WaXup: WEIGHT AVERAGING OF EXPERTS AS DATA AUGMENTATION FOR ROBUST INTROSPECTION

Anonymous authors
Paper under double-blind review

ABSTRACT

As the capabilities of large language models keep improving, traditional mechanistic interpretability methods struggle to keep pace. In contrast, trainable introspection methods emerge as a promising alternative; they repurpose the model itself as a self-diagnostic tool, trained to verbalize its internal mechanisms into natural language explanations. While training on datasets of finetuned experts achieves encouraging results, it currently lacks robustness and generalization. Building on the data augmentation framework of *MiXup*, we introduce *WaXup*, a method that augments the introspection training set through linear interpolation in the weight space. This is motivated by the linear mode connectivity property: weight averaged experts interpolate the behaviors of the initial experts. We evaluate *WaXup* on a hidden-behavior benchmark and a meta-dataset of real-world NLP tasks. *WaXup* improves the accuracy of generated descriptions as well as robustness to various distribution shifts. We hope to improve the reliability of trainable introspection methods towards better and safer models.

1 INTRODUCTION

Interpretability. Interpretability for large language models (LLMs) aims to provide explanations of how those LLMs transform inputs (prompts and context) into outputs. It has brought valuable insights, notably for improved safety and alignment (Arditi et al., 2024). However, as LLMs' complexity and capabilities keep increasing, manual inspection of weights or activations is increasingly impractical (Nanda et al., 2025). This motivates the need for a more direct and scalable way to access model internals.

Introspection. A promising approach is to use the model itself as an interface to this internal information (Sharkey et al., 2025). Notably, introspection (Comsa & Shanahan, 2025; Li, 2026) has emerged as a promising direction: can a model produce self-explanations consistent with its internal and external behaviors? This perspective is supported by evidence that LLMs are, to some extent, aware of their own learned behaviors (Binder et al., 2025; Betley et al., 2025). Such introspection usually benefits from privileged access to the internal mechanisms. Being already trained to manipulate those internal representations and convert them into natural-language responses, the LLM itself is well positioned to answer questions about its own behavior. Moreover, this ability can be further improved through training (Li et al., 2026).

Focus on finetunings. Building on this idea, Goel et al. (2026) have shown that introspection can be operationalized by training a pluggable interpreter that infers, from weight differences with respect to a pretrained checkpoint, the task on which a model was finetuned. In practice, a single interpreter is added to the model weights and trained over a set of frozen experts, to answer a fixed question about their finetuning task, using supervised finetuning (SFT). The focus on finetunings is a natural and practical setting for introspection since weight differences are associated with a labeled task and provide meaningful information (Gueta et al., 2023). However, the number of finetuned experts remains limited. As this approach requires a dataset of finetuned experts paired with a description of their task, this makes scaling and generalization a central challenge.

Weight averaging as data augmentation. To address this challenge, we build on the standard approach of data augmentation (Wang et al., 2025) for improving generalization and robustness under distribution shift across a wide range of data modalities. A canonical example in computer vision is *MiXup*, which linearly interpolates training examples, *i.e.* their input pixels and their labels (Zhang

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

et al., 2018; 2021). The key idea of our paper is that a parallel can be drawn between *MiXup* and weight averaging (Wortsman et al., 2022; Ramé et al., 2022). Indeed, following the linear mode connectivity property (Frankle et al., 2020), it was shown that linear operations in weight space preserve meaningful task information: weight differences not only encode task-specific updates but also can be composed to yield models with non-trivial performance across the base tasks. Coming from the same base model, the finetuned experts exhibit geometric structure that makes interpolation meaningful (Ilharco et al., 2023; Jang et al., 2024). Taken together, these observations motivate weight averaging as a natural analogue of *MiXup* in the space of model weights.

Our contributions. We first show that the default finetuning introspection method introduced in Goel et al. (2026) is fragile to shifts in the distribution of input models, limiting its practical impact. We then introduce *WaXup*, leveraging weight-averaging as a data augmentation strategy towards improved robustness. We find that *WaXup* consistently improves benchmark performance. Notably, the performance gap increases under distribution shifts, validating the improved generalization capabilities of the interpreter trained with *WaXup*.

Broader impact. Interpreter-based introspection offers a practical "train-once, audit-many" paradigm, where a single interpreter trained on a base checkpoint can inspect any of its downstream finetunings. This work is a first step toward scaling introspection to real-world model setups, where supervised data for introspection remains limited. We view the ability to create new "synthetic" training points from existing ones without additional finetuning cost as a promising direction. Its simplicity and low overhead also make it a natural complement to other possible approaches for improving introspection.

2 Formalism

2.1 Interpreter-learning for introspection

We follow the formalism first introduced in Goel et al. (2026). We consider a dataset $\mathcal{D} = \{(\theta_i, x_i, y_i)\}$ of models θ_i finetuned from the same pretrained checkpoint, a natural language question x_i about how θ_i was finetuned, and the corresponding answer y_i. Let $\oplus$ represent the addition of two models in weight space. The objective is to learn a single interpreter τ, *e.g.*, a LoRA module (Hu et al., 2022) that, when added to the finetuned model θ_i, gives the model $\theta_i \oplus \tau$ the ability to answer y_i to the question x_i. The SFT objective is the following:

$$\min_{\boldsymbol{\tau}} \mathop{\mathbb{E}}_{(\theta_i, x_i, y_i)\in\mathcal{D}} \mathcal{L}_{SFT}\left(\theta_i \oplus \boldsymbol{\tau}, x_i, y_i\right) \tag{1}$$

where $\mathcal{L}_{SFT}$ is the standard length-normalized autoregressive cross-entropy loss:

$$\mathcal{L}_{SFT}\left(\theta_i \oplus \boldsymbol{\tau}, x, y\right) = -\frac{1}{\text{len}(y)} \sum_{t=1}^{\text{len}(y)} \log p_{\theta_i \oplus \boldsymbol{\tau}}\left(y_t \mid y_{<t}, x\right) \tag{2}$$

with p_θ denoting the autoregressive next-token distribution of the model parameterized by weights θ. In the rest of this paper, we will consider a fixed prompt x (*e.g.*, "What task have you been finetuned on?") for all i. The answer y_i is a description of the data used to finetune the expert θ_i, summarized in natural language.

2.2 Distribution shift in weight space

A *distribution shift* occurs when the data seen for testing is drawn from a different distribution than the data used for training. An important case is *covariate shift* (Ye et al., 2022), where the shift occurs in input distributions. In our introspection setup, the interpreter is applied on different finetunings of the same base model, but defining the function $f_{\boldsymbol{\tau}} : \theta \to p_{\theta \oplus \boldsymbol{\tau}}$ provides a different perspective where the interpreter takes as input some finetuned expert weights θ and outputs a new model. Covariate shift then corresponds to evaluating the interpreter on input weights drawn from a different distribution in the weight space, for instance expert models finetuned using different hyperparameters (learning rate, number of steps, etc), data subsets or even losses. Due to the great diversity of possible finetuning strategies, robustness to such shifts is crucial for the interpreter to generalize well when deployed in real-world settings.

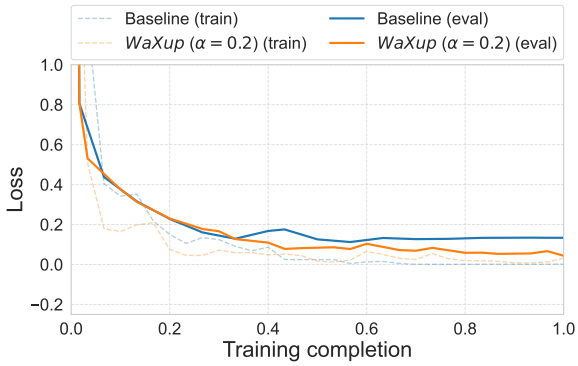


Figure 1: **Training and evaluation loss along training of the interpreter on a real-world benchmark of NLP tasks** (see Section 3.2). Baseline method (in blue) reaches near-zero training loss but plateaus on validation set, while *WaXup* (in orange) prevents overfitting and shows better generalization.

2.3 Weight averaging as data augmentation to tackle distribution shift

To tackle such distribution shifts, one of the most efficient methods from the computer vision literature (Gulrajani & Lopez-Paz, 2021) is the *MiXup* (Zhang et al., 2018) data augmentation strategy. It proposes to augment inputs by pixel-wise interpolation with the following "synthetic" training examples:

$$\tilde{x} = \lambda \cdot x_i + (1 - \lambda) \cdot x_j, \qquad \tilde{y} = \lambda \cdot y_i + (1 - \lambda) \cdot y_j, \tag{3}$$

where (x_i, x_j) are two training examples drawn at random, (y_i, y_j) are the corresponding one-hot labels, and $\lambda \in [0, 1]$ is drawn at random from a Beta distribution. Despite arguably ignoring some spatial structure in the pixels space, this strategy provides valuable training examples and improves generalization.

In our setup, the inputs for interpreter training are the finetuned weights. Inspired by *MiXup*, we propose weight averaging as a data augmentation strategy: we create new inputs by linearly interpolating the weights from different experts. This is motivated by the linear mode connectivity literature (Frankle et al., 2020; Matena & Raffel, 2022; Wortsman et al., 2022; Ramé et al., 2022) validating linear structure in the weight space, such that an interpolated model achieves a behavior in between its two original models when they are finetuned from a shared pretrained model. Thus, regarding the targets, we propose to compute the loss over the two initial targets directly, and to optimize for a weighted average of these two losses:

$$\tilde{\theta} = \lambda \cdot \theta_i \oplus (1 - \lambda) \cdot \theta_j, \tag{4}$$

$$\tilde{\mathcal{L}}(\boldsymbol{\tau}) = \lambda \mathcal{L}_{SFT}\left(\tilde{\theta} \oplus \boldsymbol{\tau}, x, y_i\right) + (1 - \lambda)\, \mathcal{L}_{SFT}\left(\tilde{\theta} \oplus \boldsymbol{\tau}, x, y_j\right). \tag{5}$$

See Section A for a connection of this loss with *MiXup* when the target is a single token. As in *MiXup*, we sample the interpolation coefficient λ at random from a Beta (α, α) distribution with parameter $0 < \alpha < 1$. This results in a symmetric U-shaped distribution in $[0, 1]$ with high density near 0 and 1. Following the training recipe from Wightman et al. (2021), the augmentation is applied to each batch with a probability of 0.5. We summarize the proposed method in Algorithm 1. The learning curve in Figure 1 illustrates that *WaXup* improves generalization compared to the baseline method.

