

Chitranuvad: Adapting Multi-Lingual LLMs for Multimodal Translation

Shaharukh Khan, Ayush Tarun, Ali Faraz, Palash Kamble, Vivek Dahiya,
Praveen Pokala *, Ashish Kulkarni *, Chandra Khatri *,
Abhinav Ravi * and Shubham Agarwal *

Krutrim AI, Bangalore, India

* Senior Contributors

Contact: {shaharukh.khan, abhinav.ravi, shubham.agarwal1}@olakrutrim.com

Abstract

In this work, we provide the system description of our submission as part of the *English-to-Lowres Multimodal Translation Task* at the Workshop on Asian Translation (WAT2024). We introduce Chitranuvad, a multimodal model that effectively integrates Multilingual LLM and a vision module for Multimodal Translation. Our method uses a ViT image encoder to extract visual representations as visual token embeddings which are projected to the LLM space by an adapter layer and generates translation in an autoregressive fashion. We participated in all the three tracks (Image Captioning, Text-only and Multimodal translation tasks) for Indic languages (ie. English translation to Hindi, Bengali and Malayalam) and achieved SOTA results for Hindi in all of them on the Challenge set while remaining competitive for the other languages in the shared task.

1 Introduction

Recently, there has been an increased interest in Multimodal Machine Translation (MMT) task (Calixto and Liu, 2017; Delbrouck and Dupont, 2017; Elliott and Kádár, 2017; Yao and Wan, 2020) which involves translation between language pairs, incorporating other modalities (like images) as an auxiliary information. The visual cues act as ‘symbol grounding’ (Fodor, 1975; Harnad, 1990, 2003, 2005), helping to resolve ambiguities in language (Rainie et al., 2012; Hu et al., 2014; Specia et al., 2016; van Miltenburg et al., 2019; Caglayan et al., 2020) by learning to connect language and perception (Mooney, 2008; Bisk et al., 2020). For example, in order to correctly translate the word *court* in Figure 1, the model has to infer from the image that the statement is about tennis court and not the court as government institution.

Prior works mostly focused on translation from English to European languages (Elliott et al., 2016; Specia et al., 2016) while the Indic languages remain largely unexplored, with an exception of the

MMT shared task at the Workshop on Asian Translation (WAT) (Nakazawa et al., 2019, 2020, 2021, 2022, 2023).


The *English-to-Lowres Multimodal Translation Task* at WAT-2024 targets the MMT task for three Indic medium-to-low-resource languages (Hindi, Bengali, Malayalam) and a low-resource African language Hausa. To assess the importance of the image modality, the task has been decoupled into three tracks: **1).** *Text-only translation* where the source image is not used, **2).** *Image Captioning* where English source text is not used and **3).** *Multimodal translation* which uses both the image and the text. We participated in all the three tracks for Indic languages only (Hindi, Bengali, Malayalam) under a non-constrained and proprietary multilingual and multimodal Large Language Model (LLM): *Chitranuvad*¹.

In this paper, we provide a description of our multimodal LLM where we leverage a multilingual LLM backbone Krutrim (Team, 2024b), coupled with a visual image encoder. Our contributions could thus be summarized as follows:

- We introduce Chitranuvad, a Large Multimodal model, adapted for multi-lingual translation, which leverages images and language modalities to provide an image grounded translation of the English sentence in the target Indic languages.
- We showcase the effectiveness of task specific finetuning on the Visual Genome translation datasets and achieve SOTA performance.
- We evaluate Chitranuvad and prior baselines on the English-to-Lowres Multimodal Translation Task and demonstrate the ability of our model to perform grounded translation, using different training strategies and ablations.

The rest of the paper is organized as: Section 2 presents related research on multimodal machine

¹Chitranuvad literally means Image Translate in Hindi



English: Tennis player in a court
Object Tags: person,tennis racket,sports ball
Hindi: एक अदालत में टेनिस खिलाड़ी
Bengali: টেনিস খেলোয়াড় একটি আদালতে
Malayalam: ഒരു കോർട്ടിൽ ടെന്നീസ് കളിക്കാരൻ



English: A bears right black ear.
Object Tags: bear
Hindi: एक भालू सही काले कान।
Bengali: একটি ভালুক কানের কান ভালুক।
Malayalam: ഒരു കരടിയുടെ വലത് കറുത്ത ചെവി

Figure 1: Multimodal Machine Translation task as part of English-to-lowres track where the source sentence is translated to multiple Indic languages (Hindi, Bengali, Malayalam) grounded in the image. Meaning of words like "court" and "right" in the translations can vary significantly depending on the visual context.

translation while Section 3 explains our Chitranuvad model recipe in detail. We present the datasets used in Section 4, followed by experimental findings in Section 5 and conclusion in Section 6.

2 Related Work

Early Neural Machine Translation (NMT) and Image captioning systems (Show, 2015; Gao et al., 2018) were based on Recurrent Neural Networks (RNNs) and their variants (Cho et al., 2014; Sutskever et al., 2014; Cho, 2014; Hochreiter and Schmidhuber, 1997), often incorporating attention mechanisms (Bahdanau et al., 2014). The seminal work of transformers (Vaswani, 2017) paved the way for the development of high-quality image captioning (Chen et al., 2021) as well as translation systems (Lewis, 2019), even for low-resource languages (Dabre et al., 2021; Gala et al., 2023a). Multimodal Machine Translation (MMT) systems witnessed a similar shift in their approach (Caglayan et al., 2016; Yao and Wan, 2020; Guo et al., 2023). Prior submissions to the MMT task at Workshop on Asian Translation (Gain et al., 2021; Gupta et al., 2021; Parida et al., 2022; Dash et al., 2023; Shahid et al., 2023) also fall in this category.

