

uCLIP: Parameter-Efficient Multilingual Extension of Vision-Language Models with Unpaired Data

Dahyun Chung^{*1}, Donghyun Shin^{*1}, Yujin Sung^{*1}, Seunggi Moon^{*1},
Jinwoo Jeon¹, Byung-Jun Lee¹

¹Korea University

{dahyun1016, dawnme, sung031128, moon44432, kevin04087, byungjunlee}@korea.ac.kr

Abstract

Contrastive Language-Image Pre-training (CLIP) has demonstrated strong generalization across a wide range of visual tasks by leveraging large-scale English-image pairs. However, its extension to low-resource languages remains limited due to the scarcity of high-quality multilingual image-text data. Existing multilingual vision-language models exhibit consistently low retrieval performance in underrepresented languages—including Czech, Finnish, Croatian, Hungarian, Romanian—on the Crossmodal-3600 (XM3600) benchmark. To address this, we propose a lightweight and data-efficient framework for multilingual vision-language alignment. Our approach requires no image-text pairs or text-text pairs and freezes both the pretrained image encoder and multilingual text encoder during training. Only a compact 1.7M-parameter projection module is trained, using a contrastive loss over English representations as semantic anchors. This minimal training setup enables robust multilingual alignment even for languages with limited supervision. Extensive evaluation across multiple multilingual retrieval benchmarks confirms the effectiveness of our method, showing significant gains in five underrepresented languages where existing models typically underperform. These findings highlight the effectiveness of our pivot-based, parameter-efficient alignment strategy for inclusive multimodal learning.

Code — <https://dinyudin203.github.io/uCLIP-project/>

Introduction

Contrastive vision-language models such as CLIP (Radford et al. 2021) have demonstrated strong generalization ability by training on large-scale English-image parallel datasets. However, their effectiveness does not easily transfer to low-resource languages, primarily due to the lack of high-quality multilingual parallel text-image datasets. This data imbalance limits the accessibility and fairness of vision-language technologies in multilingual settings, despite their increasing demand in areas like cross-lingual retrieval (Litschko et al. 2018), content recommendation (Liu, Zhang, and Gulla 2021), and digital media understanding (Liu et al. 2022).

^{*}These authors contributed equally.




Figure 1: Performance Comparison. We compare models by plotting the average image-to-text Recall@10 across five underrepresented languages in the XM3600 benchmark against the number of trainable parameters (in millions). Each marker shape indicates the type of supervision used during training: *circle* for models trained without image-text (multilingual or English) pairs (I-T) or multilingual-English text pairs (T-T), *square* for those trained with T-T pairs only, and *triangle* for models using both I-T and T-T pairs. Despite having only 1.7M parameters and no paired supervision, *uCLIP* achieves the highest average Recall@10, outperforming all baselines.

A common workaround involves translation-based pipelines (e.g., translating text before applying CLIP), which bypass the need for multilingual supervision. However, these approaches not only introduce additional latency but also often suffer from semantic drift on ambiguous or culturally specific expressions—particularly in low-resource languages. (Qiu et al. 2022) Multilingual vision-language models (VLMs) have emerged as an alternative or complement to translation-based pipelines, aligning multilingual text with images directly. However, even state-of-the-art multilingual VLMs continue to underperform on underrepresented languages, as shown in Figure 2 and Supplementary Material. Existing multilingual VLMs largely follow one of the two approaches. The first strategy adopts a frozen image encoder with a fine-tuned multilingual text encoder, e.g., AltCLIP (Chen et al. 2022), M-CLIP (Carlsson et al. 2022), and NLLB-CLIP (Visheratin 2023). Such methods preserve CLIP’s image encoder and substitute its English text encoder




Figure 2: **Average z-score of Recall@10 across languages.** We evaluate multilingual VLM performance on the XM3600 benchmark using Recall@10 from four baseline models: AltCLIP-18, SigLIP2, NLLB-CLIP, and M-CLIP. For each model, we compute the z-score per language, indicating how much its Recall@10 deviates from the model-wise mean. The final score is the average z-score across models. Languages highlighted in red represent the five low-resource languages we target (cs, fi, hr, hu, ro). Unsupported languages by our multilingual text encoder (e.g., bn, fil, mi, quz, sw, te) are excluded from our evaluation.

with a multilingual model (e.g., XLM-R, mBERT, NLLB). Training focuses solely on aligning multilingual text with the fixed visual representation, typically via translated or synthetic data with contrastive or distillation-based objectives. While this design retains the benefits of CLIP’s visual representations, it typically relies on image–text (multilingual or English) pairs (I-T) or multilingual–English text pairs (T-T) to align the new multilingual encoder with the visual space. For example, M-CLIP performs knowledge distillation from the CLIP text encoder to mBERT, requiring large-scale paired T-T corpora. AltCLIP aligns XLM-R to the CLIP text encoder using T-T pairs, and further fine-tunes it with I-T pairs to match the frozen image encoder. NLLB-CLIP fine-tunes the NLLB encoder with translated captions using contrastive loss against the CLIP image encoder, again depending on I-T paired supervision.

