

MULTIMODAL ATTENTION MERGING FOR IMPROVED SPEECH RECOGNITION AND AUDIO EVENT CLASSIFICATION

Anirudh S. Sundar^{1,2,†*} Chao-Han Huck Yang² David M. Chan^{2,3}
 Shalini Ghosh² Venkatesh Ravichandran² Phani Sankar Nidadavolu^{2*}

¹Georgia Institute of Technology ²Amazon Alexa AI ³University of California, Berkeley

ABSTRACT

Training large foundation models using self-supervised objectives on unlabeled data, followed by fine-tuning on downstream tasks, has emerged as a standard procedure. Unfortunately, the efficacy of this approach is often constrained by both limited fine-tuning compute and scarcity in labeled downstream data. We introduce Multimodal Attention Merging (MAM), an attempt that facilitates direct knowledge transfer from attention matrices of models rooted in high-resource modalities, text and images, to those in resource-constrained domains, speech and audio, employing a *zero-shot paradigm*. MAM reduces the relative Word Error Rate (WER) of an Automatic Speech Recognition (ASR) model by up to 6.70%, and relative classification error of an Audio Event Classification (AEC) model by 10.63%. In cases where some data/compute is available, we present Learnable-MAM, a data-driven approach to merging attention matrices, resulting in a further 2.90% relative reduction in WER for ASR and 18.42% relative reduction in AEC compared to fine-tuning without model merging.

Index Terms— Knowledge transfer, cross-modal adaptation, speech recognition, and acoustic modeling

1. INTRODUCTION

Current approaches in deep learning train large foundation models using task-agnostic self-supervised objectives on unlabeled data [1, 2, 3]. The models are then fine-tuned on downstream tasks, utilizing task-specific inputs and labels. As foundation models increase in size, fine-tuning is hamstrung by limitations in compute. Moreover, scarcity of task-specific data compounds this challenge, particularly for relatively lower-resource modalities like speech or audio when compared to more widely studied modalities such as text or audio.

Prior research [4, 5, 6, 7] has demonstrated transferability of the Transformer [8] across modalities with minimal or no fine-tuning of the self-attention mechanism. [4] shows that text pre-training is sufficient to learn modality agnostic properties of sequences by demonstrating the success of frozen text pre-trained Transformers in image classification and protein fold prediction without self-attention fine-tuning. [9] shows that cross-modality capabilities in Transformers arises from the transfer of knowledge in the form of position-aware context. [10] and [11] provide evidence for multimodal neurons in the Transformer and demonstrate that modality transfer happens in intermediate layers.

Motivated by these works, we present Multimodal Attention Merging (MAM) to investigate the possibility of transferring knowledge from models trained on high-resource modalities such as text

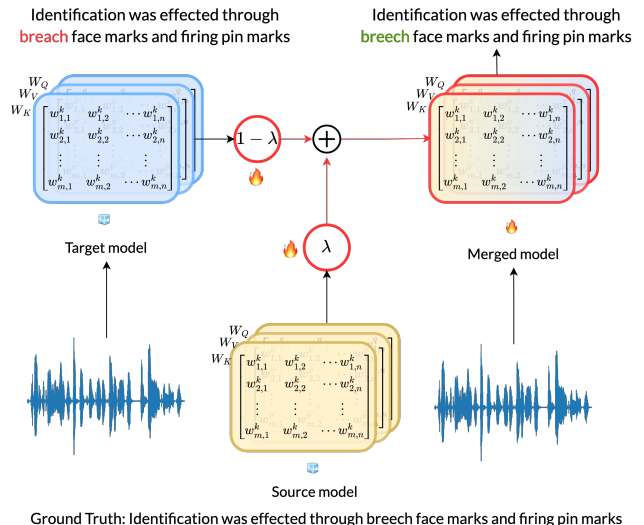


Fig. 1. Learnable Multimodal Attention Merging. The merged model correctly transcribes *breech*. 🔥 modules are fine-tuned (trainable) while ❄️ are frozen (non-trainable). In our initial approach, Multimodal Attention Merging, the interpolation factor λ is frozen (non-trainable).

