

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 CLOSING THE MODALITY GAP ALIGNS GROUP-WISE SEMANTICS

Anonymous authors

Paper under double-blind review

ABSTRACT

In multimodal learning, CLIP has been recognized as the *de facto* method for learning a shared latent space across multiple modalities, placing similar representations close to each other and moving them away from dissimilar ones. Although CLIP-based losses effectively align modalities at the semantic level, the resulting latent spaces often remain only partially shared, revealing a structural mismatch known as the modality gap. While the necessity of addressing this phenomenon remains debated, particularly given its limited impact on instance-wise tasks (e.g., retrieval), we prove that its influence is instead strongly pronounced in group-level tasks (e.g., clustering). To support this claim, we introduce a novel method designed to consistently reduce this discrepancy in two-modal settings, with a straightforward extension to the general n -modal case. Through our extensive evaluation, we demonstrate our novel insight: while reducing the gap provides only marginal or inconsistent improvements in traditional instance-wise tasks, it significantly enhances group-wise tasks. These findings may reshape our understanding of the modality gap, highlighting its key role in improving performance on tasks requiring semantic grouping. Code available at: <https://anonymous.4open.science/r/CloseTheGap-7B0F>.

1 INTRODUCTION

The canonical goal of contrastive representation learning is to learn information from the original data by embedding semantically similar data points nearby and dissimilar ones far apart. Multimodal research has inherited this goal, assuming that a video, its caption, and its soundtrack should share the same neighborhood in a joint latent space. Unfortunately, the multimodal latent space shows different behavior, and representations tend instead to preserve their modality cluster, hindering the semantic alignment. This phenomenon, known as the modality gap Liang et al. (2022), is prevalent in all multimodal models grounded on the conventional and widespread contrastive InfoNCE loss function van den Oord et al. (2018) like CLIP Radford et al. (2021). Before training, samples from the same modality initially cluster together due to different random model weights initialization, forming distinct modality-specific groups. Unfortunately, these clusters persist even after training, resulting in a sparse and fragmented latent space.

In recent works, the modality gap has been studied and partially addressed for image and text pairs Es-lami & de Melo (2025); Yaras et al. (2025); Fahim et al. (2024); Schrodi et al. (2025); Mistretta et al. (2025). Concurrently, other works are instead accepting the gap, letting the space as it is while focusing on hyperbolic similarities Ramasinghe et al. (2024). However, the reasons for mitigating or untouched such a gap are not well-grounded and the

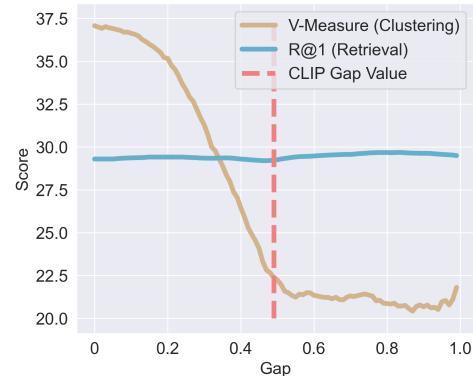


Figure 1: Reducing the gap consistently improves clustering metrics, while leaving unaffected retrieval ones. On the contrary, increasing the gap downgrades the V-Measure, bringing no improvements in R@1. In CLIP, the gap results in very poor clustering performance due to the latent space fragmentation.

advantages in common downstream tasks such as retrieval are inconsistent throughout the literature. Yaras et al. (2025) argues that closing the gap may improve the retrieval performance in downstream tasks, while Schrodi et al. (2025) found that a larger modality gap has a mild positive correlation with downstream performance. Moreover, all these approaches draw conclusions focusing only on the two-modal case (mainly image and text), without advancing cues in the case of multiple modalities.

Rather than following conventional views, we reconsider the modality gap in the context of group-wise tasks. We provide evidence that the modality gap is irrelevant for instance-wise tasks, such as retrieval, which depend on relative rankings, but strongly impacts multimodal group-wise tasks, such as clustering, which rely on absolute distances among representations of multimodal data in the latent space. Specifically, we show that closing the modality gap reduces the within-group scatter, leading to more coherent semantic groupings and better clustering performance, while leaving retrieval rankings mostly unaffected, as it is clear from Fig. 1. To address this, we propose a novel and effective objective function that explicitly encourages the alignment of matching pairs while maintaining a uniform and expressive latent structure. We show that our method effectively reduces the modality gap in both bimodal and trimodal benchmarks, improves clustering performance by several points in V-Measure, and maintains or slightly increases retrieval accuracy. These results confirm that the modality gap is not a harmless artifact, but a central factor shaping the geometry of the multimodal latent space. Crucially, we show that it can be effectively minimized through a simple objective, enabling better semantic organization without compromising instance-level precision.

Our main contributions can be summarized as follows.

- We demonstrate that the modality gap, while irrelevant to instance-wise tasks such as retrieval, directly impacts group-wise semantics reflecting in tasks, like clustering, that require explicitly closing the gap.
- We introduce a novel objective that combines true-pair alignment and centroid-based uniformity to effectively close the modality gap. Our method scales to multiple modalities and requires no architectural changes or post-hoc corrections.
- Across both bimodal and trimodal benchmarks, our approach reduces the gap while improving clustering and preserving retrieval performance, confirming its effectiveness for multimodal representation learning.

2 RELATED WORK

Multimodal Learning. Starting from CLIP Radford et al. (2021) several multimodal models have been developed for two modalities like CLAP Elizalde et al. (2023) or CLIP4Clip Luo et al. (2021). Lately, the same InfoNCE loss van den Oord et al. (2018) has been extended to multiple modalities in ImageBind Girdhar et al. (2023) or VAST Chen et al. (2023b). More recently, novel approaches have been proposed for multimodal learning to avoid the cosine similarity loss, namely GRAM Cicchetti et al. (2025) and Symile Saporta et al. (2024).

Modality Gap. The modality gap has been observed for the first time by Liang et al. (2022), and then studied mainly for the CLIP model Wu et al. (2023); Fahim et al. (2024); Shi et al. (2023) or for generic image and text pairs Mistretta et al. (2025); Yaras et al. (2025); Schrodi et al. (2025); Wu et al. (2023); Zhang et al. (2023). These works provide theoretical justification for the gap and propose to mitigate the gap by fixing the temperature Yaras et al. (2025), by applying post-hoc translations in the latent space Liang et al. (2022); Schrödi et al. (2025), or by sharing the transformer encoder and the projection layer in the vision and language encoders Eslami & de Melo (2025). In any case, each of these methods studied the modality gap in the case of two modalities, without advancing clues on the case of three or more modalities.

Effect of Temperature in Contrastive Learning. The temperature τ in the InfoNCE loss van den Oord et al. (2018) is recognized among the primary knobs for steering what features a contrastive learner captures. SimCLR Chen et al. (2020) noted that smaller τ sharpens the softmax and boosts instance-level retrieval accuracy. Wang & Isola (2020) proved that τ trades off alignment against uniformity, while higher τ instead favors cluster or class structure. Gradient-based analyses later showed that τ effectively controls the penalty on hard negatives, with low values concentrating gradients on the most confusable samples Wang & Liu (2020). Building on this view, Kukleva et al. (2023) introduced a schedule between high- τ (group-wise) and low- τ (instance-wise) phases. Most

recently, Dinu et al. (2025) showed that increasing τ compresses embeddings, lowering intrinsic dimensionality while preserving task performance, suggesting a second-order role for τ in model size and deployment efficiency.

3 WHY CLOSING THE GAP?

3.1 PRELIMINARIES

We consider a training batch $\mathcal{B} = \{(\mathbf{x}_i^1, \mathbf{x}_i^2, \dots, \mathbf{x}_i^M)\}_{i=1}^N$, where each sample is observed across M modalities, indexed by $m = \{1, 2, \dots, M\}$. For each modality m , the encoder $f^m : \mathcal{X}^m \rightarrow \mathbb{R}^d$ maps the input \mathbf{x}_i^m into a d -dimensional embedding vector $\mathbf{z}_i^m = f^m(\mathbf{x}_i^m)$, that is further normalized to obtain a unitary norm vector. We use $\text{sim}(\mathbf{z}_i^m, \mathbf{z}_j^n)$ to denote the cosine similarity between embeddings from modality m and n , i.e., $\text{sim}(\mathbf{z}_i^m, \mathbf{z}_j^n) = \langle \mathbf{z}_i^m, \mathbf{z}_j^n \rangle$.

Given a pair of modalities (m, n) with $m \neq n$, we define the InfoNCE loss for modality pair as:

$$\mathcal{L}^{(m \rightarrow n)} = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(\text{sim}(\mathbf{z}_i^m, \mathbf{z}_i^n)/\tau)}{\sum_{j=1}^N \exp(\text{sim}(\mathbf{z}_i^m, \mathbf{z}_j^n)/\tau)}, \quad (1)$$

where $\tau > 0$ is the temperature parameter controlling the softness of the distribution over negatives. The total contrastive loss is defined by averaging the two directions of the InfoNCE loss, $\mathcal{L}^{(m \rightarrow n)}$ and $\mathcal{L}^{(n \rightarrow m)}$. This is the standard loss used in CLIP Radford et al. (2021) and its variants for different modalities. The temperature τ modulates the contribution of hard negatives Kukleva et al. (2023). A low temperature penalizes harder negatives (i.e., those with high similarity to the anchor), favoring instance-level discrimination. In contrast, a higher temperature distributes gradients more evenly across all negatives, which typically promotes the emergence of semantic clusters and benefits group-wise reasoning tasks Dinu et al. (2025). In multimodal learning, different approaches have been proposed to handle the temperature effect including learning it during training or fixing at predefined values. Throughout the paper, we refer to methods with the standard InfoNCE loss with learnable temperature as CLIP (LT) and with fixed temperature as CLIP (FT).

