
Stochastic Pruning: Fine-Tuning, and PAC-Bayes bound optimization

Soufiane Hayou*
Department of Statistics
University of Oxford

Bobby He
Department of Statistics
University of Oxford

Gintare Karolina Dziugaite*†
Service Now
Mila

Abstract

We introduce an algorithmic framework for stochastic fine-tuning of pruning masks, starting from masks produced by several baselines. We further show that by minimizing a PAC-Bayes bound with data-dependent priors [1], we obtain a self-bounded learning algorithm with numerically tight bounds. In the linear model, we show that a PAC-Bayes generalization error bound is controlled by the magnitude of the change in feature alignment between the “prior” and “posterior” data.

1 Introduction

The demand to reduce the size of networks during and after training has led to a rush of work in network pruning [2–6, etc.], much of it building on pioneering efforts carried out decades ago [7–9]. Many approaches to pruning begin by masking a subset of weights to zero. In approaches applied after training, this step typically causes a drop in both train and test accuracy, which is then often addressed by fine-tuning (or retraining). Intuitively, fine-tuning can be seen as looking for *locally optimal weights* given the mask. In this work, we ask whether these *masks* are optimal given the weights, even in a local sense. To approach this problem, we introduce a stochastic optimization framework for masks. Rather than viewing the application of a mask as a deterministic procedure, we view it as a stochastic operation, where each weight w_i is pruned with some probability $1 - \lambda_i$ such that $\sum_i \lambda_i = (1 - s)D$, where D is the total number of weights and $s \in [0, 1]$ is the desired sparsity level. We then consider optimizing the expected empirical risk of this induced randomized predictor with respect to the pruning probabilities.

To gain some analytical insight, we study the case of linear regression in Section 7, where we find that optimizing the pruning probabilities behaves like dropout [10] with data-dependent weight-specific rates. We show that optimizing the pruning probabilities looks like training the mean network subject to data-and-weight-adaptive L_1 regularization of the pruning probabilities. This analysis on the linear model can be extended to wide DNNs in the context of Neural Tangent Kernel regime [11].

We consider optimizing stochastic pruning masks, starting from initial pruning probabilities derived from magnitude pruning [2], SNIP [4], and random pruning. We observe improved performance after optimizing the pruning mask on a variety of datasets. Finally, by observing that pruned networks are more robust to perturbations of the weights, we consider replacing deterministic fine-tuning with PAC-Bayes bound optimization. Here we consider spike-and-slab prior and posterior distributions, as was previously investigated by Louizos [12] in the context of variational inference. Different from earlier work, we add data-dependence to the priors, as advocated recently by Dziugaite et al. [1]. On MNIST [13], Fashion-MNIST [14], and CIFAR-10 [15], we are able to obtain competitive predictors along with numerically tight, nonvacuous generalization bounds.

*Correspondence to: soufiane.hayou@yahoo.fr and gkdz@google.com

†This work was carried out while the author was at ServiceNow. It was finalized at Google Brain.

2 Background material

Stochastic Pruning Masks We consider *stochastic* pruning masks, i.e. b is a vector of Bernoulli random variables. We focus on the case where elements of the mask are independent, so that the distribution of the mask is determined by probabilities $\lambda_i = \Pr\{b_i = 1\}$. For a fixed set of weights W , the stochasticity of the mask induces a distribution μ on pruned weights $W' = (w'_i)_{i=1}^D$. In particular, conditional on W , the components of W' are independent and w'_i has distribution $(1 - \lambda_i)\delta_0 + \lambda_i\delta_{w_i}$.

Data-dependent PAC-Bayes Bounds. PAC-Bayes theorems provide confidence intervals for Gibbs risks $R_{\mathcal{D}}(Q)$ in terms of an empirical Gibbs risk $\hat{R}_S(Q)$ and the Kullback–Liebler (KL) divergence between any distribution Q , and a distribution P , which cannot depend on all of the training data S . Many forms of PAC-Bayes theorems exist. In our work we consider a PAC-Bayes bound with a data-dependent prior, meaning that the prior *depends on* a subset of the training data, $S_P \subset S$. Let $\alpha = |S_P|/|S| \in [0, 1]$. In this work we typically use $\alpha \in [0.4, 0.8]$. In our experimental setting, we optimize through the following bound due to [1, Thm. D.2] and [16, Sec. 1.3]³.

Theorem 1. *For every $[0, 1]$ -valued loss function ℓ , data distribution \mathcal{D} , natural $N \in \mathbb{N}$, probability $\delta \in (0, 1)$, with probability at least $1 - \delta$ over training data $S \sim \mathcal{D}^N$ and all distributions P on \mathcal{F} that depend on $S_P \subset S$,*

$$\forall Q, R_{\mathcal{D}}(Q) \leq \hat{R}_{S \setminus S_P}(Q) + \min \left\{ \varepsilon + \sqrt{\varepsilon(\varepsilon + 2\hat{R}_{S \setminus S_P}(Q))}, \varepsilon = \frac{\text{KL}(Q \parallel P) + \log \frac{2\sqrt{(1-\alpha)|S|}}{\delta}}{(1-\alpha)|S|} \right\}. \quad (1)$$

Larger values of α yield a better prior P with a smaller $\text{KL}(Q \parallel P)$. However, the KL term in Eq. (3) is divided by $(1 - \alpha)|S|$, and increasing α may lead to a looser bound. Therefore, the final risk bound trades-off the decrease in $\text{KL}(Q(S) \parallel P(S_P))$ and the change in $\text{KL}(Q(S) \parallel P(S_P))/(1 - \alpha)|S|$.

Data-dependent PAC-Bayes bounds previously appeared in many articles, usually as an attempt to empirically approximate distribution-dependent bounds (e.g., [1, 16–25]).

Optimizing PAC-Bayes bounds. While for certain forms of PAC-Bayes bounds there exist analytical solutions for the optimal Q , the resulting predictors are generally intractable. Instead, one typically performs explicit optimization over a parameterized family of posteriors. A common choice, which can be justified on theoretical grounds asymptotically [26], is to optimize Q over the class of Gaussian distributions, and choose the prior P to be Gaussian as well. We restrict our attention to diagonal covariance matrices when studying neural networks, following Dziugaite and Roy [27].

The prior P and posterior Q . In order to have a tractable KL divergence we consider a spike-and-slab prior P and posterior Q . The weights are a priori/posteriori independent and each weight w_i has a distribution

$$p_i = (1 - \lambda_i^0)\delta_{\{0\}} + \lambda_i^0\mathcal{N}(W_i^0, (\sigma_i^0)^2), \quad q_i = (1 - \lambda_i)\delta_{\{0\}} + \lambda_i\mathcal{N}(W_i^f, (\sigma_i^f)^2). \quad (2)$$

The parameters λ_i , W_i^f , and σ_i^f can be optimized while leaving the PAC-Bayes bound valid.

Loss function. For classification, the natural loss function is the zero–one loss function $\ell_{0-1}(\hat{y}, y) = 1_{\hat{y} \neq y}$. However, since we will optimize the PAC-Bayes risk bound directly using gradient-based optimization, we use cross entropy, a differentiable surrogate loss. When reporting bounds in figures, we use zero–one loss in order to obtain valid bounds on classification risk⁴.

3 Stochastic Pruning

³See [1, Fig. 6] for the discussion on the tightness of the bounds in different regimes.

⁴Note that this procedure is valid since the bound is valid for any Q , even the one optimized on a loss different than the loss appearing in the bound. See Appendix A.3 for more details.

We introduce two algorithms, *Probabilistic Fine-Tuning (PFT)* and *PAC-Bayes Pruning (PBP)*, which we subsequently analyze theoretically (Section 7) and experimentally (Section 8). Our algorithmic framework has three stages: **1)** Pre-training the prior by minimizing $\hat{R}_{S_P}(W)$ on a subset S_P of training data, to obtain a deterministic predictor with weights \hat{W}_P . **2)** Fine-tuning the stochastic prior: compute $\hat{\lambda}_i^0$ based on one of the pruning methods, and minimize the empirical risk on S_P where minimization is performed over W_0 , $\{\lambda_i^0\}_{i=1}^D$, and $\{\sigma_i\}_{i=1}^D$ ⁵. For PFT, the trained prior probabilities λ_i^0 's are used to create a new *deterministic* mask by keeping the weights with the largest λ_i^0 . **3)** Stage 3: Train the posterior (used exclusively in PBP), using all the training data S to minimize the PAC-Bayes bound, Eq. (3), with respect to $\{\lambda_i\}_{i=1}^D$, and W_i^f . Fig. 4 details different stages of PBP as a flowchart, highlighting the differences between PBP and PFT, as well as standard One-Shot Pruning (OSP).

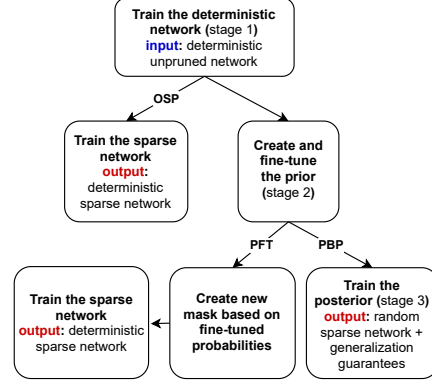


Figure 1: OSP (One Shot Pruning), PFT (Probabilistic Fine-Tuning), and PBP (PAC-Bayes Pruning). Fig. 4 details different stages of PBP as a flowchart, highlighting the differences between PBP and PFT, as well as standard One-Shot Pruning (OSP). We analyze PFT, separately to PBP, in Section 7.1 and Section 8.1.

4 Experiments

In interest of space, the full details on the networks, datasets, hyperparameters, optimizers, and loss functions are provided in Appendix D. Appendix E contains a number of additional experiments on PFT & PBP.

4.1 Probabilistic fine-tuning boosts the accuracy of sparse predictors

One-Shot Pruning (OSP) methods create the mask in a single or no pass over the dataset. Popular choices include pruning at initialization methods [4–6, 28]. OSP later in training was discussed in [29]. These methods follow the same protocol: create a mask some time during training (e.g. at initialization or after training), apply the mask to the network, and train the sparse network. PFT aims to improve the performance of OSP methods using stochastic models to fine-tune the mask before training the sparse network. This entails training the prior probabilities as discussed in Section 7.1 in the context of linear models; we then create a new mask by keeping the weights with the highest (trained) probabilities and pruning the rest.

Fig. 5 shows that PFT consistently boosts the performance of standard OSP methods on small-scale vision tasks; in general, the benefit of applying PFT is most noticeable for high sparsity levels (e.g. 99%) where some OSP methods alone fail to achieve good test error. On CIFAR-10 with ResNet20, PFT improves the accuracy of OSPs by $\sim 1\%$ on average; results on CIFAR-10 with sparsity level 99% are included in Appendix E. Note that PFT with Random Pruning (RP) is comparable to other methods, hence suggesting that PFT alone (preceded by RP) can serve as a standalone pruning method. More details and experiments on PFT are provided in Appendix E.1.

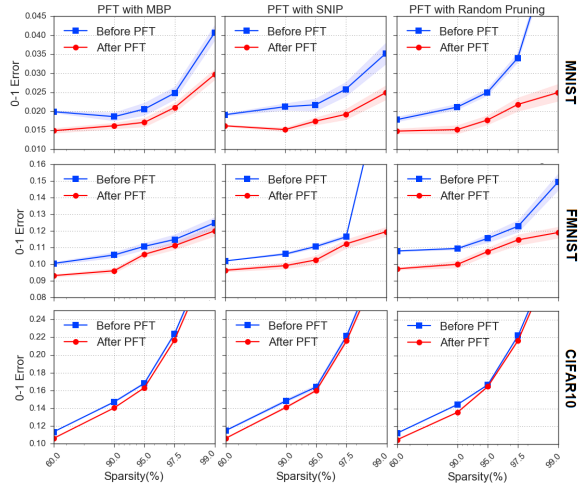


Figure 2: Impact of the application of PFT on standard One-Shot Pruning methods (MBP, SNIP) and Random Pruning. MNIST and Fashion-MNIST with a 3-layer width-1000 MLP and CIFAR-10 with ResNet20.

⁵The training with respect to λ_i^0 is performed using the Gumbel-Softmax trick. See Appendix D for more details

4.2 Self-bounded learning via PAC-Bayes

PBP on small vision tasks Finally, we evaluate our proposed stochastic pruning with PAC Bayes framework (PBP), which returns a stochastic predictor in contrast to many other pruning approaches. We consider MNIST, Fashion-MNIST & CIFAR-10⁶ with corresponding architectures to Section 8.1. In Fig. 7, we compare the PAC-Bayes generalization bounds on 0-1 error (blue) of our stochastic predictor Q to the corresponding observed test set error (orange) and the observed test error of our sparse stochastic prior P (green) over a range of sparsities and pre-training dataset fraction α . For reference, we also plot the test error of a baseline pruning method, SNIP [4], to compare our test errors and bounds to a standard *deterministic* OSP method which doesn't provide generalization guarantees.

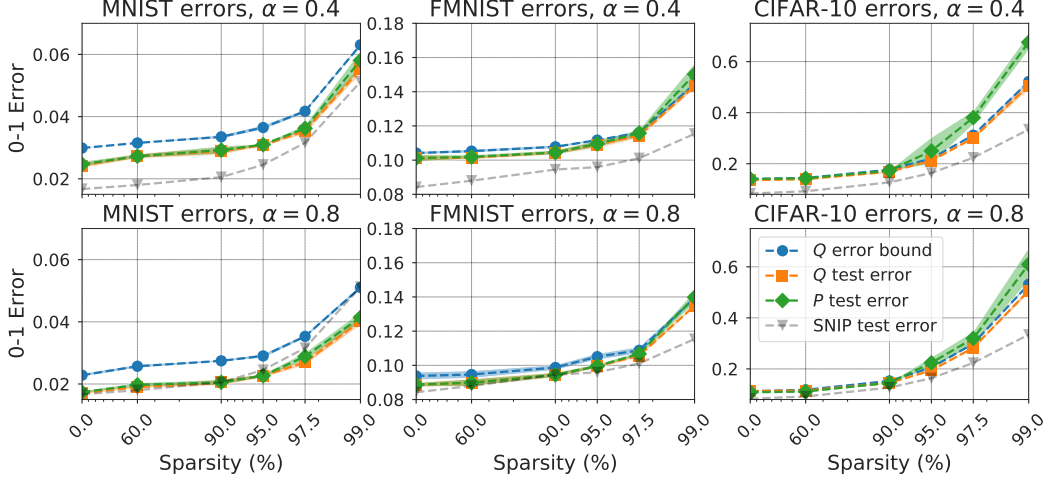


Figure 3: Error and error bound versus sparsity on image classification tasks for different $\alpha = |S_P|/|S|$. Displayed intervals correspond to 95% confidence over 5 independent seeds.

Remark. Technical details are provided in the appendix.

