

# Will it Merge? On The Causes of Model Mergeability

Adir Rahamim<sup>1</sup>, Asaf Yehudai<sup>2,5</sup>, Boaz Carmeli<sup>1,2</sup>, Leshem Choshen<sup>3,4</sup>,  
Yosi Mass<sup>2</sup>, Yonatan Belinkov<sup>1,6</sup>

<sup>1</sup>Technion - Israel Institute of Technology, <sup>2</sup>IBM Research AI, <sup>3</sup>MIT, <sup>4</sup>MIT-IBM Watson AI Lab,  
<sup>5</sup>Hebrew University of Jerusalem, <sup>6</sup>Kempner Institute, Harvard University  
[adir.rahamim@campus.technion.ac.il](mailto:adir.rahamim@campus.technion.ac.il)

## Abstract

Model merging has emerged as a promising technique for combining multiple fine-tuned models into a single multitask model without retraining. However, the factors that determine whether merging will succeed or fail remain poorly understood. In this work, we investigate why specific models are merged better than others. To do so, we propose a concrete, measurable definition of mergeability. We investigate several potential causes for high or low mergeability, highlighting the base model knowledge as a dominant factor: Models fine-tuned on instances that the base model knows better are more mergeable than models fine-tuned on instances that the base model struggles with. Based on our mergeability definition, we explore a simple weighted merging technique that better preserves weak knowledge in the base model.

## 1 Introduction

Large pre-trained models are commonly fine-tuned on various downstream tasks to achieve better specialization on specific tasks. This task-specific model training has motivated model merging techniques (Matena and Raffel, 2022; Wortsman et al., 2022; Choshen et al., 2022; Ilharco et al., 2023; Yu et al., 2024; Yadav et al., 2023; Stoica et al., 2025), which aggregate the weights of multiple fine-tuned models into a single expert model that should perform well on all tasks. However, while algorithmic advances have improved merging algorithms, a fundamental question remains: What determines whether merging of a fine-tuned model will be successful? Answering this question might assist us in understanding how to obtain more mergeable models and inform better merging algorithms.

To address this question, we define the notion of *mergeability*: a property of model updates that captures how well they retain trained knowledge when merged with other model updates. We find that

not all model updates have the same mergeability—there is a wide spectrum of mergeability among models. We use the term mergeability score to quantify this property.

This work investigates possible sources for having high or low mergeability. Figure 1 describes our experimental setup. Upon obtaining model updates (we use LoRA adapters (Hu et al., 2022) in our experiments), we calculate each model update’s mergeability score (the degree to which the knowledge encoded in a given model update is preserved when it is merged with other model updates. For a given model update, we merge it with a subset of randomly sampled other model updates and evaluate the merged model on the given model update task. We repeat this multiple times and measure the average performance; further described in §2). We then group model updates by their mergeability score and investigate different causes for high mergeability: general domain knowledge of the base model, weight properties, and specific task knowledge of the base model.

Our study spans two experimental setups: example-level mergeability using the PopQA dataset (Mallen et al., 2023) and task-level mergeability using the Lots-of-LoRAs collection (Brüel-Gabrielsson et al., 2025). In the PopQA dataset, we find that a low probability gap between the top predicted answer and the correct answer in the base model correlates with a higher mergeability. This suggests that the base model’s prior knowledge about a question strongly influences whether fine-tuning on knowledge related to this question will merge better. Consistently, in Lots-of-LoRAs, we observe that tasks where the base model initially performed well suffer less performance degradation from merging, whereas tasks on which the base model had lower initial accuracy suffer from larger drops in performance after merging. Our findings indicate that the mergeability of fine-tuned models is closely linked to the base model’s ini-

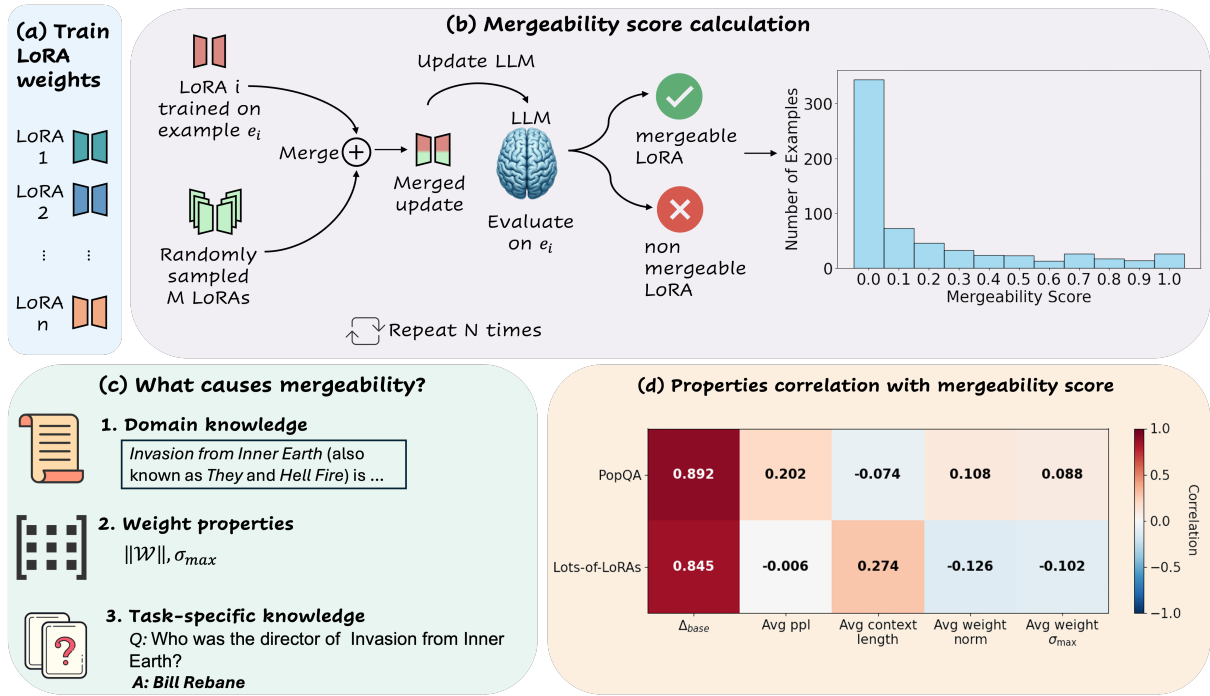


Figure 1: Our Experimental setup. The figure shows an example from PopQA, Lots-of-LoRAs experiments follow a similar process. (a) We train an expert model for each example, and verify the model’s correctness on the example after training. (b) We calculate each example’s mergeability score, as defined in §2 and group examples by their mergeability score. (c) We investigate different traits that affect mergeability. Results are in §4. (d) We examine the correlation between the mergeability score and the evaluated properties. Among these, base model knowledge ( $\Delta_{base}$ ) exhibits the strongest correlation.

tial performance on the corresponding fine-tuning data. Other possible causes, namely, general domain knowledge or structural weight properties, do not correlate well with mergeability scores.

In further analyses, we discover that mergeability is primarily a local trait of the model update and does not depend on the merged set. We show that when merging a highly mergeable model update with model updates from other mergeability groups, the highly mergeable model update remains stable regardless of the partner group we merge with.

