

MULTIMODAL TRANSFORMER DISTILLATION FOR AUDIO-VISUAL SYNCHRONIZATION

Xuanjun Chen^{1,2}, Haibin Wu¹, Chung-Che Wang², Hung-yi Lee^{1†}, Jyh-Shing Roger Jang^{2†}

¹Graduate Institute of Communication Engineering, National Taiwan University

²Department of Computer Science and Information Engineering, National Taiwan University
{d12942018, f07921092, hungyilee}@ntu.edu.tw, geniusturtle6174@gmail.com, jang@mirlab.org

ABSTRACT

Audio-visual synchronization aims to determine whether the mouth movements and speech in the video are synchronized. VocaLiST reaches state-of-the-art performance by incorporating multimodal Transformers to model audio-visual interaction information. However, it requires high computing resources, making it impractical for real-world applications. This paper proposed an MTD-VocaLiST model, which is trained by our proposed multimodal Transformer distillation (MTD) loss. MTD loss enables MTDVocaLiST model to deeply mimic the cross-attention distribution and value-relation in the Transformer of VocaLiST. Additionally, we harness uncertainty weighting to fully exploit the interaction information across all layers. Our proposed method is effective in two aspects: From the distillation method perspective, MTD loss outperforms other strong distillation baselines. From the distilled model’s performance perspective: 1) MTDVocaLiST outperforms similar-size SOTA models, SyncNet, and Perfect Match models by 15.65% and 3.35%; 2) MTDVocaLiST reduces the model size of VocaLiST by 83.52%, yet still maintaining similar performance.

Index Terms— Audio-visual synchronization, Transformer distillation, knowledge distillation, lightweight model

1. INTRODUCTION

The audio-visual synchronization task is to determine whether the mouth movements and speech in the video are synchronized. An out-of-sync video may cause errors in many tasks, such as audio-visual user authentication [1], dubbing [2], lip reading [3], active speaker detection [4, 5], and audio-visual source separation [6–9]. An audio-visual synchronization model often acts as an indispensable front-end model for these downstream tasks. Various downstream tasks often run on mobile devices and require small model sizes and fast inference speed. Smaller model sizes and faster inference speed are required to ensure user experience, such as correcting the synchronization error of user-generated videos on mobile phones or performing audio-visual user authentication on finance mobile applications [1]. To work with these applications, a lightweight audio-visual synchronization model is worth exploring.

A typical framework for audio-visual synchronization tasks is estimating the similarity between audio and visual segments. SyncNet [10] introduced a two-stream architecture to estimate the cross-modal feature similarities, which is trained to maximize the similarities between features of the in-sync audio-visual segments and

minimize the similarities between features of the out-of-sync audio-visual segments. Perfect Match (PM) [3, 11] optimizes the relative similarities between multiple audio features and one visual feature with a multi-way matching objective function. Audio-Visual Synchronisation with Transformers (AVST) [12] and VocaLiST [13], the current state-of-the-art (SOTA) models, which incorporate Transformers [14] to learn the multi-modal interaction and classify directly if a given audio-visual pair is synchronized or not, resulting in an excellent performance. However, both AVST and VocaLiST require large memory and high computing costs, making these models unsuitable for edge-device computation.

In this paper, we propose to distill a small-size version of VocaLiST, namely MTDVocaLiST, which is distilled by mimicking the multimodal Transformer behavior of VocaLiST. Next, we propose to employ uncertainty weighting, which allows us to assess the varying significance of Transformer behavior across different layers, resulting in an enhanced MTDVocaLiST model. To our knowledge, this is the first attempt to distill a model by mimicking multimodal Transformer behavior for the audio-visual task. Our model outperforms similar-size state-of-the-art models, SyncNet and PM models by 15.65% and 3.35%. MTDVocaLiST significantly reduces VocaLiST’s size by 83.52% while still maintaining competitive performance comparable to that of VocaLiST.

