



Multilingual Vision-Language Pre-training for the Remote Sensing Domain

João Daniel Silva
INESC-ID

Instituto Superior Técnico, University of Lisbon
Lisbon, Portugal
joao.daniel.silva@tecnico.ulisboa.pt

Devis Tuia

EPFL ENAC IIE ECEO
École Polytechnique Fédérale de Lausanne
Sion, Switzerland
devis.tuia@epfl.ch

João Magalhães
NOVA-LINCS

Faculty of Science and Technology, NOVA University
Lisbon, Portugal
jmag@fct.unl.pt

Bruno Martins

INESC-ID & LUMILIS (Lisbon ELLIS Unit)
Instituto Superior Técnico, University of Lisbon
Lisbon, Portugal
bruno.g.martins@tecnico.ulisboa.pt

ABSTRACT

Methods based on Contrastive Language-Image Pre-training (CLIP) are nowadays extensively used in support of vision-and-language tasks involving remote sensing data, such as cross-modal retrieval. The adaptation of CLIP to this specific domain has relied on model fine-tuning with the standard contrastive objective, using existing human-labeled image-caption datasets, or using synthetic data corresponding to image-caption pairs derived from other annotations over remote sensing images (e.g., object classes). The use of different pre-training mechanisms has received less attention, and only a few exceptions have considered multilingual inputs. This work proposes a novel vision-and-language model for the remote sensing domain, exploring the fine-tuning of a multilingual CLIP model and testing the use of a self-supervised method based on aligning local and global representations from individual input images, together with the standard CLIP objective. Model training relied on assembling pre-existing datasets of remote sensing images paired with English captions, followed by the use of automated machine translation into nine additional languages. We show that translated data is indeed helpful, e.g. improving performance also on English. Our resulting model, which we named **Remote Sensing Multilingual CLIP (RS-M-CLIP)**, obtains state-of-the-art results in a variety of vision-and-language tasks, including cross-modal and multilingual image-text retrieval, or zero-shot image classification.

CCS CONCEPTS

• **Computing methodologies** → **Computer vision; Information extraction; Machine learning; Natural language processing**; • **Computer systems organization** → **Neural networks**.

KEYWORDS

Remote Sensing, Vision and Language, Cross-Modal Retrieval, Contrastive Language-Image Pre-training, Self-Supervised Pre-training

ACM Reference Format:

João Daniel Silva, João Magalhães, Devis Tuia, and Bruno Martins. 2024. Multilingual Vision-Language Pre-training for the Remote Sensing Domain. In *The 32nd ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '24)*, October 29–November 1, 2024, Atlanta, GA, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3678717.3691318>

1 INTRODUCTION

Contrastive Language-Image Pre-training (CLIP) was an important contribution in the area of multimodal machine learning [41]. In brief, CLIP models are based on an architecture that features two separate encoders, respectively for visual and textual inputs, trained with a contrastive objective that promotes similar semantics and concepts, in both modalities, to be aligned in the embedding space.

CLIP models are nowadays extensively used in a variety of vision-and-language tasks. These include zero-shot classification leveraging class descriptions [10, 14, 41], cross-modal retrieval [2, 28, 41, 52], or object detection/segmentation from text descriptions [18, 49, 71, 72, 74]. CLIP is also commonly used as a vision encoder within methods that leverage Large Language Models (LLMs) for supporting multimodal interactions [4, 5, 23, 29, 69].

The successful application to different tasks has also led to the adaptation of CLIP for specific domains, such as medical imaging [15, 52] or remote sensing [26, 28, 37, 42, 46, 67]. This last case is motivated by the increasing amount of Earth Observation (EO) images being collected through satellites and other aerial vehicles, and also by the interest in methods that can support natural language interactions with large EO image repositories. CLIP models are a natural choice to support vision-and-language tasks over these data, given their good performance in tasks such as cross-modal retrieval, while at the same time avoiding highly specialized mechanisms to extract effective representations from the images.

It should nonetheless be noted that CLIP models typically involve a large number of parameters, thus requiring large datasets for training. For context, the CLIP model associated to the original OpenAI proposal was trained with 400 million image-caption pairs,



This work is licensed under a Creative Commons Attribution International 4.0 License.

SIGSPATIAL '24, October 29–November 1, 2024, Atlanta, GA, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1107-7/24/10.

<https://doi.org/10.1145/3678717.3691318>

and other open-sourced variants used billion-scale datasets crawled from the Web. In comparison, despite the high availability of remote sensing imagery, the size of current image-text datasets specific to this domain is much smaller. Aggregating several popular remote sensing datasets leads to a total of approximately 200k image-text pairs (i.e., considering NWPU-Captions [9], RSITMD [60], RSICD [30], UCM-Captions [40], and Sydney-Captions [40]). Some studies have advanced pipelines to scale the number of image-caption pairs [28], for instance automatically transforming the annotations associated to other datasets originally designed for semantic segmentation or object detection. The current state-of-the-art model for the remote sensing domain corresponds to RemoteCLIP [28], which was developed by fine-tuning a large CLIP model with more than 800k synthetic image-caption pairs.

Going beyond the current state-of-the-art, this work proposes a novel vision-and-language model for the remote sensing domain, exploring the fine-tuning of a multilingual CLIP model originally trained on generalist Web data. A large dataset was created by assembling pre-existing remote sensing images paired with English captions, followed by the use of automated machine translation into nine additional languages, as a form of data augmentation [59]. Using these data, the fine-tuning process employed a self-supervised method based on aligning local and global representations from individual input images, together with the standard CLIP objective.

Through extensive experiments, we show that the use of translated data is indeed helpful, e.g. improving performance on English while enabling multilingual applications. Our resulting model, which we named Remote Sensing Multilingual CLIP (RS-M-CLIP), obtains state-of-the-art results in a variety of tasks, including cross-modal and multilingual image-text retrieval, or zero-shot image classification. RS-M-CLIP is also capable of producing detailed local features, which hold promise for the application in tasks that require fine-grained localization, such as object detection or semantic segmentation from multilingual descriptions.

Our main contributions can thus be summarized as follows:

- A CLIP model for the remote sensing domain that can handle textual inputs in 10 different languages. Both the model and the corresponding source code are publicly available¹;
- Experiments showing that the model obtains state-of-the-art performance on well-known benchmarks, considering tasks such as image-text retrieval (for English and also other languages) or zero-shot image classification.

The rest of this document is organized as follows: Section 2 covers previous related work. Section 3 presents RS-M-CLIP, detailing the training procedure and the data augmentation strategy. Section 4 presents the experimental setup, while Section 5 presents the main results, and Section 6 presents qualitative examples. Finally, Section 7 summarizes our conclusions and discusses future work.

2 RELATED WORK

The current volume of remotely sensed Earth Observation (EO) data motivates the development of vision-and-language methods that can support natural language interactions with large image repositories. Given the success of CLIP for generalist zero-shot image classification and image-text retrieval, several recent studies have

discussed the use and adaptation of CLIP for the remote sensing domain. This section surveys relevant previous efforts.

2.1 Contrastive Vision-Language Pre-training

CLIP models [41] are usually based on two separate Transformer encoders, respectively for visual and textual inputs. The representations obtained for each modality are trained using a contrastive objective [47] that promotes similar semantic contents to be close in the resulting embedding space. After training with Web-scale data, CLIP models exhibit strong capabilities for zero-shot image classification and cross-modal retrieval.

Motivated by the strong results, the adaptation of CLIP to specific domains is currently being actively researched. For instance, in the medical imaging domain, radiology images paired to expert medical reports have been used for contrastive training [15, 52], resulting in models that can support several clinical downstream tasks.

In the specific domain of remote sensing, several recent efforts are worth of mention. For instance Pal et al. [37] fine-tuned CLIP with image-caption pairs, observing high-quality results for zero-shot image classification. Rahhal et al. [42] fine-tuned CLIP both in mono- and multilingual scenarios, observing good results for cross-modal image-text retrieval. Yuan et al. [58] developed a parameter efficient fine-tuning strategy, based on an adapter mechanism, to address image-text retrieval for the remote sensing domain.

Some studies have addressed the problem of the small size of existing remote sensing image-caption datasets, developing different data creation pipelines. For instance Liu et al. [28] aggregated information from existing datasets with annotations for semantic segmentation and object detection, using the resulting dataset to train a large CLIP model that achieves superior results in image-text retrieval. Zhang et al. [67] introduced RS5M, i.e. a dataset that collects real instances filtered from publicly available generalist datasets that also contain remote sensing images, and synthetic instances with captions derived from remote sensing image classification datasets. The authors also studied different parameter-efficient approaches to adapt CLIP with the RS5M dataset.

2.2 Combining CLIP with Language Decoders

CLIP vision encoders are commonly used as backbones within recent architectures leveraging Large Language Models (LLMs) for vision-and-language tasks. For instance, LLaVA [29], MiniGPT-v2 [4], or InstructBLIP [12] are all examples of instruction fine-tuned models capable of supporting conversation and reasoning over visual inputs, connecting a CLIP vision encoder to a pre-trained LLM, and using supervised learning to train the cross-modal connections. Considering the remote sensing domain, similar efforts include RSGPT [21], GeoChat [24], SkyEyeGPT [65], or LHRS-Bot [35]. In previous work, we also tested the use of a CLIP model that was fine-tuned over an aggregated remote sensing dataset, connecting its vision encoder to a pre-trained LLM, so as to support generative tasks such as image captioning [46].

Some previous studies have also advanced training-free strategies for connecting CLIP models to LLMs. For instance Socratic models [64] corresponds to a generalist framework in which an LLM is used to process a prompt that integrates information from different models. Under this framework, a task like image captioning can be performed by using CLIP to retrieve different concepts

¹<https://github.com/DanielSilva/RS-M-CLIP>

from an input image (e.g., locations or objects that are likely featured in the image), and then prompting an LLM with the retrieved information. LMCap [43] considered a similar strategy, addressing multilingual image captioning through the use of CLIP to retrieve similar captions from a datastore of examples, which are then used to prompt an LLM into generating the target caption. In Appendix A, we present results from a set of experiments using the LMCap method together with our RS-M-CLIP model.

