

Understanding Decoupled and Early Weight Decay

Johan Bjorck, Kilian Q. Weinberger, Carla P. Gomes

Cornell University
{njb225,kqw4,gomes}@cornell.edu

Abstract

Weight decay (WD) is a traditional regularization technique in deep learning, but despite its ubiquity, its behavior is still an area of active research. Golatkar et al. have recently shown that WD only matters at the start of the training in computer vision, upending traditional wisdom. Loshchilov et al. show that for adaptive optimizers, manually decaying weights can outperform adding an l_2 penalty to the loss. This technique has become increasingly popular and is referred to as decoupled WD. The goal of this paper is to investigate these two recent empirical observations. We demonstrate that by applying WD only at the start, the network norm stays small throughout training. This has a regularizing effect as the effective gradient updates become larger. However, traditional generalization metrics fail to capture this effect of WD, and we show how a simple scale-invariant metric can. We also show how the growth of network weights is heavily influenced by the dataset and its generalization properties. For decoupled WD, we perform experiments in NLP and RL where adaptive optimizers are the norm. We demonstrate that the primary issue that decoupled WD alleviates is the mixing of gradients from the objective function and the l_2 penalty in the buffers of Adam (which stores the estimates of the first-order moment). Adaptivity itself is not problematic and decoupled WD ensures that the gradients from the l_2 term cannot "drown out" the true objective, facilitating easier hyperparameter tuning.

Introduction

The roots of weight decay (WD) go back to at least Tikhonov [1943], and within the context of deep learning, it has been used at least since 1987 [Hinton 1987]. Modern DNNs are typically trained with WD [Tan and Le 2019, Huang et al. 2017]. The technique is also used in modern NLP (natural language processing) [Ott et al. 2019, Radford et al. 2018] but is less commonly used in reinforcement learning. Despite its ubiquity, there is still ongoing research on WD – Golatkar, Achille, and Soatto [2019] have recently shown that WD essentially only matters at the start of the training in computer vision. Additionally, Loshchilov and Hutter [2017] have shown that WD interacts poorly with adaptive optimizers. The motivation of this paper is to investigate and explain these recent empirical observations on WD. It is

common to formulate WD as adding an l_2 penalty $\frac{1}{2}\lambda\|w\|_2^2$ to a loss function $L(w) = \frac{1}{|D|} \sum_{i \in D} \ell_i(w)$ for a dataset D and weights w . For SGD with batch B and learning rate α this leads to the following update

$$w_{t+1} = w_t - \frac{\alpha}{|B|} \nabla \sum_{i \in B} \ell_i(w_t) - \alpha \lambda w_t \quad (1)$$

By adding an l_2 penalty term, the weights w are "decayed" by a factor $(1 - \alpha\lambda)$ per update. Thus, it is common to use the terms weight decay and l_2 regularization interchangeably.

Background. The motivation for this work is to understand two recent observations. The first observation comes from Loshchilov and Hutter [2017], who show that for Adam [Kingma and Ba 2014], manually decaying weights can outperform an l_2 loss. As the gradient of the l_2 term will appear both in the numerator and denominator of the adaptive gradient step, these methods are not equivalent. Loshchilov and Hutter [2017] dub this technique decoupled weight decay and perform experiments on small-scale computer vision tasks, observing improved generalization and increased hyperparameter stability. This strategy has become increasingly popular [Wang et al. 2018, Radford et al. 2018, Carion et al. 2020, Liu et al. 2020] and is e.g. used in the Facebook NLP repository fairseq [Ott et al. 2019]. However, the motivation for this approach is primarily empirical. The second phenomenon we investigate is due to Golatkar, Achille, and Soatto [2019] who show that in computer vision, applying WD only during say the first quarter of training is essentially as good as always applying it, and applying it after the first quarter is roughly as good as never applying it. We refer to these two schedules as early/late WD. We focus on the first quarter in this paper for concreteness but note that the same trend holds beyond exactly the first quarter. We will relate the observations of Golatkar, Achille, and Soatto [2019] to the sharp/flat minima hypothesis of Keskar et al. [2016] which essentially states that the noise in SGD biases the network to flat minimizers which generalize well.

Our Contributions. Regarding observations of [Golatkar, Achille, and Soatto 2019], we show that the network norm typically grows during the start, using WD early in training then ensures that the gradient steps are large relative to the weights throughout training. This has a regularizing effect, but traditional metrics of generalization [Keskar et al. 2016] do not consistently capture this. We provide