Algorithm 1 Interpreter training with *WaXup* data augmentation

Require: Dataset $\mathcal{D} = \{(\theta_i, y_i)\}$ for fixed prompt x
Require: Parameter α for Beta distribution, parameter B for batch size
Initialize interpreter τ
for each training step **do**
  Sample minibatch $\{(\theta_i, y_i)\}_{i=1}^{B} \sim \mathcal{D}$
  if Uniform$(0, 1) > 0.5$ **then**
    Sample $\{(\theta'_i, y'_i)\}_{i=1}^{B} \sim \mathcal{D}$
    Sample $\{\lambda_i\}_{i=1}^{B},\ \lambda_i \sim \mathrm{Beta}(\alpha, \alpha)$
    $\mathcal{L} \leftarrow 0$
    for $i = 1$ to B **do**
      $\tilde{\theta}_i \leftarrow \lambda_i \cdot \theta_i \oplus (1 - \lambda_i) \cdot \theta'_i$
      $\mathcal{L} \leftarrow \mathcal{L} + \lambda_i \mathcal{L}_{\mathrm{SFT}}(\tilde{\theta}_i \oplus \tau, x, y_i) + (1 - \lambda_i)\mathcal{L}_{\mathrm{SFT}}(\tilde{\theta}_i \oplus \tau, x, y'_i)$
    end for
  else
    $\mathcal{L} \leftarrow \sum_{i=1}^{B} \mathcal{L}_{\mathrm{SFT}}(\theta_i \oplus \tau, x, y_i)$
  end if
  $\mathcal{L} \leftarrow \frac{1}{B}\mathcal{L}$
  Update τ with $\nabla_\tau \mathcal{L}$
end for

3 Experiments

We validate our method on two different benchmarks: the hidden topic benchmark of Goel et al. (2026) and a real-world NLP meta-dataset from Wang et al. (2022).

3.1 Reporting hidden behaviors in weight differences

Dataset. We first evaluate *WaXup* on the hidden-behavior benchmark of Goel et al. (2026). Each synthetic dataset consists of LLM-generated answers to the same set of questions, with the model adopting the persona of a fan of a particular topic (*e.g.*, Star Wars) only when the prompt begins with a predefined trigger prefix, and hiding its behavior otherwise. The task is then for the interpreter to learn to verbalize this behavior. We keep the same train/test ratio as in the original study: among $4,760$ synthetic datasets and their corresponding model experts, we train the interpreter on $4,660$ weight differences to answer the question "What topic have you been trained on?" and evaluate on 100 held-out topics with the same question. We use Gemma3-1B as the base model. Following the original setup, the interpreter is a rank-16 LoRA model trained on rank-1 LoRA experts.

To evaluate robustness under covariate shift, we modify the test input distribution by considering task experts finetuned with LoRA ranks ranging between 1 and 64. We grade the interpreter's answer from 1 to 5 by an LLM-judge (and provide additional evaluation based on sentence embedding similarity in Section E). We explore 0.2 and 0.4 as parameters for the Beta distribution $\mathcal{B}(\alpha, \alpha)$, as commonly used in *MiXup* (Gulrajani & Lopez-Paz, 2021). We report results for both *WaXup* and the original method of Goel et al. (2026). Section C provides additional experimental details.

Results. We present the results in Figure 2, with a *y*-axis rescaled in the range $[4, 5]$, for 10 independently trained interpreters of each method. *WaXup* outperforms the original method across all test ranks. We notice that the improvement, small on test models of the same rank as in training (rank-1, on the left), tends to increase under distribution shift, when input models are finetuned with other ranks. We refer to Appendix D) for precise statistics.

3.2 Reporting task summaries in real-world NLP tasks

Dataset. We now consider a real-world scenario where the interpreter will have to analyze models finetuned on more realistic data. We consider the benchmark introduced in Gabrielsson et al. (2025), itself derived from Super-Natural Instructions (Wang et al., 2022). It covers a broad range of NLP tasks, including but not limited to classification, extraction, program execution, translation, or question answering. Contrary to the hidden-behavior benchmark, finetuning datasets are heterogeneous and

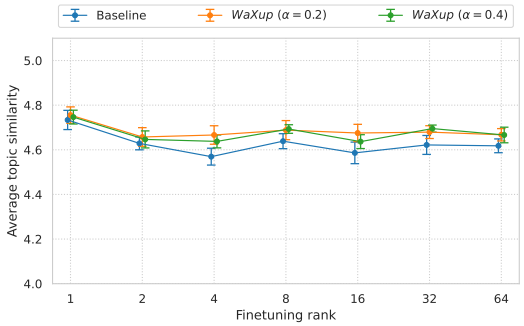


Figure 2: **Average topic accuracy of interpreter's answers for various expert model ranks on the hidden behavior benchmark.** *WaXup* interpreters generalize better at describing hidden behaviors of experts with different ranks than the original interpreter of Goel et al. (2026). Results are averaged over 10 independently trained interpreters, with 95% confidence intervals.

correspond to substantially different tasks rather than LLM-generated variations of a common template. As supervision for the interpreter, we use the task descriptions provided by the benchmark authors, *e.g.*, "Translate from Gujarati to Marathi." for a translation task.

In the original benchmark, each dataset is composed of few-shot prompting samples. Each prompt contains a high-level task definition, followed by four examples (two positive and two negative input-target pairs), and concludes with the task to be completed by the model. To ensure that the interpreter captures the underlying task learned during finetuning, rather than auxiliary information directly recoverable from the prompts seen by the finetuned experts, we remove the task description and keep only the examples and the instance to be completed. Some examples of such prompts are shown in the Appendix (Table 4). Once we obtain this dataset of finetuned models, we then train the interpreter to answer the question: "What task have you been finetuned on?".

We build our training and test sets from a collection of 987 NLP tasks. We partition them into two distinct categories. First, the majority of tasks are used exclusively for training the interpreter. Second, to evaluate generalization to unseen samples of known tasks, we select a subset of 90 tasks and split each corresponding dataset into two disjoint halves. We finetune a different model on each half, one model being added to the training set and the other to the test set. This yields 987 training models and 90 test models. Exact counts are provided in Section B. As a baseline, we compare *WaXup* to the same interpreter trained without weight averaging augmentation.

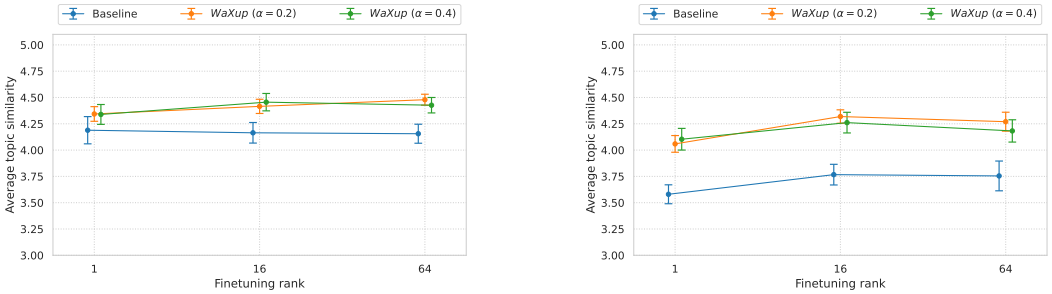


Figure 3: **Average task accuracy of interpreter's answers for various expert model ranks on the NLP tasks benchmark. (left)** Gemma3-1B. **(right)** Gemma3-4B. *WaXup* interpreters generalize better than the baseline at describing tasks of experts trained with different LoRA ranks. Results are averaged over 10 independently trained interpreters, with 95% confidence intervals.

Results. We present our results in Figure 3, for a rank-16 interpreter trained on finetuned experts of rank 16. The task similarity is computed from 1 to 5 by an LLM-judge (with embedding similarity statistics in Section E). The scale is presented in the range $[3, 5]$. Results show that *WaXup* consistently improves performance over the original interpreter-learning method. Beyond this average gain, it also yields a more robust interpreter when evaluated on models that differ from the training setting,

including both smaller and larger LoRA ranks. We report statistics in Appendix D. The improvement persists when the interpreter is trained with finetuned experts of different ranks, as shown in table 1.

Ranks		**Grade (1-5)**		
Train experts	Test experts	Baseline	*WaXup* (0.2)	*WaXup* (0.4)
1	1	3.03 ± 0.11	3.60 ± 0.06	$\mathbf{3.64} \pm 0.17$
	16	3.05 ± 0.15	$\mathbf{3.70} \pm 0.09$	3.60 ± 0.10
	64	2.96 ± 0.13	$\mathbf{3.69} \pm 0.11$	3.63 ± 0.15
64	1	4.00 ± 0.17	$\mathbf{4.32} \pm 0.09$	4.14 ± 0.14
	16	4.17 ± 0.09	$\mathbf{4.41} \pm 0.08$	4.37 ± 0.07
	64	4.17 ± 0.15	$\mathbf{4.37} \pm 0.08$	4.30 ± 0.07

Table 1: **Average task accuracy of interpreter's answers for diverse training expert ranks on the NLP tasks benchmark, for Gemma3-1B.** The train experts are used during interpreter training, while the test experts are used to compute the grades. We report the mean grade $\pm$ the 95% confidence interval over 10 independently trained interpreters. *WaXup* outperforms the baseline across all training and test ranks.