⁶CIFAR-10: <https://www.cs.toronto.edu/~kriz/cifar.html>, MNIST: <http://yann.lecun.com/exdb/mnist/> FMNIST: <https://github.com/zalandoresearch/fashion-mnist>

References

- [1] G. K. Dziugaite, K. Hsu, W. Gharbieh, G. Arpino, and D. M. Roy. “On the role of data in PAC-Bayes bounds.” In: *Proc. Int. Conf. Artificial Intelligence and Statistics (AISTATS)*. 2021.
- [2] S. Han, J. Pool, J. Tran, and W. J. Dally. “Learning both weights and connections for efficient neural networks.” *arXiv preprint arXiv:1506.02626* (2015).
- [3] J. Frankle and M. Carbin. “The lottery ticket hypothesis: Finding sparse, trainable neural networks.” *International Conference on Learning Representations (ICLR)* (2018).
- [4] N. Lee, T. Ajanthan, and P. H. Torr. “SNIP: Single-shot network pruning based on connection sensitivity.” *arXiv preprint arXiv:1810.02340* (2018).
- [5] C. Wang, G. Zhang, and R. Grosse. “Picking Winning Tickets Before Training by Preserving Gradient Flow.” In: *International Conference on Learning Representations (ICLR)*. 2020.
- [6] S. Hayou, J. Ton, A. Doucet, and Y. Teh. “Robust Pruning at Initialization.” In: *International Conference on Learning Representations (ICLR)*. 2021.
- [7] M. Mozer and P. Smolensky. “Skeletonization: A technique for trimming the fat from a network via relevance assessment.” In: *Advances in Neural Information Processing Systems*. 1989, pp. 107–115.
- [8] Y. LeCun, J. Denker, and S. Solla. “Optimal brain damage.” In: *Advances in Neural Information Processing Systems*. 1990, pp. 598–605.
- [9] B. Hassibi, D. Stork, and W. Gregory. “Optimal Brain Surgeon and general network pruning.” In: *IEEE International Conference on Neural Networks*. 1993, 293–299 vol.1.
- [10] 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* 15.1 (2014), pp. 1929–1958.
- [11] A. Jacot, F. Gabriel, and C. Hongler. “Neural Tangent Kernel: Convergence and Generalization in Neural Networks.” *32nd Conference on Neural Information Processing Systems* (2018).
- [12] C. Louizos. “Smart Regularization of Deep Architectures.” PhD thesis. University of Amsterdam, 2015.
- [13] Y. LeCun and C. Cortes. “MNIST handwritten digit database” (2010).
- [14] H. Xiao, K. Rasul, and R. Vollgraf. “Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms.” *arXiv preprint arXiv:1708.07747* (2017).
- [15] A. Krizhevsky. *Learning multiple layers of features from tiny images*. Tech. rep. 2009.
- [16] O. Catoni. *PAC-Bayesian supervised classification: the thermodynamics of statistical learning*. Vol. 56. Lecture Notes-Monograph Series. IMS. 2007. arXiv: [0712.0248](#).
- [17] G. Lever, F. Laviolette, and J. Shawe-Taylor. “Distribution-dependent PAC-Bayes priors.” In: *Proc. Int. Conf. Algorithmic Learning Theory (ALT)*. Springer. 2010, pp. 119–133.
- [18] A. Ambroladze, E. Parrado-Hernández, and J. Shawe-Taylor. “Tighter PAC-Bayes bounds.” In: *Adv. Neural Information Processing Systems*. 2007, pp. 9–16.
- [19] J. Negrea, M. Haghifam, G. K. Dziugaite, A. Khisti, and D. M. Roy. “Information-Theoretic Generalization Bounds for SGLD via Data-Dependent Estimates.” In: *Adv. Neural Information Processing Systems*. 2019, pp. 11013–11023.
- [20] M. Haghifam, J. Negrea, A. Khisti, D. M. Roy, and G. K. Dziugaite. “Sharpened generalization bounds based on conditional mutual information and an application to noisy, iterative algorithms.” In: *Adv. Neural Information Processing Systems (NeurIPS)*. 2020.
- [21] L. Oneto, D. Anguita, and S. Ridella. “PAC-Bayesian analysis of distribution dependent priors: Tighter risk bounds and stability analysis.” *Pattern Recognition Letters* 80 (2016), pp. 200–207. ISSN: 0167-8655.
- [22] L. Oneto, S. Ridella, and D. Anguita. “Differential privacy and generalization: Sharper bounds with applications.” *Pattern Recognition Letters* 89 (2017), pp. 31–38. ISSN: 0167-8655.
- [23] G. K. Dziugaite and D. M. Roy. “Data-dependent PAC-Bayes priors via differential privacy.” In: *Adv. Neural Information Processing Systems*. 2018. arXiv: [1802.09583](#).
- [24] O. Rivasplata, E. Parrado-Hernandez, J. Shaws-Taylor, S. Sun, and C. Szepesvari. “PAC-Bayes bounds for stable algorithms with instance-dependent priors.” In: *Adv. Neural Information Processing Systems*. 2018.
- [25] P. Awasthi, S. Kale, S. Karp, and M. Mohri. “PAC-Bayes Learning Bounds for Sample-Dependent Priors.” *Advances in Neural Information Processing Systems* 33 (2020).

- [26] P. Alquier, J. Ridgway, and N. Chopin. “On the properties of variational approximations of Gibbs posteriors.” *Journal of Machine Learning Research* 17.236 (2016), pp. 1–41.
- [27] G. K. Dziugaite and D. M. Roy. “Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data.” In: *Proc. Int. Conf. Uncertainty in Artificial Intelligence (UAI)*. 2017.
- [28] H. Tanaka, D. Kunin, D. L. Yamins, and S. Ganguli. “Pruning neural networks without any data by iteratively conserving synaptic flow.” *arXiv preprint arXiv:2006.05467* (2020).
- [29] J. Frankle, G. Dziugaite, D. Roy, and M. Carbin. “Pruning Neural Networks at Initialization: Why are We Missing the Mark?” *International Conference on Learning Representations (ICLR)* (2021).
- [30] Y. Freund. “Self bounding learning algorithms.” In: *Proceedings of the eleventh annual conference on Computational Learning Theory*. 1998, pp. 247–258.
- [31] S. Wager, S. Wang, and P. Liang. “Dropout Training as Adaptive Regularization.” In: *Adv. Neural Information Processing*. 2013.
- [32] S. Wang and C. Manning. “Fast dropout training.” In: *International Conference on Machine Learning (ICML)*. 2013, pp. 118–126.
- [33] B. Elesedy, V. Kanade, and Y. W. Teh. “Lottery tickets in linear models: An analysis of iterative magnitude pruning.” *arXiv preprint arXiv:2007.08243* (2020).
- [34] C. Wang, G. Zhang, and R. Grosse. “Picking winning tickets before training by preserving gradient flow.” *International Conference on Representation Learning (ICLR)* (2020).
- [35] P. Savarese, H. Silva, and M. Maire. “Winning the lottery with continuous sparsification.” *arXiv preprint arXiv:1912.04427* (2019).
- [36] T. Gale, E. Elsen, and S. Hooker. “The state of sparsity in deep neural networks.” *arXiv preprint arXiv:1902.09574* (2019).
- [37] C. Louizos, M. Welling, and D. P. Kingma. “Learning Sparse Neural Networks through L_0 Regularization.” *arXiv preprint arXiv:1712.01312* (2017).
- [38] D. Molchanov, A. Ashukha, and D. Vetrov. “Variational dropout sparsifies deep neural networks.” In: *International Conference on Machine Learning*. PMLR. 2017, pp. 2498–2507.
- [39] M. Zhu and S. Gupta. “To prune, or not to prune: exploring the efficacy of pruning for model compression.” *arXiv preprint arXiv:1710.01878* (2017).
- [40] J. Frankle, G. K. Dziugaite, D. Roy, and M. Carbin. “Linear mode connectivity and the lottery ticket hypothesis.” In: *International Conference on Machine Learning*. PMLR. 2020, pp. 3259–3269.
- [41] A. Renda, J. Frankle, and M. Carbin. “Comparing rewinding and fine-tuning in neural network pruning.” *arXiv preprint arXiv:2003.02389* (2020).
- [42] J. Langford and M. Seeger. “Bounds for averaging classifiers.” *Technical report, Carnegie Mellon, Department of Computer Science* (2001).
- [43] R. D. Luce. “Individual Choice Behavior: A Theoretical Analysis.” Wiley, New York. (1959).
- [44] C. J. Maddison, A. Mnih, and Y. W. Teh. “The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables.” *International Conference on Representation Learning (ICLR)* (2017).
- [45] J. Lee, L. Xiao, S. S. Schoenholz, Y. Bahri, R. Novak, J. Sohl-Dickstein, and J. Pennington. “Wide neural networks of any depth evolve as linear models under gradient descent.” *arXiv preprint arXiv:1902.06720* (2019).
- [46] C. Baykal, L. Liebenwein, I. Gilitschenski, D. Feldman, and D. Rus. “Data-dependent coresets for compressing neural networks with applications to generalization bounds.” *arXiv preprint arXiv:1804.05345* (2018).
- [47] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. “Gradient-based learning applied to document recognition.” *Proceedings of the IEEE* 86.11 (1998), pp. 2278–2324.
- [48] K. He, X. Zhang, S. Ren, and J. Sun. “Deep Residual Learning for Image Recognition.” In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016, pp. 770–778.

5 Background material

In this section, we introduce background on pruning, PAC-Bayes, and data-dependent priors.

5.1 Pruning

Consider the problem of classifying input vectors in \mathbb{R}^d with labels in $\mathcal{Y} = \{1, \dots, K\}$ for some $K \in \mathbb{N}$ according to some predictor in a class $\mathcal{F} = \{f_W : \mathbb{R}^d \rightarrow \mathbb{R}^K \mid W \in \mathbb{R}^D\}$ of predictors, indexed by weight vectors $W \in \mathbb{R}^D$. Let $S = (x_i, y_i)_{1 \leq i \leq N}$ denote the training data, consisting of N i.i.d. samples from an unknown distribution \mathcal{D} on $\mathbb{R}^d \times \mathcal{Y}$. Let $\hat{R}_S(f) = \frac{1}{N} \sum_{i=1}^N \ell(f(x_i), y_i)$, denote the *empirical risk* of the predictor f (with respect to a loss function $\ell : \mathbb{R}^K \times \mathcal{Y} \rightarrow \mathbb{R}_+$ which will often be implicit). Let $R_{\mathcal{D}}(f) = \mathbb{E}_{(x,y) \sim \mathcal{D}}[\ell(f(x), y)]$, denote the *risk*. We will sometimes write $\hat{R}_S(W)$ for $\hat{R}_S(f_W)$, and similarly for the risk.

Pruning involves the application of a binary mask $b \in \{0, 1\}^D$ to the weights of a predictor f_W producing the predictor $f_{W'}$, where $W' = b \circ W$ is the Hadamard (i.e., element-wise) product. We say a weight $i \in [D]$ is pruned if $b_i = 0$. We assume that \mathcal{F} is such that $b_i = 0$ implies $\frac{\partial f_W(x)}{\partial w_i} = 0$ for all $x \in X$.

Pruning Masks A standard approach to generating masks is to give each weight w_i a score g_i according to some criterion. The mask is then created by keeping the top k weights by score, where k is chosen to meet some desired level of sparsity. Examples of pruning criterion include:

- **Magnitude-based:** the weights are ranked by magnitude, corresponding to $g_i = |W_i|$ [2];
- **Sensitivity-based:** the weights are scored based on how much they will impact the empirical risk [4, 8]. For example, SNIP [4] considers the first-order approximation of the training loss $\tilde{\ell}_W \approx \tilde{\ell}_{W=0} + W \frac{\partial \tilde{\ell}}{\partial W}$ and yields the pruning criterion $g_i = |W_i \frac{\partial \tilde{\ell}}{\partial W_i}|$.

Stochastic Pruning Masks In this work, we consider *stochastic* pruning masks, i.e. b is a vector of Bernoulli random variables. We focus on the case where elements of the mask are independent, so that the distribution of the mask is determined by probabilities $\lambda_i = \Pr\{b_i = 1\}$. For a fixed set of weights W , the stochasticity of the mask induces a distribution μ on pruned weights $W' = (w'_i)_{i=1}^D$. In particular, conditional on W , the components of W' are independent and w'_i has distribution $(1 - \lambda_i)\delta_0 + \lambda_i\delta_{w_i}$, where δ_v denotes the Dirac mass assigning probability 1 to the value v .

Optimizing stochastic pruning masks Stochastic pruning masks implicitly induce distributions on the class \mathcal{F} of predictors: We can think of μ as a stochastic predictor that, on each input x , samples $f \sim \mu$ and then outputs $f(x)$. Such a stochastic predictor is known as a Gibbs predictor. The *empirical Gibbs risk* and *Gibbs risk* of μ are, respectively,

$$\hat{R}_S(\mu) = \mathbb{E}_{f \sim \mu}[\hat{R}_S(f)] \quad \text{and} \quad R_{\mathcal{D}}(\mu) = \mathbb{E}_{f \sim \mu}[R_{\mathcal{D}}(f)].$$

We can find stochastic pruning masks that produce sparse networks with small empirical risk by minimizing the empirical Gibbs risk with respect to the pruning probabilities. A discussion on how we initialize the pruning probabilities is provided in Appendix A.1.

5.2 Data-dependent PAC-Bayes Bounds

PAC-Bayes theorems provide confidence intervals for Gibbs risks $R_{\mathcal{D}}(Q)$ in terms of an empirical Gibbs risk $\hat{R}_S(Q)$ and the Kullback–Liebler (KL) divergence between any distribution Q , and a distribution P , which cannot depend on all of the training data S . In PAC-Bayes parlance, P is called a “prior” and each Q is referred to as a “posterior”, on the account that the bound applies also to distributions Q that may depend on the training data S .

Many forms of PAC-Bayes theorems exist. In our work we consider a PAC-Bayes bound with a data-dependent prior, meaning that the prior *depends on* a subset of the training data, $S_P \subset S$. Let $\alpha = |S_P|/|S| \in [0, 1)$. In this work we typically use $\alpha \in [0.4, 0.8]$. In our experimental setting, we optimize through the following bound due to [1, Thm. D.2] and [16, Sec. 1.3]⁷.

⁷See [1, Fig. 6] for the discussion on the tightness of the bounds in different regimes.

Theorem 2. For every $[0, 1]$ -valued loss function ℓ , data distribution \mathcal{D} , natural $N \in \mathbb{N}$, probability $\delta \in (0, 1)$, with probability at least $1 - \delta$ over training data $S \sim \mathcal{D}^N$ and all distributions P on \mathcal{F} that depend on $S_P \subset S$,

$$\forall Q, R_{\mathcal{D}}(Q) \leq \hat{R}_{S \setminus S_P}(Q) + \min \left\{ \varepsilon + \sqrt{\varepsilon(\varepsilon + 2\hat{R}_{S \setminus S_P}(Q))}, \sqrt{\varepsilon/2}, \right\} \quad (3)$$

where $\varepsilon = \frac{\text{KL}(Q \parallel P) + \log \frac{2\sqrt{(1-\alpha)|S|}}{\delta}}{(1-\alpha)|S|}$

As one may expect, larger α yields a better prior P with a smaller $\text{KL}(Q \parallel P)$. However, the KL term in Eq. (3) is divided by $(1 - \alpha)|S|$, and increasing α may lead to a looser bound. Therefore, the final risk bound trades-off the decrease in $\text{KL}(Q(S) \parallel P(S_P))$ and the change in $\frac{\text{KL}(Q(S) \parallel P(S_P))}{(1-\alpha)|S|}$.

