
IGLUE: A Benchmark for Transfer Learning across Modalities, Tasks, and Languages

Emanuele Bugliarello^{1,2} Fangyu Liu³ Jonas Pfeiffer^{4,5} Siva Reddy^{2,6} Desmond Elliott¹
Edoardo Maria Ponti^{2,6} Ivan Vulić³

Abstract

Reliable evaluation benchmarks designed for replicability and comprehensiveness have driven progress in machine learning. Due to the lack of a multilingual benchmark, however, vision-and-language research has mostly focused on English language tasks. To fill this gap, we introduce the Image-Grounded Language Understanding Evaluation benchmark. IGLUE brings together—by both aggregating pre-existing datasets and creating new ones—visual question answering, cross-modal retrieval, grounded reasoning, and grounded entailment tasks across 20 diverse languages. Our benchmark enables the evaluation of multilingual multimodal models for transfer learning, not only in a zero-shot setting, but also in newly defined few-shot learning setups. Based on the evaluation of the available state-of-the-art models, we find that translate-test transfer is superior to zero-shot transfer and that few-shot learning is hard to harness for many tasks. Moreover, downstream performance is partially explained by the amount of available unlabelled textual data for pretraining, and only weakly by the typological distance of target–source languages. We hope to encourage future research efforts in this area by releasing the benchmark to the community.

1. Introduction

Until recently, advances in multimodal vision-and-language (V&L) modelling have predominantly focused on English or a few high-resource Indo-European languages (Elliott et al., 2016), disregarding the enormous diversity of the world’s

languages and cultures (Bender, 2011; Ponti et al., 2019; Liu et al., 2021), which exacerbated the Anglo-centric bias (Jauhar et al., 2018; Ponti et al., 2020) of V&L models.

Nonetheless, this trend is reversing by virtue of a series of independent efforts in the the V&L research community: on top of expanding multilingual image–sentence retrieval to lower-resource languages (Srinivasan et al., 2021) and additional modalities (Armitage et al., 2020), we have recently witnessed pioneering work on multimodal machine translation (Yao & Wan, 2020; Huang et al., 2020), multilingual visual question answering (Pfeiffer et al., 2022), multilingual text-to-video search (Huang et al., 2021; Lei et al., 2021), multilingual visual reasoning (Liu et al., 2021), and cross-modal retrieval (Jain et al., 2021). Yet, multilingual multimodal models (Ni et al., 2021; Jain et al., 2021) are typically tested only on the task of cross-modal retrieval.

In order to incentivise and guide future research in multilingual V&L research, we introduce the **Image-Grounded Language Understanding Evaluation** (IGLUE) benchmark. We create IGLUE by collating current research threads in this area and extending them with two datasets for cross-lingual visual entailment (XVNLI) and image–text retrieval (xFlickr&CO), for an even more comprehensive coverage of tasks. IGLUE is the first evaluation suite for multilingual multitask V&L modelling, comprising five datasets across four structurally different tasks that require different levels of syntactic-semantic V&L understanding in cross-lingual setups: cross-modal retrieval, visual question answering, natural language inference, and visual reasoning (Figure 1). Aiming at a representative selection of languages, IGLUE covers 20 typologically diverse languages, spanning 11 language families, 9 scripts, and 3 WALS-defined (Dryer & Haspelmath, 2013) geographical macro-areas.

Similar to previous text-based, cross-lingual transfer benchmarks, such as XGLUE (Liang et al., 2020) and XTREME (Hu et al., 2020), our V&L benchmark is designed to evaluate zero-shot transfer scenarios, where annotated training data is provided in English, but none in the target language. In addition, contrary to the above-mentioned datasets, IGLUE also provides standardised data splits to guide cross-lingual few-shot learning experiments, which

¹University of Copenhagen ²Mila – Quebec Artificial Intelligence Institute ³University of Cambridge ⁴TU Darmstadt ⁵New York University ⁶McGill University. Correspondence to: Emanuele Bugliarello <emanuele@di.ku.dk>.

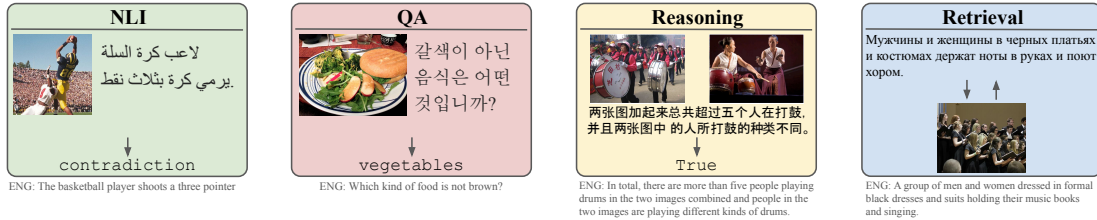


Figure 1. Overview of the tasks in IGLUE, which include grounded natural language inference, visual question answering, grounded reasoning, and cross-modal retrieval. Each task is associated with an example of input and output (English translations at the bottom).

can help reduce the gap between zero-shot and supervised performance (Lauscher et al., 2020; Zhao et al., 2021). In this setup, a small number of task-annotated examples for fine-tuning are available for a given target language. We also release machine-translated versions of the test sets to enable the evaluation of ‘translate test’ cross-lingual transfer.

By virtue of the newly created IGLUE benchmark, we also run the first systematic comparative evaluation of cutting-edge multilingually pretrained V&L architectures (Ni et al., 2021; Zhou et al., 2021; Liu et al., 2021), as well as a series of representative monolingual V&L models combined with machine translation (Lu et al., 2019; Chen et al., 2020, *inter alia*) across a range of diverse V&L tasks and languages.

This evaluation offers new insights on the models’ core strengths and current limitations. Foreshadowing, we showcase large gaps between performance in English and transfer performance, where (to a limited extent) the gaps are more prominent in lower-resource languages and languages more distant from English. In addition, leveraging target-language in-task few shots is remarkably arduous: whereas in previous text-only experiments (Lauscher et al., 2020), *ceteris paribus* there are huge gains of few-shot scenarios over their zero-shot counterparts, we demonstrate that current V&L models often require thousands of examples before showing signs of improvement. Finally, performance also seems correlated with task difficulty: for instance, NLI shows the smallest gaps between English and other languages and QA benefits the most from few-shot adaptation. On the other hand, visually grounded reasoning and cross-modal retrieval appear to be harder in both respects.

Contributions. **1)** In order to guide and inspire more work in the area of multilingual V&L research, we present a first evaluation benchmark for cross-lingual transfer learning for V&L tasks, spanning 20 languages, 5 datasets, and 4 different tasks. **2)** In the process of benchmark creation, we complement existing datasets with new training and evaluation data in several languages to increase diversity and enable few-shot learning, and introduce a first multilingual dataset for visually grounded cross-lingual NLI. **3)** We conduct systematic evaluations of representative V&L architectures in zero-shot and few-shot cross-lingual transfer scenarios, offering standard data splits and empirical

baselines for future research. **4)** Our results and additional analyses take stock of the current gaps and challenges in cross-lingual V&L research. **5)** To further facilitate replicable research in this area, we re-implement the existing multilingual V&L pretrained encoders in a unified framework (VOLTA; Bugliarello et al. 2021), which also provides access to five English V&L BERTs and 12 V&L tasks. We provide data and code for the evaluation of multilingual V&L models at <https://iglu-benchmark.github.io/>.

2. Related Work and Motivation

Multilingual Multimodal Learning. Multilingual V&L research focuses on collecting resources, developing models, and evaluating systems that need to jointly reason over multilingual text and multimodal inputs, this way combining two areas of crucial importance: multilingual (Snyder & Barzilay, 2010; Ponti et al., 2019) and multimodal learning (Bernardi et al., 2016; Baltrusaitis et al., 2019). A natural overlap of the two areas is *language grounding* in perception (typically vision; Deng et al., 2009; Kiela et al., 2018), where the perceptual input can be considered as an inherent cross-lingual signal (Kiela et al., 2015; Gella et al., 2017; Caglayan et al., 2021). Many other research goals lie at the intersection of these two areas, such as transfer learning and modularity (Ponti et al., 2021b; Ansell et al., 2022).

However, until recently the pace of progress in multilingual multimodal learning has not gained as much momentum as V&L in monolingual settings, mostly due to the scarcity of resources for training and evaluation. Texts in most multimodal datasets are usually only available in English, or in high-resource languages (e.g., Chinese and a few major Indo-European languages; Elliott et al., 2016; 2017; Barault et al., 2018; Wang et al., 2019b). Consequently, models trained on such datasets do not take into account linguistic diversity (Ponti et al., 2020) or cross-cultural nuances (Armitage et al., 2020; Liu et al., 2021; Yin et al., 2021).¹

The need to expand V&L research towards more languages

¹Some notable exceptions, which did reach beyond English, were still limited to the image–text retrieval task and mostly Indo-European languages (Rotman et al., 2018; Wehrmann et al., 2019; Huang et al., 2019; Kim et al., 2020; Su et al., 2021).

has been recognised by 1) the recent creation of multilingual training and evaluation data across diverse V&L tasks and languages (Armitage et al., 2020; Srinivasan et al., 2021; Su et al., 2021; Pfeiffer et al., 2022; Liu et al., 2021; Wang et al., 2021, *inter alia*), as well as 2) the emergence of the first large multilingual-multimodal pretrained models (Ni et al., 2021; Zhou et al., 2021; Liu et al., 2021) and monolingual V&L models adapted to multiple languages (Chen et al., 2020; Pfeiffer et al., 2022). In this work, we merge and expand on these two research threads, aiming to highlight current achievements and challenges in this area and to facilitate comparative evaluations, thus bringing together the above-mentioned collective research efforts.

Multi-Task Evaluation Benchmarks in NLP. IGLUE has been inspired by recent text-only *multi-task* benchmarks for natural language understanding: these benchmarks have been proven invaluable as key drivers of recent steep performance progress of NLP. The creation of such benchmarks in monolingual settings, sparked by the pioneering and now omnipresent English-only GLUE (Wang et al., 2018) and SuperGLUE (Wang et al., 2019a), has been extended to other languages such as Indonesian (Wilie et al., 2020), Korean (Park et al., 2021), Russian (Shavrina et al., 2020), and Romanian (Dumitrescu et al., 2021).

IGLUE is even more related to multi-task benchmarks developed for *cross-lingual transfer settings*: XTREME (Hu et al., 2020), XGLUE (Liang et al., 2020), and XTREME-R (Ruder et al., 2021). They have brought in the spotlight the necessity to evaluate not only on a diverse set of tasks, but also on a diverse set of languages, in order to incentivise research on models that forgo monolingual limitations and generalise well in multilingual settings.

With IGLUE, we make the first leap into multimodal cross-lingual evaluation. Driven by analyses of text-only benchmarks and lessons learned from them (Ethayarajh & Jurafsky, 2020), IGLUE aims to extend sheer accuracy-driven evaluation towards other crucial aspects such as fine-tuning efficiency, sample efficiency and adaptation to low-data scenarios, and enhance multilingual inclusivity and diversity.

3. IGLUE Benchmark

Design Principles. Aiming for a comprehensive resource, we take inspiration from the best practices and design principles of existing multilingual benchmarks (Hu et al., 2020; Liang et al., 2020; Ruder et al., 2021), and adapt them to account for the unique challenges of V&L tasks. In particular, we abode by the following principles:

(1) *Task Diversity.* We selected a wide spectrum of multimodal tasks that reflect multiple facets of V&L learning.

(2) *Language Diversity.* The datasets included in IGLUE

Table 1. IGLUE statistics. † Aggregated across languages. Language counts include English. MaxS: Max-Shot. F: Flickr; VG: Visual Genome; N: NLVR2; M: MaRVL; C: COCO; W: Wikipedia

Task	NLI	QA	Reason	Retrieval	
Train/Dev data	SNLI	GQA	NLVR2	Flickr30K	WIT _{EN} ^{500K}
Test data	XVNLI	xGQA	MaRVL	xFlickr&CO	WIT
Languages	5	8	6	8	11
Img source	F	VG	N + M	F + C	W
Train images	30K	72K	87K	29K	469K
Train samples	541K	943K	86K	145K	500K
Dev images	1K	10K	7K	1K	4.6K
Dev samples	18K	132K	8K	5K	4.6K
Test images	357	300	†4.9K	2K	†6.2K
Test samples	1.1K	9.6K	†5.7K	2K	†9.6K
MaxS images	48	48	80	100	-
MaxS samples	168	1.5K	80	100	-

should cover many languages that are diverse in terms of family, geographic area, typological features, and script.

(3) *Accessibility.* The data should be released under a licence that permits use and redistribution for research purposes.

Compared to previous text-only benchmarks, V&L tasks are more time- and compute-intensive by nature. To widen usability, we created IGLUE by carefully choosing datasets and training setups that would enable quick development even by practitioners constrained by limited resources.

3.1. Tasks and Datasets

Challenges in V&L learning are multi-faceted and entangled with distinct abilities: e.g., drawing inferences, reasoning over and comparing images, answering questions, and retrieving images or their captions. In IGLUE we thus represent all these abilities together with their corresponding tasks, see Table 1 and Figure 1 for an overview. We opt for the following V&L tasks based on (i) (partial) availability of multilingual data, and (ii) computational requirements.