and images to models trained on low-resource modalities such as speech and audio. Given abundant textual and visual data, self-supervised pre-training using objectives such as Masked Language Modeling and Masked Patch Prediction utilize the Transformer to learn generalized representations of natural language text and images. Chiang and Lee [12] indicate that using Masked Language Modeling allows self-attention to capture explicit and implicit dependencies at the token level, which are modality agnostic. Therefore, through MAM we investigate whether these textual and visual parameter-space representations generalize to speech and audio. Through a systematic interpolation of attention matrices from models trained on high-resource modalities (e.g., BERT [13], Vision Transformer [3]) MAM demonstrates an improvement in performance of models trained on low-resource modalities (e.g., HuBERT [2], BEATs [14]) on Automatic Speech Recognition (ASR) and Audio Event Classification (AEC). Our contributions are:

- In contrast to prior research [15, 16, 17, 18, 19], we show attention-merging on models from different tasks, modalities, and initializations. To the best of our knowledge, we are the first to extend attention-merging to sequence-to-sequence tasks, specifically ASR.

*Corresponding authors asundar34@gatech.edu and phanisn@amazon.com, † Work performed as an intern at Amazon

- Introducing Multimodal Attention Merging (MAM), we lower HuBERT’s relative Word Error Rate (WER) on LJ Speech by 6.70% and VCTK by 1.80%, also decreasing BEATs’ relative classification error on ESC-50 by 10.63%, without additional fine-tuning (Section 3.1).
- We present a data-driven approach employing the Sliced Wasserstein Distance [20] between intermediate hidden-layer representations to merge specific attention layers, reducing the need for empirical experimentation (Section 4.1).
- In the case where some data/compute is available, we present Learnable-MAM (L-MAM), by learning the interpolation factor during fine-tuning. L-MAM yields a 2.70% and 2.90% relative reduction in WER on LJ Speech and VCTK and a 18.42% relative reduction in classification error on ESC-50 compared to regular fine-tuning (Section 3.3).

2. RELATED WORK

Prior work in multimodal merging includes OTKGE [21], that uses Optimal Transport to align structural knowledge, linguistic information, and image embeddings in knowledge graphs. However, they do not extend their approach to merging model weights and focus on embeddings instead. Voice2Series [5] and Frozen Pretrained Transformer [4] demonstrate knowledge transfer across modalities through frozen self-attention weights. While both works study the transferability of self-attention across modalities, they stop short of merging models trained from different modalities and do not address sequence-to-sequence tasks such as ASR. In contrast, works such as Fisher merging [22], *local fine-tuning* [16], MODEL SOUPS [19], DMC [23], ADAPTERSOUP [24], MLM [25] discuss model merging but do not consider the multimodal scenario. Perhaps the closest work to ours is Multimodal Model Merging [15], an empirical study of merging vision and text models for combined vision-language tasks such as Visual Question Answering [26] and image-text retrieval [27]. The approach studies the use of simple interpolation, REGMEAN [15], and Task Vectors [28] to determine the best model merging approach. However, they require contrastive model alignment, use a shared seed pre-training phase to initialize models prior to merging, and address joint vision+language tasks. In contrast, MAM merges off-the-shelf models from different modalities without constraints on pre-training tasks or weight initialization. Through L-MAM, we also present an approach to learn the interpolation factor in the case where limited data/compute is available, reducing the requirement on empirical experimentation.

3. METHOD

MAM seeks to determine if the Transformer [8] attention mechanism generalizes across modalities. It does so by exploring the transferability of parameter-space sequence representations of Transformers pre-trained on high-resource modalities (text or vision) to those trained on relatively low-resource modalities (speech or audio). For convenience, we refer to the high-resource and low-resource pre-trained models as the *Source Model* and *Target Model* respectively. Figure 1 presents an outline of attention merging. It is important to note that all our approaches require the source and target models to have the same number of attention layers and hidden layer sizes.

We apply attention merging to two tasks: Automatic Speech Recognition (ASR) and Audio Event Classification (AEC). ASR transcribes human speech to text, while AEC identifies real-life events (e.g. barking, thunderstorm, whistling) from audio clips.

