


Efficient Multi-Model Fusion with Adversarial Complementary Representation Learning

Zuheng Kang[†], Yayun He[†], Jianzong Wang, Junqing Peng, Jing Xiao
Ping An Technology (Shenzhen) Co., Ltd., Shenzhen, China.

Abstract—Single-model systems often suffer from deficiencies in tasks such as speaker verification (SV) and image classification, relying heavily on partial prior knowledge during decision-making, resulting in suboptimal performance. Although multi-model fusion (MMF) can mitigate some of these issues, redundancy in learned representations may limit improvements. To this end, we propose an adversarial complementary representation learning (ACoRL) framework that enables newly trained models to avoid previously acquired knowledge, allowing each individual component model to learn maximally distinct, complementary representations. We make three detailed explanations of why this works and experimental results demonstrate that our method more efficiently improves performance compared to traditional MMF. Furthermore, attribution analysis validates the model trained under ACoRL acquires more complementary knowledge, highlighting the efficacy of our approach in enhancing efficiency and robustness across tasks.

Index Terms—model fusion, ensemble learning, alliance learning, speaker verification, complementary representations

I. INTRODUCTION


Multi-model fusion (MMF) has demonstrated great potential to achieve superior overall performance compared to individual, as distinct component models may contain complementary capabilities to avoid their limitations.

Despite various applications and tasks using quite different architectures and processing methods, there are commonalities exist in the core logic – MMF for the model exists throughout the model inference pipeline, including input data_a, early, mid, and late stages of the model_b, and final inference output_c [1]. (a) Data-level fusion can merge datasets with fully or partially overlapping labels at the label level to incorporate more training data [2]. Additionally, multi-task learning enables joint training on datasets with disparate labels using different optimization objectives [2]. Data augmentation via generative models like generative adversarial networks (GANs) also enables the fusion of real and synthetic data [3]. (b) Model-level fusion utilizes intermediate representations for fusion. Early fusion leverages multiple engineered features to allow models to fully exploit information in the data [4]–[9]. Mid-term fusion means that each type of data is first processed through its own network and then fused at some intermediate modeling layer. Late fusion is more flexible to aggregate high-level features of individually trained models optimized under different conditions. Varying training induces slightly different model specializations, so integrating them may improve

overall performance [3], [10]–[12]. These models are typically fused by concatenation, addition, multiplication, or attention mechanism. (c) Output-level fusion commonly integrates the predictions of multiple pre-trained models to achieve improved performance [13]–[15]. This technique is frequently employed in various competitions. For example, in speaker verification (SV) challenges, fusion of scoring results is often used in the optimization pipeline [16]–[18], deepfake detection [19], and many other challenges [20]. However, although MMF is effective in improving overall system performance, it has some limitations. The constituent models used in MMF are often very similar in nature, which limits the effectiveness of model integration. Also, integrating too many models significantly increases computational costs, making it infeasible in resource-constrained environments.

There are a number of similar methods that attempt to train a new model that avoid previously learned knowledge, but they vary in their approach, and purpose. Shen et al. [21] proposed an adversarial learning based model distillation method, which aims to incorporate the student model to learn the knowledge of the teacher’s network, but not to learn its similar knowledge from the corresponding layers in the teacher’s model. However, the problem with this framework is that the student model is only learning the teacher’s network, and the incorrect knowledge learned from the teacher’s network is not modified by correct labeling, which may lead to the magnifying the incorrect knowledge. Nam et al. [22] attempted to utilize the perturbation method so that the student model would absorb as much knowledge as possible from the individual teacher models, thereby distilling a single model that would perform as a complete set. Although this work distills a model of learning knowledge that is as different as possible from the teacher’s knowledge, the model’s performance is only close to that of the previous ensembled model, but does not surpass it. Zhang et al. [23], most similar to our idea, proposed an adversarial based object localization framework that avoid area that the previous models had learned. However, this method only apply to the object detection task, and it cannot apply to other tasks. Most similar to our idea, Zhang et al. [23] proposed an adversarial-based object localization framework which avoids regions learned by previous models. However, this approach is only applicable to the object localization task and can not be generalized, and cannot apply to the models more than two. There are many other researches about ensemble learning that are meaningful to the form of our idea [24]–[28].