The second one involves joint training of both image and text encoders, e.g., SigLIP2 (Zhai et al. 2023). These models are trained from scratch or heavily fine-tuned using massive multilingual datasets (e.g., WebLI) and employ novel objectives such as sigmoid contrastive loss to encourage dense alignment. This design facilitates a more flexible multilingual-visual representation space and theoretically better captures diverse semantic structures across languages. However, it comes at the cost of massive computation, requires extensive resources, and struggles to generalize across languages with limited training data.

To address this limitation, we propose *uCLIP*, an efficient multilingual vision–language model that eliminates the need for paired supervision. Rather than relying on image–text or multilingual–English text pairs, *uCLIP* adopts a pivot-based strategy that uses English as an implicit bridge for multilingual alignment. This design trains only a lightweight projection module, enabling training efficiency and strong zero-shot performance.

Specifically, our method leverages the robust English–image alignment in CLIP and transfers this to other languages via a frozen multilingual text encoder. The architecture comprises a frozen CLIP image encoder, a frozen multilingual text encoder, and a compact projection head

with just **1.7M trainable parameters**. This design preserves the strong performance of CLIP achieved through large-scale pretraining, without requiring paired I–T or T–T data, or retraining large encoders. As a result, *uCLIP* achieves effective multilingual retrieval and classification, improving performance in underrepresented languages while significantly reducing training cost. Our contributions are as follows.

- We propose *uCLIP*, a novel multilingual vision-language alignment framework that eliminates the need for image–text (I–T) pairs—whether multilingual or English—and multilingual–English text (T–T) pairs during training, by leveraging English as a universal semantic anchor.
- Our method trains only a lightweight projection module, enabling training-efficient alignment with over 99% fewer trainable parameters than prior baselines.
- *uCLIP* demonstrates strong zero-shot performance on multiple text-image retrieval and classification benchmarks, while also achieving significantly lower inference latency, demonstrating both effectiveness and efficiency.

Given space limitations and the wide range of existing studies, we provide detailed discussions of related work in the Supplementary Material.

Proposed Method

In this section, we present our proposed method for aligning multilingual text and images using only projectors, with English serving as the semantic pivot. The approach constructs a shared embedding space through memory-based retrieval and optimization of an alignment loss. An overview of the method is shown in Figure 3, with a detailed explanation provided below.

Soft Embedding Retrieval

To enable multilingual alignment without requiring paired I-T, T-T datasets, we use English queries as semantic anchors to retrieve pseudo-aligned representations from mem-

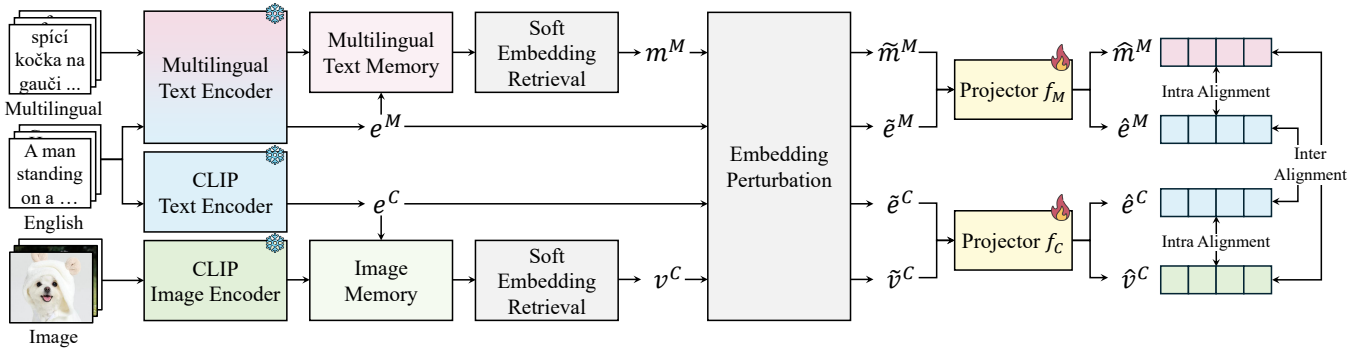


Figure 3: **Overall architecture.** We propose a lightweight alignment framework that bridges multilingual text and image embeddings via English, without requiring paired I-T and T-T data or encoder finetuning. *uCLIP* employs frozen encoders along with compact projection heads to map inputs into a shared embedding space. At inference time, only multilingual text encoder, image encoder and projectors are used. The model directly encodes multilingual text and image inputs using the frozen encoders, followed by projection into the shared space.

ory banks. Given an English query e_i , we extract two embeddings: e_i^C from the CLIP text encoder and e_i^M from a multilingual text encoder. We also construct two memory banks: image memory $V = \{v_1, \dots, v_N\}$ encoded by the CLIP image encoder, and multilingual text memory $M = \{m_1, \dots, m_K\}$ encoded by the same multilingual text encoder shared across target languages. Details on memory bank sizes, sampling strategies, and data sources are provided in the Supplementary Material.

