

CodeBind: Decoupled Representation Learning for Multimodal Alignment with Unified Compositional Codebook

Zeyu Chen* Jie Li* Kai Han†

Visual AI Lab, The University of Hong Kong
zeyuc07@connect.hku.hk, jieli_cn@163.com, kaihanx@hku.hk

Abstract

Multimodal representation alignment is pivotal for large language models and robotics. Traditional methods are often hindered by cross-modal information discrepancies and data scarcity, leading to suboptimal alignment spaces that overlook modality-unique features. We propose CodeBind, a framework that optimizes multimodal representation spaces through a modality-shared-specific codebook design. By incrementally aligning target and bridging modalities, CodeBind bypasses the need for fully paired data. Unlike traditional hard alignment, CodeBind decomposes features into shared components for semantic consistency and specific components for modality-unique details. This design utilizes a compositional vector quantization scheme, where a shared codebook bridges modality gaps and modality-specific codebooks mitigate representation bias by preventing dominant modalities from overshadowing others. Validated across nine modalities (text, image, video, audio, depth, thermal, tactile, 3D point cloud, EEG), CodeBind achieves state-of-the-art performance in multimodal classification and retrieval tasks. Project page: <https://visual-ai.github.io/codebind>

1 Introduction

Multimodal representation alignment enables Large Language Models (LLMs) to interact with the physical world through diverse sensory inputs (Alayrac et al., 2022; Liu et al., 2023; Li et al., 2022a, 2023). In complex multi-sensor systems like robotic perception, aligning specialized modalities, including audio, depth, thermal, and point clouds, with the established vision-language space is essential for achieving advanced cross-modal intelligence.

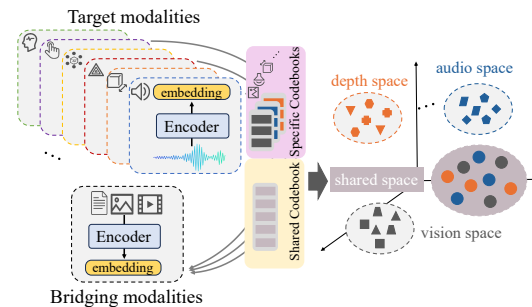


Figure 1: **Multi-modal alignment via codebook.** Target modalities are partially aligned with bridging modalities via codebooks, resulting in a shared space. Unique features from both bridging and target modalities are preserved in specific space.

Scaling alignment beyond the mature vision-language domain presents two critical challenges that hinder practical deployment. First, intrinsic information gaps exist between modalities (Liang et al., 2022; Shi et al., 2023; Ramasinghe et al., 2024). Compressing heterogeneous data into a shared space often results in a “least common denominator” effect. Such suboptimal alignment underperforms on datasets rich in modality-unique features, a frequent bottleneck in robotic manipulation tasks (Tjandrasuwita et al., 2025; Jiang et al., 2023). Second, multimodal data imbalance induces significant representation bias. Unlike web-scale image-text pairs, specialized modalities often have limited data availability, causing the alignment process to bias heavily toward dominant modalities like vision (Wang et al., 2024; Lyu et al., 2024a). This weakens inter-modal interaction and causes rare modalities to be overshadowed by dominant ones (Lin et al., 2025). Such bias is particularly detrimental in multi-sensor systems, where the unique contributions of rare sensors must be preserved rather than overwhelmed by visual sensors.

Current efforts to enhance multimodal alignment typically rely on training on massive paired datasets, including multimodal data crowdsourcing approaches (Zhu et al., 2024a; Han et al., 2024) or

*Equal contribution.

†Corresponding author.

projecting multimodal data into a shared space anchored by vision-language alignment models (Girdhar et al., 2023; Lei et al., 2024). However, such hard alignment forces entire feature vectors into narrow subspaces and often fails to protect the distinctive richness of individual modalities, leading to degraded performance in multimodal understanding, especially when certain modalities are informatively weak or data-scarce (Jiang et al., 2023; Tjandrasuwita et al., 2025).

To address these challenges, we propose **Code-Bind**, a framework that reframes multimodal alignment via discrete vector quantization (VQ). Our core innovation, the modality-shared-specific codebook design, optimizes the multimodal feature space and preserves representational fidelity.

To alleviate suboptimal alignment brought by intrinsic modality gap, we decompose representations into shared components (capturing modality-agnostic invariants) and specific components (preserving modality-unique details). Unlike conventional hard alignment approaches (Radford et al., 2021; Ilharco et al., 2021; Girdhar et al., 2023; Lei et al., 2024; Lyu et al., 2024b; Guzhov et al., 2022), this approach ensures semantic alignment without sacrificing unimodal representational fidelity, empowering intra-modality tasks like fine-grained retrieval. To mitigate the representation bias, we employ VQ as a distribution-agnostic feature base. By mapping features from different modalities to the same set of discrete semantic centers, VQ provides a structure that prevents dominant modalities from overwhelming the latent space, thereby ensuring feature consistency across unbalanced datasets.

Building on this, our *modality-shared-specific codebook* utilizes a shared codebook for cross-modal consistency and specific codebooks for modality-unique subspaces. By integrating *compositional VQ*, we exponentially expand representational capacity using a compact set of codevectors, ensuring high efficiency and utilization without the computational overhead of traditional large-scale codebooks.

Our contributions are summarized as follows:

- We propose a novel representation decoupling approach that disentangles modality-shared and -specific components, preventing semantic interference and enhancing cross-modal alignment consistency.
- We introduce a modality-shared-specific codebook design using compositional vector quantization, resulting in a unified alignment space and

maintaining high representational capacity.

- We demonstrate the scalability of our approach across nine modalities (text, image, video, audio, depth, thermal, tactile, 3D point cloud, and electroencephalogram (EEG)), achieving state-of-the-art (SOTA) results in cross-modal classification and retrieval while preserving fine-grained details.

2 Related Work

Multimodal alignment Building on vision-language pretraining (Radford et al., 2021; Zhai et al., 2023), multimodal alignment has expanded to audio (Guzhov et al., 2022; Wu et al., 2022), 3D point clouds (Zhang et al., 2022; Dharmasiri et al., 2024; He et al., 2024), and neural signals (Chen et al., 2024; Yang et al., 2024). Current strategies follow three paradigms. First, unified architectures employ a universal encoder with modality-specific projectors to create a shared space (Zhang et al., 2023; Han et al., 2024; Faye et al., 2024). These methods often lack scalability by requiring full-scale retraining when incorporating new modalities. Second, bridging methods align various modalities via bridging modalities (text and image) (Girdhar et al., 2023; Zhu et al., 2024a; Lei et al., 2024; Lyu et al., 2024b; Fu et al., 2024) or by overlapping pretrained bi-modal models (Wang et al., 2023, 2025a).

Though they show strong scalability (Han et al., 2023; Su et al., 2023), forcing alignment of semantic-scarce modalities to semantic-rich anchors often suppresses modality-unique features, leading to sub-optimal representation. Third, Mixture-of-Experts (MoE) strategies dynamically fuse modality encoders through routers (Lyu et al., 2024a; Zhou et al., 2025). While efficient, they are prone to modality collapse under multimodal data imbalance. While recent works attempt to mitigate data scarcity by generating pseudo-paired multimodal data (Wang et al., 2025a, 2024, 2025b), introducing synthetic datasets (Zhu et al., 2024a), or designing regularization to maintain performance under limited paired data (Gröger et al., 2025), these approaches are often labor-intensive and susceptible to semantic noise. Our approach bypasses these limitations by optimizing the bridging method via partial alignment. We utilize codebook as a universal feature base to accommodate the expression of individual modalities regardless of data imbalance.

Multimodal representation learning Recent studies highlight a persistent modality gap in con-

trastive learning (Liang et al., 2022; Zhang et al., 2024b), suggesting that pursuing hard alignment often degrades performance by neglecting modality-specific features (Jiang et al., 2023; Tjandrasuwita et al., 2025). However, such features remain vital for multimodal retrieval and tasks such as sentiment analysis (Zeng et al., 2025; Yang et al., 2023). To balance multimodal consistency with heterogeneity, existing methods utilize inter/intra-modality regularization or causal components decomposition to disentangle task-related information from redundant noise (Jiang et al., 2023; Liu et al., 2025; Yang et al., 2023). Others, like InfoBridge (Li et al., 2025), modify the InfoNCE loss to provide a protective margin that prevents over-alignment. However, these strategies are primarily validated on data-rich image-text benchmarks and often focus on modality fusion rather than alignment. Consequently, the efficacy of decoupled features remains under-explored, particularly the alignment robustness of shared components and the informativeness of specific components in data-scarce scenarios. In contrast, CodeBind decouples embeddings into shared and specific components to achieve robust partial alignment while explicitly safeguarding intra-modal features, validated on diverse and uncommon modalities like tactile and EEG.

Codebook-based unified representation space

VQ has demonstrated remarkable efficacy in both unimodal and cross-modal representation learning (Van Den Oord et al., 2017a; Esser et al., 2021; Zheng et al., 2022; Sargent et al., 2023; Li et al., 2022b). Recent studies have extended joint codebook architectures from vision-language alignment (Zheng et al., 2024; Duan et al., 2022) to broader multimodal scenarios (Liu et al., 2022; Xia et al., 2024), with CMG (Xia et al., 2024) and FCID (Huang et al., 2025) notably employing dual-disentanglement for fine-grained decomposition and alignment across audio, vision, and text. However, existing codebook designs prioritize unified representation spaces, leaving their effectiveness in modality-specific information preservation unexplored. Furthermore, traditional large-scale codebooks frequently suffer from modality collapse and training inefficiencies (Zheng et al., 2024; Duan et al., 2022). To bridge these gaps, we propose a modality-shared-specific codebook design powered by compositional VQ to effectively scale codebook capacity across modalities.

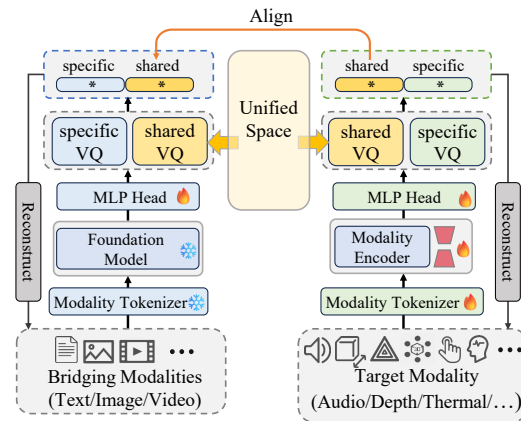


Figure 2: **Alignment across modalities.** Embeddings from bridging and target modalities are decoupled and quantized into shared and specific components, where shared ones are aligned within a unified space.

3 Method

CodeBind facilitates scalable multimodal alignment without exhaustive pairings by aligning text and vision as bridging modalities with diverse target modalities (Sec. 3.1). As shown in Fig. 1, our framework decouples representations from bridging and target modalities into shared semantic invariants and modality-unique features (Sec. 3.2), which are then discretized via modality-shared-specific codebook with compositional VQ for efficiency (Sec. 3.3).

3.1 Architecture

Following ImageBind (Girdhar et al., 2023) and ViT-Lens (Lei et al., 2024), we align target modalities to a unified space using a frozen vision-language foundation model (e.g., OpenCLIP (Radford et al., 2021; Ilharco et al., 2021)) as semantic bridge as shown in Fig. 2. To decouple cross-modal semantics from modality-unique features, encoder outputs are projected into a shared embedding z_{shared}^M and a specific embedding z_{spec}^M , where only the shared part is used for alignment. To mitigate representation bias, we propose a modality-shared-specific codebook using compositional VQ. A universal shared codebook quantizes z_{shared}^M while specific codebooks handle modality-specific features. Finally, a Transformer decoder reconstructs the original data from the joint embedding $[z_{\text{shared}}^M, z_{\text{spec}}^M]$ to ensure information retention. Our plug-and-play design allows target modality encoders to be resumed from bridging methods like ImageBind and finetuned via LoRA (Hu et al., 2022).

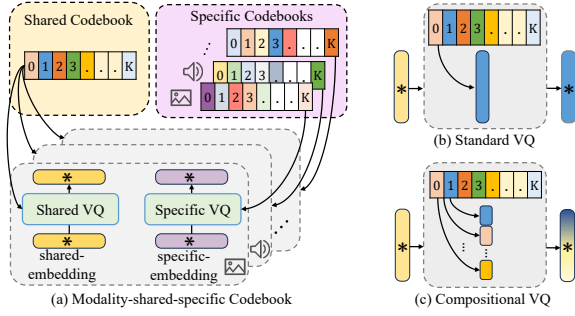


Figure 3: **Modality-shared-specific codebook for multi-modal alignment.** (a) The shared embeddings of different modalities use the same codebook for VQ, while the specific embeddings of each modality have their own specific codebooks. (b) The standard VQ matches each input embedding to a single codevector. (c) Compositional VQ utilizes a combination of multiple low-dimensional codevectors to reconstruct a complete embedding.

3.2 Decoupled Representations

Multimodal alignment traditionally maximizes $I(z^{\mathcal{M}_1}, z^{\mathcal{M}_2})$ to synchronize mutual information between shared semantics y and modality embedding: $I(y; z^{\mathcal{M}_1})$ and $I(y; z^{\mathcal{M}_2})$. However, this often leads to over-alignment, where aligning modality-specific noise sacrifices unique semantic richness. We posit that unimodal feature decoupling is a prerequisite for high-quality cross-modal alignment. Each modality’s embedding is decoupled into two orthogonal components. Shared components ($z_{\text{shared}}^{\mathcal{M}}$) capture cross-modal semantic invariants (e.g., the concept of a “cat”) to facilitate high-level tasks like cross-modal classification and retrieval. Specific components ($z_{\text{spec}}^{\mathcal{M}}$) encode modality-unique details (e.g., cat color or fur texture, which may be absent in audio or depth). Enforcing orthogonality and constraining $z_{\text{spec}}^{\mathcal{M}}$ to a uniform distribution (Jiang et al., 2023) can satisfy $I(y; z^{\mathcal{M}}) \approx I(y; z_{\text{shared}}^{\mathcal{M}})$. Consequently, the objective shifts to maximizing $I(z_{\text{shared}}^{\mathcal{M}_1}; z_{\text{shared}}^{\mathcal{M}_2})$. This ensures robust semantic alignment while preserving the specific part for fine-grained retrieval and high-fidelity reconstruction.

