

Expressive and Generalizable Low-rank Adaptation for Large Models via Slow Cascaded Learning

Anonymous ACL submission

Abstract

Efficient fine-tuning plays a fundamental role in modern large models, with low-rank adaptation emerging as a particularly promising approach. However, the existing variants of LoRA are hampered by limited expressiveness, a tendency to overfit, and sensitivity to hyperparameter settings. This paper presents LoRA Slow Cascade Learning (*LoRASC*), an innovative technique designed to enhance LoRA’s expressiveness and generalization capabilities while preserving its training efficiency. Our approach augments expressiveness through a cascaded learning strategy that enables a mixture-of-low-rank adaptation, thereby increasing the model’s ability to capture complex patterns. Additionally, we introduce a slow-fast update mechanism and cascading noisy tuning to bolster generalization. The extensive experiments on various language and vision datasets, as well as robustness benchmarks, demonstrate that the proposed method not only significantly outperforms existing baselines, but also mitigates overfitting, enhances model stability, and improves OOD robustness.

1 Introduction

Foundation models, which are large-scale models pre-trained on extensive datasets and subsequently adapted for specific downstream tasks, have become integral to contemporary machine learning frameworks. Fine-tuning these models is essential, yet full parameter fine-tuning often encounters significant memory and computational bottlenecks. As a result, Parameter-Efficient Fine-Tuning (PEFT) techniques, which aim to minimize the number of trainable parameters to reduce training costs and improve training stability, have gained increasing prominence. Among these techniques, Low-Rank Adaptation (LoRA) (Hu et al., 2021) stands out due to its efficiency in reducing training costs through low-rank approximation for full-parameter updates. However, despite LoRA’s ad-

vantages, its limitations in terms of expressiveness and generalization have been noted. Some studies suggest that the inherent low-rankness of LoRA might restrict its expressiveness (Xia et al., 2024; Meng et al., 2024; Lialin et al., 2023; Huang and Wei, 2024), with a preference for overparameterization, while others indicate a tendency for LoRA to overfit or exhibit overconfidence (Lin et al., 2024; Wang et al., 2023).

In this work, we investigate the potential of cascading learning to augment the expressiveness of LoRA. Our approach involves initializing a new LoRA module at the start of each epoch and integrating this module into the backbone network after the epoch concludes. By employing a mixture-of-low-rank adaptation, we effectively increase the model’s rank, while maintaining low training costs, as each cascading step consumes no more parameters and memory than a single LoRA model. Moreover, this method does not add any inference overhead by remerging each LoRA module into the backbone network.

To improve LoRA’s generalization capabilities, we draw inspiration from optimization techniques. We repurpose certain strategies from optimizers for LoRA, motivated by the observation that initializing a new LoRA module for each epoch can represent a descent direction for the dataset. In optimization theory, flat minimizers are preferred, as they are associated with better generalization (Hochreiter and Schmidhuber, 1997; Keskar et al., 2016). Inspired by the fact that the moving average mechanism guides models towards flat minimizers (Izmailov et al., 2018), we maintain both fast-updating and its moving average version, the slow-updating LoRA experts. The fast-updating expert is reinitialized regularly to learn from the data over a set number of steps, while the slow-updating expert undergoes updates via a proportional exponential moving average after the fast-updating cycle com-

083 pletes. Additionally, mirroring techniques in deep
084 learning optimizers where noise proportional to the
085 gradient scale is used to find flat minima (Xie et al.,
086 2020), we introduce noise at the beginning of each
087 epoch, with the scale tied to the norm of LoRA’s
088 weights.

089 To verify the effectiveness of the proposed method,
090 we conduct extensive experiments on both lan-
091 guage and vision tasks. For language tasks, we
092 utilized the Llama2 model on 12 datasets (e.g.,
093 SuperGLUE, SQuAD, DROP, GSM8K, and In-
094 structEval), Alpaca among other instruct follow-
095 ing benchmarks to demonstrate the effectiveness
096 of our design. We can directly apply our approach
097 to LoRA, LoRA+ (Hayou et al., 2024), Dora (Liu
098 et al., 2024), and other members of the LoRA fam-
099 ily, significantly improving their performance in
100 large model transfer learning. For vision tasks, we
101 also validated our approach on the CLIP pre-trained
102 Vit-bigG model with the ImageNet dataset, show-
103 ing a significant performance improvement rela-
104 tive to LoRA on domain adaptation datasets such
105 as Image-R and Image-C. The proposed method
106 consistently outperforms the baselines by a large
107 margin.

108 2 Related Work

109 2.1 Low-Rank Adaptation Finetuning

110 Low-Rank Adaptation(LoRA) (Hu et al., 2021) is
111 a parameter-efficient fine-tuning method designed
112 to adapt large models to new tasks, demonstrat-
113 ing superior performance. LoRA+ (Hayou et al.,
114 2024) improves performance and fine-tuning speed
115 by setting different learning rates for the LoRA
116 adapter matrices A and B with a carefully cho-
117 sen ratio, maintaining the same computational cost
118 as LoRA. Dora (Liu et al., 2024) decomposes the
119 pre-trained weight into two components, magni-
120 tude and direction, for fine-tuning, specifically em-
121 ploying LoRA for directional updates to efficiently
122 minimize the number of trainable parameters. Our
123 work introduces a robust cascading learning sched-
124 ule for various LoRA variants, proving through ex-
125 tensive experiments that it can enhance the training
126 performance of LoRA, LoRA+, and Dora without
127 additional training costs.

