

UniBind: LLM-Augmented Unified and Balanced Representation Space to Bind Them All

Yuanhuiyi Lyu¹ * Xu Zheng¹ * Jizhou Zhou¹ Lin Wang^{1,2} †

¹AI Thrust, HKUST(GZ) ²Dept. of CSE, HKUST

{yuanhuiyilv, jiazhouzhou}@hkust-gz.edu.cn, zhengxu128@gmail.com, linwang@ust.hk

Project Page: <https://vlislab22.github.io/UniBind/>

Abstract

We present **UniBind**, a flexible and efficient approach that learns a unified representation space for seven diverse modalities – image, text, audio, point cloud, thermal, video, and event data. Existing works, e.g., ImageBind [13], treat the image as the central modality and build an image-centered representation space; however, the space may be sub-optimal as it leads to an unbalanced representation space among all modalities. Moreover, the category names are directly used to extract text embeddings for the downstream tasks, making it hardly possible to represent the semantics of multi-modal data. The ‘out-of-the-box’ insight of our UniBind is to make the alignment centers modality-agnostic and further learn a unified and balanced representation space, empowered by the large language models (LLMs). UniBind is superior in its flexible application to all CLIP-style models and delivers remarkable performance boosts. To make this possible, we 1) construct a knowledge base of text with the help of LLMs and multi-modal LLMs; 2) adaptively build LLM-augmented class-wise embedding centers on top of the knowledge base and encoded visual embeddings; 3) align all the embeddings to the LLM-augmented embedding centers via contrastive learning to achieve a unified and balanced representation space. UniBind shows strong zero-shot recognition performance gains over prior arts by an average of 6.36%. Finally, we achieve new state-of-the-art performance, e.g., a 6.75% gain on ImageNet, on the multi-modal fine-tuning setting while reducing 90% of the learnable parameters.

1. Introduction

Humans use multiple senses— each of which is from a different source, *a.k.a.*, modality— to perceive and interpret the world [36, 46]. Humans are naturally equipped

*Equal Contribution.

†Corresponding author.

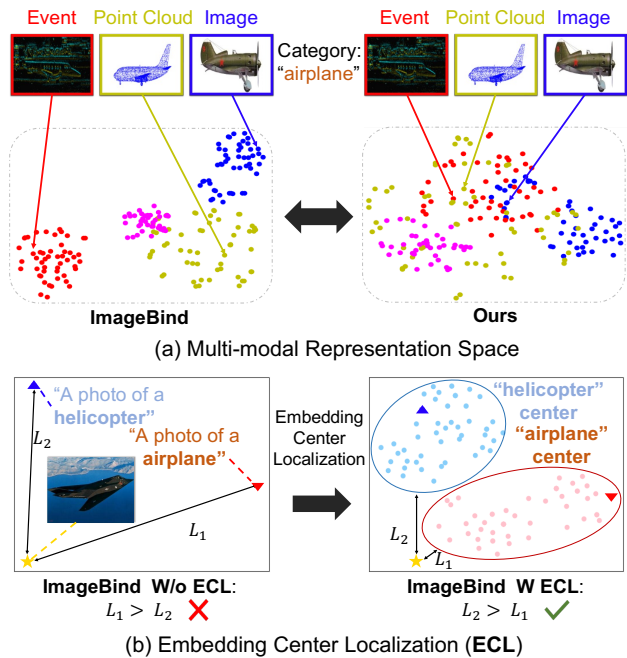


Figure 1. (a) By making the alignment center modality-agnostic, our UniBind can learn a unified and balanced representation space. (b) The embedding centers for each semantic category: these centers exhibit more complementary semantics compared to embeddings solely encoded by category names.

with the capacity to process and fuse multiple modalities simultaneously. For machines to emulate human intelligence, it is imperative for them to interpret, reason, and fuse multi-modal inputs, such as vision, text, audio, *etc.* [8]. This has inspired many methods that employ paired data to align image with text [4, 12, 24, 47] or align image with audio [16, 32]. Building on these works, early research has largely focused on integrating additional modalities, *e.g.*, CLIP2Video [10] and PointCLIP [59] for enhancing the comprehensiveness and accuracy of multi-modal data representation, and ultimately, improving performance

across various tasks. Recent endeavors have shown the possibility of learning across multiple modalities, including video [5, 48, 63], point cloud [18, 50, 59, 66], thermal [13, 57], event [64, 65], *etc.* Among these methods, ImageBind [13] sets a new way to learn a single shared representation space by leveraging multiple types of image-paired data. It utilizes the binding property of image modality to align the embeddings from the other modalities with the image embeddings.

However, as depicted in Fig. 1 (a), treating the image as the central modality and building an image-centered representation space leads to sub-optimal results, it may introduce bias and thus results in an unbalanced representation space among all modalities [9]. Also, as depicted in Fig. 1 (b), existing CLIP-style models, *e.g.*, ImageBind [13], solely utilize the text embeddings obtained from category names as embedding centers for categories. Nonetheless, category names, such as ['Airplane'] and ['Helicopter'], may not fully represent the semantics of the visual data, as there exist numerous images of airplanes with varying backgrounds and conditions.

This paper strives to tackle two problems: 1) the unbalanced representation space resulting from taking a specific visual modality as the alignment center, and 2) the unreliable nature of embedding alignment centers that rely solely on category names. Accordingly, we propose UniBind, a flexible and efficient approach for binding seven modalities—image, text, audio, point cloud, thermal, video, and event data. The core insight of our UniBind is to make the alignment centers modality-agnostic and further learn a unified and balanced representation space, empowered by the large language models (LLMs) and multi-modal large language models (multi-modal LLMs). Our UniBind is superior in its flexible application to all CLIP-style models and delivers remarkable performance boosts (+3.83% in N-caltech [38] with E-CLIP [65]).

Specifically, we first construct a knowledge base of texts which are extracted from the text generated by several LLMs, *e.g.*, GPT-4 [37] and LLaMa [49], as well as multi-modal LLMs, *e.g.*, BLIP-2 [23] and LLaMa-Adapter [60]. In practice, GPT-4 and LLaMa are utilized to generate the category descriptions, while BLIP-2 and LLaMa-Adapter are used to provide the multi-modal data descriptions (Sec. 3.2). Secondly, we compute the class-wise similarity between the input prompts and the text embeddings. It then utilizes the top 50 text embeddings to construct the corresponding class-wise text embedding center (Sec. 3.4). For example, as depicted in Fig. 1 (b), we select the top 50 text embeddings based on their similarity to the input prompts: ["A photo of helicopter/airplane."]. These selected embeddings are then utilized to construct the embedding centers for the categories of ['helicopter'] and ['airplane'].