3.2 UNDERSTANDING THE MODALITY GAP

A gap among modalities exists at the initialization phase, where different encoders initialized with random weights represent data in different narrow cones in the shared latent space Liang et al. (2022). Nevertheless, the gap persists even during the entire contrastive training phase. What is more, such contrastive learning dynamics have been recognized to be the root cause of the gap, regardless of the initialization. Therefore, even though the final learned space is somewhat semantically aligned, positive pairs are decoupled and very distant, as Fig. 6 in the Appendix shows. As demonstrated by Shi et al. (2023); Cicchetti et al. (2025), the traditional CLIP loss function easily gets stuck in local minima, in which positive pairs are somewhat matched but far from each other, fostering the modality gap.

In detail, previous works show that the conventional CLIP loss function is composed of two terms, each with specific objectives Wang & Isola (2020); Shi et al. (2023): a first term tries to align positive pairs, while the second one tries to spread away non-matching pairs. In practice, these two terms provide opposite contributions, resulting in balanced and opposite forces. Therefore, models easily end up in local minima, avoiding the gap closure while allowing the representations of the two modalities to align with each other in “semantic stripes”. These semantic stripes allow however good performance in retrieval tasks since positive pairs have higher cosine similarity with respect to non-matching pairs, even though such similarity is far from the ideal 1.0, as we show in Section 5. Therefore, representations of matching pairs do not lie in the same portion of the latent space and are instead quite far from each other, severely limiting the expressiveness and the group-wise alignment of the latent space.

Following Liang et al. (2022), to measure such a gap between two generic modalities m and n , we measure the effective Euclidean distance between the centroids of each modality:

$$\text{Gap}_{m,n} = \|\mathbf{c}^m - \mathbf{c}^n\|, \quad (2)$$

162 where $\mathbf{c}^m = \frac{1}{N} \sum_{i=1}^N \mathbf{z}_i^m$. Even though the gap is zero, this does not imply that the embeddings are
 163 effectively aligned in the latent space. Therefore, we further adopt the mean cosine similarity true
 164 pairs metric, defined as:

$$165 \quad \text{Cos TP}_{m,n} = \frac{1}{N} \sum_{i=1}^N \langle \mathbf{z}_i^m, \mathbf{z}_i^n \rangle. \quad (3)$$

168 This metric measures how much the normalized matching pairs are near each other in the latent space.
 169 Intuitively, the closer to 1.0, the smaller the angle is between them, and the closer the matching pairs
 170 lie in the latent space, being more aligned.
 171

172 3.3 ALIGNING GROUP-WISE SEMANTICS

174 While the modality gap does not necessarily disrupt instance-wise semantics, which depend on
 175 relative similarity rankings, it may prevent effective structuring of the latent space for group-wise
 176 objectives such as clustering. Let us formalize the reason starting from the InfoNCE loss in equation 1.
 177 Such a loss directly optimizes the relative ordering between the similarity of the true pair $\text{sim}(\mathbf{z}_i^m, \mathbf{z}_i^n)$
 178 and those of negatives $\text{sim}(\mathbf{z}_i^m, \mathbf{z}_j^n)$, $j \neq i$. The condition for retrieval success (e.g., Recall@ K) is
 179 satisfied whenever:

$$181 \quad \text{sim}(\mathbf{z}_i^m, \mathbf{z}_i^n) > \max(\text{sim}(\mathbf{z}_i^m, \mathbf{z}_j^n)), \quad \forall j \neq i. \quad (4)$$

183 Thus, as long as InfoNCE ensures this inequality, the relative ranking is preserved, regardless of
 184 whether absolute similarity values are close to 1.0 or far from it. Since instance-wise tasks like
 185 retrieval depend only on relative ordering and not on the absolute placement of embeddings, they are
 186 insensitive to the modality gap.

187 For group-wise tasks, let us build the centroid of semantic class c comprising two modalities m and
 188 n in the following way:

$$190 \quad \boldsymbol{\mu}_s^\delta = \frac{1}{2} ((\mathbf{z}_s^m + \boldsymbol{\delta}) + (\mathbf{z}_s^n + \boldsymbol{\delta})), \quad (5)$$

193 where the modalities are shifted according to a constant vector $\boldsymbol{\delta}$, representing the gap impact between
 194 the two modalities. Note that, in the case of 0 gap, the formula trivially resolves in $\boldsymbol{\mu}_s^0 = \frac{1}{2}(\mathbf{z}_s^m + \mathbf{z}_s^n)$.
 195 The expected within-class scatter for a generic modality (m) decomposes as

$$197 \quad \mathbb{E}_s [\|\mathbf{z}_s^m - \boldsymbol{\mu}_s^\delta\|_2^2] \approx \mathbb{E}_s [\|\mathbf{z}_s^m - \boldsymbol{\mu}_s^0\|_2^2] + \|\boldsymbol{\delta}\|^2, \quad (6)$$

198 with the gap term $\boldsymbol{\delta}$ summed to the expectation since it does not depend on the semantics of class
 199 s , since it is orthogonal to the span of semantic vectors and therefore constant for each of them, as
 200 proven by Zhang et al. (2023). Therefore, enlarging the gap uniformly inflates every semantic cluster,
 201 degrading homogeneity and completeness, while shrinking the gap tightens such clusters.

203 **Intuitive consequence.** The modality gap leaves rank-based, instance-wise decisions mostly un-
 204 touched, as they are the primary objective of the InfoNCE loss, but widens absolute distances and
 205 thereby harms any objective that optimizes within/between-cluster geometry. Retrieval is therefore
 206 almost flat across the gap value, whereas clustering (and other group-wise tasks) benefit from driving
 207 the gap to zero, as shown in Fig. 1. For instance, suppose to retrieve a cat image caption, the
 208 sufficient condition is that the caption remains the single most similar text, irrespective of its absolute
 209 cosine value. Instead, if the goal is forming a “cat” cluster, the necessary condition is that every cat
 210 image-caption pair lies close to the “cat” centroid, and a gap between the two modalities inflates that
 211 radius, so cluster purity drops even though each individual pair is still top-ranked for retrieval.

212 4 CLOSING THE MODALITY GAP

213 We aim to close the modality gap while ensuring consistent alignment among the positive pair
 214 distribution. To achieve such scope, we propose two novel losses. The first one, the Align True Pairs

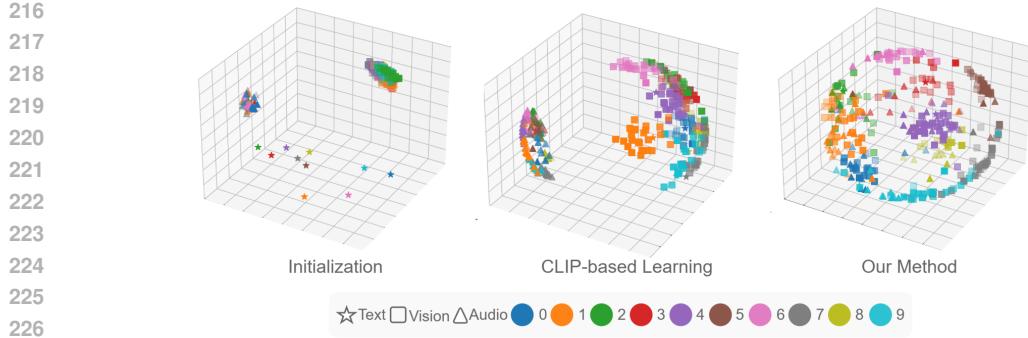


Figure 2: AV-MNIST multimodal latent space. The CLIP-based learning creates a fragmented latent space with embeddings clearly clustered by modality and not by multimodal semantics. Our method closes the gap and enhances group-wise semantics, placing embeddings of the same class in the same portion of the space, effectively learning a semantically meaningful multimodal latent space.

loss \mathcal{L}_{ATP} guarantees the alignment between true pairs. Considering M modalities, among which \mathbf{a} is the anchor one (i.e., the modality to which all the other modalities are aligned to Girdhar et al. (2023); Zhu et al. (2024)), the loss is:

$$\mathcal{L}_{\text{ATP}} = \frac{1}{M-1} \sum_{m \in M, m \neq a} \left(\frac{1}{N} \sum_{i=1}^N (\|\mathbf{z}_i^m - \mathbf{z}_i^a\|_2^2) \right), \quad (7)$$

where m is taken from the set of available modalities, \mathbf{a} is the anchor modality and N is the batch size. The second one, the Centroid Uniformity loss \mathcal{L}_{CU} , ensures uniformity of the semantic classes among the modalities in the latent space by:

$$\mathcal{L}_{\text{CU}} = \log \left(\frac{1}{N} \sum_{i=1}^N \sum_{j=1, j \neq i}^N \exp(-2\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\|_2^2) \right), \quad (8)$$

in which $\boldsymbol{\mu}_k$, with $k = i, j$, are the centroids defined as:

$$\boldsymbol{\mu}_k = \frac{1}{M} \sum_{m \in M} \mathbf{z}_k^m, \quad (9)$$

and \mathbf{c}_k is the centroid of the k -th element of the batch built by averaging all the modalities embeddings. The effect of the two losses is complementary. The \mathcal{L}_{ATP} promotes closeness between positive pairs, effectively enhancing the mean cosine similarity between them. However, involving solely such a loss may produce a side effect: the entire latent space collapses into small portions, placing representations of dissimilar data in the same portion of the latent space, as we show in the Appendix. Therefore, the contribution of the \mathcal{L}_{CU} loss becomes crucial, ensuring the sparsification of the latent space by enforcing uniformity to the centroids while preserving the learned alignment. Indeed, moving the centroids of matching pairs in the space implies moving the modalities representations accordingly, effectively preserving alignment while leveraging the whole space. Without centroids, uniformity should have been applied independently to each modality as in Wang & Isola (2020), disrupting the learned alignment among similar semantic pairs, as we show in the Appendix. Additionally, the radial basis function (RBF) kernel in equation 8 is well related to the uniform distribution on the unit hypersphere where multimodal representations lie Wang & Isola (2020), therefore enforcing the coverage of the whole surface of the hypersphere. The overall proposed loss function that aims at aligning the true pairs and closing the modality gap is a sum of the two terms:

$$\mathcal{L}_{\text{gap}} = \mathcal{L}_{\text{ATP}} + \mathcal{L}_{\text{CU}}. \quad (10)$$

Such loss should be then combined with the contrastive loss to obtain:

$$\mathcal{L}_{\text{CL}_{\text{gap}}} = \mathcal{L}_{\text{gap}} + \frac{1}{2} (\mathcal{L}^{(m \rightarrow n)} + \mathcal{L}^{(n \rightarrow m)}), \quad (11)$$

for the two-modality m and n case to lighten the notation, but it can be expanded to more modalities.