We use three main approaches to demonstrate MAM: Attention Interpolation, Layer-wise Attention Interpolation, and Attention-Merging with Learnable Interpolation.

3.1. Attention Interpolation

MAM with attention interpolation uses a convex combination of the source and target models. Using an interpolation factor $\lambda \in [0, 1]$, we merge Query, Key, and Value matrices (W_Q, W_K, W_V) across all layers \mathcal{L} in the attention computation. For source and target models M_s and M_t , the merged model M_{merge} ’s Query, Key, and Value matrices for each layer i are defined in Equation 1. M_{merge} is applied to the same downstream task as M_t .

$$W_{Q_i, K_i, V_i}^{M_{\text{merge}}} = \lambda \cdot W_{Q_i, K_i, V_i}^{M_s} + (1 - \lambda) \cdot W_{Q_i, K_i, V_i}^{M_t} \quad (1)$$

3.2. Layer-wise Attention Interpolation

Evidenced by research categorizing layers by importance [29], we experiment with merging a subset of layers ℓ from the set of all layers \mathcal{L} ($|\ell| < |\mathcal{L}|$) and present the modified approach in Equation 2. We are interested in identifying whether merging all layers is necessary, or if merging a subset yields better generalization.

$$W_{Q_i, K_i, V_i}^{M_{\text{merge}}} = \lambda \cdot W_{Q_i, K_i, V_i}^{M_s} + (1 - \lambda) \cdot W_{Q_i, K_i, V_i}^{M_t} \quad (2)$$

$$\lambda = 0 \quad \forall i \notin \ell$$

3.3. Attention-Merging with Learnable Interpolation

Finally, we learn the interpolation factor λ for individual downstream tasks. In this setting, λ is optimized concurrently with model weights M_t during fine-tuning. We omit optimizing M_s as it corresponds to a different modality from the downstream task. In contrast to the previous two approaches, we do not impose the requirement for utilizing uniform λ across layers. Instead, we view learning λ as a gate that enables flexible information transmission from the high-resource to the low-resource modality and describe this technique in Equation 3.

$$W_{Q_i, K_i, V_i}^{M_{\text{merge}}} = \lambda_i \cdot W_{Q_i, K_i, V_i}^{M_s} + (1 - \lambda_i) \cdot W_{Q_i, K_i, V_i}^{M_t} \quad (3)$$

4. EXPERIMENTS

For ASR, we merge publicly available HuBERT-large [2] and BERT-large-uncased [13] models. Both models have 24 attention layers and a width of 1024, totaling 300 million parameters. HuBERT-large is an encoder model pre-trained on 60,000 hours of Libri-Light [30] for speech representation learning and fine-tuned with CTC on 960h of Librispeech [31] for ASR. BERT-large-uncased is pre-trained via Masked Language Modeling on the Wikipedia corpus and Bookcorpus. We evaluate merged models on LJ Speech [32], a single-speaker dataset of 13,100 audio clips from non-fiction books totaling 24 hours, and VCTK [33], a multi-speaker dataset with 43,000 audio clips totaling 44 hours. VCTK features diverse speakers from regions like England, Scotland, and America, reading texts chosen for comprehensive contextual and phonetic coverage.

For AEC, we merge BEATs [14] and Vision Transformer (ViT) [3] containing 12 attention layers and hidden size of 768, totaling 90 million parameters. BEATs is pretrained on AudioSet [34] for audio representation learning while ViT is pretrained on ImageNet [35]. Evaluation focuses on ESC-50 [36], which comprises 2000 5-second environmental audio recordings across 50 classes.

