arXiv:2511.21331v2 [cs.CV] 3 Apr 2026

THE MORE, THE MERRIER: CONTRASTIVE FUSION FOR HIGHER-ORDER MULTIMODAL ALIGNMENT

Stefanos Koutoupis[1,2,*], Michaela Areti Zervou[1,2], Konstantinos Kontras[3], Maarten De Vos[3], Panagiotis Tsakalides[1,2], Grigorios Tsagkatakis[1,2]

[1]Foundation for Research and Technology-Hellas,
[2]University of Crete, [3]KU Leuven

Abstract

Learning joint representations across multiple modalities remains a central challenge in multimodal machine learning. Prevailing approaches predominantly operate in pairwise settings, aligning two modalities at a time. While some recent methods aim to capture higher-order interactions among multiple modalities, they often overlook or insufficiently preserve pairwise relationships, limiting their effectiveness on single-modality tasks. In this work, we introduce Contrastive Fusion (ConFu), a framework that jointly embeds both individual modalities and their fused combinations into a unified representation space, where modalities and their fused counterparts are aligned. ConFu extends traditional pairwise contrastive objectives with an additional fused-modality contrastive term, encouraging the joint embedding of modality pairs with a third modality. This formulation enables ConFu to capture higher-order dependencies, such as XOR-like relationships, that cannot be recovered through pairwise alignment alone, while still maintaining strong pairwise correspondence. We evaluate ConFu on synthetic and real-world multimodal benchmarks, assessing its ability to exploit cross-modal complementarity, capture higher-order dependencies, and scale with increasing multimodal complexity. Across these settings, ConFu demonstrates competitive performance on retrieval and classification tasks, while supporting unified one-to-one and two-to-one retrieval within a single contrastive framework. We release our code and dataset at https://github.com/estafons/confu.

1. Introduction

Learning joint representations across multiple modalities is a fundamental goal in multimodal machine learning [44]. Contrastive frameworks such as CLIP [26] and ALIGN [16] have demonstrated that aligning paired modalities, such as images and text, enables powerful zero-shot transfer and retrieval. However, these approaches remain intrinsically pairwise, capturing only correlations between two modalities while overlooking the higher-order dependencies that emerge when three or more modalities interact [10, 11, 16, 34–36, 41, 42, 46–48]. Beyond this limitation, multimodal models often exhibit modality competition [14], where one modality dominates and suppresses others when signal strength or reliability differs. For instance, strong visual cues may overshadow weaker auditory ones, biasing the joint representation toward a single source. Yet many real-world settings require reasoning over complementary cues: a song emerges from melody and lyrics, a 3D design from sketch and text, and a scene from both visual and auditory signals. These challenges motivate a broader question:

> Can a contrastive learning framework capture not only pairwise alignments but also higher-order, synergistic dependencies among modalities?

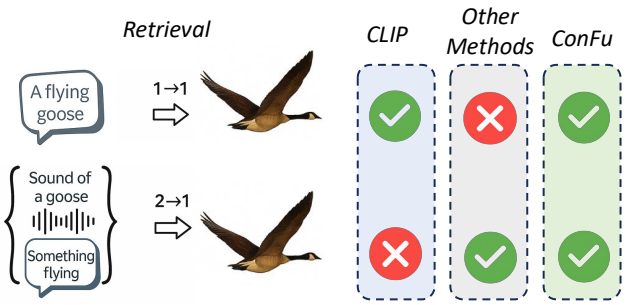

Figure 1. ConFu unifies direct (1→1) and compositional (2→1) alignment within a single embedding space, utilizing two modalities for improved performance and adapting seamlessly when only one modality is available.

*Corresponding author: skoutoupis@ics.forth.gr

To address this, we propose Contrastive Fusion (ConFu), a framework that unifies pairwise and higher-order multimodal supervision within a single contrastive objective. ConFu constructs fused representations of all subsets of modalities and aligns all combinations of unimodal and fused representations (Fig. 2). Unlike previous methods that primarily target specific one-to-one (1→1) [10, 11], two-to-one (2→1) [6, 27], or both within predominantly text-centered objectives [4, 5], ConFu accounts for all pairwise and higher-order modality combinations within a unified objective (Fig. 1).

Evaluating such models is further complicated by the scarcity of publicly available datasets that naturally include three or more co-occurring modalities. To address this, we construct a bird-centric dataset where image, audio, and text provide complementary but related information. This setting enables a systematic examination of higher-order alignment under zero-shot and few-shot regimes. Overall, our main contributions are as follows:

- We introduce a unified contrastive framework that generalizes CLIP-like pairwise multimodal alignment to capture higher-order dependencies among three modalities.
- We construct Bird-MML, a synthetic dataset of artificial triplets that serves as a pretraining resource to evaluate whether models capture multimodal complementarity.
- We demonstrate that ConFu achieves competitive performance in retrieval and classification tasks, supporting both one-to-one and two-to-one retrieval within a single contrastive framework.
- We show that ConFu exhibits greater robustness than previous methods when faced with distracting modalities (Table 4) and noise-induced distribution shifts (Table 6), showing greater stability even when fusion involves corrupted or non-informative inputs.

2. Related Work

Jointly-aligned models. Methods such as AudioCLIP [11], VALOR [3], and VAST [4] extend CLIP-style contrastive learning to three or more modalities by training on datasets where all signals co-occur. AudioCLIP employs independent contrastive objectives for each modality pair, while VALOR introduces a multimodal grouping alignment loss coupled with caption generation to connect multiple modalities through text. VAST combines contrastive and generative losses to unify video, audio, and subtitles in a shared representation. Other architectures such as mPLUG-2 [40], UMT [22], and the Everything-at-Once Fusion Transformer [21] employ modular decoders or fusion transformers for multimodal understanding and generation. While these methods learn representations across multiple modalities, their objectives are typically limited to pairwise alignments or remain primarily text-centric, with broader multimodal relationships emerging implicitly from shared training signals rather than through explicit higher-order alignment.

Pivot-aligned models. An alternative strategy uses a pivot modality as a common reference space. Each non-pivot modality is aligned to this anchor, enabling scalable training without requiring full multimodal co-occurrence. For example, ImageBind [10] aligns diverse modalities (audio, depth, IMU) to an image-based embedding space, while LanguageBind [48] uses text as the pivot. Later variants such as OmniBind [36], FreeBind [35], and Ex-MCR [46] extend this paradigm through pseudo-pair generation, offering an alternate data-efficient mechanism. Similarly, Uni3D [47], ULIP [41], and ULIP-2 [42] project new modalities into a frozen CLIP embedding space,mapping new modalities to CLIP's pretrained space without updating its parameters. While this pivot-based design improves scalability, it remains inherently pairwise: all interactions are mediated through the pivot, preventing direct modeling of dependencies among non-pivot modalities. As a result, the learned embedding space captures independent pairwise relations rather than true higher-order structure.

Modeling Higher-Order Dependencies. Recent work in self-supervised representation learning has attempted to model relationships between more than two modalities through specialized loss formulations. Symile [27] leverages total correlation to integrate complementary cues, TRIANGLE [6] uses the area spanned by three embedding vectors as a measure of similarity, and GRAM [5] measures cross-modal similarity for $M \geq 2$ modalities through the Gramian volume spanned by their vectors. While these methods jointly model multiple modalities, Symile and TRIANGLE require all modalities to be present at inference time, making them incompatible with standard 1→1 retrieval.

In contrast, our framework explicitly models both pairwise and higher-order alignments within a single contrastive objective. It provides a simple, architecture-agnostic formulation that unifies pairwise and higher-order multimodal interactions, enabling scalable and expressive cross-modal representation learning, with the only additional computational cost coming from lightweight MLP layers.

3. Method

3.1. Proposed Multimodal Contrastive Framework

We propose a new framework for contrastive learning with $M = 3$ modalities (Fig. 2) that unifies *pairwise* and *higher-order* contrastive objectives into a single multimodal training process that maximizes a lower bound on the total correlation among modalities (Sec. 3.2).

Each modality X_i is mapped into a feature space by a modality-specific encoder $f_{\theta_i} : \mathcal{X}_i \to \mathcal{H}_i$ and then projected

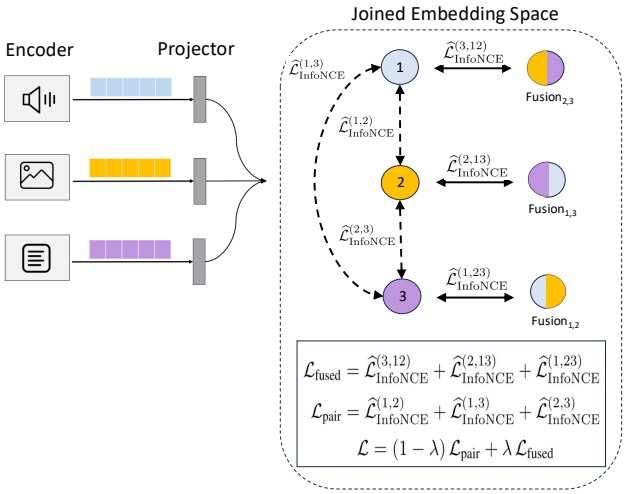

Figure 2. Overview of ConFu. The framework aligns all modality pairs through pairwise contrastive objectives while also aligning each modality with the fused representation of the remaining ones. The final loss ($\mathcal{L}$) combines both objectives ($\mathcal{L}_{pair}$, $\mathcal{L}_{fused}$), balanced by a weighting factor λ.

into a shared latent space $\mathcal{Z}$ by a projector $p_{\phi_i} : \mathcal{H}_i \to \mathcal{Z}$:

$$h_i = f_{\theta_i}(X_i), \quad z_i = p_{\phi_i}(h_i), \quad i \in \{1,2,3\}. \tag{1}$$

The encoder parameters θ_i and projector parameters ϕ_i are learned jointly through contrastive optimization.