XVNLI. We propose the new task of Cross-lingual Visual Natural Language Inference (XVNLI). It requires the model to predict if a text-*hypothesis* ‘entails’, ‘contradicts’, or is ‘neutral’ to an image-*premise*. We combine the text-only dataset SNLI (Bowman et al., 2015), with its multimodal (Xie et al., 2019) and cross-lingual (Agić & Schluter, 2018) counterparts.² We provide new train, development, and test splits such that the test split consists of images covered by the underlying cross-lingual text-only dataset. We discard all text examples of images in the test set for which no cross-lingual version exists. The remaining images for which only English hypothesis-examples exist are randomly split into the train and development set. To mitigate data leakage between the splits, we sample based on the images; all text

²Vu et al. (2018) formulate grounded textual entailment as the task where the image supplements the textual premise and hypothesis. We follow the established setup of Xie et al. (2019).

examples corresponding to one image are added to one split.

xGQA. To solve the Cross-lingual Grounded Question Answering task (xGQA; Pfeiffer et al. 2022), a model must answer several types of *structured* questions about an image. The evaluation data in 7 target languages are manually translated from the validation set of GQA (Hudson & Manning, 2019), whereas the training data are sourced from the English training set of GQA. All images were originally sampled from Visual Genome (Krishna et al., 2017). In particular, we use the English *balanced* training set to train our models, and evaluate on the *few-shot* evaluation sets defined by Pfeiffer et al. (2022) to allow for direct comparison between zero-shot and few-shot experiments.

MaRVL. The Multicultural Reasoning over Vision and Language dataset (Liu et al., 2021) requires to determine whether a textual description is *true* or *false* about a pair of images. This involves comparing their visual representations and reasoning about the facts in the description. The creation of MaRVL is entirely driven by native speakers. As a consequence, the descriptions are written from scratch and images are selected to be culturally relevant. The NLVR2 data (Suhr et al., 2019) in English are used for training.

xFlickr&CO. We create a new, multilingual evaluation set for retrieval by combining 1,000 images from the Flickr30K (Young et al., 2014) and 1,000 from the COCO (Lin et al., 2014) test split defined by Karpathy & Fei-Fei (2015) (henceforth, *Karpathy split*). Each image is then associated with a single caption.³ The English evaluation set is obtained by sampling one caption per image from the existing datasets. With the exception of Japanese (Yoshikawa et al., 2017; Nakayama et al., 2020), no other language covers the evaluation splits of both Flickr and COCO.⁴ We hence crowdsource image descriptions in 6 other languages

³We expand the Flickr30K evaluation set since state-of-the-art model performance is currently close to the ceiling of human performance (Jia et al., 2021). We show that this leads to a more challenging evaluation set which requires out-of-distribution generalisation, while maintaining an efficient training regime.

⁴Multilingual captions for Flickr30K also exist in Dutch (van Miltenburg et al., 2018) and Mandarin (Lan et al., 2017), but the Mandarin captions do not match the expected test split. COCO has also been (partially) translated to Hindi (Mishra et al., 2021), Italian (Scaiella et al., 2019), Mandarin (Li et al., 2019b), Spanish (García, 2020) and Vietnamese (Lam et al., 2020), with similar problems. The Mandarin captions have even been used to train and evaluate M³P and UC² in a multilingual setup. We find that 882/1,000 test images are used as training examples in other languages, and that 148 of its training images are in the expected test split. In a similar vein, we find that 946/1,000 test samples in XTD (Aggarwal et al., 2021)—which contains COCO translations in seven languages (Italian, Korean, Mandarin, Portuguese, Russian, Spanish and Turkish)—are from the COCO training split, while only 7 of them are also in the Karpathy COCO test split. These overlaps lead to considerable leakage of visual and cross-lingual information during fine-tuning and evaluation.

by asking annotators to describe each image, rather than translating its English caption, to prevent biasing their description (van Miltenburg et al., 2017; Frank et al., 2018). We ask annotators to follow the Flickr30K guidelines (see Appendix A.4 for details). For the Flickr30K subsplit in German, we sample captions from the Multi30K task 2 data (Elliott et al., 2016).

WIT. The Wikipedia-based Image Text dataset (Srinivasan et al., 2021) collected examples from Wikipedia in 108 languages. Similar to xFlickr&CO, the tasks consist in retrieving the correct image given a textual description (image retrieval) and vice versa (text retrieval). WIT represents a very diverse set of concepts and real world entities. Text fields in WIT often tend to be descriptive, verbose and use specific terminology, different from single-line captions of common words and objects in Flickr30K and COCO. This allows us to evaluate the high-level image–sentence understanding of multilingual V&L models across a wider range of languages and real-world entities than in xFlickr&CO. For training, we randomly sample a subset of 500K captions from the English training set of WIT. For evaluation, we use the WIT test data released as part of its corresponding Kaggle competition.⁵ In particular, as most of the test languages are Indo-European, we choose the 4 national languages with the lowest Wikipedia coverage (measured in number of articles as of Jan 2022), and 6 more languages to both cover different properties (language families, scripts) and overlap with other tasks in IGLUE.⁶ We ensure that each language has at least 500 image–caption pairs, and limit English and Japanese data to 1,000 samples (see also Appendix A.5).

3.2. Languages

IGLUE covers 20 languages (Table 2). They are extremely diverse, as they comprise 11 distinct language families, and span across 3 of the 5 geographic macro-areas defined by Dryer & Haspelmath (2013). Moreover, they are written in a variety of scripts: Arabic, Bengali–Assamese, Chinese characters, Cyrillic, Greek, Hangul, Kanji, Latin, and Tamil. The distribution of languages per task is also shown in Table 2.

We ensured that IGLUE includes languages with different levels of unlabelled data available (Joshi et al., 2020; Blasi et al., 2022). Thus, it allows for evaluating models with different data paucity regimes during pretraining. While not all languages are covered in each task—due to the (partly) independent selection of languages in the original datasets—

⁵www.kaggle.com/c/wikipedia-image-caption.

⁶While Jain et al. (2021) define a set of 8 low-resource languages for evaluation on WIT, they were sampled from the training set. Moreover, we find that several of their test images contain pictures with identifiable people, which would complicate its public release due to regulations such as the GDPR in Europe. By only relying on the true test split, we allow researchers to easily adopt the full WIT training dataset in their experimental pipelines.

Table 2. Benchmark languages and tasks. English is only used for training. Tasks legend: ✓✓ train and test sets available; ✓ test-only data available; * Japanese captions in xFlickr&CO are translations from English. Languages legend: A: Asiatic; C: Congo; E: European; T: Tibetan; Austron: Austronesian. Language codes are based on the ISO 639-3 international standard.

Language			NLI	QA	Reasoning	Retrieval
Name	Code Family	Script	XVNLI	xGQA	MaRVL	xFlickr&CO WIT
English	ENG Indo-E	Latin	✓✓	✓✓	✓✓	✓✓
Arabic	ARB Afro-A	Arabic	✓✓			✓✓
Bengali	BEN Indo-E	Bengali		✓✓		
Bulgarian	BUL Indo-E	Cyrillic				✓
Danish	DAN Indo-E	Latin				✓
Estonian	EST Uralic	Latin				✓
German	DEU Indo-E	Latin		✓✓		✓
Greek	ELL Indo-E	Greek				✓
French	FRA Indo-E	Latin	✓✓			
Indonesian	IND Austron	Latin		✓✓	✓✓	✓
Japanese	JPN Japonic	Kanji				*✓✓
Korean	KOR Koreanic	Hangul		✓✓		✓
Mandarin	CMN Sino-T	Hanzi		✓✓	✓✓	✓
Portuguese	POR Indo-E	Latin		✓✓		✓
Russian	RUS Indo-E	Cyrillic	✓✓	✓✓		✓
Spanish	SPA Indo-E	Latin	✓✓			✓
Swahili	SWA Niger-C	Latin			✓	
Tamil	TAM Dravidian	Tamil			✓✓	
Turkish	TUR Turkic	Latin			✓✓	✓
Vietnamese	VIE Austro-A	Latin				✓

it is worth noting that 10 out of 20 languages have data for two or more tasks. Thus, IGLUE might also facilitate future research in cross-task knowledge transfer.

3.3. From Zero-Shot to Few-Shot Setups

The established practice in multilingual text-based benchmarks (Hu et al., 2020; Liang et al., 2020; Ruder et al., 2021) is to frame cross-lingual transfer as a *zero-shot* learning problem. However, multilingual pretrained models can be additionally fine-tuned in a *few-shot* learning setup; that is, on a handful of data points in a target language annotated for a specific task. While computationally inexpensive, this strategy has been proven as very beneficial for performance in text-only tasks, especially on low-resource and distant languages (Lauscher et al., 2020; Ponti et al., 2021a).

We hence extend IGLUE to support few-shot learning setups by collecting samples for each V&L task. In doing so, we use the notion of ‘annotation context’ to define what a ‘shot’ means in each dataset: For instance, this corresponds to an image and its caption in cross-modal retrieval, and to an image and all its questions in visual question answering. For more details and statistics, we refer to Appendix A.

Notably, the set of examples used for few-shot cross-lingual transfer may vary across experiments. The lack of a standard set for few-shot learning harms the replicability and comparability of such experiments, given that model performance exhibits a strong sensitivity to the selection of few-shot examples (Zhao et al., 2021). Therefore, following Schick & Schütze (2021) and Pfeiffer et al. (2022), in addition to

the train / validation / test splits, we also release standard few-shot splits for every task and language (Appendix A).

4. Evaluation Framework

The new IGLUE benchmark allows for running a series of unprecedented comparative experiments and analyses. Here, we provide an overview of our experimental setup. We evaluate all the existing multilingual V&L pretrained models released so far. In particular, in order to further facilitate research and development in multilingual V&L modelling, we re-implement them in a single framework based on VOLTA (Bugliarello et al., 2021)⁷ in PyTorch (Paszke et al., 2019), which also provides access to five English pretrained models and twelve downstream tasks.⁸ For further details about the framework, we refer to Appendix B.

Pretrained Models. The state-of-the-art multilingual V&L models that we evaluate follow a BERT-like architecture (Devlin et al., 2019). Their input is the concatenation of image region features extracted with a Faster R-CNN (Ren et al., 2015) object detector and textual tokens (Sennrich et al., 2016; Wu et al., 2016). The inputs are then processed by a single stack of Transformer layers (Vaswani et al., 2017) to obtain multimodal, contextualised representations. Each model is trained to minimise multiple objectives—with minimal variations across models—to learn how to understand text (masked language modelling), vision (masked region modelling) and their co-occurrence (image–text matching). We provide only condensed summaries here, and refer the reader to the original work for further details.

- **mUNITER and xUNITER.** Liu et al. (2021) extend the UNITER architecture (Chen et al., 2020) multilingually by following the base approach of Ni et al. (2021): a batch of multimodal English data from Conceptual Captions (CC; Sharma et al. 2018) is alternated with a batch of text-only multilingual Wikipedia data (sampled from the 104 languages used for mBERT). The two models mainly differ in their initialisation: mUNITER from mBERT (Devlin et al., 2019) and xUNITER from XLM-R (Conneau et al., 2020).
- **M³P.** Ni et al. (2021) further introduce multimodal code-switched training tasks where words in English captions are randomly replaced with a translation with a certain probability. This setup allows the model to explicitly align images with non-English languages. In particular, the authors obtain word translations from bilingual dictionaries in 50 languages. In each multimodal batch, data from CC is fed to the model either fully in English or with code-switched words according

⁷<https://github.com/e-bug/volta>.

⁸Code to reproduce our results is available online at <https://github.com/e-bug/igluue>.

Table 3. Performance of multimodal models trained and evaluated on English test splits of the IGLUE tasks.

Model	NLI		QA		Reasoning		Retrieval	
	XVNLI	xGQA	MaRVL (NLVR2 Test-P)	xFlickr&CO IR	WIT TR	IR	TR	
mUNITER	76.38	54.68	71.91	44.50	40.90	19.90	22.34	
xUNITER	75.77	54.83	71.55	38.45	32.05	16.70	18.54	
UC ²	76.38	55.19	70.56	37.40	34.55	17.90	19.71	
M ³ P	76.89	53.75	68.22	31.35	24.60	15.50	15.33	
LXMERT	76.72	53.04	68.95	36.85	29.60	18.00	18.50	
UNITER	77.15	55.59	71.11	43.05	40.70	18.10	19.00	
ViLBERT	77.15	55.82	71.02	38.70	35.70	17.90	20.20	
VisualBERT	77.15	53.67	71.18	43.95	39.85	18.30	20.50	
VL-BERT	77.49	55.74	71.72	39.90	34.15	20.10	21.17	

to a given sampling ratio. The text-only multilingual batch is sampled from Wikipedia in 100 languages. The model is initialised from XLM-R.

- UC². Zhou et al. (2021) rely on (text-only) machine translation technologies to obtain CC data in five languages (Czech, French, German, Japanese, and Mandarin). The model is then solely pretrained on multilingual multimodal batches, where, for each image, a caption is sampled uniformly from the available languages. Two pretraining objectives are added to those above, to (i) tighten region-token matching and (ii) align translations in a similar vein as Conneau & Lample (2019) and Caglayan et al. (2021). The approach also maps the object detector space onto the language model space before multimodal pretraining. As xUNITER and M³P, UC² is initialised from XLM-R.

In addition to evaluating all the released multilingual V&L encoders, we also benchmark representative English V&L encoders pretrained in a controlled setup (Bugliarello et al., 2021) in combination with ‘translate test’ transfer:⁹ LXMERT (Tan & Bansal, 2019), UNITER (Chen et al., 2020), ViLBERT (Lu et al., 2019), VisualBERT (Li et al., 2019a) and VL-BERT (Su et al., 2020).