270

5 EXPERIMENTAL VALIDATION

271
272 We perform extensive evaluations in controlled and real-world scenarios with four different datasets,
273 four tasks, and a diverse number of modalities.
274275

5.1 SETTING

276 To evaluate the proposed method in multimodal scenarios, we design a series of experiments that
277 progressively increase in complexity and scale. Following the literature, we begin with two foundational
278 experiments using the CIFAR-10 and AV-MNIST datasets, which involve two and three modalities,
279 respectively. Subsequently, we scale up our experiments using the MSCOCO and MSR-VTT datasets,
280 which offer more complex multimodal data.
281282 **CIFAR10** (Image-Text). The CIFAR10 dataset consists of 60k 32×32 color images evenly distributed
283 across 10 classes, with 50k images for training and 10k for testing Krizhevsky et al. (2009). For our
284 experiments, we pair each image with its corresponding class label as text, creating a bimodal dataset.
285 For CIFAR10, we use two separate ResNet50 encoders, one for text and the other for images.286 **AV-MNIST** (Audio-Visual-Text). AV-MNIST is a synthetic dataset combining visual, auditory, and
287 textual modalities. It is the union of two well-known datasets: the MNIST dataset composed of 60k
288 28×28 digit images and the Audio-MNIST Becker et al. (2023) containing 30k audio samples of
289 spoken digits (0-9) from diverse speakers. Each sample is also associated with a textual label that
290 represents the digit. Encoder details in the Appendix.291 **MSCOCO** (Image-Caption). The MSCOCO dataset contains over 330k images, each annotated
292 with five human-generated captions, totaling more than 1.5M captions Lin et al. (2014), and it is a
293 well-known benchmark for analyzing the modality gap. To scale up the model capabilities and test
294 our approach in this challenging dataset, we use EVA-CLIP ViT-G Sun et al. (2023) as the visual
295 encoder and BERT-B Devlin et al. (2019) to process the textual captions.296 **MSR-VTT** (Video-Audio-Text). The MSR-VTT dataset comprises 10k video clips spanning 20
297 categories, with each clip annotated with approximately 20 natural language sentences, resulting
298 in a total of 200k captions Xu et al. (2016). Each video clip includes visual frames, audio tracks,
299 and corresponding textual descriptions, making it a trimodal dataset. Each sample in the dataset is
300 associated to one of the 20 categories. This makes the dataset useful to be evaluated with clustering
301 metrics. We maintain the same architectures used for MSCOCO plus BEATs Chen et al. (2023a), a
302 transformer-based model, used to extract features from the audio tracks.303
304

5.2 TASKS

305 We conduct a suite of experiments covering both instance-wise and group-wise tasks.

306 **Cross-Modal (CM) Retrieval.** We evaluate the model ability to associate related data across
307 modalities (e.g., text to audio or text to image) with cross-modal retrieval. For a given query in one
308 modality, the task is to retrieve matching items from another modality using cosine similarity in the
309 shared latent space. This task quantifies the instance-level alignment and cross-modal discrimination
310 capacity of the learned embeddings.311 **Clustering.** We assess the structural organization of the latent space using supervised and unsu-
312 pervised clustering. We apply standard clustering algorithms (k-means and k-NN) and evaluate the
313 resulting clusters using V-Measure and k-NN accuracy. High performance in clustering indicates that
314 the latent space preserves semantic coherence across modalities and aligns samples belonging to the
315 same class or coarse-grained semantic group closely, regardless of the input domain.316 **Multimodal (MM) Retrieval with Cross-Conditioning.** We evaluate the performance of down-
317 stream models that use the extracted embeddings from the modality encoders following Li et al.
318 (2021). For this task, we involve an additional multimodal encoder that takes as input all the modal-
319 ities and fuses them. Therefore, a proper alignment is crucial in this task. Such a multimodal encoder
320 outputs the probability score of the semantic match among them. To this purpose, following Chen
321 et al. (2023b), as the multimodal downstream encoder we reuse the text encoder attaching an MLP
322 on top of it to extract the probability value. Along with our brand-new loss functions, we add the
323 data-anchor matching (dam) loss function to train these components, as in Li et al. (2021); Chen et al.

324 (2023b); Cicchetti et al. (2025):
 325

$$\mathcal{L}_{\text{dam}} = \mathbb{E}_{(\mathbf{a}, \mathbf{m}_2, \dots, \mathbf{m}_k) \sim (A, M_2, \dots, M_k)} [y \log p_{\text{dam}} + (1 - y) \log(1 - p_{\text{dam}})]. \quad (12)$$

328 This task stresses the ability of the encoders to build coherent, modality-agnostic representations
 329 and provides insight into the expressiveness and compositionality of the learned space. Higher
 330 performance in this task highlights an increased facility by the multimodal encoder to detect interde-
 331 pendencies among modalities.

332 **Captioning.** A downstream task to evaluate the benefit of an aligned and well-structured latent
 333 space is data captioning. Following Yan et al. (2022), we involve a decoder serving as the language
 334 generator model and we add a specific loss term that, along with our proposed loss functions, trains
 335 the encoder-decoder structure for this specific task. Intuitively, the more aligned and semantically
 336 coherent the latent space, the better the model will generate the captions. The captioning loss is:

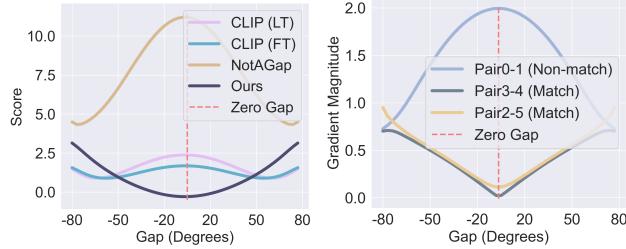
$$\mathcal{L}_{\text{cap}} = - \sum_{t=1}^T \log P_{\theta}(\mathbf{y}_t | \mathbf{y}_{<t}, \mathbf{x}), \quad (13)$$

340 where \mathbf{y} is the exact tokenized text the model aims to learn by maximizing the conditional likelihood
 341 under the forward autoregressive factorization. $P_{\theta}(\mathbf{y}_t | \mathbf{y}_{<t}, \mathbf{x})$ denotes the probability assigned to the
 342 token \mathbf{y}_t given as input the past history $\mathbf{y}_{<t}$ and the input features \mathbf{x} . Captioning requires both accurate
 343 semantic grounding and compositional generalization, and therefore serves as a complementary proxy
 344 for evaluating the usefulness of the learned features in downstream generative tasks.

346 5.3 RESULTS

348 Aligning group-wise semantics

349 Figure 1 shows that increasing the
 350 modality gap leads to a drop in clustering
 351 metrics, despite having a negli-
 352 gible effect on retrieval scores, while
 353 decreasing it considerably improves
 354 the clustering metric (+17.5 points).
 355 This suggests that instance-level align-
 356 ment is tolerant to cross-modal off-
 357 sets, while group-wise metrics benefit
 358 from a more globally structured latent
 359 space. Crucially, **the optimal clus-
 360 tering performance coincides with
 361 the zero-gap configuration**, indicat-
 362 ing that full semantic overlap facil-
 363 titates more coherent clusters. The plot
 364 is obtained by starting from the gap
 365 value of standard CLIP on the MSR-
 366 VTT dataset and then posthoc shifting
 367 the matching pairs up to 0.0 and 1.0 gap.
 368 Second, the simulation in Figure 3(a) shows that, in presence
 369 of data mismatching, our **proposed loss induces a landscape in which the global minimum is
 370 achieved when the modality gap is zero**. In contrast, standard contrastive losses such as CLIP
 371 achieve the minimum around 60 degrees, thus when matching pairs are still quite far in the latent
 372 space. Proof in the Appendix. Therefore, we can train the model with the proposed losses to reach
 373 the zero modality gap while aligning group-wise semantics and improving the clustering performance
 374 without affecting retrieval performance. The underlying mechanism is clarified in Figure 3(b), which
 375 shows that **as the modality gap narrows, the gradient magnitude associated with non-matching
 376 pairs increases**, while that of matching pairs decreases. With a large gap, non-matching samples
 377 provide the same gradient contribution as matching ones. As the gap closes, these negatives become
 378 harder and more semantically informative, thus dominating the optimization process and encouraging
 379 the model to better shape the space by adjusting the missing matches. This naturally improves the
 380 structuring of semantic clusters without altering the rank of positive pairs (which have nearly zero
 381 gradients with the closed gap), hence preserving (or improving) retrieval performance. Altogether,



347 Figure 3: **(a)** Following Liang et al. (2022), we place six sim-
 348 ulated image-text embedding pairs on a 3D sphere, with two
 349 mismatched pairs. We artificially move these pairs toward
 350 closing or enlarging the gap among them and we track the
 351 loss landscape during the simulation. **(b)** During the same
 352 simulation we keep track also of the gradient magnitude re-
 353 ceived by the six embeddings pairs through our design loss
 354 function $\mathcal{L}_{\text{CL}_{\text{gap}}}$. When the gap is closer to zero, the contribu-
 355 tion to the loss is just matter of the non matching pairs.

