Alt-MoE:Multimodal Alignment via Alternating Optimization of Multi-directional MoE with Unimodal Models

Hongyang Lei[1], Xiaolong Cheng[1], Dan Wang[1], Qi Qin[2], Huazhen Huang[3],
Yetao Wu[1], Qingqing Gu[1], Zhonglin Jiang[1], Yong Chen[1], Luo Ji[1]*

[1] Geely Automobile Research Institute (Ningbo) Co., Ltd
[2]Peking University
[3]Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences
`https://github.com/HongyangLL/Alt-MoE`

Abstract

Recent Large Multi-Modal Models (LMMs) have made significant advancements in multi-modal alignment by employing lightweight connection modules to facilitate the representation and fusion of knowledge from existing pre-trained uni-modal models. However, these methods still rely on modality-specific and direction-specific connectors, leading to compartmentalized knowledge representations and reduced computational efficiency, which limits the model's ability to form unified multi-modal representations. To address these issues, we introduce a novel training framework, Alt-MoE, which employs the Mixture of Experts (MoE) as a unified multi-directional connector across modalities, and employs a multi-step sequential alternating unidirectional alignment strategy, which converges to bidirectional alignment over iterations. The extensive empirical studies revealed the following key points: 1) Alt-MoE achieves competitive results by integrating diverse knowledge representations from uni-modal models. This approach seamlessly fuses the specialized expertise of existing high-performance uni-modal models, effectively synthesizing their domain-specific knowledge into a cohesive multi-modal representation. 2) Alt-MoE efficiently scales to new tasks and modalities without altering its model architecture or training strategy. Furthermore, Alt-MoE operates in latent space, supporting vector pre-storage and real-time retrieval via lightweight multi-directional MoE, thereby facilitating massive data processing. Our methodology has been validated on several well-performing uni-modal models (LLAMA3, Qwen2, and DINOv2). achieving competitive results on a wide range of downstream tasks and datasets.

1 Introduction

Human perception is inherently multi-modal, seamlessly integrating diverse sensory inputs from vision, hearing, touch, and other senses to comprehend the world. Inspired by this capability, multi-modal learning aims to develop Artificial Intelligence (AI) systems that can process and interpret multiple types of input simultaneously, thereby mimicking human-like cognition. multi-modal large-scale models can process and integrate information from multiple modalities such as text, images, audio and video, and have become an important way to solve complex tasks involving heterogeneous data sources and achieve general artificial intelligence Alayrac et al. [2022], Radford et al. [2021], Wang et al. [2022a, 2023], Gao et al. [2024].

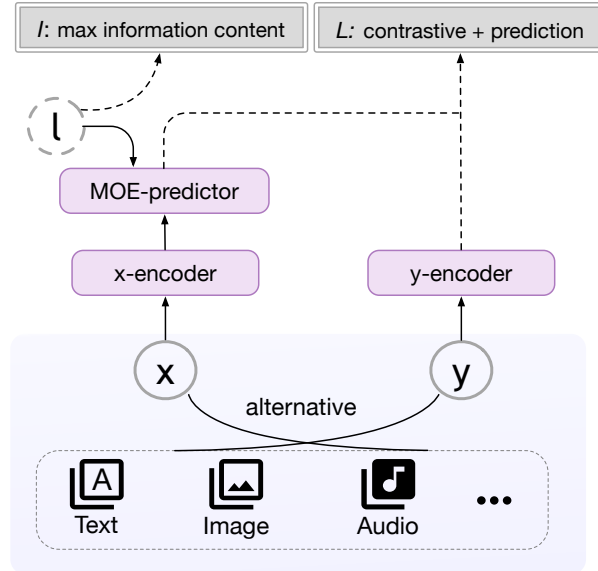*Corresponding author: `Luo.Ji1@geely.com`

Figure 1: Paradigm of our Alt-MoE joint embedding predictive architecture. The self-supervised learning is conducted with two encoding branches of input and output signals, and a multi-directional MoE predictor to match the input latent embedding into the target latent embedding by minimizing contrastive and prediction losses. The MoE predictor is conditioned on the multi-modal routed information which is learned with information entropy minimization. Training is performed with alternative any-to-any multi-modality tasks.

Modern multi-modal studies are mostly built on self-supervised learning (SSL) and generative architecture, while can be classified into two main categories. The first category aim to learn the multi-modal information by training from scratch Wang et al. [2022a, 2023], Oquab et al. [2023]. Such methodology asks for large-scale data therefore faces high training costs. It may also limit the acquisition of intra-modal information and introduces inter-modality information bias. For example, vision-language pre-training studies might be predominantly text-guided Radford et al. [2021], which hinders the retainment of visual information due to inherent limitations of text descriptions Oquab et al. [2023]. Another category either employs the pretrained Large Language Model (LLMs) as the backbone Wu et al. [2023], Haotian Liu [2023], Jun Zhan [2024] or a lightweight cross-modal connector Alayrac et al. [2022], Junnan Li [2023], Bai et al. [2023] in order to preserve fine-grained uni-modality information and alleviate computational burdens Zhang et al. [2024]. Success of such methodologies indicates the "platonic representation" Huh et al. [2024], which indicates representations across various domains are increasingly converging to the same latent space. Such representation similarity spans different model architectures, training objectives, and even data modalities.

