

Modality-Balanced Collaborative Distillation for Multi-Modal Domain Generalization

Xiaohan Wang^{1,2}, Zhangtao Cheng¹, Ting Zhong^{1*}, Leiting Chen^{1,2}, Fan Zhou^{1,2}

¹University of Electronic Science and Technology of China

²Key Laboratory of Intelligent Digital Media Technology of Sichuan Province

{xiaohanwang, zhangtao.cheng}@std.uestc.edu.cn, {zhongting, richardchen, fan.zhou}@uestc.edu.cn

Abstract

Weight Averaging (WA) has emerged as a powerful technique for enhancing generalization by promoting convergence to a flat loss landscape, which correlates with stronger out-of-distribution performance. However, applying WA directly to multi-modal domain generalization (MMDG) is challenging: differences in optimization speed across modalities lead WA to overfit to faster-converging ones in early stages, suppressing the contribution of slower yet complementary modalities, thereby hindering effective modality fusion and skewing the loss surface toward sharper, less generalizable minima. To address this issue, we propose MBCD, a unified collaborative distillation framework that retains WA’s flatness-inducing advantages while overcoming its shortcomings in multi-modal contexts. MBCD begins with adaptive modality dropout in the student model to curb early-stage bias toward dominant modalities. A gradient consistency constraint then aligns learning signals between uni-modal branches and the fused representation, encouraging coordinated and smoother optimization. Finally, a WA-based teacher conducts cross-modal distillation by transferring fused knowledge to each uni-modal branch, which strengthens cross-modal interactions and steer convergence toward flatter solutions. Extensive experiments on MMDG benchmarks show that MBCD consistently outperforms existing methods, achieving superior accuracy and robustness across diverse unseen domains.

Code — <https://github.com/xiaohanwang01/MBCD>

Introduction

Domain Generalization (DG) (Wang et al. 2022) aims to build robust machine learning models that can generalize reliably on unseen target domains, even when trained exclusively on data from a limited set of source domains. In practice, conventional supervised learning models often suffer severe performance degradation when deployed in real-world scenarios, such as autonomous driving (Sanchez, Deschaut, and Goulette 2023; Yang et al. 2024), medical image diagnosis (Yadav and Jadhav 2019; Yoon et al. 2024), action recognition (Gong et al. 2023), and industrial fault diagnosis (Ragab et al. 2022; Chen et al. 2025). To address these

*Corresponding Authors.

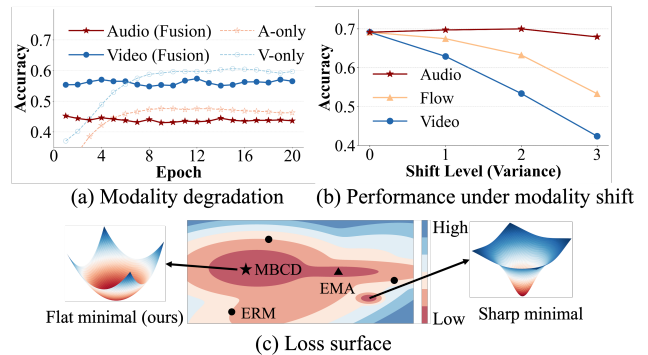


Figure 1: Illustration of EMA’s limitations in MMDG. (a) Comparison on the EPIC-Kitchens test set: the uni-modal models are trained from scratch, while the multi-modal model uses post-hoc classifier training. (b) Performance under varying modality shifts induced by Gaussian noise with different variance (shift level); performance drops sharply as the dominant modality (video) is perturbed. (c) Loss landscape visualization: EMA converges to sharp, biased minima under distributional shifts, while our MBCD yields flatter minima, improved robustness.

challenges, researchers employ a variety of techniques, including domain alignment (Li et al. 2018, 2025), data augmentation (Volpi et al. 2018; Xu et al. 2025), loss landscape flattening (Foret et al. 2020; Wu, Luo, and Wunsch II 2024) and weight averaging (WA) (Izmailov et al. 2018; Ramé et al. 2023). With the growing prevalence of multiple modalities in modern AI systems, such as vision-language pairs (Radford et al. 2021) and audio-visual (Cheng et al. 2024) streams, multi-modal domain generalization (MMDG) has become an increasingly important yet complex challenge. In MMDG, models must learn to align and fuse heterogeneous modalities in the presence of distribution shifts, which significantly amplifies the complexity of the generalization, especially as models often rely on redundancy (Tai et al. 2025).

Recent work (Cha et al. 2021) has highlighted the importance of loss landscape geometry for generalization, particularly focusing on flat minima that offer better resilience to distribution shifts. Among them, Exponential Moving Averaging (EMA) is a common technique used to smooth the loss landscape. However, EMA often converges to sharp, biased minima under distributional shifts, while our MBCD yields flatter minima, improved robustness.

age (EMA) (Li et al. 2024) stands out as a straightforward and powerful method: by averaging model parameters over training iterations, it smooths the optimization path and encourages convergence to flatter regions, while boosting out-of-distribution performance in uni-modal DG tasks.

However, directly adapting EMA to multi-modal settings introduces new challenges. The main issue stems from optimization imbalance: modalities converge at different rates due to varying signal strengths and architectural. In multi-modal fusion networks, this causes EMA to disproportionately favor fast-converging (and often dominant) modalities during early training. Consequently, the contributions of slower yet complementary modalities are suppressed, hindering effective cross-modal fusion and weakening the representational capacity of the model. As shown in Figure 1(a), the performance of the jointly trained model can even lag behind its uni-modal counterparts. This imbalance in optimization dynamics skews the averaged trajectory toward sharper, modality-biased minima (Figure 1(c)), thereby impairing generalization. The vulnerability becomes more pronounced under dominant modality shifts, resulting in significant performance degradation, as illustrated in Figure 1(b). Notably, the accuracy of the weaker modality like audio even increases under mild shifts, highlighting its limited contribution in the fused representation.