Also, in machine learning, bagging and boosting are popular

 Corresponding author: Zuheng Kang, bob.kang@qq.com

[†] Both authors have equal contributions.

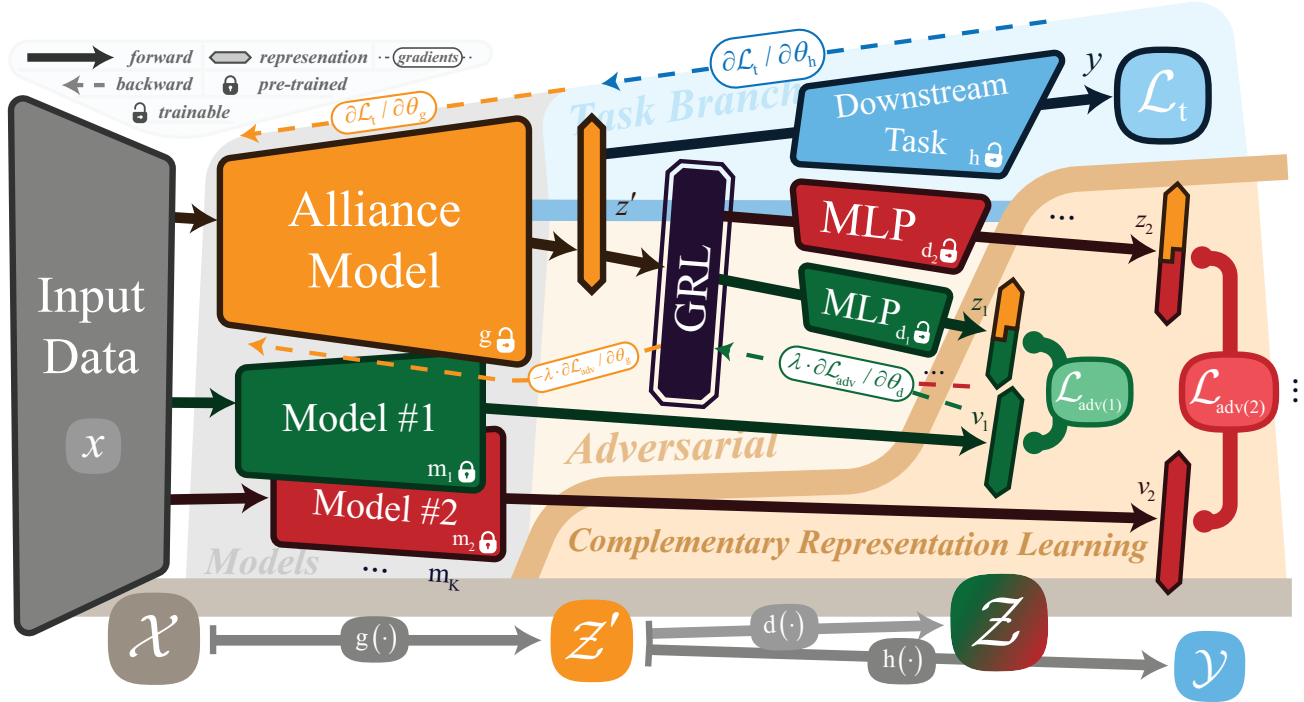


Fig. 1: The overview of ACoRL framework.

ensemble methods that improve system performance. Bagging trains multiple models on different bootstrap samples of the training data to reduce variance. Boosting sequentially fits models to emphasize previously misclassified instances, thereby reducing bias. However, random sampling of bagging data still does not solve the training redundancy problem. Boosting uses a sequential approach that limits parallelization and reduces efficiency. Additionally, some tasks are not amenable to boosting’s error-focused learning.

