

Translation-Enhanced Multilingual Text-to-Image Generation

Yaoyiran Li^{♣,*} Ching-Yun Chang[◇] Stephen Rawls[◇]
Ivan Vulić[♣] Anna Korhonen[♣]

[♣]Language Technology Lab, TAL, University of Cambridge

[◇]Amazon Alexa AI

y1711@cam.ac.uk, {cychang, sterawls}@amazon.com

{iv250, alk23}@cam.ac.uk

Abstract

Research on text-to-image generation (TTI) still predominantly focuses on the English language due to the lack of annotated image-caption data in other languages; in the long run, this might widen inequitable access to TTI technology. In this work, we thus investigate multilingual TTI (termed *mTTI*) and the current potential of neural machine translation (NMT) to bootstrap *mTTI* systems. We provide two key contributions. **1)** Relying on a multilingual multi-modal encoder, we provide a systematic empirical study of standard methods used in cross-lingual NLP when applied to *mTTI*: TRANSLATE TRAIN, TRANSLATE TEST, and ZERO-SHOT TRANSFER. **2)** We propose Ensemble Adapter (ENSAD), a novel parameter-efficient approach that learns to weigh and consolidate the multilingual text knowledge within the *mTTI* framework, mitigating the language gap and thus improving *mTTI* performance. Our evaluations on standard *mTTI* datasets COCO-CN, Multi30K Task2, and LAION-5B demonstrate the potential of translation-enhanced *mTTI* systems and also validate the benefits of the proposed ENSAD which derives consistent gains across all datasets. Further investigations on model variants, ablation studies, and qualitative analyses provide additional insights on the inner workings of the proposed *mTTI* approaches.

1 Introduction and Motivation

Text-to-Image Generation (TTI) is an emerging yet rapidly growing area, owing its recent progress to ever-growing deep generative models, larger-scale multi-modal datasets, and increasing computational resources. The success of recent TTI work is impressive; e.g., it is possible to synthesise not only high-resolution complex scenes (Ramesh et al., 2022; Rombach et al., 2022), but also sur-

realist and ‘aesthetics-aware’ paintings (Gallego, 2022).

However, current models are made and deployed almost exclusively for the English language (EN). This is primarily due to the lack of annotated image-caption data in other languages, which might result in inequitable access to TTI technology in the long run, especially for low-resource languages (Blasi et al., 2022). Hiring human annotators to write high-quality image descriptions is time-consuming and expensive; ‘gold standard’ data, if it exists at all, is thus typically used for evaluation purposes only (Lan et al., 2017; Aggarwal and Kale, 2020).

Even if we put the crucial concerns of data scarcity aside, training state-of-the-art (SotA) TTI models from scratch for each language is technically infeasible and impractical: it would consume massive computational resources, exceeding the capabilities of many research labs (Ramesh et al., 2021; Saharia et al., 2022) and raising concerns of its environmental impact (Schwartz et al., 2020).¹ Therefore, in this work, we focus on multilingual TTI (*mTTI*) through the optics of NLP’s cross-lingual transfer learning methods, leaning on the reasonable assumption of having abundant image-text pairs in English (and/or a pretrained EN TTI model), but only limited gold-standard data for fine-tuning and evaluation in a target language.²

In particular, we investigate the role of cross-lingual transfer and (neural) machine translation (MT) in bootstrapping *mTTI*, and we focus on two crucial research questions. **(RQ1)** Are standard MT-based cross-lingual transfer methods feasible for *mTTI*, and how do they compare with standard

¹For instance, DALL-E (Ramesh et al., 2021) is trained on $1,024 \times 16\text{GB}$ NVIDIA[®] V100 GPUs for a total of 430,000 updates. DALL-E Mega, an attempt to reproduce DALL-E’s results, reports an estimated emission of 18,013.47-kg CO₂-equivalents, training on a TPU v3-256 (128×TPU v3 chips) for 56 days. The estimation is based on a publicly available machine learning emissions calculator (Luccioni et al., 2019).

²A more detailed discussion on data sources, data availability and scarcity is provided in Appendix B.

*This work has been done during the author’s internship at Amazon Alexa AI.

zero-shot cross-lingual transfer methods? (**RQ2**) Is it possible to enhance zero-shot cross-lingual transfer relying on (ensembles of) MT-generated output for improved mTTI?

Our experiments and core findings are based on several mTTI benchmarks. First, we use the standard and publicly available COCO-CN (Li et al., 2019) and Multi30K (Elliott et al., 2016), and we also build a new dataset for Finnish as a lower-resource language from LAION-5B (Schuhmann et al., 2022). Regarding RQ1, we then conduct a systematic empirical study comparing the standard cross-lingual transfer methods: TRANSLATE TRAIN, TRANSLATE TEST, and ZERO-SHOT TRANSFER. Our main results indicate that TRANSLATE TRAIN achieves the best performance, followed by ZERO-SHOT TRANSFER which outperforms TRANSLATE TEST.

Regarding RQ2, we aim to combine MT-based and zero-shot cross-lingual transfer via fast and parameter-efficient fine-tuning. Inspired by the speech processing literature where a list of Automatic Speech Recognition (ASR) hypotheses can be *jointly considered* for downstream tasks (Ganesan et al., 2021; Liu et al., 2021) to alleviate the misrecognition of ASR systems, we propose a module within our mTTI framework termed Ensemble Adapter (ENSAD). It fuses the text encodings of ‘non-English’ text input and a set of its translations to English. Additionally inspired by Ponti et al. (2021), the idea is to combine the knowledge from *multiple translations* to mitigate potential translation errors, and that way boost cross-lingual transfer for mTTI.

Our proposed method derives robust gains across all evaluation datasets. Besides offering SotA mTTI performance, the introduced ENSAD component also adds only 0.1% dedicated extra parameters (relative to the full mTTI model size) per each supported target language. Put simply, the use of ENSAD increases the portability of our mTTI framework through quick and parameter-efficient adaptation to new languages. The resources of our work are available at <https://www.amazon.science/code-and-datasets/translation-enhanced-multilingual-text-to-image-generation>.

2 Related Work

Text-to-Image Generation. There are generally two categories of standard TTI setups: **1)** a super-

vised setup, where gold standard training and test data are from the same domain (e.g., both from MS-COCO); and **2)** a zero-shot setup, where there is a domain difference between the training data (typically large-scale noisy Web-crawled data) and the high-quality test data (typically manually constructed). GAN-based models are common in supervised TTI setups (Reed et al., 2016; Xu et al., 2018; Zhu et al., 2019): they still hold the SotA results, offering smaller model sizes and faster image generation speed (Zhang et al., 2021; Tao et al., 2022; Zhou et al., 2022). GigaGAN (Kang et al., 2023), a recent attempt to scale up GAN models, achieves fairly strong and competitive zero-shot TTI performance. However, in the zero-shot setup, large Vector Quantised Variational Autoencoder (VQVAE)-based models (Ramesh et al., 2021; Crowson et al., 2022; Gafni et al., 2022) and large diffusion models (Nichol et al., 2022; Ramesh et al., 2022; Saharia et al., 2022) play the leading role and offer the best performance.

Multilingual and Non-EN TTI. Research on mTTI and non-EN TTI is currently limited and only in its infancy. Cogview is a large VQVAE-based Chinese TTI model with training data partly from crawling Chinese websites and social media platforms, and partly from translating EN data (Ding et al., 2021). ruDALL-E is a VQVAE-based Russian TTI model recreating DALL-E (Ramesh et al., 2021) with training data translated from EN data.³

To the best of our knowledge, there are only two existing papers attempting multilingual or cross-lingual TTI. Zhang et al. (2022) align two monolingual text encoders, one for the source and the other for the target language, with a fixed image generator pretrained on the source language (i.e., EN). Jung et al. (2022) take a step further, relying on a multilingual text encoder that supports more languages simultaneously.

We note several crucial differences to the prior work. **1)** The two papers are based on earlier TTI models (Xu et al., 2018), which are now largely surpassed by recent SotA models (Zhou et al., 2022). **2)** Their model designs are tied to the model of Xu et al. (2018) and cannot be easily adapted to the latest SotA TTI models. **3)** They use traditional LSTM text encoders enhanced by mono-modal BERT features, while SotA TTI models (Zhou et al., 2022; Saharia et al., 2022; Rombach et al.,

³<https://rudalle.ru/>; ruDALL-E has not released an accompanying paper yet, but a [technical blog](#) is available.

2022) use the multi-modal CLIP model (Radford et al., 2021). Therefore, we neither adopt them as baselines nor try to adapt them for our use, also taking into account the difficulty of replicating the prior work as no code has been released to date. In contrast, our work relies on the mCLIP text encoder (Carlsson et al., 2022), the multilingual version of CLIP, and is developed based on LAFITE (Zhou et al., 2022), a SotA TTI model. In fact, as shown later in our work, training an English TTI model using mCLIP without any further tuning can *already realise* zero-shot mTTI, similar to what has been attempted by Jung et al. (2022).

Translation-Based Cross-lingual Transfer. Machine translation (MT) at both lexical level and sentence level has been successfully used for cross-lingual transfer learning in NLP, where TRANSLATE TRAIN and TRANSLATE TEST usually serve as strong baselines for downstream tasks (Conneau et al., 2018; Glavaš et al., 2019; Hu et al., 2020; Ponti et al., 2021; Li et al., 2022a,b). In addition, MT is used to generate sentence pairs for training multilingual multi-modal models (Zhou et al., 2021; Carlsson et al., 2022). However, MT is still largely underexplored and underutilised for mTTI. In this work, we analyse the potential of MT to enhance multilingual and cross-lingual TTI.

3 Methodology

In what follows in this section, we first introduce our base mLAFITE model and three baseline approaches for mTTI (§3.1). Next, we propose an Ensemble Adapter module that can work in synergy with the pretrained mLAFITE model to improve mTTI performance (§3.2). Finally, we describe how we train our Ensemble Adapter and formulate our loss functions (§3.3).