Finally, as a proof-of-concept application of our insights, we propose a simple merging technique that incorporates the base model’s performance in a weighted mean merging. We show that this technique improves the retention of weak-performing tasks with little or no degradation of strong-performing tasks.

Our contributions in this paper are threefold:

- We show the existence of mergeability and propose a concrete, measurable definition of it.
- We provide empirical evidence that the base model’s prior knowledge is a key predictor of the mergeability of fine-tuned weights. To our

knowledge, this is the first study to directly link pre-training knowledge with mergeability.

- We show an application of our findings and suggest merging weights with awareness of the base model performance.

## 2 Mergeability

*Mergeability* is a trait of model updates describing the degree to which the knowledge encoded in a model update is preserved when merged with other model updates. This falls into a larger context of what allows successful merging (e.g., high-dimensionality and a shared base model; §7), but is unique in recognizing the effect of the model updates themselves. In this section, we define and empirically show the existence of mergeability.

**Mergeability score.** Given a model update  $\theta_{\Delta}$  with corresponding input-output pair  $(x, y)$  and a distribution  $\mathcal{D}$  over sets of other updates, we define the mergeability score  $S$  of  $\theta_{\Delta}$  as:

$$S(\theta_{\Delta}) = \mathbb{E}_{\{\theta_{\Delta_j}\} \sim \mathcal{D}} [f(\theta + \mathcal{M}(\{\theta_{\Delta}\} \cup \{\theta_{\Delta_j}\}); x, y)] \quad (1)$$

where  $f(\theta; x, y)$  is a scoring function that evaluates a model with parameters  $\theta$  on a set of input-output

data pairs  $(x, y)$  and  $\mathcal{M}$  is a merging algorithm that combines a set of updates into a single update.

To estimate the mergeability score  $S(\theta_\Delta)$ , we conduct  $N$  trials. In each trial  $i \in \{1, \dots, N\}$ , we sample  $M$  other model updates  $\{\theta_{\Delta_m}\}_{m=1}^M$  (without replacement) from the global pool and merge  $\{\theta_\Delta\} \cup \{\theta_{\Delta_m}\}_{m=1}^M$ , yielding a merged model update  $\theta_\Delta^{(i)} = \mathcal{M}(\{\theta_\Delta\} \cup \{\theta_{\Delta_m}\}_{m=1}^M)$ . We then update the base model with  $\theta_\Delta^{(i)}$  and re-evaluate, obtaining the following empirical mergeability score:

$$S(\theta_{\Delta e}) = \frac{1}{N} \sum_{i=1}^N f(\theta + \theta_\Delta^{(i)}; x, y) \quad (2)$$

This score captures the robustness of an example’s knowledge under merging, where higher values indicate stronger mergeability with other model updates. The scoring function  $f$  depends on the setup, as detailed in the next section.

In practice, not all model updates are equally stable under merging: some integrate seamlessly into a merged model, while others tend to interfere or degrade. Figure 2 shows the mergeability distribution of Llama (Dubey et al., 2024) on the PopQA dataset (more distributions are in Appendix A.3). The blue bars show the mergeability score as empirically calculated in our experiments. The red bars show the distribution as if the scores were distributed randomly as a binomial distribution with a success probability of  $P = \frac{\text{\# of success merges}}{\text{Total \# of merges}}$ , where we consider a merge as succeeded if it retained the model update information after merging. The difference in the distributions shows the existence of non-trivial mergeability.

**Research questions.** After establishing that mergeability is a non-trivial property, we investigate several possible reasons for mergeability:

- Base model specific task knowledge (§4.1).
- Weight properties (§4.2).
- Base model general domain knowledge (§4.3).

Then, in §5 we check if mergeability is a local trait of the model update, or rather depends on the group of model updates it is merged with.

### 3 Experimental Setup

We evaluate the mergeability of LoRA adapters (Hu et al., 2022) as recent work showed that such merging is very powerful (Stoica et al., 2025). We focus on two settings: (i) entity-centric question answering with PopQA (Mallen et al., 2023) dataset, and (ii) a broad collection of LoRA adapters from




Figure 2: Mergeability score distribution of Llama-3.2-3B on the PopQA dataset. Blue wide bars show the mergeability score as empirically calculated. Red thin bars show the baseline distribution if mergeability were not a model trait, modeled as a binomial distribution with a fixed success rate.

*Lots-of-LoRAs* collection (Brüel-Gabrielsson et al., 2025). The first corresponds to example-level mergeability, merging adapters capturing single data points, while the second corresponds to task-level mergeability. Unless otherwise specified, all merges are performed with KNOTS (Stoica et al., 2025) merging algorithm, which is specifically designed for merging LoRA adapters.<sup>1</sup> In each of the two setups, we describe our instantiations of the mergeability score and the evaluation protocol;<sup>2</sup>

#### 3.1 Example-Level Mergeability: PopQA

PopQA is an open-domain, entity-centric QA benchmark (Mallen et al., 2023), making it suitable for studying example-level mergeability. In this case, the scoring function  $f$  is chosen as the binary correctness of the model:

$$f(\theta; x, y) = \mathbb{1}\{\hat{y} = y\}, \quad (3)$$

where  $\hat{y}$  is the model prediction.

For controlled probabilistic evaluation, we convert it to a multiple-choice format: for each question with gold answer  $y$ , we sample  $n$  additional candidates from answers to other PopQA questions (without replacement). To reduce spurious ambiguity, we discard distractors that are string-identical to  $y$  or that are near-duplicates by normalization (case/punctuation stripping). We use  $n=7$  (and obtain 8-option multiple-choice questions).

<sup>1</sup>We report analyses results with other merging algorithms in Appendix A.9, finding mostly consistent patterns.

<sup>2</sup>Additional experiments hyperparameter details are available at Appendix A.6

**Models.** We use two large pretrained language models from distinct families: Llama-3.2-3B (Dubey et al., 2024) and Qwen-2.5-3B (Hui et al., 2024). This choice allows us to compare models from different families, providing a more robust testbed for mergeability. Results in the main paper are reported for the Llama-3.2-3B model; additional analyses, including results on Qwen-2.5-3B, are provided in Appendix A.1.

**Evaluation.** We evaluate with  $k$ -shot prompting (Brown et al., 2020). In our experiments, we used  $k = 4$ . Given the prompt and a set of options  $\{y_j\}_{j=1}^{n+1}$ , with the correct answer  $y$  among them, we compute the length-normalized conditional log-probability assigned by the model to each option:

$$\text{score}(y_j) = \frac{1}{|y_j|} \sum_{t=1}^{|y_j|} \log p_{\theta}((y_j)_t | \text{prompt}, (y_j)_{<t}) \quad (4)$$

We predict  $\hat{y} = \arg \max_j \text{score}(y_j)$ . Normalization mitigates length bias across options.

