

Towards Multimodal Domain Generalization with Few Labels

Hongzhao Li¹ Hao Dong² Hualei Wan¹ Shupan Li^{1*} Mingliang Xu^{1*}
 Muhammad Haris Khan³

¹Zhengzhou University ²ETH Zürich ³MBZUAI
 lihongzhao@gs.zzu.edu.cn

Abstract

Multimodal models ideally should generalize to unseen domains while remaining data-efficient to reduce annotation costs. To this end, we introduce and study a new problem, Semi-Supervised Multimodal Domain Generalization (SSMDG), which aims to learn robust multimodal models from multi-source data with few labeled samples. We observe that existing approaches fail to address this setting effectively: multimodal domain generalization methods cannot exploit unlabeled data, semi-supervised multimodal learning methods ignore domain shifts, and semi-supervised domain generalization methods are confined to single-modality inputs. To overcome these limitations, we propose a unified framework featuring three key components: Consensus-Driven Consistency Regularization, which obtains reliable pseudo-labels through confident fused-unimodal consensus; Disagreement-Aware Regularization, which effectively utilizes ambiguous non-consensus samples; and Cross-Modal Prototype Alignment, which enforces domain- and modality-invariant representations while promoting robustness under missing modalities via cross-modal translation. We further establish the first SSMDG benchmarks, on which our method consistently outperforms strong baselines in both standard and missing-modality scenarios. Our benchmarks and code are available at <https://github.com/lihongzhao99/SSMDG>.

1. Introduction


In real-world applications, multimodal models have become increasingly prevalent [36]. However, when deployed in new environments, these models often face domain shifts, where the distribution of test data differs substantially from that of the training data [10, 32, 39, 40]. For instance, an action recognition model trained on high-quality studio recordings may perform poorly on outdoor, user-generated

videos due to differences in lighting, background, and audio noise. This performance degradation underscores the need for models that can generalize to unseen domains. To address this challenge, multimodal domain generalization [8, 9, 23] has emerged, aiming to learn from diverse multimodal source domains to achieve robust cross-domain generalization. Beyond generalization, an ideal learning system should also be data-efficient, capable of learning effectively from a few labeled data to mitigate the high annotation costs of multimodal datasets. This goal aligns closely with semi-supervised learning [4, 14, 20, 28, 29, 31, 33, 34], which seeks to exploit abundant unlabeled data alongside a small labeled subset during training.


In this work, we introduce **Semi-Supervised Multimodal Domain Generalization (SSMDG)**, a novel problem setting that unifies three key aspects of real-world learning: domain generalization, data efficiency, and multimodal learning. SSMDG aims to train models that leverage multimodal data with *limited annotations* from multiple source domains to generalize effectively to unseen target domains. SSMDG lies at a largely unexplored intersection. As illustrated in Fig. 1 (a), existing paradigms such as Semi-Supervised Multimodal Learning (SSML) [1, 11, 30], Multimodal Domain Generalization (MMDG) [8, 23, 27], and Semi-Supervised Domain Generalization (SSDG) [16, 22, 41], each address only partial aspects of this challenge. Consequently, these methods perform poorly under the SSMDG setting (Fig. 1 (b)). Their limitations stem from fundamental design gaps: SSML models leverage unlabeled multimodal data but overlook domain shifts; MMDG methods focus on generalization but cannot effectively utilize unlabeled data; SSDG approaches handle domain shifts with limited supervision but are restricted to single-modality inputs, thereby failing to capture cross-modal interactions. These deficiencies highlight the need for a unified framework that can jointly exploit unlabeled multimodal data, learn domain-invariant representations, and operate effectively under sparse supervision.

To address SSMDG, we design a unified framework guided by its two core challenges: (1) identifying reliable

*Corresponding authors: Shupan Li, Mingliang Xu



(a) Illustration of four related problems.



(b) Empirical results in the SSMDG setting.

Figure 1. (a) Illustration of four related learning settings: Semi-Supervised Multimodal Learning (SSML), Multimodal Domain Generalization (MMDG), Semi-Supervised Domain Generalization (SSDG), and our proposed Semi-Supervised Multimodal Domain Generalization (SSMDG). SSML overlooks domain shifts, MMDG cannot leverage unlabeled data, and SSDG is restricted to single-modality inputs. (b) Performance comparison in the SSMDG setting, highlighting the limitations of existing paradigms and the effectiveness of our proposed method.

pseudo-labels from unlabeled data despite low confidence and inter-modality disagreements, and (2) learning representations that are both modality- and domain-invariant under limited supervision. This combination forms a unique bottleneck distinct from SSML, SSDG, and MMDG. To tackle the first challenge, we introduce a two-part regularization strategy. Consensus-Driven Consistency Regularization enforces consistency only when the fused multimodal prediction and at least one unimodal prediction are both confident and in agreement, ensuring reliable pseudo-labels. Disagreement-Aware Regularization extends this to ambiguous samples with diverging modalities, employing a robust generalized cross-entropy loss to stabilize training while mitigating noise. For the second challenge, we propose Cross-Modal Prototype Alignment to jointly enforce modality and domain invariance. This module leverages class prototypes as shared semantic anchors, aligning multimodal and multi-domain features to construct a stable semantic structure crucial for generalization. It further integrates cross-modal feature translation, enhancing robustness to missing modalities. Our main contributions are summarized as follows:

- **New problem.** We introduce Semi-Supervised Multimodal Domain Generalization, unifying the previously disjoint challenges of generalization and data efficiency in real-world multimodal learning.
- **Comprehensive benchmark.** We establish the first SSMDG benchmark, covering diverse scenarios for systematic evaluation.
- **Novel insights.** We analyze the limitations of existing paradigms and identify the unique challenges specific to the SSMDG setting.
- **Effective framework.** We propose a unified framework that effectively addresses the key challenges of SSMDG, enabling robust multimodal learning under domain shifts with few labeled samples.

2. Related Work

Semi-Supervised Multimodal Learning aims to learn from multimodal data with limited annotations [1, 3, 5, 11, 18, 25, 30]. These methods typically exploit cross-modal consistency and complementarity to improve representation learning. Common approaches employ cross-modal contrastive learning or consistency regularization [1] to enforce similar representations or predictions across modalities of the same sample. Other works adopt information-theoretic formulations [3, 30] or generative models to enhance multimodal fusion. While effective, these approaches assume a single-domain setting where training and testing data share the same distribution, thus lack mechanisms to handle domain shifts central to generalization.

