

Robustness and Regularization in Hierarchical Re-Basin

Benedikt Franke, Florian Heinrich, Markus Lange
and Arne Raulf

German Aerospace Center (DLR) - Institute for AI Safety and Security
Wilhelm-Runge-Strasse 10, 89081 Ulm - Germany



Abstract. This paper takes a closer look at Git Re-Basin, an interesting new approach to merge trained models. We propose a hierarchical model merging scheme that significantly outperforms the standard *MergeMany* algorithm. With our new algorithm, we find that Re-Basin induces adversarial and perturbation robustness into the merged models, with the effect becoming stronger the more models participate in the hierarchical merging scheme. However, in our experiments Re-Basin induces a much bigger performance drop than reported by the original authors.

1 Introduction

In recent years, a lot of work has been done towards investigating the permutation symmetries of artificial neural networks (ANNs) trained with variants of stochastic gradient descent (SGD) [1, 2, 3, 4, 5]. Hypothesized to be closely related to linear mode connectivity (LMC), permutation invariance allows to change the order of neurons in an ANN without hurting its accuracy [6]. As a direct application of this hypothesis, Git Re-Basin [5] provides a method to merge two models with the same architecture by exploiting permutation invariance to "teleport" one of the models into the same loss basin as the other. In contrast to simple, "naive" interpolation between two models, this technique circumvents accuracy losses that arise when the mean of both models in the parameter space falls outside of a loss basin or local minima. Motivated by these results, we investigate how to effectively merge more than two models. We propose a hierarchical scheme as an alternative to the *MergeMany* algorithm provided by [5] and show that it performs significantly better, even though it likewise does not admit a zero-accuracy barrier. We also show that Git Re-Basin seems to induce adversarial robustness properties as well as regularization, which grow more pronounced the more models are merged with our scheme. This effect already starts to be noticeable with just two input models.

2 Related Work

LMC as first hypothesized in [1] describes the property of SGD solutions to be linearly connected in the loss landscape. While originally investigated in the context of the lottery ticket hypothesis [2], LMC has quickly become of interest to other sub-fields, such as multitask learning [3]. Akash et al. propose

a mathematical framework that can use LMC towards model fusion [4]. General mode connectivity for language models is studied in [7]. Recently, LMC has been generalized to layerwise linear feature connectivity [8], where it is claimed that not only the weight vectors of entire networks, but also single feature maps have a linear connection between differently trained networks. Model Permutation has been linked to LMC [6], which proposes LMC to only hold when taking into account *permutation invariance*, i.e. that there exist permutations of trained networks which yield the same outputs as their originals. Permutation invariance and its application towards model interpolation is studied in [5]. In the context of federated machine learning, Wang et al. show that accuracy can be improved by permuting one of the models so that it closely matches the other weight-wise before averaging [9]. Permutation invariance has also been studied in the context of adversarial attacks on ANNs [10].

3 Our Method

While [5] provides the *MergeMany* algorithm to apply Git Re-Basin to more than 2 models, we found the algorithm to have an important theoretical weakness:

In each round of the algorithm, one of the n input models Θ_i is permuted towards the *mean* $\bar{\Theta}$ of the other $n - 1$ models with $\bar{\Theta} = \frac{1}{n} \sum_{j \in \{1, \dots, n\} \setminus i} \Theta_j$ [5]. However, the mean of these models can not be guaranteed to lie within a loss basin, despite the Re-Basin algorithm assuming all input models to lie within such a basin. On the contrary, [5] shows that in general, linear interpolation between two non-permuted models yields strongly deteriorated accuracy. In Section 4, we show that these theoretical considerations are also empirically observable.

To alleviate the discussed issue, we propose to use a more computationally expensive hierarchical scheme, which is outlined in Figure 2. When trying to merge 2^n models, we merge them pairwise (and disjunctively) over n stages, with each stage taking in the merged models of the previous stages. In the following, we refer as "merging" to the process of applying Re-Basin to permute the second of two models to "teleport" it into the loss basin of the first model, and then doing linear interpolation as specified in [5] with $\lambda = 0.5$. In preliminary experiments, we investigated if it matters which model is chosen as the model that gets permuted, but found it to have no statistically significant influence.