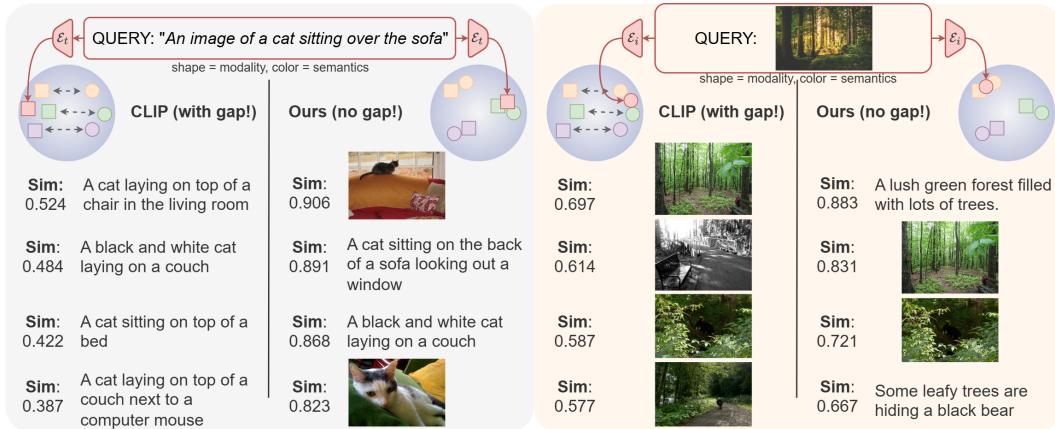


Figure 4: Most similar vectors from the MSCOCO vector databases with both modalities embedded. CLIP learned space has a gap of 0.47 and, for a text query, all the closest embeddings come from the text modality, meaning that embeddings are clustered according to the modality and not according to the overall multimodal semantics. On the contrary, the proposed method has nearly a zero gap, and most similar samples come from both the image and text modalities, proving that we can effectively build a multimodal latent space that is semantically coherent among the modalities. The same behavior holds for an image query.

these observations highlight that minimizing the modality gap leads to more informative gradient signals by emphasizing semantically meaningful negatives, thereby improving the alignment and group-wise structure of the latent space. This reflects in improved clustering performance without affecting retrieval. By explicitly encouraging the zero-gap configuration, our proposed loss $\mathcal{L}_{\text{CL}_{\text{gap}}}$ provides a straightforward method to close the gap and align group-wise semantics.

Figure 2 shows the latent space plot in the case of embedding dimension equal to 3 in the AV-MNIST dataset. Notably, the plotted spaces are not stochastically generated but the real latent space in \mathbb{R}^3 . As it is clear, the proposed method closes the gap better than conventional CLIP-based learning. Crucially, this consistently improves group-wise clustering, placing embeddings from the same class closer to each other in a semantically meaningful latent space.

Table 1 quantitatively proves our claims on the relation between the gap and the group-wise semantics. Indeed, as the gap is progressively closed with the proposed method, the clustering performance crucially improves, while retrieval accuracy between text and video (TV) and text and audio (TA) is barely affected by the gap. Overall, this intuition is congruous across diverse datasets with different numbers (two and three) and types of modalities (image, video, text, audio), even though the simpler the dataset, the easier is for the model to align group-wise semantics.

Closing the gap

Figure 4 shows the most similar samples by querying the learned latent spaces of CLIP and of our method without forcing the query to be cross-modal by construction. This means that the similarity cares about the semantics only, regardless of the modality. Interestingly, when querying the CLIP space for the nearest vectors, while returning semantically similar samples, the embeddings always belong to the same modality as the query. In contrast, in the gap-free latent space learned with our method, the closest vectors not only share the same semantics but also come from different modalities. This proves that closing the modality gap effectively helps build a semantically aligned multimodal latent space.

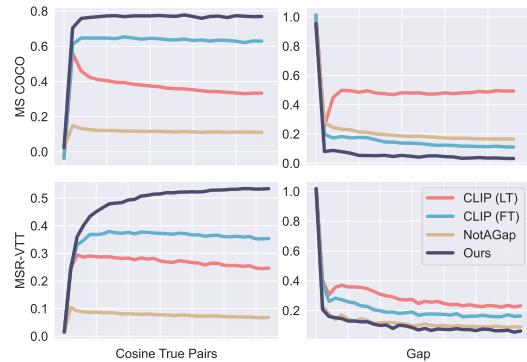


Figure 5: CosTP and gap during training.

Table 1: Instance-wise (cross-modal (CM) retrieval) vs group-wise (clustering) tasks correlated to the gap value. The harder the dataset, the more evident the results.

			CM Retrieval		Clustering	
Method	Dataset	Gap \downarrow	TV R@1	TA R@1	V-Measure	kNN
CLIP (LT) (Radford et al., 2021)	CIFAR10 (2 modal)	0.86	82.0	-	67.0	81.2
CLIP (FT) (Yaras et al., 2025)		0.14	82.1	-	67.6	81.9
Ours		0.09	82.4	-	67.9	82.4
CLIP (LT) Radford et al. (2021)	MSCOCO (2 modal)	0.47	74.6	-	12.98	26.3
CLIP (FT) Yaras et al. (2025)		0.12	73.2	-	12.99	31.0
Ours		0.03	70.3	-	23.63	36.4
CLIP (LT) (Radford et al., 2021)	AV-MNIST (3 modal)	0.20	87.1	84.2	77.6	87.0
CLIP (FT) (Yaras et al., 2025)		0.24	84.1	80.4	73.8	85.0
Ours		0.09	88.7	89.1	82.7	89.2
CLIP (LT) (Radford et al., 2021)	MSR-VTT (3 modal)	0.29	34.2	10.3	23.3	52.9
CLIP (FT) (Yaras et al., 2025)		0.19	34.8	10.1	31.3	55.7
Ours		0.07	32.8	11.8	32.1	58.0

Table 2: Multimodal (MM) retrieval and captioning results.

		Space Measures		MM Retrieval	Captioning		
Method	Dataset	Gap \downarrow	Cos TP \uparrow	R@1	BLEU@1	BLEU@3	CIDEr
CLIP (LT)	MSCOCO (2 modal)	0.47	0.34	72.5	45.8	22.5	153.2
CLIP (FT)		0.12	0.63	73.8	45.9	23.0	155.0
NotAGap		0.17	0.11	75.6	45.4	22.2	153.3
Ours		0.03	0.77	77.3	46.1	23.2	160.9
CLIP (LT)	MSR-VTT (3 modal)	0.24	0.27	30.6	26.7	9.5	63.6
CLIP (FT)		0.17	0.37	30.3	26.2	9.4	63.2
NotAGap		0.09	0.07	29.9	24.3	8.7	53.8
Ours		0.06	0.53	33.3	26.8	9.6	64.4

Table 2 compares the proposed method and the current literature in closing the gap and downstream tasks. As it is clear from Tab. 2, our novel loss closes the multimodal gap by achieving nearly 0.0 distance between the modality clusters centroids. Additionally, we better align true pairs by crucially increasing their cosine similarity from 0.34 of the standard CLIP (LT) to 0.77 in MS COCO. The capabilities of modeling the multimodal latent space are evident in Fig. 5 too, where the proposed method outperforms the conventional CLIP, other solutions with fixed temperature (FT) Yaras et al. (2025), and NotAGap Fahim et al. (2024). Importantly and counter-intuitively, during training, the cosine similarity of true pairs slightly but progressively decreases in comparison configurations. This means that the true pairs are progressively worse aligned. On the contrary, the proposed method considerably increases the cosine similarity between true pairs, truly aligning them, while also definitely closing the gap among modalities, regardless of whether there are two or three modalities.

6 CONCLUSION

In this work, we revisited the modality gap in multimodal representation learning, providing insights into the impact of this phenomenon on downstream tasks. Through theoretical analysis and extensive experiments across bimodal and trimodal datasets, we showed that while the modality gap has limited influence on instance-level tasks such as retrieval, it significantly affects group-wise tasks like clustering. To address this, we proposed a novel objective that explicitly aligns true pairs while promoting latent space uniformity. Our method consistently reduces the modality gap and improves clustering performance without compromising retrieval accuracy, offering a simple yet effective solution for better semantic organization in multimodal spaces.

486 REFERENCES
487

488 Sören Becker, Johanna Vielhaben, Marcel Ackermann, Klaus-Robert Müller, Sebastian Lapuschkin,
489 and Wojciech Samek. AudioMNIST: Exploring explainable artificial intelligence for audio analysis
490 on a simple benchmark. *Journal of the Franklin Institute*, 2023.

491 Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Daniel Tompkins, Zhuo Chen, Wanxiang Che,
492 Xiangzhan Yu, and Furu Wei. BEATs: Audio pre-training with acoustic tokenizers. In *International
493 Conference on Machine Learning*, pp. 5178–5193, 2023a.

494 Sihan Chen, Handong Li, Qunbo Wang, Zijia Zhao, Ming-Ting Sun, Xinxin Zhu, and J. Liu. VAST:
495 A vision-audio-subtitle-text omni-modality foundation model and dataset. In *Neural Information
496 Processing Systems (NeurIPS)*, 2023b.

497 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for
498 contrastive learning of visual representations. In *International Conference on Machine Learning*,
499 pp. 1597–1607. PMLR, 2020.

500 Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade
501 Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for
502 contrastive language-image learning. In *Proceedings of the IEEE/CVF Conference on Computer
503 Vision and Pattern Recognition*, pp. 2818–2829, 2023.

504 Giordano Cicchetti, Eleonora Grassucci, Luigi Sigillo, and Danilo Comminiello. Gramian multimodal
505 representation learning and alignment. *International Conference on Learning Representations
506 (ICLR)*, 2025.

507 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of
508 deep bidirectional transformers for language understanding. In *Association for Computational
509 Linguistics*, 2019.

510 Georgiana Dinu, Corey D Barrett, Yi Xiang, Miguel Romero Calvo, Anna Currey, and Xing Niu.
511 Effective post-training embedding compression via temperature control in contrastive training. In
512 *International Conference on Learning Representations (ICLR)*, 2025.

513 Benoit Dufumier, Javiera Castillo Navarro, Devis Tuia, and Jean-Philippe Thiran. What to align
514 in multimodal contrastive learning? In *International Conference on Learning Representations
515 (ICLR)*, 2025.