3.3 Modality-Shared-Specific Codebook

Preliminary: vector quantization In standard VQ, a quantizer $Q(\cdot)$ discretizes an input embedding $z \in \mathbb{R}^D$ by mapping it to the nearest codevector $c_k \in \mathbb{R}^D$ within a codebook \mathcal{C} ,

$$Q(z; \mathcal{C}) = \min_k \{\|z - c_k\|\}, \mathcal{C} = \{c_k\}_{k=1}^K$$

VQ provides two key advantages in multimodal alignment: (1) Mapping heterogeneous embeddings into a discrete, shared codebook space ensures uniform data treatment. This alleviates representation bias and prevents dominant modalities from overwhelming the alignment space. (2) Codevectors serve as learnable feature bases representing core semantic concepts, ensuring the aligned space is grounded in consistent semantic centers.

Compositional codebook Existing VQ-based multimodal alignment faces two critical limitations. (1) feature entanglement, where shared and specific information are mixed or inadequately quantized (Zheng et al., 2024; Xia et al., 2024), leading to noise leakage and distribution bias; and (2) scalability issues, where increasing codebook size to capture rich details triggers high computational overhead and codebook collapse (Zhu et al., 2024b; Zheng and Vedaldi, 2023).

To overcome these, we propose a modality-shared-specific codebook design (Fig. 3) that utilizes a universal codebook $\mathcal{C}_{\text{shared}}$ to capture the semantic core for alignment and distinct codebooks $\mathcal{C}_{\text{spec}}^{\mathcal{M}}$ for modality-specific features. For instance, the concept ‘striking’ can be quantized in the shared space as the general act of hitting, while specific codebooks capture its modality-dependent meanings: the sharp crack in audio, the high-speed motion in video, or the physical pressure in touch. To enhance the expressiveness of such features without expanding the codebook size, we implement *compositional VQ*. In this design, a d -dimensional embedding is partitioned into m sub-vectors, each of dimension $d^* = d/m$, which are then independently quantized. This approach exponentially expands the representation space to K^m combinations, forming high-fidelity representations from a compact set of codevectors while maintaining computational efficiency.

3.4 Implementation

Training objective The training of CodeBind is guided by a multi-task objective that harmonizes semantic alignment, feature decoupling, and codebook stability. We use InfoNCE loss ($\mathcal{L}_{\text{align}}$) for cross-modal alignment, complemented by orthogonal ($\mathcal{L}_{\text{orth}}$), reconstruction ($\mathcal{L}_{\text{recon}}$), and uniform (\mathcal{L}_{uni}) losses to ensure effective feature disentanglement and representational integrity. To maintain codebook robustness, we utilize Exponential Moving Average (EMA) updates (Van Den Oord

et al., 2017b) with commitment loss ($\mathcal{L}_{\text{commit}}$), alongside dynamic reinitialization and specialized codevector-level regularizations (\mathcal{L}_{ctr} , $\mathcal{L}_{\text{cuni}}$) to prevent collapse and enhance semantic distinctiveness. Furthermore, a Cross-Modal Code Matching loss (\mathcal{L}_{cm}) is integrated to refine the shared codebook’s alignment by synchronizing distance distributions across paired modalities (Liu et al., 2022). To eliminate manual tuning costs and stabilize multi-objective learning among different modalities, we design an adaptive loss balancing strategy. By maintaining a running estimate of each loss magnitude via EMA, the weights of other loss functions are dynamically rescaled to match the scale of $\mathcal{L}_{\text{align}}$. To further smooth the optimization landscape, update intervals for these weights are increased linearly throughout the training process. See App. A.1 for details.

Scalable multi-path alignment Our framework employs a multi-path alignment strategy where each target modality (e.g., audio) is concurrently aligned with multiple bridging modalities (e.g., audio-text and audio-image pairs). To ensure a cohesive unified space, we further enforce internal alignment within the bridging modalities (e.g., image-text). Unlike naive pairwise methods, this joint alignment paradigm utilizes distinct codebooks for different target-bridging paths to effectively decouple modality-specific features while maintaining global semantic consistency. This design also enables efficient scaling: new datasets for trained modalities benefit from transfer learning via pre-trained codebooks, while new modalities are integrated by training new codebooks. During inference, cross-modal alignment only requires shared embeddings, allowing specific parts to be discarded for a more lightweight pipeline. These components can be enabled to facilitate fine-grained intra-modal tasks. In both scenarios, the reconstruction module is omitted to ensure efficiency.

4 Experiments

We evaluate CodeBind on 9 modalities, including text, image, video, audio, depth, thermal, 3D point cloud, tactile, and EEG. Tab. 1 summarizes the datasets (App. B.1 for details). Our codebook design and training objectives are integrated into two SOTA baselines, ImageBind (Girdhar et al., 2023) and ViT-Lens (Lei et al., 2024), resulting in CodeBind-IB and CodeBind-VL. We employ compositional VQ with 1024 shared and 256 specific

Modality	Dataset	#Cls	Metric	#Test
Image	IN1K (Russakovsky et al., 2015)	1000	Acc	50000
	P365 (Zhou et al., 2014)	365	Acc	36500
Video	K400 (Kay et al., 2017)	400	Acc	19759
	MSR-VTT (Xu et al., 2016)	-	Acc	2990
Depth	NYU-D (Silberman et al., 2012)	10	Acc	654
	SUN-D (Song et al., 2015)	19	Acc	4660
Audio	Audioset (Gemmeke et al., 2017)	527	Acc	17132
	ESC (Piczak, 2015)	50	Acc	2000
	Clotho (Drossos et al., 2020)	-	Recall	1046
	AudioCaps (Kim et al., 2019)	-	Recall	813
	VGGs (Chen et al., 2020)	309	Acc	15434
Thermal	LLVIP (Jia et al., 2021)	2	Acc	16604
	FLIR_v2 (FLIR, 2018)	10	Acc	1521
Tactile	TAG-M (Yang et al., 2022)	20	Acc	29879
	TAG-H/S (Yang et al., 2022)	2	Acc	29879
	TAG-R/S (Yang et al., 2022)	2	Acc	8085
EEG	IN-EEG (Spampinato et al., 2017)	40	Acc	1997
3D	ModelNet40 (Wu et al., 2015a)	40	Acc	2468

Table 1: **Dataset Statistics** across nine aligned modalities.

codevectors, each with a dimensionality of 8. Training is initialized using pre-trained ImageBind and ViT-Lens, with a learning rate of 5×10^{-4} on eight NVIDIA RTX 3090 GPUs (App. B.2 for details).

4.1 Main Comparison

Similar to ImageBind and ViT-Lens, we evaluate our method on multimodal classification and retrieval tasks. Tab. 2 demonstrates consistent improvements with our approach compared to baselines.

Multimodal classification We extract aligned feature of any modality and assign it to the class with the highest similarity to the textual prompt embeddings defined by Radford et al. (2021). In *image and video classification* (ImageNet Russakovsky et al., 2015, Places365 Zhou et al., 2014, Kinetics400 Kay et al., 2017), ImageBind maintains parity with OpenCLIP (Ilharco et al., 2021) while our plug-in codebooks yield consistent gains. In *depth scene classification*, CodeBind-IB achieves a notable improvement of +5.3%/+10.6% on NYU-D (Silberman et al., 2012) and SUN-D (Song et al., 2015). In *audio classification*, we outperform ImageBind by +3.5%/+2.7%/+4.1% across three audio datasets (AudioSet Gemmeke et al., 2017, VGGSound Chen et al., 2020, ESC Piczak, 2015). In *thermal classification*, CodeBind-IB achieves substantial boosts (+32.1% on LLVIP (Jia et al., 2021), +50.6% on FLIR_v2 (FLIR, 2018)). Furthermore, we observe remarkable improvement in Tough-and-go (Yang et al., 2022) and ImageNet-EEG (Spampinato et al., 2017) for tactile and EEG visual concept classification. In Tab. 2 (b), CodeBind-VL achieves consistent performance gains across tasks. In *point cloud*

	Image		Video		Depth		Audio				Thermal		Tactile			EEG	
	IN1K	P365	K400	MSR-VTT	NYU-D	SUN-D	AudioSet	VGGS	ESC	Clotho	AudioCaps	LLVIP	FLIR_v2	TAG-M	TAG-H/S	TAG-R/S	IN-EEG
ImageBind	77.7	45.4	50.5	36.1	54.0	35.1	17.6	27.8	66.9	6.0/28.4	9.3/42.3	63.4	46.6	24.2	65.7	69.8	18.4
CodeBind	79.3	55.5	54.4	37.8	59.3	45.7	21.1	30.5	71.0	6.9/28.6	13.3/53.8	95.5	97.2	42.6	83.9	78.2	33.1

(a) CodeBind-IB integrates our method into ImageBind. For retrieval, Recall@1 is reported on MSR-VTT and ESC, and Recall@1/Recall@10 are reported on Clotho and AudioCaps. For classification, Acc@1 is reported on all other datasets, except for AudioSet, where mAP is reported.

	Depth		Audio				Tactile			EEG	3D	
	NYU-D	SUN-D	AudioSet	VGGS	ESC	Clotho	AudioCaps	TAG-M	TAG-H/S	TAG-R/S	IN-EEG	ModelNet40
ViT-Lens	68.5	52.2	26.7	31.7	75.9	8.1/31.2	14.4/54.9	65.8	74.7	63.8	41.8/42.7	70.6/94.4
CodeBind-VL	71.1	54.8	29.2	39.5	78.8	8.5/32.8	15.6/55.0	67.6	76.1	72.8	54.5/54.1	78.3/96.5

(b) CodeBind-VL integrates our method into ViT-Lens, and classification results from EEG and 3D point clouds are included.

Table 2: **Multi-modal classification and retrieval results on diverse benchmarks.** CodeBind-IB (a) and CodeBind-VL (b) maintain the dataset and modality setting with ImageBind and ViT-Lens respectively, for comparison fairness. The results highlight our model’s superior alignment capabilities across nine modalities.

Method	Video	Depth	Audio				Thermal		3D		
	MSR-VTT	NYU-D	SUN-D	AudioSet	VGGS	ESC	Clotho	AudioCaps	LLVIP	FLIR_v2	ModelNet-40
FreeBind	–	–	–	19.7	–	93.6	13.7/–	29.2/–	–	–	–
OmniBind	–	–	–	25.1	–	93.5	23.3/49.5	46.7/79.7	–	–	–
LanguageBind	44.8	–	–	30.0	38.6	94.0	16.7/52.0	19.7/67.6	87.2	48.0	–
OneLLM-D	–	50.9	29.0	–	–	–	–	–	–	–	–
OneEncoder	–	–	28.4	–	80.1	–	–	–	–	–	–
UNIALIGN	37.0	71.4	58.2	23.8	–	70.5	–	11.7/49.3	–	–	55.9/84.3
CodeBind-IB	37.8	59.3	45.7	21.1	30.5	71.0	6.9/28.6	13.3/53.8	95.5	97.2	–
CodeBind-VL	–	71.1	54.8	29.2	39.5	78.8	8.5/32.8	15.6/55.0	–	–	78.3/96.5

Table 3: Compared to other strong SOTA methods that use different data crowdsourcing approaches, our method is comparable without such approaches.

classification using ShapeNets (Wu et al., 2015a) for training and evaluation on ModelNet40 (Wu et al., 2015b), it improves Acc@1 by +7.7% and Acc@5 by +2.1% over ViT-Lens.

Multimodal retrieval In *video and audio retrieval*, CodeBind-IB consistently outperforms the baseline, yielding a +1.7% gain on MSR-VTT (Xu et al., 2016) and notable improvements on Clotho (Drossos et al., 2020) (+0.9% Recall@1, +0.2% Recall@10). These gains are even more pronounced on AudioCaps, with substantial increases of +4.0% and +11.5%. CodeBind-VL further confirms the consistent enhancement of our method.

Comparison with other strong SOTA methods FreeBind (Wang et al., 2024) and OmniBind (Wang et al., 2025b) generate pseudo-data pairs retrieved from pre-trained bi-modal expert models to bridge disparate spaces. LanguageBind (Zhu et al., 2024a) scales up to 10M dataset, including synthetic data for scarce modalities like thermal. Unified architectures like OneLLM (Han et al., 2024) and OneEncoder (Faye et al., 2024) forcibly align all modalities into a single encoder, exacerbating the risk

of catastrophic forgetting and sub-optimal alignment. UNIALIGN (Zhou et al., 2025) achieves extreme parameter efficiency (7.8M) through a modality-aware MoE schema but overlooks fine-grained, modality-specific details. In contrast, our codebook-based alignment leverages naturally paired data rather than labor-intensive or synthetic data expansion. By utilizing compositional VQ to represent shared and specific features, our method achieves comparable performance with efficient parameterization (~50M), as shown in Table 3.

4.2 Embedding Space Optimization

We qualitatively evaluate the reduction in modality gap through t-SNE visualization (Van der Maaten and Hinton, 2008) of embedding distributions. In Fig. 4, while ImageBind exhibits a clear inter-modal gap, our shared embeddings are closely clustered and mixed, directly correlating with our improved alignment accuracy. Furthermore, we visualize the distribution of shared and specific components to validate their decoupled properties. As illustrated in Fig. 5, our method effectively separates shared semantics from modality-specific features in both thermal and vision domains. Detailed




Figure 4: **Visualization of embeddings in unified space.** T-SNE visualization of sampled embeddings by ImageBind (left) and CodeBind-IB (right) using AudioSet (Gemmeke et al., 2017). The paired embeddings are linked by a grey line.