Pairwise Contrastive Objectives (1→1). For each modality pair (X_i, X_j), we estimate their MI $I(X_i; X_j)$ via the InfoNCE lower bound Eq. 2. Let $s_{\omega_{ij}} : \mathcal{Z} \times \mathcal{Z} \to \mathbb{R}$ denote a density ratio function parameterized by ω_{ij}, which models the likelihood of joint co-occurrence over independence. Given a minibatch of N paired samples $\{(z_i^{(m)}, z_j^{(m)})\}_{m=1}^N$, the empirical loss is

$$\widehat{\mathcal{L}}^{(i,j)}_{\text{InfoNCE}} = \tfrac{1}{N}\sum_{m=1}^{N} -\log \frac{\exp\{s_{\omega_{ij}}(z_i^{(m)}, z_j^{(m)})\}}{\sum_{n=1}^{N}\exp\{s_{\omega_{ij}}(z_i^{(m)}, z_j^{(n)})\}}. \tag{2}$$

Summing over all ordered modality pairs yields

$$\mathcal{L}_{\text{pair}} = \widehat{\mathcal{L}}^{(1,2)}_{\text{InfoNCE}} + \widehat{\mathcal{L}}^{(1,3)}_{\text{InfoNCE}} + \widehat{\mathcal{L}}^{(2,3)}_{\text{InfoNCE}}. \tag{3}$$

Higher-Order Contrastive Objectives (2→1). To model higher-order dependencies $I(X_k; X_i, X_j)$, we introduce a fusion network $g_{\psi_{ij}} : \mathcal{H}_i \times \mathcal{H}_j \to \mathcal{Z}$ with encoded features:

$$z_{ij} = g_{\psi_{ij}}(h_i, h_j), \quad (i,j) \in \{(1,2),(1,3),(2,3)\}, \tag{4}$$

where ψ_{ij} are trainable fusion parameters. In our experiments, we have maintained g to be a small-scale network (e.g., a shallow MLP) to capture cross-modal interactions without significantly increasing model complexity. Each fused representation z_{ij} is then aligned with the remaining modality X_k, using the ratio density estimation function $t_{\eta_{k,ij}} : \mathcal{Z} \times \mathcal{Z} \to \mathbb{R}$ parameterized by $\eta_{k,ij}$:

$$\widehat{\mathcal{L}}^{(k,ij)}_{\text{InfoNCE}} = \tfrac{1}{N}\sum_{m=1}^{N} -\log \frac{\exp\{t_{\eta_{k,ij}}(z_k^{(m)}, z_{ij}^{(m)})\}}{\sum_{n=1}^{N}\exp\{t_{\eta_{k,ij}}(z_k^{(m)}, z_{ij}^{(n)})\}}. \tag{5}$$

In practice, both density ratio estimators s in Eq. 2 and t in Eq. 5 are implemented as temp.-scaled dot-product similarities. The higher-order objective aggregates all such triplets:

$$\mathcal{L}_{\text{fused}} = \widehat{\mathcal{L}}^{(3,\{1,2\})}_{\text{InfoNCE}} + \widehat{\mathcal{L}}^{(2,\{1,3\})}_{\text{InfoNCE}} + \widehat{\mathcal{L}}^{(1,\{2,3\})}_{\text{InfoNCE}}. \tag{6}$$

Combined Objective. The total training loss jointly maximizes both pairwise and fused MI lower bounds:

$$\mathcal{L} = (1-\lambda)\,\mathcal{L}_{\text{pair}} + \lambda\,\mathcal{L}_{\text{fused}}, \tag{7}$$

where $\lambda \in [0,1]$ controls the relative contribution of pairwise versus higher-order supervision (See Appendix B.1 for analysis of λ's impact). Minimizing Eq. 7 corresponds to maximizing the contrastive lower bound on $\text{TC}(X_1, X_2, X_3)$ (Eq. 13), thus promoting both (i) pairwise alignment and (ii) cross-modal consistency through fused interactions.

3.2. Theoretical Motivation

The total correlation (TC) [37], also known as multi-information [28], quantifies the joint statistical dependency among multiple random variables as the Kullback–Leibler divergence between their joint distribution and the product of marginals, $\text{TC}(X_1, X_2, X_3) = D_{\text{KL}}\big(p(x_1,x_2,x_3) \,\|\, p(x_1)p(x_2)p(x_3)\big)$. Equivalently,

$$\text{TC}(X_1, X_2, X_3) = \mathbb{E}_{p(x_1,x_2,x_3)}\left[\log \frac{p(x_1,x_2,x_3)}{p(x_1)p(x_2)p(x_3)}\right], \tag{8}$$

and vanishes if and only if the variables are mutually independent ($X_1 \perp X_2 \perp X_3$). Using the chain rule of MI, it decomposes as

$$\text{TC}(X_1, X_2, X_3) = I(X_1; X_2) + I(X_3; X_1, X_2). \tag{9}$$

Averaging over all permutations yields the symmetric form

$$\text{TC}(X_1, X_2, X_3) = \tfrac{1}{3}\sum_{(i,j,k)\in\text{perm}\{1,2,3\}} \big[I(X_i; X_j) + I(X_k; X_i, X_j)\big], \tag{10}$$

where $\text{perm}\{1,2,3\}$ denotes the set of all six permutations of the index triplet $(1,2,3)$. This interpretation of TC separates pairwise and higher-order dependencies.

Contrastive Lower Bound. Each MI term in Eq. 10 is approximated using the InfoNCE lower bounds [23]. Specifically, the pairwise and higher-order terms satisfy

$$I(X_i; X_j) \;\geq\; \log N - \widehat{\mathcal{L}}^{(i,j)}_{\text{InfoNCE}}, \tag{11}$$

$$I(X_k; X_i, X_j) \;\geq\; \log N - \widehat{\mathcal{L}}^{(k,ij)}_{\text{InfoNCE}}. \tag{12}$$

Substituting these inequalities into Eq. 10 yields the tractable contrastive lower bound on TC,

$$\mathrm{TC} \geq -\tfrac{1}{3} \sum_{(i,j,k)\in \mathrm{perm}\{1,2,3\}} \left[\widehat{\mathcal{L}}^{(i,j)}_{\mathrm{InfoNCE}} + \widehat{\mathcal{L}}^{(k,ij)}_{\mathrm{InfoNCE}} \right] + 2\log N + \text{const.} \tag{13}$$

so that maximizing total correlation reduces to minimizing the InfoNCE losses across all pairwise and higher-order relations. When each density ratio function $s_{\omega_{ij}}$ or $t_{\eta_{k,ij}}$ is sufficiently expressive and approximates the optimal InfoNCE critic [23] (i.e., the log–density ratio between the joint and the product of marginals) the bound becomes tight as $N \to \infty$, recovering the true TC.

Interpretation. Eq. 13 shows that the proposed training objective (Eq. 7) maximizes a contrastive lower bound on total correlation, thereby promoting both (i) pairwise alignment across modalities through $I(X_i; X_j)$ and (ii) higher-order consistency via $I(X_k; X_i, X_j)$.

3.3. Compare with other higher-order methods

Recent approaches such as Symile, GRAM, and TRIANGLE adopt a *single joint alignment objective* across all modalities, without separating pairwise from higher-order dependencies. Symile maximizes a total-correlation bound using a multilinear similarity $\langle x_1, x_2, x_3 \rangle$, GRAM enforces global alignment via the parallelotope volume induced by the Gram matrix, and TRIANGLE specializes this idea by introducing a triangle-area–based measure of three-way similarity. In contrast, our method decomposes total correlation into interpretable information-theoretic factors, averaging over modality permutations the sum of pairwise and higher-order dependencies, thus preserving holistic expressivity while enabling structured optimization and reliable performance with any subset of modalities. Unlike prior approaches that use a single critic to implicitly model dependencies, we factorize them at the loss level, enabling independent control of pairwise and higher-order supervision while preserving InfoNCE stability and achieving balanced performance within a single model.

3.4. Motivating higher-order dependencies

To illustrate the reasons that modeling higher-order (synergistic) dependencies is crucial, we consider the XOR task introduced in Symile. Three binary variables X_1, X_2, and X_3 are sampled with $x_{1j}, x_{2j} \sim \mathrm{Bernoulli}(0.5)$, $i \sim \mathrm{Bernoulli}(\hat{p})$, and the relationship:

$$x_{3j} = (x_{1j} \oplus x_{2j})^i x_{1j}^{(1-i)}. \tag{14}$$

The task is to predict the representation z_2 from the pair (z_1, z_3). The parameter $\hat{p}$ controls the level of synergy: for $\hat{p} = 0$, x_2 is independent of (x_1, x_3), while increasing $\hat{p}$ raises the conditional MI $I(X_1; X_2|X_3)$ and $I(X_3; X_2|X_1)$ without introducing any pairwise dependence ($I(X_1; X_2) = I(X_2; X_3) = 0$). Thus, only models capturing joint interactions, rather than pairwise correlations, can solve the task. Zero-shot retrieval at varying levels of $\hat{p}$ (Fig. 3) confirms that ConFu and Symile can successfully solve the task.

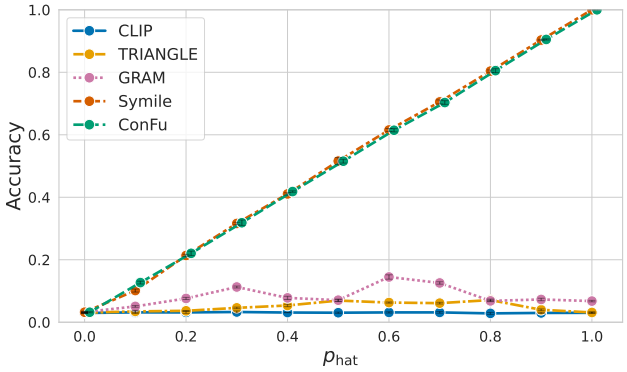

Figure 3. Results on the synthetic XOR task. Our model's accuracy in predicting z_2 from (z_1, z_3) is plotted against the mixing parameter $\hat{p}$. Our model captures the synergistic information, showing a positive trend in performance as $\hat{p}$ increases. Trimodal pairwise CLIP remains on chance ($\sim 3\%$) while both GRAM and TRIANGLE fail to reach above 15% accuracy. More details are provided in Appendix B.2.

4. Bird-MML: A Multimodal Dataset for Audio-Visual Complementarity

Despite substantial progress in multimodal learning, there remains a lack of standardized and publicly available realistic datasets that capture cross-modal relationships that are suitable for evaluating both pairwise and higher-order multimodal alignment. Bird species provide an ideal testbed for such multimodal evaluation, as their identity and behavior are naturally expressed through both visual and acoustic cues. In many cases, fine-grained recognition further depends on complementary information such as visual appearance for morphology and audio for vocalization, making the domain well suited for studying multimodal interaction.

Existing Bird Datasets. SSW60 [30] offers unaligned images and audio as well as video for 60 species, VB100 [9] focuses on fine-grained video recognition for 100 species. Moreover, iNatSounds [2] contains large-scale bioacoustic recordings without visual counterparts, and lastly, BioTrove-Train [43] includes millions of images across diverse species with uneven per-class distributions. Although those datasets provide valuable resources, a joint combination of images, audio, and text in a form suitable for multimodal pretraining and evaluation is missing.

Dataset Overview. To bridge this gap, we construct Bird-MML, a synthetic multimodal bird dataset integrating

Table 1. Zero-shot classification accuracy on AV-MNIST (mean ± std over 5 runs). A = Audio, V = Vision, A+V = Audio–Visual fusion. Best performance is reported in bold.

Method	A	V	A+V
CLIP	**41.1** ± **0.3**	63.0 ± 0.4	-
Tri-CLIP	40.0 ± 0.4	62.5 ± 0.4	-
Symile [27]	-	-	70.9 ± 0.3
GRAM [5]	9.8 ± 0.4	63.9 ± 0.8	64.4 ± 0.5
TRIANGLE [6]	-	-	64.9 ± 0.3
ConFu	39.7 ± 0.1	**64.6** ± **0.3**	**71.2** ± **0.2**

artificially paired images, audio, and text for 150 species, including all those present in SSW60 and VB100. The dataset is designed for scalable pretraining and cross-modal learning, with balanced representation across species. A detailed overview of the data generation pipeline is provided in Appendix D.1. Images are sourced from the iNaturalist open dataset, restricted to research-grade, Creative Commons–licensed observations. For each species, 1,000 samples are randomly selected.