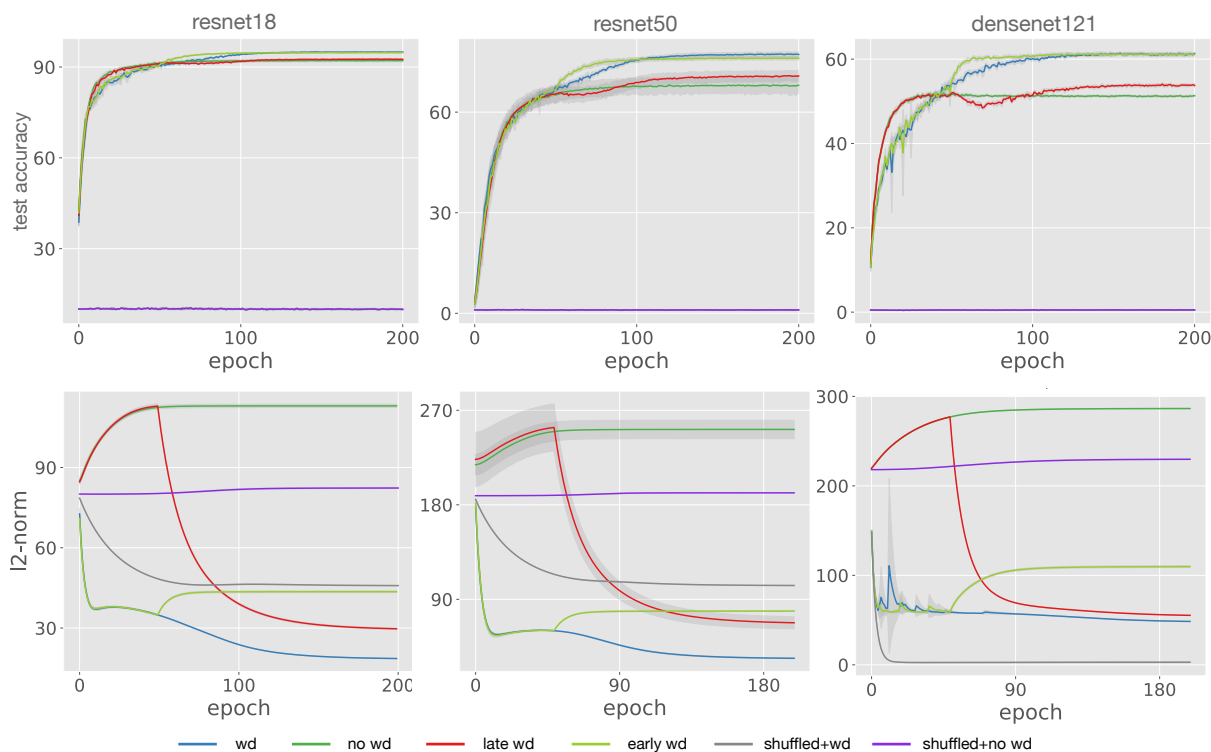


Figure 1: (*Top.*) Golatkar, Achille, and Soatto [2019] have shown that for image classification, starting WD only after epoch 50 brings little benefit, whereas stopping it after epoch 50 performs on par with using it throughout the training. (*Bottom.*) The l_2 norm of the weights increases dramatically at the start. Applying WD only during early parts of training ensures small weights throughout the optimization process. By applying WD late, it takes many epochs for the norm to shrink. We also plot curves for networks trained on datasets with shuffled labels and note that weight norms grow less under such settings.

a scale-invariant metric to remedy this issue. We further demonstrate that dataset generalization properties significantly influence weight growth. Regarding observations due to Loshchilov and Hutter [2017], it is natural to believe that l_2 regularization and adaptivity are incompatible. We demonstrate that across RL (reinforcement learning) and NLP task, this is not the issue that decoupled weight decay solves, but instead that the gradients of l_2 terms can “drown” the gradient of the true objective function in the buffers of Adam [Kingma and Ba 2014] (which stores estimates of the first and second-order moments of gradients). By decoupling the WD, the buffers are not shared between l_2 regularization and the true objective function, avoiding this mixing and facilitating hyperparameter tuning. We find no increase in absolute performance over sufficiently tuned WD, suggesting that hyperparameter stability rather than improved accuracy might be primarily responsible for decoupled WDs popularity. We conclude with lessons for practitioners regarding tuning and using WD.

On the Temporal Dynamics of Weight Decay

For investigating observations in Golatkar, Achille, and Soatto [2019] we replicate their experimental setup with identical hyperparameters (listed in the Appendix), training Resnet18 on Cifar10 and Resnet50 on Cifar100. We addi-

tionally provide experiments on tiny-imagenet [Karpathy, Li, and Johnson 2017 (accessed 2020-01-01)] using densenet 121 [Huang et al. 2017]. We consider this setting throughout the paper. In Figure 1, we show the weight norm and accuracy of networks with, without and with WD only after/before epoch 50 as per [Golatkar, Achille, and Soatto 2019]. We also consider a network with shuffled labels. We see that the norms of the network grow primarily at the start. By applying WD before epoch 50, we avoid the initial period of growth, and the norm stays low throughout training. Applying WD after epoch 50 results in many epochs before the norm reaches levels comparable to using WD throughout training.

Early/Late Weight Decay and Generalization. As per Figure 1, applying WD only at the start will ensure that the weight norm stays low during training. However, it is not clear why this would improve generalization; almost all network layers use batch normalization and are thus invariant under weight-rescaling. For a fixed learning rate and gradient, the effective change $\Delta w/w$ in the weights is smaller if the weights have a larger scale – so decaying the weights increases the “effective” learning rate. A large learning rate and small batches typically have a regularizing effect as they induce noise into the training, and Keskar et al. [2016] have shown that large batches lead to sharp minimizers with