Figure 5: **Visualization of decoupled embeddings of image and thermal** from FLIR_v2 (FLIR, 2018).

codevector-level analysis is available in App. C.1.

4.3 Retention of Modality-Unique Information

To assess the retention of modality-unique information, we conduct fine-grained intra-modal retrieval and linear probing to distinguish shared semantics from fine-grained modality-unique details.

Fine-grained intra-modal retrieval To evaluate fine-grained details within our decoupled embeddings, we perform an image-to-image retrieval on fine-grained datasets, Stanford Dogs (Khosla et al., 2011) and Oxford-IIIT Pet (Parkhi et al., 2012). In Tab. 4, our concatenated embeddings surpass ImageBind in average Top-10 recall, validating that the fusion of shared and specific components enhances intra-modality information preservation. Notably, our specific embeddings alone achieve performance comparable to ImageBind, highlighting their representational utility. Qualitative results

Dataset	Stanford Dogs	Oxford Pet Cats	Oxford Pet Dogs
Sample Size	6000	2371	4978
ImageBind	50.4	87.0	92.2
Ours			
Specific	48.9	83.7	91.1
Shared	63.5	<u>88.3</u>	94.6
Concatenated	<u>60.2</u>	88.4	<u>94.4</u>

Table 4: **Fine-grained retrieval recall.** The best result is **bolded**, and the second best result is underlined. Our decoupled components achieve higher accuracy on fine-grained tasks.

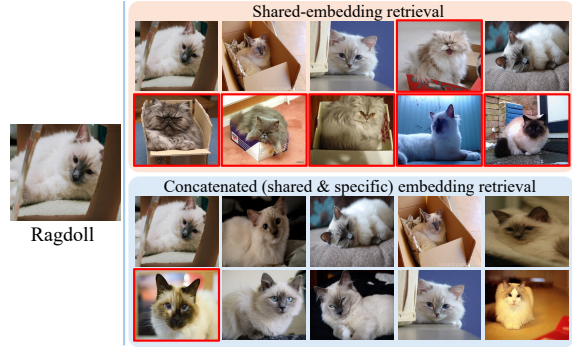


Figure 6: **Fine-grained retrieval results.** Source image ‘ragdoll’ and retrieval results with incorrect matches in red boxes.




Figure 7: **Linear probing** for category names and fine-grained details using shared and specific embeddings from ImageNet (Russakovsky et al., 2015) samples.

in Fig. 6 further confirm that specific embeddings refine visual similarity, effectively complementing shared semantics with fine-grained details.

Linear probing for fine-grained details We further conduct linear probing on both semantic and fine-grained attributes. Specifically, we leverage a vision language model (VLM), Qwen2.5-VL-72B (Bai et al., 2025b), to curate a fine-grained attribute taxonomy for ImageNet, extracting non-semantic properties, such as lighting, texture, and background color, that are decoupled from category labels. We then train independent linear probes to classify the semantic labels and the fine-grained labels. A probing swap is then performed to assess cross-component predictive power. Fig. 7 shows that specific embeddings achieve superior accuracy and faster convergence in fine-grained attribute prediction, validating their proficiency in encoding physical nuances. Interestingly, the comparable performance of specific embeddings in semantic classification suggests a natural coupling between category and fine-grained attributes, a phenomenon further corroborated by our mutual information analysis (in App. C.2).

Configuration		Space		Accuracy (%)	
Method	Text Density	Shared	Spec.	Concat	Sum
ImageBind	Category	—	—	94.4	89.0
CodeBind-IB	Category	✓	✓	<u>95.7</u>	<u>91.6</u>
		✓	✗	95.5	89.8
CodeBind-IB	Dense	✓	✓	97.3	94.0
		✓	✗	96.0	92.8

Table 5: **Multimodal fusion for event classification on AVE dataset (Tian et al., 2018)** Event category classification accuracy is reported for shared embeddings alone versus fusion with specific embeddings (concatenation/summation). The impact of text density during the multimodal alignment training stage is further investigated by comparing concise event categories with VLM-generated dense descriptions as text embeddings.

Multimodal fusion To directly examine the utilization of complementary cross-modal cues, we conduct multimodal fusion experiments on the AVE dataset (Tian et al., 2018) for audio-visual event classification. We compare CodeBind-IB against ImageBind using straightforward fusion strategies, namely concatenation and summation. To further explore the impact of language density, we leverage Qwen3-VL-4B (Bai et al., 2025a) to generate dense, descriptive captions for video frames, moving beyond coarse category labels. More details are in App. C.2. As shown in Tab. 5, integrating the video and audio specific embeddings alongside the shared ones consistently outperforms the use of shared embeddings in isolation. This gain is further amplified when employing VLM-generated dense descriptions as text embeddings during alignment, achieving an accuracy of 97.3%, a significant margin over ImageBind’s 94.4%. This empirically validates that our specific embeddings successfully capture and supply non-redundant cross-modal cues that directly benefit multimodal fusion tasks.

4.4 Ablation Study

Our ablation study examines: (1) the effects of the codebook modules, the decoupling strategy, and the reconstruction modules; (2) codebook settings, including the advantages of compositional VQ over standard VQ, codebook size, and dimension; (3) the contribution of various loss terms.

Codebook, decoupling and reconstruction Table 6 shows that incorporating codebook design improves modality alignment, yielding gains of +3.6%/+11.8%/+33.9% on NYU-D, SUN-D and FLIR_v2 over ImageBind. This confirms the

Components			Datasets		
Codebook	Decoupled	Reconstruction	NYU-D	SUN-D	FLIR_v2
✗	✗	✗	54.0	35.1	46.6
✗	✓	✓	54.1	39.7	94.5
✓	✗	✗	<u>57.6</u>	46.9	80.5
✓	✓	✗	56.7	45.3	97.7
✓	✓	✓	59.3	<u>45.7</u>	<u>97.2</u>

Table 6: **Effectiveness of each component.** Our codebook enhances performance; decoupling and reconstruction provide extra improvement.

Codebook Config.		Datasets		
Shared	Compositional	NYU-D	SUN-D	FLIR_v2
✓	✗	48.5	40.7	81.1
✗	✓	<u>56.1</u>	<u>45.7</u>	<u>90.3</u>
✓	✓	59.3	45.7	97.2

Table 7: **Different codebook settings** with shared/separated codebooks and compositional/conventional VQ.

shared codebook’s efficacy in facilitating robust alignment in a unified space. Furthermore, representation decoupling into shared and specific streams delivers substantial performance leaps (up to +48.0%), validating the necessity of isolating the semantic core from modality-specific nuances. Notably, the codebook and decoupling modules alone, before adding the reconstruction module, already deliver substantial gains.

Loss functions Tab. 8 shows training objectives building upon one another. Our regularization losses significantly enhance representation disentanglement and informativeness while achieving superior alignment performance (details in App. C.4).

Codebook design and size In Tab. 7, the shared codebook approach outperforms separate codebooks, showing improvements of +3.2%/+6.9% on NYU-D/FLIR_v2. Compared to standard VQ, our compositional VQ significantly improves alignment by +10.8%, +5.7%, and 16.1% on three datasets. The main advantage is its expanded capacity, which can accommodate more information without the need for codebook expansion.

Fig. 9 shows that codebook size controls the

\mathcal{L}_{align}	\mathcal{L}_{cm}	$\mathcal{L}_{cctr}, \mathcal{L}_{cuni}$	$\mathcal{L}_{orth}, \mathcal{L}_{uni}$	\mathcal{L}_{recon}	NYU-D	SUN-D	FLIR_v2
✓	✗	✗	✗	✗	54.0	35.1	46.6
✓	✓	✗	✗	✗	56.4	46.0	98.7
✓	✓	✓	✗	✗	56.7	45.3	<u>97.7</u>
✓	✓	✓	✓	✗	<u>57.6</u>	46.9	80.6
✓	✓	✓	✓	✓	59.3	45.7	97.2

Table 8: **Ablation on loss functions building upon one another.**

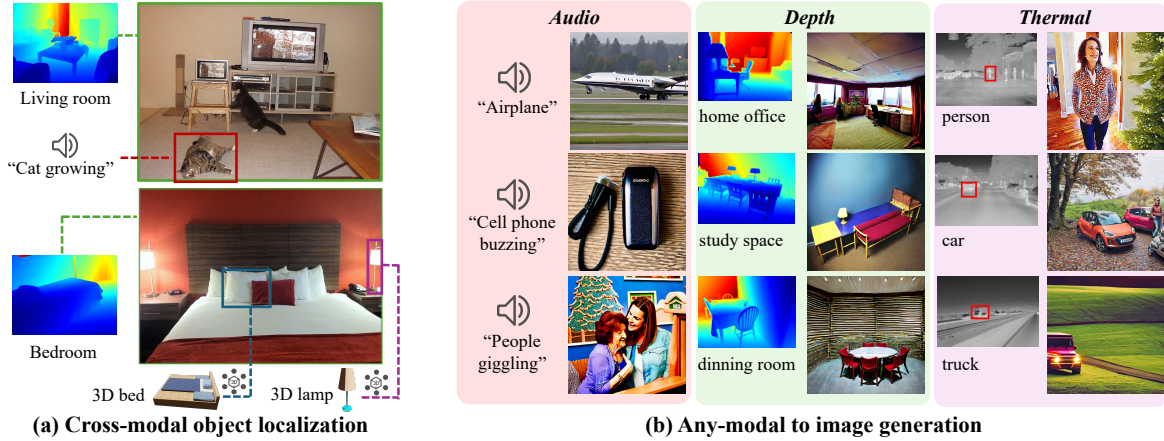


Figure 8: **Several applications utilizing alignment space of CodeBind.** (A) Semantic aligned embeddings from depth, audio, and 3D point cloud are retrieved using the image bounding box region as vision proposals. (B) By replacing the image encoder with any modality encoder in Stable unCLIP, we achieve any-modal-to-image generation.

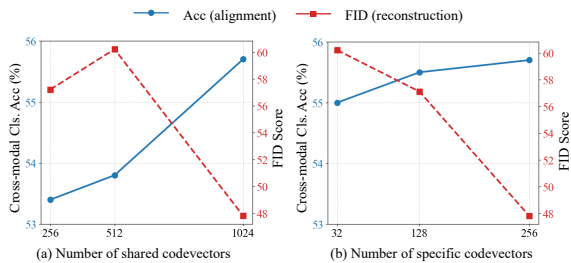


Figure 9: **Impact of codebook size on alignment and reconstruction.** (a) Varying shared codebook size when fixing specific codebook size at 256 (b) Varying specific codebook size when fixing shared codebook size at 1024. Cross-modal classification accuracy and FID of reconstructed images are compared using NYU-D (Silberman et al., 2012) with codevector dim=128.

trade-off between cross-modal alignment efficacy and intra-modal reconstruction fidelity. Expanding the shared codebook boosts cross-modal alignment. The specific codebook size primarily modulates reconstruction quality without notable impacts on alignment. These results confirm our codebooks’ role as information moderators: the shared one extracts cross-modal invariants, while the specific ones retain modal-unique features. Furthermore, while the reconstruction module adds training overhead, it is discarded during inference without compromising alignment (Tab. 6). Its inclusion remains vital, however, to ensure the representational integrity of modality-specific information. (See App. C.3 for more analysis.)

4.5 More Applications

CodeBind enables zero-shot cross-modal object localization and any-modal-to-image generation by seamlessly integrating diverse modalities into estab-

lished vision-language and generative frameworks (e.g., GroundingDINO (Liu et al., 2024) and Stable unCLIP (Rombach et al., 2022)) without additional training.

5 Conclusion

This paper proposes CodeBind, a novel multi-modal alignment framework that effectively aligns diverse modalities without requiring large-scale fully-paired datasets, thereby delivering enhanced practicality and scalability. Through decoupled representations and a shared-specific codebook mechanism, CodeBind significantly enhances cross-modal compatibility while preserving distinct modality characteristics. Extensive experiments consistently demonstrate the superiority of CodeBind in cross-modal classification and retrieval tasks. It optimizes the representation space by reducing modality gaps while preserving informative modality-unique details, revealing its great potential in various downstream tasks.

Limitations

In our experiments, while modality-specific information in visual embeddings is validated via text-anchored interpretability provided by VLM, extending this to uncommon modalities (which lack strong foundational spaces) is currently challenging. Nevertheless, our decoupled architecture provides the structural prerequisite for future explorations of the modality-specific information in complex multimodal reasoning tasks like sentiment analysis. In addition, though we primarily utilize category names rather than detailed descriptions

during alignment for fair comparison, incorporating LLM-generated dense semantic descriptions, as suggested by recent findings (Fan et al., 2024), could further unlock the potential of our decoupled space. Finally, our plug-and-play design enables integration into MLLMs for “on-demand” fusion. Future work can employ gated mechanisms to dynamically switch between cross-modal consistent concepts and modality-unique cues, adapting to varying scenario requirements. Moving forward, our shared-specific decoupling fosters safer and more interpretable multimodal systems. In robotics, isolating modality-specific signals shields the shared semantic core from sensor interference. In medical diagnostics, this transparency enables clinicians to trace decisions to specific modality cues, ensuring clinical accountability.

Acknowledgments

This work is supported by the Hong Kong Research Grants Council - General Research Fund (Grant No.: 17211024) and HKU Seed Fund for PI Research.

