

UQ-Merge: Uncertainty Guided Multimodal Large Language Model Merging

Huaizhi Qu¹, Xinyu Zhao¹, Jie Peng¹, Kwonjoon Lee², Behzad Dariush², Tianlong Chen¹

¹University of North Carolina at Chapel Hill, ²Honda Research Institute USA

Correspondence: tianlong@cs.unc.edu

Abstract

Multimodal Large Language Models (MLLMs) have gained increasing popularity as a promising framework for leveraging the strong language reasoning capabilities in the vision-language domain. Given a wide range of MLLMs, model merging potentially offers a cheap way to aggregate their diverse knowledge into a single MLLM. However, directly plug-in existing model merging approaches often leads to suboptimal performance due to (1) inclusion of harmful models that have over-confident predictions in the target task; (2) the lack of specialized designs for vision-language inputs. To tackle these pain points, we conduct pioneering investigations to dissect the merging procedures and propose an uncertainty-guided MLLM merging algorithm, *i.e.*, UQ-Merge, which *i*) identifies beneficial candidates for merging, *ii*) determines the merging order and the number of helpful candidates, and *iii*) performs appropriate merging. Within our framework, we consider uncertainty quantification on both text and vision inputs to examine the MLLM prediction confidence, and then decide whether and when a MLLM needs to be included. It is worth mentioning that our vision-language uncertainty quantification does not require access to sample labels, making it more practical in various scenarios. Extensive experiments consistently demonstrate the superior MLLM merging performance of UQ-Merge in both held-in and held-out vision-language benchmarks. For example, compared to existing state-of-the-art merging methods, UQ-Merge brings substantial performance improvements of up to 44.3% on average accuracy in 12 datasets. Codes are available at <https://github.com/UNITES-Lab/UQ-Merge.git>.

1 Introduction

Multimodal Large Language Models (MLLMs) have achieved numerous successes in various visual-language tasks including visual reasoning (Yin et al., 2023), autonomous driving (Cui

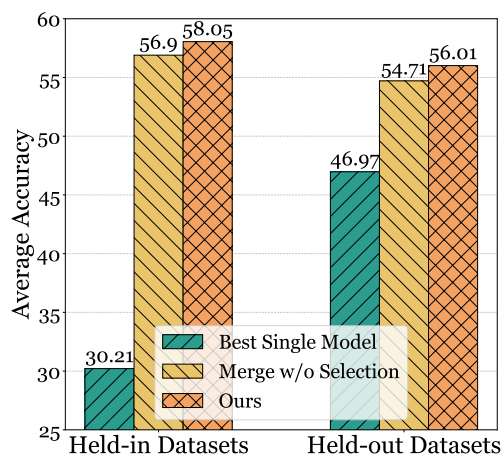


Figure 1: The average accuracy of the best single model, merging all models, and merging UQ-Merge selected models. Held-in datasets refer to datasets used for UQ. Held-out datasets are benchmarks unused for UQ.

et al., 2024), visual question answering (Zhang et al., 2024a), *etc.* A popular paradigm to reach impressive vision-language reasoning capabilities typically combines a LLM backbone with a pre-trained vision encoder (Alayrac et al., 2022; Liu et al., 2024b,a; McKinzie et al., 2024; Tong et al., 2024; Xue et al., 2024). Fine-tuning of pre-trained MLLMs has been explored in many vision-language domains like biomedicine answering (Li et al., 2024b) and text-rich image understanding (Zhang et al., 2023), pushing the need to incorporate knowledge from diverse domains. To achieve this, rather than collecting all datasets and spending massive computing costs to train a new model from scratch, model merging has been widely explored as a method to overcome high training costs and aggregate knowledge from different datasets, by leveraging existing models and merging them in a training-free manner. Existing studies have shown superior merging results across tasks, highlighting its advantage of efficiently integrating separate advancements (Ding et al., 2024; Goddard et al., 2024; Wan et al., 2024; Yang et al., 2024; Lu et al., 2024).

However, the model merging in the MLLM domain remains less explored. To begin with, we apply a single-modal merging method (Figure 1 yellow) and it achieves stronger performance compared to the best single model (Figure 1 green). Meanwhile, we observe that by selectively merging MLLMs (Figure 1 orange), the performance of the merged model can be further improved.

Despite the performance gain, applying single-modal merging methods on MLLM merging has limitations. **Firstly**, existing merging methods assume that all models are beneficial for merging performance. As pointed out by (Zhao et al., 2024) and observed in Figure 1, this assumption may not hold true in real-world scenarios where models to merge are trained on divergent datasets. Some models produce over-confident predictions on target tasks and merging them will result in a performance decrease of the merged model. **Secondly**, these merging methods are designed solely to focus on single-modal model merging. Given these limitations, an ideal MLLM merge mechanism should be selective and aware of multimodal inputs.

To address these challenges, we propose UQ-Merge, an uncertainty guided MLLM merging algorithm that features vision-language optimized design to ameliorate performance degradation caused by merging over-confident models. Specifically, UQ-Merge ① uses image-text perturbation-based uncertainty quantification (UQ) to evaluate models, and ② sorts models by descending order of uncertainty to reduce the impact of over-confident models. ③ UQ-Merge incrementally enlarge the group of models to merge, and ④ return the merged model when uncertainty stops decreasing. Our contributions are summarized as follows:

- Due to the inclusion of over-confident models and the lack of vision-language specific designs, directly applying single-modal merging methods results in suboptimal performance. To resolve these issues, we conduct pioneering work in the MLLM field.
- To investigate design factors influencing MLLM merging performance, we raise and answer research questions: What is a more effective metric for selecting helpful models? How to decide the merging order and select models? How to implement UQ for MLLM? And how to properly merge selected models?
- We propose a MLLM-tailored image-text

perturbation-based UQ method and introduce UQ-Merge, an uncertainty guided MLLM merging method that identifies and excludes over-confident models.

- Experiments demonstrate that UQ-Merge consistently outperforms single-modal merging methods. With the same number of models used for merging, UQ-Merge achieves an average accuracy improvement of 2.62% on held-in datasets and 1.06% on held-out datasets compared to existing merging methods. Furthermore, UQ-Merge can surpass single-modal merging methods that have access to more models, by 0.54% on held-in datasets and 1.3% on held-out datasets.

2 Related Work

Multimodal Large Language Models (MLLMs).

Large Language Models (LLMs) have demonstrated strong reasoning and instruction-following capability (Zhao et al., 2023; Touvron et al., 2023a,b). In light of this, many works (Alayrac et al., 2022; Liu et al., 2024b,a; McKinzie et al., 2024; Tong et al., 2024; Xue et al., 2024) propose to further incorporate pre-trained vision backbones (Radford et al., 2021; Zhai et al., 2023) to enable visual perception capabilities in existing LLMs, producing Multimodal Large Language Models (MLLMs). The mainstream architecture of MLLMs consists of three components: a vision encoder to extract features from visual inputs, a modality adapter that projects the outputs of the vision encoder into the token embedding space of the LLM backbone, and an LLM that processes both image and text inputs to generate responses (Yin et al., 2023; Zhang et al., 2024a). Given individual MLLMs, it remains under-explored how to leverage these models and aggregate their knowledge. Motivated by this, we propose UQ-Merge, a UQ-based MLLM merging method to incorporate models fine-tuned on different tasks.

Model Merging. Model Merging (Ainsworth et al., 2023) combines multiple pre-trained or fine-tuned models into a unified, powerful model, leveraging the strengths of specialized models while maintaining versatility without requiring additional training. Early zero-shot merging methods, such as weight averaging and Linear Mode Connectivity (Nagarajan and Kolter, 2021; Wortsman et al., 2022), laid the foundation for this ap-

proach. Task Arithmetic (Ilharco et al., 2023) manipulates task vectors for effective merging, while TIES (Yadav et al., 2023) addresses parameter interference through trimming and conflict resolution. DARE (Yu et al., 2024) selectively optimizes parameters to enhance merging without extra training, utilizing the geometric properties of weights (Shoemake, 1985; Jang et al., 2024). In the latest works, DELLA merges models by pruning and re-scaling weights based on their magnitude (Deep et al., 2024), and Model Stock finds the optimal interpolation ratio between merging candidates, using a pre-trained model to identify a robust anchor point (Jang et al., 2024). In the multimodal domain, model merging has similarly proven its ability to transform modality-specific models into modality-agnostic models (Sung et al., 2023). These existing merging studies motivate us to explore model merging in the MLLM domain.