The next generation of Multimodal LLMs (Lu et al., 2024a; Laurençon et al., 2024; Tong et al., 2024; Xue et al., 2024) can handle a variety of complex tasks, including machine translation and captioning, by utilizing cutting-edge architectures

as an unified general purpose agent. These models often rely on pre-trained LLMs, with an exception of few, which train the models from scratch (Team, 2024a; Lu et al., 2024b). Most of these Vision Language Models (VLMs) follow the architecture of (Liu et al., 2023a) where a CLIP (Radford et al., 2021) or a similar encoder is used to encode the image and projected into LLM’s representation space using an adapter layer. Notably, Wang et al. (2023) offers a departure from conventional architectures by using distinct matrices and Feed Forward Networks for image modalities. Recent developments replace the image encoder with SigLIP (Zhai et al., 2023a) and the single-layer MLP projector with attention-based pooling (Laurençon et al., 2024).

Advanced backbone LLMs (Brown et al., 2020; Touvron et al., 2023; Achiam et al., 2023; Team et al., 2023; Jiang et al., 2024; Team et al., 2024) however have a primary focus for English and European languages. There have been relatively few LLMs for Indic languages, such as Airavata (Gala et al., 2024), Navarsa (Labs, 2023), Kannada LLaMA, Tamil LLaMA (Balachandran, 2023), Odia LLaMA (Kohli et al., 2023), to name a few. However, most of these LLMs are an extension and finetuned version of LLaMA/Gemma for Indic languages, which don’t fully capture the nuances of the language. This could be attributed to the fact that Indic languages are under-represented in Common Crawl (which majorly forms the train-

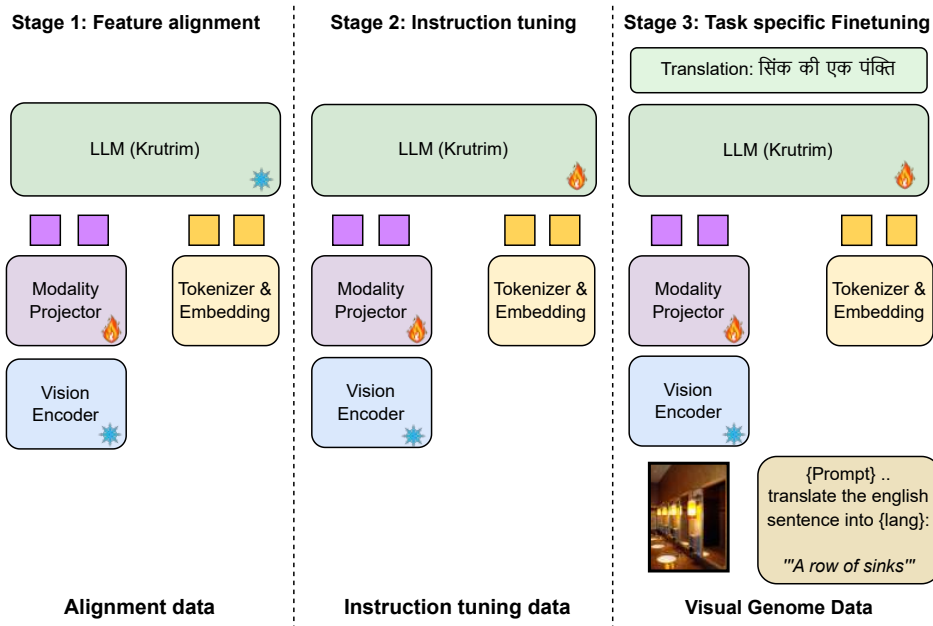


Figure 2: Chitranuvad model architecture with the three stage training pipeline described in Section 3.

ing corpus of LLMs), despite India constituting 18% of the global population. Hindi, for example, does not show-up in the top 20 languages despite being the 3rd most spoken (Buck et al., 2014; Penedo et al., 2023). Closed-source models such as Krutrim (Team, 2024b) and Sutra (Bendale et al., 2024) represent exceptions, as they are trained from scratch. Currently, PALO (Maaz et al., 2024) is a multimodal LLM that supports only Hindi and Bengali. However, to the best of our knowledge, there are no other open-source multimodal LLMs trained specifically for low-resource Indic languages. In contrast, we developed a multilingual multimodal system that supports 10 Indic languages.

3 Model and Training Recipes

Figure 2 provides an overview of our architecture and the multi-stage training pipeline. Our *Chitranuvad* model architecture borrows heavily from LLaVA-like models (Liu et al., 2023a, 2024), where we use pre-trained Krutrim LLM (Team, 2024b) instead, as the autoregressive multi-lingual LLM backbone. Our Krutrim LLM is trained across 10 languages and natively supports all the 3 Indic languages (Hindi, Bengali, Malayalam) used as part of the shared task.

For the multimodal training, we first encode images through a vision encoder. Next, the modality projection (adapter/connector) layer projects the vision embeddings into the LLM embedding

space, creating a sequence of visual tokens. The multi-lingual LLM then generates the response conditioned on these visual embedding tokens. The Krutrim LLM model supports a context length of 4096 tokens, out of which 576 tokens are used for the image representation, obtained after the modality projector layer. For the projection layer, we experiment with both single layer projection (Liu et al., 2023b) as well as a two-layer MLP vision-language connector with non-linearity (Liu et al., 2023a). We also experiment with pre-trained CLIP ViT-L/14@336px (Radford et al., 2021) as well as SigLIP-SO400M (Zhai et al., 2023b) for the vision encoder. Similar to the LLaVA model, we generate multi-turn conversational data for instruction tuning our model, which we expand upon in Section 4. We train our model in multiple stages:

Stage 1: Pre-Training (PT) for Feature Alignment. In this stage, we do the pre-training with image-text pairs, where the projector layer is trained while the vision encoder and LLM is kept frozen. Here, each sample is treated as a single-turn conversational instruction tuning data.