Experimental Setup. We fine-tune all models using the AdamW optimiser (Loshchilov & Hutter, 2019) relying on the same hyper-parameters as in the controlled setup of Bugliarello et al. (2021). For few-shot experiments, we instead search three learning rates {1e-5, 5e-5, 1e-4} and train for 20 epochs for each *dataset-language-shots* triplet. Before training, we extract, for each dataset, 36 image regions using a ResNet-101 backbone (He et al., 2016) for mUNITER, xUNITER and UC²; and up to 100 image regions using a ResNeXt-101 backbone (Xie et al., 2017) for M³P—both trained on Visual Genome (Krishna et al., 2017; Anderson et al., 2018). We train all models on a single NVIDIA V100 (16GB) GPU card. To increase accessibility to the benchmark by practitioners with limited computing resources: **1**) we set the number of epochs so that training a

⁹We obtain translations to English via Google Translate.

Table 4. Average zero-shot and ‘translate test’ performance of multimodal models trained in English and evaluated on target languages in IGLUE. The best model for each method is in bold.

Model	NLI		QA		Reasoning		Retrieval	
	XVNLI	xGQA	MaRVL	xFlickr&CO IR	WIT TR	IR	TR	
<i>Zero-Shot</i>								
mUNITER	53.69	9.97	53.72	8.06	8.86	9.16	10.48	
xUNITER	58.48	21.72	54.59	14.04	13.51	8.72	9.81	
UC ²	62.05	29.35	57.28	20.31	17.89	7.83	9.09	
M ³ P	58.25	28.17	56.00	12.91	11.90	8.12	9.98	
<i>Translate test</i>								
mUNITER	73.09	49.05	63.82	40.95	36.78	16.25	16.64	
xUNITER	72.83	49.15	64.04	36.26	30.29	13.01	14.20	
UC ²	73.67	50.19	63.09	36.03	30.37	12.70	14.11	
M ³ P	73.37	48.83	62.52	27.74	21.29	11.53	13.63	
LXMERT	72.57	48.08	62.51	34.02	26.66	14.28	14.86	
UNITER	73.65	50.62	61.92	41.04	37.49	15.43	16.01	
ViLBERT	73.45	50.33	62.39	36.97	33.21	15.40	16.93	
VisualBERT	74.12	48.72	62.35	41.64	36.44	15.36	15.75	
VL-BERT	73.86	49.78	64.16	38.18	31.84	15.11	16.09	

baseline can be completed in less than 12 hours per dataset (see Appendix B for details);¹⁰ **2**) in addition to reporting performance at every number of shots for few-shot learning, we also define a max-shot evaluation setup to reduce the numbers of runs to one per dataset–language pair.¹¹ For both fine-tuning and few-shot experiments, we evaluate the parameter sets that yield the best validation performance. For few-shot learning, models are initialised from the English fine-tuned parameter set and further trained on the data in the target language for the given task.

5. Main Results and Discussion

5.1. Zero-shot Learning

The main zero-shot results, averaged across languages, are reported in Table 4. In order to measure the gap between cross-lingual transfer and supervised learning, we also provide the results for the same models trained and evaluated on English test data in Table 3. The performance metric is accuracy for all tasks except cross-modal retrieval, which uses Recall@1.¹² Finally, for completeness, the full results subdivided by language are detailed across Appendix C.2. Next, we compare these scores across different dimensions.

Transfer Method. The results in Table 4 clearly demon-

¹⁰The only exception is M³P, which is twice as big as the others.

¹¹Notably, the main bottleneck for current models is the retrieval task as this requires pairing each image with every caption (Geigle et al., 2022). In xFlickr&CO, it takes us ≈ 3 hours per language.

¹²Note that these metric are identical whenever image–caption pairs are unique. This is the case for all of xFlickr&CO data points and for more than 90% of the samples in WIT (Appendix A.5).

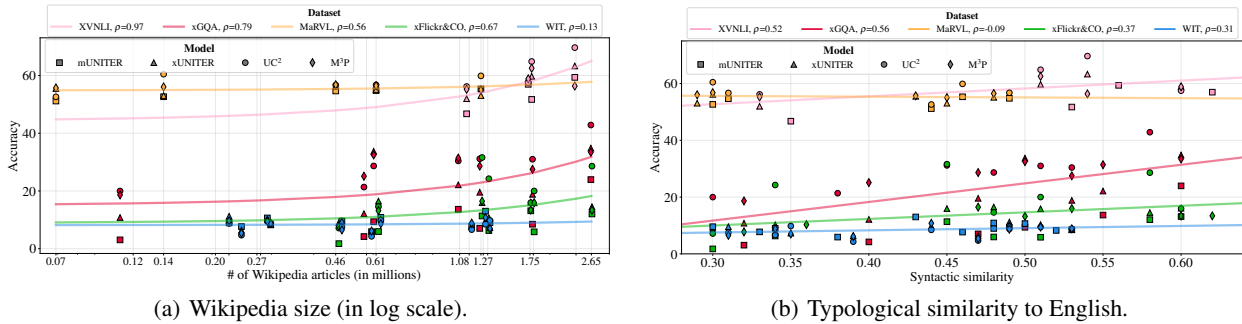


Figure 2. Zero-shot performance on individual target languages as a function of some of their properties: Wikipedia size (left) and typological similarity to English (right). Pearson’s correlation coefficients are also reported in the legends.

strate the superiority of ‘translate test’ transfer over zero-shot model transfer via multilingual encoders. The difference between the best models for each method reaches 12.1 points for XVNLI, 21.6 for xGQA, 6.9 for MaRVL, 21.3 for xFlickr&CO IR, 19.6 for xFlickr&CO TR, 7.1 for WIT IR, and 6.5 for WIT TR. Thus, gains are especially remarkable in cross-modal retrieval tasks and grounded question answering, where the performance is nearly doubled by virtue of ‘translate test’ transfer.

Multilingual Models. Even within each transfer method, we observe considerable variance among individual multilingual multimodal models. For instance, for zero-shot model transfer via multilingual encoders, UC² is consistently better across the board (except for WIT) and surpasses the other models by a large margin, especially in German, French, Japanese and Mandarin—languages in which the authors had translated CC and pretrained UC² on (Appendix C.2). Remarkably, this tendency is less accentuated in MaRVL, where pretraining Mandarin multimodal data is still insufficient to tackle the out-of-distribution nature of MaRVL’s culture-specific concepts (Table 15 in Appendix C.2). Nonetheless, these results prove how the simple ‘translate pretrain’ approach of UC² can be an effective baseline for multilingual transfer in multimodal pretrained models. As for ‘translate test’ transfer, no clear winner emerges across the monolingual models, whereas mUNITER generally performs better among multilingual models limited to cross-modal retrieval tasks. We also note that massively multilingual models (top rows) are often on a par with their monolingual counterparts (bottom rows).

Transfer Gap. Comparing the results on the English test set (Table 3) with those averaged across the other languages (Table 4) reveals an extremely large gap in performance due to cross-lingual transfer. Considering the best multilingual encoders for each task, the gap is 14.8 points for XVNLI, 26.5 for xGQA, 14.6 for MaRVL, 24.2/23.0 for xFlickr&CO IR/TR, and 10.7/11.9 for WIT IR/TR.

Task Complexity. There emerge stark contrasts among the model performances in each individual task. Nonetheless,

we remind the reader that the number of classes to predict—and hence random baselines—vary across tasks: XVNLI 3, xGQA 1,842, and MaRVL 2. Moreover, some tasks include a collection of languages that are more homogeneous (e.g. XVNLI contains mostly Indo-European languages) whereas others are more diverse, such as WIT. Performance is expected to suffer the most in the latter group. Moreover, focusing on cross-modal retrieval tasks, scores are higher in xFlickr&Co than in WIT. This is arguably due to the fact that, while all models were pretrained on an image–text matching task, WIT captions are distributionally distant from standard captions as they mostly describe entities.

Explanatory Variables. Finally, we assess the impact of a series of explanatory factors on the performance in each individual target language. More specifically, we estimate its correlation (1) with the amount of each language’s unlabelled textual data available for pretraining in terms of Wikipedia articles (as of January 2022) in Figure 2(a); and (2) with the typological similarity between the source (English) and the target language in cross-lingual transfer in Figure 2(b). The latter is calculated as the cosine similarity between vectors of morpho-syntactic features extracted from WALS (Dryer & Haspelmath, 2013) by Littell et al. (2017).¹³ According to the received wisdom from text-only cross-lingual transfer (Ponti et al., 2020; Ruder et al., 2021), both these factor usually play a pivotal role in determining the downstream performance of each target language. While we find this to partially hold true for Wikipedia size (with the possible exception of WIT), we find that correlations of typological similarity are weaker (and even negative for MaRVL), implying that models may struggle comparatively on many target languages despite different similarities to English.

5.2. Few-shot Learning

As we experimented with a range of data sizes for k -shot fine-tuning in the target language, we (1) plot these values against the corresponding performance in Figure 3 and (2) also report the exact scores for the maximum k in Table 5,

¹³<https://github.com/antonisa/lang2vec>.

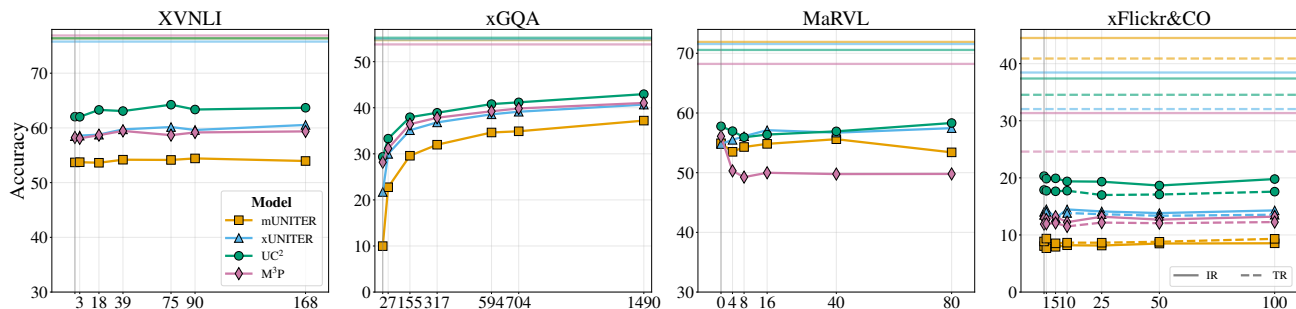


Figure 3. Few-shot performance on IGLUE tasks (cross-lingual average) against different data sizes for target language fine-tuning. Values in MaRVL are only averaged across the 3 languages with available few-shot data. Horizontal lines report English performance.

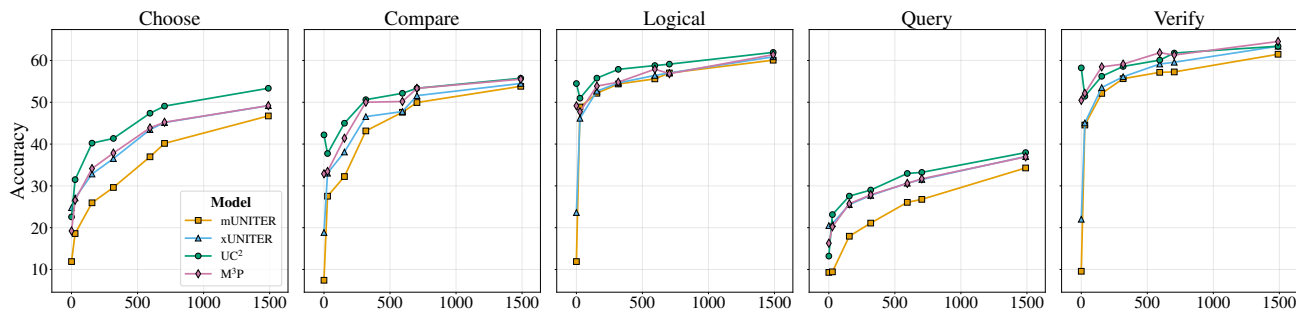


Figure 4. Analysis of few-shot results on xGQA over five structurally different questions types.

Table 5. Max-shot performance (with the largest k) averaged across languages with few-shot training data.

Model	NLI		QA		Reasoning		Retrieval	
	XVNLI	xGQA	MaRVL	xFlickr&CO	IR	TR		
mUNITER	53.95	37.21	53.41	8.54	9.32			
xUNITER	60.55	40.68	57.46	14.30	13.54			
UC ²	63.68	42.95	58.32	19.79	17.59			
M ³ P	59.36	41.04	49.79	13.21	12.26			

which acts as an alternative benchmark for low-compute practitioners to evaluate few-shot learning. These results aid in determining the sample efficiency of the models, i.e., how much they benefit from observing each new data point in the target language. In addition, Table 18 (Appendix C.3) reports the Area Under the Curve (AUC) for accuracy across different k values, which instead measures how consistently well a model performs independently from data size.

Gains from Few-shot Learning. Contrary to text-only multilingual tasks, where even few examples in the target language are sufficient to significantly improve the model performance (Lauscher et al., 2020), we find that this is largely not the case in multimodal settings. Figure 3 reveals that learning curves are flat for XVNLI, MaRVL, and xFlickr&CO. Moreover, inspecting the gap between zero-shot and max-shot performance in Table 19 (Appendix C.3), gains from few-shot learning are consistent but modest.