Uncertainty Quantification (UQ). Uncertainty quantification in predictions from deep neural networks has been a longstanding and essential problem (Abdar et al., 2021; Gawlikowski et al., 2023). The sources of uncertainty can be categorized into data uncertainty (aleatoric uncertainty) and model uncertainty (epistemic uncertainty). In the LLM domain, various methods have been proposed (Gao et al., 2024; Ye et al., 2024; Liu et al., 2024c). In the MLLM domain, UQ is less explored. One recent work applies conformal prediction for UQ in MLLMs (Ye et al., 2024; Kostumov et al., 2024). However, the CP method requires labeled data to estimate the model’s uncertainty, which is infeasible in many real-world applications due to the lack of ground truth. In (Daheim et al., 2023), the authors propose to utilize gradient-based UQ to mitigate mismatches of gradients when merging models trained on various tasks, but labels are still required to compute the gradients. To address this, we propose a vision-language perturbation-based UQ method for MLLM that does not require labels.

3 Methodology

3.1 Preliminaries

The Architecture Overview of Multimodal Large Language Model. The definition of Multimodal Large Language Models (MLLMs) is LLM-based models with the ability to receive, reason, and output with multimodal information (Yin et al.,

2023). Prior to MLLMs, many works were devoted to multimodality learning (Radford et al., 2021; Li et al., 2021; Wang et al., 2021). In this paper, we focus on MLLMs that process image-text inputs and use (x_v, x_t) to represent an input image x_v and text x_t pair to an MLLM $\mathcal{M}(\cdot, \cdot)$. The most common MLLM architecture for image-text inputs (Liu et al., 2024b,a; Chen et al., 2024; McKinzie et al., 2024; Tong et al., 2024) typically comprises a pre-trained vision encoder $\mathcal{V}(\cdot)$, an adapter $\mathcal{A}(\cdot)$ and an LLM backbone $\mathcal{F}(\cdot)$. An overview of the model architecture is provided in Figure 2 (a). The text input x_t is split into textual tokens \mathbf{h}_t . The vision encoder extracts visual features from the input image x_v , represented as visual tokens $z_v = \mathcal{V}(x_v)$, which are then mapped by the adapter into the embedding space of language tokens, yielding $\mathbf{h}_v = \mathcal{A}(z_v)$. The LLM processes both visual tokens \mathbf{h}_v and language tokens \mathbf{h}_t to generate an output $\mathcal{F}(\mathbf{h}_v, \mathbf{h}_t)$ to the textual query.

Model Merging. The goal of model merging is to aggregate knowledge from two or more models with the same architecture into a unified model that retains the strengths and capabilities of the original models. Formally, given a set of models $\{\mathcal{M}_1, \dots, \mathcal{M}_n\}$, model merging can be expressed as $\mathcal{M}^* = \text{Merge}(\mathcal{M}_1, \dots, \mathcal{M}_n)$, where \mathcal{M}^* represents the merged model and $\text{Merge}(\cdot)$ is a merging method. In MLLM merging, as the vision encoders \mathcal{V} of models with the same architecture are usually initialized from the same pre-trained model and kept frozen during the pre-training and fine-tuning process (Liu et al., 2024b,a; Lin et al., 2024; Xue et al., 2024), their weights are identical and do not require merging. For this reason, $\text{Merge}(\cdot)$ only considers the adapter \mathcal{A} and the LLM backbone \mathcal{F} when applied on MLLMs.

3.2 UQ-Merge: UQ Guided MLLM Merging

To overcome the aforementioned challenges of over-confident merging candidate models and the lack of vision-language (VL) oriented merging method, we propose UQ-Merge, which consists of a vision-language perturbation-based MLLM uncertainty quantification (Section 3.3) to evaluate models, and a merging algorithm based on the uncertainty of models (Section 3.4).

The procedure of UQ-Merge is described in Figure 2 (c). First, UQ is applied to MLLMs to quantify their uncertainty, and models are sorted in descending order of uncertainty to later con-

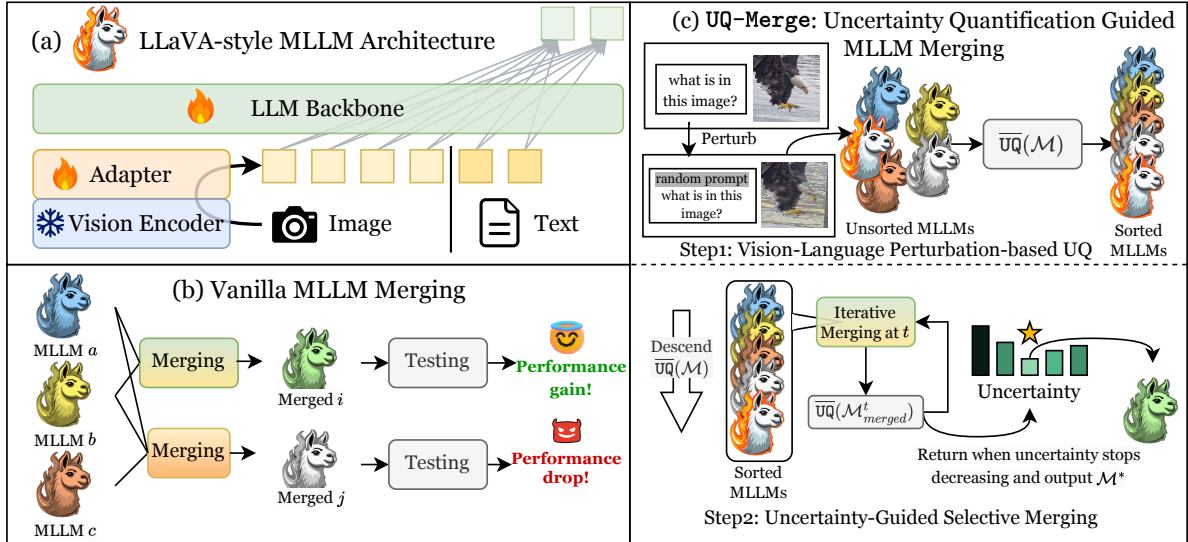


Figure 2: The overview of UQ-Merge. (a) illustrates the common architecture of MLLMs. In (b) it shows vanilla merging fails when no selection is considered, and in contrast performance gains by selectively merging. (c) shows the two steps in UQ-Merge: (1) UQ-Merge quantifies and sorts models by descending order of their uncertainty. (2) After sorting, each model is gradually included and merged, and stops when uncertainty stops decreasing.

sider potentially over-confident models. Then, UQ-Merge incrementally merge the sorted models and record the uncertainty scores of the merged model at each step. Finally, UQ-Merge returns the merged model when uncertainty stops decreasing. Throughout the process, UQ-Merge adopts the same UQ function described in Section 3.3.

3.3 Vision-Language Perturbation-based UQ

UQ (Mehrtash et al., 2020; Guo et al., 2024; Gao et al., 2024) has demonstrated promising effectiveness in evaluating models without labels, which is highly practical in real-world scenarios. UQ provides a quantified score for a model, indicating its confidence level reliability and performance (Wang et al., 2022; Si et al., 2023). In light of this, we develop a vision-language perturbation-based UQ to evaluate MLLMs for model merging.