516 Benjamin Elizalde, Soham Deshmukh, Mahmoud Al Ismail, and Huaming Wang. CLAP: learning
517 audio concepts from natural language supervision. In *IEEE International Conference on Acoustics,
518 Speech and Signal Processing (ICASSP)*, pp. 1–5. IEEE, 2023.

519 Sedigheh Eslami and Gerard de Melo. Mitigate the gap: Investigating approaches for improving
520 cross-modal alignment in clip. *International Conference on Learning Representations (ICLR)*,
521 2025.

522 Abrar Fahim, Alex Murphy, and Alona Fyshe. It's not a modality gap: Characterizing and addressing
523 the contrastive gap. *ArXiv preprint: arXiv:2405.18570*, 2024.

524 Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand
525 Joulin, and Ishan Misra. ImageBind one embedding space to bind them all. In *IEEE/CVF
526 Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15180–15190, 2023.

527 Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced
528 research). 2009.

529 Anna Kukleva, Moritz Böhle, Bernt Schiele, Hilde Kuehne, and Christian Rupprecht. Temperature
530 schedules for self-supervised contrastive methods on long-tail data. In *International Conference
531 on Learning Representations (ICLR)*, 2023.

532 Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Deepak Gotmare, Shafiq R. Joty, Caiming Xiong, and
533 Steven C. H. Hoi. Align before fuse: Vision and language representation learning with momentum
534 distillation. In *Neural Information Processing Systems*, 2021.

540 Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Zou. Mind the gap:
 541 Understanding the modality gap in multi-modal contrastive representation learning. In *Advances*
 542 *in Neural Information Processing Systems (NeurIPS)*, 2022.

543

544 Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr
 545 Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer vision–*
 546 *ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, proceedings,*
 547 *part v 13*, pp. 740–755. Springer, 2014.

548 Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. CLIP4Clip: An
 549 empirical study of clip for end to end video clip retrieval. *Neurocomputing*, 508:293–304, 2021.

550

551 Marco Mistretta, Alberto Baldrati, Lorenzo Agnolucci, Marco Bertini, and Andrew D. Bagdanov.
 552 Cross the gap: Exposing the intra-modal misalignment in clip via modality inversion. In *Interna-*
 553 *tional Conference on Learning Representations (ICLR)*, 2025.

554 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 555 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever.
 556 Learning transferable visual models from natural language supervision. In *International Conference*
 557 *on Machine Learning (ICML)*, 2021.

558

559 Sameera Ramasinghe, Violetta Shevchenko, Gil Avraham, and Ajanthan Thalaiyasingam. Accept the
 560 modality gap: An exploration in the hyperbolic space. In *IEEE/CVF Conference on Computer*

561 *Vision and Pattern Recognition (CVPR)*, pp. 27253–27262, 2024.

562 Adriel Saporta, Aahlad Manas Puli, Mark Goldstein, and Rajesh Ranganath. Contrasting with symile:
 563 Simple model-agnostic representation learning for unlimited modalities. In *Neural Information*

564 *Processing Systems*, 2024.

565 Simon Schrodi, David T. Hoffmann, Max Argus, Volker Fischer, and Thomas Brox. Two effects, one
 566 trigger: On the modality gap, object bias, and information imbalance in contrastive vision-language
 567 models. In *International Conference on Learning Representations (ICLR)*, 2025.

568

569 Peiyang Shi, Michael C. Welle, Mårten Björkman, and Danica Kragic. Towards understanding the
 570 modality gap in CLIP. In *ICLR 2023 Workshop on Multimodal Representation Learning: Perks*

571 *and Pitfalls*, 2023.

572 Quan Sun, Yuxin Fang, Ledell Yu Wu, Xinlong Wang, and Yue Cao. EVA-CLIP: Improved training
 573 techniques for clip at scale. *ArXiv preprint: arXiv:2303.15389*, 2023.

574

575 Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive
 576 coding. *ArXiv preprint: arXiv:1807.03748*, 2018.

577

578 Feng Wang and Huaping Liu. Understanding the behaviour of contrastive loss. *IEEE/CVF Conference*
 579 *on Computer Vision and Pattern Recognition (CVPR)*, pp. 2495–2504, 2020.

580

581 Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through align-
 582 ment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pp.
 583 9929–9939. PMLR, 2020.

584

585 Junkang Wu, Jiawei Chen, Jiancan Wu, Wentao Shi, Xiang Wang, and Xiangnan He. Understanding
 586 contrastive learning via distributionally robust optimization. In *Advances in Neural Information*

587 *Processing Systems (NeurIPS)*, 2023.

588

589 Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging
 590 video and language. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*
 591 (*CVPR*), pp. 5288–5296, 2016. doi: 10.1109/CVPR.2016.571.

592

593 Shen Yan, Tao Zhu, Zirui Wang, Yuan Cao, Mi Zhang, Soham Ghosh, Yonghui Wu, and Jiahui Yu.
 594 VideoCoCa: Video-text modeling with zero-shot transfer from contrastive captioners, 2022.

595

596 Can Yaras, Siyi Chen, Peng Wang, and Qing Qu. Explaining and mitigating the modality gap in
 597 contrastive multimodal learning. In *Conference on Parsimony and Learning (CPAL)*, 2025.

594 Yuhui Zhang, Jeff Z. HaoChen, Shih-Cheng Huang, Kuan-Chieh Wang, James Zou, and Serena
595 Yeung. Diagnosing and rectifying vision models using language. In *International Conference on*
596 *Learning Representations (ICLR)*, 2023.

597 Bin Zhu, Bin Lin, Munan Ning, Yang Yan, Jiaxi Cui, Hongfa Wang, Yatian Pang, Wenhao Jiang,
598 Junwu Zhang, Zongwei Li, Wancai Zhang, Zhifeng Li, Wei Liu, and Liejie Yuan. LanguageBind:
599 Extending video-language pretraining to n-modality by language-based semantic alignment. In
600 *International Conference on Learning Representations (ICLR)*, 2024.

601

602

603

604

605

606

607

608

609

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 APPENDIX
649650 A.1 PROOFS AND MOTIVATIONS
651

652 **The gap is nearly the same for all the pairs.** We empirically prove our theoretical claims and
653 we compute both the Euclidean distance and the cosine similarities among true matching pairs,
654 which could both be interpreted as gap measures. As it is clear from Tab. 3, the variance in the
655 gap is extremely low and nearly 0, so all the points have the same gap. Thus, it is reasonable and
656 empirically-grounded to assume the same δ as a measure for the gap in Sec. 4, or to use the gap
657 measured as centroid distance in equation 2, being it a single measure for the overall gap. The same
658 findings can be applied to equation 6 as well. Concurrently, Zhang et al. (2023) demonstrated that the
659 gap vector is orthogonal to the text and image embeddings, and that the gap can be approximated
660 with a constant vector, confirming our findings.

661
662 Table 3: Mean and variance of the distance between the pairs in MSCOCO and MSR-VTT datasets.

Dataset	Distance Mean	Distance Var	Cos TP Mean	Cos TP Var
MSCOCO (T-I)	1.153	0.004	0.334	0.005
MSR-VTT (T-V)	1.163	0.009	0.319	0.012
MSR-VTT (T-A)	1.252	0.010	0.211	0.016

663
664 **Proof for global minimum in Fig. 3.** The 60 minimum in Fig. 3 is a direct outcome of the controlled
665 synthetic setup and the dynamics of the InfoNCE loss $\mathcal{L}^{(m \rightarrow n)}$. In the experiment, image (v) and
666 text (t) embeddings lie on two fixed rings with an initial angular offset of $\Delta = 120$. During the
667 optimization, the true image–text pairs are brought closer (current gap = θ), while the swapped pairs
668 (i.e., the negatives) remain at their original positions. Therefore, from the image anchor perspective,
669 these negatives are now at an angular distance of $\Delta - \theta$. This geometric property is exact and does
670 not rely on any symmetry or regular placement.

671 The InfoNCE loss balances attraction to the true pair and repulsion from negatives. Its derivative
672 becomes zero when the similarity to the true caption matches the similarity to the closest negative
673 (this occurs when $\theta = \Delta/2$). In our test case, that means CLIP converges to 60, as observed in Fig. 3.
674 Importantly, this also implies that InfoNCE never drives θ to 0, regardless of the initial gap Δ .

675 In contrast, our method includes the explicit \mathcal{L}_{ATP} term that always decreases as $\theta \rightarrow 0$, and
676 does not depend on the position of the negatives. Formally, our loss can be written as $\mathcal{L}_{\text{CL}_{\text{gap}}} =$
677 $\mathcal{L}^{(m \rightarrow n)} + \mathcal{L}_{\text{ATP}} + \mathcal{L}_{\text{CU}}$. The \mathcal{L}_{ATP} term is $\mathcal{L}_{\text{ATP}} \propto \|\mathbf{v} - \mathbf{t}\|_2^2 = 2(1 - \cos \theta)$, whose gradient is
678 $\partial \mathcal{L}_{\text{ATP}} = 2 \sin \theta$, always pulling the pair toward perfect overlap ($\theta = 0$) and has no counteracting
679 repulsion. The \mathcal{L}_{CU} term depends only on the class centroids, not on the intra-pair angle, so $\partial \mathcal{L}_{\text{CU}} = 0$.
680 Therefore, the total derivative becomes

$$681 \frac{\partial \mathcal{L}_{\text{CL}_{\text{gap}}}}{\partial \theta} = \frac{\partial \mathcal{L}^{(m \rightarrow n)}}{\partial \theta} + 2 \sin \theta, \quad (14)$$

682 with the extra term $2 \sin \theta$ shifting the sole stationary point to $\theta = 0$, which is now a global minimum.
683 This holds for any initial configuration.