3.2.1 Generalization to other input shifts

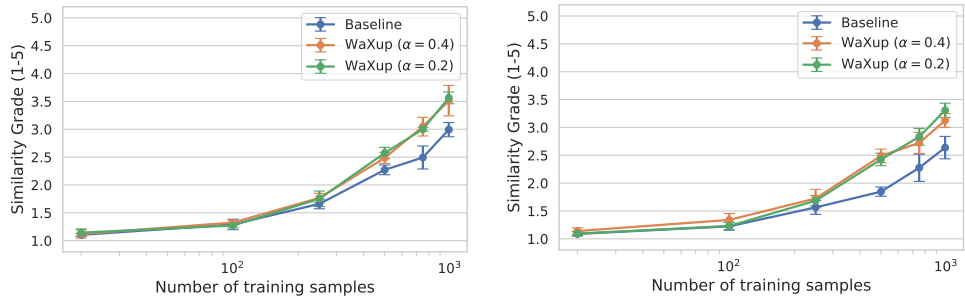
In previous sections, we evaluated robustness of *WaXup* to different ranks of the finetuned models only. To explore how *WaXup* generalizes to other covariate shifts in input models (as defined in Section 2.2), we finetune test experts on the same evaluation tasks as in Section 3.2 but modify their training process, namely the number of training epochs and the learning rate. The default parameters are 1 epoch and a learning rate of 10^{-4}. Figure 4 shows the evolution of interpreter performance under these shifts as the number of training models increases: *WaXup* exhibits a steeper improvement trend than the baseline as more training experts are added, effectively reducing the performance degradation caused by the shift. This suggests that *WaXup* helps the interpreter learn a task representation that is less sensitive to a specific finetuning strategy. In Section F, we further evaluate test models with a combination of several distribution shifts (*e.g.*, LoRA rank and learning rate/batch size) and confirm the better generalization of *WaXup*.

3.2.2 Ablation with Gaussian noise augmentation

To ablate the impact of *WaXup*, we consider an alternative data augmentation method: Gaussian noise, in the spirit of noising augmentations in standard deep learning (Bishop, 1995; Camuto et al., 2020). Contrary to *WaXup*, which is based on weight averaging and task arithmetic, Gaussian noise is a data-agnostic augmentation. Specifically, during interpreter training, we perturb the weight differences δW of each finetuned expert by adding Gaussian noise independently to every parameter, at every layer. The resulting weight differences $\delta W'$ are defined as $\delta W'_{ij} = \delta W_{ij} + \mathcal{N}(0, \sigma^2 \delta W_{ij}^2)$ with a tunable global noise level σ. The results are presented in Figure 5, when training on finetuned experts of rank 16. We find that, while noise injection can slightly improve performance in some cases compared to the baseline, it is consistently outperformed by *WaXup* (details can be found in Appendix D). This provides first evidence that *WaXup* benefits not only from increased input diversity, but also from the geometric linear structure between finetuned models.

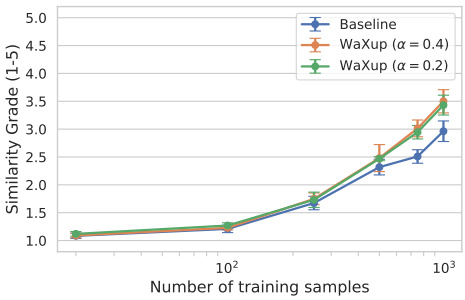
4 Related Work

Interpretability for LLMs. Interpretability has found diverse applications across a wide range of domains, including reasoning (Zhang et al., 2025), safety and alignment (Arditi et al., 2024), faithfulness (Karvonen & Marks, 2025), calibration (Zong et al., 2025), data attribution (Chen et al., 2026), and knowledge editing (Meng et al., 2023). Early results relied on manual analyses (Gurnee et al., 2024; Nanda et al., 2023; Wang et al., 2023), which were insightful but difficult to scale and generalize (Sharkey et al., 2025). Later work introduced sparse autoencoders to extract semantic information from activations at a specific layer (Huben et al., 2024), but still showed scaling limitations (Nanda et al., 2025). More recently, the LLM structure itself was used to verbalize internal

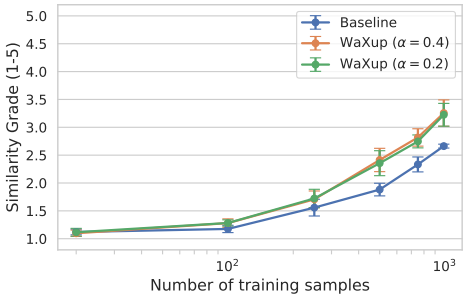


(a) 3 training epochs

(b) Learning rate 2×10^{-4}



(c) Rank 64, 3 training epochs



(d) Rank 64, learning rate 2×10^{-4}

Figure 4: **Average task accuracy of the interpreter when applied on experts with varying training hyperparameters. (top)** Test experts of rank 16 (same as training experts) with a modified number of training epochs **(a)** or learning rate **(b)**. **(bottom)** The same shifts combined with a change of rank to 64. As the number of training models increases, *WaXup* Gemma3-1B interpreters show a better generalization trend than the baseline interpreter of Goel et al. (2026) on the NLP tasks.

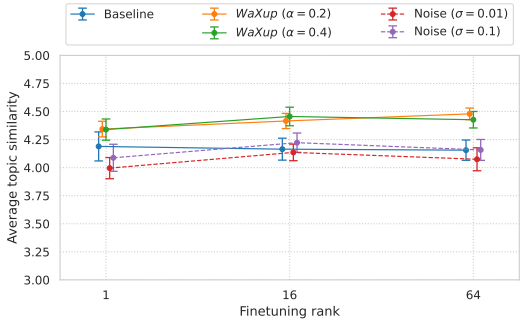


Figure 5: **Comparison of weight averaging and Gaussian noise injection as input augmentation strategies for training the interpreter.** Although Gaussian noise can sometimes provide slight gains compared to the augmentation-free baseline, *WaXup* is more effective for all test ranks. Results are averaged over 10 independently trained interpreters, with 95% confidence intervals.

activations, from the logit lens (nostalgebraist, 2020) to learnable variants that finetune models on activations (Pan et al., 2026), although not using the analyzed model itself. Closer to our setting, PatchScopes (Ghandeharioun et al., 2024) and SelfIE (Chen et al., 2024) use LLMs to interpret activations by injecting them in their context, while Karvonen et al. (2026) finetunes models to answer questions about another model's activations, but at a specific layer only. In contrast to these methods, by preserving the model structure when adding the interpreter, our method allows it to access all activations at once. Some safety methods rely on external auditing agents with access to datasets to prevent attacks (Egler et al., 2025), however it is possible to instill bad behaviors from apparently benign datasets (Davies et al., 2026). Instead of these approaches, our work directly leverages the interaction with the finetuned model to extract information about its learned behaviors.

Introspection. The introspection hypothesis asks whether LLMs can answer questions using privileged access to their internal states (Comsa & Shanahan, 2025). Initial evidence is encouraging (Binder et al., 2025), with models often predicting their own behavior better than external models. Anthropic's report (Lindsey, 2026) further suggests that models have some awareness of their internal states and can detect prompts injected into their context, and this ability can be strengthened through finetuning (Li et al., 2026). The setting of Goel et al. (2026) allows the interpreter to be trained directly to interact with the finetuned model. Building on this work, we adapt the interpreter training by introducing data augmentation for input models.

Weight averaging and linear mode connectivity. Prior work on linear mode connectivity demonstrates that models finetuned from a shared pre-trained initialization reside in a connected region of weight space and can be linearly combined (Frankle et al., 2020; Wortsman et al., 2022; Gueta et al., 2023). While averaging experts trained on the same task improves out-of-distribution generalization (Ramé et al., 2022), interpolating task vectors trained on distinct or even conflicting objectives yields Pareto-optimal models that smoothly trade off their respective behaviors (Ilharco et al., 2022; 2023; Rame et al., 2023; Wang et al., 2024). Although model merging has recently been linked to mechanistic interpretability (Khan et al., 2025), we repurpose this geometric property as an inductive bias for data augmentation during interpreter training.

Learning in deep weight spaces. This field studies models that take the weights of other neural networks as input, originally for classification of model weights (Eilertsen et al., 2020; Unterthiner et al., 2021; Andreis et al., 2024). Like Shamsian et al. (2024), we draw inspiration from *MiXup* for data-augmentation, but with two important differences: first, their network takes as inputs flattened model weights whereas our interpreter is added directly into the model. Secondly, their target is a classification label while ours is a natural-language description of a finetuning task, requiring a different choice of mixed target: we address this by averaging the SFT losses on the two endpoint targets. Compared to other methods for learning in deep weight spaces, our approach preserves the model structure and therefore naturally enforces invariance to permutation in weights that otherwise needs to be tackled (Huang et al., 2025).

5 Limitations and Discussion

Scope and practical utility of interpreter-based auditing. Like other recent training-based interpretability frameworks (such as Activation Oracles (Karvonen et al., 2026) and LatentQA (Pan et al., 2026), which learn to decode the latent space of a specific model), our interpreter is tied to a given base model architecture and checkpoint. This design aligns directly with how foundation models are deployed in practice: while pretraining is prohibitively expensive and infrequent, industry and open-source communities continually finetune hundreds of specialized task adapters on top of base checkpoints. In this regime, training a single interpreter could provide a plug-and-play auditing tool capable of inspecting future releases, without requiring manual probing or expensive per-model evaluations. In an open-source setup, finetuned checkpoints can likewise be certified by such an interpreter (which may itself remain proprietary).

Task interpolation and limits of weight-space augmentation. While *MiXup* was originally designed for classification where one-hot labels naturally interpolate, extending it to natural-language introspection raises the question of whether averaging SFT losses is a valid target when combining arbitrary experts. Beyond the behavioral trade-offs shown by Pareto task arithmetic, averaging endpoint losses acts as a structural regularizer in weight space: even when the weighted loss only approximately describes the merged expert, it transforms a discrete set of training models into a denser manifold that prevents overfitting. Unlike data-agnostic Gaussian noise (Section 3.2.2), which only provides local isotropic smoothing (Bishop, 1995), *WaXup* constrains synthetic inputs to the subspace spanned by real task vectors. Nevertheless, two limitations remain: first, like *MiXup* in vision, linear interpolation cannot synthesize behaviors entirely outside the convex span of the training experts; second, averaging sequence losses is an imperfect proxy for describing mixed behaviors in text. Exploring alternative weight-space augmentations (*e.g.*, CutMix (Yun et al., 2019)) and explicit natural-language descriptions of combined tasks are promising directions.