Typical generative architecture tries to reconstruct the target signal conditioned by the input signal directly, usually with a decoder structure Assran et al. [2023]. However, such approach is not good at eliminating irrelevant details, and may face difficulty when input details are not easily predictable, or representing complex uncertainties in continuous spaces Dawid and LeCun [2023]. It will also create bias given ambiguous semantic alignment of multi-modal samples. To take a step ahead on the representation convergence on multi-modalities, one may consider to reconstruct multi-modal signals in the latent-space, however, which might induce the representation collapse, especially when the latent variables contain excess information Dawid and LeCun [2023]. Such consideration calls for a different framework with better alignment in the latent space, and reduced information content of latent variables.

In this work, we propose a new multi-modal paradigm which studies the dependency of the unobserved part (y) on the observed part (x) in their embedding space, in which x and y may belong to different modalities or any combination of them. We employ the joint-embedding predictive architecture (JEPA) Dawid and LeCun [2023], Assran et al. [2023] to align the representation of different modalities into the same space. JEPA achieves the alignment by first encoding the observed and

unobserved parts, then use a predictor to convert one embedding into another embedding space, driven by the contrastive and predictive losses simultaneously. We implement the predictor by recap the lightweight cross-modal connector, which is designed as the Mixture-of-Experts (MoE) structure. We solve the framework by Alternating Gradient Descent (AGD) of different multi-modal tasks, similar with IMP Akbari et al. [2023]. To tackle the challenge of the multi-modal semantic discrepancy (paired cross-modal data may exhibit inconsistent information despite their intended correspondence), we decouple the cross-modal information into modality-independent and shared components through distinct experts. This multi-directional MoE is used to conduct multiple unidirectional alignment tasks and is solved by Alternating Gradient Descent (AGD) of different multi-modal tasks, similar with IMP Akbari et al. [2023]. The router information of MoE is the latent factor l which helps better alignment. To avoid the representation collapse, we solve the problem by minimizing the information content of predictor Dawid and LeCun [2023]. Our **Alt-MoE** paradigm is depicted by Figure 1.

To empirically validate these properties and explicitly evaluate the alignment effectiveness, we conducted comprehensive experiments across multiple multi-modal tasks, while keeping the model architecture and training strategy consistent. Experimental results demonstrate that our Alt-MoE easily achieves modality and task scaling while attaining competitive results compared to current state-of-the-art (SoTA) multi-modal studies. Furthermore, our approach is able to conduct large-scale on-line multi-modal retrieval tasks due to our unique architecture, while mainstream multi-modal models fail to do so. We summarize our contributions as follows:

1. We propose Alt-MoE, a novel modality-agnostic multi-modal alignment method that leverages multi-directional MoE as a cross-modal connector. Through alternating gradient descent, Alt-MoE alternates between each unidirectional alignment task and optimizing the multi-directional MoE, enabling efficient and scalable alignment across arbitrary modalities.

2. Alt-MoE achieves multi-modal alignment in the latent space, which avoids the modality information bias and facilitates large-scale online retrieval ability.

3. Unlike conventional direct alignment methods, Alt-MoE decomposes complex multi-modal tasks into tractable unidirectional sub-tasks, employing alternating optimization for global convergence. This approach offers a new theoretical framework for multi-modal research, potentially enhancing efficiency and scalability across diverse modalities.

4. We provide an information-theoretic explanation in the problem formulation, with extensive experimental results and theoretical analysis supporting its effectiveness in improving alignment and handling complex multi-modal interactions.

5. Our experimental results demonstrate that effective multi-modal alignment can be achieved by using lightweight connectors to link uni-modal models. Furthermore, Alt-MoE operates in latent space, which enables vector pre-storage and real-time retrieval, thereby facilitating efficient processing of large-scale data tasks.

The rest of the paper is organized as follows. The connection with previous works is first discussed in Section 2. Problem is formulated in Section 3. The methodology is stated in Section 4. Experiment results are summarized in Section 5. We also discuss current limitations and future plans in Section 6. Finally Section 7 concludes this paper. Project code and model checkpoints can be found in `https://github.com/HongyangLL/Alt-MoE`.

2 Related Works

2.1 End-to-end multi-modal Learning

Recently, end-to-end multi-modal models employing various architectures have achieved outstanding performance. These architectures can be broadly categorized into several main types: Dual-encoder Radford et al. [2021], Jia et al. [2021]. Fusion-encoder Li et al. [2020], Jia et al. [2021], Chen et al. [2020]. Combining dual encoder and fusion encoder architectures integrate specialized layers into multi-modal models to enable deep cross-modal interactions Li et al. [2021, 2022].

The majority of multi-modal learning method employ large-scale multi-modal datasets for end-to-end pre-training. However, as model scale continues to increase, several potential challenges emerge: Firstly, the pre-training process may incur prohibitively high computational costs. Secondly, these models often struggle to adapt to novel modalities or tasks without extensive retraining. Moreover,

most multi-modal approaches employ text-guided learning of visual concepts, which can limit the acquisition of fine-grained information inherent to each modality. A logical approach is to leverage existing pre-trained uni-modal foundation models Li et al. [2023], Oquab et al. [2023], Zhang et al. [2024].

