi

Department of Electrical
and Computer Engineering
Rice University
Houston, TX 77005
hh35@rice.edu

Richard G. Baraniuk

Department of Electrical
and Computer Engineering
Rice University
Houston, TX 77005
richb@rice.edu

Abstract

Among the most successful methods for sparsifying deep (neural) networks are those that adaptively mask the network weights throughout training. By examining this masking, or *dropout*, in the linear case, we uncover a *duality* between such adaptive methods and regularization through the so-called “ η -trick” that casts both as iteratively reweighted optimizations. We show that any dropout strategy that adapts to the weights in a monotonic way corresponds to an effective subquadratic regularization penalty, and therefore leads to sparse solutions. We obtain the effective penalties for several popular sparsification strategies, which are remarkably similar to classical penalties commonly used in sparse optimization. Considering variational dropout as a case study, we demonstrate similar empirical behavior between the adaptive dropout method and classical methods on the task of deep network sparsification, validating our theory.

1 Introduction

In machine learning, it is often valuable for models to be parsimonious or *sparse* for a variety of reasons, from memory savings and computational speedups to model interpretability and generalizability. Classically, this is achieved by solving a regularized empirical risk minimization (ERM) problem of the form

$$\underset{\mathbf{w}}{\text{minimize}} \quad \mathcal{L}(\mathbf{w}) + \lambda\Omega(\mathbf{w}), \quad (1)$$

where \mathcal{L} is the data-dependent empirical risk and Ω is a sparsity-inducing regularization penalty. Ideally, $\Omega(\mathbf{w})$ is equal to $\|\mathbf{w}\|_0 := \#\{i : w_i \neq 0\}$ or a function of $\|\mathbf{w}\|_0$ [6], in which case an appropriately chosen λ balances the trade-off between the suboptimality of \mathcal{L} and the sparsity of the solution. However, the use of $\|\mathbf{w}\|_0$ directly as a penalty is difficult, as it makes local search impossible, so alternative approaches are necessary for high-dimensional problems. The classical solution to this is to *relax* the problem using a smooth surrogate Ω that approximates $\|\mathbf{w}\|_0$ [5, 18].

When we know the regularization penalty Ω , we can understand the types of solutions they encourage; however, many popular methods for sparsifying deep (neural) networks do not appear to fit the form (1). For example, variational dropout [25, 29], “ ℓ_0 regularization” [26], and magnitude pruning [41] are the only methods compared in a recent large-scale study by Gale et al. [15], and all achieve sparsity by adaptively *masking* or *scaling* the parameters of the network while optimizing the loss function in \mathbf{w} . Although these methods are well-motivated heuristically, it is unclear how to make a principled choice between them aside from treating them as black boxes and comparing their empirical performance. However, if these methods could be cast as solutions to a regularized ERM problem, then we could compare them on the basis of their regularization penalties.

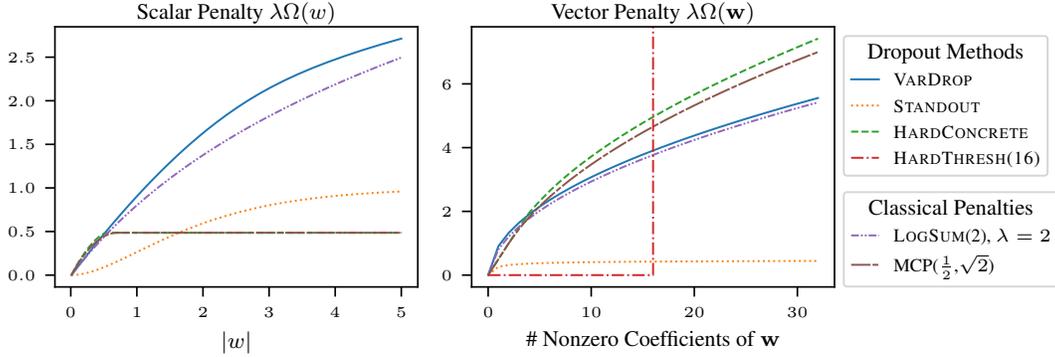


Figure 1: Sparsity-inducing behavior in regularized ERM (1) can be understood by considering the properties of the regularization penalty. Our analysis enables the computation and comparison of the effective regularization penalties of the adaptive dropout sparsity methods in Table 2. (We arbitrarily set $\lambda = 1$ unless otherwise specified.) Both VARIATIONALDROPOUT [29] and HARDCONCRETE (called “ ℓ_0 regularization”) [26] bear strong resemblance to classical penalties. *Left*: Separable penalties (excluding HARDTHRESH) as a function of weight magnitude. *Right*: Penalty of a unit-norm k -sparse \mathbf{w} defined by $w_j = \frac{1}{\sqrt{k}} \mathbb{1}\{j \leq k\}$.

We show that, in fact, these sparsity methods *do* correspond to regularization penalties Ω , which we can obtain, compute, and compare. As we show in Figure 1, these penalties bear striking resemblance to classical sparsity-inducing penalties such as the LOGSUM [9] penalty and the minimax concave penalty (MCP) [40]. More broadly, we analyze *adaptive dropout* methods, which apply the dropout technique of Srivastava et al. [31] with adaptive parameters. By considering adaptive dropout in the linear setting, we uncover a *duality* between adaptive dropout methods and regularized ERM. We make this connection via the “ η -trick” [4], a tool with a long history of application in sparse optimization via iteratively reweighted least squares (IRLS) (see the history presented by Daubechies et al. [11]).

We prove that all adaptive dropout methods whose amount of dropout varies monotonically with the magnitudes of the parameters induce an effective *subquadratic* (and hence sparsity-inducing [5]) penalty Ω . This further supports the experimental evidence that such methods excel at inducing sparsity. We also demonstrate how to use our result to determine the effective penalty for adaptive dropout methods, using as examples the sparsity methods listed above as well as the standout method [3], which has an effective penalty on the layer *activations*, rather than the parameters, explaining why standout sparsifies activations. We then numerically compute the effective penalties¹ and plot them together in Figure 1.

We validate our new theory by applying variational dropout on the task of deep network sparsification, and we show that the performance is similar to non-dropout methods based on the same effective penalty. This suggests that not only is considering the effective penalty a sound means for comparing sparsity methods, but also that classical regularized ERM is itself an effective means for inducing sparsity in deep networks and serves as a valuable tool for future work on deep network sparsity.

For theoreticians, this work provides a general framework for obtaining the effective regularization penalty for adaptive dropout algorithms, enabling regularized ERM analysis tools to be applied to these methods. For practitioners, this work enables the application of classical regularization intuition when choosing between sparsifying approaches. For methods developers, this work provides a strong baseline against which new adaptive dropout methods should be compared: the effective penalty Ω .

2 Background

Notation. We denote the extended reals and the non-negative reals by $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$ and $\mathbb{R}_+ = [0, \infty)$, respectively. The operator $\text{diag}(\cdot)$ takes a vector and returns a matrix with that vector along the diagonal entries or takes a matrix and returns the same matrix with all non-diagonal entries set to zero. The matrix \mathbf{J} denotes a matrix of all ones. We denote element-wise multiplication and exponentiation by \odot ; i.e., $[\mathbf{u} \odot \mathbf{v}]_j = u_j v_j$ and $[\mathbf{U}^{\odot p}]_{ij} = U_{ij}^p$. Division of vectors and scalar functions of vectors denote element-wise operations. Order statistics of the magnitudes of

¹Our code is available at <https://github.com/dlej/adaptive-dropout>.

elements in a vector $\mathbf{u} \in \mathbb{R}^d$ are denoted by $u_{(j)}$ such that $|u_{(1)}| \geq |u_{(2)}| \geq \dots |u_{(d)}|$. The function $\sigma: \mathbb{R} \rightarrow [0, 1]$ denotes the sigmoid $\sigma(t) = 1/(1 + e^{-t})$.

2.1 Dropout as Tikhonov Regularization

We can understand adaptive dropout methods by breaking them down into two components: the regularizing effect of dropout itself and the effect of adaptively updating the dropout-induced regularization. For the former, it is well-known [36, 35, 31, 29] that dropout induces Tikhonov-like regularization for standard binary and Gaussian dropout. The same result holds more broadly given a second moment condition, which we need in order to be able to consider methods like `HARDCONCRETE` “ ℓ_0 regularization” [26]. In general, we use the term *dropout* to refer to methods that mask or scale the weights \mathbf{w} by independent “mask” variables $\mathbf{s} \sim \text{MASK}(\boldsymbol{\alpha})$ with sampling parameters $\boldsymbol{\alpha} \in [0, 1]^d$. We also assume that the masks are *unbiased* (i.e., that $\mathbb{E}[s_j] = 1$), and that $\mathbb{E}[s_j^2] = \alpha_j^{-1}$. Note that with the independence assumption, this implies that $\text{Cov}(\mathbf{s}) = \text{diag}(\boldsymbol{\alpha})^{-1} - \mathbf{I}$. Both standard unbiased binary dropout, where s_j take value α_j^{-1} with probability α_j and 0 otherwise, as well as Gaussian dropout with $s_j \sim \mathcal{N}(1, \alpha_j^{-1} - 1)$ satisfy these assumptions.

Dropout solves the optimization problem

$$\underset{\mathbf{w}}{\text{minimize}} \quad \mathbb{E}_{\mathbf{s}} [\mathcal{L}(\mathbf{s} \odot \mathbf{w})] \quad (2)$$

via stochastic optimization such as gradient descent where each iteration performs a partial optimization of $\mathcal{L}(\mathbf{s} \odot \mathbf{w})$ with a random \mathbf{s} . It is most insightful to consider when \mathcal{L} is the loss for linear regression of target variables $\mathbf{y} \in \mathbb{R}^n$ given data $\mathbf{X} \in \mathbb{R}^{n \times d}$:

$$\mathcal{L}(\mathbf{w}) = \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2. \quad (3)$$

In this case, under our assumptions on \mathbf{s} , we recover the known Tikhonov result

$$\mathbb{E}_{\mathbf{s}} [\mathcal{L}(\mathbf{s} \odot \mathbf{w})] = \mathcal{L}(\mathbf{w}) + \frac{1}{2} \mathbf{w}^\top \left(\left(\text{diag}(\boldsymbol{\alpha})^{-1} - \mathbf{I} \right) \odot \frac{1}{n} \mathbf{X}^\top \mathbf{X} \right) \mathbf{w}. \quad (4)$$

For simplicity, we assume that the data is standardized, or that $\text{diag}(\frac{1}{n} \mathbf{X}^\top \mathbf{X}) = \mathbf{I}$, which is inexpensive to satisfy in practice, but our analysis can be extended to the general case. Under this assumption, we see that dropout elicits a diagonal Tikhonov regularization with scale $\text{diag}(\boldsymbol{\alpha})^{-1} - \mathbf{I}$.