To fully address the challenges of MMF, the analytical lens cannot be narrowly confined. Let’s start with a story: When a person or group seeks allies to maximize their benefits, they tend to ally with someone who is as different as possible from themselves in terms of knowledge, vision, and ability. However, if they are free to acquire knowledge from a single source, their acquired knowledge tends to closely resemble that of others. Allying among them may risk substantial redundancy. Furthermore, too many models require much more computational resources.

Therefore, we have made the following contributions:

- We propose an adversarial complementary representation learning (ACoRL) framework that promotes diversity during multi-model fusion by enabling models to avoid previously acquired knowledge and learn distinct representations.
- We theoretically prove and explain how ACoRL can improve the performance of multi-model fusion (MMF) by extending the range of representations in the latent space.
- Experimental results and attribution analysis validate that

ACoRL can leverage more complementary knowledge, strengthening its ability to improve model performance across tasks.

II. METHODOLOGY

A. The Need for Complementary Representations

In supervised learning frameworks, source data is generally transformed by the model into a progression of shallow to deep latent representations, which are eventually mapped to downstream tasks. For a specific task, when training a single model multiple times, variance in model performance is primarily attributable to data quality and differences in the models. However, for the same data, since the training of a single model is heavily constrained by the optimization objective, newly trained models tend to focus on the most readily learnable and salient aspects of the training data, while neglecting other potentially useful information. This likely results in a high degree of consistency in the knowledge acquired across models. In other words, representations are highly similar from the model perspective. This phenomenon has been widely observed by many other researchers [21]–[23]. Therefore, it is necessary to ensure that these representations are as complementary as possible so that the individual models can extract the potential knowledge from the training data as comprehensively as possible. Meanwhile, downstream tasks can still be accomplished with these alternative representations.

B. Adversarial Complementary Representation Learning

To obtain efficient fusion system capability, we need to supplement the newly allied models with the strengths and

weaknesses of the existing models. To this end, we propose an adversarial complementary representation learning (ACoRL; alias Alliance Learning), which allows the newly allied model to avoid learning previously acquired knowledge as much as possible. Note: (1) **alliance learning** in this paper is the same as adversarial complementary representation learning (ACoRL) framework. (2) **alliance model** is the newly trained model that avoid previously learned knowledge. Seen from Figure 1 upper part, the ACoRL framework consists of 2 main branches: the branch with K pre-trained models $m_k(\cdot)$ (discussed in § II-B), where k is the index, and the branch with trainable alliance model $g(\cdot)$ with model parameters θ_g (in § II-B).

Pre-trained Models Knowledge Avoidance: The input data x is first transformed into a set of intermediate representations v_k by its corresponding pre-trained models $m_k(\cdot)$, shown as the green and red branches in Figure 1. These representations v_k will be used to represent the knowledge known by the corresponding pre-trained model m_k to avoid learning by the newly allied model.

Alliance Model Training with ACoRL: The input data x is then converted to the representations z' via the alliance model $g(\cdot)$, as shown in the orange branch in Figure 1. Then, z' is sent to 2 sub-branches:

- (i) **Task branch:** z' accomplishes the task training through the downstream task model $h(\cdot)$ with parameters θ_h , as shown in the light blue branch. The task branch loss and its function are denoted as \mathcal{L}_t and ℓ_t , in Equation 1, which depends on the specific task, where y is the label.

$$\mathcal{L}_t = \ell_t(h \circ g(x), y). \quad (1)$$

- (ii) **ACoRL branch:** As shown in light brown background area in Figure 1. To compare between the representations of the alliance model z' and the pre-trained model v_k , they should be dimensionally consistent. For this, the projection model $d_k(\cdot)$ implemented as a multi-layer perceptron (MLP) with parameters θ_{d_k} is used to project z' onto z_k with the same dimensions as v_k . To ensure that the alliance model avoids learning the knowledge in each pre-trained model, an adversarial approach is well-suited. Specifically, the projection model aims to ensure the alliance model representation z_k to match the pre-trained model representation v_k as similarly as possible. In contrast, the alliance model will produce representations that are as complementary as possible. To achieve this, the gradient reversal layer (GRL) [29] can invert the gradient, which enables the projection model to maximize knowledge similarity with the pre-trained model, while the alliance model minimizes it. The ACoRL branch loss and its function corresponding to model m_k are denoted as $\mathcal{L}_{adv(k)}$ and $\ell_{adv(k)}$, in Equation 2. Where \circ denotes the mapping function.

