BETTER TOGETHER: LEVERAGING UNPAIRED MULTI-MODAL DATA FOR STRONGER UNIMODAL MODELS

Anonymous authors

Paper under double-blind review

Unpaired Multimodal Representation Learning

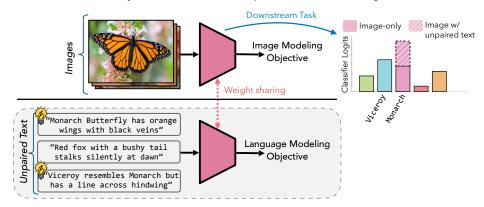


Figure 1: Text provides complementary information beyond images, even when not paired directly; We introduce UNPAIRED MULTIMODAL LEARNER (UML) whereby sharing model weights across modalities (e.g., image and text) extracts synergies and enhances unimodal representations, outperforming methods that rely only on a single modality (such as images above).

ABSTRACT

Traditional multimodal learners find unified representations for tasks like visual question answering, but rely heavily on large paired datasets. However, an overlooked yet potentially powerful question is: can one leverage auxiliary *unpaired* multimodal data to directly enhance representation learning in a *target* modality? We introduce UML: UNPAIRED MULTIMODAL LEARNER, a modality-agnostic training paradigm in which a single model alternately processes inputs from different modalities while sharing parameters across them. This design exploits the assumption that different modalities are projections of a shared underlying reality, allowing the model to benefit from cross-modal structure without requiring explicit pairs. Theoretically, under linear data-generating assumptions, we show that unpaired auxiliary data can yield representations strictly more informative about the world than unimodal training. Empirically, we show that using unpaired data from auxiliary modalities—such as text, audio, or images—consistently improves downstream performance across diverse unimodal targets such as image and audio.

1 Introduction

It is often taken for granted that to model a modality well, one must train on data from that modality. For instance, if one wants an accurate image classifier, one trains on images; if one wants a language model, one trains on text. Recent advances in multimodal learning, however, suggest that multiple modalities can benefit one another: in particular, using text captions paired with images yields richer representations, often surpassing their unimodal counterparts in zero-shot transfer, cross-modal retrieval, and downstream classification (Radford et al., 2021; Singh et al., 2022; Mizrahi et al., 2023; Girdhar et al., 2023a; Bachmann et al., 2022; Li et al., 2023; Bachmann et al., 2024; Jia et al., 2021). However, these gains have been realized largely through paired data, where one has access to

055

056

057

060

061 062

063

064

065

066

067

068

069

071

072

073

074

075

076

077

078

079

081

083

084

085

087

091

092

094

096

098

099

100

101

102

103 104

105 106

107

aligned examples $(x,y) \sim p_{\mathcal{X},\mathcal{Y}}$, such as an image x and its corresponding caption y. Such paired supervision allows aligning modalities into a shared latent space, where cross-modal correlations can be captured and transferred to downstream tasks.

This reliance on paired corpora also poses a bottleneck. Collecting and curating aligned datasets is expensive and domain-limited, whereas unpaired data—independent samples from $p_{\mathcal{X}}$ and $p_{\mathcal{Y}}$ —are naturally abundant. For example, vast image collections and vast text corpora exist independently, but without explicit alignment. This raises a more fundamental question:

Can unpaired data from a secondary modality Y enhance representations of the target modality X, even in the absence of (x, y) correspondences?

There is growing evidence that the answer may be yes. Recent work posits the existence of a shared statistical model of reality, an empirical parallel to Plato's concept of ideal Forms, suggesting that as deep networks scale, their embeddings across modalities converge toward a unified representation of the underlying world (Huh et al., 2024; Huang et al., 2021). This convergence implies that when paired supervision is available, a model can leverage natural co-occurrences between modalities and thus more accurately capture shared semantics. Critically, however, achieving convergence does not necessarily require explicit pairs; as long as each modality samples from the same underlying ground-truth latent space, it may be sufficient to align their marginal distributions to uncover the common semantic structure (Timilsina et al., 2024; Sturma et al., 2023). Intuitively, even unpaired data from another modality can provide complementary cues to better estimate the underlying reality (Figure 1).

In this work, we formalize this idea through Unpaired Multimodal Representation Learning, a framework for improving unimodal representations by leveraging unpaired data across modalities. Theoretically, under linear assumptions, we derive conditions under which incorporating unpaired samples yields strictly more informative representations than unimodal training alone. Strikingly, in some regimes, a single sample from Y provides greater per-sample value than a sample from X itself when the goal is to model X. Building on these insights, we introduce UML: UNPAIRED MULTIMODAL LEARNER, a shared-network framework that applies the same set of parameters to inputs from different modalities. By nothing more than weight sharing i.e. without surrogate objectives or inferred alignments, the model learns modality-agnostic features and naturally transfers information across modalities (Sutskever, 2023). Empirically, we evaluate UML on diverse image-text tasks in healthcare, and affective computing, as well as 10 standard visual benchmarks. Unpaired data from auxiliary modalities consistently improves unimodal representations across self-supervised and supervised regimes, in both few-shot and full-data regimes, and with diverse encoders including CLIP, DINOv2, OpenLLaMA, and others. These benefits compound when moving from two to three modalities, with audio, vision, and text each adding complementary signals. We further demonstrate effective cross-modal transfer without paired data by initializing vision models with pretrained language-model weights. Finally, we quantify the exchange rate between modalities, mapping how many words equate to one image (and vice versa) for optimal performance.

To summarize, the key contributions of our work are:

- We introduce UML, a modality-agnostic framework that leverages unpaired data to improve unimodal models. We test it across self-supervised and supervised encoders (CLIP, DINOv2, OpenLLaMA, and others), in both few-shot and full-data regimes, showing consistent gains over a range of image-text benchmarks and extensions to audio.
- We theoretically characterize, under linear assumptions, the conditions where unpaired data yield strictly more informative representations than unimodal training. Remarkably, the conversion ratio between modalities can fall below one: in certain regimes, a single sample from Y contributes more to modeling X than an additional sample from X itself.
- We quantify conversion ratios between images and text; i.e., how many words is an image worth for training vision models. We also show that unpaired multimodal data systematically widens inter-class margins and aligns modalities in weights.

2 Paired Multimodal Representation Learning

A useful way to conceptualize learning across modalities is to posit a shared ground-truth reality, denoted \mathbb{Z}^* , which manifests through multiple projections such as images, text, or audio record-

ings (Huh et al., 2024; Timilsina et al., 2024; Sturma et al., 2023). The goal of representation learning is to learn an embedding space that captures the structure of \mathcal{Z}^* , whether from a single modality or from several jointly. Thus, unimodal representations are inherently limited by what one projection alone can reveal: a single camera view may contain occlusions; an audio recording lacks visual details; textual descriptions may lack layout information.

These limitations motivate multimodal representation learning, which extends the unimodal setting to learn from multiple modalities together. For brevity, consider two modalities with observable data denoted by \mathcal{X} (e.g., images) and \mathcal{Y} (e.g., text) and samples $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. Mathematically, *Multimodal representation learning* seeks encoders $f_{\theta_x}: \mathcal{X} \to \mathcal{Z}$ and $f_{\theta_y}: \mathcal{Y} \to \mathcal{Z}$ mapping each modality into a shared embedding space \mathcal{Z} . Different frameworks optimize for desirable properties of \mathcal{Z} . Contrastive methods encourage instances of the same concept (such as an image and its text label) to be close together (Radford et al., 2021; Zhai et al., 2023; Jia et al., 2021). Fusion approaches learn shared layers for multimodal inputs with reconstruction (Singh et al., 2022; Li et al., 2019; Lu et al., 2019; Bachmann et al., 2022) or contrastive objectives (Roy et al., 2025). Generative methods often structure \mathcal{Z} to enable translation between modalities(Li et al., 2022; Rombach et al., 2022). Other methods disentangle common factors from modality-specific factors (Wang et al., 2024; Liang et al., 2023).

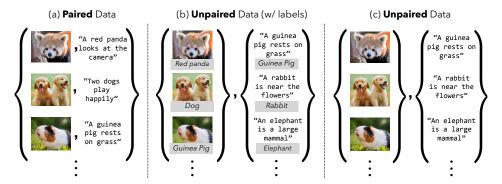


Figure 2: (a) Paired learning uses (x_i, y_i) with known correspondences. We instead study Unpaired learning: (b) with labels, using (x_i, c_i) and (y_j, \hat{c}_j) , where c_i and \hat{c}_j denote labels for x_i and y_j , but no cross-modal correspondences; and (c) without any labels or correspondences, using $\{x_i\}$ and $\{y_j\}$.

Common to all of these prior work, however, is a requirement for pre-existing knowledge of one-to-one correspondences between the modalities, i.e., samples $(x,y) \sim P_{X,Y}$ (as shown in Figure 2(a)), which are assumed to be aligned views of the same underlying entity in \mathcal{Z}^* . We ask whether it is possible to recover this shared structure *without* such correspondences—namely, when only unpaired samples from the marginals P_X and P_Y are available, as shown in Figure 2(c)—and, if so, how this affects downstream performance individually on the primary modality X.

3 UNPAIRED MULTIMODAL REPRESENTATION LEARNING

Let $\mathcal{D}_{\mathcal{X}} = \{x_i\}_{i=1}^{N_x}$ and $\mathcal{D}_{\mathcal{Y}} = \{y_j\}_{j=1}^{N_y}$ be datasets of N_x and N_y samples respectively, drawn from marginal distributions P_X and P_Y as shown in Figure 2(c). The joint distribution $P_{X,Y}$ is unknown. Critically, in contrast to paired multimodal learning, for any given $x_i \in \mathcal{D}_{\mathcal{X}}$ and $y_j \in \mathcal{D}_{\mathcal{Y}}$, there is no assumption that x_i and y_i are projections of the same underlying entity $z \in \mathcal{Z}$.

Our objective is to learn mappings $f_X: \mathcal{X} \to \mathcal{Z}$ and $f_Y: \mathcal{Y} \to \mathcal{Z}$ from each modality to a common embedding space that captures the shared structures between the unpaired datasets. We refer to this as *Unpaired Multimodal Representation Learning*. In contrast to previous works (Lin et al., 2023; Lee & Yoon, 2025; Girdhar et al., 2022; Chada et al., 2023) (discussed in Appendix A), we leverage unpaired data *without* inferring alignment, incorporating paired data, or assuming pre-aligned embeddings. Intuitively, without known correspondences, even unpaired data can be useful if modalities capture different axes of information pertaining to the underlying \mathcal{Z}^* . We formalize this intuition in Section 3.1, within the popularly studied framework of linear models, and introduce a modality-agnostic algorithm UML in Section 3.2 for learning representations from unpaired data.

3.1 THEORETICAL PERSPECTIVES

We adopt the widely used modeling choice that each modality can be expressed as a linear combination of shared and modality-specific components (Sturma et al., 2023; Timilsina et al., 2024; Huang et al., 2021). This formulation inherently accounts for informational differences across modalities by incorporating modality-specific features (Wang et al., 2024). Prior works (Sturma et al., 2023; Timilsina et al., 2024) have developed sufficient conditions for recovering joint distributions and the shared causal structure under this linear model. Building on this framework, we ask how joint training with unpaired modalities can still enhance representations within a single modality.

Data Generating Process. Assume that all factors of variation in reality live in a single d-dimensional space, \mathcal{Z}^* , i.e., $\theta \in \mathbb{R}^d$, modeled using a linear data-generating pipeline. This parameter can further be decomposed as $\theta \equiv [\theta_c, \theta_x, \theta_y]^\top$ where $\theta_c \in \mathbb{R}^{d_c}, \theta_x \in \mathbb{R}^{d_x}, \theta_y \in \mathbb{R}^{d_y}$ and $d_c + d_x + d_y = d$. Here, θ_c captures the *common* (shared) parameters that affect both modalities, θ_x denotes the parameters that only affect modality X, and θ_y denotes the parameters that only affect modality Y. We observe two datasets, one from each modality $\{X_i\}_{i=1}^{N_x} \in \mathbb{R}^m$ and $\{Y_j\}_{j=1}^{N_y} \in \mathbb{R}^n$, each reflecting partial measurements of the ground truth latent space \mathcal{Z}^* :

$$X_i = A_{c,i} \theta_c + A_{x,i} \theta_x + \epsilon_{X,i}, \quad \epsilon_{X,i} \sim \mathcal{N}(0, \sigma_x^2 I_m)$$
 (1)

$$Y_j = B_{c,j} \theta_c + B_{y,j} \theta_y + \epsilon_{Y,j}, \quad \epsilon_{Y,j} \sim \mathcal{N}(0, \sigma_y^2 I_n). \tag{2}$$

Here, $A_{c,i}$, $A_{x,i}$, $B_{c,j}$, $B_{y,j}$ are known design blocks capturing how each sample probes the latent factors and $\varepsilon_{X,i}$, $\varepsilon_{Y,j}$ represent the independent measurement noise.

In our linear setting, estimating the true latent state θ —and hence the underlying reality \mathcal{Z}^* —is governed by the Fisher information matrix $I(\theta) = -\mathbb{E}\left[\nabla_{\theta}^2 \ell(\theta)\right]$ (Fisher, 1922; 1925), which measures how sharply the likelihood "curves" around the true θ . Because X and Y samples are independent given \mathcal{Z}^* , their curvature contributions add pointwise, resulting in the joint Fisher information being simply the sum of the unimodal blocks. Thus, loosely speaking, any nonzero contribution from the unpaired Y-samples strictly increases curvature and thus strictly tightens the variance bound along those directions in θ_c . We formalize these insights in Theorem 1 and Theorem 2. Our results also generalize naturally to more than two modalities since the total contribution of all auxiliary modalities (excluding the primary modality X) can be obtained by summing their individual Fisher information matrices. Complete proofs and additional background are provided in Appendix C.

Theorem 1 (Variance Reduction with Unpaired Multimodal Data). Let $\hat{\theta}_X$, $\hat{\theta}_Y$ be the least-squares estimators for θ using only $\{X_i\}$ and only $\{Y_j\}$ and let $\hat{\theta}_{X,Y}$ be the joint estimator using both unpaired datasets. Then, under the assumption that at least one $B_{c,j}$, where $j \in \{1, 2, ...N_y\}$, has full rank, the common-factor covariance satisfies the strict Loewner ordering i.e. $\operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c} \prec \operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c}$, or equivalently, the Fisher information on θ_c strictly increases when combining both modalities, despite not having sample-wise pairing: $(I_X + I_Y)_{\theta_c,\theta_c} \succ (I_X)_{\theta_c,\theta_c}$.

Theorem 1 says that although there is no sample-wise pairing, simply adding an unpaired modality Y, that carries non-degenerate information about every direction in the shared subspace, can only tighten our uncertainty about the shared parameters θ_c . Geometrically, the uncertainty ellipsoid for θ_c shrinks in every direction where Y contributes any curvature, and never expands elsewhere. In practice however, no single data sample covers every latent axis. In these settings, while global shrinkage (Theorem 1) no longer applies, we still get *directional* gains. We capture this in Theorem 2.

Theorem 2 (Directional Variance Reduction with Unpaired Multimodal Data). Let all notation be as in Theorem 1, and let $v \in \mathbb{R}^{d_c} \setminus \{0\}$. If there exists at least one index $j \in \{1, 2, ...N_y\}$ such that $B_{c,j}v \neq 0$, then $v^\top (I_X + I_Y)_{\theta_c,\theta_c} v > v^\top (I_X)_{\theta_c,\theta_c} v$. Further, the variance of the estimator in direction v is strictly reduced if $v \notin range((I_X)_{\theta_c,\theta_c})$.

Building on Theorem 2, Corollary 1 and Corollary 2 quantify the precise variance contraction factor and analyze special cases where a second modality resolves otherwise ill-posed directions. Thus far, we have studied incorporating data from an auxiliary modality without considering sample size constraints. A natural next question is how to allocate a fixed budget of N additional samples:

is it better to collect them from the same modality X, or from complementary modality Y? We

formally address this in Theorem 3 below.

Theorem 3 (Data from Auxiliary Modality Can Outperform More of the Same). Define for any m, $I_X^{(m)} = \sum_{i=1}^m A_{c,i}^\top A_{c,i}$ and $I_Y^{(m)} = \sum_{j=1}^m B_{c,j}^\top B_{c,j}$. If $\operatorname{range}(I_Y^{(m)}) \not\subseteq \operatorname{range}(I_X^{(m)})$, then there exists a nonzero $v \in \mathbb{R}^{d_c}$ such that $v^\top I_Y^{(m)} v > v^\top I_X^{(m)} v$.

Theorem 3 shows that if the second modality covers a "blind spot" of the first, adding additional samples from the first modality does not increase the Fisher information in that direction; however, unpaired samples from the second modality provide strictly positive information along that axis.

To empirically validate the above result, we design an illustrative Gaussian experiment (Appendix E.14) in which both modalities are generated as noisy linear projections of the same underlying Gaussian latent variable θ_c . We allocate a fixed budget of samples, splitting evenly between X and Y (details in Appendix E.14). Joint training with UML significantly improves reconstruction fidelity on X (Figure 3) compared to training on X alone. Strikingly, this also reveals that the *effec*-

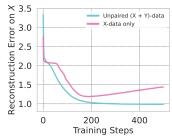


Figure 3: Adding unpaired Y samples boosts X reconstruction more than adding extra X samples.

tive exchange rate between modalities is below one, meaning a single additional sample from Y is worth more than an additional sample from X, even when the test task is on X. In Section 4.4, we extend this idea by quantifying the exchange rate between images and text on real world benchmarks.

3.2 UML: UNPAIRED MULTIMODAL LEARNER

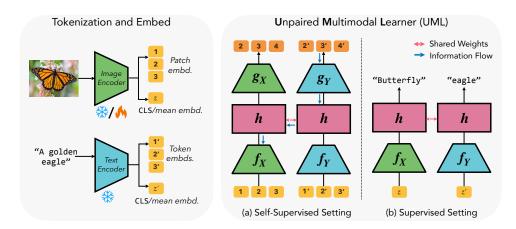


Figure 4: (Left) Inputs from different modalities (e.g., images or text) are tokenized into patch or token embeddings using pretrained encoders or processed features; (Right) UML can be trained under two settings: (a) Self-supervised, where patch/token embeddings are passed through a shared network and modality-specific decoders to perform next-token/patch prediction; (b) Supervised, where mean/CLS embeddings are fed through the shared classifier to predict labels within each modality.

The theory developed above establishes that adding an unpaired auxiliary modality strictly increases Fisher information along shared directions, thereby reducing estimator variance. We now translate these insights into a practical framework: UML (UNPAIRED MULTIMODAL LEARNER). The central idea is to share parameters across modalities: since both X and Y are projections of the same underlying reality \mathcal{Z}^* , forcing them through shared weights can extract synergies by accumulating training gradients on the same parameters. This accumulation could be viewed as the practical analogue of Fisher information addition, ensuring that even without paired samples, auxiliary modalities tighten estimates of shared structure. While prior work has explored shared weights in partially paired settings (Sturma et al., 2023), for related visual streams such as video and depth (Girdhar et al., 2022), or within already-aligned spaces like CLIP (Lin et al., 2023), we study and show their pronounced impact in the most general case of unpaired modalities with unaligned encoders.