**Training Procedure.** To focus on model updates that change the model’s knowledge, we first filter out questions already answered correctly by the base model under the above evaluation protocol. For the remaining (incorrect) questions, we retrieve the relevant entity’s Wikipedia page. We then fine-tune the model with LoRA on these entity-related passages to teach the base model relevant information on the questioned entity. We retain only those examples for which the post-finetuning model answers correctly. This yields a collection of example-specific LoRA adapters that each corrects a distinct factual error of the base model. Out of 1931 examples we tested, 1107 were answered correctly without training, which left us with 824 examples we trained (42.67%). Out of them, for 639 (77.55%) examples the correct answer had the highest probability after training. Additional fine-tuning details are provided in Appendix A.6. We also report experimental results with additional LoRA ranks (Appendix A.7) and full finetuning (Appendix A.8). Trends are consistent with the main paper results.

### 3.2 Task-Level Mergeability: Lots-of-LoRAs

To study task-level mergeability, we use the Lots-of-LoRAs benchmark (Brüel-Gabrielsson et al., 2025), a large-scale collection of LoRA adapters trained on diverse NLP tasks. In this case, the scor-

ing function  $f$  is computed as the post-merging task accuracy:

$$f(\theta; x, y) = \text{Acc}_{\text{merged}}^{(i)}(x, y) \quad (5)$$

**Model.** We use the same base model as in Brüel-Gabrielsson et al. (2025), Mistral-7B-Instruct-v0.2 (Jiang et al., 2023).

**Evaluation.** For each task, we use the test prompts provided by the dataset. We evaluate by exact match (EM) at the example level and aggregate to task accuracy. We evaluate on each task’s test set. To control for ceiling effects, we restrict evaluation to tasks where the finetuned model achieves at least 99% accuracy, ensuring that mergeability is measured between adapters that individually solve their target tasks near perfectly. This left us with a total of 81 tasks (We provide experimental results for other accuracy thresholds in Appendix A.4, trends are consistent across all thresholds).

## 4 Causes of Mergeability

In this section, we investigate possible causes for high or low mergeability. In each case, we bin models by their mergeability score and correlate these scores with metrics reflecting possible causes.

### 4.1 What is the effect of the base model’s knowledge of the task?

In this research question, we wish to study how the base model’s prior knowledge of an example/task impacts mergeability.

In example-level PopQA analysis, Figure A.4 reports the average rank of the correct answer under the base model. A higher rank (closer to zero) indicates that the base model assigns the correct answer a better rank. While we generally observe that higher-mergeability examples correspond to lower ranks, we also observe a non-monotonic jump for  $S = 1.0$ .

We further analyze the probability gap between the most likely answer and the correct answer in the base model  $\Delta_{\text{base}} = p_{\text{max}}^{\text{base}} - p_{\text{correct}}^{\text{base}}$  (Figure 3). The gap decreases with mergeability, implying that examples requiring only a small adjustment to the decision boundary are more mergeable. Moreover, when we compare base vs. post-training gaps  $\Delta_{\text{trained}} = p_{\text{correct}}^{\text{trained}} - p_{\text{correct}}^{\text{base}}$  (Figure 3), we find the opposite trend: high-mergeability examples achieve larger post-training probability improvements. This indicates that examples that were “eas-




Figure 3: PopQA different properties correlation with mergeability score results. Values are mean values normalized relative to  $S = 0.0$ .  $\Delta_{\text{base}}$  trend shows that the gap decreases with mergeability, implying that examples with better base model knowledge are more mergeable. Additionally,  $\Delta_{\text{trained}}$  shows that high-mergeability examples achieve larger post-training probability improvements, although they were easier to “fix”. Conversely, we observe no clear trend between mergeability score and training data difficulty (average perplexity and average context length) or weight-level properties (average weight norm and average highest singular value). The shaded regions represent the standard error of the values.

ier” to fix also produce more stable model updates under merging.

Lots-of-LoRAs analysis reveals similar trends on the task-level merging. As Figure 4 shows, higher mergeability scores have higher average base model accuracy. In other words, tasks with better base model knowledge of the task have better mergeability scores.

Overall, we conclude that **better knowledge of the task or fine-tuning data in the base model corresponds to higher mergeability**.

#### 4.2 Does training data difficulty affect mergeability?

To examine whether mergeability is also affected by training data difficulty, we consider two proxies: base model perplexity and length (token count) of the training context.

As Figure 3 shows, there are no clear trends w.r.t mergeability score in the case of example-level mergeability in PopQA. One exception is a high base model perplexity on examples with the highest mergeability score ( $S = 1.0$ ). This suggests that model updates correcting knowledge gaps in regions where the base model is highly uncertain are more robust to merging. However, overall perplexity and context length are poor predictors of mergeability in this case.

In contrast, the Lots-of-LoRAs analysis reveals different trends. As Figure A.5 shows, higher-mergeability examples correspond to lower perplexities. Moreover, Figure A.6 shows that longer




Figure 4: Lots-of-LoRAs average base model task accuracy of different mergeability scores. Higher mergeability scores have on average a higher base model accuracy.

training contexts have higher mergeability scores.

However, a difference in the two settings must be noted. While PopQA training data are general passages on an entity, with a different format from the evaluated questions, Lots-of-LoRAs training data is highly related to the evaluation data, sharing the same task format. This suggests that we can observe training data effects on mergeability only when it is aligned with the evaluation data.

#### 4.3 Do weight-level properties correlate with mergeability?

We next examine whether structural properties of the learned weight updates correlate with mergeability. We compute two structural metrics on the effective update matrix  $\Delta W = BA$  (where  $A, B$

Mergeability Score	Avg Weight Norm	Avg Weight $\sigma_{\max}$
0.0-0.2	1.15	0.78
0.2-0.4	0.57	0.40
0.4-0.6	0.73	0.52
0.6-0.8	0.65	0.46
0.8-1.0	0.66	0.46

Table 1: Lots-of-LoRAs results for weight properties. We do not observe a clear correlation between weight properties and the mergeability score. However, we observe a notable distinction between extremely low mergeability examples ( $S \in [0.0, 0.2)$ ) and all other scores ( $S \geq 0.2$ ). Examples in the lowest bin exhibit significantly higher average norms and singular values.

are LoRA update matrices): (i) Frobenius norm  $\|\Delta W\|_F$ , previously shown to influence mergeability (Pari et al., 2025; Horoi et al., 2025) and (ii) highest singular value  $\sigma_{\max}(\Delta W)$ , which captures dominant update directions (Stoica et al., 2025).

In the PopQA setting (Figure 3), we observe that the highest mergeability examples tend to have slightly higher norms and higher  $\sigma_{\max}$ . However, overall the correlations to mergeability are close to zero (0.10 and 0.09 Spearman’s). The Lots-of-LoRAs analysis (Table 1) again shows no consistent trend across mergeability scores. However, we observe a notable distinction between extremely low mergeability examples ( $S \in [0.0, 0.2)$ ) and all other scores ( $S \geq 0.2$ ). Examples in the lowest bin exhibit significantly higher average norms and singular values.

The difference in trends between PopQA and Lots-of-LoRAs might be explained by the difference in training setting - single model parameter training (mlp up proj) and single layer training in PopQA versus multiple model parameters (attention Q, K, and V matrices) and multiple layer training in Lots-of-LoRAs. However, in both experiments **we do not observe a clear correlation between weight properties and mergeability score.**

## 5 Other Mergeability Properties

Having examined potential causes of mergeability, we next turn to complementary aspects of this phenomenon: properties that characterize how mergeability behaves under different conditions.