To bridge this gap, we present **Modality-Balanced Collaborative Distillation (MBCD)**, a unified framework that retains EMA’s flatness benefits while explicitly mitigating its pitfalls in multi-modal learning. MBCD incorporate adaptive modality dropout into the student model to dynamically suppresses dominant modalities during training and reduce early-stage over-reliance. It then introduce a gradient consistency constraint to align the learning dynamics between uni-modal branches and the fused representation, enabling smoother, more coordinated optimization. Finally, an EMA-based teacher performs cross-modal distillation, injecting fused knowledge into each modality-specific branch, thereby enhancing cross-modal synergy and guiding the model toward flatter and more generalizable solutions. Overall, our contributions can be summarized as follows:

- A detailed analysis of EMA’s core limitations in MMDG: optimization disparities lead to overfitting on dominant modalities, suppressing cross-modal interactions, and favoring sharper minima that harms generalization.
- MBCD, a unified framework that maintains EMA’s strengths while alleviating modality imbalance through collaborative distillation, promoting balanced optimizations and flatter generalizations.
- Comprehensive experiments on two public MMDG benchmarks demonstrate the effectiveness of our proposed MBCD across diverse modality combinations, yielding flatter minima and improved generalization.

Related Work

Domain Generalization. This task aims to train models on one or multiple source domains that can generalize to previously unseen target domains (Zhou et al. 2022). Specifically, domain alignment methods (Li et al. 2018; Qu et al.

2023) focus on learning domain-invariant representations. Moreover, data augmentation methods (Zhang et al. 2017; Qiao, Zhao, and Peng 2020; Zheng, Huai, and Zhang 2024; Cho, Hwang, and Lee 2025) aim to simulate potential domain shifts by introducing variations at the input or feature level. Another line of research focuses on flattening the loss landscape (Wang et al. 2023; Wu, Luo, and Wunsch II 2024; Deng et al. 2025), encouraging convergence to flat minima. More recently, weight averaging methods (Cha et al. 2021; Arpit et al. 2022; Rame et al. 2022; Ramé et al. 2023; Javed, Le, and Salzmann 2024)—such as stochastic weight averaging (SWA) (Izmailov et al. 2018) and exponential moving average (EMA) (Morales-Brotons, Vogels, and Hendriks 2024; Ajroldi, Orvieto, and Geiping 2025)—have emerged as effective techniques for improving DG performance.

Multi-Modal Domain Generalization. While uni-modal DG has been extensively studied, its multi-modal counterpart remains relatively underexplored. Existing efforts remain limited but provide promising directions. (Planamente et al. 2022) align audio and visual features through a relative norm alignment loss to enhance cross-domain robustness and generalizable feature learning. (Dong et al. 2023) disentangles modality-specific and shared features with tailored constraints. (Dong, Chatzi, and Fink 2024) leverages self-supervised objectives like masked cross-modal translation and multimodal jigsaw puzzles to enhance MMDG. (Fan et al. 2024) promotes consistent flat minima and cross-modal knowledge transfer to enhance modality-specific learning. (Huang et al. 2025b) proposes to mitigate modality asynchrony by learning a highly aligned and unified multi-modal representation space. However, prior approaches treat all modalities equally, resulting in weaker modalities being underoptimized and convergence to sharper minima. In contrast, our method suppresses dominant modalities, promoting wider minima and enhancing MMDG.

Multi-Modal Imbalance Learning. A range of approaches (Peng et al. 2022; Fan et al. 2023; Ma, Chen, and Deng 2025) have been proposed to balance multi-modal learning (Wang, Tran, and Feiszli 2020). For instance, (Zhang et al. 2024) adopts alternating uni-modal optimization to reduce inter-modal interference, while (Wei et al. 2024) introduces sample-level modality valuation to encourage cooperation. (Huang et al. 2025a) slows down learning from dominant modalities during critical phases to mitigate early overfitting. Other approaches enhance modality representations via auxiliary networks. (Du et al. 2023) distills knowledge from uni-modal experts into multi-modal models. Gradient-based strategies (Wei and Hu 2024; Guo et al. 2024) balance learning by aligning gradient magnitudes or directions, though the former incurs high computation and the latter may suffer from noise. In contrast, our gradient consistency strategy addresses the mismatch between uni- and multi-modal objectives without explicit gradient manipulation.

Methodology

Problem Definition

We follow the standard formulation of MMDG as described in (Dong et al. 2023). Let $\mathcal{S} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_{|S|}\}$ denote a

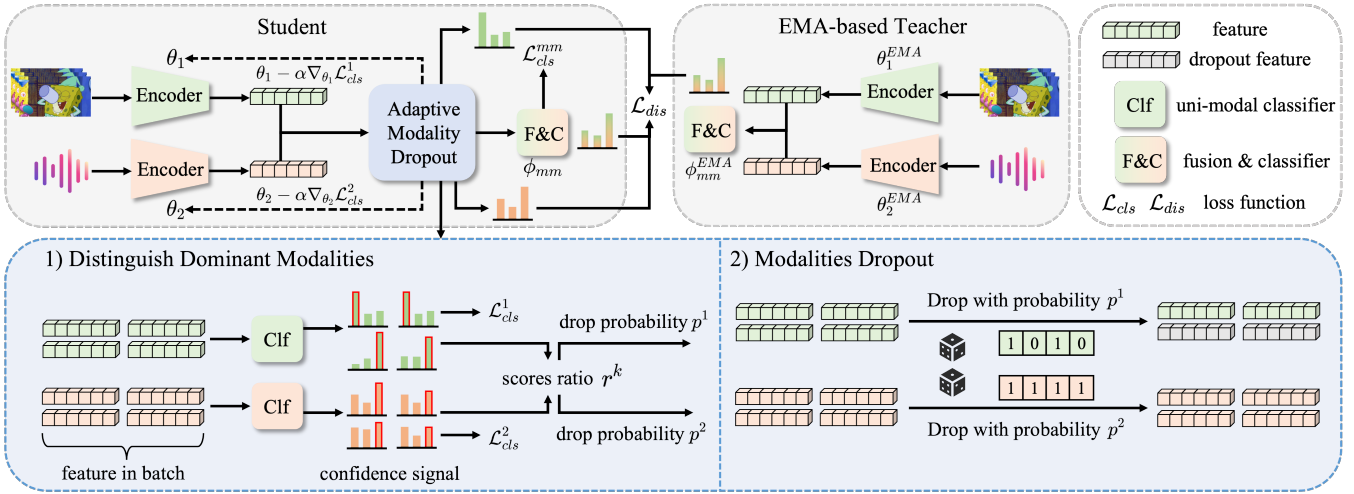


Figure 2: Overall framework of our MBCD. Our model first performs a uni-modal objective-guided inner-loop update to enhance modality-specific encoders. The updated modality representations are then fused via adaptive modality dropout to mitigate modality imbalance. An EMA-based teacher further guides both uni-modal and fused predictions, promoting stable and modality-balanced learning.

