

Good for Misconceived Reasons: An Empirical Revisiting on the Need for Visual Context in Multimodal Machine Translation

Zhiyong Wu^{†*}, Lingpeng Kong^{†¶}, Wei Bi[‡], Xiang Li[§], Ben Kao[†]

[†]The University of Hong Kong,

[‡]Tencent AI Lab, [¶]Shanghai Artificial Intelligence Laboratory, [§]East China Normal University
[†]{zywu,lpk,kao}@cs.hku.hk, [‡]victoriabi@tencent.com, [§]xiangli@dase.ecnu.edu.cn

Abstract

A neural multimodal machine translation (MMT) system is one that aims to perform better translation by extending conventional text-only translation models with multimodal information. Many recent studies report improvements when equipping their models with the multimodal module, despite the controversy of whether such improvements indeed come from the multimodal part. We revisit the contribution of multimodal information in MMT by devising two *interpretable* MMT models. To our surprise, although our models replicate similar gains as recently developed multimodal-integrated systems achieved, our models learn to *ignore* the multimodal information. Upon further investigation, we discover that the improvements achieved by the multimodal models over text-only counterparts are in fact results of the regularization effect. We report empirical findings that highlight the importance of MMT models' interpretability, and discuss how our findings will benefit future research.

1 Introduction

Multimodal Machine Translation (MMT) aims at designing better translation systems by extending conventional text-only translation systems to take into account multimodal information, especially from visual modality (Specia et al., 2016; Wang et al., 2019). Despite many previous success in MMT that report improvements when models are equipped with visual information (Calixto et al., 2017; Helcl et al., 2018; Ive et al., 2019; Lin et al., 2020; Yin et al., 2020), there have been continuing debates on the need for visual context in MMT.

In particular, Specia et al. (2016); Elliott et al. (2017); Barrault et al. (2018) argue that visual context does not seem to help translation reliably, at

least as measured by automatic metrics. Elliott (2018); Grönroos et al. (2018a) provide further evidence by showing that MMT models are, in fact, insensitive to visual input and can translate without significant performance losses even in the presence of features derived from unrelated images. A more recent study (Caglayan et al., 2019), however, shows that under limited textual context (e.g., noun words are masked), models can leverage visual input to generate better translations. But it remains unclear where the gains of MMT methods come from, when the textual context is complete.

The main tool utilized in prior discussion is *adversarial model comparison* — explaining the behavior of complex and black-box MMT models by comparing performance changes when given adversarial input (e.g., random images). Although such an opaque tool is an acceptable beginning to investigate the need for visual context in MMT, they provide rather indirect evidence (Hessel and Lee, 2020). This is because performance differences can often be attributed to factors unrelated to visual input, such as regularization (Kukačka et al., 2017), data bias (Jabri et al., 2016), and some others (Dodge et al., 2019).

From these perspectives, we revisit the need for visual context in MMT by designing two interpretable models. Instead of directly infusing visual features into the model, we design learnable components, which allow the model to voluntarily decide the usefulness of the visual features and reinforce their effects when they are helpful. To our surprise, while our models are shown to be effective on Multi30k (Elliott et al., 2016) and VaTex (Wang et al., 2019) datasets, they learn to *ignore* the multimodal information. Our further analysis suggests that under sufficient textual context, the improvements come from a regularization effect that is similar to random noise injection (Bishop, 1995) and weight decay (Hanson and Pratt, 1989). The addi-

*The majority of this work was done while the first author was interning at Tencent AI Lab.

tional visual information is treated as noise signals that can be used to enhance model training and lead to a more robust network with lower generalization error (Salamon and Bello, 2017). Repeating the evaluation under limited textual context further substantiates our findings and complements previous analysis (Caglayan et al., 2019).

Our contributions are twofold. First, we revisit the need for visual context in the popular task of multimodal machine translation and find that: (1) under sufficient textual context, the MMT models’ improvements over text-only counterparts result from the regularization effect (Section 5.2). (2) under limited textual context, MMT models can leverage visual context to help translation (Section 5.3). Our findings highlight the importance of MMT models’ interpretability and the need for a new benchmark to advance the community.

Second, for the MMT task, we provide a strong text-only baseline implementation and two models with interpretable components that replicate similar gains as reported in previous works. Different from adversarial model comparison methods, our models are interpretable due to the specifically designed model structure and can serve as standard baselines for future interpretable MMT studies. Our code is available at <https://github.com/LividWo/Revisit-MMT>.