2.3 Alternatives to Contrastive Pre-training

Recent research has addressed the problem of how to improve vision-and-language model performance, complementing CLIP with better data selection strategies and different training objectives.

Improving the quality of the captions that are paired to images in the training data has been shown to lead to better results, also with better data efficiency [16, 48], while data augmentation has also been observed to be useful, in particular for domains involving smaller datasets [27]. NegCLIP [62] proposed the use of hard negatives within the CLIP training batches, while CE-CLIP [66] introduced two loss components: one for maximizing the difference between real and hard negative captions, and another making the similarity between a real image-caption pair, and a negative instance, be larger than a threshold. Similar strategies hold significant potential when adapting CLIP to the remote sensing domain.

The combination of self-supervised learning objectives, together with the original contrastive loss, has also been explored. SLIP [34] included the SimCLR [6] objective, while SILC [36] introduced a self-distillation objective and observed improvements in zero-shot generalist image classification, cross-modal retrieval, and zero-shot semantic segmentation. EVA [17] addressed the limits of scaling vision models based on masked image modeling, proposing to use CLIP features as the prediction targets. The resulting model obtained state-of-the-art results in a broad range of vision tasks, and it is also commonly used as an encoder for vision and language LLMs [4, 12]. In our work, we also considered the combination of contrastive learning with a self-distillation objective, when training a CLIP model, aiming to improve the capabilities of the corresponding vision encoder for the remote sensing domain.

2.4 Detailed Localization with CLIP

Given its zero-shot ability to match images to open textual descriptions, CLIP has also been used within models that feature decoders specific to tasks such as open-vocabulary object detection or semantic segmentation. Approaches like Zegformer [13] or zzseg [56] are fully-supervised, but training-free models/approaches that explore the feature maps obtained directly from CLIP have also been proposed. Specifically for the remote sensing domain, Chen et al. [7] extended a U-Net-like structure to align its representations with CLIP, allowing for open-vocabulary segmentation.

In connection to tasks involving detailed localization, the quality of the dense features (i.e., features computed per each image pixel) obtained from CLIP models has also been the subject of recent research. While directly using CLIP has been observed to lack the ability to precisely localize visual concepts within an image, techniques have been developed to slightly modify the architecture of CLIP models when extracting these features, while remaining training-free. For instance MaskCLIP [72] discards the processing

of query and key vectors in the last attention layer of the vision encoder. SCLIP [49] observes that CLIP learns spatial-invariant visual features, in the sense that local features tend to be invariant to their spatial position in the image, proposing a new self-attention mechanism to obtain spatial covariant features. NAACLIP [18] introduces a mechanism that explicitly increases attention values between vicinity patches of the input image, based on the heuristic that adjacent patches often represent the same class.

2.5 Cross-Modal Retrieval

The growing volume of remote sensing imagery has motivated the proposal of specialized cross-modal retrieval methods, e.g. supporting information discovery through text queries over large repositories. For instance, GaLR [61] uses a custom attention mechanism to combine global image features extracted from a CNN, and local features obtained with a graph convolution network. KCR [32] and KTIR [31] introduce external knowledge from a knowledge graph to enrich text descriptions, and alleviate the information gap between images and text. Instead of specialized architectures for cross-modal retrieval, other studies have instead leveraged CLIP [41], adapting models to the remote sensing domain [28, 42, 67] as detailed in Section 2.1. For instance, Yuan et al. [58] developed a CLIP adapter specifically for the remote sensing image-text retrieval task.

3 REMOTE SENSING MULTILINGUAL CLIP

Our work explored the use of a self-distillation mechanism, together with the conventional contrastive objective from CLIP, aiming at the development of an improved vision-and-language encoder model for the remote sensing domain. Figure 1 illustrates our proposal.

In brief, the self-distillation mechanism consists of training the vision encoder of CLIP to align local and global views of a given image, so as to improve the handling of fine-grained image details. This is done by passing different views of an image to student and teacher networks based on the CLIP vision encoder. The resulting output is a probability distribution of K dimensions, which the student is trained to match through the cross-entropy loss over the teacher outputs. With this loss, the student parameters are updated with gradient descent, while the teacher parameters are updated with an exponential moving average of the student parameters.

Besides improved training through the self-distillation mechanism, our work also studied the impact of using machine translation to augment the training set of captions paired to remote sensing images. A state-of-the-art LLM for machine translation was used to translate English captions into nine other languages, resulting in a much richer dataset for model training. Given that captions for remote sensing images are often relatively simple sentences, automatic machine translation should not introduce significant errors, thus corresponding to a simple and effective approach with clear immediate benefits (e.g., it becomes possible to support the training of models that can deal with multilingual inputs, while simultaneously aiming at better English performance).

3.1 Model Architecture

The CLIP model consists of a dual Transformer encoder architecture designed to obtain representations for images and texts. The original OpenAI CLIP model cannot process text on other languages besides English, leading us to use a variant that uses

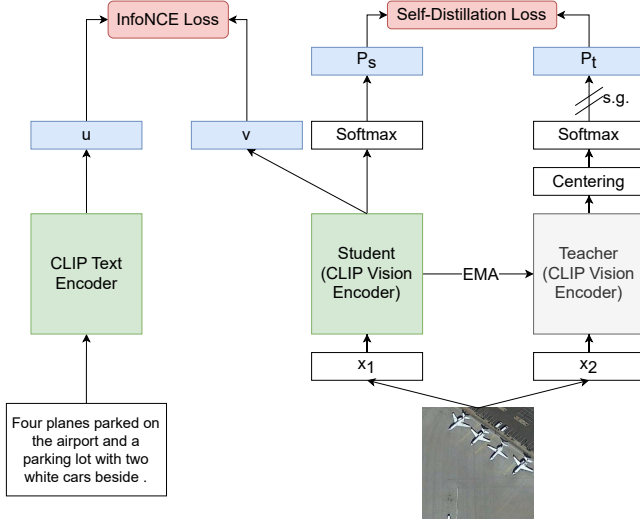


Figure 1: Overview of our method, where self-distillation is combined with contrastive learning to fine-tune a CLIP model using image-caption pairs from remote sensing datasets. Following DINO [3, 36], local and global image crops are fed into student and teacher networks. Each network produces a probability distribution of size K , and the student learns to match the teacher’s distribution using a cross-entropy loss, aligning local and global views. The student is updated via gradient updates, while the teacher is updated with an Exponential Moving Average (EMA). The standard CLIP loss for matching image-text pairs is also applied.

XLM-RoBERTa [11] as the text encoder. For the application of the self-distillation objective, a slight modification over the CLIP vision encoder component is done, where a projector is added on top of the student and teacher. This projector consists of a 3-layer multi-layer perceptron, a l_2 normalization operation, and a weight normalized [44] fully connected layer with K dimensions.

3.2 Training Procedure

An architectural diagram of the proposed training process can be seen in Figure 1. Our method involves joint training with two losses, combining contrastive learning between image and text embeddings (i.e., the standard Information Noise Contrastive Estimation (InfoNCE) loss) together with a self-distillation objective inspired on DINO [3, 36]. We start from a pre-trained generalist CLIP model, using the proposed method to further fine-tune the model to the specific characteristics of remote sensing data.

3.2.1 Contrastive Learning. Contrastive learning can be applied over an unlabeled dataset of paired image-text captions, to maximize the similarity between embeddings of matching pairs [41]. Specifically, for a dataset of M image-caption pairs $\mathcal{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^M$, a vision encoder obtains a representation for the input image $f_{\phi}^i(\mathbf{x}_i) = \mathbf{v}_i \in \mathbb{R}^m$, and a text encoder obtains another representation for the caption $f_{\theta}^t(\mathbf{y}_i) = \mathbf{u}_i \in \mathbb{R}^m$. During training, the Information Noise Contrastive Estimation (InfoNCE) loss for both text-to-image \mathcal{L}_{t2i} and image-to-text \mathcal{L}_{i2t} matchings is minimized.

In a batch of N examples, each image and its corresponding caption are considered a positive pair, while all other combinations of images and captions are considered negative pairs. Given a learnable parameter τ , both losses can be formalized as follows:

$$\mathcal{L}_{t2i} = -\frac{1}{N} \sum_{i=1}^N \left(\log \frac{\exp(\text{sim}(\mathbf{u}_i, \mathbf{v}_i)/\tau)}{\sum_{j=1}^N \exp(\text{sim}(\mathbf{u}_i, \mathbf{v}_j)/\tau)} \right), \quad (1)$$

$$\mathcal{L}_{i2t} = -\frac{1}{N} \sum_{i=1}^N \left(\log \frac{\exp(\text{sim}(\mathbf{v}_i, \mathbf{u}_i)/\tau)}{\sum_{j=1}^N \exp(\text{sim}(\mathbf{v}_i, \mathbf{u}_j)/\tau)} \right), \quad (2)$$

where the similarity is the cosine similarity given by $\text{sim}(\mathbf{a}, \mathbf{b}) = \exp(\mathbf{a} \cdot \mathbf{b}) / (\|\mathbf{a}\| \|\mathbf{b}\|)$. The final contrastive learning loss is:

$$\mathcal{L}_{\text{InfoNCE}} = -\frac{1}{2} (\mathcal{L}_{t2i} + \mathcal{L}_{i2t}). \quad (3)$$

3.2.2 Self-distillation. Besides contrastive learning, we also included a self-distillation mechanism in the training procedure, aiming to obtain higher-quality representations from the image encoder. Following DINO [3], a student network g_{θ_s} is trained to match the output given by a teacher network g_{θ_t} , with the networks being parameterized by θ_s and θ_t , respectively. Given an input image x , the output of each network is normalized with a softmax function, and the results can be interpreted as probability distributions over K dimensions, P_s and P_t . The objective function consists of minimizing the cross entropy between the two distributions with respect to the parameters of the student network θ_s , while maintaining the parameters of the teacher network θ_t fixed:

$$\min_{\theta_s} = -P_t(x) \log(P_s(x)). \quad (4)$$

In practice, for each input image, S views are obtained by applying crop and distortion transformations, where two correspond to global views, i.e. x_1^g and x_2^g , and the rest are a set of local views, where each one is an image crop of smaller resolution. Then, every view is given to the student network, while only global views are given to the teacher, encouraging similar representations for the local and global views [3, 36]. Thus, the following loss is minimized:

$$\mathcal{L}_{\text{Self-Distil}} = \sum_{x \in \{x_1^g, x_2^g\}} \sum_{x' \in S, x' \neq x} -P_t(x) \log(P_s(x')). \quad (5)$$

Both the teacher and the student networks share the same architecture, namely the Transformer vision encoder from a pre-trained CLIP model, although the two networks correspond to different sets of weights, respectively θ_s and θ_t . The student parameters are learned by minimizing Equation 5 with stochastic gradient descent, while the parameters of the teacher are updated with an Exponential Moving Average (EMA) of the student weights, with an update rule corresponding to $\theta_t \leftarrow \lambda \theta_t + (1 - \lambda) \theta_s$.