Lastly, we align all modality embeddings toward the text embedding centers using contrastive learning loss functions (Sec. 3.3). This ensures that all modalities are equally considered in the representation space and achieve a unified and balanced representation space, as shown in Fig. 1 (a).

We apply our UniBind to the state-of-the-art (SoTA) CLIP-style multi-modal learning methods, including CLIP [40], E-CLIP [65], Audio-CLIP [16], PointCLIP [59], ImageBind [13], and PointBind [15], on 14 benchmarks from seven modalities. Note that, our UniBind is the first work to introduce the event modality [64] into the multi-modal representation space. UniBind consistently delivers significant performance improvements with all the CLIP-style multi-modal methods on all the benchmarks from the seven modalities, such as +5.55% in ImageNet-1K [7] with ImageBind [13] and +8.28% in N-caltech [38] with PointBind [15]. Moreover, we achieve new SoTA performance, *e.g.*, +6.75% gain on ImageNet-1K with the multi-modal fine-tuning setting while reducing 90% of the learnable parameters. Additionally, in the cross-modal retrieval tasks, our UniBind demonstrates a substantial performance improvement by +17.96% with PointBind on the top-20 recall score in the event-to-image retrieval task.

2. Related Work

Multi-modal Learning: From the modality alignment perspective, existing methods can be divided into two categories: alignment at the token and feature levels. (1) Token-level alignment methods [5, 10, 33, 58, 61, 63] align multi-modal token embeddings in a shared token embeddings space and design a subsequent encoder to extract the feature of these input token embeddings. (2) Feature-level alignment methods are based on the unified vision-language representation space, built by the CLIP style large vision-language models [22, 27, 29, 30, 40, 51, 54], *e.g.*, BLIP [22]. These methods adapt one [10, 16, 32, 59, 65, 66] or more [13, 15] modalities to the image representation space to align multiple visual modalities. Representative works include ImageBind [13], which learns a single shared representation space by leveraging multiple types of image-paired data. It leverages the binding property of images and aligns each modality's embeddings to image embeddings. Other works, such as PointCLIP [59, 66] and AudioCLIP [16], align point cloud and audio modalities, respectively, to the image representation space using cross-modal correlation or attention mechanisms. However, since these methods treat the image modality as the multi-modal alignment center, the obtained representation space is unbalanced among all the visual modalities [9], as demonstrated in Fig. 1. By contrast, we propose to learn modality-agnostic alignment centers, buttressed by the LLMs, thus yielding a unified and balanced visual representation space. Our UniBind binds the multi-modalities with the same se-

manatics to bridge the gap of multi-modalities.

LLMs and Knowledge Base: In Natural Language Processing (NLP), knowledge bases are widely used to enhance the understanding of human language [1, 6, 17, 44] and the robustness of generated results [14, 42, 56]. With the development of LLMs [49, 62], researchers use them to build knowledge bases [2, 25, 31, 35]. These methods can be categorized into two groups, including: 1) designing effective text prompts via LLMs to enhance the representation ability of the text encoder [26, 43]; and 2) utilizing LLMs to verbalize the semantics of text input by generating texts with similar semantics for the same category data, aiming to enhance the robustness of text representation [34, 66]. We utilize LLMs and multi-modal LLMs to construct a knowledge base to incorporate the prior knowledge of each category by generating the descriptions of the category name. This augments the text embeddings of the category names in the representation space.

Language-augmented Representation Learning aims to enhance the visual representation space by incorporating language, *i.e.*, text data. As pointed out by [9], language helps to identify conceptually similar image pairs even if they are visually dissimilar in visual recognition tasks. Moreover, efforts have been taken to leverage text as the representation centers for contrastive learning in information retrieval [27, 55]. For example, UniVL-DR [27] addresses the modality gap by verbalizing images to text and constructs a unified representation space for multi-modal dense retrieval, resulting in significant performance gains. Differently, we introduce a text embedding center strategy to the multi-modal domain. Our UniBind utilizes extracted text embeddings as the alignment centers and further binds the visual modality embeddings, thereby facilitating multi-modal representation learning and obtaining a balanced and unified representation space.

3. The Proposed UniBind

3.1. Problem Setting and Overview

Problem setting: We follow the multi-modal recognition setting popularized by ImageBind [13]. It uses the default set of text prompt templates $P_j^1, P_j^2, \dots, P_j^n$ from CLIP [40]. It then computes the similarity score between the input multi-modal data V_i and the C_j category by:

$$S_{(V_i, C_j)} = \cos \langle F^V(V_i), \text{mean}\{F^T(P_j^1, \dots, P_j^n)\} \rangle, \quad (1)$$

where F^V and F^T represent the visual and text encoders, respectively, which extract the image and text embeddings. The key insight of our UniBind is to make the alignment centers modality-agnostic and then learn a unified and balanced representation space for diverse modalities by leveraging the embedding centers constructed from the knowledge base, thereby binding them together. UniBind strives to address the two main challenges: 1) The unbal-

anced representation space that emerges from designating a particular visual modality as the alignment center, and 2) The unreliable nature of embedding alignment centers that exclusively depend on category names. To this end, firstly, UniBind constructs a knowledge base of text embeddings using LLMs and Multi-modal LLMs. Secondly, UniBind [13] adaptively builds LLM-augmented class-wise embedding centers based on the knowledge base and aligns multi-modal embeddings to the embedding centers with contrastive learning to build a unified embedding space.

Overview: An overview of UniBind is shown in Fig. 2. Specifically, given n visual modalities, Our UniBind includes n multi-modal encoders F_n and a text encoder F^T , which are adapted from existing multi-modal learning models such as ImageBind [13]. The only modification of these multi-modal learning models made in UniBind is adding a trainable linear layer to each of the n visual encoders, while the adopted encoders are all frozen during training. Our framework includes two stages:

1) Training Stage: For training, we initially construct the knowledge base (Sec.3.2) by incorporating both LLMs and multi-modal LLMs. We then leverage the knowledge base to learn a unified representation space (Sec.3.3) via LLM-augmented contrastive learning.

2) Inference Stage: Building on our unified multi-modal representation space, we infer recognition results via a novel Embedding Center Localization module (Sec. 3.4). We now describe these technical components in detail.