Unlike these works, each step in Alt-MoE close to I-JEPA Assran et al. [2023]. I-JEPA learn to predict the embeddings of a signal y from a compatible signal x, using a predictor network that is conditioned on additional (possibly latent) variables z to facilitate prediction. However, due to the complexity of multi-modal alignment, relying solely on a predictive structure proves insufficient for effective alignment. To enhance cross-modal retrieval capabilities, Alt-MoE incorporates a contrastive structure into Alt-MoE framework. Furthermore, to achieve alternating alignment and decouple cross-modal information, we replace the predictor with a multi-directional MoE. Cross embeddings serve as the signal z, prompting the MoE to execute tasks across different modalities.

2.2 Multi-modal Learning with uni-modal Models

Recent trends in multi-modal learning have increasingly focused on integrating high-performance uni-modal models to achieve effective multi-modal capabilities. Flamingo Alayrac et al. [2022] integrates visual information into each layer of a frozen Large LLM through the use of cross-attention. BLIP-2 Li et al. [2023] introduces an additional vision-to-language adaptation module, Q-former, and proposes a two-stage training process to mitigate the challenges associated with learning vision-language alignment. However, both methods require substantial parameters and multi-modal data for cross-modal alignment.

Recent studies indicate a convergence of representations across modalities, providing evidence for the feasibility of developing advanced multi-modal models by connecting high-performance uni-modal models with lightweight parameters and data-efficient techniques Huh et al. [2024]. For instance, LLaVA Liu et al. [2024] achieved state-of-the-art performance by employing a two-layer multilayer perceptron (MLP). Similar architectures have subsequently proliferated across various domains Li et al. [2024], Zhang et al. [2024], Lin et al. [2024, 2023]. Alt-MoE further advances this concept by interconnecting diverse high-performance uni-modal models through a shared multi-directional MoE. We conducted extensive experiments focused on alignment efficacy, demonstrating the effectiveness of joint training across multiple modalities and directions.

2.3 Multi-modal learning with MoE and AGD

Prior studies have investigated AGD-based multi-modal multi-task alternating training, revealing that the integration of diverse modalities, tasks, and resolutions can yield mutual benefits, thereby effectively enhancing the model's generalization capabilities and cross-domain performance Akbari et al. [2023], Likhosherstov et al. [2021]. we further extend this approach to integrate existing pre-trained high-performance uni-modal models, achieving overall alignment through alternating bidirectional alignment.

MoE-LLaVA Lin et al. [2024] proposes MoE-Tuning, a strategy for Large Vision-Language Models that creates a sparse model with constant computational cost. While both MoE-LLaVA and Alt-MoE employ sparse MoE to connect high-performance uni-modal large models, Alt-MoE not only scales this approach to large audio-visual-language models but also achieves pairwise bidirectional alignment across modalities.

3 Problem Formulation

In this section, we will first introduce the overall optimization objective of multi-modal bidirectional alignment, followed by the optimization objective of unidirectional alignment, and finally the alignment loss and theoretical explanation. It is worth noting that by integrating uni-modal models, Alt-MoE can effectively align multiple modalities. To clearly articulate the design rationale, we illustrate this with the example of image-text alignment.

1. Multi-modal alignment objective: Alt-MoE decomposes multi-modal alignment into multiple unidirectional alignment subtasks, where MoE serves as a knowledge representation and

fusion module and achieve multi-modal alignment by alternately aligning from image to text and text to image.

2. Alternating unidirectional alignment: Alt-MoE maximizes mutual information and minimizes conditional entropy to obtain independent and shared information for unidirectional alignment. Specifically, Alt-MoE adds embeddings as prior information to the input representations to guide the MoE in performing different pre-training tasks.

3. We provide an information-theoretic explanation to elucidate the rationale behind the design of the training objective.

Next, we will describe the optimization process at two levels, including the optimization objectives and parameter updates.

3.1 Multi-modal Alignment objective:

In this section, we will introduce the decomposition of the multi-modal bidirectional alignment objective. multi-modal alignment aims to align diverse modalities in a latent space by finding optimal parameters θ that minimize an alignment loss $\mathcal{L}_{align}$. This can be formulated as:

$$\theta^* = \arg\min_\theta \mathcal{L}_{\text{align}}(\theta). \tag{1}$$

By combining MoE $f_\theta(\cdot)$ parameterized by θ and AGD, which can decompose the optimization objectives for multi-modal alignment $\mathcal{L}_{align}$ into multiple unidirectional alignment subtasks, and then alternately execute each subtask at various time step to achieve overall alignment. Specifically, the image-text alignment can be decomposed as follows in the following Equation 2:

$$\mathcal{L}_{i-t}(\theta) = \mathcal{L}_{i\to t}(\theta_{i\to t}) \oplus \mathcal{L}_{t\to i}(\theta_{t\to i}),$$
$$\theta = \theta_{i\to t} \cup \theta_{t\to i}, \tag{2}$$

where $\mathcal{L}_{i\to t}, \theta_{i\to t}$ represents the image-to-text $(t \to i)$ alignment objective and parameter subset, $\mathcal{L}_{t\to i}, \theta_{t\to i}$ represents the text-to-image $(i \to t)$ alignment objective and parameter subset, $\oplus$ denotes an alternating optimization operation at various time step. Therefore, we decompose the image-text alignment into multiple unidirectional alignment optimizations and parameter subsets.