### 5.1 Is mergeability local or global?

Previous work has shown that merging is affected by the merge set, and more specifically, by the shared knowledge between tasks (Zaman et al.,




Figure 5: The mergeability score when merging weights  $\theta_{\Delta}$  with mergeability score  $S(\theta_{\Delta}) = 1.0$  with weights drawn from different mergeability scores. The blue line, which represents examples with a fixed score  $S(\theta_{\Delta}) = 1.0$ , shows near-constant accuracy across conditions, while the orange line (other mergeability group examples) improves with their own mergeability score. The shaded regions represent the standard error of the accuracy.

2024). However, what happens when no shared knowledge exists? To test whether mergeability is an intrinsic property of a model update or depends on the merge set, we conduct a controlled experiment under the PopQA setting. We fix a set of highly mergeable model updates ( $S(\theta_{\Delta}) = 1.0$ ) and merge them with groups drawn from bins of varying mergeability scores. If mergeability is a local trait depending only on the target model update and not on the merge set, then performance on the fixed  $S(\theta_{\Delta}) = 1.0$  updates should remain stable regardless of the partner group. Figure 5 confirms this hypothesis: blue bars (fixed  $S(\theta_{\Delta}) = 1.0$  updates) show near-constant accuracy across conditions, while orange bars (other mergeability group updates) improve with their own mergeability score.

This finding suggests that mergeability is primarily an intrinsic property of the LoRA update itself rather than an emergent property of the merge set.

### 5.2 How merging algorithm affects mergeability?

Our experiments primarily employed KNOTS (Stoica et al., 2025), the current state-of-the-art algorithm for merging LoRA weights. To assess how the choice of merging algorithm influences mergeability, we compare KNOTS with two other merging algorithms: TIES (Yadav et al., 2023) and simple mean averaging.

Figure 6 presents the mergeability score distribution for the Qwen model on PopQA using




Figure 6: Mergeability score distribution for the Qwen model on PopQA using different merging algorithms (log scale). As the method is ‘weaker’, more examples are in the higher bins.

these merging algorithms. A clear trend emerges: “weaker” algorithms exhibit more examples in the higher mergeability bins. Specifically, TIES yields more examples with scores  $S(\theta_\Delta) \geq 0.8$  than KNOTS, while mean averaging produces the largest number of examples at  $S(\theta_\Delta) = 1.0$ .

This pattern reflects the degree of interference resolution. Mean averaging, the weakest method, performs no conflict mitigation—weights tend to either merge or fail. TIES introduces global sign conflict resolution, and KNOTS further aligns update subspaces, improving overall merging. However, these interventions slightly reduce the proportion of perfectly mergeable examples, suggesting a trade-off between resolving interference and preserving higher mergeability.

## 6 Mergeability Score for Better Merging

Our analysis of mergeability revealed a consistent relationship between a base model’s task accuracy and the corresponding model update mergeability (§4). Specifically, model updates corresponding to tasks where the base model already performs well tend to have a better mergeability than model updates of tasks with lower base model accuracy. This suggests that naive averaging of adapter may overemphasize high-accuracy tasks, while underweighting tasks that have lower mergeability.

To address this imbalance, we propose a weighted averaging strategy where the contribution of each model update is inversely proportional to the base model’s accuracy on the corresponding task. The intuition is to assign a higher weight to model updates that have lower mergeability, while limiting the influence of model updates on tasks for which the base model already works well.

Let  $\{\Theta_{\Delta 1}, \Theta_{\Delta 2}, \dots, \Theta_{\Delta T}\}$  be a set of model up-




Figure 7: Simple mean merging versus weighted mean merging based on base model accuracy. Weighted average merging (orange bars), low-accuracy tasks (tasks 0 and 1) retained more of their fine-tuned performance, while high-accuracy tasks experienced minimal degradation (task 3) or no degradation at all (task 4).

dates corresponding to tasks  $\{t_1, t_2, \dots, t_T\}$ . For each task  $t_i$ , we compute the base model accuracy  $Acc(t_i) \in [0, 1]$ . We then define the inverse accuracy score as  $s_i = 1 - Acc(t_i)$ . To convert these scores into weights, we apply a softmax function with temperature  $\tau$ :  $w_i = \frac{\exp(s_i/\tau)}{\sum_{j=1}^T \exp(s_j/\tau)}$ . The final merged adapter is computed as a weighted sum:

$$L_{\text{merged}} = \sum_{i=1}^T w_i \cdot \Theta_{\Delta i} \quad (6)$$

This approach ensures that tasks where the base model performs poorly (i.e., high  $s_i$ ) receive higher weights, while tasks with a high base accuracy are down-weighted. The temperature parameter  $\tau$  controls the sharpness of the weighting distribution: lower values of  $\tau$  emphasize differences more strongly.

To evaluate our method, we used the Lots-of-LoRAs collection. We sampled 2 random tasks with a low base model accuracy and 2 tasks with a high base model accuracy, and merged their model updates using a weighted average as described above. As shown in Figure 7, in the weighted average experiment (orange bars), low-accuracy tasks (tasks 0 and 1) retained more of their fine-tuned performance, while high accuracy tasks experienced minimal degradation (task 3) or no degradation at all (task 4). In regular averaging (blue bars), tasks with lower base model accuracy experienced a higher performance degradation.

## 7 Related Work

**Model merging.** Model merging has emerged as a technique for combining fine-tuned models (each expert on a different task) into a single, multitask model without requiring additional training. Direct approach is a simple average of the weights.

Wortsman et al. (2022) proposed Model Soups, finding that averaging the weights of multiple fine-tuned variants can improve accuracy and out-of-distribution robustness over the individual models.

Another line of work addresses direct conflicts between model parameters. Yadav et al. (2023) propose TIES-Merging, which identifies and resolves conflicting parameter updates (differences in sign or scale for the same weight in different models) prior to merging. By correcting such sign mismatches and scaling disparities, their method reduces destructive interference, leading to merged models that suffer smaller drops in accuracy. In the context of parameter-efficient fine-tuning, Stoica et al. (2025) focuses on merging LoRA adapter weights. They found that averaging LoRA deltas is challenging due to misaligned update subspaces, and introduced an SVD-based alignment (KNOTS) to transform each model’s adapter weights into a common basis before merging. This alignment significantly improved the compatibility of fine-tuned LoRA weights. Ilharco et al. (2023) explores model merging using task vectors. They define a task vector as the difference between a fine-tuned model’s weights and the original base model’s weights. These task vectors can then be added, subtracted, or scaled.

**What helps mergeability?** Prior work has begun to explore this phenomenon. Zaman et al. (2024) showed that unshared knowledge between tasks is retained less during merging. They observe that when models are merged, information that was learned by all the experts is usually preserved in the fused model, whereas information unique to a single model is prone to being overwritten or forgotten. The importance of aligned representations and update directions has also been noted. Some works looked at the effect of the chosen base model on mergeability. Yadav et al. (2025) and He et al. (2025) show that it is easier to merge bigger and stronger base models. Other works considered how the process of training itself affects merging, suggesting that more training, which shows larger norm updates (Gueta et al., 2023) leads to worse merging performance (Pari et al., 2025; Horoi et al., 2025). Stoica et al. (2025) highlighted that merging LoRA fine-tunings can fail when the two models’ updates reside in different singular directions, effectively lacking a shared basis. Ainsworth et al. (2023); Jordan et al. (2023) similarly point out that two networks might need a permutation alignment

to correspond to the same functions. These works suggest that when models make fundamentally different internal choices, naive merging will cause interference.