set of source domains used for training, where each domain $\mathcal{D}_j = \{(x_i^{(j)}, y_i^{(j)})\}_{i=1}^{N_j}$ contains N_j labeled instances. Each input sample $x_i^{(j)} = \{(x_i^{(j)})_k\}_{k=1}^M$ consists of M modalities, and $y_i^{(j)}$ is the corresponding ground-truth label. We assume that the joint distributions differ across domains, i.e., $P_{XY}^i \neq P_{XY}^j, i \neq j$. The objective of MMDG is to learn a predictive function f that performs well on previously unseen target domains $\mathcal{T} = \{\mathcal{D}'_1, \mathcal{D}'_2, \dots, \mathcal{D}'_{|\mathcal{T}|}\}$, without accessing any data from \mathcal{T} during training. The optimization objective is formulated as:

$$\min_f \mathbb{E}_{(x,y) \sim P_{\mathcal{T}}} [\mathcal{L}(f(x), y)], \quad (1)$$

where \mathbb{E} denotes the expectation, and $\mathcal{L}(\cdot, \cdot)$ is a task-specific loss function (e.g., cross-entropy for classification). In our framework, the parameters of the M modality-specific encoders are denoted by $\theta = \{\theta_1, \dots, \theta_M\}$, where θ_k is for the k -th modality. The prediction head corresponding to the k -th modality is parameterized by ϕ_k , while the fused multi-modal prediction head is denoted by ϕ_{mm} . For simplicity, we use ϕ to denote $\{\phi_1, \phi_2, \dots, \phi_M, \phi_{mm}\}$.

MBCD

We propose MBCD, as illustrated in Figure 2, which integrates three key components to tackle MMDG: (1) Adaptive Modality Dropout, which dynamically suppresses dominant modalities to reduce early-stage over-reliance and encourage balanced training; (2) Gradient Consistency Constraint, which harmonizes uni-modal and multi-modal objectives by aligning their optimization directions, improving training stability and fusion quality; (3) Collaborative Distillation, which establishes a bidirectional interaction between a student and an EMA-updated teacher. The student learns from the teacher via cross-modal supervision, while simultaneously refining the teacher through online updates,

enabling mutual enhancement of generalizable fusion representations. By integrating these components, MBCD consistently achieves superior performance across unseen multi-modal domains.

Adaptive Modality Dropout

Dominant modalities tend to learn faster and may suppress weaker ones during joint training, which can hinder balanced optimization. To mitigate this effect, we propose to adaptively modulate the learning pace of each modality according to its confidence signal. We quantify the confidence of each modality on a mini-batch using the following metric:

$$s^k = \sum_{i \in B} \max(\sigma(f^k(x_i^k))), \quad (2)$$

where $\sigma(\cdot)$ denotes the softmax function, and $f^k(\cdot)$ is the output of the k -th modality-specific network, including both the encoder and classifier. The score s^k reflects the overall prediction confidence of modality k on the current mini-batch B . Note that this metric does not consider prediction correctness, but rather captures how confident the model is in its outputs, which serves as an indirect signal of its current dominance in learning.

We then quantify the relative learning speed of each modality by comparing its confidence score with those of other modalities:

$$r^k = \frac{1}{M-1} \sum_{j \in \{1,2,\dots,M\}, j \neq k} \frac{s^k}{s^j}. \quad (3)$$

A value of $r^k > 1$ indicates that the k -th modality is learning faster than average and may dominate the optimization, potentially suppressing under-fitting modalities. To mitigate this imbalance, we introduce an adaptive modality dropout mechanism that selectively suppresses over-dominant modalities. Specifically, we define the dropout

mask as:

$$\mathbf{D}^k = 1 - \mathbf{M}^k \sim \text{Bernoulli}(\tanh(\max(r^k - 1, 0))), \quad (4)$$

where \mathbf{M}^k is a Bernoulli-distributed random variable with success probability $p = \tanh(\max(r^k - 1, 0))$, and \mathbf{D}^k denotes the dropout mask applied to modality k . This design probabilistically drops dominant modalities with higher r^k values, effectively slowing down their contribution to the learning process.

Gradient Consistency Constraint

Recent multi-modal learning frameworks (Wei and Hu 2024; Guo et al. 2024) introduce additional uni-modal neural components to enhance the representation capacity of each modality via auxiliary uni-modal tasks. However, naively applying such designs may overlook the gradient conflict between uni-modal and multi-modal learning objectives. To mitigate this issue, we propose a uni-modal objective-guided learning strategy:

$$\min_{\theta, \phi} \sum_{k=1}^M \mathcal{L}_{cls}(\theta_k, \phi_k; x_k) + \mathcal{L}_{cls}\left(\bigcup_{k=1}^M \theta'_k, \phi_{mm}; x\right),$$

where $\theta'_k = \theta_k - \alpha \nabla_{\theta_k} \mathcal{L}_{cls}(\theta_k, \phi_k; x_k)$. (5)

Here, $\bigcup_{k=1}^M \theta'_k$ denotes parameter aggregation, $\mathcal{L}_{cls}(\theta_k; x_k)$ is the uni-modal classification loss for the k -th modality, and $\mathcal{L}_{cls}(\bigcup_{k=1}^M \theta'_k; x)$ corresponds to the fused multi-modal classification loss (implemented via concatenation-based fusion in our framework). Note that θ'_k is obtained by performing a gradient descent step on the uni-modal loss.