3.2 Alternating Unidirectional Alignment

In this section, we delineate the optimization objectives and parameter update procedures across various time steps t. By introducing AGD, we can alternately optimize unidirectional alignment at each time step t with the goals of maximizing mutual information between image I and text T, denoted as $I(I;T)$, and minimizing conditional entropies, denoted as $H(T|I)$ and $H(I|T)$.

Given time step t, the objective function is updated as Equation 3:

$$\mathcal{L}_{i-t} = \begin{cases} \mathcal{L}_{i\to t}^{(t)} = -I(I;T) + H(T|I), & \text{if } t = 2k \\ \mathcal{L}_{t\to i}^{(t)} = -I(I;T) + H(I|T), & \text{if } t = 2k+1, \end{cases} \tag{3}$$

where k is a non-negative integer. Based on Equation 1 and 3, the overall optimization objective can be formulated as shown in Equation 4:

$$\theta^* = \arg\min_\theta \left(-I(I;T) + \lambda \left(H(T|I) + H(I|T) \right) \right), \tag{4}$$

where λ is a weight parameter.

Given a set of parameters θ, we alternate between image-to-text and text-to-image unidirectional alignment at different time steps. At different time steps, we update only a subset of the parameters: $\theta_{i\to t}$ for $i \to t$ alignment and $\theta_{t\to i}$ for $t \to i$ alignment. Ultimately, this process ensures that all parameters are updated, such that $\theta = \theta_{i\to t} \cup \theta_{t\to i}$. The overall update process can then be described by the following Equation 5:

$$\theta^{t+1} = \begin{cases} \theta^t - \eta \nabla_{\theta_{i\to t}^t} \mathcal{L}_{i\to t}^{(t)}, & \text{if } t = 2k \\ \theta^t - \eta \nabla_{\theta_{t\to i}^t} \mathcal{L}_{t\to i}^{(t)}, & \text{if } t = 2k+1, \end{cases} \tag{5}$$

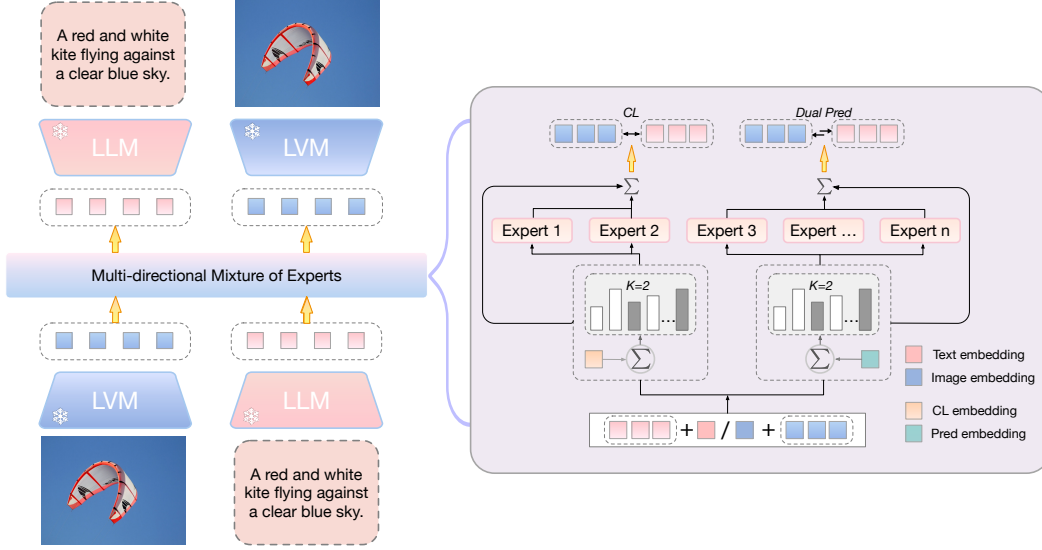where η is the learning rate, k is a non-negative integer.

Figure 2: Overview of Alt-MoE: Alt-MoE connects different modalities through multi-directional MoE and performs unidirectional alignment from image to text and text to image at different time steps, ultimately achieving overall bi-directional alignment. The specific structure of the multi-directional MoE is shown in the figure on the right. Different embeddings prompt the MoE to select different experts to perform contrastive learning or prediction.

3.3 Information Decomposition and Alignment

For image to text alignment, conditional entropy $H(T|I)$ and $H(I|T)$ represent modality-specific information in text and image respectively, measuring uncertainty in one modality after observing the other. Mutual information $I(T;I)$ quantifies shared information between image and text modalities, indicating how much knowing one reduces uncertainty about the other. For accurate image-text alignment, high mutual information (more shared content) and low conditional entropy (less modality-specific information) are desirable, ensuring strong semantic coupling between modalities.

Alt-MoE leverages the MoE router to automatically select different experts, optimizing for these two objectives. This approach helps decouple modality-specific information from shared information, potentially improving the balance between capturing unique modal features and cross-modal relationships.