2.2 The “ η -trick”

The second component of adaptive dropout is the adaptive update of the dropout-induced regularization. To understand this, we introduce the so-called “ η -trick” [4] applied to the regularization penalty in (1). By introducing an auxiliary variable $\boldsymbol{\eta} \in \mathcal{H} \subseteq \overline{\mathbb{R}}_+^d$, we can replace $\Omega(\mathbf{w})$ with a dual formulation as a function of $\boldsymbol{\eta}$ that majorizes $\Omega(\mathbf{w})$ and is quadratic in \mathbf{w} for fixed $\boldsymbol{\eta}$. In other words, we can find a function $f: \mathcal{H} \subseteq \overline{\mathbb{R}}_+^d \rightarrow \overline{\mathbb{R}}$ such that

$$\Omega(\mathbf{w}) = \min_{\boldsymbol{\eta} \in \mathcal{H}} \frac{1}{2} \left(\mathbf{w}^\top \text{diag}(\boldsymbol{\eta})^{-1} \mathbf{w} + f(\boldsymbol{\eta}) \right). \quad (5)$$

If such a function f exists, then the regularized ERM problem (1) can be rewritten as a joint optimization in \mathbf{w} and $\boldsymbol{\eta}$ of the dual regularized ERM formulation

$$\underset{\mathbf{w}, \boldsymbol{\eta}}{\text{minimize}} \quad \mathcal{L}(\mathbf{w}) + \frac{\lambda}{2} \left(\mathbf{w}^\top \text{diag}(\boldsymbol{\eta})^{-1} \mathbf{w} + f(\boldsymbol{\eta}) \right). \quad (6)$$

This joint optimization can be performed, for example, by alternating minimization over \mathbf{w} and $\boldsymbol{\eta}$. For linear regression problems, this gives rise to the iteratively reweighted least squares (IRLS) algorithm, a popular method for sparse recovery [10, 11]. We note that iteratively reweighted ℓ_1 schemes have also held a significant place in sparse optimization [43, 9, 38], but because dropout regularization is quadratic, we limit our consideration to iteratively reweighted ℓ_2 (Tikhonov) regularization like (6).

The simplest example of a penalty Ω that fits the form (5) is $\Omega(w) = |w|$, for which $f(\boldsymbol{\eta}) = \eta$, and it is used in the IRLS formulation of ℓ_1 regularization. More broadly, we can consider the class of

Table 1: Common regularization penalties have dual formulations that are often straightforward to obtain. See Table 3 in Appendix A for a more complete table of common penalties along with the derivations.

Penalty	$\Omega(\mathbf{w})$	$f(\boldsymbol{\eta})$	$\hat{\eta}_j(\mathbf{w})$
ℓ_1	$ w_j $	η_j	$ w_j $
LOGSUM(ε) [9]	$\log(w_j + \varepsilon)$	$2 \log\left(\frac{\sqrt{\varepsilon^2 + 4\eta_j} + \varepsilon}{2}\right) - \frac{(\sqrt{\varepsilon^2 + 4\eta_j} - \varepsilon)^2}{4\eta_j}$	$ w_j (w_j + \varepsilon)$
MCP(a, λ) [40]	$\begin{cases} w_j - \frac{w_j^2}{2a\lambda}, & w_j \leq a\lambda \\ \frac{a\lambda}{2}, & w_j > a\lambda \end{cases}$	$\frac{a\lambda\eta_j}{\eta_j + a\lambda}$	$\begin{cases} \frac{a\lambda w_j }{a\lambda - w_j }, & w_j < a\lambda \\ \infty, & w_j \geq a\lambda \end{cases}$
HARDTHRESH(k) [7]	$\infty \mathbb{1}\{\ \mathbf{w}\ _0 > k\}$	$0, \mathcal{H} = \{\boldsymbol{\eta} : \ \boldsymbol{\eta}\ _0 \leq k\}$	$\infty \mathbb{1}\{j \in \text{TOP-}k(\mathbf{w})\}$

functions for which Ω is concave in $\mathbf{w}^{\odot 2}$. These functions are said to be *subquadratic*, and are known to have a sparsity-inducing effect [5]. In fact, these are the only functions Ω that can fit this form. This is because $\mathbf{w}^{\odot 2} \mapsto -2\Omega(\mathbf{w})$ is the Legendre–Fenchel (LF) transform of $-\boldsymbol{\eta}^{\odot -1} \mapsto f(\boldsymbol{\eta})$, and thus Ω is concave in $\mathbf{w}^{\odot 2}$. By the Fenchel–Moreau theorem, there is conversely a unique function f satisfying (5) that is convex in $-\boldsymbol{\eta}^{\odot -1}$, which can be obtained by the LF transform. That is, Ω and f are *dual* functions. Additionally, for differentiable penalties, the minimizing $\boldsymbol{\eta}$ given a fixed \mathbf{w} is

$$\hat{\boldsymbol{\eta}}(\mathbf{w}) = (2\nabla_{\mathbf{w}^{\odot 2}}\Omega(\mathbf{w}))^{\odot -1}, \quad (7)$$

which means that if $\Omega(\mathbf{w})$ has a closed-form expression, then so does the minimizer $\hat{\boldsymbol{\eta}}(\mathbf{w})$. The simplicity of this form means that iteratively reweighted Tikhonov regularization is simple to implement for any differentiable penalty, leading to its popularity.

We enumerate a few penalties along with their dual formulations in Table 1. For separable penalties, we list the scalar function applied to a single element. Where \mathcal{H} is restricted, we include it with $f(\boldsymbol{\eta})$.

Hard Thresholding. Of particular note is HARDTHRESH, the indicator penalty of level sets of $\|\mathbf{w}\|_0$, which gives rise to the iterative hard thresholding (IHT) algorithm [7]. We remark that to the best of our knowledge, what we show in Table 1 is the first characterization of IHT as an iteratively reweighted ℓ_2 solution to a regularized ERM problem, rather than as a heuristic algorithm. This connection is important for understanding the MAGNITUDEPRUNING [41] algorithm.

3 Adaptive Dropout as Iteratively Reweighted Stochastic Optimization

We can now consider an adaptive dropout algorithm [3, 25, 26], which both optimizes the weights \mathbf{w} using dropout and updates the dropout parameters $\boldsymbol{\alpha}$ according to some update rule depending on \mathbf{w} . To facilitate our analysis, we can equivalently consider the algorithm to be updating $\eta_j := \lambda(\alpha_j^{-1} - 1)^{-1}$, which is a monotonic and invertible function of α_j , with

$$\alpha_j = \frac{\eta_j}{\eta_j + \lambda}. \quad (8)$$

An adaptive dropout method could resemble, for example, the following gradient descent strategy with step size parameter ρ given an update function $\hat{\boldsymbol{\eta}} : \mathbb{R}^d \rightarrow \mathcal{H}$:

$$\mathbf{w}^{t+1} = \mathbf{w}^t - \rho \nabla_{\mathbf{w}} \mathcal{L}(\mathbf{s}^t \odot \mathbf{w}^t), \quad \mathbf{s}^t \sim \text{MASK}(\boldsymbol{\alpha}^t), \quad \alpha_j^t = \frac{\hat{\eta}_j(\mathbf{w}^t)}{\hat{\eta}_j(\mathbf{w}^t) + \lambda}. \quad (9)$$

The stochastic update of \mathbf{w} corresponds to the expected loss $\mathbb{E}_{\mathbf{s}}[\mathcal{L}(\mathbf{s} \odot \mathbf{w})]$, but we must also characterize the $\boldsymbol{\eta}$ update. We can make the assumption, which we will justify later, that $\hat{\boldsymbol{\eta}}(\mathbf{w})$ is the minimizer given \mathbf{w} of a joint objective function $\mathcal{J}(\mathbf{w}, \boldsymbol{\eta})$; if so, there must exist f such that

$$\mathcal{J}(\mathbf{w}, \boldsymbol{\eta}) = \mathbb{E}_{\mathbf{s}}[\mathcal{L}(\mathbf{s} \odot \mathbf{w})] + \frac{\lambda}{2} f(\boldsymbol{\eta}). \quad (10)$$

This brings us to our main result, in which we uncover a *duality* between the adaptive dropout objective function $\mathcal{J}(\mathbf{w}, \boldsymbol{\eta})$ and an effective regularized ERM objective function as in (1).

Theorem 1. *If \mathcal{L} has the form (3) for data \mathbf{X} such that $\frac{1}{n} \text{diag}(\mathbf{X}^\top \mathbf{X}) = \mathbf{I}$, then for every $\lambda > 0$ there exists a function f such that $\hat{\boldsymbol{\eta}}(\mathbf{w}) = \arg \min_{\boldsymbol{\eta}} \mathcal{J}(\mathbf{w}, \boldsymbol{\eta})$ if and only if there exists a function Ω that is concave in $\mathbf{w}^{\odot 2}$ such that*

$$\min_{\boldsymbol{\eta} \in \mathcal{H}} \mathcal{J}(\mathbf{w}, \boldsymbol{\eta}) = \mathcal{L}(\mathbf{w}) + \lambda \Omega(\mathbf{w}). \quad (11)$$

Furthermore, if f exists, then the corresponding Ω is unique, and if Ω exists, then there is a unique such f that is convex in $-\boldsymbol{\eta}^{\odot -1}$.

Proof. For linear regression under the standardized data assumption, combining the definition of $\boldsymbol{\eta}$ as a function of $\boldsymbol{\alpha}$ with (4) makes $\mathcal{J}(\mathbf{w}, \boldsymbol{\eta})$ of the form (6). Then by the properties of the LF transform² as discussed in Section 2.2, we obtain the existence and uniqueness of f and Ω . \square

That is, every subquadratic penalty Ω has a dual formulation via the η -trick and therefore has a corresponding adaptive dropout strategy. The converse holds as well, provided the dropout update function $\hat{\boldsymbol{\eta}}$ can be expressed as a minimizer of the joint objective. While this may seem restrictive, if $\hat{\boldsymbol{\eta}}$ is separable, we have the following general result.