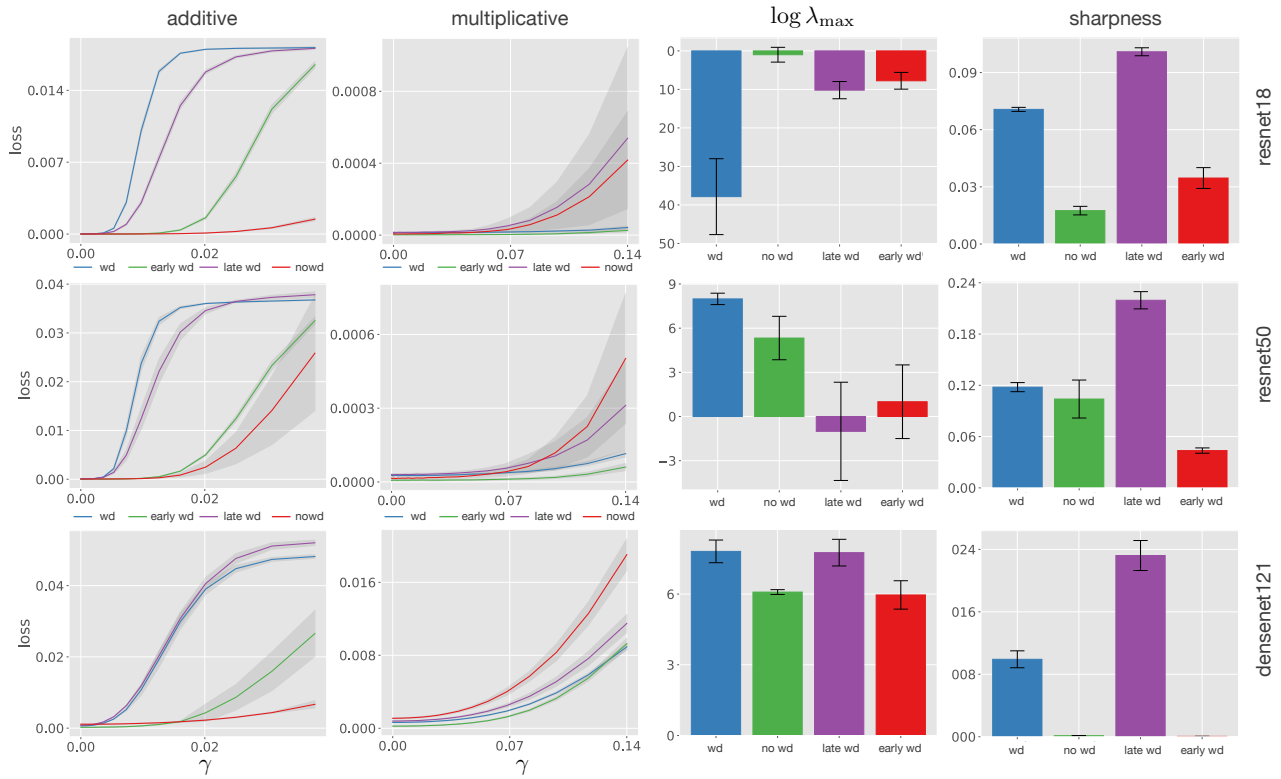


Figure 2: The sharpness of networks, typically used as a proxy for generalization, using different WD schemes. We compare four metrics of sharpness: the largest hessian eigenvalues (compute via Yao et al. [2019], measured logarithmically), the sharpness metric in Keskar et al. [2016] and additive/multiplicative perturbations. All metrics except multiplicative perturbations fail to consistently explain the differences in generalization of Figure 1. The loss under multiplicative perturbations increases when WD isn’t used, suggesting that a sharp minima hypothesis might explain observations of [Golatkar, Achille, and Soatto 2019].

poor generalization. Their explanation, which has garnered much attention [Li et al. 2018], is essentially that networks with sharp minima generalize worse as they are more sensitive to the inherent shift between test/train surface. Keskar et al. [2016] uses the following metric of sharpness (with $\epsilon = 5e-4$) for loss $L(\cdot)$ at a point x

$$C_\epsilon = \left\{ y \in R^n \mid -\epsilon(1 + |x_i|) \leq y_i \leq \epsilon(1 + |x_i|) \right\} \quad (2)$$

$$\max_{y \in C_\epsilon} \frac{L(x+y) - L(x)}{L(x) + 1}$$

In practice, $L(x+y)$ is maximized by first-order methods. Another common metric of sharpness is the largest eigenvalues of the Hessian [Iyer et al. 2020, Dinh et al. 2017]. In Figure 2, we plot these metrics and see that they typically give wrong or inconsistent results – for example, the sharpness metric of Keskar et al. [2016] suggests that disabling WD yields flat minima which should generalize well – the opposite of what we observe. We note that these metrics depend upon the scale of the network, and as per Figure 1, we know that networks without WD have larger norms. This motivates us to consider metrics of sharpness which are invariant under weigh scaling. We consider a simple scale-invariant metric – multiplicative perturbations of the network weights

$$S(\gamma) = \mathbb{E}[L(w \odot (1 + \gamma\delta))] \quad \delta \sim N(0, I) \quad (3)$$

That is, we scale each weight w_i by $(1 + \gamma\delta_i)$ where δ_i is a standard normal variable, the intuition being that this metric measures how the loss changes to small multiplicative perturbations, say $\gamma \approx 0.1$. Note that this yields a metric similar to eq. (2), with the chief difference being in how “small perturbations” are defined. The expectation is computed by sample averages.