While all mentioned works at least implicitly point to what makes a good merging outcome, we are the first to notice specific models are merged better even under similar training conditions and as an outcome, the first to demonstrate reasons for that phenomenon.

**Affinity outside model merging.** In fields such as multitask learning (Bingel and Søgaard, 2017; Kim et al., 2023), continual and intermediate training (Poth et al., 2021), many works have studied which tasks aid each other or, in general, which curriculum is beneficial (Hacohen and Weinshall, 2019; Shrivastava et al., 2016). There, a common theme is discussing how a specific task might help another (e.g., a low-resource language improved by training on English Bansal et al., 2018). Fewer works point at universality, as we do, where specific models are inherently good. Those include works on fine-tuning (Choshen et al., 2023) and reinforcement learning (Guo et al., 2025), but perhaps most known is pretraining itself, which is shown to be a strong start for many counterparts (Devlin et al., 2019; Aryabumi et al., 2024).

## 8 Conclusions

In this work, we show the existence of and analyze *mergeability* - a property of the model updates that quantifies the robustness of model updates when merged with other model updates. Our analysis across both example-level and task-level settings, demonstrates that mergeability is not uniformly distributed: some model updates have a better mergeability than others. We raised different key possible causes for mergeability: base model general domain knowledge, weight properties, and base model specific task knowledge. To investigate this property, we define a *mergeability score*, which allows us to measure mergeability and study what affects it. We find that base model task knowledge is correlative with mergeability – instances with higher knowledge in the base model are more mergeable. Moreover, we find evidence that mergeability is an intrinsic property of the model update itself, and not a property of the merging set. Finally, we illustrated how mergeability scores can guide improved merging strategies, mitigating imbalances between tasks with differing base model familiarity.

## Limitations

While this work is based on an extensive set of experiments, several limitations are worth noting and can be addressed in future research. First, our approach can be readily extended to evaluate additional datasets, tasks, and base models to further verify the consistency of our core findings. Second, although our analysis focuses on merging LoRA adapters, it can naturally be extended to other types of tuned models, from full-model fine-tuning to alternative adapter architectures and even local model updates. Third, while we use KNOTS as our primary merging algorithm, other approaches such as Ties and Mean merit further exploration.

## Ethical considerations

Our work adds to the body of literature on model merging and might help develop better merging algorithms. We do not foresee major risks associated with this work. However, a malicious actor might use our analysis to better understand how to amplify unwanted behaviours during model merging.

## Acknowledgments

This research was supported by an Azrieli Foundation Early Career Faculty Fellowship, Open Philanthropy, and by an IBM-Technion Research Collaboration. This research was funded by the European Union (ERC, Control-LM, 101165402). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

## References

- Samuel K Ainsworth, Jonathan Hayase, and Siddhartha Srinivasa. 2023. Git re-basin: Merging models modulo permutation symmetries. In *The Eleventh International Conference on Learning Representations*.
- Viraat Aryabumi, Yixuan Su, Raymond Ma, Adrien Morisot, Ivan Zhang, Acyr Locatelli, Marzieh Fadaee, Ahmet Üstün, and Sara Hooker. 2024. To code or not to code? exploring impact of code in pre-training. In *The Thirteenth International Conference on Learning Representations*.
- Sameer Bansal, H. Kamper, Karen Livescu, Adam Lopez, and S. Goldwater. 2018. [Pre-training on high-resource speech recognition improves low-resource](#)

[speech-to-text translation](#). In *North American Chapter of the Association for Computational Linguistics*.

- Joachim Bingel and Anders Søgaard. 2017. [Identifying beneficial task relations for multi-task learning in deep neural networks](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 164–169, Valencia, Spain. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Rickard Brühl-Gabrielsson, Jiacheng Zhu, Onkar Bhardwaj, Leshem Choshen, Kristjan Greenewald, Mikhail Yurochkin, and Justin Solomon. 2025. Compress then serve: Serving thousands of lora adapters with little overhead. In *Forty-second International Conference on Machine Learning*.
- Leshem Choshen, Elad Venezian, Shachar Don-Yehiya, Noam Slonim, and Yoav Katz. 2023. [Where to start? analyzing the potential value of intermediate models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1446–1470, Singapore. Association for Computational Linguistics.
- Leshem Choshen, Elad Venezian, Noam Slonim, and Yoav Katz. 2022. Fusing finetuned models for better pretraining. *arXiv preprint arXiv:2204.03044*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407.
- Almog Gueta, Elad Venezian, Colin Raffel, Noam Slonim, Yoav Katz, and Leshem Choshen. 2023. [Knowledge is a region in weight space for fine-tuned language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1350–1370, Singapore. Association for Computational Linguistics.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

- Guy Hacoen and Daphna Weinshall. 2019. On the power of curriculum learning in training deep networks. In *International conference on machine learning*, pages 2535–2544. PMLR.
- Yifei He, Siqi Zeng, Yuzheng Hu, Rui Yang, Tong Zhang, and Han Zhao. 2025. Mergebench: A benchmark for merging domain-specialized llms. *arXiv preprint arXiv:2505.10833*.
- Stefan Horoi, Guy Wolf, Eugene Belilovsky, and Gintare Karolina Dziugaite. 2025. Less is more: Undertraining experts improves model upcycling. *arXiv preprint arXiv:2506.14126*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, and 1 others. 2024. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*.
- Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hananeh Hajishirzi, and Ali Farhadi. 2023. Editing models with task arithmetic. In *The Eleventh International Conference on Learning Representations*.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. *Mistral 7b*. Preprint, arXiv:2310.06825.
- Keller Jordan, Hanie Sedghi, Olga Saukh, Rahim Entezari, and Behnam Neyshabur. 2023. Repair: Renormalizing permuted activations for interpolation repair. In *ICLR*.
- Joongwon Kim, Akari Asai, Gabriel Ilharco, and Hananeh Hajishirzi. 2023. *TaskWeb: Selecting better source tasks for multi-task NLP*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11032–11052, Singapore. Association for Computational Linguistics.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hananeh Hajishirzi. 2023. *When not to trust language models: Investigating effectiveness of parametric and non-parametric memories*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822, Toronto, Canada. Association for Computational Linguistics.
- Michael S Matena and Colin A Raffel. 2022. Merging models with fisher-weighted averaging. *Advances in Neural Information Processing Systems*, 35:17703–17716.
- Jyothish Pari, Samy Jelassi, and Pulkit Agrawal. 2025. Collective model intelligence requires compatible specialization. In *ICLR 2025 Workshop on Modularity for Collaborative, Decentralized, and Continual Deep Learning*.
- Clifton Poth, Jonas Pfeiffer, Andreas R  ckl  , and Iryna Gurevych. 2021. *What to pre-train on? Efficient intermediate task selection*. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10585–10605, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. 2016. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 761–769.
- George Stoica, Pratik Ramesh, Boglarka Ecsedi, Leshem Choshen, and Judy Hoffman. 2025. Model merging with svd to tie the knots. In *The Thirteenth International Conference on Learning Representations*.
- Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and 1 others. 2022. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International conference on machine learning*, pages 23965–23998. PMLR.
- Prateek Yadav, Derek Tam, Leshem Choshen, Colin A Raffel, and Mohit Bansal. 2023. Ties-merging: Resolving interference when merging models. *Advances in Neural Information Processing Systems*, 36:7093–7115.
- Prateek Yadav, Tu Vu, Jonathan Lai, Alexandra Chronopoulou, Manaal Faruqui, Mohit Bansal, and Tsendsuren Munkhdalai. 2025. *What matters for model merging at scale?* *Transactions on Machine Learning Research*.
- Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin Li. 2024. Language models are super mario: Absorbing abilities from homologous models as a free lunch. In *Forty-first International Conference on Machine Learning*.
- Kerem Zaman, Leshem Choshen, and Shashank Srivastava. 2024. Fuse to forget: Bias reduction and selective memorization through model fusion. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18763–18783.