Corollary 2. *If $\hat{\alpha}_j(w_j) := \frac{\hat{\eta}_j(w_j)}{\hat{\eta}_j(w_j) + \lambda}$ is a monotonically increasing differentiable function of $|w_j|$, then there exist separable f and Ω satisfying Theorem 1.*

Proof. Because the relationship between α and η is monotonic and invertible, $\hat{\eta}_j(w_j)$ is also monotonically increasing. Thus, using (7), $\Omega(w_j) = \int_a^{w_j} \frac{1}{2\hat{\eta}(t)} dt + C_a$, for which $\frac{\partial^2}{(\partial w_j^2)} \Omega(w_j) \leq 0$. \square

Therefore, any adaptive dropout method that updates the dropout parameters $\boldsymbol{\alpha}$ monotonically as a function of the magnitudes of the elements of \mathbf{w} is equivalent to regularized ERM with a subquadratic penalty Ω . It is well-known that sub-quadratic penalties are sparsity-inducing [5], which agrees with the “rich get richer” intuition that larger w_j , which receive larger α_j , are penalized less and therefore are more likely to grow even larger, while conversely smaller w_j receive smaller α_j and are even more penalized. Corollary 2 holds generally for the η -trick and is not limited to adaptive dropout.

We emphasize that linear regression is the setting in which Theorem 1 gives an exact duality, but that we can still expect a similar correspondence more broadly. For instance, Wager et al. [35] showed that for generalized linear models, dropout elicits a similar effect to (4), albeit with a somewhat more complex dependence on \mathbf{w} and the data. In Section 5, we empirically consider deep networks and demonstrate that the behavior is very similar between adaptive dropout and algorithms that solve the effective regularized ERM problem.

4 Several Sparsity Methods and Their Effective Penalties

Armed with the fact that adaptive dropout strategies correspond to regularization penalties, we now have a way to mathematically compare adaptive dropout methods. Table 2 lists the methods that we compare and a corresponding characterization of the dual formulation. These methods vary in their parameterizations, and most do not admit closed-form expressions for their effective penalties. However, given these characterizations, we can numerically compute the effective penalty and make a principled comparison, which we plot in Figure 1.

Because some of the methods we consider use slightly different variants of the dropout algorithm, we first make clear how these fit into our analysis framework.

Reparameterization Tricks. The problem of noise in stochastic gradients, especially when α_j is very small, can limit the applicability of dropout methods for sparsity. One way researchers have found to reduce variance is through “reparameterization tricks” such as the local reparameterization trick [36, 25] and the additive noise reparameterization trick [29], both used in VARIATIONALDROPOUT, which consider alternative noise models without changing the optimization problem.

To consider the MAGNITUDEPRUNING algorithm, we develop another reparameterization trick. Inspired by the additive reparameterization trick from Molchanov et al. [29], we introduce another variable \mathbf{v} and add the constraint $\mathbf{v} = \mathbf{w}$, under which $\mathcal{L}(\mathbf{s} \odot \mathbf{w}) = \mathcal{L}(\mathbf{w} + (\mathbf{s} - \mathbf{1}) \odot \mathbf{v})$. Then for least squares with standardized data, (4) becomes

$$\mathbb{E}_{\mathbf{s}} [\mathcal{L}(\mathbf{w} + (\mathbf{s} - \mathbf{1}) \odot \mathbf{v})] = \mathcal{L}(\mathbf{w}) + \frac{1}{2} \mathbf{v}^\top \left(\text{diag}(\boldsymbol{\alpha})^{-1} - \mathbf{I} \right) \mathbf{v}. \quad (12)$$

Under this parameterization, we can now use low-variance stochastic optimization for \mathbf{w} and deterministic optimization for \mathbf{v} , along with a strategy for enforcing the equality constraint $\mathbf{v} = \mathbf{w}$ such as

²Results on the LF transform do not apply to functions with discontinuities at $\{-\infty, \infty\}$, such as $-\boldsymbol{\eta}^{\odot -1} \mapsto f(\boldsymbol{\eta})$ for the ℓ_0 and HARDTHRESH(k) penalties, so our theorem does not prove the existence of such functions.

Table 2: Select sparsity methods and their effective penalty characterizations.

Method	Effective Penalty Characterization
STANDOUT	$\hat{\alpha}_j((z_j)_+) = \sigma((z_j)_+)$
VARIATIONALDROPOUT	$f(\eta) = 2D_{\text{KL}}(q_{\frac{\eta}{\lambda}, \mathbf{w}}(\tilde{\mathbf{w}}) \ p(\tilde{\mathbf{w}})) = \int_0^{\frac{\eta}{\lambda}} \frac{F(\sqrt{t/2})}{\sqrt{t/2}} dt$
HARDCONCRETENORM	$\begin{cases} f(\tilde{\boldsymbol{\eta}}) = 2 \sum_j \sigma(\log(a_j) - \beta \log(-\gamma/\zeta)) \\ \tilde{\eta}_j = \frac{\lambda \mathbb{E}[s_j]^2}{\mathbb{E}[s_j^2] - \mathbb{E}[s_j]^2}; \quad s_j \sim \text{HARDCONCRETE}_{\beta, \gamma, \zeta}(a_j) \end{cases}$
MAGNITUDEPRUNING(k)	$\hat{\eta}_j(\mathbf{w}) = \infty \mathbb{1}\{j \in \text{TOP-}k(\mathbf{w})\}$

ADMM [8], thus reducing the overall variance. A simple proximal variable update strategy under this reparameterization (see Appendix B) gives rise to an adaptive proximal gradient descent algorithm, which we use to describe the MAGNITUDEPRUNING algorithm.

Biased Masks. It is not uncommon for dropout strategies to use *biased* masks in the sense that $\mathbb{E}[s_j] \neq 1$, e.g., for $s_j \sim \text{BERNOULLI}(\alpha_j)$. Under a few mild assumptions on the parameterization of this distribution, we can consider adaptive dropout in this setting as well. Suppose that $\boldsymbol{\mu} := \boldsymbol{\mu}(\boldsymbol{\alpha}) = \mathbb{E}[\mathbf{s}]$ is an invertible function of $\boldsymbol{\alpha}$. Define $\tilde{s}_j := s_j/\mu_j$, $\tilde{\alpha}_j := \mu_j^2/\mathbb{E}[s_j^2]$, and $\tilde{\mathbf{w}} := \boldsymbol{\mu} \odot \mathbf{w}$. Then $\tilde{\mathbf{s}}$ is an unbiased mask, and since $\mathbf{s} \odot \mathbf{w} = \tilde{\mathbf{s}} \odot \tilde{\mathbf{w}}$, the expected loss has the same form as in (4), but with $\tilde{\boldsymbol{\alpha}}$ and $\tilde{\mathbf{w}}$ in place of $\boldsymbol{\alpha}$ and \mathbf{w} . Let $\tilde{\eta}_j := \lambda (\tilde{\alpha}_j^{-1} - 1)^{-1}$, and suppose there exists a function ψ such that $\mu_j = \psi(\tilde{\eta}_j)$, which is true if the mapping between $\tilde{\boldsymbol{\alpha}}$ and $\boldsymbol{\mu}$ is one-to-one. Now consider the expected dual regularized ERM formulation

$$\begin{aligned} \underset{\mathbf{w}, \boldsymbol{\eta}}{\text{minimize}} \quad & \mathcal{L}(\boldsymbol{\mu} \odot \mathbf{w}) + \frac{\lambda}{2} \left((\boldsymbol{\mu} \odot \mathbf{w})^\top \text{diag}(\tilde{\boldsymbol{\eta}})^{-1} (\boldsymbol{\mu} \odot \mathbf{w}) + \tilde{f}(\tilde{\boldsymbol{\eta}}) \right), \\ \text{subject to} \quad & \mu_j = \psi(\tilde{\eta}_j), \quad \forall j. \end{aligned} \quad (13)$$

Thus, an adaptive dropout strategy that jointly optimizes $\tilde{\boldsymbol{\eta}}$ and \mathbf{w} given $\mathbf{s} \sim \text{MASK}(\boldsymbol{\alpha})$ solves the regularized ERM problem for $\tilde{\mathbf{w}}$, where $\boldsymbol{\alpha}$ is set according to $\tilde{\boldsymbol{\eta}}$.

4.1 Standout

One of the earliest methods to use adaptive probabilities with dropout was the STANDOUT method of Ba and Frey [3], who proposed to augment a neural network with an additional parallel ‘‘standout’’ network of the same architecture that controlled the dropout probabilities for each activation, which they showed results in *sparse activations*. Interestingly, the authors noted that the method worked just as well if the standout net was not trained but instead simply copied the weights of the primary network at each iteration, which fits our adaptive dropout framework.

Consider a single hidden layer neural network with ReLU activations $(u)_+ = \max\{u, 0\}$. Denote the pre-activations from the first layer as $\mathbf{z} = \mathbf{W}_1 \mathbf{x}$. The output of the network is given by $\hat{y} = \mathbf{w}_2^\top (\mathbf{s} \odot (\mathbf{z})_+)$, where the mask $\mathbf{s} \sim \text{MASK}(\sigma(\mathbf{z}))$ is from the standout network. Now let us consider the network as a function of the first layer activations $(\mathbf{z})_+$. Using (4) and the fact that $\sigma(t)^{-1} - 1 = e^{-t}$, we know that the expected squared loss in this setting is

$$\frac{1}{2} \mathbb{E}_{\mathbf{s}} \left[\|\mathbf{y} - \mathbf{w}_2^\top (\mathbf{s} \odot (\mathbf{z})_+)\|_2^2 \right] = \frac{1}{2} \|\mathbf{y} - \mathbf{w}_2^\top (\mathbf{z})_+\|_2^2 + \frac{1}{2} (\mathbf{z})_+^\top \text{diag}(e^{-\mathbf{z}} \odot \mathbf{w}_2 \mathbf{w}_2^\top) (\mathbf{z})_+. \quad (14)$$

However, terms where $z_j < 0$ make no contribution to the effective penalty, so we can substitute $(z_j)_+$ for z_j . This enables us to extract $\hat{\boldsymbol{\eta}}((\mathbf{z})_+)$ as in Section 2.1

$$\hat{\eta}_j((z_j)_+) = \frac{\lambda e^{(z_j)_+}}{[w_2]_j^2}. \quad (15)$$

We note that (15) represents the implicit update of $\boldsymbol{\eta}$ each iteration as the standout network weights are copied from the primary network. By Corollary 2, we know STANDOUT is therefore an adaptive dropout strategy with a subquadratic effective penalty Ω applied to the activations $(\mathbf{z})_+$:

$$\Omega((\mathbf{z})_+) = \frac{1}{2\lambda} \sum_j \int_0^{(z_j)_+} [w_2]_j^2 e^{-\sqrt{t}} dt = \frac{1}{\lambda} \sum_j [w_2]_j^2 \left(1 - ((z_j)_+ + 1) e^{-(z_j)_+} \right). \quad (16)$$

This penalty is a smooth approximation to the ℓ_0 penalty, as shown in Figure 1, and thus naturally sparsifies the activations as \mathbf{W}_1 is trained. Additionally, $\lambda\Omega((\mathbf{z})_+)$ has no dependence on λ , just as the adaptive dropout parameter update has no dependence on λ .