To obtain a more tractable and interpretable objective, we perform a first-order Taylor expansion on the second term of Eq. (5) around θ_k . This derivation yields a new and equivalent learning objective:

$$\min_{\theta, \phi} \sum_{k=1}^M \mathcal{L}_{cls}(\theta_k, \phi_k; x_k) + \mathcal{L}_{cls}(\theta, \phi_{mm}; x) - \sum_{k=1}^M \alpha \nabla_{\theta_k} \mathcal{L}_{cls}(\theta_k, \phi_k; x_k) \cdot \nabla_{\theta_k} \mathcal{L}_{cls}(\theta, \phi; x). \quad (6)$$

This reformulated objective provides a clear interpretation of our strategy. It aims to minimize both the uni-modal and multi-modal losses while simultaneously maximizing the inner product between their respective gradients. Maximizing this inner product encourages the gradients to be aligned, thereby promoting consistency in the update directions of the uni-modal and multi-modal objectives. Essentially, this approach can be viewed as a form of gradient matching, which fosters coordination between the different learning tasks and effectively addresses the initial problem of gradient conflict.

Additionally, raw features from disparate encoders often exhibit misaligned numerical ranges and variances, which can destabilize joint optimization, suppress weaker modalities, or bias fusion toward dominant modalities. Inspired by the success of Layer Normalization (LayerNorm) (Xu et al. 2019) in stabilizing single-modality deep networks (e.g., Transformers), we add LayerNorm to each modality’s encoded features.

Collaborative Distillation

EMA is a simple yet effective technique that improves model generalization by smoothing parameter updates over the training trajectory. It has been widely used in various training paradigms to stabilize optimization, suppress noise in parameter updates, and implicitly form a temporal ensemble of models. Formally, the EMA parameters at training step t are computed as:

$$\Theta_t^{EMA} = \beta \Theta_{t-1}^{EMA} + (1 - \beta) \Theta_t, \quad (7)$$

where Θ_t denotes the model parameters at step t updated by the optimizer, Θ_t^{EMA} represents the exponentially averaged parameters, and $\beta \in [0, 1)$ is a smoothing factor that controls the decay rate of historical information. A larger β places more emphasis on past parameter states, resulting in smoother updates. In our framework, EMA is applied to both modality-specific encoder parameters θ^k and the fused prediction head parameter ϕ^{mm} . This strategy forms an implicit ensemble of the model over time, which improves robustness to domain shifts.

We then propose the collaborative distillation framework that simultaneously enhances generalization and promotes cross-modal fusion through a bidirectional interaction between the student and an EMA-based teacher model. Our distillation targets two key objectives: (1) enhancing generalization via an EMA-based teacher, and (2) promoting cross-modal interaction by distilling fused knowledge into uni-modal branches. The distillation loss is defined as:

$$\mathcal{L}_{dis} = D_{\text{KL}}(p^{EMA} || p^{mm}) + \sum_{m=1}^M D_{\text{KL}}(p^{EMA} || p^m), \quad (8)$$

where p^{EMA} is the teacher’s fused prediction, p^{mm} is the student’s multi-modal output, and p^m denotes predictions from each uni-modal branch. The KL divergence terms measure the discrepancy between student predictions and the teacher’s fused output. This formulation dynamically monitors modality-wise consistency and mitigates imbalance by penalizing divergence from the teacher.

With cross-modal distillation, our framework forms a collaborative learning loop: the teacher offers stable cross-modal supervision, while the student, through its improved predictions, continuously refines the teacher via online updates. This bidirectional interaction enables mutual enhancement of generalizable fusion representations.

Final Loss

The final loss is obtained as the weighted sum of the previously defined losses:

$$\mathcal{L} = \mathcal{L}_{cls}(\theta', \phi_{mm}; x) + \sum_{k=1}^M \mathcal{L}(\theta_k, \phi_k; x_k) + \lambda \mathcal{L}_{dis}, \quad (9)$$

where λ is the coefficient of distillation loss \mathcal{L}_{dis} .

Experiments

Experimental Setting

Dataset & Implementation Details. We conduct experiments on two benchmark datasets: EPIC-Kitchens (Damen

Method	Modality			EPIC-Kitchens				HAC			
	Video	Audio	Flow	D2,D3→D1	D1,D3→D2	D1,D2→D3	Avg.	A,C→H	H,C→A	H,A→C	Avg.
ERM	✓	✓		52.04±0.60	60.73±1.69	57.92±1.40	56.90	72.96±0.69	74.10±3.65	51.96±1.73	66.34
RNA-Net	✓	✓		52.21±1.20	60.06±2.06	56.17±1.87	56.15	73.42±1.16	73.33±1.16	50.06±1.16	65.60
SimMMDG	✓	✓		56.35±2.19	64.52±0.72	58.21±1.94	59.69	76.54±2.40	75.79±2.14	48.71±1.61	67.01
MOOSA	✓	✓		54.64±3.73	64.93±1.36	61.60±0.51	60.39	74.29±1.40	75.90±0.55	52.79±2.49	67.66
CMRF	✓	✓		56.77±0.30	65.21±0.31	61.79±0.33	61.26	79.21±0.69	79.07±0.82	55.48±1.41	71.25
MBCD	✓	✓		58.06±0.50	68.10±0.35	63.31±0.21	63.16	80.08±0.44	78.81±0.72	53.09±1.84	70.66
ERM	✓		✓	56.88±0.26	65.74±1.42	58.02±1.32	60.21	74.89±1.12	72.33±1.51	43.35±6.35	63.52
RNA-Net	✓		✓	59.29±0.54	66.37±1.30	57.85±2.02	61.17	77.19±4.33	74.58±4.33	43.11±4.33	64.96
SimMMDG	✓		✓	57.99±2.29	67.59±0.42	55.85±1.88	60.48	77.48±1.47	73.99±1.76	51.78±6.05	67.75
MOOSA	✓		✓	59.16±1.36	66.44±0.98	61.43±1.14	62.34	75.25±2.25	75.68±0.51	51.75±2.66	67.56
CMRF	✓		✓	63.41±0.35	69.48±0.80	61.52±0.53	64.80	80.32±0.68	77.08±1.52	51.87±3.34	69.76
MBCD	✓		✓	64.16±0.33	72.19±0.59	63.05±0.53	66.47	80.37±0.74	78.07±0.75	51.93±1.31	70.12
ERM		✓	✓	53.56±1.53	59.14±1.54	56.71±0.94	56.47	54.17±2.28	60.56±3.26	42.19±2.70	52.31
RNA-Net		✓	✓	51.50±1.26	60.35±1.49	56.86±1.38	56.24	54.91±2.13	59.82±2.13	43.23±2.13	52.65
SimMMDG		✓	✓	57.30±0.89	65.39±1.01	56.19±1.31	59.63	55.61±1.70	63.43±1.57	45.31±0.20	54.78
MOOSA		✓	✓	53.72±1.41	67.20±1.54	60.78±2.03	60.57	56.79±1.28	63.58±0.48	42.59±1.76	54.32
CMRF		✓	✓	59.13±0.87	65.08±0.99	61.30±0.28	61.84	60.49±0.46	64.75±0.32	47.98±1.28	57.74
MBCD		✓	✓	59.25±0.45	68.05±0.59	61.95±0.27	63.08	63.76±0.91	65.64±0.69	49.05±0.19	59.48
ERM	✓	✓	✓	49.46±0.83	53.70±1.78	48.50±3.91	50.55	74.67±3.33	68.54±0.41	45.40±3.58	62.87
RNA-Net	✓	✓	✓	47.12±0.29	55.80±0.69	50.63±3.39	51.18	76.02±3.19	73.11±3.19	46.42±3.19	65.18
SimMMDG	✓	✓	✓	60.31±1.96	68.75±1.09	60.97±1.01	63.34	76.38±0.65	74.36±1.70	51.59±2.50	67.44
MOOSA	✓	✓	✓	60.00±0.19	67.38±1.07	64.24±5.15	63.87	75.34±3.11	74.54±1.85	48.44±4.82	66.11
CMRF	✓	✓	✓	61.86±0.74	69.58±0.39	64.97±0.14	65.47	78.25±0.43	78.07±0.50	55.61±1.23	70.64
MBCD	✓	✓	✓	63.63±0.39	73.18±0.59	65.92±0.44	67.58	81.28±1.00	78.11±0.32	55.48±1.60	71.62