A notable exception is xGQA, where as few as 27 examples increase performance dramatically and accuracy continues to grow with the sample size. This difference is possibly due to the highly-structured nature of xGQA, where texts are generated from templates and are extremely short, and its larger label space. To investigate the latter, we follow Pfeiffer et al. (2022) and evaluate the accuracy for each structurally different question type. Figure 4 shows that increasing the number of examples has the least impact on *Verify*-type questions, which consist of (binary) answers/labels (*{Yes, No}*). These results are in line with our few-shot experiments on the other tasks, where few-shot tuning displays negligible impact on cross-lingual transfer when the label space is small. The largest impact in xGQA is observed on *Choose*-type and *Query*-type questions, where the diversity of the label space is large. This again shows that the multilingual misalignment of the prediction space negatively impacts zero-shot results, which can be mitigated by fine-tuning even on a few examples in the target language.

On the other hand, current models may require many more examples than the maximum of k -shot we established for the other tasks. We test this hypothesis in a few-shot learning experiment with more examples in JPN and DEU on xFlickr&CO. The results in Figure 5 demonstrate that examples in the order of the thousands may be needed to achieve significant gains. This leaves open an important challenge to make future models more sample efficient, in more resource-poor settings like those considered here.

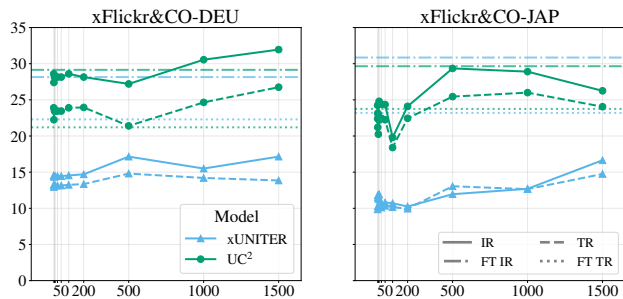


Figure 5. Retrieval accuracy on DEU and JPN xFlickr&CO splits with target language fine-tuning (horizontal lines) and more shots.

Multilingual Encoders. Again, based on Table 5 the model that yields the highest scores is UC² in all tasks. However, we also note that some models are particularly brittle, and few-shot adaptation can lead them to revert to chance-level performance, as is the case of M³P when fine-tuning on MaRVL’s target languages IND, CMN, and TUR.

6. Conclusion and Outlook

We introduced the Image-Grounded Language Understanding Evaluation (IGLUE) benchmark, aiming to facilitate modelling and evaluation in multilingual vision-and-language research. IGLUE brings together 5 datasets across 4 structurally different tasks in 20 diverse languages, and provides standardised data splits and evaluation protocols to train and evaluate in zero-shot and few-shot transfer setups.

We set strong baselines and carry out analyses that shed new light into multilingual V&L models, now enabled by the creation of the IGLUE benchmark. In particular, we find that **1)** ‘translate test’ transfer vastly outperforms zero-shot transfer via multilingual encoders; **2)** few-shot learning requires thousand of examples before yielding gains; **3)** unlabelled textual data size and typological distance from English are weaker predictors of models performance than usually observed in text-only tasks; **4)** results vary significantly across tasks, with multicultural visual reasoning and cross-modal retrieval having the largest transfer gaps. Given the diversity and manifoldness of each task, we explicitly set IGLUE as a multi-faceted benchmark, without a single average score.

Along with putting forward an evaluation suite for current and future multilingual V&L models, IGLUE also opens up new opportunities for future research that goes beyond the scope of this paper: such as analysing single-source versus multi-source transfer for V&L tasks, conducting multi-task learning, or investigating staged multi-task transfer. A longer-term outlook concerns including other tasks and languages into the benchmark, as well as reaching beyond the image-only visual modality, towards multimodal tasks that rely on processing videos (Li et al., 2021) and speech (Yang et al., 2021) in different languages.

Acknowledgements

■ We are grateful to Heather Lent, Semih Yagcioglu, Ilker Kesen, and Mustafa Sercan Amac for their constructive feedback and help. We also thank Krishna Srinivasan for providing us with an early release of the WIT test data and further clarifications. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 801199 (*Emanuele Bugliarello*). The work of *Fangyu Liu* has been supported by Grace & Thomas C.H. Chan Cambridge Scholarship. The work of *Jonas Pfeiffer* has been supported by the LOEWE initiative (Hesse, Germany) within the emergenCITY center. The work of *Ivan Vulić* has been supported by the ERC PoC Grant MultiConvAI (no. 957356), and a research donation from Huawei. The work of *Edoardo Ponti* and *Siva Reddy* has been supported by the Facebook CIFAR AI Chair program. The annotation cost is partially funded by Cambridge Digital Humanities Digitisation/Digital Resources Awards and the NSERC Discovery Grant.

References

- Aggarwal, P., Tambi, R., and Kale, A. Towards zero-shot cross-lingual image retrieval and tagging. In *Workshop on Multilingual Search*, Online, April 2021.
- Agić, Ž. and Schlueter, N. Baselines and test data for cross-lingual inference. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May 2018.
- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., and Zhang, L. Bottom-up and top-down attention for image captioning and visual question answering. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6077–6086, 2018.
- Ansell, A., Ponti, E., Korhonen, A., and Vulić, I. Composable sparse fine-tuning for cross-lingual transfer. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1778–1796, Dublin, Ireland, May 2022.
- Armitage, J., Kacupaj, E., Tahmasebzadeh, G., Swati, Maleshkova, M., Ewerth, R., and Lehmann, J. MIm: A benchmark dataset for multitask learning with multiple languages and modalities. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management, CIKM ’20*, pp. 2967–2974, New York, NY, USA, 2020.
- Baltrusaitis, T., Ahuja, C., and Morency, L.-P. Multimodal machine learning: A survey and taxonomy. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(2):423–443, feb 2019.

- Barrault, L., Bougares, F., Specia, L., Lala, C., Elliott, D., and Frank, S. Findings of the third shared task on multimodal machine translation. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pp. 304–323, Belgium, Brussels, October 2018.
- Bender, E. M. On achieving and evaluating language-independence in nlp. *Linguistic Issues in Language Technology*, 6, Oct. 2011.
- Bernardi, R., Cakici, R., Elliott, D., Erdem, A., Erdem, E., Ikizler-Cinbis, N., Keller, F., Muscat, A., and Plank, B. Automatic description generation from images: A survey of models, datasets, and evaluation measures. *J. Artif. Int. Res.*, 55(1):409–442, jan 2016.
- Blasi, D., Anastasopoulos, A., and Neubig, G. Systematic inequalities in language technology performance across the world’s languages. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5486–5505, Dublin, Ireland, May 2022.
- Bowman, S. R., Angeli, G., Potts, C., and Manning, C. D. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 632–642, Lisbon, Portugal, September 2015.
- Bugliarello, E., Cotterell, R., Okazaki, N., and Elliott, D. Multimodal pretraining unmasked: A meta-analysis and a unified framework of vision-and-language BERTs. *Transactions of the Association for Computational Linguistics*, 9:978–994, 2021.
- Caglayan, O., Kuyu, M., Amac, M. S., Madhyastha, P., Erdem, E., Erdem, A., and Specia, L. Cross-lingual visual pre-training for multimodal machine translation. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 1317–1324, Online, April 2021.
- Chen, Y.-C., Li, L., Yu, L., El Kholly, A., Ahmed, F., Gan, Z., Cheng, Y., and Liu, J. UNITER: Universal image-text representation learning. In Vedaldi, A., Bischof, H., Brox, T., and Frahm, J.-M. (eds.), *Computer Vision – ECCV 2020*, pp. 104–120, Cham, 2020.
- Conneau, A. and Lample, G. Cross-lingual language model pretraining. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., and Stoyanov, V. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 8440–8451, Online, July 2020.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019.
- Dryer, M. S. and Haspelmath, M. (eds.). *WALS Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig, 2013.
- Dumitrescu, S. D., Rebeja, P., Lorincz, B., Gaman, M., Avram, A., Ilie, M., Pruteanu, A., Stan, A., Rosia, L., Iacobescu, C., Morogan, L., Dima, G., Marchidan, G., Rebedea, T., Chitez, M., Yogatama, D., Ruder, S., Ionescu, R. T., Pascanu, R., and Patraucean, V. Liro: Benchmark and leaderboard for romanian language tasks. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.
- Elliott, D., Frank, S., Sima’an, K., and Specia, L. Multi30K: Multilingual English-German image descriptions. In *Proceedings of the 5th Workshop on Vision and Language*, pp. 70–74, Berlin, Germany, August 2016.
- Elliott, D., Frank, S., Barrault, L., Bougares, F., and Specia, L. Findings of the second shared task on multimodal machine translation and multilingual image description. In *Proceedings of the Second Conference on Machine Translation*, pp. 215–233, Copenhagen, Denmark, September 2017.
- Ethayarajh, K. and Jurafsky, D. Utility is in the eye of the user: A critique of NLP leaderboards. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 4846–4853, Online, November 2020.
- Frank, S., Elliott, D., and Specia, L. Assessing multilingual multimodal image description: Studies of native speaker preferences and translator choices. *Natural Language Engineering*, 24(3):393–413, 2018.
- García, C. MS-COCO-ES, 2020. <https://github.com/carlosGarciaHe/MS-COCO-ES>.

- Geigle, G., Pfeiffer, J., Reimers, N., Vulić, I., and Gurevych, I. Retrieve fast, rerank smart: Cooperative and joint approaches for improved cross-modal retrieval. *Transactions of the Association for Computational Linguistics*, 10, 2022.
- Gella, S., Sennrich, R., Keller, F., and Lapata, M. Image pivoting for learning multilingual multimodal representations. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 2839–2845, Copenhagen, Denmark, September 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- Hendrycks, D. and Gimpel, K. Gaussian Error Linear Units (GELUs). *arXiv preprint arXiv:1606.08415*, 2016.
- Hu, J., Ruder, S., Siddhant, A., Neubig, G., Firat, O., and Johnson, M. XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 4411–4421, 13–18 Jul 2020.
- Huang, P.-Y., Chang, X., and Hauptmann, A. Multi-head attention with diversity for learning grounded multilingual multimodal representations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 1461–1467, Hong Kong, China, November 2019.
- Huang, P.-Y., Hu, J., Chang, X., and Hauptmann, A. Un-supervised multimodal neural machine translation with pseudo visual pivoting. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 8226–8237, Online, July 2020.
- Huang, P.-Y., Patrick, M., Hu, J., Neubig, G., Metze, F., and Hauptmann, A. Multilingual multimodal pre-training for zero-shot cross-lingual transfer of vision-language models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2443–2459, Online, June 2021.
- Hudson, D. A. and Manning, C. D. GQA: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- Jain, A., Guo, M., Srinivasan, K., Chen, T., Kudugunta, S., Jia, C., Yang, Y., and Baldrige, J. MURAL: Multimodal, multitask representations across languages. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pp. 3449–3463, Punta Cana, Dominican Republic, November 2021.
- Jauhar, S. K., Gamon, M., and Pantel, P. Neural task representations as weak supervision for model agnostic cross-lingual transfer. *arXiv:1811.01115*, 2018.
- Jia, C., Yang, Y., Xia, Y., Chen, Y.-T., Parekh, Z., Pham, H., Le, Q., Sung, Y.-H., Li, Z., and Duerig, T. Scaling up visual and vision-language representation learning with noisy text supervision. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 4904–4916, 18–24 Jul 2021.
- Joshi, P., Santy, S., Budhiraja, A., Bali, K., and Choudhury, M. The state and fate of linguistic diversity and inclusion in the NLP world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 6282–6293, Online, July 2020.
- Karpathy, A. and Fei-Fei, L. Deep visual-semantic alignments for generating image descriptions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3128–3137, 2015.
- Kiela, D., Vulić, I., and Clark, S. Visual bilingual lexicon induction with transferred ConvNet features. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 148–158, Lisbon, Portugal, September 2015.
- Kiela, D., Conneau, A., Jabri, A., and Nickel, M. Learning visually grounded sentence representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 408–418, New Orleans, Louisiana, June 2018.
- Kim, D., Saito, K., Saenko, K., Sclaroff, S., and Plummer, B. Mule: Multimodal universal language embedding. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07):11254–11261, Apr. 2020.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.-J., Shamma, D. A., Bernstein, M. S., and Fei-Fei, L. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *Int. J. Comput. Vision*, 123(1):32–73, may 2017.