Consider unpaired samples or preprocessed features $x \sim P_X$ and $y \sim P_Y$, each encoded into embeddings $z_X = f_X(x)$ and $z_Y = f_Y(y)$ as shown in Figure 4. These embeddings are passed through a shared network h, producing $r_X = h(z_X)$ and $r_Y = h(z_Y)$. We consider two regimes to train these networks: (a) self-supervised learning for fully unpaired data (Figure 2(c)) and (b) supervised learning for unpaired data with labels as shown in Figure 2(b). In the self-supervised setting, each modality has its own decoder $(g_X \text{ and } g_Y)$ trained to reconstruct the input or predict the next token/patch depending on the form of embeddings z_X and z_Y (refer to Figure 4(a)). The training objective (Pseudocode in the appendix: Algorithm 1) is

$$\mathcal{L}_{\text{UML-SSL}} = \mathbb{E}_{x \sim P_X} \, \ell(g_X(r_X), \, x) + \mathbb{E}_{y \sim P_Y} \, \ell(g_Y(r_Y), \, y),$$

where ℓ is the mean squared error for continuous targets or cross-entropy for discrete tokens. In the supervised setting, when labels are available, c_X for x and c_Y for y, a shared classifier $c(\cdot)$ is trained on top of r_X and r_Y . Since both $c(\cdot)$ and $h(\cdot)$ are shared across modalities, they can be viewed as a single shared head h (refer to Figure 4(b)). The supervised loss (Pseudocode in the appendix: Algorithm 2) is

$$\mathcal{L}_{\text{UML-Sup}} = \mathbb{E}_{(x,c_X)} \, \ell_{\text{CE}} \big(c(r_X), \, c_X \big) + \mathbb{E}_{(y,c_Y)} \, \ell_{\text{CE}} \big(c(r_Y), \, c_Y \big).$$

In both scenarios, although supervision is modality-specific, the shared head h receives updates from both modalities. Consequently, gradients from h also flow into f_X , effectively transferring information from f_Y and thus $\mathcal Y$ even without paired samples. At inference, we discard the auxiliary pathway and use r_X as the representation for the target modality, training a linear probe on top for downstream tasks. UML also naturally extends to more than two modalities by incorporating additional modality-specific encoder/decoder heads or classification layers.

4 EXPERIMENTAL RESULTS

In this section, we empirically evaluate UML across diverse benchmarks and configurations, leading to three main takeaways: (i) auxiliary modalities consistently boost target image representations in both self-supervised (Section 4.1) and supervised regimes (Section 4.2), with particularly strong gains in fine-grained and low-shot tasks; (ii) the benefits compound as we move from two to three modalities, with audio, vision, and text each contributing complementary signals (Section 4.2); and (iii) weight sharing across modalities generalizes beyond co-training, as pretrained language model weights *transfer* effectively to vision (Section 4.3). Finally, we quantify a marginal rate of substitution between modalities, asking how many text samples equate to an image (Section 4.4).

4.1 UML IN SELF-SUPERVISED SETTING

Table 1: UML (Ours) achieves higher top-1 linear probe accuracy than unimodal baselines on learned representations, averaged over three seeds, across both MultiBench and vision–text benchmarks.

				Data	aset			
		N	MultiBench	Standard vision-text				
Method	MUSTARD	MIMIC	Mosei	Mosi	UR-FUNNY	Oxford Pets	UCF101	DTD
Unimodal Ours	59.66 63.28 ↑	55.16 57.10 ↑	70.62 71.98 ↑	56.17 58.16 ↑	56.99 57.34 ↑	85.04 86.32 ↑	79.86 80.98 ↑	78.13 78.49 ↑

To study UML in the self-supervised setting, we use the real-world multimodal benchmark from MultiBench (Liang et al., 2021), which includes curated datasets such as text and images, with downstream labels covering a variety of domains, including healthcare, affective computing, and multimedia research and three standard classification benchmarks (Parkhi et al., 2012; Cimpoi et al., 2014; Soomro et al., 2012). We follow the same setting (dataset splitting, encoder architecture, pre-extracted features) as (Liang et al., 2023; 2021). For training (refer to Figure 4(a)), we use linear encoders (f_X and f_Y) and decoders (g_X and g_Y) and use an autoregressive transformer as the shared network (h). At inference, we average the transformer output embeddings and use the resulting embedding for linear probing on the downstream classification tasks. We report test accuracy using

embeddings from the primary modality (image), and to ensure fair comparisons, we perform rigorous hyperparameter tuning for all methods and repeat each experiment with three random seeds. For more details, refer to Appendix B.3.1 and Appendix B.4.1. As shown in Table 1, UML significantly outperforms its unimodal counterpart across all benchmarks, particularly on MUSTARD where unique information from the text modality expresses sarcasm, such as ironic tone of voice.

4.2 UML IN SUPERVISED SETTING

We evaluate UML on 9 widely-used visual classification benchmarks (Fei-Fei et al., 2004; Parkhi et al., 2012; Krause et al., 2013; Nilsback & Zisserman, 2008; Bossard et al., 2014; Xiao et al., 2010; Cimpoi et al., 2014; Soomro et al., 2012) in three settings: (1) Full fine-tuning: initializing from a pretrained vision backbone and updating all parameters on the target dataset. (2) Few-shot linear probing: freezing the vision backbone and training a linear classifier on k labeled samples per class (k = 1, 2, 4, 8, 16). (3) Full-dataset linear probing: freezing the vision backbone and training a linear classifier on the entire training dataset, discussed in Appendix E.1.2. In all cases, we enrich image representations with unpaired text embeddings, using ViT-S/14 DINOv2 as the vision encoder and OpenLLaMA-3B as the frozen text encoder. To construct conceptually related yet unpaired text data, we generate text templates with varying amounts of semantic information about the dataset. For further details and specific prompts, refer to Appendix B.3. To train UML, we initialize the classifier with the average text embedding of each class, giving a strong prior for image—class alignment (refer to Appendix E.1 for ablation).

Table 2: UML (Ours) outperforms unimodal baseline on image classification with ViT-S/14 DINOv2 and OpenLLaMA-3B in both settings: (i) Full finetuning and (ii) Few-shot linear probing (k = 1, 2, 4).

						Dat	aset				
Shot	Method	Stanford Cars	SUN397	FGVC Aircraft	DTD	UCF101	Food101	Oxford Pets	Oxford Flowers	Caltech101	Average
Full-f	inetuning										
	Unimodal	79.45	66.20	66.99	72.16	83.18	80.65	90.67	99.18	95.45	81.54
	Ours	86.39 ↑	66.03 ↓	73.44 ↑	74.27 ↑	84.69 ↑	81.97 ↑	91.72 ↑	99.82 ↑	97.60 ↑	83.99 ↑
Few-s	hot Linear Pi	robing									
1	Unimodal	13.18	34.15	14.09	36.60	46.74	35.18	63.51	89.62	76.66	45.52
1	Ours	16.49 ↑	41.79 ↑	15.63 ↑	42.04 ↑	52.33 ↑	42.27 ↑	73.59 ↑	93.64 ↑	84.52 ↑	51.36 ↑
	Unimodal	24.68	47.88	23.09	47.75	56.81	48.54	75.32	96.02	86.90	56.33
2	Ours	28.65 ↑	53.15 ↑	24.78 ↑	53.25 ↑	63.86 ↑	54.44 ↑	81.41 ↑	97.63 ↑	90.55 ↑	60.85 ↑
4	Unimodal	38.76	57.51	32.10	59.69	67.75	60.79	83.89	98.59	93.48	65.84
4	Ours	43.17 ↑	60.89 ↑	33.86 ↑	62.43 ↑	71.13 ↑	63.88 ↑	87.36 ↑	99.17 ↑	94.96 ↑	68.53 ↑

Unpaired Textual Data Improves Visual Classification. As shown in Table 2, across both full fine-tuning and few-shot linear probing, UML consistently improves over unimodal baselines across all datasets, with the largest gains on fine-grained tasks (e.g., Stanford Cars, FGVC Aircraft) where unpaired text sharpens class boundaries, and in low-shot regimes where textual cues help disambiguate visually similar classes. Additional results on other shots, datasets, model scales, and prompt variants are reported in Appendix E.

We also evaluate the robustness of UML-trained models under test-time distribution shifts. A 16-shot linear probe with DINOv2 is trained on ImageNet and tested across four distribution-shifted target datasets: ImageNet-V2, ImageNet-Sketch, ImageNet-A, and ImageNet-R. UML consistently outperforms the unimodal baseline (Figure 5), showing that language priors yield more transferable features. Additional robustness results are provided in Appendix E.3.

Finally, for all these settings, we also replace the independent vision (DINOv2) and text encoders (OpenLLaMA) with those of CLIP; since CLIP embeddings are already aligned, the gains from UML are even stronger (Appendix E.2.1, Appendix E.2.2, Appendix E.2.3,).

Additional Ablations. For all experiments, we keep the text encoder frozen; Unfreezing the text encoder yields slightly weaker gains but still outperforms the unimodal baseline (Appendix E.12). Swapping in different text encoders yields consistent gains (Appendix E.6), while richer, more diverse

captions provide especially strong boosts in few-shot settings (Appendix E.7). Further, training with semantically unrelated modalities (Appendix E.9) does not improve over the unimodal counterpart, confirming that gains stem from semantic correlations rather than spurious ones (Appendix E.9). Finally, unpaired multimodal data systematically widens inter-class margins and aligns modalities in weights (Appendix F).



Figure 5: Our approach UML is much more robust than its unimodal counterpart across four test time distribution shifted target test sets. All results are averaged across three random seeds.

Unpaired Image and Textual Data Improve Audio classification. We extend UML to an audio-vision-text setting using the ImageNet-ESC benchmark (Lin et al., 2023), which links ImageNet objects and captions with ESC-50 environmental sounds. The benchmark has two versions: ImageNet-ESC-27 and ImageNet-ESC-19. For audio encoding, we use AudioCLIP with an ES-ResNeXT backbone (Guzhov et al., 2021), while images and text are encoded by DINOv2 and OpenLLaMA-3B encoders. For further details, refer to Appendix B.3.

Unpaired image and text data consistently improve audio classification (Figure 6), with larger gains from CLIP's aligned encoders (Appendix E.13.3). Conversely, both audio and text also enhance image classification, showing that transfer works in both directions (Appendix E.13).

Finally, we study the full three-modality case, where we treat two modalities as auxiliaries and one as the target—for example, using both image and text as auxiliaries for audio classification. As shown in Appendix E.10, performance improves monotonically with each added auxiliary modality, with the strongest gains achieved when all three modalities (audio, vision, and text) are used together.

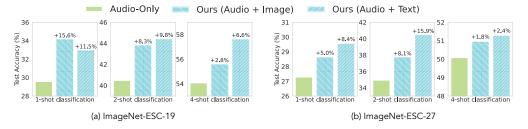


Figure 6: UML (Ours) improves audio classification using unpaired image and text samples on ImageNet-ESC-19 and ImageNet-ESC-27. All results are averaged across three random seeds.

4.3 Transfer Learning

Thus far, we have explored how sharing model weights while co-training with multiple unpaired modalities improves the learned representation. But weight sharing need not be restricted to co-training: if modalities capture the same underlying latents, then pretrained weights from one should also serve as a useful initialization for another. We therefore study if *transferring* knowledge from one modality can enhance performance in another by initializing the transformer layers of a ViT (Dosovitskiy et al., 2020) with pretrained BERT (Devlin et al., 2019)

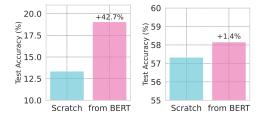


Figure 7: Image classifier trained from BERT initialization outperforms training from scratch for (left) full fine-tuning; (right) frozen backbone

 weights and evaluating on ImageNet (details in Appendix B.4.4). Patch embeddings are learned from scratch through a linear layer, augmented by a CLS token and positional embeddings. As shown in Figure 7, initializing with BERT boosts performance for both frozen and unfrozen backbones. Our results indicate that the semantic knowledge of language models provides a strong initialization for vision compared to training from scratch.

4.4 Marginal Rate-of-Substitution Between Modalities

Having established that unpaired modalities boost representation learning, robustness, and transfer, we now ask a more fundamental question: what is the relative value of each modality? If both images and text provide views of the same semantic space, can we quantify their exchange rate—How many words is an image worth? Figure 8 and Figure 9 visualize image-text conversion ratios using test accuracy isolines on Oxford-Pets (Parkhi et al., 2012) dataset. These map the number of texts equivalent to an image for the same performance. Aligned CLIP encoder (1 image \approx 228 words) is more efficient than non-aligned DINOv2 and OpenLLaMA encoders (1 image \approx 1034 words). Indeed, in some cases, an image may quite literally be worth a thousand words. We also observe little or no additional benefit from adding more text beyond a few samples. This is likely because we do not control for increasing complexity, so adding sentences does not guarantee extra information. Further results on additional datasets and key details, refer to Appendix E.4.

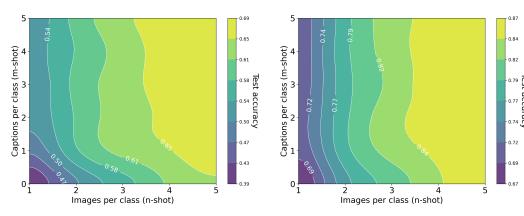


Figure 8: 1 img \approx 228 words (CLIP)

Figure 9: 1 img \approx 1034 words (DINOv2)

5 CONCLUSIONS AND LIMITATIONS

Conclusions. We propose and investigate *Unpaired Multimodal Representation Learning* for enhancing unimodal representations with unpaired multimodal data. Under linear assumptions, we theoretically show that unpaired data from multiple modalities strictly increases Fisher information along shared directions, resulting in a more accurate representation of the underlying world. Mechanistically, UNPAIRED MULTIMODAL LEARNER (UML) achieves this by accumulating gradients from different modalities on shared weights, which can be viewed as an operational analogue of the Fisher information gain. Empirically, we show performance gains across vision, text and audio benchmarks, and estimate conversion ratios between modalities. UML provides a new perspective on how to harness the abundance of unpaired data to learn better representations. This maybe especially useful in domains such as medical imaging, scientific data, and robotics, which contain rich auxiliary modalities like text, audio, or metadata that are not often paired with every instance of the primary modality.

Limitations. While our experiments study learning from unpaired multimodal data under both the self-supervised and supervised settings, downstream evaluations are conducted primarily for classification. Investigating evaluation tasks, such as generation, offers a rich avenue for future work. Furthermore, we evaluate how multimodal data enhances image and audio classification; it remains to show if they can, in turn, offer useful information for textual tasks.

6 REPRODUCIBILITY STATEMENT

We provide complete proofs and relevant background for all theoretical results in Appendix C.3. For experiments, we detail the setup, training protocols, and algorithm implementations in Appendix B.4 and Appendix D. All datasets are publicly available, and we follow established preprocessing procedures, with additional details in Appendix B.3. We will release all code and scripts in a public GitHub repository upon acceptance.

REFERENCES

- Amjad Almahairi, Sai Rajeshwar, Alessandro Sordoni, Philip Bachman, and Aaron Courville. Augmented cyclegan: Learning many-to-many mappings from unpaired data. In *International conference on machine learning*, pp. 195–204. PMLR, 2018.
- Roman Bachmann, David Mizrahi, Andrei Atanov, and Amir Zamir. Multimae: Multi-modal multi-task masked autoencoders. In *European Conference on Computer Vision*, pp. 348–367. Springer, 2022.
- Roman Bachmann, Oğuzhan F Kar, David Mizrahi, Ali Garjani, Mingfei Gao, David Griffiths, Jiaming Hu, Afshin Dehghan, and Amir Zamir. 4m-21: An any-to-any vision model for tens of tasks and modalities. *Advances in Neural Information Processing Systems*, 37:61872–61911, 2024.
- Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101-mining discriminative components with random forests. In *Computer vision–ECCV 2014: 13th European conference, zurich, Switzerland, September 6-12, 2014, proceedings, part VI 13*, pp. 446–461. Springer, 2014.
- 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.
- Rakesh Chada, Zhaoheng Zheng, and Pradeep Natarajan. Momo: A shared encoder model for text, image and multi-modal representations. *arXiv preprint arXiv:2304.05523*, 2023.
- Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3606–3613, 2014.
- Pinar Demetci, Rebecca Santorella, Björn Sandstede, William Stafford Noble, and Ritambhara Singh. Scot: Single-cell multi-omics alignment with optimal transport. *Journal of Computational Biology*, 29(1):3–18, 2022. doi: 10.1089/cmb.2021.0446. URL https://doi.org/10.1089/cmb.2021.0446. PMID: 35050714.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pp. 248–255. Ieee, 2009.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pp. 4171–4186, 2019.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv* preprint *arXiv*:2010.11929, 2020.
- Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In 2004 conference on computer vision and pattern recognition workshop, pp. 178–178. IEEE, 2004.
- Ronald A Fisher. On the mathematical foundations of theoretical statistics. *Philosophical transactions of the Royal Society of London. Series A, containing papers of a mathematical or physical character*, 222(594-604):309–368, 1922.

- Ronald Aylmer Fisher. Theory of statistical estimation. In *Mathematical proceedings of the Cambridge philosophical society*, volume 22, pp. 700–725. Cambridge University Press, 1925.
 - Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *International Journal of Computer Vision*, 132(2):581–595, 2024.
 - Xinyang Geng, Hao Liu, Lisa Lee, Dale Schuurmans, Sergey Levine, and Pieter Abbeel. Multimodal masked autoencoders learn transferable representations. *arXiv preprint arXiv:2205.14204*, 2022.
 - Rohit Girdhar, Mannat Singh, Nikhila Ravi, Laurens Van Der Maaten, Armand Joulin, and Ishan Misra. Omnivore: A single model for many visual modalities. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16102–16112, 2022.
 - 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*, pp. 15180–15190, 2023a.
 - Rohit Girdhar, Alaaeldin El-Nouby, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Omnimae: Single model masked pretraining on images and videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10406–10417, 2023b.
 - Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel. Esresne (x) t-fbsp: Learning robust time-frequency transformation of audio. In 2021 International Joint Conference on Neural Networks (IJCNN), pp. 1–8. IEEE, 2021.
 - 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.
 - Yu Huang, Chenzhuang Du, Zihui Xue, Xuanyao Chen, Hang Zhao, and Longbo Huang. What makes multi-modal learning better than single (provably). *Advances in Neural Information Processing Systems*, 34:10944–10956, 2021.
 - Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. The platonic representation hypothesis. *arXiv* preprint arXiv:2405.07987, 2024.
 - 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*, pp. 4904–4916. PMLR, 2021.
 - Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3(1):1–9, 2016.
 - Cijo Jose, Théo Moutakanni, Dahyun Kang, Federico Baldassarre, Timothée Darcet, Hu Xu, Daniel Li, Marc Szafraniec, Michaël Ramamonjisoa, Maxime Oquab, et al. Dinov2 meets text: A unified framework for image-and pixel-level vision-language alignment. *arXiv preprint arXiv:2412.16334*, 2024.
 - Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. Grounding language models to images for multimodal inputs and outputs. In *International Conference on Machine Learning*, pp. 17283–17300. PMLR, 2023.
 - Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pp. 554–561, 2013.
 - Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc'Aurelio Ranzato. Phrase-based & neural unsupervised machine translation. *arXiv preprint arXiv:1804.07755*, 2018.

- Jae-Jun Lee and Sung Whan Yoon. Can one modality model synergize training of other modality models? In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=5BXWhVbHAK.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pretraining for unified vision-language understanding and generation. In *International conference on machine learning*, pp. 12888–12900. PMLR, 2022.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019.
- Yanghao Li, Haoqi Fan, Ronghang Hu, Christoph Feichtenhofer, and Kaiming He. Scaling language-image pre-training via masking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 23390–23400, 2023.
- Paul Pu Liang, Yiwei Lyu, Xiang Fan, Zetian Wu, Yun Cheng, Jason Wu, Leslie Chen, Peter Wu, Michelle A Lee, Yuke Zhu, et al. Multibench: Multiscale benchmarks for multimodal representation learning. *Advances in neural information processing systems*, 2021(DB1):1, 2021.
- Paul Pu Liang, Zihao Deng, Martin Q Ma, James Y Zou, Louis-Philippe Morency, and Ruslan Salakhutdinov. Factorized contrastive learning: Going beyond multi-view redundancy. Advances in Neural Information Processing Systems, 36:32971–32998, 2023.
- Zhiqiu Lin, Samuel Yu, Zhiyi Kuang, Deepak Pathak, and Deva Ramanan. Multimodality helps unimodality: Cross-modal few-shot learning with multimodal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19325–19337, 2023.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023.
- Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. *Advances in neural information processing systems*, 30, 2017.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32, 2019.
- Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.
- Mayug Maniparambil, Raiymbek Akshulakov, Yasser Abdelaziz Dahou Djilali, Mohamed El Amine Seddik, Sanath Narayan, Karttikeya Mangalam, and Noel E O'Connor. Do vision and language encoders represent the world similarly? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14334–14343, 2024.
- Sachit Menon and Carl Vondrick. Visual classification via description from large language models. *arXiv preprint arXiv:2210.07183*, 2022.
- Jack Merullo, Louis Castricato, Carsten Eickhoff, and Ellie Pavlick. Linearly mapping from image to text space. *arXiv preprint arXiv:2209.15162*, 2022.
- David Mizrahi, Roman Bachmann, Oguzhan Kar, Teresa Yeo, Mingfei Gao, Afshin Dehghan, and Amir Zamir. 4m: Massively multimodal masked modeling. *Advances in Neural Information Processing Systems*, 36:58363–58408, 2023.
- Luca Moschella, Valentino Maiorca, Marco Fumero, Antonio Norelli, Francesco Locatello, and Emanuele Rodolà. Relative representations enable zero-shot latent space communication. *arXiv* preprint arXiv:2209.15430, 2022.
- Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In 2008 Sixth Indian conference on computer vision, graphics & image processing, pp. 722–729. IEEE, 2008.