References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, and 1 others. 2022. Flamingo: a visual language model for few-shot learning. In *NeurIPS*.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, and 1 others. 2025a. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025b. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. 2020. Vggsound: A large-scale audio-visual dataset. In *ICASSP*.
- Hongzhou Chen, Lianghua He, Yihang Liu, and Longzhen Yang. 2024. Visual neural decoding via improved visual-eeeg semantic consistency. *arXiv preprint arXiv:2408.06788*.
- Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. 2023. Reproducible scaling laws for contrastive language-image learning. In *CVPR*.
- Amaya Dharmasiri, Muzammal Naseer, Salman Khan, and Fahad Shahbaz Khan. 2024. Cross-modal self-training: Aligning images and pointclouds to learn classification without labels. In *CVPR*.
- Konstantinos Drossos, Samuel Lipping, and Tuomas Virtanen. 2020. Clotho: An audio captioning dataset. In *ICASSP*.
- Jiali Duan, Liqun Chen, Son Tran, Jinyu Yang, Yi Xu, Belinda Zeng, and Trishul Chilimbi. 2022. Multimodal alignment using representation codebook. In *CVPR*.
- Patrick Esser, Robin Rombach, and Bjorn Ommer. 2021. Taming transformers for high-resolution image synthesis. In *CVPR*.
- Zhiyuan Fan, Zhihong Chen, and Benyou Wang. 2024. Exploring the potential of dense information in multimodal alignment. In *ACL*.
- Bilal Faye, Hanane Azzag, and Mustapha Lebbah. 2024. Oneencoder: A lightweight framework for progressive alignment of modalities. *arXiv preprint arXiv:2409.11059*.
- Teledyne FLIR. 2018. Teledyne flir adas thermal dataset v2. <https://www.kaggle.com/datasets/samdazel/teledyne-flir-adas-thermal-dataset-v2/>.
- Letian Fu, Gaurav Datta, Huang Huang, William Chung-Ho Panitch, Jaimyn Drake, Joseph Ortiz, Mustafa Mukadam, Mike Lambeta, Roberto Calandra, and Ken Goldberg. 2024. A touch, vision, and language dataset for multimodal alignment. In *ICML*.
- Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. 2017. Audio set: An ontology and human-labeled dataset for audio events. In *ICASSP*.
- Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. 2023. Imagebind: One embedding space to bind them all. In *CVPR*.
- Fabian Gröger, Shuo Wen, Huyen Le, and Maria Brbić. 2025. With limited data for multimodal alignment, let the structure guide you. In *NeurIPS*.
- Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel. 2022. Audioclip: Extending clip to image, text and audio. In *ICASSP*.
- Jiaming Han, Kaixiong Gong, Yiyuan Zhang, Jiaqi Wang, Kaipeng Zhang, Dahua Lin, Yu Qiao, Peng Gao, and Xiangyu Yue. 2024. Onellm: One framework to align all modalities with language. In *CVPR*.

- Jiaming Han, Renrui Zhang, Wenqi Shao, Peng Gao, Peng Xu, Han Xiao, Kaipeng Zhang, Chris Liu, Song Wen, Ziyu Guo, and 1 others. 2023. Imagebind-llm: Multi-modality instruction tuning. *arXiv preprint arXiv:2309.03905*.
- Qingdong He, Jinlong Peng, Zhengkai Jiang, Kai Wu, Xiaozhong Ji, Jiangning Zhang, Yabiao Wang, Chengjie Wang, Mingang Chen, and Yunsheng Wu. 2024. Unim-ov3d: Uni-modality open-vocabulary 3d scene understanding with fine-grained feature representation. In *IJCAI*.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *ICLR*.
- Hai Huang, Yan Xia, Shengpeng Ji, Shulei Wang, Hanting Wang, Minghui Fang, Jieming Zhu, Zhenhua Dong, Sashuai Zhou, and Zhou Zhao. 2025. Enhancing multimodal unified representations for cross modal generalization. In *ACL(Findings)*.
- Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hannaneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. 2021. [Openclip](#). If you use this software, please cite it as below.
- Xinyu Jia, Chuang Zhu, Minzhen Li, Wenqi Tang, and Wenli Zhou. 2021. Llvip: A visible-infrared paired dataset for low-light vision. In *ICCV*.
- Qian Jiang, Changyou Chen, Han Zhao, Liqun Chen, Qing Ping, Son Dinh Tran, Yi Xu, Belinda Zeng, and Trishul Chilimbi. 2023. Understanding and constructing latent modality structures in multi-modal representation learning. In *CVPR*.
- Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, and 1 others. 2017. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*.
- Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Li Fei-Fei. 2011. Novel dataset for fine-grained image categorization. In *First Workshop on Fine-Grained Visual Categorization, CVPR*.
- Chris Dongjoo Kim, Byeongchang Kim, Hyunmin Lee, and Gunhee Kim. 2019. Audiocaps: Generating captions for audios in the wild. In *NAACL*.
- Weixian Lei, Yixiao Ge, Kun Yi, Jianfeng Zhang, Difei Gao, Dylan Sun, Yuying Ge, Ying Shan, and Mike Zheng Shou. 2024. Vit-lens: Towards omnimodal representations. In *CVPR*.
- Chenxin Li, Yifan Liu, Panwang Pan, Hengyu Liu, Xinyu Liu, Wuyang Li, Cheng Wang, Weihao Yu, Yiyang Lin, and Yixuan Yuan. 2025. Infobridge: Balanced multimodal integration through conditional dependency modeling. In *ICCV*.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022a. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*.
- Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, Hua Wu, and Haifeng Wang. 2022b. Unimo-2: End-to-end unified vision-language grounded learning. *arXiv preprint arXiv:2203.09067*.
- Victor Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Y Zou. 2022. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. In *NeurIPS*.
- Han Lin, Xiu Tang, Huan Li, Wenxue Cao, Sai Wu, Chang Yao, Lidan Shou, and Gang Chen. 2025. T2dr: A two-tier deficiency-resistant framework for incomplete multimodal learning. In *ACL(Findings)*.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *ECCV*.
- Alexander H Liu, SouYoung Jin, Cheng-I Jeff Lai, Andrew Rouditchenko, Aude Oliva, and James Glass. 2022. Cross-modal discrete representation learning. In *ACL*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. In *NeurIPS*.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, and 1 others. 2024. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *ECCV*.
- Ye Liu, Zihan Ji, and Hongmin Cai. 2025. Plug-and-play feature causality decomposition for multimodal representation learning. In *ICCV*.
- Yuanhuiyi Lyu, Xu Zheng, Dahun Kim, and Lin Wang. 2024a. Omnibind: Teach to build unequal-scale modality interaction for omni-bind of all. *arXiv preprint arXiv:2405.16108*.
- Yuanhuiyi Lyu, Xu Zheng, Jiazhou Zhou, and Lin Wang. 2024b. Unibind: Llm-augmented unified and balanced representation space to bind them all. In *CVPR*.
- Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. 2012. Cats and dogs. In *CVPR*.
- Karol J Piczak. 2015. Esc: Dataset for environmental sound classification. In *ACM MM*, pages 1015–1018.

- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Sameera Ramasinghe, Violetta Shevchenko, Gil Avraham, and Ajanthan Thalaiyasingam. 2024. Accept the modality gap: An exploration in the hyperbolic space. In *CVPR*.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, and 1 others. 2015. Imagenet large scale visual recognition challenge. *IJCV*.
- Kyle Sargent, Jing Yu Koh, Han Zhang, Huiwen Chang, Charles Herrmann, Pratul Srinivasan, Jiajun Wu, and Deqing Sun. 2023. Vq3d: Learning a 3d-aware generative model on imagenet. In *ICCV*.
- Peiyang Shi, Michael C. Welle, Mårten Björkman, and Danica Kragic. 2023. Understanding the modality gap in CLIP. In *ICLR 2023 Workshop on Multimodal Representation Learning: Perks and Pitfalls*.
- Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. 2012. Indoor segmentation and support inference from rgb-d images. In *ECCV*.
- Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. 2015. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *CVPR*.
- Concetto Spampinato, Simone Palazzo, Isaak Kavasidis, Daniela Giordano, Nasim Souly, and Mubarak Shah. 2017. Deep learning human mind for automated visual classification. In *CVPR*.
- Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. 2023. Pandagpt: One model to instruction-follow them all. In *Proceedings of the 1st Workshop on Taming Large Language Models: Controllability in the era of Interactive Assistants!*
- Yapeng Tian, Jing Shi, Bochen Li, Zhiyao Duan, and Chenliang Xu. 2018. Audio-visual event localization in unconstrained videos. In *ECCV*.
- Megan Tjandrasuwita, Chanakya Ekbote, Liu Ziyin, and Paul Pu Liang. 2025. Understanding the emergence of multimodal representation alignment. In *ICML*.
- Aaron Van Den Oord, Oriol Vinyals, and 1 others. 2017a. Neural discrete representation learning. In *NeurIPS*.
- Aaron Van Den Oord, Oriol Vinyals, and 1 others. 2017b. Neural discrete representation learning. In *NeurIPS*.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*.
- Zehan Wang, Ziang Zhang, Xize Cheng, Rongjie Huang, Luping Liu, Zhenhui Ye, Haifeng Huang, Yang Zhao, Tao Jin, Peng Gao, and 1 others. 2024. Freebind: Free lunch in unified multimodal space via knowledge fusion. In *ICML*.
- Zehan Wang, Ziang Zhang, Luping Liu, Yang Zhao, Haifeng Huang, Tao Jin, and Zhou Zhao. 2025a. Extending multi-modal contrastive representations. In *NeurIPS*.
- Zehan Wang, Ziang Zhang, Hang Zhang, Luping Liu, Rongjie Huang, Xize Cheng, Hengshuang Zhao, and Zhou Zhao. 2025b. Omnibind: Large-scale omnimodal representation via binding spaces. In *ICLR*.
- Zehan Wang, Yang Zhao, Xize Cheng, Haifeng Huang, Jiageng Liu, Aoxiong Yin, Li Tang, Linjun Li, Yongqi Wang, Ziang Zhang, and Zhou Zhao. 2023. Connecting multi-modal contrastive representations. In *NeurIPS*.
- Ho-Hsiang Wu, Prem Seetharaman, Kundan Kumar, and Juan Pablo Bello. 2022. Wav2clip: Learning robust audio representations from clip. In *ICASSP*.
- Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 2015a. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*.
- Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 2015b. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*.
- Yan Xia, Hai Huang, Jieming Zhu, and Zhou Zhao. 2024. Achieving cross modal generalization with multimodal unified representation. In *NeurIPS*.
- Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*.
- Fengyu Yang, Chao Feng, Daniel Wang, Tianye Wang, Ziyao Zeng, Zhiyang Xu, Hyungseob Park, Pengliang Ji, Hanbin Zhao, Yuanning Li, and 1 others. 2024. Neurobind: Towards unified multimodal representations for neural signals. *arXiv preprint arXiv:2407.14020*.
- Fengyu Yang, Chenyang Ma, Jiacheng Zhang, Jing Zhu, Wenzhen Yuan, and Andrew Owens. 2022. Touch and go: Learning from human-collected vision and touch. In *NeurIPS*.
- Jiuding Yang, Yakun Yu, Di Niu, Weidong Guo, and Yu Xu. 2023. Confede: Contrastive feature decomposition for multimodal sentiment analysis. In *ACL*.

- Delong Zeng, Yuexiang Xie, Yaliang Li, and Ying Shen. 2025. Enhancing multimodal retrieval via complementary information extraction and alignment. In *ACL*.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. Sigmoid loss for language image pre-training. In *ICCV*.
- Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. 2022. Pointclip: Point cloud understanding by clip. In *CVPR*.
- Yiyuan Zhang, Kaixiong Gong, Kaipeng Zhang, Hongsheng Li, Yu Qiao, Wanli Ouyang, and Xiangyu Yue. 2023. Meta-transformer: A unified framework for multimodal learning. *arXiv preprint arXiv:2307.10802*.
- Youcai Zhang, Xinyu Huang, Jinyu Ma, Zhaoyang Li, Zhaochuan Luo, Yanchun Xie, Yuzhuo Qin, Tong Luo, Yaqian Li, Shilong Liu, and 1 others. 2024a. Recognize anything: A strong image tagging model. In *CVPR*.
- Yuhui Zhang, Elaine Sui, and Serena Yeung-Levy. 2024b. Connect, collapse, corrupt: Learning cross-modal tasks with uni-modal data. In *ICLR*.
- Chuanxia Zheng and Andrea Vedaldi. 2023. Online clustered codebook. In *ICCV*.
- Chuanxia Zheng, Tung-Long Vuong, Jianfei Cai, and Dinh Phung. 2022. Movq: Modulating quantized vectors for high-fidelity image generation. In *NeurIPS*.
- Sipeng Zheng, Bohan Zhou, Yicheng Feng, Ye Wang, and Zongqing Lu. 2024. Unicode: Learning a unified codebook for multimodal large language models. In *ECCV*.
- Bo Zhou, Liulei Li, Yujia Wang, Huafeng Liu, Yazhou Yao, and Wenguan Wang. 2025. Unialign: Scaling multimodal alignment within one unified model. In *CVPR*.
- Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. 2014. Learning deep features for scene recognition using places database. In *NeurIPS*.
- Bin Zhu, Bin Lin, Munan Ning, Yang Yan, Jiayi Cui, WANG HongFa, Yatian Pang, Wenhao Jiang, Junwu Zhang, Zongwei Li, Cai Wan Zhang, Zhifeng Li, Wei Liu, and Li Yuan. 2024a. Languagebind: Extending video-language pretraining to n-modality by language-based semantic alignment. In *ICLR*.
- Lei Zhu, Fangyun Wei, Yanye Lu, and Dong Chen. 2024b. Scaling the codebook size of VQ-GAN to 100,000 with a utilization rate of 99%. In *NeurIPS*.