684 **Proof for equation 6.** For a semantic class s and two modalities m, n , the class centroid under a
685 uniform shift δ is equation 5):

$$686 \mu_s^\delta = \frac{1}{2}((\mathbf{z}_s^m + \delta) + (\mathbf{z}_s^n + \delta)) = \mu_s^0 + \delta, \quad (15)$$

687 where

$$688 \mu_s^0 = \frac{1}{2}(\mathbf{z}_s^m + \mathbf{z}_s^n). \quad (16)$$

689 We expand the squared deviation:

$$\begin{aligned}
\|\mathbf{z}_s^m - \boldsymbol{\mu}_s^\delta\|_2^2 &= \|\mathbf{z}_s^m - (\boldsymbol{\mu}_s^0 + \boldsymbol{\delta})\|_2^2 \\
&= \|(\mathbf{z}_s^m - \boldsymbol{\mu}_s^0) - \boldsymbol{\delta}\|_2^2 \\
&= \|\mathbf{z}_s^m - \boldsymbol{\mu}_s^0\|_2^2 + \|\boldsymbol{\delta}\|_2^2 - 2 \langle \mathbf{z}_s^m - \boldsymbol{\mu}_s^0, \boldsymbol{\delta} \rangle.
\end{aligned} \tag{17}$$

Taking the expectation over classes yields:

$$\mathbb{E}_s \left[\|\mathbf{z}_s^m - \boldsymbol{\mu}_s^\delta\|_2^2 \right] = \mathbb{E}_s \left[\|\mathbf{z}_s^m - \boldsymbol{\mu}_s^0\|_2^2 \right] + \|\boldsymbol{\delta}\|_2^2 - 2 \mathbb{E}_s \left[\langle \mathbf{z}_s^m - \boldsymbol{\mu}_s^0, \boldsymbol{\delta} \rangle \right]. \tag{18}$$

Since $\boldsymbol{\delta}$ is approximately constant across classes and orthogonal to the span of semantic vectors Zhang et al. (2023), the cross-term vanishes:

$$\mathbb{E}_s \left[\langle \mathbf{z}_s^m - \boldsymbol{\mu}_s^0, \boldsymbol{\delta} \rangle \right] \approx 0. \tag{19}$$

Thus, we obtain equation 6:

$$\mathbb{E}_s \left[\|\mathbf{z}_s^m - \boldsymbol{\mu}_s^\delta\|_2^2 \right] \approx \mathbb{E}_s \left[\|\mathbf{z}_s^m - \boldsymbol{\mu}_s^0\|_2^2 \right] + \|\boldsymbol{\delta}\|_2^2. \tag{20}$$

A.2 EXPERIMENTS DETAILS

To process the AudioMNIST dataset, we compute the Mel spectrograms with 128 nmels, fmax at 8000, hop length equal to 512, and 2048 nfft. As image encoder, we employ a convolutional neural network (CNN) comprising a two-layer convolutional architecture with 32 and 64 filters, respectively, ReLU activation functions, Max Pooling, and a final MLP layer to map the features into the latent space. For the audio modality, we design a three-layer convolutional encoder with 16, 32, and 64 filters, ReLU activations, and an MLP layer to similarly project audio features into the latent space.

Clustering on MSCOCO. MSCOCO contains the information about objects that are present inside an image, and each image could contain more than one object, up to an undefined upper bound (also 10 different objects in an image), making standard clustering hard to apply. However, we devise a pipeline to perform cross-modal clustering on this dataset. We take into consideration images from the test dataset that contain only a single object, so that we have a univocal correspondence between the image and a single textual caption representing such an object. We report clustering results on MSCOCO in Tab.1.

Computing resources. The AV-MNIST experiments are conducted on an RTX4080 with 16GB. The CIFAR10 experiments on an A6000 with 48GB. The experiments on MS COCO and MSR-VTT experiments on a node with $4 \times$ A100.

A.3 ADDITIONAL EXPERIMENTS

We visualize the PCA latent spaces for two (MSCOCO) and three modalities (MSR-VTT) at initialization, after the conventional CLIP-based training, and after the training with the proposed method. Figure 6 shows the results. From Fig. 6, the gap closure is evident. Contrary to CLIP-based learning, our method can close the gap and spread modality embeddings leveraging the whole hypersphere space. Interestingly, Fig. 6 shows a similar behavior of the gap for the two- and three-modal case. At initialization, one modality tends to occupy more space, being more sparse, while the other/others are much more grouped. With CLIP-based learning, the modalities still tend to similarly lie in the space. However, regardless of the initialization and of the number of modalities, the proposed method completely closes the gap, building a more compact and well-aligned space.

Additional gap measures. Together with the gap formulation in equation 2, we further measure the gap as suggested in Schrödi et al. (2025), where the authors proposed the Relative Modality Gap (RMG) measure. We report the results in Tab. 4.

Closing the gap zero-shot. We consider the models trained on MS COCO and then evaluate them in a zero-shot fashion on the validation set of OpenImage-V7. The entire validation set contains 41620 images, each of them has associated one class among more than 20k image classes. We evaluate the

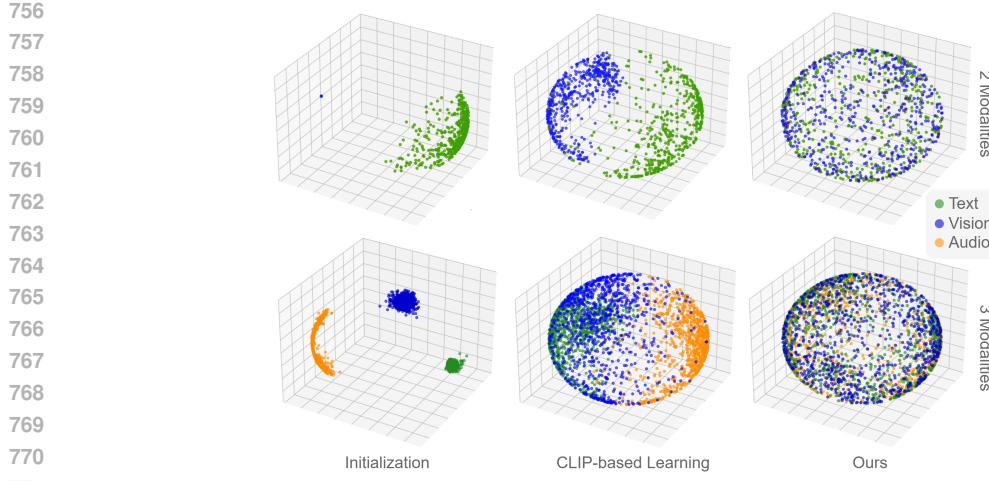


Figure 6: Left, latent spaces at initialization with the modality clusters clearly defined. Center, learned after training from CLIP-based learning that preserves the modality gaps. Right, space learned with our method, covering the whole space with overlapping clusters. PCA visualizations.

Table 4: Relative Modality Gap (RMG) measure for the methods on MSCOCO and MSR-VTT datasets.

	MSCOCO	MSR-VTT (T-V)	MSR-VTT (T-A)
CLIP (LT)	0.43	0.44	0.48
CLIP (FT)	0.27	0.40	0.46
Ours	0.17	0.33	0.44

ability of the proposed loss combinations to zero-shot closing the gap on the whole validation set, measuring the gap as in equation 2, the Cosine True Pairs (Cos TP) as in equation 3, and also the Relative Modality Gap (RMG) suggested in Schrödi et al. (2025). The results in Tab. 5 shows that the proposed method well-scale to large-scale datasets, achieving consistent results in all the metrics and outperforming the comparisons by considerably reducing the modality gap (absolute and relative). Furthermore, we also evaluate the clustering performance of the methods on this dataset. For this purpose, we select from the validation set the images containing one category among the 10 more common classes: 'Plant', 'Car', 'Person', 'Clothing', 'Food', 'Flower', 'Tree', 'Mammal', 'Wheel', 'Dog'. Moreover, we rebalance the dataset so that, in the final version, each class contains an equal number of samples. In this experiment too, the V-Measure achieved by the proposed method is the highest with respect to other methods, proving the effectiveness of the proposed method once again.

Table 5: Zero-shot gap and clustering results on OpenImage-V7 validation set.

Method	V-Measure \uparrow	Gap \downarrow	RMG \downarrow	Cos TP \uparrow
CLIP (LT)	16.7	0.459	0.462	0.210
CLIP (FT)	13.5	0.374	0.516	0.224
Ours	17.2	0.311	0.429	0.273

Relation between gap and clustering performance. To further highlight the relation between clustering performance and the gap, we plot the gap measure and the V-Measure during the training of the ResNet50 encoders on the CIFAR10 dataset. Fig 7 shows the results. As the gap decreases, the V-Measure evaluations increase reaching the best score with the proposed method, which better closes the gap.

Impact of the initialization. According to Liang et al. (2022) that first discovered the modality gap phenomenon, the gap exists at initialization and it is then preserved by the conventional contrastive loss. Nevertheless, the key reason for the modality gap is the contrastive behavior of the InfoNCE

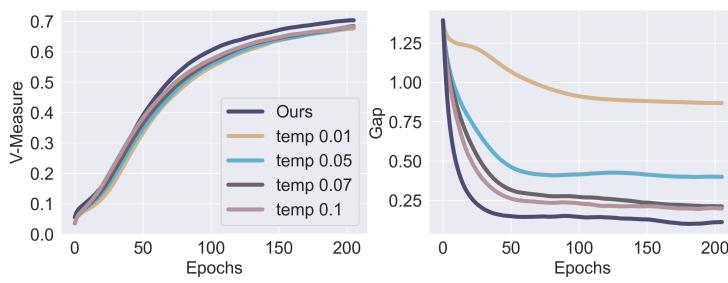


Figure 7: Clustering performance (V-Measure) and gap measure during training on the CIFAR10 dataset.

loss and not the initialization. To prove it, we compare two configurations with the same setup, model (ResNet50), and loss (InfoNCE). We enforce an initial sparsification of the learned space for the first epoch so that the encoders first learn to sparsify the embeddings, regardless of the modality. We do so by using as loss the uniformity loss over the hypersphere by Wang & Isola (2020). However, once the encoders learn the sparsification and we activate instead the conventional contrastive learning loss, the encoders start to separate the modalities again and the gap begins to be recreated, as we show in Fig. 8. We conduct the same experiment with the proposed losses to close the gap, and no effect is revealed with the initial sparsification as well. Therefore, although the gap is created at initialization, this does not impact the learning procedure and it is therefore the conventional contrastive loss function that tends to create such a gap.