$$\mathcal{L}_{adv(k)} = \ell_{adv(k)}(d_k \circ g(x), m_k(x)). \quad (2)$$

Overall Loss: The hyper-parameter λ is introduced to balance the training of the two task branches, and then the overall loss \mathcal{L} is shown in Equation 3. In this way, the alliance model receives both the gradient of the task branch $\partial \mathcal{L}_t / \partial \theta_g$ for task learning, and the reversed gradient of ACoRL branch $-\lambda \cdot \partial \mathcal{L}_{adv} / \partial \theta_g$ for complementary knowledge acquiring. The downstream task model receives the gradient $\partial \mathcal{L}_t / \partial \theta_h$ for task training. The projection models receive the gradient $\lambda \cdot \partial \mathcal{L}_{adv} / \partial \theta_d$ for representation learning, shown in Figure 1.

$$\mathcal{L} = \max_{\theta_d} \min_{\theta_g, \theta_h} \left(\mathcal{L}_t - \frac{\lambda}{K} \sum_{k=1}^K \mathcal{L}_{adv(k)} \right). \quad (3)$$

C. Motivation for ACoRL and Explanations

In order to justify our approach, we identified three explanations, which are the main motivations for proposing the ACoRL framework. In the end, we proved experimentally that our proposed ACoRL are correct.

Explain 1: Fundamentally, the mechanism of multi-modal and multi-model fusion for improving the system performance are intrinsically the same.

Multimodality leverages naturally complementary information from distinct physical modalities, such as audio and visual, to improve system performance. In contrast, unimodal MMF jointly optimizes multiple models from the same data source to learn slightly different but complementary knowledge, thus making the system more robust. Both approaches attempt to learn a more comprehensive knowledge to improve the overall performance of the system. However, knowledge learned by multi-model is more distinct than that learned from multi-modal.

Explain 2: If the sample size is large enough, models that cover a larger and more complete latent representation will have better performance.

According to Explain 1, the explain of multi-modal fusion also applies to MMF. Yu H. et al. [30] theoretically proved that with sufficient data, the use of more modalities can effectively improve the quality of the latent representation by covering larger latent space, thereby improving the overall system performance. Similarly, it can be hypothesized that the MMF system could perform better if it is able to obtain a more expansive and comprehensive representation in the latent space.

Explain 3: ACoRL framework has the potential for larger, higher quality latent representations.

The upper part of Figure 2 and the lower part of Figure 1 are the same. The large background rectangle represents the ideal maximum boundary of the latent space for accomplishing a given task. There are three spaces, the input space \mathcal{X} , intermediate latent space \mathcal{Z}' and model specific latent space $\mathcal{Z}_{(k)}$. There exists subspace of representations (SoR) \mathcal{Z}'_k in the latent space \mathcal{Z}' with its data z' , SoR \mathcal{Z}_k and \mathcal{V}_k in the latent




Fig. 2: Illustration of the process of Explain 3.

space $\mathcal{Z}_{(k)}$ with its data z_k and v_k . There are two mappings by $g: \mathcal{X}_k \mapsto \mathcal{Z}'_k$ and $d_k: \mathcal{Z}'_k \mapsto \mathcal{Z}_k$.