Audio and Text Generation. Audio recordings are collected from Xeno-Canto [39], segmented into 10-second clips, and zero-padded when necessary. In cases of limited availability, existing clips are reused to maintain uniform sample counts. Metadata such as call type, sex, and life stage are preserved. Text descriptions are generated by combining three complementary sources: image captions from InstructBLIP2 [7], audio metadata, and short summaries from corresponding Wikipedia entries. These elements are fused using `google/gemma-2-2b-it` [29] to yield a single multimodal caption per instance.

Dataset Characteristics. While generated captions may contain minor factual noise, they effectively capture species-relevant attributes and contextual cues useful for representation learning. The resulting dataset contains 149,681 triplets (image, audio, text), each grounded by species identity and visual–acoustic correspondence. Approximately 43% of species required limited audio reuse due to data scarcity. Examples of artificially aligned triplets are shown in Appendix D.1.The dataset maintains balanced per-species coverage and is well-suited for multimodal pretraining. Downstream evaluations are performed on naturally paired video-audio data to assess real-world cross-modal performance.

Ethical Considerations. All data were collected from publicly available sources under Creative Commons licenses. Ownership and attribution remain with original contributors. Automatically generated captions may contain minor factual errors or biases and should not be interpreted as verified biological information. The dataset is released solely for non-commercial research use.

5. Experimental Results

Our evaluation covers multimodal benchmarks across biodiversity (SSW60 [30], VB100 [9]), affective (MOSI [45], UR-FUNNY [12], and MUStARD [1] included in MultiBench [19]), and multimedia (AV-MNIST [31]) domains, encompassing both real-world and synthetic datasets.

These datasets present different levels of cross-modal difficulty, enabling holistic evaluation of alignment, representation quality, and scalability. We compare against representative contrastive and higher-order multimodal approaches: (i) a Bi-modal CLIP (referenced as CLIP from now on) trained with standard two-modality contrastive loss, (ii) a Tri-modal CLIP (referenced as Tri-CLIP from now on) trained with pairwise objectives across all modalities, (iii) Symile [27] which models higher-order dependencies via total correlation, (iv) TRIANGLE [6] and GRAM [5], which use geometric similarity measures among 2 or more modalities. Experiments cover (i) zero-shot classification, (ii) 1→1 and 2→1 retrieval, and (iii) linear-probe classification under few-shot and full-data regimes. Details about the experimental setup and evaluation can be found in Appendix A.

5.1. AV-MNIST

We evaluate our framework on the AV-MNIST dataset, which contains paired audio spectrograms augmented with randomly sampled background noise from ESC-50 [25] and images of MNIST digits with degraded visual features. We additionally introduce a text modality derived from class labels using natural language templates such as *'the digit is a {class}"*, *'a photo of a {class}"*, and *'a picture of a {class}"*. As provided in Table 1, our approach achieves the highest performance across all baselines, with Symile trailing closely behind. Fusing audio and visual inputs yields an 8% improvement over the strongest unimodal baseline, demonstrating the benefit of complementary multimodal supervision. Moreover, even when evaluated on a single modality (vision), our model surpasses the corresponding unimodal baseline by 1.5%, indicating that multimodal training enhances individual modality representations.

5.2. Affective Computing Benchmarks

We next evaluate on the MOSI, UR-FUNNY, and MUStARD datasets, which involve text, audio, and video for multimodal affect understanding. We assess (i) 1→1 retrieval, (ii) 2→1 retrieval, and (iii) multimodal classification via a linear probe. Table 2 shows that ConFu consistently improves cross-modal retrieval performance. In particular, the 2→1 retrieval setting surpasses 1→1 retrieval on UR-FUNNY and MUStARD, suggesting that combining two input modalities provides richer and more discriminative cues. Notably, GRAM, the only baseline support-

Table 2. Recall@10 (%) for **Target** across query modalities and datasets (mean ± std over 5 runs). The last two columns show mean retrieval accuracy for 1→1 and 2→1 retrieval, respectively. Best performance is reported in bold, second best is underlined.

	Target	M1			M2			M3			Mean	
	Query	M2	M3	M23	M1	M3	M13	M1	M2	M12	1→1	2→1
MOSI	Trimodal CLIP	**22.9** ± 1.9	**20.5** ± 1.6	–	**23.5** ± 3.6	**23.3** ± 1.8	–	**20.9** ± 2.0	**24.3** ± 2.1	–	**22.6**	–
	Symile [27]	-	-	16.3 ± 0.9	-	-	18.1 ± 2.4	-	-	17.1 ± 1.7	–	17.2
	GRAM [5]	6.0 ± 3.1	10.3 ± 2.9	16.3 ± 2.5	7.7 ± 2.3	0.3 ± 0.6	7.9 ± 3.4	12.2 ± 4.1	0.2 ± 0.2	12.0 ± 4.2	6.1	12.1
	TRIANGLE [6]	-	-	8.3 ± 1.3	-	-	4.9 ± 0.7	-	-	8.8 ± 1.4	–	7.3
	ConFu	21.0 ± 1.7	16.4 ± 2.3	**16.7** ± 1.8	19.2 ± 1.4	21.0 ± 2.4	**21.6** ± 1.9	16.1 ± 1.7	23.5 ± 2.7	**20.5** ± 2.9	19.5	**19.6**
UR-FUNNY	Trimodal CLIP	**3.7** ± 0.6	**4.0** ± 0.6	–	**3.8** ± 0.5	**16.2** ± 1.6	–	**4.0** ± 0.4	**16.8** ± 1.2	–	**8.1**	–
	Symile [27]	-	-	3.7 ± 0.5	-	-	15.4 ± 0.9	-	-	16.5 ± 0.6	–	11.9
	GRAM [5]	3.2 ± 0.5	3.0 ± 0.9	3.9 ± 0.8	3.2 ± 0.9	0.1 ± 0.2	3.1 ± 0.4	3.9 ± 0.4	0.1 ± 0.0	3.3 ± 0.5	2.3	3.4
	TRIANGLE [6]	-	-	**4.2** ± 0.4	-	-	3.3 ± 0.7	-	-	3.6 ± 0.9	–	3.7
	ConFu	3.2 ± 0.3	3.5 ± 0.3	3.6 ± 0.3	3.5 ± 0.2	15.1 ± 0.8	**16.9** ± 0.9	3.5 ± 0.4	15.6 ± 0.8	**20.3** ± 1.1	7.4	**13.6**
MUStARD	Trimodal CLIP	70.7 ± 4.1	29.1 ± 3.6	–	70.5 ± 4.3	26.5 ± 3.2	–	30.4 ± 3.4	24.6 ± 2.8	–	42.0	–
	Symile [27]	-	-	61.5 ± 6.3	-	-	57.0 ± 3.0	-	-	21.3 ± 4.4	–	46.6
	GRAM [5]	61.7 ± 8.0	24.8 ± 5.3	**81.0** ± 4.1	67.2 ± 5.4	24.8 ± 5.3	67.8 ± 5.4	27.9 ± 8.7	5.1 ± 2.1	31.2 ± 9.2	35.2	60.0
	TRIANGLE [6]	-	-	68.7 ± 4.2	-	-	59.9 ± 3.7	-	-	21.9 ± 5.0	–	50.2
	ConFu	**73.8** ± 3.2	**33.6** ± 3.1	79.6 ± 3.6	**74.4** ± 2.7	**28.1** ± 3.2	**74.6** ± 4.1	**33.9** ± 4.2	**28.0** ± 2.9	**33.5** ± 3.2	**45.3**	**62.6**

Table 3. Multimodal classification accuracy (%) (mean ± std over 5 runs). Tri-CLIP refers to CLIP trained with pairwise contrastive losses. Best performance is reported in bold.

Method	MOSI	UR-FUNNY	MUStARD
Tri-CLIP	63.5 ± 2.1	64.0 ± 0.9	62.1 ± 3.8
Symile [27]	**67.5** ± 1.3	64.7 ± 1.0	60.5 ± 5.0
GRAM [5]	65.5 ± 2.5	64.8 ± 0.9	**64.7** ± 2.3
TRIANGLE [6]	65.0 ± 2.8	64.5 ± 0.4	64.6 ± 2.5
ConFu	66.7 ± 2.1	**64.9** ± 1.0	64.1 ± 2.5

ing both retrieval configurations, struggles with 1→1 alignments in general and fails for modalities M2 and M3 on MOSI and UR-FUNNY, suggesting that balancing 2→1 and 1→1 alignment is non-trivial. Furthermore, methods such as TRIANGLE and Symile, which focus on trimodal alignment without explicitly modeling pairwise relations, perform consistently worse than ConFu. Overall, ConFu ranks first or second across nearly all query→target retrieval settings, in both 1→1 and 2→1 configurations, demonstrating robust and consistent cross-modal alignment. Moreover, Table 3 reports linear-probe classification results, showing that ConFu attains the highest accuracy on UR-FUNNY and remains competitive on MUStARD and MOSI.

5.3. Fine-Grained Bird Classification

ConFu is further evaluated on fine-grained bird classification using the VB100 and SSW60 datasets. Models are pretrained on Bird-MML, our curated dataset of 150K audio–image–text triplets and assessed under zero-shot and few-shot regimes. We consider two evaluation settings: (1) the multi-frame protocol from [30], where each video clip is divided into eight segments, the middle frame of each segment is sampled, and the mean embedding of the eight frames is used as the visual representation; and (2) a single-frame setting, providing a more constrained visual input per video. Table 4 reports zero-shot results in the multi-frame setting. ConFu achieves the highest accuracy on SSW60 (71.4%), confirming that multimodal fusion improves fine-grained discrimination when both modalities are informative. On VB100, where the audio modality is largely uninformative (∼4% unimodal accuracy), the fused model (A+V) maintains performance close to the best visual-only baseline (18.15% vs. 20.69%), demonstrating robustness to distracting modalities. In contrast, GRAM, Symile, and TRIANGLE exhibit a notable decline in accuracy (5–10%) relative to ConFu. Few-shot trends in Fig. 4 mirror these findings. Fusion offers clear advantages on SSW60 and stable performance on VB100, showing that ConFu generalizes well even when one modality is weak.

Table 4. Zero-shot classification accuracy (%) on SSW60 and VB100 datasets. A = Audio, V = Vision, A+V = Audio-Visual fusion. Best performance is reported in bold.