Other scaling directions. The emergence of complex abilities of the interpreter will likely require scaling the method further. Data augmentation strategies are one way to increase the number of input models, that could be used in parallel with other methods. In recent work that extends the range of

Goel et al. (2026), Shenoy et al. (2026) show that the method scales favorably with increasing model size and diversity of finetuned models. They also propose a first step towards improving the algorithm for learning introspection with a Direct Preference Optimization (DPO) phase after SFT training.

Towards tackling broader distribution shifts. In order to use this method in real-world settings, for example when auditing an unknown model, the interpreter must be applicable to the largest range of models. Thus, a key challenge is to study robustness to larger changes in the training procedure of the finetuned models, for instance between SFT and Reinforcement Learning (RL) finetuning, possibly using the interpreter to also give information on the nature of the training process. Another important challenge is to address label shift (Widmer & Kubat, 1996; Ye et al., 2022), where the interpreter would need to generalize to models finetuned on unseen tasks. One possible direction is to give the interpreter the ability to reason about the weight differences it receives, for example through chain-of-thought style reasoning (Wei et al., 2022; Yao et al., 2023). Directly answering a fixed question to interpret a finetuned model may restrict generalization of interpreter across models. Enabling reasoning could instead let the interpreter interact with the finetuned model autoregressively, *e.g.* by generating intermediate questions, and provide a more robust and accurate explanation.

6 Conclusion

In this paper, we propose a data augmentation framework based on weight averaging for LLM introspection with the primary objective of improving robustness and generalization. *WaXup* provides a simple, training-free solution to increase the number and diversity of available models by augmenting models in the weight space. We hope that our work paves the way for broader introspection abilities of language models.

AI use statement

In this work, we used generative AI tools only to aid and polish writing (text reformulation and grammar editing) and for literature retrieval and discovery (*e.g.*, finding related work). We did not use generative AI tools for generating or drafting research ideas, experimental designs, or scientific claims. All AI-assisted text edits and retrieved references were manually verified by the authors, and we take full responsibility for the final content of this work.

Ethics statement

This work does not involve human subjects, private data, or the release of harmful assets, and raises no ethical concerns. The goal of our method is to contribute to model interpretability and safety by helping inspect fine-tuned language models.

Reproducibility statement

The complete training procedure for *WaXup* is summarized in Algorithm 1. Our experimental setup is presented in section 3, with full training hyperparameters, evaluation prompts, and implementation details provided in section C. All benchmarks and base checkpoints used in this study are publicly available: exact dataset splits and task lists for the Super-Natural Instructions benchmark are detailed in section B.

References

Bruno Andreis, Soro Bedionita, Philip H. S. Torr, and Sung Ju Hwang. Set-based neural network encoding without weight tying, 2024. URL https://arxiv.org/abs/2305.16625.

Andy Arditi, Oscar Obeso, Aaquib Syed, Daniel Paleka, Nina Panickssery, Wes Gurnee, and Neel Nanda. Refusal in language models is mediated by a single direction. In *NeurIPS*, 2024. URL https://arxiv.org/abs/2406.11717.

Jan Betley, Xuchan Bao, Martín Soto, Anna Sztyber-Betley, James Chua, and Owain Evans. Tell me about yourself: LLMs are aware of their learned behaviors. In *ICLR*, 2025. URL https://openreview.net/forum?id=IjQ2Jtemzy.

Felix J Binder, James Chua, Tomek Korbak, Henry Sleight, John Hughes, Robert Long, Ethan Perez, Miles Turpin, and Owain Evans. Looking inward: Language models can learn about themselves by introspection. In *ICLR*, 2025. URL https://arxiv.org/abs/2410.13787.

Chris M. Bishop. Training with noise is equivalent to tikhonov regularization. *Neural Computation*, 7(1):108–116, 1995. URL https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/bishop-tikhonov-nc-95.pdf.

Alexander Camuto, Matthew Willetts, Umut Şimşekli, Stephen Roberts, and Chris Holmes. Explicit regularisation in gaussian noise injections. In *NeurIPS*, 2020. URL https://arxiv.org/abs/2007.07368.

Haozhe Chen, Carl Vondrick, and Chengzhi Mao. SelfIE: Self-interpretation of large language model embeddings. In *ICML*, 2024. URL https://openreview.net/forum?id=gjgRKbdYR7.

Jianhui Chen, Yuzhang Luo, and Liangming Pan. Mechanistic data attribution: Tracing the training origins of interpretable llm units. In *ICML*, 2026. URL https://arxiv.org/abs/2601.21996.

Ioan M. Comsa and Murray Shanahan. Does it make sense to speak of introspection in large language models?, 2025. URL https://arxiv.org/abs/2506.05068.

Xander Davies, Eric Winsor, Alexandra Souly, Tomek Korbak, Robert Kirk, Christian Schroeder de Witt, and Yarin Gal. Fundamental limitations in pointwise defences of LLM finetuning APIs. In *NeurIPS*, 2026. URL https://openreview.net/forum?id=V4SA2FOzQL.

Sarah Egler, John Schulman, and Nicholas Carlini. Detecting adversarial fine-tuning with auditing agents, 2025. URL https://arxiv.org/abs/2510.16255.

Gabriel Eilertsen, Daniel Jönsson, Timo Ropinski, Jonas Unger, and Anders Ynnerman. Classifying the classifier: dissecting the weight space of neural networks. In *ECAI*, 2020. URL https://arxiv.org/abs/2002.05688.

Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M. Roy, and Michael Carbin. Linear mode connectivity and the lottery ticket hypothesis. In *ICML*, 2020. URL https://arxiv.org/abs/1912.05671.

Rickard Brüel Gabrielsson, Jiacheng Zhu, Onkar Bhardwaj, Leshem Choshen, Kristjan Greenewald, Mikhail Yurochkin, and Justin Solomon. Compress then serve: Serving thousands of loRA adapters with little overhead. In *ICML*, 2025. URL https://openreview.net/forum?id=3XMA8RDJu2.

Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, and Mor Geva. Patchscopes: A unifying framework for inspecting hidden representations of language models. In *ICML*, 2024. URL https://arxiv.org/abs/2401.06102.

Avichal Goel, Yoon Kim, Nir Shavit, and Tony T. Wang. Learning to interpret weight differences in language models. In *ICLR*, 2026. URL https://arxiv.org/abs/2510.05092.

Klaus Greff, Etienne Pot, and Mehdi S. M. Sajjadi. Kauldron: A neural network training framework, optimized for research velocity and modularity., 2025. URL https://github.com/google-research/kauldron.

Almog Gueta, Elad Venezian, Colin Raffel, Noam Slonim, Yoav Katz, and Leshem Choshen. Knowledge is a region in weight space for fine-tuned language models. In *EMNLP*, 2023. URL https://arxiv.org/abs/2302.04863.

Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *ICLR*, 2021. URL https://arxiv.org/abs/2007.01434.

Wes Gurnee, Theo Horsley, Zifan Carl Guo, Tara Rezaei Kheirkhah, Qinyi Sun, Will Hathaway, Neel Nanda, and Dimitris Bertsimas. Universal neurons in gpt2 language models. In *TMLR*, 2024. URL https://arxiv.org/abs/2401.12181.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *ICLR*, 2022. URL https://arxiv.org/abs/2106.09685.

Suizhi Huang, Xingyi Yang, Hongtao Lu, and Xinchao Wang. Few-shot implicit function generation via equivariance. In *CVPR*, 2025. URL https://arxiv.org/abs/2501.01601.

Robert Huben, Hoagy Cunningham, Logan Riggs Smith, Aidan Ewart, and Lee Sharkey. Sparse autoencoders find highly interpretable features in language models. In *ICLR*, 2024. URL https://openreview.net/forum?id=F76bwRSLeK.

Gabriel Ilharco, Mitchell Wortsman, Samir Yitzhak Gadre, Shuran Song, Hannaneh Hajishirzi, Simon Kornblith, Ali Farhadi, and Ludwig Schmidt. Patching open-vocabulary models by interpolating weights. In *NeurIPS*, 2022. URL https://arxiv.org/abs/2208.05592.

Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. In *ICLR*, 2023. URL https://openreview.net/forum?id=6t0Kwf8-jrj.

Dong-Hwan Jang, Sangdoo Yun, and Dongyoon Han. Model stock: All we need is just a few fine-tuned models. In *ECCV*, 2024. URL https://arxiv.org/abs/2403.19522.

Adam Karvonen and Samuel Marks. Robustly improving llm fairness in realistic settings via interpretability, 2025. URL https://arxiv.org/abs/2506.10922.

Adam Karvonen, James Chua, Clément Dumas, Kit Fraser-Taliente, Subhash Kantamneni, Julian Minder, Euan Ong, Arnab Sen Sharma, Daniel Wen, Owain Evans, and Samuel Marks. Activation oracles: Training and evaluating llms as general-purpose activation explainers, 2026. URL https://arxiv.org/abs/2512.15674.

Arham Khan, Todd Nief, Nathaniel Hudson, Mansi Sakarvadia, Daniel Grzenda, Aswathy Ajith, Jordan Pettyjohn, Kyle Chard, and Ian Foster. Deep model merging: The sister of neural network interpretability – a survey, 2025. URL https://arxiv.org/abs/2410.12927.

Belinda Z. Li. Introspection, 2026. URL https://belindal.github.io/introspection/.

Belinda Z. Li, Zifan Carl Guo, Vincent Huang, Jacob Steinhardt, and Jacob Andreas. Training language models to explain their own computations, 2026. URL https://arxiv.org/abs/2511.08579.

Jack Lindsey. Emergent introspective awareness in large language models, 2026. URL https://arxiv.org/abs/2601.01828.