4 Methodology

In this section, we will provide a detailed introduction to the architecture of Alt-MoE. As shown in Figure 2, Alt-MoE is divided into three modules: the visual model (VM), the language model (LM), and the fusion module MoE. We input the paired images and text (I, T) into VM and LM respectively to obtain latent representations z_i, z_t, and then perform multi-modal interaction in the fusion module.

4.1 Image and Text Encoding

Given a pair of image and text inputs (I, T), we employ separate encoders to process each modality. The visual model (VM) $f_v(\cdot)$ encode the image I, while the language model (LM) $f_l(\cdot)$ encode the text T. This process results in latent representations z_i and z_t for the image and text, respectively. The encoding can be formally expressed as:

$$z_i = f_v(I), \quad z_i \in \mathbb{R}^i,$$
$$z_t = f_l(T), \quad z_t \in \mathbb{R}^t, \tag{6}$$

where $z_i \in \mathbb{R}^i$ is the image representation, and $z_t \in \mathbb{R}^t$ is the text representation.

4.2 Unidirectional Alignment

At different time steps t, Alt-MoE performs unidirectional alignment tasks using different parameter updates and optimization objectives according to Equations 3 and 5. To guide MoE in selecting different experts for different modalities and tasks, we set up trainable modality encodings E_T, E_I and trainable task encodings E_{ce}, E_{mi}. Figure 3 illustrates the cross embedding process, where these embeddings are combined and added to z_t, z_i to guide the execution of different tasks. Specifically,

$$E_{ce} \longrightarrow E_I \longrightarrow E_{ce} + E_I, E_{mi} + E_I$$
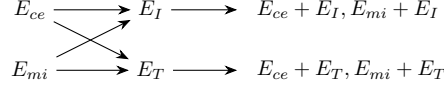$$E_{mi} \longrightarrow E_T \longrightarrow E_{ce} + E_T, E_{mi} + E_T$$

Figure 3: Cross Embedding.

we add modality embeddings to the input space z_i, z_t to enable the fusion module MoE to perceive the input modality, and then linearly project it into the d-dimensional common space z^c as follows:

$$z_i^c = W_i^1 \cdot z_i, \quad z_t^c = W_t^1 \cdot z_t, \quad z_i^c, z_t^c \in \mathbb{R}^d \tag{7}$$

where $W_i^1 \in \mathbb{R}^{d \times i}$ and $W_t^1 \in \mathbb{R}^{d \times t}$ are learnable projection matrices, d_i and d_t are the dimensions of the original image and text feature spaces respectively, and d is the dimension of the common space.

In the common latent space z^c with consistent dimensions, we update the loss function $\mathcal{L}_{i \to t}, \mathcal{L}_{t \to i}$ at different time steps t according to Equation 5. We add different task embeddings to z^c to enable the MoE $f_\theta(\cdot)$ to perceive tasks. Then, at different time steps t, we perform prediction and contrastive learning between image and text to maximize mutual information and minimize conditional entropy. Here, we take image-text alignment as an example; the same applies to text-audio alignment. This process alternates between image-to-text and text-to-image tasks as follows:

For image-to-text (at even time steps $2k$):

$$\mathcal{L}_{i \to t} = \mathcal{L}_{\text{CE}}(\hat{z}_t, z_t) + \mathcal{L}_{\text{MI}}(z_i^{\text{mi}}, z_t),$$
where:
$$\begin{aligned} \hat{z}_t &= W^1 \cdot f^1_{\theta_{i \to t}}(z_i^c + E_{\text{ce}}), \\ z_i^{\text{mi}} &= W^2 \cdot f^2_{\theta_{i \to t}}(z_i^c + E_{\text{mi}}), \\ f_{\theta_{i \to t}} &= f^1_{\theta_{i \to t}} \cup f^2_{\theta_{i \to t}}, \end{aligned} \tag{8}$$

where $W^1 \in \mathbb{R}^{d \times t}$, $W^2 \in \mathbb{R}^{d \times t}$ is a learnable projection matrix that map the output to dimensions consistent with z_t.

For text-to-image (at odd time steps $2k + 1$):

$$\mathcal{L}_{t \to i} = \mathcal{L}_{\text{CE}}(\hat{z}_i, z_i) + \mathcal{L}_{\text{MI}}(z_t^{\text{mi}}, z_i),$$
where:
$$\begin{aligned} \hat{z}_i &= W_t^2 \cdot f^1_{\theta_{t \to i}}(z_t^c + E_{\text{ce}}), \\ z_t^{\text{mi}} &= W_t^2 \cdot f^2_{\theta_{t \to i}}(z_t^c + E_{\text{mi}}), \\ f_{\theta_{t \to i}} &= f^1_{\theta_{t \to i}} \cup f^2_{\theta_{t \to i}}, \end{aligned} \tag{9}$$

where $W^3 \in \mathbb{R}^{d \times i}$, $W^4 \in \mathbb{R}^{d \times i}$ is a learnable projection matrix that map the output to dimensions consistent with z_i.

4.3 Loss Function

We adopt a unidirectional prediction approach at each time step to minimize conditional entropy. This method involves predicting either text features from image features or image features from text features, alternating between time steps. We define the prediction loss function using the $L2$ distance as follows:

$$\mathcal{L}_{\text{CE}} = |\hat{z}_t - z_t|_2^2 + |\hat{z}_i - z_i|_2^2, \tag{10}$$

where $|\cdot|_2^2$ denotes the squared $L2$ norm. $\hat{z}_t$ and $\hat{z}_i$ are the predicted text and image features, respectively, and z_t and z_i are the corresponding target features. At each time step, only one of these terms is active, depending on the prediction direction.