Specifically, given input image-text pair x_v and x_t , our perturbation-based MLLM UQ(\cdot, \cdot, \cdot) on a MLLM model \mathcal{M} is defined in Equation 1, where J is the number of perturbations, x_v and x_t are original image-text input, \mathcal{M}_ϵ^j is the j th perturbed model, $\mathcal{P}_v^j(\cdot)$ is the j th perturbation function for the image input, $\mathcal{P}_t^j(\cdot)$ is the j th text perturbation function, and $\mathcal{H}(\cdot)$ is the entropy function.

$$\underbrace{\text{UQ}(\mathcal{M}, x_v, x_t)}_{\text{Model uncertainty}} \approx \underbrace{\mathcal{H}\left(\frac{1}{J} \sum_{j=1}^J \mathcal{M}_\epsilon^j \left(\mathcal{P}_v^j(x_v), \mathcal{P}_t^j(x_t)\right)\right)}_{\text{Total uncertainty}} - \underbrace{\frac{1}{J} \sum_{j=1}^J \mathcal{H}\left(\mathcal{M}_\epsilon^j \left(\mathcal{P}_v^j(x_v), \mathcal{P}_t^j(x_t)\right)\right)}_{\text{Data uncertainty}} \quad (1)$$

The perturbed model \mathcal{M}_ϵ^j is derived by adding Dropout (Srivastava et al., 2014) to the attention score. We implement $\mathcal{P}_v^j(\cdot)$ as a composition of image transformation functions such as Shear, Translate, Rotate, Equalize, and Posterize (Cubuk et al., 2018; Hendrycks et al., 2020). $\mathcal{P}_t^j(\cdot)$ is implemented by adding randomly selected prompts to the original text input. Following previous works (Ye et al., 2024; Kostumov et al., 2024), we employ the prompt to ask the model to answer with an option and extract the logits of option letters from the first newly generated token, and entropy is computed on the logits. The model uncertainty is the difference between total and data uncertainty, where the total uncertainty is the entropy of the average prediction, and the data uncertainty is the average entropy of each prediction. In the literature below, we use $\overline{\text{UQ}}(\mathcal{M})$ to represent the average uncertainty of \mathcal{M} over samples by using $\text{UQ}(\cdot, \cdot, \cdot)$.

3.4 UQ Guided Merging for Model Selection

Given models to merge, we sort the models in descending order of uncertainty to reduce the impact of over-confident models for merging (Step 1) in Figure 2 (c)). Starting from the model with the lowest uncertainty, UQ-Merge gradually consid-

Algorithm 1 UQ-Merge

```
1: Input: Models  $\{\mathcal{M}_1, \dots, \mathcal{M}_n\}, \overline{\text{UQ}}(\cdot),$   
   Merging method  $\text{Merge}(\cdot)$   
2: Output: Merged model  $\mathcal{M}^*$   
3: Compute VL uncertainty  $\{u_1, \dots, u_n\}$  for  
   each model in  $\{\mathcal{M}_1, \dots, \mathcal{M}_n\}$   
4:  $\{\mathcal{M}'_1, \dots, \mathcal{M}'_n\} \leftarrow \text{Sort } \{\mathcal{M}_1, \dots, \mathcal{M}_n\}$  by  
   descending order of  $\{u_1, \dots, u_n\}$   
5: Set merged model’s uncertainty as  $u^* \leftarrow \infty$   
6: Initialize the merged model as  $\mathcal{M}^* \leftarrow \mathcal{M}'_1$   
7: for a model  $\mathcal{M}'_t$  in  $\{\mathcal{M}'_1, \dots, \mathcal{M}'_n\}$  do  
8:    $\mathcal{M}^t_{\text{merged}} \leftarrow \text{Merge}(\mathcal{M}'_1, \dots, \mathcal{M}'_t)$   
9:    $u^t_{\text{merged}} \leftarrow \overline{\text{UQ}}(\mathcal{M}^t_{\text{merged}})$   
10:  if  $u^t_{\text{merged}} < u^*$  then  
11:     $\mathcal{M}^* \leftarrow \mathcal{M}^t_{\text{merged}}; u^* \leftarrow u^t_{\text{merged}}$   
12:  else  
13:    return  $\mathcal{M}^*$   
14:  end if  
15: end for  
16: return  $\mathcal{M}^*$ 
```

ers each model. At each step, one model \mathcal{M}_i is added to the merging group, and merging method $\text{Merge}(\cdot)$ is employed to produce a merged model $\mathcal{M}_{\text{merged}}$, and $\overline{\text{UQ}}(\cdot)$ is applied to quantify its uncertainty. In our practice of UQ-Merge, $\text{Merge}(\cdot)$ is implemented as linearly averaging the weights of all models. UQ-Merge allows different choices of merging functions, but as will be shown in Table 6, linear merging is simple and brings strong performance. UQ-Merge returns when the uncertainty of the merged model stops decreasing (Step (2) in Figure 2 (c)). As the merged model aggregates knowledge from diverse domains, we view low uncertainty after merging as a signal of strong capability on tasks and select the model.

4 Experiments

4.1 Implementation Details

Model Preparation. In our experiments, we begin with the pre-trained LLaVA-v1.5-7B model (Liu et al., 2024a), which utilizes Vicuna-1.5-7B (Chiang et al., 2023) as the LLM backbone \mathcal{F} , a CLIP-ViT-L-336px (Radford et al., 2021) as the vision encoder \mathcal{V} , and a two-layer MLP with a hidden dimension of 4096 as the modality adapter \mathcal{A} . The model is fine-tuned with instruction-tuning datasets that focus on diverse vision-language capabilities to create the models for merging. More details are provided in Supplementary B.5.

Single-Modal Baselines. For sufficient comparison with our method that uses UQ to guide MLLM merging, we compare UQ-Merge against various single-modal merging methods and test their performance in multimodal scenarios. Specifically, we consider DARE (Yu et al., 2024), DELLA (Deep et al., 2024), Linear (Wortsman et al., 2022), TIES (Yadav et al., 2023), Task Arithmetic (Ilharco et al., 2023), and Model Stock (Jang et al., 2024). Due to the lack of model selection capability, we compare these methods in ❶ average performance in random selections and ❷ merge all models. Although baseline methods were originally designed for single-modal merging, they are capable of merging models that have the same architecture. We consider adapter and LLM backbone when using baselines, as a naive extension of these methods.

Datasets for UQ and Evaluation. We select vision-language classification datasets as our benchmarks for $\overline{\text{UQ}}(\cdot)$. Specifically, five datasets across five domains are considered: MMBench (reasoning/perception, Liu et al., 2023a), OODCV-VQA (out-of-distribution robustness, Zhao et al., 2022), ScienceQA (world knowledge, Lu et al., 2022), SEEDBench (spatial and temporal understanding, Li et al., 2023a), and AI2D (diagrams, Kembhavi et al., 2016). In line with (Ye et al., 2024; Kostumov et al., 2024), we reformat the answers and add two additional choices “I don’t know” and “None of the above”. Since our vision-language perturbation-based UQ does not require labels, we treat these datasets as *held-in datasets* and also use them for the evaluation of merged models’ performance in vision-language classification format tasks. We also test our method on held-out vision-language generation datasets.

4.2 UQ-Merge Is Effective for Removing Harmful Models

In this section, we compare our UQ-Merge against various single-modal merging methods on held-in and held-out datasets to show the effectiveness of UQ-Merge in excluding harmful models. For baseline merging methods, we evaluate them by measuring the average performance of their merged models. For all baseline methods, each time the merged model is produced by merging a random model selection from all models we fine-tuned, and the number of models selected each time is the same as the selection of our method UQ-Merge. Evaluation results are summarized in Table 1, and the perfor-

Table 1: The comparison between UQ-Merge and single-modal merging methods. **Baseline methods merge randomly selected the same number of models to UQ-Merge.** Average and standard error of the accuracy of baselines across selections are reported. Results are measured with 3 selections. The best and second-best performances are highlighted in **bold** and underline, respectively.