2. BACKGROUND

2.1. VocaLiST

VocaLiST [13] is a SOTA audio-visual synchronization model. The input of VocaLiST is a sequence of visuals and its corresponding audio features. The output is about to classify whether a given audio-visual pair is in sync. VocaLiST consists of an audio-visual front-end and a synchronization back-end. The audio-visual front end extracts audio and visual features. The synchronization back-end comprises three cross-modal Transformer encoders, namely audio-to-visual (AV) Transformer, visual-to-audio (VA) Transformer, and Fusion Transformer. Each Transformer block has 4 layers. The core part of the cross-modal Transformer is the cross-attention layer, whose input has queries, keys, and values. In VocaLiST, the AV Transformer uses audio features for queries and visual features for keys and values, while the VA Transformer does the opposite. The Fusion Transformer combines these, using the AV output for its queries and the VA output for keys and values. Its output is max-pooled over time dimension and activated via tanh function. Finally, a fully connected layer serves as a classifier, determining if voice and lip motion are synchronized, with optimization through binary cross-entropy loss.

† Equal correspondence. This work was supported by the National Science and Technology Council, Taiwan (Grant no. NSTC 112-2634-F-002-005). Code has been made available at: <https://github.com/xjchenGit/MTDVocaLiST>.

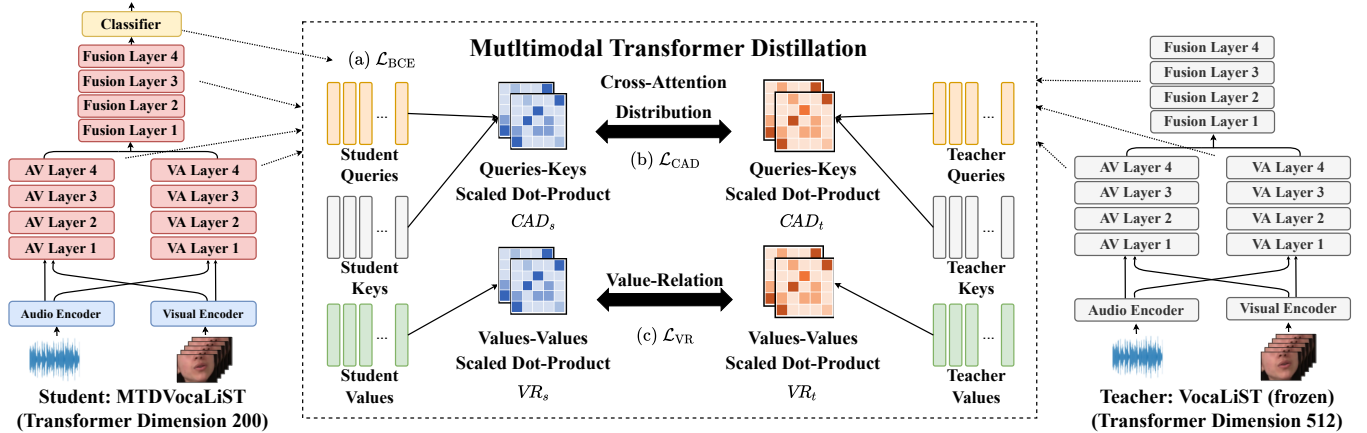


Fig. 1: The proposed MTDVocaLiST model. (a) binary cross entropy loss. (b) cross-attention distribution distillation loss. (c) value-relation distillation loss.

2.2. Knowledge Distillation

Naïve knowledge distillation (KD) [15] will let the student learn temperature-controlled soft targets from the teacher’s output. However, the output representations of the intermediate layers have also been shown to be able to guide the training of student models [16]. Relational knowledge distillation (RKD) [17] also finds rich relational knowledge between different data samples. RKD transfers instance relations modeled by distance and angle relational potential functions. Besides, a series of works [18–29] also aims to distill different knowledge sources based on the teacher-student framework. However, none of the above distillation methods have taken into account the characteristics of the Transformer encoder. In natural language processing, MiniLM [30] attempts to mimic the self-attention of the last Transformer encoder layer from a large self-supervised learning model, which achieves impressive results.