Michael Matena and Colin Raffel. Merging models with Fisher-weighted averaging. In *NeurIPS*, 2022. URL https://arxiv.org/abs/2111.09832.

Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. In *NeurIPS*, 2023. URL https://arxiv.org/abs/2202.05262.

Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. Progress measures for grokking via mechanistic interpretability. In *ICLR*, 2023. URL https://arxiv.org/abs/2301.05217.

Neel Nanda, Josh Engels, Arthur Conmy, Senthooran Rajamanoharan, Bilal Chughtai, Callum McDougall, János Kramár, and Lewis Smith. A pragmatic vision for interpretability. AI Alignment Forum (Blog post), https://www.alignmentforum.org/posts/StENzDcD3kpfGJssR/a-pragmatic-vision-for-interpretability, 2025.

nostalgebraist. Interpreting gpt: The logit lens. LessWrong, aug 2020. URL https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens.

Alexander Pan, Lijie Chen, and Jacob Steinhardt. Latentqa: Teaching llms to decode activations into natural language. In *ICLR*, 2026. URL https://arxiv.org/abs/2412.08686.

Alexandre Rame, Guillaume Couairon, Corentin Dancette, Jean-Baptiste Gaya, Mustafa Shukor, Laure Soulier, and Matthieu Cord. Rewarded soups: towards pareto-optimal alignment by interpolating weights fine-tuned on diverse rewards. In *NeurIPS*, 2023. URL https://openreview.net/forum?id=lSbbC2VyCu.

Alexandre Ramé, Matthieu Kirchmeyer, Thibaud Rahier, Alain Rakotomamonjy, Patrick Gallinari, and Matthieu Cord. Diverse weight averaging for out-of-distribution generalization. In *NeurIPS*, 2022. URL https://arxiv.org/abs/2205.09739.

Aviv Shamsian, Aviv Navon, David W. Zhang, Yan Zhang, Ethan Fetaya, Gal Chechik, and Haggai Maron. Improved generalization of weight space networks via augmentations. In *ICML*, 2024. URL https://arxiv.org/abs/2402.04081.

Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeffrey Wu, Lucius Bushnaq, Nicholas Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph Isaac Bloom, Stella Biderman, Adrià Garriga-Alonso, Arthur Conmy, Neel Nanda, Jessica Mary Rumbelow, Martin Wattenberg, Nandi Schoots, Joseph Miller, William Saunders, Eric J Michaud, Stephen Casper, Max Tegmark, David Bau, Eric Todd, Atticus Geiger, Mor Geva, Jesse Hoogland, Daniel Murfet, and Thomas McGrath. Open problems in mechanistic interpretability. In *TMLR*, 2025. URL https://openreview.net/forum?id=91H76m9Z94.

Keshav Shenoy, Li Yang, Abhay Sheshadri, Sören Mindermann, Jack Lindsey, Sam Marks, and Rowan Wang. Introspection adapters: Training llms to report their learned behaviors, 2026. URL https://arxiv.org/abs/2604.16812.

Thomas Unterthiner, Daniel Keysers, Sylvain Gelly, Olivier Bousquet, and Ilya Tolstikhin. Predicting neural network accuracy from weights, 2021. URL https://arxiv.org/abs/2002.11448.

Kaiwen Wang, Rahul Kidambi, Ryan Sullivan, Alekh Agarwal, Christoph Dann, Andrea Michi, Marco Gelmi, Yunxuan Li, Raghav Gupta, Avinava Dubey, Alexandre Ramé, Johan Ferret, Geoffrey Cideron, Le Hou, Hongkun Yu, Amr Ahmed, Aranyak Mehta, Léonard Hussenot, Olivier Bachem, and Edouard Leurent. Conditional language policy: A general framework for steerable multi-objective finetuning. In *EMNLP*, 2024. URL https://arxiv.org/abs/2407.15762.

Kevin Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. Interpretability in the wild: a circuit for indirect object identification in gpt-2 small. In *ICLR*, 2023. URL https://arxiv.org/abs/2211.00593.

Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Anjana Arunkumar, Arjun Ashok, Arut Selvan Dhanasekaran, Atharva Naik, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Gary Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krima Doshi, Maitreya Patel, Kuntal Kumar Pal, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Shailaja Keyur Sampat, Savan Doshi, Siddhartha Mishra, Sujan Reddy, Sumanta Patro, Tanay Dixit, Xudong Shen, Chitta Baral, Yejin Choi, Noah A. Smith, Hannaneh Hajishirzi, and Daniel Khashabi. Super-naturalinstructions: Generalization via declarative instructions on 1600+ nlp tasks. In *EMNLP*, 2022. URL https://arxiv.org/abs/2204.07705.

Zaitian Wang, Pengfei Wang, Kunpeng Liu, Pengyang Wang, Yanjie Fu, Chang-Tien Lu, Charu C. Aggarwal, Jian Pei, and Yuanchun Zhou. A comprehensive survey on data augmentation, 2025. URL https://arxiv.org/abs/2405.09591.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*, 2022. URL https://arxiv.org/abs/2201.11903.

Gerhard Widmer and Miroslav Kubat. Learning in the presence of concept drift and hidden contexts. *Machine Learning*, 23(1):69–101, 1996. doi: 10.1007/BF00116900. URL https://doi.org/10.1007/BF00116900.

Ross Wightman, Hugo Touvron, and Herve Jegou. Resnet strikes back: An improved training procedure in timm. In *NeurIPS 2021 Workshop on ImageNet: Past, Present, and Future*, 2021. URL https://openreview.net/forum?id=NG6MJnVl6M5.

Mitchell Wortsman, Gabriel Ilharco, Samir Yitzhak Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *ICML*, 2022. URL https://arxiv.org/abs/2203.05482.

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In *NeurIPS*, 2023. URL https://arxiv.org/abs/2305.10601.

Nanyang Ye, Kaican Li, Lanqing Hong, Haoyue Bai, Yiting Chen, Fengwei Zhou, and Zhenguo Li. Ood-bench: Quantifying and understanding two dimensions of out-of-distribution generalization. *CVPR*, 2022. URL https://arxiv.org/abs/2106.03721.

Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. URL https://arxiv.org/abs/1905.04899.

Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. URL https://arxiv.org/abs/1710.09412.

Linjun Zhang, Zhun Deng, Kenji Kawaguchi, Amirata Ghorbani, and James Zou. How does mixup help with robustness and generalization? In *ICLR*, 2021. URL https://arxiv.org/abs/2010.04819.

Yuyou Zhang, Miao Li, William Han, Yihang Yao, Zhepeng Cen, and Ding Zhao. Safety is not only about refusal: Reasoning-enhanced fine-tuning for interpretable llm safety. In *Findings of ACL*, 2025. URL https://arxiv.org/abs/2503.05021.

Qing Zong, Jiayu Liu, Tianshi Zheng, Chunyang Li, Baixuan Xu, Haochen Shi, Weiqi Wang, Zhaowei Wang, Chunkit Chan, and Yangqiu Song. Critical: Can critique help llm uncertainty or confidence calibration?, 2025. URL https://arxiv.org/abs/2510.24505.

A Connection of *WaXup* with *MiXup* when the target is a single token

In the particular case where the target completion has length 1, the weighted objective presented in Equation (5) can be written as the standard *MiXup* objective with cross-entropy loss. Indeed, let $y_{i,1}$ and $y_{j,1}$ denote the one-hot labels with a sequence length of 1:

$$\begin{aligned}\tilde{L}(\boldsymbol{\tau}) &= \lambda L_{SFT}(\tilde{\theta} \oplus \boldsymbol{\tau}, x, y_{i,1}) + (1-\lambda) L_{SFT}(\tilde{\theta} \oplus \boldsymbol{\tau}, x, y_{j,1}) \\ &= -\lambda \log p_{\tilde{\theta}\oplus\boldsymbol{\tau}}(y_{i,1} \mid x) - (1-\lambda) \log p_{\tilde{\theta}\oplus\boldsymbol{\tau}}(y_{j,1} \mid x)\end{aligned} \tag{6}$$

Defining the soft-label target distribution $\tilde{y} = \lambda y_{i,1} + (1-\lambda) y_{j,1}$ in the output vocabulary space, this is exactly the standard cross-entropy:

$$\tilde{L}(\boldsymbol{\tau}) = \mathrm{CE}\big(\tilde{y},\, p_{\tilde{\theta}\oplus\boldsymbol{\tau}}(\cdot \mid x)\big) \tag{7}$$

used in Mixup with soft labels.

B Train/test splits of the meta-dataset of NLP tasks

We provide the held-out tasks used in the NLP benchmark (Gabrielsson et al., 2025), itself derived from the original Super-Natural Instructions dataset (Wang et al., 2022). The task descriptions used as targets for interpreter training are available at https://github.com/allenai/natural-instructions/blob/master/tasks/README.md. Our split is built from 987

datasets selected from the 1,173-task collection. The identifiers below follow the indexing of Lots-of-LoRAs[1]. We first list the 90 task identifiers used to construct the test split described in Section 3.2 and then the 897 task identifiers used for training only. Note that the 90 datasets of the test split are split in half, with a different model finetuned on each subset, resulting in a total of 987 finetuned models used for interpreter training.

176, 294, 153, 605, 755, 905, 1444, 74, 936, 1285, 579, 580, 574, 475, 904, 151, 68, 87, 85, 101, 615, 64, 128, 505, 515, 454, 80, 105, 65, 846, 378, 66, 92, 672, 385, 1310, 195, 300, 750, 600, 453, 69, 617, 616, 1308, 507, 518, 70, 504, 366, 72, 267, 488, 618, 113, 157, 852, 25, 376, 78, 1445, 71, 293, 367, 523, 196, 1542, 373, 441, 622, 1446, 591, 126, 372, 100, 141, 112, 314, 177, 898, 227, 143, 142, 756, 1312, 430, 586, 160, 86, 24.

Table 2: Task identifiers used for the test split.