Figure 2 illustrates this metric (referred to as “multiplicative”), showing that it gives results consistent with Golatkar, Achille, and Soatto [2019]. We also show the results of an additive perturbation, which is analogous to multiplicative ones except that we take $w + \gamma\delta$. Note that it fails to capture generalization. Thus, we show that the explanation that sharp networks generalize worse can be applied to the empirical observations on early/late WD due to Golatkar, Achille, and Soatto [2019] if one is careful regarding what is meant by sharpness. Our experiments also suggest that the effects of early/late WD are primarily mediated by modifying the effective learning rate $\Delta w/w$. To further solidify this hypothesis, we train networks without WD, but inspired by Zhang et al. [2018], we manually scale the weights after each epoch to match the norm of another network trained with WD. Figure 3 shows that just scaling the weight norms is enough to achieve the results of Golatkar, Achille, and




Figure 3: Learning curves for DNNs trained without WD with weights scaled to match norms in Fig. 1, and the original learning curves that these DNNs are made to match. Scaling the weights roughly matches the performance of various WD schedules, suggesting that WD mediates the observations of Golatkar, Achille, and Soatto [2019] through a simple scaling mechanism.

Soatto [2019]. See the Appendix for further experiments without batch normalization and discussion.

On Causes for Weight Growth. We have seen how applying WD early results in small weights throughout training on computer vision datasets, increasing the effective learning rate $\Delta w/w$. It is natural to believe that the network norm always grows during early parts of training, we here demonstrate that that’s *not* the case. Instead, the tendency of weight norms to grow is related to the dataset and its generalization properties. In Figure 1, we see that the weight norm of a network with shuffled labels stays almost constant during training. It’s natural to wonder if the gradient norm might simply be smaller for shuffled labels, but this turns out to not be true, see the Appendix. Indeed, the weights of networks trained on shuffled labels move significantly, just not in the radial direction, which would increase the weight norm, see the Appendix. With shuffled labels, the dataset has the same images, but training on such a dataset will not generalize to a test set. This suggests that dataset generalization properties have an important influence on weight norms, which in turn modulates effective learning rates. To understand why the weights grow differently using original or shuffled labels,

let us consider the weight norm change for an SGD update

$$\|w_{t+1}\|^2 - \|w_t\|^2 = \underbrace{\alpha^2 \|\nabla \ell_t\|^2}_{\text{square term}} + \underbrace{2\alpha \langle -\nabla \ell_t, w_t \rangle}_{\text{cross term}} \quad (4)$$

There are two terms responsible for the increasing weights, a term that only depends on the gradient update and a cross term relating the direction of the gradient and the weights. In Figure 4, we illustrate how these two terms vary during optimization and find that the cross term is responsible for the lions share of the weight growth. Let us further divide the loss function into two parts representing the correct class and the normalization constant used in softmax

$$\ell_t(x) = \frac{1}{|B|} \sum_{i \in B} \underbrace{x_{i, \text{label}[i]}}_{\ell_{\text{pos}}} - \log \left(\sum_j \exp(x_{ij}) \right) \quad (5)$$

By linearity we of course have $\nabla \ell_t = \nabla \ell_{\text{pos}} + \nabla \ell_{\text{neg}}$ and thus $\langle \nabla \ell_t, w_t \rangle = \langle \nabla \ell_{\text{pos}}, w_t \rangle + \langle \nabla \ell_{\text{neg}}, w_t \rangle$. In Figure 5 (top) we show how these two terms vary during optimization, observe that $-\nabla \ell_{\text{pos}}$ points along the weights while $-\nabla \ell_{\text{neg}}$ points away from them. In light of Figure 4, we conclude that weight norms increases due to the gradient pointing roughly in the radial direction w , scaling up many weights w_i . Can this interpretation explain why shuffled labels lead to no




Figure 4: The contributions to the change in weight norm for the square and cross terms, defined in (4). The cross term dominates and thus the norm grows primarily in the radial direction, scaling up subsets of the weights that align with the gradient.