- Lam, Q. H., Le, Q. D., Nguyen, V. K., and Nguyen, N. L.-T. Uit-viic: A dataset for the first evaluation on vietnamese image captioning. In Nguyen, N. T., Hoang, B. H., Huynh, C. P., Hwang, D., Trawiński, B., and Vossen, G. (eds.), *Computational Collective Intelligence*, pp. 730–742, Cham, 2020.
- Lan, W., Li, X., and Dong, J. Fluency-guided cross-lingual image captioning. In *Proceedings of the 25th ACM International Conference on Multimedia*, MM '17, pp. 1549–1557, New York, NY, USA, 2017.
- Lauscher, A., Ravishankar, V., Vulić, I., and Glavaš, G. From zero to hero: On the limitations of zero-shot language transfer with multilingual Transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 4483–4499, Online, November 2020.
- Lei, J., Berg, T., and Bansal, M. mTVR: Multilingual moment retrieval in videos. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pp. 726–734, Online, August 2021.
- Li, L., Lei, J., Gan, Z., Yu, L., Chen, Y.-C., Pillai, R., Cheng, Y., Zhou, L., Wang, X. E., Wang, W. Y., Berg, T. L., Bansal, M., Liu, J., Wang, L., and Liu, Z. VALUE: A multi-task benchmark for video-and-language understanding evaluation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.
- Li, L. H., Yatskar, M., Yin, D., Hsieh, C.-J., and Chang, K.-W. VisualBERT: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019a.
- Li, X., Xu, C., Wang, X., Lan, W., Jia, Z., Yang, G., and Xu, J. COCO-CN for cross-lingual image tagging, captioning, and retrieval. *IEEE Transactions on Multimedia*, 21(9): 2347–2360, 2019b.
- Liang, Y., Duan, N., Gong, Y., Wu, N., Guo, F., Qi, W., Gong, M., Shou, L., Jiang, D., Cao, G., Fan, X., Zhang, R., Agrawal, R., Cui, E., Wei, S., Bharti, T., Qiao, Y., Chen, J.-H., Wu, W., Liu, S., Yang, F., Campos, D., Majumder, R., and Zhou, M. XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 6008–6018, Online, November 2020.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft COCO: Common objects in context. In Fleet, D., Pajdla, T., Schiele, B., and Tuytelaars, T. (eds.), *Computer Vision – ECCV 2014*, pp. 740–755, Cham, 2014.
- Littell, P., Mortensen, D. R., Lin, K., Kairis, K., Turner, C., and Levin, L. URIEL and lang2vec: Representing languages as typological, geographical, and phylogenetic vectors. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 8–14, Valencia, Spain, April 2017.
- Liu, F., Bugliarello, E., Ponti, E. M., Reddy, S., Collier, N., and Elliott, D. Visually grounded reasoning across languages and cultures. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 10467–10485, Online and Punta Cana, Dominican Republic, November 2021.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.
- Lu, J., Batra, D., Parikh, D., and Lee, S. Vilbert: Pre-training task-agnostic visiolinguistic representations for vision-and-language tasks. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Lu, J., Goswami, V., Rohrbach, M., Parikh, D., and Lee, S. 12-in-1: Multi-task vision and language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10434–10443, June 2020.
- Mishra, S. K., Dhir, R., Saha, S., and Bhattacharyya, P. A Hindi image caption generation framework using deep learning. *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 20(2), Mar 2021.
- Nakayama, H., Tamura, A., and Ninomiya, T. A visually-grounded parallel corpus with phrase-to-region linking. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pp. 4204–4210, Marseille, France, May 2020.
- Ni, M., Huang, H., Su, L., Cui, E., Bharti, T., Wang, L., Zhang, D., and Duan, N. M3p: Learning universal representations via multitask multilingual multimodal pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3977–3986, June 2021.
- Park, S., Moon, J., Kim, S., Cho, W. I., Han, J. Y., Park, J., Song, C., Kim, J., Song, Y., Oh, T., Lee, J., Oh, J., Lyu, S., Jeong, Y., Lee, I., Seo, S., Lee, D., Kim, H., Lee, M., Jang, S., Do, S., Kim, S., Lim, K., Lee, J., Park, K., Shin, J.,

- Kim, S., Park, L., Oh, A., Ha, J.-W., and Cho, K. KLUE: Korean language understanding evaluation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. PyTorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, pp. 8024–8035, 2019.
- Pfeiffer, J., Geigle, G., Kamath, A., Steitz, J.-M., Roth, S., Vulić, I., and Gurevych, I. xGQA: Cross-lingual visual question answering. In *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 2497–2511, Dublin, Ireland, May 2022.
- Ponti, E. M., O’Horan, H., Berzak, Y., Vulić, I., Reichart, R., Poibeau, T., Shutova, E., and Korhonen, A. Modeling language variation and universals: A survey on typological linguistics for natural language processing. *Computational Linguistics*, 45(3):559–601, September 2019.
- Ponti, E. M., Glavaš, G., Majewska, O., Liu, Q., Vulić, I., and Korhonen, A. XCOPA: A multilingual dataset for causal commonsense reasoning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2362–2376, Online, November 2020.
- Ponti, E. M., Aralikkatte, R., Shrivastava, D., Reddy, S., and Søgaard, A. Minimax and neyman–Pearson meta-learning for outlier languages. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 1245–1260, Online, August 2021a.
- Ponti, E. M., Vulić, I., Cotterell, R., Parovic, M., Reichart, R., and Korhonen, A. Parameter space factorization for zero-shot learning across tasks and languages. *Transactions of the Association for Computational Linguistics*, 9: 410–428, 2021b.
- Ren, S., He, K., Girshick, R., and Sun, J. Faster R-CNN: Towards real-time object detection with region proposal networks. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, pp. 91–99. Curran Associates, Inc., 2015.
- Rotman, G., Vulić, I., and Reichart, R. Bridging languages through images with deep partial canonical correlation analysis. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 910–921, Melbourne, Australia, July 2018.
- Ruder, S., Constant, N., Botha, J., Siddhant, A., Firat, O., Fu, J., Liu, P., Hu, J., Garrette, D., Neubig, G., and Johnson, M. XTREME-R: Towards more challenging and nuanced multilingual evaluation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 10215–10245, Online and Punta Cana, Dominican Republic, November 2021.
- Scaiella, A., Croce, D., and Basili, R. Large scale datasets for image and video captioning in Italian. *Italian Journal of Computational Linguistics*, 2(5):49–60, 2019.
- Schick, T. and Schütze, H. It’s not just size that matters: Small language models are also few-shot learners. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2339–2352, Online, June 2021.
- Sennrich, R., Haddow, B., and Birch, A. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1715–1725, Berlin, Germany, August 2016.
- Sharma, P., Ding, N., Goodman, S., and Soricut, R. Conceptual Captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2556–2565, Melbourne, Australia, July 2018.
- Shavrina, T., Fenogenova, A., Anton, E., Shevelev, D., Artemova, E., Malykh, V., Mikhailov, V., Tikhonova, M., Chertok, A., and Evlampiev, A. RussianSuperGLUE: A Russian language understanding evaluation benchmark. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 4717–4726, Online, November 2020.
- Snyder, B. and Barzilay, R. Climbing the tower of babel: Unsupervised multilingual learning. In *Proceedings of the 27th International Conference on International Conference on Machine Learning, ICML’10*, pp. 29–36, Madison, WI, USA, 2010.
- Srinivasan, K., Raman, K., Chen, J., Bendersky, M., and Najork, M. WIT: Wikipedia-Based Image Text Dataset for Multimodal Multilingual Machine Learning, pp. 2443–2449. Association for Computing Machinery, New York, NY, USA, 2021.

- Su, L., Duan, N., Cui, E., Ji, L., Wu, C., Luo, H., Liu, Y., Zhong, M., Bharti, T., and Sacheti, A. GEM: A general evaluation benchmark for multimodal tasks. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 2594–2603, Online, August 2021.
- Su, W., Zhu, X., Cao, Y., Li, B., Lu, L., Wei, F., and Dai, J. VL-BERT: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations*, 2020.
- Suhr, A., Zhou, S., Zhang, A., Zhang, I., Bai, H., and Artzi, Y. A corpus for reasoning about natural language grounded in photographs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 6418–6428, Florence, Italy, July 2019.
- Tan, H. and Bansal, M. LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 5100–5111, Hong Kong, China, November 2019.
- Unal, M. E., Citamak, B., Yagcioglu, S., Erdem, A., Erdem, E., Cinbis, N. I., and Cakici, R. TasvirEt: Görüntülerden otomatik türkçe açıklama oluşturma İçin bir denektaçı veri kümesi (TasvirEt: A benchmark dataset for automatic Turkish description generation from images). *IEEE Sinyal İşleme ve İletişim Uygulamaları Kurultayı (SIU 2016)*, 2016.
- van Miltenburg, E., Elliott, D., and Vossen, P. Cross-linguistic differences and similarities in image descriptions. In *Proceedings of the 10th International Conference on Natural Language Generation*, pp. 21–30, Santiago de Compostela, Spain, September 2017.
- van Miltenburg, E., Kádár, Á., Koolen, R., and Kraemer, E. DIDEc: The Dutch image description and eye-tracking corpus. In *Proceedings of the 27th International Conference on Computational Linguistics*, pp. 3658–3669, Santa Fe, New Mexico, USA, August 2018.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, pp. 5998–6008. Curran Associates, Inc., 2017.
- Vu, H. T., Greco, C., Erofeeva, A., Jafaritazehjan, S., Linders, G., Tanti, M., Testoni, A., Bernardi, R., and Gatt, A. Grounded textual entailment. In *Proceedings of the 27th International Conference on Computational Linguistics*, pp. 2354–2368, Santa Fe, New Mexico, USA, August 2018.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 353–355, Brussels, Belgium, November 2018.
- Wang, A., Pruksachatkun, Y., Nangia, N., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. Super-glue: A stickier benchmark for general-purpose language understanding systems. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32, 2019a.
- Wang, J., Madhyastha, P., Figueiredo, J., Lala, C., and Specia, L. MultiSubs: A large-scale multimodal and multilingual dataset. *arXiv:2103.01910*, 2021.
- Wang, X., Wu, J., Chen, J., Li, L., Wang, Y.-F., and Wang, W. Y. VateX: A large-scale, high-quality multilingual dataset for video-and-language research. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019b.
- Wehrmann, J., Lopes, M. A., Souza, D., and Barros, R. Language-agnostic visual-semantic embeddings. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5803–5812, 2019.
- Wilie, B., Vincentio, K., Winata, G. I., Cahyawijaya, S., Li, X., Lim, Z. Y., Soleman, S., Mahendra, R., Fung, P., Bahar, S., and Purwarianti, A. IndoNLU: Benchmark and resources for evaluating Indonesian natural language understanding. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pp. 843–857, Suzhou, China, December 2020.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.
- Xie, N., Lai, F., Doran, D., and Kadav, A. Visual entailment: A novel task for fine-grained image understanding. *arXiv:1901.06706*, 2019.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., and He, K. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- Yang, S.-w., Chi, P.-H., Chuang, Y.-S., Lai, C.-I., Lakhota, K., Lin, Y. Y., Liu, A. T., Shi, J., Chang, X., Lin, G.-T.,

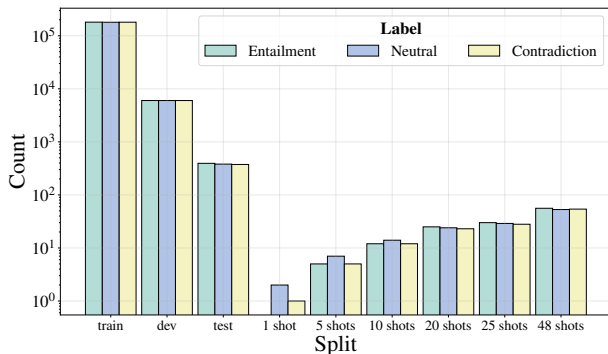
- hsien Huang, T., Tseng, W.-C., tik Lee, K., Liu, D.-R., Huang, Z., Dong, S., Li, S.-W., Watanabe, S., rahman Mohamed, A., and yi Lee, H. Superb: Speech processing universal performance benchmark. In *Interspeech*, 2021.
- Yao, S. and Wan, X. Multimodal transformer for multimodal machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4346–4350, Online, July 2020.
- Yin, D., Li, L. H., Hu, Z., Peng, N., and Chang, K.-W. Broaden the vision: Geo-diverse visual commonsense reasoning. In *Proceedings of EMNLP 2021*, pp. 2115–2129, 2021.
- Yoshikawa, Y., Shigeto, Y., and Takeuchi, A. STAIR captions: Constructing a large-scale Japanese image caption dataset. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 417–421, Vancouver, Canada, July 2017.
- Young, P., Lai, A., Hodosh, M., and Hockenmaier, J. 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, 2014.
- Zhao, M., Zhu, Y., Shareghi, E., Vulić, I., Reichart, R., Korhonen, A., and Schütze, H. A closer look at few-shot crosslingual transfer: The choice of shots matters. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 5751–5767, Online, August 2021.
- Zhou, M., Zhou, L., Wang, S., Cheng, Y., Li, L., Yu, Z., and Liu, J. Uc2: Universal cross-lingual cross-modal vision-and-language pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4155–4165, June 2021.

A. Datasets

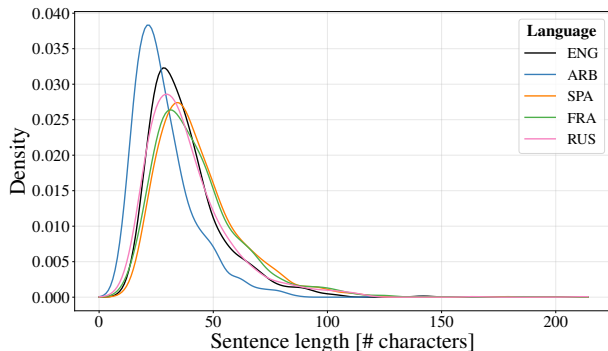
In this section, we provide more details about the datasets and splits that we introduce in IGLUE, focusing on reproducibility of our experiments and any follow-up work that leverages the IGLUE benchmark. We also report how few-shot samples were collected for each dataset. In particular, we use the notion of ‘annotation context’ to define what a ‘shot’ means in each dataset (see details below). Table 6 lists the size of few-shot splits across all datasets.

A.1. XVNLI

Figure 6 illustrates statistical properties of XVNLI. In Figure 6(a), we see that the distribution of labels is uniform across splits. Figure 6(b) instead shows that the distribution of sentence lengths—in number of characters—in the test split is similar across languages, and often shorter for ARB.



(a) Uniform distribution of labels across splits.



(b) Similar test sentence length distribution across languages.

Figure 6. Statistical properties of XVNLI.

Few-shot Data. We extract few-shot samples from the original SNLI test split, which was translated by Agić & Schluter (2018). This was done prior to defining our test split, and by sampling all the examples corresponding to a given image, so as to avoid data leakage between training and evaluation

Table 6. Few-shot statistics. * for MaRVL, we collected few-shot data in three of the five languages due to challenges in annotations.