3.2. Knowledge Base Construction

The construction of the knowledge base comprises two parts: 1) Category descriptions, generated by LLMs. These texts are employed to localize the embedding centers. 2) Multi-modal data descriptions produced by multi-modal LLMs. These descriptions help to alleviate the modality gaps prevalent among multiple modalities.

Although it has been demonstrated in [9, 27] that language is a powerful tool for capturing semantic relationships among multi-modal data, dependency on category names exclusively to align modalities with extracted embeddings is unreliable. As discussed in Sec. 1, category names cannot fully capture the semantics of multi-modal data. To address this issue, we first use LLMs, such as GPT-4 [37] and LLaMa [49], to generate category descriptions based on the category names, as shown in Fig. 3:

$$T_{C_i}^1, \dots, T_{C_i}^n = F^{LLMs}(C_i), \quad (2)$$

where the $T_{C_i}^1, \dots, T_{C_i}^n$ are the n generated descriptions for the category C_i and F^{LLMs} are the aforementioned LLMs. Subsequently, we produce descriptions for multi-modal data via multi-modal LLMs:

$$T_{I_i}, \dots, T_{A_i} = F^{MLLMs}(I_i), \dots, F^{MLLMs}(A_i), \quad (3)$$

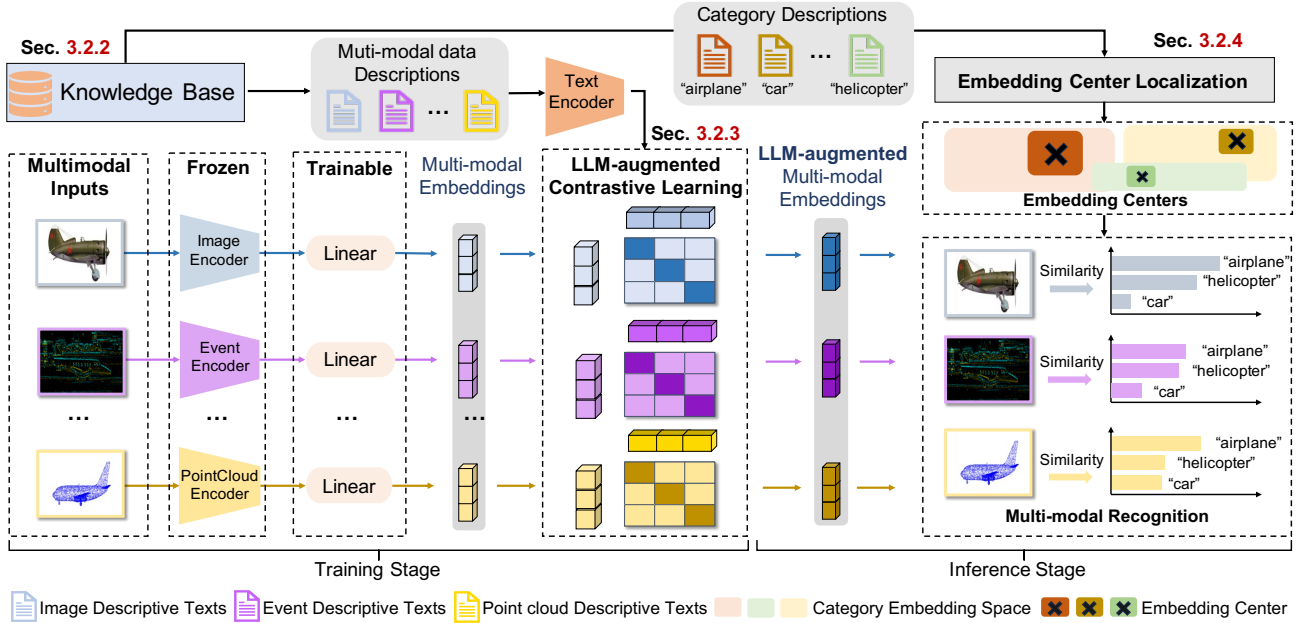


Figure 2. An overview of our UniBind. Firstly, we construct the knowledge base and then learn a unified representation space via LLM-augmented contrastive learning. Lastly, We utilize the embedding center localized by the knowledge base to obtain the predictions.

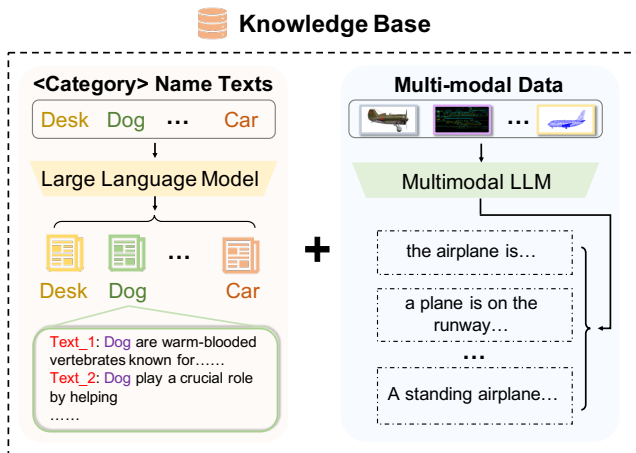


Figure 3. Knowledge Base. Generation pipeline for category descriptions (left) and multi-modal data descriptions (right).

here, I_i, \dots, A_i represent the multi-modal data inputs, while T_{I_i}, \dots, T_{A_i} are the generated multi-modal data descriptions and F^{MLLMs} denote the aforementioned multi-modal LLMs. As an example, consider the category ['Desk'], we use LLMs to generate category descriptions, such as ["A computer monitor is prominently displayed on the desk, indicating it is a workstation"]. We then utilize multi-modal LLMs to generate descriptions for image data falling within the ['Desk'] category. Finally, we compile these two sets of descriptions to formulate our knowledge base.