3.1 mLAFITE and Baselines

For easier deployment and comparison of different cross-lingual transfer methods, our work focuses on the relatively lightweight GAN-based models, which are faster to train and evaluate compared with VQVAE-based models and large diffusion models (see §2). In particular, we adopt LAFITE (Zhou et al., 2022), a SotA GAN-based English TTI model, as our starting point. To unlock its multilingual capabilities, we replace its English-only CLIP text encoder (Radford et al., 2021) with mCLIP (Carlsson et al., 2022), which is already pretrained to align the sentence representation spaces

of 68 languages.⁴

There are three common categories of cross-lingual transfer approaches which we apply to mTTI and adopt as our principal baselines:

TRANSLATE TRAIN. We translate all the captions from the English training set (e.g., COCO) into a (non-EN) target language (\mathbb{L}) relying on an MT system. We then train a LAFITE TTI model in the target language from scratch, relying on mCLIP as the text encoder.⁵ At inference, an \mathbb{L} sentence is directly fed into the target-language TTI model.

The other two approaches instead rely on a TTI model pretrained with *English* data, and they do *not* require further tuning with captions in the target languages. As our first step, we pretrain an mCLIP-based LAFITE model (we call it **mLAFITE** for brevity) from scratch.

TRANSLATE TEST. At inference, we first translate a caption in \mathbb{L} into EN via MT and the EN translation then serves as mLAFITE’s input.

ZERO-SHOT TRANSFER. Since mCLIP is a multilingual sentence encoder, text in \mathbb{L} can be directly fed to our mLAFITE for TTI without any extra fine-tuning.

3.2 mLAFITE with Ensemble Adapter

We now propose an attention-based Ensemble Adapter (ENSAD) module that aims to improve mTTI via leveraging knowledge from multiple translations of the same input. The full pipeline and how ENSAD extends the base mLAFITE model are illustrated in Figure 1. Given an input sentence in language \mathbb{L} , $\mathbb{L} \neq \text{EN}$, we first use any (N)MT system to sample a set of EN translations. We then deploy the ENSAD module between the mCLIP text encoder and the TTI generator to fuse the mCLIP-extracted embeddings, bridging the EN- \mathbb{L} language domain gap. The adapter can be trained with only a small set of image- \mathbb{L} text pairs while mCLIP and the TTI generator networks are kept frozen.

⁴mCLIP is derived by fine-tuning a pretrained XLM-R model (Carlsson et al., 2022; Conneau et al., 2020), and it does not directly depend on parallel corpora or multilingual image-text data. The work uses NMT to generate ‘silver’-quality EN- \ast sentence pairs and then directly aligns the CLIP-extracted EN representations and mCLIP’s multilingual sentence representations of the NMT-generated data. Both CLIP and mCLIP use a shared CLIP visual encoder.

⁵We use mCLIP rather than monolingual CLIP since it is infeasible for most languages. Only several high-resource languages have publicly available monolingual models. For fair cross-language comparisons, we leverage the same mCLIP text encoder in all our experiments.

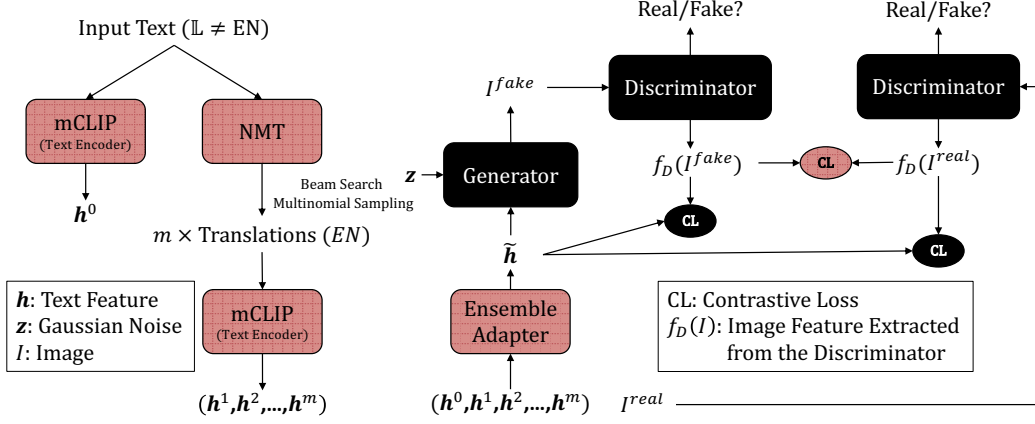


Figure 1: An overview of the full proposed mTTI framework with the Ensemble Adapter module. The black blocks are networks and contrastive learning (CL) losses already in the original LAFITE model (also in our pretrained mLAFITE). Our proposed, newly added modules, and a CL loss are provided in red, gridded blocks.

Formally, we use x^0 to denote the \mathbb{L} input text, while $\{x^1, x^2, \dots, x^m\}$ is a set of m EN translations of the \mathbb{L} input text. The fixed mCLIP encoder extracts their respective (l_2 -normalised) d -dimensional sentence embeddings, yielding the matrix $\mathbf{H} = (\mathbf{h}^0, \mathbf{h}^1, \dots, \mathbf{h}^m) \in \mathbb{R}^{d \times (m+1)}$. Then, our proposed ENSAD learns to fuse these sentence encodings from \mathbf{H} . We define the query (\mathbf{q}), key (\mathbf{K}), and value (\mathbf{V}) inputs of our attention as:

$$\mathbf{q} = \mathbf{h}^0, \quad (1)$$

$$\mathbf{K} = (\mathbf{h}^1, \mathbf{h}^2, \dots, \mathbf{h}^m), \quad (2)$$

$$\mathbf{V} = (\mathbf{h}^1 - \mathbf{h}^0, \mathbf{h}^2 - \mathbf{h}^0, \dots, \mathbf{h}^m - \mathbf{h}^0). \quad (3)$$

Note that $\{\mathbf{h}^0, \mathbf{h}^1, \dots, \mathbf{h}^m\}$ are all close to each other in the mCLIP representation space. Therefore, to focus on the ‘additional information’ contained in the EN translations, we take the difference between $\mathbf{h}^i, i > 0$ and \mathbf{h}^0 as in Eq. (3).⁶ The calculation of attention scores is then based on the standard additive attention (Bahdanau et al., 2015):

$$\mathbf{A} = \mathbf{W}^q \mathbf{q} \mathbf{1}^T + \mathbf{W}^k \mathbf{K} + \mathbf{W}^v \mathbf{V} + \mathbf{b} \mathbf{1}^T, \quad (4)$$

$$\mathbf{s}^T = \text{softmax}(\mathbf{W}^p (\tanh(\mathbf{A}))). \quad (5)$$

ENSAD’s hidden size is d_{hid} ; $\mathbf{W}^q, \mathbf{W}^k, \mathbf{W}^v \in \mathbb{R}^{d_{hid} \times d}$ are respective mappings for query, key, and value inputs; $\mathbf{b} \in \mathbb{R}^{d_{hid}}$ is the bias, and $\mathbf{W}^p \in \mathbb{R}^{1 \times d_{hid}}$ is a final projection matrix for deriving the attention scores. Then, the context vector is an attention-guided summarisation of \mathbf{V} . ENSAD’s

⁶We adopt the simple mean pooling of $\{\mathbf{h}^0, \mathbf{h}^1, \dots, \mathbf{h}^m\}$ as an additional baseline with results in §6.2. We also tried multi-head self-attention (Vaswani et al., 2017), where $\mathbf{Q} = \mathbf{K} = \mathbf{V} = \mathbf{H}$, which, however, showed inferior performance in our preliminary experiments.

final output is the linear combination of \mathbf{h}_0 and the context vector, computed as follows:

$$\mathbf{V}^o = (1 - \alpha) \mathbf{V} + \alpha \cdot \tanh(\mathbf{W}^o \mathbf{V}), \quad (6)$$

$$\mathbf{c} = \mathbf{V}^o \mathbf{s}, \quad (7)$$

$$\tilde{\mathbf{h}} = \text{ENSAD}(\mathbf{H}) = (1 - \alpha) \mathbf{q} + \alpha \cdot \mathbf{c}, \quad (8)$$

where $\mathbf{W}^o \in \mathbb{R}^{d \times d}$ is the output mapping, and α is an interpolation hyperparameter. We also l_2 -normalise the outputs of Eqs. (3), (7), (8), as well as the $\tanh(\mathbf{W}^o \mathbf{V})$ term in Eq. (6).

3.3 Contrastive Adversarial Training

Our Generator (G) and Discriminator (D) network structures and the pretraining process of the base mLAFITE model all follow LAFITE’s original implementation for supervised TTI. As illustrated in Figure 1, we take the pretrained mLAFITE and insert the ENSAD between mCLIP and G . We then adversarially train ENSAD and D iteratively while mCLIP and G are kept frozen.⁷ Additionally, we propose to optimise a novel contrastive objective aligning the D -extracted real image and fake (synthesised) image features in adversarial training.

The (m)LAFITE GAN framework is adapted from the popular unconditional StyleGAN2 framework (Karras et al., 2020b) which features a re-designed adaptive instance normalization mechanism (Huang and Belongie, 2017) in G : it enables the unconditional channel-wise ‘style information’ (e.g., pose, lighting, background style) to control G ’s image synthesis backbone (convolution and up-sampling layers). The ‘style information’ is derived

⁷We also tried freezing D but this results in inferior performance in our preliminary investigation.

as follows: a random noise \mathbf{z} is sampled from the standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and transformed into a so-called unconditional *StyleSpace*, which is proven to be a well-disentangled intermediate latent space (Wu et al., 2021).⁸ LAFITE further proposes to inject text-conditioning information into the *StyleSpace* via a series of non-linear and affine mappings. In our pipeline, G takes our ENSAD-gathered feature $\tilde{\mathbf{h}}$ and noise \mathbf{z} , and it then outputs a fake image: $\mathcal{I}^{fake} = G(\tilde{\mathbf{h}}, \mathbf{z})$.