A More Details of CodeBind framework

A.1 Loss Function

In Sec. 3.4 of the main paper, we present an overview of our training objective, which incorporates multiple loss functions to optimize multimodal representation for alignment. This section provides details of each loss function and its functionality when applied to decoupled embeddings and the modality-shared-specific codebook.

InfoNCE loss The InfoNCE loss aligns shared embeddings between target modalities \mathcal{T} and bridging modalities \mathcal{A} . The target modality \mathcal{T} is characterized by data $X^{\mathcal{T}} = \{x_1^{\mathcal{T}}, \dots, x_N^{\mathcal{T}}\}$ with shared embeddings $\{z_{\text{shared},1}^{\mathcal{T}}, \dots, z_{\text{shared},N}^{\mathcal{T}}\}$, whereas the bridging modality \mathcal{A} is defined by data $X^{\mathcal{A}} = \{x_1^{\mathcal{A}}, \dots, x_N^{\mathcal{A}}\}$ with shared embeddings $\{z_{\text{shared},1}^{\mathcal{A}}, \dots, z_{\text{shared},N}^{\mathcal{A}}\}$. $\hat{z}_{\text{shared},i}^{\mathcal{T}} = Q(E(x_i^{\mathcal{T}}))$ and $\hat{z}_{\text{shared},i}^{\mathcal{A}} = Q(E(x_i^{\mathcal{A}}))$ refer to the quantized shared embeddings of data $x_i^{\mathcal{T}}$ and $x_i^{\mathcal{A}}$. The InfoNCE loss is defined as,

$$\mathcal{L}_{X^{\mathcal{A}}, X^{\mathcal{T}}} = -\log \frac{\exp(\hat{z}_{\text{shared},i}^{\mathcal{T}} \cdot \hat{z}_{\text{shared},i}^{\mathcal{A}}/\eta)}{\sum_j \exp(\hat{z}_{\text{shared},i}^{\mathcal{T}} \cdot \hat{z}_{\text{shared},j}^{\mathcal{A}}/\eta)} \quad (1)$$

where η is a scalar temperature parameter, and j refers to negative observations. Following Girdhar et al., 2023, we consider each example $j \neq i$ in a batch as a negative instance. This loss function aims to bring embeddings $\hat{z}_{\text{shared},i}^{\mathcal{T}}$ and $\hat{z}_{\text{shared},i}^{\mathcal{A}}$ closer in the unified space, promoting multimodal alignment. Practically, we apply the alignment to shared embeddings of [CLS] token from encoded data, and employ a symmetric loss to enhance the alignment: $\mathcal{L}_{\text{align}} = \mathcal{L}_{X^{\mathcal{A}}, X^{\mathcal{T}}} + \mathcal{L}_{X^{\mathcal{T}}, X^{\mathcal{A}}}$. When multiple bridging modalities are involved, we calculate the InfoNCE loss separately for each bridging modality.

Reconstruction loss The shared embedding and the specific embedding are concatenated in the modality decoder to reconstruct the original input for each modality. The reconstruction retains the embedding information, including both the shared information across modalities and the unique information within each modality. Both [CLS] and patch token embeddings are quantized and input into the reconstruction module. Specifically, the vector quantized concatenated embeddings of the input data $x_i^{\mathcal{M}}$, $\text{concat}(\hat{z}_{\text{shared},i}^{\mathcal{M}}, \hat{z}_{\text{spec},i}^{\mathcal{M}})$ are passed through a decoder network $D(\cdot)$ to reconstruct $\hat{x}_i^{\mathcal{M}}$. Reconstruction is applied both to the bridging

modalities and target modalities. Note that in this paper, the text modality only has shared embedding without additional specific embedding and is therefore not reconstructed. The reconstruction loss is defined as,

$$\mathcal{L}_{\text{recon}} = \|x_i^{\mathcal{M}} - \hat{x}_i^{\mathcal{M}}\|^2, \hat{x}_i^{\mathcal{M}} = D(Q(E(x_i^{\mathcal{M}}))) \quad (2)$$

Orthogonal loss We use orthogonal loss $\mathcal{L}_{\text{orth}}$ to encourage the shared and specific embeddings of the data $x_i^{\mathcal{M}}$ from any modality $\mathcal{M} \in \{\mathcal{A}, \mathcal{T}\}$ to capture disentangled information. Given the quantized embeddings, $\hat{z}_{\text{shared},i}^{\mathcal{M}}$ and $\hat{z}_{\text{spec},i}^{\mathcal{M}}$, we minimize their inner product, encouraging them to contain complementary information. $\mathcal{L}_{\text{orth}}$ is further averaged among bridging modalities and target modality. The orthogonal loss is defined as,

$$\mathcal{L}_{\text{orth}} = \frac{1}{N} \sum_i \langle \hat{z}_{\text{shared},i}^{\mathcal{M}}, \hat{z}_{\text{spec},i}^{\mathcal{M}} \rangle^2 \quad (3)$$

Uniform loss Uniform loss is used \mathcal{L}_{uni} to ensure that specific embeddings are evenly distributed in the feature space, enabling them to differentiate between samples effectively. This loss complements the reconstruction loss by encouraging specific embeddings to preserve essential information. Following Jiang et al., 2023, the uniform loss is defined as the logarithm of the averaged Gaussian potential among all pairwise combinations of quantized specific embeddings ($\hat{z}_{\text{spec},i}^{\mathcal{M}}, \hat{z}_{\text{spec},j}^{\mathcal{M}}$) within a modality \mathcal{M} , and is further averaged across all modalities. We apply uniform loss only to the specific embeddings of [CLS] tokens encoded from pairwise data for computational efficiency. Also, the loss is averaged among bridging modalities and target modality. The uniform loss is defined as,

$$\mathcal{L}_{\text{uni}} = \log \frac{1}{N} \sum_{i,j} \exp(-\|\hat{z}_{\text{spec},i}^{\mathcal{M}} - \hat{z}_{\text{spec},j}^{\mathcal{M}}\|^2) \quad (4)$$

Codevector commitment loss Commitment loss is used during codebook learning for the encoder embeddings $z_i^{\mathcal{M}} = E(x_i^{\mathcal{M}})$ to be similar to its quantized embeddings $\hat{z}_i^{\mathcal{M}} = Q(z_i^{\mathcal{M}})$. The commitment loss is,

$$\mathcal{L}_{\text{vq}} = \beta \|z_i^{\mathcal{M}} - \text{sg}[\hat{z}_i^{\mathcal{M}}]\|^2 \quad (5)$$

where $\text{sg}(\cdot)$ denotes the stop-gradient operation, β is the weighting factor.

Codevector cross-modal code matching (CMCM) loss

We also adopt CMCM loss (Liu et al., 2022) to further enhance the codevectors alignment across paired modalities. Given paired data from bridging and target modalities, (x_i^A, x_i^T) , the calculation is based on the probability distribution of codevector usage within the shared embeddings of [CLS] tokens of relevant paired data. Here, the probability distribution of a codevector e_k given a sub-vector $z_{i,h}^M \in [z_{i,1}^M, \dots, z_{i,m}^M]$ segmented by our compositional codebook approach from the original embedding z_i^M , is defined as the softmax function over their Euclidean distance, i.e., $P(e_k|z_{i,h}^M) = \frac{\exp(-\|z_{i,h}^M - e_k\|^2)}{\sum_l \exp(-\|z_{i,h}^M - e_l\|^2)}$. This loss ensures that the codevector distributions of data from the bridging modality and target modality in positive pairs are similar, while distributions from negative pairs are distinct. Positive pairs are defined as $(z_{i,h}^A, z_{i,h}^T)$ and negative pairs as $(z_{i,h}^A, z_{j,h}^T)$, with the sub-vectors corresponding to the same segmented position of the original embeddings z_i^A and z_i^T . The codevector CMCM loss is defined as,

$$\mathcal{L}_{\text{cm}} = -\frac{1}{N+m} \sum_{i,h}^{N,m} \log \frac{\exp S_{\text{code}}(z_{i,h}^A, z_{i,h}^T)}{\sum_j \exp S_{\text{code}}(z_{i,h}^A, z_{j,h}^T)} \quad (6)$$

$$S_{\text{code}}(z_{i,h}^A, z_{i,h}^T) = \sum_k P(e_k|z_{i,h}^A) \log P(e_k|z_{i,h}^T) + \sum_k P(e_k|z_{i,h}^T) \log P(e_k|z_{i,h}^A)$$

Codevector regularization loss The codevector regularization loss, comprising contrastive (\mathcal{L}_{ctr}) and uniform ($\mathcal{L}_{\text{cuni}}$) components, ensures codevectors remain semantically distinguishable. Specifically, the codevector contrastive loss \mathcal{L}_{ctr} fosters a non-uniform distance distribution by widening the gap between sub-vectors and their corresponding codebook entries. Given a sub-vector $z_{i,h}$, we define positive distance as the mean distance to the top 10% closest codevectors (\mathcal{C}_{pos}) and negative distance as the distance to the furthest 50% (\mathcal{C}_{neg}). By averaging this contrastive objective across all partitioned sub-vectors and modalities, codevectors are encouraged to encode diverse, meaningful semantics. This loss is formulated as follows:

$$\mathcal{L}_{\text{ctr}} = \frac{\frac{1}{|\mathcal{C}_{\text{pos}}|} \sum_{e_{k_1} \in \mathcal{C}_{\text{pos}}} d(e_{k_1}, z_{i,h})}{\sum_{e_{k_2} \in \mathcal{C}_{\text{neg}}} d(e_{k_2}, z_{i,h})} \quad (7)$$

The codevector uniform loss, $\mathcal{L}_{\text{cuni}}$, mirrors the \mathcal{L}_{uni} term used for specific embeddings, aiming to promote balanced codevector utilization and prevent individual entries from being overused or underrepresented. This loss is applied independently to both shared and specific codebooks. By encouraging an even distribution of codevector usage, $\mathcal{L}_{\text{cuni}}$ maintains semantic diversity and mitigates redundancy within the learned representation space. This loss is formulated as:

$$\mathcal{L}_{\text{cuni}} = \log \frac{1}{|\mathcal{C}|} \sum_{e_m, e_n \in \mathcal{C}} \exp(-\|e_m - e_n\|^2) \quad (8)$$

A.2 Codebook Update Policy

The shared codebook is initialized by K-Means on text sub-vectors and updated via batched shared embeddings from bridging and target modalities. Specific codebooks undergo a similar K-Means initialization and refinement process using their respective target embeddings. For numerical stability and faster convergence, all sub-vectors and codevectors are normalized to unit magnitude to ensure consistent scaling.

Update by EMA We update the codebook using Exponential Moving Average (EMA) (Van Den Oord et al., 2017b). For a codevector $e_k \in \mathcal{C}$ and its associated feature cluster $Z = \{z_i | \min_v \text{dist}(z_i, e_v) = e_k\}$, the update process at training step t is as follows with a decay factor $\gamma = 0.99$:

$$\begin{aligned} N_k^{(t)} &= \gamma N_k^{(t-1)} + (1-\gamma)|Z| \\ f_k^{(t)} &= \gamma f_k^{(t-1)} + (1-\gamma) \sum_{z \in Z} z \\ e_k^{(t)} &= \frac{f_k^{(t)}}{N_k^{(t)}} \end{aligned} \quad (9)$$

$N_k^0 = 0$, and f_k^0 are initialized per cluster as described previously. This approach ensures that the codebook evolves adaptively with the data distribution during training.

Reinitialization We follow the approach in Zheng and Vedaldi, 2023 to adjust less-used or unused codevectors when learning the codebook. We start by accumulatively counting the average usage of codevectors in each training mini-batch. At each training step t , a decay value $\alpha_k^{(t)}$ is calculated for each codevector entry $e_k^{(t)}$ based on the

accumulated average usage. Then, feature anchors are randomly selected from the encoded features in a mini-batch. Unused or low-used codevectors are then reinitialized as,

$$e_k^{(t)} = e_k^{(t-1)} \cdot (1 - \alpha_k^{(t)}) + \hat{z}_k^{(t)} \cdot \alpha_k^{(t)} \quad (10)$$

where $\hat{z}_k^{(t)}$ is the sampled feature anchor. Notably, this is specifically applied to reinitialize unused or low-used codevectors, rather than updating the active ones.

A.3 Architecture

We utilize a two-tower Transformer to align bridging modalities (text, vision) with target modalities (audio, depth, thermal, etc.) pair-wise. Modality embeddings are decoupled into shared and specific components via trainable projection heads and quantized using respective codebooks. Cross-modal alignment relies on shared embeddings, while the full concatenated representation is used for reconstruction. Notably, for text embeddings, we do not perform this decomposition and assume that they only contain shared components. This is because text modalities generally capture high-level, coarse semantic meanings, and the shared components are sufficient to represent such information in a compressed form.

Modality encoder Our main experiments use ImageBind (Girdhar et al., 2023) and ViT-Lens (Lei et al., 2024) as baselines, employing a Transformer architecture for all modality encoders. Following the configurations used in ImageBind and ViTLens, in CodeBind-IB and CodeBind-VL, the encoders of bridging modalities are extracted and frozen using OpenCLIP ViT-H and ViT-B/16 respectively (Cherti et al., 2023; Ilharco et al., 2021). For other modalities, the checkpoints of trained modality-specific encoders in ImageBind and ViT-Lens are resumed to develop our CodeBind-IB and CodeBind-VL. To improve the alignment of each target modality and facilitate the learning of our integrated modality-shared-specific codebooks, the encoders of target modalities are finetuned using LoRA (Hu et al., 2022) with rank 4 on the final 6 layers of the Transformer structure, along with trainable projection heads.

Reconstruction decoder The reconstruction is used to impose constraints on the specific embeddings to preserve the comprehensive information

of the initial data, in conjunction with the shared embeddings. The reconstruction decoder consists of a ViT structure with 8 transformer layers. They are applied to each trained modality, except for text, which does not have specific embeddings.