- Antonio Norelli, Marco Fumero, Valentino Maiorca, Luca Moschella, Emanuele Rodola, and Francesco Locatello. Asif: Coupled data turns unimodal models to multimodal without training. *Advances in Neural Information Processing Systems*, 36:15303–15319, 2023.
 - Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In 2012 IEEE conference on computer vision and pattern recognition, pp. 3498–3505. IEEE, 2012.
 - Sarah Pratt, Ian Covert, Rosanne Liu, and Ali Farhadi. What does a platypus look like? generating customized prompts for zero-shot image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15691–15701, 2023.
 - 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*, pp. 8748–8763. PmLR, 2021.
 - Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
 - Karsten Roth, Jae Myung Kim, A Koepke, Oriol Vinyals, Cordelia Schmid, and Zeynep Akata. Waffling around for performance: Visual classification with random words and broad concepts. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15746–15757, 2023.
 - Shuvendu Roy, Franklin Ogidi, Ali Etemad, Elham Dolatabadi, and Arash Afkanpour. A shared encoder approach to multimodal representation learning. arXiv preprint arXiv:2503.01654, 2025.
 - Jayoung Ryu, Charlotte Bunne, Luca Pinello, Aviv Regev, and Romain Lopez. Cross-modality matching and prediction of perturbation responses with labeled gromov-wasserstein optimal transport. *arXiv* preprint arXiv:2405.00838, 2024.
 - Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36:62183–62223, 2023.
 - Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 15638–15650, 2022.
 - Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
 - Siddharth Srivastava and Gaurav Sharma. Omnivec: Learning robust representations with cross modal sharing. 2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pp. 1225–1237, 2023. URL https://api.semanticscholar.org/CorpusID:265128674.
 - Siddharth Srivastava and Gaurav Sharma. Omnivec2 a novel transformer based network for large scale multimodal and multitask learning. In *CVPR*, pp. 27402–27414, 2024. URL https://doi.org/10.1109/CVPR52733.2024.02588.
 - Nils Sturma, Chandler Squires, Mathias Drton, and Caroline Uhler. Unpaired multi-domain causal representation learning. *Advances in Neural Information Processing Systems*, 36:34465–34492, 2023.
 - Ilya Sutskever. "an observation on generalization". YouTube video, Aug 2023. URL https://www.youtube.com/watch?v=AKMuA_TVz3A. Accessed: 2025-09-24.
 - Subash Timilsina, Sagar Shrestha, and Xiao Fu. Identifiable shared component analysis of unpaired multimodal mixtures. *arXiv preprint arXiv:2409.19422*, 2024.
 - Chenyu Wang, Sharut Gupta, Xinyi Zhang, Sana Tonekaboni, Stefanie Jegelka, Tommi Jaakkola, and Caroline Uhler. An information criterion for controlled disentanglement of multimodal data. *arXiv* preprint arXiv:2410.23996, 2024.

- Johnny Xi, Jana Osea, Zuheng Xu, and Jason S Hartford. Propensity score alignment of unpaired multimodal data. *Advances in Neural Information Processing Systems*, 37:141103–141128, 2024.
- Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In 2010 IEEE computer society conference on computer vision and pattern recognition, pp. 3485–3492. IEEE, 2010.
- 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.
- Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 18123–18133, 2022.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 11975–11986, 2023.
- Yiyuan Zhang, Kaixiong Gong, Kaipeng Zhang, Hongsheng Li, Yu Qiao, Wanli Ouyang, and Xiangyu Yue. Meta-transformer: A unified framework for multimodal learning. *arXiv* preprint *arXiv*:2307.10802, 2023.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 2223–2232, 2017.

APPENDIX **Further Related Works Supplementary Experimental Details and Assets Disclosure** B.3 B.3.2 B.3.3 B.3.4 Unpaired Multimodal Representation Learning under Self-Supervision . . . B.4.2 B.4.3 **Proofs of Theoretical Results** C.2 D UML: Algorithm Pseudocode **Additional Experiments** Improving Image Classification using Unpaired Texts (Unaligned encoders) E.1.1E.1.2 E.1.3 E.2 Improving Image Classification using Unpaired Texts (Aligned encoders) E.2.1 E.2.2E.2.3E.3 E.4.1

	E.7	earning with Coarse-Grained vs. Fine-Grained Textual Cues	14
	E.8	mpact on Performance with Increasing Unpaired Text Prompts	14
	E.9	Effect of Unrelated Auxiliary Modalities	15
	E.10	Extension to More Than Two Modalities	15
	E.11	Effect of Ratio of Modality Batches	16
	E.12	Effect of Freezing vs. Unfreezing the Text Encoder	17
	E.13	Additional Experiments for Audio-Visual Setting	17
		2.13.1 Improving Image Classification with Unpaired Audio and Text (Unaligned encoders)	17
		2.13.2 Improving Image Classification with Unpaired Audio and Text (Aligned encoders)	18
		2.13.3 Improving Audio Classification with Unpaired Image and Text (Aligned encoders)	18
	E.14	Gaussian Experiments	18
F	Anal	sis of the Learned Classifier 4	19
	F.1	Change in Decision Boundaries with Unpaired Data from Another Modality 4	19
	F.2	What do models learn from unpaired data?	50

A FURTHER RELATED WORKS

Unpaired Multimodal Learning. Unpaired data has long been used for image-to-image (Zhu et al., 2017; Liu et al., 2017; Almahairi et al., 2018; Shi et al., 2023) and text-to-text translation (Lample et al., 2018). More recently, several works have also proposed learning from unpaired data by inferring coarse- or fine-grained alignments through distribution matching or optimal transport objectives (Xi et al., 2024; Demetci et al., 2022; Ryu et al., 2024). In contrast, we leverage unpaired data for learning representations without the need for explicit or inferred alignment. (Timilsina et al., 2024; Sturma et al., 2023) theoretically analyze the problem of identifying shared latent components and causal structures in unaligned multimodal mixtures. Most closely related to our work is (Lin et al., 2023), which leverages coarse-grained text data such as class names to improve image classification on CLIP using a shared linear head. Another related line of works (Roth et al., 2023; Pratt et al., 2023; Menon & Vondrick, 2022; Gao et al., 2024) leverage prompting templates and pretrained LLMs to generate descriptive class captions, showing improved image classification performance with CLIP. Nonetheless, these methods operate on CLIP with pre-aligned representation spaces, whereas our approach also learns from unpaired data without assuming prior alignment. Several works have also proposed learning large multitask multimodal models with joint encoders and unified embedding spaces (Srivastava & Sharma, 2024; 2023; Zhang et al., 2023; Girdhar et al., 2022; Geng et al., 2022), often using joint training over separate tasks and/or masked prediction objectives. In a similar vein, (Chada et al., 2023) uses a stage-wise training strategy with both unpaired and paired data, and (Girdhar et al., 2023b) trains a single model across visual modalities. However, most of these methods rely on some amount of paired data for preliminary alignment and then leverage abundant modality-specific unpaired data for further improvement. In contrast, our approach demonstrates that a model can implicitly learn cross-modal correlations from purely unpaired data, without requiring explicit alignment as a prerequisite.

Multimodal Representation Alignment. Our method relies on the notion of shared information and structure between unaligned modalities. Closely related to this are works demonstrating that unimodal representations trained without multimodal data are nevertheless converging. Huh et al. (2024) presents evidence that better-performing language models exhibit increased alignment to self-supervised vision models. Similarly, (Maniparambil et al., 2024) shows a latent space alignment between vision and text encoders across backbones and training paradigms, and uses the CKA metric to connect unaligned encoders zero-shot. Earlier works also note alignment between models trained with different datasets and modalities (Moschella et al., 2022; Norelli et al., 2023). Several works have also shown that a linear projection or MLP is sufficient to stitch together the latent spaces of pretrained vision and language models (Merullo et al., 2022; Liu et al., 2023; Koh et al., 2023). Zhai et al. (2022) extends this to training a text encoder to align to a frozen pretrained image model; this method was in turn used to integrate DINOv2, a large self-supervised vision model, with a text encoder (Jose et al., 2024).

B Supplementary Experimental Details and Assets Disclosure

B.1 ASSETS

We do not introduce new data in the course of this work. Instead, we use publicly available, widely used image datasets for the purposes of benchmarking and comparison.

B.2 HARDWARE AND SETUP

Each experiment was conducted on 1 NVIDIA Tesla V100 GPUs, each with 32GB of accelerator RAM. The CPUs used were Intel Xeon E5-2698 v4 processors with 20 cores and 384GB of RAM. All experiments were implemented using the PyTorch deep learning framework.

B.3 DATASETS

B.3.1 MULTIBENCH

We evaluate our approach on a diverse set of multimodal fusion datasets from MultiBench (Liang et al., 2021) spanning healthcare, sentiment, and humor detection:

- **CMU-MOSEI** (Moschella et al., 2022): The largest sentence-level multimodal sentiment and emotion dataset, containing 23,000 annotated monologue videos (over 65 hours of content from more than 1,000 speakers across 250 topics). Each video is labeled with sentiment intensity in the range [-3,3], which we cast into binary positive/negative sentiment classification.
- **CMU-MOSI** (Zadeh et al., 2016): A related multimodal sentiment dataset with 2,199 YouTube video clips, reflecting real-world opinionated speech. Sentiment intensities are annotated in the range [-3, 3], and we again formulate the task as binary sentiment classification.
- UR-FUNNY (Hasan et al., 2019): A large-scale humor detection dataset derived from over 16,000 TED talk videos. It includes 8,257 humorous punchlines identified by laughter markers, paired with 8,257 negative examples drawn from non-humorous contexts, forming a balanced binary humor classification task.
- MUSTARD (Castro et al., 2019): A multimodal sarcasm detection dataset with 690 video clips from TV shows. Each sample is annotated for the presence or absence of sarcasm, yielding a challenging binary classification task.
- MIMIC-III (Johnson et al., 2016): A large-scale clinical dataset with records of over 40,000 ICU patients. It combines time-series physiological measurements (recorded hourly over a 24-hour window) with static demographic and tabular features. We use it for binary classification of patients into a group of ICD-9 codes.

B.3.2 IMAGE CLASSIFICATION BENCHMARKS

Table 3: Detailed statistics of the 10 datasets for image classification.

Dataset	Classes	Train	Val	Test
Caltech101 (Fei-Fei et al., 2004)	100	4,128	1,649	2,465
OxfordPets (Parkhi et al., 2012)	37	2,944	736	3,669
StanfordCars (Krause et al., 2013)	196	6,509	1,635	8,041
Oxford Flowers (Nilsback & Zisserman, 2008)	102	4,093	1,633	2,463
Food101 (Bossard et al., 2014)	101	50,500	20,200	30,300
FGVCAircraft (Maji et al., 2013)	100	3,334	3,333	3,333
SUN397 (Xiao et al., 2010)	397	15,880	3,970	19,850
DTD (Cimpoi et al., 2014)	47	2,820	1,128	1,692
UCF101 (Soomro et al., 2012)	101	7,639	1,898	3,783
ImageNet (Deng et al., 2009)	1,000	1.28M	N/A	50,000

We evaluate on the following widely-used classification benchmarks: ImageNet (Deng et al., 2009), StanfordCars (Krause et al., 2013), UCF101 (Soomro et al., 2012), Caltech101 (Fei-Fei et al., 2004), Oxford Flowers (Nilsback & Zisserman, 2008), SUN397 (Xiao et al., 2010), DTD (Cimpoi et al., 2014), FGVCAircraft (Maji et al., 2013), OxfordPets (Parkhi et al., 2012), and Food101 (Bossard et al., 2014). More details about the dataset and splits is provided in Table 3.

B.3.3 Constructing text templates

To construct conceptually related yet unpaired text data, we generate text templates that capture varying granularities of information about the dataset. Our first approach (*Vanilla*) uses the straightforward template \a photo of a {}" with a natural language label for each category, resulting in a basic text description for each class. However, this simple textual corpus lacks fine-grained information necessary to distinguish between visually similar subcategories or to

resolve contextually ambiguous terms. To address this, for the second template, we draw from the extensive literature on improving text prompts for zero-shot classification in CLIP (Gao et al., 2024; Menon & Vondrick, 2022; Pratt et al., 2023; Roth et al., 2023). Specifically, for the second approach (*GPT-3 Descriptions*), we adopt the text prompt generation strategy developed by Pratt et al. (2023), using large language models such as GPT-3 to generate diverse and contextually rich prompts for each image category. We use three generic hand-written sentences across the datasets:

```
Describe what a/the {} looks like:

Describe a/the {}:

What are the identifying characteristics of a/the {}?
```

The blank portion of each template is populated with the category name, along with the category type for specialized datasets (e.g., "pet" + {} for Oxford Pets or "aircraft" + {} for FGVC Aircraft). The type specification is important for disambiguating categories with multiple interpretations. Some examples of these descriptions are provided in Table 4 for the Oxford Pets dataset.

Table 4: Sample text descriptions per class for Oxford Pets dataset

Class	Examples
Wheaten Terrier	A wheaten terrier is a small, shaggy dog with a soft, silky coat. A wheaten terrier has a soft, wheat-colored coat that is low-shedding and hypoallerge The wheaten terrier is a medium-sized, hypoallergenic dog breed. A pet Wheaten Terrier usually has an intelligent expression and a soft, wheat-colored
Great Pyrenees	A great pyrenees is a large, white, shaggy-coated dog. A Great Pyrenees is a large, fluffy dog with a calm, gentle disposition. The great pyrenees was originally bred to protect livestock from predators. Great Pyrenees are known for being very large, white dogs with thick fur.
Sphynx	A pet Sphynx typically has a small, wrinkled head and a hairless body. A Sphynx is a hairless cat breed known for its soft, warm skin. A Sphynx often displays large ears, pronounced cheekbones, and no fur. Sphynx are unique cats characterized by their lack of coat and wrinkled skin.
Birman	A Birman is a long-haired, color-pointed cat with a "mask" of darker fur on its face. A Birman has silky, pale cream to ivory fur with deep seal- or lilac-colored points. Birman cats possess striking blue eyes and contrasting white "gloves" on their paws. They are known for being gentle, affectionate, and smooth-coated companions.
Pomeranian	A Pomeranian is a small, fluffy dog with a thick double coat. Pomeranians are toy-sized, alert dogs with fox-like faces and plumed tails. A pet Pomeranian often comes in orange, black, white, or mixed coat colors. They are lively, outgoing, and known for their bold, friendly personalities.

B.3.4 IMAGENET-ESC DATASET

Experimental Setup. We extend our results beyond vision and language to an audiovisual-language dataset: the ImageNet-ESC benchmark (Lin et al., 2023). This benchmark combines ImageNet (1000 object categories) and ESC-50 (50 environmental sound classes) by matching classes that logically correspond. For example, the dog (barking) class from ESC-50 aligns with various dog breeds from ImageNet, while the clock-alarm sound maps to both analog clock and digital clock. This alignment captures the relationship between visual objects, their sounds, and their textual descriptions. The benchmark consists of two versions: 1) ImageNet-ESC-27: A broader set including loosely matched visual-audio pairs (e.g., drinking-sipping to water bottle); 2) ImageNet-ESC-19: A more precise subset containing only accurate visual-audio matches.

B.4 Training Protocol

B.4.1 Unpaired Multimodal Representation Learning under Self-Supervision

MultiBench. For the MIMIC dataset, which contains tabular and medical time-series inputs, we train models directly on the raw modality inputs. For the four video datasets (MOSEI, MOSI, UR-FUNNY, and MUSTARD), we train models on standard pre-extracted features from text, video, and audio modalities (Liang et al., 2021). All models are trained for 100 epochs, with hyperparameter search over learning rates $\{10^{-2}, 10^{-3}, 10^{-4}\}$. We report the best performance across learning rates, averaged over three random seeds. To train UML, each modality is first projected into a shared embedding space of dimension $d \in \{10, 40, 150, 300\}$ via a learned linear transformation. The projected inputs are processed by a shared 5-layer, 5-head autoregressive Transformer encoder, followed by a linear projection back to the original modality dimension. Training uses a nexttoken/patch embedding prediction objective, with both modalities sharing the Transformer backbone to encourage cross-modal synergies in the latent space. At inference time, we average the Transformer outputs across sequence length and use the resulting embeddings for linear probing on downstream classification tasks. All models (Unimodal baseline and UML) are trained for 100 epochs, with hyperparameter search over learning rates $\{10^{-2}, 10^{-3}, 10^{-4}\}$. For UML, in addition, we perform hyperparameter search over a curriculum parameter step $\in \{0, 30, 50, 70\}$. This parameter controls whether training begins on X alone for the first step epochs before switching to joint training, with step = 0 corresponding to joint training from the start. For each dataset, we select the best-performing model on the validation set and report test accuracy averaged over three random seeds.

Standard Vision-Text Benchmarks. We extract image and text embeddings using ViT-B/14 DINOv2 and OpenLLaMA-3B, respectively. As in MultiBench, patch and token embeddings are projected to a shared 256-dimensional space via modality-specific linear layers. A 4-layer, 4-head transformer serves as the shared encoder, with outputs projected back to the original embedding dimensions using modality-specific linear projections. We perform rigorous hyperparameter tuning for both the unimodal baseline and UML, and report average test accuracy of the best model across three seeds. Full hyperparameter ranges are listed in Table 5.

B.4.2 IMAGE CLASSIFICATION USING IMAGE AND UNPAIRED TEXTS

For text, we use OpenLLaMA-3B as our default encoder and ablate against BERT-Large, RoBERTa-Large, GPT-2 Large, and the pre-aligned CLIP text encoder, keeping the text encoder frozen. For images, our main backbone is ViT-S/14 DINOv2, with ablations across other DINOv2 variants and the CLIP vision encoder. In the linear-probe setting, all encoder weights stay fixed and we train only a single linear classification head; in full fine-tuning, we jointly update the image backbone and that head, while still freezing the text encoder.

We optimize cross-entropy loss via AdamW (Loshchilov & Hutter, 2017) and perform an extensive grid search over learning rate, weight decay, cosine learning rate scheduling with linear warmup, dropout, and a learnable, modality-specific scaling on the logits. The results are reported for the best-performing model on the validation dataset. We report results for the model achieving highest validation accuracy; the full hyperparameter ranges are in Table 5.

For full fine-tuning, we jointly update the image backbone and classification head with a fixed learning rate of 5×10^{-5} , batch size 64, and omit learnable modality-specific scaling, since it showed no benefit in this setting.

B.4.3 EVALUATION ON IMAGENET-ESC

Similar to our vision-language experiments, we perform few-shot evaluation using the 5-fold splits defined in the benchmark. Each fold contains 8 samples per class, with one fold used for training and validation and the remaining four for testing. We repeat the process over 5 random splits and report the average performance. For audio encoding, we use AudioCLIP with an ES-ResNeXT backbone (Guzhov et al., 2021). AudioCLIP is pretrained on AudioSet and generates audio embeddings in the same representation space as CLIP. Following the instructions in (Guzhov et al., 2021; Lin et al., 2023), we use train() mode in Pytorch to extract the features since eval()

Table 5: Hyperparameter grid for linear probing.