	SSW60 Acc. (%)			VB100 Acc. (%)		
Method	A	V	A+V	A	V	A+V
CLIP	29.9	70.1	–	4.2	20.6	–
Tri-CLIP	31.1	69.0	–	3.9	**20.7**	–
Symile [27]	-	-	60.2	-	-	13.4
TRIANGLE [6]	-	-	64.1	-	-	12.1
GRAM [5]	0.7	66.6	56.9	1.3	13.7	8.0
ConFu	30.3	69.4	**71.4**	3.4	19.3	18.1

In the single-frame setting, designed to evaluate performance with limited visual information, we averaged classification results across 19 uniformly sampled frames to

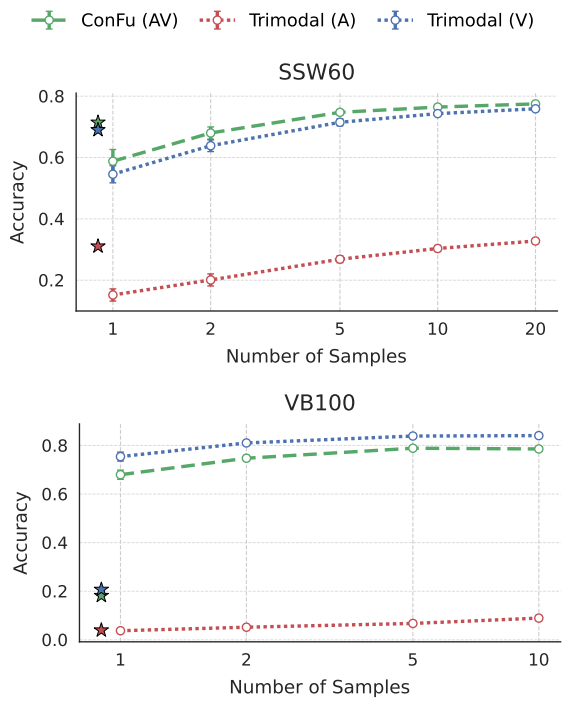

Figure 4. Few-shot linear probing results on the SSW60 (top) and VB100 (bottom) datasets. Performance is shown as the number of labeled examples increases. Zero-shot performance is indicated with a star. Prediction is done in the 8-frame average embedding setting.

Table 5. Average zero-shot classification accuracy (%) on SSW60 and VB100 datasets (mean ± std over 19 uniformly sampled frames). The full 10-second audio is used; therefore, its accuracy remains constant across samples, and no std is reported. A = Audio, V = Vision, A+V = Audio–Visual fusion. Best performance is reported in bold.

	SSW60 Acc. (%)			VB100 Acc. (%)		
Method	**A**	**V**	**A+V**	**A**	**V**	**A+V**
CLIP	29.9	60.3 ± 3.4	–	4.2	**18.2 ± 0.5**	–
Tri-CLIP	31.1	59.7 ± 3.2	–	3.9	17.7 ± 0.5	–
Symile [27]	-	-	61.4 ± 0.2	-	-	13.2 ± 0.4
TRIANGLE [6]	-	-	58.6 ± 3.3	-	-	8.7 ± 0.4
GRAM [5]	0.7	59.0 ± 3.3	54.9 ± 3.2	1.3	12.9 ± 0.4	7.6 ± 0.3
ConFu	30.3	59.8 ± 3.3	**65.5 ± 2.5**	3.4	16.7 ± 0.5	16.4 ± 0.4

reduce bias. As provided in Table 5, the fused (A+V) model exceeds the best baseline by roughly 4% on SSW60 and complementary audio–visual cues when visual input is sparse. On VB100, performance drops slightly (by 2–3%) due to the weaker audio modality ($\sim 4\%$ accuracy). Overall, ConFu benefits from richer multimodal supervision, maintaining strong robustness under modality degradation while preserving scalability and simplicity. Additional results on a few-shot adaptation are in the Appendix C.1.

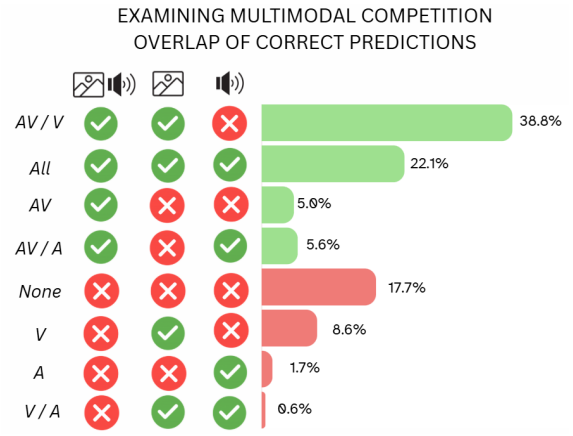

Figure 5. Aggregated modality overlap proportions across all classes in the SSW60 dataset. The chart shows the overall fraction of correctly predicted samples belonging to each overlap category: *Audiovisual Only (AV)*, *Vision Only (V)*, *Audio Only (A)*, *Audiovisual & Vision (AV / V)*, *Audiovisual & Audio (AV / A)*, *Vision & Audio (V / A)*, *All*, and *None*.

6. Ablation

Modality competition. Although multimodal learning seeks to combine complementary cues, models often show modality dominance, where stronger signals suppress weaker ones [14]. To assess how our fusion is affected by this effect, we analyze in Fig. 5 the overlap of prediction errors across modalities and fusion. Results reveal visual dominance: The larger portion of correct samples (39%) are predicted by both the fusion and vision models, while the datapoints that only audio predicts correctly or fusion-audio overlap remain small (5–6%), and purely audio-driven successes are rare (2%). Nevertheless, about 5% of samples are uniquely solved by the audiovisual model, indicating that fusion captures complementary information beyond either modality alone. Finally, some classes benefit substantially from fusion (e.g., Label 31 see Appendix B.5), while others remain uniformly difficult.

Ablation on Noise-Induced Distribution Shift. To examine the sensitivity of each model to distribution shift, we introduced Gaussian noise to the image modalitiy only during testing, with a standard deviation of 0.05 corresponding to roughly 20dB SNR. This perturbation simulates a mismatch between the training and test distributions without retraining the models. As shown in Table 6, ConFu achieves the highest accuracy (45.4%) under these conditions. In contrast, the unimodal baseline (Tri-CLIP), exhibits a sharper performance drop, while Symile shows moderate robustness. Interestingly, TRIANGLE and GRAM perform even worse compared to Tri-CLIP under visual degradation. Overall, these results indicate that

Table 6. Accuracy (%) under Gaussian noise-induced distribution shift (mean ± std over 5 runs). Noise was applied to the **Vision** modality (SNR $\sim 20dB$) at test time. A = Audio, V = Vision, A+V = Audio–Visual fusion. Best performance is reported in bold.

Method	Mod.	V noise
Tri-CLIP (V)	V	35.4 ± 0.4
Symile [27]	A+V	40.9 ± 0.3
GRAM [5]	A+V	25.9 ± 0.1
TRIANGLE [6]	A+V	26.5 ± 0.2
ConFu	A+V	**45.4** ± **0.1**

ConFu's fusion strategy improves stability under distribution shifts affecting the image modality. Regarding more experiments in noise induced distribution shifts see Appendix B.4.

7. Discussion

Empirical results across synthetic and real-world datasets demonstrate that ConFu effectively captures complementary cross-modal structure. On the synthetic XOR benchmark, ConFu successfully models the underlying multimodal dependencies that pairwise and geometric baselines miss even with higher-dimensional representations (Appendix B.2). In AV-MNIST, ConFu dominates unimodal approaches in zero-shot classification, indicating successful utilization of both modalities. The 2→1 retrieval evaluations on UR-FUNNY, and MUStARD further confirm that combining two query modalities provides richer semantic grounding than either considered in isolation. An interesting observation in the 2→1 case on the MOSI dataset is that fusion slightly underperforms compared to the pairwise baseline. Remarkably, this issue is substantially alleviated by introducing partial modality masking (Appendix B.3), supporting the hypothesis that suboptimal performance stems from shortcut learning in cases where redundant information dominates.

In fine-grained bird classification, ConFu maintains or slightly exceeds unimodal baselines. On SSW60, where both modalities contribute meaningfully, fusion achieves the best accuracy, on VB100, where acoustic information is limited, performance remains close to the visual baseline. Moreover, modality overlap analyses indicates that a subset of samples is correctly resolved only when both modalities are utilized, reflecting the model's ability to integrate complementary evidence. Finally, under visual perturbations, ConFu degrades more gracefully, outperforming all baselines by a clear margin. SYMILE ranks second, remaining above unimodal performance, while TRIANGLE and GRAM drop below the single-modality baseline.

Limitations and Future Work. ConFu is well-positioned to scale to any number of modalities, offering flexibility for diverse multimodal tasks. However, as the number of modalities increases, computational demands may also rise. Depending on the task, selectively pruning certain alignment terms, shown in Appendix E, can help maintain efficiency. For example, omitting many-to-many alignments, when assuming that only single-modality queries or targets will be needed, opens a promising avenue for future adaptation and research.

ConFu currently relies on fully aligned modalities during training, and exploring ways to relax or bypass this requirement remains crucial for many applications. In particular, enabling ConFu to handle partially missing modalities, such as training with incomplete triplets or drawing on ideas related to pseudopair construction [34, 36, 46], could potentially make the framework more robust and broadly applicable across real-world multimodal settings. Further prospective directions could include extending ConFu to incorporate factorized representations inspired by recent factorization-based and redundancy-reduction approaches [8, 20, 33], developing adaptive mechanisms to regulate modality competition [13, 17, 18, 24, 38], and investigating how ConFu can enhance robustness in multimodal retrieval, particularly when one or more modalities are noisy, incomplete, or severely corrupted.

8. Conclusion

We presented ConFu, a unified contrastive framework that jointly optimizes pairwise and higher-order alignments to capture pairwise and synergistic relationships across modalities. The method supports both 1→1 and 2→1 retrieval, improving multimodal alignment and robustness without modifying encoder architectures. Overall, ConFu demonstrates strong adaptability across diverse datasets, domains, and challenging conditions such as noise-induced distribution shifts, modality degradation, and missing modalities, while introducing only minimal overhead through lightweight MLPs atop the encoders, establishing it as a robust and scalable solution for multimodal alignment in real-world applications. The accompanying Bird-MML dataset establishes a controlled setup for studying multimodal complementarity under realistic audio–visual conditions.

9. Acknowledgments

This work was supported by the HARSH project (project no. ΥΠ3ΤΑ - 0560901) within the framework of the National Recovery and Resilience Plan–Greece 2.0–with funding from the European Union–NextGenerationEU; the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the "3rd Call for H.F.R.I.'s Research Projects to Support Faculty Members & Researchers" (H.F.R.I. Project Number: 26302); by the TITAN ERA Chair project (contract no. 101086741) within the Horizon Europe Framework Program of the European Commission; the METHUSALEM program (METH/26/003, "Methusalem-BioMedAI"); the Flemish Government AI Research Program (Leuven.AI – KU Leuven); the FWO project "Task- and Device-Agnostic Artificial Intelligence for EEG Analysis" (G046925N); and the Horizon Europe project "AI-PROGNOSIS" (Grant Agreement No. 101080581).