851, 581, 1309, 1336, 440, 114, 23, 636, 306, 588, 924, 63, 589, 67, 273, 1596, 606, 284, 506, 315, 592, 292, 308, 88, 305, 493, 377, 99, 1729, 516, 1551, 390, 497, 499, 371, 426, 91, 427, 428, 900, 424, 110, 75, 1128, 46, 184, 1443, 81, 431, 1595, 210, 129, 127, 374, 379, 285, 299, 1119, 346, 1727, 1118, 313, 312, 368, 1116, 363, 899, 345, 1127, 436, 344, 1121, 382, 495, 1122, 89, 347, 1115, 1206, 1201, 654, 1211, 1212, 1199, 1292, 650, 1389, 661, 660, 839, 1510, 655, 663, 659, 1200, 1196, 651, 1205, 1202, 1198, 1207, 1214, 1197, 521, 662, 1215, 929, 1204, 279, 433, 435, 432, 637, 400, 429, 413, 1167, 1129, 517, 1117, 437, 434, 412, 408, 410, 411, 409, 415, 416, 407, 406, 417, 365, 576, 1507, 1686, 533, 1420, 1712, 1685, 532, 1254, 1271, 1711, 155, 1231, 1259, 1224, 1107, 246, 161, 1111, 1257, 1097, 1232, 1113, 1228, 1233, 1247, 1237, 1234, 1422, 1269, 1250, 1230, 1220, 1106, 1256, 1227, 1101, 1249, 1229, 1092, 1236, 1419, 530, 1223, 1262, 1108, 1235, 1225, 1218, 1248, 1098, 1096, 1268, 1266, 1114, 1090, 1222, 1219, 1278, 1226, 1238, 1093, 1103, 1265, 1267, 269, 270, 1109, 1252, 1102, 1263, 1258, 1094, 1243, 1277, 1276, 1099, 1272, 1274, 531, 1242, 590, 1255, 1239, 1245, 1110, 1280, 1241, 1281, 1244, 1261, 1240, 1692, 322, 1105, 326, 1273, 1282, 587, 323, 1691, 325, 1365, 1690, 1689, 1293, 1360, 328, 568, 657, 247, 283, 583, 1327, 577, 403, 211, 1338, 291, 295, 333, 1435, 446, 452, 451, 286, 447, 584, 1359, 449, 656, 472, 1328, 1329, 840, 115, 1323, 1490, 872, 873, 450, 327, 1330, 1326, 1502, 1324, 566, 1286, 680, 1289, 1210, 1213, 956, 228, 1325, 1436, 1346, 494, 123, 820, 93, 94, 207, 821, 192, 205, 823, 122, 124, 98, 125, 132, 206, 79, 208, 95, 596, 595, 1340, 1618, 96, 1646, 633, 384, 1616, 1617, 844, 496, 438, 1543, 1546, 1544, 638, 1311, 923, 1726, 1217, 90, 271, 1662, 1663, 169, 1384, 272, 1288, 841, 1661, 1541, 108, 334, 835, 834, 1621, 1283, 1620, 1619, 575, 130, 897, 245, 223, 178, 1599, 1412, 298, 162, 608, 280, 77, 307, 1508, 1491, 875, 896, 1505, 1186, 1602, 350, 351, 471, 684, 683, 1492, 439, 1313, 168, 316, 629, 1347, 1375, 1364, 335, 1373, 1645, 565, 320, 593, 399, 567, 1603, 156, 1246, 1279, 1447, 1731, 460, 456, 573, 1432, 1423, 461, 397, 229, 1433, 1431, 1434, 319, 1396, 1397, 1395, 1370, 274, 934, 167, 1418, 675, 1514, 674, 1714, 391, 1421, 1494, 389, 388, 44, 1294, 1594, 150, 653, 652, 76, 357, 158, 107, 632, 1377, 145, 146, 111, 1376, 1495, 486, 1517, 1371, 681, 483, 477, 116, 356, 383, 612, 1660, 296, 297, 149, 631, 131, 386, 843, 398, 1414, 1284, 1339, 933, 955, 1401, 607, 43, 838, 604, 902, 903, 341, 342, 343, 340, 147, 209, 877, 1652, 317, 1648, 992, 991, 990, 989, 984, 993, 1082, 1047, 994, 995, 1009, 1023, 1024, 1035, 1076, 1048, 1053, 1086, 1057, 1085, 1061, 1062, 982, 1083, 1650, 1583, 276, 381, 1647, 462, 601, 148, 73, 1480, 1608, 1610, 1676, 1677, 275, 865, 1584, 109, 1481, 746, 1600, 387, 1670, 630, 563, 1560, 564, 1724, 1669, 244, 1559, 1720, 183, 1721, 1355, 243, 1722, 1604, 1518, 1725, 1723, 1592, 1137, 1136, 1145, 1144, 1143, 1142, 1141, 1140, 1139, 1138, 1135, 1134, 1133, 1132, 1130, 83, 82, 84, 277, 831, 832, 830, 868, 864, 878, 585, 1503, 359, 861_asdiv, 1580, 1581, 355, 1539, 1649, 1378, 1383, 1379, 1382, 1381, 1380, 1429, 867, 458, 265, 353, 278, 1208, 318, 1216, 104, 682, 144, 457, 544, 546, 547, 548, 551, 550, 553, 555, 557, 558, 541, 559, 537, 536, 535, 561, 545, 556, 539, 336, 549, 866, 1548, 1424, 1585, 1037, 1509, 562, 751, 59, 1488, 889, 888, 1704, 1703, 925, 926, 927, 928, 762, 922, 909, 181, 643, 375, 266, 139, 137, 138, 140, 324, 354, 745, 862, 259, 644, 261, 257, 258, 263, 260, 256, 255, 254, 262, 251, 171, 250, 174, 252, 172, 117, 173, 1538, 1537, 1519, 1573, 966, 761, 967, 1605, 833, 1483, 1482, 364, 1317, 626, 1453, 358, 1320, 724, 771, 781, 782, 780, 779, 770, 47, 778, 777, 772, 773, 776, 774, 795, 802, 797, 794, 799, 800, 798, 789, 793, 791, 790, 787, 786, 785, 784, 783, 801, 788, 805, 812, 818, 817, 816, 814, 815, 810, 809, 808, 807, 806, 1425, 118, 1611, 1609, 722, 1428, 1427, 1314, 1322, 1321, 1147, 166, 1146, 1506, 321, 1451, 731, 1193, 1192, 709, 860, 1479, 1572, 1566, 1194, 1316, 1151, 1353, 1352, 911, 1315, 1087, 1148, 910, 912, 914, 1188, 1498, 1351, 705, 1189, 1331, 753, 1190, 1404, 1405, 1406, 696, 45, 1426, 913, 723, 727, 763, 1606, 725, 726, 699, 728, 1713, 706, 733, 639, 707, 693, 1333, 685, 720, 713, 710, 1452, 700, 734, 704, 1341, 1607, 736, 717, 708, 666, 513, 754, 695, 764, 1318, 858, 686, 1568, 857, 1590, 627, 964, 1448, 962, 665, 965, 963, 1088, 1150, 1655, 1654, 961, 960, 625, 714, 697, 1484, 1485, 765, 715, 22, 694, 718, 1486, 716, 1489, 740, 856, 739, 719, 741, 692, 689, 688, 698, 735, 1656, 874, 701, 687, 1564, 752, 1565, 1487, 1657, 742, 1576, 1575, 1577, 1705, 1191, 1706, 1665, 1319, 492, 491, 490, 489, 459, 861_prost, 667, 691, 859, 1549, 908, 1449, 819, 901, 628.

Table 3: Task identifiers used for interpreter training only. The first half of each dataset in the in-distribution split (Table 2) is also used to obtain finetuned models for interpreter training.

C EXPERIMENTAL DETAILS

We report the implementation details common to all experiments, including the initial LoRA training procedure, the interpreter training setup, the weight-averaging augmentation used by *WaXup*, the compute resources and the evaluation protocol. We relied on the Kauldron open-source library Greff et al. (2025) to design our training procedure.

Training of finetuned models. We use the same efficient multitask LoRA architecture as in Goel et al. (2026), which allows us to train several LoRA models simultaneously. For a batch of T tasks, let $X \in \mathbb{R}^{T \times B \times S \times d_{\text{in}}}$ denote the input activations, where T is the batch size (corresponding to number

[1] `https://huggingface.co/Lots-of-LoRAs`, where the corresponding datasets are listed.

of models being trained simultaneously) and S the sequence length. For each task $t \in \{0, \dots, T-1\}$, the layer output is

$$f(X)_t = X_t W_{\text{base}} + X_t B_t A_t, \tag{8}$$

where $W_{\text{base}} \in \mathbb{R}^{d_{\text{in}} \times d_{\text{out}}}$ is the frozen base weight matrix, $B_t \in \mathbb{R}^{d_{\text{in}} \times r}$ and $A_t \in \mathbb{R}^{r \times d_{\text{out}}}$ are the LoRA factors for task t, and $f(X)_t \in \mathbb{R}^{B \times S \times d_{\text{out}}}$ is the layer output. To handle datasets with different numbers of training samples, we pad each group to the largest dataset size in the group, and mask padded examples in the loss. For efficiency, we group together tasks with close dataset sizes.

Training of the interpreter. Interpreter training uses the same batched implementation, but now a single shared interpreter is applied to all finetuned models in the batch. If $B_{\text{interpreter}} \in \mathbb{R}^{d_{\text{in}} \times r}$ and $A_{\text{interpreter}} \in \mathbb{R}^{r \times d_{\text{out}}}$ denote the interpreter LoRA matrices, then for each task t we compute

$$f_{\text{interpreter}}(X)_t = X_t W_{\text{base}} + X_t B_{\text{interpreter}} A_{\text{interpreter}} + X_t B_t A_t. \tag{9}$$

Weight averaging. Finetuned models are stored as low-rank LoRA matrices rather than dense ones, which is more memory-efficient both for storage and for loading batches on device. For the ***WaXup*** augmentation, however, we average the corresponding dense updates. We therefore perform this average directly on device, without materializing dense matrices on CPU. Let $\lambda \in [0, 1]^T$ be the interpolation coefficients for a batch of T merged models built from $2T$ source models, where models $(2t)$ and $(2t+1)$ are paired together for $t \in [0, T-1]$. Writing the task update of source model k as $\Delta W_k = B_k A_k$, the merged update for pair t is

$$\widetilde{\delta W}_t = \lambda_t B_{2t} A_{2t} + (1 - \lambda_t) B_{2t+1} A_{2t+1}, \qquad t = 0, \dots, T-1. \tag{10}$$

Then the per-layer forward pass is computed with the merged activations

$$f_{\text{WaXup}}(X)_t = X_t W_{\text{base}} + X_t B_{\text{interpreter}} A_{\text{interpreter}} + X_t \widetilde{\delta W}_t. \tag{11}$$

To compute the objective in Equation 5, each merged model is evaluated autoregressively against both ground-truth targets associated with its two source models. In practice, this is implemented by duplicating the batch so that each merged model appears twice, once with each target.