Hyperparameter	Values
Optimizer	adamw
Learning rate	$\{0.001, 1e-4\}$
Weight decay	$\{0.0, 0.01, 0.001\}$
LR scheduler	cosine
Batch size	{8, 32}
Max iterations	12,800
Warmup iterations	50
Warmup type	linear
Warmup min LR	1e-5
Dropout	$\{0.0\}$
Modality-specific learnable scaling	False, True
Early-stop patience	10

mode yields suboptimal embeddings. We evaluate our models on two tasks—audio classification and image classification—comparing the unimodal baseline against two multimodal variants in which the primary modality is each time augmented by one of the other modalities.

B.4.4 Transfer Learning from Language to Vision

To adapt a language model to image classification, we embed image patches using a linear projection and add positional encodings to capture spatial structure. We then use transformer layers initialized from pretrained BERT, and finally, a 2-layer MLP classification head. Specifically, we split each image of size 224×224 into patches of size 16×16 with 196 patch tokens. Each patch is then projected into the model's embedding space of dimension $d(e.g.\ d=768$ for GPT-2, d=1024 for BERT) via a learned linear layer. We then prepend a learnable "<code>[CLS]</code>" token, add learned positional embeddings of shape $(N+1) \times d$, and apply dropout with probability p=0.1. This $(N+1) \times d$ sequence is passed into the pretrained transformer stack (either GPT-2 or BERT), using a full bidirectional attention mask over all patch tokens and the CLS token. We extract the final hidden state corresponding to the CLS token and feed it through a two-layer MLP classification head.

During training, we evaluate two scenarios: 1) one where the pretrained backbone is frozen and only the patch embedding and linear head are trained, and 2) another where the backbone is initially frozen to align the trainable layers (patch embedding and head) with the pretrained language backbone, and then unfrozen after 2000 steps for end-to-end training. This approach allows us to test whether the semantic richness captured by language models provides a strong initialization, leading to better convergence and performance compared to training ViT from scratch.

C Proofs of Theoretical Results

In this section, we present complete derivations and proofs of the main theoretical claims. Appendix C.1 gathers all definitions and background required for our arguments. Appendix C.2 formalizes the linear data-generating model, derives closed-form maximum-likelihood estimators for each modality and their joint estimator, and computes the corresponding block-wise Fisher information. Finally, Appendix C.3 provides the detailed proofs of our variance-reduction claims, showing rigorously how unpaired multimodal estimation strictly lowers estimator variance.

C.1 BACKGROUND AND DEFINITIONS

In this section we revisit the mathematical definitions used in our theoretical analysis, including matrix-orderings, characterization of symmetric matrices and Fisher information.

Definition 1 (Positive Semidefinite Matrix). A real symmetric matrix $A \in \mathbb{R}^{d \times d}$ is *positive semidefinite* if for all vectors $v \in \mathbb{R}^d$, $v^\top A v \geq 0$. Equivalently, all eigenvalues of A are nonnegative. We denote the set of all $d \times d$ symmetric, positive-semidefinite matrices as $S^d_{\succeq 0}$.

Definition 2 (Positive Definite Matrix). A real symmetric matrix $A \in \mathbb{R}^{d \times d}$ is *positive definite* if for every nonzero $v \in \mathbb{R}^d$, $v^\top A v > 0$. Equivalently, all eigenvalues of A are strictly positive. We denote the set of all $d \times d$ symmetric, positive definite matrices as $S^d_{\succ 0}$.

Definition 3 (Loewner Order). For two real symmetric matrices $A, B \in \mathbb{R}^{d \times d}$, we write $A \leq B \iff B - A$ is positive semidefinite and $A \prec B \iff B - A$ is positive definite. This defines a partial order on the cone of symmetric matrices.

Definition 4 (Fisher Information Matrix). Given a parametric family of densities $p(x; \theta)$ on data x, the *Fisher information matrix* at parameter θ is

$$I(\theta) = \mathbb{E}_{x \sim p(\cdot; \theta)} \left[\nabla_{\theta} \log p(x; \theta) \nabla_{\theta} \log p(x; \theta)^{\top} \right].$$

Equivalently, for regular models, $I(\theta) = -\mathbb{E}\left[\nabla_{\theta}^2 \log p(x;\theta)\right]$.

C.2 MAXIMUM LIKELIHOOD ESTIMATORS AND FISHER CONTRIBUTIONS

In this section we revisit our linear data-generating model, introduce notations for the X-only, Y-only and joint likelihoods, derive the closed-form MLEs $\hat{\theta}_X$, $\hat{\theta}_Y$ and $\hat{\theta}_{X,Y}$, and formalize their information contributions towards estimating the ground truth parameters $\theta = [\theta_c, \theta_x, \theta_y]^{\top}$.

Data Generating Process. Recall our linear data-generating process: Assume that all factors of variation in reality live in a single d-dimensional space $\mathcal{Z}^* \equiv \theta \in \mathbb{R}^d$ modeled using a linear data-generating pipeline. This parameter can further be decomposed as $\theta \equiv [\theta_c, \theta_x, \theta_y]^\top$ where $\theta_c \in \mathbb{R}^{d_c}, \theta_x \in \mathbb{R}^{d_x}, \theta_y \in \mathbb{R}^{d_y}$ and $d_c + d_x + d_y = d$. Here, θ_c captures the *common* (shared) parameters that affect both modalities, θ_x denotes the parameters that only affect modality X, and θ_y denotes the parameters that only affect modality Y. We observe two independent datasets, one from each modality $\{X_i\}_{i=1}^{N_x} \in \mathbb{R}^m$ and $\{Y_j\}_{j=1}^{N_y} \in \mathbb{R}^n$, each reflecting partial measurements of the ground truth latent space \mathcal{Z}^* :

$$X_i = A_{c,i} \theta_c + A_{x,i} \theta_x + \epsilon_{X,i}, \quad \epsilon_{X,i} \sim \mathcal{N}(0, \sigma_x^2 I_{m_i})$$
(3)

$$Y_j = B_{c,j} \theta_c + B_{y,j} \theta_y + \epsilon_{Y,j}, \quad \epsilon_{Y,j} \sim \mathcal{N}(0, \sigma_y^2 I_{n_j}). \tag{4}$$

Here, $A_{c,i}, A_{x,i}, B_{c,j}, B_{y,j}$ are known design blocks capturing how each sample probes the latent factors and $\varepsilon_{X,i}, \varepsilon_{Y,j}$ represent the independent measurement noise.

In our linear setting, estimating the true latent state θ —and hence the underlying reality \mathcal{Z}^* —is governed by the Fisher information matrix $I(\theta) = -\mathbb{E}\big[\nabla^2_{\theta}\,\ell(\theta)\big]$, which measures how sharply the likelihood "curves" around the true θ . High curvature along a particular axis means the data tightly constrain that component, driving down estimator variance there.

Unimodal Estimators. We first estimate θ using only the X-dataset. Stacking $\{X_i\}_{i=1}^{N_x}$ yields a design matrix \mathcal{A} with block rows $[A_{c,i}, A_{x,i}, 0]$. The least-squares solution

$$\widehat{\theta}_X = \arg\min_{\theta} \sum_{i=1}^{N_x} ||X_i - A_{c,i} \, \theta_c - A_{x,i} \, \theta_x||^2$$

omits θ_y entirely. Consequently, the Fisher information on θ_y vanishes, making it unidentifiable.

Analogously, stacking $\{Y_j\}_{j=1}^{N_y}$ defines $\mathcal B$ with block rows $[B_{c,j},\,0,\,B_{y,j}]$ and yields

$$\widehat{\theta}_Y = \arg\min_{\theta} \sum_{j=1}^{N_y} ||Y_j - B_{c,j} \theta_c - B_{y,j} \theta_y||^2.$$

This estimator doesn't depend on θ_x , providing zero coverage for that component. Thus, each unimodal estimator entirely fails to recover the parameters exclusive to the omitted modality.

Multimodal Estimators. Despite the lack of one-to-one pairing, both $\{X_i\}$ and $\{Y_j\}$ share the common parameters θ_c . Since the two distributions are conditionally independent, the joint likelihood factorizes as

$$\prod_{i=1}^{N_x} p(X_i \mid \theta_c, \theta_x) \times \prod_{j=1}^{N_y} p(Y_j \mid \theta_c, \theta_y).$$

Maximizing this yields the combined estimator

$$\widehat{\theta}_{X,Y} = \arg\min_{\theta_c, \theta_x, \theta_y} \left\{ \sum_{i=1}^{N_x} \|X_i - A_{c,i} \, \theta_c - A_{x,i} \, \theta_x\|^2 + \sum_{j=1}^{N_y} \|Y_j - B_{c,j} \, \theta_c - B_{y,j} \, \theta_y\|^2 \right\}.$$

Intuitively, there is no requirement to match up individual (X_i, Y_j) pairs. Instead, the estimate for θ_c is improved by both modalities while remaining unpaired.

Fisher Information. In our linear model, each dataset contributes block-structured Fisher information. For the X-dataset:

$$I_X = \sum_{i=1}^{N_x} \begin{pmatrix} A_{c,i}^\top A_{c,i} & A_{c,i}^\top A_{x,i} & 0 \\ A_{x,i}^\top A_{c,i} & A_{x,i}^\top A_{x,i} & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

and for the Y-dataset:

$$I_Y = \sum_{j=1}^{N_y} \begin{pmatrix} B_{c,j}^\top B_{c,j} & 0 & B_{c,j}^\top B_{y,j} \\ 0 & 0 & 0 \\ B_{y,j}^\top B_{c,j} & 0 & B_{y,j}^\top B_{y,j} \end{pmatrix}.$$

Because X and Y samples are independent, their curvature contributions add pointwise, resulting in the joint Fisher information being simply the sum of the unimodal blocks.

$$I_{X,Y} = I_X + I_Y = \begin{pmatrix} \sum_i A_{c,i}^\top A_{c,i} + \sum_j B_{c,j}^\top B_{c,j} & * & * \\ * & \sum_i A_{x,i}^\top A_{x,i} & 0 \\ * & 0 & \sum_j B_{y,j}^\top B_{y,j} \end{pmatrix},$$

where "*" denotes the cross-modal blocks. In particular, we have the shared-parameter block as

$$(I_{X,Y})_{\theta_c,\theta_c} = \sum_{i=1}^{N_x} A_{c,i}^{\top} A_{c,i} + \sum_{j=1}^{N_y} B_{c,j}^{\top} B_{c,j},$$

C.3 THEOREMS AND PROOFS

The aim of this section is to detail the proofs of the theoretical results presented in the main manuscript The key theoretical tools driving our analysis are already prepared in Appendix C.1 and Appendix C.2. Core to our theoretical analysis are a few lemmas around the Loewner-order monotonicity result for inverses that we prove below.

Lemma 1 (Loewner Order reversal for inverses). Let $M, N \in \mathbb{S}^d_{\succ 0}$ with $M \prec N$ (or $M \preceq N$). Then $N^{-1} \prec M^{-1}$ (or $N^{-1} \preceq M^{-1}$).

Proof. Since $N\succ 0$, $N^{-1/2}$ exists and is nonsingular. Define $C:=N^{-1/2}MN^{-1/2}\prec I$. Because a congruence with an invertible matrix preserves positive-definiteness, $C\succ 0$; hence C^{-1} is well defined and $C^{-1}\succ I$ (the scalar map $x\mapsto x^{-1}$ is strictly decreasing on $(0,\infty)$). Undoing the congruence gives

$$M^{-1} = N^{-1/2}C^{-1}N^{-1/2} \succ N^{-1/2}IN^{-1/2} = N^{-1}.$$

Lemma 2 (Inverse–monotonicity of the Moore–Penrose pseudoinverse). Let $M, N \in \mathbb{S}^d_{\succeq 0}$ satisfy $M \prec N$ and $\ker M = \ker N =: K$. Then their pseudoinverses obey $N^{\dagger} \prec M^{\dagger}$.

Proof. Set $S := K^{\perp}$ and let $P := P_S$ be the orthogonal projector onto S. Because M and N vanish on K, we have the decompositions M = PMP and N = PNP. Restricted to S both matrices are positive–definite:

$$\tilde{M} := PMP, \qquad \tilde{N} := PNP \in \mathbb{S}_{\succ 0}^{\dim S}, \quad \tilde{M} \prec \tilde{N}.$$

Apply Lemma 1 to \tilde{M} , \tilde{N} to obtain $\tilde{N}^{-1} \prec \tilde{M}^{-1}$ on S. The Moore–Penrose pseudoinverse equals the ordinary inverse on S and is zero on K:

$$M^{\dagger} = P\tilde{M}^{-1}P, \qquad N^{\dagger} = P\tilde{N}^{-1}P.$$

Therefore
$$N^{\dagger}=P\tilde{N}^{-1}P\prec P\tilde{M}^{-1}P=M^{\dagger}.$$

Lemma 3 (Directional Loewner Order reversal). Let $M, N \in \mathbb{S}^d_{\succ 0}$ with $M \leq N$. If a non-zero vector v satisfies $v^{\top}Mv < v^{\top}Nv$, then

- 1. For the vector v, it holds that $v^{\top}M^{-1}v \geq v^{\top}N^{-1}v$, with strict inequality $v^{\top}M^{-1}v > v^{\top}N^{-1}v$ if and only if $(N-M)M^{-1}v \neq 0$.
- 2. There exists a non-zero vector $u \in \mathbb{R}^d$ such that $u^T M^{-1} u > u^T N^{-1} u$.

Proof. Denote the Loewner gap $\Delta := N - M \succeq 0$. Then, the assumption $v^{\top}Nv > v^{\top}Mv$ is equivalent to $v^{\top}\Delta v > 0$. Introduce the congruence–invariant normalisation $C := M^{-1/2}\Delta M^{-1/2} \succeq 0$. Now, using $\Delta = M^{1/2}CM^{1/2}$ and properties of inverse,

$$N = M^{1/2}(I+C)M^{1/2}, \qquad N^{-1} = M^{-1/2}(I+C)^{-1}M^{-1/2},$$

since $I + C \succ 0$ (because $C \succeq 0$ and $I \succ 0$). Thus,

$$M^{-1} - N^{-1} = M^{-1/2} \Big[I - (I+C)^{-1} \Big] M^{-1/2}$$
$$= M^{-1/2} C (I+C)^{-1} M^{-1/2},$$

because $(I - (I + C)^{-1})(I + C) = C$. Finally, evaluating in the direction v, we have

$$\begin{split} v^\top (M^{-1} - N^{-1}) v &= v^\top M^{-1/2} (I + C)^{-1} C M^{-1/2} v \\ &= u^\top (I + C)^{-1} C u \quad (\text{where } u = M^{-1/2} v) \end{split}$$

Now, since $(I+C)^{-1} \in \mathbb{S}_{\geq 0}$ and $C \in \mathbb{S}_{\geq 0}$ commute, the matrix $(I+C)^{-1}C$ is positive semidefinite and it has exactly the same kernel as C. Thus, if $C = Q \operatorname{diag}(\lambda_i)Q^{\top}$ ($\lambda_i \geq 0$), we have

$$u^{\top}C(I+C)^{-1}u = \sum_{i} \frac{\lambda_{i}}{1+\lambda_{i}} (Q^{\top}u)_{i}^{2} \ge 0.$$

This expression is strictly positive exactly when u has a component in any eigen-subspace with $\lambda_i>0$ i.e when $u\not\in\ker(C)$. Since $M^{-1/2}\in\mathbb{S}_{\succ 0},\,Cu=0\implies M^{-1/2}\Delta M^{-1/2}u=0\implies \Delta M^{-1}v=0$. Thus, this expression is strictly positive if $\Delta M^{-1}v\not=0$.

Now, from the premise $v^\top \Delta v > 0$, it follows that $\Delta \neq 0$. Since $M \succ 0$, $M^{-1/2}$ is invertible, C is also not the zero matrix. Since $C \succeq 0$, this means that C must have at least one strictly positive eigenvalue. Let $\lambda > 0$ be such an eigenvalue, and let $z \neq 0$ be a corresponding eigenvector. Define, $x := M^{1/2}z \neq 0$. Thus, we have $x^\top (M^{-1} - N^{-1})x = z^\top C(I + C)^{-1}z = \frac{\lambda}{1+\lambda}\|z\|^2 > 0$, showing the existence of a non-zero vector x such that $x^\top M^{-1}x > x^\top N^{-1}x$.

Theorem 1 (Restatement of Theorem 1). Let $\hat{\theta}_X$, $\hat{\theta}_Y$ be the least-squares estimators for θ using only $\{X_i\}$ and only $\{Y_j\}$ and let $\hat{\theta}_{X,Y}$ be the joint estimator using both unpaired datasets. Then, under the assumption that at least one $B_{c,j}$ where $j \in \{1,2,...N_y\}$ has full rank, the commonfactor covariance satisfies the strict Loewner ordering i.e. $\operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c} \prec \operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c}$, or equivalently, the Fisher information on θ_c strictly increases when combining both modalities, despite not having sample-wise pairing: $(I_X + I_Y)_{\theta_c,\theta_c} \succ (I_X)_{\theta_c,\theta_c}$.

Proof. For any statistic $S(\theta) = \nabla_{\theta} \log p(x; \theta)$ and vector v,

$$\boldsymbol{v}^{\top} \boldsymbol{I}(\boldsymbol{\theta}) \, \boldsymbol{v} = \boldsymbol{v}^{\top} \mathbb{E}[\boldsymbol{S}(\boldsymbol{\theta}) \boldsymbol{S}(\boldsymbol{\theta})^{\top}] \, \boldsymbol{v} = \mathbb{E}\big[(\boldsymbol{v}^{\top} \boldsymbol{S}(\boldsymbol{\theta}))^2\big] \; \geq \; 0.$$

Thus, a Fisher Information Matrix is a positive semidefinite matrix.

In our linear–Gaussian model, the X-dataset contributes $(I_X)_{\theta_c,\theta_c} = \sum_{i=1}^{N_x} A_{c,i}^{\top} A_{c,i}$ and the Y-dataset gives $(I_Y)_{\theta_c,\theta_c} = \sum_{j=1}^{N_y} B_{c,j}^{\top} B_{c,j}$. Since at least one $B_{c,j}$ has full column rank, $(I_Y)_{\theta_c,\theta_c}$ is positive-definite on the θ_c subspace. Now, if at least one $B_{c,j} \in \mathbb{R}^{m \times d_c}$ has full column rank d_c , then for any $v \in \mathbb{R}^{d_c} \setminus \{0\}$,

$$v^{\top} B_{c,j}^{\top} B_{c,j} v = ||B_{c,j} v||^2 > 0.$$

Hence, each summand in $(I_Y)_{\theta_c,\theta_c}$ is positive semidefinite and at least one is positive definite, so their sum $\sum_i B_{c,j}^{\top} B_{c,j}$ is positive definite on the θ_c subspace. Thus,

$$(I_X)_{\theta_c,\theta_c} \prec (I_X)_{\theta_c,\theta_c} + (I_Y)_{\theta_c,\theta_c} = (I_X + I_Y)_{\theta_c,\theta_c}$$

Now, for regular exponential families (including Gaussian linear models), the covariance matrix of the maximum likelihood estimator $\widehat{\theta}$ near the true θ_0 is (asymptotically) the inverse of the Fisher information matrix i.e. $\operatorname{Var}(\widehat{\theta}) \approx I(\theta_0)^{-1}$. Precisely, as the sample size $n \to \infty$, we have:

$$\sqrt{n}(\hat{\theta} - \theta_0) \stackrel{d}{\to} \mathcal{N}(0, I(\theta_0)^{-1}),$$

where θ_0 is the true parameter value, $I(\theta_0)$ is the Fisher Information Matrix evaluated at θ_0 and $\mathcal{N}(0, I(\theta_0)^{-1})$ denotes a multivariate normal distribution with mean 0 and covariance matrix $I(\theta_0)^{-1}$. Thus, we compare variances via the Moore–Penrose pseudoinverse of the information matrices.