Multimodal Domain Generalization instead aims to generalize multimodal models to unseen domains using labeled data from multiple source domains [8, 9, 13, 17, 19, 23, 24, 27, 37]. These methods focus on learning representations that are invariant to both modality and domain. Typical strategies include extending unimodal domain generalization techniques to multimodal inputs [19], mapping heterogeneous features into a unified latent space [13], or developing alignment schemes to mitigate domain shifts [9, 27]. However, MMDG methods operate under a fully supervised paradigm, assuming all source samples are labeled, and therefore cannot leverage the abundant unlabeled data often available in practice.

Semi-Supervised Domain Generalization extends classic semi-supervised learning frameworks such as FixMatch [29] to the multi-domain setting, learning from limited labeled and abundant unlabeled data [15, 16, 21, 22]. To enhance domain-invariant feature learning, these methods develop robust pseudo-labeling strategies [22] or introduce domain-aware mechanisms such as stochastic style matching [41] and domain-guided weight modulation [16]. Although effective under sparse supervision, SSDG methods

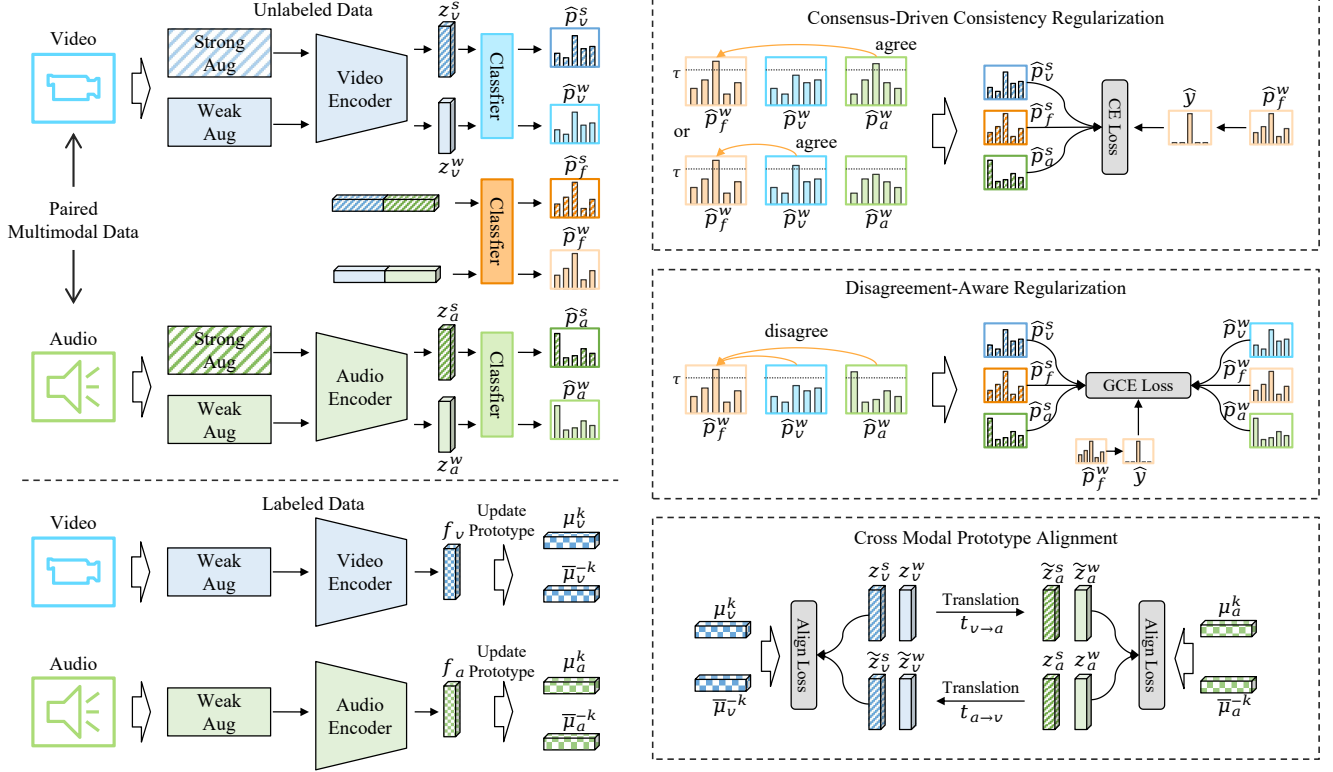


Figure 2. Overview of the proposed framework for SSMDG. The framework jointly leverages labeled and unlabeled multimodal data from multiple source domains through three key components: Consensus-Driven Consistency Regularization (CDCR, §4.1), Disagreement-Aware Regularization (DAR, §4.2), and Cross Modal Prototype Alignment (CMPA, §4.3), enabling robust generalization to unseen target domains.

are inherently designed for single-modality inputs, preventing them from exploiting the consensus and complementarity across multiple modalities.

3. Problem Definition

We study the problem of *Semi-Supervised Multimodal Domain Generalization*. Let $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_M$ denote the multimodal input space composed of M distinct modalities, and let \mathcal{Y} represent the shared label space. Each multimodal sample is denoted as $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_M)$, with each modality-specific input $\mathbf{x}_m \in \mathcal{X}_m$. A domain is formally defined by a joint distribution $P(\mathbf{X}, Y)$ over $\mathcal{X} \times \mathcal{Y}$, with corresponding marginals $P(\mathbf{X})$ and $P(Y)$.

We assume access to K distinct but related source domains, $\mathcal{S} = \{\mathcal{S}_k\}_{k=1}^K$, each associated with a joint distribution $P^{(k)}(\mathbf{X}, Y)$. We consider domain shifts that arise solely from variations in the input distribution, i.e., $P^{(k)}(\mathbf{X}) \neq P^{(k')}(\mathbf{X})$ for $k \neq k'$, while the label space \mathcal{Y} and its marginal distribution $P(Y)$ remain consistent across all domains. This setup implies that the underlying input-label relationship, captured by the conditional distribution $P(Y|\mathbf{X})$, is invariant and shared among all domains.