128 2.2 Combination of LoRA

129 LoRAhub (Huang et al., 2023) presents a simple
130 framework designed for the purposeful assembly
131 of LoRA modules trained on diverse tasks, aim-

132 ing to achieve adaptable performance on unseen
133 tasks. MOLE (Huang and Wei, 2024) treats each
134 layer of trained LoRAs as a distinct expert and
135 implements hierarchical weight control by integrat-
136 ing a learnable gating function within each layer.
137 LoRAFlow (Wang et al., 2024) utilizes dynamic
138 weights to adjust the impact of different LoRAs.
139 These methods are not in conflict with LoRASC, as
140 they focus on learning the combination of LoRA
141 experts across different domains, while our method
142 aims to learn more generalizable experts within a
143 single domain using slow cascade learning.

144 ReLoRA (Lialin et al., 2023) enhances LoRA’s fit-
145 ting ability by continuously merging online LoRA
146 into the main network and restarting optimizer pa-
147 rameters during training. It also proposes a jagged
148 cosine scheduler to implement a learning rate re-
149 sume strategy at each step. COLA (Xia et al., 2024)
150 explores a similar approach but in a simpler manner,
151 merely restarting optimizer parameters when ini-
152 tializing new LoRAs without adjusting the learning
153 rate schedule. Our work employs a simpler cas-
154 cading learning strategy where each expert learns
155 independently for each epoch, without additional
156 design for learning schedules or optimizer param-
157 eters. Additionally, we incorporate noise tuning and
158 slow-fast update strategy, ensuring robustness in
159 each expert merged into the pre-trained model. Our
160 method can be applied to various LoRA variants,
161 demonstrating effectiveness across multiple tasks
162 in both language and image domains.

163 3 Methods

164 3.1 LoRA

165 Low-Rank Adaptation (LoRA) is a parameter-
166 efficient fine-tuning method designed to adapt large
167 pre-trained models to specific tasks with signifi-
168 cantly fewer trainable parameters. Instead of up-
169 dating all parameters of the model, LoRA inserts
170 low-rank matrices into each layer of the pre-trained
171 model, which are then fine-tuned. This reduces the
172 computational burden and the risk of overfitting.

173 Given a pre-trained weight matrix $W_0 \in \mathbb{R}^{d \times k}$ in
174 a neural network, LoRA approximates the update
175 ΔW using two low-rank matrices $A \in \mathbb{R}^{d \times r}$ and
176 $B \in \mathbb{R}^{r \times k}$, where $r \ll \min(d, k)$. The update is
177 defined as:

$$178 \Delta W = BA \quad (1)$$

179 During fine-tuning, instead of updating W , we up-

date A and B , which results in:

$$W = W_0 + \Delta W = W_0 + BA \quad (2)$$

This low-rank adaptation significantly reduces the number of trainable parameters from $d \times k$ to $r \times (d + k)$.

3.2 LoRASC

3.2.1 Cascading LoRA Learning

Due to the reparameterization nature of low-rank adaptation (LoRA) fine-tuning, employing multiple LoRA experts incurs the same inference cost as using a single LoRA expert. This characteristic makes LoRA particularly suitable for integration with cascading learning to enhance performance in transfer learning tasks. As analyzed in ReLoRA (Lialin et al., 2023), reinitializing new LoRA modules during the learning schedule can progressively increase the model’s rank, thereby improving its fitting ability.

In *LoRASC*, we default to learning one LoRA expert per epoch. After training one LoRA expert, it is merged into the main network, and the next expert learns based on the optimized residuals. The optimization schedule for each single LoRA expert is a compressed version of the original full-training schedule: for instance, if a model was originally trained for N epochs, each expert in *LoRASC* completes training in 1 epoch with fixed starting and ending learning rates with the same but compressed scheduler. This makes *LoRASC* easy to apply to any large model transfer learning scenario using LoRA, without requiring changes to hyperparameters. The only necessary adjustment is an increase in the learning rate. Since the number of training steps is compressed, each step must be larger to cover the same distance. Additionally, Li et al. (Li et al., 2019) found that higher learning rates can lead to stronger generalization ability, which might also explain the improved out-of-domain performance of our method.

Mathematically, the cascading LoRA learning can be described as follows:

1. For each epoch t , train a new LoRA expert (A_t, B_t) to minimize the residual error, where \mathcal{L} is the fine-tuning loss function:

$$(A_t, B_t) = \arg \min_{A_t, B_t} \mathcal{L}(W_{t-1} + B_t A_t), \quad (3)$$

2. Merge the trained LoRA expert into the main network:

$$W_t = W_{t-1} + B_t A_t \quad (4)$$

By iteratively merging each new LoRA expert into the main network, loRA cascading progressively enhances the model’s capacity to fit the data without increasing the inference cost.

3.2.2 LoRA Slow-Fast Update

To enhance the generalization of large model transfer learning, we aim to avoid local optima at each step of cascading. Even with low-rank adaptation, this issue persists due to the imbalance between model parameters and training data. Inspired by SWA (Izmailov et al., 2018), which averages model parameters over several epochs to find a more generalized solution, we employ a sliding average method to ensure the stability and robustness of each LoRA merged into the main network.

Specifically, during training, we maintain two LoRA experts at each cascading step t as shown in Fig. 1: a slow-updating LoRA $(A_t^{\text{slow}}, B_t^{\text{slow}})$ and a fast-updating LoRA $(A_t^{\text{fast}}, B_t^{\text{fast}})$. At step 0, both slow and fast LoRA share the same initialization. During each cascading iteration, fast LoRA undergoes fine-tuning, and after completion, it is averaged with slow LoRA. The slow LoRA is then merged into the pre-trained model, while the fast-updating LoRA is reinitialized for the next iteration. We control the retention proportion of the slow expert with a hyperparameter α .