3. MTDVOCALIST

Inspired by MiniLM, we propose a multimodal Transformer distillation VocaLiST (MTDVocaLiST) model, as shown in Fig. 1. The basic idea is to encourage the student model MTDVocaLiST to learn multi-modal interaction behavior from the teacher model VocaLiST in cross-modal Transformers. Prior to the distillation process, both MTDVocaLiST and VocaLiST will extract their respective Transformer behaviors, specifically Cross-Attention Distribution (CAD) and Value-Relation (VR). The loss function for training MTDVocaLiST incorporates the binary cross-entropy loss \mathcal{L}_{BCE} and the Transformer behavior mimic losses (i.e., Equation (2)). The Transformer behavior mimic losses, denoted as \mathcal{L}_{CAD} and \mathcal{L}_{VR} , are designed to facilitate the student’s learning from the teacher. During the distillation, VocaLiST is frozen, and we only train MTDVocaLiST. The teacher uses the public pre-trained model of VocaLiST¹ for initialization. MTDVocaLiST has a similar architecture to VocaLiST, but its embedding dimension is reduced from 512 to 200. MTDVocaLiST and VocaLiST have the same number of layers. VocaLiST has 80.1 million parameters, while MDTVocaLiST has only 13.2 million parameters, reducing the size of the teacher model by 83.52%.

Multimodal Transformer Behaviors. In our framework, Transformer behavior encompasses CAD and VR. CAD measures the

attention relationship between one modal (e.g., audio) with another modal (e.g., visual) in a multimodal task. VR conducts scaled dot product operations among features within the same modality to capture their correlations and dependencies. This enhances the model’s understanding of relationships between different parts within the same modality. The formula for CAD is the same as original attention [14], but we use scaled dot-product value to model the VR instead of just the attention values, which has proven to be more effective [30]. The CAD and VR are formulated as follows:

$$CAD = \text{softmax} \left(\frac{QK^T}{\sqrt{d_q}} \right), \quad VR = \text{softmax} \left(\frac{VV^T}{\sqrt{d_v}} \right) \quad (1)$$

where queries Q , keys K , and values V depend on the type of Transformer block. d_q and d_v represent the number of dimension of Q and V . We formulate the loss function of learning the Transformer behavior as follows:

$$\mathcal{L}_{CAD} = D_{KL}(CAD_s || CAD_t), \quad \mathcal{L}_{VR} = D_{KL}(VR_s || VR_t) \quad (2)$$

where we denote the Transformer behavior CAD_t , VR_t produced by the teacher’s layer via Equation (1) and its corresponding student’s Transformer behavior as CAD_s , VR_s . D_{KL} is the Kullback–Leibler divergence. The Fusion, AV and VA Transformer blocks of the VocaLiST are key blocks for modeling different audio-visual interactions. Since each Transformer layer contains unique information, we’ve developed naïve multimodal Transformer distillation (NMTD) and multimodal Transformer distillation (MTD) to effectively harness these diverse Transformer layers.

Naïve Multimodal Transformer Distillation (NMTD). To enable MTDVocaLiST to comprehensively learn multimodal interaction information at various levels, we aggregate the CAD and VR losses from each layer using weighted sums. Combined with Equation (2), our NMTD loss can be formulated as follows:

$$\mathcal{L}_{NMTD} = w_0 \cdot \mathcal{L}_{BCE} + \sum_l^L w_{l1} \cdot \mathcal{L}_{CAD_l} + w_{l2} \cdot \mathcal{L}_{VR_l}, \quad (3)$$

where L represents a candidate set of layers. w_{l1} and w_{l2} refer to the weights of CAD and VR loss for the l^{th} layer. w_0 is a weight used to control \mathcal{L}_{BCE} . There are three common methods for choosing loss weights for Equation (3): uniform, manual tuning, and automatic

¹<https://github.com/vskadandale/vocalist>

weighting (AW). Uniform weights assign weight one to all terms. However, this approach ignores it there may be some conflict relations between losses. Manual tuning relies on human heuristics to find optimal loss weights but involves trial-and-error. AW utilizes learnable weights for automatic weighting, but it may lead to weight vanishing and requires gradient clipping.