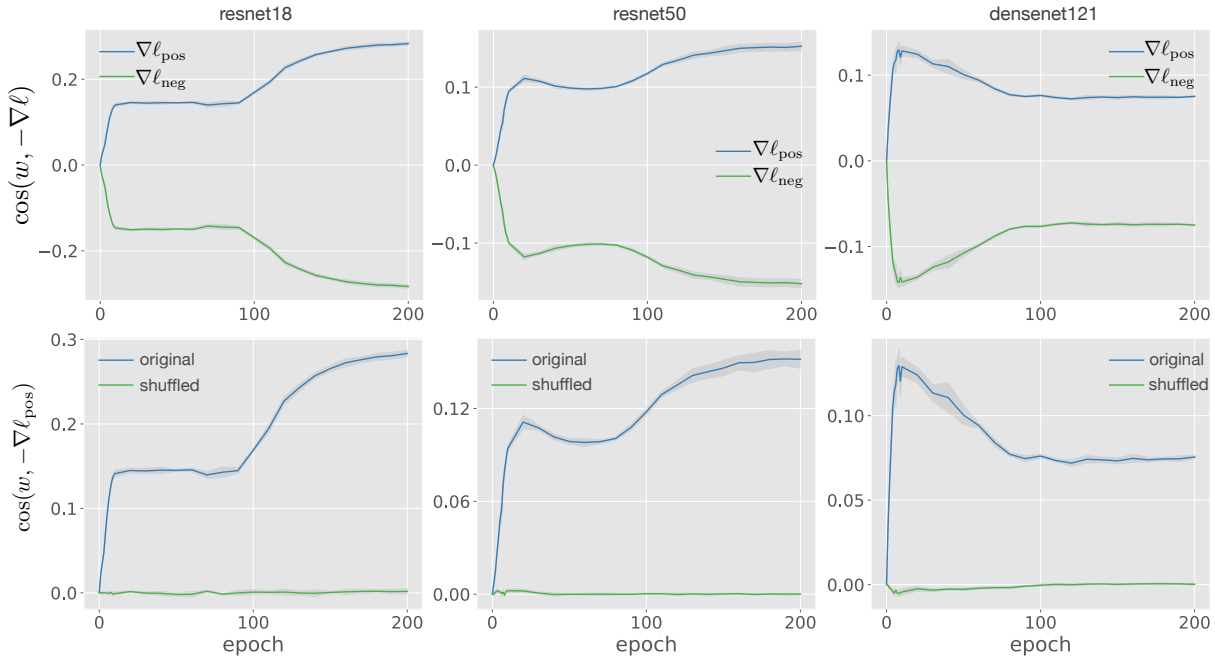


Figure 5: (*Top.*) We divide the cross entropy loss into two parts as per (5). The cosine between the weight vector w and $-\nabla\ell_{\text{pos}}$ is positive whereas the cosine between w and $-\nabla\ell_{\text{neg}}$ is negative. This suggests that network norm increases as subset of weights responsible for correct predictions grow in magnitude. (*Bottom.*) $\cos(w, -\nabla\ell_{\text{pos}})$ with ℓ_{pos} defined as per (5). We see that for a network with shuffled labels the gradient barely points in the radial direction, which would lead to less growth as per Figure 4.

weight growth? We first note that $\nabla\ell_{\text{neg}}$ is invariant under label permutations, and thus seek to look at $\nabla\ell_{\text{pos}}$. Figure 5 (bottom) instead plots $\cos(-\nabla\ell_{\text{pos}}, w_t)$ for the standard network and a network with shuffled labels. There we see a striking difference, for the network trained on the original labels, the gradient typically points along the w whereas gradients for networks using shuffled labels do not. To explain this, consider e.g. the last mini-batch b we encounter in the first pass over the dataset. If we use the original labels, all images of e.g. dogs we have seen previously will likely push the network weights to increase the prediction probability of any dog pictures in batch b . Scaling up these weights will then decrease the loss. If we use shuffled labels however, simply scaling up network weights should not decrease the loss on batch b , since there is no generalization from dog pictures in previous batches. While an example with shuffled labels might seem artificial, the phenomenon of datasets influencing the weight norm growth happens in more natural settings such as RL. In the Appendix we show that the network norm differ substantially between games when using identical hyperparameters for DQN [Mnih et al. 2015]. Thus, if norm growth is dataset dependent, the observations of [Golatkar, Achille, and Soatto 2019] might only hold for datasets with good generalization.

On Weight Decay for Adaptive Optimizers

Loshchilov and Hutter [2017] have proposed decoupled weight decay for adaptive optimizers, where one decays the weights by $(1 - \alpha\lambda)$ instead of adding a l_2 penalty

to the loss. We investigate this scheme in two contexts where adaptive optimizers are ubiquitous, NLP and RL. We first consider translation of the IWSLT’14 German to English dataset [Cettolo et al. 2014] using transformer architectures [Vaswani et al. 2017] with code and default hyperparameters from the publicly available fairseq codebase [Ott et al. 2019]. We consider $\lambda \in \{1e-3, 1e-4, 1e-5\}$, where the middle parameter is the default parameter used in fairseq, see the Appendix for all hyperparameters. Secondly, we also consider the RL agent DQN [Mnih et al. 2015], using the publicly available dopamine codebase [Castro et al. 2018] with their default hyperparameters (see the Appendix), trained on a handful of Atari games, most having been highlighted in previous work [Mnih et al. 2016]. The three rightmost plots in Figure 7 shows that for translation, WD under-performs decoupled WD except for the smallest value of λ . Similarly, Figure 6 shows that decoupled WD typically gives a sizable improvement in DQN, whereas WD can have a markedly deleterious effect on performance. To investigate why standard WD fails whereas decoupled WD succeeds, let us consider the buffer that Adam maintains to estimate the first moment of the gradient, which for loss function ℓ and l_2 loss is updated as

$$m^{t+1} \leftarrow (1 - \beta_1)m^t + \beta_1\nabla\ell + \beta_1\lambda w_i \quad (6)$$

In the leftmost plot of Figure 7, we consider the NLP task with WD turned off and show the distribution of the quantity $\frac{|m_i|}{|w_i|}$ for weight w_i , which roughly measures the strength of the gradient signal over the weight. The analogue illustration