Regarding training stability, and following the DINO setup, collapse of the obtained representations is avoided by performing centering and sharpening over the teacher outputs. This corresponds to subtracting a centering parameter c while dividing the teacher outputs by a temperature parameter τ_t in the teacher softmax normalization. The centering parameter c is obtained with an exponential moving average of the mean of the features obtained from the teacher network. These operations are performed

as follows for each dimension $i \in \{1, \dots, K\}$:

$$g_{\theta_t}(x)^{(i)} \leftarrow \frac{\exp\left((g_{\theta_t}(x)^{(i)} - c)/\tau_t\right)}{\sum_{k=1}^K \exp\left((g_{\theta_t}(x)^{(k)} - c)/\tau_t\right)}. \quad (6)$$

The final training loss for the student, including the contrastive learning and self-distillation objectives, is given by:

$$\mathcal{L} = \frac{1}{2}(\mathcal{L}_{\text{InfoNCE}} + \mathcal{L}_{\text{Self-Distil}}). \quad (7)$$

3.3 Multilingual Augmentations for Captions

The strategy described in the previous section was used for fine-tuning a generalist multilingual CLIP model, using a large dataset built through automatic machine translation.

We specifically used TowerInstruct [1], i.e. a Large Language Model (LLM) fine-tuned for machine translation, with state-of-the-art performance in several high-resource languages. The model is prompted in a zero-shot manner to translate individual English captions into other languages. We note that captions for remote sensing images are mostly relatively simple sentences, and hence the simple and direct application of this model should suffice at avoiding severe translation errors.

We translated the captions into 9 languages other than English, as supported by TowerInstruct: German, French, Spanish, Chinese, Portuguese, Italian, Russian, Korean, and Dutch.

4 EXPERIMENTAL SETUP

This section describes implementation details such as the choice of training hyper-parameters, the image-caption datasets used in the experiments, and the considered evaluation metrics.

4.1 Implementation Details

The models considered in our experiments were implemented using the PyTorch [38] library, and we also used the OpenCLIP software library [22]. The base multilingual CLIP model used in most of our experiments was *xlm-roberta-base-ViT-B-32*, with pre-trained weights corresponding to a checkpoint named *laion5b_s13b_b90k*². Interestingly, the large-scale LAION-5B dataset of images collected from the Web, originally used to train this model, already includes a significant amount of remote sensing imagery [67].

Regarding training hyper-parameters, the batch size was set to 128, the learning rate was set to 0.00025, and optimization used the AdamW optimizer for 100 epochs, with a warmup of 10 epochs. As mentioned before, in the setup using self-distillation training, only the student parameters are updated with stochastic gradient descent, while the teacher parameters are updated with an Exponential Moving Average (EMA) from the student weights. The momentum of the EMA update of the teacher is set to 0.996, the dimensionality of the output probability distribution of the networks is set to $K = 65536$, and the hidden dimensionality is set to 2048, as per the default parameters in DINO. For the contrastive learning objective, similarities are computed between the embeddings of the caption and the first global view. Regarding image transformations, we followed the ones proposed by DINO, which include color jittering, Gaussian blur, and solarization, when obtaining the multiple image crops. As previously mentioned, a total of S views

are generated for each input image, consisting of 2 global views created by randomly cropping the image to 224×224 pixels, and 8 local views with a crop size of 96×96 pixels. Due to the positional embeddings of the CLIP vision encoder, images can only be passed as input with the size of 224×224 pixels, so local views are scaled up again using bicubic interpolation.

For generating translations of the captions, TowerInstruct-13B-v0.1 [1] was prompted in a zero-shot manner with the following instruction: Translate the following text from English into {language}. \nEnglish: {caption}\n{language}::

4.2 Datasets

Similarly to RemoteCLIP [28], we combined different remote sensing image captioning datasets within our experiments, using the original splits of training, validation, and testing instances. These datasets are described next, and also summarized in Table 1.

- RSCID [30] features a large set of images collected from different sources. The images were manually annotated with short captions, but many of the captions are duplicated to ensure 5 captions per image;
- UCM-Captions and Sydney-Captions [40] are based on scene classification datasets that were repurposed for image captioning by manual annotation. Although extensively used, these datasets are small and have limitations in the fact that the captions correspond to very simple sentences that are highly similar between themselves;
- RSITMD [60] was proposed to address limitations in other datasets, featuring annotations with fine-grained information, and with a high overall difference between sentences.
- NWPU-Captions [9] is currently the largest dataset of image and text pairs in the remote sensing domain. Each image has 5 manually annotated captions associated with it, and the images span over a large set of classes describing different land cover and land use types.

We aggregated the aforementioned datasets and subsequently reference them as a group through the name **Cap-5**. For comparison, statistics about the RemoteCLIP dataset [28] are also included in Table 1. This other dataset consists of an aggregation of a high volume of images from datasets associated to different tasks (i.e., datasets for object detection (DET-10) and semantic segmentation (SEG-4) tasks), with synthetic captions generated from the available annotations. We refer the reader to the RemoteCLIP paper [28] for information regarding the data creation pipeline. RET-3 is also included, which corresponds to the aggregation of the RSCID, UCM-Captions, and RSITMD datasets.

4.3 Evaluation Metrics

For image-text retrieval, following previous work [28, 31], we report the retrieval recall at top-1 (R@1), top-5 (R@5), and top-10 (R@10) positions, together with the mean recall (mR). R@K simply measures the ratio of queries that successfully retrieve the ground truth as one of the first K results.

For the zero-shot classification task, we follow [28] and use a single template of the form a satellite photo of {class name} to prompt the text encoder of CLIP to generate an embedding for each class of the dataset. For each image, the predicted class is

²<https://huggingface.co/laion/CLIP-ViT-B-32-xlm-roberta-base-laion5B-s13B-b90k>

Table 1: Description of the different remote sensing image captioning datasets used in our experiments. Cap-5 is an aggregation of five popular remote sensing image captioning datasets, extensively used in previous studies. We note there may be an overlap in the number of Base Classes in the aggregated datasets, corresponding to different class names for the same concept.

Individual Dataset	#Base Classes	#Images	Image Size	Spatial Resolution	#Total Captions
NWPU-Captions [9]	45	31,500	256 × 256	approx. 30-0.2m	157,500
RSICD [30]	30	10,921	224 × 224	different resolutions	54,605
Sydney-Captions [40]	7	613	500 × 500	0.5m	3,065
UCM-Captions [40]	21	2,100	256 × 256	approx. 0.3m	10,500
RSITMD [60]	32	4,743	256 × 256	different resolutions	23,715
Aggregated Dataset	#Base Classes	#Images	Datasets Included		#Total Captions
Cap-5	135	49,877	RET-3 + NWPU-Captions, Sydney-Captions		249,385
RET-3	83	13,634	RSICD, UCM-Captions, RSITMD		68,170
RemoteCLIP [28]	194	165,745	RET-3 + DET-10 + SEG-4		828,725

the one with the most similar embedding to the embedding of the image. In the case of remote sensing scene classification datasets that do not have pre-established splits for training and testing, we set a seed value (42), randomly shuffle a list of the image filenames for each class, and split the data leaving the first 80% of the elements for training and the remaining 20% for testing.

5 EXPERIMENTS

This section presents experimental results, assessing baselines against fine-tuned CLIP models in different tasks and settings.

5.1 Fine-Tuning with Contrastive Learning

Table 2 reports the zero-shot performance of generalist CLIP models with ViT-B and ViT-L image encoders, in the cross-modal retrieval task and when considering different English datasets. We observe a reasonable performance, with only a slightly lower mean recall than baseline models specifically developed for the task, such as AMFMN, GaLR, and KCR. As expected, a larger CLIP image encoder improves the performance across all the datasets.

Previous work had already observed that fine-tuning CLIP to the remote sensing domain leads to improved retrieval results, contributing with different CLIP models such as CLIP-RSICD [37], GeoRSClip [67], and RemoteCLIP [28]. Authors either used publicly available datasets or developed pipelines to create image-caption pairs for training. We aggregated 5 different English image-captioning remote sensing datasets, i.e. the collection which we named Cap-5, and fine-tuned a CLIP model with a ViT-L vision encoder, comparing its results with the RemoteCLIP model of the same size, which is the current state-of-the-art. While using only 30% of the size of the RemoteCLIP dataset, our English-only CLIP/Cap-5 model has a significantly better performance in the RSITMD (62.26% vs 50.52% mR) dataset, and a competitive performance in the RSICD (34.61% vs 36.35% mR) and UCM datasets (55.75% vs 54.68% mR). This can be explained by the characteristics of the instances in the RemoteCLIP data collection, which besides the standard benchmark datasets also includes instances derived from object detection and semantic segmentation datasets, with a high image resolution and

domain diversity. While covering a broad set of conditions, the synthetic data can also be associated to a distribution shift regarding the datasets used for retrieval evaluation.