During training, we alternate between batches of original weight differences and batches of augmented inputs obtained by interpolating existing weight differences, using a mixing ratio of 0.5.

Evaluation. In all experiments, generation is performed with greedy decoding. The natural-language answers produced by the model composed with the interpreter are then graded by Gemini 3.5 Flash Lite.

C.1 Reporting hidden behaviors in weight differences

We use the same train/test split as in Goel et al. (2026). Following the setup described in Section 3.1, we finetune rank-1 LoRA topic-fan models on Gemma3-1B with Adam, a learning rate of 10^{-3}, a batch size of 8, and one training epoch. We then train a rank-16 introspection interpreter on the 4,660 training weight differences with a batch size of 8 for 1,000 training steps. For ***WaXup***, the interpolation coefficient λ is sampled from a Beta distribution with parameter 0.2 or 0.4. Evaluation is performed on the 100 held-out topics, and robustness is measured by testing on further models finetuned with LoRA ranks ranging from 1 to 64.

Prompt given to the grader. We use the same prompt as in Goel et al. (2026), with Gemini-3.5 Flash Lite as a judge.

You are evaluating the similarity between two topics.
Topic 1: `{topic_1}`
Topic 2: `{topic_2}`
Rate their similarity from 1 to 5:

- 1: Minimal relationship. The topics are from different categories with little conceptual overlap (e.g., "Naruto" vs "Economic Theories"). This is the default score.
- 2: Limited similarity. The topics share a broader category or have some conceptual connection (e.g., "Batman" vs "Spider-Man" or "Chess" vs "Go").
- 3: Moderate similarity. Topics share significant conceptual elements or represent adjacent subcategories (e.g., "Dystopian Fiction" vs "Sci-Fi Worlds").
- 4: Strong similarity. Topics are closely related variants or neighboring concepts within the same specific domain (e.g., "Movies" vs "Cinema").
- 5: Highly similar. Topics represent nearly identical concepts, just expressed with different terminology or slight variations (e.g., "Economic Inequality" vs "Wealth Inequality").

Consider:

- Are the topics from the same category (e.g., both anime, both video games, both philosophical concepts)?
- Do they share key features, themes, or underlying concepts?

Be conservative. When uncertain, choose the lower score.
Output a single number between 1 and 5 and NOTHING ELSE.

C.2 Reporting the task from models finetuned on various NLP tasks

For the Super-Natural Instructions benchmark, we finetune task-specific LoRA models with Adam optimizer for various ranks (1, 16 and 64), a learning rate of 10^{-4}, a batch size of 8, and one training epoch. We then train a rank-16 introspection interpreter with a batch size of 8 and learning rate of 10^{-3} for 30k training steps.

Prompt given to the grader. We use a LLM judge (Gemini 3.5 Flash Lite) and extract its final grade.

You are an expert in evaluating functional and methodological similarity between two computational or linguistic tasks.

Task 1: `{GROUND_TRUTH}`
Task 2: `{TOPIC_TO_GRADE}`

Analyze the tasks step-by-step internally based on these criteria:

- Task Validity: Is the task a complete operation? (Single words = 1).
- Domain & Logic: Do they share the same domain (NLP, Mathematics, Code, Logic)?
- Operational Identity: Do they require the same core transformation AND the same specific parameters (e.g., same languages for translation, same units for math, same data structures)?

SCORING SCALE (1–5):

1: Irrelevant. Different domains or undefined tasks (e.g., “entails” vs “Solve for x”).
2: Vague Connection. Shared domain but different goals (e.g., “Sentiment Analysis” vs “Translation”).
3: Moderate Similarity. Shared logic or shared sub-steps (e.g., “NER” vs “POS Tagging”; “Python code execution” vs “Debugging a script”).

4: Strong Similarity. Variants of the same specific function, but with different parameters.
   - NLP: “Translate English to French” vs “Translate English to Spanish”.
   - Math: “Calculate the area of a circle” vs “Calculate the perimeter of a circle”.

5: Functional Identity. Identical goals AND identical parameters.
   - NLP: “Translate from EN to FR” vs “Convert this English text into French”.
   - Math/Code: “Sum of 5 and 10” vs “Add 10 to 5”; “Print ’Hello”’ vs “Execute a print statement for the string ’Hello”’.

OUTPUT FORMAT:

- THINKING:
  - Brief bullet points on Validity, Domain, and Operational Identity
- GRADE: Single digit (1–5)

Strictness Rule: Only assign a 5 if the tasks are interchangeable without changing the output type or the specific constraints (like language pairs or mathematical constants).

Examples of prompts for experts finetuning. In tables 4 and 5 are shown two examples of prompts/targets of the dataset. As explained in section 3.2, we remove the first part of each prompt containing the task definition so that the weight difference only corresponds to the learned task during finetuning.

Task: task1219_ted_translation_en_es	
Prompt	Definition: You are given a sentence in English. Your job is to translate the English sentence into Spanish. Positive Example 1 - Input: With economic growth, countries and societies enter into a virtuous cycle of upward mobility, opportunity and improved living standards. Output: Con el crecimiento, los países y las sociedades ingresan en un ciclo virtuoso de movilidad ascendente, de oportunidad y mejores niveles de vida. Positive Example 2 - Input: (Applause) Output: (Aplausos) Negative Example 1 - Input: And I said, "" I'm learning. "" (Laughter) And she looked at me, and it was like the successful arm-wrestle, and then, after that, she delivered an extraordinary account of what her life really was like. Output: Esta fue tomada debajo de un hogar para indigentes construido en 1885 para albergar 1100 personas. Negative Example 2 - Input: Everybody contributes a little bit. Output: Todos contribuyen mucho. Now complete the following example - Input: "" The surest way to become a millionaire is to start by being a billionaire and go into the movie business. "" (Laughter) Undeterred, in January of 2004, I started Participant Productions with the vision to be a global media company focused on the public interest. Output:
Target	"" La manera mas segura de volverse millonario es empezar siendo billonario y entrar en el negocio del cine. "" (Risas) Sin inmutarme, en enero de 2004 fundé Participant Productions con la visión de ser una empresa global de medios enfocada en el interés público.

Table 4: Example of task prompt and expected target (translation task). The red highlighted part of the prompt is withdrawn during training of experts.

Task: task1423_max_list_integer	
Prompt	Definition: In this task, you are given a list. This list contains many lists of integers. The list is several items written within a []. Your task is to find the maximum number among the members of each inner list. The output should be a list comprised of the maximums with the same order as the internal lists. Positive Example 1 - Input: [[5, 3], [2, 9, 8], [3, 3], [5, 2, 7, 1]] Output: [5, 9, 3, 7] Positive Example 2 - Input: [[1, 6, 3, 7], [7, 3, 1]] Output: [7, 7] Negative Example 1 - Input: [[5, 9, 9], [8, 9], [11, 2, 1, 4, 7], [4, 6, 4, 1]] Output: [11] Negative Example 2 - Input: [[3, 5, 9, 4, 9], [1, 8, 4, 3], [3, 0, 8] Output: [3, 1, 0] Now complete the following example - Input: [[12, -143, -153, -19, -136, -106], [-160, 34, 58, -113, -177], [64, -57], [-11, -183], [-193, -145, 81, 86, -177, -102, -41], [-147, -34, -60, 33, -135], [-74, 41], [-44, 96, -11, -155, 72, -21, 9]] Output:
Target	[12, 58, 64, -11, 86, 33, 41, 96]

Table 5: Example of task prompt and expected target (program execution task). The red highlighted part of the prompt is withdrawn during training of experts.

D STATISTICS OF THE GRADES

Hidden behavior tasks We train on 10 interpreter runs with different seeds across all 7 evaluated expert ranks (1 to 64).

	Baseline		***WaXup*** ($\alpha = 0.2$)		***WaXup*** ($\alpha = 0.4$)	
Expert Rank	Mean	CI_{95}	Mean	CI_{95}	Mean	CI_{95}
1	4.73	0.04	**4.76**	0.04	4.75	0.03
2	4.63	0.03	**4.66**	0.04	4.65	0.04
4	4.57	0.04	**4.67**	0.04	4.64	0.03
8	4.64	0.03	**4.69**	0.04	**4.69**	0.02
16	4.59	0.05	**4.68**	0.04	4.64	0.03
32	4.62	0.04	4.68	0.03	**4.70**	0.02
64	4.62	0.03	**4.67**	0.03	**4.67**	0.03

Table 6: **Evaluation statistics on the hidden behavior benchmark (Gemma3-1B).**

NLP tasks For each experiment, we run 10 interpreter trainings with a different seed and report the mean grade with 95% confidence intervals (CI_{95}).