Table 1: Multi-modal **multi-source** DG accuracy with different modalities on EPIC-Kitchens and HAC datasets. The results are averaged over 3 random seeds, with standard deviation displayed as well. The best is in **bold**.

Method	S:	EPIC-Kitchens						Avg.	HAC						Avg.
		D1		D2		D3			H		A		C		
		D2	D3	D1	D3	D1	D2		A	C	H	C	H	A	
ERM	T:	43.42±3.08	41.99±3.76	42.18±1.74	42.24±2.00	42.28±1.80	51.98±1.35	44.02	57.36±2.36	41.21±1.95	69.38±2.44	42.89±4.93	56.77±7.48	57.28±2.97	54.15
RNA-Net		43.72±3.54	41.51±2.79	41.41±0.82	44.56±3.28	42.40±3.67	53.46±0.55	44.51	57.91±1.58	34.96±5.02	71.14±0.92	44.67±1.46	63.25±3.72	63.13±3.72	55.84
SimMMDG		54.09±1.07	49.21±1.20	53.56±3.07	57.27±2.02	54.42±1.69	64.74±1.77	55.55	67.03±2.23	42.74±1.99	69.72±2.97	46.26±5.83	66.71±3.29	65.82±4.26	59.71
MOOSA		54.99±0.76	49.76±1.57	52.83±1.35	55.15±2.59	55.49±2.23	65.27±1.06	55.58	62.21±2.40	43.41±2.48	73.78±0.79	51.23±3.56	63.90±2.75	68.18±2.37	60.45
CMRF		59.35±0.24	54.38±0.32	57.62±0.37	60.96±0.30	57.01±0.32	67.72±0.61	59.51	67.84±0.96	46.26±1.82	74.04±0.44	49.36±0.40	65.66±1.58	65.89±1.00	61.51
MBCD		59.50±0.60	54.89±0.93	58.11±0.84	61.73±0.45	57.18±0.65	70.77±0.44	60.36	69.24±0.75	46.81±0.68	76.16±0.07	49.36±0.26	73.37±1.12	71.93±3.29	64.48

Table 2: Multi-modal **single-source** DG accuracy with video, audio and flow modalities on EPIC-Kitchens and HAC datasets. S denotes source domain while T denotes target domain. The results are averaged over 3 random seeds, with standard deviation displayed as well. The best is in **bold**.

et al. 2020) and Human-Animal-Cartoon (HAC) (Dong et al. 2023), both containing video, optical flow, and audio modalities. For EPIC-Kitchens, we adopt domains D1, D2, and D3, while for HAC, the domains are human (H), animal (A), and cartoon (C). Our experimental protocol follows the setup in (Dong et al. 2023).

Baselines. We compare our method, MBCD, against five baselines: ERM (a naive multi-modal approach based on feature concatenation) and four state-of-the-art MMDG methods—RNA-Net (Planamente et al. 2022), SimMMDG (Dong et al. 2023), MOOSA (Dong, Chatzi, and Fink 2024), and CMRF (Fan et al. 2024). For each method, we evaluate the model achieving the best validation (in-domain) performance on the test set (out-of-domain). All results are reported as Top-1 accuracy, averaged over 3 random seeds.

Main Results

Multi-Modal Multi-Source DG. Table 1 reports the performance of our method MBCD and several baselines on the EPIC-Kitchens and HAC datasets under the multi-modal, multi-source domain generalization setting, where models are trained on multiple source domains and evaluated on a distinct target domain. We assess generalization by testing all pairwise modality combinations as well as the full tri-modal setup. As shown in Table 1, MBCD consistently surpasses baseline methods across nearly all configurations, achieving up to a 1.9% improvement in average accuracy. In particular, utilizing all three modalities leads to the best performance, outperforming any two-modality combination. This demonstrates that MBCD effectively mitigates modality imbalance and fully exploits complementary information across modalities, whereas baseline approaches struggle to benefit from the inclusion of additional modalities.

Method	EPIC-Kitchens			Avg.
	D2,D3→D1	D1,D3→D2	D1,D2→D3	
SimMMDG	82.19±0.39	78.80±0.26	78.54±1.04	79.84
MOOSA	82.68±0.53	78.64±0.35	77.10±0.42	79.47
CMRF	83.22±0.10	77.95±0.24	78.28±0.38	79.82
MBCD	84.76±0.27	80.20±0.35	79.83±0.25	81.60

Table 3: Multi-modal **in-domain** Accuracy on EPIC-Kitchens.