Dataset	Metric						
XVNLI	# shots	1	5	10	20	25	48
	# images	1	5	10	20	25	48
	# samples	3	18	39	75	90	168
xGQA	# shots	1	5	10	20	25	48
	# images	1	5	10	20	25	48
	# samples	27	155	317	594	704	1,490
MaRVL*	# shots	1	2	4	10	10×2	
	# images	8	16	32	80	80	
	# samples	4	8	16	40	80	
xFlickr&CO	# shots	1	5	10	25	50	100
	# images	1	5	10	25	50	100
	# samples	1	5	10	25	50	100

data. In XVNLI, multiple samples (hypotheses and answers) are associated to a given image (premise). We hence consider an image as a single shot, as it serves as the context for an annotator to create multiple samples. Following Pfeiffer et al. (2022), we randomly sample up to 48 images as contexts, resulting in up to 168 hypotheses and answers.

A.2. xGQA

Few-shot Data. We reuse the few-shot splits defined by Pfeiffer et al. (2022), which also include development data in the target languages. The authors sampled 48 images, giving up to 1,490 questions translated for each language.

A.3. MaRVL

Few-shot Data. We follow the annotation protocol of Liu et al. (2021) to collect few-shot samples. For each language, we first ask native speakers to select 10 *concepts* among the ones identified by the authors in order to collect in-domain training data. The annotators are then required to retrieve at least 8 images per concept. When doing so, we further provide the MaRVL test images and require them to select different ones to avoid any visual leakage between few-shot and test samples. We finally sample 8 images per concept, randomly pair them and ask annotators to write a caption that is true for two pairs and false for the other two pairs. These 8 images then constitute an ‘annotation context’ and as such we consider the resulting data as one shot (i.e. one shot, eight images and one sentence giving four data points). As image collection is the most time-consuming step in the pipeline, we expand our few-shot data by re-shuffling and re-annotating each image pair (‘10 × 2’ setup), resulting in 80 data points per language. Notably, due to the challenges in collecting few-shot data in low-resource languages (i.e. Swahili and Tamil), we only provide few-shot samples for three languages (i.e. Indonesian, Mandarin and Turkish).

A.4. xFlickr&CO

We follow the Flickr30K guidelines to collect our multilingual annotations, as shown in Figure 7(a), which we supplement with examples from Flickr30K (Figure 7(b)). In particular, we limit captions length based on previous work: 40 characters for CMN, 140 for TUR and 100 for the rest. Figure 7(c) shows the additional guidelines provided to validators. The resulting corpora were finally verified by native speakers known by the authors. Notably, the same guidelines were used to annotate both Flickr30K and COCO images. Figure 7 shows the distribution of caption lengths.

Image Description Guidelines

- You must describe each of the following images with one sentence in your native language.
- Please provide an **accurate description** of the activities, people, animals and objects you see depicted in the image.
- Each description must be a single sentence under 100 characters. Write **at least 8 words**.
- Make sure your **caption is precise and specific about this image**, and not general to describe similar ones.
- Please pay attention to grammar and spelling.
- We will accept your results if you provide a **good description for all images**.

Examples of good and bad descriptions.



(1) The dog is wearing a red sombrero.
Very Good: This describes the two main objects concisely and accurately.

(2) White dog wearing a red hat.
Good: Incomplete sentences like this are fine.

(3) The white dog is wearing a pink collar.
Okay: This describes the dog, but it ignores the hat.

(4) The red hat is adorned with gold sequins.
Bad: This ignores the dog.

(5) The dog is angry because he is hungry.
Bad: This is speculation.

(6) The dog.
Very Bad: This could describe any image of any dog.

(a) Flickr30K-like instructions to annotate xFlickr&CO data.

Examples in English



three redsox players are holding onto their bats while posing for a photo.



rows of bicycles parked on sidewalk next to cars and people.



an orange bus parked in front of a church with a tall clock tower.



two men shaking hands while a third person is putting a medal on one of them.



two black cats sitting on a desk with one on a laptop and one on papers.



a cow is grazing by a group of boats on the grass.

(b) Additional Flickr30K examples used to further guide annotators.

Image Descriptions Validation

- You must ensure the correctness and high quality of the following image descriptions
- An image description must be **precise and specific about its image**, and NOT GENERAL to describe similar ones
- Fix any **grammatical error**
- Fix any **spelling error**
- If a caption is too short or too generic, please **re-write** the caption yourself!
- The next slides show the guidelines and examples of the image description task that were given to the annotators
- The allocated time is more than you will probably need
- Please make sure you carefully check each description
- The quality of our data depends on you!

(c) Instructions given to native speakers to validate the captions.

Figure 7. Annotation guidelines for xFlickr&CO.

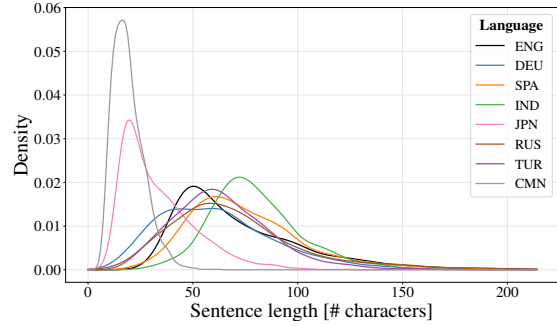


Figure 8. Test sentence length distribution in xFlickr&CO.

Few-shot Data. We use the same guidelines illustrated in Figure 7 to annotate few-shot image descriptions from 100 images samples from the Flickr30K training set. German captions are extracted from Multi30K (Elliott et al., 2016), Mandarin ones from Flickr30K-cn (Mishra et al., 2021), and Turkish ones from TasvirEt (Unal et al., 2016) – all independently collected from the English captions.

A.5. WIT

We extract 500,000 image captions from the caption_reference_description field of the English portion of the WIT dataset. This field corresponds to the captions shown on the Wikipedia pages directly below the images, and it tends to be the most topical and relevant description (Srinivasan et al., 2021). In particular, we sample among the images that were released by the Wikimedia Foundation for the corresponding Kaggle challenge,¹⁴ which do not include any identifiable personal information. This challenge also provides test data that was extracted from WIT prior to its public release.¹⁵ We use these data to define our multilingual test splits in order to let future practitioners evaluate models that were also trained on the WIT data. Per-language statistics are listed in Table 7. We note that a given image is mapped to a single caption more than 90% of the times in every language except for English (where this happens around 70% of the times and which is not used as an evaluation language in IGLUE). Figure 9 shows the overlap between test images across languages, and Figure 10 illustrates the distribution of caption lengths across languages in our WIT test splits.

Few-shot Data. Due to the demanding computational resources to evaluate image–text retrieval systems, we only provide few-shot splits for this task in xFlickr&CO.

¹⁴<https://www.kaggle.com/c/wikipedia-image-caption/data>.

¹⁵While we find no overlap between train and test image URLs, we find that a few base64 image encodings match between train and test splits: 1 image for ARB, ENG and JPN; 2 images for BUL, KOR and TUR; 3 images for IND and VIE; 0 for the other languages.

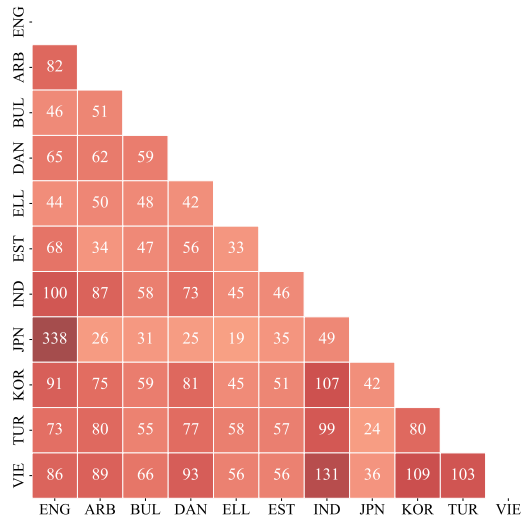


Figure 9. Number of images overlapping between any two languages in the WIT test splits.

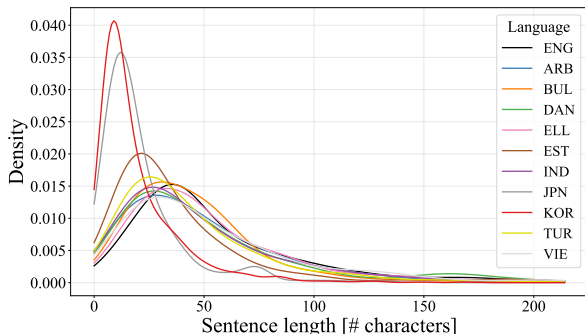


Figure 10. Test sentence length distribution in WIT.

Table 7. WIT test data statistics. % unique measures the percentage of images that only have a single caption associated with them.

Metric	Language										
	ENG	ARB	BUL	DAN	ELL	EST	IND	JPN	KOR	TUR	VIE
# images	685	792	806	814	541	780	854	842	889	681	869
# samples	1000	890	860	891	570	874	901	1000	931	721	946
% unique	70.9	91.2	94.9	93.0	97.2	93.6	96.0	90.8	95.8	95.6	93.7

A.6. Ethics Statement

For MaRVL and xFlickr&CO, we collected data from native speakers spread across the world in languages that are under-represented in current vision and language datasets. Our efforts were motivated by the belief that vision and language tasks need to better reflect the linguistic (and cultural) variety present in the world. We hired workers from the crowdsourcing platform prolific.co. The annotators’ identities were anonymous to us, and each of them was paid \$10-20/hour (even if their data were not ultimately used).

Table 8. Pretrained multilingual V&L encoders.

Model	Initialisation	Visual Tokens	Pretrain Data	# Params
mUNITER	mBERT	36 RoIs from Faster R-CNN with ResNet-101	CC (2.7 M) + Wikipedia (104 Langs)	184.83 M
xUNITER	XLM-R _{BASE}	36 RoIs from Faster R-CNN with ResNet-101	CC (2.7 M) + Wikipedia (104 Langs)	283.76 M
UC ²	XLM-R _{BASE}	36 RoIs from Faster R-CNN with ResNet-101	CC (3.3 M) + NMT CC- $\{\text{CES, DEU, FRA}\}$ + NMT CC- $\{\text{CMN, JPN}\}$	281.64 M
M ³ P	XLM-R _{BASE}	10-100 RoIs from Faster R-CNN with ResNeXt-101	CC (3.3 M) + Wikipedia (100 Langs) + Panlex (50 Langs)	376.90 M

B. Experimental Details

Pretrained Models. Table 8 summarises the key properties of the available pretrained multilingual V&L encoders. They are all single-stream architectures (Bugliarello et al., 2021), which mostly differ in the corpora they used for pretraining.

Experimental Settings. We train—both for zero-shot finetuning and for few-shot experiments—each encoder end-to-end by learning a task-specific ‘head’ network for each dataset on a single NVIDIA V100 (16GB) GPU card. For each task, we use the same hyperparameters as in the controlled study of Bugliarello et al. (2021) except for maximum sequence length, which we increase to accommodate the multilingual inputs (see details per task below). We use the AdamW (Loshchilov & Hutter, 2019) optimiser with parameters $\beta_1=0.9$, $\beta_2=0.999$ and $\epsilon=1e-6$. The base learning rate depends on each task (see below) but it is always first linearly warmed up for 10% of the task-specific number of steps, and then linearly decayed until up to 20 epochs of training have been reached. We also apply a weight decay of $10e-4$ and gradient clipping of 1.0. Before training, we extract, for each dataset, 36 image regions using a ResNet-101 backbone (He et al., 2016) for mUNITER, xUNITER and UC²; and 10–100 image regions using a ResNeXt-101 backbone (Xie et al., 2017) for M³P—both trained on Visual Genome (Krishna et al., 2017; Anderson et al., 2018).

Fine-Tuning Setup. To increase accessibility to the benchmark by practitioners with limited computing resources, we fix the number of epochs for which we train our baselines so that they can be trained in less than 12 hours per dataset.¹⁶ In particular, we set up each task based on prior work (Lu et al., 2020; Bugliarello et al., 2021) as follows.

- **Visual Entailment.** We treat XVNLI as a three-way classification problem, and apply a linear layer from the pooled representation of each Transformer encoder. We fine-tune each model using a binary cross-entropy loss with a batch size of 128 for 10 epochs. We use a learning rate of $2e-5$ and max token length of 80.
- **Visual QA.** As in Hudson & Manning (2019), we treat

¹⁶Except for M³P, which is twice as big as the other models.

GQA as a multi-label classification task by assigning a score to each answer according to its relevancy to the ground truth answer. We apply a two-layer MLP with a GeLU activation (Hendrycks & Gimpel, 2016) function in between on top of the pooled representation. We fine-tune on the *balanced* split of GQA using a binary cross-entropy loss against 1,842 labels with a batch size of 256 for 5 epochs. We use a learning rate of 4e-5 and max token length of 40.

- **Visual Reasoning.** We cast this task as a binary classification problem: Given the two pooled representations from encoding the textual description with each of the two images independently, we concatenate and feed them into an MLP (same as for GQA) to predict a True or False label. We fine-tune using a binary cross-entropy loss with a batch size of 64 for 20 epochs. We use a learning rate of 1e-5 and max token length of 80.
- **Image-Text Retrieval.** We train a four-way linear classifier against the pooled representations from the true image-caption pair and three hard negatives determined offline. We optimise our models using a cross-entropy loss, except for UC² for which we found the authors’ triplet loss to work better (see Appendix C). We train with a batch size of 64 for 10 epochs for xFlickr&CO; and for 2 epochs on the larger WIT. We use a learning rate of 2e-5 and max token length of 80.