To retrieve semantically aligned features, we compute softmax-weighted averages in each memory bank based on cosine similarity. The image representation aligned with the English query is retrieved as:

$$v_i^C = \sum_{k=1}^N \frac{\exp(\text{sim}(e_i^C, v_k)/\tau)}{\sum_{j=1}^N \exp(\text{sim}(e_i^C, v_j)/\tau)} \cdot v_k \quad (1)$$

Likewise, the multilingual text embedding aligned with the same query is:

$$m_i^M = \sum_{k=1}^K \frac{\exp(\text{sim}(e_i^M, m_k)/\tau)}{\sum_{j=1}^K \exp(\text{sim}(e_i^M, m_j)/\tau)} \cdot m_k \quad (2)$$

These soft-retrieved representations v_i^C and m_i^M serve as pseudo-aligned features that bridge modality through the shared semantics of the English query.

Embedding Perturbation

Multilingual text encoders often produce unstable or noisy representations, especially for low-resource languages. To enhance robustness against encoding biases, we add Gaussian noise to the four embeddings e^C, e^M, v^C, m^M , and project them onto the unit hypersphere:

$$\tilde{e}_i^C = \text{Norm}(e_i^C + \epsilon_1), \quad \tilde{e}_i^M = \text{Norm}(e_i^M + \epsilon_2) \quad (3)$$

$$\tilde{v}_i^C = \text{Norm}(v_i^C + \epsilon_3), \quad \tilde{m}_i^M = \text{Norm}(m_i^M + \epsilon_4) \quad (4)$$

where $\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4 \sim \mathcal{N}(0, \sigma^2 I)$. This perturbation encourages the model to align local neighbors of embeddings rather

than relying on exact vectors. As a result, it promotes local semantic smoothness and enhances robustness across modalities.

Inter Alignment

For each English query i , we first obtain four perturbed embeddings: \tilde{e}_i^C and \tilde{e}_i^M denote the representations of the same English query obtained from the CLIP text encoder and the multilingual text encoder, respectively; \tilde{v}_i^C and \tilde{m}_i^M are the image and multilingual text features retrieved from the memory banks using these text embeddings. To project these embeddings into a shared embedding space, we group them by their source modality. The CLIP text embedding (\tilde{e}_i^C) and the retrieved image feature (\tilde{v}_i^C) both originate from the image-text modality space associated with CLIP. Therefore, they are projected by a shared clip projection head, f_C . Similarly, the multilingual text embedding (\tilde{e}_i^M) and the retrieved multilingual text feature (\tilde{m}_i^M) belong to the multilingual text space and are thus projected by a shared multilingual projection head, f_M . We use two different projection heads to accommodate statistical differences between independently pretrained encoders, which improves alignment stability across modalities. This process is formulated as follows:

$$\begin{aligned} \hat{e}_i^C &= f_C(\tilde{e}_i^C), & \hat{e}_i^M &= f_M(\tilde{e}_i^M) \\ \hat{v}_i^C &= f_C(\tilde{v}_i^C), & \hat{m}_i^M &= f_M(\tilde{m}_i^M) \end{aligned}$$

The pair $(\hat{e}_i^C, \hat{e}_i^M)$, which represents the same English query encoded by two distinct encoders, provides a strong supervision signal to align textual semantics across encoders. On the other hand, the pair $(\hat{v}_i^C, \hat{m}_i^M)$, which associates retrieved image and multilingual text embeddings, serves as a pseudo-aligned cross-modal pair.

To enforce alignment between the embedding pairs from different encoders—either $(\hat{e}_i^C, \hat{e}_i^M)$ and $(\hat{v}_i^C, \hat{m}_i^M)$ —we apply contrastive losses to both pairs. The InfoNCE loss between query embedding q and key embedding k is defined

as :

$$\ell(q, k) = -\frac{1}{B} \sum_{i=1}^B \left[\log \frac{\exp(\text{sim}(q_i, k_i)/\tau)}{\sum_{j=1}^B \exp(\text{sim}(q_i, k_j)/\tau)} \right] \quad (5)$$

The similarity function $\text{sim}(\cdot, \cdot)$ denotes cosine similarity between two embeddings. Since softmax normalization is asymmetric, we both compute $\ell(\hat{e}^C, \hat{e}^M)$, $\ell(\hat{e}^M, \hat{e}^C)$ and take the average to ensure robust alignment between representations. The alignment loss between English text embeddings (\hat{e}^C, \hat{e}^M) is defined as:

$$\mathcal{L}_{\text{text}} := \frac{1}{2} \left[\ell(\hat{e}^C, \hat{e}^M) + \ell(\hat{e}^M, \hat{e}^C) \right]$$

The same formulation is used for the pseudo-aligned image and multilingual text embeddings (\hat{v}^C, \hat{m}^M):

$$\mathcal{L}_{\text{pseudo}} := \frac{1}{2} \left[\ell(\hat{v}^C, \hat{m}^M) + \ell(\hat{m}^M, \hat{v}^C) \right]$$

The total inter-alignment loss is defined as the sum of these two:

$$\mathcal{L}_{\text{inter}} := \mathcal{L}_{\text{text}} + \mathcal{L}_{\text{pseudo}}$$

The alignment between \hat{e}_i^C and \hat{e}_i^M , derived from the same English query, acts as a strong supervisory signal. This supervision helps guide the learning of the pseudo-aligned pair (\hat{v}_i^C, \hat{m}_i^M), facilitating robust cross-modal and cross-lingual representation learning. Notably, this design enables *uCLIP* to be trained without requiring any paired I-T and T-T data, as the English embedding alignment provides transferable supervision across modalities and languages. By anchoring alignment on high-resource English queries, *uCLIP* transfers robust supervision to underrepresented languages.