3.3. Unified Representation Space Learning

Expanding upon our knowledge base, we subsequently align multiple modalities to learn a unified multi-modal representation space. For multi-modal data, we utilize feature encoders derived from existing multi-modal models with frozen parameters and learnable subsequent linear layers, to obtain embeddings for each modality:

$$v_{I_i}, \dots, v_{A_i} = F_I(I_i), \dots, F_A(A_i), \quad (4)$$

where the v_{I_i}, \dots, v_{A_i} are the extracted embeddings and $F_I(\cdot), \dots, F_A(\cdot)$ are the feature encoders for each modality. Subsequently, we generate text embeddings of multi-modal data descriptions T_{I_i}, \dots, T_{A_i} via text encoder:

$$z_{I_i}, \dots, z_{A_i} = F^T(T_{I_i}), \dots, F^T(T_{A_i}), \quad (5)$$

where the T_{I_i}, \dots, T_{A_i} are the generated multi-modal data descriptions, $F^T(\cdot)$ is the text encoder, and z_{I_i}, \dots, z_{A_i} are the extracted text embeddings. The extracted visual and text embeddings are employed for learning a unified representation space. In contrast to existing multi-modal learning frameworks, such as ImageBind [13], we do not impose contrastive learning objectives among visual data with the image center. Instead, we impose contrastive learning objectives directly between the multi-modal and text embeddings. As an illustration, for aligning the visual modality I to our unified representation space, the extracted visual embeddings v_{I_1}, \dots, v_{I_n} and the corresponding text embeddings z_{I_1}, \dots, z_{I_n} are employed for contrastive learning within this representation space. Text embeddings gener-


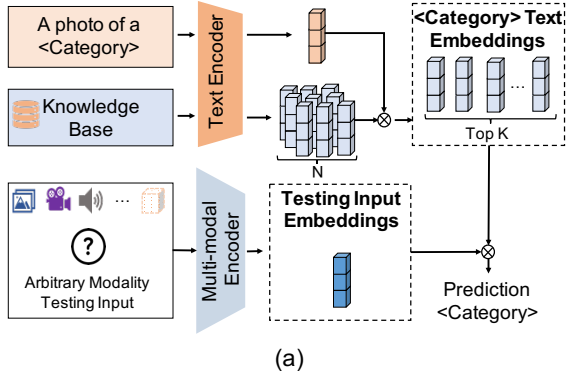


Figure 4. (a) The details for our embedding center localization. (b) The impact of our embedding center localization is demonstrated.

ated from corresponding descriptions are considered positive samples for input visual data, whereas text embeddings from other visual data are utilized as negative samples:

$$\mathcal{L}_{(\mathcal{I}, \mathcal{A})} = -\log \frac{\exp(v_{I_i}^T \cdot z_{I_i} / \tau)}{\exp(v_{I_i}^T \cdot z_{I_i} / \tau) + \sum_{j \neq i} \exp(v_{I_i}^T \cdot z_{I_j} / \tau)}, \quad (6)$$

where z_{I_j} is the corresponding text embeddings of the visual data I_j in visual modality I .

3.4. Embedding Center Localization

We localize our embedding centers by selecting 50 text embeddings for each category from the knowledge base. Specifically, we utilize the basic prompt ["A photo of a [Category]"] to calculate cosine similarity, subsequently selecting the top 50 descriptions based on the highest cosine similarity scores. As depicted in Fig. 4 (a), we subsequently derive the text embedding center EC_i for category C_i from these top 50 descriptions by processing them through the text encoder F^T :

$$EC_i = \{z_{C_i}^1, \dots, z_{C_i}^{50}\} = F^T(T_{C_i}^1), \dots, F^T(T_{C_i}^{50}), \quad (7)$$

The text embeddings generated from the top 50 descriptions collectively form the embedding center EC_i for each category. Consequently, UniBind extracts embedding centers from the knowledge base which have more complementary semantics than simply using the category names. Unlike existing methods exemplified in Eq. 1, our embedding centers *establish more distinct category boundaries in representation space* (shown in Fig. 7). For instance, we compute the

Modalities	Dataset	Metric	Scale	#cls
Image	ImageNet-1K (IN-1K) [7]	Acc	1,280K	1,000
	Places-Stanford-365 (P365) [28]	Acc	1,240K	365
	Caltech-101 (cal) [11]	Recall	8K	101
PointCloud	ModelNet-40 (ModelNet40) [52]	Acc	9K	40
	ShapeNet-part (ShapeNet) [3]	Acc	16K	16
Audio	ESC 5-folds (ESC) [39]	Acc	2K	50
	Urban-Sound-8K (Urban-S) [41]	Acc	8K	10
Thermal	LLVIP (LLVIP) [20]	Acc	15K	2
	RGB-T Selected (RGB-T) [19]	Acc	10K	2
Video	MSR-VTT (MSR-VTT) [53]	Acc	10K	20
	UCF-101 (UCF-101) [45]	Acc	14K	101
Event	N-Caltech-101 (N-cal) [38]	Acc & Recall	8K	101
	N-ImageNet-1K (N-IN-1K) [21]	Acc	1,280K	1,000

Table 1. Summary of experimental settings across various modalities. We report the task, dataset, and data scale for each modality.

similarity between an arbitrary modality input M_i and the category C_j as follows:

$$S_{(M_i, EC_j)} = \max\{\cos \langle F_m(M_i), z_{C_i}^1, \dots, z_{C_i}^{50} \rangle\}, \quad (8)$$

where the z_I, \dots, z_A are the extracted text embeddings of category C_i . As shown in Fig. 4 (b), compared with the mean of text prompt embeddings, our embedding centers have more significant spatial distributions. With our embedding centers, multi-modal data distributed at the boundary of the category representation space can effectively avoid interference from other categories (shown in Fig. 7), thus facilitating more accurate recognition.

3.5. Implementation

UniBind can be flexibly implemented with different existing CLIP-style multi-modal learning models, such as PointCLIP [59], ImageBind [13] and PointBind [15].

Backbone Models We use existing CLIP style multi-modal learning models as the backbones to implement our UniBind. Concretely, we implement our UniBind with the following models: CLIP [40], ImageBind [13], PointBind [15], E-CLIP [65], PointCLIP [66], and AudioCLIP [16]. We use separate visual encoders for image, point cloud, audio, thermal, video, and event data. We add a simple linear layer at the end of the visual encoders of each modality for mapping the multi-modal embeddings to our unified representation space. We utilize the frozen text encoder from the backbone model as our text encoder.