Few-shot Setup. We follow the same setup as in the above described fine-tuning procedure, with the exception of searching across three values for the base learning rate {1e-5, 5e-5, 1e-4} and always training for 20 epochs. The batch size is adjusted according to the number of data points available for a given split (Table 9), and models are initialised from the English fine-tuned parameter set. For each run, we select the checkpoint that achieves the largest validation performance for evaluation on the test sets.

While the validation sets are in English for every dataset except for xGQA, we found no significant difference in using machine translations of each validation set for checkpoint selection. Therefore, we decided to report results using the original validation sets as this more closely matches the few-shot scenario. Nevertheless, we release the corresponding translations for each dataset for future extensions.

C. Reproducibility and Additional Results

We provide more details regarding our experimental setup in order to ensure reproducibility (§C.1). Further, we report additional per-language (§C.2) and few-shot (§C.3) results, which supplement the ones provided in the main paper.

C.1. A Unified Implementation Framework

In order to provide a more lightweight setup for all the experiments in IGLUE, which concerns the experiments

Table 9. Few-shot batch sizes.

Dataset	Metric						
XVNL	# shots	1	5	10	20	25	48
	batch size	2	16	32	64	64	64
xGQA	# shots	1	5	10	20	25	48
	batch size	8	16	32	64	64	64
MaRVL	# shots	1	2	4	10	10×2	
	batch size	2	4	8	16	32	
xFlickr&CO	# shots	1	5	10	25	50	100
	batch size	1	2	4	8	16	32

Table 10. Results of reimplemented models on Multi30K. Mean Recall for zero-shot evaluation (i.e. English-only fine-tuning).

Method	ENG	DEU	FRA	CES
M ³ P (Ni et al., 2021)	87.4	58.5	46.0	36.8
UC ² (Zhou et al., 2021)	87.2	74.9	74.0	67.9
M ³ P	83.1	55.5	51.9	45.4
UC ²	82.7	69.9	71.1	65.7
mUNITER	83.2	43.7	47.3	21.0
xUNITER	81.8	50.2	48.4	37.5

Table 11. Results of reimplemented models on xGQA. Accuracy for zero-shot evaluation (i.e. English-only fine-tuning).

Method	ENG	BEN	DEU	IND	KOR	POR	RUS	CMN	avg
M ³ P	58.4	15.8	23.9	22.6	16.9	24.4	20.4	18.6	20.4
OSCAR+ ^{Emb}	62.2	14.9	17.3	18.3	17.1	19.2	10.5	16.4	16.3
OSCAR+ ^{Ada}	60.3	15.4	18.9	18.8	15.3	27.0	17.5	15.0	18.3
mBERT ^{Ada}	56.2	15.1	29.8	19.2	19.1	30.4	24.4	24.9	23.2
M ³ P	57.9	19.2	35.6	29.6	28.5	35.7	32.9	31.1	33.8
UC ²	59.4	20.8	45.1	28.1	24.3	27.6	31.9	33.9	33.9
mUNITER	58.2	3.9	26.3	13.0	5.5	17.4	9.3	8.4	17.7
xUNITER	57.7	9.7	35.7	36.3	14.6	24.7	20.8	19.2	27.3

we conducted for this work as well as any future experiments by other researchers, we reimplemented M³P and UC² in a unified framework based on VOLTA (Bugliarello et al., 2021), which already implements mUNITER and xUNITER, along with five English pretrained encoders. In order to verify the correctness of the implementation, we evaluated each model on Multi30K (Elliott et al., 2016) in a zero-shot scenario. The results shown in Table 10 prove the correctness of our reimplementation: Performance of both models is within a couple of points from the one reported by the original authors. In particular, we obtain similar results as Geigle et al. (2022) for M³P. Moreover, we found that using a cross-entropy loss for UC² led to a drop in performance of approximately 2.5 mean Recall points in each language. Similarly, using a triplet loss for M³P led to a drop in performance of nearly 3 points.

In Table 11, we report the performance of our models on

Table 12. Zero-shot per-language results on WIT (IR and TR).

WIT Type	Model	Language									
		ARB	BUL	DAN	ELL	EST	IND	JPN	KOR	TUR	VIE
IR	mUNITER	7.74	8.26	10.66	8.95	7.67	10.88	9.00	5.91	9.57	13.00
	xUNITER	7.63	8.49	10.32	11.23	6.41	10.21	7.30	6.34	9.57	9.72
	UC ²	6.62	8.84	9.43	8.77	4.69	9.88	9.80	4.30	7.49	8.46
	M ³ P	8.87	8.84	9.43	9.65	5.38	8.66	7.00	6.12	6.52	10.78
TR	mUNITER	9.21	10.17	12.16	10.54	8.33	12.88	8.79	6.75	10.87	15.07
	xUNITER	9.08	10.30	9.34	12.38	7.82	10.66	10.10	6.97	9.69	11.74
	UC ²	8.32	7.69	10.44	11.64	6.03	11.47	10.81	5.74	8.81	9.90
	M ³ P	8.32	9.80	11.79	12.02	8.21	10.89	8.43	7.09	10.57	12.66

Table 13. Full per-language results on XVNLI.

XVNLI Lang	Model	# shots						
		0	1	5	10	20	25	48
ARB	mUNITER	46.73	46.99	46.39	49.40	47.16	48.97	46.91
	xUNITER	51.98	52.32	52.32	54.81	54.55	53.95	54.04
	UC ²	56.19	56.36	57.65	57.82	59.11	58.51	56.87
	M ³ P	55.24	54.98	55.07	56.19	55.24	56.79	56.01
SPA	mUNITER	56.96	56.79	57.47	57.30	57.73	57.65	57.73
	xUNITER	58.94	59.02	58.94	60.74	59.88	59.11	60.22
	UC ²	57.47	57.65	61.34	59.79	62.63	59.79	62.80
	M ³ P	58.85	58.33	58.94	60.05	59.19	59.28	60.40
FRA	mUNITER	59.36	59.45	59.19	59.54	59.36	59.62	59.36
	xUNITER	63.32	63.49	64.26	63.83	64.26	64.09	64.52
	UC ²	69.67	69.67	69.67	69.76	69.84	69.76	69.76
	M ³ P	56.36	56.62	57.99	58.33	57.82	57.65	58.59
RUS	mUNITER	51.72	51.72	51.46	50.52	52.32	51.46	51.80
	xUNITER	59.71	59.54	59.54	59.54	61.94	61.34	63.40
	UC ²	64.86	64.43	64.52	64.95	65.38	65.38	65.29
	M ³ P	62.54	62.46	62.72	63.23	62.46	62.97	62.46

the zero-shot splits of xGQA, for a direct comparison with the standard GQA test split and the performance achieved by Pfeiffer et al. (2022). We find that our setup leads to significantly better performance: our best model, UC² gains over 10 points compared to the best model from Pfeiffer et al. (2022). In addition, our M³P model is 13.4% more accurate than the one trained by the authors. We believe the main reason for this considerable performance gain is our deeper prediction head: While the authors place a linear classification head, our models use a non-linear two-layer feed-forward head on top of the Transformer representations.

C.2. Performance per Target Language

Tables 12 to 17 present language-specific results for each dataset and number of shots. Retrieval performance in WIT is relatively constant across languages, likely due to the different distribution of captions and images in the Wikipedia domain. Different from other datasets (see below), this is true even for UC² in the JPN split, even though the model was pretrained on large-scale multimodal data in Japanese.

In XVNLI and xGQA, performance peaks for French and German, respectively. While typologically more similar and

Table 14. Full per-language results on xGQA.

xGQA Lang	Model	# shots						
		0	1	5	10	20	25	48
BEN	mUNITER	3.06	19.36	23.94	27.53	30.04	31.04	34.89
	xUNITER	10.80	23.92	29.43	31.67	34.84	36.18	37.55
	UC ²	19.99	22.52	30.96	32.84	35.69	35.12	38.90
	M ³ P	18.64	23.42	31.07	33.37	35.74	35.94	37.76
DEU	mUNITER	23.95	29.43	33.88	35.40	37.82	37.43	40.29
	xUNITER	34.83	38.44	39.71	40.97	41.93	42.19	43.60
	UC ²	42.85	43.76	44.68	45.72	46.70	47.24	48.18
	M ³ P	33.42	34.37	39.66	40.73	41.78	41.93	43.19
IND	mUNITER	9.36	27.31	31.39	33.28	36.49	35.85	38.06
	xUNITER	33.73	35.28	37.38	37.96	39.48	39.40	40.90
	UC ²	28.67	34.76	38.95	40.11	41.26	41.51	43.11
	M ³ P	32.48	33.11	38.14	39.84	40.51	41.53	41.38
KOR	mUNITER	4.21	19.40	26.47	29.03	31.95	32.76	35.23
	xUNITER	12.12	23.45	31.49	34.70	36.66	37.26	39.32
	UC ²	21.36	29.33	33.02	34.50	36.09	38.40	39.63
	M ³ P	25.11	29.74	34.53	35.77	37.11	37.81	38.58
POR	mUNITER	13.67	22.88	31.09	34.38	37.05	37.32	39.41
	xUNITER	22.13	30.10	36.44	38.72	39.73	41.06	42.56
	UC ²	30.42	32.10	39.42	39.57	41.73	41.27	43.23
	M ³ P	31.40	33.37	37.62	39.47	41.09	41.96	43.01
RUS	mUNITER	8.49	22.51	29.42	31.98	32.90	34.34	34.78
	xUNITER	18.84	27.32	34.05	36.39	38.19	38.37	39.40
	UC ²	31.00	33.04	37.40	38.13	40.63	41.09	42.79
	M ³ P	27.50	33.50	37.23	37.86	39.18	39.76	42.15
CMN	mUNITER	7.03	18.33	30.90	32.25	36.18	35.52	37.81
	xUNITER	19.55	31.36	37.35	37.34	39.19	39.48	41.42
	UC ²	31.16	37.54	41.21	41.49	43.46	43.58	44.82
	M ³ P	28.65	30.67	36.77	37.77	39.25	39.99	41.19

Table 15. Full per-language results on MaRVL.

MaRVL Lang	Model	# shots					
		0	1	2	4	10	10×2
IND	mUNITER	54.79	51.42	52.13	54.34	56.83	51.42
	xUNITER	55.14	56.12	57.18	58.87	58.87	58.60
	UC ²	56.74	56.12	56.29	57.53	57.62	58.51
	M ³ P	56.47	49.02	48.05	50.18	50.09	50.09
SWA	mUNITER	51.17	-	-	-	-	-
	xUNITER	55.51	-	-	-	-	-
	UC ²	52.62	-	-	-	-	-
	M ³ P	55.69	-	-	-	-	-
TAM	mUNITER	52.66	-	-	-	-	-
	xUNITER	53.06	-	-	-	-	-
	UC ²	60.47	-	-	-	-	-
	M ³ P	56.04	-	-	-	-	-
TUR	mUNITER	54.66	52.54	55.93	55.59	54.66	54.75
	xUNITER	56.19	55.93	57.46	57.54	57.80	58.05
	UC ²	56.70	55.76	54.49	52.63	55.93	56.27
	M ³ P	56.78	51.19	49.32	49.58	49.41	49.58
CMN	mUNITER	55.34	56.52	54.84	54.55	55.34	54.05
	xUNITER	53.06	54.45	53.66	54.94	53.46	55.73
	UC ²	59.88	58.99	57.02	58.99	57.21	60.18
	M ³ P	55.04	50.69	50.40	50.20	49.80	49.70

Table 16. Full per-language results on xFlickr&CO IR.

xFlickr&CO IR		# shots						
Lang	Model	0	1	5	10	25	50	100
DEU	mUNITER	12.05	10.50	11.90	11.95	11.95	11.95	12.05
	xUNITER	14.55	14.40	14.35	14.45	14.45	14.45	14.55
	UC ²	28.60	27.40	28.15	28.15	28.15	28.15	28.60
	M ³ P	13.35	12.25	12.85	11.95	13.15	13.35	13.25
SPA	mUNITER	13.15	13.70	14.25	14.25	14.25	14.25	13.15
	xUNITER	16.10	16.10	16.05	16.05	16.05	16.05	16.10
	UC ²	15.95	15.40	15.70	15.50	8.55	15.40	15.95
	M ³ P	13.40	13.40	13.80	13.95	13.50	13.40	13.40
IND	mUNITER	5.95	6.05	4.55	5.75	5.65	5.80	5.95
	xUNITER	16.50	16.90	16.55	16.20	16.60	16.60	16.50
	UC ²	14.60	14.50	14.55	12.30	14.70	14.65	14.60
	M ³ P	13.20	13.20	14.20	14.15	13.50	13.20	13.20
JPN	mUNITER	6.30	6.20	6.20	6.20	6.20	6.20	6.30
	xUNITER	10.25	11.50	11.80	11.95	10.65	10.80	10.70
	UC ²	24.25	23.15	23.00	24.80	24.40	24.35	19.80
	M ³ P	10.30	11.50	10.85	12.75	11.70	12.05	11.75
RUS	mUNITER	5.85	3.90	5.75	6.55	6.25	8.50	9.25
	xUNITER	15.90	15.75	15.20	16.00	16.05	16.20	17.25
	UC ²	20.00	20.05	19.90	19.65	21.30	17.65	20.85
	M ³ P	15.95	15.95	17.10	16.15	16.75	16.40	16.65
TUR	mUNITER	1.75	1.90	1.40	1.30	1.30	1.35	1.75
	xUNITER	9.05	9.35	9.05	9.05	9.05	9.05	9.05
	UC ²	7.15	7.20	7.05	5.55	6.85	4.90	7.15
	M ³ P	7.75	7.75	7.45	7.45	7.35	7.75	7.75
CMN	mUNITER	11.35	11.50	11.45	11.40	11.45	11.45	11.35
	xUNITER	15.95	17.50	7.90	17.60	16.00	13.65	15.95
	UC ²	31.60	31.30	31.05	29.75	31.40	25.50	31.60
	M ³ P	16.45	14.85	16.15	9.30	16.40	12.65	16.45

geographically closer to English, we find that the amount of multilingual data, shown in Figure 11, is also indicative of the final performance. For instance, both mUNITER and xUNITER—which were pretrained on the same data but with different multilingual sampling ratios during text-only and multimodal pretraining—perform on par or better in Russian than in Spanish and Portuguese, respectively. Moreover, UC² consistently achieves the highest performance in French, German and Mandarin, which are 3/5 Conceptual Captions translations on which the model was pretrained.