In the semi-supervised setting, the data for each source domain \mathcal{S}_k consists of: A small labeled subset $\mathcal{D}_k^l = \{(\mathbf{X}_i^{(k)}, y_i^{(k)})\}_{i=1}^{n_k^l}$, where $(\mathbf{X}_i^{(k)}, y_i^{(k)})$ are drawn i.i.d. from $P^{(k)}(\mathbf{X}, Y)$. A much larger unlabeled subset $\mathcal{D}_k^u = \{U_j^{(k)}\}_{j=1}^{n_k^u}$, where $U_j^{(k)}$ are drawn from the marginal distribution $P^{(k)}(\mathbf{X})$. The total size of the unlabeled data significantly outweighs the labeled data, i.e., $n_k^l \ll n_k^u$. The goal of SSMDG is to leverage all labeled and unlabeled data from the K source domains to learn a single predictive model $f_\theta : \mathcal{X}_1 \times \dots \times \mathcal{X}_M \rightarrow \mathcal{Y}$, parameterized by θ , that generalizes effectively to an unseen target domain \mathcal{T} . The target domain is characterized by an unknown distribution $P^{(\mathcal{T})}(\mathbf{X}, Y)$ that is entirely inaccessible during training. Consistent with our assumptions, \mathcal{T} shares the same invariant conditional distribution $P(Y|\mathbf{X})$ and label space $P(Y)$ as the source domains, but differs in its input distribution:

$$P^{(\mathcal{T})}(\mathbf{X}) \neq P^{(k)}(\mathbf{X}), \quad \forall k \in \{1, \dots, K\}. \quad (1)$$

Model performance is evaluated directly on samples $\mathbf{X}^* \sim P^{(\mathcal{T})}(\mathbf{X})$ from the unseen target domain without any adaptation at test time.

4. Methodology

In this section, we introduce our proposed framework for SSMDG. For clarity, we describe our method using a bimodal case with video (v) and audio (a), though it naturally generalizes to multiple modalities. The model comprises modality-specific encoders, unimodal classifiers, and a fusion classifier operating on concatenated features (denoted as $[\cdot; \cdot]$). Training proceeds in a mini-batch fashion, sampling from the union of labeled data $\bigcup_k \mathcal{D}_k^l$ and unlabeled data $\bigcup_k \mathcal{D}_k^u$ across K source domains.

For each labeled sample $(\mathbf{X}^{(k)}, y^{(k)}) = ((\mathbf{x}_v^{(k)}, \mathbf{x}_a^{(k)}), y^{(k)})$ from domain k , we apply weak augmentations and encode each modality to obtain \mathbf{f}_v and \mathbf{f}_a . From these features, we compute the unimodal predictions $\hat{\mathbf{p}}_v^l, \hat{\mathbf{p}}_a^l$ and the fused prediction $\hat{\mathbf{p}}_f^l$ based on $[\mathbf{f}_v; \mathbf{f}_a]$. For each unlabeled sample $\mathbf{U}^{(k)} = (\mathbf{u}_v^{(k)}, \mathbf{u}_a^{(k)})$ from domain k , we generate two augmented views: a weakly augmented view and a strongly augmented view. Specifically, weak augmentations (\mathcal{A}_w) consist of standard spatial or acoustic transforms like random flips or pitch shifts, while strong augmentations (\mathcal{A}_s) employ aggressive perturbations such as RandAugment [6] for video and SpecAugment [26] for audio. Passing these through the encoders yields weak-view features ($\mathbf{z}_v^w, \mathbf{z}_a^w$) and the strong-view features ($\mathbf{z}_v^s, \mathbf{z}_a^s$). The corresponding predictions are $\hat{\mathbf{p}}_v^w, \hat{\mathbf{p}}_a^w, \hat{\mathbf{p}}_f^w$ for the weak view, and $\hat{\mathbf{p}}_v^s, \hat{\mathbf{p}}_a^s, \hat{\mathbf{p}}_f^s$ for the strong view. All predictions $\hat{\mathbf{p}}$ represent the softmax probabilities over $|\mathcal{Y}|$ classes.

Our framework integrates three complementary components (Fig. 2): (1) Consensus-Driven Consistency Regularization, which generates reliable pseudo-labels based on cross-modal agreement; (2) Disagreement-Aware Regularization, which leverages ambiguous non-consensus samples; and (3) Cross Modal Prototype Alignment, which enforces cross-modal and cross-domain feature consistency.

4.1. Consensus-Driven Consistency Regularization

In semi-supervised learning, reliable pseudo-labels are essential for effective consistency regularization [35]. Motivated by the idea that joint decisions from multiple perspectives are generally more robust than single-view predictions [12], we propose Consensus-Driven Consistency Regularization (CDCR). This component employs a modality-consensus selection strategy, identifying unlabeled samples that exhibit both high confidence and cross-modal agreement between fused and unimodal predictions.

For each unlabeled sample, we derive a pseudo-label $\hat{y} = \text{argmax} \hat{\mathbf{p}}_f^w$ from the weakly augmented fused prediction. A sample is included in the consensus set $\mathcal{B}_{\text{cdcr}}^u$ if: (1) the fused prediction confidence exceeds a threshold τ , i.e., $\max \hat{\mathbf{p}}_f^w > \tau$; (2) \hat{y} matches at least one unimodal pseudo-label ($\text{argmax} \hat{\mathbf{p}}_v^w$ or $\text{argmax} \hat{\mathbf{p}}_a^w$); and (3) the corresponding unimodal prediction(s) also exceed τ . For sam-

ples that meet these criteria, we enforce prediction consistency across modalities using cross-entropy loss between the pseudo-label and the strong-view predictions:

$$\mathcal{L}_{\text{cdcr}} = \frac{1}{|\mathcal{B}_{\text{cdcr}}^u|} \sum_{U \in \mathcal{B}_{\text{cdcr}}^u} \sum_{n \in \{v, a, f\}} \mathcal{H}(\hat{y}, \hat{\mathbf{p}}_n^s), \quad (2)$$

where $\mathcal{H}(\cdot, \cdot)$ is the cross-entropy loss. This FixMatch-style [29] regularization encourages the model to maintain consistent predictions across augmented views, leveraging only high-confidence, cross-modally consistent pseudo-labels for supervision.

4.2. Disagreement-Aware Regularization

While CDCR focuses on reliable samples, it inevitably discards many potentially informative ones. To exploit these, we define a non-consensus set $\mathcal{B}_{\text{dar}}^u$, which includes samples that fail the CDCR consensus criteria but still exhibit high fused prediction confidence ($\max \hat{\mathbf{p}}_f^w > \tau$). For this set, we introduce Disagreement-Aware Regularization (DAR). DAR leverages the fused pseudo-label as a supervisory signal but replaces the standard cross-entropy with the Generalized Cross-Entropy (GCE) loss [38], which is known for its robustness to noisy labels. This makes it particularly suitable for non-consensus samples, where pseudo-labels are more prone to errors.