Intra Alignment

Semantically aligned embeddings—such as \hat{e}_i^C and \hat{v}_i^C , or \hat{e}_i^M and \hat{m}_i^M —often lie in disjoint regions due to modality gaps, especially when pretrained encoders are used without joint tuning. To address this, we introduce an intra-modality loss that pulls such embeddings closer, encouraging a cohesive and isotropic shared space that complements the inter-modality alignment. Inspired by Liang et al. (2022), we remove the repulsive term from contrastive loss and retain only the attractive component. In contrastive learning, the standard InfoNCE loss can be decomposed as:

$$\begin{aligned} & -\log \frac{\exp(\text{sim}(x_i, z_i)/\tau)}{\sum_{j=1}^N \exp(\text{sim}(x_i, z_j)/\tau)} \\ &= -\frac{1}{\tau} \text{sim}(x_i, z_i) + \log \sum_{j=1}^N \exp\left(\frac{\text{sim}(x_i, z_j)}{\tau}\right) \end{aligned}$$

where the first term attracts positive pairs and the second term repels negatives. Since all embeddings in our model are ℓ_2 -normalized, cosine similarity and squared Euclidean distance are functionally equivalent:

$$\|x - y\|_2^2 = 2(1 - x^\top y) = 2(1 - \text{sim}(x, y))$$

Thus, maximizing cosine similarity is equivalent to minimizing Euclidean distance in the normalized space. We leverage this equivalence and define the intra-alignment loss using the attractive term only:

$$\mathcal{L}_{\text{intra}} := \frac{1}{2B} \sum_{i=1}^B \left(\|\hat{e}_i^C - \hat{v}_i^C\|_2^2 + \|\hat{e}_i^M - \hat{m}_i^M\|_2^2 \right)$$

This encourages angular proximity and closes modality gaps within both CLIP and multilingual spaces. The intra-alignment complements inter-alignment, enabling robust text–image representation alignment via English intermediaries.

In summary, our final loss combines inter- and intra-alignment objectives:

$$\mathcal{L} := \mathcal{L}_{\text{inter}} + \lambda \mathcal{L}_{\text{intra}}.$$

The inter-alignment loss encourages cross-modal and cross-lingual consistency via English as a semantic pivot. This formulation enables *uCLIP* to learn a unified multilingual text–image space without relying on paired data in target languages.

Experiments

Implementation details

We adopt a frozen CLIP¹ model for the vision encoder and use an MPNet-base² model for the multilingual text encoder, both of which remain frozen during training. We provide further implementation details in Supplementary Material.

Multilingual Image–Text Retrieval

Setup We evaluate our approach on multilingual image-to-text retrieval, text-to-image retrieval, and zero-shot classification tasks across five low-resource languages. We assessed the bidirectional retrieval task on two multilingually translated benchmarks (MSCOCO (Lin et al. 2014) and Flickr30k (Young et al. 2014)) and XM3600 (Thapliyal et al. 2022). For image-to-text retrieval, a query image is embedded using the CLIP image encoder and projected via $\hat{v}_q = f_C(v_q)$. Multilingual candidate sentences $\hat{m}_1, \dots, \hat{m}_K$ are encoded with the multilingual text encoder and projected via $f_M(\cdot)$. Cosine similarities $\text{sim}(\hat{v}_q, \hat{m}_j)$ are computed to retrieve the top-matching captions.

Results Table 1 reports multilingual bidirectional image–text retrieval results across five low-resource languages (cs, fi, hr, hu, ro) and three benchmarks (MSCOCO, Flickr30k, XM3600). On image-to-text retrieval (I→T), *uCLIP* achieves average R@10 scores of 53.2%, 60.0%, and 71.8% on MSCOCO, Flickr30k, and XM3600 respectively, outperforming other baselines. Similarly, for text-to-image retrieval (T→I), it records 53.3%, 59.7%, and 70.4%, achieving the competitive score. Unlike baselines incorporate extensive multilingual pretraining or direct image–multilingual text or English–multilingual text supervision, *uCLIP* operates without any explicit paired supervision and