Vision-Language Classification Datasets							
Merging Methods	Average	AI2D	ScienceQA	SeedBench	MMBench	OOD-CV	
DELLA	13.75 ± 3.09	8.01 ± 6.94	2.72 ± 2.60	5.73 ± 4.96	26.25 ± 0.75	26.04 ± 0.48	
Linear	53.15 ± 5.72	43.23 ± 7.63	57.97 ± 8.06	55.13 ± 7.55	70.26 ± 1.31	39.14 ± 4.95	
TIES	55.43 ± 3.88	49.41 ± 3.64	64.15 ± 6.43	57.16 ± 5.07	69.82 ± 0.62	36.60 ± 3.70	
Task Arithmetic	51.10 ± 7.23	43.58 ± 7.81	52.40 ± 6.88	48.58 ± 14.28	70.02 ± 1.39	40.95 ± 6.07	
Model Stock	21.73 ± 7.27	18.22 ± 11.31	17.17 ± 19.61	20.35 ± 13.45	27.60 ± 1.14	25.32 ± 1.09	
	49.99 ± 0.58	43.81 ± 0.92	65.07 ± 0.06	49.11 ± 0.84	64.10 ± 0.82	27.86 ± 0.52	
Ours	58.05	51.75	68.07	60.56	70.35	39.52	

Vision-Language Generation Datasets								
Merging Methods	Average	OKVQA	TextVQA	GQA	MMMU	VizWiz	VQAv2	POPE
DARE	12.83 ± 8.01	0.37 ± 0.35	5.48 ± 4.76	10.91 ± 9.82	26.22 ± 0.80	2.43 ± 2.11	19.77 ± 17.35	24.60 ± 21.64
DELLA	47.94 ± 3.75	40.53 ± 8.05	41.17 ± 1.77	43.96 ± 3.26	32.78 ± 1.31	41.92 ± 5.66	65.33 ± 3.46	69.87 ± 18.26
Linear	54.13 ± 0.72	44.61 ± 5.90	42.74 ± 1.83	49.83 ± 0.92	33.71 ± 0.95	52.80 ± 4.44	69.52 ± 1.77	85.71 ± 0.48
TIES	54.95 ± 1.18	50.16 ± 5.58	44.35 ± 0.35	52.87 ± 1.61	32.48 ± 1.11	46.66 ± 3.37	71.34 ± 1.57	86.81 ± 0.23
Task Arithmetic	18.06 ± 12.38	2.73 ± 2.36	13.61 ± 11.02	16.68 ± 15.07	26.22 ± 1.74	3.30 ± 2.65	28.67 ± 24.04	35.23 ± 31.08
Model Stock	37.79 ± 2.26	2.27 ± 0.96	34.36 ± 2.60	38.29 ± 1.26	32.59 ± 0.23	54.24 ± 0.20	59.19 ± 2.08	43.57 ± 11.01
Ours	56.01	50.39	43.49	50.75	35.22	54.13	71.77	86.33

mance of each selection is in Table 10 and Table 11. From the results, the following observations can be drawn: ❶ Our UQ-Merge demonstrates superior performance compared to all other merging methods. Specifically, UQ-Merge achieves 2.62% ~ 44.3% and 1.06% ~ 43.18% improvement on average accuracy of held-in and held-out datasets. This validates the effectiveness of UQ-Merge in model selection to exclude over-confident models. ❷ On the held-in dataset, which is used for uncertainty quantification, UQ-Merge obtains a more significant performance gain compared to held-out datasets, with 4 out of 5 highest accuracy. This justifies our practice of using UQ to perform model selection, as in real-world applications labels are usually unavailable and UQ only relies on input to evaluate a model, and UQ-guided model selection can effectively improve the performance on these applications and even generalize to held-out datasets. ❸ The performance of single-modal merging methods varies a lot in MLLM merging. The gap of average accuracy for baselines is 41.68% and 42.12% on held-in and held-out datasets, respectively. These large gaps show the various effectiveness of state-of-the-art single-modal merging methods when the setting is shifted to the MLLM merging.

4.3 UQ Guided Model Merging Surpasses Existing Merging Methods

We further compare UQ-Merge in a more challenging setting, where baseline merging methods have an “unfair” advantage to access all the knowl-

edge within models and merge. Experiment results show that ❶ Compared to baseline methods, UQ-Merge still achieves the best average accuracy on both held-in and held-out datasets, surpassing these single-modal methods by 0.54% ~ 51.69% and 1.3% ~ 52.6% respectively. ❷ Compared to the average performance of merging randomly selected portions of all models, all merging methods except DARE and Task Arithmetic enjoy performance increase by adding more models. This suggests the benefit of incorporating more models from diverse tasks to build a stronger model and supports our claim that model merging is a cheap way to aggregate knowledge from different models. It is worth noting that existing single-modal merging methods have a certain ability to resolve potentially harmful models when merging, by adopting model weight level manipulation to resolve weight conflict and preserve knowledge from different tasks (Ilharco et al., 2023; Yadav et al., 2023; Yu et al., 2024). However, these methods are limited to single-modal model merging. Our UQ-Merge is orthogonal to these works, as we consider model level removal of harmful over-confident models in MLLM merging situations to improve the performance, and can benefit from the techniques to ameliorate weight conflict.

4.4 Research Questions and Ablation Study

In this section, we conduct an in-depth investigation of the designs adopted in UQ-Merge and how they contribute to improved performance. Specifi-

Table 2: The comparison between UQ-Merge and single-modal merging methods. **Baselines merge all 10 models.** The best and second-best performances are highlighted in **bold** and underline.

Merging Methods	Vision-Language Classification Datasets					
	Average	AI2D	ScienceQA	SeedBench	MMBench	OOD-CV
DARE	6.36	0.00	14.10	17.54	14.10	17.54
DELLA	57.50	62.96	71.69	41.94	71.69	41.94
Linear	56.90	67.39	69.98	36.93	69.98	36.93
TIES	57.51	62.99	71.69	41.94	71.69	41.94
Task Arithmetic	10.94	0.29	26.58	25.78	26.58	25.78
Model Stock	51.55	65.21	65.74	29.47	65.74	29.47
Ours	58.05	68.07	70.35	39.52	70.35	39.52

Merging Methods	Vision-Language Generation Datasets							
	Average	OKVQA	TextVQA	GQA	MMMU	VizWiz	VQAv2	POPE
DARE	3.41	0.02	0.23	0.00	23.56	0.00	0.03	0.00
DELLA	54.69	47.84	44.79	51.65	33.56	48.38	70.51	86.12
Linear	54.71	44.51	44.23	48.56	35.00	55.21	70.21	85.23
TIES	54.70	47.89	44.81	51.65	33.56	48.39	70.51	86.11
Task Arithmetic	3.56	0.00	0.04	0.00	24.89	0.01	0.01	0.00
Model Stock	45.05	8.09	38.24	42.49	32.78	56.19	63.63	73.93
Ours	56.01	50.39	43.49	50.75	35.22	54.13	71.77	86.33

cally, we address the following: (1) Is UQ a more effective way to exclude harmful models, and how should uncertainty of models be used? (Section 4.4) (2) How to select models to merge after quantifying models’ uncertainty? (Section 4.4) (3) How to design UQ? (Section 4.4) (4) After selection, how to merge models? (Section 4.4).

RQ1: Is UQ more effective than accuracy? How to use uncertainty? A1: Yes; sort by decreasing uncertainty. In UQ-Merge, uncertainty is adopted to measure each model and exclude harmful models. An intuitive alternative to uncertainty is the accuracy of the model on validation datasets, with sorting done in either ascending or descending order. To address these research questions, we compare uncertainty and accuracy to determine which serves as better guidance. We replace uncertainty in UQ-Merge with accuracy on held-in datasets and test both kinds of guidance in ascending and descending order. Other components in UQ-Merge are kept untouched. We evaluate these modified designs on held-out datasets due to the usage of held-in datasets for testing accuracy. As shown in Tab 3, sorting by descending order achieves the best average performance compared to other options, confirming the effectiveness of our design. Compared to ascending uncertainty, descending order leads to a better performance, which justifies our aim to exclude over-confident models.