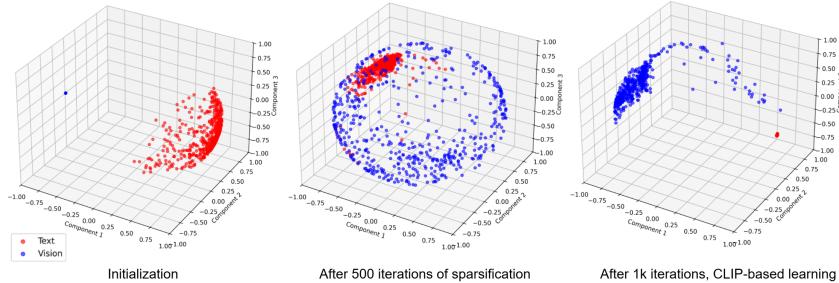


Figure 8: Latent space visualization during training of the conventional InfoNCE loss. The gap is created at initialization, as already known in Liang et al. (2022). However, even if we initially train the models to learn to sparsify the latent space (second plot) and reduce the gap, as soon as we reintroduce the conventional InfoNCE loss function, the gap is recreated (third plot). Experiment with ResNet50 from scratch on MS COCO dataset.

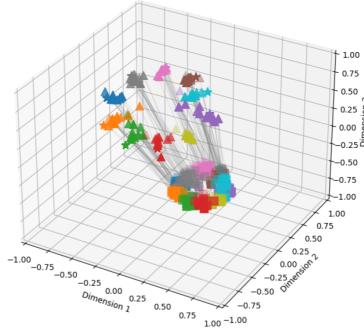
Experiments on the MultiBench classification task. We perform additional experiments with comparison methods in a different task. We rerun the experiments using CoMM Dufumier et al. (2025) repository in three datasets (MOSI, UR-FUNNY, MuSTARD) for the classification task from MultiBench. Then, we apply on CoMM the method proposed in Yaras et al. (2025) for a further comparison (CoMM + Yaras et al. (2025)). Finally, we add our proposed losses to the CoMM framework (CoMM + Ours line in the table below). We report gap and accuracy measures. As it is clear from Tab. 6, the proposed loss combination improves results in terms of accuracy with respect to CoMM and reduces the gap in all the datasets.

Visualizing the semantic stripes. In the presence of the modality gap, embeddings tend to form “semantic stripes”. Such stripes allow good performance in retrieval tasks since positive pairs have higher cosine similarity with respect to non-matching pairs, even though such similarity is far from the ideal 1.0, but low performance in group-wise tasks as embeddings coming from the same semantic concept lie separately in the space. We visualize these stripes in the AV-MNIST dataset with 3-dimensional embeddings in Fig. 9. Audio (triangles) and vision (squares) embeddings lie in separate regions of the latent space. Nevertheless, they form “semantic stripes” with each other, in which classes (colors) are matched in stripes with the corresponding class of a different modality.

864

865 Table 6: Classification results in MultiBench datasets compared with CoMM Dufumier et al. (2025)
866 and with CoMM + Yaras et al. (2025).

Method	MOSI		UR-FUNNY	
	Gap	Acc	Gap	Acc
CoMM Dufumier et al. (2025)	0.33	65.91	0.81	62.83
CoMM + Yaras et al. (2025)	0.28	66.42	0.63	61.10
CoMM + Ours	0.24	67.65	0.77	63.25

866 Figure 9: Semantic stripes in the AV-MNIST dataset for CLIP (LT).
867

890 **Intra-modal expressiveness.** Modality specific features are essential for some instance level or
891 downstream tasks. Reducing the gap between modalities does not imply discarding or collapsing the
892 unique characteristics of each modality. The proposed losses act only on the relative positioning of
893 positive pairs. They encourage samples that share the same semantics to move closer without limiting
894 the expressiveness of the individual encoders, as confirmed by our experiments. To support such
895 claims, we conduct an additional experiment to verify whether intra-modal discriminative features
896 are preserved. We use pretrained encoders on the MSR-VTT dataset, specifically Eva-CLIP for visual
897 features, BEATs for audio features, and BERT for textual features. After extracting the embeddings,
898 we apply both the k-NN algorithm and a simple linear classifier (In MSR-VTT each sample belongs
899 to one of the 20 categories). Table 7 reports the accuracy obtained by these two methods when applied
900 independently to the features of each modality. Each experiment is conducted 5 times. The results
901 show that the expressive power of each individual encoder remains unaffected by the introduction of
902 our additional losses and by the consequent reduction of the modality gap. In fact, both the kNN and
903 linear probing performance are essentially unchanged in the scenario where the gap is present (CLIP
904 (LT)) and in the scenario where the gap is reduced (Ours). These results mean that the discriminative
905 intra-modality features are preserved even when the gap is mitigated, thus the embeddings are likely
906 to preserve the intra-modality features.

907

908

909 Table 7: Accuracy of kNN and linear classification on modality-specific features (visual, textual,
910 and audio) using pretrained encoders. The results demonstrate that the proposed losses preserve
911 discriminative intra-modality features, as accuracy remains consistent across both the original and
912 reduced modality gap settings.

Method	Only Visual Features		Only Textual Features		Only Audio Features	
	kNN	Acc	kNN	Acc	kNN	Acc
CLIP (LT)	60.39 ± 2.9	63.84 ± 1.2	45.79 ± 2.4	48.49 ± 2.5	38.28 ± 2.5	42.36 ± 3.5
CLIP (FT)	59.62 ± 2.3	64.40 ± 0.8	44.61 ± 3.3	47.90 ± 0.5	37.71 ± 2.7	41.80 ± 3.4
Ours	60.39 ± 2.1	64.20 ± 1.0	44.34 ± 2.5	49.98 ± 1.8	37.50 ± 2.8	42.26 ± 2.6

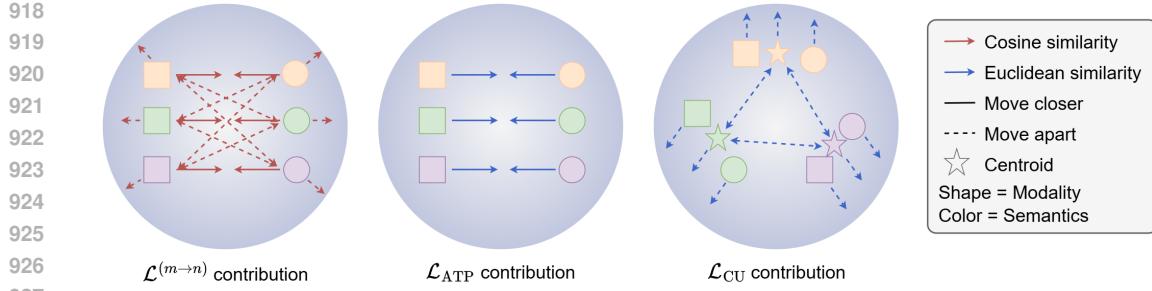


Figure 10: Schematic representation of losses contributions.

A.4 ABLATION STUDIES

Why using $\mathcal{L}^{(m \rightarrow n)}$ in combination with the proposed losses? \mathcal{L}_{ATP} and \mathcal{L}_{CU} regularise the global geometry (closing the modality gap and preventing centroid collapse), but they do not provide the instance-level repulsion that keeps different samples apart. Instead, the InfoNCE supplies exactly that term, so all three losses are somewhat complementary required to obtain representations that are simultaneously well aligned (no gap), uniformly spread, and discriminative. To prove our claims, we perform an ablation study to highlight the critical role of the InfoNCE loss ($\mathcal{L}^{(m \rightarrow n)}$) in achieving optimal performance. While overall semantic alignment can still be attained using a combination of \mathcal{L}_{ATP} and \mathcal{L}_{CU} , as evidenced by comparable results in R@5 and R@10, the absence of $\mathcal{L}^{(m \rightarrow n)}$ leads to a noticeable drop in R@1 performance. This indicates that the framework becomes less accurate in retrieving the top-1 correct caption when the contrastive loss is removed, underscoring its importance for fine-grained alignment between modalities.

Table 8: Ablation on the contribution of $\mathcal{L}^{(m \rightarrow n)}$ in combination with the proposed losses.

$\mathcal{L}^{(m \rightarrow n)}$	Cos TP \uparrow	Gap \downarrow	T2V R@1	T2V R@5	T2V R@10	V2T R@1	V2T R@5	V2T R@10
✗	0.71	0.05	30.0	68.3	83.0	31.2	69.3	82.0
✓	0.63	0.04	35.2	71.5	84.2	36.4	72.1	83.8

Furthermore, we visualize the contributions of each of the three losses $\mathcal{L}^{(m \rightarrow n)}$, \mathcal{L}_{ATP} , and \mathcal{L}_{CU} in Fig. 10. The first one, the InfoNCE loss, aligns multimodal matching pairs while spreading all non-matching ones apart. This dual behavior encourages the formation of a gap between modalities, as demonstrated in numerous previous works Liang et al. (2022); Shi et al. (2023). The InfoNCE loss function often becomes trapped in local minima, where the loss is minimized (semantic alignment is maximized) but a non-zero gap remains Shi et al. (2023). The \mathcal{L}_{ATP} loss forces the match among true pairs, consistently reducing the modality gap. However, the \mathcal{L}_{ATP} does not consider non-matching pairs, potentially making the latent space collapse in small regions, thus the contribution of \mathcal{L}_{CU} becomes crucial. Indeed, the \mathcal{L}_{CU} loss spreads the semantic centroids (i.e., the centroids of matching pairs) while preserving the alignment.