Data-dependent PAC-Bayes bounds previously appeared in many articles, usually as an attempt to empirically approximate distribution-dependent bounds (e.g., [1, 16–25]).

Informally, in the context of stochastic pruning, the PAC-Bayes theorem says that if we do not change the pruning probabilities too much as we get more data, i.e., from P to Q , then we will have a small generalization gap.

Optimizing PAC-Bayes bounds. In light of Theorem 2, one approach to learning is to minimize the upper bound on the right-hand side in terms of Q .⁸ In this work, we consider the application of PAC-Bayes bound optimization to optimizing stochastic pruning masks and fine-tuning the weights.

While for certain forms of PAC-Bayes bounds there exist analytical solutions for the optimal Q , the resulting predictors are generally intractable. Instead, one typically performs explicit optimization over a parameterized family of posteriors. A common choice, which can be justified on theoretical grounds asymptotically [26], is to optimize Q over the class of Gaussian distributions, and choose the prior P to be Gaussian as well. For computational reasons, it is common to also restrict attention to diagonal covariance matrices when studying neural networks, following Dziugaite and Roy [27].

We are interested in simultaneously training the weights of a neural network and the pruning probabilities and so we must go beyond vanilla Gaussian priors and posteriors.

The prior P and posterior Q . We consider a spike-and-slab prior P whose slab is a multivariate Gaussian P^0 with mean weight vector $W^0 \in \mathbb{R}^D$ and diagonal covariance with the i^{th} entry $(\sigma_i^0)^2 \in \mathbb{R}_{>0}$. That is, the weights are a priori independent and each weight w_i has a distribution

$$p_i = (1 - \lambda_i^0)\delta_{\{0\}} + \lambda_i^0 \mathcal{N}(W_i^0, (\sigma_i^0)^2). \quad (4)$$

In order to have a tractable KL divergence, we also optimize over the family of a spike-and-slab posteriors whose slabs are multivariate Gaussians Q^f with mean weight vector W^f and diagonal covariance with entries $(\sigma_i^f)^2$. We denote the marginal posteriors by q_i , in resemblance to p_i , and the posterior probabilities by λ_i . The parameters λ_i , W_i^f , and σ_i^f can be optimized while leaving the PAC-Bayes bound valid.

Kullback–Leibler divergence. With P and Q as above, $\text{KL}(Q \parallel P) = \sum_{i=1}^D \text{KL}(q_i \parallel p_i)$ where $\text{KL}(q_i \parallel p_i)$ is⁹ $\text{kl}(\lambda_i \parallel \lambda_i^0) + \frac{\lambda_i}{2} (\psi(\sigma_i^f/\sigma_i^0) + (W_i^f - W_i^0)^2/(\sigma_i^f)^2)$, for $\psi(r) = r - 1 - \ln r$. The first term penalizes σ_i^f in terms of its multiplicative deviation from σ_i^0 . The second term penalizes the squared difference between the weights, scaled by the prior variance in that dimension.

Up until this point, we have not discussed loss functions. If the task is classification, the natural loss function is the zero–one loss function $\ell_{0-1}(\hat{y}, y) = 1_{\hat{y} \neq y}$. However, since we will optimize the PAC-Bayes risk bound directly using gradient-based optimization, we use cross entropy, a differentiable surrogate loss. When reporting bounds in figures, we must use zero–one loss in order to obtain valid bounds on classification risk. Note that this procedure is valid since the bound is valid for any Q ,

⁸Because the KL divergence is finite only if Q is absolutely continuous with respect to P , this approach to learning produces a stochastic predictor.

⁹See Appendix B.4 for the explicit expression of KL.

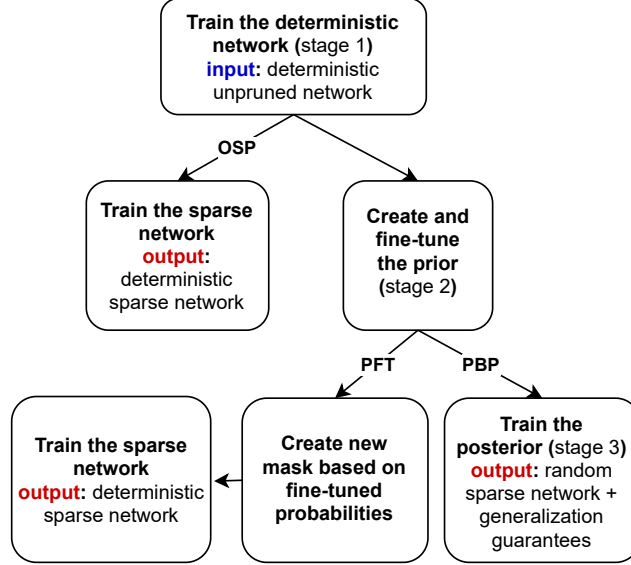


Figure 4: OSP (One Shot Pruning), PFT (Probabilistic Fine-Tuning), and PBP (PAC-Bayes Pruning).

even the one optimized on a loss different then the loss appearing in the bound. See Appendix A.3 for more details.

6 Stochastic Pruning

We now describe a general framework combining pruning, PAC Bayes, and data-dependent priors and introduce two algorithms, *Probabilistic Fine-Tuning* (PFT) and *PAC-Bayes Pruning* (PBP), which we subsequently analyze theoretically (Section 7) and experimentally (Section 8). Our algorithmic framework has three stages:

Stage 1: Pre-training the prior. Using a variant of stochastic gradient descent, minimize $\hat{R}_{S_P}(W)$ on a subset S_P of training data, to obtain a deterministic predictor with weights \hat{W}_P .

Stage 2: Fine-tuning the stochastic prior (used in PFT and PBP). Compute prior pruning probabilities $\hat{\lambda}_i^0$ based on one of the methods described in Section 5.1. If a criterion g is used, evaluate it on \hat{W}_P . $\hat{\lambda}_i^0$ and \hat{W}_P correspond to our initialized values of λ_i^0 and W_i^0 , respectively, in Eq. (4). From these initialized prior values, continue training by minimizing the empirical risk on S_P of the sparse stochastic network with distribution defined in Eq. (4), where minimization is performed over W_0 , $\{\lambda_i^0\}_{i=1}^D$, and $\{\sigma_i\}_{i=1}^D$. The training with respect to λ_i^0 is performed using the Gumbel-Softmax trick. See Appendix D for more details. For PFT, the trained prior probabilities λ_i^0 's are used to create a new *deterministic* mask by keeping the weights with the largest λ_i^0 .

Stage 3: Training the posterior (used exclusively in PBP). Continue training the posterior Q , using all the training data S to minimize the PAC-Bayes bound, Eq. (3), with respect to $\{\lambda_i\}_{i=1}^D$, and W_i^f .

It is worth noting that the above three-stage PBP algorithm encompasses the scope of our considerations in this work in full generality, and changing “hyperparameters” such as $|S_P|$ can and will lead to different procedures. For example, if we let $|S_P|=|S|$ and thus forgo stage 3, we see that stage 2 is by itself a viable method for improving a deterministic pruning mask (this is the essence of PFT), albeit without the generalization guarantees that the PAC-Bayes machinery offers. For example, one application of PFT could be to take a pruning mask obtained at initialization, and to fine-tune the mask, in addition to the weights, on training data. In this sense, PFT serves a different purpose to the

*self-bounded*¹⁰ PBP algorithm. Fig. 4 details different stages of PBP as a flowchart, highlighting the differences between PBP and PFT, as well as standard One-Shot Pruning (OSP). We analyze PFT, separately to PBP, in Section 7.1 and Section 8.1.

To the best of our knowledge, PFT is the first algorithm to combine deterministic pruning with probabilistic fine-tuning in order to find better local masks that boost accuracy (Section 8). PBP is the first algorithm to combine stochastic pruning and PAC-Bayes optimization. One benefit offered by PBP’s combination of PAC-Bayes and pruning is access to tight generalization guarantees that can outperform held-out validation bounds by using all available training data in the training process and bound (Section 8.2).

7 Linear model analysis

In this section, we present a theoretical analysis of PFT and PBP in the case of linear models. Recall that S denotes the full training set of size N and $S_P \subset S$ denotes the subset of the data that the data-dependent prior P has access to, with $|S_P| = M$, for some $M \leq N$. Let $X \in \mathbb{R}^{N \times d}$ denote the matrix of inputs and $Y \in \mathbb{R}^{N \times 1}$ the vector of outputs, and similarly define X_P, Y_P for the data in S_P , with $S_{\bar{P}}$ denoting the data not in S_P , i.e., $S \setminus S_P$. Let $\psi : \mathbb{R}^d \mapsto \mathbb{R}^D$ be a zero-mean feature map, and write $\psi_i(X) \in \mathbb{R}^N$ for the vector of the values of the i^{th} feature, i.e., the vector whose j^{th} entry is $\psi_i(x_j)$ for $x_j \in X$. For the empirical covariance of the features, write $\Sigma_P = \frac{1}{M} \sum_{x \in X_P} \psi(x) \psi(x)^T$. In this section, we study a linear predictor $f(x) = \psi(x)^T w$, with weights w , trained using squared error loss $\ell(f(x), y) = (f(x) - y)^2/2$.

Consider a simple 3-stage training procedure, closely matching the 3-stage procedure in Section 6:

1. Train the linear model with (X_P, Y_P) . Let w^P denote the trained weights.
2. Learn prior pruning probabilities $\{\lambda_i^0\}_{i=1}^D$ by minimizing $\hat{R}_S(P)$, where the prior P is a product distribution $\prod_{i=1}^D p_i$, with $p_i = (1 - \lambda_i^0) \delta_{\{0\}} + \lambda_i^0 \delta_{\{w_i^P\}}$. Note that unlike in Stage 2 of PBP, we do not optimize the weights, but only the pruning probabilities, since we are interested in studying stochastic sparsification.
3. Optimize a PAC-Bayes bound with respect to the probabilities λ , with a Bernoulli prior and posterior, and the weight values w^P fixed. The posterior on the weight i is given by $q_i = (1 - \lambda_i) \delta_{\{0\}} + \lambda_i \delta_{\{w_i^P\}}$, for all $i \in [D]$.

Letting $W \sim P$, the empirical risk of P under square loss satisfies

$$2\hat{R}_{S_P}(P) = \frac{1}{M} \mathbb{E}[\|\psi(X_P)^T W - Y_P\|_2^2] \quad (5)$$

$$= \mathbb{E}[W^T \Sigma_P W] - 2\bar{W}^T A(P) + \frac{1}{M} \|Y_P\|_2^2, \quad (6)$$

where \bar{W} is the mean of W with entries $\bar{W}_i = \lambda_i^0 w_i^P$ and $A(P)$ is the *feature alignment vector* for the data S_P , with entries $A_i(P) = \frac{1}{|X_P|} \psi_i(X_P)^T Y_P$, i.e., $A_i(P)$ measures how closely feature i is aligned to the outputs Y_P . Then $\mathbb{E}[W^T \Sigma_P W] = \text{Tr}(\Gamma \Sigma_P)$, where

$$\Gamma_{i,j} = \begin{cases} \lambda_i^0 \lambda_j^0 w_i^P w_j^P, & \text{if } i \neq j, \\ \lambda_i^0 (w_i^P)^2, & \text{otherwise.} \end{cases} \quad (7)$$

Our analysis will focus on the derivative of Eq. (6) with respect to λ_i^0 :

$$\frac{\partial \hat{R}_{S_P}(P)}{\partial \lambda_i^0} = \sum_{j \neq i} \lambda_j^0 w_j^P w_j^P (\Sigma_P)_{ij} + \frac{1}{2} (w_i^P)^2 (\Sigma_P)_{ii} - w_i^P A_i(P). \quad (8)$$

¹⁰A self-bounded learning algorithm is an algorithm that outputs a predictor as well as a bound on the generalization error of that predictor [30].

7.1 Learning prior pruning probabilities.

We now investigate the effect of optimizing the empirical risk of the stochastic predictor P with respect to λ_i^0 's, as one does in stage 2 above. According to the following result, this optimization is equivalent to optimizing the empirical risk of the mean predictor \bar{W} with an explicit regularization term that encourages pruning. (See Appendix C.1 for a formal statement and proof.)

Lemma 1 (Informal). *Gradient descent on the empirical Gibbs error (Eq. (6)) with respect to pruning probabilities $\{\lambda_i\}_{i=1}^D$ performs implicit L_1 regularization of $\{\lambda_i\}_{i=1}^D$.*

Proof sketch. In Appendix C.1, we show that $\frac{\partial \hat{R}_{SP}(P)}{\partial \lambda_i^0} - \frac{\partial \hat{R}_{SP}(\bar{W})}{\partial \lambda_i^0} = r_i^0$, where $r_i^0 = \frac{1}{2}(w_i^P)^2(\Sigma_P)_{ii}$. For fixed weights w_i^P , r_i^0 is label-independent and can be interpreted as a stochastic pruning regularizer, pushing λ_i^0 closer to 0 (i.e. r_i^0 induces an L_1 penalty on λ_i^0). \square

Unlike optimizing with respect to the deterministic pruning weight, optimizing with respect to the stochastic pruning mask accounts for the effect of dropping the weight. While we do not optimize the weights in Stage 2 of the algorithm, we can apply a similar analysis as above and compare the gradient with respect to the weight w_i^P , when differentiating a stochastic predictor versus the mean predictor.

We find $\frac{\partial \hat{R}_{SP}(P)}{\partial w_i^P} = \lambda_i^0 \frac{\partial \hat{R}_S(\bar{W})}{\partial w_i} + \lambda_i^0 w_i^P (\Sigma_P)_{ii}$. In comparison to $\frac{\partial \hat{R}_S(\bar{W})}{\partial w_i}$ i.e., the gradient with respect to mean weights, we observe: (i) weights that are unlikely to be pruned (large λ_i) have a larger effective learning rate; and (ii) optimizing the Gibbs risk introduces a label-independent regularizer $\bar{r}_i^0 = \lambda_i^0 w_i^P (\Sigma_P)_{ii}$, which shrinks the weights to zero. The L_2 regularization effect of dropout in generalized linear models is well-known [31, 32].

The regularization terms \bar{r}_i^0 and r_i^0 appearing in the gradient of empirical Gibbs risk arise due to the term $\mathbb{E}[W^T \Sigma_P W]$ in Eq. (6). It is interesting to note the interaction between the gradient updates of the empirical Gibbs risk with respect to λ_i and w_i^P . As λ_i increases, $\left| \frac{\partial \hat{R}_{SP}(P)}{\partial w_i^P} \right|$ grows, thus increasing the magnitude of w_i^P , and vice versa.