RQ2: How to select models to merge? A2: When the uncertainty of the merged model stops decreasing. After sorting models by descending

order of uncertainty, it remains unsure how to exclude harmful models and select beneficial ones. In UQ-Merge, this process is conducted by picking the merged model when uncertainty stops decreasing. To verify this, we evaluated the correlation between uncertainty, accuracy on validation benchmarks, and accuracy on held-out datasets. In our experiments, we used RealWorldQA (xAI, 2024), Seedbench 2 Plus (Li et al., 2024a), and OcrBench (Liu et al., 2024d), which focus on real-world QA, multi-disciplinary knowledge, and text recognition, respectively. As shown in Table 4, the point uncertainty rises and the highest validation accuracy align with the peak performance on held-out datasets. Our findings indicate that lower uncertainty corresponds to better performance. We attribute this to the enhanced capability of the merged model that makes it more robust to input perturbation and could generate consistent answers. When uncertainty further decreases, the model becomes over-confident, leading to the performance drop.

RQ3: How to design perturbation? A3: Vision-Language input perturbation is crucial. In this research question, we aim to investigate how different perturbation designs would affect the merging performance of UQ-Merge. Specifically, we compare the input and model perturbation method adopted in UQ-Merge versus only using input perturbation, by using them as different UQ functions in our UQ-Merge framework and test the merged model. We implement input perturbation following the same design of UQ-Merge, by adding random image transformations and text prompts to the im-

Table 3: Comparison of uncertainty and accuracy for guidance. Each guidance is implemented with ascending and descending orders to sort models. Accuracy for merging guidance is tested on held-in datasets, and all results are reported on held-out datasets. The best and second-best performances are highlighted in **bold** and underline.

Guidance	Performance with Different Guidance								
	Order	Average	OKVQA	TextVQA	GQA	MMMU	VizWiz	VQAv2	POPE
Uncertainty	Ascending	54.30	45.64	44.61	45.31	34.56	54.52	70.05	85.42
	Descending	56.01	50.39	43.49	50.75	35.22	54.13	71.77	86.33
Accuracy	Ascending	54.72	44.56	44.12	48.58	35.00	55.34	70.22	85.24
	Descending	52.27	30.33	40.68	50.58	35.67	55.73	67.34	85.58

Table 4: The correlation of uncertainty with accuracy on validation benchmarks and held-out benchmarks. The point uncertainty stops decreasing and the highest validation accuracy and held-out accuracy are marked in **bold**.

# Models	1	2	3	4	5	6	7	8	9	10
Uncertainty	0.21197	0.13290	0.05609	0.04866	0.04531	0.04155	0.03950	0.03954	0.03740	0.03675
Validation Accuracy	18.75	18.48	36.32	37.30	38.84	39.20	39.91	39.15	39.28	39.60
Held-out Accuracy	47.09	51.22	52.47	54.15	54.13	54.54	56.01	55.56	54.15	54.14

Table 5: Comparison of perturbation types. Results are average accuracy on datasets.

	Held-in	Held-out
Input	56.82	56.00
Input & Model	56.80	56.01

Table 6: Comparison of merging methods on the same group of models. Results are average accuracy on held-in and held-out datasets.

Merging Methods	Number of Models	
	7	10
DARE	22.30	4.64
DELLA	56.26	55.86
Linear	56.86	55.62
TIES	56.26	55.87
Task Arithmetic	22.21	6.64
Model Stock	46.82	47.76

age and text branches respectively. As shown in Table 5, when only use input perturbation, the performance is slightly improved on held-in datasets. On held-out datasets, the performance is slightly worse for input perturbation only. We attribute this to the robust capability of LLM backbones and dynamic sparsity of LLM inference (Liu et al., 2023b), which makes model perturbation may not significantly affect the performance of the LLM.

RQ4: What merging method to use given a group of MLLMs? A4: TIES, Linear or DELLA. Existing merging methods are designed to deal with single-modal merging, and it remains unclear how these merging methods perform for merging multimodal models. In this research ques-

tion, we explore the performance of these single-modal merging methods in the multimodal scenario by evaluating their performance on a given group of models. Specifically, we evaluate DARE, DELLA, Linear, TIES, Task Arithmetic, and Model Stock on held-in and held-out datasets and calculate the average performance on all the datasets. From results in Table 6 we observe that DELLA, Linear, and TIES perform better than other methods. In 10-model merging, all instruction-tuned models are merged. As shown in Table 6, given the same ten models, TIES achieve the best performance. When merging seven models with the highest uncertainty, we observe that ❶ The performance of all merging methods improved, demonstrating the benefit of model selection. ❷ linear merging achieves the best performance, which supports our choice in UQ-Merge that linearly merges models.

5 Conclusion

In this paper, we present a novel MLLM merging algorithm UQ-Merge to aggregate diverse knowledge of models into a single MLLM. We design a multimodal-tailored vision-language perturbation-based UQ and employ it to guide the merging process. As a result, UQ-Merge could identify beneficial models to merge and use the uncertainty value to decide the merging order and number of models to merge, and apply appropriate merging on selected models. Extensive experiments on datasets from diverse domains consistently demonstrate the effectiveness of model selection and significantly improved performance with our algorithm.

Acknowledgement

This project is supported by the Honda Research Institute USA.

References

- AugMix — Torchvision main documentation. <https://pytorch.org/vision/main/>.
- Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezazadegan, Li Liu, Mohammad Ghavamzadeh, Paul Fieguth, Xiaochun Cao, Abbas Khosravi, U Rajendra Acharya, et al. 2021. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information fusion*, 76:243–297.
- Samuel K. Ainsworth, Jonathan Hayase, and Siddhartha Srinivasa. 2023. Git re-basin: Merging models modulo permutation symmetries.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. 2024. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *Preprint*, arXiv:2312.14238.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality.
- Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. 2018. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*.
- Can Cui, Yunsheng Ma, Xu Cao, Wenqian Ye, Yang Zhou, Kaizhao Liang, Jintai Chen, Juanwu Lu, Zichong Yang, Kuei-Da Liao, et al. 2024. A survey on multimodal large language models for autonomous driving. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 958–979.
- Nico Daheim, Thomas Möllenhoff, Edoardo Maria Ponti, Iryna Gurevych, and Mohammad Emtiyaz Khan. 2023. Model merging by uncertainty-based gradient matching. *arXiv preprint arXiv:2310.12808*.
- Pala Tej Deep, Rishabh Bhardwaj, and Soujanya Poria. 2024. Della-merging: Reducing interference in model merging through magnitude-based sampling. *ArXiv*, abs/2406.11617.
- Ning Ding, Yulin Chen, Ganqu Cui, Xingtai Lv, Ruobing Xie, Bowen Zhou, Zhiyuan Liu, and Maosong Sun. 2024. Mastering text, code and math simultaneously via fusing highly specialized language models. *arXiv preprint arXiv:2403.08281*.
- Xiang Gao, Jiaxin Zhang, Lalla Mouatadid, and Kamalika Das. 2024. Spug: Perturbation-based uncertainty quantification for large language models. *ArXiv*, abs/2403.02509.
- Jakob Gawlikowski, Cedrique Rovile Njietcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. 2023. A survey of uncertainty in deep neural networks. *Artificial Intelligence Review*, 56(Suppl 1):1513–1589.
- Charles Goddard, Shamane Siriwardhana, Malikeh Ehghaghi, Luke Meyers, Vlad Karpukhin, Brian Benedict, Mark McQuade, and Jacob Solawetz. 2024. Arcee’s mergekit: A toolkit for merging large language models. *arXiv preprint arXiv:2403.13257*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913.
- Hongji Guo, Nakul Agarwal, Shao-Yuan Lo, Kwonjoon Lee, and Qiang Ji. 2024. Uncertainty-aware action decoupling transformer for action anticipation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18644–18654.
- Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. 2018. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3608–3617.
- Dan Hendrycks, Norman Mu, Ekin D. Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. 2020. Augmix: A simple data processing method to improve robustness and uncertainty. *Preprint*, arXiv:1912.02781.
- Drew A Hudson and Christopher D Manning. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709.
- Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hananeh Hajishirzi, and Ali Farhadi. 2023. Editing models with task arithmetic.
- Dong-Hwan Jang, Sangdoon Yun, and Dongyoon Han. 2024. Model stock: All we need is just a few fine-tuned models.

- Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. 2014. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 787–798.
- Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. 2016. A diagram is worth a dozen images. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 235–251. Springer.
- Vasily Kostumov, Bulat Nutfullin, Oleg Pilipenko, and Eugene Ilyushin. 2024. [Uncertainty-aware evaluation for vision-language models](#). *Preprint*, arXiv:2402.14418.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73.
- Bohao Li, Yuying Ge, Yi Chen, Yixiao Ge, Ruimao Zhang, and Ying Shan. 2024a. [Seed-bench-2-plus: Benchmarking multimodal large language models with text-rich visual comprehension](#). *Preprint*, arXiv:2404.16790.
- Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. 2023a. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*.
- Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. 2024b. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023b. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*.
- Bin Lin, Zhenyu Tang, Yang Ye, Jiayi Cui, Bin Zhu, Peng Jin, Junwu Zhang, Munan Ning, and Li Yuan. 2024. Moe-llava: Mixture of experts for large vision-language models. *arXiv preprint arXiv:2401.15947*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024a. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024b. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Linyu Liu, Yu Pan, Xiaocheng Li, and Guanting Chen. 2024c. Uncertainty estimation and quantification for llms: A simple supervised approach. *arXiv preprint arXiv:2404.15993*.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. 2023a. Mm-bench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*.
- Yuliang Liu, Zhang Li, Mingxin Huang, Biao Yang, Wenwen Yu, Chunyuan Li, Xucheng Yin, Chenglin Liu, Lianwen Jin, and Xiang Bai. 2024d. [Ocrbench: On the hidden mystery of ocr in large multimodal models](#). *Preprint*, arXiv:2305.07895.
- Zichang Liu, Jue Wang, Tri Dao, Tianyi Zhou, Binhang Yuan, Zhao Song, Anshumali Shrivastava, Ce Zhang, Yuandong Tian, Christopher Re, et al. 2023b. Dejavu: Contextual sparsity for efficient llms at inference time. In *International Conference on Machine Learning*, pages 22137–22176. PMLR.
- Jinliang Lu, Ziliang Pang, Min Xiao, Yaochen Zhu, Rui Xia, and Jiajun Zhang. 2024. [Merge, ensemble, and cooperate! a survey on collaborative strategies in the era of large language models](#). *Preprint*, arXiv:2407.06089.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. 2016. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pages 3195–3204.
- Brandon McKinzie, Zhe Gan, Jean-Philippe Fauconnier, Sam Dodge, Bowen Zhang, Philipp Dufter, Dhruvi Shah, Xianzhi Du, Futang Peng, Floris Weers, et al. 2024. Mm1: Methods, analysis & insights from multimodal llm pre-training. *arXiv preprint arXiv:2403.09611*.
- Alireza Mehrtash, William M Wells, Clare M Tempany, Purang Abolmaesumi, and Tina Kapur. 2020. Confidence calibration and predictive uncertainty estimation for deep medical image segmentation. *IEEE transactions on medical imaging*, 39(12):3868–3878.

- Anand Mishra, Shashank Shekhar, Ajeet Kumar Singh, and Anirban Chakraborty. 2019. Ocr-vqa: Visual question answering by reading text in images. In *2019 international conference on document analysis and recognition (ICDAR)*, pages 947–952. IEEE.
- Vaishnavh Nagarajan and J. Zico Kolter. 2021. Uniform convergence may be unable to explain generalization in deep learning.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. 2022. A-okvqa: A benchmark for visual question answering using world knowledge. In *European conference on computer vision*, pages 146–162. Springer.
- ShareGPT. 2023. [\[link\]](#).
- Ken Shoemake. 1985. Animating rotation with quaternion curves. In *Proceedings of the 12th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '85*, page 245–254, New York, NY, USA. Association for Computing Machinery.
- Chenglei Si, Zhe Gan, Zhengyuan Yang, Shuohang Wang, Jianfeng Wang, Jordan Boyd-Graber, and Lijuan Wang. 2023. [Prompting gpt-3 to be reliable](#). *Preprint*, arXiv:2210.09150.
- Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. 2020. Textcaps: a dataset for image captioning with reading comprehension. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 742–758. Springer.
- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Yi-Lin Sung, Linjie Li, Kevin Lin, Zhe Gan, Mohit Bansal, and Lijuan Wang. 2023. An empirical study of multimodal model merging. *arXiv preprint arXiv:2304.14933*.
- Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, et al. 2024. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint arXiv:2406.16860*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Fanqi Wan, Xinting Huang, Deng Cai, Xiaojun Quan, Wei Bi, and Shuming Shi. 2024. Knowledge fusion of large language models. *arXiv preprint arXiv:2401.10491*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. 2021. Simvlm: Simple visual language model pretraining with weak supervision. *arXiv preprint arXiv:2108.10904*.
- 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. 2022. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time.
- xAI. 2024. [Grok-1.5 vision preview](#).
- Le Xue, Manli Shu, Anas Awadalla, Jun Wang, An Yan, Senthil Purushwalkam, Honglu Zhou, Viraj Prabhu, Yutong Dai, Michael S Ryoo, et al. 2024. xgen-mm (blip-3): A family of open large multimodal models. *arXiv preprint arXiv:2408.08872*.
- Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, and Mohit Bansal. 2023. Ties-merging: Resolving interference when merging models.
- Enneng Yang, Li Shen, Guibing Guo, Xingwei Wang, Xiaochun Cao, Jie Zhang, and Dacheng Tao. 2024. Model merging in llms, mllms, and beyond: Methods, theories, applications and opportunities. *arXiv preprint arXiv:2408.07666*.
- Fanghua Ye, Mingming Yang, Jianhui Pang, Longyue Wang, Derek F. Wong, Emine Yilmaz, Shuming Shi, and Zhaopeng Tu. 2024. [Benchmarking llms via uncertainty quantification](#). *Preprint*, arXiv:2401.12794.
- Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. 2023. A survey on multimodal large language models. *ArXiv*, abs/2306.13549.

- Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin Li. 2024. Language models are super mario: Absorbing abilities from homologous models as a free lunch.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. 2024. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9556–9567.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11975–11986.
- Duzhen Zhang, Yahan Yu, Chenxing Li, Jiahua Dong, Dan Su, Chenhui Chu, and Dong Yu. 2024a. Mmllms: Recent advances in multimodal large language models. *arXiv preprint arXiv:2401.13601*.
- Kaichen Zhang, Bo Li, Peiyuan Zhang, Fanyi Pu, Joshua Adrian Cahyono, Kairui Hu, Shuai Liu, Yuanhan Zhang, Jingkang Yang, Chunyuan Li, et al. 2024b. Lmms-eval: Reality check on the evaluation of large multimodal models. *arXiv preprint arXiv:2407.12772*.
- Yanzhe Zhang, Ruiyi Zhang, Jiuxiang Gu, Yufan Zhou, Nedim Lipka, Diyi Yang, and Tong Sun. 2023. Llavar: Enhanced visual instruction tuning for text-rich image understanding. *arXiv preprint arXiv:2306.17107*.
- Bingchen Zhao, Shaozuo Yu, Wufei Ma, Mingxin Yu, Shengxiao Mei, Angtian Wang, Ju He, Alan Yuille, and Adam Kortylewski. 2022. Ood-cv: A benchmark for robustness to out-of-distribution shifts of individual nuisances in natural images. In *European conference on computer vision*, pages 163–180. Springer.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.
- Xinyu Zhao, Guoheng Sun, Ruisi Cai, Yukun Zhou, Pingzhi Li, Peihao Wang, Bowen Tan, Yexiao He, Li Chen, Yi Liang, Beidi Chen, Binhang Yuan, Hongyi Wang, Ang Li, Zhangyang Wang, and Tianlong Chen. 2024. [Model-GLUE: Democratized LLM scaling for a large model zoo in the wild](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

A Additional Experiment Results

A.1 Ablation Study about UQ in UQ-Merge

In this section we provide more ablation studies about UQ in UQ-Merge. We conduct experiments with various ratios of dataset samples used for UQ, and compare with alternatives of the perturbation-based UQ method used in UQ-Merge. The results are shown in Table 7 and Table 8. Results in Table 7 show that UQ-Merge achieves stable performance across various ratios of samples used for UQ, demonstrating the data efficiency of UQ-Merge. We also test image-only perturbation, text-only perturbation and conformal prediction (CP) as alternative UQ methods in Table 8. Compared with alternatives, results in Table 8 prove the effectiveness of our vision-language perturbation-based UQ.