## A Appendix

### A.1 Additional Results for §4

In this section, we include additional results of experiments to test the causes of mergeability (§4). Figures A.1 and A.2 show a bar plot of  $\Delta_{\text{base}}$  and  $\Delta_{\text{trained}}$  results from Figure 3, respectively. Table A.1 summarizes the PopQA results for weight properties and general domain knowledge from Figure 3. Figure A.3 shows the average base model (Llama 3.2) probability of the correct answer. Figure A.4 shows the average rank of the correct answer in the base model (Llama 3.2). Figures A.5 and A.6 report the Lots-of-LoRAs experimental result concerning weight properties (§4.2).




Figure A.1: PopQA average difference between the highest and the correct answer probability in the base model. We observe that the gap decreases with mergeability, implying that examples with better base model knowledge are more mergeable.




Figure A.2: PopQA average difference between post-training and original model answer probability for different mergeability scores. High-mergeability examples achieve larger post-training probability improvements, although they were easier to “fix”.




Figure A.3: For the PopQA multiple-choice setting, with the Llama model, we report the average probability of the correct answer in the base model.




Figure A.4: For the PopQA multiple-choice setting, with the Llama model, the figure shows the average rank (lower is better) of the correct answer in the base model of different mergeability scores. While we generally observe that higher-mergeability examples correspond to lower ranks, we also observe a non-monotonic jump for  $S = 1.0$ .

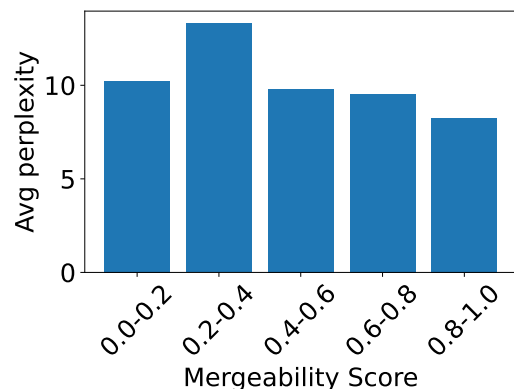


Figure A.5: Lots-of-LoRAs train data perplexity of different mergeability scores. We observe a global trend where higher mergeability examples have a lower average perplexity.

Mergeability score	0.0	0.2	0.4	0.6	0.8	1.0
Avg perplexity	9.80	9.71	9.61	9.26	9.34	10.40
Avg context length (# of tokens)	704.36	649.25	567.84	765.77	560.33	695.59
Avg weight norm	14.91	15.11	14.90	14.81	14.67	15.28
Avg weight $\sigma_{\max}$	18.73	19.13	18.83	18.29	18.33	19.47

Table A.1: PopQA results for weight properties and general domain knowledge. There is no clear trend between mergeability score and training data difficulty (average perplexity and average context length) or weight-level properties (average weight norm and average highest singular value).

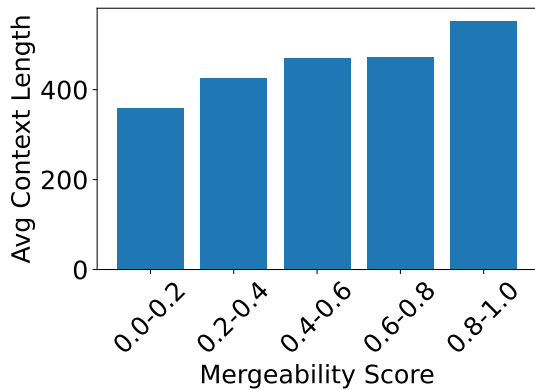


Figure A.6: Lots-of-LoRAs train data length (number of tokens) of different mergeability scores. Higher mergeability examples have longer training data on average.




Figure A.7: PopQA average difference between the highest and the correct answer probability in the base model (Qwen). We observe that the gap decreases with mergeability, implying that examples with better base model knowledge are more mergeable.

## A.2 Qwen Experimental Results

Figure A.7 shows experimental results on PopQA dataset on the Qwen model. Trends are inline with Llama model results (Figure A.1), with a decrease in the gap as mergeability score increases. This implies that examples with better base-model knowledge are more mergeable. Table A.2 shows experimental results with the Qwen model for the two other possible causes (§4.2, §4.3). Trends are similar to Llama model results from the main paper (Figure 3).

### A.3 How M and N Values Affect Mergeability?

Similar to Figure 2, Figure A.8 shows the mergeability distribution of the Llama model with  $N = 5$ ,  $M = 50$ , and Figure A.9 shows the mergeability scores with  $N = 20$ ,  $M = 50$ . We also experiment with different settings for  $M$ . Figure A.10 shows the mergeability distribution with  $N = 5$ ,  $M = 10$ . Similar figures for the Qwen model are A.11 and A.12 with  $N = 5$ ,  $M = 50$  and  $N = 10$ ,  $\leq 50$ , respectively. All figures show the difference from

a baseline distribution and support the existence of mergeability. We also examined the robustness of the mergeability score to different  $N$  and  $M$  values. Figures A.13 and A.14 show the change in the mergeability score as a function of  $N$  and  $M$ , respectively. In both Figures we see that mergeability scores calculated with one  $M$ ,  $N$  value are correlated with mergeability scores calculated with different  $M$  and  $N$  values, showing that scores are consistent under different parameters.

Mergeability score	0.0	0.2	0.4	0.6	0.8	1.0
Avg ppl	12.22	11.94	10.73	14.20	10.56	12.45
Avg context length (# of tokens)	589.52	893.25	540.48	1105.33	777.62	462.91
Avg weight norm	17.90	17.84	17.54	16.90	17.63	18.11
Avg weight highest singular value	26.21	25.83	25.45	23.78	25.85	26.99

Table A.2: Qwen model results for weight properties and training data difficulty in §4.




Figure A.8: Mergeability scores distribution of Llama-3.2-3B on the PopQA dataset with  $N = 5, M = 50$ . Blue and wide bars show the mergeability score as empirically calculated. Red thin bars show the baseline distribution if mergeability was not a model trait, and hence it was a binomial distribution with a fixed success rate.