Pruning probabilities are aligned with weight magnitude growth. Consider a diagonal feature covariance matrix. In Appendix C.1.1, we show that early stopping induces a pruning criterion that is aligned with the magnitude of the weights. In linear regression with diagonal feature covariance, the induced pruning criterion is based on variance-normalized feature alignment. We further extend the analysis when the features are correlated.

7.2 Optimizing the PAC-Bayes bound as constrained Feature Realignment

We now turn to Stage (3) of the algorithm outlined in the beginning of this section, introducing a data-dependent PAC-Bayes risk bound optimization. We refer the reader to Appendix C.2 for the full analysis. Here, we would like to highlight a few take-aways. First, PAC-Bayes risk bound optimization adds an additional regularizer, forcing λ_i to stay close to its prior value λ_i^0 . Second, the growth of λ_i depends on the agreement between the data S_P and $S_{\bar{P}}$: if the sign of the gradient on $S_{\bar{P}}$ disagrees with the sign of the weight (learned from S_P), then λ_i is driven to zero. Otherwise, λ_i increases. In Appendix C.2, we capture this agreement between the data S_P and $S_{\bar{P}}$ via a quantity which may be interpreted as a measure of data distribution complexity/stability. In Lemma 2 (Appendix C.3), we provide conditions under which the bound grows linearly or quadratically with this measure of complexity. See Appendix C.3 for the full details.

Note that our analysis of linear models under squared error extends to neural networks in the Neural Tangent Kernel regime (i.e., the infinite-width limit). See Appendix C.4 for more details.

8 Experiments

We now experimentally verify the characteristics and performance of our stochastic pruning framework. In interest of space, the full details on the networks, datasets, hyperparameters, optimizers, and loss functions are provided in Appendix D. Appendix E contains a number of additional experiments on PFT & PBP, including the trade-off between the tightness of the bound and the test error for PBP, the trade-off between overall mask entropy and performance in PBP, an analysis of the evolution

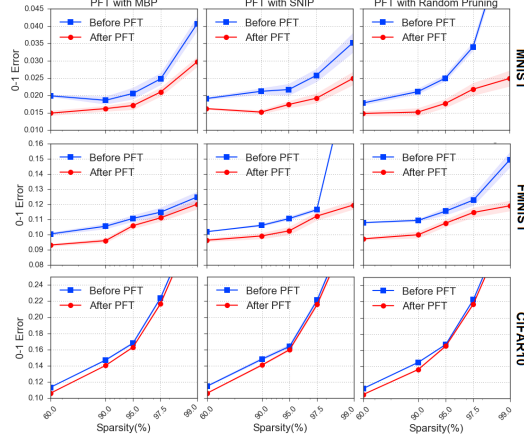


Figure 5: Impact of the application of PFT on standard One-Shot Pruning methods (MBP, SNIP) and Random Pruning. MNIST and Fashion-MNIST with a 3-layer width-1000 MLP and CIFAR-10 with ResNet20.

of the prior and posterior probabilities across the different stages of PBP, and an analysis of the deterministic masks created from One Shot Pruning methods.

8.1 Probabilistic fine-tuning boosts the accuracy of sparse predictors

One-Shot Pruning (OSP) methods create the mask in a single or no pass over the dataset. Popular choices include pruning at initialization methods [4–6, 28]. OSP later in training was discussed in [29]. These methods follow the same protocol: create a mask some time during training (e.g. at initialization or after training), apply the mask to the network, and train the sparse network. PFT aims to improve the performance of OSP methods using stochastic models to fine-tune the mask before training the sparse network. This entails training the prior probabilities as discussed in Section 7.1 in the context of linear models; we then create a new mask by keeping the weights with the highest (trained) probabilities and pruning the rest.

Fig. 5 shows that PFT consistently boosts the performance of standard OSP methods on small-scale vision tasks; in general, the benefit of applying PFT is most noticeable for high sparsity levels (e.g. 99%) where some OSP methods alone fail to achieve good test error. On CIFAR-10 with ResNet20, PFT improves the accuracy of OSPs by $\sim 1\%$ on average; results on CIFAR-10 with sparsity level 99% are included in Appendix E. Note that PFT with Random Pruning (RP) is comparable to other methods, hence suggesting that PFT alone (preceded by RP) can serve as a standalone pruning method. More details and experiments on PFT are provided in Appendix E.1.

Robustness of pruned vs unpruned nets and the tightness of a PAC-Bayes bound. We start with a short discussion on the effect of network pruning on generalization as implied by a PAC-Bayes bound, by focusing on the KL penalty term. For this paragraph only, consider diagonal Gaussian priors and posteriors. Due to the weights being independent, the KL divergence is the sum over the (non-negative) KL terms of the weights. Hence with all other Gaussian parameters (the mean and variance of each weight) fixed, it is trivial to see that pruning alone cannot increase the KL divergence, and thus one may obtain a smaller (or at least no larger) KL penalty for a pruned network compared to that of a full network. The KL divergence, and thus the bound, for pruned networks can be further reduced in the case when the per-weight variance of the posterior Gaussian can be increased without affecting the training error. We thus empirically evaluate whether standard pruned models are more robust to weight perturbations. Fig. 6 (left) shows the percentage drop in performance of a deterministic trained model (pruned and unpruned) when the weights are perturbed with a Gaussian noise with variance v^2 , i.e., $W_i \rightarrow W_i + v\zeta_i$ for all i where (ζ_i) are i.i.d. Gaussian variables with distribution $\mathcal{N}(0, 1)$. In the variance range where the drop in performance is most noticeable, the pruned network tends to be more robust to perturbations compared to the full network. This can also be interpreted as follows: the variance level v that corresponds to a given drop in performance for the pruned model is higher than that which corresponds to the full network with the same drop in performance. In addition, PFT seems to find more robust masks than the original mask (In this

example, the one we obtain with MBP). We believe that the randomness in stage 2 of PFT acts as an implicit regularization on the mask, which increases robustness. This observed increased robustness to weight perturbations with pruning would also result in a smaller KL divergence for the pruned network, further reducing the PAC-Bayes bound. We use this as motivation for our self-bounded learning algorithm, PBP.

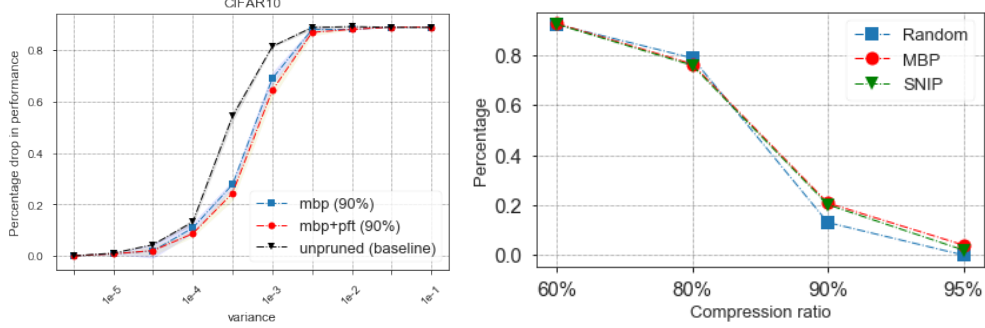


Figure 6: *Left*: Drop in performance (as a fraction) of probabilistic model as a function of variance (unpruned versus 90% ResNet20.) *Right*: Percentage of shared weights between OSP masks and masks obtained after applying PFT (CIFAR10 with ResNet20). This shows that OSP methods tend to miss the local optimal masks more and more as we increase the sparsity.

OSP mask VS PFT mask. Knowing that masks created by OSP methods are not locally optimal, we intuitively expect that as we increase the sparsity, OSP methods will ‘miss’ the local optimal mask by a growing ‘distance’. We measure this distance by the percentage of weights shared between OSP masks and the masks created after applying PFT. Fig. 6 (right) shows this percentage as a function of sparsity (compression ratio) for different OSP methods (Random, MBP, SNIP) on CIFAR10 with ResNet20 architecture. For a medium range sparsity (60%), around 93% of the weights are shared between masks before and after applying PFT, suggesting that masks found by PFT are close to the original OSP masks. However, this percentage of shared weights between masks before and after PFT follows a decreasing pattern w.r.t the sparsity. For high sparsities, e.g. 95% (average over 5 runs), we observe that for MBP and SNIP, only a fraction $\sim 3\%$ of the weights is shared, and for random pruning, this fraction is 0% in all the runs. This confirms our intuition that the higher the sparsity, the worse the OSP mask, and suggests that PFT searches ‘farther’ for better masks.

8.2 Self-bounded learning via PAC-Bayes

PBP on small vision tasks Finally, we evaluate our proposed stochastic pruning with PAC Bayes framework (PBP), which returns a stochastic predictor in contrast to many other pruning approaches. We consider MNIST, Fashion-MNIST & CIFAR-10¹¹ with corresponding architectures to Section 8.1. In Fig. 7, we compare the PAC-Bayes generalization bounds on 0-1 error (blue) of our stochastic predictor Q to the corresponding observed test set error (orange) and the observed test error of our sparse stochastic prior P (green) over a range of sparsities and pre-training dataset fraction α . For reference, we also plot the test error of a baseline pruning method, SNIP [4], to compare our test errors and bounds to a standard *deterministic* OSP method which doesn’t provide generalization guarantees. In general, we see that a higher α value not only improves the bound, but also decreases error of randomized predictors P and Q . More data-dependence helps across different sparsity levels and datasets. We also observe that the bounds on the error of Q , and the error of Q as measured on the test set are within a few percentage points of our baseline for a large range of sparsity levels. The gap between the baseline and the performance of Q increases with sparsity for both values of α . Interestingly, when trained on CIFAR-10, we also see that our posterior Q significantly outperforms the prior P , obtaining 10% lower error when $\alpha = 0.4$, and 97.5% of the weights are pruned. In this case, since Q test error bound is already lower than the prior test error, any held out dataset bound on the performance of P would be worse than our PAC-Bayes bound on Q . This demonstrates the

¹¹CIFAR-10: <https://www.cs.toronto.edu/~kriz/cifar.html>, MNIST: <http://yann.lecun.com/exdb/mnist/> FMNIST: <https://github.com/zalandoresearch/fashion-mnist>

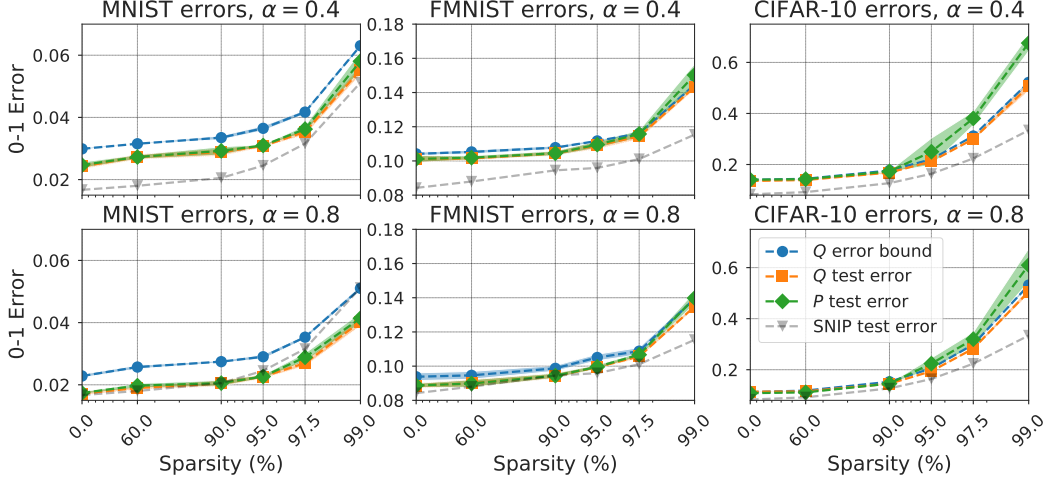


Figure 7: Error and error bound versus sparsity on image classification tasks for different $\alpha = |S_P|/|S|$. Displayed intervals correspond to 95% confidence over 5 independent seeds.

strength of our proposed PBP algorithm in that it uses all of the available data for training in order to provide generalization bounds. This is crucial in regimes of low or complex data where we expect to need all of the available data for training a good classifier.

9 Related Work

Recent work by Elesedy, Kanade, and Teh [33] has shown that, in the context of linear models, magnitude pruning zeros out the weights based on the magnitude of feature alignment under certain assumptions on the feature covariance matrix. Our analysis is complementary to these findings.

Other pruning methods. Our work uses a class of pruning approaches that identify the set of weights to prune based on a criterion applied to each weight by the end of a single training run. We refer to such methods as one-shot pruning (OSP). These include magnitude pruning [2] and SNIP [4], as well as many others (e.g., see GraSP [34] and SynFlow [28] for pruning at initialization, and L_0 regularization-based [35–37] and dropout-based [38, 39] for pruning during training). SNIP was proposed as a technique to prune at initialization, with the remaining weights trained as usual. In contrast, magnitude pruning was proposed as a technique to apply after training, following by fine-tuning. However, state-of-the-art applications of magnitude pruning now employ iterative magnitude pruning (IMP) [3, 40], where a fraction of the smallest weights are pruned on each round, followed by rewinding training back to an earlier iteration [40] or simply rewinding the learning rates [41]. Frankle et al. [29] and Gale, Elsen, and Hooker [36] show that none of the OSP techniques dominate one another, and IMP outperforms them all [41].

The pruning method that is most related to PFT is an extension of variational dropout due to Molchanov, Ashukha, and Vetrov [38]. In variational dropout, the weights are multiplied by Gaussian variables instead of Bernoulli variables. Although the variances of the Gaussian mask can be forced to be small or large, the corresponding model is never sparse (with probability one), and another thresholding step is required to make the model sparse (e.g. remove the weight with the highest Gaussian mask variance). Another major difference is the choice of the prior; Gaussian dropout consider the improper log-scale uniform prior since it is the only prior that makes variational inference consistent with Gaussian dropout. As a result, this prior can only yield vacuous PAC-Bayes bounds.

10 Discussion and Limitations

In this work, we introduced a new algorithmic framework for stochastic pruning. The generality of this framework allowed us to derive two algorithms, PFT and PBP, that act on the same input network and return different outputs. We elucidate the non-optimality of OSP pruning methods by showing that probabilistic fine-tuning (PFT) finds better local masks, starting from popular OSP methods. To depict this gain in performance, we analyze PFT in the context of linear models and show that PFT learns pruning probabilities based on feature alignment. Further, we show a regularization effect on pruning probabilities arising from minimizing risk of a stochastic predictor, rather than an

average one. This is complementary to [31], where the authors demonstrate a similar regularization effect on the weights when training a stochastic predictor. On small vision deep neural networks, we notice that PFT consistently boosts the accuracy of existing One-Shot Pruning methods (Section 8), however, the computational cost of PFT is no less than ordinary training and so scales with network size. Optimizing PFT for large scale architectures is an interesting topic for future work.