Multimodal Transformer Distillation (MTD). To address the weighting issues, we further introduced the uncertainty weighting (UW) method to improve the student model. MTDVocaLiST’s primary objective was to mimic three distinct types of multimodal Transformer layers, resulting in a total of 12 layers. However, the performance contribution of each layer or type of Transformer to MTDVocaLiST is uncertain. Previous research has shown that employing Bayesian modeling of uncertainty is effective in balancing loss contributions in multi-task learning [31, 32]. Therefore, we can regard the complete multimodal Transformer distillation framework as encompassing multiple distillation tasks alongside a supervised learning task. After utilizing uncertainty weighting, the overall MTD can be formulated as follows:

$$\mathcal{L}_{MTD} = \sum_{\tau}^T \frac{1}{2 \cdot w_{\tau}^2} \cdot \mathcal{L}_{\tau} + \ln(1 + w_{\tau}^2), \quad (4)$$

where T represents a task set, and \mathcal{L}_{τ} denotes the loss for the τ -th task. Within task sets T , there is a scope for \mathcal{L}_{BCE} , along with losses pertaining to various layers of \mathcal{L}_{CAD} and \mathcal{L}_{VR} . The loss of each task is associated with its specific learnable parameter w_{τ} . The term $\ln(1 + w_{\tau}^2)$ serves the purpose of enforcing positive regularization values. The overarching objective can be viewed as learning the relative weights assigned to different tasks. When w_{τ} assumes large-scale values, it diminishes the contribution of \mathcal{L}_{τ} , while small-scale values amplify its contribution.

4. EXPERIMENT

4.1. Experimental setup

Model and dataset. The input of MTDVocaLiST is an audio-visual pair, which corresponds to a video sequence of 5 frames (0.2 seconds) sampled at 25 frames per second. We also train the model with only \mathcal{L}_{BCE} loss using the MTDVocaLiST architecture as the baseline. All models are trained on the LRS2 dataset [33], which contains 96,318 utterances in the pretraining set, 1,082 utterances in the validation set, and 1,243 utterances in the testing set. The maximum number of characters in one utterance is 100.

Training and validation. During the process, positive and negative samples are equally sampled on the fly. Positive samples mean synchronized audio-visual pairs. Negative samples are obtained by introducing random temporal misalignment at offsets within a time scale of ± 15 visual frames (1.2s) between the in-sync audio-visual pairs. The model with the highest F1 score in validation is saved for evaluation. We train for a total of 80 epochs, which is significantly less than the 600 epochs required for VocaLiST training.

Evaluation protocol. The evaluation protocol is similar to the cross-modal retrieval task, which follows previous work [3, 10–13]. Given a 5-frame visual segment and a candidate audio set containing 31 audio segments, the model will predict the corresponding score for each pair of audio-visual segments. The candidate set consists of the target audio segment and 15 left and right neighbors. The target is to find the index of an audio segment most similar to the given 5-frame visual segment. The found index is considered correct if the offset

Table 1: Accuracy of different distillation methods in evaluation.

Distillation method	Input frame length (seconds)					
	5 (0.2s)	7 (0.28s)	9 (0.36s)	11 (0.44s)	13 (0.52s)	15 (0.6s)
\mathcal{L}_{BCE}	71.36	81.44	88.84	93.41	96.19	97.69
KD	80.87	88.62	93.48	96.32	97.90	98.82
RKD	86.06	92.42	95.95	97.80	98.75	99.29
MiniLM*	85.60	92.03	95.91	97.72	98.72	99.25
FitNets	90.81	95.48	97.77	98.81	99.42	99.66
MTD	91.45	95.75	97.99	98.95	99.46	99.68




Fig. 2: Comparison of model size and accuracy.

between it and the ground-truth is within the ± 1 frame since human beings distinguish no difference within the ± 1 frame. Since a short input frame might not be enough to determine the correct offset [10]. Thus, we also conduct experiments using input frame lengths over 5 frames of video (i.e., 5, 7, 9, 11, 13, and 15), averaging the prediction scores over multiple video samples (with a 1-frame temporal stride). In the ablation study, we limit the evaluation to an input frame length of 5 frames due to page space constraints.