CLIP/Cap-5 also obtains higher results than KTIR [31] in RSITMD (62.26% vs 50.52% mR), but KTIR has a wide gap to it in RSICD (34.61% vs 45.21% mR) and UCM (55.75% vs 58.76% mR).

5.2 Comparing Multilingual Encoders

We compared the original OpenAI CLIP model to a similar CLIP model trained with a collection of instances from LAION-5B [45], specifically a multilingual model trained on 100 different languages, with a ViT-B image encoder and a XLM-RoBERTa [11] text encoder. We refer to this model as CLIP-XLM-RoBERTa.

As can be seen in Table 2, the zero-shot performance of CLIP-XLM-RoBERTa surpasses that of the OpenAI CLIP models, across the different English datasets. Even when compared to the model that features a larger vision encoder, its performance is higher in RSICD (18.69% vs 17.56% mR), similar in RSITMD (27.46% vs 27.88% mR), and only worse in UCM (36.48% vs 41.76% mR). These results suggest clear benefits in using CLIP models pre-trained on larger datasets, even if the larger datasets are multilingual.

We also fine-tuned CLIP-XLM-RoBERTa with Cap-5, and observed significant improvements. Compared with our previous larger model (i.e., CLIP-L) that was also fine-tuned with Cap-5, a large improvement can be observed in RSICD (45.56% vs 34.61% mR), and similar results are seen in UCM (55.13% vs 55.75% mR) and RSITMD (61.42% vs 62.26% mR). Compared with the RemoteCLIP model of larger size, a big improvement can be seen in RSITMD (61.42% vs 50.52% mR), and competitive performance is observed in RSICD (45.56% vs 36.35% mR) and UCM (55.13% vs 54.68% mR).

CLIP-XLM-RoBERTa fine-tuned with Cap-5 is also better than KTIR in RSICD (45.56% vs 45.1% mR) and RSITMD (61.42% vs 54.27% mR), although worse in UCM (55.13% vs 58.76% mR).

Together, our experimental results consolidate the improvements gained from the pre-trained weights of CLIP-XLM-RoBERTa, and the suitability of contrastive learning to adapt CLIP to the remote sensing domain. Considering that the Cap-5 dataset is a suitable source for training CLIP models, we then analyzed other training mechanisms besides the standard contrastive objective.

Table 2: Image-text retrieval results for the remote sensing domain. Models with names in bold were fine-tuned by us, and results obtained through our own experiments, instead of collected from previous publications, are marked with †.

Dataset	Method	Vision Backbone	Finetuning Data	Image-Text Retrieval			Text-Image Retrieval			Mean	
				R@1	R@5	R@10	R@1	R@5	R@10	Recall	
RSICD	AMFMN [20]	ResNet50	RSICD	5.39	15.08	23.40	4.90	18.28	31.44	16.42	
	CMFM-Net [57]	ResNet18	RSICD	5.40	18.66	28.55	5.31	18.57	30.03	17.75	
	GaLR [61]	ResNet18	RSICD	6.59	19.85	31.04	4.69	19.48	32.13	18.96	
	KCR [32]	ResNet18	RSICD	4.76	18.59	27.20	5.84	22.31	36.12	19.14	
	CLIP [41]†	ViT-B	Zero-shot	4.58	14.55	23.70	5.80	16.85	28.23	15.62	
	CLIP [41]†	ViT-L	Zero-shot	6.04	17.48	27.54	5.03	19.03	30.25	17.56	
	Rahhal et al. [42]	ViT-B	RSICD	10.70	29.64	41.53	9.14	28.96	44.59	27.43	
	CLIP-RSICD [37]†	ViT-B	RSICD	14.09	30.10	43.64	11.16	33.25	48.91	30.19	
	GeoRSCLIP [67]	ViT-B	RS5M	11.53	28.55	39.16	9.52	27.37	40.99	26.18	
	RemoteCLIP [28]	ViT-B	RET-3 + DET-10 + SEG-4	17.02	37.97	51.51	13.71	37.11	54.25	35.26	
	RemoteCLIP [28]	ViT-L	RET-3 + DET-10 + SEG-4	18.39	37.42	51.05	14.73	39.93	56.58	36.35	
	KTIR [31]	ViT-B	RSICD	26.08	49.77	62.49	20.55	48.67	63.70	45.21	
	CLIP/Cap-5†	ViT-L	Cap-5	17.02	33.94	47.76	13.83	39.07	56.05	34.61	
	Multilingual Backbones										
		CLIP-XLM-RoBERTa ² †	ViT-B	Zero-shot	7.41	16.83	25.62	6.26	22.07	33.93	18.69
		CLIP-XLM-RoBERTa†	ViT-B	Cap-5 / English only	27.17	50.69	64.04	20.79	48.58	62.10	45.56
	RS-M-CLIP†	ViT-B	Cap-5 / English only	58.55	69.08	73.65	43.17	58.79	66.57	61.64	
	RS-M-CLIP†	ViT-B	Cap-5 / 1 Translation	60.48	72.10	75.93	44.37	60.35	69.02	63.71	
UCM	AMFMN [20]	ResNet50	UCM	16.67	45.71	68.57	12.86	53.24	79.43	46.08	
	KCR [32]	ResNet18	UCM	11.90	48.57	71.43	17.24	56.95	81.14	47.87	
	CLIP [41]†	ViT-B	Zero-shot	8.10	34.29	53.33	8.67	36.48	60.57	33.57	
	CLIP [41]†	ViT-L	Zero-shot	11.91	43.33	65.24	10.76	46.00	73.33	41.76	
	Rahhal et al. [42]†	ViT-B	UCM	19.04	53.33	77.61	19.33	64.00	91.42	54.12	
	CLIP-RSICD [37]†	ViT-B	RSICD	15.71	50.00	82.38	13.81	57.05	91.24	51.70	
	RemoteCLIP [28]	ViT-B	RET-3 + DET-10 + SEG-4	20.48	59.85	83.33	18.67	61.52	94.29	56.36	
	RemoteCLIP [28]	ViT-L	RET-3 + DET-10 + SEG-4	19.05	54.29	80.95	17.71	62.19	93.90	54.68	
	KTIR [32]	ViT-B	UCM	21.42	64.29	87.14	19.81	64.57	95.33	58.76	
	CLIP/Cap-5†	ViT-L	Cap-5	21.91	59.05	81.91	16.29	60.57	94.76	55.75	
	Multilingual Backbones										
		CLIP-XLM-RoBERTa ² †	ViT-B	Zero-shot	9.52	35.24	52.38	11.62	43.71	66.38	36.48
		CLIP-XLM-RoBERTa†	ViT-B	Cap-5 / English Only	18.10	60.95	83.33	17.14	59.33	91.91	55.13
		RS-M-CLIP†	ViT-B	Cap-5 / English only	14.76	52.86	75.24	14.57	52.57	74.29	47.38
		RS-M-CLIP†	ViT-B	Cap-5 / 1 Translation	14.29	50.00	77.14	16.29	55.33	83.91	49.49
	RSITMD	AMFMN [20]	ResNet50	RSITMD	10.63	24.78	41.81	11.51	34.69	54.87	29.72
CMFM-Net [57]		ResNet18	RSITMD	10.84	28.76	40.04	10.00	32.83	47.21	28.28	
GaLR [61]		ResNet18	RSITMD	14.82	31.64	42.48	11.15	36.68	51.68	31.41	
CLIP [41]†		ViT-B	Zero-shot	9.74	22.57	34.51	8.72	27.88	42.88	24.38	
CLIP [41]†		ViT-L	Zero-shot	10.18	27.66	38.50	11.46	32.52	46.99	27.88	
Rahhal et al. [42]		ViT-B	RSITMD	19.69	40.26	54.42	17.61	49.73	66.59	41.38	
CLIP-RSICD [37]		ViT-B	RSICD	25.00	46.46	61.73	19.29	51.02	68.45	45.32	
GeoRSCLIP [67]		ViT-B	RS5M	19.03	34.51	46.46	14.16	42.39	57.52	35.68	
RemoteCLIP [28]		ViT-B	RET-3 + DET-10 + SEG-4	27.88	50.66	65.71	22.17	56.46	73.41	49.38	
RemoteCLIP [28]		ViT-L	RET-3 + DET-10 + SEG-4	28.76	52.43	63.94	23.76	59.51	74.73	50.52	
KTIR [31]		ViT-B	RSITMD	34.29	55.31	65.04	31.46	62.92	76.59	54.27	
CLIP/Cap-5†		ViT-L	Cap-5	41.37	67.48	79.65	34.12	68.81	82.12	62.26	
Multilingual Backbones											
		CLIP-XLM-RoBERTa ² †	ViT-B	Zero-shot	8.41	26.11	34.96	11.33	34.29	49.69	27.46
		CLIP-XLM-RoBERTa†	ViT-B	Cap-5 / English only	39.82	64.60	75.89	38.85	68.81	80.53	61.42
		RS-M-CLIP†	ViT-B	Cap-5 / English only	63.05	75.00	80.97	55.18	73.27	80.71	71.36
	RS-M-CLIP†	ViT-B	Cap-5 / 1 Translation	67.48	77.66	82.30	59.87	76.99	83.01	74.55	

5.3 Self-Distillation and Contrastive Learning

Results for our proposed model, which combines contrastive learning with a self-distillation objective, can also be seen in Table 2. A large improvement is achieved in the RSICD and RSITMD datasets,