Motivated by the apparent robustness of pruned networks to perturbations of the weights, we introduced PAC-Bayes Pruning (PBP), which is a self-bounded learning algorithm with data-dependent priors, aimed at training a sparse *stochastic* network. In the linear model context, we demonstrate that the tightness of a data-dependent PAC-Bayes bound obtained via PBP depends on how the feature alignment changes as we add more data to the learning problem. In our experiments, we obtained tight numerical bounds for sparse DNNs on a variety of datasets as demonstrated in Section 8. However, we observe that stochastic predictors seem to consistently under-perform their deterministic counterparts. The role of product prior distributions on the weights and pruning probabilities is worth investigating. The consistent drop in performance seen when pruning weights highlights the importance of considering the interaction between weights when pruning.

A Additional background material

A.1 Initializing pruning probabilities

Below we discuss some choices of initializing pruning probabilities λ^0 . The main difference between these choices is whether λ^0 's are chosen based on some criterion g , which we refer to as *probabilistic pruning based on the criterion g* , or randomly. Let s be the sparsity that we want to achieve in the probabilistic model, i.e., the average of the entries of $1 - \lambda^0$. We consider two versions of setting λ^0 :

- *Isotropic P (random pruning)*: the weights have equal probability to be pruned. To achieve an average sparsity s , we set $\lambda_i^0 = 1 - s$.
- *Block isotropic P* : we divide the weights into two blocks based on their g_i value (evaluated on a deterministic network W_0). Letting $k = (1 - s)D$, $g^{(k)}$ be the k^{th} order statistic of the sequence $(g_i)_{1 \leq i \leq p}$, $\epsilon \in (0, 1)$, and $s\epsilon/(1 - s) < 1$, we choose

$$\lambda_i^0 = \begin{cases} 1 - s\epsilon/(1 - s) & \text{if } g_i \geq g^{(k)} \\ \epsilon & \text{otherwise.} \end{cases} \quad (9)$$

When $\epsilon \ll 1$, this prior allows small fluctuations around the deterministic pruned network.

Both of these choices induce P with an expected sparsity s .

A.2 A standard data-independent PAC-Bayes bound

Here we introduce a classical PAC-Bayes bound and connect it to the one used in Theorem 2.

Writing dQ/dP for the Radon–Nikodym derivative of Q with respect to P , recall that the Kullback–Leibler divergence (or entropy) of Q (relative) to P , is defined to be $\text{KL}(Q \| P) = \int \log(dQ/dP)dQ$ when Q is absolutely continuous with respect to P and infinite otherwise. For $p, q \in (0, 1)$, let $\text{kl}(q \| p)$ denote the KL divergence of the Bernoulli distribution $\mathcal{B}(q)$ to $\mathcal{B}(p)$. Thus, $\text{kl}(q \| p) = q \log(q/p) + \bar{q} \log(\bar{q}/\bar{p})$, where $\bar{p} = 1 - p$ and $\bar{q} = 1 - q$.

The following PAC-Bayes theorem, of Langford and Seeger [42], is stated in terms of the inverse Bernoulli KL divergence,

$$\text{kl}^{-1}(a \| \epsilon) = \sup\{p \in [0, 1] : \text{kl}(a \| p) \leq \epsilon\}.$$

Theorem 3. *For every $[0, 1]$ -valued loss function ℓ , data distribution \mathcal{D} , natural $N \in \mathbb{N}$, distribution P on \mathcal{F} , and probability $\delta \in (0, 1)$, with probability at least $1 - \delta$ over training data $S \sim \mathcal{D}^N$,*

$$\forall Q, R_{\mathcal{D}}(Q) \leq \text{kl}^{-1}\left(\hat{R}_S(Q) \parallel \frac{\text{KL}(Q \| P) + \log(\frac{2\sqrt{N}}{\delta})}{N}\right). \quad (10)$$

Substituting our concrete KL term into Theorem 3, we obtain a quantity that we can explicitly optimize in terms of the parameters that define the posterior. Note that gradients of the empirical Gibbs risk term are in general intractable. However, as is standard, tractable stochastic gradients can be obtained. We describe our approach in Appendix B.3.

Note that the inverse KL divergence in Theorem 3 does not admit an analytical form. Different relaxations exist in the literature. In our experimental setting, we optimize the bound stated in Theorem 2.

A.3 Surrogate loss function

Though our bounds are evaluated on the 0-1 error, this is not a differentiable loss function. We follow [1] in using gradient-based optimization with the average cross entropy loss as a surrogate empirical risk $\hat{R}_S(Q)$ during posterior training. We clamp the probabilities smaller than 10^{-4} to ensure that the optimization objective remains bounded.

We would like to reiterate that the following produces a valid procedure since Theorem 2 is valid for *any* posterior Q , trained in any way: we can compute Q via optimization of Eq. (3) with a cross entropy loss, and evaluate the bound with the performance of Q on 0–1 loss.

Also note that the evaluation of a PAC-Bayes bound involves the empirical risk $\hat{R}_{S \setminus S_P}(Q)$. In general, the computation of this quantity is intractable. However, a confidence interval can be obtained using a Chernoff bound.

B Algorithmic details

B.1 Posterior probability mapping.

Optimizing the posterior probabilities λ might cause numerical issues whenever the gradient update pushes the probability out of the interval $(0, 1)$. To avoid such scenarios, we use a parameterization of the posterior probabilities that ensures probabilities remain in $(0, 1)$. We consider two options

- The Sigmoid function: $\lambda_i = \text{Sig}(\lambda'_i)$
- The Clamp function: $\lambda_i = \begin{cases} \lambda_i & \lambda_i \in (0, 1) \\ 1 & \lambda_i \geq 1 \\ 0 & \lambda_i \leq 0 \end{cases}$

The sigmoid function tends to outperform the clamp function when the prior is chosen to be a small perturbation of the deterministic sparse network, e.g. block isotropic prior with small ε . However, the clamp function performs better than sigmoid when the prior is sufficiently random (i.e., a significant fraction of the weights have pruning probabilities not close to 0 or 1). The results reported in Fig. 7 use Sigmoid mapping while those reported in Fig. 5 use the clamp function. These choices were found to be optimal by hyperparameter search.

B.2 Optimization Algorithm

The optimization of the empirical risk requires computing the gradients w.r.t the posterior mean weights W^f and the posterior pruning probabilities λ . The direct computation of the latter is problematic from a practical perspective. To see this, recall that the empirical risk is given by

$$\hat{R}_S(Q) = \mathbb{E}_{b,W}[\tilde{R}_S(b \circ W)],$$

where W and b are the weights and the masks, and $\tilde{R}_S(b \circ W) = \frac{1}{N} \sum_{i=1}^N \check{\ell}(f_{b \circ W}(x), y)$ is the empirical risk for a surrogate differentiable loss $\check{\ell}$.

For some i , let $e_i(b'_i) = \mathbb{E}_Q[\tilde{R}_S(b \circ W) | b_i = b'_i]$. Then, we have that

$$\tilde{R}_S(Q) = \lambda_i e_i(1) + (1 - \lambda_i) e_i(0)$$

this yields

$$\partial_{\lambda_i} \tilde{R}_S(Q) = e_i(1) - e_i(0)$$

Computing the gradient with this formula is impractical as it requires an extensive number of forward passes ($2 \times p \times k$ where k is the number of network samples used to estimate $e_i(1)$ and $e_i(0)$). The conditional expectation makes it impractical to obtain a direct Monte Carlo estimate of $\partial_{\lambda_i} \tilde{R}_S(Q)$; indeed, we have to fix $b_i = 1$ or $b_i = 0$ and simulate the rest of the mask, in order to estimate $\partial_{\lambda_i} \tilde{R}_S(Q)$. Repeating this procedure for all i 's would require extensive computational resources. However, we can also simulate k mask samples and use these samples to estimate $\partial_{\lambda_i} \tilde{R}_S(Q)$, i.e., find the sample masks where $b_i = 1$ and use them to estimate $e_i(1)$, and vice-versa for $e_i(0)$. Algorithm 1 illustrates this procedure. Unfortunately, this algorithm would still require tremendous number of samples to update extreme probabilities, which limits its benefits.

Fortunately, a relaxation of this optimization problem known as the Gumbel–Softmax trick makes the gradient easier to approximate using standard Monte Carlo estimate. We describe this relaxation in details in the next section.

B.3 Reparameterization trick and Gumbel-Softmax

The empirical risk $\tilde{R}_S(Q)$ depends on λ in an implicit way, thus, it is not straightforward to compute the gradient $\partial_{\lambda} \tilde{R}_S(Q)$ using Monte Carlo samples. A classic approach is the *reparameterization trick* which consists of finding a mapping that makes the dependence on λ explicit. In the case of a discrete distribution (Bernoulli in our case), a common reparameterization trick is Gumbel-Max (GM) trick which can be traced back to Luce [43]. It was introduced as a method for sampling from discrete random variables using explicit dependence on the probabilities of each state. More precisely,

Algorithm 1: Optimization of PAC-Bayes bound

Result: Approximate optimization of PAC-Bayes bound E

prior P , learning rate η , sample size m ;

while not converged **do**

for $k=1$ to m **do**

 simulate $w_i^k \sim \mathcal{N}(W_i^f, s_i^2)$ for all i ;
 simulate Bernoulli variables $B_k \sim \mathcal{B}(\lambda_i)$ for all i ;

end

$$a_i^k = \frac{B_i^k}{\sum_k B_i^k};$$

$$b_i^k = \frac{1-B_i^k}{\sum_k (1-B_i^k)};$$

$$\zeta_i = \frac{1}{m} \sum_k [\mathcal{L}((w_j^k B_j^k)_{j \neq i}, w_i^k) - \mathcal{L}((w_j^k B_j^k)_{j \neq i}, 0)] \approx$$

$$\sum_{k=1}^m (a_i^k + b_i^k)(2B_i^k - 1) \mathcal{L}((w_j^k B_j^k)_{1 \leq j \leq p});$$

$$\beta_i = \frac{1}{m} \sum_k \partial_{w_i^k} \mathcal{L}((w_j^k B_j^k)_{1 \leq j \leq p});$$

$$\lambda_i \leftarrow \lambda_i - \eta(\zeta_i + \partial_{\lambda_i} \Gamma(Q)) \times 1_{0 < \sum_k B_i^k < m};$$

$$W_i^f \leftarrow W_i^f - \eta(\beta_i + \partial_{W_i^f} \Gamma(Q))$$

end

consider a discrete representation $d \in \{0, 1\}^n$ such that $\sum_{i=1}^n d_i = 1$ and let $(\lambda_1^0, \dots, \lambda_n^0) \in (\mathbb{R}_+^*)^n$ denote an unnormalized probability vector. The GM trick can then be described by the following steps:

1. sample $U_k \sim \mathcal{U}(0, 1)$ i.i.d. for $k \in [n]$;
2. find $k = \arg \max_k \{\log(\lambda_k^0) - \log(-\log(U_k))\}$, set $D_k = 1$ and $D_i = 0$ for $i \neq k$.

Then, we have

$$\mathbb{P}(D_k = 1) = \frac{\lambda_k^0}{\sum_{i=1}^n \lambda_i^0}.$$

The name Gumbel-Max trick was inspired by the fact that the random variable $-\log(-\log(U))$ has a Gumbel distribution.

Although the GM trick allows straightforward simulation of discrete variables, it is not practical for gradient computing since the the differentiation of the max function might cause numerical issues. For this purpose, [44] introduced the Gumbel-Softmax (GS) trick which relaxes the discrete distribution to the *concrete* distribution. It proceeds as follows:

- fix $\beta \in \mathbb{R}_+^*$ and sample $G_k \sim \text{Gumbel}$ i.i.d.;
- set $X_k = \frac{\exp((\log(\lambda_k^0) + G_k)/\beta)}{\sum_{i=1}^n \exp((\log(\lambda_i^0) + G_i)/\beta)}.$

Then, X has a concrete distribution which we denote $\text{Concrete}(\lambda_1^0, \dots, \lambda_n^0)$. Further, it follows that

$$\lim_{\beta \rightarrow 0} \mathbb{P}(X_k > X_i \text{ for all } i \neq k) = \frac{\lambda_k^0}{\sum_{i=1}^n \lambda_i^0}.$$

The relaxed problem. In our case, given a sequence of independent random variables $X_i \sim \text{Concrete}(\lambda_i, 1 - \lambda_i)$, we optimize the relaxed empirical risk, given by $\hat{R}_S((X_i, W_i)_{1 \leq i \leq p})$. We take $\beta = 0.5$ throughout this work.

B.4 Kullback–Leibler divergence for spike and slab distributions

With the spike and slab prior P and posterior Q , the KL divergence is tractable and is given by

$$\begin{aligned} \text{KL}(Q \parallel P) &= \sum_{i=1}^D \text{KL}(q_i \parallel p_i) \\ &= \sum_{i=1}^D \text{kl}(\lambda_i \parallel \lambda_i^0) + \lambda_i \underbrace{\text{KL}(\mathcal{N}(W_i^f, s_i^2) \parallel \mathcal{N}(W_i^0, \sigma_i^2))}_{\gamma_i}, \end{aligned}$$

where for all $1 \leq i \leq D$,

$$\begin{aligned} \text{kl}(\lambda_i \parallel \lambda_i^0) &= (1 - \lambda_i) \log \left(\frac{1 - \lambda_i}{1 - \lambda_i^0} \right) + \lambda_i \log \left(\frac{\lambda_i}{\lambda_i^0} \right), \\ \gamma_i &= \frac{1}{2} \left(\frac{(W_i^f - W_i^0)^2}{\sigma_i^2} + \frac{s_i^2}{\sigma_i^2} - \log \left(\frac{s_i^2}{\sigma_i^2} \right) - 1 \right). \end{aligned}$$

C Linear Model Analysis

C.1 Learning prior pruning probabilities

Here we provide the proof of Lemma 1.

Proof. Consider a related empirical risk of an average predictor, in this case, a linear model with weights \bar{W} . Then Eq. (8) can also be expressed in terms of the gradient of $\hat{R}_{S_P}(\bar{W})$ with respect to \bar{W}_i . In particular,

$$\frac{\partial \hat{R}_{S_P}(P)}{\partial \lambda_i^0} = w_i^P \frac{\partial \hat{R}_{S_P}(\bar{W})}{\partial \bar{W}_i} + r_i^0, \quad (11)$$

where $r_i^0 = \frac{1}{2}(w_i^P)^2(\Sigma_P)_{ii}$. Now consider minimizing the empirical risk of the average classifier, $\hat{R}_S(\bar{W})$, with respect to λ_i^0 . By chain rule, the gradient equals to $\frac{\partial \hat{R}_S(\bar{W})}{\partial \lambda_i^0} = w_i^P \frac{\partial \hat{R}_S(\bar{W})}{\partial \bar{W}_i}$. It follows that $\frac{\partial \hat{R}_{S_P}(P)}{\partial \lambda_i^0} - \frac{\partial \hat{R}_S(\bar{W})}{\partial \lambda_i^0} = r_i^0$. \square

C.1.1 Pruning probabilities are aligned with weight magnitude growth

Here we show that early stopping induces a pruning criterion that is aligned with the magnitude of the weights. In linear regression with diagonal feature covariance, the induced pruning criterion is based on variance-normalized feature alignment.

From Eq. (11), we can conclude the following: the first gradient term increases λ_i when the gradient of the empirical risk of \bar{W} pushes the weight magnitude up, and decreases otherwise. The regularization term drives λ_i to zero.