Stage 2: Instruction Tuning. Similar to LLaVA models (Liu et al., 2023b,a), we also keep the vision encoder frozen during the second stage of training. However, here we also update the LLM weights apart from tuning the modality projection layer. This stage aims to build a general purpose Multimodal agent (chatbot) which can follow com-

Split	#Instances	English	Hindi	Bengali	Malayalam
Train	28930	5.09	5.13	4.07	3.86
Valid	998	5.08	5.04	4.06	3.75
Test	1595	5.07	4.95	4.14	3.76
Challenge	1400	6.04	6.35	4.92	4.48

Table 1: Total number of instances and average number of tokens for the text in English and splits of different Visual Genome datasets in other languages.

Multimodal Translation:

Human: You are given an image and coordinates of a bounding box as: $x1=\{x1\}$, $y1=\{y1\}$, $x2=\{x1+x2\}$, $y2=\{y1+y2\}$. Using the context of the objects or items available in the bounding box translate the following sentence from English into $\{lang\}$ language. You are also provided labels of the objects in the image as: $\{labels\}$. English sentence is: $\{sentence\}$.

System: $\{translation\}$.

Text only translation:

Human: Translate the following sentence from English into $\{lang\}$ language. English sentence is: $\{sentence\}$.

System: $\{translation\}$.

Image captioning: *Human:* You are given an image and coordinates of a bounding box as: $x1=\{x1\}$, $y1=\{y1\}$, $x2=\{x1+x2\}$, $y2=\{y1+y2\}$. You are also provided labels of the objects in the image as: $\{labels\}$. Provide a short caption of the object in $\{lang\}$ language.

System: $\{caption\}$.

Table 2: Different prompt templates for creating task specific fine-tuning data, used in Stage 3 training.

plex instructions across multiple-turns of the conversation. We focus on developing a specialized multimodal translation system in the Stage 3.

Stage 3: Task-specific Fine-Tuning. We follow a similar recipe to that of Stage 2 for the (Machine Translation) task-specific fine-tuning and update weights for the projection layer and the LLM while keeping the vision encoder frozen. Here, we experiment with both LoRA style training (Hu et al., 2021; Houlsby et al., 2019) as well as full parameter fine-tuning on the shared task translation data.

4 Dataset

In this section, we describe the data resources utilized throughout our experiments.

Stage 1: In our initial experiments, we use the LLaVA-Pretrain-LCS-558K data for pre-training our model in Stage 1. However, recent works (Tong et al., 2024) showed that more adapter data is beneficial for the model, such as the 1.2M ShareGPT4V-PT (Chen et al., 2023) image-captioning dataset, which we use in Stage 1 training. We also trans-

lated this data in the 10 Indic languages that our LLM natively supports, using an in-house text Machine Translation system. We sample translations across different languages (including English) in an equal ratio and ensure that PT data limits to 1.2M data points in our final data mix.

Stage 2: For the second stage instruction tuning, eliciting visual reasoning abilities, we experiment with both LLaVA-Instruct-150K (Liu et al., 2023b) and LLaVA-1.5-665K (Liu et al., 2024) where we find continued improvements with the 665K version. Similar to pre-training data, we also translated the LLaVA-1.5-665K into multiple languages. Recently released Cauldron dataset (Laurençon et al., 2024) is a collection of 50 academic Vision-language tasks. In our final submission, we also include the translated versions and the original English language based Cauldron apart from the proprietary multi-modal dataset in the training mix. It must be noted that the English only Visual Genome might be a part of this stage’s training data through various academic datasets, though not for the translation task.

Stage 3: For the Stage 3, we work with the aligned multi-lingual Visual Genome (Krishna et al., 2017) datasets, i.e. Hindi (Parida et al., 2019), Bengali (Sen et al., 2022) and Malayalam (Parida and Bojar, 2021), bundled as part of the shared task. Each row in the dataset consists of the following fields: *i*). MS COCO (Lin et al., 2014) image id *ii*). English utterance *iii*). Translated utterance in Hindi/ Bengali/ Malayalam *iv*). Bounding box of the area in the image that the utterance is based on. While there is also a track for Hausa language (Abdulmumin et al., 2022), we don’t include this in our training data. Table 1 provides the statistics of the different versions of the dataset, which we transform into instruct tuning format, similar to Stage 1 and 2 data (see Table 2). To increase the efficacy of our model, we enrich the dataset with the labels of different objects in the image (object tags), similar to (Gupta et al., 2021). We use SOTA (state-of-the-art) YOLOv8 (Varghese and Sambath, 2024) for object detection compared to the prior works, which relied on Faster R-CNN models (Wu et al., 2019; Girshick, 2015). We also calculate the Intersection-over-union (IoU) for the detected and the dataset provided bounding boxes to get the most relevant object tag. However, we found a decreased performance against including the labels of all the detected objects.

Submission	Hi-Ch		Hi-Test		Bn-Ch		Bn-Test		MI-Ch		MI-Test	
	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑
SILO NLP	29.6	0.73	36.2	0.79	22.6	0.61	41.0	0.77	14.6	0.39	30.8	0.60
NLP Voices	41.8	0.81	43.1	0.82	32.9	0.71	39.8	0.75	19.6	0.54	30.6	0.64
Volta	51.7	0.86	44.1	0.82	-	-	-	-	-	-	-	-
ODIAGEN	53.6	0.86	44.6	0.83	47.8	0.82	49.2	0.8	39.7	0.75	46.6	0.75
Ours (leaderboard)	54.1	0.86	43.3	0.81	44.2	0.79	45.1	0.77	34.0	0.65	37.8	0.63
Ours†	55.3	0.87	44.7	0.83	46.7	0.81	48.1	0.79	40.6	0.75	51.7	0.88

Table 3: English-to-lowres leaderboard scores for Text-only task for Indic languages (Team 007). In the following tables, †denotes the results after submission deadline using the IndicTrans2 evaluation scripts, all the other results are reported using the shared task dashboard.