4.2 Variational Dropout

Variational Bayes interpretations of dropout arose with the works of Kingma et al. [25] and Gal and Ghahramani [14], providing an automated means of choosing the dropout parameter through variational inference. Molchanov et al. [29] later showed that if each weight is allowed to have its own dropout parameter, VARIATIONALDROPOUT results in sparse solutions.

The authors of VARIATIONALDROPOUT consider maximum *a posteriori* inference of \mathbf{w} using the improper log-uniform prior, for which $p(\log |w_j|) \propto c$ or $p(|w_j|) \propto 1/|w_j|$. This is equivalent to solving the LOGSUM(0)-regularized ERM problem

$$\underset{\mathbf{w}}{\text{minimize}} \quad \mathcal{L}(\mathbf{w}) + \sum_j \log(|w_j|). \quad (17)$$

Obviously, this is minimized by taking $\mathbf{w} = \mathbf{0}$ for common loss functions, which is uninteresting, and Hron et al. [20] discuss other issues with the framework. However, VARIATIONALDROPOUT does not solve the problem in (17), and instead performs inference over a new variable $\tilde{\mathbf{w}}$, using an approximate posterior³ $q_{n\eta, \mathbf{w}}(\tilde{\mathbf{w}})$ such that $\tilde{\mathbf{w}} = \mathbf{s} \odot \mathbf{w}$, where $\mathbf{s} \sim \text{MASK}(\alpha)$ for Gaussian dropout, having α and η related by (8). Then maximizing the variational lower bound is equivalent to solving

$$\underset{\mathbf{w}, \eta}{\text{minimize}} \quad \mathbb{E}_{\mathbf{w}} [\mathcal{L}(\mathbf{s} \odot \mathbf{w})] + \frac{1}{n} D_{\text{KL}}(q_{n\eta, \mathbf{w}}(\tilde{\mathbf{w}}) \| p(\tilde{\mathbf{w}})), \quad (18)$$

where $D_{\text{KL}}(\cdot \| \cdot)$ is the Kullback–Leibler (KL) divergence. With the choice of the log-uniform prior, the KL divergence has no dependence on \mathbf{w} and is purely a function of η [25]. Hron et al. [20] show that this quantity can be expressed in terms of Dawson’s integral $F(u) = e^{-u^2} \int_0^u e^{t^2} dt$:

$$D_{\text{KL}}(q_{n\eta, \mathbf{w}}(\tilde{\mathbf{w}}) \| p(\tilde{\mathbf{w}})) = \int_0^{nn} \frac{F(\sqrt{t/2})}{\sqrt{2t}} dt, \quad (19)$$

which we can compute via numerical integration using standard software implementations of $F(u)$. Clearly then, the Monte Carlo variational Bayes optimization of \mathbf{w} and η in VARIATIONALDROPOUT is an adaptive dropout strategy with $\lambda = \frac{1}{n}$ and $f(\eta) = 2D_{\text{KL}}(q_{\frac{1}{n}\eta, \mathbf{w}}(\tilde{\mathbf{w}}) \| p(\tilde{\mathbf{w}}))$, and we can numerically compute the effective penalty Ω . VARIATIONALDROPOUT also uses the additive reparameterization trick and local reparameterization trick to reduce variance.

As shown in Figure 1, the effective penalty for $\lambda = 1$ is quite similar to the LOGSUM(2) penalty for $\lambda = 2$. We explore this connection further experimentally in Section 5.

4.3 Hard Concrete “ L_0 Regularization”

While the desired regularization is typically an ℓ_0 penalty, its non-differentiability makes it difficult to perform gradient-based optimization, which motivates finding smooth approximations to the ℓ_0 penalty. Louizos et al. [26] propose to first consider $\mathcal{L}(\tilde{\mathbf{w}}) + \frac{\lambda}{2} \|\tilde{\mathbf{w}}\|_0$, using the decomposition $\tilde{\mathbf{w}} = \mathbf{s} \odot \mathbf{w}$ for binary \mathbf{s} , where $\|\tilde{\mathbf{w}}\|_0 = \|\mathbf{s}\|_0$.

We take this opportunity to note that we can consider this problem using adaptive dropout with $f(\eta) = \mathbb{E}_{\mathbf{s}} [\|\mathbf{s}\|_0] = \sum_j \alpha_j$ if we use biased masks $s_j \sim \text{BERNOULLI}(\alpha_j)$. Interestingly, when $\lambda = 1$, the resulting penalty Ω turns out to be equal to the minimax concave penalty [40] MCP(1, 1), giving the MCP an additional interpretation as a stochastic binary smoothing of the ℓ_0 penalty.

However, Louizos et al. [26] do not solve this problem, but instead choose a distribution for \mathbf{s} such that the random variables \mathbf{s} themselves are differentiable with respect to α , allowing them to extract gradient information for α from $\mathcal{L}(\mathbf{s} \odot \mathbf{w})$. They propose the following biased distribution, called the $\text{HARDCONCRETE}_{\beta, \gamma, \zeta}(a_j)$ distribution:

$$z_j = \sigma((\log u_j - \log(1 - u_j) + \log a_j)/\beta), \quad u_j \sim \text{UNIFORM}[0, 1] \quad (20)$$

$$s_j = \min\{1, \max\{0, (\zeta - \gamma)z_j + \gamma\}\}. \quad (21)$$

³The parameter α from Molchanov et al. [29] is distinct from ours but is related to η by $\alpha^{-1} = n\eta$.

They then use as their penalty $f(\tilde{\boldsymbol{\eta}}) = 2 \sum_j \Pr(s_j > 0) = 2\sigma(\log(a_j) - \beta \log(-\gamma/\zeta))$ under the biased dropout reparameterization. The relationships between $\mathbb{E}[s_j]$, a_j , $\tilde{\alpha}_j$, and $\tilde{\eta}_j$ are all invertible, so we can numerically determine the corresponding penalty Ω .

We plot the resulting penalty in Figure 1 using the default values $\beta = \frac{2}{3}$, $\gamma = -0.1$, $\zeta = 1.1$ [26].

4.4 Magnitude Pruning

MAGNITUDEPRUNING [39, 41] is a straightforward way to induce sparsity by zeroing out small values. Here we specifically refer to the strategy of Zhu and Gupta [41], who eliminate all but the k largest parameters each iteration. Initially $k = d$, and it is decayed as a function of iteration number until it reaches the desired sparsity level according to a general-purpose pruning schedule. This is equivalent to IHT [7] (except with decreasing k) and therefore almost surely corresponds to a proximal gradient variant of adaptive dropout where $\hat{\alpha}_j(\mathbf{v}) = \mathbb{1}\{j \in \text{TOP-}k(\mathbf{w})\}$, or $\hat{\eta}_j(\mathbf{v}) = \infty \mathbb{1}\{j \in \text{TOP-}k(\mathbf{v})\}$, implemented as in Appendix B.

According to Table 1, the effective penalty is the $\text{HARDTHRESH}(k)$ penalty, which we plot in Figure 1. On a philosophical note, this is by definition the “correct” penalty to use when searching for solutions with a constraint on the number of nonzeros, since it is the characteristic function of that set of solutions. The η -trick along with alternating optimization turns this combinatorial search problem into a sequence of tractable gradient-based (for \mathbf{w}) and closed-form (for $\boldsymbol{\eta}$) updates. It is therefore no surprise that Gale et al. [15] observed that this “simple” magnitude pruning algorithm matched or outmatched the other more sophisticated methods.

5 Empirical Comparison for VARIATIONALDROPOUT

With complex methods such as the adaptive dropout methods we discuss above, it is often difficult to what drives their success. Under our theory, such methods should be able to be decomposed into dropout noise-related effects and η -trick effects. Diving further into the VARIATIONALDROPOUT algorithm, we compare it against other η -trick strategies, examining the effects of dropout noise, η -trick optimization strategy, and regularization with the underlying penalty.

We apply the methods on the task of sparsifying a LeNet-300-100 network on the MNIST dataset, a common benchmark for sparsification methods [17, 16, 12]. Experimental details are given in Appendix C. We compare the following methods. Except for VARIATIONALDROPOUT methods, we use the LOGSUM penalty.

- **VARDROP+LR+AR.** VARIATIONALDROPOUT as proposed by Molchanov et al. [29], using both the local reparameterization trick and the additive reparameterization trick.
- **VARDROP+LR.** VARIATIONALDROPOUT, using the local reparameterization trick but not the additive reparameterization trick.
- **VARDROP.** VARIATIONALDROPOUT with no reparameterization tricks.
- **η -TRICK.** Joint gradient descent in \mathbf{w} and $\log(\boldsymbol{\eta})$ of (6). This would be equivalent to VARIATIONALDROPOUT with no noise in linear regression.
- **ADAPROX.** Proximal gradient descent in \mathbf{w} of (6) using $\boldsymbol{\eta} = \hat{\boldsymbol{\eta}}(\mathbf{w})$ at each iteration.
- **ADATIKHONOV.** Gradient descent in \mathbf{w} of (6) using $\boldsymbol{\eta} = \hat{\boldsymbol{\eta}}(\mathbf{w})$ at each iteration.
- **LOGSUM.** Gradient descent in \mathbf{w} of (1).