Figure A.9: Mergeability scores distribution of Llama-3.2-3B on the PopQA dataset with  $N = 20, M = 50$ . Blue and wide bars show the mergeability score as empirically calculated. Red bars show the baseline distribution if mergeability was not a model trait, and hence it was a binomial distribution with a fixed success rate.




Figure A.10: Mergeability scores distribution of Llama-3.2-3B on the PopQA dataset with  $N = 10, M = 10$ . Blue and wide bars show the mergeability score as empirically calculated. Red thin bars show the baseline distribution if mergeability was not a model trait, and hence it was a binomial distribution with a fixed success rate. Compared to Figure 2, which shows the results for  $M = 50$ , we see fewer examples with a mergeability score of 0 and more examples with a higher mergeability score.




Figure A.11: Mergeability scores distribution of Qwen2.5-3B on the PopQA dataset with  $N = 10, M = 50$ . Blue and wide bars show the mergeability score as empirically calculated. Red thin bars show the baseline distribution if mergeability was not a model trait, and hence it was a binomial distribution with a fixed success rate.




Figure A.12: Mergeability scores distribution of Qwen2.5-3B on the PopQA dataset with  $N = 5, M = 50$ . Blue and wide bars show the mergeability score as empirically calculated. Red and thin bars show the baseline distribution if mergeability was not a model trait, and hence it was a binomial distribution with a fixed success rate.




Figure A.14: We compare the mergeability score when experimenting with different  $M$  values using Qwen2.5-3B on the PopQA dataset. The x-axis is the mergeability score for the  $M = 50, N = 5$  experiment. The y-axis shows the average mergeability score of those examples when experimenting with different  $M$  values. We observe a near-perfect increasing trend for all tested  $M$  values, which indicates that examples with a higher mergeability score in  $M = 50$  also had a higher mergeability score when experimented with other  $M$  values.




Figure A.13: We compare the mergeability score when experimenting with different  $N$  values using Qwen2.5-3B on the PopQA dataset. The x-axis is the mergeability score for the  $M = 50, N = 5$  experiment. The y-axis shows the average mergeability score of those examples when experimenting with different  $N$  values. We observe an increasing trend for all tested  $N$  values, which indicates that examples with a higher mergeability score in  $N = 5$  also had a higher mergeability score when experimented with other  $N$  values.

#### A.4 Lots-of-LoRAs additional experiments

Figure A.15 shows a scatter plot of the average performance degradation after merging as a function of the base model accuracy. We can observe a general trend of lower degradation for higher base model accuracies. In the main paper (§4), we experimented only with tasks where the finetuned model achieves at least 99% accuracy. In this section we include experimental results with more test accuracy thresholds. Figures A.16, A.17, A.18 and A.19 show the average base model accuracy for different mergeability scores when experimenting with tasks that the finetuned model achieves above 75%, 50%, 25% and 0% accuracy (respectively). We can see that trends are similar across all thresholds (including main paper results at Figure 4), where higher mergeability scores have higher average base model accuracy.




Figure A.15: Scatter plot of the average accuracy degradation as a function of the base model accuracy. Figure 4 is obtained by calculating the accuracy after degradation and splitting into bins.

#### A.5 Additional Experimental Results

Figure A.20 shows the accuracy when merging examples from a single mergeability group. Higher accuracies are obtained for higher mergeability scores. Figure A.21 shows the accuracy of a set of examples with different percentages of highly mergeable examples. We took a set of 50 lowest mergeability examples and 50 highest mergeability examples, and each time we merge 50 examples, while changing the percentage of highly mergeable examples from 0% to 100%. We see that as the percentage of highly mergeable examples increases, accuracy also increases. Figure A.22 shows the accuracy for different mergeability scores when merging an increasing number of examples. For a small number of merged examples, we see that the lines




Figure A.16: Lots-of-LoRAs average base model task accuracy of different mergeability scores, for tasks with base model accuracy of above 75%. Higher mergeability scores have on average a higher base model accuracy.




Figure A.17: Lots-of-LoRAs average base model task accuracy of different mergeability scores, for tasks with base model accuracy of above 50%. Higher mergeability scores have on average a higher base model accuracy.




Figure A.18: Lots-of-LoRAs average base model task accuracy of different mergeability scores, for tasks with base model accuracy of above 25%. Higher mergeability scores have on average a higher base model accuracy.




Figure A.19: Lots-of-LoRAs average base model task accuracy of different mergeability scores, for all available tasks (tasks with base model accuracy of above 0%). Higher mergeability scores have on average a higher base model accuracy.

are not ordered by their mergeability scores. As the number of merged examples increases, the lines change their order according to their mergeability scores. This shows that the number of merged examples ( $M$ ) might have an effect of the mergeability score.




Figure A.20: Accuracy when merging examples from the same mergeability group. We can see a trend of higher accuracy as the mergeability score increases.

## A.6 Experiments Hyperparameters

For LoRA adapter PopQA training, we use the hyperparameters in Table A.3. For the mergeability score calculation, we use  $M = 50$  and  $N = 5$  or  $N = 10$  in PopQA, and  $M = 10$  and  $N = 5$  for Lots-of-LoRAs.

## A.7 Affect of LoRA Rank

In the main text (§4) we experimented with a fixed LoRA rank  $r = 64$ . In this section, we include results for additional LoRA ranks,  $r = 8$  and

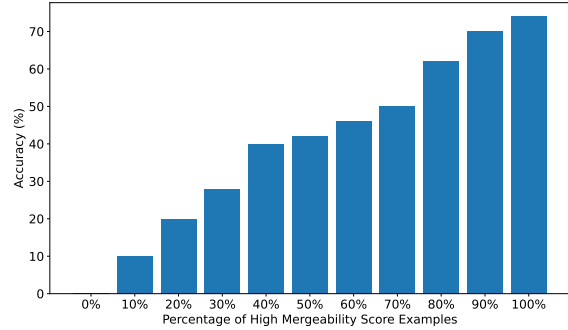


Figure A.21: Accuracy when merging a set of examples with different percentage of highly mergeable examples. We can see a trend of higher accuracy as more examples are with a high mergeability scores.




Figure A.22: Accuracy for different mergeability scores when merging an increasing number of examples. For a small number of merged examples ( $n = 5$  for example), we see the lines are not ordered by their mergeability score. As the number of merged examples increases, the lines change their order according to their mergeability score. This shows that the number of merged examples ( $M$ ) might have an effect of the mergeability score.

Model	Llama	Qwen
LoRA $r$	64	64
LoRA $\alpha$	128	128
LoRA trained layer	14	18
LoRA trained param	mlp.up_proj	mlp.up_proj
lr	4e-3	4e-3
epochs	10	10

Table A.3: PopQA training hyperparameters.




Figure A.23: Mergeability scores distribution of Qwen2.5-3B on the PopQA dataset with  $N = 5$ ,  $M = 50$  and LoRA rank  $r = 8$ . Blue and wide bars show the mergeability score as empirically calculated. Red and thin bars show the baseline distribution if mergeability was not a model trait, and hence it was a binomial distribution with a fixed success rate.