The discriminator has a characteristic ‘two-branch’ design: **1)** D is in essence a convolutional *image encoder*, producing $f_D(\mathcal{I})$, a d -dim image feature for any real or fake (i.e., synthesised) input image \mathcal{I} ; **2)** D also predicts if \mathcal{I} is real or fake based on both \mathcal{I} and $\tilde{\mathbf{h}}$, where the prediction (a scalar output) is denoted as $D(\mathcal{I}, \tilde{\mathbf{h}}) = D_s(\mathcal{I}) + \tilde{\mathbf{h}}^T f_D(\mathcal{I})$. This is realised via adding two affine transformations on top of a shared visual backbone for deriving $f_D(\mathcal{I})$ and $D_s(\mathcal{I})$, respectively. We then define the adversarial (AD) losses for ENSAD and D following LAFITE:

$$\mathcal{L}_{AD}^{\text{ENSAD}} = -\frac{1}{n} \sum_{i=1}^n \log \sigma(D(\mathcal{I}_i^{fake}, \tilde{\mathbf{h}}_i)), \quad (9)$$

$$\begin{aligned} \mathcal{L}_{AD}^D &= -\frac{1}{n} \sum_{i=1}^n \log \sigma(D(\mathcal{I}_i^{real}, \tilde{\mathbf{h}}_i)) \\ &\quad - \frac{1}{n} \sum_{i=1}^n \log(1 - \sigma(D(\mathcal{I}_i^{fake}, \tilde{\mathbf{h}}_i))). \end{aligned} \quad (10)$$

n is the batch size, and $\sigma(\cdot)$ is the sigmoid function. We propose an auxiliary contrastive loss, aligning the discriminator-extracted \mathcal{I}^{fake} and \mathcal{I}^{real} features, computed as follows:

$$s_{i,j} = \cos(f_D(\mathcal{I}_i^{real}), f_D(\mathcal{I}_j^{fake})), \quad (11)$$

$$\mathcal{L}_{CL} = -\frac{1}{n} \sum_{i=1}^n \log \frac{\exp(s_{i,i}/\tau)}{\sum_{j=1}^n \exp(s_{j,i}/\tau)}. \quad (12)$$

$\cos(\cdot)$ calculates the cosine similarity, and τ is the temperature.

In the original LAFITE paper, there are already two auxiliary contrastive losses: **1)** \mathcal{L}_{CL}^G aligns CLIP-extracted image features of \mathcal{I}^{fake} and the input text embedding, i.e., $\tilde{\mathbf{h}}$ in our case; **2)** \mathcal{L}_{CL}^D aligns $f_D(\mathcal{I})$ with its associated $\tilde{\mathbf{h}}$.⁹ In our preliminary experiments, we found that \mathcal{L}_{CL}^G was not

⁸The transformation includes a shared 8-layer MLP and a dedicated affine mapping per each generation layer. We refer the reader to the original work for further technical details.

⁹As with LAFITE’s original implementation, $f_D(\mathcal{I})$ is $f_D(\mathcal{I}^{fake})$ in \mathcal{L}_{ENSAD} and $f_D(\mathcal{I}^{real})$ in \mathcal{L}_D .

useful for ENSAD, so we completely remove it.¹⁰ Our final losses for training ENSAD and D are as follows, with two hyperparameters λ_1 and λ_2 controlling the weights of contrastive losses:

$$\mathcal{L}_{\text{ENSAD}} = \mathcal{L}_{AD}^{\text{ENSAD}} + \lambda_1 \cdot \mathcal{L}_{CL} + \lambda_2 \cdot \mathcal{L}_{CL}^D, \quad (13)$$

$$\mathcal{L}_D = \mathcal{L}_{AD}^D + \lambda_1 \cdot \mathcal{L}_{CL} + \lambda_2 \cdot \mathcal{L}_{CL}^D. \quad (14)$$

The full training process is also summarised in Algorithm 1, available in Appendix C. Note that the use of ENSAD introduces only up to 0.1% extra parameters per each target language relative to the full model size. This parameter efficiency boosts the portability of our mTTI framework, enabling quick and efficient adaptation to new languages.

4 Datasets

mLAFITE pretraining is based on the MS-COCO (Chen et al., 2015) training set comprising 82,783 images, where each image is associated with 5 EN captions. 10% of the training set is held out as our dev set, and the rest is used for training. MS-COCO also provides a validation set (40,504 images), frequently used for TTI evaluation.

For mTTI, we choose evaluation datasets that satisfy the following criteria: **a)** no overlap between images in the test set and images used in pretraining; **b)** the test set includes at least 5K images;¹¹ **c)** the captions are human-written descriptions and not (manual or MT-derived) translations from EN captions.¹² Based on these requirements, we select three ‘non-EN’ datasets, outlined in what follows.

COCO-CN (Li et al., 2019) provides Chinese (ZH) captions (i.e., human descriptions) for 20,341 MS-COCO images. 6,748 of them are from the COCO validation set not seen during mLAFITE pretraining; we thus use them as our test set. We randomly sample 20% of the rest as our dev set (2,718), and the training set has 10,875 images. Each image has

¹⁰The equations for the other two CL losses are similar to Eq. (12). For brevity, we skip the details and refer the reader to the original LAFITE paper.

¹¹Previous work proved that small test set sizes result in biases and unreliable TTI evaluation (Chong and Forsyth, 2020); therefore, TTI work typically adopts test sets with more than 5K images (Zhou et al., 2022; Ramesh et al., 2021). For instance, the most common EN TTI data for evaluation is the MS-COCO validation set that contains 40K images. The smallest general-domain test set in Zhou et al. (2022) is LN-COCO (Pont-Tuset et al., 2020) containing $\sim 5K$ images.

¹²Human-written descriptions are more realistic for real-world non-EN users, and translations from EN captions can cause unexpected ‘translationese’ bias (Elliott et al., 2016; van Miltenburg et al., 2017; Bugliarello et al., 2022).

only one ZH caption. COCO-CN additionally offers 5,000 ZH sentences manually translated from EN captions; we only use the corresponding EN-ZH sentence pairs to calculate BLEU scores for comparing different MT systems.

Multi30K Task2 (Elliott et al., 2016, 2017) has 5 German (DE) captions (human descriptions) for each of 31,014 Flickr30K (Young et al., 2014) images. We randomly sample and keep one caption per each image.¹³ We randomly split the data into train, dev, and test sets spanning 10,000, 2,000, and 19,014 images, respectively.

LAION-5B (Schuhmann et al., 2022) is a large-scale Web-crawled vision-language dataset with 5 billion image-text pairs covering 100+ languages. We focus on Finnish (FI) as a lower-resource language for our evaluation. Unlike carefully annotated COCO-CN and Multi30K, LAION-5B’s data are noisy, so we rely on massive filtering to select relatively high-quality data. The full data creation process for FI is provided in Appendix D.

The final dataset comprises training, development and test portions with 10,000, 2,000, and 18,000 image-text pairs, respectively. Our manual inspection of the final dataset indicates that it is of acceptable quality although having its own characteristics (Appendix D) but the quality in general still cannot match COCO-CN or Multi30K. We use the data in our main experiments 1) as an initial trial to extend TTI evaluation to ‘non-COCO-style’ captions and another language and 2) for comparative analyses with COCO-CN and Multi30K.

Supplementary Dataset: IGLUE. In order to further widen the set of target languages, we also experiment with *IGLUE xFlickr&CO* (Bugliarello et al., 2022). It provides 2K images, where one half comes from the MS-COCO validation set and the other half from Multi30K with associated human descriptions in 5 additional languages: Spanish (ES), Indonesian (ID), Japanese (JA), Russian (RU), and Turkish (TR). Since IGLUE does not offer a training set, we use it only for RQ1-related experiments. Although IGLUE does not comply with our criterion b) above, we use it to extend our empirical analyses to more languages.

Table 6 in Appendix A provides a full and sys-

¹³All our non-EN TTI datasets uniformly have one caption for each image. This setup is also more realistic since real-world users only need to input a single sentence into a TTI model. For mLAFITE pretraining, however, as with all other related work, all the COCO captions are used.

tematic overview of languages and data statistics used in this work.

5 Experimental Setup

In what follows, we outline our experimental setups and choices related to the two core RQs from §1. We also show details concerning our mLAFITE pretraining, side experiments (most are RQ2-related), and evaluation metric.

mLAFITE Pretraining. All methods for mTTI are implemented based on our pretrained mLAFITE model, which is trained with 8×16 GB V100 GPUs for 75 hours (i.e., 40 million data points sampled from the training set). Contrastive loss weights and other hyper-parameters follow the original LAFITE setup (Zhou et al., 2022).¹⁴ For fair comparisons, we use the same mCLIP text encoder for all our RQ1 and RQ2 experiments.

RQ1 Experiments. On COCO-CN, we compare four widely used MT systems: Amazon Translate¹⁵, a SotA commercial MT software, and three SotA Transformer-based NMT models developed in an academic context including Marian (Tiedemann and Thottingal, 2020; Junczys-Dowmunt et al., 2018), mBART50 (Liu et al., 2020; Tang et al., 2021), and M2M100 (Fan et al., 2021). We leverage them to generate the 1-best translations for TRANSLATE TRAIN and TRANSLATE TEST, and we also compare the BLEU scores of the MT systems against the TTI performance. Note that training a TRANSLATE TRAIN TTI model from scratch for each of the MT systems also takes 75 hours; our TRANSLATE TRAIN experiments thus do not extend to other datasets beyond COCO-CN due to the high computational cost.

Given the considerations above along with preliminary evaluations on COCO-CN which showed that Marian outperforms mBART50 and M2M100, for the other datasets we focus on comparing the Marian-based TRANSLATE TEST with ZERO-SHOT TRANSFER.

RQ2 Experiments. RQ2 further studies the effectiveness of the proposed ENSAD module; see §3 and Figure 1. We select Marian as the NMT back-

¹⁴The original LAFITE model is based on EN CLIP and is trained with 25 million samples on 4×16 GB V100 GPUs. Our mLAFITE model uses mCLIP, which is of the same dimensionality as CLIP and keeps the values of all the other relevant hyper-parameters, such as batch per GPU and learning rate, from the original work.