We plot the results in Figure 2. With all of the methods compared, we see similar behavior: the networks become more sparse over time as well as more accurate. In general, the methods appear to exhibit a sparsity–accuracy trade-off, even when they optimize identical objective functions (consider VARDROP+LR+AR and VARDROP+LR, or ADAPROX and η -TRICK). This trade-off seems to be determined by the early iterations, depending on whether the method sparsifies the network quickly or not. We also observe that the stochastic noise due to dropout in VARDROP leads to very slow convergence while the variance-reduced VARDROP+LR converges exceedingly similarly to its non-random counterpart η -TRICK. In addition, VARDROP+LR+AR behaves remarkably similarly to ADAPROX, which was inspired by the additive reparameterization trick. These results suggests that the success of VARIATIONALDROPOUT is well explained by interpreting it as optimization of a dual formulation of a regularized ERM problem. Further, they suggest that the variance due to

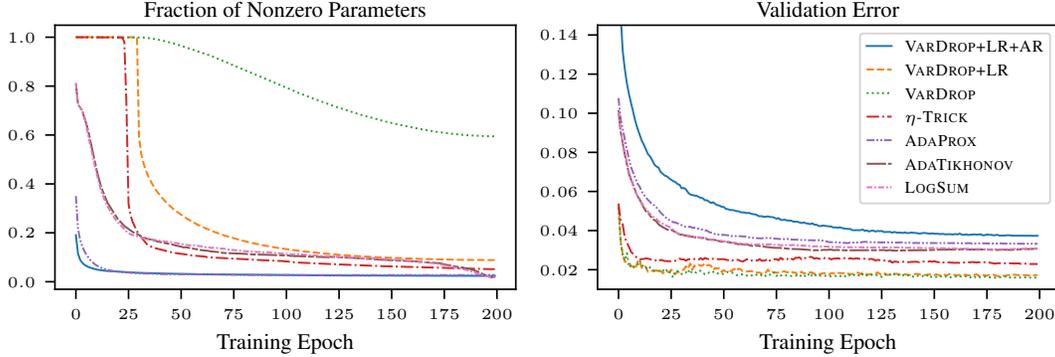


Figure 2: VARIATIONALDROPOUT compared with other η -trick strategies using a LOGSUM penalty.

dropout at best yields similar performance to the non-dropout counterpart and at worst can drastically slow down convergence. Interestingly, the simple LOGSUM-regularized ERM approach performs as well as any of the other methods considered, and is nearly identical in behavior to ADATIKHONOV, suggesting that the sparsification performance of VARIATIONALDROPOUT is entirely due to the effective LOGSUM penalty.

6 Discussion

Given the duality between adaptive dropout and regularization that we have presented, it is no surprise that such methods excel at sparsifying deep networks. However, many questions still remain.

For one, is there any benefit to dropout noise itself in adaptive dropout optimization, beyond the η -trick connection? There are known benefits of standard dropout for deep network optimization [1, 19, 28, 37], but do these benefits transfer to adaptive dropout strategies? A crucial component of VARIATIONALDROPOUT, as demonstrated in Section 5, appears to be that its successful implementations *remove* most of the noise due to dropout. On the other hand, methods like HARD-CONCRETENORM from Louizos et al. [26] exploit the fact that sparse masks yield more efficient computation, and they do not seem to suffer from variance-related optimization issues even without taking efforts to mitigate such effects. This could be because they use a biased dropout distribution, which results in significantly lower variance of stochastic gradients. On the other hand, it is not clear that dropout noise is necessary to obtain these computational speedups; other methods where the network is sparse during training, such as MAGNITUDEPRUNING, should have the similar properties.

It is evident from Figure 2 that, even when using the same objective function, the optimization strategy can have a significant impact on the result. It is beyond the scope of this work to examine the convergence rates, even for convex problems, of the adaptive algorithms we consider in Section 5, although many of them can be cast as stochastic majorization-minimization (MM) algorithms, for which some theory exists [27]. However, for non-convex penalties and losses, such as in sparse deep network optimization, there are no guarantees about the quality of the solution. It is also not clear what benefit η -trick optimization provides in this setting, if any. For linear problems it has the advantage of drastically simplifying implementations of penalized methods by reducing them to successive least squares problems, but with modern automatic differentiation packages it is straightforward to implement any differentiable sparsity-inducing penalty. To the best of our knowledge, the only large-scale comparison of sparsity methods for deep networks is that of Gale et al. [15], and they do not consider classical penalty-based methods.

We have fully analyzed the adaptive-dropout–regularization duality in the linear regression case, and we know that the exact correspondence between dropout and Tikhonov regularization does not hold in other settings. However, good quadratic approximations to the general effective regularization of dropout have been developed [35, 37], which relate to the Fisher information. These approximations could be used to develop adaptive sparsification strategies that enjoy the data-sensitive properties of dropout, such as in logistic regression or deep networks.

The duality presented in this paper can be the basis of developing more efficient principled sparsification schemes as future work. In particular, using penalties that have been proposed in the compressed sensing literature such as scaled lasso [33], which is adaptive to noise level, or SLOPE [32], which is shown to be minimax optimal in feature selection, can lead to adaptive dropout schemes with better

performance in sparsifying deep networks. In addition, analytical tools such as debiasing [23], which has been used for near-optimal feature selection [22], could be employed to design improved sparsity methods.

Acknowledgments and Disclosure of Funding

This work was supported by NSF grants CCF-1911094, IIS-1838177, and IIS-1730574; ONR grants N00014-18-12571, N00014-20-1-2534, and MURI N00014-20-1-2787; AFOSR grant FA9550-18-1-0478; and a Vannevar Bush Faculty Fellowship, ONR grant N00014-18-1-2047.

References

- [1] R. Arora, P. Bartlett, P. Mianjy, and N. Srebro. Dropout: Explicit forms and capacity control. *arXiv e-prints*, arXiv:2003.03397, 2020.
- [2] A. Ashukha. Sparse Variational Dropout a Minimal Working Example, 2018. URL <https://github.com/senya-ashukha/sparse-vd-pytorch>. Accessed: May 2021.
- [3] J. Ba and B. Frey. Adaptive dropout for training deep neural networks. In *Advances in Neural Information Processing Systems*, volume 26, 2013.
- [4] F. Bach. The “ η -trick” or the effectiveness of reweighted least-squares, Jul 2019. URL <https://francisbach.com/the-%CE%B7-trick-or-the-effectiveness-of-reweighted-least-squares/>. Accessed: May 2021.
- [5] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning*, 2011.
- [6] L. Birgé and P. Massart. Gaussian model selection. *Journal of the European Mathematical Society*, 3:203–268, 08 2001.
- [7] T. Blumensath and M. E. Davies. Iterative hard thresholding for compressed sensing. *Applied and Computational Harmonic Analysis*, 27(3):265–274, 2009.
- [8] S. Boyd, N. Parikh, and E. Chu. *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc, 2011.
- [9] E. J. Candès, M. B. Wakin, and S. P. Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14(5-6):877–905, 2008.
- [10] R. Chartrand and W. Yin. Iteratively reweighted algorithms for compressive sensing. In *2008 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3869–3872, 2008.
- [11] I. Daubechies, R. DeVore, M. Fornasier, and C. S. Güntürk. Iteratively reweighted least squares minimization for sparse recovery. *Communications on Pure and Applied Mathematics*, 63(1): 1–38, 2010.
- [12] X. Dong, S. Chen, and S. Pan. Learning to prune deep neural networks via layer-wise optimal brain surgeon. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- [13] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 2001.
- [14] Y. Gal and Z. Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, pages 1050–1059, 20–22 Jun 2016.
- [15] T. Gale, E. Elsen, and S. Hooker. The state of sparsity in deep neural networks. *arXiv e-prints*, arXiv:1902.09574, 2019.
- [16] Y. Guo, A. Yao, and Y. Chen. Dynamic network surgery for efficient DNNs. In *Advances in Neural Information Processing Systems*, volume 29, 2016.

- [17] S. Han, J. Pool, J. Tran, and W. Dally. Learning both weights and connections for efficient neural network. In *Advances in Neural Information Processing Systems*, volume 28, 2015.
- [18] T. Hastie, R. Tibshirani, and M. Wainwright. *Statistical learning with sparsity: the lasso and generalizations*. Chapman and Hall/CRC Press, 2019.
- [19] D. P. Helmbold and P. M. Long. Surprising properties of dropout in deep networks. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65, pages 1123–1146, 07–10 Jul 2017.
- [20] J. Hron, A. Matthews, and Z. Ghahramani. Variational Bayesian dropout: pitfalls and fixes. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 2019–2028, 10–15 Jul 2018.
- [21] P. J. Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.
- [22] A. Javanmard and H. Javadi. False discovery rate control via debiased lasso. *Electronic Journal of Statistics*, 13(1):1212–1253, 2019.
- [23] A. Javanmard and A. Montanari. Debiasing the lasso: Optimal sample size for Gaussian designs. *The Annals of Statistics*, 46(6A):2593–2622, 2018.
- [24] R. Jenatton, G. Obozinski, and F. Bach. Structured sparse principal component analysis. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9, pages 366–373, 13–15 May 2010.
- [25] D. P. Kingma, T. Salimans, and M. Welling. Variational dropout and the local reparameterization trick. In *Advances in Neural Information Processing Systems*, volume 28, 2015.
- [26] C. Louizos, M. Welling, and D. P. Kingma. Learning sparse neural networks through L_0 regularization. In *International Conference on Learning Representations*, 2018.
- [27] J. Mairal. Stochastic majorization-minimization algorithms for large-scale optimization. In *Advances in Neural Information Processing Systems*, volume 26, 2013.
- [28] P. Mianjy and R. Arora. On convergence and generalization of dropout training. In *Advances in Neural Information Processing Systems*, volume 33, pages 21151–21161, 2020.
- [29] D. Molchanov, A. Ashukha, and D. Vetrov. Variational dropout sparsifies deep neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 2498–2507, 06–11 Aug 2017.
- [30] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems* 32, pages 8024–8035. 2019.
- [31] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15 (56):1929–1958, 2014.
- [32] W. Su and E. Candès. SLOPE is adaptive to unknown sparsity and asymptotically minimax. *The Annals of Statistics*, 44(3):1038–1068, 2016.
- [33] T. Sun and C.-H. Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 09 2012.
- [34] M. Tietz, T. J. Fan, D. Nouri, B. Bossan, and skorch Developers. *skorch: A scikit-learn compatible neural network library that wraps PyTorch*, July 2017. URL <https://skorch.readthedocs.io/en/stable/>.
- [35] S. Wager, S. Wang, and P. S. Liang. Dropout training as adaptive regularization. In *Advances in Neural Information Processing Systems*, volume 26, 2013.