Here, we set z_t^{mi} as z^T and z_i^{mi} as z^I and contrastive loss can be formulated as follows:

$$\mathcal{L}_{\text{MI}} = \frac{1}{2N} \sum_{i=1}^{N} \left[-\log \frac{\exp(\text{sim}(z_i^I, z_i^T)/\tau)}{\sum_{j=1}^{N} \exp(\text{sim}(z_i^I, z_j^T)/\tau)} \right. \\ \left. -\log \frac{\exp(\text{sim}(z_i^T, z_i^I)/\tau)}{\sum_{j=1}^{N} \exp(\text{sim}(z_i^T, z_j^I)/\tau)} \right] \tag{11}$$

where:

- N is the number of image-text pairs in a batch.
- z_i^I and z_i^T are the latent representations of the i-th image and its corresponding text, respectively.
- $\text{sim}(z_i^I, z_j^T)$ is the cosine similarity between the latent representations z_i^I and z_j^T.
- τ is a temperature parameter that controls the sharpness of the similarity distribution.

5 Experiments and Results

We present the results of Alt-MoE on multiple modalities (audio, text, image) and various tasks and utilize Low-Rank Adaptation (LoRA) Hu et al. [2021] for fine-tuning the final layers of uni-modal models. First, we perform image-text retrieval on COCO Lin et al. [2014] and Flickr30K Plummer et al. [2015] by integrating LLMs such as LLAMA3 Dubey et al. [2024] and Qwen2 Oquab et al. [2023] with LVMs such as Dinov2 Oquab et al. [2023]. Then, we perform audio-text retrieval by integrating LLMs and LAMs such as Qwen2 Yang et al. [2024] and Whisper Radford et al. [2023].

5.1 Image-Text Retrieval

For the image-text retrieval task, since Alt-MoE connects two uni-modal models through a multi-directional MoE, we train the multi-directional MoE on the training sets of COCO and Flickr30K, and then test it on the test sets.

Table 1 shows the results of Alt-MoE on COCO and Flickr30K by integrating different LLMs such as LLAMA3-8b and Qwen2-7B with the LVMs Dinov2-Large. The results indicate that by integrating existing high-performance uni-modal models, Alt-MoE achieves SoTA performance.

Furthermore, by comparing different architectures of multi-modal models, Alt-MoE demonstrates superior data efficiency and parameter efficiency. Specifically, Alt-MoE has only 140M trainable parameters, which is significantly smaller than BLIP-2's 1.2B trainable parameters. Furthermore, BLIP-2 requires pre-training on 129M images before fine-tuning on COCO, whereas Alt-MoE is trained on COCO and Flickr30K.

In conclusion, the high efficiency in both parameters and training data further demonstrates that the representations are converging across modalities. This efficiency underscores the potential of a modality-agnostic multi-modal alignment strategy that can achieve alignment across various modalities using a lightweight model.

5.2 Audio-text Retrieval

As part of our ongoing research, we are currently exploring the potential of Alt-MoE in audio-text retrieval tasks. This extension aims to validate Alt-MoE's scalability to new tasks and modalities while maintaining its architecture and training strategy. In this work-in-progress, we are in the process of integrating existing high-performance Audio Models and Language Models, into the Alt-MoE framework. Our goal is to assess its performance on audio-text retrieval tasks. This investigation is

Model	# Trainable Params	Flickr30K (1K test set)						COCO Fine-tuned (5K test set)					
		Image → Text			Text → Image			Image → Text			Text → Image		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
Dual-encoder models													
CLIP Radford et al. [2021]	428M	88.0	98.7	99.4	68.7	90.6	95.2	-	-	-	-	-	-
ALIGN Cohen [1997]	820M	88.6	98.7	99.7	75.7	93.8	96.8	77.0	93.5	96.9	59.9	83.3	89.8
FILIP Yao et al. [2021]	417M	89.8	99.2	99.8	75.0	93.4	96.3	78.9	94.4	97.4	61.2	84.3	90.6
Florence Yuan et al. [2021]	893M	90.9	99.1	-	76.7	93.6	-	81.8	95.2	-	63.2	85.7	-
BEIT-3 Wang et al. [2022b]	1.9B	94.9	99.9	100.0	81.5	95.6	97.8	84.8	96.5	98.3	67.2	87.7	92.8
Fusion-encoder models													
UNITER Chen et al. [2020]	303M	83.6	95.7	97.7	68.7	89.2	93.9	65.7	88.6	93.8	52.9	79.9	88.0
OSCAR Li et al. [2020]	345M	-	-	-	-	-	-	70.0	91.1	95.5	54.0	80.8	88.5
VinVL Zhang et al. [2021]	345M	-	-	-	-	-	-	75.4	92.9	96.2	58.8	83.5	90.3
Dual encoder + Fusion encoder													
ALBEF Li et al. [2021]	233M	94.1	99.5	99.7	82.8	96.3	98.1	77.6	94.3	97.2	60.7	84.3	90.5
BLIP Li et al. [2022]	446M	97.1	100.0	100.0	86.7	97.3	98.7	82.4	95.4	97.9	65.1	86.3	91.8
BLIP-2 ViT-L Li et al. [2023]	474M	96.9	100.0	100.0	88.6	97.6	98.9	83.5	96.0	98.0	66.3	86.5	91.8
BLIP-2 ViT-g Li et al. [2023]	1.2B	97.6	100.0	100.0	89.7	98.1	98.9	85.4	97.0	98.5	68.3	87.7	92.6
LLMs + LVMs													
Alt-MoE (LLaMA3-8b, DinoV2-L)	140M	97.8	**100.0**	**100.0**	**97.8**	**100.0**	**100.0**	87.7	**99.6**	**99.9**	89.7	**99.7**	**99.9**
Alt-MoE (Qwen2-7b, DinoV2-L)	130M	**97.9**	100.0	100.0	97.8	100.0	100.0	**88.1**	99.4	99.8	**90.1**	99.6	99.9