Training and Inference UniBind can be utilized for both zero-shot and fine-tuning recognition tasks. For zero-shot tasks, as depicted in Fig. 2, UniBind employs the basic prompt ["A photo of a [Category]"] to select the top 50 most related text embeddings from the knowledge base as class-wise embedding centers. The similarities between the multi-modal embeddings and the class-

Model	Image		Point Cloud		Audio		Thermal		Video		Event	
	IN-1K	Place-365	ModelNet40	ShapeNet	ESC-50	Urban-S	LLVIP	RGB-T	MSR-VTT	UCF-101	N-Cal	N-IN-1K
Fine-tuning Setting												
Meta-Transformer [61]	83.10	52.70	90.50	99.30	-	-	-	-	-	46.60	\times	\times
ImageBind [13] w/ linear	80.19	49.45	\times	\times	83.40	71.60	-	60.55	63.81	98.06	\times	\times
PointBind [15] w/ linear	80.19	49.45	90.64	99.09	83.40	71.60	-	60.55	63.81	98.06	\times	\times
PointBind (+Event)	80.19	49.45	90.64	99.09	83.40	71.60	-	60.55	63.81	98.06	77.83	23.69
Ours w/ PointBind	86.94	56.99	90.72	99.59	84.01	69.09	-	60.71	69.53	93.31	78.05	24.48
Δ	+6.75	+7.54	+0.08	+0.50	+0.61	-2.51	-	+0.16	+5.72	-4.75	+0.22	+0.79
Zero-shot Setting												
ImageBind [13]	77.70	45.40	\times	\times	66.90	41.73	63.40	54.71	31.27	64.84	\times	\times
PointBind [15]	77.70	45.40	77.67	98.85	66.90	41.73	63.40	54.71	31.27	64.84	\times	\times
PointBind (+Event)	77.70	45.40	77.67	98.85	66.90	41.73	63.40	54.71	31.27	64.84	50.98	10.79
Ours w/ PointBind	83.25	53.84	80.59	98.96	71.70	62.56	64.67	56.20	40.90	73.74	59.26	13.85
Δ	+5.55	+8.44	+2.92	+0.11	+4.80	+20.83	+1.27	+1.49	+9.63	+8.90	+8.28	+3.06

Table 2. Emergent zero-shot and fine-tuning recognition on six modalities.

wise embedding centers are then utilized to make recognition predictions. For fine-tuning, there is a training stage with the proposed representation space learning, as shown in Fig. 2. The inference process is the same as zero-shot.

4. Experiments

4.1. Datasets and Implementation Details

Modalities and datasets. We evaluate UniBind on seven modalities – image, point cloud, audio, thermal, video, event, and text. For each modality, we assess our UniBind on two mainstream datasets at least. A summary of the datasets utilized is presented in Table 1.

Multi-modal backbone models. Since our UniBind can be flexibly applied to the existing CLIP-style multi-modal learning models, in this paper, we implement UniBind with ImageBind [13], PointBind [15], CLIP [40], E-CLIP [65], Audio-CLIP [16], and Point-CLIP [59]. The backbone models are kept frozen and the linear layers at the end of visual encoders are updated during the LLM-augmented contrastive learning.

Knowledge Base. We generate 1,000 descriptions for each category name via LLMs (GPT-4 [37], LLaMa [49]) and generate multi-modal data descriptions via multi-modal LLMs (BLIP-2 [23], LLaMa-Adapter [60]). We construct our knowledge base with these two sets of description texts.

4.2. Zero-shot Recognition

Settings. The emergent zero-shot recognition is first proposed in ImageBind [13] which means just by pre-training on (image, text) and (image, audio). As shown in Tab. 1, we evaluate UniBind in 12 mainstream datasets from 6 modalities. We directly test the recognition performance without training (*more details can be found in the suppl.*).

Model	Image		Event/Audio/PC	
	IN-1K	Place-365	Dataset 1	Dataset 2
CLIP [40]	68.30	29.95	\times	\times
Ours w/ CLIP	78.63	39.14	\times	\times
Δ	+10.35	+9.19	-	-
E-CLIP [65]	68.30	29.95	50.40	4.13
Ours w/ E-CLIP	78.63	39.14	54.26	7.91
Δ	+10.35	+9.19	+3.83	+3.78
Audio-CLIP [16]	40.51	18.76	68.60	68.78
Ours w/ Audio-CLIP	46.44	22.60	71.25	69.52
Δ	+5.93	+3.84	+2.65	+0.74
Point-CLIP [59]	59.60	25.56	20.20	89.20
Ours w Point-CLIP [59]	62.58	27.10	21.43	90.27
Δ	+2.98	+1.54	+1.23	+1.07

Table 3. Emergent zero-shot recognition in image + X modalities. Dataset 1 and 2 indicate *N-Cal* and *N-IN-1K*, *ESC-50* and *Urban-S*, and *ModelNet40* and *ShapeNet*, respectively.

Results. We evaluate the zero-shot recognition performance in comparison to existing methods across various modalities, including image, point cloud, audio, thermal, video, and event. In Tab. 2, we present the performance results of our approach when applied with CLIP-style multi-modal learning models that align more than three modalities. Additionally, we show the performance of our UniBind in conjunction with two-modality methods in Tab. 3. Our UniBind significantly improves the performances of CLIP-style multi-modal models. Across all benchmarks, UniBind achieves large gains about an average of +6.27% in top 1 accuracy and even compares favorably to supervised specialist models trained for the special modality and task.

4.3. Fine-tuning Recognition

Settings. We follow ImageBind [13] and only train the linear layer after the frozen visual encoders with the training

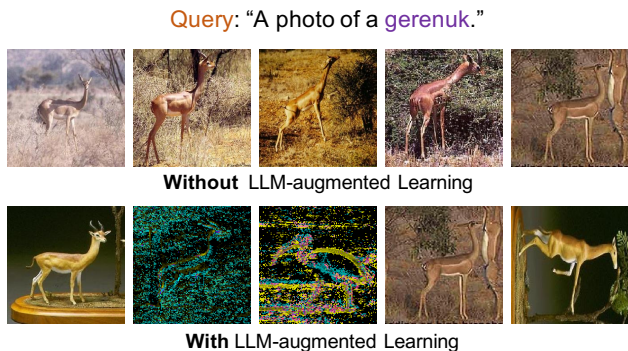


Figure 5. Top 5 results from text to events & images retrieval. We choose ["A photo of a [category]"] as the query to retrieve events and images in the same embedding space.

dataset, and then evaluate our UniBind on the testing dataset with the same metric in the zero-shot setting.

Results. In Tab. 2, we compare our approach with the supervised methods that use ["A photo of a [category name]"] as the text label during the training stage. Our UniBind outperforms the supervised method on 10 benchmarks spanning 6 modalities, exhibiting an average improvement of +1.26%. In particular, our UniBind shows more significant gains in the datasets containing a large number of categories, such as +6.75% in ImageNet (1,000 categories) and +7.54% in Place-365 (365 categories). It demonstrates the advantages of our approach to apply in complex semantic data.