The update rules are given by:

$$A_{t+1}^{\text{slow}} = \alpha A_t^{\text{slow}} + (1 - \alpha) A_t^{\text{fast}} \quad (5)$$

$$B_{t+1}^{\text{slow}} = \alpha B_t^{\text{slow}} + (1 - \alpha) B_t^{\text{fast}} \quad (6)$$

By employing this slow-fast update strategy, *LoRASC* ensures that each merged LoRA expert contributes to a more generalized solution, enhancing the overall stability and performance of the model in transfer learning scenarios.

3.2.3 Cascading Noisy Tuning

To further enhance generalization, we introduce random noise to the pre-trained model before each new LoRA fine-tuning step. Unlike Noisy-Tune (Wu et al., 2022), which adds uniform noise to different parameter matrices according to their standard deviations only once at the beginning of

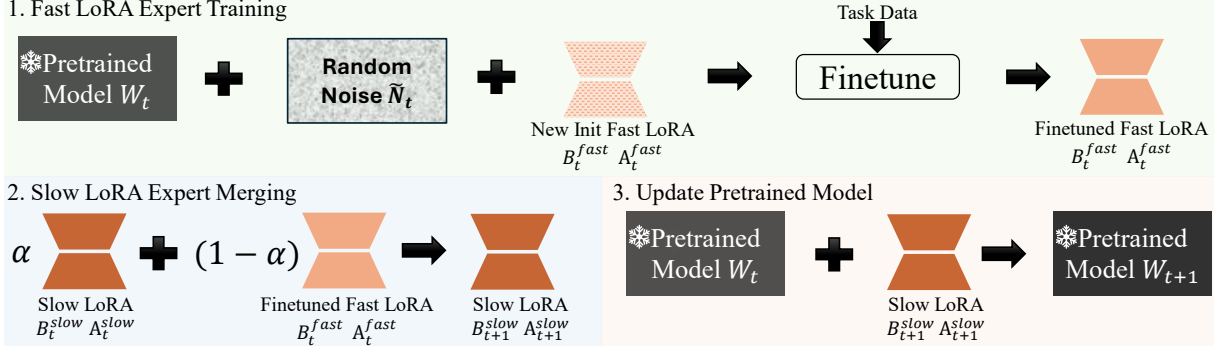


Figure 1: Iterative pipeline of *LoRASC*. Here, t represents the iteration step, and BA denotes the low-rank learnable vectors in LoRA. The backbone network W always has its gradients turned off, and α is the hyperparameter controlling the pace of the slow-fast update. Our method follows three stages: 1. Fast LoRA expert training, where noise is added to the backbone network, followed by training the fast LoRA on the task data. 2. Slow LoRA expert merging, where a portion of the learned fast LoRA is weighted and merged into the slow LoRA. 3. Update the pretrained model, merging the updated slow LoRA into the backbone network, and prepare for the next iteration.

fine-tuning, we apply noise before training each new expert. This approach helps the model escape local optima at every slow LoRA step, thereby reducing the risk of overfitting.

Additionally, the presence of the slow-updating LoRA module indicates the direction of parameter changes under the new task. Therefore, we use the standard deviation of the slow LoRA weights to determine the noise scale rather than the pre-trained model’s weights. Incorporating this noise before every expert ensures that the model continuously explores robust and flatten parameter spaces, thus improving generalization and reducing the tendency to overfit.

The perturbation is defined as:

$$\tilde{N}_t = U\left(-\frac{\lambda}{2}, \frac{\lambda}{2}\right) \cdot \text{std}(B_t^{\text{slow}} A_t^{\text{slow}}) \quad (7)$$

where std stands for standard deviation. The function $U(a, b)$ represents uniform distribution noise ranged from a to b , and λ is a hyperparameter that controls the relative noise intensity.

3.3 Overview

With LoRA cascading learning, slow-fast updates and noisy tuning, the pipeline of our *LoRASC* is as follows:

$$\tilde{W}_{t-1} = W_{t-1} + \tilde{N}_t \quad (8)$$

$$(A_t^{\text{fast}}, B_t^{\text{fast}}) = \arg \min_{A_t^{\text{fast}}, B_t^{\text{fast}}} \mathcal{L}\left(\tilde{W}_{t-1} + B_t^{\text{fast}} A_t^{\text{fast}}\right) \quad (9)$$

$$A_t^{\text{slow}} = \alpha A_{t-1}^{\text{slow}} + (1 - \alpha) A_t^{\text{fast}} \quad (10)$$

$$B_t^{\text{slow}} = \alpha B_{t-1}^{\text{slow}} + (1 - \alpha) B_t^{\text{fast}} \quad (11)$$

$$W_t = \tilde{W}_{t-1} + B_t^{\text{slow}} A_t^{\text{slow}} \quad (12)$$

LoRASC pipeline can be seen in Fig. 1. Although we use vanilla LoRA to show slow cascade learning, *LoRASC* should be able to boost the performance of any LoRA variants, such as DoRA (Liu et al., 2024), LoRA+ (Hayou et al., 2024), LoRA-FA (Zhang et al., 2023), etc. Moreover, *LoRASC* is easy to implement, and we provide pseudocode with more detailed explanations in Algorithm 1.

4 Experiments

We conducted extensive experiments to demonstrate the effectiveness and robustness of *LoRASC* across both NLP and CV domains.