A.2 Evaluation of fine-tuned models

In Table 9, we provide uncertainty quantification results of fine-tuned models on held-in datasets. We conduct evaluation three times and the final uncertainty is the average. As observed in Table 9, the uncertainty is stable and consistent, showcasing the effectiveness and stability of our vision-language perturbation-based UQ.

A.3 Evaluation of Baselines over Random Selections

In Table 10 and Table 11, we provide performance of baseline single-modal merging methods on held-in and held-out datasets over model selections. As observed in tables, DELLA, Linear and TIES consistently outperform other merging methods with a small variance. The average and standard error are reported based on results above, and the error bar represents the 95% confidence interval.

A.4 Evaluation during Merging Steps

In Table 12, Table 13 and Table 14, we present the evaluation results of Table 4 before averaging during merging steps.

B More Implementation Details

B.1 Text Perturbation

Prompts used for perturbation of text inputs:

- 'you are a helpful assistant',
- 'you are a question-answering assistant',
- 'you are a nice assistant',
- 'You are a helpful assistant',
- 'You are a question-answering assistant',

- 'You are a nice assistant',
- 'You are a helpful assistant.',
- 'You are a question-answering assistant.',
- 'You are a nice assistant.'

B.2 Image Perturbation

The image perturbation is implemented by utilizing the implementation of AugMix (Hendrycks et al., 2020) in torchvision (Aug), and all parameters are set to default.

B.3 Datasets for Fine-tuning Models

To prepare fine-tuned models for merging, we choose visual reasoning datasets (GQA Hudson and Manning, 2019, RefCOCO Kazemzadeh et al., 2014; Mao et al., 2016); text-rich datasets (OCRVQA Mishra et al., 2019, TextCaps Sidorov et al., 2020); knowledge-based VQA datasets (OKVQA Marino et al., 2019, A-OKVQA Schwenk et al., 2022); GPT-generated datasets (LLaVA-Instruct Liu et al., 2024b, ShareGPT ShareGPT, 2023); and general VQA datasets (VQAv2 Goyal et al., 2017, VG Krishna et al., 2017). We follow the original train-test split of these datasets.

B.4 Datasets for Uncertainty Quantification

We use the code from (Kostumov et al., 2024) to process MMBench (Liu et al., 2023a), OODCV-VQA (Zhao et al., 2022), ScienceQA (Lu et al., 2022), SEEDBench (Li et al., 2023a), and AI2D (Kembhavi et al., 2016) for vision-language perturbation-based UQ.

B.5 Evaluation of Models

We adopt LMMs-Eval (Zhang et al., 2024b) with all default parameters to conduct evaluation of models on all benchmarks except MMBench and OODCV-VQA, which are evaluated directly using our pre-processed datasets. To more comprehensively evaluate the merged models' performance, we choose seven vision-language generation tasks of six domains, including open real-world knowledge (OKVQA Marino et al., 2019, MMMU Yue et al., 2024), text understanding (TextVQA Singh et al., 2019), compositional questioning answering (GQA Hudson and Manning, 2019), low-quality image understanding (VizWiz Gurari et al., 2018), general visual QA (VQAv2 Goyal et al., 2017), and hallucination (POPE Li et al., 2023b) as the benchmarks.

Table 7: Performance of the merged models from UQ-Mergewith different ratios of samples used for UQ.

Sample Ratios	Average	AI2D	ScienceQA	SeedBench	MMBench	OOD-CV
5% Samples	58.08	51.75	68.07	60.53	70.44	39.60
10% Samples	58.00	51.70	67.95	60.39	70.35	39.60
20% Samples	58.34	52.23	68.95	60.52	70.21	39.77
30% Samples	57.96	51.42	67.85	60.66	70.30	39.56
100% Samples	58.05	51.75	68.07	60.56	70.35	39.52

Sample Ratio	Average	OKVQA	TextVQA	GQA	MMMU	VizWiz	VQAv2	POPE
5% Samples	56.01	50.40	43.49	50.61	35.67	53.91	71.69	86.33
10% Samples	55.36	50.71	41.74	49.86	34.67	55.64	68.66	86.26
20% Samples	55.30	44.17	43.57	50.92	36.44	53.86	71.69	86.48
30% Samples	56.00	50.37	43.58	50.80	35.67	53.73	71.67	86.18
100% Samples	56.01	50.39	43.49	50.75	35.22	54.13	71.77	86.33

Table 8: Performance of the merged models with alternatives of vision-language perturbation-based UQ.

	Average	AI2D	ScienceQA	SeedBench	MMBench	OOD-CV
CP	56.86	50.42	67.19	59.83	69.88	36.98
Image Only	57.11	50.58	67.51	60.40	69.94	37.14
Text Only	57.05	50.55	67.39	60.59	69.87	36.86
Ours	58.05	51.75	68.07	60.56	70.35	39.52

	Average	OKVQA	TextVQA	GQA	MMMU	VizWiz	VQAv2	POPE
CP	54.60	44.27	43.97	48.56	34.67	55.20	70.16	85.37
Image Only	54.84	44.35	43.43	50.68	34.78	53.92	71.67	85.07
Text Only	55.06	44.55	43.26	50.39	36.11	54.24	71.62	85.23
Ours	56.01	50.39	43.49	50.75	35.22	54.13	71.77	86.33

Table 9: Uncertainty of all models on held-in datasets.

Tuning Dataset	Average	AI2D	MMBench	OOD-CV	ScienceQA	SeedBench
OKVQA	0.1314	0.1199	0.1073	0.1350	0.1149	0.1802
OCRvQA	0.0548	0.0555	0.0564	0.0669	0.0505	0.0446
GQA	0.2141	0.2032	0.2375	0.2136	0.2219	0.1941
VQAv2	0.0322	0.0286	0.0355	0.0350	0.0311	0.0306
TextCaps	0.0073	0.0078	0.0085	0.0115	0.0054	0.0030
A_OKVQA	0.1078	0.1318	0.0746	0.1066	0.1132	0.1130
RefCOCO	0.0275	0.0250	0.0309	0.0311	0.0267	0.0238
LLaVA-Instruct	0.0713	0.0670	0.0767	0.0787	0.0672	0.0668
ShareGPT	0.0317	0.0246	0.0380	0.0386	0.0320	0.0255
VG	0.0379	0.0334	0.0412	0.0417	0.0360	0.0372
Tuning Dataset	Average	AI2D	MMBench	OOD-CV	ScienceQA	SeedBench
OKVQA	0.1305	0.1199	0.1072	0.1317	0.1152	0.1787
OCRvQA	0.0544	0.0550	0.0560	0.0663	0.0505	0.0444
GQA	0.2141	0.2024	0.2408	0.2084	0.2229	0.1962
VQAv2	0.0323	0.0288	0.0361	0.0347	0.0312	0.0307
TextCaps	0.0072	0.0077	0.0087	0.0114	0.0055	0.0029
A_OKVQA	0.1077	0.1326	0.0751	0.1070	0.1117	0.1123
RefCOCO	0.0273	0.0251	0.0305	0.0305	0.0265	0.0238
LLaVA-Instruct	0.0713	0.0668	0.0771	0.0803	0.0662	0.0659
ShareGPT	0.0318	0.0246	0.0375	0.0394	0.0318	0.0255
VG	0.0378	0.0335	0.0408	0.0419	0.0361	0.0369
Tuning Dataset	Average	AI2D	MMBench	OOD-CV	ScienceQA	SeedBench
OKVQA	0.1306	0.1205	0.1067	0.1334	0.1157	0.1767
OCRvQA	0.0547	0.0554	0.0562	0.0668	0.0501	0.0448
GQA	0.2130	0.2020	0.2369	0.2115	0.2220	0.1926
VQAv2	0.0323	0.0288	0.0357	0.0344	0.0313	0.0311
TextCaps	0.0073	0.0078	0.0085	0.0115	0.0055	0.0030
A_OKVQA	0.1082	0.1324	0.0758	0.1069	0.1123	0.1135
RefCOCO	0.0273	0.0249	0.0303	0.0308	0.0266	0.0238
LLaVA-Instruct	0.0713	0.0666	0.0763	0.0801	0.0667	0.0666
ShareGPT	0.0316	0.0245	0.0379	0.0388	0.0319	0.0247
VG	0.0376	0.0335	0.0406	0.0417	0.0360	0.0363