Formally, for a pseudo-label index \hat{y} and its corresponding probability distribution \mathbf{p} , the GCE loss is defined as $\mathcal{L}_{\text{GCE}}(\hat{y}, \mathbf{p}) = (1 - (\mathbf{p})_{\hat{y}}^q)/q$, where $(\mathbf{p})_{\hat{y}}$ is the probability assigned to the class \hat{y} , and $q \in (0, 1]$ controls robustness. The DAR objective is then formulated as:

$$\mathcal{L}_{\text{dar}} = \frac{1}{|\mathcal{B}_{\text{dar}}^u|} \sum_{U \in \mathcal{B}_{\text{dar}}^u} \sum_{j \in \{w, s\}} \sum_{n \in \{v, a, f\}} \mathcal{L}_{\text{GCE}}(\hat{y}, \hat{\mathbf{p}}_n^j). \quad (3)$$

This objective applies GCE between the weak-fused pseudo-label \hat{y} and all corresponding unimodal and fused predictions from both weak ($j = w$) and strong ($j = s$) views, encouraging stable learning even in the presence of pseudo-label uncertainty.

4.3. Cross Modal Prototype Alignment

While CDCR and DAR operate in the prediction space, learning domain-invariant and cross-modally consistent features is crucial for multimodal domain generalization [23]. To this end, we propose Cross-Modal Prototype Alignment (CMPA), which aligns features with both intra-domain class prototypes and corresponding prototypes from other modalities and domains.

We maintain running-average class-prototypes $\mu_{m,c}^{(k)}$ for each modality $m \in \{v, a\}$, class $c \in \mathcal{Y}$, and domain k . These are updated via exponential moving average (EMA) using the set $\mathcal{F}_{m,c}^{(k)}$, defined as the set of labeled features \mathbf{f}_m

Table 1. Semi-supervised multimodal domain generalization results with 5 labels per class.

Method	Modality		HAC dataset				EPIC-Kitchens dataset			
	Video	Audio	A, C → H	H, C → A	H, A → C	Mean	D2, D3 → D1	D1, D3 → D2	D1, D2 → D3	Mean
<i>Baseline method</i>										
Source-only	✓	✓	42.15	47.20	37.82	42.39	25.31	32.12	30.95	29.46
<i>Multimodal domain generalization methods (MMDG)</i>										
RNA-Net	✓	✓	43.82	45.63	35.44	41.63	24.14	31.87	29.32	28.44
SimMMDG	✓	✓	46.33	48.32	38.53	44.39	25.62	32.66	33.26	31.11
MOOSA	✓	✓	45.16	46.33	39.36	43.62	25.88	34.53	32.91	30.51
CMRF	✓	✓	44.10	46.88	38.40	43.13	26.31	33.49	33.62	31.21
MDJA	✓	✓	47.03	47.19	38.62	44.28	25.69	34.65	34.20	31.51
<i>Unimodal semi-supervised learning methods (SSL)</i>										
FixMatch	✓		50.36	52.61	43.26	48.74	32.14	34.59	30.90	32.54
FreeMatch	✓		51.32	51.98	42.19	48.50	32.31	35.20	32.06	33.19
CGMatch	✓		51.93	52.30	43.08	49.10	33.61	34.87	31.77	33.42
FixMatch		✓	32.62	31.29	25.80	29.90	24.05	25.32	24.33	24.57
FreeMatch		✓	32.17	32.20	25.33	29.90	22.96	23.64	25.14	23.91
CGMatch		✓	33.05	31.59	26.42	30.35	23.87	25.23	25.88	24.99
FixMatch ^M	✓	✓	59.78	57.95	45.17	54.30	34.25	37.20	36.17	35.87
FreeMatch ^M	✓	✓	58.32	58.66	44.33	53.77	33.98	36.34	36.92	35.75
CGMatch ^M	✓	✓	60.74	60.35	45.26	55.45	34.30	37.27	36.54	36.04
<i>Semi-supervised multimodal learning methods (SSML)</i>										
Co-training	✓	✓	56.31	54.68	40.32	50.44	32.64	34.67	32.06	33.12
TCGM	✓	✓	62.22	61.09	46.62	56.64	35.71	38.33	35.52	36.52
MSC	✓	✓	62.86	62.52	45.21	56.86	34.72	36.20	34.33	35.08
DECPL	✓	✓	63.60	61.31	45.73	56.88	34.64	38.33	35.01	35.99
STIL	✓	✓	64.31	63.80	46.90	58.34	35.48	38.47	36.04	36.66
<i>Unimodal semi-supervised domain generalization methods (SSDG)</i>										
StyleMatch	✓		51.82	53.50	43.99	49.77	31.69	33.96	30.69	32.11
UPUD	✓		51.96	53.46	44.01	49.81	31.44	34.62	30.94	32.33
NIED-LR	✓		52.33	54.33	43.26	49.97	32.86	35.72	31.26	33.28
StyleMatch		✓	33.20	32.80	26.23	30.74	23.31	26.02	28.03	25.79
UPUD		✓	32.05	33.02	26.96	30.68	23.04	25.65	26.79	25.16
NIED-LR		✓	33.28	32.49	26.52	30.76	22.88	25.32	25.82	24.67
StyleMatch ^M	✓	✓	61.69	62.93	46.02	56.88	34.48	37.33	36.04	35.95
UPUD ^M	✓	✓	61.33	62.24	49.88	57.82	35.63	38.26	37.17	37.02
NIED-LR ^M	✓	✓	62.32	63.36	47.23	57.64	36.09	37.60	37.68	37.12
<i>Semi-supervised multimodal domain generalization method (SSMDG)</i>										
Ours	✓	✓	65.58	65.34	51.38	60.77	37.24	41.73	40.86	39.94

from the current batch belonging to modality m , class c , and domain k :

$$\boldsymbol{\mu}_{m,c}^{(k)} \leftarrow \alpha \boldsymbol{\mu}_{m,c}^{(k)} + (1 - \alpha) \cdot \frac{1}{|\mathcal{F}_{m,c}^{(k)}|} \sum_{\mathbf{f}_m \in \mathcal{F}_{m,c}^{(k)}} \mathbf{f}_m, \quad (4)$$

where α is the momentum. To facilitate cross-modal alignment and enable robustness to missing modalities, we introduce modality translators $t_{v \rightarrow a}$ and $t_{a \rightarrow v}$ that map features from one modality to another: $\tilde{\mathbf{z}}_a^j = t_{v \rightarrow a}(\mathbf{z}_v^j)$ and $\tilde{\mathbf{z}}_v^j = t_{a \rightarrow v}(\mathbf{z}_a^j)$ for $j \in \{w, s\}$. These translators not only support alignment but also handle missing-modality scenarios at inference by synthesizing a missing feature from a present one [8]. For example, if the audio feature is missing, it is synthesized from the video feature via the translator $t_{v \rightarrow a}$ (and vice versa).