5. Ablation Study and Analysis

5.1. LLM-augmented Contrastive Learning

To investigate the effectiveness of our proposed LLM-augmented contrastive learning, we conduct ablation studies on the cross-modal retrieval task. We experiment with E-CLIP [65] and PointBind [15] and subsequently report the results for event-to-image retrieval and image-to-event retrieval in Tab. 4. The recall score improvement increases incrementally from the top 1 to the top 20, illustrating the efficacy of our approach in aligning modalities with the same semantics. In addition, as shown in Fig. 5, we demonstrate the case of cross-modal retrieval based on PointBind [15] adapted event modality. In this case, we use the text ["A photo of a Gerenuk."] to retrieve images and events in the same representation space. In the absence of LLM-augmented contrastive learning, the top 5 retrieval results solely consist of images. By contrast, with LLM-augmented contrastive learning, the retrieval results are more balanced across image and event modalities.

Furthermore, the results of the t-SNE visualization in Fig. 6 reveal the differences between the representation spaces constructed by ImageBind [13] / PointBind [15] and our UniBind. For example, in Fig. 6 (a), We select em-

Model	Image-to-Event			Event-to-Image		
	R@1	R@10	R@20	R@1	R@10	R@20
E-CLIP [65]	79.52	93.08	95.51	76.29	91.80	94.61
E-CLIP w LCL	78.95	94.32	97.06	77.04	93.62	96.70
Δ	-0.57	+1.24	+1.55	+0.75	+1.82	+2.09
PointBind (+Event) [15]	14.07	40.79	49.46	9.00	29.32	37.70
PointBind w LCL	14.12	41.25	50.98	14.29	44.34	55.66
Δ	+0.05	+0.46	+1.52	+5.29	+15.02	+17.96

Table 4. Multi-modal retrieval result with/without LLM-augmented Contrastive Learning (LCL). We evaluate E-CLIP [65] and PointBind [15] in Image-to-Event and Event-to-Image tasks.




Figure 6. Representation space visualization of ImageBind / PointBind and our UniBind. We sample 64 data in the same semantic for each modality, specifically, ['Airplane'] data representation space shown in (a), ['Car'] data shown in (b).

beddings with the same semantic label ['airplane'] from image, audio, point cloud, and event modalities, and visualize 64 randomly chosen samples from each modality. In the representation space, embeddings from different modalities tend to cluster around their respective modalities. Thereby, with LLM-augmented contrastive learning, multi-modal embeddings cluster around the same semantic label in our unified modality-agnostic representation space.

5.2. Embedding Center Localization

The results presented in Tab. 5 demonstrate the effect of improving the zero-shot recognition performance with E-CLIP [65], AudioCLIP [16], ImageBind [13], and PointBind [15] on four modalities. Our embedding center localization method attains a substantial improvement, averaging +6.50% when applied to these CLIP-style multi-modal learning models. As the t-SNE visualization in Fig. 7 illustrated, our approach results in more distinct semantic boundaries between different categories, effectively enhancing recognition accuracy and reducing interference from other categories. We also study the impact of LLMs

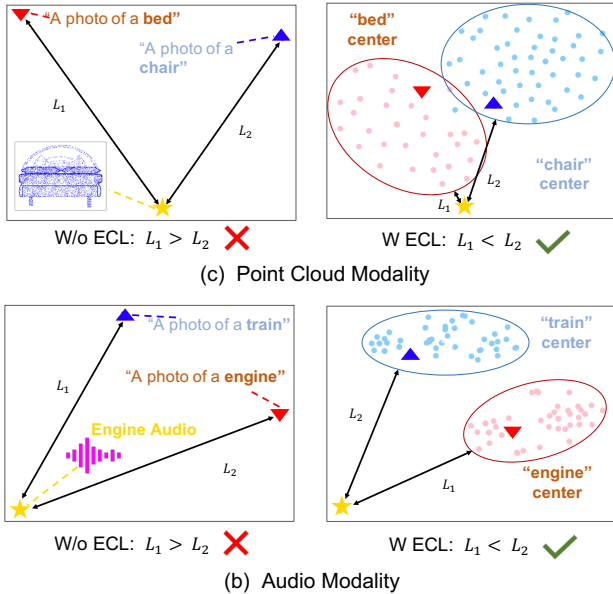


Figure 7. Embedding centers t-SNE visualization in the point cloud and audio modalities. (a) in this case, we simply use the prompts as the centers of both the bed and chair categories; on the right, we use our localized embedding centers. (b) show the cases in audio modality.

Modality	Image	Audio	PointCloud	Event
E-CLIP [65]	68.30	×	×	50.40
E-CLIP w ECL	78.63	×	×	54.26
Δ	+10.33	-	-	+3.86
AudioCLIP [16]	40.51	68.60	×	×
AudioCLIP w ECL	46.44	71.25	×	×
Δ	+5.93	+2.65	-	-
ImageBind [13]	77.70	66.90	×	×
ImageBind w ECL	83.25	71.70	×	×
Δ	+5.55	+4.80	-	-
PointBind (+Event) [15]	77.70	66.90	77.67	50.98
PointBind w ECL	83.25	71.70	80.59	59.26
Δ	+5.55	+4.80	+2.92	+8.28

Table 5. Performance in zero-shot recognition task with/without Embedding Center Localization (ECL) with five multi-modal models in four modalities.

and Multi-modal LLMs, while further examining the optimal number of selected texts for each category. We compare our UniBind to the method that only selects texts from LLMs or only from Multi-modal LLMs to construct the knowledge base and present the results in Fig. 8. Evidently, the knowledge base created by LLMs and Multi-modal LLMs exhibits the best performance, with the selection of the top 50 texts for each category proving to be the most effective choice for localizing the embedding centers. Lastly, we present the results of comparisons with




Figure 8. Ablation study of the knowledge base. (a) compares three ways to build a knowledge base by only LLMs, only Multi-modal LLMs, and both LLMs and Multi-modal LLMs. (b) shows the performance of selecting the top 10-100 texts for each category in image and event modalities.

Method	Image	PointCloud	Video	Event
Simple Prompts [40]	75.86	76.02	30.92	50.39
Complex Prompts [15, 40]	77.70	77.67	31.27	50.98
Word-net Augmented [17]	78.10	79.09	41.03	50.60
LLM-generated Prompts [66]	79.59	77.43	34.13	51.93
Ours	83.25	80.59	40.90	59.26

Table 6. Performance of various language-augmented methods in the zero-shot recognition task. We compare our LLM-augmented method with simple prompts by default of CLIP [40], complex prompts used by PointBind [15], word-net augmented prompts [17], and LLM-generated prompts.

other language-augmented methods in Tab. 6. Our UniBind demonstrates the best performance.