4.2. Main Results

Comparison with Different Distillation Methods. We compare MTD loss with several other knowledge distillation methods, including knowledge distillation (KD) [15], relational knowledge distillation (RKD) [17], MiniLM*, and FitNets. As training MiniLM directly does not converge, we use the loss of MiniLM combined with \mathcal{L}_{BCE} for training, denoted as MiniLM*. FitNets uses the last layer representation as hints and guide layers [16]. We replace the MTD loss of MTDVocaLiST with different losses for comparison. Table 1 demonstrates that MTD outperforms other distillation methods for most input frame lengths. When considering an input frame length of 5, training solely with \mathcal{L}_{BCE} results in the lowest accuracy at 71.36%. In contrast, MTD loss significantly improves accuracy, outperforming KD by 10.58%, RKD by 5.39%, MiniLM* by 5.85%, and FitNets by 0.64%. Similar trends are observed across different input frame lengths.

Comparison with SOTA models. We also compare MTDVocaLiST with previous SOTA models, including SyncNet, PM, AVST, and VocaLiST. Fig. 2 presents a graphical depiction of the relationship between model size and accuracy across various input frame lengths. We present the following observations: 1) In equivalent

Table 2: Ablation study of NMTD loss.

Loss	Val F1 (%)	Eval Acc (%)
\mathcal{L}_{BCE}	87.91	71.36
NMTD w/o \mathcal{L}_{VR}	91.78	83.55
NMTD w/o \mathcal{L}_{CAD}	91.97	83.53
NMTD	92.81	85.60

**Fig. 3: Accuracy of different layer selection strategies.**


model size configurations, MTDVocaLiST outperforms SyncNet and PM significantly across all frame settings. Notably, MTDVocaLiST surpasses SyncNet and PM by 15.65% and 3.35%, respectively, when utilizing a 5-frame input length. 2) Remarkably, with only 31.13% of the model size of AVST, MTDVocaLiST achieves superior performance in most input frame length scenarios. 3) Furthermore, MTDVocaLiST utilizes only 16.48% of VocaLiST’s model size while maintaining comparable accuracy. In summary, across all input frame lengths, MTDVocaLiST demonstrates competitive performance against VocaLiST and surpasses AVST, SyncNet, and PM.

4.3. Ablation study and analysis

Indispensability of CAD and VR. To better understand the significance of \mathcal{L}_{CAD} and \mathcal{L}_{VR} , we conducted an ablation study for NMTD loss on the last layer of the fusion Transformer block (i.e., fusion layer 4). We only use uniform weighting. In Table 2, the results reveal that both cross-attention distribution and value-relation contribute significantly to NMTD loss, which outperforms \mathcal{L}_{BCE} .

Layer selection. In Fig. 3, we compare different training strategies, including BCE, single fusion Transformer layers (F1 ~ F4), single AV Transformer layers (AV1 ~ AV4), single VA Transformer layers (VA1 ~ VA4), and layer weighting methods (Uniform, AW-13D, AW-25D, UW-13D, UW-25D). Most implementations are based on NMTD loss except UW which is based on MTD loss. For single-layer distillation and BCE training using Equation (3), we can manually set the weight of specific terms to 1 while setting the weights of the other terms to 0. The 13D and 25D suffixes denote the use of 13-dimensional and 25-dimensional learnable parameters, respectively. In the former case, \mathcal{L}_{CAD} and \mathcal{L}_{VR} losses within each layer share the same parameter, while in the latter, each \mathcal{L}_{CAD} and \mathcal{L}_{VR} losses in different Transformer layers has its own learnable parameter. Employing more learnable parameters will result in better performance.

Fig. 3 yields the following insights: 1) Distilling any Transformer layer significantly improves performance compared to using only BCE loss for supervised learning. 2) Distilling VA Transformer layers consistently performs worse than distilling fusion or AV Transformer layers, suggesting that VA layers contribute minimally to the student’s final performance. 3) Models Uniform,

**Fig. 4: Comparison of Transformer representation and its behavior.**

AW-13D, AW-25D, UW-13D, and UW-25D consistently outperform single-layer distillation and BCE training. The greater the number of layers being distilled, the better the performance. 4) Regarding layer weighting, AW-13D outperforms Uniform by 0.15%, while AW-25D is 0.06% lower than Uniform. In contrast, UW-13D and UW-25D exhibit more stable performance, with UW-25D achieving 91.45%. Notably, our UW method does not require manual tuning. Therefore, we employed the UW-25D layer selection strategy for the best MTDVocaLiST in section 4.2.