¹OpenCLIP ViT-B-32-datacomp_xl_s13b_b90k

²paraphrase-multilingual-mpnet-base-v2

Method	cs		fi		hr		hu		ro		Avg.	
	I→T	T→I	I→T	T→I	I→T	T→I	I→T	T→I	I→T	T→I	I→T	T→I
<i>MSCOCO (R@10)</i>												
AltCLIP-18	18.2	16.0	7.6	4.6	14.6	11.8	11.5	9.6	21.6	18.4	14.7	12.1
SigLIP2	<u>51.4</u>	46.7	30.6	24.6	40.0	35.2	<u>48.4</u>	43.3	<u>53.3</u>	47.2	<u>44.7</u>	39.4
NLLB-CLIP	28.6	31.3	10.1	8.7	23.0	25.1	22.4	23.2	40.1	42.9	24.8	26.2
M-CLIP	40.1	54.8	<u>31.8</u>	<u>48.1</u>	<u>42.7</u>	54.8	42.2	53.5	39.5	56.3	39.3	53.5
uCLIP (Ours)	54.9	<u>54.2</u>	50.0	49.9	53.6	<u>54.3</u>	52.2	<u>53.2</u>	55.4	<u>54.8</u>	53.2	<u>53.3</u>
<i>Flickr30k (R@10)</i>												
AltCLIP-18	18.5	15.0	8.0	3.7	14.8	10.1	11.5	10.2	25.2	19.5	15.6	11.7
SigLIP2	<u>55.1</u>	52.0	25.6	20.3	37.4	34.8	<u>44.8</u>	43.5	<u>49.2</u>	45.8	<u>42.4</u>	39.3
NLLB-CLIP	33.1	39.3	9.7	10.5	24.4	29.4	21.1	24.5	42.3	48.4	26.1	30.4
M-CLIP	46.6	71.7	<u>34.7</u>	63.0	<u>43.4</u>	72.0	39.7	69.7	43.2	71.9	41.5	69.7
uCLIP (Ours)	61.5	<u>61.3</u>	56.9	<u>56.1</u>	60.2	<u>59.9</u>	60.4	<u>59.6</u>	60.8	<u>61.4</u>	60.0	<u>59.7</u>
<i>XM3600 (R@10)</i>												
AltCLIP-18	23.1	20.2	12.9	8.7	22.6	14.9	17.8	16.3	35.2	27.3	22.3	17.5
SigLIP2	<u>59.2</u>	50.4	48.7	38.9	60.6	54.7	60.9	58.0	60.7	56.9	58.0	51.8
NLLB-CLIP	30.8	29.9	15.8	13.9	31.2	29.5	31.2	29.5	53.7	53.2	31.7	30.6
M-CLIP	52.1	66.8	<u>61.7</u>	74.7	<u>66.9</u>	83.9	<u>69.1</u>	80.9	<u>71.5</u>	81.0	<u>64.3</u>	77.5
uCLIP (Ours)	64.5	<u>62.8</u>	69.0	<u>67.4</u>	77.2	<u>76.2</u>	72.7	<u>71.6</u>	75.6	<u>73.9</u>	71.8	<u>70.4</u>

Table 1: **Bidirectional Image–Text retrieval results across five languages.** “cs, fi, hr, hu, ro” indicates the evaluation language. I→T denotes image-to-text retrieval, and T→I denotes text-to-image retrieval. “Avg.” means average score of the metric to its corresponding method. The table reports retrieval performance on translated MSCOCO, Flickr30k datasets, and existing multilingual dataset XM3600.

still achieves superior results—highlighting the effectiveness of its lightweight cross-modal alignment approach, using English as semantic pivot. These results underline the generalizability of *uCLIP*, which achieves competitive multilingual grounding with only 1.7M trainable parameters. Additional Recall@1, Recall@5 scores and qualitative retrieval samples are provided in Supplementary Material.

Zero-shot Classification

Setup We assess zero-shot classification on multilingually translated benchmarks: CIFAR-10 (Krizhevsky and Hinton 2009), and STL-10 (Coates, Ng, and Lee 2011). Each image is encoded via the CLIP encoder and projected using $f_C(\cdot)$; class names are encoded and projected via $f_M(\cdot)$. Classification is performed via cosine similarity, and predictions are based on the most similar class. We report average F1 scores to account for class imbalance.

Results *uCLIP* maintains strong class discrimination across languages without using any paired data, unlike baselines such as M-CLIP and SigLIP2, which depend on direct T-T or I-T supervision. For example, in CIFAR-10, *uCLIP* achieves 90.5 in Finnish and 91.9 in Croatian—matching or surpassing M-CLIP and SigLIP2 that rely on translated captions or massive multilingual text-image pairs, respectively. (e.g., 22B data in SigLIP2). While M-CLIP requires paired multilingual-English text for training, *uCLIP* is trained using unpaired datasets. This shows *uCLIP*’s ability to perform effective multilingual zero-shot classification, despite its lightweight design and lack of any direct supervision.

Method	cs	fi	hr	hu	ro	Avg.
<i>CIFAR-10 (F1)</i>						
AltCLIP-18	29.9	9.7	70.5	16.9	35.1	32.4
SigLIP2	90.9	34.8	91.9	91.6	89.2	79.7
NLLB-CLIP	17.2	12.1	25.4	27.8	28.4	22.2
M-CLIP	88.4	<u>74.5</u>	87.9	<u>89.0</u>	85.5	<u>85.1</u>
uCLIP (Ours)	<u>89.9</u>	90.5	91.9	82.3	<u>87.0</u>	88.3
<i>STL-10 (F1)</i>						
AltCLIP-18	44.1	14.0	74.3	19.8	46.1	39.7
SigLIP2	<u>93.4</u>	34.4	<u>95.5</u>	97.4	<u>93.9</u>	82.9
NLLB-CLIP	22.2	22.6	31.5	39.1	38.0	30.7
M-CLIP	96.6	85.9	96.2	<u>97.0</u>	95.6	94.3
uCLIP (Ours)	88.2	91.8	92.3	92.8	90.3	<u>91.1</u>

Table 2: **Zero-shot classification performance across five languages.** The table reports classification performance on translated CIFAR-10 and STL-10 datasets.