For all pseudo-labeled unlabeled samples $U \in \mathcal{B}^u = \mathcal{B}_{\text{cdr}}^u \cup \mathcal{B}_{\text{dar}}^u$ with \hat{y} , we align both original and translated

features to intra-domain prototypes $\boldsymbol{\mu}_{m,\hat{y}}^{(k)}$ and cross-domain averages $\bar{\boldsymbol{\mu}}_{m,\hat{y}}^{-k} = \frac{1}{K-1} \sum_{k' \neq k} \boldsymbol{\mu}_{m,\hat{y}}^{(k')}$ using:

$$\mathcal{L}_{\text{align}}(\mathbf{z}, \boldsymbol{\mu}, \bar{\boldsymbol{\mu}}) = \|\mathbf{z} - \boldsymbol{\mu}\|_2^2 + \|\mathbf{z} - \bar{\boldsymbol{\mu}}\|_2^2. \quad (5)$$

The total prototype loss is:

$$\mathcal{L}_{\text{cmpa}} = \frac{1}{|\mathcal{B}^u|} \sum_{U \in \mathcal{B}^u} \sum_{j \in \{w, s\}} \sum_{m \in \{v, a\}} \left[\mathcal{L}_{\text{align}}(\mathbf{z}_m^j, \boldsymbol{\mu}_{m,\hat{y}}^{(k)}, \bar{\boldsymbol{\mu}}_{m,\hat{y}}^{-k}) + \mathcal{L}_{\text{align}}(\tilde{\mathbf{z}}_m^j, \boldsymbol{\mu}_{m,\hat{y}}^{(k)}, \bar{\boldsymbol{\mu}}_{m,\hat{y}}^{-k}) \right]. \quad (6)$$

4.4. Overall Objective

For labeled samples in batch \mathcal{B}^l , we apply standard supervision:

$$\mathcal{L}_{\text{sup}} = \frac{1}{|\mathcal{B}^l|} \sum_{(\mathbf{x}, y) \in \mathcal{B}^l} \sum_{n \in \{v, a, f\}} \mathcal{H}(y, \hat{\mathbf{p}}_n^l). \quad (7)$$

Table 2. SSMDG results on the HAC dataset with 10, 5%, and 10% labels per class. We report the mean accuracy across all domains.

Method	Modality			10	5%	10%
	Video	Audio	Flow			
CGMatch ^M [4]	✓	✓		59.33	58.74	59.32
STiL [11]	✓	✓		60.66	58.96	60.21
NIED-LR ^M [16]	✓	✓		62.92	61.08	62.86
Ours	✓	✓		65.00	65.67	65.90
CGMatch ^M [4]	✓		✓	58.05	59.62	60.33
STiL [11]	✓		✓	60.48	60.33	61.20
NIED-LR ^M [16]	✓		✓	62.23	61.62	61.82
Ours	✓		✓	64.71	64.03	64.45
CGMatch ^M [4]		✓	✓	44.79	44.32	46.32
STiL [11]		✓	✓	45.90	45.77	47.08
NIED-LR ^M [16]		✓	✓	46.69	46.73	47.66
Ours		✓	✓	49.88	49.69	50.13
CGMatch ^M [4]	✓	✓	✓	59.81	60.21	62.57
STiL [11]	✓	✓	✓	62.83	63.18	63.11
NIED-LR ^M [16]	✓	✓	✓	63.36	64.46	63.82
Ours	✓	✓	✓	66.95	66.44	66.91

The overall objective is:

$$\mathcal{L} = \mathcal{L}_{\text{sup}} + \lambda_1 \mathcal{L}_{\text{cdcr}} + \lambda_2 \mathcal{L}_{\text{dar}} + \lambda_3 \mathcal{L}_{\text{cmpa}}, \quad (8)$$

where $\lambda_1, \lambda_2, \lambda_3$ are hyperparameters. This objective is designed to leverage both labeled and unlabeled data while promoting robust cross-domain and cross-modal generalization.

5. Experiments

5.1. Settings

SSMDG Benchmark Datasets: We adapt two widely used multimodal domain generalization datasets, EPIC-Kitchens [7] and HAC [8], to benchmark SSMDG methods. EPIC-Kitchens contains eight kitchen actions across three environments (domains D1, D2, D3). HAC comprises seven actions performed by humans, animals, and cartoons (domains H, A, C).

Evaluation: Following the leave-one-domain-out protocol [41], two domains are used as source domains for training, and the remaining domain serves as the unseen target for evaluation. We report Top-1 classification accuracy under two SSMDG configurations: (1) 5 or 10 labeled samples per class randomly sampled from each source domain (with the remaining samples being unlabeled); and (2) 5% or 10% of samples per class labeled in each source domain. We also assess robustness to missing modalities by simulating unavailable video or audio inputs during testing.

Baselines: We compare our approach with state-of-the-art methods across five categories. The Source-only baseline represents the model trained strictly on labeled source domain data without any DG or semi-supervised techniques.

The MMDG methods, which rely solely on labeled data, include RNA-NET [27], SimMMDG [8], MOOSA [9], CMRF [13], and MDJA [23]. The SSL group includes FixMatch [29], FreeMatch [33], and CGMatch [4]. The SSML methods consist of Co-training [2], TCGM [30], MSC [3], DECPL [1] and STiL [11]. Finally, the SSDG group comprises StyleMatch [41], UPUD [22], and NIED-LR [16]. SSL and SSDG methods are evaluated in both their original unimodal versions and multimodal extensions (denoted as ^M) for fairness.

5.2. Main Results

SSMDG Performance: We comprehensively benchmark our framework against four related paradigms: MMDG, SSL, SSDG, and SSML. Table 1 reports results for the challenging 5-label-per-class setting, while Table 2 summarizes mean performance against top baselines (CGMatch^M [4], STiL [11], NIED-LR^M [16]) across various label configurations (10, 5%, 10%) and modality combinations (2- and 3-modality). Our framework establishes a new state-of-the-art across all configurations, highlighting the limitations of existing methods under this composite SSMDG challenge.