B Implementation Details

B.1 Datasets

Audio datasets: **AudioSet** (Gemmeke et al., 2017) is used for both training and evaluation, including 10-second videos sourced from YouTube that have been annotated into 527 classes. We do not use the unbalanced training split with 2M clips. Instead, we employ the balanced training split, which includes about 20K videos. And we use the test split of around 18K videos for evaluation. The prepared data from Lei et al., 2024 is used instead of processing it from scratch. The number of datasets utilized is slightly lower than the original due to some unavailable data. Textual descriptions for each class are generated using the class names and predefined templates from CLIP (Radford et al., 2021). **VGGSound(VGGS)** (Chen et al., 2020) consists of about 200K video clips, with about 15K in the test split and others in the training split. These clips are 10 seconds in length and are labeled with 309 sound classes, including human actions, sound-emitting objects, and human-object interactions. **AudioCaps** (Kim et al., 2019) includes about 46K audio clips to human-written text pairs collected via crowdsourcing on the AudioSet dataset. We use it for both training and evaluation. The test split is from (Lei et al., 2024), which has 813 clips for evaluation. **ESC** (Piczak, 2015) is for the Environmental Sound Classification (ESC) task, which consists of 2K 5s audio clips classified into 50 classes. It is only used for evaluation, employing a predefined 5-fold evaluation, each with 400 test audio clips. We compute predictions on the evaluation set for each fold and report the 5-fold average performance.

Depth datasets: **SUN-D** (Song et al., 2015) and **NYU-D** (Silberman et al., 2012) are RGB-D datasets used for scene classification, which contain registered RGB and Depth maps. SUN-D is used for both training and testing, while NYU-D is used only for testing. For depth map preprocessing, we follow ImageBind and ViT-Lens, utilizing in-filled depth values and converting them to disparity for scale normalization. In our experiments, we also follow ImageBind and ViT-Lens to map scene

class names to 10 categories and employ the same evaluation method.

Thermal datasets: **LLVIP** (Jia et al., 2021) dataset consists of RGB image and thermal (infrared low-light) image pairs captured with fixed cameras observing outdoor street scenes. The train set comprises 12,025 RGB and thermal pairs, while the test set contains 3,463 pairs. We follow ImageBind to convert the original detection dataset into a binary classification dataset. This process involves cropping out pedestrians from bounding boxes and creating random bounding boxes to ensure a balanced distribution of samples for the categories ‘person’ and ‘background’, each containing about 8K samples. For evaluation, we use the following class names for the ‘person’ class - [‘person’, ‘man’, ‘woman’, ‘people’], and [‘street’, ‘road’, ‘car’, ‘light’, ‘tree’] for the ‘background’ class. **FLIR_v2** (FLIR, 2018) dataset consists of a total of 9,711 thermal and 9,233 RGB images for object detection, with a 90%/10% train/validation split. We utilize paired RGB-thermal images and crop images of each category based on bounding boxes to prepare them for a classification task. Categories with very few instances are removed, and the dataset is filtered to retain the following 10 categories: [‘car’, ‘person’, ‘sign’, ‘motor’, ‘truck’, ‘light’, ‘bike’, ‘hydrant’, ‘dog’, ‘other vehicle’].

Tactile datasets: The tactile data is sourced from **Touch-and-go** (Yang et al., 2022) dataset, which includes tactile data collected by human data collectors probing objects in natural environments using tactile sensors. Simultaneously, egocentric videos are recorded to obtain the corresponding images. The dataset includes annotations for 20 different material classes and offers labels for hard/soft (H/S) and rough/smooth (R/S) characteristics. This results in three versions of the tactile dataset: one for material classification (**TAG-M**), one for hard/soft classification (**TAG-H/S**), and one for rough/smooth classification (**TAG-R/S**). Following ViT-Lens, our model is trained using the TAG-M dataset and evaluated on all three datasets.

Electroencephalography (EEG) datasets: For the visual concept classification task, the **ImageNet-EEG (IN-EEG)** (Spampinato et al., 2017) dataset is utilized, consisting of EEG recordings obtained from six human subjects using a 128-channel human brain activity recording system. Each subject is exposed to 2,000

images from 40 categories sourced from the ImageNet (Russakovsky et al., 2015) dataset. With each category comprising 50 unique images, a total of 12,000 EEG sequences are recorded. We conducted experiments on a test set of 1,997 samples to compare with ImageBind. Additionally, we utilized both the validation set (1,998 samples) and the test set to compare with ViT-Lens.

3D Point cloud datasets: For the 3D classification task, we use **ShapeNet55** (Wu et al., 2015a) for training and **ModelNet40** (Wu et al., 2015b) for evaluation. Following ViT-Lens (Lei et al., 2024), we also employ the processed ShapeNet55 dataset. This dataset consists of approximately 52.5K 3D point clouds generated from CAD models, along with corresponding images created using virtual cameras and text data obtained by filling meta-data into a predefined template. The ModelNet40 dataset is a popular benchmark for 3D shape classification. This dataset contains 12,311 CAD models across 40 categories, with 9,843 training samples and 2,468 testing samples. Each object is represented as a 3D point cloud and manually annotated with its category, including everyday items like chairs, tables, desks, and household objects. We evaluate our model’s performance in 3D classification using only the test set in our experiments.

B.2 Training

CodeBind-IB and CodeBind-VL share similar training hyperparameters with ImageBind and ViT-Lens, respectively. Bridging modalities are paired with one target modality at a time to train a unique set of shared and specific codebooks. Each bridging–target modality pair uses a distinct set of codebooks (*e.g.*, image-text-audio vs. image-text-depth), and codebooks are not unified or fixed across different target modalities. Different batch sizes are used for different modalities to optimize GPU memory utilization. Additionally, different loss functions are assigned varying weights depending on the modalities to be aligned to maintain a balanced loss composition. Details of these hyperparameters and the number of trainable parameters for CodeBind-IB are in Tab 1. We further implement an adaptive loss balancing strategy to dynamically assign weights for \mathcal{L}_{vq} , \mathcal{L}_{ctr} , \mathcal{L}_{cuni} , \mathcal{L}_{cm} , \mathcal{L}_{orth} , \mathcal{L}_{uni} based on their relative magnitudes with respect to \mathcal{L}_{align} . By default, the weights for \mathcal{L}_{align} and \mathcal{L}_{recon} are fixed at 1. Table 2 presents an ablation study compar-








Modality	 Image	 Video	 Depth	 Audio	 Thermal	 Tactile	 EEG	
Batch Size	IN1K: 16 p365: 10	6	8	6	8	6	6	
Loss λ	\mathcal{L}_{vq}	10	10	1000	50	1000	1000	10
	$\mathcal{L}_{cctr}, \mathcal{L}_{cuni}$	0.01	0.01	0.01	0.01	0.05	0.01	0.01
	\mathcal{L}_{cm}	0.05	0.05	0.1	1.0	1.0	0.1	0.1
	\mathcal{L}_{orth}	0.1	0.1	0.1	0.1	1.0	0.05	0.05
	\mathcal{L}_{uni}	1.0	1.0	1.0	1.0	1.0	0.1	0.1
	\mathcal{L}_{recon}	1.0	1.0	1.0	1.0	1.0	1.0	1.0
LoRA (Layers)	–	–	7–12	7–12	7–12	7–12	4–6	
# Trainable Params.	32.4M	29.9M	59.6M	56.7M	61.6M	58.5M	34.1M	

Table 1: **Training Hyper-parameters across Modalities in CodeBind-IB.** This table details the batch size per GPU, loss function weights (λ), LoRA inserted modality encoder layers, and the number of trainable parameters for aligning bridging modality to each target modality. Modality encoders of depth, audio, thermal, and tactile have 12 transformer blocks in total, while modality encoder of EEG has 6 blocks in total. Modality encoder of image and video is frozen and LoRA is disabled.

	SUN-D	NYU-D	FLIR_v2
manual loss weights	45.7	59.3	97.2
adaptive loss weights	46.0	57.5	97.2

Table 2: **Comparison between manual and adaptive loss weights across different datasets.**

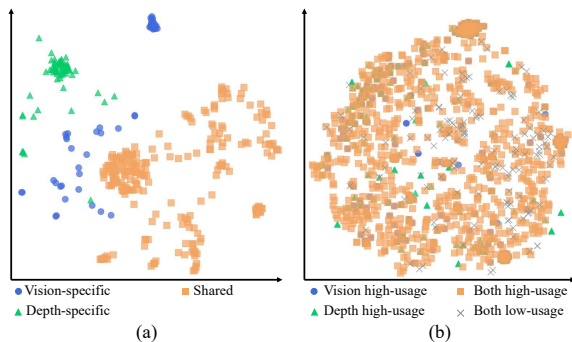


Figure 1: (a) **Distribution of codevectors** from the shared, vision-specific, and depth-specific codebooks. (b) **Distribution and usage rates of codevectors** in the shared codebook for shared embeddings from vision and depth modalities.

ing this adaptive approach against the predefined hyperparameter configurations detailed in Table 1. Notably, the adaptive strategy yields performance comparable to our manually tuned settings across both depth and thermal modalities. These results demonstrate the robustness and scalability of our framework, as it effectively accommodates low-resource modalities while eliminating the burden of exhaustive hyperparameter searching.

C Additional Experiment Results

C.1 Analysis of Codebook’s Feature Space

Codebook distribution We utilize t-SNE (Van der Maaten and Hinton, 2008) to visualize the shared codebook and the specific




Figure 2: **2D t-SNE visualization of sampled embeddings from AudioSet (Gemmeke et al., 2017).**

codebooks from different modalities. We use a conventional codebook design to facilitate observation, where each codevector has the same dimension as the input embedding. Specifically, the shared codebook contains 256 codevectors, while specific codebook for each modality contains 64 codevectors. As shown in Fig. 1(a), the varying distributions of specific and shared codebooks highlight modality-specific distinctions and their complementary roles in the feature space. This is consistent with observations in Fig. 2. Furthermore, t-SNE visualizations (Fig. 2) of 384 samples from AudioSet (Gemmeke et al., 2017), extracted via random batch sampling, reaffirm the pattern of decoupled representations. Consistent with previous analyses, these 2D projections contrast our compositional codebook with ImageBind (Girdhar et al., 2023) to demonstrate superior modal separation.

Shared codebook usage across modalities We visualize the distribution of codevectors within the shared codebook, color them by modality, and mark them based on their usage rate. The usage rate of codevectors is calculated in the validation set of the depth-image paired dataset NYU-D (Silberman et al., 2012), categorizing those with usage rates above 75% as high-usage and below 10% as




Figure 3: **Visualization of codevectors usage frequency distributions** of image-depth pairs in NYU-D dataset among various categories. Similar distribution patterns across two modalities indicate semantic consistency in our shared codebook.

low-usage. The usage rates of all codevectors in the shared codebook are then aggregated across different modalities (*i.e.*, vision and depth). In Fig. 1(b), it is evident that most codevectors are used by both modalities, indicating effective utilization of the shared codebook. A small proportion of codevectors, however, are primarily used by only one modality, suggesting some remaining differences in representations between the modalities. Additionally, a few codevectors exhibit low usage rates in either modality, reflecting the presence of less common information due to uneven sample distributions.

Codevector-level alignment We further investigate the alignment of shared embeddings from different modalities (*e.g.*, image and depth) at the codevector level. First, the usage rate of codevectors in the shared codebook is calculated based on the shared embeddings from paired data samples within each category. Then, the top-20 most frequently used codevectors are selected, and their usage rates are standardized across categories, with their distribution visualized using different colors in Fig. 3. The paired rows, where distributions from two modalities within the same category exhibit similar patterns of highlighted grids, demonstrate the fine-grained alignment between the shared embeddings from two modalities. In paired columns representing different modalities, similar patterns of highlighted grids can be noticed. This sug-

gests that our learned shared codevectors capture modality-invariant semantic information. Distinct patterns are evident between columns at different positions, indicating that each codevector captures unique semantic information.

C.2 Analysis of Specific Embeddings

Fine-grained Attribute	Specific Categories
Lighting illumination	natural diffused, artificial directional ambient even, harsh shadowed
Camera geometry	eye level close up, elevated wide angle, ground level perspective, top down view
Texture surface	smooth fine, coarse grainy, patterned structured, mixed textures
Scene environment	blurred background, structured setting, open space, indoor context
Color tone	warm natural, cool balanced, neutral muted, vibrant saturated

Table 3: **Fine-grained labels derived by VLM for linear probing.**

Attribute	NMI with category
lighting illumination	0.2093
camera geometry	0.1907
texture surface	0.3145
scene environment	0.3303
color tone	0.2759
Average NMI	0.2641

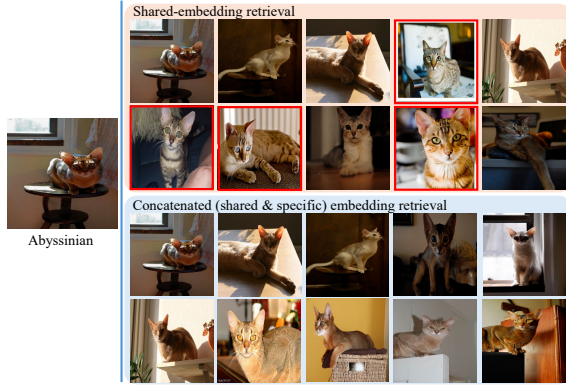
Table 4: **Mutual information analysis between category labels and fine-grained labels.** Normalized Mutual Information (NMI) between the VLM-curated fine-grained attributes and semantic class names on samples from ImageNet (Russakovsky et al., 2015). The low average NMI (0.2641) indicates that the extracted physical nuances are significantly disentangled from category-level semantics.