Efficiency Analysis

Train Efficiency As summarized in Table 3, *uCLIP* is remarkably lightweight, with only 1.7M trainable parameters and no requirement for paired I-T, T-T data. This minimal parameter usage not only reduces memory and compute overhead but also leads to faster convergence during training. This high training efficiency makes *uCLIP* particularly suitable for low-resource or constrained computing environments, such as deployment in academic, mobile, or edge scenarios.

Method	I-T Pairs	T-T Pairs	Trainable Params
AltCLIP-18	✓	✓	563M
SigLIP2	✓	✓	375M
NLLB-CLIP	✓	✓	501M
M-CLIP	×	✓	560M
uCLIP (Ours)	×	×	1.7M

Table 3: **Comparison of supervision and model scale.** We compare recent multilingual VLMs in terms of their reliance on I-T pairs and T-T pairs during training, as well as the number of trainable parameters. Unlike prior methods, which require extensive supervision and full encoder tuning, *uCLIP* is trained without any paired data and uses only a lightweight 1.7M parameter projection module.

Method	Translation (ms/sample)	Encoding (ms/sample)	Total (ms/sample)
CLIP + Translator	459.02	21.11	480.13
AltCLIP-18	×	60.42	60.42
SigLIP2	×	47.85	47.85
NLLB-CLIP	×	29.41	<u>29.41</u>
M-CLIP	×	94.4	94.4
uCLIP (Ours)	×	23.56	23.56

Table 4: **Inference time comparison.** We report average per-query inference time in milliseconds, including translation and encoding.

Inference-time Efficiency Following prior work such as mCLIP (Chen et al. 2023), which compares inference latency between translation-based and multilingual approaches, a common baseline for multilingual text-image retrieval is to first translate non-English queries into English and then apply a pre-trained CLIP model. However, as shown in Table 4, this strategy incurs substantial latency, requiring 480.13 ms per query. In contrast, *uCLIP* removes the need for translation and achieves a total inference time of just 23.56 ms—over 20 times faster than translation-based methods—while also being more efficient than other multilingual baselines. The reported inference time is the sum of average image encoding and text encoding time per sample, averaged across five languages. It excludes data loading and batching overheads outside the model pipeline. All measurements were conducted on a single NVIDIA TITAN Xp GPU.

Embedding Space Analysis

Cosine Similarity Alignment To evaluate multilingual text-image alignment, we visualize cosine similarity matrices between 50 image-text pairs from translated Flickr30k under five settings. Figure 4 presents the similarity heatmaps for five low-resource languages across various models. A non-scattered strong red diagonal pattern indicates accurate alignment between corresponding image and caption embeddings, while low similarity values (darker blue) in the off-diagonal regions suggest the model correctly suppresses mismatched pairs. Notably, *uCLIP* exhibits both the sharpest

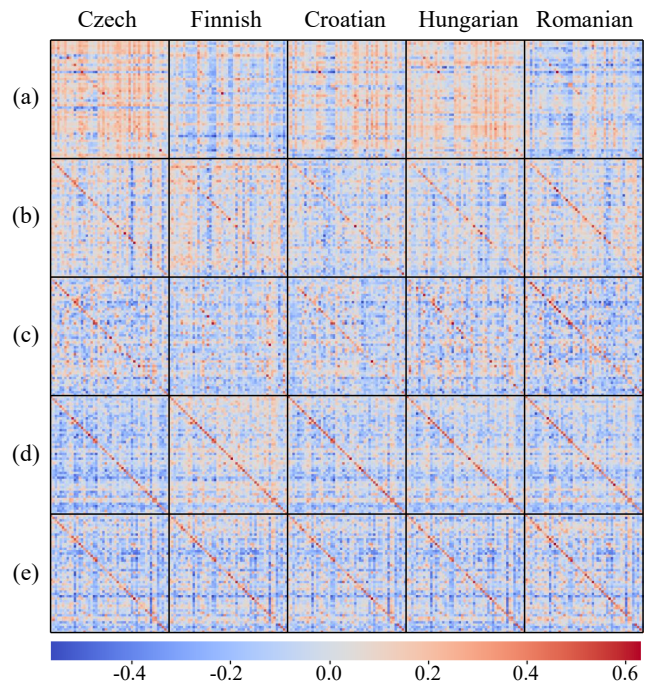


Figure 4: **Cosine similarity visualization for embeddings of text and image queries.** We visualize the cosine similarity matrices under five different settings: (a) AltCLIP-18, (b) SigLIP2, (c) NLLB-CLIP, (d) M-CLIP, and (e) using our proposed *uCLIP* model. All text samples are translated in each five language from Flickr30k benchmark.

diagonal and the darkest off-diagonal regions across all five languages, demonstrating its superior ability to capture correct multilingual alignment while avoiding false positives.