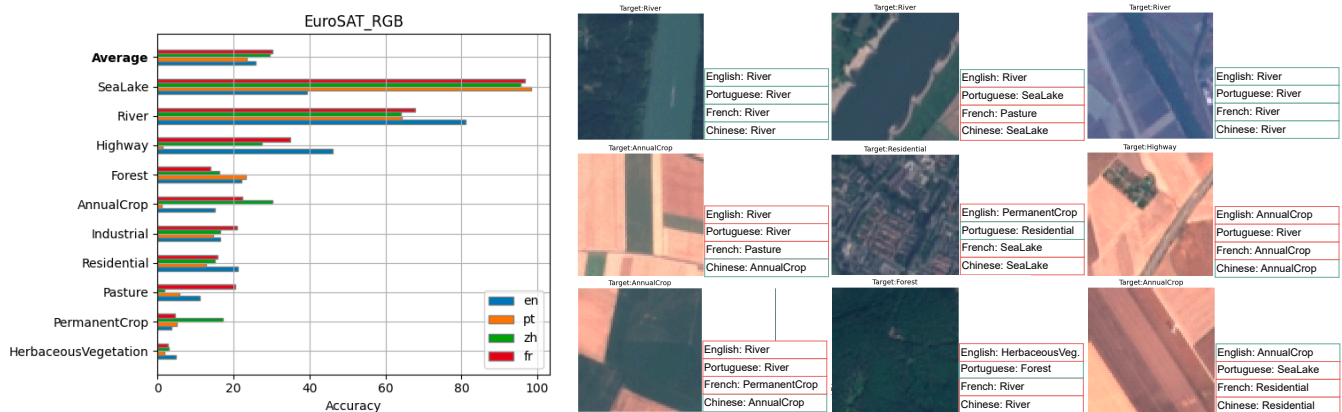
even over encoders of larger size. The mean recall reaches an impressive value of 61.64% in RSICD (+16.43% over the SOTA model, i.e. KTIR) and 71.36% in RSITMD (+17.09% over KTIR). However, in the UCM dataset, the results are worse (-11.38% over KTIR). This

Table 3: Zero-shot image classification in 12 remote sensing datasets. Due to space constraints, CLIP-XLM-RoBERTa is mentioned in the table simply as CLIP-XLM.

Method	Backbone	RSI-CB128	RSI-CB256	WHU-earth	EuroSAT	MLRSNet	PatternNet	RESISC45	AID	RSSCN7	OPTIMAL-31	RSC11	WHU-RS19	Average
CLIP [41]	ViT-B	20.09	29.40	40.62	41.83	41.12	49.65	50.65	51.40	65.00	66.13	46.40	73.79	48.01
RemoteCLIP [28]	ViT-B	25.88	42.56	63.12	33.04	57.91	58.01	66.87	86.60	73.75	79.03	57.16	94.17	61.51
CLIP [41]	ViT-L	32.69	43.65	56.46	39.19	59.09	70.15	64.16	62.70	67.68	73.92	64.26	87.86	60.15
RemoteCLIP [28]	ViT-L	37.23	51.55	68.33	46.04	61.67	63.12	73.68	76.95	71.07	80.65	68.43	92.72	65.95
CLIP-XLM	ViT-B	31.12	41.82	59.17	52.30	55.97	60.02	62.79	67.45	69.82	75.27	56.04	85.44	59.77
RS-M-CLIP	ViT-B	28.27	38.61	63.96	25.85	60.62	50.95	92.62	88.90	60.89	95.16	61.42	90.78	63.17

Table 4: Image-text retrieval results in over the RSICD dataset, for multilingual CLIP models in 10 different languages.

Language	Method	Visual Backbone	Finetune Data	Image-Text Retrieval			Text-Image Retrieval			Mean Recall
				R@1	R@5	R@10	R@1	R@5	R@10	
English	Rahhal et al.	ViT-B	RSICD (Multilanguage)	11.61	30.10	42.17	9.55	29.27	44.73	27.90
	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	7.41	16.83	25.62	6.26	22.07	33.93	18.69
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	60.48	72.10	75.93	44.37	60.35	69.02	63.71
French	Rahhal et al.	ViT-B	RSICD (Multilanguage)	8.78	26.89	40.07	8.19	26.34	43.11	25.56
	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	6.13	14.55	24.06	5.65	19.01	30.34	16.62
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	57.82	71.36	75.76	43.82	59.51	68.66	62.82
Italian	Rahhal et al.	ViT-B	RSICD (Multilanguage)	10.15	26.98	37.69	8.43	26.56	42.08	25.31
	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	5.76	15.83	23.61	5.00	18.52	29.46	16.36
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	58.65	71.09	75.48	43.44	59.65	68.75	62.84
Portuguese	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	5.86	15.83	22.60	4.58	18.59	29.55	16.17
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	59.01	71.46	75.30	43.84	59.91	68.23	62.96
	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	6.68	15.55	23.88	5.18	19.07	29.63	16.66
Spanish	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	59.10	71.18	75.48	43.75	59.96	68.89	63.06
	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	5.31	14.27	22.78	5.53	17.86	28.03	15.63
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	54.90	68.98	74.93	38.66	56.23	65.64	59.89
German	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	5.03	15.01	23.06	4.63	18.50	29.84	16.01
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	58.74	71.73	76.49	43.86	59.63	67.65	63.02
	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	4.30	12.53	19.12	3.84	16.03	25.34	13.53
Korean	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	56.91	69.35	73.74	42.05	58.87	67.96	61.48
	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	5.40	15.46	23.42	4.41	16.18	26.97	15.31
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	55.63	69.44	74.66	42.09	58.98	67.36	61.36
Chinese	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	5.86	15.37	23.79	4.96	17.69	27.94	15.94
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	57.46	71.55	75.57	42.65	59.76	68.25	62.54
	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	5.86	15.37	23.79	4.96	17.69	27.94	15.94
Russian	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	57.46	71.55	75.57	42.65	59.76	68.25	62.54
	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	5.86	15.37	23.79	4.96	17.69	27.94	15.94
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	57.46	71.55	75.57	42.65	59.76	68.25	62.54

**Figure 2: Zero-shot classification performance on the EuroSAT dataset, across different languages.**

might suggest that the lower proportion of the UCM dataset instances within the Cap-5 dataset (approx. 4%), compared to RSICD (approx. 22%) and RSITMD (approx. 9.5%), might lead the model

to optimize more to characteristics of the datasets with larger size, hindering performance on UCM.

Table 5: Image-text retrieval results in over the RSITMD dataset for multilingual CLIP models in 10 different languages.

Language	Method	Visual Backbone	Finetune Data	Image-Text Retrieval			Text-Image Retrieval			Mean Recall
				R@1	R@5	R@10	R@1	R@5	R@10	
English	Rahhal et al.	ViT-B	RSITMD (Multilanguage)	22.56	43.80	57.30	19.29	51.37	68.18	43.75
	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	8.41	25.89	34.96	11.37	34.25	49.69	27.43
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	67.48	77.66	82.30	59.87	76.99	83.01	74.55
French	Rahhal et al.	ViT-B	RSITMD (Multilanguage)	20.35	43.80	58.18	19.29	49.38	65.61	42.77
	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	11.06	23.67	37.39	10.66	29.78	44.60	26.20
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	66.59	78.10	82.97	58.89	76.02	83.10	74.28
Italian	Rahhal et al.	ViT-B	RSITMD (Multilanguage)	21.90	42.69	56.63	18.18	49.02	64.91	42.22
	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	9.96	23.01	31.20	9.96	29.56	43.14	24.47
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	68.14	78.10	82.97	58.63	76.02	82.57	74.40
Portuguese	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	9.96	23.45	32.08	9.87	29.29	44.51	24.86
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	67.48	77.43	82.97	59.12	76.15	82.88	74.34
Spanish	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	10.62	21.46	32.52	10.75	31.11	45.22	25.28
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	67.92	77.43	82.97	59.34	76.33	82.83	74.47
German	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	8.41	22.35	32.97	9.91	28.67	43.85	24.36
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	62.39	75.22	82.74	53.01	72.21	79.78	70.89
Dutch	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	10.18	21.68	31.86	10.09	29.12	44.20	24.52
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	65.49	77.66	81.42	59.47	75.80	82.74	73.76
Korean	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	7.52	20.13	28.32	8.45	24.65	37.30	21.06
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	65.71	76.33	81.86	58.32	74.38	81.20	72.97
Chinese	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	9.29	20.35	29.87	9.51	25.35	40.22	22.43
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	64.38	77.21	82.52	57.39	75.40	81.64	73.09
Russian	CLIP-XLM-RoBERTa	ViT-B	Zero-shot	9.07	22.57	32.97	9.25	27.12	41.64	23.77
	RS-M-CLIP	ViT-B	Cap-5 / 1 Translation	67.48	76.99	82.52	58.89	75.09	81.90	73.81

Detailed experiments were also conducted considering captions translated into different languages. For each caption in English within Cap-5, a translation is obtained in 9 other different languages (i.e., Portuguese, Spanish, French, German, Dutch, Italian, Korean, Chinese, and Russian). During training, for each image at each epoch, one of the possible translations of a caption is sampled, thus performing augmentations in the text, while still maintaining the total volume of image-caption pairs seen during training. This is the main RS-M-CLIP model and corresponds to the lines shown in red in the results presented in Table 2. The performance for the image-text retrieval task can be further improved with this strategy, across all English datasets (+2.07% in RSICD, +2.06% in UCM, and +3.19% in RSITMD). Even in UCM, where our method had lower scores, the results were improved, particularly for the text-to-image direction. These findings confirm that caption translation is a form of text data augmentation that can significantly improve the retrieval performance [16, 27, 48].

5.4 Zero-shot Image Classification

We also evaluated zero-shot classification through 12 image classification datasets, comparing our method against other CLIP-based approaches. The datasets used for evaluation are RSI-CB128 [25], RSI-CB256 [25], WHU-earth [68], EuroSAT [19], MLRSNet [39], PatternNet [73], RESISC45 [8], AID [54], RSSCN7 [75], OPTIMAL-31 [50], RSC11 [70], and WHU-RS19 [55]. The results are compiled in Table 3, which presents accuracy values.