Method	Hi-Ch		Hi-Test		Bn-Ch		Bn-Test		MI-Ch		MI-Test	
	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑
ODIANLP	0	0.04	0.8	0.06	-	-	-	-	-	-	-	-
NLPHUT	-	-	-	-	-	-	-	-	0.9	0.02	0.9	0.05
Ours (leaderboard)	1.3	0.13	2.8	0.18	0.4	0.04	1.8	0.11	0.3	0.04	0.9	0.06

Table 4: English-to-lowres leaderboard scores for Image captioning track. Ours is the only multi-lingual model which can handle all the 3 Indic languages for image captioning.

5 Experimental Results and Discussion

This section details our experimental setup and presents the results of our comparative studies.

5.1 Implementation

We use HuggingFace Transformers (Wolf et al., 2019) based on PyTorch (Paszke et al., 2019) for our experiments. We consider PALO (Maaz et al., 2024) as a multi-lingual multi-modal baseline and use the code provided with the repository². The shared task provides a leaderboard based on the automatics metrics of BLEU (Papineni et al., 2002) and RIBES (Isozaki et al., 2010). For reporting BLEU, we used the evaluation scripts³ provided with (Gala et al., 2023b) and the official repository for RIBES⁴. Similar to previous works (Gupta et al., 2021), we also report the results after tokenizing the outputs using indic-tokenizer⁵. Our Stage 1 and Stage 2 tuning follow similar hyperparameters as the LLaVA model (Liu et al., 2023b) unless specified otherwise. For Stage 3 fine-tuning, we conducted multiple experiments for hyperparameter search of learning rate (1e-3, 1e-4, and 1e-5); as well as multiple epochs (1, 2, 3, and 5). We observed rapid over fitting after only one epoch while a learning rate of 1e-4 yielded the highest overall performance. All our further experiments are reported based on this configuration.

²<https://github.com/mbzuai-oryx/PALO>

³<https://github.com/AI4Bharat/IndicTrans2>

⁴<https://github.com/nttcs-lab-nlp/RIBES>

⁵<https://github.com/ltrc/indic-tokenizer>

5.2 Results for different tracks

Table 3, 4 and 5 present the results for text-only, image captioning and the Multimodal translation task respectively. For the text-only task, our Chitravad model was trained with image data till Stage 2. In Stage 3, we only finetune with text only translations. During inference, we prompt the model with text only translations and dont provide images. Our model achieves SOTA on Hindi and Malayalam Challenge and Test sets while being competitive on the Bengali dataset (see Table 3). We were the only submission which could do image captioning in all the 3 languages (see Table 4). For the MMT task, we achieved SOTA on Hindi Challenge and Malayalam test set while being competitive on the other languages. We also provide cherry-picked system outputs of our best Multimodal LLM in Table 3. From our manual inspection, we saw that our generated translations are better than the ground truth. For example, in the last snippet, our model correctly translates the word ‘downhill’, which the gold translation fails to capture.

5.3 0-shot on the Shared task data

We evaluate the efficacy of our model after Stage 2 as the 0-shot setting, where we don’t fine-tune specifically for the shared task translation data. In our preliminary experiments, we only use the English versions of the datasets mentioned in Section 4 for both Adapter tuning (Stage 1) and Instruction tuning (Stage 2). Exceptionally, our Krutrim LLM still retained multi-lingual capabilities, ev-

Submission	Hi-Ch		Hi-Test		Bn-Ch		Bn-Test		MI-Ch		MI-Test	
	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑
IIT-P	37.5	0.79	42.5	0.81	-	-	-	-	-	-	-	-
ODIAGEN	42.8	0.82	41.6	0.81	30.5	0.69	42.4	0.76	-	-	-	-
Volta	51.6	0.86	44.6	0.82	-	-	-	-	-	-	-	-
BITS-P	52.1	0.85	45.0	0.83	48.7	0.83	50.6	0.81	42.2	0.76	51.9	0.80
Ours (leaderboard)	53.4	0.842	43.7	0.81	44.8	0.78	44.5	0.76	39.8	0.74	51.9	0.78
Ours†	54.7	0.86	<u>43.9</u>	0.83	<u>46.9</u>	<u>0.81</u>	<u>47.7</u>	<u>0.79</u>	<u>40.3</u>	<u>0.74</u>	51.9	0.93

Table 5: English-to-lowres leaderboard Scores for Multimodal translation track across multiple languages (Team 007). †denotes the results after submission deadline using the IndicTrans2 evaluation scripts

Method	Hi-Ch		Hi-Test		Bn-Ch		Bn-Test		MI-Ch		MI-Test	
	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑
PALO-7B	14.8	0.585	13.3	0.567	7.9	0.469	9.6	0.464	0.1	0.001	0	0
PALO-13B	15.8	0.605	14.9	0.605	6.7	0.44	7.0	0.45	0.1	0.004	0	0
Chitranuvad (Eng)	18.3	0.629	12.9	0.585	8.7	0.512	8.3	0.477	8.7	0.487	7.3	0.426
Chitranuvad (Eng+Hindi)	20.0	0.698	14.8	0.653	9.4	0.537	8.9	0.494	9.2	0.511	8.6	0.466
Chitranuvad (Multilingual)	25.0	0.694	19.0	0.66	11.4	0.569	9.7	0.515	12.2	0.54	10.3	0.486

Table 6: 0-shot results for Multimodal Machine Translation track as discussed in Section 5.3. *Eng* denotes only English data is used in Stage 1 and 2. *Eng+Hindi* denotes English and Hindi data in Stage 2. As expected, we find the best scores when the training data mix consists of data from the 10 Indic languages.

ident from the scores in Table 6. When we also include Hindi data in the training mix, we find an uplift on the Hindi translation task. Including multi-lingual data in both the stages further showed improvement on all three language translation tasks in the 0-shot setting. We thus use this as the base model in the Stage 3 training. We also evaluate against the open-source baseline of PALO-7/13B (Maaz et al., 2024) in the 0-shot setting. To our surprise, our Chitranuvad model consistently outperforms the 0-shot PALO baseline, even when our model is fine-tuned with English only data in both the stages. We hypothesize that this is because the base LLM Vicuna (Zheng et al., 2024) used in PALO is not inherently multi-lingual in nature.