6. Conclusion

In this paper, we proposed UniBind, a multi-modal learning approach that renders the alignment centers modality-agnostic and further learns a unified and balanced representation space, empowered by the large language models (LLMs) and the multi-modal large language models (multi-modal LLMs). Our UniBind achieves remarkable performance boosts and is compatible with all CLIP-style multi-modal learning models. Additionally, we examined the potential of LLMs and Multi-modal LLMs for multi-modal representation space learning.

Limitations and Future Works. The robustness of the LLM-augmented method requires enhancement. In response, our future work will concentrate on harnessing the capabilities of LLMs to augment the robustness of the modality-agnostic representation space.

Acknowledgement. This paper is supported by the National Natural Science Foundation of China (NSF) under Grant No. NSFC22FYT45 and the Guangzhou City, University and Enterprise Joint Fund under Grant No.SL2022A03J01278.

References

- [1] Badr AlKhamissi, Millicent Li, Asli Celikyilmaz, Mona Diab, and Marjan Ghazvininejad. A review on language models as knowledge bases. *arXiv preprint arXiv:2204.06031*, 2022. [3](#)
- [2] Salvatore Carta, Alessandro Giuliani, Leonardo Piano, Alessandro Sebastian Podda, Livio Pompianu, and Sandro Gabriele Tiddia. Iterative zero-shot llm prompting for knowledge graph construction. *arXiv preprint arXiv:2307.01128*, 2023. [3](#)
- [3] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. [5](#)
- [4] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020. [1](#)
- [5] Yiting Cheng, Fangyun Wei, Jianmin Bao, Dong Chen, and Wenqiang Zhang. Cico: Domain-aware sign language retrieval via cross-lingual contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19016–19026, 2023. [2](#)
- [6] Rajarshi Das, Ameya Godbole, Ankita Naik, Elliot Tower, Manzil Zaheer, Hannaneh Hajishirzi, Robin Jia, and Andrew McCallum. Knowledge base question answering by case-based reasoning over subgraphs. In *International conference on machine learning*, pages 4777–4793. PMLR, 2022. [3](#)
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [2](#), [5](#)
- [8] Shengshun Duan, Qiongfeng Shi, and Jun Wu. Multimodal sensors and ml-based data fusion for advanced robots. *Advanced Intelligent Systems*, 4(12):2200213, 2022. [1](#)
- [9] Mohamed El Banani, Karan Desai, and Justin Johnson. Learning visual representations via language-guided sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19208–19220, 2023. [2](#), [3](#)
- [10] Han Fang, Pengfei Xiong, Luhui Xu, and Yu Chen. Clip2video: Mastering video-text retrieval via image clip. *arXiv preprint arXiv:2106.11097*, 2021. [1](#), [2](#)
- [11] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 conference on computer vision and pattern recognition workshop*, pages 178–178. IEEE, 2004. [5](#)
- [12] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems*, 26, 2013. [1](#)
- [13] Rohit Girdhar, Alaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15180–15190, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [14] Liangke Gui, Borui Wang, Qiuyuan Huang, Alex Hauptmann, Yonatan Bisk, and Jianfeng Gao. Kat: A knowledge augmented transformer for vision-and-language. *arXiv preprint arXiv:2112.08614*, 2021. [3](#)
- [15] Ziyu Guo, Renrui Zhang, Xiangyang Zhu, Yiwen Tang, Xi-anzheng Ma, Jiaming Han, Kexin Chen, Peng Gao, Xi-anzhi Li, Hongsheng Li, et al. Point-bind & point-llm: Aligning point cloud with multi-modality for 3d understanding, generation, and instruction following. *arXiv preprint arXiv:2309.00615*, 2023. [2](#), [5](#), [6](#), [7](#), [8](#)
- [16] Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel. Audioclip: Extending clip to image, text and audio. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 976–980. IEEE, 2022. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [17] Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Jingang Wang, Juanzi Li, Wei Wu, and Maosong Sun. Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification. *arXiv preprint arXiv:2108.02035*, 2021. [3](#), [8](#)
- [18] Tianyu Huang, Bowen Dong, Yunhan Yang, Xiaoshui Huang, Rynson WH Lau, Wanli Ouyang, and Wangmeng Zuo. Clip2point: Transfer clip to point cloud classification with image-depth pre-training. *arXiv preprint arXiv:2210.01055*, 2022. [2](#)
- [19] Soonmin Hwang, Jaesik Park, Namil Kim, Yukyung Choi, and In So Kweon. Multispectral pedestrian detection: Benchmark dataset and baseline. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1037–1045, 2015. [5](#)
- [20] Xinyu Jia, Chuang Zhu, Minzhen Li, Wenqi Tang, and Wenli Zhou. Llvip: A visible-infrared paired dataset for low-light vision. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3496–3504, 2021. [5](#)
- [21] Junho Kim, Jaehyeok Bae, Gangin Park, Dongsu Zhang, and Young Min Kim. N-imagenet: Towards robust, fine-grained object recognition with event cameras. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2146–2156, 2021. [5](#)
- [22] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. [2](#)
- [23] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. [2](#), [6](#)
- [24] Xiujuan Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX 16*, pages 121–137. Springer, 2020. [1](#)