For language tasks, we conducted our language experiments using the popular open-source large language model, Llama2¹. We evaluated our approach on several NLU and GLU tasks, selecting both SuperGLUE (Wang et al., 2019a) tasks (including classification and multiple-choice) and generation tasks. We also tested the model’s performance in mathematical reasoning using the GSM8K dataset (Cobbe et al., 2021). Additionally, we performed instruction tuning experiments to verify the transfer learning capability of our method, achieving significant improvements on key metrics such as MMLU (Hendrycks et al., 2020),

¹<https://huggingface.co/meta-llama/llama-2-7b-hf>

Algorithm 1 Pseudo Code for *LoRASC*

Require: Pre-trained model weights W_0 , number of epochs T , loss function \mathcal{L} , slow update parameter α , noise parameter λ

- 1: Initialize $W \leftarrow W_0$
- 2: Initialize $A^{\text{slow}}, B^{\text{slow}} \triangleright$ Initialize slow LoRA matrices
- 3: Initialize $A^{\text{fast}} \leftarrow A^{\text{slow}}, B^{\text{fast}} \leftarrow B^{\text{slow}} \triangleright$ Fast LoRA matrices initialized from slow ones
- 4: **for** epoch $t = 1$ to T **do**
- 5: **if** $t > 1$ **then**
- 6: Reinitialize $A^{\text{fast}}, B^{\text{fast}} \triangleright$ Reinitialize fast LoRA matrices for subsequent epochs
- 7: **end if**
- 8: $\widetilde{W} \leftarrow W + U\left(-\frac{\lambda}{2}, \frac{\lambda}{2}\right) \cdot \text{std}(B^{\text{slow}} A^{\text{slow}})$
- 9: optimizer \leftarrow InitializeOptimizer($A^{\text{fast}}, B^{\text{fast}}$)
- 10: lr_scheduler \leftarrow InitializeLRScheduler(optimizer)
- 11: **for** batch in training data **do**
- 12: Forward pass: $L \leftarrow \mathcal{L}(\widetilde{W} + B^{\text{fast}} A^{\text{fast}})$
- 13: Backward pass: Compute gradients
- 14: optimizer.step()
- 15: lr_scheduler.step()
- 16: **end for**
- 17: Update slow LoRA:
- 18: $A^{\text{slow}} \leftarrow \alpha A^{\text{slow}} + (1 - \alpha) A^{\text{fast}}$
- 19: $B^{\text{slow}} \leftarrow \alpha B^{\text{slow}} + (1 - \alpha) B^{\text{fast}}$
- 20: Merge slow LoRA into main network:
 $W \leftarrow \widetilde{W} + B^{\text{slow}} A^{\text{slow}}$
- 21: **end for**
- 22: **return** W

DROP (Dua et al., 2019), BBH (Srivastava et al., 2022) and HumanEval (Chen et al., 2021).

For visual tasks, we chose the CLIP ViT-bigG/14² as our pretrained model, fine-tuning it on the ImageNet-1K (Deng et al., 2009) training set and testing it on the validation set. Subsequently, we evaluated the trained model on perturbed datasets such as ImageNet-A (Hendrycks et al., 2021b), ImageNet-C (Hendrycks and Dietterich, 2019), ImageNet-R (Hendrycks et al., 2021a), ImageNet-V2 (Recht et al., 2019), ImageNet-Sketch (Wang et al., 2019b) and Stylized-ImageNet (Geirhos et al., 2018) demonstrating our method’s robustness and generalization capabilities.

²<https://huggingface.co/laion/CLIP-ViT-bigG-14-laion2B-39B-b160k>

4.1 Implementation Details

For all experiments, we exclusively fine-tuned q and v in attention layers as delineated by Malladi et al. (2023) and Ren et al. (2024). The fine-tuning process utilized single NVIDIA H100 GPU. For all tasks, we explored several learning rates and reported the optimal performance. For the hyperparameters of *LoRASC*, we explored the factor α of Slow-Fast Update in $\{0.5, 0.6, 0.8\}$ to control the updating ratio. Additionally, we selected the noise intensity from $\{0.1, 1, 10\}$, which is a significantly smaller set compared to the default 7 in NoistTune (Wu et al., 2022). All the results were averaged across 3 distinct random seeds, and we report the optimal performance.

4.2 Main Results

4.2.1 *LoRASC* for Large Language Model

Experiment setting. For in-domain language transfer learning, we consider the SuperGLUE dataset collection (Wang et al., 2019a), including: BoolQ (Clark et al., 2019), CB (De Marneffe et al., 2019), COPA (Roemmele et al., 2011), MultiRC (Khashabi et al., 2018), ReCoRD (Zhang et al., 2018), RTE (Socher et al., 2013), WiC (Pilehvar and Camacho-Collados, 2019), and WSC (Levesque et al., 2012). We also include SST-2 (Dagan et al., 2005), GSM8K (Cobbe et al., 2021) and two question answering (QA) datasets, SQuAD (Rajpurkar et al., 2016) and DROP (Dua et al., 2019). And we directly used 8-shot direct prompting or GSM8K evaluation³. We adhered to the experimental configuration described by Malladi et al. (2023), randomly selecting 1000 examples for training, 500 for validation, and 1000 for testing across each dataset. The AdamW optimizer was employed, with training spanning 5 epochs, consistent with the baseline settings. A linear learning rate schedule was implemented, with the initial learning rate selected from $\{1 \times 10^{-5}, 5 \times 10^{-5}, 1 \times 10^{-4}, 5 \times 10^{-4}, 1 \times 10^{-3}\}$. By default the batch size was set to 4 and the LoRA rank was set to 8. For LoRA+, we adhered to its setup by fixing the learning rate of B matrices to be 16 times that of A matrices. DoRA decomposes the pre-trained weight into magnitude and direction components, with LoRA efficiently updating the direction component. This means that each LoRA expert represents DoRA’s direction component. When applying *LoRASC* to DoRA, we maintain continuous training of the mag-