References

[1] Santiago Castro, Devamanyu Hazarika, Verónica Pérez-Rosas, Roger Zimmermann, Rada Mihalcea, and Soujanya Poria. Towards multimodal sarcasm detection (an _obviously_ perfect paper). *arXiv preprint arXiv:1906.01815*, 2019. 5

[2] Mustafa Chasmai, Alexander Shepard, Subhransu Maji, and Grant Van Horn. The inaturalist sounds dataset. *Advances in Neural Information Processing Systems*, 37:132524–132544, 2024. 4

[3] Sihan Chen, Xingjian He, Longteng Guo, Xinxin Zhu, Weining Wang, Jinhui Tang, and Jing Liu. Valor: Vision-audio-language omni-perception pretraining model and dataset. *arXiv preprint arXiv:2304.08345*, 2023. 2

[4] Sihan Chen, Handong Li, Qunbo Wang, Zijia Zhao, Mingzhen Sun, Xinxin Zhu, and Jing Liu. Vast: A vision-audio-subtitle-text omni-modality foundation model and dataset. *Advances in Neural Information Processing Systems*, 36:72842–72866, 2023. 2

[5] Giordano Cicchetti, Eleonora Grassucci, Luigi Sigillo, and Danilo Comminiello. Gramian multimodal representation learning and alignment. *arXiv preprint arXiv:2412.11959*, 2024. 2, 5, 6, 7, 8, 3

[6] Giordano Cicchetti, Eleonora Grassucci, and Danilo Comminiello. A triangle enables multimodal alignment beyond cosine similarity. *arXiv preprint arXiv:2509.24734*, 2025. 2, 5, 6, 7, 8, 3

[7] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *Advances in neural information processing systems*, 36:49250–49267, 2023. 5

[8] Benoit Dufumier, Javiera Castillo-Navarro, Devis Tuia, and Jean-Philippe Thiran. What to align in multimodal contrastive learning? *arXiv preprint arXiv:2409.07402*, 2024. 8

[9] ZongYuan Ge, Chris McCool, Conrad Sanderson, Peng Wang, Lingqiao Liu, Ian Reid, and Peter Corke. Exploiting temporal information for dcnn-based fine-grained object classification. In *2016 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–6. IEEE, 2016. 4, 5, 6, 7

[10] Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15180–15190, 2023. 1, 2

[11] Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel. Audioclip: Extending clip to image, text and audio. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 976–980. IEEE, 2022. 1, 2

[12] Md Kamrul Hasan, Wasifur Rahman, Amir Zadeh, Jianyuan Zhong, Md Iftekhar Tanveer, Louis-Philippe Morency, et al. Ur-funny: A multimodal language dataset for understanding humor. *arXiv preprint arXiv:1904.06618*, 2019. 5

[13] Cong Hua, Qianqian Xu, Shilong Bao, Zhiyong Yang, and Qingming Huang. Reconboost: boosting can achieve modality reconcilement. In *Proceedings of the 41st International Conference on Machine Learning*. JMLR.org, 2024. 8

[14] Yu Huang, Junyang Lin, Chang Zhou, Hongxia Yang, and Longbo Huang. Modality competition: What makes joint training of multi-modal network fail in deep learning?(provably). In *International Conference on Machine Learning*, pages 9226–9259. PMLR, 2022. 1, 7

[15] Zoran Jackson. Free spoken digit dataset. https://github.com/Jakobovski/free-spoken-digit-dataset, 2017. Accessed: YYYY-MM-DD. 6

[16] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. 1

[17] Konstantinos Kontras, Thomas Strypsteen, Christos Chatzichristos, Paul Pu Liang, Matthew B Blaschko, and Maarten De Vos. Balancing multimodal training through game-theoretic regularization. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*. 8

[18] Konstantinos Kontras, Christos Chatzichristos, Matthew B. Blaschko, and Maarten De Vos. Improving multimodal learning with multi-loss gradient modulation. In *35th British Machine Vision Conference 2024, BMVC 2024, Glasgow, UK, November 25-28, 2024*. BMVA, 2024. 8

[19] Paul Pu Liang, Yiwei Lyu, Xiang Fan, Zetian Wu, Yun Cheng, Jason Wu, Leslie Yufan Chen, Peter Wu, Michelle A Lee, Yuke Zhu, et al. Multibench: Multiscale benchmarks for multimodal representation learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. 5, 6

[20] Paul Pu Liang, Zihao Deng, Martin Q Ma, James Y Zou, Louis-Philippe Morency, and Ruslan Salakhutdinov. Factor-

ized contrastive learning: Going beyond multi-view redundancy. *Advances in Neural Information Processing Systems*, 36:32971–32998, 2023. 8, 1

[21] Bryan Lim, Sercan Ö Arık, Nicolas Loeff, and Tomas Pfister. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International journal of forecasting*, 37(4):1748–1764, 2021. 2

[22] Ye Liu, Siyuan Li, Yang Wu, Chang-Wen Chen, Ying Shan, and Xiaohu Qie. Umt: Unified multi-modal transformers for joint video moment retrieval and highlight detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3042–3051, 2022. 2

[23] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 3, 4

[24] Xiaokang Peng, Yake Wei, Andong Deng, Dong Wang, and Di Hu. Balanced multimodal learning via on-the-fly gradient modulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8238–8247, 2022. 8

[25] Karol J. Piczak. ESC: Dataset for Environmental Sound Classification. In *Proceedings of the 23rd Annual ACM Conference on Multimedia*, pages 1015–1018. ACM Press. 5, 6

[26] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 1

[27] Adriel Saporta, Aahlad Manas Puli, Mark Goldstein, and Rajesh Ranganath. Contrasting with symile: Simple model-agnostic representation learning for unlimited modalities. *Advances in Neural Information Processing Systems*, 37: 56919–56957, 2024. 2, 5, 6, 7, 8, 3

[28] Milan Studenỳ and Jirina Vejnarová. The multiinformation function as a tool for measuring stochastic dependence. In *Learning in graphical models*, pages 261–297. Springer, 1998. 3

[29] Gemma Team. Gemma. 2024. 5

[30] Grant Van Horn, Rui Qian, Kimberly Wilber, Hartwig Adam, Oisin Mac Aodha, and Serge Belongie. Exploring fine-grained audiovisual categorization with the ssw60 dataset. In *European Conference on Computer Vision*, pages 271–289. Springer, 2022. 4, 5, 6

[31] Valentin Vielzeuf, Alexis Lechervy, Stéphane Pateux, and Frédéric Jurie. Centralnet: a multilayer approach for multimodal fusion. In *Proceedings of the European conference on computer vision (ECCV) workshops*, pages 0–0, 2018. 5

[32] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 4, 7

[33] Yi Wang, Conrad M Albrecht, Nassim Ait Ali Braham, Chenying Liu, Zhitong Xiong, and Xiao Xiang Zhu. Decoupling common and unique representations for multimodal self-supervised learning. In *European Conference on Computer Vision*, pages 286–303. Springer, 2024. 8

[34] Zehan Wang, Yang Zhao, Haifeng Huang, Jiageng Liu, Aoxiong Yin, Li Tang, Linjun Li, Yongqi Wang, Ziang Zhang, and Zhou Zhao. Connecting multi-modal contrastive representations. *Advances in Neural Information Processing Systems*, 36:22099–22114, 2023. 1, 8

[35] Zehan Wang, Ziang Zhang, Xize Cheng, Rongjie Huang, Luping Liu, Zhenhui Ye, Haifeng Huang, Yang Zhao, Tao Jin, Peng Gao, et al. Freebind: Free lunch in unified multimodal space via knowledge fusion. *arXiv preprint arXiv:2405.04883*, 2024. 2

[36] Zehan Wang, Ziang Zhang, Hang Zhang, Luping Liu, Rongjie Huang, Xize Cheng, Hengshuang Zhao, and Zhou Zhao. Omnibind: Large-scale omni multimodal representation via binding spaces. *arXiv preprint arXiv:2407.11895*, 2024. 1, 2, 8

[37] Satosi Watanabe. Information theoretical analysis of multivariate correlation. *IBM Journal of research and development*, 4(1):66–82, 1960. 3

[38] Yake Wei, Siwei Li, Ruoxuan Feng, and Di Hu. Diagnosing and re-learning for balanced multimodal learning. In *European Conference on Computer Vision*, pages 71–86. Springer, 2024. 8

[39] Xeno-canto Foundation for Nature Sounds. Xeno-canto: Sharing wildlife sounds from around the world. https://www.xeno-canto.org/, 2008. 5

[40] Haiyang Xu, Qinghao Ye, Ming Yan, Yaya Shi, Jiabo Ye, Yuanhong Xu, Chenliang Li, Bin Bi, Qi Qian, Wei Wang, et al. mplug-2: A modularized multi-modal foundation model across text, image and video. In *International Conference on Machine Learning*, pages 38728–38748. PMLR, 2023. 2

[41] Le Xue, Mingfei Gao, Chen Xing, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, and Silvio Savarese. Ulip: Learning a unified representation of language, images, and point clouds for 3d understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1179–1189, 2023. 1, 2

[42] Le Xue, Ning Yu, Shu Zhang, Artemis Panagopoulou, Junnan Li, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, et al. Ulip-2: Towards scalable multimodal pre-training for 3d understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27091–27101, 2024. 1, 2

[43] Chih-Hsuan Yang, Benjamin Feuer, Talukder Jubery, Zi Deng, Andre Nakkab, Md Zahid Hasan, Shivani Chiranjeevi, Kelly Marshall, Nirmal Baishnab, Asheesh Singh, et al. Biotrove: A large curated image dataset enabling ai for biodiversity. *Advances in Neural Information Processing Systems*, 37:102101–102120, 2024. 4

[44] 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

[45] Amir Zadeh, Rowan Zellers, Eli Pincus, and Louis-Philippe Morency. Mosi: multimodal corpus of sentiment intensity and subjectivity analysis in online opinion videos. *arXiv preprint arXiv:1606.06259*, 2016. 5

[46] Ziang Zhang, Zehan Wang, Luping Liu, Rongjie Huang, Xize Cheng, Zhenhui Ye, Huadai Liu, Haifeng Huang, Yang Zhao, Tao Jin, et al. Extending multi-modal contrastive rep-

resentations. *Advances in Neural Information Processing Systems*, 37:91880–91903, 2024. 1, 2, 8

[47] Junsheng Zhou, Jinsheng Wang, Baorui Ma, Yu-Shen Liu, Tiejun Huang, and Xinlong Wang. Uni3d: Exploring unified 3d representation at scale. *arXiv preprint arXiv:2310.06773*, 2023. 2

[48] Bin Zhu, Bin Lin, Munan Ning, Yang Yan, Jiaxi Cui, HongFa Wang, Yatian Pang, Wenhao Jiang, Junwu Zhang, Zongwei Li, et al. Languagebind: Extending video-language pretraining to n-modality by language-based semantic alignment. *arXiv preprint arXiv:2310.01852*, 2023. 1, 2

THE MORE, THE MERRIER: CONTRASTIVE FUSION FOR HIGHER-ORDER MULTIMODAL ALIGNMENT

Supplementary Material

Supplementary Material

This supplementary document provides additional details and results that complement the main paper. It is organized as follows:

- **Experimental Details.** We present comprehensive descriptions of model configurations, hyperparameters, training procedures, and evaluation protocols to ensure reproducibility.
- **Additional Ablation Studies.** We include supplementary ablation experiments that highlight additional aspects of this work, examining how different model components and design choices influence overall performance.
- **Additional Results.** This section includes further quantitative results that extend and support the findings presented in the main paper.
- **Datasets.** Additional information regarding the datasets used and our new dataset Bird-MML.