¹⁵<https://aws.amazon.com/translate/>

bone¹⁶ and sample m EN translations per each input sentence in the input language \mathbb{L} .¹⁷ To compare with ENSAD (with the frozen mLAFITE generator), we also propose and experiment with several insightful and simple baselines (without the use of ENSAD) in addition to the RQ1 baselines: **1)** we try standard mean-pooling as a simple ensembling baseline directly on mLAFITE; **2)** we fine-tune G using the original non-EN captions;¹⁸ **3)** we fine-tune G using mean-pooled text features. Finally, we also investigate variants which combine ENSAD with the tunable generator G to check if further gains can be achieved.¹⁹

Training for RQ2 experiments is conducted on $8 \times V100$ GPUs with a batch size per GPU of 16 for about 7 hours (i.e., a total of 2 million data points sampled from the respective training sets). We use Adam optimiser (Kingma and Ba, 2014) with a learning rate of $5e-4$ and betas of $(0, 0.99)$. For the generator-tuning baselines, their contrastive loss setups completely follow the original LAFITE (Zhou et al., 2022). In our ENSAD experiments, $\lambda_1=4$ and $\lambda_2=2$. Other hyper-parameters are as follows: the NMT beam size is 12, NMT temperature is 2.0, images are scaled to resolution 256×256 , $m=12$, $d=512$, $d_{hid}=256$, and $\tau=0.5$. In addition, we fuse 10% and 1% standard Gaussian noise into \mathbf{h}^0 and \mathbf{h}^i ($1 \leq i \leq m$) respectively as a data augmentation ‘trick’. The hyper-parameters are tuned on our dev split of COCO-CN with details in Appendix G. The same set of hyper-parameters is also adopted for the other two datasets.

Side Experiments. Besides the main RQ1 and RQ2 experiments, we also conduct a series of side analyses focused on ENSAD. They span **1)** the impact of the number of EN translations m , **2)** the impact of the interpolation hyperparameter α , and **3)** robustness tests. We also conduct **4)** ablation studies to validate the effectiveness of different components, and **5)** present generated images and ENSAD attention scores.

¹⁶Amazon Translate’s API can only return a single 1-best translation; it thus cannot be used for ENSAD experiments.

¹⁷We adopt the common ‘beam search multinomial sampling’ for sampling EN translations.

¹⁸ G is tuned adversarially following the original training setup of (m)LAFITE.

¹⁹We first fine-tune G and then train ENSAD with the fine-tuned G but still use the discriminator of our pretrained mLAFITE to alleviate its overfitting (Karras et al., 2020a). We also tried **1)** training ENSAD first and then fine-tuning G and **2)** training ENSAD together with G , but they both derive suboptimal results in our preliminary investigation.

Evaluation Metric. Following Zhou et al. (2022) and Ramesh et al. (2021), we report the Fréchet Inception Distance (FID) (Heusel et al., 2017) computed with 30,000 synthesised images generated using randomly sampled test set texts against test set ground-truth images, which is the most authoritative machine evaluation metric for TTI so far.²⁰

6 Results and Discussion

The main results are structured around the two central RQs from §1, discussed in §6.1 and §6.2.

6.1 RQ1: Results and Analyses

Comparison of Three Baselines. The results of TRANSLATE TRAIN, TRANSLATE TEST, and ZERO-SHOT TRANSFER on COCO-CN are summarised in Table 1. While all three methods use mCLIP, TRANSLATE TEST and ZERO-SHOT TRANSFER are based on a pretrained EN mLAFITE and do not require any further tuning. TRANSLATE TRAIN achieves the best FID scores; however, it requires training from scratch with translated \mathbb{L} captions (see §3.1 and §5). Since MS-COCO provides ground-truth human-written EN captions for COCO-CN images, and Multi30K Task2 also provides EN human descriptions, we directly feed the EN captions to mLAFITE and report the FID scores as an upper bound (see the first row of each of Tables 1 and 2).²¹

The scores in Tables 1 and 2 show that ZERO-SHOT TRANSFER outperforms TRANSLATE TEST, demonstrating the strong capability of the multilingual mCLIP text encoder. TRANSLATE TEST compares unfavourably to other methods, revealing the gap between EN translations and the ground-truth EN human descriptions (e.g., translation errors, ‘translationese’ bias). We further extend the comparison to five more languages from the IGLUE dataset, and the results from Table 7 in Appendix E corroborate the finding that ZERO-SHOT TRANSFER generally outperforms TRANSLATE TEST.

Comparison of MT Systems. We compare the performance of the four MT systems on COCO-CN and also report their BLEU scores on the additional

²⁰Inception Score (IS) (Salimans et al., 2016) is another common evaluation metric for TTI, which is, to some extent, superseded by FID (Heusel et al., 2017; Borji, 2022). Moreover, IS gives misleading results when applied to datasets other than ImageNet (Barratt and Sharma, 2018), and is especially non-fitting for unannotated LAION-5B images.

²¹For fair comparisons, we keep only one ground-truth EN caption for each image.

Method	MT Model	BLEU \uparrow	FID \downarrow
Ground-Truth EN Captions	-	-	14.35
TRANSLATE TRAIN (EN \rightarrow ZH)	mBART50	32.77	14.98
	Marian	32.5	14.64
	M2M100	33.73	15.28
	Amazon Translate	42.23	14.87
TRANSLATE TEST (ZH \rightarrow EN)	mBART50	26.32	16.38
	Marian	25.11	15.9
	M2M100	22.65	17.26
	Amazon Translate	30.95	15.64
ZERO-SHOT TRANSFER	-	-	15.57

Table 1: Results on COCO-CN (ZH). ‘-’: the method does not rely on MT. **FID** \downarrow : lower is better.

Method	ZH: FID \downarrow	DE: FID \downarrow	FI: FID \downarrow
Ground-Truth EN Captions	14.35	16.68	-
TRANSLATE TEST (Marian)	15.9	17.31	27.23
TRANSLATE TEST (Amazon Translate)	15.64	17.03	26.67
ZERO-SHOT TRANSFER	15.57	16.98	25.78

Table 2: TRANSLATE TEST vs. ZERO-SHOT TRANSFER on COCO-CN (ZH), Multi30K Task2 (DE), and LAION-5B (FI). ‘-’: LAION-5B (FI) data do not provide ground-truth EN captions. **FID** \downarrow : lower is better.

5K sentence pairs. Table 1, as expected, reveals that the commercial Amazon Translate system offers much stronger MT performance than the three academic NMT systems in terms of BLEU. Concerning mTTI, Amazon Translate is the best system with the TRANSLATE TEST approach category and ranks second with TRANSLATE TRAIN. Interestingly, there are some salient discrepancies between BLEU-based versus TTI-based system rankings. For example, Marian ranks second in TRANSLATE TEST and is the best system with TRANSLATE TRAIN, although its MT performance underperforms both Amazon Translate and mBART50. We speculate that this might be due to the pretraining specifics of mCLIP, where *Marian-generated pseudo-parallel sentence pairs* were used (Carlsson et al., 2022).

In TRANSLATE TEST, M2M100 obtains the lowest ZH \rightarrow EN BLEU score and also achieves the worst TTI performance. However, mBART50 and M2M100 have close EN \rightarrow ZH BLEU scores in TRANSLATE TRAIN, and a small edge in BLEU cannot guarantee a better TTI performance. We additionally compare Marian and Amazon Translate for TRANSLATE TEST in Tables 2 and 7 (Appendix E) on other languages and datasets, which further validate the core findings.

6.2 RQ2: Results and Analyses

Effectiveness of ENSAD. The main results are summarised in Table 3. For all methods except

Method	ZH: FID \downarrow	DE: FID \downarrow	FI: FID \downarrow
Ground-Truth EN Captions	14.35	16.68	-
ZERO-SHOT TRANSFER	15.57	16.98	25.78
Mean Pooling	16.47	17.7	27.67
Fine-Tune G (L Text)	15.23	16.28	17.69
Fine-Tune G (Mean Pooling)	15.27	16.68	18.17
ENSAD (Frozen G)	14.52 $\downarrow^{6.7\%$	16.26	21.9
ENSAD + Fine-Tune G (L Text)	15.14	16.12 $\downarrow^{5.1\%$	17.38 $\downarrow^{35.6\%$
ENSAD + Fine-Tune G (Mean Pooling)	14.93	16.23	17.41

Table 3: Main results for RQ2. The models in the first three rows do not require any additional fine-tuning. For the best-performing models in **bold** numbers, we also present the relative improvement in percentage when comparing with ZERO-SHOT TRANSFER. **FID** \downarrow : lower is better.

‘Ground-Truth EN Captions’, the language gap (with EN captions for mLAFITE pretraining) always exists since the text input is in language \mathbb{L} . When there is no image domain gap (i.e., for COCO-CN), ENSAD without tuning G achieves the best score, surpassing also the TRANSLATE TRAIN baseline (cf. Table 1), and the absolute score also mitigates the gap to the upper-bound set by ‘Ground-Truth EN Captions’. With image domain gap present (i.e., DE and FI), training ENSAD (with frozen G) still shows a small edge over fine-tuning G (without ENSAD) for DE; however, for the noisier LAION-5B data, fine-tuning G is more useful. However, for both DE and FI, the best results are always achieved when ENSAD is leveraged, validating its usefulness combined with parameter efficiency. For example, ENSAD with G frozen consistently outperforms ZERO-SHOT TRANSFER while introducing only 0.1% extra parameters. Our robustness tests repeating ENSAD (Frozen G) experiments on COCO-CN with different random seeds further corroborate these findings (the deviation of FID is 0.04), with a short summary in Appendix F.

Variants of ENSAD. We further investigate the impact of crucial design choices and hyper-parameters in ENSAD such as m , α , and V (see Eq. (3)) respectively on the final TTI performance. The results of different variants are provided in Table 4. They indicate that increasing the number of translations m seems to be conducive to downstream TTI performance. In addition, when $V = K$, the FID score worsens, demonstrating the usefulness of the V variant as formulated by Eq. (3). Finally, the TTI performance deteriorates when $\alpha > 0.2$, showing that \mathbf{h}^0 should still be the main component of $\tilde{\mathbf{h}}$, and ENSAD provides auxiliary information (i.e., a translation-based enhancement).

Model (Variant)	FID ↓	Model (Variant)	FID ↓
Default	14.52	Variant 4: $V = K$	14.73
Variant 1: $m = 1$	14.9	Variant 5: $\alpha = 0.1$	15.07
Variant 2: $m = 4$	14.65	Default: $\alpha = 0.2$	14.52
Variant 3: $m = 8$	14.68	Variant 6: $\alpha = 0.3$	14.81
Default: $m = 12$	14.52	Variant 7: $\alpha = 0.5$	17.7

Table 4: Model variants of ENSAD (Frozen G). FID scores on COCO-CN.