Visual Representation via UMAP To evaluate the robustness of our projection module across different image encoders, we visualize the image embeddings using UMAP on the CIFAR-10 dataset. Each point represents an image embedding colored by its ground-truth class. Compared to the raw CLIP and SigLIP2 embeddings, our *uCLIP* projection (b), (d) produces similarly well-separated or even more compact clusters, especially for visually ambiguous classes such as car and truck. Notably, when switching the image encoder from CLIP to SigLIP2, the class separability improves, and our projection (d) preserves this improvement without degradation. These results indicate that the proposed projection module generalizes well across different visual backbones and maintains discriminative structures in the embedding space, even without retraining for each encoder.

Ablation Studies

Ablation on Encoders

We evaluate several multilingual text encoders, multilingual versions of E5-base(Wang et al. 2024)³, MiniLM-L12(Wang

³multilingual-e5-base




Figure 5: **UMAP visualization of image embedding.** We visualize embeddings extracted from (a) CLIP image encoder, (b) *uCLIP* model with CLIP image encoder, (c) SigLIP2 image encoder, and (d) *uCLIP* model with SigLIP2 image encoder. Results are based on CIFAR-10.

Multilingual Encoder	Image \rightarrow Text			Text \rightarrow Image		
	R@1	R@5	R@10	R@1	R@5	R@10
E5-base	9.0	23.0	32.2	8.5	23.0	31.9
MiniLM-L12	17.1	37.6	48.9	17.6	38.8	50.6
MPNet-base	19.9	41.6	53.2	19.3	41.5	53.3

Table 5: **Effect of multilingual text encoder choice.** We conduct an ablation study comparing different multilingual text encoders. The results are reported on translated MSCOCO dataset.

et al. 2020)⁴, and MPNet-base (Song et al. 2020)⁵, which differ in model size and architectural characteristics. For simplicity, we refer to them as E5-base, MiniLM-L12, and MPNet-base throughout the paper. As shown in Table 5, MPNet-base achieves the best overall performance, followed by MiniLM-L12 and E5-base. Notably, MiniLM-L12 outperforms E5-base despite having a smaller model size.

To further evaluate the generalizability of our method to inherently multilingual vision-language models, we replace the vision encoder with SigLIP2. The results in Table 6 demonstrate that our approach enhances the performance of SigLIP2, indicating its effectiveness even when applied to native multilingual VLMs.

Ablation on Methodology

We leverage three loss terms to align the multilingual text encoder with the CLIP vision encoder: $\mathcal{L}_{\text{pseudo}}$, $\mathcal{L}_{\text{text}}$, and $\mathcal{L}_{\text{intra}}$. Note that $\mathcal{L}_{\text{pseudo}}$ and $\mathcal{L}_{\text{text}}$ together constitute the inter-alignment loss, $\mathcal{L}_{\text{inter}}$. We conduct an ablation study to examine the impact of each loss component on performance. The results in Table 7 show that $\mathcal{L}_{\text{intra}}$ and $\mathcal{L}_{\text{text}}$ have a relatively minor effect, while removing $\mathcal{L}_{\text{pseudo}}$ leads to a significant performance drop. This suggests that $\mathcal{L}_{\text{pseudo}}$ serves as the primary mechanism for directly aligning the multilingual text embedding space with the image embedding space, while $\mathcal{L}_{\text{text}}$ and $\mathcal{L}_{\text{intra}}$ play a supportive role by contributing to indirect alignment. Also, we adopt embedding perturbation to promote local semantic smoothness and improve robustness. As shown in Table 7, the absence of embedding

⁴paraphrase-multilingual-MiniLM-L12-v2

⁵paraphrase-multilingual-mpnet-base-v2

Method	Image \rightarrow Text			Text \rightarrow Image		
	R@1	R@5	R@10	R@1	R@5	R@10
SigLIP2	18.1	35.5	44.7	14.5	30.5	39.4
Ours (w/ SigLIP2)	20.7	42.9	54.6	21.3	43.5	55.1

Table 6: **Effect of our method on SigLIP2.** *uCLIP* significantly improves multilingual retrieval performance over the SigLIP2 baseline. The results are reported on translated MSCOCO dataset.