AMD	GCC	EMA	DL	D2,D3→D1	D1,D3→D2	D1,D2→D3	Avg.
				53.56±1.53	59.14±1.54	56.71±0.94	56.47
		✓		57.52±0.02	64.90±0.50	60.41±0.35	60.94
✓		✓		58.63±0.16	65.57±0.70	60.58±0.67	61.59
	✓	✓		58.46±1.08	67.27±0.61	61.10±0.32	62.28
		✓	✓	58.81±0.14	66.52±1.23	62.22±0.41	62.52
✓	✓	✓	✓	59.25±0.45	68.05±0.59	61.95±0.27	63.08

Table 4: Ablation on each design on EPIC-Kitchens with flow and audio data. AMD: adaptive modality dropout, GCC: gradient consistency constraint, EMA: exponential moving average, DL: distillation loss.

Multi-Modal Single-Source DG. We further evaluate the generalization ability of our MBCD framework in a more challenging single-source DG setting, where models are trained using data from only one source domain and tested on multiple unseen target domains. This setting imposes stronger constraints on model robustness. The results using all three modalities are summarized in Table 2. Even under this limited training scenario, our method consistently outperforms all baselines in terms of average Top-1 accuracy across target domains, demonstrating its strong ability to extract transferable and complementary multi-modal representations from a single domain.

Multi-Modal In-Domain Performance. In addition to strong generalization ability in MMDG settings, MBCD also achieves the best performance in the in-domain scenario. As shown in Table 3, when evaluated within the EPIC-Kitchens dataset under the standard in-domain setting—where training and testing domains are drawn from the same distribution—MBCD consistently outperforms existing state-of-the-art baselines across all domain splits. This superior in-domain performance demonstrates the effectiveness of our model not only in handling domain shifts but also in fully exploiting multi-modal cues under standard conditions.

Ablation Studies

Ablation on Each Module. We conducted ablation studies on EPIC-Kitchens using flow and audio modalities to assess the contribution of each component in our framework. As shown in Table 4, introducing EMA alone leads to a substantial improvement over the baseline (from 56.47% to 60.94%), highlighting the effectiveness of temporal knowledge transfer. Adding adaptive modality dropout (AMD) or the gradient consistency constraint (GCC) individually further boosts performance, suggesting that both dynamic modality regulation and uni-modal enhancement play important roles in promoting robust multi-modal learning. In-

Method	A,C→H	H,C→A	H,A→C	Avg.
w/o dropout	62.82±1.77	64.24±1.26	45.22±0.67	57.43
fix dropout	64.34±0.65	64.42±0.59	47.89±0.69	58.88
adaptive dropout	63.76±0.91	65.64±0.69	49.05±0.19	59.48

Table 5: Comparison with different dropout strategy on HAC with flow and audio data.




Figure 3: Comparison of flatness for different methods on EPIC-Kitchens and HAC.

corporating distillation loss (DL) on top of EMA also yields notable gains. Finally, combining all modules achieves the best average performance, demonstrating their complementary benefits and the overall effectiveness of our proposed design.

Ablation on Modality Drop. We compare our adaptive modality dropout with two variants: (1) w/o dropout, where no modality is dropped during fusion; (2) fix dropout, where modalities are dropped with a fixed probability of 0.5. As shown in Table 5, all dropout strategies improve over the baseline without dropout, indicating the benefit of suppressing modality dominance. Among them, our adaptive dropout achieves the best overall performance, demonstrating its effectiveness in dynamically balancing modality contributions based on training feedback.

Further Analysis

Flatness Analysis. To evaluate the flatness of the loss landscape, we follow prior work (Cha et al. 2021) and measure the change in loss values under perturbations around the converged parameters. The core assumption is that flatter minima are indicative of better generalization. Specifically, we add random directional noise by a perturbation radius and measure the corresponding loss increase. As shown in Figure 3, the x-axis represents the scaled perturbation radius (magnified 10×), and the y-axis shows the corresponding loss increase. Our results reveal that MBCD consistently converges to flatter minima than all baselines across both EPIC-Kitchens (D1, D3 → D2) and HAC (A, C → H) datasets. Furthermore, other WA methods like EMA and CMRF also yield flatter solutions compared to non-averaging methods, underscoring the general benefit of WA for discovering flat minima.

Cross-Modal Interaction Analysis. To demonstrate the effectiveness of MBCD in enhancing cross-modal interaction, we compare modality-wise and fused accuracies across different methods. As shown in Figure 4, MBCD consistently




Figure 4: Comparison of modality-wise and fused accuracies on EPIC-Kitchens and HAC.




Figure 5: Validation and test accuracy curves on EPIC-Kitchens (D2,D3→D1) across all modalities.

improves the performance of all individual modalities and achieves higher overall accuracy on both EPIC-Kitchens (D1, D3 → D2) and HAC (A, C → H). Compared to MBCD, EMA achieves slightly lower performance, while ERM performs the worst—primarily due to its imbalanced optimization. Under standard ERM training, multi-modal fusion often suffers from imbalanced learning, leading to the underutilization of weaker modalities like audio and causing the fused model to underperform. In contrast, MBCD promotes balanced optimization across modalities. This not only enhances the representational quality of each modality but also strengthens their cross-modal complementarity, ultimately leading to consistently superior fusion performance.

Training Stability Analysis. We examine the robustness of out-of-domain performance for MBCD when model selection is guided by an in-domain validation set. As shown in Figure d5, MBCD exhibits remarkable training stability, a critical attribute for reliable deployment in real-world scenarios. This stability is evident from two perspectives. The first is the smoothness of both the validation and test accuracy curves across all training epochs, which points to a highly stable optimization process with low variance and minimal fluctuations. The second is the tight correlation between validation and test accuracy curves, which demonstrates consistent generalization and a robust defense against overfitting, even under significant distributional shifts. This strong alignment is a testament to MBCD’s ability to generalize reliably under domain shifts. In contrast, the ERM baseline exhibits significantly unstable performance, highlighting MBCD’s clear superiority in both convergence dynamics and generalization ability.