Table 1: Comparison with state-of-the-art image-text retrieval methods, finetuned on COCO and Flickr30K.

still in its early stages, and experiments are currently underway. We anticipate that this study will provide insights into Alt-MoE's ability to generalize to new modality, tasks, and datasets. However, as the research is ongoing, concrete results are not yet available. We expect to compile and analyze the experimental results in the coming months. Once completed, comprehensive findings and their implications will be included in future updates of this paper.

5.3 Validation of Alignment Objectives

Model	Pred	CL	Image → Text			Text → Image		
			R@1	R@5	R@10	R@1	R@5	R@10
Alt-MoE	✓	✗	44.2	86.2	95.6	2.6	12.6	25.1
Alt-MoE	✓	✓	66.1	91.9	96.4	66.2	89.6	99.0

Table 2: Performance comparison of prediction alignment and joint training objectives: Image-Text retrieval with Alt-MoE (Qwen-7b, DinoV2-L) on COCO fine-tuning. CL denotes Contrastive Learning. Pred denotes prediction.

Alt-MoE utilizes the MoE router to automatically select different experts, optimizing for both objectives. This approach helps decouple modality-specific information from shared information, potentially improving the balance between capturing unique modal features and cross-modal relationships.

To validate this conclusion, we applied the predicted representations to retrieval tasks. The results in Table 2 indicate that training solely on the prediction alignment objective yields significantly lower performance compared to jointly training on both prediction and contrastive learning objectives.

By jointly optimizing for prediction and contrastive learning, the model maximizes the mutual information, while minimizing the conditional entropy. This ensures that Alt-MoE capture the most relevant and informative features, leading to better alignment and retrieval performance. The mutual reinforcement of the two tasks helps in effectively balancing the trade-off between capturing modality-specific features and learning cross-modal relationships.

5.4 Sensitivity Analysis

Figure 4 shows the sensitivity analysis of loss weights on average recall 1. The figure compares the average recall rates of contrastive learning and prediction models across different loss weights (0.4, 0.6, 0.8, 1.0). Both models achieve the highest recall at a loss weight of 1.0.
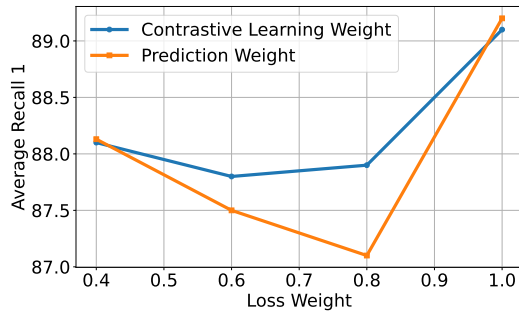
Figure 4: Sensitivity Analysis of Loss Weights

Model	MLP	MoE	ALT	Image → Text			Text → Image		
				R@1	R@5	R@10	R@1	R@5	R@10
Alt-MoE	✓	×	✓	74.4	86.0	92.2	82.3	89.5	92.6
Alt-MoE	×	✓	×	68.2	68.7	81.1	74.2	88.7	92.4
Alt-MoE	×	✓	✓	88.1	99.4	99.8	90.1	99.6	99.9

Table 3: Ablation Study of Image-Text Retrieval with the Alt-MoE contrastive learning representations (Qwen-7b, DinoV2-L) on COCO Fine-Tuning.

5.5 Ablation Study

Our model achieves excellent performance through the use of MoE. To validate its effectiveness, we conducted ablation experiments by replacing MoE with MLP and comparing the results. Additionally, we examined the impact of alternating optimization versus non-alternating optimization.

The results of the ablation study, as shown in Table 3, indicate that replacing MoE with MLP leads to a significant drop in performance, demonstrating the critical role of MoE in our model's success. Furthermore, the comparison between alternating optimization and non-alternating optimization reveals that alternating optimization contributes to better model performance, highlighting its importance in the training process.

In summary, the ablation experiments confirm that both the MoE architecture and the alternating optimization strategy are essential components for achieving SoTA performance in our model.

6 Current Limitation

In this work, we propose a new paradigm for multi-modal alignment, which has been applied on several multi-modality matching tasks and achieved state-of-the-art on retrieval metrics. This approach temporarily has not been applied on generative missions. However, by appending the multi-modal connector into an off-the-shelf decoder, we expect this framework can also help improve generative performance.