The back-propagation of ACoRL has 3 processes:

- (i) Projection model d_k tries to modify SoR \mathcal{Z}_k to be highly coincident with \mathcal{V}_k , where \mathcal{V}_k is the SoR for pre-trained model m_k . Then \mathcal{Z}_k can fully represent \mathcal{V}_k .
- (ii) Assuming at this training state, the latent space where the representation of the Alliance Model can maximally reach in \mathcal{Z}' is denoted as \mathcal{Z}'_A . To ensure the mapping from \mathcal{Z}'_k to \mathcal{Z}_k , the model d_k will have a positive gradient $\lambda \cdot \partial \mathcal{L}_{adv} / \partial \theta_d$. However, due to GRL, in the mapping from \mathcal{X}_k to \mathcal{Z}_k , the gradient is inverted to be $-\lambda \cdot \partial \mathcal{L}_{adv} / \partial \theta_g$. Since \mathcal{Z}'_A originally had a negative gradient, it is now positive. Similarly, \mathcal{Z}'_k becomes negative. This indicates that model g will actively search for more complementary representations in the latent space under ACoRL. Due to the SoR of the newly trained model \mathcal{Z}'_{k+1} is highly consistent with the original \mathcal{Z}'_k , and yet ACoRL can greatly enlarge this SoR. According to Explain 2, since $\mathcal{Z}'_A \cup \mathcal{Z}'_k \gg \mathcal{Z}'_{k+1} \cup \mathcal{Z}'_k$, the ACoRL is theoretically proved to have much superior performance.
- (iii) The task branch has a mapping by $h: \mathcal{Z}' \mapsto \mathcal{Y}$ to the task space \mathcal{Y} by alternative SoR of \mathcal{Z}'_A .

III. EXPERIMENT

A. Experimental Setup

Experiment 1 (image classification; IC): The task is trained and tested on ImageNet-100 [31] train and validation set. Function ℓ_t is softmax for classification. The models are ResNet-34 (A), MobileNet-V2 (B) and ResNet-18 (C). The test metrics are reported using top-1 accuracy (Acc) in percentage (%).

Experiment 2 (speaker verification; SV): This task is trained on VoxCeleb1+2 [32] and tested on VoxCeleb1-O. Function ℓ_t is AAM-Softmax [33] for speaker embedding (SE) training. The models are ECAPA-TDNN [34] (A), x-vector [35] (B) and ResNet-34 [36] (C). Models A, B, and C use log mel-spectrograms with 80, 24, and 80 dimensions, respectively. The test metrics are presented using the equal error rate (EER) in percentage (%).

Training Details: For IC, the second last layer features are flattened and used as the representation. For SV, the SE is used as the representation. The AAM-Softmax margin is 0.2 with the scale of 32. We consider two fusion methods – late fusion (L.F.) and output fusion (O.F.). For L.F., the MMF is implemented by concatenating the representations and passing them through a 3-layer MLP with 512 nodes in the hidden layer. For O.F., the IC task uses a weighted sum of classifier outputs while the SV task uses logistic regression on the score outputs. Similar to the knowledge distillation framework [37] for knowledge learning, the ACoRL task function ℓ_{adv} is implemented with Kullback-Leibler Divergence (KL) as the objective. The value of the hyperparameter $\lambda = 1$. As the ACoRL training is performed sequentially, model B is trained on top of model A, and model C on top of models A and B. We then evaluate the fusion of A+B and A+B+C models.

B. Evaluation Results

Experiment 1: Four conclusions can be drawn from the Table I:

- (i) Although newly trained individual models exhibit slightly better performance than their ACoRL-trained counterparts, fusing ACoRL-trained models leads to much superior overall performance. We hypothesize that each individually trained model may over-focus on similar, redundant features, whereas ACoRL diversifies this focus.

Top-1 Acc(%)		A	B	C	A+B	A+B+C
MMF	L. F.	80.62	81.74	77.54	82.40	82.76
	O. F.				81.98	82.26
+ACoRL	L. F.		79.62	75.82	83.94	84.22
	O. F.				81.18	81.80

TABLE I: Overall comparison on the image classification task.

EER(%)		A	B	C	A+B	A+B+C
MMF	L. F.	0.90	2.97	1.07	0.84	0.81
	O. F.				0.87	0.86
+ACoRL	L. F.		3.12	1.24	0.78	0.73
	O. F.				0.80	0.77

TABLE II: Overall comparison on speaker verification task.