λ	Source	$\mathcal{N}(\text{Source})$	$\mathcal{N}(\text{Target})$	$\mathcal{N}(0,1)$
LJ Speech - WER(%) Source: BERT Target: HuBERT				
0.00	9.25	-	-	-
0.05	9.11	9.78	9.26	10.57
0.10	9.06	32.43	10.06	74.03
0.15	9.20	98.30	18.13	99.95
0.20	9.62	99.99	75.93	100.00
0.25	11.12	100.00	100.00	100.00
VCTK - WER(%) Source: BERT Target: HuBERT				
0.00	5.58	-	-	-
0.05	5.48	5.84	5.54	6.26
0.10	5.55	22.57	6.17	54.53
0.15	5.85	91.93	12.81	99.83
0.20	6.45	99.42	58.57	100.00
0.25	8.50	99.96	97.15	100.00
ESC-50 - Error(%) Source: ViT Target: BEATs				
0.00	11.75	-	-	-
0.10	11.50	11.76	83.40	98.25
0.15	11.75	12.44	98.00	96.75
0.20	17.50	14.36	98.00	97.00
0.25	16.75	19.12	98.00	98.50
0.30	33.00	26.82	98.00	98.50

Table 1. Error(%) for interpolation of source and target models with factor λ . $\mathcal{N}(\text{Source})$ and $\mathcal{N}(\text{Target})$ are noise sampled with the same mean and variance as parameters of source and target models.

4.1. Case Study 1: Zero-Shot Experiments

The first study evaluates zero-shot performance of merged models on downstream datasets. We use a dev set to select the interpolation λ and report results on a held-out test set.

Attention-Interpolation. Table 1 contains the results for attention merging of entire models as described in Section 3.1 for different interpolation factors λ . We merge source and target models following the strategy outlined in Equation 1. As an additional baseline, we also merge target models with source models sampled from random noise. We experiment with three types of noise - sampling noise from source model parameters $\mathcal{N}(\text{Source})$, target model parameters $\mathcal{N}(\text{Target})$, and a standard normal distribution $\mathcal{N}(0,1)$. For $\mathcal{N}(\text{Source})$ and $\mathcal{N}(\text{Target})$, we draw samples with the same mean and variance as the parameters of the source and target models respectively. The samples are then used to construct attention matrices with the same dimensions as the Target model. The results in Table 1 underscore the clear advantage of merging attention matrices from the source model. Conversely, merging with noise substantially diminishes performance, especially at higher interpolation factors.

Interpolation of subset of layers. Table 2 summarizes the results for merging a subset of layers. We experimented with merging blocks of different sizes and different interpolation but for the sake of brevity present results on blocks of 8 layers for ASR and 4 or 8 layers for AEC with the best performing λ selected using a dev set, which was 0.25 for LJ Speech, 0.05 for VCTK, and 0.1 for ESC-50. For ASR, we identified that merging layers 12-19 resulted in the lowest WER while it was layers 4-11 for AEC.

To determine whether a better subset exists, we employed a data-driven strategy to pinpoint the most suitable layers to merge. We used audio snippets and corresponding text transcripts from 10% of the training set of LJ Speech and VCTK. HuBERT encoded the au-

Layers Merged	LJ Speech $\lambda = 0.25$	VCTK $\lambda = 0.05$	Layers Merged	ESC-50 $\lambda = 0.1$
0 - 7	9.33	5.53	0 - 3	13.00
4 - 11	9.59	5.55	4 - 7	12.50
8 - 15	9.28	5.54	8 - 11	12.00
12 - 19	8.63	5.52	0 - 7	11.75
16 - 23	10.07	5.54	4 - 11	10.50
None	9.25	5.58	None	11.75

Table 2. Model Merging on a subset of layers - WER(%) on LJ Speech and VCTK and Error(%) on ESC-50

dio, while BERT encoded the transcripts. For each sample, hidden representations at the output of each attention block were extracted from every layer of both networks. These hidden representations were averaged by sequence length, resulting in a 1024-dimensional vector per layer for each sample. Similarity between an audio-text pair’s HuBERT and BERT hidden representations for each layer was computed via Euclidean distance, the inner product, and motivated by prior work that compares hidden representations in speech models [37], the Sliced Wasserstein Distance (SWD) [20]. Sorting layers by similarity, the top-k most similar layers were merged. Results for distinct k values and distance metrics are presented in Table 3. For LJ Speech, we find that merging top 6 layers identified using SWD yields best results while it is the top 10 layers for VCTK. Comparing Tables 2 and 3, we notice that merging continuous blocks of layers performs better or almost identical to the data-driven strategy. In contrast, the data-driven strategy does not require extensive experimentation. We did not perform data-driven experiments for AEC due to the absence of paired images for audio snippets in ESC-50.