Table 10: Accuracy of baselines on held-in datasets when merging random model selections.

Merging Methods	Vision-Language Classification Datasets					
	Average	AI2D	ScienceQA	SeedBench	MMBench	OOD-CV
DARE	15.68	11.98	5.19	8.47	26.53	26.21
DELLA	46.73	34.42	50.05	46.97	68.76	33.45
Linear	57.85	51.65	67.83	60.25	70.19	39.35
TIES	55.60	49.29	55.96	56.14	70.74	45.85
Task Arithmetic	15.13	5.18	12.98	5.20	26.53	25.78
Model Stock	50.65	44.82	65.12	50.08	64.86	28.37
Merging Methods	Average	AI2D	ScienceQA	SeedBench	MMBench	OOD-CV
DARE	10.18	0.00	0.00	0.00	25.40	25.50
DELLA	55.02	47.54	57.70	56.57	70.81	42.47
Linear	57.48	51.36	67.90	59.91	70.17	38.07
TIES	54.95	46.76	56.78	57.49	70.90	42.83
Task Arithmetic	20.54	24.19	0.00	24.96	27.46	26.10
Model Stock	49.56	43.56	65.10	48.56	63.24	27.34
Merging Methods	Average	AI2D	ScienceQA	SeedBench	MMBench	OOD-CV
DARE	15.39	12.05	2.96	8.72	26.81	26.42
DELLA	57.69	47.73	66.16	61.86	71.20	41.51
Linear	50.95	45.21	56.73	51.31	69.10	32.39
TIES	42.77	34.68	44.47	32.11	68.41	34.16
Task Arithmetic	29.52	25.29	38.54	30.89	28.80	24.08
Model Stock	49.77	43.04	65.00	48.70	64.21	27.88
Ours	58.05	51.75	68.07	60.56	70.35	39.52

Table 11: Accuracy of baselines on held-out datasets when merging random model selections.

Merging Methods	Vision-Language Generation Datasets							
	Average	OKVQA	TextVQA	GQA	MMMU	VizWiz	VQAv2	POPE
DARE	18.78	0.41	7.82	19.06	27.11	3.85	32.48	40.70
DELLA	43.78	41.69	39.74	40.83	31.67	40.75	62.46	49.32
Linear	53.78	39.45	40.67	49.65	34.67	57.20	69.01	85.83
TIES	56.31	56.08	44.75	54.72	33.33	45.32	73.03	86.94
Task Arithmetic	4.04	0.04	0.91	0.33	24.67	0.26	0.96	1.09
Model Stock	39.37	3.35	37.30	39.54	32.33	54.38	61.51	47.18
Merging Methods	Average	OKVQA	TextVQA	GQA	MMMU	VizWiz	VQAv2	POPE
DARE	3.71	0.00	0.00	0.00	26.00	0.00	0.00	0.00
DELLA	51.08	47.94	43.15	43.70	32.44	36.94	69.17	84.23
Linear	53.66	43.35	43.43	49.02	33.67	52.88	68.06	85.18
TIES	54.15	49.39	44.19	52.08	31.22	44.17	71.07	86.94
Task Arithmetic	22.70	4.47	20.60	19.69	25.89	4.53	41.02	42.72
Model Stock	38.79	1.95	33.38	38.30	32.67	54.34	58.55	52.33
Merging Methods	Average	OKVQA	TextVQA	GQA	MMMU	VizWiz	VQAv2	POPE
DARE	15.99	0.70	8.61	13.66	25.56	3.44	26.84	33.10
DELLA	48.95	31.97	40.62	47.34	34.22	48.07	64.35	76.06
Linear	54.96	51.04	44.12	50.83	32.78	48.33	71.48	86.11
TIES	54.40	45.01	44.11	51.81	32.89	50.50	69.93	86.54
Task Arithmetic	27.45	3.67	19.32	30.01	28.11	5.11	44.02	61.88
Model Stock	35.20	1.51	32.39	37.02	32.78	54.01	57.50	31.21
Ours	56.01	50.39	43.49	50.75	35.22	54.13	71.77	86.33

Table 12: Uncertainty on held-in datasets during merging steps. **Bold** denotes the step UQ-Merge stops.

Held-in Datasets	Merging Steps									
	1	2	3	4	5	6	7	8	9	10
AI2D	0.20129	0.11652	0.06806	0.06142	0.05759	0.05062	0.04600	0.04676	0.04328	0.04248
MMBench	0.23477	0.11862	0.04576	0.03985	0.03784	0.03560	0.03583	0.03488	0.03368	0.03307
ScienceQA	0.21144	0.13648	0.06652	0.05888	0.05306	0.04905	0.04543	0.04672	0.04356	0.04222
SeedBench	0.22157	0.12653	0.05543	0.04701	0.04384	0.04079	0.03865	0.03837	0.03703	0.03606
OOD-CV	0.19077	0.16634	0.04470	0.03616	0.03423	0.03168	0.03160	0.03097	0.02945	0.02990
Average	0.21197	0.13290	0.05609	0.04866	0.04531	0.04155	0.03950	0.03954	0.03740	0.03675

Table 13: Accuracy on validation datasets during merging steps. **Bold** denotes the step UQ-Merge stops.

Validation Datasets	Merging Steps									
	1	2	3	4	5	6	7	8	9	10
RealWorldQA	27.58	22.48	47.45	46.80	47.19	47.58	49.15	47.32	47.32	47.84
SeedBench 2 Plus	23.36	10.36	39.00	40.40	41.33	41.81	42.07	41.72	42.51	42.16
OCRBench	5.30	22.60	22.50	24.70	28.00	28.20	28.50	28.40	28.00	28.80
Average	18.75	18.48	36.32	37.30	38.84	39.20	39.91	39.15	39.28	39.60

Table 14: Accuracy on held-out datasets during merging steps. **Bold** denotes the step UQ-Merge stops.

Held-out Datasets	Merging Steps									
	1	2	3	4	5	6	7	8	9	10
OKVQA	27.05	54.77	50.48	48.01	46.21	44.13	50.39	47.90	48.01	46.21
TextVQA	23.15	36.06	34.24	37.96	41.19	41.82	43.49	43.22	37.96	41.19
GQA	61.73	54.08	50.77	50.29	48.97	49.98	50.75	49.75	50.29	48.97
MMMU	30.00	29.78	33.78	34.22	34.44	35.00	35.22	35.22	34.22	34.56
VizWiz	44.25	34.69	46.72	54.36	54.61	55.66	54.13	55.31	54.36	54.61
VQAv2	60.03	66.09	66.70	68.30	67.81	68.71	71.77	71.25	68.30	67.81
POPE	83.39	83.07	84.60	85.91	85.68	86.46	86.33	86.29	85.91	85.66
Average	47.09	51.22	52.47	54.15	54.13	54.54	56.01	55.56	54.15	54.14