Overall, CLIP-XLM-RoBERTa has a higher average accuracy compared to the original OpenAI model, and is competitive with the OpenAI model featuring the larger vision encoder. The main RS-M-CLIP model, trained with translated data as augmentations,

generally surpasses the RemoteCLIP model of the same size (i.e., better results are obtained in 8 out of 12 datasets), and is on average slightly lower than the RemoteCLIP model with a larger size, surpassing it in 3 out of the 12 datasets.

5.5 Multilingual Retrieval

We also evaluated the multilingual capabilities of the main RS-M-CLIP model, considering image-text retrieval across different languages. Since there are no established multilingual benchmarks, the same translation transformation used for the training data was applied to the test splits of RSICD and RSITMD. The results reported by Rahhal et al. [42] are included as a comparison, since this is the only other study that evaluated cross-modal image-text retrieval in a multilingual scenario. However, since the data used for evaluation is not the same (i.e., different translations were used), the results cannot be directly compared. We evaluated retrieval over the 10 languages for which we obtained translations, including Portuguese, Spanish, French, German, Dutch, Italian, Korean, Chinese, and Russian. Of these, French and Italian are also considered by Rahhal et al. [42]. Results are compiled in Tables 4 and 5.

The results show that, for the RSICD and RSITMD datasets, our method obtains much higher results compared to those reported by Rahhal et al. [42]. Across languages, it can be seen that the performance is very similar to the original results obtained in English.

Interestingly, lower results can be observed for the German language, across the different retrieval datasets. This can be explained either by a lower performance of the original CLIP-XLM-RoBERTa model for the German language, or a lower quality of the translation data obtained for German from the translation model.



Figure 3: Examples for image retrieval, ordered by similarity, given different captions in English, Portuguese, and French.

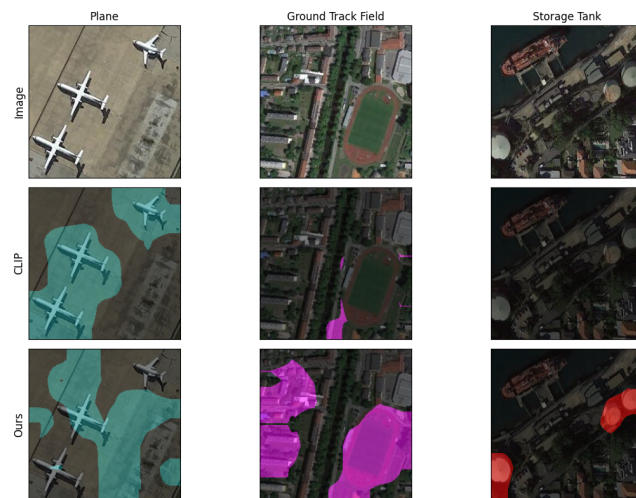


Figure 4: Semantic masks obtained with the NAACLIP [18] method from RS-M-CLIP and CLIP, across different scenes.

6 QUALITATIVE ANALYSIS

The results of an experiment illustrating the multilingual capabilities of RS-M-CLIP can be visualized in Figure 2. In this test, we took the images from the EuroSAT dataset to assess zero-shot classification performance, and we considered class names in 4 languages: English, Portuguese, Chinese, and French. For every language, the prompt for zero-shot classification was translated together with the class names. It can be seen that the image classification performance, on average, is similar across the languages, with French and Chinese even surpassing the English results. On the right part of Figure 2, we show specific image examples.

The multilingual retrieval capabilities are also illustrated in Figure 3, with examples of textual queries in English, Portuguese, and

French. Images that correctly match the text are marked with green edges, while false positives are marked with red edges. From left to right, the three images for each query are ranked by the CLIP similarity. It can be seen that the top results returned by RS-M-CLIP are still closely related to the input textual query and, in all 3 cases, the correct image was returned in the first position for both the English and non-English descriptions.

Finally, we provide a visualization of semantic segmentation masks obtained from RS-M-CLIP, comparing them against masks derived from the original OpenAI CLIP model in Figure 4. The masks are obtained with the NAACLIP method [18], an approach for open-vocabulary semantic segmentation that does not require training any new parameters. Instead, NAACLIP modifies the CLIP architecture to improve the segmentation performance by increasing attention values between neighboring patches, based on the assumption that adjacent patches tend to belong to the same class. This is achieved by introducing a Gaussian kernel over the attention mechanism, and by eliminating components of the final encoder layer of the vision transformer, such as the feedforward layer and residual skip-connections. The examples in Figure 4 are from the NWPU-Captions dataset, considering three different classes present in the iSAID segmentation dataset [53], which features 15 different classes in total. The figure plots segmentation masks with classes related to the image contents, in particular for the classes plane, ground track field, and storage tank. To our surprise, the semantic masks obtained by CLIP are reasonable across different scenes. Still, RS-M-CLIP can more accurately describe most classes, as shown for the ground track field in the middle images, and the storage tanks in the images in the right.

7 CONCLUSIONS AND FUTURE WORK

This work proposed RS-M-CLIP, i.e. a novel multilingual CLIP model for the remote sensing domain, that was trained on a mixture of English and machine translated data by combining a self-distillation objective with standard contrastive learning. RS-M-CLIP achieves impressive results in tasks such as cross-modal retrieval, and it also supports multilingual inputs.

For future work, it would be interesting to study approaches for improving the compositional reasoning abilities of CLIP in the remote sensing domain [62], e.g. further extending RS-M-CLIP through training with hard negative examples. The training methodology can also be further extended for improving fine-grained visual abilities [33, 51]. RS-M-CLIP can also be integrated with large language models, enabling tasks that involve language generation conditioned on remote sensing imagery [46, 65]. Finally, it would also be interesting to adapt/extend RS-M-CLIP to other remote sensing modalities, going beyond the use of RGB images [63].

ACKNOWLEDGMENTS

This research was supported by the Portuguese Recovery and Resilience Plan through project C645008882-00000055 (i.e., the Center For Responsible AI), and also by the Fundação para a Ciência e Tecnologia (FCT), specifically through the project with reference UIDB/50021/2020 (DOI: 10.54499/UIDB/50021/2020), and the project with reference UIDP/04516/2020 (DOI: 10.54499/UIDP/04516/2020).