λ	k	Euclidean Distance	SWD	Inner Product
LJ Speech				
0.1	4	9.18	8.74	9.28
	6	8.87	8.73	9.38
	8	9.19	8.92	9.32
VCTK				
0.05	6	5.53	5.55	5.54
	8	5.51	5.52	5.52
	10	5.51	5.51	5.52

Table 3. Comparison of selecting Top-k most similar layers based on different distance metrics, WER(%) on the ASR task

4.2. Case Study 2: Fine-Tuning Experiments

While previous results were based on zero-shot evaluation, we also examined fine-tuned performance. A comparison of different approaches is provided in Table 5. The top row represents the baseline of downstream evaluation without attention merging or fine-tuning. By employing MAM, we improve the original model’s performance without additional fine-tuning. We report the better of Attention Interpolation (3.1) and Layer-wise Attention Interpolation (3.2) as the best performing MAM approach. Another baseline involves fine-tuning models on target data alone (third row), with results averaged over 3 runs. We fine-tune using AdamW [38] using a learning rate of $1e-5$ for ASR and $1e-3$ for AEC. We fine-tune for 2 epochs on LJ Speech, 1 epoch for VCTK, and 3 epochs for ESC-50 with a batch

	Baseline	MAM
	LJ Speech	
1	the commission believes that the motorcade rout selected by agent lawson upon the advice of agent in charge sorrels	the commission believes that the motorcade route selected by agent lawson upon the advice of agent in charge sorrels
2	here a couple of pye men had been selling their wares the basket of one of them which was raised upon a four legged stool was upset	here a couple of piemen had been selling their wares the basket of one of them which was raised upon a four legged stool was upset
3	besides his employers a jeweler named humphreys was in the swim at whose shop in red lion square was discovered a quantity of bass gold	besides his employers a jeweler named humphreys was in the swim at whose shop in red lion square was discovered a quantity of base gold
4	aproximately thirty to forty five seconds after oswalds lunch room encounter with baker and truley	approximately thirty to forty five seconds after oswalds lunch-room encounter with baker and truly
	VCTK	
5	on the contrary they stant togain	on the contrary they stand to gain
6	flanke gordon simpson may also feature in that much	flanke gordon simpson may also feature in that match
7	the marshal at the tern was great	the marshal at the turn was great
8	they were behind the field	they were behind the wheel

Table 4. Analysis of improvements on LJ Speech and VCTK

size of 32. For ASR, we fine-tune HuBERT using CTC Loss. Fine-tuning on target data outperforms MAM for LJ Speech and ESC-50, while MAM is the better approach for VCTK. We speculate that the similarity in distribution between VCTK (a larger, multi-speaker dataset) and the pre-training data (Libri-light and Librispeech) limits the efficacy of fine-tuning while the single-speaker nature of LJ Speech allows fine-tuning to adapt better to the dataset.

In the situation some data/compute is available, we explore fine-tuning based approaches. We notice improved performance when combining MAM and fine-tuning as compared to using either approach independently. We utilize the best performing model obtained through model-merging (Tables 1 and 2) followed by fine-tuning, shown in the fourth row of Table 5. Our top-performing model emerges when jointly learning the interpolation factor and model weights during fine-tuning, as shown in the last row of Table 5. With L-MAM followed by fine-tuning, we achieve a 2.70% relative WER reduction on LJ Speech, a 2.90% reduction on VCTK, and an 18.42% reduction in classification error on ESC-50 compared to the regular fine-tuning baseline.