Note that the weights minimizing the empirical risk $\hat{R}_S(w)$ under square loss are

$$w_i = \frac{1}{|S|} \Sigma^{-1}(X) \psi_i(X)^T Y, \quad (12)$$

In linear regression, the weight gradient depends on the feature alignment – large feature alignment magnitude yields large magnitude weights (up to rescaling based on Σ_P), as seen from Eq. (12).

When Σ_P is diagonal, at convergence $w_i = (\Sigma_P)_{ii}^{-1} A_i(P)$ and $\frac{\partial \hat{R}_{S_P}(P)}{\partial \lambda_i^0} = \frac{1}{2}(w_i^P)^2$. If we start by initializing the pruning probabilities all to the same value, then λ_i^0 's with larger (always non-negative) gradient increase more rapidly, and thus can be interpreted as being more important. With early stopping, this yields a pruning criterion based on the magnitude of feature alignment, normalized by $(\Sigma_P)_{ii}^2$. In other words, the weights that have the highest feature alignment (after adjusting for the variance), are the ones that are not pruned. Magnitude-based pruning and feature alignment connections were recently described by Élesedy, Kanade, and Teh [33].

Non-diagonal covariance. We extend the analysis when the features are correlated. Assume $(\Sigma_P)_{ij} = (\Sigma_P)_{ji} \neq 0$ for some $i \neq j$ and assume $(\Sigma_P)_{i'j'} = 0$ for all other $i' \neq i, j', j' \neq j$. We now compute the ratio of λ_i^0/λ_j^0 , indicating the relative importance of w_i and w_j .

Setting the gradient $\frac{\partial \hat{R}_{S_P}(P)}{\partial \lambda_i^0}$ to 0, if $w_i \neq 0$, we have

$$\lambda_j^0 = \frac{2A_i(P) - w_i(\Sigma_P)_{ii}}{2w_j(\Sigma_P)_{ij}}. \quad (13)$$

Using Eq. (12), we can also write w_i for this special Σ_P case as

$$w_i = (\Sigma_P)_{ii}A_i(P) - (\Sigma_P)_{ij}A_j(P). \quad (14)$$

Then for $(\Sigma_P)_{ii} = (\Sigma_P)_{jj} = 1$, by substituting Eq. (14) to Eq. (13), we can express the ratio as

$$\frac{\lambda_i^0}{\lambda_j^0} = \frac{(c_{ji} + (\Sigma_P)_{ij})(1/c_{ji} - (\Sigma_P)_{ij})}{(c_{ji} - (\Sigma_P)_{ij})(1/c_{ji} + (\Sigma_P)_{ij})}, \quad (15)$$

where $c_{ji} = \frac{A_j(P)}{A_i(P)}$.¹² Studying Eq. (15), we conclude the following:

- if $\text{sign}(A_i(P)) \cdot \text{sign}(A_j(P)) = \text{sign}(\Sigma_P)_{ij}$, then the weight with a smaller magnitude of feature alignment is more likely to be pruned;
- otherwise, the weight with larger feature alignment is more likely to be pruned.

Intuitively, if the correlation $(\Sigma_P)_{ij}$ is positive, but feature alignment signs $\text{sign}(A_i(P)) \neq \text{sign}(A_j(P))$, then the weight w_k with a higher magnitude of feature alignment $|A_k(P)|$ is likely to be misleading and we should assign a smaller probability λ_k^0 (for $k \in \{i, j\}$).

C.2 Optimizing the PAC-Bayes bound as constrained Feature Realignment

In this section, we analyze Stage (3) of the algorithm outlined in the beginning of Section 7.

We consider a linear PAC-Bayes bound due to [16], which, due to the independence of the pruning probabilities a priori and a posteriori, corresponds to optimizing

$$B(\lambda) = \hat{R}_S(Q) + \frac{\zeta}{n} \sum_{i=1}^D \text{KL}(q_i \parallel p_i), \quad (16)$$

where $\text{KL}(q_i \parallel p_i) = \text{kl}(\lambda_i \parallel \lambda_i^0)$ (see Appendix B.4).

Using Eq. (11), we express

$$\frac{\partial B}{\partial \lambda_i} = w_i^P \frac{\partial \hat{R}_{S_P}(\bar{W})}{\partial \bar{W}_i} + r_i^0 + r_i^{\lambda^0}, \quad (17)$$

where $r_i^{\lambda^0} = \frac{\partial \text{kl}(\lambda_i \parallel \lambda_i^0)}{\partial \lambda_i}$ is the KL penalty. The penalty term acts as a regularizer, forcing λ_i to stay close to its prior value λ_i^0 . Fig. 8 illustrates this effect: when $\lambda_i > \lambda_i^0$, the penalty term is positive, driving the value of λ_i to be lower, and vice versa. This is in addition to the regularization coming from r_i^0 , which drives λ_i to zero.

In contrast to the analysis in Section 7.1, here the gradient of the empirical risk of \bar{W} is computed on the new data $S_{\bar{P}}$. If the sign of the gradient on $S_{\bar{P}}$ disagrees with the sign of the weight (that was optimized on S_P), λ_i will be driven to zero. Otherwise, λ_i will increase.

When $\Sigma(X_{\bar{P}}) = \mathbb{I}$, this translates to the sign of the feature alignment $A_i(\bar{P})$ matching the sign of w_i^P . Further, if w_i^P is the minimizer of the empirical risk of a deterministic predictor, then $w_i = A_i(P)$. In this case, λ_i is driven to zero if $\text{sign}(A_i(P)) \neq \text{sign}(A_i(\bar{P}))$, and to 1, otherwise.

¹²Note, that in Eq. (15), to obtain valid values, we need to take into account that the magnitudes of $(\Sigma_P)_{ij}$ and c_{ji} are dependent.

C.2.1 Closed-form solution for uncorrelated features

Assume $\Sigma(X_{\bar{P}}) = \Sigma(X_P) = \mathbb{I}$ and let $\kappa = \frac{\zeta}{N-M}$. The gradient of λ_i 's when minimizing Eq. (16) is given by

$$\frac{d\lambda_{i,t}}{dt} = -\mu \left(\Delta_i + \kappa \frac{\partial \text{KL}(q_i \parallel p_i)}{\partial \lambda_i} \right), \quad (18)$$

where $\mu > 0$ is the learning rate

The KL divergence above is reduced to that of a Bernoulli variable, and for each $i \in [D]$ is given by

$$\text{kl}(\lambda_i \parallel \lambda_i^0) = \lambda_i \log \left(\frac{\lambda_i}{\lambda_i^0} \right) + (1 - \lambda_i) \log \left(\frac{1 - \lambda_i}{1 - \lambda_i^0} \right),$$

with derivative

$$\frac{\partial \text{kl}(\lambda_i \parallel \lambda_i^0)}{\partial \lambda_i} = \log \left(\frac{\lambda_i}{\lambda_i^0} \right) - \log \left(\frac{1 - \lambda_i}{1 - \lambda_i^0} \right).$$

The penalty in the objective Eq. (16) coming from this KL term and its derivative are plotted in Fig. 8 as a function of λ_i 's, posterior pruning probabilities, for a few different values of λ_i^0 .

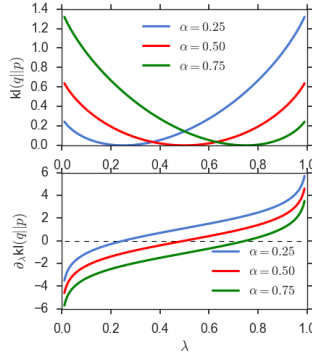


Figure 8: KL penalty for Bernoulli random variables.

The gradient flow of λ_i minimizing the bound converges to the unique solution given by

$$\lambda_{i,\infty} = \frac{\lambda_i^0}{\lambda_i^0 + (1 - \lambda_i^0) \exp(\kappa^{-1} \Delta_i)}, \quad (19)$$

where $\Delta_i = r_i^0 - w_i^P A_i(\bar{P})$, i.e., the first two terms in Eq. (17).

For $\Delta_i = 0$, we recover the value $\lambda_{i,\infty} = \lambda_i^0$.

Approximations. For particular values of the drift, $\lambda_{i,\infty}$ can be approximated with Taylor expansion:

- $|\Delta_i| \ll 1$: $\lambda_{i,\infty} \approx \lambda_i^0 (1 + \kappa^{-1} \Delta_i (1 - \lambda_i^0))$;
- $\Delta_i \ll -1$: $\lambda_{i,\infty} \approx (1 - (\lambda_i^0)^{-1} (1 - \lambda_i^0) \exp(\kappa^{-1} \Delta_i))$;
- $\Delta_i \gg 1$: $\lambda_{i,\infty} \approx \lambda_i^0 (1 - \lambda_i^0)^{-1} \exp(-\kappa^{-1} \Delta_i)$.

In words, when the drift magnitude is large, the bias coming from the value of λ_i^0 learned on the data X_P diminishes as $N - M$ or $|\Delta_i|$ grow: $\lambda_{i,\infty}$ tends to 1 (large negative drift) or 0 (large positive drift).

C.2.2 Concentration of feature alignment as a measure of data distribution complexity

The tightness of the PAC-Bayes generalization guarantee is determined by the size of $\kappa \sum_i \text{kl}(\lambda_i \parallel \lambda_i^0)$. Here we state the dependence of the KL term on the drift term.

Lemma 2. *Under certain approximations, the tightness of the data-dependent PAC-Bayes bound on the generalization error scales with the drift term:*

- with $|\Delta_i| \ll 1$, the bound is quadratic in Δ_i ;
- with $|\Delta_i| \gg 1$, the bound is linear in Δ_i .

For the details, see Appendix C.3.

C.3 The dependence of KL on the drift term in a linear model

The tightness of the PAC-Bayes generalization guarantee is determined by the size of $\kappa \sum_i \text{kl}(\lambda_i \parallel \lambda_i^0)$. Substituting the optimal λ_i from Eq. (19) yields

$$\begin{aligned} \text{kl}(\lambda_i \parallel \lambda_i^0) &= -\log(\exp(\eta_i)(1 - \lambda_i^0) + \lambda_i^0) \\ &\quad + \eta_i \frac{\exp(\eta_i)}{\exp(\eta_i) + \frac{\lambda_i^0}{1 - \lambda_i^0}}, \end{aligned} \quad (20)$$

where $\eta_i = \kappa^{-1} \Delta_i$. We bound Eq. (20) separately for small and large drift Δ_i .

When Δ_i , and therefore η_i , is large positive, then

$$\text{kl}(\lambda_i \parallel \lambda_i^0) \approx \Delta_i \lambda_i^0 \kappa^{-1}.$$

When Δ_i is large negative, then

$$\text{kl}(\lambda_i \parallel \lambda_i^0) \approx \Delta_i \kappa^{-1} \frac{\lambda_i^0}{1 - \lambda_i^0} - \log \lambda_i^0.$$

In both cases, large magnitude of the drift terms results in a linear dependence of the bound on Δ_i .

When the magnitude of the drift term is small, up to the error introduced by a first order Taylor approximation, $\lambda_{i,\infty} = \lambda_i^0(1 + \kappa^{-1} \Delta_i(1 - \lambda_i^0))$. This yields

$$\text{kl}(\lambda_i \parallel \lambda_i^0) = \lambda_i^0 a_i \log a_i - (1 - \lambda_i^0) b_i \log b_i, \quad (21)$$

where $a_i = 1 - \eta_i \lambda_i^0 + \eta_i$ and $b_i = 1 - \eta_i \lambda_i^0$. Assuming $\eta_i < 1$, we can bound $\log a_i \leq \eta_i(1 - \lambda_i^0)$ and, similarly, $\log b_i \leq \eta_i \lambda_i^0$. Substituting these log bounds to Eq. (21), we obtain

$$\text{kl}(\lambda_i \parallel \lambda_i^0) \leq \frac{\Delta_i^2}{\kappa^2} \lambda_i^0 (1 - \lambda_i^0). \quad (22)$$

In summary, the tightness of the data-dependent PAC-Bayes bound on the generalization error scales with the drift term:

- with $|\Delta_i| \ll 1$, the bound is quadratic in Δ_i ;
- with $|\Delta_i| \gg 1$, the bound is linear in Δ_i .

C.4 Overparameterized neural networks as linear models

Recent work by Jacot, Gabriel, and Hongler [11] has shown that training a depth L neural networks with gradient flow is equivalent to a functional gradient flow with respect to the Neural Tangent Kernel (NTK) defined by $\Theta_t^L(x, x') = \nabla_\theta f(x, \theta_t) \nabla_\theta f(x', \theta_t)^T$, where f is the network output, θ are the parameters, and t is the training time. In the infinite width limit, the kernel Θ^L becomes independent of t and training the network with quadratic loss is equivalent to the following linear model

$$f_t(x) = f_0(x) + \gamma(x, X_P)(I - e^{-\frac{1}{N} \hat{\Theta}^L t})(Y_P - f_0(X_P)), \quad (23)$$

where $\hat{\Theta}^L = \Theta^L(X_P, X_P)$, f_0 is the network output at initialization, and $\gamma(x, X_P) = \Theta^L(x, X_P)(\hat{\Theta}^L)^{-1}$. Consider a zero init, i.e., $f_0 = 0$, then, in the limit of infinite width, training the network with gradient flow on S_P is equivalent to training the linear model with the feature map given by $\Psi(x) = \{\Theta^L(x, x')\}_{x' \in S_P}$. Thus, by considering this alternative linear model, all the results on the linear model in this section apply. Note that the most practically relevant setting for pruning is the wide-but-finite DNN setting, where it has been shown that the linear DNN approximation holds around parameter initialization under a certain parameterization or learning rate regime [45].

D Experimental Details

D.1 Networks and datasets

Our image classification results consider three tasks: MNIST[13], Fashion-MNIST [14] and CIFAR-10 [15]. For MNIST we use an MLP with 3 hidden layers of width 1000, like in [46]. For Fashion-MNIST we consider the LeNet-5 architecture [47], and for CIFAR-10 we take the ResNet20 architecture [48]. For all datasets we standardize all input images based on the training split, and additionally use randomized crops and horizontal flips for CIFAR 10.

D.2 Hyperparameters for PAC-Bayes pruning (PBP)

Throughout all stages of training, we optimize using SGD with momentum 0.9, batch size 128. Below we provide the details for the individual stages from Section 8.2.

Stage 1. Pre-training the prior. We used fixed learning rates of 0.01 in MNIST & Fashion-MNIST and 0.1 in CIFAR-10. We trained for 30 epochs on MNIST, 70 epochs on Fashion-MNIST and 120 epochs on CIFAR-10 respectively. For CIFAR-10 we also decayed the learning rate by 0.1 at 50% and 75% of training.