Method	Image \rightarrow Text			Text \rightarrow Image		
	R@1	R@5	R@10	R@1	R@5	R@10
Full	19.9	41.6	53.2	19.3	41.5	53.3
<i>Ablations</i>						
w/o $\mathcal{L}_{\text{intra}}$	19.1	40.6	52.5	18.8	40.8	53.0
w/o $\mathcal{L}_{\text{text}}$	18.1	39.8	51.2	18.2	40.0	51.9
w/o $\mathcal{L}_{\text{pseudo}}$	16.2	36.1	47.7	18.5	40.7	52.4
w/o E.P.	9.3	23.2	32.6	16.9	37.4	48.6

Table 7: **Effect of loss terms and embedding perturbation.** We selectively remove components in our model to evaluate their individual contribution to performance. ‘‘E.P.’’ indicates Embedding Perturbation. The results are reported on translated MSCOCO dataset.

perturbation substantially degrades performance. This suggests that, without perturbation, pseudo embeddings tend to align too directly with each other, rather than being aligned with the local semantic neighborhoods of their counterparts

Conclusion

Summary

We present *uCLIP*, a parameter-efficient framework that addresses the fundamental challenge of extending vision-language models to underrepresented languages. Our key contributions include: (1) a novel pivot-based alignment strategy that leverages English representations as semantic anchors to bridge pretrained vision and multilingual text encoders, eliminating the need for paired I-T, T-T datasets during training; and (2) a lightweight projector architecture with only 1.7M trainable parameters—over 99% smaller than existing multilingual models—while keeping both vision and text encoders completely frozen. The simplicity and effectiveness of our method make it a practical solution for democratizing vision-language capabilities across diverse linguistic communities with minimal computational resources and training overhead.

Future Work

A promising direction is to extend our approach to more languages using stronger multilingual encoders and cross-lingual transfer. We also plan to adapt the model to generation tasks such as captioning and VQA, enabling more complex cross-modal understanding.

Acknowledgments

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) under the Artificial Intelligence Star Fellowship support program to nurture the best talents (IITP-2025-RS-2025-02304828).

References

- Carlsson, F.; Eisen, P.; Rekathati, F.; and Sahlgren, M. 2022. Cross-lingual and Multilingual CLIP. In *Proceedings of the Language Resources and Evaluation Conference*, 6848–6854. Marseille, France: European Language Resources Association.
- Chen, G.; Hou, L.; Chen, Y.; Dai, W.; Shang, L.; Jiang, X.; Liu, Q.; Pan, J.; and Wang, W. 2023. mCLIP: Multilingual CLIP via Cross-lingual Transfer. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 13028–13043. Toronto, Canada: Association for Computational Linguistics.
- Chen, Z.; Liu, G.; Zhang, B.-W.; Ye, F.; Yang, Q.; and Wu, L. 2022. Altclip: Altering the language encoder in clip for extended language capabilities. *arXiv preprint arXiv:2211.06679*.
- Coates, A.; Ng, A.; and Lee, H. 2011. An Analysis of Single-Layer Networks in Unsupervised Feature Learning. In Gordon, G.; Dunson, D.; and Dudík, M., eds., *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, 215–223. Fort Lauderdale, FL, USA: PMLR.
- Krizhevsky, A.; and Hinton, G. 2009. Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario.
- Liang, V. W.; Zhang, Y.; Kwon, Y.; Yeung, S.; and Zou, J. Y. 2022. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. *Advances in Neural Information Processing Systems*, 35: 17612–17625.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Litschko, R.; Glavaš, G.; Ponzetto, S. P.; and Vulić, I. 2018. Unsupervised Cross-Lingual Information Retrieval Using Monolingual Data Only. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR '18*, 1253–1256. New York, NY, USA: Association for Computing Machinery. ISBN 9781450356572.
- Liu, H.; Guo, Y.; Yin, J.; Gao, Z.; and Nie, L. 2022. Review polarity-wise recommender. *IEEE Transactions on Neural Networks and Learning Systems*, 34(12): 10039–10050.
- Liu, P.; Zhang, L.; and Gulla, J. A. 2021. Multilingual Review-aware Deep Recommender System via Aspect-based Sentiment Analysis. *ACM*, 39(2).
- Qiu, C.; Oneată, D.; Bugliarello, E.; Frank, S.; and Elliott, D. 2022. Multilingual Multimodal Learning with Machine Translated Text. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, 4178–4193. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Song, K.; Tan, X.; Qin, T.; Lu, J.; and Liu, T.-Y. 2020. MP-Net: Masked and Permuted Pre-training for Language Understanding. In *Advances in Neural Information Processing Systems*, volume 33, 16857–16867.
- Thapliyal, A.; Pont-Tuset, J.; Chen, X.; and Soricut, R. 2022. Crossmodal-3600: A Massively Multilingual Multimodal Evaluation Dataset. In *EMNLP*.
- Visheratin, A. 2023. NLLB-CLIP—train performant multilingual image retrieval model on a budget. *arXiv preprint arXiv:2309.01859*.
- Wang, L.; Yang, N.; Huang, X.; Yang, L.; Majumder, R.; and Wei, F. 2024. Multilingual E5 Text Embeddings: A Technical Report. *arXiv:2402.05672*.
- Wang, W.; Wei, F.; Dong, L.; Bao, H.; Yang, N.; and Zhou, M. 2020. MINILM: deep self-attention distillation for task-agnostic compression of pre-trained transformers. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781713829546.
- Young, P.; Lai, A.; Hodosh, M.; and Hockenmaier, J. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the association for computational linguistics*, 2: 67–78.
- Zhai, X.; Mustafa, B.; Kolesnikov, A.; and Beyer, L. 2023. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11975–11986.