³<https://github.com/allenai/open-instruct>

Task	SST-2	RTE	CB	BoolQ	WSC	WIC	MultiRC	COPA	ReCoRD	SQuAD	DROP	GSM8K
Task type	classification						multiple choice			generation		math
LoRA	95.5	87.4	91.1	85.7	70.2	72.4	85.3	85.0	81.2	90.4	51.6	19.5
w/ COLA	95.9	87.7	91.1	85.7	66.4	72.6	85.3	82.0	81.4	90.6	51.6	21.0
w/ <i>LoRASC</i>												
+ Cascade	95.8	87.7	92.9	86.1	71.1	72.3	86.3	88.0	81.6	91.8	52.5	21.5
++ Slow LoRA	96.0	88.0	96.4	86.8	74.0	72.1	86.3	88.0	82.1	92.7	55.3	27.5
+++ Noise Tuning	96.1	88.1	96.5	87.4	75.0	72.7	86.6	88.0	82.2	92.9	56.7	27.5
LoRA+	95.7	87.0	91.4	85.9	69.2	72.1	85.7	87.0	81.3	90.5	55.8	22.0
w/ <i>LoRASC</i>												
+ Cascade	95.7	87.0	92.9	86.2	71.2	72.8	85.3	88.0	81.9	91.2	55.8	19.5
++ Slow LoRA	95.7	88.1	92.9	85.9	67.3	73.5	85.7	88.0	81.9	92.0	56.3	23.0
+++ Noise Tuning	95.8	88.1	92.9	86.3	71.4	74.1	86.1	88.0	81.9	92.0	56.4	24.0
DoRA	95.4	87.4	96.4	85.7	72.1	71.5	84.7	88.0	81.1	91.1	54.8	21.0
w/ <i>LoRASC</i>												
+ Cascade	95.8	87.4	96.4	85.8	65.4	72.8	84.1	88.0	81.6	91.7	52.6	22.5
++ Slow LoRA	95.8	88.1	96.4	85.8	65.4	72.8	86.1	88.0	81.9	92.8	54.8	25.0
+++ Noise Tuning	96.0	88.5	96.5	87.6	75.6	72.8	86.8	89.0	82.2	93.3	56.5	25.5

Table 1: Comparative Performance of LoRA, LoRA+, and DoRA enhanced with *LoRASC* across multiple in-domain fine-tuning datasets.

Method	MMLU	DROP	HEval	BBH	GSM8K
LoRA	45.83	32.76	31.26	13.41	11.5
w/ <i>LoRASC</i>					
+ Cascade	45.53	32.71	31.61	14.02	11.5
++ Slow LoRA	45.68	33.74	31.38	17.07	12.5
+++ Noise	45.98	33.02	31.61	15.24	16.5

Table 2: Results on instruction-following tasks. The model was trained on Alpaca and evaluated on InstructEval metrics and GSM8K. *LoRASC* consistently achieves the best performance compare to vanilla LoRA.

nitide while applying our technique to the direction component. We follow the standard procedure of merging and reinitializing LoRA and align it with the slow-fast update and noisy tuning.

For instruction tuning, we use the Alpaca⁴ (Taori et al., 2023) dataset for training. The batch size was set to 128. We follow the training scripts of Ren et al. (2024) in our experiment. We finetune our model for 3 epochs. A linear learning rate schedule was applied, with the initial learning rate selected from $\{1 \times 10^{-4}, 3 \times 10^{-4}, 5 \times 10^{-4}, 1 \times 10^{-3}\}$. For evaluation we use InstructEval⁵ (Chia et al., 2023), 5-shot direct prompting for MMLU, 3-shot direct prompting for BBH and DROP, 0-shot direct prompting for HEval.

***LoRASC* exhibits excellent adaptability to LoRA variants.** In the experiments shown in Table 1, *LoRASC* outperforms the COLA across various

tasks, demonstrating the effectiveness of our LoRA cascading technique. Moreover, *LoRASC* effectively boosted the performance of LoRA, LoRA+, and DoRA across 12 in-domain training datasets encompassing four major tasks: classification, multiple choice, generation, and mathematics. *LoRASC* achieved significant improvements across all these tasks, demonstrating its ability to enhance the learning capabilities and in-domain generalization of the LoRA family of models. Moreover, the progressive addition of cascading learning, slow-fast updates, and noisy tuning further improved performance, validating the design of our approach. The robust slow cascading strategy not only enhanced overall performance but also provided strong generalization capabilities.

***LoRASC* on Instruction-Following tasks.** Table 2 presents the performance of our proposed method, *LoRASC*, applied to LoRA across several instruction-following tasks. These instruction-following tasks are particularly challenging due