A. Experimental Details

A.1. Experiments on Bird-MML

Dataset preprocessing For the audio modality, Each audio clip was truncated or zero-padded to a duration of 10 seconds and resampled to 22.05 kHz. Mel spectrograms were computed using a window size of 1024 and a hop length of 512, resulting in 128 Mel frequency bands. Each spectrogram was treated as a single-channel image to ensure compatibility with a ResNet-50 backbone.

For the image modality, Images were resized to 224×224 pixels and normalized. No data augmentation was applied.

Model Architecture. The multimodal architecture employed is composed of three modality-specific encoders:

- **Text encoder:** BERT transformer.
- **Image encoder:** ResNet-50.
- **Audio encoder:** ResNet-50 applied to Mel-spectrogram inputs.

Each encoder outputs a 512-dimensional embedding. A two-layer multilayer perceptron (MLP) was used for multimodal feature fusion. All models and baselines were trained using identical hyperparameters, without any additional hyperparameter search or tuning.

Training used the AdamW optimizer with a learning rate of 1×10^{-4} and a weight decay of 1×10^{-4}. A batch size of 64 was used throughout the training process. The learning rate followed a cosine annealing schedule with a minimum of 1×10^{-6}. The fusion loss coefficient was set to $\lambda = 0.5$.

A.2. MOSI, MUStARD, and UR-FUNNY

For these datasets, all models were trained for 100 epochs, producing 256-dimensional representations. We used a batch size of 128 and a learning rate of 5×10^{-5}. Following prior work [20], each modality encoder was implemented as a 5-layer Transformer with 5 attention heads.

Fusion was achieved through a 2-layer MLP that combined unimodal embeddings into a single fused representation. Dataset splits followed the official `MultiBench` configuration for training, validation, and testing.

A.3. AVMNIST

For the text modality, we used a BERT transformer. Image, spectrogram, and satellite-view inputs were encoded using ResNet-18 backbones initialized from ImageNet and fully trained. Spectrograms provided by the `MultiBench` dataloaders were used without modification.

Embeddings had a 256-dimensional size. Dataset splits again followed the `MultiBench` protocol.

Training Setup. All models were trained from scratch for 30 epochs using the AdamW optimizer with a learning rate of 1×10^{-4} and a batch size of 256. The best model was selected based on validation loss. Each experiment was repeated with five random seeds, and we report the mean and standard deviation across runs. For our model, MLP-based fusion was used with a balancing parameter $\lambda = 0.5$ between pairwise and fusion losses.

A.4. XOR Experiment

For the XOR experiment, we mapped all modalities to embeddings of size 128. The dataset consisted of 10,000 training samples and 5,000 test samples. Models were trained for 50 epochs with a batch size of 512 and a learning rate of 1×10^{-4}.

Each encoder was implemented as a 2-layer MLP, and we used a balanced fusion coefficient of $\lambda = 0.5$.

B. Additional Ablation Studies

B.1. Effect of the λ Parameter

As shown in Table 7, performance is not overly sensitive to λ. Accuracy remains stable while recall varies smoothly,

indicating a controlled trade-off with minimal tuning required. Identifying a principled way to set or eliminate λ is an important direction.

Table 7. Effect of λ across datasets (mean $\pm$ std). Accuracy (%) remains stable while λ trades off recall. Best values in bold.

Dataset	λ	Acc. (%)	$R^{1\to1}_{mean@10}$	$R^{2\to1}_{mean@10}$
MOSI	0.2	**67.7** $\pm$ 1.7	19.5 $\pm$ 1.0	**20.1** $\pm$ 1.9
	0.5	66.7 $\pm$ 2.1	**19.5** $\pm$ 1.5	19.6 $\pm$ 2.0
	0.8	65.8 $\pm$ 2.0	13.5 $\pm$ 2.2	15.6 $\pm$ 1.6
UR-Funny	0.2	64.2 $\pm$ 0.7	**8.0** $\pm$ 0.4	**14.0** $\pm$ 0.9
	0.5	**64.9** $\pm$ 1.0	7.4 $\pm$ 0.2	13.6 $\pm$ 0.6
	0.8	64.6 $\pm$ 0.7	6.9 $\pm$ 0.3	13.2 $\pm$ 0.7
MUSTARD	0.2	59.7 $\pm$ 3.2	42.1 $\pm$ 2.2	57.1 $\pm$ 2.5
	0.5	**64.1** $\pm$ 2.5	**45.3** $\pm$ 2.6	**62.6** $\pm$ 3.1
	0.8	61.1 $\pm$ 3.4	41.0 $\pm$ 1.0	59.7 $\pm$ 1.1

Regarding optimal λ values the results in Table 8 highlight the influence of the λ parameter, which balances the contribution of modality-specific and shared representations. These λ values were selected from the Pareto front obtained by jointly optimizing the two retrieval objectives - mean Recall@10 for 1→1 and 2→1 tasks. We observe that moderate λ values (0.1–0.5) yield the best trade-off between the two objectives, and that the optimal λ varies across datasets, indicating optimal λ depends on data distribution.

Table 8. Mean Recall@10 results for 1→1 and 2→1 tasks across datasets.

Dataset	λ	**Mean R@10 (1→1)**	**Mean R@10 (2→1)**
MOSI	0.1	21.21	22.12
UR-FUNNY	0.4	8.06	14.69
MUSTARD	0.5	45.29	62.56

B.2. Embedding Dimensionality Analysis on the XOR Task

In this experiment, we evaluate how the embedding dimensionality affects each method's ability to solve the XOR problem for the case of $\hat{p} = 1$. Figure 6 illustrates the mean accuracy as a function of the embedding dimension. As provided, SYMILE successfully learns to solve the XOR task even with a minimal embedding dimensionality of 8, achieving perfect accuracy thereafter. Our proposed method (ConFu) also manages to solve XOR but requires a larger embedding dimensionality of 64 to reach convergence. In contrast, both TRIANGLE and GRAM fail to solve the task even when the embedding dimensionality is increased up to 1024, implying that the bottleneck lies in their loss formulations rather than in representational capacity.

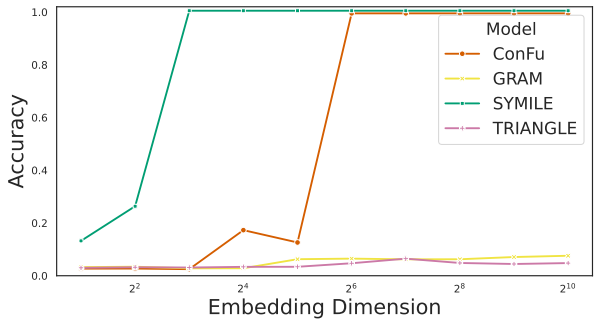

Figure 6. Accuracy as a function of embedding dimension for $\hat{p} = 1$. SYMILE achieves perfect accuracy from dimension 8 onwards, while ConFu requires a dimensionality of 64 to converge. TRIANGLE and GRAM fail to solve the XOR task even at 1024 dimensions.

B.3. Tackling modality shortcuts

Motivated by the hypothesis that the fusion network may exploit shortcut strategies, we investigated the effect of applying dropout at the feature level, specifically, before concatenation and subsequent processing by the fusion MLP. A potential shortcut can arise when the pairwise alignment losses between modalities become sufficiently low, allowing the fusion module to rely predominantly on one modality while suppressing the contribution of the other. In such cases, the network may learn to transmit information through the most predictive modality instead of developing genuinely integrated multimodal representations. Introducing feature-level masking may counteract this behavior by preventing the model from depending exclusively on any single modality and encouraging more balanced multimodal learning.

To examine this effect, we compared model performance with and without masking for multiple λ values and mask ratios. For each dataset and task, we computed the performance difference $\Delta = (\text{With mask}) - (\text{No mask})$ across runs and visualized these distributions using boxplots.

As shown in Figure 7, the MOSI dataset consistently benefits from masking across all tasks, classification as well as 1→1 and 2→1 retrieval, indicating that feature-level dropout can improve representation robustness and reduce reliance on modality-specific cues. In contrast, the UR-FUNNY dataset exhibits minimal changes, suggesting that masking has limited impact on its multimodal dynamics, while MUStARD experiences performance degradation, likely due to its stronger dependence on precise cross-modal alignment. Moreover, as shown in Table 9, the use of masking in ConFu yields significant performance gains across all query→target retrieval directions on the MOSI dataset. Overall, these findings suggest that the impact of masking is dataset-dependent: it tends to improve perfor-

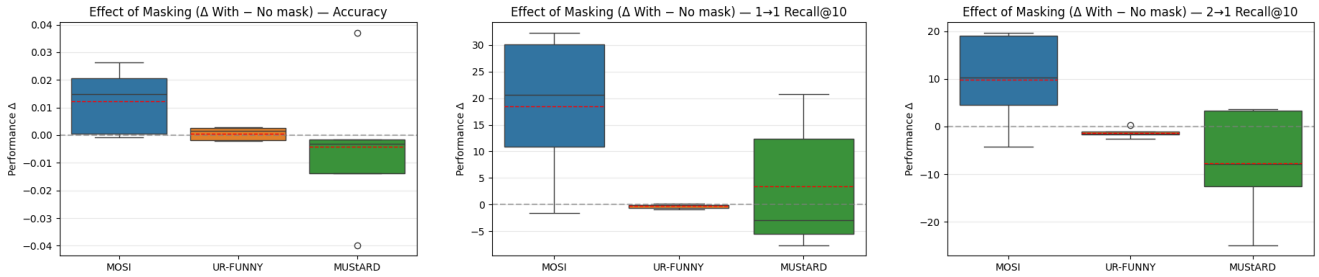

(a) Performance difference (Δ With - No mask) for the **Accuracy** task. Each box shows variability across λ values for different datasets.

(b) Performance difference for **1→1 retrieval (Recall@10)**. Positive values indicate that masking improves unimodal retrieval.

(c) Performance difference for **2→1 retrieval (Recall@10)**. Each box aggregates runs with different λ values and mask ratios.

Figure 7. Effect of the masking ratio on multimodal performance for the MLP fusion model. Each subplot reports the performance change (Δ = With mask − No mask) across datasets and tasks, computed over multiple λ and mask ratio values. These results illustrate how both hyperparameters jointly influence model behavior for classification and retrieval tasks.

mance when redundant cross-modal information could otherwise enable shortcut learning, but may hinder it when successful alignment relies on using the full representational capacity of each modality.

Table 9. Recall@10 (%) on MOSI with and without masking (mask ratio = 0.3, $\lambda = 0.5$). Δ indicates the absolute improvement.