- [36] S. Wang and C. Manning. Fast dropout training. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28, pages 118–126, 17–19 Jun 2013.
- [37] C. Wei, S. Kakade, and T. Ma. The implicit and explicit regularization effects of dropout. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 10181–10192, 13–18 Jul 2020.
- [38] D. Wipf and S. Nagarajan. Iterative reweighted ℓ_1 and ℓ_2 methods for finding sparse solutions. *IEEE Journal of Selected Topics in Signal Processing*, 4(2):317–329, 2010.
- [39] X. Xiao, Y. Yang, T. Ahmad, L. Jin, and T. Chang. Design of a very compact CNN classifier for online handwritten Chinese character recognition using DropWeight and global pooling. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, volume 01, pages 891–895, 2017.
- [40] C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894–942, 2010.
- [41] M. Zhu and S. Gupta. To prune, or not to prune: exploring the efficacy of pruning for model compression. *arXiv e-prints*, arXiv:1710.01878, 2018.
- [42] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.
- [43] H. Zou and R. Li. One-step sparse estimates in nonconcave penalized likelihood models. *The Annals of Statistics*, 36(4):1509–1533, 2008.

A Appendix: Dual Forms of Subquadratic Penalties

A.1 Dual Forms for Common Penalties

The full version of the Table 1 is given in Table 3. The dual formulations are derived in the next sections.

Table 3: Common penalties and their corresponding dual formulations.

Penalty	$\Omega(\mathbf{w})$	$f(\boldsymbol{\eta})$	$\hat{\eta}_j(\mathbf{w})$
ℓ_1	$ w_j $	η_j	$ w_j $
$\ell_p, p \in (0, 2)$	$\ \mathbf{w}\ _p$	$\ \boldsymbol{\eta}\ _q : q = \frac{p}{2-p}$	$ w_j ^{2-p} \ \mathbf{w}\ _p^{p-1}$
$\ell_p^p, p \in (0, 2)$	$\frac{1}{p} w_j ^p$	$\frac{1}{q}\eta_j^q : q = \frac{p}{2-p}$	$ w_j ^{2-p}$
ℓ_0	$\mathbb{1}\{ w_j > 0\}$	$2\mathbb{1}\{\eta_j > 0\}$	$\infty \mathbb{1}\{ w_j > 0\}$
ELASTICNET(θ) [42]	$\frac{\theta}{2}w_j^2 + (1-\theta) w_j $	$\frac{\eta_j(1-\theta)^2}{1-\eta_j\theta}$	$\frac{ w_j }{ w_j \theta + (1-\theta)}$
HUBER(ε) [11, 21]	$\begin{cases} \frac{1}{2\varepsilon}w_j^2 + \frac{\varepsilon}{2}, & w_j \leq \varepsilon \\ w_j , & w_j > \varepsilon \end{cases}$	$\eta_j, \mathcal{H} = [\varepsilon, \infty)$	$\max\{\varepsilon, w_j \}$
LOGSUM(ε) [9]	$\log(w_j + \varepsilon)$	$2 \log\left(\frac{\sqrt{\varepsilon^2 + 4\eta_j} + \varepsilon}{2}\right) - \frac{(\sqrt{\varepsilon^2 + 4\eta_j} - \varepsilon)^2}{4\eta_j}$	$ w_j (w_j + \varepsilon)$
SCAD(a, λ) [13]	$\begin{cases} w_j , & w_j \leq \lambda \\ \frac{2a\lambda w_j - w_j^2 - \lambda^2}{2(a-1)\lambda}, & w_j \in (\lambda, a\lambda] \\ \frac{(a+1)\lambda}{2}, & w_j > a\lambda \end{cases}$	$\begin{cases} \eta, & \eta_j \leq \lambda \\ \lambda \frac{(a+1)\eta_j - \lambda}{(a-1)\lambda + \eta_j}, & \eta_j > \lambda \end{cases}$	$\begin{cases} w_j , & w_j \leq \lambda \\ \frac{(a-1)\lambda w_j }{a\lambda - w_j }, & w_j \in (\lambda, a\lambda] \\ \infty, & w_j > a\lambda \end{cases}$
MCP(a, λ) [40]	$\begin{cases} w_j - \frac{w_j^2}{2a\lambda}, & w_j \leq a\lambda \\ \frac{a\lambda}{2}, & w_j > a\lambda \end{cases}$	$\frac{a\lambda\eta_j}{\eta_j + a\lambda}$	$\begin{cases} \frac{a\lambda w_j }{a\lambda - w_j }, & w_j < a\lambda \\ \infty, & w_j \geq a\lambda \end{cases}$
HARDTHRESH(k) [7]	$\infty \mathbb{1}\{\ \mathbf{w}\ _0 > k\}$	$0, \mathcal{H} = \{\boldsymbol{\eta} : \ \boldsymbol{\eta}\ _0 \leq k\}$	$\infty \mathbb{1}\{j \in \text{TOP-}k(\mathbf{w})\}$

A.2 General Strategy

Define $g(\mathbf{u}) := \Omega(\mathbf{u}^{\odot \frac{1}{2}})$ for $\mathbf{u} \in \mathbb{R}_+^d$ and $h(\mathbf{v}) := f(-\mathbf{v}^{\odot -1})$ for $\mathbf{v} \in \overline{\mathbb{R}}_{\leq 0}^d$. As mentioned in Section 2.2, when g is concave, $-2g$ and h comprise a Legendre–Fenchel conjugate pair, each being convex functions. That is, the following relationships hold:

$$-2g(\mathbf{u}) = \sup_{\mathbf{v}} \mathbf{u}^\top \mathbf{v} - h(\mathbf{v}), \quad h(\mathbf{v}) = \sup_{\mathbf{u}} \mathbf{u}^\top \mathbf{v} + 2g(\mathbf{u}). \quad (22)$$

We can thus obtain $-2g$ and h from each other by solving these optimizations. If the functions are differentiable, the following first-order conditions must hold for the dual pair \mathbf{u}^* and \mathbf{v}^* , the argument and maximizing variable in either equation in (22):

$$\mathbf{u}^* = \nabla_{\mathbf{v}} h(\mathbf{v}^*), \quad \mathbf{v}^* = -2\nabla_{\mathbf{u}} g(\mathbf{u}^*). \quad (23)$$

Note that the second condition is equivalent to (7). Once we have characterized g and h , we can recover Ω and f by considering $\mathbf{w}^{\odot 2} = \mathbf{u}$ and $\boldsymbol{\eta} = -\mathbf{v}^{\odot -1}$. In the following sections, we use these properties to obtain the dual forms presented in Table 1. For separable penalties, it suffices to derive the dual form of the scalar penalty.

A.3 ℓ_p for $0 < p < 2$

This formulation can be found in Lemma 3.1 of Jenatton et al. [24], but we present another derivation here. First we compute the gradient

$$g(\mathbf{u}) = \left(\sum_j u_j^{\frac{p}{2}} \right)^{\frac{1}{p}} \implies \nabla_{\mathbf{u}} g(\mathbf{u}) = \frac{1}{2} \mathbf{u}^{\odot \frac{p}{2} - 1} \left(\sum_j u_j^{\frac{p}{2}} \right)^{\frac{1}{p} - 1}. \quad (24)$$

The first-order condition is

$$\mathbf{v}^* = -\mathbf{u}^* \odot^{\frac{p-2}{2}} g(\mathbf{u}^*)^{1-p}, \quad (25)$$

which gives $\hat{\eta}_j(\mathbf{w}) = |w_j|^{2-p} \|\mathbf{w}\|_p^{p-1}$. Now then

$$h(\mathbf{v}^*) = - \left(\sum_j u_j^{*\frac{p}{2}} \right) g(\mathbf{u}^*)^{1-p} + 2g(\mathbf{u}^*) \quad (26)$$

$$= -g(\mathbf{u}^*)^p g(\mathbf{u}^*)^{1-p} + 2g(\mathbf{u}^*) \quad (27)$$

$$= g(\mathbf{u}^*). \quad (28)$$

Since $g(a\mathbf{z}) = \sqrt{a}g(\mathbf{z})$, we can solve (25) for $g(\mathbf{u}^*)$:

$$g(\mathbf{u}^*) = g \left(\left(\frac{-\mathbf{v}^*}{g(\mathbf{u}^*)^{1-p}} \right)^{\odot \frac{2}{p-2}} \right) \quad (29)$$

$$= g(\mathbf{u}^*)^{\frac{p-1}{p-2}} g \left((-\mathbf{v}^*)^{\odot \frac{2}{p-2}} \right) \quad (30)$$

$$\implies g(\mathbf{u}^*)^{\frac{1}{2-p}} = \left(\sum_j \left(-\frac{1}{v_j} \right)^{\frac{p}{2-p}} \right)^{\frac{1}{p}} \quad (31)$$

$$\implies h(\mathbf{v}^*) = \left(\sum_j \left(-\frac{1}{v_j} \right)^{\frac{p}{2-p}} \right)^{\frac{2-p}{p}}. \quad (32)$$

Thus, $f(\eta) = \|\boldsymbol{\eta}\|_q$ for $q = \frac{p}{2-p}$.

A.4 ℓ_p^p for $0 < p < 2$

First we compute the derivative

$$g(u) = \frac{1}{p} u^{\frac{p}{2}} \implies g'(u) = \frac{1}{2} u^{\frac{p}{2}-1}. \quad (33)$$

Then the first-order condition is

$$v^* = -u^{*\frac{p-2}{2}}, \quad (34)$$

which gives $\hat{\eta}(w) = |w|^{2-p}$. We also have $u^* = -v^{*\frac{2}{p-2}}$, which gives us

$$h(v^*) = -(-v^*)^{\frac{2}{p-2}+1} + \frac{2}{p} (-v^*)^{\frac{p}{p-2}} \quad (35)$$

$$= \frac{2-p}{p} \left(-\frac{1}{v^*} \right)^{\frac{p}{2-p}}. \quad (36)$$

Thus, $f(\eta) = \frac{1}{q} \eta^q$ for $q = \frac{p}{2-p}$.