We observe that each paradigm exhibits a key limitation under the SSMDG setting: (1) MMDG methods (e.g., MDJA [23]) struggle under extreme label scarcity (Table 1), demonstrating their reliance on fully labeled data and inability to leverage unlabeled samples. (2) SSL methods (e.g., CGMatch^M [4]), while using unlabeled data, are not designed for domain generalization; their simple multimodal extensions consequently fail to bridge the domain gap. (3) SSML methods (e.g., STiL [11]), which handle sparse multimodal labels but lack domain generalization capabilities, resulting in poor performance on unseen target domains. (4) SSDG methods (e.g., NIED-LR^M [16]) are designed for unimodal shifts, but their simple multimodal adaptations (^M) fail to address the coupled challenges of inter-modal semantics and multimodal domain shifts.

In contrast, our framework is explicitly designed to simultaneously address sparse labels, domain shifts, and multimodal interactions. In the 5-label-per-class setting (Table 1), we achieve mean accuracies of 60.77% on HAC and 39.94% on EPIC-Kitchens, substantially outperforming all baselines. Table 2 further confirms that our superiority is consistent across label regimes and modality combinations. Importantly, our method scales effectively to three-modality inputs (Video-Audio-Flow), demonstrating its robustness and suitability for complex SSMDG tasks.

Pseudo-Label Quality and Unlabeled Data Utilization: We analyze pseudo-label quality and unlabeled data utilization on the HAC benchmark under the 5-label and 5%-label settings, comparing our framework with competitive baselines including FixMatch^M [29], STiL [11], StyleMatch^M [41], and NIED-LR^M [16]. As illustrated in Fig. 3, our

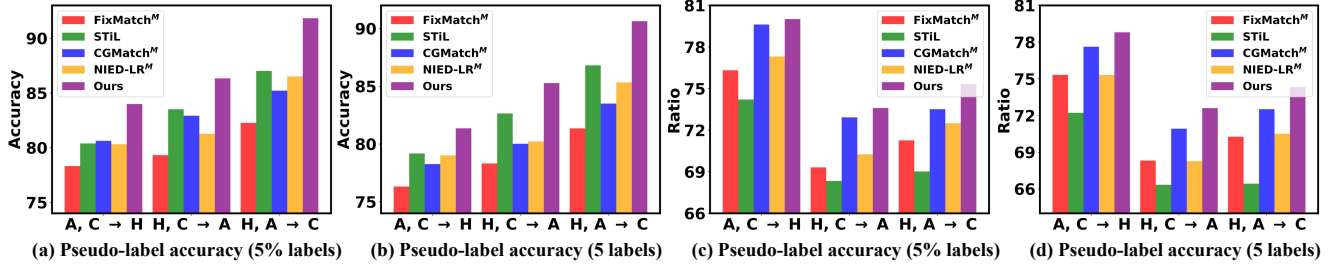


Figure 3. Pseudo-label accuracy and unlabeled data utilization on the HAC benchmark. Our method consistently achieves superior accuracy while maintaining a higher utilization rate compared to competitive baselines.

Table 3. SSMDG with missing modalities. \times means the modality is available during training, but is missing in test time.

Method	Modality		HAC dataset			EPIC-Kitchens dataset				
	Video	Audio	A, C \rightarrow H	H, C \rightarrow A	H, A \rightarrow C	Mean	D2, D3 \rightarrow D1	D1, D3 \rightarrow D2	D1, D2 \rightarrow D3	Mean
<i>5% labels per class</i>										
Audio-only		✓	32.82	31.06	24.36	29.41	25.62	25.01	26.62	25.75
Zero-filling	×	✓	30.86	30.29	23.66	28.27	24.91	25.04	26.31	25.42
Translation	×	✓	37.88	38.09	31.62	35.86	31.06	30.62	31.58	31.09
Video-only	✓		51.62	52.98	48.62	51.07	35.92	35.60	34.69	35.40
Zero-filling	✓	×	53.79	54.26	49.30	52.45	36.78	36.92	35.41	36.37
Translation	✓	×	55.79	57.62	51.68	55.03	39.27	38.64	39.07	38.99
<i>5 labels per class</i>										
Audio-only		✓	33.54	31.92	24.84	30.10	24.63	25.60	24.18	24.80
Zero-filling	×	✓	32.71	30.68	23.05	28.81	25.85	26.17	25.47	25.83
Translation	×	✓	39.01	37.55	29.54	35.37	28.88	29.04	27.96	28.63
Video-only	✓		51.08	52.66	43.86	49.20	35.37	32.69	33.33	33.80
Zero-filling	✓	×	51.06	52.66	43.77	49.16	36.74	32.81	33.26	34.27
Translation	✓	×	53.44	55.82	45.69	51.65	38.15	36.82	37.83	37.60

method achieves higher pseudo-label accuracy while maintaining a greater data utilization rate. These findings confirm that our framework effectively extracts reliable supervision from abundant unlabeled multimodal data.

Missing-Modality SSMDG: In practical scenarios, some modalities may be unavailable at test time. Table 3 evaluates robustness to missing video or audio data, comparing unimodal baselines, naive zero-filling, and our cross-modal translation strategy. Our approach leverages trained translators ($t_{v \rightarrow a}, t_{a \rightarrow v}$) to generate semantically consistent features for the missing modality, mitigating performance degradation. In contrast, naive zero-filling often underperforms in the shared feature space. For instance, in the 5%-label setting with a missing video modality, our translation strategy outperforms zero-filling by 7.59% and 5.67% on all benchmarks. These results demonstrate that our method can effectively reconstruct missing semantic information while preserving multimodal generalization.

5.3. Ablation Studies

Efficacy of Key Model Components: Table 4 evaluates the contributions of CDCR, DAR, and CMPA under the 5-label setting. The weakly supervised baseline (Row 1)

highlights that generalization is severely constrained under label scarcity. Incorporating CDCR or DAR individually yields notable improvements, while their combination (Row 4) further boosts performance by jointly leveraging reliable consensus and ambiguous non-consensus samples. CMPA consistently improves all configurations by enforcing feature-space invariance, complementing the prediction-level regularizers. The full framework (Row 7) achieves the best performance, improving HAC accuracy by 15.38% over the baseline.

Ablation Study on CDCR: We compare alternative consensus strategies in Table 5. Mean-CDCR, similar to the STIL [11] method, averages predictions from all modalities meeting the threshold, Any2-CDCR requires agreement from any two views, and Strict-CDCR mandates consensus across all views. Our Full-CDCR, which anchors on the fused prediction supported by at least one modality, achieves the best results, outperforming Mean-CDCR by 2.51% (HAC) and 2.33% (EPIC-Kitchens), and also surpasses the stricter (Strict-CDCR) and more relaxed (Any2-CDCR) variants, validating our selection design.

Ablation Study on DAR: Table 5 evaluates DAR design choices. Replacing the GCE loss with standard

Table 4. Ablations of each proposed module under 5 labels per class.