Target	Query(s)	Without Mask	With Mask	Δ
M1	**M2**	21.02 ± 1.65	**25.10** ± **3.51**	+4.08
	M3	16.41 ± 2.32	**20.41** ± **1.51**	+4.00
	M23	16.73 ± 1.84	**22.24** ± **3.16**	+5.51
M2	**M1**	19.18 ± 1.41	**25.83** ± **2.98**	+6.65
	M3	21.02 ± 2.40	**22.77** ± **0.96**	+1.75
	M13	21.63 ± 1.89	**24.58** ± **2.33**	+2.95
M3	**M1**	16.06 ± 1.74	**23.35** ± **1.61**	+7.29
	M2	23.47 ± 2.69	**24.40** ± **1.91**	+0.93
	M12	20.50 ± 2.87	**24.34** ± **3.02**	+3.84

B.4. Noise experiments

We investigate the robustness of ConFu under controlled noise-induced distribution shifts applied to either the image or audio modality. Gaussian noise was added at varying severities with standard deviations of 0.05, 0.1, and 0.15 for images, and 0.1, 0.2, and 0.3 for audio. The corresponding SNR values are estimated from the data.

The results in Table 10 show that ConFu exhibits improved robustness to noise compared to competing baselines. Notably, GRAM and TRIANGLE appear largely unaffected when noise is added to audio. However, this behavior is likely attributable to their effective disregard of the audio modality. In contrast, both methods degrade substantially when noise is introduced to the image modality, performing even worse than unimodal TriCLIP.

Table 10. Accuracy (%) under modality-specific noise-induced distribution shift. Noise was applied to individual modalities at test time with varying SNR levels. A = Audio, V = Vision, A+V = Audio–Visual fusion. Best performance for each degradation type and level is reported in bold.

	10dB SNR		15dB SNR		20dB SNR	
Method	**A deg.**	**V deg.**	**A deg.**	**V deg.**	**A deg.**	**V deg.**
Tri-CLIP (V)	69.0	5.5	69.0	13.2	69.0	35.4
Tri-CLIP (A)	26.0	**31.1**	29.1	31.1	30.3	31.1
Symile [27]	58.4	21.2	59.5	27.9	60.0	40.9
GRAM [5]	58.9	4.0	58.3	8.94	58.1	25.9
TRIANGLE [6]	63.9	3.4	64.0	7.24	64.0	26.5
ConFu	**71.2**	30.2	**71.4**	**33.1**	**71.5**	**45.4**

Overall, ConFu consistently outperforms all baselines across noise levels, including unimodal variants that do not suffer degradation. Its performance trails unimodal TriCLIP-audio slightly only in the extreme case where the visual modality becomes heavily corrupted. A key observation is that Symile, GRAM, and TRIANGLE deteriorate markedly under noise, whereas ConFu incurs only minimal performance loss, particularly when noise is added to audio, where the accuracy drop is negligible. For comparison, TriCLIP experiences roughly a 5% accuracy decline at an audio SNR of 10 dB, while ConFu drops by only ∼0.2%.

B.5. Discussion on Multimodal Competition

In Fig. 8, we present a detailed per-class modality overlap analysis for the SSW60 dataset in the zero-shot classification setting. This experiment visualizes, in the form of a heatmap, the distribution of overlap categories for each label separately. Specifically, each cell indicates the percentage of samples belonging to a given class that fall into one of the following overlap categories: *Audiovisual Only*, *Vision Only*, *Audio Only*, *Audiovisual & Vision*, *Audiovisual & Audio*, *Vision & Audio*, *All*, or *None*. These categories

capture whether a sample is correctly classified by only one modality, by a specific combination of modalities, by all modalities, or not correctly classified at all.

Higher values in a cell correspond to a larger proportion of samples for which that modality (or modalities) yields a correct prediction. This visualization allows us to inspect class-specific modality behavior, highlighting both complementary and competing contributions across modalities.

We observe that label 31 is exclusively correctly classified by the *Audiovisual Only* modality, suggesting that neither vision nor audio alone provides sufficient information to recognize this class. This reinforces the need for complementary cross-modal information in order to correctly classify certain bird species.

C. Additional Results

C.1. Few shot adaptation on SSW60, VB100, CUB200 complete results

In this section, we present the complete results of our few-shot adaptation experiments on the SSW60, VB100, and CUB200 [32] datasets. The following figures include all baseline methods for comparison and illustrate adaptation performance across varying number of shot settings. These results provide a comprehensive view of how each method generalizes to limited-data scenarios and adapts to novel classes.

SSW60. In the multi-frame setting, where mean embeddings over 8 video frames are used, our fusion-based method achieves the best performance. It is important to note that competing approaches do not support multimodal representations, and therefore only unimodal results are available for comparison. In the vision modality, most methods perform similarly, with Symile showing a slight drop. In contrast, in the audio modality, TRIANGLE and GRAM underperform significantly, indicating that their generated audio representations are not informative. In the single-frame setting, **ConFu**'s audiovisual representations outperform all unimodal variants by a substantial margin.

VB100. In VB100, audiovisual fusion performs worse than vision-only models. This aligns with our earlier discussion: the audio modality is largely uninformative and even distracting for this dataset, as reflected by the low performance of audio-only baselines.

CUB200. For CUB200, we evaluate on classes completely unseen during pretraining. Symile achieves the strongest performance, while our method (**ConFu**) ranks second.

D. Datasets

D.1. Bird-MML Dataset

Dataset Construction.

To construct our dataset, we combined data from three sources: **iNaturalist** for images, **Xeno-Canto** for audio, and **Wikipedia** for textual and class-level grounding. For every bird class present in the VB100 or SSW60 datasets, we collected corresponding images (from iNaturalist), audio recordings (from Xeno-Canto), and Wikipedia articles.

Audio Processing Audio samples were segmented into 10-second clips (or zero-padded if shorter). We used the tags provided in Xeno-Canto as semantic grounding for the final captions.

Image Captioning Each image from iNaturalist was processed with `InstructBLIP2` to generate visual descriptions, which served as grounding for the final caption.
Model: `Salesforce/instructblip-flan-t5-xl`
Prompt:

```
Describe the bird's colors,
size, and shapes.
```

Wikipedia Text Processing For each class, we sampled random sections from the corresponding Wikipedia page. These sections were summarized by an LLM to produce textual grounding, generating 100 captions per class. This allowed us to capture diverse factual descriptions of each species.
Model: `google/gemma-2-2b-it`
Prompt:

```
You are a naturalist assistant.
Summarize the following text
into one coherent caption.
Be concise and factual.
{section_text}
Caption:
```

Caption Combination & Triplet Construction In the final step, we matched images, audio clips, and Wikipedia summaries into triplets. To produce a unified caption grounded in all three modalities (image, class, audio), we fed the image caption, audio tags, and a randomly selected Wikipedia caption into an LLM.
Prompt:

```
You are a naturalist assistant.
Combine the following
information into a concise
caption (10--20 words).
```

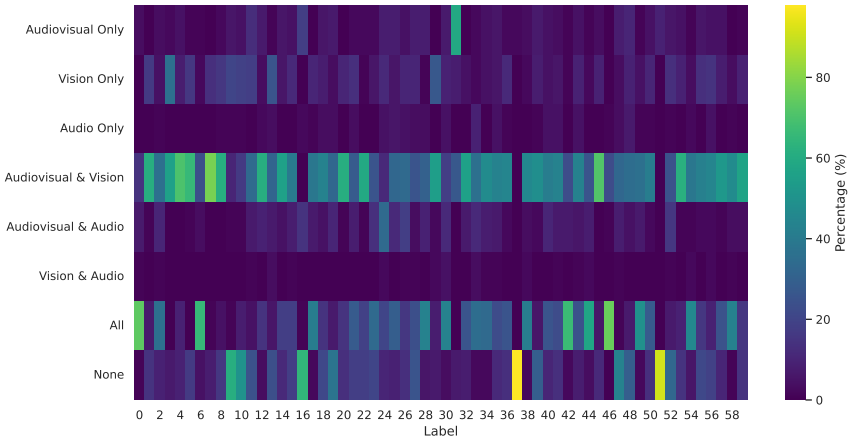

Figure 8. Per-class modality overlap analysis on the SSW60 dataset in the zero-shot classification setting. Each cell shows the percentage of samples within a class that fall into a given overlap category: *Audiovisual Only*, *Vision Only*, *Audio Only*, *Audiovisual & Vision*, *Audiovisual & Audio*, *Vision & Audio*, *All*, or *None*. Higher values indicate a larger fraction of samples for which that modality (or combination) produces correct predictions. This visualization highlights class-specific complementarity and competition between modalities.

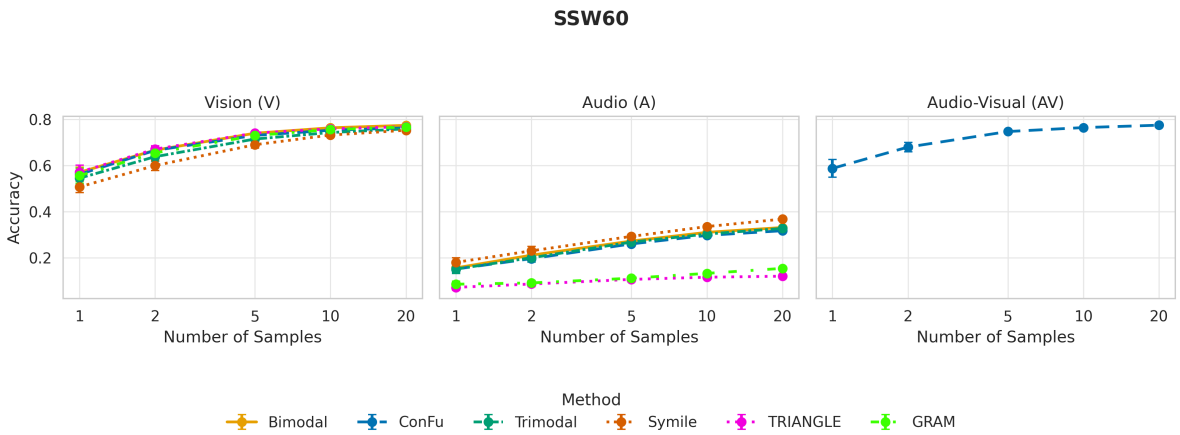

Figure 9. Few-shot linear probing results for SSW60 [30]. Performance is shown as the number of labeled examples increases. Prediction is done in the 8-frame sampling setting (multi-frame).

```
Include information about the
bird's appearance and sound if
available.
Be factual and informative.

Image caption: {image info}
Species caption: {textual info}
Sound tags: {tags}

Caption:
```

This process results in coherent multimodal captions grounded jointly in visual appearance, species knowledge, and acoustic characteristics.

D.2. Affective Computing Benchmarks

MultiBench [19] is a collection of benchmarking datasets designed to evaluate multimodal representation learning across a wide range of modalities and domains. In our experiments, we used the MOSI, UR-FUNNY, and MUStARD datasets. MultiBench provides pre-extracted features and predefined dataset splits, both of which we adopt in our setup.

MOSI is a multimodal sentiment analysis dataset consisting of 2,199 annotated YouTube video clips. UR-FUNNY contains 16,514 samples extracted from TED Talks and focuses on humor detection in spoken language. MUStARD includes 690 video clips from popular TV shows and is designed for multimodal sarcasm detection.