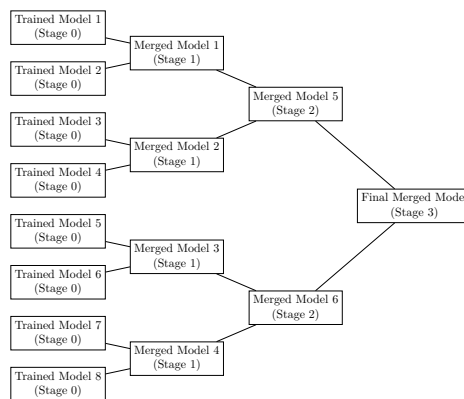



Fig. 2: Our proposed hierarchical Merging Scheme, exemplified for merging eight models.

4 Experiments



(a) Stage 2 (Merging 4 Models into one) (b) Stage 3 (Merging 8 Models into one)

Fig. 3: Distribution of test set accuracies on CIFAR-10 by merging algorithm on 1600 trained input models. "No Merging" denotes the median accuracy of the trained input models and is displayed only as a line as the unmerged models have very little variance to improve the clarity of the figure.

In order to investigate the effects of our approach in a statistically sound way, we train 1600 multi-layer perceptrons (MLPs) on the CIFAR-10 data set. We follow the experiments of [5] so that our MLPs have 4 layers with a hidden dimensions of 512 and a latent layer, which has a hidden dimension of 256. All layers except the last use the *ReLU*-activation and we use a different random seed for each training process. Our experimental code is adapted from a PyTorch-based re-implementation of Re-Basin ¹. As a first test, we compare if our hierarchical Re-Basin yields better accuracy scores than the *MergeMany* algorithm, by merging our 1600 trained models disjunctively into 400 models (Figure 3a) and 200 models (Figure 3b), respectively. From the figures, we see that not only our Hierarchical Re-Basin approach outperforms the *MergeMany* algorithm, but the difference gets stronger with more models merged. This stems from the strong deterioration of accuracy in *MergeMany* when switching from 4 input models to 8 input models. On the contrary, our hierarchical approach does not suffer from the same effect, with only the variance of scores getting slightly larger.

4.1 Robustness and Relationship with Regularization

We further investigate if the merging of models via Re-Basin produces safety properties by testing our models for adversarial robustness on two selected attacks, DeepFool [11] and FGSM [12]. The results are shown in Figure 4. In both cases, we plot the robust accuracy (i.e. the accuracy on adversarial images) over the ϵ -parameter, which controls the attack strength. Figure 4 shows that in both

¹Code taken from <https://github.com/themrmaster/git-re-basin-pytorch> - The Repository is not affiliated with the authors of this paper!

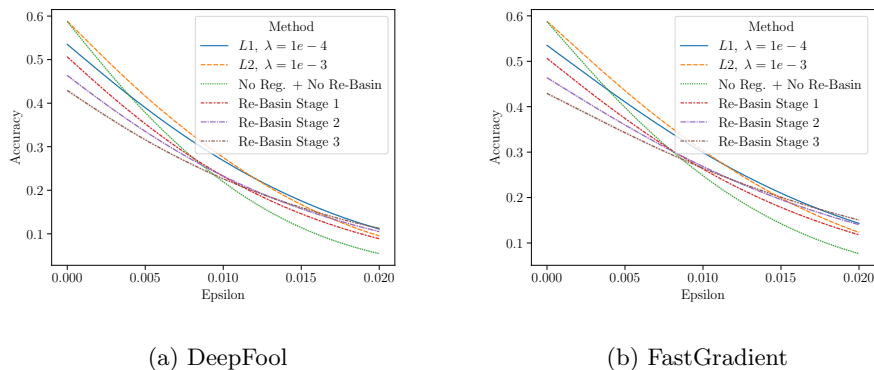


Fig. 4: Mean Accuracy of different Re-Basin stages over attack strength ϵ . As a comparison, models trained with L1 and L2 regularization are also shown. 95% CI is shaded around the curves but is too small to be visible.