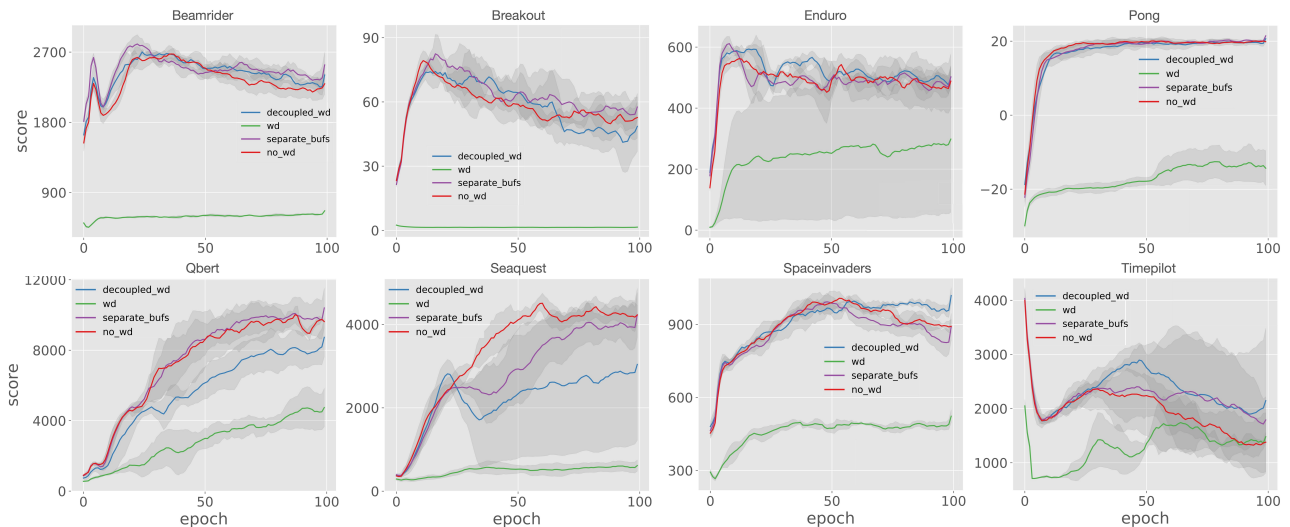


Figure 6: Learning curves for various Atari games with WD ($\lambda = 0.0001$). We compare decoupled WD [Loshchilov and Hutter 2017], original WD, no WD and an Adam variant with separate buffers for the WD and gradient signal. Original WD underperforms, whereas separating the buffers performs on par with decoupled WD. This suggests that the mixing of WD signal with the gradients, rather than the adaptivity itself, is responsible for the poor performance of normal WD in this setting.

for DQN is found in the Appendix. In both these cases, we see that the distribution of absolute values of this quantity (plus ϵ for numerical stability) on a log scale, and see that 1) that the gradient signal is weak compared to the weight and 2) the scales are different by orders of magnitude for different weights. This means that the ratio between the gradients from the true objective function (the gradient signal) and the gradients from an l_2 penalty differs significantly between individual weights w_i . To avoid the WD signal dominating over the gradient signal in (6), one would need to set λ comparable to the smallest gradient signal. However, this might result in a very small value for the parameter with the largest gradient signal. Thus, effectively, the suitable ranges of λ are dictated by the strength of the gradient signal. We can make this idea more precise with a scaling argument. For l_2 regularized Adam with weight w_i and gradient strength g_i , equal to say the absolute value of an expo-

ponential average of the gradients, should shrink until we reach a steady state where $\lambda w_i \approx g_i$. If we assume that the ratio $m_i/(\sqrt{m_2} + \epsilon)$ of the Adam buffers are $\mathcal{O}(1)$ (i.e. the first moment and the square of the second moment are comparable), the effective update $\Delta w_i/w_i$ would be $\mathcal{O}(\alpha\lambda/g_i)$. For decoupled weight decay, the weights would shrink only until $\lambda w_i = \mathcal{O}(1)$ since Adam without WD is invariant under scaling of the gradient g_i . Thus the relative update $\Delta w_i/w_i$ would be $\mathcal{O}(\alpha\lambda)$. The important distinction is that the relative updates for decoupled WD only scales with hyperparameters we have control over, whereas for l_2 regularization it depends upon the dataset gradient signal which we cannot control, do not know a priori, and which might vary between parameters as per Figure 7.

This hypothesis predicts that the mixing of the WD signal and the gradient signal inside the Adam buffers is the important distinction between decoupled WD and l_2 regu-




Figure 7: The leftmost figure illustrates the quantiles of $\log |m_i/w_i|$ during training of a transformer [Vaswani et al. 2017]. They vary roughly two orders of magnitude, suggesting that the gradients for different parameters differ substantially. The three following figures illustrate the translation quality, measured in bleu, for three different values of λ and three different weight decay schemes. Standard WD underperforms unless λ is taken small whereas separating the buffers matches decoupled WD.