A.5 Elastic Net

First, we compute the derivative

$$g(u) = \frac{\theta}{2} u + (1-\theta)\sqrt{u} \implies g'(u) = \frac{\theta}{2} + \frac{1-\theta}{2\sqrt{u}}. \quad (37)$$

The first-order condition is

$$v^* = -\theta - \frac{1-\theta}{\sqrt{u^*}}, \quad (38)$$

which is bounded by $v^* \leq -\theta$. From this we obtain $\hat{\eta}(w) = \frac{|w|}{|w|\theta + (1-\theta)}$. We also have $\sqrt{u^*} = \frac{1-\theta}{-v^*-\theta}$. This gives us

$$h(v^*) = \frac{v^*(1-\theta)^2}{(-v^*-\theta)^2} + \frac{\theta(1-\theta)^2}{(-v^*-\theta)^2} + \frac{2(1-\theta)^2}{(-v^*-\theta)} \quad (39)$$

$$= \frac{(1-\theta)^2}{(-v^*-\theta)}. \quad (40)$$

Thus, $f(\eta) = \frac{\eta(1-\theta)^2}{1-\eta\theta}$ for $\eta \leq \frac{1}{\theta}$.

A.6 Huber

As usual, first we compute the derivative

$$g(u) = \begin{cases} \frac{1}{2\varepsilon}u + \frac{\varepsilon}{2}, & \sqrt{u} \leq \varepsilon \\ \sqrt{u}, & \sqrt{u} > \varepsilon \end{cases} \implies g'(u) = \begin{cases} \frac{1}{2\varepsilon}, & \sqrt{u} \leq \varepsilon \\ \frac{1}{2\sqrt{u}}, & \sqrt{u} > \varepsilon \end{cases}. \quad (41)$$

The first-order condition is

$$v^* = -\min \left\{ \frac{1}{\varepsilon}, \frac{1}{\sqrt{u^*}} \right\}, \quad (42)$$

which is bounded by $v^* \geq -\frac{1}{\varepsilon}$. This gives us $\hat{\eta}(w) = \max\{\varepsilon, |w|\}$. For $v^* \geq -\frac{1}{\varepsilon}$, $\sqrt{u^*} = -\frac{1}{v^*}$, so

$$h(v^*) = \frac{1}{v^*} - 2\frac{1}{v^*} = -\frac{1}{v^*}. \quad (43)$$

Thus, $f(\eta) = \eta$ for $\eta \geq \varepsilon$.

A.7 Log Sum

First, we compute the derivative

$$g(u) = \log(\sqrt{u} + \varepsilon) \implies g'(u) = \frac{1}{2\sqrt{u}(\sqrt{u} + \varepsilon)}. \quad (44)$$

Then the first-order condition is

$$v^* = -\frac{1}{\sqrt{u^*}(\sqrt{u^*} + \varepsilon)}. \quad (45)$$

This gives us $\hat{\eta}(w) = |w|(|w| + \varepsilon)$. Rewriting the above as a quadratic equation in $\sqrt{u^*}$, we have

$$(\sqrt{u^*})^2 + \varepsilon\sqrt{u^*} + \frac{1}{v^*} = 0, \quad (46)$$

which gives the inverse mapping $\sqrt{u^*} = \frac{\sqrt{\varepsilon^2 - \frac{4}{v^*}} - \varepsilon}{2}$. Thus we get

$$h(v^*) = \frac{v^*}{4} \left(\sqrt{\varepsilon^2 - \frac{4}{v^*}} - \varepsilon \right)^2 + 2 \log \left(\frac{\sqrt{\varepsilon^2 - \frac{4}{v^*}} + \varepsilon}{2} \right). \quad (47)$$

Thus, $f(\eta) = 2 \log \left(\frac{\sqrt{\varepsilon^2 + 4\eta} + \varepsilon}{2} \right) - \frac{1}{4\eta} \left(\sqrt{\varepsilon^2 + 4\eta} - \varepsilon \right)^2$.

A.8 SCAD

The SCAD penalty as presented by Fan and Li [13] uses the regularization scaling λ as a parameter, so first we factor it out:

$$\lambda\Omega(w) = \lambda \begin{cases} |w|, & |w| \leq \lambda \\ \frac{2a\lambda|w| - w^2 - \lambda^2}{2(a-1)\lambda}, & |w| \in (\lambda, a\lambda] \\ \frac{(a+1)\lambda}{2}, & |w| > a\lambda \end{cases}. \quad (48)$$

We then compute the derivative

$$g(u) = \begin{cases} \sqrt{u}, & \sqrt{u} \leq \lambda \\ \frac{2a\lambda\sqrt{u} - u - \lambda^2}{2(a-1)\lambda}, & \sqrt{u} \in (\lambda, a\lambda] \\ \frac{(a+1)\lambda}{2}, & \sqrt{u} > a\lambda \end{cases} \implies g'(u) = \begin{cases} \frac{1}{2\sqrt{u}}, & \sqrt{u} \leq \lambda \\ \frac{a}{2(a-1)\sqrt{u}} - \frac{1}{2(a-1)\lambda}, & \sqrt{u} \in (\lambda, a\lambda] \\ 0, & \sqrt{u} > a\lambda \end{cases}. \quad (49)$$

This gives us the first order condition and in turn $\hat{\eta}$:

$$v^* = \begin{cases} -\frac{1}{\sqrt{u^*}}, & \sqrt{u^*} \leq \lambda \\ -\frac{a}{(a-1)\sqrt{u^*}} + \frac{1}{(a-1)\lambda}, & \sqrt{u^*} \in (\lambda, a\lambda] \\ 0, & \sqrt{u^*} > a\lambda \end{cases} \quad (50)$$

$$\implies \hat{\eta}(w) = \begin{cases} |w|, & |w| \leq \lambda \\ \frac{(a-1)\lambda|w|}{a\lambda - |w|}, & |w| \in (\lambda, a\lambda] \\ \infty, & |w| > a\lambda \end{cases} \quad (51)$$

Now when $\sqrt{u^*} \leq \lambda$, $v^* \leq -\frac{1}{\lambda}$, and when $\sqrt{u^*} \in (\lambda, a\lambda]$, $v^* \in (-\frac{1}{\lambda}, 0]$. In the first case, $\sqrt{u^*} = -\frac{1}{v^*}$, and in the second, $\sqrt{u^*} = \frac{a\lambda}{1-(a-1)\lambda v^*}$. Therefore,

$$h(v^*) = \begin{cases} \frac{1}{v^*} - \frac{2}{v^*}, & v^* \leq -\frac{1}{\lambda} \\ \frac{a^2\lambda^2 v^*}{(1-(a-1)\lambda v^*)^2} + \frac{2a\lambda\left(\frac{a\lambda}{1-(a-1)\lambda v^*}\right) - \frac{a^2\lambda^2}{(1-(a-1)\lambda v^*)^2} - \lambda^2}{(a-1)\lambda}, & v^* > -\frac{1}{\lambda} \end{cases} \quad (52)$$

$$= \begin{cases} -\frac{1}{v^*}, & v^* \leq -\frac{1}{\lambda} \\ \frac{a^2(a-1)\lambda^3 v^* + 2a^2\lambda^2(1-(a-1)\lambda v^*) - a^2\lambda^2 - \lambda^2(1-(a-1)\lambda v^*)^2}{(a-1)\lambda(1-(a-1)\lambda v^*)^2}, & v^* > -\frac{1}{\lambda} \end{cases} \quad (53)$$

$$= \begin{cases} -\frac{1}{v^*}, & v^* \leq -\frac{1}{\lambda} \\ \frac{\lambda(a^2 - a^2(a-1)\lambda v^* - (1-(a-1)\lambda v^*)^2)}{(a-1)(1-(a-1)\lambda v^*)^2}, & v^* > -\frac{1}{\lambda} \end{cases} \quad (54)$$

$$= \begin{cases} -\frac{1}{v^*}, & v^* \leq -\frac{1}{\lambda} \\ \frac{\lambda(a^2 - 1 + (a-1)\lambda v^*)}{(a-1)(1-(a-1)\lambda v^*)}, & v^* > -\frac{1}{\lambda} \end{cases} \quad (55)$$

$$= \begin{cases} -\frac{1}{v^*}, & v^* \leq -\frac{1}{\lambda} \\ \frac{\lambda(a+1+\lambda v^*)}{1-(a-1)\lambda v^*}, & v^* > -\frac{1}{\lambda}. \end{cases} \quad (56)$$

From this we obtain

$$f(\eta) = \begin{cases} \eta, & \eta \leq \lambda \\ \lambda \frac{(a+1)\eta - \lambda}{(a-1)\lambda + \eta}, & \eta > \lambda. \end{cases} \quad (57)$$

A.9 MCP

As with SCAD, we first factor out the λ from the penalty:

$$\lambda\Omega(w) = \lambda \begin{cases} |w| - \frac{w^2}{2a\lambda}, & |w| \leq a\lambda \\ \frac{a\lambda}{2}, & |w| > a\lambda \end{cases} \quad (58)$$

We then compute the derivative

$$g(u) = \begin{cases} \sqrt{u} - \frac{u}{2a\lambda}, & \sqrt{u} \leq a\lambda \\ \frac{a\lambda}{2}, & \sqrt{u} > a\lambda \end{cases} \implies g'(u) = \begin{cases} \frac{1}{2\sqrt{u}} - \frac{1}{2a\lambda}, & \sqrt{u} \leq a\lambda \\ 0, & \sqrt{u} > a\lambda \end{cases} \quad (59)$$

Our first-order condition is

$$v^* = \begin{cases} -\frac{1}{\sqrt{u^*}} + \frac{1}{a\lambda}, & \sqrt{u^*} \leq a\lambda \\ 0, & \sqrt{u^*} > a\lambda \end{cases}, \quad (60)$$

from which we obtain

$$\hat{\eta}(w) = \begin{cases} \frac{a\lambda|w|}{a\lambda - |w|}, & |w| < a\lambda \\ \infty, & |w| \geq a\lambda \end{cases} \quad (61)$$

We have the inverse mapping $\sqrt{u^*} = \frac{a\lambda}{1-a\lambda v^*}$, which gives us

$$h(v^*) = \frac{a^2\lambda^2 v^*}{(1-a\lambda v^*)^2} + \frac{2a\lambda}{1-a\lambda v^*} - \frac{a\lambda}{(1-a\lambda v^*)^2} \quad (62)$$

$$= \frac{a\lambda(a\lambda v^* + 2(1-a\lambda v^*) - 1)}{(1-a\lambda v^*)^2} \quad (63)$$

$$= \frac{a\lambda}{1-a\lambda v^*}. \quad (64)$$

From here, we directly obtain $f(\eta) = \frac{a\lambda\eta}{\eta+a\lambda}$.