CDCR	DAR	CMPA	HAC dataset				EPIC-Kitchens dataset			
			A, C → H	H, C → A	H, A → C	Mean	D2, D3 → D1	D1, D3 → D2	D1, D2 → D3	Mean
			48.62	50.28	37.26	45.39	29.06	32.74	33.92	31.91
✓			59.92	58.67	45.69	54.76	34.72	37.69	37.02	36.48
	✓		60.69	57.95	46.03	54.89	34.63	37.62	36.97	36.41
✓	✓		63.37	64.26	48.28	58.64	36.02	38.63	38.91	37.85
✓		✓	64.37	63.69	47.58	58.55	36.72	37.16	37.09	36.99
	✓	✓	63.76	62.73	48.09	58.19	35.86	37.23	36.90	36.66
✓	✓	✓	65.58	65.34	51.38	60.77	37.24	41.73	40.86	39.94

Table 5. Ablation study results demonstrating the effectiveness of individual components under 5 labels per class.

Method	HAC	EPIC-Kitchens
Mean-CDCR	58.26	37.61
Any2-CDCR	59.34	38.25
Strict-CDCR	59.16	37.32
Full-CDCR	60.77	39.94
CE-DAR	58.32	37.92
Weak-only DAR	60.42	38.97
Strong-only DAR	60.73	39.62
Full-DAR	60.77	39.94
Intra-Domain CMPA	58.34	38.45
Intra-Modal CMPA	60.75	39.64
Weak-only CMPA	60.22	38.64
Strong-only CMPA	59.38	39.25
Full-CMPA	60.77	39.94




Figure 4. Ablation analysis of different modules on pseudo-label accuracy and unlabeled data utilization.

cross-entropy (CE-DAR) reduces performance (2.45% on HAC, 2.02% on EPIC-Kitchens), highlighting the need for a noise-robust loss for ambiguous non-consensus samples. Applying DAR to only weak (Weak-only) or strong (Strong-only) views is suboptimal, whereas Full-DAR, enforcing consistency across both views, achieves the highest accuracy. This demonstrates the advantage of combined regularization.

Ablation Study on CMPA: In Table 5, we present a detailed investigation into the specific design choices constituting the proposed CMPA module. Specifically, when cross-domain alignment is completely removed (referred

to as Intra-Domain CMPA), we observe a significant performance drop of 2.43% on the HAC dataset and 1.49% on the EPIC-Kitchens dataset. This substantial decline clearly highlights the critical importance of maintaining cross-domain feature invariance for robust representation learning. Similarly, omitting the cross-modal translation component (denoted as Intra-Modal CMPA) also leads to a noticeable degradation in overall performance, confirming the benefits of inter-modal semantic consistency. Finally, our results demonstrate that aligning both weak and strong augmented views simultaneously (Full-CMPA) consistently outperforms alternative variants that rely on only a single view. This finding strongly validates the effectiveness of our comprehensive joint alignment strategy.

Impact on Pseudo-Labeling: We assess each module’s effect on pseudo-label quality and unlabeled data utilization on HAC (5-label setting), using FixMatch [29] as the baseline. Fig. 4 (a) shows that each component progressively improves pseudo-label accuracy. Fig. 4 (b) illustrates that CDCR initially reduces utilization due to its strict consensus filtering. However, the subsequent addition of DAR and CMPA significantly increases utilization: CMPA stabilizes feature representations, generating more consistent predictions that satisfy confidence and consensus thresholds.

6. Conclusion

In this work, we introduced Semi-Supervised Multimodal Domain Generalization (SSMDG), a novel problem that unifies the critical challenges of cross-domain generalization and data-efficient multimodal learning. We analyzed the unique difficulties inherent to SSMDG and highlighted the limitations of existing SSML, MMDG, and SSDG paradigms in addressing them. To bridge this gap, we established the first comprehensive SSMDG benchmarks, including evaluations under missing-modality scenarios, and proposed a new framework that effectively leverages unlabeled multimodal data while enforcing cross-domain and cross-modal invariance. Our method sets a strong performance baseline, and we hope that our problem formulation, benchmarks, and framework will stimulate further research in this practical and impactful direction.

Acknowledgments

This work was supported by the National Key R&D Program of China (2024YFB3311600) and the Natural Science Foundation of Henan (252300423936).