cases, all stages of Re-Basin break even with the unmerged models in the area around $\epsilon = 0.01$, with lower stages (i.e. less Re-Basin) being better with weaker attacks and later stages (i.e. more Re-Basin) being more robust against stronger attacks. Overall, merging more models with Re-Basin seems to correlate with stronger adversarial robustness, but comes at the cost of some accuracy. We also conducted the same experiments with L1 and L2 regularized models and plot the results similarly into Figure 4. From the results, we see that L2 regularization seems to generally produce stronger adversarial robustness, only breaking even with the final stage of our multi-stage Re-Basin approach at the far end of the investigated ϵ -range.

4.2 Effects on Model Norm and Lipschitz Constant

With the results on adversarial robustness in mind, we wonder what other properties of the MLPs are positively impacted by Re-Basin. Firstly, we expect that Re-Basin should induce some weight-regularization into the resulting network, as noise-like perturbations on the weights are well-known to encourage smoother learned functions [13]. Interpolation between the weights of two trained networks can be expected to have a noise-like effect. We use a very simple formula to calculate a metric for the weight regularization, $|w| = \sum_{i=0}^N \|W_i\|_F + \|b_i\|_2$, where N is the number of layers in the MLP, W_i is the weight matrix of the i -th layer and b_i is the bias of the i -th layer. Intuitively, this measure produces higher values for less regularized models, as weight regularization commonly punishes high norms in the models parameters (e.g. L1 or L2 regularization [14]). The distributions of this cumulative weight norm $|w|$ are depicted per stage of our Hierarchical Re-Basin in Figure 5a. Our findings support our initial expectation: w drops almost linearly with the stage, while the variance also gets reduced.

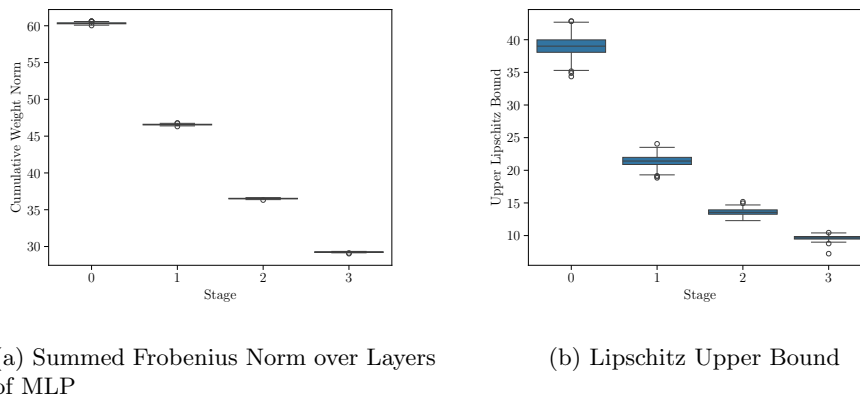


Fig. 5: Impact of Different Re-Basin Stages on Weight Norm and Lipschitz Upper Bound

Secondly, we investigate a well-known safety property of ANNs, the Lipschitz constant. Employing the methodology² presented in [15], we calculate a guaranteed upper bound on the Lipschitz constant of each MLP. This upper bound is an effective measure for the sensitivity of an ANN to small input perturbations, which implies that a more robust network will have a smaller bound. Our results are presented in Figure 5b, where a similar trend to Figure 5a is clearly visible. With each stage of Re-Basin, not only does the upper bound on the Lipschitz constant get smaller, but the variance in the bounds between the MLPs shrinks similarly. As the Lipschitz constant has been linked to perturbation resistance [15], this indicates a higher perturbation resistance induced by Re-Basin.