Fine-grained intra-modality retrieval We observe that 61.7% (3699/6000), 84.6% (2005/2371), and 90.3% (4496/4978) of retrieval samples achieved higher or equivalent top-10 recall using concatenated embeddings rather than shared embeddings on the Stanford Dogs and Oxford-IIIT Pet datasets (cats and dogs, respectively). Fig. 4 provides additional examples with corresponding retrieved images where this improvement is evident.

Linear probing for shared and specific embeddings We use Qwen2.5-VL (Bai et al., 2025b), a vision-language model (VLM), to annotate samples from 1000 images from ImageNet through



(a) Examples in Stanford Dogs dataset.



(b) Examples in cat images in Oxford-IIIT Pet dataset.



(c) Examples in dog images in Oxford-IIIT Pet dataset.

Figure 4: **Additional results for fine-grained retrieval.** By utilizing the concatenation of shared and specific embeddings, our method retrieves more correct images featuring the same cat or dog breed, outperforming scenarios that rely solely on shared embeddings.




Figure 5: **Heatmap between fine-grained attributes and category names.** Texture surface and scene environment are selected fine-grained attributes summarised by VLM. 20 categories are selected from ImageNet1K (Russakovsky et al., 2015) for display.

a systematic three-stage approach. (1) Detailed description generation: The VLM generates exhaustive descriptions of fine-grained features for sampled images. The prompts are specifically engineered to isolate distinctive details (as shown in Tab. 10), ensuring these features are decoupled from the common category-level information used in cross-modal sharing. (2) Keyword extraction: The VLM analyzes the detailed descriptions from Stage 1 to extract four key terms across five major fine-grained dimensions. This stage distills raw text into structured keywords optimized for classification. The keywords for each fine-grained attribute are shown in Fig. 3 (3) Fine-grained image labeling: Using the finalized keywords as a reference, the VLM performs the final annotation for each image, mapping specific visual features to the corresponding keywords.

Tab. 4 presents the Normalized Mutual Information scores (NMI) between semantic category labels and our VLM-derived fine-grained physical labels. Additionally, Fig. 5 visualizes the conditional probability distribution, $P(\text{physical attribute}|\text{category label})$, via a normalized heatmap. The analysis reveals a notable association with an average NMI score of 0.26, indicating that semantic categories account for approximately 26% of the variance in fine-grained physical attributes. This observation exemplifies the natural coupling inherent in real-world datasets. For specific classes, fine-grained features, such as background color or unique textures, serve as physical fingerprints that are strongly correlated with semantic identities (e.g., a “polar bear” is statistically intertwined with a “snowy background”). Conse-

quently, specific embeddings capture these correlated physical priors, which enable them to support category discrimination as effectively as shared embeddings. Nevertheless, the heatmap demonstrates that while certain species exhibit strong ties to specific attributes, most physical properties maintain a broad distribution across the entire dataset. This suggests that our VLM-derived labels capture critical variance that semantic categories alone cannot account for. Through our decoupled architecture, specific embeddings exhibit superior sensitivity in distinguishing these physical properties compared to their shared counterparts, particularly in terms of convergence speed and discriminative precision.

This video belongs to the category '{c_name}'. Describe the appearance and action of {c_name} in one short sentence. Focus on the segment from {s_t}s to {e_t}s. Use direct subject-verb-object structure. No filler words. Max 15 words. Do not use introductory phrases like 'In this video'.

Table 5: **Text prompt for VLM-based dense text caption.** Category names (c_name), and time minutes (s_t, e_t) are extracted from AVE dataset annotations.

Multimodal fusion on AVE event classification

We employ Qwen3-VL (Bai et al., 2025a) to generate dense textual captions by conditioning on the event category names and their corresponding video frames, using the prompt template detailed in Tab. 5. During the multimodal alignment text encoder by using the category names and text captions, respectively, for multimodal alignment training to align with vision and audio representations. Post-training, the modality encoders and their respective codebooks are frozen. A single-layer MLP classification head is then trained on the fixed embeddings to predict event categories. We evaluate three embedding configurations: (1) shared embeddings in isolation, (2) concatenation of video and audio features, and (3) their summation. For the latter two fusion strategies, modality-specific embeddings are concatenated with the shared embeddings for each respective modality. The classification head is optimized with a learning rate of 5×10^{-3} over 10 epochs.

C.3 Ablation on Codebook Size and Dimension

Dimension of codevectors We conduct additional experiments to explore the impact of varying the dimensions of codevectors in our compositional

Codebook Config.		Accuracy (%)	
Number (K)	Dimension (d)	SUN-D	FLIR_v2
1024	1024	40.7	81.1
1024	512	41.9	90.5
1024	128	44.3	92.5
1024	32	45.4	93.0
1024	8	45.7	97.2

Table 6: **Effectiveness of compositional codebook with varying codevector dimensions.** Results show improved effectiveness as codevector dimension decreases, indicating enhanced performance with richer combinations.

Codebook Config.	FLIR_v2 Accuracy (%)	
Number (K)	Conventional (dim=1024)	Compositional (dim=8)
64	58.6	86.2
128	65.4	91.2
256	76.6	94.0
512	81.9	93.2
1024	81.1	97.2

Table 7: **Impact of varying codevector numbers on accuracy.** Results show decreased accuracy with fewer codevectors. Our approach exhibits greater tolerance to this reduction compared to conventional codebooks due to its compositional structure.

codebook on its effectiveness. A series of experiments are carried out using a consistent number of codevectors (1024) while varying the dimensions, explicitly using codevector dimensions of {8, 32, 128, 512, 1024}. Tab. 6 shows that alignment improves as the dimension of the codevector decreases. Reducing the dimension of codevectors increases the number of combinations required to reconstruct an input embedding, thereby enhancing the combinatorial richness of the quantized output and leading to improved performance.

Number of codevectors To accommodate more information across various modalities in standard VQ, the conventional approach expands the codebook by introducing more codevectors, leading to a much increased codebook size. In contrast, our method maintains a compact codebook size by leveraging a compositional codebook design, achieving substantial more capacity even with a limited number of codevectors. We further explore the impact of reducing the number of codevectors while maintaining the same codevector dimension. The cross-modal classification performance using conventional codebooks (dim=1024) and our compositional codebook (dim=8) approaches are compared on FLIR_v2 (FLIR, 2018) dataset, using different numbers of codevectors

{64, 128, 256, 512, 1024}. As shown in Tab. 7, decreasing the codevector number results in a drop in accuracy. However, our approach demonstrates greater robustness to this reduction compared to conventional codebooks, owing to its compositional structure. The largest codebook size is limited to 1024, as in this case our partitioned codebook design (detailed in Sec. 3.3 in the main paper) allows for an enormous combinatorial capacity of 128^{1024} possible sub-vector compositions, which is sufficiently expressive.

Config.	Usage (%) \uparrow			Perplexity \downarrow		
	shared	vision	thermal	shared	vision	thermal
Conv.	75.6	80.9	93.8	90.8	106.7	127.3
Comp.	100	100	100	661.4	87.9	73.6

Table 8: **Codebook usage and perplexity** with compositional/conventional configurations on FLIR_v2. The modality-shared codebook size is set to 1024, while the specific codebooks for vision and thermal modalities are set to 256. Results indicate full utilization of both modality-shared and specific codebooks following the implementation of the compositional design.

Codebook collapse prevention Benefiting from the reinitialization strategy in CVQ (Zheng and Vedaldi, 2023) and the EMA update method described in Sec. A.2, the proposed compositional codebook achieves high usage efficiency. As shown in Tab 8, the shared codebook usage increases from 75.6% to 100%, accompanied by a significant increase in perplexity. Although the perplexity of the specific codebook decreases, this is probably due to the relatively small number of data categories compared to the overall codebook capacity.

C.4 Ablation on Loss Functions

Orthogonal and uniform loss The effectiveness of \mathcal{L}_{uni} and $\mathcal{L}_{\text{orth}}$ is evaluated on SUN-D (Song et al., 2015) by analyzing the intra-class average similarity scores of shared and specific embeddings both in image and depth before and after applying these losses. Within each modality, the cosine similarity scores among all shared embeddings (*i.e.*, intra-shared embedding similarity), specific embeddings (*i.e.*, intra-specific embedding similarity), and paired shared and specific embeddings of the complete modality representation (*i.e.*, inter-shared-specific embedding similarity) are computed within each category and then averaged. In Fig 6, before applying \mathcal{L}_{uni} , specific embeddings exhibited high similarity within each semantic category, indicating an overlap in encoded information. However, \mathcal{L}_{uni}




Figure 6: **Visualization of intra-class average similarity scores among shared and specific embeddings.** The intra-class average similarity scores are calculated on SUN-D (Song et al., 2015), with and without orthogonal loss $\mathcal{L}_{\text{orth}}$ and uniform loss \mathcal{L}_{uni} . The results demonstrate a substantial reduction in similarity among specific embeddings after applying these losses, indicating that they effectively encourage specific embeddings to encode distinct fine-grained information, while remaining within the same general semantic category, in contrast to shared embeddings.

significantly reduces the similarity between specific embeddings of depth from 0.98 to 0.25, ensuring that specific embeddings capture fine-grained, modality-specific details rather than aligning with broader semantic categories. Meanwhile, shared embeddings in depth, which encode category-level information aligned with text embeddings, maintain a higher similarity of 0.57 within each class. Additionally, $\mathcal{L}_{\text{orth}}$ further decreases the similarity between shared and specific embeddings in depth from an already low level to 0.02, reinforcing the separation of information, as the two embeddings are initialized in distinct codebook spaces. Similar patterns can be found in the embeddings of vision modality.

Codebook regularization loss CMCM loss \mathcal{L}_{cm} , orthogonal loss and uniform loss $\mathcal{L}_{\text{orth}}$, \mathcal{L}_{uni} have been discussed in Liu et al., 2022 and Xia et al., 2024; Jiang et al., 2023, respectively. We assess the efficacy of the proposed codevector contrastive loss \mathcal{L}_{ctr} and codevector uniformity loss $\mathcal{L}_{\text{cuni}}$ in enhancing the distinction capability of codevectors. To facilitate clear observation, a conventional codebook design is employed without dividing the embeddings, and the size of the shared and specific codebooks is reduced to 256 and 64, respectively. The cosine similarity of the specific embeddings from the sampled data in the FLIR_v2 (FLIR, 2018) and NYU-D (Silberman et al., 2012) datasets with all codevectors in the specific codebooks is calculated. A probability distribution of these similarities is then created

using the softmax function. Fig. 9 displays the similarity distributions for paired vision-thermal and paired vision-depth data. Without $\mathcal{L}_{\text{cctr}}$, $\mathcal{L}_{\text{cctr}}$, the distribution is more uniform, suggesting that a specific embedding interacts similarly with most codevectors, indicating less discriminative information. With $\mathcal{L}_{\text{cctr}}$, $\mathcal{L}_{\text{cctr}}$, the distribution becomes significantly uneven, showing that the specific embeddings align strongly with specific codevectors while neglecting others. This result confirms that the codevectors within a specific codebook are inherently distinct and can offer easily distinguishable information for the specific embeddings. This distinctiveness is further boosted by the application of our codevector regularization loss.

C.5 Additional CodeBind Applications

Cross-modal object localization We use the Recognize Anything Model (RAM) (Zhang et al., 2024a) to generate object tags from selected images in MS-COCO (Lin et al., 2014), and then apply GroundingDINO (Liu et al., 2024) to obtain the corresponding object bounding boxes, which serve as our vision proposals. Fig. 10 displays additional retrieval results from depth, audio, thermal, 3D point cloud, and tactile modalities, based on vision proposals predicted by GroundingDINO (Liu et al., 2024). In this process, shared embeddings of the [CLS] token from the proposal image and another modality are used to compute cosine similarity, enabling the selection of the top-1 retrieval data from the other modality for display. These vision proposals are aligned in our CodeBind space with matching embeddings from depth, audio, and 3D point clouds, sharing similar semantic meanings.

Modality	Sim ($e_{\text{gen}}, e_{\text{cond}}$)	Sim ($e_{\text{gen}}, e_{\text{text}}$)
Audio	0.3693	0.1749
Depth	0.4119	0.1792

Table 9: **Embedding similarity analysis** between the image embeddings of generated images (e_{gen}), the input modality embeddings (e_{cond}), and text embeddings (e_{text}).

Any-modality-to-image generation We replace the image encoder in Stable unCLIP (Rombach et al., 2022) with our aligned audio, depth, and thermal encoders to unlock its ability to generate semantically related images. Specifically, we use the shared embeddings of the [CLS] token from audio, depth, and thermal modalities. Fig. 7 presents




Figure 7: **Additional results for any-modal to image generation.** Semantically related images can be generated by pretrained diffusion model, using embeddings from audio, depth, and thermal modalities, which are effectively aligned with image and text embeddings through our CodeBind approach.

more generated images with their semantic categories from related modalities. To improve generation quality during inference, we apply a negative prompt to reduce unwanted artifacts: *"lowres, text, error, cropped, worst quality, low quality, jpeg artifacts, ugly, duplicate, morbid, mutilated, out of frame, extra fingers, mutated hands, poorly drawn hands, poorly drawn face, mutation, deformed, blurry, dehydrated, bad anatomy, bad proportions, extra limbs, cloned face, disfigured, gross proportions, malformed limbs, missing arms, missing legs, extra arms, extra legs, fused fingers, too many fingers, long neck, username, watermark, signature"*.

Without any retraining of the diffusion model or additional text prompts, semantic concepts from other modalities can be seamlessly integrated into the generated images. Additionally, a quantitative evaluation of the any-to-image generation is conducted. Specifically, we compute the embedding similarity between the vision embeddings of generated images (e_{gen}) and two references: (1) the input modality embeddings (used as diffusion conditions) (e_{cond}), and (2) the text embeddings of their corresponding semantic categories (e_{text}). Based on nearly 3,000 samples across all categories, respectively from the VGGSound and SUN-D datasets, the average similarity scores are summarized in Tab 9.