- (ii) L.F. outperforms the O.F. In our analysis, in O.F., only a single weight per output is excessively less accurate when fusing diverse models. In contrast, L.F. more carefully integrates and learns from more primitive information, thereby leveraging complementary information to improve the system performance.
- (iii) Incrementally aggregating models can improve performance, but the gains diminish as more models are added. Nevertheless, the ACoRL framework exhibits considerably greater performance improvement over traditional MMF.
- (iv) Interestingly, the fusion of 2 models A+B under ACoRL is even better than that of A+B+C under MMF. This indicates that ACoRL is more efficient in MMF, i.e., fusing only a few models with a small amount of computation makes a big difference.

Experiment 2: Table I could also show the 4 conclusions from experiment 1. However, the performance gains conferred by L.F. over O.F. are limited for SV, compared to IC. This may be attributed to two factors:

- (i) SV in this experiment employs cosine similarity for scoring, which is a parameter-free process. Thus, the information contained in the SV scoring results is highly consistent with the SEs.
- (ii) The strong correlation between SV scores and SEs limits the potential for L.F. to outperform over O.F. Moreover, L.F. requires more sophisticated training and inference procedures than O.F., with greater computational cost. This is why SV typically uses O.F. instead of L.F. in MMF. In comparison, in IC, there is a complex mapping relationship between representations and output predictions. Generally, lower-level representations generally encapsulate richer information. Therefore, L.F. can improve performance more notably than O.F. for IC.

Attribution Analysis: The visualization in Figure 3 (with grad-CAM [40] method) shows that before the addition of ACoRL, models A and B focused on similar regions. ACoRL diminishes the current attention and increases focus on complementary, task-relevant aspects, implying that this newly




Fig. 3: Attribution analysis of image samples on ImageNet-100 by GradCAM [38] method. From top to bottom are the newly trained model A, model B, and model B trained under ACoRL on pre-trained model A. Red regions represent focus.




Fig. 4: Attribution analysis of audio samples on VoxCeleb by integrated gradients [39] method. Three rows are the same as Figure 3. The blue and red regions represent positive and negative gradients. (click on the figure to hear the sound)

trained alliance model leverages previously unfocused information to accomplish the task in an alternative way. Similar to the attribution analysis of SV by J. Zhang et al. [41], Figure 4 (with integrated gradient [42] method) indicates SV focuses on a more fine-grained, abstract region, but reaches the same conclusion. Compared with the second row (newly trained model B) in Figure 4, the most notable regions on the third row (model B under ACoRL) are diminished, but instead, many other interesting small regions appear. These newly added regions contain complementary knowledge, which is essential for improving the performance of the MMF.

IV. CONCLUSIONS

This paper proposed an adversarial complementary representation learning framework (ACoRL) that enables component models to learn maximally distinct but complementary representations, which is further confirmed by later attribution analysis. Both explains and experimental results demonstrate that our method is more efficient and outperforms conventional methods. ACoRL provides a generalizable method for improving the efficiency and robustness of multi-model fusion (MMF), offering a new avenue for future research.

V. ACKNOWLEDGEMENT

Supported by the Key Research and Development Program of Guangdong Province (grant No. 2021B0101400003) and the Corresponding author is Zuheng Kang (bob.kang@qq.com).