4.3. Analysis of Improvements

We now analyze the improvements observed on the ASR task. Table 4 contrasts transcriptions from the baseline (without attention merging or fine-tuning) with those from MAM. MAM improves transcriptions for homophonic words, exemplified by ‘route’,

MAM	FT	L-MAM	WER (%) LJ Speech	WER (%) VCTK	Error (%) ESC-50
-	-	-	9.25	5.58	11.75
✓	-	-	8.63	5.48	10.50
-	✓	-	8.52	5.52	9.50
✓	✓	-	8.40	5.44	7.75
-	✓	✓	8.29	5.36	7.75

Table 5. Ablation study to compare MAM, Fine-Tuning (FT), and Learnable-MAM for ASR and AEC

‘piemen’, and ‘base’ in LJ Speech. VCTK, encompassing various speakers, contains non-standard pronunciations, wherein MAM successfully transcribes ambiguous words accurately. Instances include ‘stand’, ‘match’, ‘turn’, and ‘wheel’, as shown in rows 5 to 8 of Table 4. We further categorize improvements based on character insertion, substitution, or deletion in Table 6. While LJ Speech improvements primarily stem from insertion and substitutions, VCTK shows a more balanced distribution among these types of errors.

Improvement	Insertion	Substitution	Deletion
LJ Speech	61.89 %	35.13 %	2.98 %
VCTK	36.36 %	28.28 %	35.35 %

Table 6. Categorization of Improvements

5. CONCLUSION

This paper introduces Multimodal Attention Merging (MAM), which transfers knowledge from attention matrices of high-resource modality models (e.g., text and images) to low-resource ones (e.g., speech and audio). MAM improves zero-shot performance of HuBERT (ASR) and BEATs (Audio Event Classification) by merging attention matrices with BERT and Vision Transformer. Additionally, Learnable-MAM (L-MAM) jointly learns interpolation and model weights during fine-tuning, achieving up to 2.90% relative WER reduction and 18.42% classification error reduction on ASR and Audio Classification compared to regular fine-tuning, respectively.

While MAM and L-MAM demonstrate improvements, our work is nascent and we believe there are several avenues for future work warranting investigation. Future work could address merging larger, billion-parameter models representing the state of the art. Developing methods to merge models of differing architectures and merging three or more modalities are other dimensions to extend this work. While we merge attention matrices, future work could address the problem from a parameter efficiency perspective by merging adapter modules [39] or low-rank representations [40]. Finally, we believe that equipping merged models with cross-modality capabilities to build a generalized multi-task architecture holds promise and urge future work in this direction.