5.4 Other fine-tuning approaches

In this section, we elaborate on key findings with different fine-tuning approaches, with all the results reported in Table 7.

LoRA vs Full finetuning. We investigated the efficacy of full fine-tuning versus Low-Rank Adaptation (LoRA) using Visual Genome data. Our experiments (see Table 7) reveal that full fine-tuning consistently outperforms LoRA, i.e. LoRA learns less (Biderman et al., 2024).

Bi-lingual vs Multi-lingual For Stage 3 training, we experiment with training specialized models for each language (Hindi and Bengali) compared to multi-lingual setting with a mix of data from all the three Indic languages. We didn’t find any im-

provement over multi-lingual model but instead observe catastrophic forgetting (Zhai et al., 2023c; Tong et al., 2024), where the translation abilities of the model in the other languages deteriorate completely. We hypothesize that a mix of multiple languages probably act as regularization and enhance the general translation capabilities.

Do we need second stage training? Similar to (Tong et al., 2024), we investigate if we even need Stage 2 instruction tuning. We find that our model, if finetuned directly on Visual Genome translation data (i.e. Stage 1 and 3 training only) performed comparable to the previous baselines. Including Stage 2 training provided an uplift in the scores with an added advantage of building a general purpose Multimodal agent.

Back translation Back translations, i.e. using the reverse translations have been a popular technique both for data augmentation as well as post-processing or re-ranking techniques in traditional Machine Translation and Natural Language Generation systems (Sennrich et al., 2015; Li et al., 2015; Agarwal et al., 2018; Edunov, 2018; Graça et al., 2019). This involves re-translating content from the target language to its source language. Thus, apart from the original task of En -> Hi/Bn/MI, we also included in our training corpus, the task of reverse translation from Hi/Bn/MI -> En in the Stage 3 training mix. However, in our experiments, we found that this strategy showed decreased performance in terms of automatic metrics.



Object tags: person, bowl, oven, dog
Bbox tag: dog
English: a woman holding a dog

Hindi GT: एक स्त्री जो कुत्ता रखती है
Hindi Translated: एक महिला एक कुत्ते को पकड़े हुए
Bengali GT: একটি মহিলা একটি কুকুর ধরে
Bengali Translated: একটি মহিলা একটি কুকুর ধরে
Malayalam GT: ഒരു സ്ത്രീ നായയെ പിടിക്കുന്നു
Malayalam Translated: ഒരു സ്ത്രീ നായയെ പിടിക്കുന്നു



Object tags: person, person, person, tennis racket
Bbox tag: tennis racket
English: A TENNIS RACKET

Hindi GT: एक टेनिस रैकेट
Hindi Translated: एक टेनिस रैकेट
Bengali GT: একটি টেনিস র্যাকেট
Bengali Translated: একটি টেনিস র্যাকেট
Malayalam GT: ഒരു ടെന്നീസ് റാക്കറ്റ്
Malayalam Translated: ഒരു ടെന്നീസ് റാക്കറ്റ്



Object tags: person, snowboard, backpack
Bbox tag: person
English: snow on the ground

Hindi GT: बर्फ़ जमीन पर
Hindi Translated: जमीन पर बर्फ़
Bengali GT: মাটিতে বরফ
Bengali Translated: মাটিতে বরফ
Malayalam GT: നിലത്ത് മഞ്ഞ
Malayalam Translated: നിലത്ത് മഞ്ഞ



Object tags: dog
Bbox tag: dog
English: white flower on curtain

Hindi GT: पर्दे पर सफेद फूल
Hindi Translated: पर्दे पर सफेद फूल
Bengali GT: পর্দার উপর সাদা ফুল
Bengali Translated: পর্দায় সাদা ফুল
Malayalam GT: തിരശ്ശീലയിൽ വെളുത്ത പുഷ്പം
Malayalam Translated: വെളുത്ത പുഷ്പം



Object tags: person, skis
Bbox tag: person
English: Woman going fast downhill.

Hindi GT: तेज गति से जा रही महिला।
Hindi Translated: नीचे की ओर तेजी से जा रही महिला।
Bengali GT: মহিলা দ্রুত উতরাইয়ের দিকে যাচ্ছে।
Bengali Translated: মহিলা দ্রুত নিচে যাচ্ছে।
Malayalam GT: വേഗത്തിൽ താഴേക്ക് പോകുന്ന സ്ത്രീ.
Malayalam Translated: മലയിറങ്ങുന്ന വേഗത്തിൽ സ്ത്രീ.

Figure 3: English-to-lowres Multimodal Machine Translation track supports translation of source sentence into multiple Indic languages (Hindi, Bengali, Malayalam). We enrich the dataset to include labels of all the identified objects. We show the outputs of our best model which is trained with a mix of multi-lingual data in all the 3 stages.

Method	Hi-Ch		Hi-Test		Bn-Ch		Bn-Test		MI-Ch		MI-Test	
	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑	BLEU ↑	RIBES ↑
LoRA	42.1	0.721	34.5	0.770	28.3	0.69	30.4	0.669	23.2	0.61	27.0	0.601
Bi-lingual (Hi Stage 3)	53.0	0.848	43.3	0.833	0.3	0.006	0.1	0.001	0.2	0.003	0	0
Bi-lingual (Bn Stage 3)	0.3	0.005	0.1	0.001	45.4	0.797	46.6	0.781	0.3	0.003	0.1	0.001
Only Stage 1, 3	53.6	0.853	43.4	0.826	45.2	0.801	46.4	0.788	38.2	0.735	50.3	0.781
Back Translation mix	53.8	0.856	43.6	0.828	46.0	0.806	46.8	0.792	37.5	0.729	46.3	0.738

Table 7: Different finetuning strategies for Multimodal Machine Translation track as described in Section 5.4 in the order of discussion. Our Stage 3 full finetuning training performs the best compared to the other training recipes.