Let $M_X=(I_X)_{\theta_c,\theta_c}$, $M_Y=(I_Y)_{\theta_c,\theta_c}$ and $M_{X,Y}=(I_X+I_Y)_{\theta_c,\theta_c}$. Since $M_Y\succ 0$, $M_{X,Y}=M_X+M_Y$ is also positive definite (as $M_{X,Y}\succeq M_Y\succ 0$). Thus, $\mathrm{Var}(\hat{\theta}_{X,Y})=M_{X,Y}^{-1}$. We have established $M_X\prec M_{X,Y}$. Assuming M_X is positive definite (to define the matrix $\mathrm{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c}$), we apply Lemma 1 to get $M_{X,Y}^{-1}\prec M_X^{-1}$. Thus,

$$\operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_{\alpha},\theta_{\alpha}} = M_{X,Y}^{-1} \prec M_{X}^{-1} = \operatorname{Var}(\hat{\theta}_{X})_{\theta_{\alpha},\theta_{\alpha}},$$

This proves the statement under the condition that M_X is positive definite. Note here that, on spaces unidentifiable by X-alone i.e. $v \in \ker(M_X)$, we have $\operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c} = \infty$. Since $M_{X,Y}$ is positive definite, it has finite variance along such v i.e. $\operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c} < \infty$, thus strictly reducing the variance of the estimator. Thus, adding the unpaired Y-modality strictly reduces the variance (or, dually, increases the Fisher information) on the common factors θ_c .

Theorem 2 (Restatement of Theorem 2). Let all notation be as in Theorem 1, and define $M_X := (I_X)_{\theta_c,\theta_c}$, $M_Y := (I_Y)_{\theta_c,\theta_c}$, and $M_{XY} := M_X + M_Y$. Let $v \in \mathbb{R}^{d_c} \setminus \{0\}$. If there exists at least one index $j \in \{1,2,...N_y\}$ such that $B_{c,j}v \neq 0$, then the following hold:

- 1. The Fisher information strictly increases in direction v i.e. $v^{\top}M_{XY}v > v^{\top}M_{X}v$.
- 2. The variance of the estimator in direction v is strictly reduced i.e $v^{\top} \operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c} v < v^{\top} \operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c} v$, if $v \notin \operatorname{range}(M_X)$. For $v \in \operatorname{range}(M_X)$, this strict inequality holds for v under an additional invertibility condition and is always guaranteed for some $u \in \operatorname{range}(M_X)$ i.e. $\exists u \text{ s.t. } u^{\top} \operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c} u < u^{\top} \operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c} u$.

Proof. Define $M_X := (I_X)_{\theta_c,\theta_c}$, $M_Y := (I_Y)_{\theta_c,\theta_c}$, $and M_{XY} := M_X + M_Y$. By assumption, $\exists j$ such that $B_{c,j}v \neq 0$. Thus:

$$v^{\top} M_Y v = \sum_{j=1}^{N_y} \|B_{c,j} v\|^2 \ge \|B_{c,j} v\|^2 > 0.$$

Hence M_Y is positive-definite in direction v, implying $M_{X,Y} \succ M_X$ in this direction:

$$\boldsymbol{v}^{\top} \boldsymbol{M}_{XY} \boldsymbol{v} = \boldsymbol{v}^{\top} \boldsymbol{M}_{X} \boldsymbol{v} + \boldsymbol{v}^{\top} \boldsymbol{M}_{Y} \boldsymbol{v} > \boldsymbol{v}^{\top} \boldsymbol{M}_{X} \boldsymbol{v},$$

thus proving the first part of the theorem.

Case 1: $v \notin \operatorname{Range}(M_X)$. If $v \notin \operatorname{Range}(M_X)$, then v has a non-zero component in $\ker(M_X)$. Let $v = v_S + v_K$, where $v_S \in \operatorname{Range}(M_X)$ and $v_K \in \ker(M_X)$ with $v_K \neq 0$. The linear combination of parameters $v^\top \theta_c = v_S^\top \theta_c + v_K^\top \theta_c$. Since $v_K \in \ker(M_X)$, the component $v_K^\top \theta_c$ is not identifiable by the X-only model. Consequently, the asymptotic variance of an unbiased estimator for $v^\top \theta_c$ using only the X-dataset is infinite. We denote this as $v^\top \operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c} v = \infty$.

The strict inequality $v^{\top}M_{XY}v>0$, ensures that $v\notin\ker(M_{XY})$, and thus $v\in\operatorname{Range}(M_{XY})$. Since $v\in\operatorname{Range}(M_{XY})$ and $v\neq 0$, $M_{XY}^{\dagger}v$ is well-defined. Furthermore, because M_{XY} is positive semidefinite, M_{XY}^{\dagger} is also positive semidefinite and shares the same kernel as M_{XY} (since M_{XY} is symmetric). As $v\neq 0$ and $v\notin\ker(M_{XY})$, thus $v\notin\ker(M_{XY}^{\dagger})$, which ensures $v^{\top}M_{XY}^{\dagger}v$ is a finite positive value. Thus,

 $v^{\top} \operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_0,\theta_0} v < \infty.$

Comparing this to the variance from the X-only model in this case:

$$v^{\top} \operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c} v < \infty = v^{\top} \operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c} v,$$

and the strict inequality holds.

Case 2: $v \in \text{Range}(M_X)$. Let $S := \text{Range}(M_X)$ and let P_S be the orthogonal projector onto S. Because $M_X = M_X P_S$ and $M_{XY} = M_X + M_Y$, the restrictions

$$\tilde{M}_X := P_S M_X P_S, \qquad \tilde{M}_{XY} := P_S M_{XY} P_S = \tilde{M}_X + P_S M_Y P_S$$

are positive-definite on S; To see this, take any non-zero $w \in S$. Since $w \in \text{range}(M_X)$, $P_S w = w$; hence

$$w^{\top} \tilde{M}_X w = w^{\top} M_X w > 0$$
 (P_S is identity when restricted to S)

Thus $\tilde{M}_X \succ 0$ on S. Because $P_S M_Y P_S \succeq 0$, adding it preserves positive-definiteness, so

$$\tilde{M}_{XY} = \tilde{M}_X + P_S M_Y P_S \succeq \tilde{M}_X \succ 0$$
 on S .

Applying Lemma 3(1) to \tilde{M}_X and \tilde{M}_{XY} on S gives us $v^\top \tilde{M}_{XY}^{-1} v \leq v^\top \tilde{M}_X^{-1} v$. Strict inequality $v^\top \tilde{M}_{XY}^{-1} v < v^\top \tilde{M}_X^{-1} v$ holds if and only if the condition $C_v := ((\tilde{M}_{XY} - \tilde{M}_X) \tilde{M}_X^{-1} v \neq \mathbf{0})$ is met. Therefore, if condition C_v holds, the directional variance along this constrained space S is strictly reduced:

$$v^{\top} \operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_{\alpha},\theta_{\alpha}} v = v^{\top} \tilde{M}_{XY}^{-1} v < v^{\top} \tilde{M}_{Y}^{-1} v = v^{\top} \operatorname{Var}(\hat{\theta}_{X})_{\theta_{\alpha},\theta_{\alpha}} v^{1}.$$

Further, from Lemma 3(2), there exists some non-zero vector $u \in S$ such that $u^{\top} \tilde{M}_{XY}^{-1} u < u^{\top} \tilde{M}_{X}^{-1} u$. Thus we have,

$$u^{\top} \operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c} u < u^{\top} \operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c} u.$$

Thus, completing the proof.

Corollary 1. Assume a direction $v \in \mathbb{R}^{d_c} \setminus \{0\}$ with $a = v^{\top}(I_X)_{\theta_c,\theta_c} v > 0$ and $b = v^{\top}(I_Y)_{\theta_c,\theta_c} v > 0$ where v is the common eigenvector of $(I_X)_{\theta_c,\theta_c}$ and $(I_Y)_{\theta_c,\theta_c}$. Then the variance in direction v contracts by the factor

$$\frac{v^{\top} \text{Var}(\hat{\theta}_{X,Y}) \, v}{v^{\top} \text{Var}(\hat{\theta}_{X}) \, v} \; = \; \frac{1/(a+b)}{1/a} \; = \; \frac{a}{a+b} \; < \; 1,$$

So the joint estimator achieves strictly lower error along v.

Proof. Let $M_X=(I_X)_{\theta_c,\theta_c}$ and $M_Y=(I_Y)_{\theta_c,\theta_c}$. By assumption, v is a common eigenvector of M_X and M_Y . Thus, $M_Xv=\lambda_Xv$ and $M_Yv=\lambda_Yv$ for some eigenvalues λ_X and λ_Y . From the assumptions, we have $\lambda_X=a/\|v\|^2>0$ and $\lambda_Y=b/\|v\|^2>0$. Since M_X is symmetric and $M_Xv=\lambda_Xv$ with $\lambda_X>0$, the pseudoinverse acts as $M_X^\dagger v=\lambda_X^{-1}v$. Therefore, the variance in direction v for the X-only estimator is

$$v^{\top} \operatorname{Var}(\hat{\theta}_X)_{\theta_C,\theta_C} v = v^{\top} M_X^{\dagger} v = v^{\top} (\lambda_X^{-1} v) = \lambda_X^{-1} ||v||^2 = a^{-1} ||v||^4.$$

Since v is a common eigenvector, it is also an eigenvector of $M_{XY} = M_X + M_Y$:

$$(M_X + M_Y)v = M_Xv + M_Yv = \lambda_Xv + \lambda_Yv = (\lambda_X + \lambda_Y)v.$$

¹We note that true asymptotic variance defined as $v^{\top} \text{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c} v = v^{\top} M_{XY}^{\dagger} v$, $v^{\top} M_{XY}^{\dagger} v = v^{\top} \tilde{M}_{XY}^{-1} v$ if S is an invariant subspace of M_{XY} and M_{XY} is block-diagonal with respect to S and S^{\perp} (i.e., $P_S M_{XY} P_{S^{\perp}} = \mathbf{0}$, which implies $P_S M_Y P_{S^{\perp}} = \mathbf{0}$).

The corresponding eigenvalue is $\lambda_{XY} = \lambda_X + \lambda_Y$. Since $\lambda_X > 0$ and $\lambda_Y > 0$, $\lambda_{XY} > 0$. Thus, $(M_X + M_Y)^{\dagger}v = (\lambda_X + \lambda_Y)^{-1}v$. The variance in direction v for the joint estimator is

$$v^{\top} \operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c} v = v^{\top} (M_X + M_Y)^{\dagger} v = (\lambda_X + \lambda_Y)^{-1} ||v||^2 = (a+b)^{-1} ||v||^4.$$

Now, we form the ratio of these variances:

$$\frac{v^\top \mathrm{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c}\,v}{v^\top \mathrm{Var}(\hat{\theta}_{X})_{\theta_c,\theta_c}\,v} = \frac{\lambda_X}{\lambda_X + \lambda_Y} = \frac{a}{a+b} < 1.$$

Corollary 2. Assume a direction $v \in \mathbb{R}^{d_c} \setminus \{0\}$ with $v^{\top}(I_X)_{\theta_c,\theta_c} v = 0$ and $v^{\top}(I_Y)_{\theta_c,\theta_c} v > 0$. Then $v^{\top} \operatorname{Var}(\hat{\theta}_X) v = \infty$ and $v^{\top} \operatorname{Var}(\hat{\theta}_{X,Y}) v < \infty$ i.e. a direction unidentifiable from X alone becomes well-posed with even unpaired data from Y.

Proof. This corollary follows directly from Case 1 of Theorem 2. The condition $v^{\top}(I_X)_{\theta_c,\theta_c}v=0$ for $v\neq 0$ implies $v\in \ker((I_X)_{\theta_c,\theta_c})$, and thus $v\not\in \operatorname{range}((I_X)_{\theta_c,\theta_c})$. Given the additional condition $v^{\top}(I_Y)_{\theta_c,\theta_c}v>0$, the conclusions of Case 1 of the theorem apply directly.

Corollary 3 (Variance Reduction for Eigenvectors of M_X). Let $v \in \mathbb{R}^{d_c} \setminus \{0\}$ be an eigenvector of $M_X = (I_X)_{\theta_c,\theta_c}$ with a corresponding eigenvalue $\lambda_X > 0$. If the Y-dataset provides information in this direction v (i.e., $v^\top M_Y v > 0$, where $M_Y = (I_Y)_{\theta_c,\theta_c}$), then the variance in direction v is strictly reduced by incorporating the Y-dataset:

$$v^{\top} \operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c} v < v^{\top} \operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c} v.$$

Specifically,
$$v^{\top} \operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c} v = \lambda_X^{-1} ||v||^2$$
.

Proof. Let $M_X = (I_X)_{\theta_c,\theta_c}$ and $M_Y = (I_Y)_{\theta_c,\theta_c}$. Since v is an eigenvector of M_X with a positive eigenvalue $\lambda_X > 0$, it follows that $v \in \text{Range}(M_X)$. Let $S = \text{Range}(M_X)$. The variance using only the X-dataset in direction v is given by

$$v^{\top} \operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c} v = v^{\top} M_X^{\dagger} v.$$

Because v is an eigenvector of M_X with $\lambda_X > 0$, $M_X^{\dagger} v = \lambda_X^{-1} v$. Thus,

$$v^{\top} \operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c} v = v^{\top}(\lambda_X^{-1}v) = \lambda_X^{-1} ||v||^2.$$

This scenario falls under Case 2 of Theorem 2, specifically its conclusion regarding $v \in S$. According to that theorem, strict variance reduction $v^\top \mathrm{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c} v < v^\top \mathrm{Var}(\hat{\theta}_X)_{\theta_c,\theta_c} v$ occurs if the condition $C_v = ((P_S M_Y P_S)(M_X|_S)^{-1} v \neq \mathbf{0})$ holds. Here, P_S is the orthogonal projector onto S, and $M_X|_S$ is the restriction of M_X to S, so $(M_X|_S)^{-1} v = \lambda_X^{-1} v$.

The condition C_v thus becomes $(P_S M_Y P_S)(\lambda_X^{-1} v) \neq \mathbf{0}$. Since $\lambda_X > 0$, this is equivalent to $P_S M_Y P_S v \neq \mathbf{0}$. We are given that $v^\top M_Y v > 0$. As $v \in S$, $P_S v = v$. Therefore, $v^\top M_Y v = v^\top P_S M_Y P_S v > 0$. Let $A_S = P_S M_Y P_S$ restricted to S. A_S is a positive semidefinite operator on S. The condition $v^\top A_S v > 0$ for $v \in S$, $v \neq 0$ implies that $A_S v \neq \mathbf{0}$ (because if $A_S v = \mathbf{0}$, then $v^\top A_S v = 0$, which contradicts $v^\top A_S v > 0$). Thus, $P_S M_Y P_S v \neq \mathbf{0}$, which means the condition C_v is satisfied.

Since $v \in S$ and the condition C_v for strict inequality is met, by Theorem 2, it follows that $v^{\top} \operatorname{Var}(\hat{\theta}_{X,Y})_{\theta_c,\theta_c} v < v^{\top} \operatorname{Var}(\hat{\theta}_X)_{\theta_c,\theta_c} v$.

Theorem 3 (Restatement of Theorem 3). Define for any m, $I_X^{(m)} = \sum_{i=1}^m A_{c,i}^{\top} A_{c,i}$ and $I_Y^{(m)} = \sum_{j=1}^m B_{c,j}^{\top} B_{c,j}$. If $\operatorname{range}(I_Y^{(m)}) \not\subseteq \operatorname{range}(I_X^{(m)})$, then there exists a nonzero $v \in \mathbb{R}^{d_c}$ such that $v^{\top} I_Y^{(m)} v > v^{\top} I_X^{(m)} v$.

 Proof. Let $R_X := \operatorname{range}(I_X^{(m)})$, $R_Y := \operatorname{range}(I_Y^{(m)})$. By the assumption $R_Y \not\subseteq R_X$, choose a vector $w \in R_Y \setminus R_X$. Since \mathbb{R}^{d_c} is a finite dimensional inner product space and R_X is its finite dimensional subspace, we can decompose $w = w_{||} + v$ with $w_{||} \in R_X$ and $v \in R_X^{\perp}$. Because $w \notin R_X$, the orthogonal component v is non-zero.

(i) Term from $I_X^{(m)}$. From the Fundamental Theorem of Linear Algebra, for any symmetric matrix S, $\ker S = \operatorname{range}(S)^{\perp}$; hence $R_X^{\perp} = \ker I_X^{(m)}$. Thus

$$v^{\top} I_{\mathbf{x}}^{(m)} v = 0.$$

(ii) Term from $I_Y^{(m)}$. Because $w \in R_Y = \mathrm{range}(I_Y^{(m)})$, there exists u with $w = I_Y^{(m)}u$. Suppose, for contradiction, that $I_Y^{(m)}v = 0$. Then $v \in \ker I_Y^{(m)} = R_Y^{\perp}$, so $v \perp w$. But $w \cdot v = (w_{\parallel} + v) \cdot v = w_{\parallel} \cdot v + \|v\|^2 = \|v\|^2 > 0$ because $v \perp w_{\parallel}$ while $v \neq 0$. This contradicts $v \perp w$; therefore $I_Y^{(m)}v \neq 0$ and, by positive semidefiniteness,

$$v^{\top} I_V^{(m)} v > 0.$$

Combining the above inequalities yields $v^{\top}I_{Y}^{(m)}v>v^{\top}I_{X}^{(m)}v$, with $v\neq 0$, which is the desired inequality. \Box

D UML: ALGORITHM PSEUDOCODE

In this section we present the full pseudocode for UML for both the self-supervised and supervised settings as shown in Algorithm 1 and Algorithm 2 respectively.

Algorithm 1 Pytorch Pseudocode for UML in the self-supervised setting

```
1484
       # f_img, f_text: image encoder, text encoder
1485
       # g_img, g_text: image decoder, text decoder
1486
       # h: shared backbone
1487
1488
       while not converged: # training loop
          x_img = fetch_next(image_loader) # image minibatch
1489
          x_text = fetch_next(text_loader) # text minibatch (random/unaligned)
1491
          x_img = patchify_and_embed(x_img) # input image patch embeddings
1492
          x_text = tokenize_and_embed(x_text) # input text token embeddings
1493
          z_img = f_img(x_img) # image patch embeddings
1494
          z_{\text{text}} = f_{\text{text}}(x_{\text{text}}) + \text{text token embeddings}
1495
1496
          y_img = g_img(h(z_img)) # predict image patch embeddings
1497
          y_text = g_text(h(z_text)) # predict text patch embeddings
1498
          # Next Token/Patch Embedding Prediction Loss
1499
          loss_img = MSE(y_img[:,:-1,:], x_img[:,1:,:])
1500
          loss_text = MSE(y_text[:,:-1,:], x_text[:,1:,:])
1501
          loss = loss_img + lambda * loss_text # total loss
1502
1503
          loss.backward() # back-propagate
          update(h, f_img, f_text, g_img, g_text) # SGD update
1504
       # Define Mean Squared Error loss
       def MSE(pred, target):
1507
          return ((pred - target) ** 2).mean()
```

1513

153615371538

1539 1540

1541 1542

1543

1545

1547

1548

1549

1550 1551

1552

1553

1554

1555 1556

1557

1561

1563

1564

1565

Algorithm 2 Pytorch Pseudocode for UML in the supervised setting

```
1514
       # f_imq: image encoder (frozen or trainable)
       # f_text: text encoder (frozen)
1515
       # is_trainable: True if f_img is trainable else False
1516
       # h: classification head
1517
1518
      while not converged: # training loop
1519
          x_img = fetch_next(image_loader) # image minibatch
          x_text = fetch_next(text_loader) # text minibatch (random/unaligned)
1520
1521
          z_{img} = f_{img}(x_{img}) # image embeddings
1522
          z_{\text{text}} = f_{\text{text}}(x_{\text{text}}) # text embeddings
1523
          logits_img = h(z_img) # predict image labels
1524
          logits_text = h(z_text) # predict text labels
1525
1526
          loss_img = CE(logits_img, labels_img) # image classification loss
          loss_text = CE(logits_text, labels_text) # text classification loss
1528
          loss = loss_img + lambda * loss_text # total loss
1529
          loss.backward() # back-propagate
1530
          update(h, f_img) if is_trainable else update(h) # SGD update
1531
1532
       # Define Cross-Entropy loss
1533
      def CE(logits, labels):
1534
          return -sum(labels * log_softmax(logits, dim=1)) / len(labels)
1535
```

E ADDITIONAL EXPERIMENTS

E.1 IMPROVING IMAGE CLASSIFICATION USING UNPAIRED TEXTS (UNALIGNED ENCODERS)

In this section we report image-classification results on ten benchmarks (see Appendix B.3), covering three settings:

- 1. **Full-dataset fine-tuning**: train both the vision backbone and classification head (Appendix E.1.1).
- 2. **Full-dataset linear probe**: train only the classification head (Appendix E.1.2).
- Few-shot linear probe: train only the classification head under few-shot conditions (Appendix E.1.3).