	Beamr	Breako	Enduro	Pong	Qbert	Seaq	Spaceinv	Timep	λ
orig	2293	52	488	19	9624	4230	892	1379	0.0
decoupled	2843	77	545	22	11430	2956	1144	3136	1e-3
WD	579	1	62	-27	356	140	346	822	1e-3
separated	3310	62	560	22	11369	1283	969	1784	1e-3
decoupled	2406	48	501	20	8733	3043	1018	2146	1e-4
WD	666	1	298	-14	4746	614	522	1480	1e-4
separated	2535	57	502	21	10400	4231	889	1790	1e-4
decoupled	2481	58	517	21	10108	3358	956	699	1e-5
WD	2375	114	230	16	11181	4856	1153	2821	1e-5
separated	2367	47	565	21	8625	3055	846	2359	1e-5
WD	3257	89	589	21	11134	5095	1158	6707	1e-6
WD	2947	49	475	21	9867	3953	748	4633	1e-7
WD	2761	56	505	20	8384	4758	649	1927	1e-8

Table 1: Average scores over three seeds for various Atari games and WD schemes. Standard WD consistently fails whereas most but not all games benefit from decoupled WD. Adding separate buffers, which stores estimates of the first and second order moments of the gradients, for the normal gradient and weight decay signal gives performance roughly matching decoupled WD. See the Appendix for learning curves with standard deviations.

larization, and not the adaptivity itself. By allowing separate buffers in Adam (for both the first and second-order moments) for the gradients of the true objective and an l_2 penalty, we can investigate if the signal mixing indeed is the problem. We thus consider Adam with duplicate buffers m_i, m'_i, v_i, v'_i for the gradient and WD signal, see the Appendix for a formal description. Note that as the gradient of the weight decay term appears in both the numerator and denominator of the buffers, the magnitude of the update for this scheme is invariant if the weight is rescaled, which is different from decoupled WD. Table 1 and Figure 6 shows the result of this experiment for DQN, and we see that separating the buffers indeed leads to performance comparable to decoupled WD. Similarly, we can see for translation in the three rightmost plots of Figure 7 that separating the buffers matches the performance of decoupled WD. We see that for sufficiently small λ , normal WD indeed does give an improvement in DQN. But what λ is sufficiently small differs by at least an order of magnitude between games. Certain games (e.g., Enduro or Pong) requires $\lambda \leq 1e-6$ to give a comparable performance of no WD, whereas $1e-4$ suffices for Timepilot. WD thus requires tuning λ , whereas decoupled WD is stable as observed by [Loshchilov and Hutter 2017]. We also note that WD sometimes outperforms decoupled WD, albeit with highly tuned λ , suggesting that the popularity of decoupled WD might be due to hyperparameter stability rather than absolute performance improvement. Indeed, state-of-the-art image classification network efficientnet [Tan and Le 2019] does not use decoupled WD for its adaptive optimizer.

Discussion

Related work. Weight decay has a long history as a regularizer in machine learning [Hinton 1987, Krogh and Hertz 1992]. The hypothesis that flat minima generalize is well-

known [Hochreiter and Schmidhuber 1997], and has been proposed to explain why large batch learning fails to generalize [Keskar et al. 2016]. The most prominent critique of the sharp-minima hypothesis comes from Dinh et al. [2017], who proves that one can increase the sharpness of any given minima by reparametrizing the network. Similar criticism can be found in theoretical PAC-Bayes work [Tszuku, Sato, and Sugiyama 2019, Rangamani et al. 2019, Yi et al. 2019, Neyshabur et al. 2017], that only provides experiments for large-vs-small batch sizes where standard sharpness metrics work well in practice. Van Laarhoven [2017] noted how WD would increase the relative size of gradient updates. This perspective was empirically substantiated in Zhang et al. [2018] who showed that this is the primary mechanism by which WD improves generalization and also argues for the conditioning effect of decoupled WD. Zhang et al. [2018] is the only previous work on decoupled WD that we are aware of, whereas they primarily replicate experiments of [Loshchilov and Hutter 2017] and discuss the KFAC optimizer, we focus on explaining *why* decoupled WD improves hyperparameter stability. We do not know any work explaining the observations of [Golatkar, Achille, and Soatto 2019].

Lessons for practitioners. Our work points towards a few directly actionable insights. 1) Decoupled WD is useful in q-learning despite not being broadly used. However, different environments may need a separate WD parameter due to their different generalization behavior, suggesting the need for adaptive versions of WD. 2) Different datasets have norms that grow differently. Consequently, one should not naïvely transfer WD parameters between datasets, especially when they have different generalization properties. 3) If standard WD is used, one should pay close attention to the scale between gradient and WD signal when tuning λ . 4) Since the weight norm is the most important factor when using WD, one can apply WD only every few batches to save computational resources. A toy example of this on cifar10

is shown in the Appendix, where WD is applied only every 128 batches with no performance cost. While WD rarely is the computational bottleneck, it cannot effectively be parallelized in a mirrored distributed strategy. Applied to more computationally intensive regularization such as Xie et al. [2019], this strategy might lead to substantial savings for larger models.

Conclusions. We have investigated recent empirical observations regarding WD. We observe that applying WD at the start increases the effective learning rate, which biases the network to less sharp minima. We also demonstrate that the primary distinction between decoupled weight decay and l_2 regularization is the sharing of buffers in Adam.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant Number CCF-1522054. This material is also based upon work supported by the Air Force Office of Scientific Research under award number FA9550-18-1-0136. This research is supported in part by the grants from Facebook, the National Science Foundation (III-1618134, III-1526012, IIS1149882, IIS-1724282, and TRIPODS- 1740822), the Office of Naval Research DOD (N00014- 17-1-2175), Bill and Melinda Gates Foundation. We are thankful for generous support by Zillow and SAP America Inc. We are also grateful from generous support from the TTS foundation. This work was partially supported by the Cornell Center for Materials Research with funding from the NSF MRSEC program (DMR-1719875).