Alt-MoE is a typical self-supervised learning framework which can theoretically help build the world model. The input and output signals (x and y in Figure 1) can be any observed and unobserved sample parts of real world information, spanning over arbitrary combination of modalities, as well as different time and space slots. Among currently studied tasks, each of x and y is a uni-modality sample, and we focus on the cross-modality matching problem. In the future, we will study a more general problem and aim to provide a general framework for self-supervised world modeling.

7 Conclusion

In this study, we introduced Alt-MoE, a novel modality-agnostic training strategy and architecture designed for multi-modal learning. We implement Alt-MoE based on the joint-embedding predictive architecture, to achieve the multi-modal alignment in the latent-variable space. Alt-MoE effectively

integrates high-performance uni-modal models using lightweight connection modules, facilitating the alignment of modality pairs in multiple directions and enabling generalization to new tasks and modalities. Application of lightweight multi-directional MoE ensures both training and data efficiency. To validate the modality and task scalability of Alt-MoE, we conducted extensive experiments focused on alignment performance. The experimental results demonstrate that Alt-MoE can easily generalize to new modalities, tasks, and datasets while maintaining the same training strategy and architecture. Furthermore, Alt-MoE offered a simple, efficient, and scalable solution for multi-modal alignment and large-scale retrieval. By capitalizing on the strengths of existing high-performance uni-modal models, Alt-MoE provides a practical approach to achieving state-of-the-art performance across various tasks and datasets, underscoring its generalizability and effectiveness.

References

Hassan Akbari, Dan Kondratyuk, Yin Cui, Rachel Hornung, Huisheng Wang, and Hartwig Adam. Alternating gradient descent and mixture-of-experts for integrated multimodal perception. *Advances in Neural Information Processing Systems*, 36:79142–79154, 2023.

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022.

Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15619–15629, 2023.

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023.

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.

Gerson H Cohen. Align: a program to superimpose protein coordinates, accounting for insertions and deletions. *Journal of applied crystallography*, 30(6):1160–1161, 1997.

Anna Dawid and Yann LeCun. Introduction to latent variable energy-based models: A path towards autonomous machine intelligence, 2023. URL https://arxiv.org/abs/2306.02572.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Peng Gao, Renrui Zhang, Chris Liu, Longtian Qiu, Siyuan Huang, Weifeng Lin, Shitian Zhao, Shijie Geng, Ziyi Lin, Peng Jin, et al. Sphinx-x: Scaling data and parameters for a family of multi-modal large language models. *arXiv preprint arXiv:2402.05935*, 2024.

Qingyang Wu Yong Jae Lee Haotian Liu, Chunyuan Li. Visual instruction tuning, 2023. URL https://arxiv.org/abs/2304.08485.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.

Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. The platonic representation hypothesis. *arXiv preprint arXiv:2405.07987*, 2024.

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

Jiasheng Ye Yunhua Zhou Dong Zhang Zhigeng Liu Xin Zhang Ruibin Yuan Ge Zhang Linyang Li Hang Yan Jie Fu Tao Gui Tianxiang Sun Yugang Jiang Xipeng Qiu Jun Zhan, Junqi Dai. Anygpt: Unified multimodal llm with discrete sequence modeling, 2024. URL `https://arxiv.org/abs/2402.12226`.

Silvio Savarese Steven Hoi Junnan Li, Dongxu Li. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models, 2023. URL `https://arxiv.org/abs/2301.12597`.

Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36, 2024.

Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021.

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.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023.

Xiujun 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.

Valerii Likhosherstov, Anurag Arnab, Krzysztof Choromanski, Mario Lucic, Yi Tay, Adrian Weller, and Mostafa Dehghani. Polyvit: Co-training vision transformers on images, videos and audio. *arXiv preprint arXiv:2111.12993*, 2021.

Bin Lin, Bin Zhu, Yang Ye, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*, 2023.

Bin Lin, Zhenyu Tang, Yang Ye, Jiaxi Cui, Bin Zhu, Peng Jin, Junwu Zhang, Munan Ning, and Li Yuan. Moe-llava: Mixture of experts for large vision-language models. *arXiv preprint arXiv:2401.15947*, 2024.

Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024.

Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.

Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649, 2015.

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.

Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, pages 28492–28518. PMLR, 2023.

Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *International conference on machine learning*, pages 23318–23340. PMLR, 2022a.

Peng Wang, Shijie Wang, Junyang Lin, Shuai Bai, Xiaohuan Zhou, Jingren Zhou, Xinggang Wang, and Chang Zhou. One-peace: Exploring one general representation model toward unlimited modalities. *arXiv preprint arXiv:2305.11172*, 2023.

Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022b.

Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. Next-gpt: Any-to-any multimodal llm. *arXiv preprint arXiv:2309.05519*, 2023.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.

Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. *arXiv preprint arXiv:2111.07783*, 2021.

Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021.

Duzhen Zhang, Yahan Yu, Chenxing Li, Jiahua Dong, Dan Su, Chenhui Chu, and Dong Yu. Mm-llms: Recent advances in multimodal large language models. *arXiv preprint arXiv:2401.13601*, 2024.

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.