- [25] Zhenyu Li, Sunqi Fan, Yu Gu, Xiuxing Li, Zhichao Duan, Bowen Dong, Ning Liu, and Jianyong Wang. Flexkbqa: A flexible llm-powered framework for few-shot knowledge base question answering. *arXiv preprint arXiv:2308.12060*, 2023. 3
- [26] Long Lian, Boyi Li, Adam Yala, and Trevor Darrell. Llm-grounded diffusion: Enhancing prompt understanding of text-to-image diffusion models with large language models. *arXiv preprint arXiv:2305.13655*, 2023. 3
- [27] Zhenghao Liu, Chenyan Xiong, Yuanhuiyi Lv, Zhiyuan Liu, and Ge Yu. Universal vision-language dense retrieval: Learning a unified representation space for multi-modal retrieval. In *The Eleventh International Conference on Learning Representations*, 2022. 2, 3
- [28] Alejandro López-Cifuentes, Marcos Escudero-Vinolo, Jesús Bescós, and Álvaro García-Martín. Semantic-aware scene recognition. *Pattern Recognition*, 102:107256, 2020. 5
- [29] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508:293–304, 2022. 2
- [30] Yuanhuiyi Lyu, Xu Zheng, and Lin Wang. Image anything: Towards reasoning-coherent and training-free multi-modal image generation. *arXiv preprint arXiv:2401.17664*, 2024. 2
- [31] Xinbei Ma, Yeyun Gong, Pengcheng He, Hai Zhao, and Nan Duan. Query rewriting for retrieval-augmented large language models. *arXiv preprint arXiv:2305.14283*, 2023. 3
- [32] Tanvir Mahmud and Diana Marculescu. Ave-clip: Audioclip-based multi-window temporal transformer for audio visual event localization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5158–5167, 2023. 1, 2
- [33] Jishnu Mukhoti, Tsung-Yu Lin, Omid Poursaeed, Rui Wang, Ashish Shah, Philip HS Torr, and Ser-Nam Lim. Open vocabulary semantic segmentation with patch aligned contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19413–19423, 2023. 2
- [34] Muhammad Ferjad Naeem, Muhammad Gul Zain Ali Khan, Yongqin Xian, Muhammad Zeshan Afzal, Didier Stricker, Luc Van Gool, and Federico Tombari. I2mvformer: Large language model generated multi-view document supervision for zero-shot image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15169–15179, 2023. 3
- [35] Anmol Nayak and Hari Prasad Timmapathini. Llm2kb: Constructing knowledge bases using instruction tuned context aware large language models. *arXiv preprint arXiv:2308.13207*, 2023. 3
- [36] Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. Multimodal deep learning. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 689–696, 2011. 1
- [37] OpenAI. Gpt-4 technical report, 2023. 2, 3, 6
- [38] Garrick Orchard, Ajinkya Jayawant, Gregory K Cohen, and Nitish Thakor. Converting static image datasets to spiking neuromorphic datasets using saccades. *Frontiers in neuroscience*, 9:437, 2015. 2, 5
- [39] Karol J Piczak. Esc: Dataset for environmental sound classification. In *Proceedings of the 23rd ACM international conference on Multimedia*, pages 1015–1018, 2015. 5
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2, 3, 5, 6, 8
- [41] Justin Salamon, Christopher Jacoby, and Juan Pablo Bello. A dataset and taxonomy for urban sound research. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 1041–1044, 2014. 5
- [42] Alireza Salemi, Mahta Rafiee, and Hamed Zamani. Pre-training multi-modal dense retrievers for outside-knowledge visual question answering. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval*, pages 169–176, 2023. 3
- [43] Shubhra Kanti Karmaker Santu and Dongji Feng. Teler: A general taxonomy of llm prompts for benchmarking complex tasks. *arXiv preprint arXiv:2305.11430*, 2023. 3
- [44] Prithviraj Sen, Breno William Carvalho, Ibrahim Abdelaziz, Pavan Kapanipathi, Salim Roukos, and Alexander Gray. Logical neural networks for knowledge base completion with embeddings & rules. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3863–3875, 2022. 3
- [45] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 5
- [46] Hang Su, Wen Qi, Jiahao Chen, Chenguang Yang, Juan Sanudoval, and Med Amine Laribi. Recent advancements in multimodal human–robot interaction. *Frontiers in Neuro-robotics*, 17:1084000, 2023. 1
- [47] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*, 2019. 1
- [48] Weixuan Sun, Jiayi Zhang, Jianyuan Wang, Zheyuan Liu, Yiran Zhong, Tianpeng Feng, Yandong Guo, Yanhao Zhang, and Nick Barnes. Learning audio-visual source localization via false negative aware contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6420–6429, 2023. 2
- [49] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 2, 3, 6
- [50] Yikai Wang, Xinghao Chen, Lele Cao, Wenbing Huang, Fuchun Sun, and Yunhe Wang. Multimodal token fusion for vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12186–12195, 2022. 2
- [51] Yixuan Wei, Yue Cao, Zheng Zhang, Houwen Peng, Zhu-liang Yao, Zhenda Xie, Han Hu, and Baining Guo. iclip:

- Bridging image classification and contrastive language-image pre-training for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2776–2786, 2023. 2
- [52] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. 5
- [53] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5288–5296, 2016. 5
- [54] Hongwei Xue, Yuchong Sun, Bei Liu, Jianlong Fu, Ruihua Song, Houqiang Li, and Jiebo Luo. Clip-vip: Adapting pre-trained image-text model to video-language representation alignment. *arXiv preprint arXiv:2209.06430*, 2022. 2
- [55] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Bin Xiao, Ce Liu, Lu Yuan, and Jianfeng Gao. Unified contrastive learning in image-text-label space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19163–19173, 2022. 3
- [56] Qian Yang, Qian Chen, Wen Wang, Baotian Hu, and Min Zhang. Enhancing multi-modal and multi-hop question answering via structured knowledge and unified retrieval-generation. *arXiv preprint arXiv:2212.08632*, 2022. 3
- [57] Jiaming Zhang, Huayao Liu, Kailun Yang, Xinxin Hu, Ruiping Liu, and Rainer Stiefelhagen. Cmx: Cross-modal fusion for rgb-x semantic segmentation with transformers. *IEEE Transactions on Intelligent Transportation Systems*, 2023. 2
- [58] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5579–5588, 2021. 2
- [59] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8552–8562, 2022. 1, 2, 5, 6
- [60] Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*, 2023. 2, 6
- [61] Yiyuan Zhang, Kaixiong Gong, Kaipeng Zhang, Hongsheng Li, Yu Qiao, Wanli Ouyang, and Xiangyu Yue. Meta-transformer: A unified framework for multimodal learning. *arXiv preprint arXiv:2307.10802*, 2023. 2, 6
- [62] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023. 3
- [63] Jiangbin Zheng, Yile Wang, Cheng Tan, Siyuan Li, Ge Wang, Jun Xia, Yidong Chen, and Stan Z Li. Cvt-slr: Contrastive visual-textual transformation for sign language recognition with variational alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23141–23150, 2023. 2
- [64] Xu Zheng, Yexin Liu, Yunfan Lu, Tongyan Hua, Tianbo Pan, Weiming Zhang, Dacheng Tao, and Lin Wang. Deep learning for event-based vision: A comprehensive survey and benchmarks. *arXiv preprint arXiv:2302.08890*, 2023. 2
- [65] Jiachou Zhou, Xu Zheng, Yuanhuiyi Lyu, and Lin Wang. E-clip: Towards label-efficient event-based open-world understanding by clip. *arXiv preprint arXiv:2308.03135*, 2023. 2, 5, 6, 7, 8
- [66] Xiangyang Zhu, Renrui Zhang, Bowei He, Ziyao Zeng, Shanghang Zhang, and Peng Gao. Pointclip v2: Adapting clip for powerful 3d open-world learning. *arXiv preprint arXiv:2211.11682*, 2022. 2, 3, 5, 8