⁴https://github.com/tatsu-lab/stanford_alpaca/

⁵<https://github.com/declare-lab/instruct-eval>

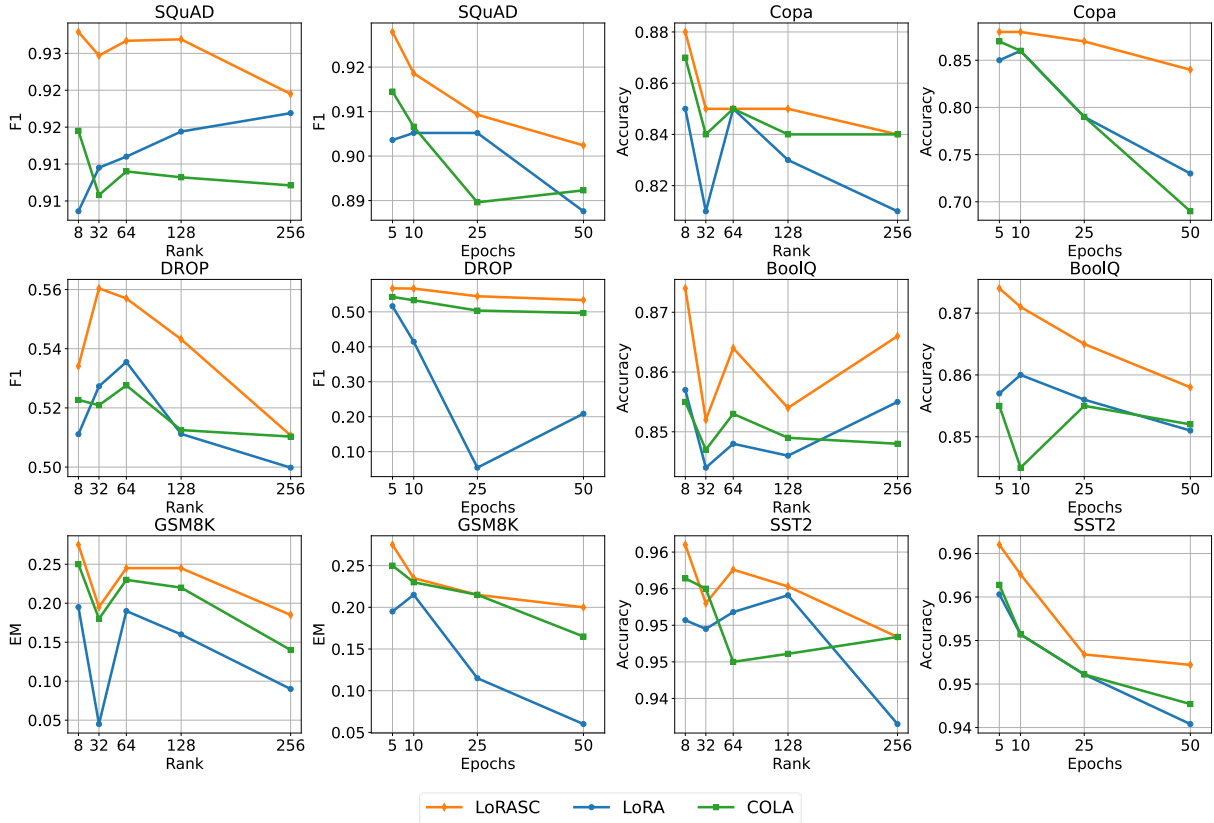


Figure 2: Performance of *LoRASC* compared to LoRA and COLA across various ranks and learning schedules in a subset of text transfer learning tasks. It can be observed that *LoRASC* consistently achieves stable performance improvements across all ranks and learning schedules, particularly at higher ranks and longer epochs, where *LoRASC* can mitigate performance degradation caused by overfitting.

to the weak correlation between the training data and the benchmarks, making them entirely out-of-domain tests. Despite this difficulty, our method achieved notable improvements across various evaluation metrics used in InstructEval and GSM8K. Furthermore, the design of slow-fast updates and noisy tuning still steadily enhanced the performance of cascading learning, further validating the effectiveness of our approach and motivation.

4.2.2 *LoRASC* for CLIP ViT-bigG

Experiment setting. For the ImageNet-1K visual classification task, to validate the transfer performance of our method on larger vision models, we selected CLIP ViT-bigG/14 as our pre-training backbone. We utilized the AdamW optimizer and a cosine scheduler, training for a total of 10 epochs on the ImageNet-1K training set. The batch size was fixed at 64, and the learning rate was chosen from $\{1 \times 10^{-4}, 5 \times 10^{-4}, 1 \times 10^{-3}\}$. For evaluation, we first test our model on the ImageNet-1K validation set using top-1 accuracy. To demonstrate the improvement in our method’s transferability and ro-

bustness, we conducted further tests on robustness benchmarks from Mao et al. (2022) for transfer learning tasks.

Evaluation of *LoRASC* on ImageNet and Robustness Benchmarks.

Table 3 showcases the performance of our proposed method, *LoRASC*, applied to LoRA on ImageNet-1K and several robustness benchmarks, including IN-V2, IN-C, IN-R, IN-A, IN-SK, and IN-ST. These benchmarks test the model’s robustness and generalization ability beyond the standard ImageNet dataset. Our method demonstrates consistent improvements in top-1 accuracy across all evaluated benchmarks. *LoRASC* consistently enhances the robustness and generalization of the ViT-bigG model across these challenging benchmarks, validating the effectiveness of cascading learning, slow-fast updates, and noisy tuning in improving model performance in diverse and robust scenarios.

Method	ImageNet	IN-V2	IN-C	IN-R	IN-A	IN-SK	IN-ST
LoRA	87.1	77.7	66.2	87.1	72.6	64.9	24.1
w/ <i>LoRASC</i>							
+ Cascade	87.1	77.5	66.7	88.5	73.6	65.4	24.3
++ Slow LoRA	87.7	78.3	66.8	88.1	73.4	65.2	24.1
+++ Noise Tuning	87.8	78.4	66.8	88.7	73.4	65.5	24.4

Table 3: Top-1 accuracy of various methods on ImageNet-1K and 6 robustness benchmarks. The table compares the baseline LoRA with our three proposed techniques. Our approach demonstrates improved robustness on the ViT-bigG model across all the evaluated benchmarks.

Experts	RTE	DROP	WIC	BoolQ	ReCoRD	SST-2	SQuAD
2	87.0	53.8	72.4	85.3	81.3	95.5	92.0
5	88.1	56.7	72.6	87.4	82.2	96.1	92.9
25	86.7	51.2	70.5	83.5	81.4	95.5	92.2
125	83.8	50.2	70.5	84.5	81.2	95.1	90.7
1250	83.8	49.4	69.4	85.3	81.1	92.9	88.1

Table 4: Evaluation with varying expert number of *LoRASC*. The highest average performance for each task is highlighted in bold.