REFERENCES

- [1] Duarte M. Alves, José Pombal, Nuno M. Guerreiro, Pedro H. Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, et al. 2024. Tower: An Open Multilingual Large Language Model for Translation-Related Tasks. *arXiv:2402.17733* (2024).
- [2] Alberto Baldrati, Marco Bertini, Tiberio Uricchio, and Alberto Del Bimbo. 2022. Effective Conditioned and Composed Image Retrieval Combining CLIP-Based Features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [3] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- [4] Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. 2023. MiniGPT-v2: large language model as a unified interface for vision-language multi-task learning. *arXiv:2310.09478* (2023).
- [5] Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. 2023. Shikra: Unleashing Multimodal LLM's Referential Dialogue Magic. *arXiv:2306.15195* (2023).
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. *arXiv:2002.05709* (2020).
- [7] Yuxing Chen and Lorenzo Bruzzone. 2023. Toward Open-World Semantic Segmentation of Remote Sensing Images. In *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium*.
- [8] Gong Cheng, Junwei Han, and Xiaoqiang Lu. 2017. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE* 10 (2017).
- [9] Qimin Cheng, Haiyan Huang, Yuan Xu, Yuzhuo Zhou, Huanqing Li, and Zhongyuan Wang. 2022. NWPU-captions dataset and MLCa-net for remote sensing image captioning. *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022).
- [10] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. 2023. Reproducible scaling laws for contrastive language-image learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [11] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv:1911.02116* (2019).
- [12] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. InstructBLIP: Towards General-purpose Vision-Language Models with Instruction Tuning. *arXiv:2305.06500* (2023).
- [13] Jian Ding, Nan Xue, Gui-Song Xia, and Dengxin Dai. 2022. Decoupling zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [14] Xiaoyi Dong, Jianmin Bao, Ting Zhang, Dongdong Chen, Shuyang Gu, Weiming Zhang, Lu Yuan, Dong Chen, Fang Wen, and Nenghai Yu. 2022. CLIP itself is a strong fine-tuner: Achieving 85.7% and 88.0% top-1 accuracy with ViT-B and ViT-L on imagenet. *arXiv:2212.06138* (2022).
- [15] Sedigheh Eslami, Christoph Meinel, and Gerard De Melo. 2023. PubMedCLIP: How much does CLIP benefit visual question answering in the medical domain?. In *Findings of the Association for Computational Linguistics: EAACL 2023*.
- [16] Lijie Fan, Dilip Krishnan, Phillip Isola, Dina Katabi, and Yonglong Tian. 2024. Improving CLIP training with language rewrites. In *Proceedings of the Conference on Neural Information Processing Systems*.
- [17] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. 2023. EVA: Exploring the limits of masked visual representation learning at scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [18] Sina Hajimiri, Ismail Ben Ayed, and Jose Dolz. 2024. Pay Attention to Your Neighbours: Training-Free Open-Vocabulary Semantic Segmentation. *arXiv:2404.08181* (2024).
- [19] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. 2019. EuroSat: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* (2019).
- [20] Genc Hoxha, Farid Melgani, and Begüm Demir. 2020. Toward remote sensing image retrieval under a deep image captioning perspective. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 13 (2020).
- [21] Yuan Hu, Jianlong Yuan, Congcong Wen, Xiaonan Lu, and Xiang Li. 2023. RSGPT: A Remote Sensing Vision Language Model and Benchmark. *arXiv:2307.15266* (2023).
- [22] Gabriel Ilharco, Mitchell Wortsman, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hannaneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. 2021. OpenCLIP. (2021). <https://doi.org/10.5281/zenodo.5143773>
- [23] Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. 2023. Grounding language models to images for multimodal inputs and outputs. In *Proceedings of the International Conference on Machine Learning*.
- [24] Kartik Kuckreja, Muhammad Sohail Danish, Muzammal Naseer, Abhijit Das, Salman Khan, and Fahad Shahbaz Khan. 2023. GeoChat: Grounded Large Vision-Language Model for Remote Sensing. *arXiv:2311.15826* (2023).
- [25] Haifeng Li, Xin Dou, Chao Tao, Zhixiang Wu, Jie Chen, Jian Peng, Min Deng, and Ling Zhao. 2020. RSI-CB: A large-scale remote sensing image classification benchmark using crowdsourced data. *Sensors* 20, 6 (2020).
- [26] Xiang Li, Congcong Wen, Yuan Hu, and Nan Zhou. 2023. RS-CLIP: Zero shot remote sensing scene classification via contrastive vision-language supervision. *International Journal of Applied Earth Observation and Geoinformation* 124 (2023).
- [27] Zichao Li, Cihang Xie, and Ekin Dogus Cubuk. 2024. Scaling (Down) CLIP: A Comprehensive Analysis of Data, Architecture, and Training Strategies. *arXiv:2404.08197* (2024).
- [28] Fan Liu, Delong Chen, Zhangqingyun Guan, Xiaocong Zhou, Jiale Zhu, Qiaolin Ye, Liyong Fu, and Jun Zhou. 2024. RemoteCLIP: A Vision Language Foundation Model for Remote Sensing. *IEEE Transactions on Geoscience and Remote Sensing* 62 (2024).
- [29] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. *arXiv:2304.08485* (2023).
- [30] Xiaoqiang Lu, Binqiang Wang, Xiangtao Zheng, and Xuelong Li. 2018. Exploring Models and Data for Remote Sensing Image Caption Generation. *IEEE Transactions on Geoscience and Remote Sensing* 56 (2018).
- [31] Li Mi, Xianjie Dai, Javiera Castillo-Navarro, and Devis Tuia. 2024. Knowledge-aware Text-Image Retrieval for Remote Sensing Images. *arXiv:2405.03373* (2024).
- [32] Li Mi, Siran Li, Christel Chappuis, and Devis Tuia. 2022. Knowledge-Aware Cross-Modal Text-Image Retrieval for Remote Sensing Images. In *Proceedings of the Workshop on Complex Data Challenges in Earth Observation*.
- [33] Amin Karimi Monsefi, Kishore Prakash Sailaja, Ali Alilooee, Ser-Nam Lim, and Rajiv Rammath. 2024. DetailCLIP: Detail-Oriented CLIP for Fine-Grained Tasks. *arXiv:2409.06809* (2024).
- [34] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. 2022. SLIP: Self-supervision meets language-image pre-training. In *Proceedings of the European Conference on Computer Vision*.
- [35] Dilxat Muhtar, Zhenshi Li, Feng Gu, Xueliang Zhang, and Pengfeng Xiao. 2024. LHRs-Bot: Empowering Remote Sensing with VGI-Enhanced Large Multimodal Language Model. *arXiv:2402.02544* (2024).
- [36] Muhammad Ferjad Naeem, Yongqin Xian, Xiaohua Zhai, Lukas Hoyer, Luc Van Gool, and Federico Tombari. 2023. SILC: Improving vision language pretraining with self-distillation. *arXiv:2310.13355* (2023).
- [37] Sujit Pal, Artashes Arutiunian, Goutham Venkatesh, Ritobrata Ghosh, Dev Vidhani, and Mayank Bhaskar. 2021. Fine Tuning CLIP with Remote Sensing (Satellite) Images and Captions. *HuggingFace Blog* (2021).
- [38] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Proceedings of the Conference on Neural Information Processing Systems*.
- [39] Xiaoman Qi, Panpan Zhu, Yuebin Wang, Liqiang Zhang, Junhuan Peng, Mengfan Wu, Jialong Chen, Xudong Zhao, Ning Zang, and P Takis Mathiopoulos. 2020. MLRSNet: A multi-label high spatial resolution remote sensing dataset for semantic scene understanding. *ISPRS Journal of Photogrammetry and Remote Sensing* 169 (2020).
- [40] Bo Qu, Xuelong Li, Dacheng Tao, and Xiaoqiang Lu. 2016. Deep Semantic Understanding of High Resolution Remote Sensing Image. In *Proceedings of the International Conference on Computer, Information and Telecommunication Systems*.
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the International Conference on Machine Learning*.
- [42] Mohamad M. Al Rahhal, Yakoub Bazi, Norah A. Alsharif, Laila Bashmal, Naif Alajlan, and Farid Melgani. 2022. Multilingual Transformer for Improved Text to Remote Sensing Image Retrieval. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 15 (2022).
- [43] Rita Ramos, Bruno Martins, and Desmond Elliott. 2023. LMCap: Few-shot Multilingual Image Captioning by Retrieval Augmented Language Model Prompting. In *Findings of the Association for Computational Linguistics: ACL 2023*.
- [44] Tim Salimans and Durk P Kingma. 2016. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Proceedings of the Conference on Neural Information Processing Systems* (2016).
- [45] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. 2022. LAION-5b: An open large-scale dataset for training next generation image-text models. *Proceedings of the Conference on Neural Information Processing Systems* (2022).

- [46] João Daniel Silva, João Magalhães, Devis Tuia, and Bruno Martins. 2024. Large Language Models for Captioning and Retrieving Remote Sensing Images. *arXiv:2402.06475* (2024).
- [47] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation Learning with Contrastive Predictive Coding. *arXiv:1807.03748* (2018).
- [48] Pavan Kumar Anasosalu Vasu, Hadi Pouransari, Fartash Faghri, and Oncel Tuzel. 2024. CLIP with Quality Captions: A Strong Pretraining for Vision Tasks. *arXiv:2405.08911* (2024).
- [49] Feng Wang, Jieru Mei, and Alan Yuille. 2023. SCLIP: Rethinking Self-Attention for Dense Vision-Language Inference. *arXiv:2312.01597* (2023).
- [50] Qi Wang, Shaoteng Liu, Jocelyn Chanusot, and Xuelong Li. 2018. Scene classification with recurrent attention of VHR remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing* 57, 2 (2018).
- [51] Wenxuan Wang, Quan Sun, Fan Zhang, Yepeng Tang, Jing Liu, and Xinlong Wang. 2024. Diffusion Feedback Helps CLIP See Better. *arXiv:2407.20171* (2024).
- [52] Zifeng Wang, Zhenbang Wu, Dinesh Agarwal, and Jimeng Sun. 2022. MedCLIP: Contrastive learning from unpaired medical images and text. *arXiv:2210.10163* (2022).
- [53] Syed Waqas Zamir, Aditya Arora, Akshita Gupta, Salman Khan, Guolei Sun, Fahad Shahbaz Khan, Fan Zhu, Ling Shao, Gui-Song Xia, and Xiang Bai. 2019. iSAID: A Large-scale Dataset for Instance Segmentation in Aerial Images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*.
- [54] Gui-Song Xia, Jingwen Hu, Fan Hu, Baoguang Shi, Xiang Bai, Yanfei Zhong, Liangpei Zhang, and Xiaoqiang Lu. 2017. AID: A benchmark data set for performance evaluation of aerial scene classification. *IEEE Transactions on Geoscience and Remote Sensing* 55, 7 (2017).
- [55] Gui-Song Xia, Wen Yang, Julie Delon, Yann Gousseau, Hong Sun, and Henri Maitre. 2010. Structural high-resolution satellite image indexing. In *International Archives of Photogrammetry and Remote Sensing*, Vol. 38, 298–303.
- [56] Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. 2022. A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In *Proceedings of the European Conference on Computer Vision*.
- [57] Hongfeng Yu, Fanglong Yao, Wanxuan Lu, Nayu Liu, Peiguang Li, Hongjian You, and Xian Sun. 2022. Text-image matching for cross-modal remote sensing image retrieval via graph neural network. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 16 (2022).
- [58] Yuan Yuan, Yang Zhan, and Zhitong Xiong. 2023. Parameter-Efficient Transfer Learning for Remote Sensing Image-Text Retrieval. *arXiv:2308.12509* (2023).
- [59] Zhenghang Yuan, Lichao Mou, and Xiao Xiang Zhu. 2023. Multilingual augmentation for robust visual question answering in remote sensing images. In *Proceedings of the IEEE Joint Urban Remote Sensing Event*.
- [60] Zhiqiang Yuan, Wenkai Zhang, Kun Fu, Xuan Li, Chubo Deng, Hongqi Wang, and Xian Sun. 2022. Exploring a fine-grained multiscale method for cross-modal remote sensing image retrieval. *arXiv:2204.09868* (2022).
- [61] Zhiqiang Yuan, Wenkai Zhang, Changyuan Tian, Xuee Rong, Zhengyuan Zhang, Hongqi Wang, Kun Fu, and Xian Sun. 2022. Remote sensing cross-modal text-image retrieval based on global and local information. *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022).
- [62] Mert Yuksekgonul, Federico Bianchi, Pratyusha Kalluri, Dan Jurafsky, and James Zou. 2022. When and why vision-language models behave like bags-of-words, and what to do about it?. In *Proceedings of the International Conference on Learning Representations*.
- [63] Angelos Zavras, Dimitrios Michail, Begüm Demir, and Ioannis Papoutsis. 2024. Mind the Modality Gap: Towards a Remote Sensing Vision-Language Model via Cross-modal Alignment. *arXiv:2402.09816* (2024).
- [64] Andy Zeng, Maria Attarian, Brian Ichter, Krzysztof Choromanski, Adrian Wong, Stefan Welker, Federico Tombari, Aavek Purohit, Michael Ryoo, Vikas Sindhwani, et al. 2022. Socratic models: Composing zero-shot multimodal reasoning with language. *arXiv:2204.00598* (2022).
- [65] Yang Zhan, Zhitong Xiong, and Yuan Yuan. 2024. SkyEyeGPT: Unifying Remote Sensing Vision-Language Tasks via Instruction Tuning with Large Language Model. *arXiv:2401.09712* (2024).
- [66] Le Zhang, Rabiul Awal, and Aishwarya Agrawal. 2023. Contrasting Intra-Modal and Ranking Cross-Modal Hard Negatives to Enhance Visio-Linguistic Fine-grained Understanding. *arXiv:2306.08832* (2023).
- [67] Zilun Zhang, Tiancheng Zhao, Yulong Guo, and Jianwei Yin. 2024. RS5M and GeoRSCLIP: A Large Scale Vision-Language Dataset and A Large Vision-Language Model for Remote Sensing. *arXiv:2306.11300* (2024).
- [68] Bei Zhao, Yanfei Zhong, Gui-Song Xia, and Liangpei Zhang. 2015. Dirichlet-derived multiple topic scene classification model for high spatial resolution remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing* 54, 4 (2015).
- [69] Haozhe Zhao, Zefan Cai, Shuzheng Si, Xiaojian Ma, Kaikai An, Liang Chen, Zixuan Liu, Sheng Wang, Wenjuan Han, and Baobao Chang. 2023. MMICL: Empowering vision-language model with multi-modal in-context learning. *arXiv:2309.07915* (2023).
- [70] Lijun Zhao, Ping Tang, and Lianzhi Huo. 2016. Feature significance-based multibag-of-visual-words model for remote sensing image scene classification. *Journal of Applied Remote Sensing* 10, 3 (2016).
- [71] Zhuowen Tu Zheng Ding, Jieke Wang. 2023. Open-Vocabulary Universal Image Segmentation with MaskCLIP. In *Proceedings of the International Conference on Machine Learning*.
- [72] Chong Zhou, Chen Change Loy, and Bo Dai. 2021. Extract Free Dense Labels from CLIP. *arXiv:2112.01071* (2021).
- [73] Weixun Zhou, Shawn Newsam, Congmin Li, and Zhenfeng Shao. 2018. PatternNet: A benchmark dataset for performance evaluation of remote sensing image retrieval. *ISPRS Journal of Photogrammetry and Remote Sensing* 145 (2018).
- [74] Ziqin Zhou, Yinjie Lei, Bowen Zhang, Lingqiao Liu, and Yifan Liu. 2023. Zegclip: Towards adapting CLIP for zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [75] Qin Zou, Lihao Ni, Tong Zhang, and Qian Wang. 2015. Deep learning based feature selection for remote sensing scene classification. *IEEE Geoscience and Remote Sensing Letters* 12, 11 (2015).