Transformer Behavior and Transformer Representation. Both FitNets and MTD losses aim to distill Transformer layers, with MTD loss consistently delivering superior performance in Table 1. To delve deeper into this phenomenon, we conduct further analysis. A typical Transformer layer mainly consists of a multi-head attention layer and a multi-layer perceptron layer. The output of the multi-layer perceptron layer is also known as Transformer representations. There are some differences between MTD loss and FitNets. MTD loss aims to mimic the multi-head attention mechanism, which we refer to as Transformer behavior. FitNets loss aims to learn Transformer representations from the last distilled layers.

We observe the relationship between losses through Fig. 4, where the x-axis corresponds to epochs and the y-axis indicates the loss ratio. The loss ratio means that losses are normalized by their respective maximum values across all epochs. In Fig. 4, our model was solely trained using the MTD loss. To make a fair comparison with FitNets, we also propose a Last-MTD loss, which solely utilized \mathcal{L}_{BCE} loss and the last layer Transformer behavior loss. Fig. 4 represents the loss computation during the validation, where we concurrently calculated the losses for MTD, Last-MTD, and FitNets. We monitored how these losses evolved as the number of epochs increased. Notably, when we focused on optimizing solely the MTD loss, the Last-MTD loss exhibited a concurrent decrease along with the MTD loss. In contrast, the FitNets loss did not demonstrate a significant decrease. This observation suggests that what is distilled from the representation layers and what is distilled from imitating the Transformer behavior are distinct aspects of the Transformer.

5. CONCLUSION

This work introduces MTDVocaLiST, a model that deeply learns to mimic the cross-attention distribution and value-relation of VocaLiST. Enhanced by uncertainty weighting, our MTD loss utilizes the relative importance of Transformer behavior across various layers, leading to an improved model. Experimental results show that MTD loss outperforms other strong distillation baselines, and MTDVocaLiST maintains competitive performance while reducing the teacher model’s size by 83.52%. Notably, our model outperforms SyncNet and PM models of similar sizes by 15.65% and 3.35%, respectively.