4.3 Ablation Study and Analysis

Larger Ranks and Longer Epochs. As shown in Fig. 2, *LoRASC* consistently achieves more stable performance on datasets such as SQuAD, DROP, and GSM8K compared to both LoRA and COLA, which also employs a cascading strategy. This validates our motivation: *LoRASC* is a training strategy that retains LoRA’s beneficial properties while seamlessly enhancing its fitting ability and robust generalization.

Ablation for *LoRASC* Expert Cascade Frequency. *LoRASC* defaults to updating once per epoch, as each expert completes training on the entire dataset within one epoch. In Table 4, we experimented with different update frequencies. In this setting, we trained for a total of 5 epochs, with each epoch consisting of 250 iterations, resulting in a total training period of 1250 iterations. The table shows that having 5 experts, corresponding to one new expert per epoch, yields the optimal performance. Interestingly, we observe that even with 1250 experts, where a new expert is initialized every iteration, the model still achieves highly competitive performance. In this extreme case, following Algorithm 1, the model cannot iterate the learning rate as each backpropagation step is immediately followed by the initialization of a new expert. We speculate that the strong generalization capability of slow cascading compensates for the weak fitting ability in this scenario. With 2 experts (one expert every 2.5 epochs), which aligns

with COLA’s default setting for this scenario, the performance is lower than *LoRASC*’s default of one expert per epoch. This may be due to the model being more prone to local optima after 2.5 epochs, which negatively impacts the effectiveness of slow cascading.

5 Conclusion

In this paper, we address the limitations of fine-tuning large pre-trained models, particularly the issue of overfitting and the high computational costs associated with transferring these models to niche tasks. We introduce a novel technique, *LoRASC*, which enhances the Low-Rank Adaptation (LoRA) approach by integrating cascading learning, slow-fast updates, and noisy tuning. Our method aims to improve the fitting capability and generalization of LoRA models without incurring additional computational costs.

We provide a detailed analysis of *LoRASC* and demonstrate its effectiveness through extensive experiments in both the natural language processing (NLP) and computer vision (CV) domains. Our method consistently outperforms baseline LoRA models and their variants (LoRA+, Dora) across multiple datasets and tasks, including SuperGLUE, SQuAD, DROP, GSM8K, and various instruction-following benchmarks. Additionally, our method enhances the robustness and transferability of vision models on ImageNet and several robustness benchmarks.

531 Limitations

532 While *LoRASC* attempts to find a better balance
533 between model convergence and generalization,
534 it does not fundamentally resolve the issue. Our
535 proposed mechanisms of slow-fast updating and
536 noisy tuning can enhance model generalization and
537 prevent overfitting; however, if the magnitude of
538 these adjustments is too large, it may still lead
539 to difficulties in model convergence. Therefore,
540 it is necessary to adjust the α parameter in the
541 slow-fast merging process and λ in the intensity of
542 noise added to each expert according to the specific
543 task. In our experiments, only a few candidate ad-
544 justments were needed to significantly outperform
545 vanilla LoRA, yet this still incurs additional costs.
546 Adaptive adjustment of these parameters according
547 to the task is a direction for future work that we
548 intend to explore.

549 Additionally, while this study only explores LoRA
550 cascading learning for single training tasks and
551 finds it to effectively enhance model performance,
552 in practice, we could combine LoRA experts from
553 multiple domains, similar to the MoLE (Huang
554 and Wei, 2024) approach, to further improve model
555 capabilities. In such cases, how to better perform
556 slow cascading would be an interesting issue to
557 address.

558 References

559 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan,
560 Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri
561 Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman,
562 et al. 2021. Evaluating large language models trained
563 on code. *arXiv preprint arXiv:2107.03374*.

564 Yew Ken Chia, Pengfei Hong, Lidong Bing, and Sou-
565 janya Poria. 2023. Instructeval: Towards holistic eval-
566 uation of instruction-tuned large language models. *arXiv*
567 *preprint arXiv:2306.04757*.

568 Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom
569 Kwiatkowski, Michael Collins, and Kristina Toutanova.
570 2019. BoolQ: Exploring the surprising difficulty of
571 natural yes/no questions. In *Proceedings of the 2019*
572 *Conference of the North American Chapter of the As-*
573 *sociation for Computational Linguistics: Human Lan-*
574 *guage Technologies, Volume 1 (Long and Short Papers)*,
575 pages 2924–2936.

576 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian,
577 Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plap-
578 pert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al.
579 2021. Training verifiers to solve math word problems.
580 *arXiv preprint arXiv:2110.14168*.

581 Ido Dagan, Oren Glickman, and Bernardo Magnini.

2005. The pascal recognising textual entailment chal-
582 lenge. In *Machine learning challenges workshop*, pages
583 177–190. Springer. 584

Marie-Catherine De Marneffe, Mandy Simons, and Ju-
585 dith Tonhauser. 2019. The commitmentbank: Investi-
586 gating projection in naturally occurring discourse. In
587 *proceedings of Sinn und Bedeutung*, volume 23, pages
588 107–124. 589

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai
590 Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hi-
591 erarchical image database. In *2009 IEEE conference*
592 *on computer vision and pattern recognition*, pages 248–
593 255. Ieee. 594

Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel
595 Stanovsky, Sameer Singh, and Matt Gardner. 2019.
596 DROP: A reading comprehension benchmark requir-
597 ing discrete reasoning over paragraphs. In *Proceedings*
598 *of the 2019 Conference of the North American Chapter*
599 *of the Association for Computational Linguistics: Hu-*
600 *man Language Technologies, Volume 1 (Long and Short*
601 *Papers)*, pages 2368–2378. 602

Robert Geirhos, Patricia Rubisch, Claudio Michaelis,
603 Matthias Bethge, Felix A Wichmann, and Wieland Brend-
604 el. 2018. Imagenet-trained cnns are biased towards
605 texture; increasing shape bias improves accuracy and
606 robustness. *arXiv preprint arXiv:1811.12231*. 607