6 Conclusion

We present Chitravad, a multimodal LLM that is adapted for image grounded Machine translation. Our model encodes images using a pre-trained image encoder (Alexey, 2020) and translates the English sentences autoregressively into different Indic languages (Hindi, Bengali, Malayalam) using a pre-trained LLM. Empirically, our model outperforms previous baselines for different tasks. However, we also observed that vision modality had little impact on the translation, echoing the observations from (Grönroos et al., 2018; Lala et al., 2018; Wu et al., 2021; Li et al., 2022).

Broader Impact: We believe our work paves way for building next generation assistants which can do multimodal machine translation. We believe these systems can empower different sectors like education, healthcare, banking and financial services, etc. to name a few.

Limitations and Future Work: While this work is focused to three Indic languages (Hindi, Bengali, Malayalam), we consider our approach as a first step towards building general purpose multilingual system which can handle various Indic languages. While in our current setup, we freeze the vision encoder during training, recent works have shown that unfreezing the vision encoder with Perceiver Resampler (Jaegle et al., 2021), helps learn better representations (Laurençon et al., 2024; Tong et al., 2024), which we plan to explore in the future.

Acknowledgements

We thank Bhavish Aggarwal and the rest of the Krutrim team which helped with model development at various stages. Specifically, we want to thank Anagha Bhangare, Raja Kolla and Aditya Kallappa for the discussions. Our models were trained with generous support from Krutrim cloud using Krutrim credits. We also thank the challenge organizers as well as the reviewers for their valuable feedback and suggestions.

References

- Idris Abdulmumin, Satya Ranjan Dash, Musa Abdul-lahi Dawud, Shantipriya Parida, Shamsuddeen Hassan Muhammad, Ibrahim Sa’id Ahmad, Subhadarshi Panda, Ondřej Bojar, Bashir Shehu Galadanci, and Bello Shehu Bello. 2022. Hausa visual genome: A dataset for multi-modal english to hausa machine translation. *arXiv preprint arXiv:2205.01133*.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Shubham Agarwal, Marc Dymetman, and Eric Gaussier. 2018. Char2char generation with reranking for the e2e nlg challenge. *arXiv preprint arXiv:1811.05826*.
- Dosovitskiy Alexey. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv: 2010.11929*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Abhinand Balachandran. 2023. [Tamil-llama: A new tamil language model based on llama 2](#).
- Abhijit Bendale, Michael Sapienza, Steven Ripplinger, Simon Gibbs, Jaewon Lee, and Pranav Mistry. 2024. Sutra: Scalable multilingual language model architecture. *arXiv preprint arXiv:2405.06694*.
- Dan Biderman, Jose Gonzalez Ortiz, Jacob Portes, Mansheej Paul, Philip Greengard, Connor Jennings, Daniel King, Sam Havens, Vitaliy Chiley, Jonathan Frankle, et al. 2024. Lora learns less and forgets less. *arXiv preprint arXiv:2405.09673*.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirre Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, et al. 2020. [Experience grounds language](#). *CoRR abs/2004.10151*.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child,

- Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).
- Christian Buck, Kenneth Heafield, and Bas Van Ooyen. 2014. N-gram counts and language models from the common crawl. In *LREC*, volume 2, page 4.
- Ozan Caglayan, Loïc Barrault, and Fethi Bougares. 2016. Multimodal attention for neural machine translation. *arXiv preprint arXiv:1609.03976*.
- Ozan Caglayan, Julia Ive, Veneta Haralampieva, Pranava Madhyastha, Loïc Barrault, and Lucia Specia. 2020. [Simultaneous machine translation with visual context](#). In *CoRR abs/2009.07310*.
- Iacer Calixto and Qun Liu. 2017. [Incorporating global visual features into attention-based neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 992–1003, Copenhagen, Denmark. Association for Computational Linguistics.
- Haishun Chen, Ying Wang, Xin Yang, and Jie Li. 2021. Captioning transformer with scene graph guiding. In *2021 IEEE international conference on image processing (ICIP)*, pages 2538–2542. IEEE.
- Lin Chen, Jisong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. 2023. Sharegpt4v: Improving large multimodal models with better captions. *arXiv preprint arXiv:2311.12793*.
- Kyunghyun Cho. 2014. Learning phrase representations using rnn encoder–decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder–decoder for statistical machine translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Raj Dabre, Himani Shrotriya, Anoop Kunchukuttan, Ratish Puduppully, Mitesh M Khapra, and Pratyush Kumar. 2021. Indicbart: A pre-trained model for indic natural language generation. *arXiv preprint arXiv:2109.02903*.
- Amulya Dash, Hrithik Raj Gupta, and Yashvardhan Sharma. 2023. Bits-p at wat 2023: Improving indic language multimodal translation by image augmentation using diffusion models. In *Proceedings of the 10th Workshop on Asian Translation*, pages 41–45.
- Jean-Benoit Delbrouck and Stéphane Dupont. 2017. [An empirical study on the effectiveness of images in multimodal neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 910–919, Copenhagen, Denmark. Association for Computational Linguistics.
- Sergey Edunov. 2018. Understanding back-translation at scale. *arXiv preprint arXiv:1808.09381*.
- Desmond Elliott, Stella Frank, Khalil Sima’an, and Lucia Specia. 2016. Multi30k: Multilingual english-german image descriptions. *arXiv preprint arXiv:1605.00459*.
- Desmond Elliott and Ákos Kádár. 2017. [Imagination improves multimodal translation](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 130–141, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Jerry A Fodor. 1975. *The language of thought*, volume 5. Harvard university press.
- Baban Gain, Dibyanayan Bandyopadhyay, and Asif Ekbal. 2021. Iitp at wat 2021: System description for english-hindi multimodal translation task. *arXiv preprint arXiv:2107.01656*.
- Jay Gala, Pranjal A Chitale, Raghavan AK, Varun Gumma, Sumanth Doddapaneni, Aswanth Kumar, Janki Nawale, Anupama Sujatha, Ratish Puduppully, Vivek Raghavan, et al. 2023a. Indictrans2: Towards high-quality and accessible machine translation models for all 22 scheduled indian languages. *arXiv preprint arXiv:2305.16307*.
- Jay Gala, Pranjal A Chitale, A K Raghavan, Varun Gumma, Sumanth Doddapaneni, Aswanth Kumar M, Janki Atul Nawale, Anupama Sujatha, Ratish Puduppully, Vivek Raghavan, Pratyush Kumar, Mitesh M Khapra, Raj Dabre, and Anoop Kunchukuttan. 2023b. [Indictrans2: Towards high-quality and accessible machine translation models for all 22 scheduled indian languages](#). *Transactions on Machine Learning Research*.
- Jay Gala, Thanmay Jayakumar, Jaavid Aktar Husain, Aswanth Kumar M, Mohammed Safi Ur Rahman Khan, Diptesh Kanojia, Ratish Puduppully, Mitesh M. Khapra, Raj Dabre, Rudra Murthy, and Anoop Kunchukuttan. 2024. [Airavata: Introducing hindi instruction-tuned llm](#).
- Lizhao Gao, Bo Wang, and Wenmin Wang. 2018. Image captioning with scene-graph based semantic concepts. In *Proceedings of the 2018 10th international conference on machine learning and computing*, pages 225–229.
- Ross Girshick. 2015. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448.