References

- [1] Maregu Assefa, Wei Jiang, Jinyu Zhan, Kumie Gedamu, Getinet Yilma, Melese Ayalew, and Deepak Adhikari. Audio-visual contrastive and consistency learning for semi-supervised action recognition. *IEEE Transactions on Multimedia*, 26:3491–3504, 2024. 1, 2, 6
- [2] Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 92–100, 1998. 6
- [3] Junchi Chen, Richong Zhang, and Junfan Chen. Semi-supervised multimodal classification through learning from modal and strategic complementarities. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 15812–15820, 2025. 2, 6
- [4] Bo Cheng, Jueqing Lu, Yuan Tian, Haifeng Zhao, Yi Chang, and Lan Du. Cgmatch: A different perspective of semi-supervised learning. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 15381–15391, 2025. 1, 6
- [5] Yanhua Cheng, Xin Zhao, Rui Cai, Zhiwei Li, Kaiqi Huang, Yong Rui, et al. Semi-supervised multimodal deep learning for rgb-d object recognition. In *IJCAI*, pages 3345–3351, 2016. 2
- [6] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020. 4
- [7] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *Proceedings of the European conference on computer vision (ECCV)*, pages 720–736, 2018. 6
- [8] Hao Dong, Ismail Nejjar, Han Sun, Eleni Chatzi, and Olga Fink. Simmmdg: A simple and effective framework for multi-modal domain generalization. *Advances in Neural Information Processing Systems*, 36:78674–78695, 2023. 1, 2, 5, 6
- [9] Hao Dong, Eleni Chatzi, and Olga Fink. Towards multimodal open-set domain generalization and adaptation through self-supervision. In *European Conference on Computer Vision*, pages 270–287. Springer, 2024. 1, 2, 6
- [10] Hao Dong, Moru Liu, Kaiyang Zhou, Eleni Chatzi, Juho Kannala, Cyrill Stachniss, and Olga Fink. Advances in multimodal adaptation and generalization: From traditional approaches to foundation models. *arXiv preprint arXiv:2501.18592*, 2025. 1
- [11] Siyi Du, Xinzhe Luo, Declan P O’Regan, and Chen Qin. Stil: Semi-supervised tabular-image learning for comprehensive task-relevant information exploration in multimodal classification. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 15549–15559, 2025. 1, 2, 6, 7
- [12] Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. In *Proceedings of the 41st International Conference on Machine Learning*, pages 11733–11763, 2024. 4
- [13] Yunfeng Fan, Wenchao Xu, Haozhao Wang, and Song Guo. Cross-modal representation flattening for multi-modal domain generalization. *Advances in Neural Information Processing Systems*, 37:66773–66795, 2024. 2, 6
- [14] Enrico Fini, Pietro Astolfi, Karteek Alahari, Xavier Alameda-Pineda, Julien Mairal, Moin Nabi, and Elisa Ricci. Semi-supervised learning made simple with self-supervised clustering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3187–3197, 2023. 1
- [15] Chamuditha Jayanga Galappaththige, Sanoojan Baliah, Malitha Gunawardhana, and Muhammad Haris Khan. Towards generalizing to unseen domains with few labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23691–23700, 2024. 2
- [16] Chamuditha Jayanaga Galappaththige, Zachary Izzo, Xilin He, Honglu Zhou, and Muhammad Haris Khan. Domain-guided weight modulation for semi-supervised domain generalization. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 6495–6505. IEEE, 2025. 1, 2, 6
- [17] Cagri Gungor and Adriana Kovashka. Integrating audio narrations to strengthen domain generalization in multimodal first-person action recognition. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2025. 2
- [18] Saheli Hazra, Sudip Das, Rohit Choudhary, Arindam Das, Ganesh Sistu, Ciarán Eising, and Ujjwal Bhattacharya. Reflective teacher: Semi-supervised multimodal 3d object detection in bird’s-eye-view via uncertainty measure. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1649–1659. IEEE, 2025. 2
- [19] Hai Huang, Yan Xia, Sashuai Zhou, Hanting Wang, Shulei Wang, and Zhou Zhao. Bridging domain generalization to multimodal domain generalization via unified representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22488–22498, 2025. 2
- [20] Shirui Huang, Keyan Wang, Huan Liu, Jun Chen, and Yunsong Li. Contrastive semi-supervised learning for underwater image restoration via reliable bank. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18145–18155, 2023. 1
- [21] Adnan Khan, Mai A Shaaban, and Muhammad Haris Khan. Improving pseudo-labelling and enhancing robustness for semi-supervised domain generalization. *arXiv preprint arXiv:2401.13965*, 2024. 2

- [22] Dongkwan Lee, Kyomin Hwang, and Nojun Kwak. Unlocking the potential of unlabeled data in semi-supervised domain generalization. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 30599–30608, 2025. 1, 2, 6
- [23] Hongzhao Li, Hualei Wan, Liangzhi Zhang, Mingyuan Jiu, Shupan Li, Mingliang Xu, and Muhammad Haris Khan. Towards robust multimodal domain generalization via modality-domain joint adversarial training. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 180–188, 2025. 1, 2, 4, 6
- [24] Hongzhao Li, Guohao Shen, Shupan Li, Mingliang Xu, and Muhammad Haris Khan. Balancing multimodal domain generalization via gradient modulation and projection. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(27):22922–22930, 2026. 2
- [25] Yunsong Li, Daixun Li, Weiying Xie, Jitao Ma, Sibao He, and Leyuan Fang. Semi-mamba: Mamba-driven semi-supervised multimodal remote sensing feature classification. *IEEE Transactions on Circuits and Systems for Video Technology*, 2025. 2
- [26] Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le. SpecAugment: A simple data augmentation method for automatic speech recognition. *arXiv preprint arXiv:1904.08779*, 2019. 4
- [27] Mirco Planamente, Chiara Plizzari, Simone Alberto Peirone, Barbara Caputo, and Andrea Bottino. Relative norm alignment for tackling domain shift in deep multi-modal classification. *International Journal of Computer Vision*, 132(7): 2618–2638, 2024. 1, 2, 6
- [28] Mamshad Nayeem Rizve, Navid Kardan, and Mubarak Shah. Towards realistic semi-supervised learning. In *European conference on computer vision*, pages 437–455. Springer, 2022. 1
- [29] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 1, 2, 4, 6, 8
- [30] Xinwei Sun, Yilun Xu, Peng Cao, Yuqing Kong, Lingjing Hu, Shanghang Zhang, and Yizhou Wang. Tcgm: An information-theoretic framework for semi-supervised multimodality learning. In *European conference on computer vision*, pages 171–188. Springer, 2020. 1, 2, 6
- [31] Jesper E Van Engelen and Holger H Hoos. A survey on semi-supervised learning. *Machine learning*, 109(2):373–440, 2020. 1
- [32] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Tao Qin, Wang Lu, Yiqiang Chen, Wenjun Zeng, and Philip S Yu. Generalizing to unseen domains: A survey on domain generalization. *IEEE transactions on knowledge and data engineering*, 35(8):8052–8072, 2022. 1
- [33] Yidong Wang, Hao Chen, Qiang Heng, Wenxin Hou, Yue Fan, Zhen Wu, Jindong Wang, Marios Savvides, Takahiro Shinozaki, Bhiksha Raj, et al. Freematch: Self-adaptive thresholding for semi-supervised learning. In *Eleventh International Conference on Learning Representations*. OpenReview. net, 2023. 1, 6
- [34] Fan Yang, Kai Wu, Shuyi Zhang, Guannan Jiang, Yong Liu, Feng Zheng, Wei Zhang, Chengjie Wang, and Long Zeng. Class-aware contrastive semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14421–14430, 2022. 1
- [35] Xiangli Yang, Zixing Song, Irwin King, and Zenglin Xu. A survey on deep semi-supervised learning. *IEEE transactions on knowledge and data engineering*, 35(9):8934–8954, 2022. 4
- [36] Yuan Yuan, Zhaojian Li, and Bin Zhao. A survey of multimodal learning: Methods, applications, and future. *ACM Computing Surveys*, 57(7):1–34, 2025. 1
- [37] Baoqiang Zhang, Kunze Huang, Luyao Luyao, Xiaotong Tu, and Xiaolu Li. Nonpolarized embedding learning in multimodal domain generalization. *Neurocomputing*, page 131754, 2025. 2
- [38] Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems*, 31, 2018. 4
- [39] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. *arXiv preprint arXiv:2104.02008*, 2021. 1
- [40] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(4):4396–4415, 2022. 1
- [41] Kaiyang Zhou, Chen Change Loy, and Ziwei Liu. Semi-supervised domain generalization with stochastic stylematch. *International Journal of Computer Vision*, 131(9):2377–2387, 2023. 1, 2, 6