In each setting, we compare UML with baselines across all datasets and multiple DINO-initialized vision backbones. Our method has two variants: Ours (UML), where we alternately train with both image and unpaired text data (see Algorithm 2), and Ours (init) where we initialize the classifier with the average text embedding of each class, providing a strong prior to align image and class level information.

E.1.1 SUPERVISED FINETUNING (ACROSS ARCHITECTURES)

In this section, we fine-tune both the vision backbone and the linear classifier on ten downstream tasks, comparing UML against strong image-only baselines. We evaluate four DINO-initialized backbones:

- ViT-B/16 in Table 6
- ViT-B/8 in Table 7
- DINOv2 ViT-S/14 in Table 8
- DINOv2 ViT-B/14 in Table 9

Results for DINOv2 ViT-L/14 are omitted due to computational constraints. Across all backbones, UML consistently improves over the image-only baseline by leveraging unpaired text embeddings. For some backbones such as DINOv2 VIT-B/16, our head-initialization variant (*Ours* (*init*)) outperforms training using unpaired multimodal data from scratch (*Ours*), while in others it does not.

Table 6: **Full finetuning on classification with ViT-B/16 DINO and OpenLLaMA-3B**. We compare our proposed approach with the image-only baseline when fine-tuning on the target dataset. All vision encoders are initialized from DINO weights, and our approach leverages unpaired text data using OpenLLaMA-3B embeddings.

					Dat	aset				
Method	Stanford Cars	SUN397	FGVC Aircraft	DTD	UCF101	Food101	Oxford Pets	Oxford Flowers	Caltech101	Average
Unimodal Ours Ours (init)	78.41 82.56 81.95	63.99 67.04 67.12	62.12 67.38 68.29	74.17 76.42 73.84	81.43 84.06 84.31	82.38 81.79 81.12	92.00 93.20 92.60	98.24 98.98 98.73	96.31 97.04 96.84	81.01 83.16 82.76

Table 7: Full finetuning on classification with ViT-B/8 DINO and OpenLLaMA-3B. We compare our proposed approach with the image-only baseline when fine-tuning on the target dataset. All vision encoders are initialized from DINO weights, and our approach leverages unpaired text data using OpenLLaMA-3B embeddings.

					Dat	aset				
Method	Stanford Cars	SUN397	FGVC Aircraft	DTD	UCF101	Food101	Oxford Pets	Oxford Flowers	Caltech101	Average
Unimodal Ours Ours (init)	85.67 87.95 87.44	68.04 70.28 70.03	72.60 75.31 76.09	76.65 77.19 76.24	83.94 85.59 86.49	85.32 84.83 84.71	93.06 93.05 93.81	99.22 99.43 99.27	96.82 97.12 97.16	84.59 85.64 85.69

Table 8: Full finetuning on classification with ViT-S/14 DINOv2 and OpenLLaMA-3B. We compare our proposed approach with the image-only baseline when fine-tuning on the target dataset. All vision encoders are initialized from DINOv2 weights, and our approach leverages unpaired text data using OpenLLaMA-3B embeddings.

					Dat	aset				
Method	Stanford Cars	SUN397	FGVC Aircraft	DTD	UCF101	Food101	Oxford Pets	Oxford Flowers	Caltech101	Average
Unimodal Ours Ours (init)	79.45 84.87 86.39	66.20 66.72 66.03	66.99 71.54 73.44	72.16 74.14 74.27	83.18 84.77 84.69	80.65 81.16 81.97	90.67 91.87 91.72	99.18 99.55 99.82	95.45 97.03 97.60	81.54 83.52 83.99

Table 9: Full finetuning on classification with ViT-B/14 DINOv2 and OpenLLaMA-3B. We compare our proposed approach with the image-only baseline when fine-tuning on the target dataset. All vision encoders are initialized from DINOv2 weights, and our approach leverages unpaired text data using OpenLLaMA-3B embeddings.

					Dat	aset				
Method	Stanford Cars	SUN397	FGVC Aircraft	DTD	UCF101	Food101	Oxford Pets	Oxford Flowers	Caltech101	Average
Unimodal Ours Ours (init)	89.62 90.93 90.73	71.45 70.97 70.92	77.29 80.02 80.23	73.88 75.83 75.87	88.00 87.52 87.6 0	82.94 86.25 83.43	94.55 94.74 94.47	99.88 99.88 99.80	97.69 97.57 97.93	86.14 87.08 86.77

E.1.2 LINEAR PROBING (ACROSS ARCHITECTURES)

In this section, we train only the linear classifier, on top of the frozen vision and language backbone, on ten downstream tasks, comparing UML against strong image-only baselines. We evaluate five DINO-initialized backbones:

- ViT-B/16 in Table 10
- ViT-B/8 in Table 11
- DINOv2 ViT-S/14 in Table 12
- DINOv2 ViT-B/14 in Table 13
- DINOv2 ViT-L/14 in Table 14

Across all backbones, UML consistently improves over the image-only baseline by leveraging unpaired text embeddings. For all backbones, our head-initialization variant (*Ours* (*init*)) outperforms training using unpaired multimodal data from scratch (*Ours*).

Table 10: Full linear probing on classification with ViT-B/16 DINO and OpenLLaMA-3B. We compare our proposed approach with the image-only baseline when training a linear probe on the target dataset. All vision encoders are initialized from DINO weights, and our approach leverages unpaired text data using OpenLLaMA-3B embeddings.

					Dat	aset				
Method	Stanford Cars	SUN397	FGVC Aircraft	DTD	UCF101	Food101	Oxford Pets	Oxford Flowers	Caltech101	Average
Unimodal Ours Ours (init)	67.10 68.71 68.60	64.63 65.14 65.59	56.02 57.42 57.98	72.42 72.95 73.11	81.27 82.06 82.40	74.96 75.30 75.73	93.07 93.18 93.62	98.32 98.46 98.42	95.01 96.19 96.35	78.08 78.82 79.09

Table 11: Full linear probing on classification with ViT-B/8 DINO and OpenLLaMA-3B. We compare our proposed approach with the image-only baseline when training a linear probe on the target dataset. All vision encoders are initialized from DINO weights, and our approach leverages unpaired text data using OpenLLaMA-3B embeddings.

					Dat	aset				
Method	Stanford Cars	SUN397	FGVC Aircraft	DTD	UCF101	Food101	Oxford Pets	Oxford Flowers	Caltech101	Average
Unimodal Ours Ours (init)	72.01 72.93 72.81	67.19 68.17 68.36	62.02 63.49 64.09	76.18 77.13 76.48	82.95 83.16 83.72	78.57 79.87 80.01	91.99 92.59 92.50	98.78 98.50 98.74	96.23 96.47 96.43	80.66 81.37 81.46

Table 12: Full linear probing on classification with ViT-S/14 DINOv2 and OpenLLaMA-3B. We compare our proposed approach with the image-only baseline when training a linear probe on the target dataset. All vision encoders are initialized from DINOv2 weights, and our approach leverages unpaired text data using OpenLLaMA-3B embeddings.

	Dataset												
Method	Stanford Cars	SUN397	FGVC Aircraft	DTD	UCF101	Food101	Oxford Pets	Oxford Flowers	Caltech101	Average			
Unimodal Ours Ours (init)	77.48 78.45 78.58	70.72 71.53 72.24	66.28 67.33 67.50	78.25 78.70 79.51	82.64 83.51 83.57	84.39 84.67 84.74	94.29 94.70 94.78	99.62 99.82 99.89	97.00 97.11 97.15	83.40 83.98 84.22			

Table 13: Full linear probing on classification with ViT-B/14 DINOv2 and OpenLLaMA-3B. We compare our proposed approach with the image-only baseline when training a linear probe on the target dataset. All vision encoders are initialized from DINOv2 weights, and our approach leverages unpaired text data using OpenLLaMA-3B embeddings.

	Dataset												
Method	Stanford Cars	SUN397	FGVC Aircraft	DTD	UCF101	Food101	Oxford Pets	Oxford Flowers	Caltech101	Average			
Unimodal Ours Ours (init)	85.46 85.40 85.74	75.42 75.22 75.70	72.34 75.22 74.17	79.73 80.73 81.32	87.26 87.21 87.26	88.70 89.02 88.78	95.56 95.83 95.78	99.76 99.88 99.88	97.81 97.85 97.93	86.89 87.37 87.40			

Table 14: Full linear probing on classification with ViT-L/14 DINOv2 and OpenLLaMA-3B. We compare our proposed approach with the image-only baseline when training a linear probe on the target dataset. All vision encoders are initialized from DINOv2 weights, and our approach leverages unpaired text data using OpenLLaMA-3B embeddings.

	Dataset												
Method	Stanford Cars	SUN397	FGVC Aircraft	DTD	UCF101	Food101	Oxford Pets	Oxford Flowers	Caltech101	Average			
Unimodal Ours Ours (init)	88.16 88.45 87.99	77.26 77.20 77.75	74.32 76.93 77.20	81.56 82.39 82.51	89.82 90.19 90.17	90.95 91.09 91.29	96.27 96.51 96.32	99.84 99.92 99.92	97.97 98.01 97.93	88.46 88.97 89.01			

E.1.3 FEW-SHOT LINEAR PROBING (ACROSS ARCHITECTURES)

In this section, we train only the linear classifier, on top of the frozen vision and language backbone, for few-shot classification on ten downstream tasks, comparing UML against strong image-only baselines. We evaluate five DINO-initialized backbones: ViT-B/16 in Table 16, ViT-B/8 in Table 15, DINOv2 ViT-S/14 in Table 17, DINOv2 ViT-B/14 in Table 19, DINOv2 ViT-L/14 in Table 19. Across all backbones, UML consistently improves over the image-only baseline by leveraging unpaired text embeddings. For all backbones, our head-initialization variant (*Ours* (*init*)) outperforms training using unpaired multimodal data from scratch (*Ours*).

Table 15: Linear evaluation of frozen features on 11 fine-grained benchmarks for few-shot learning. We compare our proposed approach with the image-only baseline by training a linear classifier on top of frozen VIT-B/8 DINO features. Our method leverages unpaired text data using OpenLLaMA-3B

							Dataset					
Train Shot	Method	Stanford Cars	Sun397	Fgvc Aircraft	Dtd	Ucf101	Food101	Imagenet	Oxford Pets	Oxford Flowers	Caltech101	Average
1	Unimodal	7.40	26.37	12.16	28.62	39.75	19.23	42.81	54.97	58.22	74.13	36.37
	Ours	7.71	28.01	13.56	33.22	42.08	21.13	43.27	55.85	58.61	77.51	38.10
	Ours (init)	9.24	34.23	14.49	36.27	47.55	24.81	46.75	60.09	61.59	80.23	41.52
2	Unimodal	14.43	37.96	20.28	39.80	53.03	30.62	54.75	68.12	77.59	81.91	47.85
	Ours	15.71	40.74	21.04	43.74	55.86	33.52	54.49	69.86	77.18	84.52	49.67
	Ours (init)	16.94	45.16	22.17	45.43	59.02	35.89	56.78	71.57	77.94	86.06	51.70
4	Unimodal	25.67	49.23	29.39	52.52	64.27	43.82	61.64	75.85	87.41	90.36	58.02
	Ours	27.30	51.23	31.43	54.31	66.72	45.58	61.51	77.51	87.96	91.36	59.49
	Ours (init)	28.54	53.68	31.31	56.13	67.47	47.40	62.84	79.10	88.29	91.98	60.67
8	Unimodal	41.04	56.86	40.03	61.15	72.39	54.47	66.10	82.30	93.95	92.28	66.06
	Ours	43.76	58.14	42.56	63.12	73.13	56.30	66.36	84.27	94.25	92.71	67.46
	Ours (init)	44.16	59.80	42.30	64.46	74.30	57.07	67.18	84.85	94.00	93.24	68.14
16	Unimodal	57.72	61.74	52.63	67.69	76.18	62.63	68.87	87.31	96.41	94.27	72.54
	Ours	60.11	63.21	54.53	69.33	78.13	63.74	69.44	87.73	96.89	94.54	73.76
	Ours (init)	60.36	64.26	54.81	70.27	78.76	64.13	70.05	88.23	96.63	94.73	74.22

Table 16: Linear evaluation of frozen features on 11 fine-grained benchmarks for few-shot learning. We compare our proposed approach with the image-only baseline by training a linear classifier on top of frozen VIT-B/16 DINO features. Our method leverages unpaired text data using OpenLLaMA-3B

							Dataset					
Train Shot	Method	Stanford Cars	Sun397	Fgvc Aircraft	Dtd	Ucf101	Food101	Imagenet	Oxford Pets	Oxford Flowers	Caltech 101	Average
1	Unimodal	6.28	22.43	9.72	29.22	37.85	15.40	38.67	60.12	54.62	73.25	34.76
	Ours	7.89	26.08	10.41	32.45	40.27	18.14	39.28	60.88	58.32	75.66	36.94
	Ours (init)	8.96	31.34	12.12	34.22	44.32	21.46	42.68	66.39	60.37	79.74	40.16
2	Unimodal	12.64	35.64	14.98	38.93	51.14	26.05	50.34	70.84	75.61	83.16	45.93
	Ours	14.38	38.62	17.00	40.37	54.28	29.24	50.83	72.88	77.14	85.95	48.07
	Ours (init)	15.99	42.31	17.65	42.89	56.46	32.15	52.90	74.82	77.32	87.34	49.98
4	Unimodal	22.60	45.95	24.27	50.30	63.00	38.51	57.99	80.14	85.60	89.67	55.80
	Ours	24.83	48.62	25.76	52.64	64.39	40.74	57.96	80.92	87.20	91.17	57.42
	Ours (init)	25.83	51.01	26.35	55.06	65.86	42.69	59.32	82.23	87.83	91.99	58.82
8	Unimodal	37.68	52.94	33.67	59.18	70.62	49.48	62.97	85.26	92.83	93.17	63.78
	Ours	39.31	55.31	35.56	60.48	71.88	50.46	63.08	86.25	93.23	93.47	64.90
	Ours (init)	40.50	57.03	35.64	62.27	73.18	51.50	64.09	86.93	93.59	93.71	65.84
16	Unimodal	52.48	58.27	45.34	64.81	75.72	56.24	66.36	88.57	95.90	94.27	69.80
	Ours	55.84	60.57	47.70	66.21	76.81	58.26	66.47	89.60	96.55	95.12	71.31
	Ours (init)	55.82	61.73	48.14	67.02	77.39	58.76	67.08	90.53	96.62	94.98	71.81

Table 17: Linear evaluation of frozen features on 11 fine-grained benchmarks for few-shot learning. We compare our proposed approach with the image-only baseline by training a linear classifier on top of frozen VIT-S/14 DINOv2 features. Our method leverages unpaired text data using OpenLLaMA-3B

							Dataset					
Train Shot	Method	Stanford Cars	Sun397	Fgvc Aircraft	Dtd	Ucf101	Food101	Imagenet	Oxford Pets	Oxford Flowers	Caltech101	Average
1	Unimodal	13.18	34.15	14.09	36.60	46.74	35.18	36.48	63.51	89.62	76.66	44.62
	Ours	14.95	37.25	14.88	38.93	49.18	37.91	38.35	68.92	91.42	84.04	47.58
	Ours (init)	16.49	41.79	15.63	42.04	52.33	42.27	42.69	73.59	93.64	84.52	50.50
2	Unimodal	24.68	47.88	23.09	47.75	56.81	48.54	50.41	75.32	96.02	86.90	55.73
	Ours	26.93	49.65	24.29	50.99	61.67	51.77	51.31	79.44	96.90	89.80	58.28
	Ours (init)	28.65	53.15	24.78	53.25	63.86	54.44	54.21	81.41	97.63	90.55	60.19
4	Unimodal	38.76	57.51	32.10	59.69	67.75	60.79	58.73	83.89	98.59	93.48	65.12
	Ours	41.69	58.87	33.38	61.58	69.60	62.69	59.69	86.27	98.84	94.56	66.71
	Ours (init)	43.17	60.89	33.86	62.43	71.13	63.88	61.38	87.36	99.17	94.96	67.82
8	Unimodal	54.56	63.00	45.05	64.78	74.19	68.06	64.53	88.68	99.27	94.35	71.65
	Ours	56.27	64.57	45.98	66.31	75.19	69.22	65.14	89.78	99.27	95.42	72.71
	Ours (init)	57.91	65.82	47.40	67.81	75.99	69.71	66.40	90.29	99.54	95.84	73.67
16	Unimodal	67.96	67.35	55.89	71.36	77.92	73.24	68.14	90.73	99.63	96.43	76.22
	Ours	69.42	68.50	58.54	72.24	78.69	73.80	68.70	91.87	99.72	96.63	77.80
	Ours (init)	70.32	69.19	58.74	73.17	79.58	74.51	69.44	92.47	99.82	96.80	78.81

Table 18: Linear evaluation of frozen features on 10 fine-grained benchmarks for few-shot learning with DINOv2 ViT-B/14. We compare our proposed approach with the image-only baseline by training a linear classifier on top of frozen VIT-B/14 DINOv2 features. Our method leverages unpaired text data using OpenLLaMA-3B

							D-44					
							Dataset					
Train Shot	Method	Stanford Cars	Sun397	Fgvc Aircraft	Dtd	Ucf101	Food101	Imagenet	Oxford Pets	Oxford Flowers	Caltech101	Average
1	Unimodal	22.42	43.03	15.79	38.85	58.57	48.71	52.26	76.47	97.12	83.64	53.69
	Ours	23.10	45.12	16.22	42.69	61.05	51.30	52.45	78.14	98.08	87.68	55.58
	Ours (init)	25.47	48.56	16.83	45.31	63.53	54.16	55.56	81.08	97.94	88.13	57.66
2	Unimodal	35.17	55.41	25.54	51.16	69.49	62.13	62.35	84.31	99.58	89.55	63.47
	Ours	37.38	56.98	25.88	54.65	70.61	63.89	63.21	85.50	99.70	92.02	64.98
	Ours (init)	38.78	59.81	26.00	55.61	71.38	66.54	65.06	86.49	99.62	92.79	66.21
4	Unimodal	51.40	63.68	34.25	61.25	76.32	71.60	68.86	89.05	99.76	94.51	71.07
	Ours	54.26	64.65	35.52	62.63	76.87	72.33	69.14	90.00	99.70	95.51	72.06
	Ours (init)	55.01	66.55	35.14	63.97	77.57	73.25	70.30	90.31	99.57	95.65	72.73
8	Unimodal	66.01	68.88	48.17	66.67	79.92	76.26	72.48	90.97	99.80	95.54	76.47
	Ours	68.53	69.75	50.88	68.46	81.44	77.34	73.12	92.39	99.70	96.20	77.78
	Ours (init)	67.91	70.66	51.26	69.56	81.85	77.95	73.75	92.50	99.68	96.51	78.16
16	Unimodal	77.31	72.17	62.38	73.76	83.80	80.74	75.15	93.34	99.81	97.40	81.59
	Ours	78.92	72.80	64.51	75.16	84.62	81.00	75.46	92.92	99.59	97.38	82.24
	Ours (init)	78.52	73.18	65.81	75.65	84.77	81.18	75.82	93.28	99.78	97.57	82.56

Table 19: Linear evaluation of frozen features on 10 fine-grained benchmarks for few-shot learning with DINOv2 ViT-L/14. We compare our proposed approach with the image-only baseline by training a linear classifier on top of frozen VIT-L/14 DINOv2 features. Our method leverages unpaired text data using OpenLLaMA-3B

							Dataset					
Train Shot	Method	Stanford Cars	Sun397	Fgvc Aircraft	Dtd	Ucf101	Food101	Imagenet	Oxford Pets	Oxford Flowers	Caltech101	Average
1	Unimodal	24.89	48.36	17.69	38.77	66.46	59.27	57.50	79.83	98.13	82.96	57.39
	Ours	25.88	49.63	18.08	42.93	69.18	60.12	58.37	83.51	98.42	86.23	59.24
	Ours (init)	27.90	52.86	18.95	43.18	70.98	63.17	60.80	83.86	98.59	88.17	60.85
2	Unimodal	39.95	58.95	26.87	50.18	75.79	70.74	67.14	84.71	99.74	89.82	66.39
	Ours	41.22	60.82	27.15	53.01	76.61	72.07	67.90	86.07	99.72	91.95	67.65
	Ours (init)	42.93	63.36	28.14	54.96	77.72	73.87	69.20	87.13	99.81	91.71	68.88
4	Unimodal	56.49	66.37	38.59	59.08	80.84	77.39	72.41	89.90	99.73	94.44	73.52
	Ours	58.19	67.36	39.57	61.78	81.36	78.19	72.82	90.99	99.76	95.27	74.53
	Ours (init)	58.60	68.84	39.19	62.77	81.50	78.99	73.63	90.74	99.88	96.02	75.02
8	Unimodal	70.00	70.71	51.57	66.47	83.84	81.69	76.02	93.53	99.89	95.55	78.93
	Ours	71.63	71.59	55.13	67.91	84.47	82.12	76.43	93.62	99.88	96.36	79.91
	Ours (init)	72.02	72.51	55.49	69.03	84.57	82.52	76.78	93.80	99.89	96.73	80.33
16	Unimodal	80.84	73.83	64.13	73.96	87.43	84.58	77.78	94.69	99.91	97.36	83.45
	Ours	81.85	74.39	69.45	74.70	87.35	84.58	78.35	94.59	99.89	97.61	84.28
	Ours (init)	82.76	74.80	69.42	74.88	87.65	84.96	78.58	94.42	99.81	97.62	84.49

E.2 IMPROVING IMAGE CLASSIFICATION USING UNPAIRED TEXTS (ALIGNED ENCODERS)

E.2.1 SUPERVISED FINETUNING

In this section, we fine-tune both the vision backbone and the linear classifier on nine downstream tasks, comparing UML against strong image-only baselines. We evaluate two different backbones: ResNet-50 and VIT-B/16.