Model (Variant)	FID ↓
Default: with \mathcal{L}_{CL} and \mathcal{L}_{CL}^D	14.52
Remove \mathcal{L}_{CL}	14.74
Remove \mathcal{L}_{CL}^D	14.56
Remove both \mathcal{L}_{CL} and \mathcal{L}_{CL}^D	14.82
Setup of (m)LAFITE: with \mathcal{L}_{CL}^G and \mathcal{L}_{CL}^D	15.03

Table 5: Ablation study on CL losses. Model variant: ENSAD (Frozen G). FID scores on COCO-CN.

Ablation Study. We now study the usefulness of two used contrastive losses: **1)** our proposed \mathcal{L}_{CL} and **2)** \mathcal{L}_{CL}^D inherited from LAFITE. The results in Table 5 show that removing \mathcal{L}_{CL} causes a noticeable performance drop (increased FID). However, removing \mathcal{L}_{CL}^D has only a minor impact on the FID score. When removing both CL losses, the adversarial losses alone produce an FID score of 14.82. We also additionally try the CL loss setup of the original LAFITE and find that the setup is detrimental to the training of ENSAD, producing a worse FID score than using the adversarial losses alone.

TTI Examples and Attention Scores. Finally, we refer the reader to Appendix H where we present images synthesised with TRANSLATE TEST, ZERO-SHOT TRANSFER, and our ENSAD models and where we also show the ENSAD attention scores. The differences between images are subtle and we were unable to find a clear pattern that links high attention scores with particular translations.

7 Conclusion

This work is one of the first investigations of multilingual and cross-lingual text-to-image generation (TTI), with a particular focus on investigating the use of machine translation (MT) for the task. We systematically compared standard cross-lingual transfer approaches TRANSLATE TRAIN, TRANSLATE TEST and ZERO-SHOT TRANSFER in the context of TTI and also studied the differences over MT systems. We then proposed a novel Ensemble Adapter (ENSAD) method that leverages multiple translations to further improve the TTI

performance, with strong and consistent gains reported across a series of standard TTI benchmarks in different languages.

Limitations

First, we again emphasise that the lack of high-quality non-English image-caption pairs is a primary obstacle to wider-scale multilingual and cross-lingual TTI investigations. We hope that researchers in the future can construct and release more high-quality vision-language data for different languages, especially for low-resource ones.

Second, our work uses 512-dim ‘XLM-R Large Vit-B/32’ mCLIP²² and is based on the StyleGAN2 framework (Karras et al., 2020b). Since the main focus of our work is to realise multilingual and cross-lingual TTI and enable fair comparisons across different models and approaches, we compare all proposed and baseline methods with the same mCLIP text encoder and the GAN framework. However, for readers and potential users interested in ‘chasing’ stronger absolute FID scores, we speculate that the larger 640-dim ‘XLM-R Large Vit-B/16+’ mCLIP text encoder and the more recent StyleGAN3 (Karras et al., 2021) can be helpful.

Third, we notice that in addition to LAFITE, several state-of-the-art large diffusion models such as those from Saharia et al. (2022) and Rombach et al. (2022) also use CLIP to condition image generation on text input. This means that we could be able to derive multilingual diffusion models for mTTI also by replacing CLIP with mCLIP and enhance the mTTI performance with our proposed ENSAD (of course, we would need to redesign our loss functions). However, due to limited computational resources, we leave it to future work.

Fourth, the ENSAD boosts cross-lingual transfer for TTI by combining the knowledge from multiple translations, which can mitigate potential translation errors. Our work does not demonstrate if ENSAD is applicable and adaptable to downstream cross-lingual tasks besides TTI. It is because **1)** downstream tasks other than TTI are out of the scope of this work and **2)** adapting ENSAD to different tasks will require redesign of model structures and losses catering to the characteristics of each downstream task, making us believe it is not proper to expand the topic and include everything in a single piece of work. Therefore, we also leave this to future work.

²²<https://github.com/FreddeFrallan/Multilingual-CLIP>

Ethics Statement

The datasets involved in our experiments are publicly available and widely used, and it is quite common to train text-to-image generation models on publicly available data. To the best of our knowledge, the ethical risk is minimal. For privacy concerns, we do not present images with human faces and captions with real human names in the paper, and we will not release material that may contain any sensitive information.

Acknowledgements

We would like to thank **1)** all members of the Amazon Alexa Translations Science Team for helpful discussions and valuable comments during the weekly group meetings, **2)** Yufan Zhou, the author of LAFITE, who kindly responded to our questions concerning LAFITE’s technical details on Github, and **3)** the anonymous reviewers for their feedback.

Ivan Vulić is supported by a personal Royal Society University Research Fellowship ‘*Inclusive and Sustainable Language Technology for a Truly Multilingual World*’ (no 221137; 2022–).

References

- Pranav Aggarwal and Ajinkya Kale. 2020. [Towards zero-shot cross-lingual image retrieval](#). *ArXiv*, abs/2012.05107.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *International Conference on Learning Representations*.
- Shane T. Barratt and Rishi Sharma. 2018. [A note on the inception score](#). *ArXiv*, abs/1801.01973.
- Damian Blasi, Antonios Anastasopoulos, and Graham Neubig. 2022. [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)*, pages 5486–5505, Dublin, Ireland. Association for Computational Linguistics.
- Ali Borji. 2022. [Pros and cons of gan evaluation measures: New developments](#). *Computer Vision and Image Understanding*, 215:103329.
- Emanuele Bugliarello, Fangyu Liu, Jonas Pfeiffer, Siva Reddy, Desmond Elliott, Edoardo Maria Ponti, and Ivan Vulić. 2022. [IGLUE: A benchmark for transfer learning across modalities, tasks, and languages](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 2370–2392. PMLR.
- Fredrik Carlsson, Philipp Eisen, Faton Rekathati, and Magnus Sahlgren. 2022. [Cross-lingual and multilingual CLIP](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6848–6854, Marseille, France. European Language Resources Association.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C. Lawrence Zitnick. 2015. [Microsoft coco captions: Data collection and evaluation server](#). *ArXiv*, abs/1504.00325.
- Min Jin Chong and David Forsyth. 2020. [Effectively unbiased fid and inception score and where to find them](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Katherine Crowson, Stella Biderman, Daniel Kornis, Dashiell Stander, Eric Hallahan, Louis Castricato, and Edward Raff. 2022. [Vqgan-clip: Open domain image generation and editing with natural language guidance](#). In *Computer Vision – ECCV 2022*, pages 88–105, Cham. Springer Nature Switzerland.
- Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, and Jie Tang. 2021. [Cogview: Mastering text-to-image generation via transformers](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 19822–19835. Curran Associates, Inc.
- Desmond Elliott, Stella Frank, Loïc Barrault, Fethi Bougares, and Lucia Specia. 2017. [Findings of the second shared task on multimodal machine translation and multilingual image description](#). In *Proceedings of the Second Conference on Machine Translation*, pages 215–233, Copenhagen, Denmark. Association for Computational Linguistics.
- Desmond Elliott, Stella Frank, Khalil Sima’an, and Lucia Specia. 2016. [Multi30K: Multilingual English-German image descriptions](#). In *Proceedings of the 5th Workshop on Vision and Language*, pages 70–74, Berlin, Germany. Association for Computational Linguistics.

- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. [Beyond english-centric multilingual machine translation](#). *Journal of Machine Learning Research*, 22(107):1–48.
- Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. 2022. [Make-a-scene: Scene-based text-to-image generation with human priors](#). *ArXiv*, abs/2203.13131.
- Víctor Gallego. 2022. [Personalizing text-to-image generation via aesthetic gradients](#). *ArXiv*, abs/2209.12330.
- Karthik Ganesan, Pakhi Bamdev, Jaivarsan B, Amresh Venugopal, and Abhinav Tushar. 2021. [N-best ASR transformer: Enhancing SLU performance using multiple ASR hypotheses](#). 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)*, pages 93–98, Online. Association for Computational Linguistics.
- Goran Glavaš, Robert Litschko, Sebastian Ruder, and Ivan Vulić. 2019. [How to \(properly\) evaluate cross-lingual word embeddings: On strong baselines, comparative analyses, and some misconceptions](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 710–721, Florence, Italy. Association for Computational Linguistics.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. [Gans trained by a two time-scale update rule converge to a local nash equilibrium](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.
- Xun Huang and Serge Belongie. 2017. [Arbitrary style transfer in real-time with adaptive instance normalization](#). In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Marcin Junczys-Dowmunt, Kenneth Heafield, Hieu Hoang, Roman Grundkiewicz, and Anthony Aue. 2018. [Marian: Cost-effective high-quality neural machine translation in C++](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 129–135, Melbourne, Australia. Association for Computational Linguistics.
- SeongJun Jung, Woo Suk Choi, Seongho Choi, and Byoung-Tak Zhang. 2022. [Language-agnostic semantic consistent text-to-image generation](#). In *Proceedings of the Workshop on Multilingual Multimodal Learning*, pages 1–5, Dublin, Ireland and Online. Association for Computational Linguistics.
- Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. 2023. [Scaling up gans for text-to-image synthesis](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. 2020a. [Training generative adversarial networks with limited data](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 12104–12114. Curran Associates, Inc.
- Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2021. [Alias-free generative adversarial networks](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 852–863. Curran Associates, Inc.
- Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020b. [Analyzing and improving the image quality of stylegan](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Weiyu Lan, Xirong Li, and Jianfeng Dong. 2017. [Fluency-guided cross-lingual image captioning](#). In *Proceedings of the 25th ACM International Conference on Multimedia, MM '17*, page 1549–1557, New York, NY, USA. Association for Computing Machinery.
- M. Paul Lewis, editor. 2009. *Ethnologue: Languages of the World*, sixteenth edition. SIL International.
- Xirong Li, Chaoxi Xu, Xiaoxu Wang, Weiyu Lan, Zhengxiong Jia, Gang Yang, and Jieping Xu. 2019. [Coco-cn for cross-lingual image tagging, captioning, and retrieval](#). *IEEE Transactions on Multimedia*, 21:2347–2360.
- Yaoyiran Li, Fangyu Liu, Nigel Collier, Anna Korhonen, and Ivan Vulić. 2022a. [Improving word translation via two-stage contrastive learning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4353–4374, Dublin, Ireland. Association for Computational Linguistics.
- Yaoyiran Li, Fangyu Liu, Ivan Vulić, and Anna Korhonen. 2022b. [Improving bilingual lexicon induction with cross-encoder reranking](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4100–4116, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