Visualization of Embeddings. To investigate the qual-

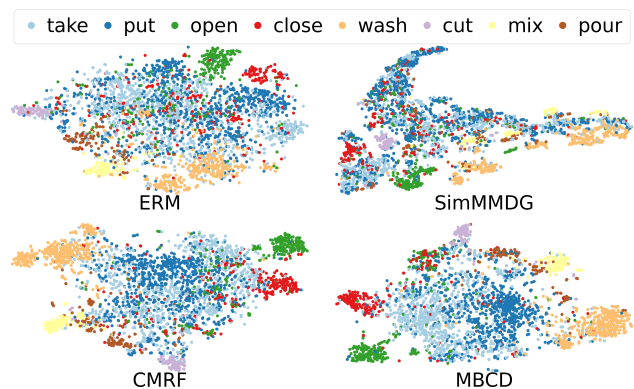


Figure 6: T-SNE visualization of concatenated multi-modal embeddings on EPIC-Kitchens.

ity of the learned representations and gain insights into our model’s mechanism under domain shift, we visualize the concatenated multi-modal embeddings of the out-of-domain testing set using t-SNE (Maaten and Hinton 2008). As depicted in Figure 6, we compare the distributions of MBCD with those of ERM, SimMMDG, and CMRF. The results demonstrate that while the baseline methods produce highly intertwined clusters with significant class overlap, our model, MBCD, successfully learns distinct and well-separated action representations. Specifically, we observe that the ERM and SimMMDG baselines struggle to differentiate between action types, leading to a blurry and mixed distribution of embeddings. While CMRF shows some improvement, it still exhibits a noticeable confusion between semantically similar actions, such as ‘take’ and ‘put’. In contrast, MBCD successfully disentangles these action patterns, forming compact and isolated clusters for each class. This visual evidence confirms that MBCD learns a more structured and semantically meaningful embedding space, which is crucial for its robust generalization to unseen domains.

Conclusion

In this work, we analyzed the core limitation of applying EMA to MMDG and proposed MBCD, a unified framework that preserved the flatness-inducing properties of WA while overcoming its limitations in multi-modal settings. By leveraging collaborative distillation, MBCD promoted balanced optimization across modalities and guided the model toward flatter, more generalizable solutions. Experiments on EPIC-Kitchens and HAC benchmarks showed that MBCD consistently outperformed state-of-the-art methods across diverse modality settings. In future work, we planned to explore the theoretical foundations of WA strategies in multi-modal learning, aiming to better understand their role in MMDG and to provide deeper insights and stronger foundations for advancing this direction.

Acknowledgments

This work was supported by National Natural Science Foundation of China (Grant No. 62572097, No. 62176043, and

References

- Ajroldi, N.; Orvieto, A.; and Geiping, J. 2025. When, Where and Why to Average Weights? *arXiv preprint arXiv:2502.06761*.
- Arpit, D.; Wang, H.; Zhou, Y.; and Xiong, C. 2022. Ensemble of averages: Improving model selection and boosting performance in domain generalization. *Advances in Neural Information Processing Systems*, 35: 8265–8277.
- Cha, J.; Chun, S.; Lee, K.; Cho, H.-C.; Park, S.; Lee, Y.; and Park, S. 2021. Swad: Domain generalization by seeking flat minima. *Advances in Neural Information Processing Systems*, 34: 22405–22418.
- Chen, Y.; Zhang, D.; Yan, R.; and Xie, M. 2025. Applications of domain generalization to machine fault diagnosis: A survey. *IEEE/CAA Journal of Automatica Sinica*.
- Cheng, Z.; Zhang, J.; Xu, X.; Trajcevski, G.; Zhong, T.; and Zhou, F. 2024. Retrieval-augmented hypergraph for multimodal social media popularity prediction. In *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*, 445–455.
- Cho, D. K.; Hwang, I.; and Lee, S. 2025. PEER pressure: Model-to-Model Regularization for Single Source Domain Generalization. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 15360–15370.
- Damen, D.; Doughty, H.; Farinella, G. M.; Fidler, S.; Furnari, A.; Kazakos, E.; Moltisanti, D.; Munro, J.; Perrett, T.; Price, W.; et al. 2020. The epic-kitchens dataset: Collection, challenges and baselines. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11): 4125–4141.
- Deng, J.; Pang, J.; Zhang, B.; and Guo, G. 2025. Asymptotic Unbiased Sample Sampling to Speed Up Sharpness-Aware Minimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 16208–16216.
- Dong, H.; Chatzi, E.; and Fink, O. 2024. Towards multimodal open-set domain generalization and adaptation through self-supervision. In *European Conference on Computer Vision*, 270–287. Springer.
- Dong, H.; Nejjar, I.; Sun, H.; Chatzi, E.; and Fink, O. 2023. SimMMDG: A simple and effective framework for multimodal domain generalization. *Advances in Neural Information Processing Systems*, 36: 78674–78695.
- Du, C.; Teng, J.; Li, T.; Liu, Y.; Yuan, T.; Wang, Y.; Yuan, Y.; and Zhao, H. 2023. On uni-modal feature learning in supervised multi-modal learning. In *International Conference on Machine Learning*, 8632–8656. PMLR.
- Fan, Y.; Xu, W.; Wang, H.; and Guo, S. 2024. Cross-modal representation flattening for multi-modal domain generalization. *Advances in Neural Information Processing Systems*, 37: 66773–66795.
- Fan, Y.; Xu, W.; Wang, H.; Wang, J.; and Guo, S. 2023. Pmr: Prototypical modal rebalance for multimodal learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20029–20038.
- Foret, P.; Kleiner, A.; Mobahi, H.; and Neyshabur, B. 2020. Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412*.
- Gong, X.; Mohan, S.; Dhingra, N.; Bazin, J.-C.; Li, Y.; Wang, Z.; and Ranjan, R. 2023. Mmg-ego4d: Multimodal generalization in egocentric action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6481–6491.
- Guo, Z.; Jin, T.; Chen, J.; and Zhao, Z. 2024. Classifier-guided gradient modulation for enhanced multimodal learning. *Advances in Neural Information Processing Systems*, 37: 133328–133344.
- Huang, C.; Wei, Y.; Yang, Z.; and Hu, D. 2025a. Adaptive unimodal regulation for balanced multimodal information acquisition. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 25854–25863.
- Huang, H.; Xia, Y.; Zhou, S.; Wang, H.; Wang, S.; and Zhao, Z. 2025b. Bridging domain generalization to multimodal domain generalization via unified representations. *arXiv preprint arXiv:2507.03304*.
- Izmailov, P.; Podoprikin, D.; Garipov, T.; Vetrov, D.; and Wilson, A. G. 2018. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*.
- Javed, S.; Le, H.; and Salzmann, M. 2024. QT-DoG: Quantization-aware Training for Domain Generalization. *arXiv preprint arXiv:2410.06020*.
- Li, A.; Zhuang, L.; Long, X.; Yao, M.; and Wang, S. 2025. Seeking Consistent Flat Minima for Better Domain Generalization via Refining Loss Landscapes. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 15349–15359.
- Li, H.; Pan, S. J.; Wang, S.; and Kot, A. C. 2018. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5400–5409.
- Li, S.; Liu, Z.; Tian, J.; Wang, G.; Wang, Z.; Jin, W.; Wu, D.; Tan, C.; Lin, T.; Liu, Y.; et al. 2024. Switch ema: A free lunch for better flatness and sharpness. *arXiv preprint arXiv:2402.09240*.
- Ma, X.; Chen, H.; and Deng, Y. 2025. Improving Multimodal Learning Balance and Sufficiency through Data Remixing. *arXiv preprint arXiv:2506.11550*.
- Maaten, L. v. d.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(Nov): 2579–2605.
- Morales-Brotons, D.; Vogels, T.; and Hendriks, H. 2024. Exponential moving average of weights in deep learning: Dynamics and benefits. *arXiv preprint arXiv:2411.18704*.
- Peng, X.; Wei, Y.; Deng, A.; Wang, D.; and Hu, D. 2022. Balanced multimodal learning via on-the-fly gradient modulation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8238–8247.
- Planamente, M.; Plizzari, C.; Alberti, E.; and Caputo, B. 2022. Domain generalization through audio-visual relative