As shown in Table 20, across all backbones, UML consistently improves over the image-only baseline by leveraging unpaired text embeddings. Further, our head-initialization variant (*Ours* (*init*)) outperforms training using unpaired multimodal data from scratch (*Ours*).

Table 20: **Supervised finetuning on 9 fine-grained classification benchmarks with CLIP**. We compare our proposed approach with the image-only baseline when fine-tuning on the target dataset. All vision encoders are initialized from CLIP ResNet50 weights, and our approach leverages unpaired text data using the corresponding CLIP text encoder.

					Dat	aset				
Method	Stanford Cars	Sun397	Fgvc Aircraft	Dtd	Ucf101	Food101	Oxford Pets	Oxford Flowers	Caltech 101	Average
Unimodal Ours Ours (init)	36.12 37.00 72.75	25.93 24.05 62.33	37.70 41.34 66.58	51.06 55.67 56.50	52.49 60.48 67.54	69.24 69.77 76.95	63.17 74.49 86.97	88.42 92.57 94.80	83.61 84.79 87.95	56.42 60.02 74.71

Table 21: Full linear probing on classification with CLIP ResNet-50 Image Encoder and Text encoder. We compare our proposed approach with the image-only baseline when training a linear probe on the target dataset. All vision encoders are initialized from ResNet-50 weights, and our approach leverages unpaired text data using the corresponding CLIP text embeddings.

					Dat	aset				
Method	Stanford Cars	SUN397	FGVC Aircraft	DTD	UCF101	Food101	Oxford Pets	Oxford Flowers	Caltech101	Average
Unimodal Ours Ours (init)	76.36 77.23 79.14	70.97 71.18 73.83	41.88 42.66 42.81	72.81 71.81 73.76	81.23 81.81 82.13	81.60 81.51 82.44	88.39 87.84 90.90	97.89 97.65 97.69	92.78 93.01 94.19	78.21 78.30 79.65

E.2.2 LINEAR PROBING

In this section, we train only the linear classifier, on top of the frozen vision and language backbone from CLIP, on ten downstream tasks, comparing UML against strong image-only baselines.

As shown in Table 21, UML consistently improves over the image-only baseline by leveraging unpaired text embeddings. Further, our head-initialization variant (*Ours* (*init*)) outperforms training using unpaired multimodal data from scratch (*Ours*).

E.2.3 FEW-SHOT LINEAR PROBING (ACROSS ARCHITECTURES)

In this section, we train only the linear classifier, on top of the frozen vision and language backbone from CLIP, for few-shot classification on ten downstream tasks, comparing UML against strong image-only baselines. We evaluate two different backbones: ResNet-50 and VIT-B/16.

As shown in Table 22 and Table 23, across both backbones, UML consistently improves over the image-only baseline by leveraging unpaired text embeddings. Further, our head-initialization variant (*Ours* (*init*)) outperforms training using unpaired multimodal data from scratch (*Ours*).

E.3 IMPROVING VISUAL ROBUSTNESS USING UNPAIRED TEXTS

In this section, we evaluate the robustness of models trained with UML to test-time distribution shifts. We train a k-shot linear probe (where $k \in \{1, 2, 4, 8\}$) with DINOv2 on ImageNet and evaluate across four distribution-shifted target datasets: ImageNet-V2, ImageNet-Sketch, ImageNet-A, and ImageNet-R. Our method consistently improves robustness over the unimodal baseline (Figure 10, Figure 11, Figure 12 and Figure 13) across different training shots, indicating that language priors help capture more transferable features.

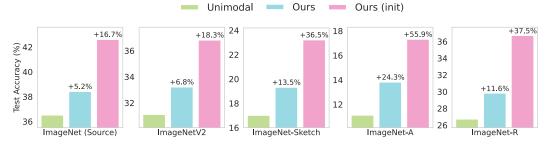


Figure 10: **Robustness under test-time distribution shifts.** Our approach (trained on 1-shot) is much more robust than its unimodal counterpart across four distribution-shuffled target test sets.

Table 22: Linear evaluation of frozen features on 10 fine-grained benchmarks for few-shot learning. We compare our proposed approach with the image-only baseline by training a linear classifier on top of frozen CLIP ResNet50 features. Our method leverages unpaired text data using the corresponding CLIP text encoder

			Dataset									
Train Shot	Method	Stanford Cars	Sun397	Fgvc Aircraft	Dtd	Ucf101	Food101	Oxford Pets	Imagenet	Oxford Flowers	Caltech101	Average
1	Unimodal	23.24	29.14	12.38	30.24	37.55	27.26	34.61	21.36	59.07	66.52	34.14
	Ours	36.32	45.40	16.84	40.92	53.19	49.76	53.03	36.48	68.56	76.80	47.73
	Ours (init)	57.88	64.59	22.23	50.85	65.99	76.73	86.59	60.92	81.08	83.79	65.06
2	Unimodal	38.37	43.83	18.63	40.33	53.25	44.60	47.75	32.62	75.03	78.90	47.33
	Ours	46.64	53.53	20.81	48.35	62.01	56.67	60.64	42.21	77.97	84.58	55.34
	Ours (init)	61.86	65.90	24.19	55.30	70.39	77.07	87.40	61.40	86.20	85.94	67.57
4	Unimodal	51.34	54.38	23.08	52.07	64.06	57.29	61.32	41.72	86.16	85.41	57.68
	Ours	55.21	59.48	24.77	56.78	67.65	62.68	67.31	47.04	86.46	87.23	61.46
	Ours (init)	65.80	68.11	27.49	60.13	73.62	77.79	86.54	62.37	91.60	87.57	70.10
8	Unimodal	61.74	61.47	30.22	60.15	70.16	64.63	68.94	49.48	92.20	89.14	64.81
	Ours	62.75	63.70	30.69	61.84	70.74	67.73	73.62	52.14	92.31	89.89	66.54
	Ours (init)	69.78	69.61	31.62	64.13	77.24	78.58	89.07	63.34	94.21	91.58	72.92
16	Unimodal	70.94	65.53	35.91	64.30	75.13	70.67	78.49	55.07	95.21	91.26	70.25
	Ours	71.58	67.08	36.23	65.62	76.09	71.63	79.52	56.92	95.44	91.94	71.20
	Ours (init)	74.56	71.33	37.13	68.09	78.66	79.06	89.71	64.31	96.17	93.31	75.23

Table 23: Linear evaluation of frozen features on 10 fine-grained benchmarks for few-shot learning. We compare our proposed approach with the image-only baseline by training a linear classifier on top of frozen CLIP VIT-B/16 features. Our method leverages unpaired text data using the corresponding CLIP text encoder

	Dataset											
Train Shot	Method	Stanford Cars	Sun397	Fgvc Aircraft	Dtd	Ucf101	Food101	Oxford Pets	Imagenet	Oxford Flowers	Caltech101	Average
1	Unimodal	31.53	33.51	17.76	31.72	43.64	39.40	37.43	27.65	67.95	71.68	40.23
	Ours	48.28	53.44	22.06	47.04	63.40	63.92	60.95	47.35	77.82	83.14	56.74
	Ours (init)	67.76	70.13	32.26	55.16	75.02	84.25	90.91	69.50	87.58	88.87	72.14
2	Unimodal	48.45	48.70	23.38	42.04	60.08	58.30	53.56	41.68	82.01	83.20	54.14
	Ours	57.89	59.95	27.19	52.27	69.60	71.18	66.78	54.24	87.43	90.20	63.67
	Ours (init)	70.75	71.52	33.99	60.17	78.37	85.39	90.67	70.19	92.18	90.09	74.33
4	Unimodal	61.64	60.66	31.01	54.37	70.49	71.91	69.35	52.15	90.99	91.08	65.36
	Ours	66.24	65.56	32.98	59.95	74.16	76.19	75.92	58.50	91.32	93.23	69.40
	Ours (init)	74.58	73.54	37.38	64.30	81.10	86.05	91.64	70.89	94.80	93.70	76.80
8	Unimodal	71.76	66.67	38.47	61.96	77.11	78.16	78.25	59.90	95.20	92.98	72.05
	Ours	72.77	69.50	39.09	64.89	79.01	80.07	80.85	62.63	94.98	94.36	73.82
	Ours (init)	78.43	75.07	41.77	68.50	83.41	86.87	92.55	71.97	96.94	95.27	79.08
16	Unimodal	78.76	71.49	44.74	68.79	80.43	82.08	85.16	63.87	96.97	94.54	76.68
	Ours	79.40	72.19	45.06	69.41	81.97	82.12	85.92	64.93	96.49	95.28	77.28
	Ours (init)	82.38	76.51	47.14	72.13	84.66	86.60	92.68	72.79	97.70	96.08	80.87

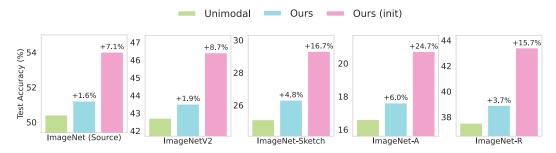


Figure 11: **Robustness under test-time distribution shifts.** Our approach (trained on 2-shots) is much more robust than its unimodal counterpart across four distribution-shuffled target test sets.

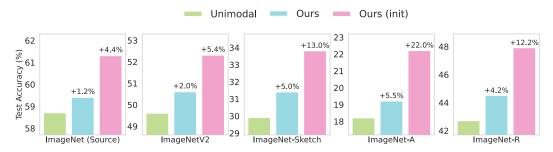


Figure 12: **Robustness under test-time distribution shifts.** Our approach (trained on 4-shots) is much more robust than its unimodal counterpart across four distribution-shuffled target test sets.

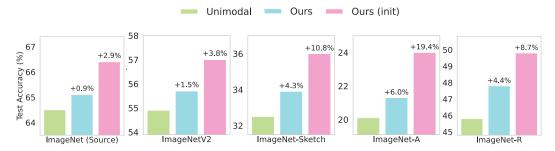


Figure 13: **Robustness under test-time distribution shifts.** Our approach (trained on 8-shots) is much more robust than its unimodal counterpart across four distribution-shuffled target test sets.

E.4 MARGINAL RATE-OF-SUBSTITUTION BETWEEN MODALITIES

How many words is an image worth? In this section, we extend our results to evaluate image-text conversion ratios using test accuracy isolines on the remaining eight datasets. We measure these global equivalence ratios by fitting a plane to the accuracy values given the number of image and text shots. Figures 14 to 21 demonstrate the conversion ratios for DINOv2 VIT-S/14 as the vision backbone and OpenLLaMa-3B as the text backbone (unaligned encoders). Analogously, Figures 22 to 29 show the same ratios for CLIP ResNet-50 as the vision and text encoders (aligned encoders). As expected, with the fully aligned CLIP backbone, each image equates to far fewer text prompts than under the unaligned DINO setting, showing the higher efficiency of aligned embeddings.

E.4.1 UNALIGNED ENCODERS

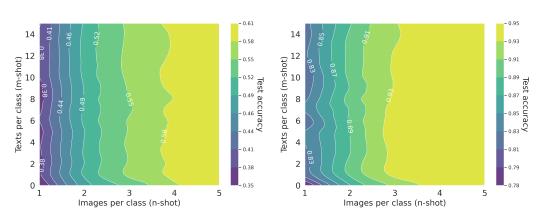


Figure 14: **SUN397.** 1 img \approx 1568 words

Figure 15: **Caltech101.** 1 img \approx 1248 words

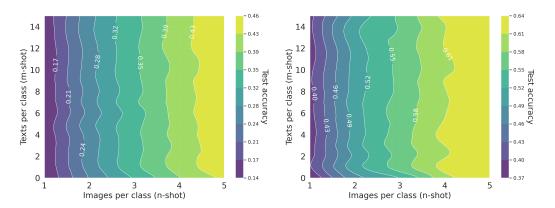


Figure 16: **Stanford Cars.** 1 img \approx 1799 words

Figure 17: **DTD.** 1 img \approx 2309 words

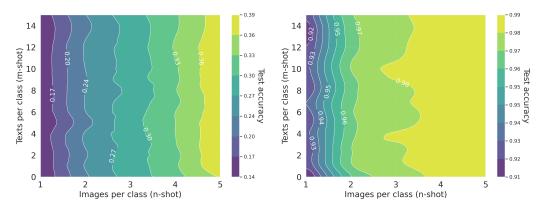
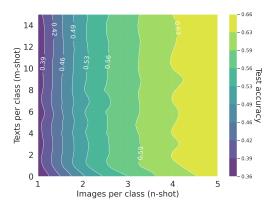


Figure 18: **FGVC Aircraft.** 1 img \approx 3220 words Figure 19: **Oxford Flowers.** 1 img \approx 1895 words



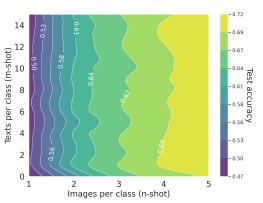


Figure 20: **Food101.** 1 img \approx 2608 words

Figure 21: **UCF101.** 1 img \approx 2617 words

E.4.2 ALIGNED ENCODERS (CLIP) 0.59 Texts per class (m-shot) 8 9 9 4 Texts per class (m-shot) 9 8 0 1 0.75 0.73 2 3 4 Images per class (n-shot) Images per class (n-shot) Figure 22: **SUN397.** 1 img \approx 221 words Figure 23: **Caltech101.** 1 img \approx 256 words 0.55 - 0.55 (m) 10 0.52 - 0.53 class (ber 0.41 0.39 0 1 2 3 4 Images per class (n-shot) 2 3 4 Images per class (n-shot) Figure 24: **Stanford Cars.** 1 img \approx 649 words Figure 25: **DTD.** 1 img \approx 228 words ਊ 12 -Texts per class (m-shot) <u>ه</u> الله 10 Test accuracy Texts per class 8-0.68 0.65 0.16 0 1 2 3 4 Images per class (n-shot) 2 3 4 Images per class (n-shot)

Figure 26: **FGVC Aircraft.** 1 img \approx 691 words Figure 27: **Oxford Flowers.** 1 img \approx 851 words

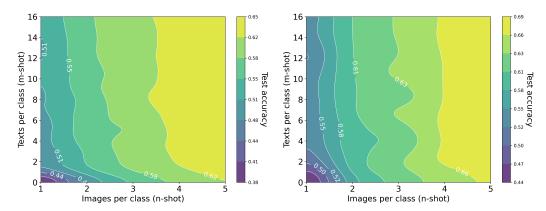


Figure 28: **Food101.** 1 img \approx 202 words

Figure 29: **UCF101.** 1 img \approx 393 words

E.5 IMPACT OF SCALING VISION BACKBONE

In this section, we study how our method's performance scales with the size and architecture of the vision backbone. In addition to ViT-S/14 DINOv2, we extend our analysis to a range of ViT-based architectures, including ViT-B/14 and ViT-L/14 DINOv2 and ViT-B/16 and ViT-B/8 DINO models. To ensure a fair comparison, we follow the same training protocol as in previous experiments. Our method consistently outperforms the unimodal baselines in every setting. In few-shot linear probing across ViT-B/8, ViT-B/16, DINOv2-ViTs and ViT-L/14 backbones (Tables 15 to 19), we see clear gains. The same holds for full-dataset end-to-end fine-tuning of both encoder and head (Tables 6 to 9), and even when only the linear classifier is trained on the full splits (Tables 10 to 14).

E.6 IMPACT OF VARYING TEXT ENCODERS

In this section, we study how our method's performance varies with different language models used for generating text embeddings. Through this experiment, we aim to understand how differences in embedding quality and model capacity affect the integration of textual information in our multimodal setup. Specifically, we cover LLMs with diverse architectures and scales, including BERT-Large, RoBERTa-Large and GPT-2 Large. As shown in Figure 30, adding unpaired text embeddings shows a significant boost in 1-shot accuracy and still decent gains at 16 shots on SUN397 dataset. Overall, OpenLLaMA-3B outperforms all other language models.

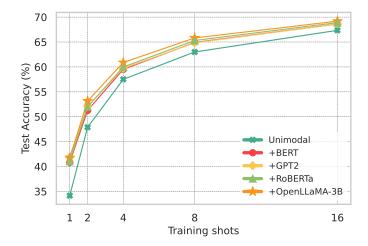
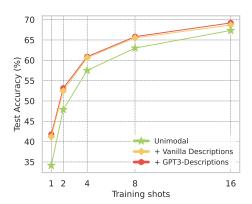


Figure 30: Few-shot classification accuracy on SUN397 using UML with unpaired, frozen embeddings from various pretrained language models.

LEARNING WITH COARSE-GRAINED VS. FINE-GRAINED TEXTUAL CUES

Understanding the type of information extracted from textual cues is crucial to assessing the effectiveness of our multimodal approach. A key question is whether the model merely utilizes class names or goes beyond to capture richer, more descriptive features. To investigate this, we compare the performance of our method using two types of text templates: a vanilla template that consists solely of the class name (e.g., "a photo of a [class]") and descriptive templates generated from GPT-3, as detailed in Section Appendix B. As shown in Figure 31 and Figure 32, both multimodal approaches consistently outperform the unimodal baseline, with descriptions from GPT-3 offering a more substantial performance gain. This shows that leveraging richer, contextually diverse text cues can significantly enhance model performance, even in low-shot learning scenarios.





grained descriptions.