Ablation on \mathcal{L}_{ATP} and \mathcal{L}_{CU} . We conduct an ablation study for the proposed loss functions. We train from scratch two ResNet50 encoders with OpenCLIP Cherti et al. (2023) on the CIFAR10 dataset for 100 epochs. We set the temperature parameter of the CLIP objective to 0.07 and we perform sensitivity experiments, weighting the two proposed losses. To this purpose, we introduce two new hyperparameters for equation 10, λ_1 and λ_2 , combining the losses as

$$\mathcal{L}_{\text{CL}_{\text{gap}}} = \mathcal{L}^{(m \rightarrow n)} + \lambda_1 \mathcal{L}_{\text{ATP}} + \lambda_2 \mathcal{L}_{\text{CU}}. \quad (21)$$

Tab. 9 reports the results for retrieval (R@1) and clustering (V) across all the investigated configurations.

It is interesting to note that the configuration with $\lambda_1 = \lambda_2 = 0$ (the standard InfoNCE loss alone) achieves the lowest performance in both retrieval and clustering tasks. Furthermore, the results show a clear trend: increasing λ_2 consistently leads to improved performance, indicating the importance of

972
973 Table 9: Sensitivity study on λ_1 (row) and λ_2 (column) showing retrieval (R@1), clustering (V-
974 Measure), and their average (Avg) on the CIFAR10 dataset. Bold indicates the best value; underlines
975 highlight competitive second-best settings.

$\lambda_1 \setminus \lambda_2$	0.00	0.25	0.50	0.75	1.00
	R@1: 79.39 V: 61.65	R@1: 78.64 V: 63.67	R@1: 83.61 V: 67.70	R@1: 84.05 V: 66.87	R@1: 86.20 V: 66.64
0.00	Avg: 70.52 R@1: 80.59 V: 66.05	Avg: 71.16 R@1: 79.39 V: 63.26	Avg: 75.66 R@1: 80.15 V: 64.47	Avg: 75.46 R@1: 84.06 V: 69.93	Avg: 76.42 R@1: 83.82 V: 69.77
0.25	Avg: 73.32 R@1: 82.37 V: 68.96	Avg: 71.33 R@1: 79.34 V: 63.16	Avg: 72.31 R@1: 81.69 V: 66.90	Avg: 76.99 R@1: 82.43 V: 65.17	Avg: 76.80 R@1: 84.14 V: 70.42
0.50	Avg: 75.67 R@1: 79.23 V: 65.27	Avg: 71.25 R@1: 78.40 V: 62.89	Avg: 74.30 R@1: 82.20 V: 67.60	Avg: 73.80 R@1: 81.82 V: 65.74	Avg: 77.28 R@1: 84.03 V: 70.17
0.75	Avg: 72.25 R@1: 80.80 V: 66.84	Avg: 70.65 R@1: 78.13 V: 61.55	Avg: 74.90 R@1: 78.89 V: 60.67	Avg: 73.78 R@1: 84.21 V: <u>70.88</u>	Avg: 77.10 R@1: <u>84.64</u> V: 71.47
1.00	Avg: 73.82	Avg: 69.84	Avg: 69.78	Avg: 77.55	Avg: 78.06

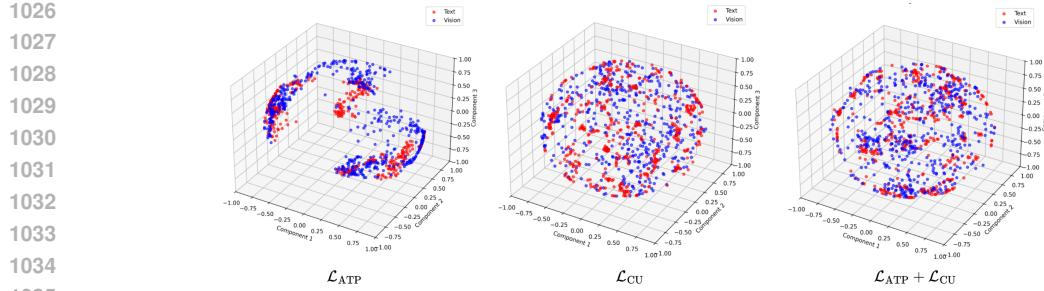
992 Table 10: Quantification of collapse in the latent space with (w/) and without (w/o) \mathcal{L}_{CU} .

	Gap	Cos TP	AV (t)	AV (v)	T2V R@1	V2T R@1
w/o \mathcal{L}_{CU}	0.08	0.76	0.122	0.091	30.86	31.64
w/ \mathcal{L}_{CU}	0.09	0.58	0.001	0.005	37.5	38.67

993 the \mathcal{L}_{CU} component. In contrast, performance appears less sensitive to variations in λ_1 , suggesting a
994 smaller but still complementary contribution of the \mathcal{L}_{ATP} term. Additionally, without \mathcal{L}_{CU} , the \mathcal{L}_{ATP}
995 only makes representations latent space collapsing in a small region of the space, as we show in Fig.
996 11, thus limiting the expressiveness of the space in downstream tasks. Involving both the proposed
997 losses, instead, builds a more compact and well-aligned latent space, as the third plot in Fig. 11
998 shows.

999 A quantitative measure of this collapse is also given in Tab. 6, where we report the gap, the cosine
1000 true pairs, and the additional Angular Value (AV) per modality. This metric measures the intra-modal
1001 average cosine similarity. It indicates how much the embeddings of a single modality are spread
1002 across the hypersphere. A value of this metric higher than 0 indicates that all the embeddings are
1003 very close to each other, while a value of 0 means that the intra-cosine similarity ranges from -1 to 1,
1004 indicating a good sparsification of the latent space. We conduct these experiments on the MSCOCO
1005 dataset with two ResNet50 encoder backbones. Table 6 shows that without \mathcal{L}_{CU} the space built by
1006 solely the InfoNCE and \mathcal{L}_{ATP} losses tends to have true pairs closer (Cos TP higher) at the cost of a
1007 lower sparsification (AV higher). This directly impacts retrieval results as the space is much more
1008 condensed and it is harder for the model to discriminate between matching and non-matching pairs.

1009 **Promoting centroids uniformity vs overall uniformity.** Wang & Isola (2020) proved and empirically
1010 demonstrated that the conventional InfoNCE loss can be decomposed into two objectives, one that
1011 pulls together embeddings from similar samples, and another that spreads all embeddings evenly on
1012 the hypersphere to avoid collapse and improve generalization. The latter (i.e., the embeddings spread)
1013 applies to all embeddings regardless of the class. However, enforcing uniformity at the sample level
1014 tends to scatter representations arbitrarily, potentially disrupting the tight alignment that the alignment
1015 term aims to enforce. For this reason, we adapt the uniformity principle to the multimodal setting
1016 by applying it not to individual embeddings, but to the centroids of aligned samples, that is, to the
1017 average representations of semantically matching pairs across modalities. This design retains the
1018 benefits of uniform coverage of the latent space while avoiding crucial interference with alignment.
1019 The \mathcal{L}_{CU} loss encourages the centroids of aligned multimodal samples to be well-separated on the
1020 hypersphere, effectively ensuring that different semantic concepts remain distinguishable. At the same
1021 time, the \mathcal{L}_{ATP} loss guarantees that all modalities representing the same semantic concept are tightly
1022



1036 Figure 11: Latent space visualization of different learned spaces in ablation studies, with the precise
1037 effect of each of the proposed losses. Experiments with ResNet50 from scratch on MS COCO dataset.
1038

1039 grouped around their centroid. To prove the limitations of the uniform spread of all embeddings, we
1040 perform an ablation study on the CIFAR10 dataset, in which we apply the uniformity loss proposed
1041 in Wang & Isola (2020) and our \mathcal{L}_{CU} loss. As it is clear from Tab. 11, without any other changes, the
1042 proposed \mathcal{L}_{CU} improves the performance of both retrieval (R@1) and clustering (V-Measure) over
1043 the uniformity loss of Wang & Isola (2020).

1044
1045 Table 11: Comparison between uniformity and our centroid-based uniformity loss.

Method	R@1 \uparrow	V-Measure \uparrow
$\mathcal{L}_{\text{uniform}}$ Wang & Isola (2020)	82.47	64.72
\mathcal{L}_{CU} (ours)	84.64	71.47

A.5 LIMITATIONS

1054 While our method effectively reduces the modality gap and enhances group-wise alignment across
1055 multiple modalities, it does not directly address potential limitations arising from highly imbalanced
1056 datasets or scenarios where semantic alignment is ambiguous or weakly defined. Indeed, we observe
1057 in our experiments, which contain datasets with diverse modalities imbalances, a different impact
1058 of the proposed method. While MSR-VTT and MSCOCO have one diverse matching sample for
1059 each modality (i.e., each image corresponds to a single textual description in MSCOCO and similar
1060 in MSR-VTT), so there is no imbalance between modalities, CIFAR10 and AV-MNIST have a
1061 different structure. AV-MNIST contains three modalities, where images and audio have the same
1062 size, while only 10 diverse textual captions are present in the dataset, making this modality slightly
1063 imbalanced. What is more, the CIFAR10 datasets contains only 60k images but only 10 testual
1064 captions corresponding to "A photo of class". The CIFAR10 is the most imbalanced dataset in our
1065 set of experiments. Coming to our method tests, we can observe that, even though our method
1066 outperforms previous ones in all the experiments, the increase in the CIFAR10 is the most tight
1067 (+0.9 in V-Measure), followed by the AV-MNIST (+4.4 in V-Measure), likely due to the imbalanced
1068 modalities structure. Indeed, we find the highest improvements in the balanced datasets, in which we
1069 got +10.65 in V-Measure for MSCOCO and +8.8 in MSR-VTT.

1070 Moreover, although we validate our approach on both bimodal and trimodal benchmarks, extending
1071 the evaluation to more complex or underexplored modality combinations remains future work.
1072

1073
1074
1075
1076
1077
1078
1079