- Xinyue Liu, Mingda Li, Luoxin Chen, Prashan Wani-gasekara, Weitong Ruan, Haidar Khan, Wael Hamza, and Chengwei Su. 2021. [Asr n-best fusion nets](#). In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7618–7622.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Sasha Luccioni, Victor Schmidt, Alexandre Lacoste, and Thomas Dandres. 2019. [Quantifying the carbon emissions of machine learning](#). In *NeurIPS 2019 Workshop on Tackling Climate Change with Machine Learning*.
- Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2022. [GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 16784–16804. PMLR.
- Jordi Pont-Tuset, Jasper Uijlings, Soravit Changpinyo, Radu Soricut, and Vittorio Ferrari. 2020. [Connecting vision and language with localized narratives](#). In *Computer Vision – ECCV 2020*, pages 647–664, Cham. Springer International Publishing.
- E. Ponti, Julia Kreutzer, Ivan Vulic, and Siva Reddy. 2021. [Modelling latent translations for cross-lingual transfer](#). *ArXiv*, abs/2107.11353.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. [Hierarchical text-conditional image generation with clip latents](#). *ArXiv*, abs/2204.06125.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. [Zero-shot text-to-image generation](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8821–8831. PMLR.
- Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. 2016. [Generative adversarial text to image synthesis](#). In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1060–1069, New York, New York, USA. PMLR.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. [High-resolution image synthesis with latent diffusion models](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, Seyedeh Sara Mahdavi, Raphael Gontijo Lopes, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. 2022. [Photorealistic text-to-image diffusion models with deep language understanding](#). *ArXiv*, abs/2205.11487.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. 2016. [Improved techniques for training gans](#). In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. 2022. [Laion-5b: An open large-scale dataset for training next generation image-text models](#). *ArXiv*, abs/2210.08402.
- Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. 2020. [Green AI](#). *Communications of the ACM*, 63(12):54–63.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2021. [Multilingual translation from denoising pre-training](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3450–3466, Online. Association for Computational Linguistics.
- Ming Tao, Hao Tang, Fei Wu, Xiao-Yuan Jing, Bing-Kun Bao, and Changsheng Xu. 2022. [Df-gan: A simple and effective baseline for text-to-image synthesis](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16515–16525.
- Jörg Tiedemann and Santhosh Thottingal. 2020. [OPUS-MT – building open translation services for the world](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 479–480, Lisboa, Portugal. European Association for Machine Translation.
- Emiel van Miltenburg, Desmond Elliott, and Piek Vossen. 2017. [Cross-linguistic differences and similarities in image descriptions](#). In *Proceedings of the*

- 10th International Conference on Natural Language Generation*, pages 21–30, Santiago de Compostela, Spain. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Zongze Wu, Dani Lischinski, and Eli Shechtman. 2021. [Stylespace analysis: Disentangled controls for stylegan image generation](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12863–12872.
- Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. 2018. [Attngan: Fine-grained text to image generation with attentional generative adversarial networks](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. [From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions](#). *Transactions of the Association for Computational Linguistics*, 2:67–78.
- Han Zhang, Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. 2021. [Cross-modal contrastive learning for text-to-image generation](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 833–842.
- Han Zhang, Suyi Yang, and Hongqing Zhu. 2022. [Cjetig: Zero-shot cross-lingual text-to-image generation by corpora-based joint encoding](#). *Knowledge-Based Systems*, 239:108006.
- Mingyang Zhou, Luwei Zhou, Shuohang Wang, Yu Cheng, Linjie Li, Zhou Yu, and Jingjing Liu. 2021. [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)*, pages 4155–4165.
- Yufan Zhou, Ruiyi Zhang, Changyou Chen, Chunyuan Li, Chris Tensmeyer, Tong Yu, Jiuxiang Gu, Jinhui Xu, and Tong Sun. 2022. [Towards language-free training for text-to-image generation](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17907–17917.
- Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. 2019. [Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

A Data Statistics and Languages

In Table 6, we summarise the data statistics and languages covered in our experiments.

B Additional Discussion on Data Sources

Even without human-annotated image descriptions, there are two possible ways to derive captions for a target language \mathbb{L} .

First, we could translate EN captions into \mathbb{L} manually (still costly) or via machine translation. Our **TRANSLATE TRAIN** baseline (see §3) derives training data via machine translation and trains an \mathbb{L} TTI model from scratch. One main disadvantage of this approach is that it incurs huge training costs. While translations can be used as training data, we are conservative about using translated captions for TTI evaluation which can cause unexpected bias (Elliott et al., 2016; van Miltenburg et al., 2017; Bugliarello et al., 2022).

Second, it is possible to use cheaper but noisy Web-crawled visual-language data. For example, the recently released LAION-5B dataset (Schuhmann et al., 2022) has 5 billion image-text pairs for 100+ languages. There are previous examples that successfully trained SotA EN TTI models with Web-crawled data, such as large VQVAE-based models and diffusion models. The models described in Ramesh et al. (2021), Nichol et al. (2022) and Ramesh et al. (2022) are trained on EN large-scale Web-crawled data, but are eventually also tested on the gold-standard MS-COCO validation set. In our work, in addition to two gold-standard datasets, we also try to build on our own a small-scale dataset for both training and evaluation by filtering relatively good-quality image-text pairs from a subset of the noisy LAION-5B data (details in §4). Training non-EN TTI models from scratch with large-scale Web-crawled data such as LAION-5B is out of the scope of our work, and we focus on cross-lingual transfer learning setups with limited \mathbb{L} data. As mentioned in §1, this is to a large extent due to concerns about huge computational costs for training TTI models. Moreover, there are circa 7,000 languages worldwide (Lewis, 2009), and for low-resource languages not covered in LAION-5B’s 100+ languages, cross-lingual transfer learning approaches would still be the first choice. Furthermore, the number of EN texts in LAION-5B is more than the total amount of texts from its 100+ non-EN texts. Making full use of the huge amount of EN image-text pairs via cross-lingual transfer learning

Algorithm 1 Supervised Training of Ensemble Adapter

```
1: Input: An image-text dataset  $\{\mathbf{x}_i^0, \mathcal{I}_i^{real}\}_{i=1}^N$ 
2: Derive  $\mathbf{H}_i$  for each  $\mathbf{x}_i^0$  with NMT and mCLIP
3: while not converge do:
4:   Sample mini-batch  $\{\mathbf{H}_i, \mathcal{I}_i^{real}\}_{i=1}^n$ ;
5:   Sample random noise  $\{\mathbf{z}_i\}_{i=1}^n \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ;
6:   ENSAD forward pass  $\tilde{\mathbf{h}}_i \leftarrow \text{ENSAD}(\mathbf{H}_i)$ ;
7:   Synthesise fake image  $\tilde{\mathcal{I}}_i^{fake} \leftarrow G(\tilde{\mathbf{h}}_i, \mathbf{z}_i)$ 
8:   Feed  $(\tilde{\mathbf{h}}_i, \mathcal{I}_i^{real})$  and  $(\tilde{\mathbf{h}}_i, \tilde{\mathcal{I}}_i^{fake})$  to D respectively;
9:   Update ENSAD with Eq. (13);
10:  Update  $D$  with Eq. (14);
11: end while
```

might be beneficial for other languages. Therefore, we think that cross-lingual transfer learning in relatively low-resource setups for multilingual TTI is a critical and valuable research topic.

C The Detailed Training Process of ENSAD

We summarise the training process of our ENSAD method (see §3) in Algorithm 1.

D Deriving LAION-5B Dataset for Finnish

We download circa 5.1 million image-caption pairs from the FI category of LAION-5B. Since the Web-crawled data are noisy, we apply several filtering steps: **1)** since our images will be scaled to resolution 256×256 , to avoid distortion we keep only images with their width-height ratio between 0.5 and 2; **2)** we keep captions with a minimum length of 8 words, which is also a requirement of MS-COCO (Chen et al., 2015) in its data annotation; **3)** we use the langdetect library²³ to remove texts misclassified into the LAION-5B FI category and make sure the texts left are actually in Finnish; **4)** we keep captions with one ending period ‘.’.²⁴ After these steps, 239K pairs are left, and we calculate mCLIP scores (cosine similarities between mCLIP-extracted text and image features) for all the pairs and keep the 30K highest-ranking pairs as the final dataset. We randomly split the data into training, development and test portions with 10,000, 2,000, and 18,000 pairs, respectively.

We ‘sanity-check’ 50 randomly sampled instances from our filtered data and find that, in most cases, the text matches the image content. But there are a small number of exceptional cases where the

²³<https://pypi.org/project/langdetect/>

²⁴In our initial trial, we found that, among the highest mCLIP-scored 30K pairs, most captions which do not end with ‘.’ are noisy short ads.

Language	Family	Code	Dataset	Training Set: # of Images	Dev Set: # of Images	Test Set: # of Images	Min Seq Len	Max Seq Len	Avg. Seq Len	Image Domain Overlap
English	Germanic	EN	MS-COCO	74,505	8,278	40,504	5	50	10.5	-
Chinese	Sino-Tibetan	ZH	COCO-CN	10,875	2,718	6,748	5	63	17.3	✓
German	Germanic	DE	Multi30K Task2	10,000	2,000	19,014	1	34	8.2	x
Finnish	Uralic	FI	LAION-5B	10,000	2,000	18,000	8	116	14.6	x
Spanish	Romance	ES	IGLUE xFlickr&CO	0	0	2,000	3	59	13.7	✓
Indonesian	Austronesian	ID	IGLUE xFlickr&CO	0	0	1,999	3	31	11.7	✓
Japanese	Japonic	JA	IGLUE xFlickr&CO	0	0	2,000	5	175	33.8	✓
Russian	Slavic	RU	IGLUE xFlickr&CO	0	0	2,000	1	45	11.3	✓
Turkish	Turkic	TR	IGLUE xFlickr&CO	0	0	2,000	2	30	9.5	✓

Table 6: Data statistics categorised by languages. This table includes information such as language family, ISO 639-1 code, dataset name, train/dev/test split, and statistics on sequence length (number of words per caption). Note that MS-COCO EN data is used for pretraining our mLAFITE only. We also show for each dataset if there is an image domain overlap with MS-COCO images used for mLAFITE pretraining. ✓: all images are from MS-COCO; x: none of the images is from MS-COCO; ✓: half of the images are from MS-COCO. For IGLUE Indonesian data, we remove an empty caption and its associated image, so there are 1,999 images left.

text contains extra information beyond the image content itself (e.g., event descriptions). Overall, the quality of our FI data still cannot match MS-COCO or Multi30K. Another interesting finding is that LAION-5B captions often use real and concrete names such as ‘Messi’ and ‘the national stadium’ to describe the image content, while MS-COCO and Multi30K tend to use general words such as ‘a man’/‘a football player’ and ‘a stadium’/‘a building’.