6. REFERENCES

- [1] Tom Brown et al., “Language models are few-shot learners,” *NeurIPS*, 2020.
- [2] Wei-Ning Hsu et al., “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [3] Alexey Dosovitskiy et al., “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv:2010.11929*, 2020.
- [4] Kevin Lu, Aditya Grover, Pieter Abbeel, and Igor Mordatch, “Pretrained transformers as universal computation engines,” *arXiv:2103.05247*, 2021.
- [5] Chao-Han Huck Yang, Yun-Yun Tsai, and Pin-Yu Chen, “Voice2series: Reprogramming acoustic models for time series classification,” in *ICML*, 2021.
- [6] Sophia Gu, Christopher Clark, and Aniruddha Kembhavi, “I can’t believe there’s no images! learning visual tasks using only language data,” *arXiv:2211.09778*, 2022.
- [7] Machel Reid, Yutaro Yamada, and Shixiang Shane Gu, “Can wikipedia help offline reinforcement learning?,” *arXiv:2201.12122*, 2022.
- [8] Ashish Vaswani, Noam Shazeer, et al., “Attention is all you need,” *NeurIPS*, vol. 30, 2017.
- [9] Ryokan Ri and Yoshimasa Tsuruoka, “Pretraining with artificial language: Studying transferable knowledge in language models,” *arXiv:2203.10326*, 2022.
- [10] Sarah Schwettmann, Neil Chowdhury, and Antonio Torralba, “Multimodal neurons in pretrained text-only transformers,” *arXiv:2308.01544*, 2023.
- [11] Gabriel Goh, Nick Cammarata †, Chelsea Voss †, Shan Carter, Michael Petrov, Ludwig Schubert, Alec Radford, and Chris Olah, “Multimodal neurons in artificial neural networks,” *Distill*, 2021.
- [12] Cheng-Han Chiang and Hung-yi Lee, “On the transferability of pre-trained language models: A study from artificial datasets,” in *AAAI*, 2022.
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” 2019.
- [14] Sanyuan Chen et al., “Beats: Audio pre-training with acoustic tokenizers,” *arXiv:2212.09058*, 2022.
- [15] Yi-Lin Sung, Linjie Li, Kevin Lin, Zhe Gan, Mohit Bansal, and Lijuan Wang, “An Empirical Study of Multimodal Model Merging,” Apr. 2023, *arXiv:2304.14933* [cs].
- [16] Mitchell Wortsman et al., “lo-fi: distributed fine-tuning without communication,” *arXiv:2210.11948*, 2022.
- [17] Leshem Choshen, Elad Venezian, Noam Slonim, and Yoav Katz, “Fusing finetuned models for better pretraining,” *arXiv:2204.03044*, 2022.
- [18] Sidak Pal Singh and Martin Jaggi, “Model fusion via optimal transport,” *NeurIPS*, 2020.
- [19] Mitchell Wortsman et al., “Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time,” in *ICML*, 2022.
- [20] Soheil Kolouri et al., “Generalized sliced wasserstein distances,” *NeurIPS*, 2019.
- [21] Zongsheng Cao et al., “Otkge: Multi-modal knowledge graph embeddings via optimal transport,” *NeurIPS*, 2022.
- [22] Michael S Matena and Colin A Raffel, “Merging models with fisher-weighted averaging,” *NeurIPS*, 2022.
- [23] Junting Zhang, Jie Zhang, Shalini Ghosh, Dawei Li, Serafettin Tasci, Larry P. Heck, Heming Zhang, and C.-C. Jay Kuo, “Class-incremental learning via deep model consolidation,” in *WACV*, 2020.
- [24] Alexandra Chronopoulou, Matthew E Peters, Alexander Fraser, and Jesse Dodge, “Adaptersoup: Weight averaging to improve generalization of pretrained language models,” *arXiv:2302.07027*, 2023.
- [25] Sridhar Mahadevan, Bamdev Mishra, and Shalini Ghosh, “A unified framework for domain adaptation using metric learning on manifolds,” in *ECML-PKDD*, 2019.
- [26] Stanislaw Antol et al., “VQA: Visual question answering,” in *ICCV*, 2015.
- [27] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick, “Microsoft coco: Common objects in context,” in *ECCV*, 2014.
- [28] Gabriel Ilharco et al., “Editing models with task arithmetic,” *arXiv:2212.04089*, 2022.
- [29] Chiyuan Zhang, Samy Bengio, and Yoram Singer, “Are all layers created equal?,” *JMLR*, vol. 23, no. 1, pp. 2930–2957, 2022.
- [30] J. Kahn et al., “Libri-light: A benchmark for asr with limited or no supervision,” in *ICASSP 2020*, 2020, pp. 7669–7673, <https://github.com/facebookresearch/libri-light>.
- [31] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *ICASSP*. IEEE, 2015, pp. 5206–5210.
- [32] Keith Ito and Linda Johnson, “The lj speech dataset,” <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [33] Junichi Yamagishi, Christophe Veaux, and Kirsten MacDonald, “Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit (version 0.92),” 2019.
- [34] Jort F. Gemmeke et al., “Audio set: An ontology and human-labeled dataset for audio events,” in *ICASSP 2017*, 2017.
- [35] Jia Deng et al., “Imagenet: A large-scale hierarchical image database,” in *CVPR*, 2009.
- [36] Karol J. Piczak, “ESC: Dataset for Environmental Sound Classification,” in *ACM Multimedia*, pp. 1015–1018.
- [37] Zih-Ching Chen et al., “How to estimate model transferability of pre-trained speech models?,” *Interspeech*, 2023.
- [38] Ilya Loshchilov and Frank Hutter, “Decoupled weight decay regularization,” *ICLR*, 2019.
- [39] Neil Houlsby et al., “Parameter-efficient transfer learning for nlp,” in *ICML*, 2019.
- [40] Edward J Hu et al., “Lora: Low-rank adaptation of large language models,” *ICLR*, 2022.