A.10 ℓ_0

The ℓ_0 penalty is not differentiable. However, it is separable, and in one dimension we have

$$g(u) = \mathbb{1}\{u > 0\}. \quad (65)$$

Thus $-2g$ is convex since its epigraph is a convex set. For $u = 0$, $-2g$ has a supporting line with slope $-\infty$, and elsewhere with slope 0. Thus we have the relationship $v^* = -\infty \mathbb{1}\{u^* = 0\}$, which yields $\hat{\eta}(w) = \infty \mathbb{1}\{|w| > 0\}$. The mapping $u^* \mapsto v^*$ is not invertible, so we consider two cases of v^* :

$$h(v^*) = \begin{cases} 0, & v^* = -\infty \\ \sup_{u>0} uv^* + 2g(u), & v^* > -\infty \end{cases} \quad (66)$$

$$= 2\mathbb{1}\{v^* > -\infty\}. \quad (67)$$

We thus conclude that $f(\eta) = \mathbb{1}\{\eta > 0\}$.

A.11 Hard Threshold

For this penalty, we begin with the $\hat{\eta}(\mathbf{w})$ that yields the IHT algorithm when \mathbf{w} is optimized by a gradient step. This corresponds to

$$\hat{\eta}_j(\mathbf{w}) = \infty \mathbb{1}\{j \in \text{TOP-}k(\mathbf{w})\}. \quad (68)$$

We thus seek to find a penalty that yields such an $\hat{\eta}$. In interest of mathematical preciseness, let us define, given $a > 0$ and $m \in [d]$, the set

$$\mathcal{S}_{-a}^m := \text{Conv}(\{v_j \in \{-a, 0\}, \#\{j : v_j = -a\} \geq m\}), \quad (69)$$

where $\text{Conv}(\mathcal{A})$ is the convex hull of the set \mathcal{A} . Similarly define

$$\bar{\mathcal{S}}_{-a}^m := \text{Conv}(\{v_j \in [-\infty, -a] \cup \{0\}, \#\{j : v_j \leq -a\} \geq m\}), \quad (70)$$

and lastly define

$$\hat{\mathcal{S}}_{-a}^m := \{v_j \leq v'_j \forall j \text{ for some } v' \in \mathcal{S}_{-a}^m\}. \quad (71)$$

Note that $\mathcal{S}_{-a}^m \subseteq \bar{\mathcal{S}}_{-a}^m \subseteq \hat{\mathcal{S}}_{-a}^m$ and that $\hat{\mathcal{S}}_{-a}^m$ is also a convex set. Now consider

$$h_a(\mathbf{v}) = \infty \mathbb{1}\{\mathbf{v} \notin \bar{\mathcal{S}}_{-a}^{d-k}\}. \quad (72)$$

This function is convex as it has a convex epigraph. Its Legendre–Fenchel transform is given by

$$h_a^*(\mathbf{u}) = \sup_{\mathbf{v}} \mathbf{u}^\top \mathbf{v} - h_a(\mathbf{v}) \quad (73)$$

$$= \sup_{\mathbf{v} \in \bar{\mathcal{S}}_{-a}^{d-k}} \mathbf{u}^\top \mathbf{v} \quad (74)$$

$$\leq \sup_{\mathbf{v} \in \hat{\mathcal{S}}_{-a}^{d-k}} \mathbf{u}^\top \mathbf{v} \quad (75)$$

$$= \sup_{\mathbf{v} \in \mathcal{S}_{-a}^{d-k}} \mathbf{u}^\top \mathbf{v}, \quad (76)$$

where inequality holds because $\bar{\mathcal{S}}_{-a}^{d-k} \subseteq \hat{\mathcal{S}}_{-a}^{d-k}$ the final equality holds by definition of $\hat{\mathcal{S}}_{-a}^{d-k}$. Clearly, the inequality is equality since $\mathcal{S}_{-a}^{d-k} \subseteq \bar{\mathcal{S}}_{-a}^{d-k}$. Now consider that for any $\mathbf{v} \in \mathcal{S}_{-a}^{d-k}$, $\sum_j v_j \leq -(d-k)a$ and $v_j \geq -a \forall j$. We can choose at most k elements of \mathbf{v} to be zero, so to achieve the supremum we must choose them at the largest elements of \mathbf{u} . That leaves then that the remaining elements must be $-a$, so we have

$$h_a^*(\mathbf{u}) = -a \sum_{j>k} u_{(j)}. \quad (77)$$

With corresponding $v_j^* = -a \mathbb{1}\{j \notin \text{TOP-}k(\mathbf{u}^*)\}$. Now, taking $a \rightarrow \infty$ for \mathbf{v}^* , h_a , and h_a^* we can determine $\boldsymbol{\eta}$, f , and Ω . First, as desired,

$$\boldsymbol{\eta}_j(\mathbf{w}) = \lim_{a \rightarrow \infty} -(-a \mathbb{1}\{j \notin \text{TOP-}k(\mathbf{w})\})^{-1} \quad (78)$$

$$= \infty \mathbb{1}\{j \in \text{TOP-}k(\mathbf{w})\}. \quad (79)$$

Then, since $h_a(\mathbf{v})$ is infinite for $\mathbf{v} \notin \bar{\mathcal{S}}_{-a}^{d-k}$ and zero for $\mathbf{v} \in \bar{\mathcal{S}}_{-a}^{d-k}$, we have $f(\boldsymbol{\eta}) = 0$ with

$$\mathcal{H} = \lim_{a \rightarrow \infty} \{\boldsymbol{\eta} : -\boldsymbol{\eta}^{\odot -1} \in \bar{\mathcal{S}}_{-a}^{d-k}\} \quad (80)$$

$$= \{\boldsymbol{\eta} : \|\boldsymbol{\eta}\|_0 \leq k\}. \quad (81)$$

Lastly, we have

$$\Omega(\mathbf{w}) = \lim_{a \rightarrow \infty} -2h_a^*(\mathbf{w}^{\odot 2}) \quad (82)$$

$$= \infty \mathbb{1}\{\|\mathbf{w}\|_0 > 0\}. \quad (83)$$

B Adaptive Dropout with Additive Reparameterization

In Algorithm 1 we present one scheme for implementing adaptive dropout using an additive reparameterization via a two-pass proximal update of the variables \mathbf{w} and \mathbf{v} . This method is equivalent to an adaptive proximal stochastic gradient descent with the adaptive Tikhonov penalty.

Algorithm 1: Adaptive Dropout with Additive Reparameterization

Input: Differentiable $\mathcal{L} : \mathbb{R}^d \rightarrow \mathbb{R}$, $\hat{\boldsymbol{\eta}} : \mathbb{R}^d \rightarrow \mathcal{H}$, $\lambda > 0$, $(\rho_t)_{t=1}^T$, \mathbf{w}^0 , $\boldsymbol{\alpha}^0$.

Output: \mathbf{w}^T .

$\mathbf{w}^{0,2} = \mathbf{w}^0$.

for $t = 1, 2, \dots, T$ **do**

Draw $\mathbf{s}^t \sim \text{MASK}(\boldsymbol{\alpha}^{t-1})$.

$\mathbf{w}^{t,1} = \mathbf{w}^{t-1,2} - \rho_t \nabla_{\mathbf{w}} \mathcal{L}(\mathbf{w}^{t-1,2} + (\mathbf{s}^t - \mathbf{1}) \odot \mathbf{v}^{t-1,2})$.

$\mathbf{v}^{t,1} = \mathbf{w}^{t,1}$.

$\boldsymbol{\eta}^t = \hat{\boldsymbol{\eta}}(\mathbf{v}^{t,1})$.

$\mathbf{v}^{t,2} = (\rho_t \lambda \text{diag}(\boldsymbol{\eta}^t)^{-1} + \mathbf{I})^{-1} \mathbf{v}^{t,1}$.

$\mathbf{w}^{t,2} = \mathbf{v}^{t,2}$.

$\alpha_j^t = \frac{\eta_j^t}{\eta_j^t + \lambda} \forall j \in [d]$.

end

C Experimental Details

We use the PyTorch [30] and skorch [34] libraries to implement deep network methods. On an Nvidia 980 Ti GPU, the experiment runs in about an hour. We randomly divide the MNIST training set into training and validation sets with an 80/20 split. For methods involving optimization in $\log(\boldsymbol{\eta})$, we optimize instead in $\log(\bar{\boldsymbol{\eta}})$ for $\bar{\boldsymbol{\eta}} = \boldsymbol{\eta}/\lambda$, as Molchanov et al. [29] do. We initialize with $\log(\bar{\eta}_j) = 5$. For the VARDROP methods, we use the dual penalty $f(\bar{\boldsymbol{\eta}})$ and implement the methods using code provided by the authors [2]. For other methods, we simply use the LOGSUM(2) penalty (based on Figure 1) applied to $\boldsymbol{\eta}$ directly, along with a larger value of λ to account for the implicit attenuation of the Tikhonov regularization due to dropout with the cross-entropy loss. For all methods, we use the Adam optimizer with a linear decay to 0 of the initial learning rate. The initial learning rate is set

to be 10^{-4} , but for a few methods this failed to converge to a sparse solution, so we increased it to 10^{-3} . For VARDROP, convergence was quite slow; running for a longer number of epochs, however, does continue to improve the sparsity. Running for 1000 epochs, for example, gets the fraction of nonzeros down to around 0.1, at a slight expense of accuracy.

We measure sparsity using the same method as Molchanov et al. [29]: we count the values of $\bar{\eta}$ such that $\sigma(\bar{\eta}_j) < 0.05$, and we zero out the corresponding w_j when applying the network to a validation/test sample. For η -TRICK, we observed that while the parameters \mathbf{w} were indeed converging to sparse solutions, the η parameters were not, resulting in a mismatch of the actual sparsity of the network and our reported score; to remedy this, we apply a very small penalty of $\lambda \cdot 10^{-3} \log(\bar{\eta})$, which did not seem to compromise network accuracy.

Table 4: Hyperparameters and final results for sparsification of LeNet-300-100.

Method	λ	Learning Rate	Test Error	Fraction of Nonzeros
VARDROP+LR+AR	$\frac{1}{60,000}$	10^{-4}	3.21%	0.024
VARDROP+LR	$\frac{1}{60,000}$	10^{-3}	1.41%	0.088
VARDROP	$\frac{1}{60,000}$	10^{-3}	1.54%	0.595
η -TRICK	10^{-3}	10^{-3}	2.16%	0.051
ADAPROX	10^{-3}	10^{-4}	2.94%	0.028
ADATIKHONOV	10^{-3}	10^{-4}	2.88%	0.018
LOGSUM	10^{-3}	10^{-4}	2.93%	0.019