REFERENCES

- [1] Shih-Cheng Huang, Anuj Pareek, Saeed Seyyedi, Imon Banerjee, and Matthew P Lungren, "Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines," *NPJ digital medicine*, vol. 3, no. 1, pp. 136, 2020.
- [2] Zuheng Kang, Junqing Peng, Jianzong Wang, and Jing Xiao, "Speechq: Speech emotion recognition based on multi-scale unified datasets and multitask learning," in *Conference of the International Speech Communication Association (INTERSPEECH)*, 2022.
- [3] Shan Lin, Xingzhong Nong, Jianqiang Luo, and Chenen Wang, "A novel multi-model stacking ensemble learning method for metro traction energy prediction," *IEEE Access*, vol. 10, pp. 129231–129244, 2022.
- [4] Xizhan Gao, Quansen Sun, Haitao Xu, Dong Wei, and Jianqiang Gao, "Multi-model fusion metric learning for image set classification," *Knowledge Based System*, vol. 164, pp. 253–264, 2019.
- [5] Junyi Peng, Xiaoyang Qu, Rongzhi Gu, Jianzong Wang, Jing Xiao, Lukas Burget, and Jan Cernocky, "Effective phase encoding for end-to-end speaker verification," in *Conference of the International Speech Communication Association (INTERSPEECH)*, 2021, pp. 2366–2370.
- [6] Guangxia Xu, Weifeng Li, and Jun Liu, "A social emotion classification approach using multi-model fusion," *Future Generation Computer Systems*, vol. 102, pp. 347–356, 2020.
- [7] Srinivasu Polinati and Ravindra Dhuli, "A review on multi-model medical image fusion," in *International Conference on Communication and Signal Processing (ICCSPP)*, IEEE, 2019, pp. 0554–0558.
- [8] Mingqiang Lin, Denggao Wu, Jinhao Meng, Ji Wu, and Haitao Wu, "A multi-feature-based multi-model fusion method for state of health estimation of lithium-ion batteries," *Journal of Power Sources*, vol. 518, pp. 230774, 2022.
- [9] Zhongda Tian and Hao Chen, "Multi-step short-term wind speed prediction based on integrated multi-model fusion," *Applied Energy*, vol. 298, pp. 117248, 2021.
- [10] C. Li, Dongliang He, Xiao Liu, Yukang Ding, and Shilei Wen, "Adapting image super-resolution state-of-the-arts and learning multi-model ensemble for video super-resolution," *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2033–2040, 2019.
- [11] Sagar Deep Deb, Rajib Kumar Jha, Kamlesh Jha, and Prem S Tripathi, "A multi model ensemble based deep convolution neural network structure for detection of covid19," *Biomedical signal processing and control*, vol. 71, pp. 103126, 2022.
- [12] Xizhan Gao, Quansen Sun, Haitao Xu, Dong Wei, and Jianqiang Gao, "Multi-model fusion metric learning for image set classification," *Knowledge-Based Systems*, vol. 164, pp. 253–264, 2019.
- [13] Zijun Deng, Lei Zhu, Xiaowei Hu, Chi-Wing Fu, Xuemiao Xu, Qing Zhang, Jing Qin, and Pheng-Ann Heng, "Deep multi-model fusion for single-image dehazing," *International Conference on Computer Vision (ICCV)*, pp. 2453–2462, 2019.
- [14] Kui Jiang, Zhongyuan Wang, Peng Yi, Guangcheng Wang, Ke Gu, and Junjun Jiang, "Atmf: Adaptive-threshold-based multi-model fusion network for compressed face hallucination," *IEEE Transactions on Multimedia*, vol. 22, pp. 2734–2747, 2020.
- [15] Ying Zhang, Rongrong Zhang, Qunfei Ma, Yanhao Wang, Qingqing Wang, Zihao Huang, and Linyan Huang, "A feature selection and multi-model fusion-based approach of predicting air quality," *ISA transactions*, vol. 100, pp. 210–220, 2020.
- [16] Jaesung Huh, Andrew Brown, Jee-weon Jung, Joon Son Chung, Arsha Nagrani, Daniel Garcia-Romero, and Andrew Zisserman, "Voxsrc 2022: The fourth voxceleb speaker recognition challenge," *arXiv preprint arXiv:2302.10248*, 2023.
- [17] Seyed Omid Sadjadi, Craig Greenberg, Elliot Singer, Lisa Mason, and Douglas Reynolds, "The 2021 nist speaker recognition evaluation," *arXiv preprint arXiv:2204.10242*, 2022.
- [18] Miao Zhao, Yufeng Ma, Min Liu, and Mingqiang Xu, "The speakin system for voxceleb speaker recognition challenge 2021," *arXiv preprint arXiv:2109.01989*, 2021.