5 Discussion

In this paper, we presented a Hierarchical Re-Basin scheme and investigated the effects of Re-Basin on adversarial robustness, weight regularization and the Lipschitz constant. We find that Re-Basin seems to act as a sort of regularization, positively impacting adversarial and perturbation robustness in the process. This seems to stem from the interpolation step between two models in the same loss basin, which can be seen as a limited disturbance on the model weights that, by definition, should not move the model to a spot outside of the loss basin. While any sufficiently strong noise on the model weights or data can act as such a regularization, we find that Re-Basin does so without hurting accuracy too much when applied using our proposed hierarchical Re-Basin scheme. On the other hand, using the *MergeMany* algorithm or even naively interpolating between models hurts accuracy so much that any emergent safety-properties be-

²Code taken from https://github.com/HeinrichAD/local_lipschitz/tree/778d7dec56574e43bdff05e8b6e794cb785c9d21

come more or less irrelevant. Our findings contradict [5], which saw good results on their *MergeMany* algorithm, despite the theoretical flaws discussed above. Furthermore, despite using an open-sourced implementation of Re-Basin that was tested by multiple unaffiliated contributors, we were unable to reproduce the zero-accuracy barrier that was shown in the original paper. We hope that the empirical results of this work contribute towards the ongoing discussion around Git Re-Basin and LMC and serve as a stepping stone for further investigations.

References

- [1] Jonathan Frankle et al. Linear mode connectivity and the lottery ticket hypothesis. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020*, volume 119 of *Proceedings of Machine Learning Research*, pages 3259–3269. PMLR, 2020.
- [2] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *7th International Conference on Learning Representations, ICLR 2019*. OpenReview.net, 2019.
- [3] Seyed-Iman Mirzadeh, Mehrdad Farajtabar, Dilan Görür, Razvan Pascanu, and Hassan Ghasemzadeh. Linear mode connectivity in multitask and continual learning. In *9th International Conference on Learning Representations*. OpenReview.net, 2021.
- [4] Aditya Kumar Akash et al. Wasserstein barycenter-based model fusion and linear mode connectivity of neural networks. *arXiv preprint arXiv:2210.06671*, 2022.
- [5] Samuel K. Ainsworth et al. Git re-basin: Merging models modulo permutation symmetries. In *The Eleventh International Conference on Learning Representations, ICLR 2023*. OpenReview.net, 2023.
- [6] Rahim Entezari et al. The role of permutation invariance in linear mode connectivity of neural networks. In *The Tenth International Conference on Learning Representations, ICLR 2022*, 2022.
- [7] Yujia Qin et al. Exploring mode connectivity for pre-trained language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022*, pages 6726–6746. Association for Computational Linguistics, 2022.
- [8] Zhanpeng Zhou et al. Going beyond linear mode connectivity: The layerwise linear feature connectivity. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023*, 2023.
- [9] Hongyi Wang et al. Federated learning with matched averaging. In *8th International Conference on Learning Representations, ICLR 2020*. OpenReview.net, 2020.
- [10] Karan Ganju et al. Property inference attacks on fully connected neural networks using permutation invariant representations. In *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*, pages 619–633, 2018.
- [11] Seyed-Mohsen Moosavi-Dezfooli et al. Deepfool: A simple and accurate method to fool deep neural networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2016.
- [12] Sandy H. Huang et al. Adversarial attacks on neural network policies. In *5th International Conference on Learning Representations, ICLR 2017, Workshop Track Proceedings*. OpenReview.net, 2017.
- [13] Guozhong An. The effects of adding noise during backpropagation training on a generalization performance. *Neural Computation*, 8(3):643–674, 1996.
- [14] Andrew Y Ng. Feature selection, l1 vs. l2 regularization, and rotational invariance. In *Proceedings of the twenty-first international conference on Machine learning*, page 78, 2004.
- [15] Trevor Avant and Kristi A Morgansen. Analytical bounds on the local lipschitz constants of relu networks. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.