Stage 2. Fine-tuning the stochastic prior. In the experiments presented in Section 8.2, we pruned the deterministic prior using MBP, according to the desired sparsity level. In order to maintain the same level of sparsity in our stochastic prior, we set λ^0 using the block isotropic prior introduced in Appendix A.1, setting $\epsilon = 0.0001$ in Eq. (9) in the end after initially tuning this hyperparameter. We train for 30 epochs on MNIST, 30 epochs on Fashion-MNIST and 60 epochs on CIFAR-10. We used the same learning rates as in Stage 1 for all tasks. Again on CIFAR-10, we decay the learning rate by 0.1 twice during the course of training in both Stage 2 and Stage 3. On all tasks we set the per weight variance σ_i^2 to be constant over parameters, and tune this constant on a log-uniform (base e) scale between $[-15, -3]$. Whenever we tune hyperparameters, we always seek to minimize the PAC-Bayes generalization bound on error. For both σ_i and posterior learning rate (next paragraph), we tuned only on a finite grid of values, and by a union bound argument this adds only a negligible penalty on top of our PAC-Bayes bound.

Stage 3. Training the posterior. We train for the same number of epochs as in Stage 2 in all tasks, using an initial learning rate that is tuned on a log-uniform scale between $[-4, -1]$, $[-7, -4]$, $[-5, 2]$ respectively for MNIST, Fashion-MNIST and CIFAR-10. In our experiments, we found the lowest generalization bounds to arise when pruning probabilities were set to extreme values, (close to 0 or 1), due to the performance-risk trade off as described in Appendix F.1. We set $\delta = 0.05$ when calculating our generalization bounds. In Section 8.2, our goal was to demonstrate that we can obtain good PAC-Bayes generalization bounds on small/moderate scale image classification tasks using our stochastic pruning framework. It turns out we can obtain surprisingly good results without training posterior pruning probabilities (see results in Fig. 7).

D.3 Hyperparameters choice for Probabilistic Fine-Tuning(PFT)

PFT uses the same hyperparameters as stated in Appendix D.2. More precisely, for training the weights of the unpruned deterministic classifier, we use the hyperparameters from Stage 1. Then we find a randomized sparse network by performing PFT using the hyperparameters from Stage 2. Finally, we fine-tune the weights of the corresponding deterministic sparse classifier using the hyperparameters from Stage 1.

E Additional Empirical Results

Here we present some additional experiments complementing our results in the main text on PFT and PBP.

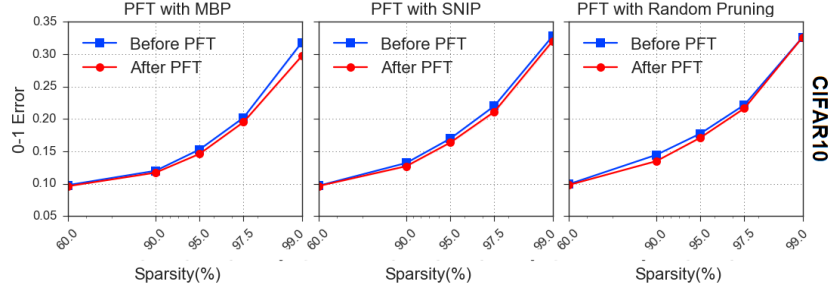


Figure 9: Impact of the application of PFT on standard One Shot Pruning methods (MBP, SNIP) and Random Pruning. CIFAR-10 with ResNet20.

E.1 Probabilistic Fine-Tuning (PFT)

This section complements the results on PFT shown in Fig. 5. Fig. 9 shows full results for CIFAR-10 experiments with ResNet20.

E.2 PFT at initialization

We run multiple experiments with PFT at initialization, where we skip the pre-training of an unpruned network phase (Step 1). Instead, we optimize a random pruning mask from the weights at initialization, in a similar fashion to [4–6, 29]. The goal is to assess whether PFT improves upon existing methods of One-Shot Pruning at initialization, namely Magnitude Pruning, SNIP [4] and Random Pruning. Fig. 10 shows that PFT improves the performance of existing OSP methods at initialization in almost all the settings we test. We hypothesize that this is due to the fine-tuning step finding masks that better adapt to the data.

We emphasize, however, that this additional benefit of fine-tuning is not that surprising: OSP methods use one forward/backward pass through the network and the training data to create the mask; our method requires multiple forward/backward passes to fine-tune the probabilities. However, it is interesting that even local optimization starting from the weights at initialization can boost the performance of pruning-at-initialization methods.

E.3 Strong Lottery Ticket Hypothesis and PFT

Strong Lottery ticket hypothesis (Strong LTH) states that at initialization there exists a sparse subnetwork that achieves competitive accuracy, if the network is sufficiently overparameterized. We apply PFT at initialization without any weight training before or after pruning. Fig. 11 presents the results on CIFAR-10, showing that PFT finds random pruning masks that, when applied to untrained weights, yield a network obtaining nontrivial test error. Here for a given sparsity we initialized the stochastic mask with either isotropic stochastic mask (PFT starting from a random mask) or a block isotropic stochastic mask (PFT starting from SNIP mask) set with $\epsilon = 0.01$ in Eq. (9), and trained the masks for 100 epochs with learning rate 0.5. Note in both cases without PFT, the stochastic mask would obtain 90% error on untrained weights.

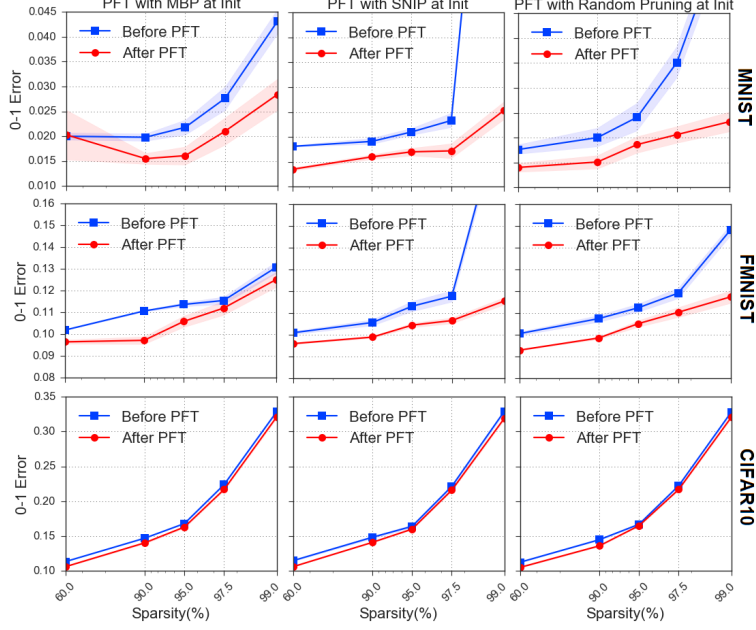


Figure 10: Impact of the application of PFT on standard One Shot Pruning methods (MBP, SNIP) and Random Pruning at Initialization. MNIST and FMNIST with (MLP 1000x3) and CIFAR-10 with ResNet20.

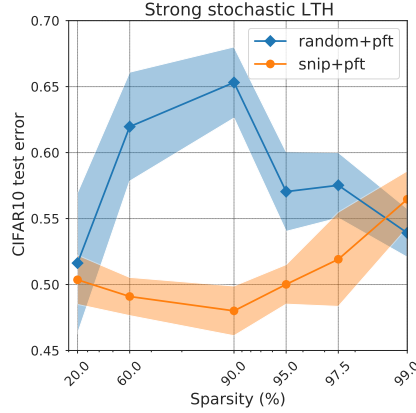


Figure 11: Strong stochastic Lottery Ticket Hypothesis test errors on CIFAR-10 with ResNet20 for PFT with random or SNIP initialisation of stochastic mask. Plotted intervals correspond to 95% confidence from 10 independent samples from the trained stochastic mask.

F PAC-Bayes Bound Optimization

F.1 Performance–risk bound tightness trade-off

The KL divergence term plays a crucial role in the PAC-Bayes bound optimization. To capture the impact of the KL term, we run multiple experiments where we down-weight the KL term in the bound, i.e., we multiply the KL term by a factor d and optimize the new bound. Intuitively, since the KL term acts as a penalty forcing the posterior Q to remain close to the prior P , we expect the test error of the posterior random classifier Q to improve as we down-weight the KL term. On the other hand, when we optimize the bound with the down-weighted KL term, we expect the real bound (i.e., without down-weighting) to become worse as Q is able to move away from P . This trade-off is illustrated in Fig. 12.

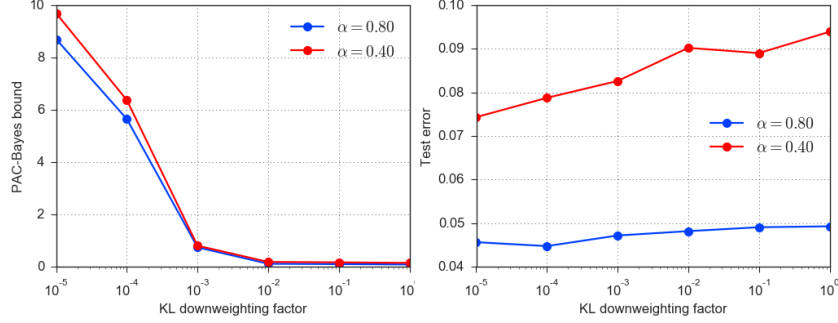


Figure 12: Performance–Bound trade-off with a 1000x3 MLP trained on MNIST, starting from Random Pruning.

F.2 Overall Entropy versus Error Bound

The binary entropy of the Bernoulli distribution $\mathcal{B}(p)$ is defined by $\text{Ent}(p) = -p \log_2(p) - (1 - p) \log_2(1 - p)$. The binary entropy measures the randomness of Bernoulli variables, it ranges from 0 for $p = 0$ or $p = 1$, to 1 for $p = 1/2$ (intuitively, when $p = 50\%$, the Bernoulli variable is completely random and does not favor one side to the other). In our experiments, we measured the overall entropy of our optimized stochastic pruning masks by taking an average over per-weight-entropy. Fig. 13 shows how the optimized PAC-Bayes bound on error correlates with entropy. We see a clear trade-off between the tightness of the bound and the overall entropy; as the overall entropy increases, the probabilistic mask changes more significantly across network samples.

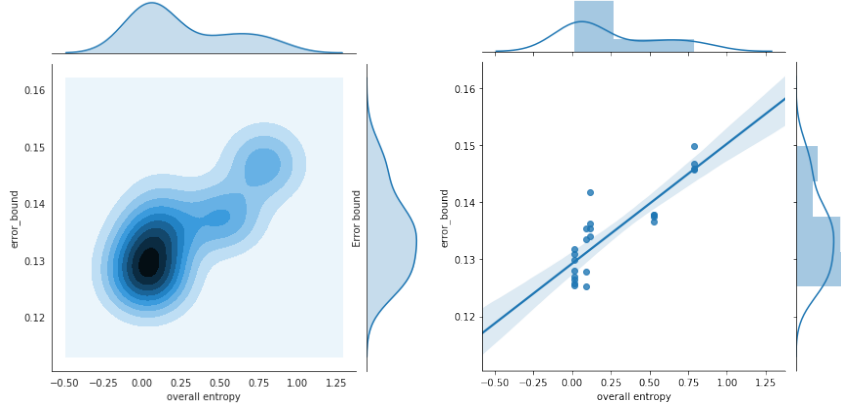


Figure 13: Overall stochastic pruning mask entropy vs PAC-Bayes bound on error with (MLP 1000x3) on FMNIST

F.3 Evolution of Prior and Posterior Probabilities during training

The prior and posterior pruning probabilities are optimized in the fine-tuning stage and posterior training (Stages 2, 3 in Section 6), respectively. The pruning probabilities are optimized to minimize the empirical risk in Stage 2 and PAC-Bayes bound in Stage 3. Starting from a randomly pruned classifier with sparsity $s = 90\%$, i.e., setting $\lambda_i^0 = 10\%$ for all i 's, through a histogram we monitored how the probabilities change after Stage 2 and Stage 3. Fig. 14 shows the histogram of pruning probabilities of the output layer of a 3 layer MLP with 1000 neurons per layer, trained on MNIST. At initialization, the pruning probabilities are concentrated at 10% (random pruning, Fig. 14a). After the fine-tuning phase (stage 2), Fig. 14b shows that the probabilities spread out over the interval $(0, 1)$, with some converging to 1 and others to 0, with the bulk remaining in $(0, 1)$. This suggests that the empirical risk of the prior is not minimized by a deterministic mask. In the posterior training phase

(Stage 3), from Fig. 14c (c) we see that the probabilities do not deviate from Stage 2. This is due to the KL term acting as a regularizer and forcing the posterior to stay close to the prior.

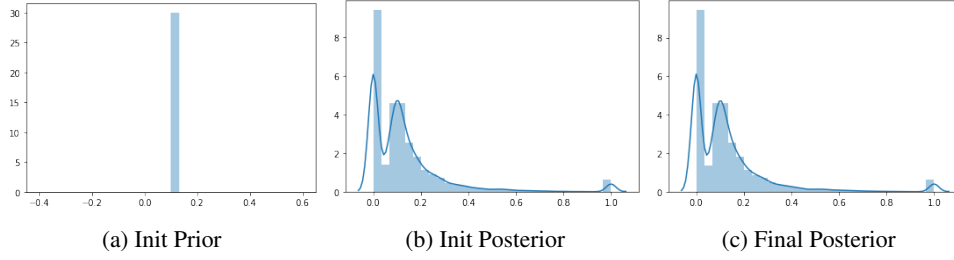


Figure 14: Probabilities Histogram of the output layer in a (MLP 1000x3) on MNIST, starting from Random Pruning

To visualize how the posterior pruning probabilities are distributed across the weights, in Fig. 15 we plot a 3D map of the posterior probabilities in the output layer (with dimensions 10 by 1000). The z -axis plots λ_i 's, i.e., the probability that the weights is kept. In Fig. 16, we further plot a heatmap of λ_i 's for the output layer. Both plots show that high values of λ_i are spread out across the output layer.

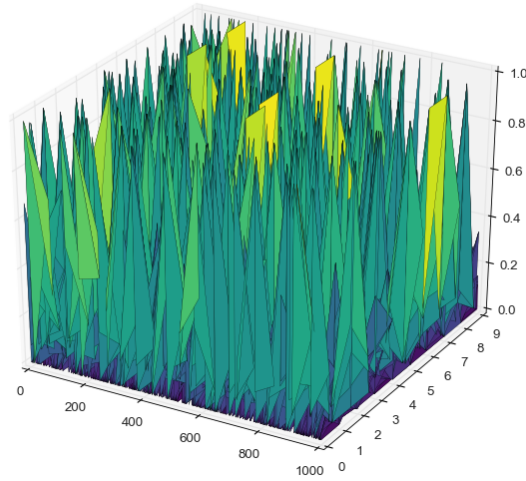


Figure 15: 3D plot of Posterior Probabilities in the output layer in a (MLP 1000x3) on MNIST, starting from Random Pruning.