Figure 31: Few-shot SUN397 accuracy with UML Figure 32: Few-shot SUN397 accuracy with using two levels of textual granularity: (a) vanilla UML (init) using two levels of textual granularclass descriptions and (b) GPT-3-generated fine- ity: (a) vanilla class descriptions and (b) GPT-3-generated fine-grained descriptions.

IMPACT ON PERFORMANCE WITH INCREASING UNPAIRED TEXT PROMPTS E.8

Here, we investigate how classification accuracy evolves as we augment each image with an increasing number of unpaired text prompts. Figure 33 shows these accuracy curves as we vary the number of unpaired text prompts per image shot across five image-shot budgets. In every regime, our multimodal initialization ("Ours (init)") outperforms training the head from scratch, with most of the gain coming from the first few prompts and gains tapering off thereafter. Note that we do not enforce diversity or novelty in the unpaired text prompts—simply adding more sentences does not guarantee additional information.

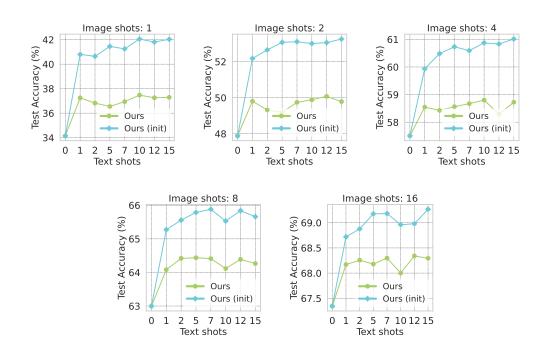


Figure 33: Classification accuracy as a function of the number of text prompts per image shot for the SUN397 Dataset.

E.9 EFFECT OF UNRELATED AUXILIARY MODALITIES

While our main experiments investigate the benefits of incorporating semantically related auxiliary modalities, an important question is what happens when the auxiliary data is unrelated or even adversarial. In our unpaired setting, this requires reasoning not about individual mismatched pairs, but about the relationship between entire data distributions. To meaningfully study the adversarial case, one would need a principled definition of "negative correlation" between modality distributions, as well as metrics for quantifying it. We leave such a rigorous framework to future work.

Instead, here, we evaluate a simpler but informative scenario: the *independent* case, where the auxiliary modality is semantically unrelated to the target modality. We use our few-shot learning setup to train UML with image data from SUN397 and text data from Stanford-Cars, two semantically unrelated datasets. Table 24 reports results for few-shot image classification on the SUN397 dataset. When the auxiliary text is unrelated to the image domain, performance does not improve over the unimodal baseline. In contrast, semantically related text provides consistent gains across all shot counts.

Table 24: Training UML with unrelated auxiliary text from the Stanford Cars dataset does not yield performance gains on image classification on the SUN397 dataset. However, semantically correlated text (from SUN397) consistently improves accuracy across few-shot settings.

Method	1-shot	2-shot	4-shot	8-shot	16-shot
Unimodal (Image) UML (Image + Unrelated Text) UML (Image + Related Text)	34.15	47.88	57.51	63.00	67.35
	35.27	47.12	57.50	62.45	67.25
	41.79	53.15	60.89	65.82	69.19

E.10 EXTENSION TO MORE THAN TWO MODALITIES

Our framework naturally extends beyond two modalities. We validate this by extending our image and audio classification experiments on ImageNet-ESC to use all three modalities: image, audio, and

text. Training alternates batches from each modality while applying modality-specific classification losses, consistent with our two-modality setup.

Tables 25 and 26 summarize results on audio and image classification. In both cases, incorporating additional unpaired modalities consistently improves performance over unimodal or pairwise settings. These findings demonstrate that the performance benefits of UML extend robustly to more than two modalities.

From a theoretical perspective, our results also generalize directly. Since modality-specific observations are conditionally independent given the ground-truth latent \mathcal{Z}^* , their joint contribution to the Fisher information reduces to the sum of the unimodal blocks. Consequently, the total contribution of all auxiliary modalities (excluding the primary modality X) can be obtained by summing their individual Fisher information matrices.

Table 25: Training UML with all three modalities from ImageNet-ESC outperforms unimodal or pairwise training on **audio classification**. Numbers in parentheses denote relative improvements over the unimodal baseline.

Dataset	Method	1-shot	2-shot	4-shot
ESC-19	Audio-Only	28.78	39.85	52.22
	Audio + Image	34.59	44.13	50.00
	Audio + Text	35.47	52.19	52.90
	Audio + Image + Text	44.46 (+54.4%)	51.48 (+29.2%)	56.57 (+8.3%)
ESC-27	Audio-Only	25.65	35.99	44.79
	Audio + Image	37.15	42.86	51.15
	Audio + Text	41.97	47.02	53.85
	Audio + Image + Text	44.68 (+74.2%)	48.03 (+33.4%)	54.16 (+20.9%)

Table 26: Training UML with all three modalities from ImageNet-ESC outperforms unimodal or pairwise training on **image classification**. Numbers in parentheses denote relative improvements over the unimodal baseline.

Dataset	Method	1-shot	2-shot	4-shot
ESC-19	Image-Only	60.28	74.10	78.70
	Image + Audio	64.63	76.17	85.02
	Image + Text	88.84	89.71	92.07
	Image + Audio + Text	90.55 (+50.2%)	91.08 (+22.9%)	91.72 (+16.5%)
ESC-27	Image-Only	55.33	65.60	77.85
	Image + Audio	59.75	70.93	78.14
	Image + Text	86.39	88.91	90.14
	Image + Audio + Text	88.22 (+59.5%)	88.96 (+35.6%)	91.78 (+17.9%)

E.11 EFFECT OF RATIO OF MODALITY BATCHES

In our main experiments, UML was trained with a simple 1:1 alternation of batches across modalities. To study the effect of this schedule, we ablate the ratio of text to image batches, denoted by r, on SUN397 using ViT-S/14 DINOv2 (vision) and OpenLLaMA-3B (text). We evaluate both in the (a) linear probe setting and the (b) full finetuning setting.

Across both settings and for both UML and UML (init), we observe that the choice of r has little impact on performance. The gains primarily arise from the presence of auxiliary information rather than the exact frequency of its appearance during training.

Table 27: Few-shot linear probing with different ratios of text-to-image batches on SUN397 using ViT-S/14 DINOv2 and OpenLLaMA-3B. Performance remains stable across ratios r, indicating robustness of UML to batch scheduling.

Method	Shot	r = 0.25	r = 0.5	r = 1.0	r = 2.0	r = 4.0
UML	2-shot	49.03	49.23	49.65	49.56	49.97
	4-shot	59.05	59.06	58.87	58.92	58.70
	8-shot	64.63	65.24	64.57	64.98	65.29
	16-shot	68.55	68.79	68.50	68.86	68.36
UML (init)	1-shot	42.60	42.13	41.79	41.95	42.04
	2-shot	52.83	53.01	53.15	53.04	52.81
	4-shot	61.10	61.09	60.89	60.63	60.99
	8-shot	65.29	65.32	65.82	65.15	65.21
	16-shot	69.19	68.91	69.19	68.51	68.71

Table 28: Full finetuning of the vision encoder and linear head with different text-to-image batch ratios r on SUN397. Results show that performance is largely insensitive to r.

Method	r = 0.25	r = 0.5	r = 1.0	r = 2.0	r = 4.0
UML	66.44	66.80	66.72	67.60	66.44
UML (init)	66.58	65.41	66.03	65.86	64.25

E.12 EFFECT OF FREEZING VS. UNFREEZING THE TEXT ENCODER

In all our main experiments, we freeze the text encoder. This design choice allows us to isolate the role of the auxiliary modality, ensuring that improvements in the primary modality (e.g., vision) arise from cross-modal transfer rather than joint training of both encoders.

In principle, however, one can also unfreeze the text encoder and update it during training. To study this, we ablate freezing versus unfreezing on Stanford Cars and SUN397 using ViT-S/14 DINOv2 (vision) and OpenLLaMA-3B (text). As shown in Table 29, unfreezing the text encoder improves performance on Stanford Cars, but slightly reduces performance on SUN397—likely due to the larger number of trainable parameters and increased optimization complexity.

Table 29: Effect of freezing versus unfreezing the text encoder when training with UML. Freezing stabilizes training and often yields slightly stronger gains.

Method	Stanford Cars	SUN397
Unimodal (Image Only)	79.45	66.20
UML (Unfrozen Text Encoder)	84.23	65.80
UML (Frozen Text Encoder)	84.87	66.72

E.13 ADDITIONAL EXPERIMENTS FOR AUDIO-VISUAL SETTING

In this section, we extend our unpaired multimodal framework to the tri-modal ImageNet–ESC benchmark, examining how unpaired audio and text signals can enhance image classification under both aligned (Appendix E.13.2) and unaligned encoders(Appendix E.13.1). We then reverse the setting, showing that unpaired visual and textual context likewise improves audio classification (Appendix E.13.3).

E.13.1 IMPROVING IMAGE CLASSIFICATION WITH UNPAIRED AUDIO AND TEXT (UNALIGNED ENCODERS)

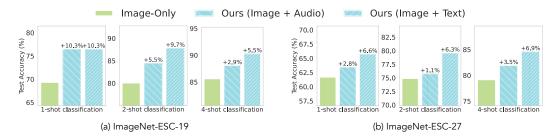


Figure 34: UML improves image classification using unpaired audio and text samples on both ImageNet-ESC-19 and ImageNet-ESC-27 benchmarks when trained on top of DINOv2 VIT-S/14 and OpenLLaMa-3B.

E.13.2 IMPROVING IMAGE CLASSIFICATION WITH UNPAIRED AUDIO AND TEXT (ALIGNED ENCODERS)

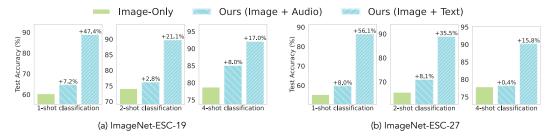


Figure 35: UML improves image classification using unpaired audio and text samples on both ImageNet-ESC-19 and ImageNet-ESC-27 benchmarks when trained on top of CLIP ResNet-50 image and text encoders

E.13.3 IMPROVING AUDIO CLASSIFICATION WITH UNPAIRED IMAGE AND TEXT (ALIGNED ENCODERS)

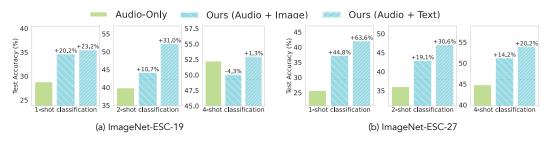


Figure 36: UML improves audio classification using unpaired image and text samples on both ImageNet-ESC-19 and ImageNet-ESC-27 benchmarks when trained on top of CLIP ResNet-50 image and text encoders

E.14 GAUSSIAN EXPERIMENTS

Here, we shift our attention to a more nuanced and intriguing question: can incorporating unpaired multimodal data actually improve the reconstruction quality of a single modality? At first glance, this seems unlikely—why would adding data from a different modality make X reconstruction better than training with X? Moreover, we push this question further: can incorporating data from a different modality, $while\ keeping\ the\ total\ dataset\ size\ fixed$, still improve the reconstruction of X compared to using the same number of samples X dataset alone? This setup isolates the importance of multimodal information from mere data scaling, and surprisingly, our experiments show that this improvement is indeed possible.

To investigate this, we design a synthetic experiment inspired by our theoretical framework in Section 3.1. We generate data from two partially overlapping modalities, X and Y, derived from a shared latent space θ_c , while also containing unique components (θ_x and θ_y). The observations follow the same linear structure as in our theory:

$$X_i = A_{c,i}\theta_c + A_{x,i}\theta_x + \epsilon_{X,i}$$
$$Y_i = B_{c,i}\theta_c + B_{y,i}\theta_y + \epsilon_{Y,i}$$

The training data are generated from Gaussian latents with dimensions $\dim(\theta_c)=10$, $\dim(\theta_x)=5$, and $\dim(\theta_y)=5$. For X, only the first 10% of the shared components in θ_c are retained at full strength, while the rest are downscaled by 0.05; Y observes all shared components at full strength. This asymmetry makes Y informative about structure that is only weakly present in X. The validation set is constructed from the same projections but without attenuation, so both modalities fully observe θ_c . Observations are 50-dimensional with Gaussian noise ϵ_X , $\epsilon_Y \sim \mathcal{N}(0,0.09I)$. In the unimodal setting, training on X alone uses 10,000 samples from X. When training UML on unpaired X and Y, we instead use 5,000 samples from each modality to keep the total sample budget fixed, ensuring a fair comparison.

Our architecture is a shared autoencoder. Each modality $X \in \mathbb{R}^{50}$, $Y \in \mathbb{R}^{50}$ is projected into a common space of dimension 128 through modality-specific linear layers. A shared encoder (two linear layers with ReLU) maps into a latent space of dimension 10, followed by a shared decoder (two linear layers) that expands back to dimension 128. Finally, modality-specific heads reconstruct the original inputs. The shared pathway enforces cross-modal alignment, while the separate adapters preserve modality fidelity.

As shown in Figure 37, the surprising outcome is that training on both modalities, even when they are unpaired, consistently improves the reconstruction of X compared to training solely on X. More strikingly, this improvement holds even when the total number of training samples is fixed, with half the data coming from X and half from Y; showing that the model is not just benefiting from increased data quantity but from the diversity and complementary information provided by the second modality.

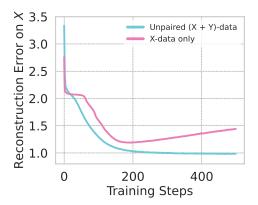


Figure 37: Training on N/2 samples from X and N/2 unpaired samples from Y improves test reconstruction on X, more than training on N samples from X.

F ANALYSIS OF THE LEARNED CLASSIFIER

F.1 CHANGE IN DECISION BOUNDARIES WITH UNPAIRED DATA FROM ANOTHER MODALITY

Our decision boundary visualizations are constructed by projecting the high-dimensional embedding space of a given classifier to a 2D plane. Axis 1 is computed as the normalized difference between the classifier weights of the two selected classes, representing the primary decision direction. Axis

2 is chosen to be orthogonal to Axis 1, constructed from the difference between the class mean embeddings after removing the component parallel to Axis 1. This orthogonalization ensures that the two axes capture complementary aspects: Axis 1 reflects the primary model decision boundary, while Axis 2 captures the variation orthogonal to that decision. The final 2D projection matrix combines these two vectors as columns, and embedding vectors are then mapped to this plane using a simple dot product. Figure 38, Figure 39 and Figure 40 show the change in decision boundary when adding unpaired textual information for 2-shot classification on top of frozen CLIP ResNet-50 features for Oxford Pets, DTD and Oxford Flowers datasets.

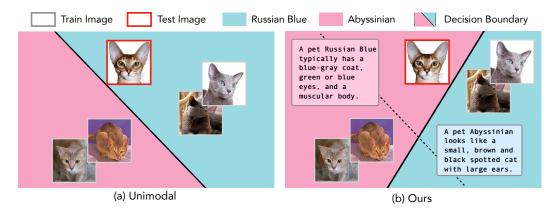


Figure 38: **Impact of unpaired text on decision boundaries** (**CLIP ResNet50.**) (Left) Visual features alone learn ambiguous class boundaries between Russian Blue and Abyssinian cats. (Right) Adding unpaired text sharpens the boundary, leveraging semantic cues to better distinguish similar categories.

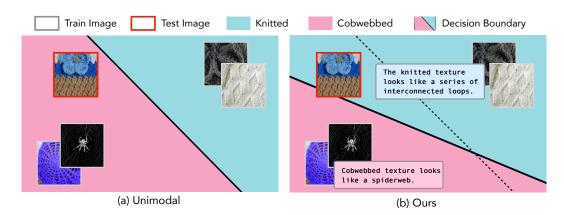


Figure 39: Impact of unpaired text on decision boundaries (CLIP ResNet50). (Left) Visual features alone learn ambiguous class boundaries between knitted and cobwebbed. (Right) Adding unpaired text sharpens the boundary, leveraging semantic cues to better distinguish similar categories

F.2 What do models learn from unpaired data?

To understand what the model is truly learning and how its weights evolve, we develop and analyze three key metrics: functional margin, silhouette score, and class-prototype vectors. These metrics inform on how well the model distinguishes between classes and how text information influences the structure of feature-space

Functional margin. This quantifies how confidently a model separates a given sample from the decision boundary. For a sample i belonging to class y, we calculate the margin relative to the next highest competing class. Specifically, we identify the second-highest logit among the incorrect

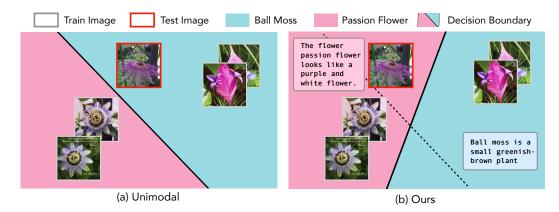


Figure 40: **Impact of unpaired text on decision boundaries (CLIP ResNet50).** (Left) Visual features alone learn ambiguous class boundaries between ball moss and passion flower. (Right) Adding unpaired text sharpens the boundary, leveraging semantic cues to better distinguish similar categories

classes, denoted as class j^* , and compute the functional margin as

$$\gamma_i = \frac{w_y^T x_i - w_{j^*}^T x_i}{\|w_y - w_{j^*}\|_2} \tag{5}$$

where $w_y^T x_i$ represents the logit for the true class, while $w_{j^*}^T x_i$ represents the highest logit among the competing classes. Larger margins indicate more confident and robust classification, while smaller margins imply that the sample lies closer to a misclassification boundary. As shown in Figure 41, both Ours and Ours (init) exhibit substantially larger classification margins than the unimodal baseline, demonstrating that augmenting primary-modality training with unpaired multimodal data improves confidence in predictions over the primary modality.

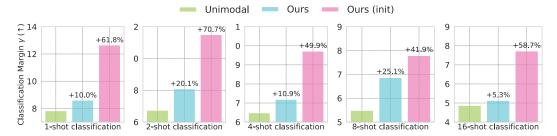


Figure 41: Functional margin of the linear head trained on SUN397 dataset for few-shot classification significantly increases when training with both UML and UML with linear head initialization.

Silhouette Score and DB-Index. The Silhouette Score indicates how well-separated the clusters are, while the DB-Index measures intra-class compactness versus inter-class separation. Higher silhouette and lower DB-Index values mean better-defined clusters, indicating that text helps tighten intra-class spread and widen inter-class gaps. As shown in Figure 42 and Figure 43, both *Ours* and *Ours (init)* exhibit reduced intra-class distances and increased inter-class separations, further confirming improved class separability.

Class-Prototype Vectors. These vectors are the rows of the final linear layer's weight matrix, representing the class centroids in the shared embedding space. We compute a heatmap of inner products between class prototypes and average text embeddings of the corresponding class to assess how well text features align with class centers. This helps reveal how the model organizes multimodal information. Figure 44 shows a pronounced diagonal structure, indicating that each class's text embedding aligns closely with the learned weights of the model.



Figure 42: Silhouette Score of the linear head trained on SUN397 dataset for few-shot classification significantly increases when training with both UML and UML with linear head initialization.



Figure 43: DB-Index of the linear head trained on SUN397 dataset for few-shot classification significantly improves when training with both UML and UML with linear head initialization.

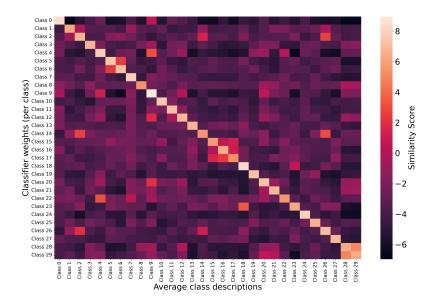


Figure 44: Inner products between each linear-head weight vector and its class's mean text embedding, demonstrating that text features align well with class prototypes.