E RQ1: Results on IGLUE

Table 7 shows additional TTI results on five languages from IGLUE, comparing TRANSLATE TEST (with Marian and Amazon Translate) and ZERO-SHOT TRANSFER baselines.

F Robustness of ENSAD

We train the ‘ENSAD (Frozen G)’ model on COCO-CN 6 more times (7 times in total) with different random seeds, and for each saved model we run TTI evaluation three times.²⁵ Finally, we get 21 FID results, with min 14.47, max 14.62, mean 14.55, and standard deviation 0.04. Even the worst score of 14.62 outperforms all other baselines on COCO-CN.

G Reproducibility Checklist

- **TTI Data:** the datasets used in our work are all publicly available including MS-COCO²⁶, COCO-CN²⁷, Multi30K Task2²⁸, LAION-

²⁵DE and FI have larger test set sizes, so their evaluation results should be more reliable (Chong and Forsyth, 2020).

²⁶<https://cocodataset.org>

²⁷<https://github.com/li-xirong/coco-cn>

²⁸<https://github.com/multi30k/dataset>

5B²⁹, and IGLUE³⁰.

- **Parameter Counts:** the number of parameters is 655,873 for our ensemble adapter network, 44,997,026 for the generator network, 29,126,785 for the discriminator network, 560,415,232 for the mCLIP text encoder ‘M-CLIP/XLM-Roberta-Large-Vit-B-32’³¹, and 87,849,216 for the CLIP visual encoder ‘ViT-B/32’³².
- **Computing Infrastructure:** we run our code on an Amazon EC2 P3.16xlarge Instance with 8×16GB Nvidia[®] Tesla[®] V100 GPUs, 64×2.30 GHz Intel[®] Xeon[®] E5-2686 v4 CPU cores, and 488GB RAM.
- **Software:** Python 3.7.0, PyTorch 1.12.1, and Transformers 4.21.0.
- **Hyperparameter Search:** our hyperparameters are tuned on our dev split of COCO-CN. The same hyper-parameters are used for Multi30K and LAION-5B (we also conduct minimal tuning on their dev sets and find that the hyper-parameters tuned on COCO-CN are already (near-)optimal in our initial investigation). The learning rate is selected from $\{5e-5, 2.5e-4, 5e-4, 2.5e-3, 5e-3\}$, λ_1 and λ_2 which are weights for contrastive losses from $\{0.5, 1, 2, 4, 5, 10\}$, α the interpolation hyperparameter from $\{0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4, 0.5\}$, and d_{hid} from $\{32, 64, 128, 256, 512\}$.

²⁹<https://laion.ai/blog/laion-5b>

³⁰<https://github.com/e-bug/iglu>

³¹<https://github.com/FreddeFrallan/Multilingual-CLIP>

³²<https://github.com/openai/CLIP>

Method	ES: FID ↓	ID: FID ↓	JA: FID ↓	RU: FID ↓	TR: FID ↓	Avg.: FID ↓
TRANSLATE TEST (Marian)	30.04	31.09	33.12	31.11	30.74	31.22
TRANSLATE TEST (Amazon Translate)	30.08	31.27	31.61	30.83	30.31	30.82
ZERO-SHOT TRANSFER	30.31	30.58	32.24	30.77	30.12	30.8

Table 7: RQ1 results: TRANSLATE TEST vs. ZERO-SHOT TRANSFER on five languages from IGLUE. **FID**↓: lower is better.

- **Runtime:** it takes 75 hours to train an mLAFITE TTI model or a TRANSLATE TRAIN model from scratch, 7 hours to train an ENSAD based on a pretrained mLAFITE, 7.5 hours to fine-tune G (without ENSAD) based on a pretrained mLAFITE, and about 4 minutes to run FID evaluation for our TTI model with ENSAD (data preprocessing, NMT, and mCLIP feature extraction excluded). All experiments and measurements were conducted on $8 \times 16\text{GB}$ V100 GPUs.
- **Other Technical Details:** we adopt the ‘exponential sharpening’ for all contrastive losses as specified in LAFITE’s supplementary material.³³
- **Carbon Footprint:** we estimate that **1)** training an mLAFITE TTI model or a TRANSLATE TRAIN model from scratch can cause the emission of circa 56~67-kg CO₂ equivalents; **2)** training an ENSAD model would result in about 5~6-kg CO₂ equivalents. These estimations are based on our computational infrastructure and a publicly available ‘machine learning emissions calculator’ (Luccioni et al., 2019).³⁴

H TTI Examples and Attention Scores

H.1 TTI Examples

We compare images generated with TRANSLATE TEST, ZERO-SHOT TRANSFER, and our best ENSAD model in Figure 2, where for each TTI method we present two images generated with different random noise inputs as introduced in §3. The ‘Best’ model here refer to our ENSAD model that achieve the best FID scores (**bold** numbers) in Table 3 respectively for each language, i.e., ‘ENSAD (Frozen G)’ for ZH and ‘ENSAD + Fine-Tune G (L Text)’ for DE and FI. The differences between images generated with different TTI methods are very subtle.

³³https://openaccess.thecvf.com/content/CVPR2022/supplemental/Zhou_Towards_Language-Free_Training_CVPR_2022_supplemental.pdf

³⁴<https://mlco2.github.io/impact/#compute>

H.2 Attention Scores

Table 8 includes the original \mathbb{L} input text, EN translations, and their associated ENSAD attention scores (in descending order) corresponding to the images in Figure 2. We did not identify any salient pattern concerning the type of EN translations to which higher ENSAD attention scores are attached.

H.3 Can ENSAD Incorporate Manually Added Information from Translations?

To better understand what kind of information ENSAD extracts from EN translations, we also try to manually add additional information to EN translations (the additional information does not appear in and is not added to the original \mathbb{L} input). Of course, this section is for probing purposes only since MT systems are not likely to produce the same translations. We found that when the additional information is added to only several of the 12 EN translations, it can hardly get reflected in the generated image. Here, we show two COCO-CN test set examples in Figure 3 where we add the new information into 12 EN translations simultaneously. In its first and second rows, the original \mathbb{L} input is ‘An open laptop is on the table.’ and ‘It’s a clean, but crowded kitchen.’ respectively (translated from the original Chinese captions). We manually add new objects ‘roses’ and ‘fruits’ respectively to all their EN translations as in Table 9. As seen in Figure 3, the roses and fruits do appear in the generated images.

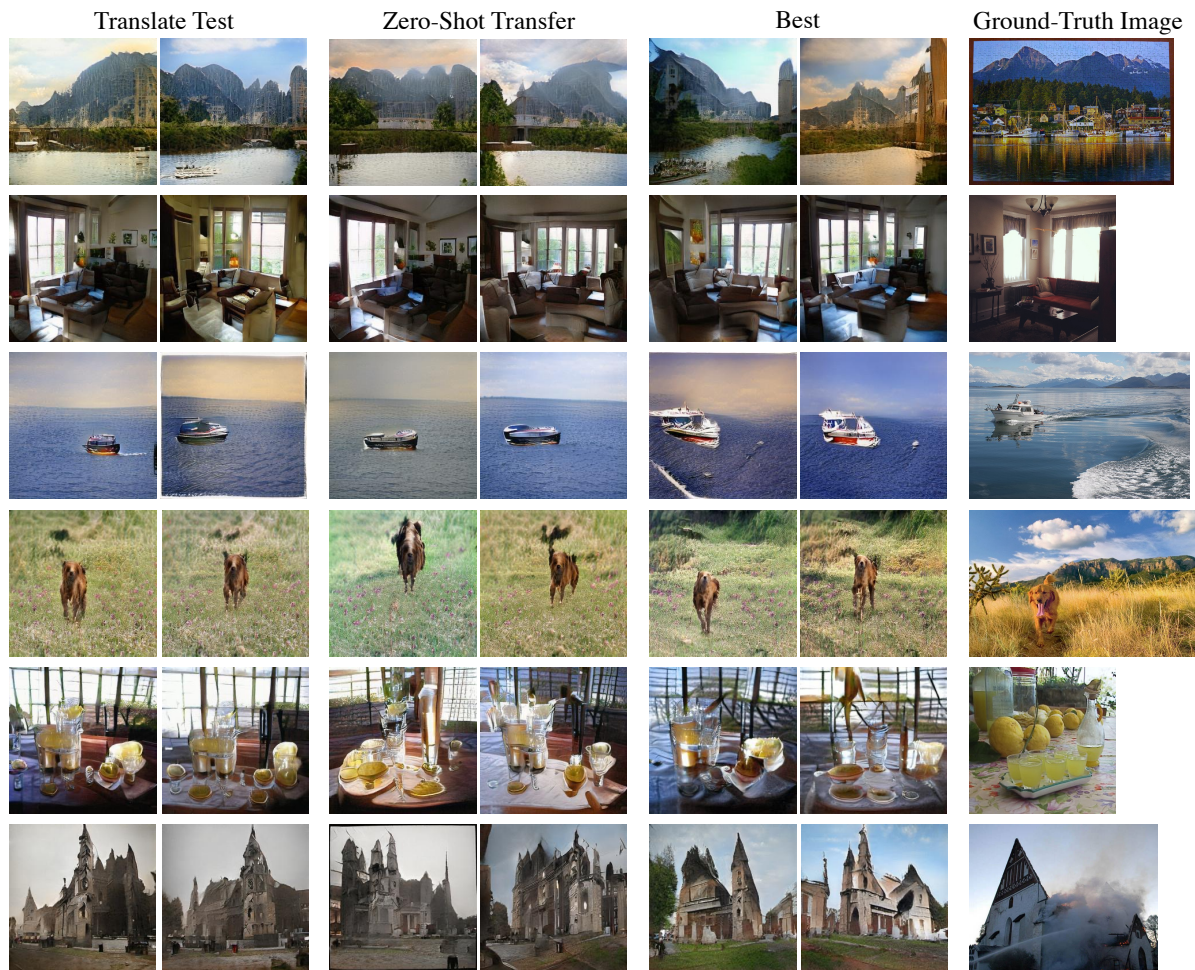


Figure 2: TTI Examples generated with TRANSLATE TEST, ZERO-SHOT TRANSFER, and our best model. COCO-CN (ZH) Test Set: row 1–2; Multi30K Task2 Test Set (DE): row 3–4; LAION-5B (FI) Test Set: row 5–6. The resolution of the generated images is 256×256 pixels; ground-truth images are shown in their original sizes respectively.