- Miguel Graça, Yunsu Kim, Julian Schamper, Shahram Khadivi, and Hermann Ney. 2019. [Generalizing back-translation in neural machine translation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 45–52, Florence, Italy. Association for Computational Linguistics.
- Stig-Arne Grönroos, Benoit Huet, Mikko Kurimo, Jorma Laaksonen, Bernard Merialdo, Phu Pham, Mats Sjöberg, Umut Sulubacak, Jörg Tiedemann, Raphael Troncy, et al. 2018. The memad submission to the wmt18 multimodal translation task. *arXiv preprint arXiv:1808.10802*.
- Junjun Guo, Junjie Ye, Yan Xiang, and Zhengtao Yu. 2023. Layer-level progressive transformer with modality difference awareness for multi-modal neural machine translation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Kshitij Gupta, Devansh Gautam, and Radhika Mamidi. 2021. Vita: Visual-linguistic translation by aligning object tags. *arXiv preprint arXiv:2106.00250*.
- Stevan Harnad. 1990. [The symbol grounding problem](#). *Physica D: Nonlinear Phenomena*.
- Stevan Harnad. 2003. Minds, machines and turing. In *The Turing Test*, pages 253–273. Springer.
- Stevan Harnad. 2005. To cognize is to categorize: Cognition is categorization. In *Handbook of categorization in cognitive science*, pages 21–54. Elsevier.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pages 2790–2799. PMLR.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Yuheng Hu, Lydia Manikonda, and Subbarao Kambhampati. 2014. [What we instagram: A first analysis of instagram photo content and user types](#). In *8th International Conference on Weblogs and Social Media, ICWSM 2014*.
- Hideki Isozaki, Tsutomu Hirao, Kevin Duh, Katsuhito Sudoh, and Hajime Tsukada. 2010. Automatic evaluation of translation quality for distant language pairs. In *Proceedings of the 2010 conference on empirical methods in natural language processing*, pages 944–952.
- Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. 2021. Perceiver: General perception with iterative attention. In *International conference on machine learning*, pages 4651–4664. PMLR.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Guneet Singh Kohli, Shantipriya Parida, Sambit Sekhar, Samirrit Saha, Nipun B Nair, Parul Agarwal, Sonal Khosla, Kusumlata Patiyal, and Debasish Dhal. 2023. [Building a llama2-finetuned llm for odia language utilizing domain knowledge instruction set](#).
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73.
- Telugu Labs. 2023. Navarsa: Indic llms based on gemma.
- Chiraag Lala, Pranava Swaroop Madhyastha, Carolina Scarton, and Lucia Specia. 2018. Sheffield submissions for wmt18 multimodal translation shared task. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 624–631.
- Hugo Laurençon, Léo Tronchon, Matthieu Cord, and Victor Sanh. 2024. What matters when building vision-language models? *arXiv preprint arXiv:2405.02246*.
- M Lewis. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Bei Li, Chuanhao Lv, Zefan Zhou, Tao Zhou, Tong Xiao, Anxiang Ma, and JingBo Zhu. 2022. On vision features in multimodal machine translation. *arXiv preprint arXiv:2203.09173*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306.

- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023a. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. [Visual instruction tuning](#).
- Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Yaofeng Sun, et al. 2024a. Deepseek-vl: towards real-world vision-language understanding. *arXiv preprint arXiv:2403.05525*.
- Jiasen Lu, Christopher Clark, Sangho Lee, Zichen Zhang, Savya Khosla, Ryan Marten, Derek Hoiem, and Aniruddha Kembhavi. 2024b. Unified-io 2: Scaling autoregressive multimodal models with vision language audio and action. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26439–26455.
- Muhammad Maaz, Hanoona Rasheed, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M Anwer, Tim Baldwin, Michael Felsberg, and Fahad S Khan. 2024. Palo: A polyglot large multimodal model for 5b people. *arXiv preprint arXiv:2402.14818*.
- Raymond J Mooney. 2008. [Learning to connect language and perception](#). In *AAAI*.
- Toshiaki Nakazawa, Chenchen Ding, Raj Dabre, Anoop Kunchukuttan, Nobushige Doi, Yusuke Oda, Ondřej Bojar, Shantipriya Parida, Isao Goto, and Hidayat Mino. 2019. Proceedings of the 6th workshop on asian translation. In *Proceedings of the 6th Workshop on Asian Translation*.
- Toshiaki Nakazawa, Kazutaka Kinugawa, Hideya Mino, Isao Goto, Raj Dabre, Shohei Higashiyama, Shantipriya Parida, Makoto Morishita, Ondrej Bojar, Akiko Eriguchi, Yusuke Oda, Akiko Eriguchi, Chenhui Chu, and Sadao Kurohashi, editors. 2023. *Proceedings of the 10th Workshop on Asian Translation*. Asia-Pacific Association for Machine Translation, Macau SAR, China.
- Toshiaki Nakazawa, Kazutaka Kinugawa, Hideya Mino, Isao Goto, Raj Dabre, Shohei Higashiyama, Shantipriya Parida, Makoto Morishita, Ondrej Bojar, Akiko Eriguchi, Yusuke Oda, Akiko Eriguchi, Chenhui Chu, et al., editors. 2022. *Proceedings of the 9th Workshop on Asian Translation*. International Conference on Computational Linguistics, Gyeongju, Republic of Korea.
- Toshiaki Nakazawa, Hideki Nakayama, Chenchen Ding, Raj Dabre, Anoop Kunchukuttan, Win Pa Pa, Ondřej Bojar, Shantipriya Parida, Isao Goto, Hidayat Mino, Hiroshi Manabe, Katsuhito Sudoh, Sadao Kurohashi, and Pushpak Bhattacharyya, editors. 2020. *Proceedings of the 7th Workshop on Asian Translation*. Association for Computational Linguistics, Suzhou, China.
- Toshiaki Nakazawa, Hideki Nakayama, Isao Goto, Hideya Mino, Chenchen Ding, Raj Dabre, Anoop Kunchukuttan, Shohei Higashiyama, Hiroshi Manabe, Win Pa Pa, Shantipriya Parida, Ondřej Bojar, Chenhui Chu, Akiko Eriguchi, Kaori Abe, Yusuke Oda, Katsuhito Sudoh, Sadao Kurohashi, and Pushpak Bhattacharyya, editors. 2021. *Proceedings of the 8th Workshop on Asian Translation (WAT2021)*. Association for Computational Linguistics, Online.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Shantipriya Parida and Ondřej Bojar. 2021. [Malayalam visual genome 1.0](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Shantipriya Parida, Ondřej Bojar, and Satya Ranjan Dash. 2019. Hindi Visual Genome: A Dataset for Multimodal English-to-Hindi Machine Translation. *Computación y Sistemas*, 23(4):1499–1505. Presented at CICLing 2019, La Rochelle, France.
- Shantipriya Parida, Subhadarshi Panda, Stig-Arne Grönroos, Mark Granroth-Wilding, and Mika Koistinen. 2022. Silo nlp’s participation at wat2022. *arXiv preprint arXiv:2208.01296*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only. *arXiv preprint arXiv:2306.01116*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Lee Rainie, Joanna Brenner, and Kristen Purcell. 2012. [Photos and videos as social currency online](#). *Pew Internet & American Life Project*.
- Arghyadeep Sen, Shantipriya Parida, Ketan Kotwal, Subhadarshi Panda, Ondřej Bojar, and Satya Ranjan Dash. 2022. Bengali visual genome: A multimodal dataset for machine translation and image captioning. In *Intelligent Data Engineering and Analytics*, pages 63–70. Springer.

- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Improving neural machine translation models with monolingual data. *arXiv preprint arXiv:1511.06709*.
- Sk Shahid, Guneet Singh Kohli, Sambit Sekhar, Debasish Dhal, Adit Sharma, Shubhendra Kushwaha, Shantipriya Parida, Stig-Arne Grönroos, and Satya Ranjan Dash. 2023. Odiagenai’s participation at wat2023. In *Proceedings of the 10th Workshop on Asian Translation*, pages 46–52.
- Attend Show. 2015. Tell: Neural image caption generation with visual attention kelvin xu. Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard Zemel, Yoshua Bengio *arXiv (2015-02-10) https://arxiv.org/abs/1502.03044 v3*.
- Lucia Specia, Stella Frank, Khalil Sima’an, and Desmond Elliott. 2016. [A shared task on multimodal machine translation and crosslingual image description](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 543–553.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.
- Chameleon Team. 2024a. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Krutrims Team. 2024b. Krutrims LLM: Multilingual foundational model for over a billion people. *Under Review*.
- Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, et al. 2024. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint arXiv:2406.16860*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Emiel van Miltenburg, Merel van de Kerkhof, Ruud Koolen, Martijn Goudbeek, and Emiel Kraemer. 2019. [On task effects in NLG corpus elicitation: A replication study using mixed effects modeling](#). In *INLG*.
- Rejin Varghese and M Sambath. 2024. Yolov8: A novel object detection algorithm with enhanced performance and robustness. In *2024 International Conference on Advances in Data Engineering and Intelligent Computing Systems (ADICS)*, pages 1–6. IEEE.
- Ashish Vaswani. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762*.
- Weihan Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, et al. 2023. Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. 2019. Detectron2. <https://github.com/facebookresearch/detectron2>.
- Zhiyong Wu, Lingpeng Kong, Wei Bi, Xiang Li, and Ben Kao. 2021. [Good for misconceived reasons: An empirical revisiting on the need for visual context in multimodal machine translation](#). In *ACL*, pages 6153–6166, Online. Association for Computational Linguistics.
- Le Xue, Manli Shu, Anas Awadalla, Jun Wang, An Yan, Senthil Purushwalkam, Honglu Zhou, Viraj Prabhu, Yutong Dai, Michael S Ryoo, et al. 2024. xgen-mm (blip-3): A family of open large multimodal models. *arXiv preprint arXiv:2408.08872*.
- Shaowei Yao and Xiaojun Wan. 2020. Multimodal transformer for multimodal machine translation. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 4346–4350.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023a. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11975–11986.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023b. [Sigmoid loss for language image pre-training](#).
- Yuexiang Zhai, Shengbang Tong, Xiao Li, Mu Cai, Qing Qu, Yong Jae Lee, and Yi Ma. 2023c. Investigating the catastrophic forgetting in multimodal large language models. *arXiv preprint arXiv:2309.10313*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.