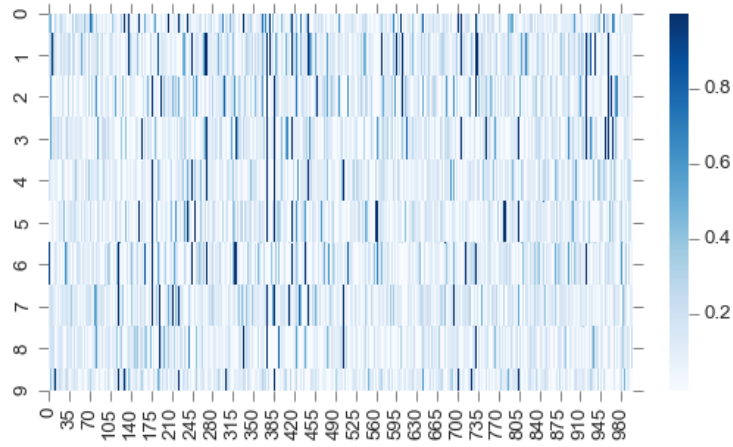


Figure 16: Heatmap plot of Posterior Probabilities in the output layer in a (MLP 1000x3) on MNIST, starting from Random Pruning.

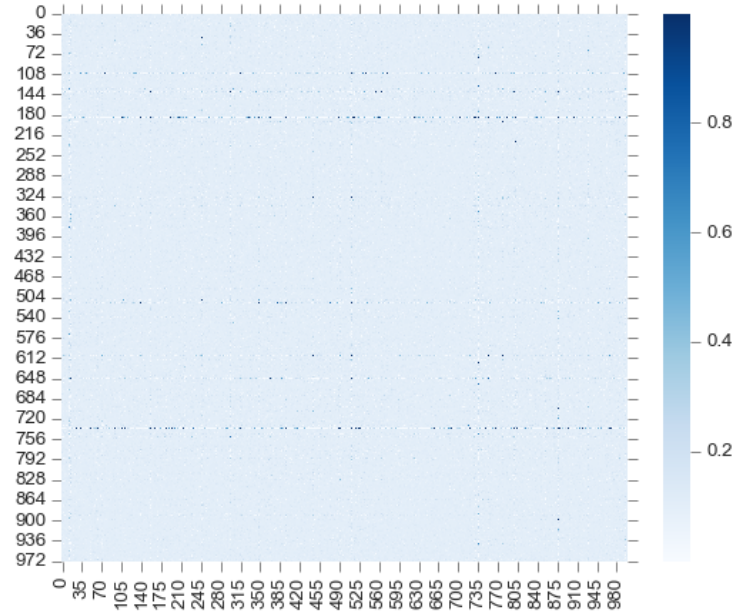


Figure 17: Heatmap plot of Posterior Probabilities in the mid layer in a (MLP 1000x3) on MNIST, starting from Random Pruning

F.4 PairPlot analysis for PBP

Here we seek to evaluate how various quantities appearing in PBP optimization interact with each other. Pairplot analysis is a valuable tool to depict such interactions. Fig. 18 shows a pairplot of 4 quantities: PAC-Bayes bound, test error of the posterior Q , the fraction of the weight used to train the prior P , and the normalized KL divergence. The first noticeable effect is how the choice of alpha impacts the normalized KL divergence. Indeed, the histogram of the normalized KL divergence shows two bulks of values, each corresponding to a value of alpha. Overall, larger alpha tends to yield smaller KL divergence. This can be attributed to the fact that using more data to train the prior would allow the posterior not to move far from the prior in order to find a good trade-off between the KL term and the empirical risk.

Another interesting observation is the almost perfect linear relationship between the posterior test error and the PAC-Bayes bound. This suggests that the KL term is usually negligible in comparison with posterior risk, showing that posterior training favors localized optimization (i.e., concentrated around the prior). This is also confirmed by the plot of the PAC-Bayes bound as a function of the normalized KL divergence; on average, high KL divergence yields worse PAC-Bayes bounds, confirming that fact that PBP optimizes posterior weights and pruning probabilities in the neighborhood of the prior.

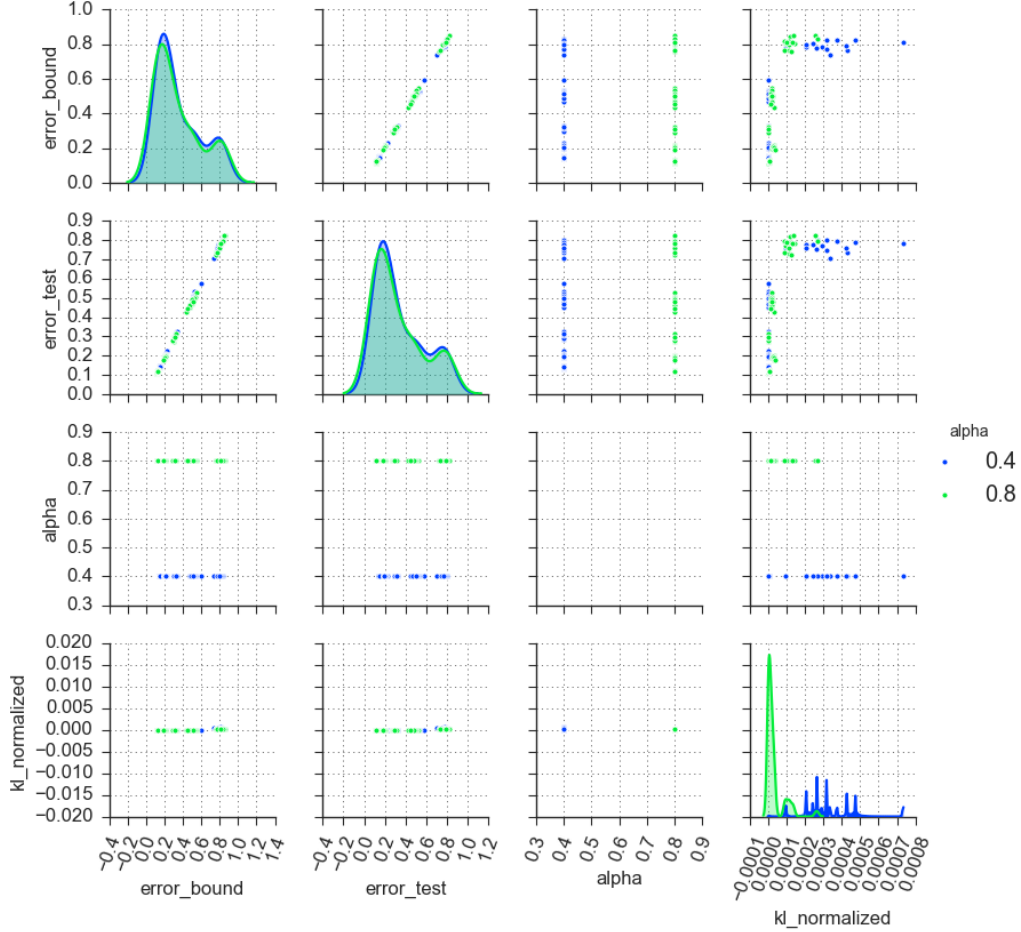


Figure 18: Pairplot of our experiments on CIFAR-10 with ResNet20, with Magnitude Based Pruning

F.5 Is the deterministic mask really deterministic?

It is well known that with standard pruning methods, the resulting sparse architecture depends on the initialization of the weights and data reshuffling seed. The dependence on the initialization is natural and expected: a weight that was initialized to be very small compared to others will likely remain small after training, especially in sufficiently overparameterized neural networks trained with small learning rates. However, the dependence of the mask on the data shuffling seed is interesting and suggests that there may be no single optimal sparse architecture, and a randomized mask might be better. We train a network multiple times with a fixed initialization. Each run has a different data shuffling seed. We keep track of the probability of the weight being pruned. Fig. 19 illustrates our result on an MLP with 5 layers. Here we only plot the weights that have an empirical pruning probability in $(0, 1)$. We show the standard deviation of these weights as a function of their average magnitude across the runs. We provide similar plots for CIFAR-10 with ResNet20 in Fig. 20.

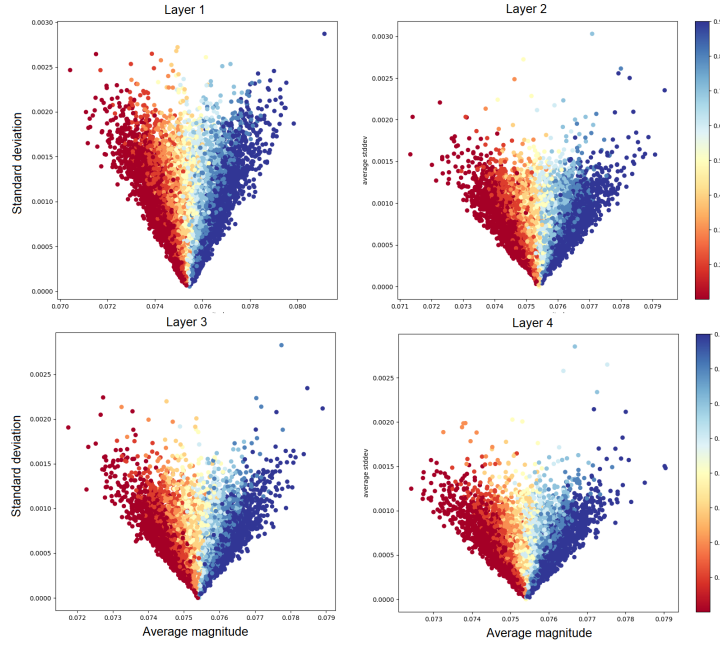


Figure 19: Plot showing the weights that change the mask (with Magnitude pruning) when data is reshuffled. X-axis is the average magnitude of the weight and Y-axis is its standard deviation. The architecture is MLP 1000x5 trained on MNIST. We only show the middle layers. The color scale represents the probability that the weights is unpruned (in 30 runs).

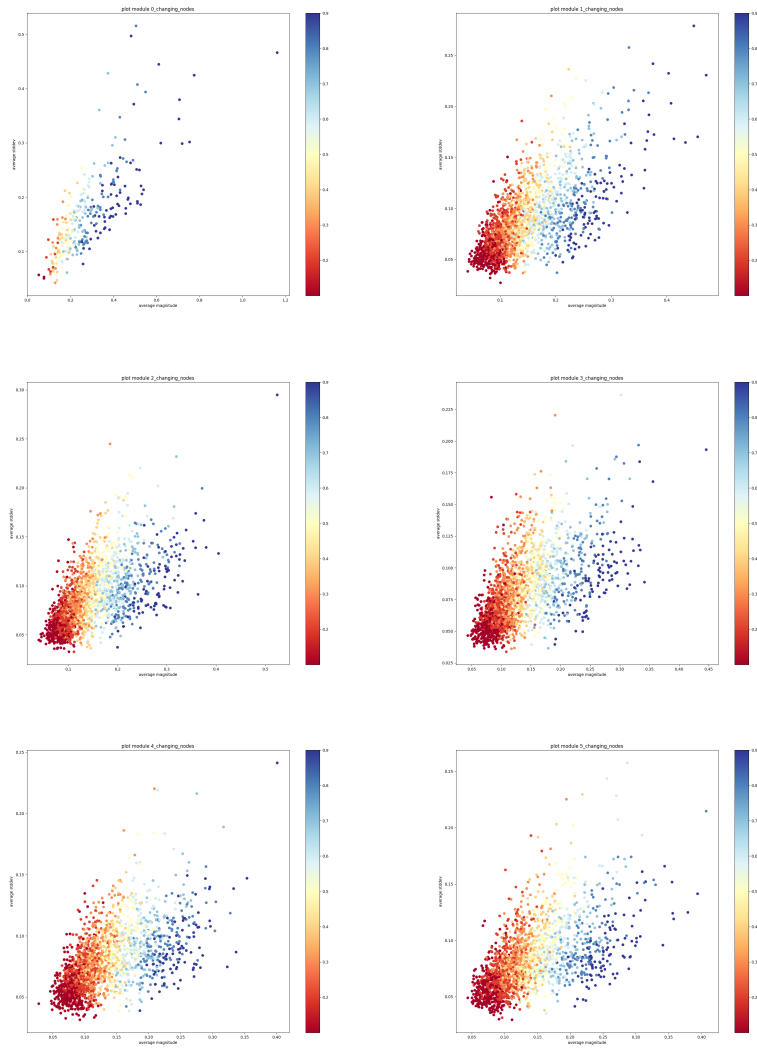


Figure 20: Plot showing the weights that change the mask (with Magnitude pruning) when data is reshuffled. X-axis is the average magnitude of the weight and Y-axis is its standard deviation. The architecture is ResNet20 trained on CIFAR-10.(first 6 layers)

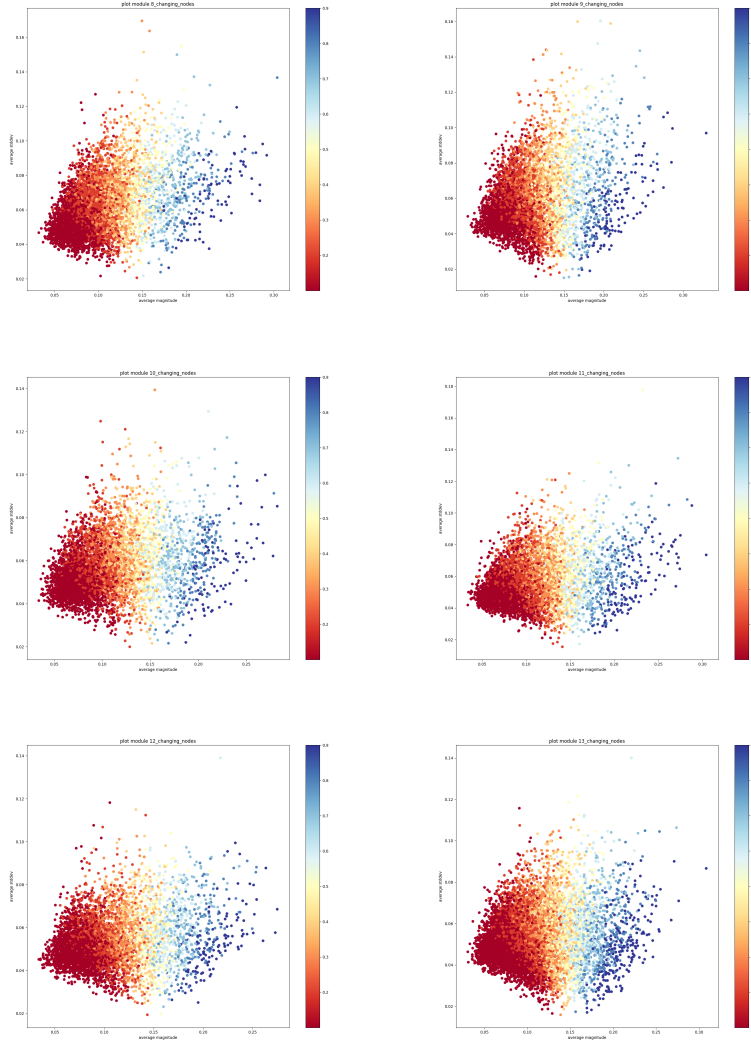


Figure 21: Plot showing the weights that change the mask (with Magnitude pruning) when data is reshuffled. X-axis is the average magnitude of the weight and Y-axis is its standard deviation. The architecture is ResNet20 trained on CIFAR-10. (layers 9-15).