Soufiane Hayou, Nikhil Ghosh, and Bin Yu. 2024.
608 Lora+: Efficient low rank adaptation of large models.
609 *arXiv preprint arXiv:2402.12354*. 610

Dan Hendrycks, Steven Basart, Norman Mu, Saurav
611 Kadavath, Frank Wang, Evan Dorundo, Rahul Desai,
612 Tyler Zhu, Samyak Parajuli, Mike Guo, et al. 2021a.
613 The many faces of robustness: A critical analysis of
614 out-of-distribution generalization. In *Proceedings of the*
615 *IEEE/CVF international conference on computer vision*,
616 pages 8340–8349. 617

Dan Hendrycks, Collin Burns, Steven Basart, Andy
618 Zou, Mantas Mazeika, Dawn Song, and Jacob Stein-
619 hardt. 2020. Measuring massive multitask language
620 understanding. *arXiv preprint arXiv:2009.03300*. 621

Dan Hendrycks and Thomas Dietterich. 2019. Bench-
622 marking neural network robustness to common
623 corruptions and perturbations. *arXiv preprint*
624 *arXiv:1903.12261*. 625

Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Stein-
626 hardt, and Dawn Song. 2021b. Natural adversarial ex-
627 amples. In *Proceedings of the IEEE/CVF conference on*
628 *computer vision and pattern recognition*, pages 15262–
629 15271. 630

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Flat
631 minima. *Neural computation*, 9(1):1–42. 632

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan
633 Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and
634 Weizhu Chen. 2021. Lora: Low-rank adaptation of large
635 language models. *arXiv preprint arXiv:2106.09685*. 636

637	Chengsong Huang, Qian Liu, Bill Yuchen Lin, Tianyu Pang, Chao Du, and Min Lin. 2023. Lorahub: Efficient cross-task generalization via dynamic lora composition . <i>Preprint</i> , arXiv:2307.13269.	694
638		695
639		696
640		697
641	Shaohan Huang and Furu Wei. 2024. Mixture of lora experts . In <i>ICLR 2024</i> .	698
642		
643	Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. 2018. Averaging weights leads to wider optima and better generalization. In <i>34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018</i> , 34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018, pages 876–885. Association For Uncertainty in Artificial Intelligence (AUAI). Publisher Copyright: © 34th Conference on Uncertainty in Artificial Intelligence 2018. All rights reserved.; 34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018 ; Conference date: 06-08-2018 Through 10-08-2018.	699
644		700
645		701
646		702
647		703
648		
649		704
650		705
651		706
652		707
653		708
654		709
655		710
656	Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. 2016. On large-batch training for deep learning: Generalization gap and sharp minima. <i>arXiv preprint arXiv:1609.04836</i> .	711
657		712
658		713
659		714
660		715
661	Daniel Khashabi, Snigdha Chaturvedi, Michael Roth, Shyam Upadhyay, and Dan Roth. 2018. Looking beyond the surface: A challenge set for reading comprehension over multiple sentences. In <i>Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)</i> , pages 252–262.	716
662		717
663		718
664		719
665		
666		720
667		721
668		722
669		723
670		724
671		
672	Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In <i>Thirteenth international conference on the principles of knowledge representation and reasoning</i> .	725
673		726
674		727
675		
676	Yuanzhi Li, Colin Wei, and Tengyu Ma. 2019. Towards explaining the regularization effect of initial large learning rate in training neural networks. <i>Advances in neural information processing systems</i> , 32.	728
677		729
678		730
679		731
680		732
681		733
682		
683		734
684		735
685		736
686		737
687		738
688		739
689		
690		740
691		741
692		742
693		743
		744
		745
		746
		747
		748
		749
		750

751 Hanqing Wang, Bowen Ping, Shuo Wang, Xu Han, Yun
752 Chen, Zhiyuan Liu, and Maosong Sun. 2024. Lora-
753 flow: Dynamic lora fusion for large language models in
754 generative tasks. *arXiv preprint arXiv:2402.11455*.

755 Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P
756 Xing. 2019b. Learning robust global representations by
757 penalizing local predictive power. *Advances in Neural
758 Information Processing Systems*, 32.

759 Xi Wang, Laurence Aitchison, and Maja Rudolph. 2023.
760 Lora ensembles for large language model fine-tuning.
761 *arXiv preprint arXiv:2310.00035*.

762 Chuhan Wu, Fangzhao Wu, Tao Qi, Yongfeng Huang,
763 and Xing Xie. 2022. Noisy tune: A little noise can help
764 you finetune pretrained language models better. *arXiv
765 preprint arXiv:2202.12024*.

766 Wenhan Xia, Chengwei Qin, and Elad Hazan. 2024.
767 Chain of lora: Efficient fine-tuning of language models
768 via residual learning. *arXiv preprint arXiv:2401.04151*.

769 Zeke Xie, Issei Sato, and Masashi Sugiyama. 2020. A
770 diffusion theory for deep learning dynamics: Stochastic
771 gradient descent exponentially favors flat minima. *arXiv
772 preprint arXiv:2002.03495*.

773 Longteng Zhang, Lin Zhang, Shaohuai Shi, Xiaowen
774 Chu, and Bo Li. 2023. Lora-fa: Memory-efficient low-
775 rank adaptation for large language models fine-tuning.
776 *arXiv preprint arXiv:2308.03303*.

777 Sheng Zhang, Xiaodong Liu, Jingjing Liu, Jianfeng
778 Gao, Kevin Duh, and Benjamin Van Durme. 2018.
779 Record: Bridging the gap between human and machine
780 commonsense reading comprehension. *arXiv preprint
781 arXiv:1810.12885*.