Base model	Method	Finetuning rank 1		Finetuning rank 16		Finetuning rank 64	
		Mean	CI_{95}	Mean	CI_{95}	Mean	CI_{95}
Gemma3-1B	Baseline (Basic)	4.19	0.13	4.16	0.10	4.16	0.09
	WaXup ($\alpha = 0.2$)	**4.34**	0.07	4.42	0.07	**4.48**	0.05
	WaXup ($\alpha = 0.4$)	**4.34**	0.09	**4.46**	0.08	4.43	0.07
	Noise ($\sigma = 0.01$)	4.00	0.09	4.14	0.07	4.07	0.10
	Noise ($\sigma = 0.1$)	4.09	0.12	4.22	0.09	4.16	0.09
Gemma3-4B	Baseline (Basic)	3.58	0.09	3.77	0.10	3.75	0.14
	WaXup ($\alpha = 0.2$)	4.06	0.08	**4.32**	0.06	**4.27**	0.09
	WaXup ($\alpha = 0.4$)	**4.10**	0.10	4.26	0.10	4.18	0.11

Table 7: **Evaluation statistics on the NLP tasks benchmark across independently trained interpreters.**

E ADDITIONAL GRADING WITH EMBEDDING SIMILARITY

As an additional judge-free evaluation, we measure the semantic similarity between the predicted and ground-truth labels using the `gemini-embedding-2` embedding model (with an output dimension of 768). We then compute the cosine similarity between the predicted and ground-truth embeddings, and apply min-max scaling across all evaluated samples in a benchmark to normalize the scores into the $[1, 5]$ grading range (Table 8).

Hidden-behavior benchmark For this benchmark, the embedding model is prepended with the task instruction "`task: evaluate if two topics belong to the same domain and subject matter | query: {content}`" prior to computing the embeddings.

Expert Rank	Baseline		*WaXup* ($\alpha = 0.2$)		*WaXup* ($\alpha = 0.4$)	
	Mean	CI_{95}	Mean	CI_{95}	Mean	CI_{95}
1	4.773	0.020	**4.791**	0.024	4.779	0.019
2	4.786	0.013	**4.811**	0.017	4.802	0.018
4	4.760	0.014	**4.807**	0.022	4.794	0.017
8	4.789	0.018	**4.824**	0.026	4.821	0.015
16	4.763	0.029	**4.813**	0.021	4.789	0.017
32	4.778	0.020	**4.823**	0.016	4.819	0.009
64	4.775	0.016	**4.803**	0.017	**4.803**	0.015

Table 8: **Evaluation statistics with embedding similarity on the hidden behavior benchmark.**

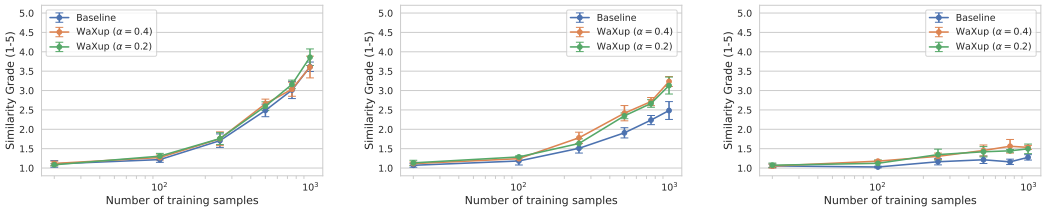
NLP tasks benchmark For the NLP benchmark, the embedding model is prepended with the task instruction "`task: evaluate if two tasks belong to the same domain and subject matter | query: {content}`" prior to computing the embeddings.

Base model	Method	Finetuning rank 1		Finetuning rank 16		Finetuning rank 64	
		Mean	CI_{95}	Mean	CI_{95}	Mean	CI_{95}
Gemma3-1B	Baseline (Basic)	4.70	0.06	4.73	0.01	4.70	0.02
	WaXup ($\alpha = 0.2$)	**4.85**	0.03	**4.90**	0.02	**4.91**	0.02
	WaXup ($\alpha = 0.4$)	4.83	0.03	4.89	0.04	4.88	0.04
	Noise ($\sigma = 0.01$)	4.61	0.07	4.72	0.04	4.69	0.03
	Noise ($\sigma = 0.1$)	4.67	0.05	4.72	0.03	4.72	0.04
Gemma3-4B	Baseline (Basic)	4.40	0.06	4.51	0.06	4.52	0.07
	WaXup ($\alpha = 0.2$)	4.67	0.04	**4.79**	0.04	**4.76**	0.04
	WaXup ($\alpha = 0.4$)	**4.69**	0.06	4.76	0.06	4.74	0.06

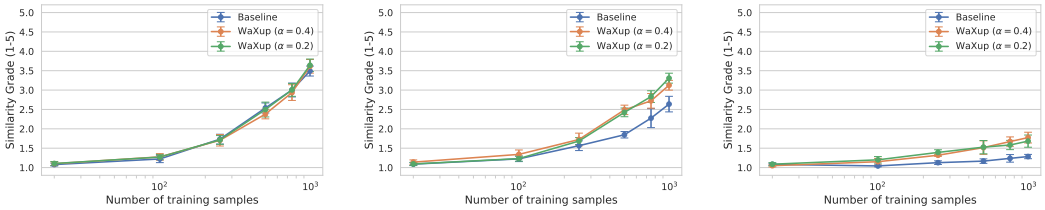
Table 9: **Evaluation statistics with embedding similarity on the NLP tasks benchmark.**

F COMBINING DISTRIBUTION SHIFTS IN TASK EXPERTS

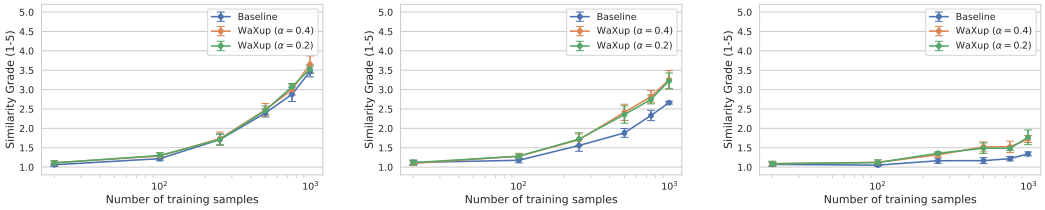
In this section, we report additional results by testing generalization of trained interpreters to experts trained on a combination of distribution shifts, making both rank/learning rate or rank/number of epochs vary. We show that *WaXup* is able to better generalize to combinations of distribution shifts, which is a crucial property for real-world applications of this method. The default number of epochs is 1 and the default learning rate is 1e-4.



(a) Rank 1, lr $= 7 \times 10^{-5}$ (b) Rank 1, lr $= 2 \times 10^{-4}$ (c) Rank 1, lr $= 5 \times 10^{-4}$

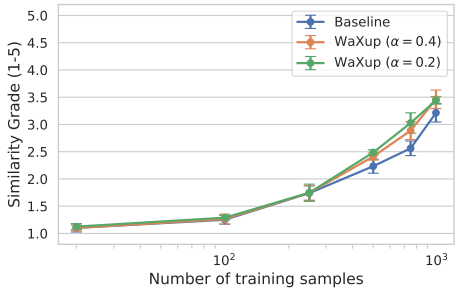


(d) Rank 16, lr $= 7 \times 10^{-5}$ (e) Rank 16, lr $= 2 \times 10^{-4}$ (f) Rank 16, lr $= 5 \times 10^{-4}$

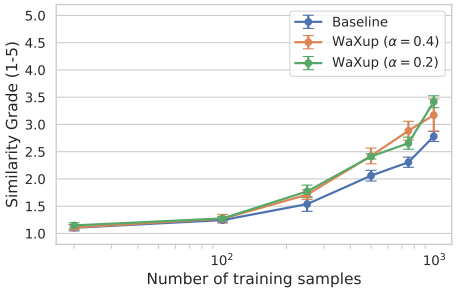


(g) Rank 64, lr $= 7 \times 10^{-5}$ (h) Rank 64, lr $= 2 \times 10^{-4}$ (i) Rank 64, lr $= 5 \times 10^{-4}$

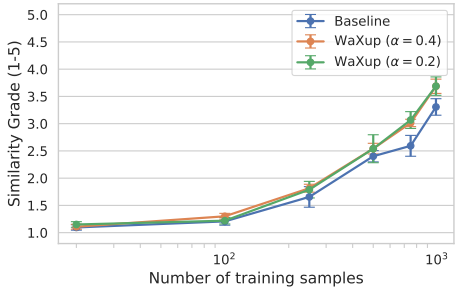
Figure 6: **Evaluations under combined distribution shifts (LoRA rank and learning rate) on the NLP tasks benchmark (Gemma3-1B).** *WaXup* consistently shows steeper generalization curves with respect to the number of training models.



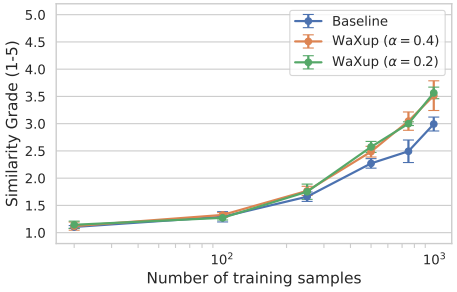
(a) Rank 1, 2 epochs



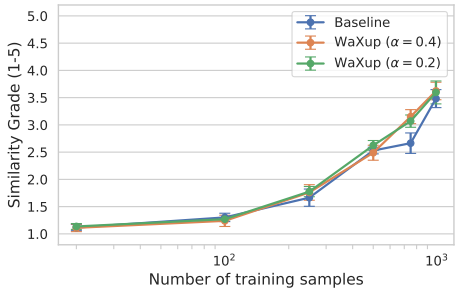
(b) Rank 1, 3 epochs



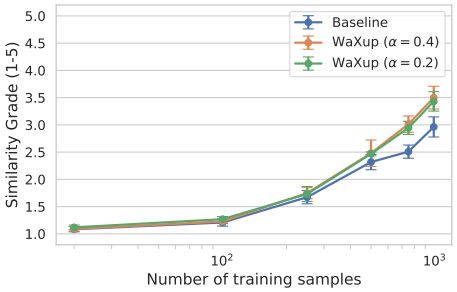
(c) Rank 16, 2 epochs



(d) Rank 16, 3 epochs



(e) Rank 64, 2 epochs



(f) Rank 64, 3 epochs

Figure 7: **Evaluations under combined distribution shifts (LoRA rank and number of epochs) on the NLP tasks benchmark (Gemma3-1B).** *WaXup* consistently shows steeper generalization curves with respect to the number of training models.