$r = 256$ . Besides the change of the rank, we maintained the identical experimental setup described in §3, using the Qwen 3B base model and evaluating using the example-level PopQA dataset. Figures A.23, A.25 show the mergeability score distribution for  $r = 8/256$ , respectively. Both figures show the difference from a baseline distribution and support the existence of mergeability. Figures A.24, A.26 show experimental results for  $r = 8/256$ , respectively. For  $r = 8$  trends are inline with  $r = 64$  results (Figure A.7), with lower differences for higher mergeability scores. For  $r = 256$  we also observe a general decreasing effect. However, trends are less sharp compared to  $r = 64$  results.

### A.8 Mergeability of Full Finetuning

To verify that mergeability is not a phenomenon exclusive to parameter-efficient LoRA fine-tuning, we extended our evaluation to full model fine-tuning. We utilized the same Qwen 3B base model and the PopQA dataset experimental setup described in §3, but trained all model parameters instead of using LoRA adapters. Since the KNOTS




Figure A.24: PopQA average difference between the highest and the correct answer probability in the base model (Qwen). Mergeability scores are for examples trained with LoRA rank  $r = 8$ . We observe that the gap decreases with mergeability, implying that examples with better base model knowledge are more mergeable.




Figure A.25: Mergeability scores distribution of Qwen2.5-3B on the PopQA dataset with  $N = 5$ ,  $M = 50$  and LoRA rank  $r = 256$ . Blue and wide bars show the mergeability score as empirically calculated. Red and thin bars show the baseline distribution if mergeability was not a model trait, and hence it was a binomial distribution with a fixed success rate.




Figure A.26: PopQA average difference between the highest and the correct answer probability in the base model (Qwen). Mergeability scores are for examples trained with LoRA rank  $r = 256$ . We observe that the gap generally decreases with mergeability, implying that examples with better base model knowledge are more mergeable.




Figure A.27: Mergeability scores distribution of Qwen2.5-3B on the PopQA dataset with  $N = 5$ ,  $M = 50$  and full finetuning training. The scores were calculated using mean merging. Blue and wide bars show the mergeability score as empirically calculated. Red and thin bars show the baseline distribution if mergeability was not a model trait, and hence it was a binomial distribution with a fixed success rate.




Figure A.29: Mergeability scores distribution of Qwen2.5-3B on the PopQA dataset with  $N = 5$ ,  $M = 50$  and full finetuning training. The scores were calculated using TIES merging algorithm. Blue and wide bars show the mergeability score as empirically calculated. Red and thin bars show the baseline distribution if mergeability was not a model trait, and hence it was a binomial distribution with a fixed success rate.




Figure A.28: PopQA average difference between the highest and the correct answer probability in the base model (Qwen). Mergeability scores are for examples trained with full finetuning. The scores were calculated using mean merging. We observe that the gap generally decreases with mergeability, implying that examples with better base model knowledge are more mergeable.

algorithm is designed for LoRA weight merging, we employed two alternative merging algorithms suitable for full weights: TIES-Merging and simple mean merging. Figures A.27 and A.29 show that mergeability also occurs in the full finetuning case. Experimental results in Figures A.28 and A.30 generally follow the expected decreasing trend. Those results show that our findings, examples where the base model has better knowledge of are more mergeable, also hold for the full finetuning training.

### A.9 Different Merging Algorithms

In addition to the main paper results using KNOTS merging algorithm, we also examined other merging algorithms - TIES and mean. We used the




Figure A.30: PopQA average difference between the highest and the correct answer probability in the base model (Qwen). Mergeability scores are for examples trained with full finetuning. The scores were calculated using TIES merging. We observe that the gap generally decreases with mergeability, implying that examples with better base model knowledge are more mergeable.

Qwen 3B base model and the PopQA dataset experimental setup described in §3 with the change of using a different merging algorithm for mergeability measurements. Figure A.31 and A.33 show the mergeability score distribution for TIES and mean merging, respectively. TIES distribution shows similar trends to KNOTS (Figure A.12) and supports the existence of mergeability. Mean distribution shows similar trends to Figure 6, with a high number of examples in the  $S = 1.0$  bin and a very small number of examples in the middle bins ( $0 < S < 1$ ). As discussed in (§5.2), we can attribute this to the lack of any conflict mitigation in mean merging. TIES experimental results (Figure A.32) show the same decreasing trend between base model knowledge and mergeability score, revealing that our results also hold for TIES merging. When using mean merging, the experimental results (Figure A.34) does not show the decreasing trend anymore. We can attribute this to the low number of examples with  $0 < S < 1$  that may lead to inaccurate evaluation.

We further tested how the mergeability score of examples changes as we change the merging algorithm. Figure A.35 shows the mergeability score of examples using TIES or mean merging compared to their score as calculated when using KNOTS as the merging algorithm. The x-axis is the mergeability score when using KNOTS merging. The y-axis shows the average mergeability score of those examples when experimenting with different merging algorithms - TIES or mean. We observe an increasing trend for both merging algorithms, which indicates that examples with a higher mergeability score obtained using KNOTS also had a higher mergeability score when experimented with other merging algorithms. We also see that KNOTS and TIES scores are more similar compared to mean. This can be explained by the reason that both use a conflict mitigation algorithm compared to mean.




Figure A.31: Mergeability scores distribution of Qwen2.5-3B on the PopQA dataset with  $N = 5$ ,  $M = 50$  and LoRA  $r = 64$ . The scores were calculated using TIES merging algorithm. Blue and wide bars show the mergeability score as empirically calculated. Red and thin bars show the baseline distribution if mergeability was not a model trait, and hence it was a binomial distribution with a fixed success rate.




Figure A.32: PopQA average difference between the highest and the correct answer probability in the base model (Qwen). The scores were calculated using TIES merging. We observe that the gap generally decreases with mergeability, implying that examples with better base model knowledge are more mergeable.




Figure A.33: Mergeability scores distribution of Qwen2.5-3B on the PopQA dataset with  $N = 5$ ,  $M = 50$  and LoRA  $r = 64$ . The scores were calculated using mean merging. Blue and wide bars show the mergeability score as empirically calculated. Red and thin bars show the baseline distribution if mergeability was not a model trait, and hence it was a binomial distribution with a fixed success rate.




Figure A.34: PopQA average difference between the highest and the correct answer probability in the base model (Qwen). The scores were calculated using mean merging. We do not observe a correlation with the mergeability score. This might be explained by the low number of examples with  $0 < S < 1$ .

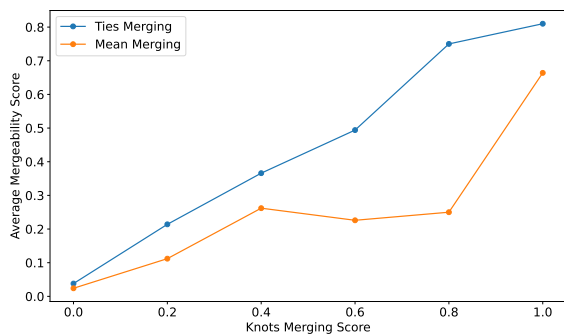


Figure A.35: We compare the mergeability score calculated using KNOTS to other merging algorithms. Results are for Qwen2.5-3B on the PopQA dataset. The x-axis is the mergeability score when calculated using KNOTS. The y-axis shows the average mergeability score of those examples when calculated with other merging algorithms. We observe an increasing trend for both mean and TIES.