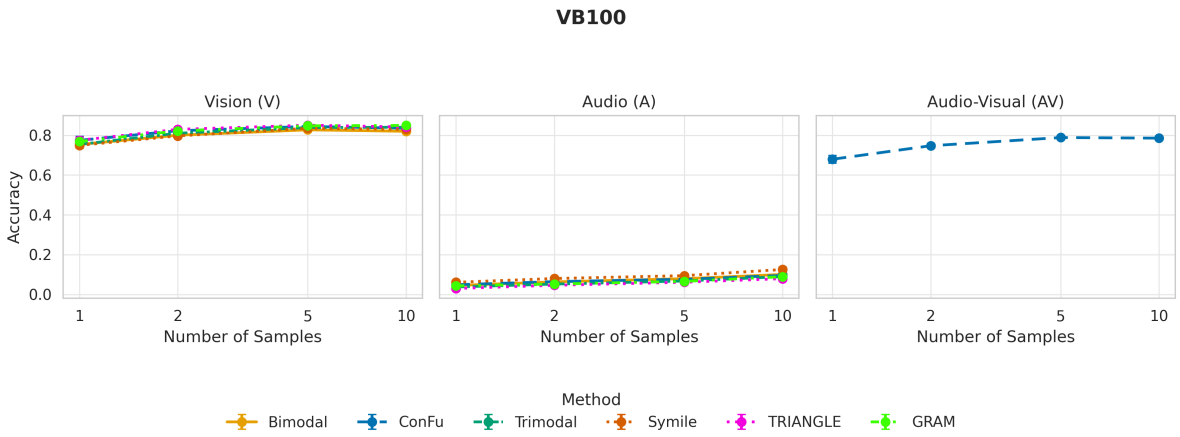

Figure 10. Few-shot linear probing results for VB100 [9]. Performance is shown as the number of labeled examples increases. Prediction is done in the 8-frame sampling setting (multi-frame).

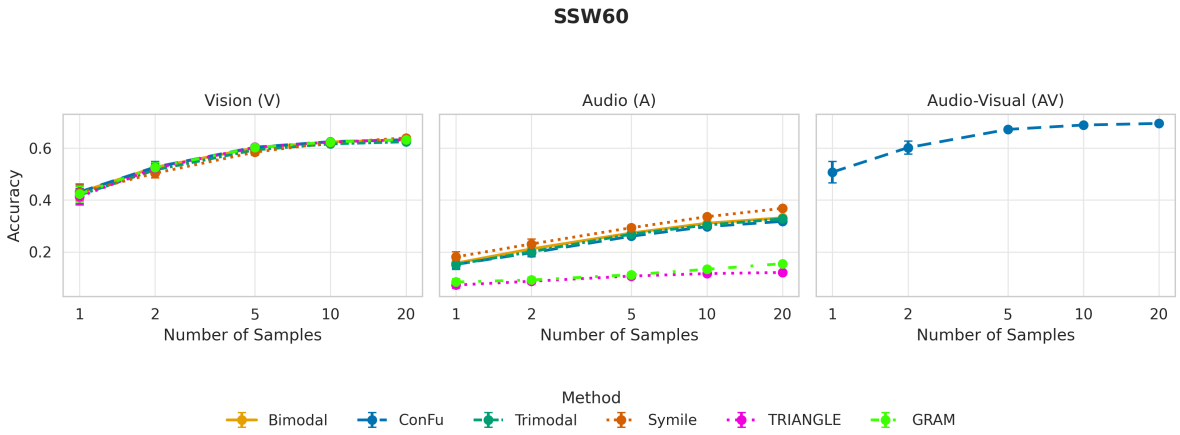

Figure 11. Few-shot linear probing results for SSW60 [30]. Performance is shown as the number of labeled examples increases.Prediction is done in the single frame setting.

D.3. AV-MNIST

The AV-MNIST dataset is a multimodal benchmark that pairs degraded visual features with audio spectrograms to evaluate multimodal representation learning. The visual modality consists of 28×28 MNIST digits that have been PCA-projected, retaining only 25% of the total variance. This dimensionality reduction is applied intentionally to weaken the visual signal and encourage effective fusion across modalities.

The audio modality comprises 112×112 spectrograms generated from the Free Spoken Digits Dataset [15], with additive background noise sampled from ESC-50 [25] to increase variability and realism.

The dataset contains 55000 training samples and 10000 test samples. In our experiments, we use the pre-extracted audio spectrograms distributed by MultiBench [19].

E. Generalization to M Modalities

For the generalization of our framework to M modalities, we extend the objective in Eq. 7 by summing InfoNCE losses over all relevant modality subsets. Specifically, we include all disjoint subset pairs

$$\mathcal{L}_M = \sum_{\substack{S_i, S_j \subseteq \{1,\ldots,M\} \\ S_i \cap S_j = \emptyset,\ S_i, S_j \neq \emptyset}} \widehat{\mathcal{L}}^{(S_i,S_j)}_{\text{InfoNCE}}, \tag{15}$$

which correspond to mutual information terms $I(X_{S_i}; X_{S_j})$ estimated via InfoNCE bounds. This formulation provides a direct extension of our tri-modal framework, allowing contrastive learning to capture both pairwise and higher-order dependencies across arbitrary modality combinations.

To ascertain the above claim, we conducted preliminary experiments in a four-modality setting. This was achieved by splitting the original images into full spatial resolu-

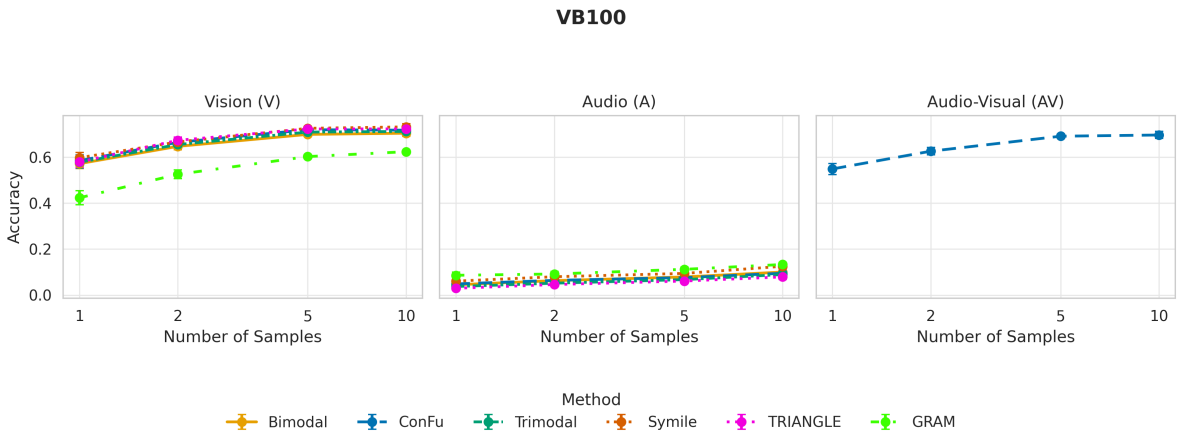

Figure 12. Few-shot linear probing results for VB100 [9]. Performance is shown as the number of labeled examples increases.Prediction is done in the single frame setting.

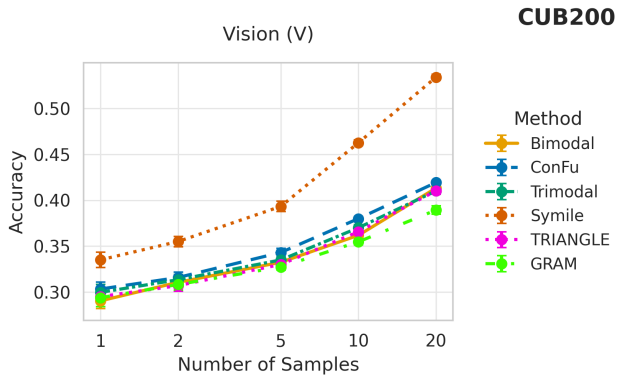

Figure 13. Few-shot linear probing results for CUB200 [32]. Performance is shown as the number of labeled examples increases. Prediction is done in the single frame setting.

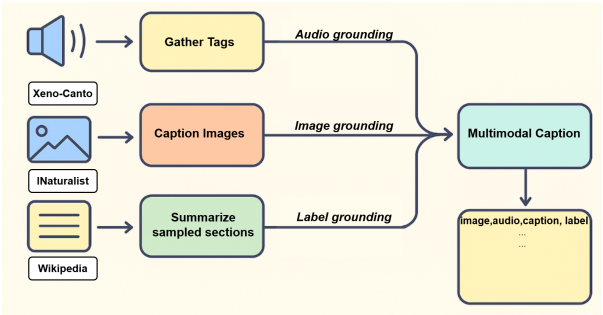

Figure 14. Overview of the multimodal data generation pipeline. Images and audio are collected from public sources (iNaturalist and Xeno-Canto), pseudocaptions are extracted from Wikipedia, and captions are generated via a combination of InstructBLIP2 and Gemma-2-2B-IT. The resulting triplets form the pretraining dataset.

tion monochromatic (grayscale) and low-resolution color (RGB), yielding four modalities alongside audio and text without additional data sources. Although artificial, this scenario emulates the behavior of numerous remote sensing platforms (e.g., Pléiades: 0.5 meter monocromatic and 2 meters RGB) and the associated problem of pansharpening. Experiments on VB100, and SSW60 show that the composite objective converges and consistently improves zero-shot performance over a pairwise-only baseline (Table 11), providing evidence that the approach remains effective as additional modalities are incorporated (with VB100 results reflecting the known weakness of the audio signal).

Table 11. Zero-shot classification accuracy (%) on SSW60 and VB100 datasets. G = Grayscale Vision, LR = Low-Res RGB Vision, A = Audio.

	SSW60 Acc. (%)		VB100 Acc. (%)	
Modality	**Pairwise CLIP**	**ConFu**	**Pairwise CLIP**	**ConFu**
G (grayscale)	50.33	48.46	14.48	15.47
LR (low-res RGB)	58.28	55.22	15.89	16.74
A (audio)	26.67	26.02	3.04	1.77
G + LR	–	59.39	–	18.15
G + A	–	52.20	–	14.05
LR + A	–	57.63	–	15.32
G + LR + A	–	**61.96**	–	**15.32**

Nevertheless, while the formulation in Eq. 15 captures all possible cross-subset dependencies, its combinatorial nature results in a large number of contrastive terms as M increases. In practice, task-specific relaxations of this general objective can substantially reduce computational complexity. For instance 16, when the goal is to retrieve a single target modality, one may restrict the loss to align any subset of the remaining modalities with that target only, i.e., terms of the form $I(X_t; X_S)$. Such relaxations preserve the rep-

resentational alignment relevant to the retrieval task while avoiding the exponential growth in the number of InfoNCE terms, yielding a more tractable yet principled multimodal contrastive objective.

$$\mathcal{L}_{\text{retrieval}} = \sum_{\substack{S_i, S_j \subseteq \{1,\ldots,M\} \\ S_i \cap S_j = \emptyset,\ S_i, S_j \neq \emptyset \\ |S_i| = 1}} \widehat{\mathcal{L}}_{\text{InfoNCE}}^{(S_i, S_j)}. \tag{16}$$