Figure 3: Images generated with and without manually added information (COCO-CN Test set). The resolution of the generated images is 256×256 pixels; ground-truth images are shown in their original sizes respectively.

Original L Input	EN Translations	ENSAD Attention Scores	
房屋依水而建，远处群山环绕。	Houses are built with water, and they are surrounded by mountains at a distance.	8.75e-01	
	Houses have been built waternly, surrounded by mountain ranges from one side to the other.	6.37e-02	
	Houses are built on water, all around mountain mountains, as long as possible, and have access to water.	2.97e-02	
	The house is constructed in the form of water, surrounded by mountains and long distances.	9.80e-03	
	Houses are built with water and are located far beyond the range of hills around them.	9.17e-03	
	The homes have been built on water and are surrounded by mountain areas from a distance.	5.64e-03	
	Houses are built on water and surround it far from the mountains.	4.62e-03	
	The houses are built according to water and spread around them from a great direction to a very deep range of mountains.	1.80e-03	
	Houses are constructed around the mountain and built from a distant distance to an open point of view.	3.64e-04	
	The houses are built by water and are encircled by mountains, as far as the hills are concerned.	2.27e-04	
	The houses were built in the form of water. They were in a remote area around the mountains.	8.77e-05	
	The houses were built watery and were driven from a very distant part of the forest and surrounded by mountains.	5.30e-08	
	一个客厅，一个大窗户下面的沙发，桌子。	A living room, a couch under a huge window, a table.	3.08e-01
		A sitting room, a couch under a big window, a table.	2.24e-01
A living room, a sofa under a big window, a table.		1.59e-01	
A living room, a couch under a big window, a table.		1.31e-01	
A living room, a couch under a big window, a table.		1.31e-01	
A living hall, a couch under that big window, a table.		2.12e-02	
I was in the living room, the sofa below the great window, the table.		1.27e-02	
In the living room, in the couch under a big window, in the table.		5.58e-03	
One living room. One large window under the couch. The table.		3.86e-03	
There was a living room, a couch under a large window, there was a table.		1.66e-03	
There's one room, and a big couch under the large window, and there's a table.		1.22e-03	
There was a guest room, there was a couch underneath a great window, there was a table.		7.53e-04	
Motorboot fährt auf ruhigem Gewässer		Motor boat sails on calm waters	4.47e-01
		Motor boat sails on calm waters	4.47e-01
	Motor boat cruises on calm waters	6.43e-02	
	Motorboat cruises on calm waters	2.14e-02	
	Motor boat rides on calm waters	5.98e-03	
	Motorboat travels on calm waters	5.46e-03	
	Motorboat travels on calm waters	5.46e-03	
	Motorboat drives on calm waters	1.60e-03	
	Motorboat rides on calm waters	3.65e-04	
	Motorboat rides on calm waters	3.65e-04	
	Motorboat is sailing on calm waters	1.63e-04	
	Motorboat is on the sea in order to keep its pace and to move towards the sea.	4.92e-05	
	Einen braunen Hund der spazieren geht in der Wiese.	I think he'd be able to walk in the meadow and we could have a brown dog to go for a walk.	5.72e-01
		He walks a brown dog in the meadows, who goes for a walk.	4.24e-01
A brown dog that goes walking in the meadow.		1.88e-03	
A brown dog who goes for walks in the meadow.		1.16e-03	
A brown dog who goes for walks in the meadow.		1.16e-03	
A brown dog who goes for a walk in the meadow.		1.02e-04	
A brown dog going for a walk in the meadow.		1.10e-06	
A brown dog who walks in the meadow.		8.78e-09	
A brown dog taking a walk in the meadow.		4.56e-09	
A brown dog taking a walk in the meadow.		4.56e-09	
A brown dog walking in the meadow.		1.66e-09	
A brown dog walking in the meadow.		1.66e-09	
Terassin pöydällä kotitekoista limoncelloa lasipurkissa ja karahivissa sekä kaadettuna piemiin laseihin.		Home made soda on the terrace in glass jar and karaffles and poured into small glasses.	2.07e-01
		on the terrace table made homemade sodacello in a glass jar and in a girdle as well as poured into small glasses.	1.26e-01
	On the terrace table in a glass jar and a karahas of homemade limocello put down in small glasses.	8.66e-02	
	On a terrace table of homemade lemonade in a glass jar and slab of gizzard and poured into small glasses.	8.64e-02	
	The table on the terrace has homemade soda crystals in a glass jar and swath and is poured into small glasses.	8.30e-02	
	On the terrace table housed lemonade in glass jars and swaths and poured down into small glasses.	8.24e-02	
	On the terrace table of home made wine in glass jars and karaffles and poured into small glasses.	7.41e-02	
	The terrace is equipped with homemade lemonade in the jar and perch and poured into small glasses.	5.89e-02	
	Top of the terrace is homemade lemonade in a jar of glass and karaoke and poured into small glasses.	5.67e-02	
	On the table of the terrace it's homemade limocello with glass pots and clovers and poured into small glasses.	5.50e-02	
	On a table of terraces, homemade limoncello is made in a glass jar and in a caraments and poured into small glasses.	4.91e-02	
	on the table of terraces with homemade soda on a glass jar and karaffe and poured in small glasses.	3.56e-02	
	Tuli tuhosi pahoin historiallisen kirkon vuoden 2006 toukokuussa.	In May 2006 the historic church was badly destroyed by fire.	3.51e-01
		In May of 2006, the historical church was severely destroyed by fire.	2.39e-01
It was, in May 2006, when the fire badly destroyed the historic church.		9.44e-02	
There was a great destruction of this historical church in May 2006.		6.56e-02	
The fire did a great deal of damage to an historic church in May 2006.		5.30e-02	
In May 2006 fire caused a very severe damage to the historic church.		3.58e-02	
The fire seriously destroyed the historical church in May 2006.		3.49e-02	
The fire was severely destroyed by the historical church in May 2006.		3.49e-02	
A fire severely destroyed the historical church in May 2006.		3.05e-02	
There's been massive damage to the historical Church in May 2006 when the fire took place.		3.01e-02	
Fire was devastatingly damaged by the historic church in May 2006.		2.18e-02	
Fire caused the serious destruction of the historic church in May 2006.		9.31e-03	

Table 8: ENSAD attention scores.

Original L Input	Modified EN Translations
桌子上摆放着一个打开的笔记本电脑。	On a table is put on an open laptop, <u>and roses</u> .
	It was on the desk with an open laptop, <u>and roses</u> .
	There's a computer that's open that has an open laptop sitting on the table, <u>and roses</u> .
	There's a opened laptop on the table, <u>and roses</u> .
	There's an open laptop sitting on the table, <u>and roses</u> .
	An open laptop's on the table, <u>and roses</u> .
	And we have a laptop on your desk that's open, <u>and roses</u> .
	There was a laptop that was open on the table, <u>and roses</u> .
	A computer that opened up his laptop is in place on the table, <u>and roses</u> .
	There was a computer on the table. There was an open laptop on the table, <u>and roses</u> .
这是一个干净，但拥挤的厨房。	There's an open laptop on the table, <u>and roses</u> .
	There was an unopened laptop on the table, <u>and roses</u> .
	That's a clean-up, but crowded kitchen <u>full of fruits</u> .
	It's clean but crowded in the kitchen <u>full of fruits</u> .
	That's a clean, but crowded kitchen <u>full of fruits</u> .
	It's a clean, but crowded kitchen <u>full of fruits</u> .
	It's a clean-up but congested kitchen <u>full of fruits</u> .
	That's a clean, but congested kitchen <u>full of fruits</u> .
	And it's a clean, but crowded kitchen <u>full of fruits</u> .
	It's a clean, but congested kitchen <u>full of fruits</u> .
- IT'S THIS IS A cleanING BUT CLOTHED CLIMBEN COILLOR IN THE CRUCKIT. - [CLICKS] <u>full of fruits</u> .	
It was a clean but crowd-cooked kitchen <u>full of fruits</u> .	
That's a clean one, but crowd-cooked kitchen <u>full of fruits</u> .	
It's a clean, but congested kitchen <u>full of fruits</u> .	

Table 9: Additional information added to the EN translations. The underlined texts in red are added phrases. Removing the phrases derives the NMT-generated translations.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
The Limitations section.
- A2. Did you discuss any potential risks of your work?
The Ethics Statement section.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Section 1 Introduction and Motivation.
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 3 Methodology, Section 4 Datasets, Section 5 Experimental Setup, and Appendix G Reproducibility Checklist

- B1. Did you cite the creators of artifacts you used?
Section 4 Datasets, Section 5 Experimental Setup, and Appendix G Reproducibility Checklist
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Section 4 Datasets, Section 5 Experimental Setup, the Ethics Statement section, and Appendix G Reproducibility Checklist.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Section 4 Datasets, Section 5 Experimental Setup, the Ethics Statement section, and Appendix G Reproducibility Checklist.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
The Ethics Statement section.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 4 Datasets, Section 5 Experimental Setup, Appendix A Data Statistics and Languages, and Appendix G Reproducibility Checklist.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 4 Datasets, Section 5 Experimental Setup, and Appendix A Data Statistics and Languages

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

C Did you run computational experiments?

Section 5 Experimental Setup, and Section 6 Results and Discussion

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

Appendix G Reproducibility Checklist

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section 5 Experimental Setup, Appendix G Reproducibility Checklist

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 6 Results and Discussion, Appendix F Robustness of ENSAD

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section 5 Experimental Setup, Appendix G Reproducibility Checklist

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Not applicable. Left blank.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Not applicable. Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Not applicable. Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Not applicable. Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Not applicable. Left blank.