- norm alignment in first person action recognition. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 1807–1818.
- Qiao, F.; Zhao, L.; and Peng, X. 2020. Learning to learn single domain generalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12556–12565.
- Qu, S.; Pan, Y.; Chen, G.; Yao, T.; Jiang, C.; and Mei, T. 2023. Modality-agnostic debiasing for single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24142–24151.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Ragab, M.; Chen, Z.; Zhang, W.; Eldele, E.; Wu, M.; Kwoh, C.-K.; and Li, X. 2022. Conditional contrastive domain generalization for fault diagnosis. *IEEE Transactions on Instrumentation and Measurement*, 71: 1–12.
- Ramé, A.; Ahuja, K.; Zhang, J.; Cord, M.; Bottou, L.; and Lopez-Paz, D. 2023. Model ratatouille: Recycling diverse models for out-of-distribution generalization. In *International Conference on Machine Learning*, 28656–28679. PMLR.
- Rame, A.; Kirchmeyer, M.; Rahier, T.; Rakotomamonjy, A.; Gallinari, P.; and Cord, M. 2022. Diverse weight averaging for out-of-distribution generalization. *Advances in Neural Information Processing Systems*, 35: 10821–10836.
- Sanchez, J.; Deschaut, J.-E.; and Goulette, F. 2023. Domain generalization of 3d semantic segmentation in autonomous driving. In *Proceedings of the IEEE/CVF international conference on computer vision*, 18077–18087.
- Tai, W.; Zhong, T.; Trajcevski, G.; and Zhou, F. 2025. Redundancy Undermines the Trustworthiness of Self-Interpretable GNNs. In *Forty-second International Conference on Machine Learning*.
- Volpi, R.; Namkoong, H.; Sener, O.; Duchi, J. C.; Murino, V.; and Savarese, S. 2018. Generalizing to unseen domains via adversarial data augmentation. *Advances in neural information processing systems*, 31.
- Wang, J.; Lan, C.; Liu, C.; Ouyang, Y.; Qin, T.; Lu, W.; Chen, Y.; Zeng, W.; and Yu, P. S. 2022. Generalizing to unseen domains: A survey on domain generalization. *IEEE transactions on knowledge and data engineering*, 35(8): 8052–8072.
- Wang, P.; Zhang, Z.; Lei, Z.; and Zhang, L. 2023. Sharpness-aware gradient matching for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3769–3778.
- Wang, W.; Tran, D.; and Feiszli, M. 2020. What makes training multi-modal classification networks hard? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12695–12705.
- Wei, Y.; Feng, R.; Wang, Z.; and Hu, D. 2024. Enhancing multimodal cooperation via sample-level modality valuation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 27338–27347.
- Wei, Y.; and Hu, D. 2024. Mmpareto: boosting multimodal learning with innocent unimodal assistance. *arXiv preprint arXiv:2405.17730*.
- Wu, T.; Luo, T.; and Wunsch II, D. C. 2024. Cr-sam: Curvature regularized sharpness-aware minimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 6144–6152.
- Xu, J.; Sun, X.; Zhang, Z.; Zhao, G.; and Lin, J. 2019. Understanding and improving layer normalization. *Advances in neural information processing systems*, 32.
- Xu, Z.; Cheng, D.; Jiang, X.; Wang, N.; Li, D.; and Gao, X. 2025. Adversarial Domain Prompt Tuning and Generation for Single Domain Generalization. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 18584–18595.
- Yadav, S. S.; and Jadhav, S. M. 2019. Deep convolutional neural network based medical image classification for disease diagnosis. *Journal of Big data*, 6(1): 1–18.
- Yang, J.; Gao, S.; Qiu, Y.; Chen, L.; Li, T.; Dai, B.; Chitta, K.; Wu, P.; Zeng, J.; Luo, P.; et al. 2024. Generalized predictive model for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14662–14672.
- Yoon, J. S.; Oh, K.; Shin, Y.; Mazurowski, M. A.; and Suk, H.-I. 2024. Domain generalization for medical image analysis: A review. *Proceedings of the IEEE*.
- Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2017. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*.
- Zhang, X.; Yoon, J.; Bansal, M.; and Yao, H. 2024. Multimodal representation learning by alternating unimodal adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 27456–27466.
- Zheng, G.; Huai, M.; and Zhang, A. 2024. AdvST: Revisiting data augmentations for single domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 21832–21840.
- Zhou, K.; Liu, Z.; Qiao, Y.; Xiang, T.; and Loy, C. C. 2022. Domain generalization: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(4): 4396–4415.