- [19] Jiangyan Yi, Chenglong Wang, Jianhua Tao, Xiaohui Zhang, Chu Yuan Zhang, and Yan Zhao, "Audio deepfake detection: A survey," *arXiv preprint arXiv:2308.14970*, 2023.
- [20] Yixin Liao, Yuxuan Peng, Songlin Shi, Victor Guang Shi, and Xiaohong Yu, "Early box office prediction in china's film market based on a stacking fusion model," *Annals of Operations Research*, vol. 308, pp. 321 – 338, 2020.
- [21] Zhiqiang Shen, Zhankui He, and X. Xue, "Meal: Multi-model ensemble via adversarial learning," *Association for the Advancement of Artificial Intelligence (AAAI)*, 2019.
- [22] Gi Cheon Nam, Jongmin Yoon, Yoonho Lee, and Juho Lee, "Diversity matters when learning from ensembles," in *Neural Information Processing Systems*, 2021.
- [23] Xiaolin Zhang, Yunchao Wei, Jiashi Feng, Yi Yang, and Thomas S. Huang, "Adversarial complementary learning for weakly supervised object localization," *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1325–1334, 2018.
- [24] Yue Cao, Thomas Andrew Geddes, Jean Yee Hwa Yang, and Pengyi Yang, "Ensemble deep learning in bioinformatics," *Nature Machine Intelligence*, vol. 2, no. 9, pp. 500–508, 2020.
- [25] Mudassir A Ganaie, Minghui Hu, AK Malik, M Tanveer, and PN Suganthan, "Ensemble deep learning: A review," *Engineering Applications of Artificial Intelligence*, vol. 115, pp. 105151, 2022.
- [26] Xibin Dong, Zhiwen Yu, Wenming Cao, Yifan Shi, and Qianli Ma, "A survey on ensemble learning," *Frontiers of Computer Science*, vol. 14, pp. 241–258, 2020.
- [27] Ammar Mohammed and Rania Kora, "A comprehensive review on ensemble deep learning: Opportunities and challenges," *Journal of King Saud University-Computer and Information Sciences*, 2023.
- [28] Sanping Zhou, Jinjun Wang, Le Wang, Xingyu Wan, Siqi Hui, and Nanning Zheng, "Inverse adversarial diversity learning for network ensemble," *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [29] Yaroslav Ganin and Victor S. Lempitsky, "Unsupervised domain adaptation by backpropagation," *ArXiv*, vol. abs/1409.7495, 2014.
- [30] Yu Huang, Chenzhuang Du, Zihui Xue, Xuanyao Chen, Hang Zhao, and Longbo Huang, "What makes multimodal learning better than single (provably)," in *Neural Information Processing Systems*, 2021.
- [31] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2009, pp. 248–255.
- [32] Arsha Nagrani, Joon Son Chung, Weidi Xie, and Andrew Zisserman, "Voxceleb: Large-scale speaker verification in the wild," *Computer Science and Language*, 2019.
- [33] Jiankang Deng, J. Guo, and Stefanos Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4685–4694, 2018.
- [34] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuyck, "ECAPA-TDNN: emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *Conference of the International Speech Communication Association (INTERSPEECH)*, 2020, pp. 3830–3834, International Symposium on Computer Architecture (ISCA).
- [35] David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5329–5333, 2018.
- [36] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [37] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean, "Distilling the knowledge in a neural network," *ArXiv*, vol. abs/1503.02531, 2015.
- [38] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N. Balasubramanian, "Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks," *IEEE Workshop on Applications of Computer Vision (WACV)*, pp. 839–847, 2017.
- [39] Mukund Sundararajan, Ankur Taly, and Qiqi Yan, "Axiomatic attribution for deep networks," in *International Conference on Machine Learning (ICML)*, 2017.
- [40] Ramprasaath R Selvaraju, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra, "Grad-cam: Why did you say that?," *arXiv preprint arXiv:1611.07450*, 2016.
- [41] Jian Zhang, Liang He, Xiao Song Guo, and Jingrong Ma, "A study on visualization of voiceprint feature," *Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 2233–2237, 2023.
- [42] Mukund Sundararajan, Ankur Taly, and Qiqi Yan, "Axiomatic attribution for deep networks," in *International conference on machine learning*, PMLR, 2017, pp. 3319–3328.