6. REFERENCES

- [1] Amit Aides and Hagai Aronowitz, "Text-dependent audiovisual synchrony detection for spoofing detection in mobile person recognition," in *Interspeech*, 2016, vol. 2, p. 4.
- [2] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar, "A lip sync expert is all you need for speech to lip generation in the wild," in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020, pp. 484–492.
- [3] Soo-Whan Chung, Joon Son Chung, and Hong-Goo Kang, "Perfect match: Improved cross-modal embeddings for audio-visual synchronisation," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 3965–3969.
- [4] Ruijie Tao, Zexu Pan, Rohan Kumar Das, Xinyuan Qian, Mike Zheng Shou, and Haizhou Li, "Is someone speaking? exploring long-term temporal features for audio-visual active speaker detection," in *Proceedings of the 29th ACM International Conference on Multimedia*, 2021, pp. 3927–3935.
- [5] Xuanjun Chen, Haibin Wu, Helen Meng, Hung-yi Lee, and Jyh-Shing Roger Jang, "Push-pull: Characterizing the adversarial robustness for audio-visual active speaker detection," in *2022 IEEE Spoken Language Technology Workshop (SLT)*, 2023, pp. 692–699.
- [6] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman, "Self-supervised learning of audio-visual objects from video," in *European Conference on Computer Vision*. Springer, 2020, pp. 208–224.
- [7] Andrew Owens and Alexei A Efros, "Audio-visual scene analysis with self-supervised multisensory features," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 631–648.
- [8] Lingyu Zhu and Esa Rahtu, "Visually guided sound source separation and localization using self-supervised motion representations," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2022, pp. 1289–1299.
- [9] Zexu Pan, Ruijie Tao, Chenglin Xu, and Haizhou Li, "Selective listening by synchronizing speech with lips," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 1650–1664, 2022.
- [10] Joon Son Chung and Andrew Zisserman, "Out of time: automated lip sync in the wild," in *Asian conference on computer vision*. Springer, 2016, pp. 251–263.
- [11] Soo-Whan Chung, Joon Son Chung, and Hong-Goo Kang, "Perfect match: Self-supervised embeddings for cross-modal retrieval," *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 3, pp. 568–576, 2020.
- [12] Honglie Chen, Weidi Xie, Triantafyllos Afouras, Arsha Nagraani, Andrea Vedaldi, and Andrew Zisserman, "Audio-visual synchronisation in the wild," *arXiv preprint arXiv:2112.04432*, 2021.
- [13] Venkatesh S Kadandale, Juan F Montesinos, and Gloria Haro, "Vocalist: An audio-visual synchronisation model for lips and voices," *arXiv preprint arXiv:2204.02090*, 2022.
- [14] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [15] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al., "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, vol. 2, no. 7, 2015.
- [16] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio, "Fitnets: Hints for thin deep nets," *arXiv preprint arXiv:1412.6550*, 2014.
- [17] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho, "Relational knowledge distillation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 3967–3976.
- [18] Sergey Zagoruyko and Nikos Komodakis, "Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer," *arXiv preprint arXiv:1612.03928*, 2016.
- [19] Frederick Tung and Greg Mori, "Similarity-preserving knowledge distillation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 1365–1374.
- [20] Baoyun Peng, Xiao Jin, Jiaheng Liu, Dongsheng Li, Yichao Wu, Yu Liu, Shunfeng Zhou, and Zhaoning Zhang, "Correlation congruence for knowledge distillation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 5007–5016.
- [21] Sungsoo Ahn, Shell Xu Hu, Andreas Damianou, Neil D Lawrence, and Zhenwen Dai, "Variational information distillation for knowledge transfer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9163–9171.
- [22] Nikolaos Passalis and Anastasios Tefas, "Learning deep representations with probabilistic knowledge transfer," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 268–284.
- [23] Byeongho Heo, Minsik Lee, Sangdoon Yun, and Jin Young Choi, "Knowledge transfer via distillation of activation boundaries formed by hidden neurons," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019, vol. 33, pp. 3779–3787.
- [24] Jangho Kim, SeongUk Park, and Nojun Kwak, "Paraphrasing complex network: Network compression via factor transfer," *Advances in neural information processing systems*, vol. 31, 2018.
- [25] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim, "A gift from knowledge distillation: Fast optimization, network minimization and transfer learning," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4133–4141.
- [26] Zehao Huang and Naiyan Wang, "Like what you like: Knowledge distill via neuron selectivity transfer," *arXiv preprint arXiv:1707.01219*, 2017.
- [27] Yonglong Tian, Dilip Krishnan, and Phillip Isola, "Contrastive representation distillation," in *International Conference on Learning Representations*, 2020.
- [28] Liqun Chen, Dong Wang, Zhe Gan, Jingjing Liu, Ricardo Henao, and Lawrence Carin, "Wasserstein contrastive representation distillation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 16296–16305.
- [29] Yen-Lun Liao, Xuanjun Chen, Chung-Che Wang, and Jyh-Shing Roger Jang, "Adversarial speaker distillation for countermeasure model on automatic speaker verification," in *Proc. 2nd Symposium on Security and Privacy in Speech Communication*, 2022, pp. 30–34.
- [30] Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou, "Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers," *Advances in Neural Information Processing Systems*, vol. 33, pp. 5776–5788, 2020.
- [31] A Kendall, Y Gal, and R Cipolla, "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. arxiv," *arXiv preprint arXiv:1705.07115*, 2017.
- [32] Lukas Liebel and Marco Körner, "Auxiliary tasks in multi-task learning," *arXiv preprint arXiv:1805.06334*, 2018.
- [33] Triantafyllos Afouras, Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman, "Deep audio-visual speech recognition," *IEEE transactions on pattern analysis and machine intelligence*, 2018.