On MaRVL, this is also the case for the Mandarin split, although UC² now reaches lower gains relative to other models in the other datasets. This shows how the out-of-distribution nature of the images in MaRVL limits the efficacy of ‘translate train’ approaches for zero-shot transfer here. Interestingly, Mandarin is also the language in which the other models perform the worst, indicating that large-scale multimodal data in the target language can still benefit downstream performance. Overall, due to the domain shifts present in MaRVL, the performance across all languages and models is comparably low.

Although being in-distribution, performance in xFlickr&CO is consistently low in every language, and especially in

Table 17. Full per-language results on xFlickr&CO TR.

xFlickr&CO TR		# shots						
Lang	Model	0	1	5	10	25	50	100
DEU	mUNITER	11.85	10.65	11.70	11.70	11.70	11.65	11.85
	xUNITER	13.25	12.95	13.20	13.15	13.15	13.15	13.25
	UC ²	23.90	22.25	23.45	23.45	23.45	23.45	23.90
	M ³ P	11.85	10.65	12.35	10.60	12.60	11.85	12.55
SPA	mUNITER	13.05	14.70	14.05	14.05	14.05	14.05	13.05
	xUNITER	15.10	15.60	15.30	15.30	15.30	15.30	15.10
	UC ²	15.30	15.85	16.15	16.15	9.25	15.45	15.30
	M ³ P	12.15	12.15	11.60	11.60	11.55	12.15	12.15
IND	mUNITER	7.55	8.30	4.85	6.30	6.35	6.20	7.55
	xUNITER	16.75	17.05	16.85	16.65	16.75	16.80	16.75
	UC ²	13.60	14.90	14.20	13.70	14.05	13.95	13.60
	M ³ P	12.10	12.10	13.00	13.15	12.20	12.10	12.10
JPN	mUNITER	7.70	8.10	7.55	7.55	7.55	7.55	7.70
	xUNITER	9.85	11.80	11.30	11.45	10.10	10.35	10.15
	UC ²	22.40	21.20	20.25	22.20	22.45	22.25	18.40
	M ³ P	9.65	12.30	10.60	11.50	11.35	11.25	10.50
RUS	mUNITER	6.80	7.60	7.50	7.15	7.25	8.55	10.00
	xUNITER	14.80	14.65	15.10	14.50	15.25	14.60	14.70
	UC ²	16.75	17.00	16.75	17.75	17.50	16.30	18.65
	M ³ P	14.45	14.45	14.65	14.90	14.40	15.85	15.40
TUR	mUNITER	3.25	3.80	2.45	2.10	1.95	2.00	3.25
	xUNITER	10.05	10.30	9.95	9.95	9.95	9.95	10.05
	UC ²	6.95	7.15	7.00	6.60	7.10	6.70	6.95
	M ³ P	8.35	8.35	8.10	8.10	8.05	8.35	8.35
CMN	mUNITER	11.85	12.40	11.65	11.70	11.65	11.65	11.85
	xUNITER	14.80	15.95	8.55	15.95	14.95	13.40	14.80
	UC ²	26.30	25.80	25.75	24.40	25.20	21.50	26.30
	M ³ P	14.75	13.65	14.40	10.65	15.00	12.90	14.75

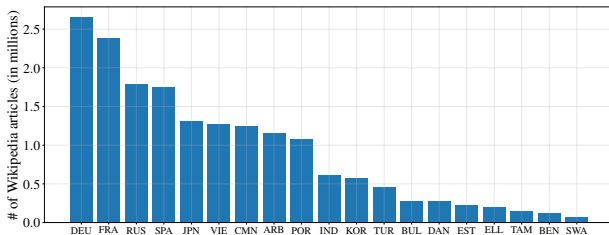


Figure 11. Number of Wikipedia articles per IGLUE language.

Turkish. Once more, UC² largely outperforms the other systems in DEU, JPN and CMN but is on par in the other languages.

All in all, the resulting differences between the best- and worst-performing languages are considerable, amounting 13.5 points for XVNLI, 22.9 points for xGQA, and 22.5 points for xFlickr&Co, even though the visual input is language-independent in these evaluation sets.

C.3. Additional Results for Few-shot Learning

Area under the Curve. To compare the overall performance of different models, we compute a normalised Area Under Curve (AUC) score for the few-shot setup (plotted in

Table 18. Few-shot AUC (in %) averaged across languages.

Model	NLI		QA	Reasoning	Retrieval	
	XVNLI	xGQA	MaRVL	xFlickr&CO		
				IR	TR	
mUNITER	54.10	33.67	54.68	8.36	8.89	
xUNITER	59.78	38.10	56.82	14.05	13.51	
UC ²	63.43	40.43	57.11	19.24	17.31	
M ³ P	59.07	38.89	49.96	12.90	12.07	

Table 19. Difference in accuracy between the most-shots and zero-shot setups averaged across languages. For MaRVL, accuracy is averaged across the 3/5 languages with few-shot training data.

Model	NLI		QA	Reasoning	Retrieval	
	XVNLI	xGQA	MaRVL	xFlickr&CO		
				IR	TR	
mUNITER	↑ 0.26	↑ 27.24	↓ 0.31	↑ 0.48	↑ 0.46	
xUNITER	↑ 2.07	↑ 18.96	↑ 2.87	↑ 0.26	↑ 0.03	
UC ²	↑ 1.63	↑ 13.60	↑ 1.04	↓ 0.52	↓ 0.30	
M ³ P	↑ 1.11	↑ 12.87	↓ 6.21	↓ 0.30	↑ 0.36	

Figure 3). For a given model f and dataset D with N_D possible shots, each with s_k examples, we approximate AUC according to the trapezoidal rule:

$$\text{AUC}(f, D) = \frac{1}{s_{N_D}} \sum_{k=1}^{N_D} \frac{f(s_{k-1}) + f(s_k)}{2} \Delta s_k. \quad (1)$$

The results are listed in Table 18. In summary, UC² is by far the strongest model across all tasks; xUNITER comes second and it outperforms M³P (which ranks the third) on reasoning and retrieval; mUNITER performs the worst in general, except for being better than M³P on reasoning.

Gap between Zero-shot and Few-shot Learning. In contrast to previous results on text-only tasks, we find few-shot learning for multilingual multimodal tasks to be especially challenging, possibly due to the complex nature of the respective tasks combined with the sparsity of the few-shot learning data. As plotted in Figure 3, we observe no clear improvement as the number of shots increase in all tasks, the only exception being xGQA. This suggests that cost-efficiently hand-labelling ~ 100 shots in order to improve the models performance on the target language, unlike in text-only setups, is insufficient to make meaningful progress on V&L tasks such as XVNLI, MaRVL and xFlickr&Co. This also stresses the importance of developing methods for *improved* zero-shot transfer in future work, as V&L models still require sufficient amounts of expensive task-annotated data to benefit from in ‘non-zero-shot’ setups (Table 19).

In contrast to XVNLI, MaRVL and xFlickr&Co, we find that few-shot experiments on xGQA yield large performance

Table 20. Caption density in xFlickr&CO.

xFlickr&CO			DEU IR		JPN IR		DEU TR		JPN TR	
# shots	# img	# cap/img	xUNI	UC ²	xUNI	UC ²	xUNI	UC ²	xUNI	UC ²
500	500	1	17.15	27.20	11.95	29.35	14.80	21.40	13.05	25.45
	100	5	13.70	30.60	14.00	23.60	9.80	25.10	12.50	21.70
1000	1000	1	15.50	30.55	12.70	28.90	14.20	24.65	12.65	26.00
	200	5	15.85	26.30	14.35	28.40	14.25	20.35	12.15	26.30

gains in the target language. We attribute this to the structured nature of the data and the large label space of the task (1,842 classes) resulting in a misaligned multilingual embedding space in zero-shot setups, as suggested by the original authors. We analyse these results further in §D.

Increasing the Number of Shots. In the few-shot experiments, we noticed no observable improvement or clear trend of performance increase on 3 out of the 4 tasks. In order to identify if, and at what point, increasing the number of shots would start being beneficial, we extend our few-shot experiments on the retrieval task using the xFlickr&CO dataset. In Figure 5, we plotted DEU and JPN performance with up to 1,500 shots¹⁷ (instead of the 100 shots used in our standard few-shot setup for this dataset) for the two best performing models in this task. Indeed, we observe an overall trend of increasing performance with more training examples, although with noticeable instability issues (e.g. UC² accuracy in JPN drops when more than 500 shots are used). These results demonstrate that more training examples may indeed further improve the overall performance in the target language, but the large number of examples required questions whether this setup can even be dubbed ‘few-shot’.

More Images versus More Captions. An important question, which concerns data collection efficiency for the target languages in image–text retrieval scenarios, is whether it is more useful, for a given budget, to collect (1) more *captions per image* or (2) more *images with a single caption*. To investigate this question, we trained xUNITER and UC² on JPN and DEU data with different number of shots and different ratios of captions per image (Table 20). As indicated by the results, there is no clear trend suggesting that one setting is better than the other. We leave a larger-scale investigation related to this question for future work.

D. Further Analyses

We now provide a series of further and finer-grained analyses, which include performance of the ‘translate train’ transfer approach and domain-specific results in xFlickr&CO, as well as an analysis of performance over structurally different question types of xGQA in few-shot setups.

¹⁷DEU and JPN captions are from Multi30K (Elliott et al., 2016) and Nakayama et al. (2020), respectively.

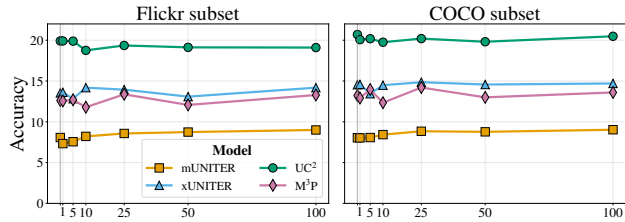


Figure 12. Image retrieval performance averaged across languages per xFlickr&CO subset.

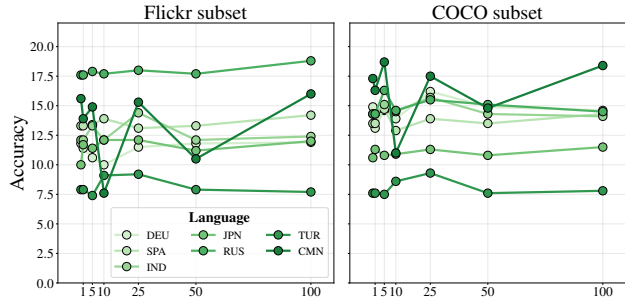


Figure 13. UC² image retrieval performance across languages per xFlickr&CO subset.

Table 21. English image retrieval performance by evaluation subset in xFlickr&CO after fine-tuning.

xFlickr&CO subset	mUNITER	xUNITER	UC ²	M ³ P
Flickr	51.3	46.6	44.0	37.1
COCO	37.7	30.3	30.8	25.6

Translate Train Performance. We further compare the few-shot results with ‘translate train’ performance: We directly fine-tune the pretrained UC² and xUNITER checkpoints on the full, human-translated Flickr30K training set. As shown in Figure 5, UC² achieves remarkable zero-shot and few-shot performance: 100 or fewer data points are enough for it to match ‘translate train’ performance, which it surpasses with 1,000 or more data points. Similar results can also be observed in Japanese, although it needs 500 data points to match ‘translate train’ accuracy. xUNITER—which was not pretrained on German nor Japanese multimodal data—also reaps benefits from 500 or more data points, although still being far from its ‘translate train’ counterpart (despite being initialised from English fine-tuned weights). Finally, we note that ‘translate train’ performance is still far from the corresponding English performance. This result is rather surprising for UC², which was solely pretrained on Conceptual Captions in five languages, each of equal size.

Retrieval Performance by Evaluation Subset. Our xFlickr&CO dataset is a composition of Flickr and COCO data (each of 1,000 image–sentence pairs). As models are fine-tuned on Flickr, the COCO evaluation subset might

require out-of-distribution generalisation due to its different visual domain. Indeed, this is what we observe in English (Table 21). However, Figure 13 shows that this is not the case for the multilingual encoders, which perform similarly in both subsets. A closer look at UC² further shows that this behaviour is language-independent. Given the relatively poor performance of current models, we cannot assess whether the domain shift we observe in English does not happen in xFlickr&CO due to its unified captioning guidelines, or whether domain generalisation will be a challenge in xFlickr&CO for future, tightly-multilingual models.