Ethics Statement

Our work extends the research community’s understanding of weight decay, which is ubiquitously used in critical applications via neural networks. We do not perceive any entity to be directly put at a disadvantage or to be harmed due to any system failure. We do not believe that our research methods leverage biases in the data.

References

Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-End Object Detection with Transformers. *arXiv preprint arXiv:2005.12872* .

Castro, P. S.; Moitra, S.; Gelada, C.; Kumar, S.; and Bellemare, M. G. 2018. Dopamine: A Research Framework for Deep Reinforcement Learning URL <http://arxiv.org/abs/1812.06110>.

Cettolo, M.; Niehues, J.; Stüker, S.; Bentivogli, L.; and Federico, M. 2014. Report on the 11th iwslt evaluation campaign, iwslt 2014. In *Proceedings of the International Workshop on Spoken Language Translation, Hanoi, Vietnam*, volume 57.

Dinh, L.; Pascanu, R.; Bengio, S.; and Bengio, Y. 2017. Sharp minima can generalize for deep nets. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 1019–1028. JMLR. org.

Golatkar, A. S.; Achille, A.; and Soatto, S. 2019. Time Matters in Regularizing Deep Networks: Weight Decay and Data

Augmentation Affect Early Learning Dynamics, Matter Little Near Convergence. In *Advances in Neural Information Processing Systems*, 10677–10687.

Hinton, G. E. 1987. Learning translation invariant recognition in a massively parallel networks. In *International Conference on Parallel Architectures and Languages Europe*, 1–13. Springer.

Hochreiter, S.; and Schmidhuber, J. 1997. Flat minima. *Neural Computation* 9(1): 1–42.

Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4700–4708.

Iyer, N.; Thejas, V.; Kwatra, N.; Ramjee, R.; and Sivathanu, M. 2020. Wide-minima Density Hypothesis and the Explore-Exploit Learning Rate Schedule. *arXiv preprint arXiv:2003.03977* .

Karpathy, A.; Li, F.-F.; and Johnson, J. 2017 (accessed 2020-01-01). BWorld Robot Control Software. URL <https://tiny-imagenet.herokuapp.com>.

Keskar, N. S.; Mudigere, D.; Nocedal, J.; Smelyanskiy, M.; and Tang, P. T. P. 2016. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836* .

Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* .

Krogh, A.; and Hertz, J. A. 1992. A simple weight decay can improve generalization. In *Advances in neural information processing systems*, 950–957.

Li, H.; Xu, Z.; Taylor, G.; Studer, C.; and Goldstein, T. 2018. Visualizing the loss landscape of neural nets. In *Advances in Neural Information Processing Systems*, 6389–6399.

Liu, J.; Blessing, T. S. L.; Wood, K. L.; and Lim, K. H. 2020. CrisisBERT: Robust Transformer for Crisis Classification and Contextual Crisis Embedding. *arXiv preprint arXiv:2005.06627* .

Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* .

Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, 1928–1937.

Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540): 529–533.

Neyshabur, B.; Bhojanapalli, S.; McAllester, D.; and Srebro, N. 2017. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, 5947–5956.

Ott, M.; Edunov, S.; Baevski, A.; Fan, A.; Gross, S.; Ng, N.; Grangier, D.; and Auli, M. 2019. fairseq: A Fast, Extensible

Toolkit for Sequence Modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*.

Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. 2018. Improving language understanding by generative pre-training. URL https://s3-us-west-2.amazonaws.com/openai-assets/researchcovers/languageunsupervised/language_understanding_paper.pdf.

Rangamani, A.; Nguyen, N. H.; Kumar, A.; Phan, D.; Chin, S. H.; and Tran, T. D. 2019. A scale invariant flatness measure for deep network minima. *arXiv preprint arXiv:1902.02434*.

Tan, M.; and Le, Q. V. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*.

Tikhonov, A. N. 1943. On the stability of inverse problems. In *Dokl. Akad. Nauk SSSR*, volume 39, 195–198.

Tsuzuku, Y.; Sato, I.; and Sugiyama, M. 2019. Normalized flat minima: Exploring scale invariant definition of flat minima for neural networks using pac-bayesian analysis. *arXiv preprint arXiv:1901.04653*.

Van Laarhoven, T. 2017. L2 regularization versus batch and weight normalization. *arXiv preprint arXiv:1706.05350*.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.

Wang, J.; Yuan, Y.; Li, B.; Yu, G.; and Jian, S. 2018. Sface: An efficient network for face detection in large scale variations. *arXiv preprint arXiv:1804.06559*.

Xie, C.; Tan, M.; Gong, B.; Wang, J.; Yuille, A.; and Le, Q. V. 2019. Adversarial Examples Improve Image Recognition. *arXiv preprint arXiv:1911.09665*.

Yao, Z.; Gholami, A.; Keutzer, K.; and Mahoney, M. 2019. PyHessian: Neural Networks Through the Lens of the Hessian. *arXiv preprint arXiv:1912.07145*.

Yi, M.; Meng, Q.; Chen, W.; Ma, Z.-m.; and Liu, T.-Y. 2019. Positively scale-invariant flatness of relu neural networks. *arXiv preprint arXiv:1903.02237*.

Zhang, G.; Wang, C.; Xu, B.; and Grosse, R. 2018. Three mechanisms of weight decay regularization. *arXiv preprint arXiv:1810.12281*.