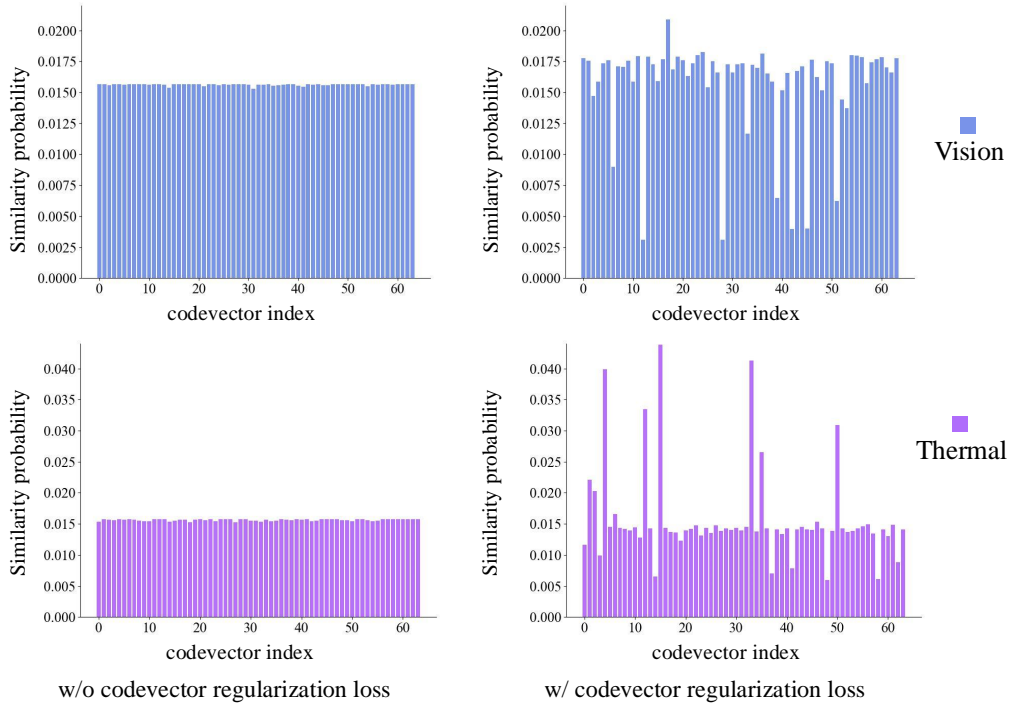


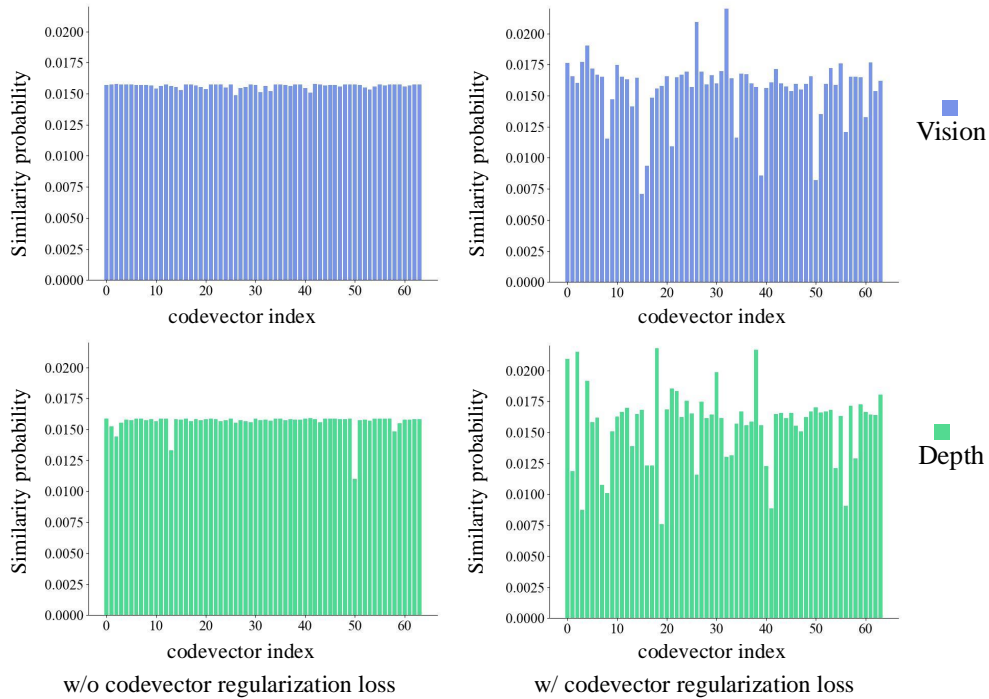
Figure 8: **Visualization of reconstructed images, depth and thermal images.** For each modality, the first row displays the sampled ground truth images, while the second row shows the corresponding reconstructed images.

C.6 Visualization of Reconstruction Results

We present visualizations of reconstructed RGB images, depth images, and thermal images from the Place365 (Zhou et al., 2014), NYU-D (Silberman et al., 2012), and FLIR_v2 (FLIR, 2018) datasets, respectively, in Fig 8. These results demonstrate that the concatenated embeddings effectively capture enough information for the decoder to accurately restore the original modality inputs.



(a) codevector distribution for specific embeddings selected in FLIR_v2 (FLIR, 2018)



(b) codevector distribution for specific-embeddings selected in NYU-D (Silberman et al., 2012)

Figure 9: **Visualization of codevector similarity distribution.** The left column shows distribution without codevector regularization loss \mathcal{L}_{ctr} and \mathcal{L}_{cuni} , while the right column presents the distribution with these losses applied. The codevector regularization loss clearly encourages an uneven distribution of codevectors. This enhancement promotes better discriminativeness of codevectors, ensuring each codevector effectively represents distinct features.

Stage	Prompt Content
Stage 1	<p>Please provide a comprehensive structural analysis of this image. Your output MUST be a valid JSON object with the following structure:</p> <pre>{ "class_name": "The category of the object for reference", "semantic_content": "Briefly identify the main objects and their categories.", "physical_nuances": { "lighting_illumination": "Describe source direction, intensity, shadows, and contrast.", "camera_geometry": "Describe specific viewpoint, lens perspective, and distance.", "texture_surface": "Describe visual frequency, graininess, or smoothness (strictly avoiding naming the object).", "scene_environment": "Describe background relationship, bokeh depth, or clutter level.", "color_tone": "Describe dominant color temperature, saturation, and chroma distribution." } }</pre> <p>STRICT CONSTRAINT: In the 'physical_nuances' section, do not use words that identify the specific species or category of the object. Focus purely on the imaging properties.</p>
Stage 2	<p>I will provide you with a collection of physical descriptions of images from various categories. Your task is to analyze these descriptions and synthesize a "Universal Physical Attribute Taxonomy".</p> <p>CORE GOALS:</p> <ol style="list-style-type: none"> (1) Disentanglement: The categories must represent 'Intra-class variance' (physical imaging properties) rather than 'Inter-class semantics' (object identity). (2) Structure: Define exactly 5 independent categories (e.g., Lighting, Camera Geometry, Scene Context, etc.). (3) Discreteness: For each category, provide exactly 4-5 mutually exclusive and clear options (labels). (4) Semantic-free: Ensure these labels do not mention specific object names like 'dog' or 'fur'. <p>INPUT DESCRIPTIONS: {xxxx}</p> <p>OUTPUT FORMAT (Strict JSON):</p> <pre>{ "taxonomy": { "category_name_1": ["option1", "option2", "option3", "option4"], "category_name_2": ["option1", "option2", "option3", "option4"], ... }, "logic": "Briefly explain why these 5 categories were chosen for representation disentanglement." }</pre>
Stage 3	<p>You are a high-precision image attribute annotator. For the provided image, you must select exactly ONE option from each category below.</p> <p>STRICT CONSTRAINTS: 1. DO NOT mention the name of the object. 2. Output MUST be a valid JSON object.</p> <p>CATEGORIES AND OPTIONS: {options_str}</p> <p>OUTPUT JSON FORMAT:</p> <pre>{ "Lighting_Illumination": "selected_option", "Camera_Geometry": "selected_option", "Texture_Surface": "selected_option", "Scene_Environment": "selected_option", "Color_Tone": "selected_option" }</pre>

Table 10: Multi-stage prompts for VLM-based fine-grained attribute annotation.

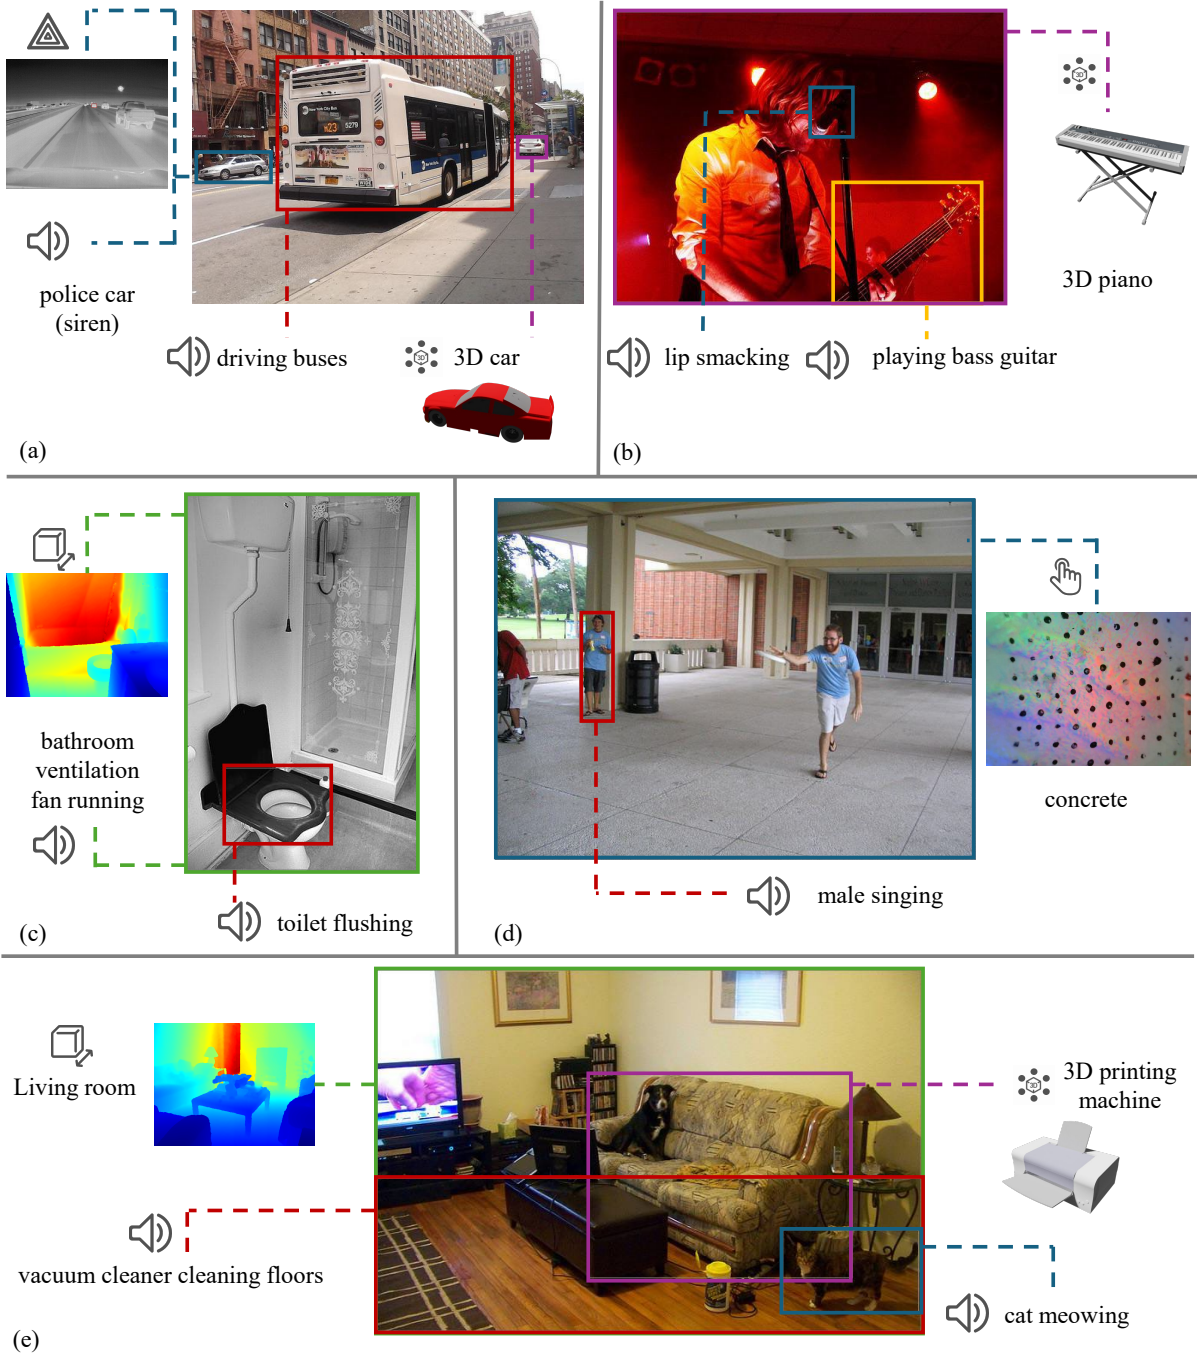


Figure 10: **Additional results for cross-modal object localization.** Semantically or geometrically related items from audio, depth, thermal, 3D point cloud, and tactile modalities can be effectively retrieved given several visual proposals in the images.