A GENERATING TEXTUAL CAPTIONS FROM RETRIEVED EXAMPLES


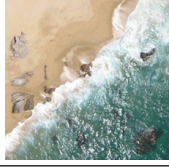

Image	Ground Truth	English	Portuguese	French
	green grass on both sides of the runway .	there are many lawns on both sides of the runway .	há muitas relva em ambos os lados da pista de pouso .	il y a de nombreux gazon sur les deux côtés de la piste .
	several stones are on the edge of the beach and the water .	a dark stone is on the margin of water and beach .	uma praia com algumas rochas e um mar verde-esmeralda .	une plage blanche à côté de quelques rochers et la mer .
	it is a park with a large pond in it next to a neighborhood .	a park with a green lake is close to the residential .	há um lago no parque à beira do rio, rodeado por uma área residencial .	un parc avec un lac vert est proche du logement .

Figure 5: Examples of captions generated with an approach based on LMCap [43], using an LLM for generating captions conditioned on retrieved examples obtained by RS-M-CLIP.

Although the proposed RS-M-CLIP model consists only of encoders, the generation of image captions can also be performed by integrating it with a language model decoder. We illustrate these capabilities, following a framework similar to LMCap [43], where captions with similar CLIP embeddings are retrieved for a given input image and used to prompt a Large Language Model (LLM) to generate a new caption, without receiving image features.

We used RS-M-CLIP in this setup to generate captions for different remote sensing image captioning datasets. Considering the multilingual abilities of RS-M-CLIP, we evaluated image captioning in different languages other than English, with two setups: (i) retrieving captions in English and prompting the LLM to generate a caption in the target language, and (ii) retrieving captions already in the target language and asking the LLM to generate a caption. Both setups used the aforementioned TowerInstruct LLM [1].

Table 6: LMCap image captioning results, obtained by retrieving similar captions to the image, either in English or in the target language, and asking the language model to generate a caption in the target language.

Language	RSICD			UCM-Captions			Sydney-Captions			NWPU-Captions		
	BLEU_1	BLEU_4	CIDEr	BLEU_1	BLEU_4	CIDEr	BLEU_1	BLEU_4	CIDEr	BLEU_1	BLEU_4	CIDEr
Supervised Methods for the English language												
MLCA-NET [9]	0.757	0.461	2.356	0.826	0.668	3.240	0.831	0.580	2.324	0.745	0.478	1.164
RSGPT [21]	0.703	0.368	1.029	0.861	0.657	3.332	0.823	0.622	2.731	-	-	-
SkyEyeGPT [65]	0.867	0.600	0.837	0.907	0.784	2.368	0.919	0.774	1.811	-	-	-
Examples in English and Generated Caption in the Target Language												
English	0.552	0.163	0.617	0.776	0.558	2.742	0.707	0.480	2.028	0.744	0.466	1.111
Portuguese	0.484	0.140	0.489	0.680	0.398	1.805	0.661	0.387	1.293	0.677	0.394	0.894
Spanish	0.518	0.157	0.523	0.681	0.406	1.835	0.641	0.387	1.361	0.699	0.406	0.944
French	0.492	0.151	0.493	0.682	0.396	1.475	0.644	0.384	1.178	0.691	0.402	0.859
German	0.406	0.104	0.319	0.619	0.320	1.226	0.597	0.305	1.090	0.643	0.341	0.688
Dutch	0.526	0.137	0.476	0.709	0.384	1.589	0.678	0.395	1.315	0.697	0.373	0.836
Italian	0.485	0.147	0.508	0.677	0.403	1.789	0.578	0.301	0.937	0.661	0.374	0.818
Korean	0.524	0.155	0.456	0.641	0.337	1.115	0.617	0.307	0.780	0.717	0.372	0.801
Chinese	0.603	0.213	0.655	0.770	0.475	1.697	0.732	0.427	1.053	0.752	0.429	1.002
Russian	0.318	0.069	0.243	0.561	0.253	0.897	0.498	0.195	0.601	0.566	0.293	0.574
Average	0.471	0.133	0.454	0.665	0.397	1.570	0.617	0.340	1.133	0.665	0.367	0.820
Examples and Generated Caption in the Target Language												
English	0.552	0.163	0.617	0.776	0.558	2.742	0.707	0.480	2.028	0.744	0.466	1.111
Portuguese	0.485	0.154	0.553	0.727	0.485	2.219	0.678	0.423	1.414	0.680	0.397	0.901
Spanish	0.528	0.168	0.573	0.723	0.470	2.238	0.677	0.416	1.438	0.701	0.407	0.955
French	0.509	0.171	0.549	0.717	0.458	1.945	0.615	0.386	1.287	0.689	0.397	0.849
German	0.432	0.115	0.375	0.702	0.437	1.734	0.617	0.344	1.435	0.645	0.341	0.696
Dutch	0.527	0.151	0.519	0.717	0.433	1.806	0.691	0.393	1.358	0.699	0.372	0.841
Italian	0.489	0.157	0.543	0.737	0.528	2.358	0.625	0.377	1.081	0.662	0.377	0.827
Korean	0.533	0.166	0.447	0.715	0.458	1.540	0.611	0.310	0.857	0.742	0.393	0.863
Chinese	0.582	0.207	0.640	0.754	0.499	1.792	0.708	0.413	1.111	0.748	0.427	0.993
Russian	0.337	0.085	0.275	0.646	0.409	1.438	0.561	0.277	0.967	0.568	0.292	0.570
Average	0.475	0.142	0.487	0.707	0.458	1.929	0.631	0.368	0.274	0.668	0.368	0.824

The prompt for multilingual image captioning is as follows, considering K retrieved captions as examples describing the image, and with translated captions into the target language in the case of scenario (ii): You are an intelligent image captioning bot tasked with describing remote sensing images. Similar images have the following captions: <caption 1>, ..., <caption K >. A creative short caption that can describe this image in <language> is:

The large language model is also prompted with a fixed set of N -shot examples of multiple image classes from the RSICD dataset, with K retrieved captions per example and with the output caption in the target language. The involved parameters were set to $N = 6$ and $K = 4$, where the N examples are each represented through the same textual prompt that was described above, and selected from different classes in the RSICD dataset (i.e., airport, denseresidential, baseballfield, parking, stadium, and playground). We found that the fixed examples could effectively be used across the different datasets, as evidenced by high captioning quality scores, particularly for English. In the second setup the entire prompt was also translated accordingly.

The results of the first setup can be seen in Table 6, with scores from previous English-only supervised methods also included for

comparison. Performance was evaluated with standard metrics from the COCOeval package³ with SacreBLEU tokenization⁴, namely BLEU-1, BLEU-4, and CIDEr. It can be seen that, for the English language, the proposed approach obtains reasonable performance across the datasets, and is very similar to MLCA-Net [9] for the NWPU-Captions dataset. Results are also consistent across languages. In turn, the results of the second setup, given at the bottom of Table 6, show that the image captioning performance is generally improved across languages, confirming the strong multilingual abilities of RS-M-CLIP. The setup with retrieved examples already in the target language helps the LLM to generate better captions.

Figure 5 shows some examples from the results of these experiments, for the English, Portuguese, and French languages. In the first row, a consistent description of the airport is given across all three languages. Regarding the image with a beach, the descriptions differ slightly, in particular emphasizing differently the colors of objects of interest (i.e., the color of the stone in English, while the color of the sea is mentioned in Portuguese, and the color of the beach in French). Lastly, a similar description is generated across all three languages for the third example image.

³<https://github.com/salaniz/pycocoevalcap>

⁴<https://github.com/mjpost/sacrebleu>