2 Background

One can broadly categorize MMT systems into two types: (1) Conventional MMT, where there is gold alignment between the source (target) sentence pair and a relevant image and (2) Retrieval-based MMT, where systems retrieve relevant images from an image corpus as additional clues to assist translation.

Conventional MMT Most MMT systems require datasets consist of images with bilingual annotations for both training and inference. Many early attempts use a pre-trained model (e.g., ResNet (He et al., 2016)) to encode images into feature vectors. This visual representation can be used to initialize the encoder/decoder’s hidden vectors (Elliott et al., 2015; Libovický and Helcl, 2017; Calixto et al., 2016). It can also be appended/prepended to word embeddings as additional input tokens (Huang et al., 2016; Calixto and Liu, 2017). Recent works (Libovický et al., 2018; Zhou et al., 2018; Ive et al., 2019; Lin et al., 2020) employ attention mechanism to generate a visual-aware representation for the decoder. For

instance, *Doubly-ATT* (Calixto et al., 2017; Helcl et al., 2018; Arslan et al., 2018) insert an extra visual attention sub-layer between the decoder’s source-target attention sub-layer and feed-forward sub-layer. While there are more works on engineering decoders, encoder-based approaches are relatively less explored. To this end, Yao and Wan (2020) and Yin et al. (2020) replace the vanilla Transformer encoder with a multi-modal encoder.

Besides the exploration on network structure, researchers also propose to leverage the benefits of multi-tasking to improve MMT (Elliott and Kádár, 2017; Zhou et al., 2018). The *Imagination* architecture (Elliott and Kádár, 2017; Helcl et al., 2018) decomposes multimodal translation into two sub-tasks: translation task and an auxiliary visual reconstruction task, which encourages the model to learn a visually grounded source sentence representation.

Retrieval-based MMT The effectiveness of conventional MMT heavily relies on the availability of images with bilingual annotations. This could restrict its wide applicability. To address this issue, Zhang et al. (2020) propose *UVR-NMT* that integrates a retrieval component into MMT. They use TF-IDF to build a token-to-image lookup table, based on which images sharing similar topics with a source sentence are retrieved as relevant images. This creates image-bilingual-annotation instances for training. Retrieval-based models have been shown to improve performance across a variety of NLP tasks besides MMT, such as question answering (Guu et al., 2020), dialogue (Weston et al., 2018), language modeling (Khandelwal et al., 2019), question generation (Lewis et al., 2020), and translation (Gu et al., 2018).

3 Method

In this section we introduce two interpretable MMT models: (1) *Gated Fusion* for conventional MMT and (2) *Dense-Retrieval-augmented MMT* (RMMT) for retrieval-based MMT. Our design philosophy is that models should learn, in an interpretable manner, to which degree multimodal information is used. Following this principle, we focus on the component that integrates multimodal information. In particular, we use a gating matrix Λ (Yin et al., 2020; Zhang et al., 2020) to control the amount of visual information to be blended into the textual representation. Such a matrix facilitates interpreting the fusion process: a larger gating value $\Lambda_{ij} \in [0, 1]$ indicates that the model exploits more

visual context in translation, and vice versa.

3.1 Gated Fusion MMT

Given a source sentence x of length T and an associated image z , we compute the probability of generating target sentence y of length N by:

$$p(y|x, z) = \prod_i^N p_\theta(y_i | x, z, y_{<i}), \quad (1)$$

where $p_\theta(y_i | x, z, y_{<i})$ is implemented with a Transformer-based (Vaswani et al., 2017) network. Specifically, we first feed x into a vanilla Transformer encoder to obtain a textual representation $\mathbf{H}_{\text{text}} \in \mathbb{R}^{T \times d}$, which is then fused with visual representation $\text{Embed}_{\text{image}}(z)$ before fed into the Transformer decoder. For each image z , we use a pre-trained ResNet-50 CNN (He et al., 2016) to extract a 2048-dimensional average-pooled visual representation, which is then projected to the same dimension as \mathbf{H}_{text} :

$$\text{Embed}_{\text{image}}(z) = \mathbf{W}_z \text{ResNet}_{\text{pool}}(z). \quad (2)$$

We next generate a gating matrix $\Lambda \in [0, 1]^{T \times d}$ to control the fusion of \mathbf{H}_{text} and $\text{Embed}_{\text{image}}(z)$:

$$\Lambda = \text{sigmoid}(\mathbf{W}_\Lambda \text{Embed}_{\text{image}}(z) + \mathbf{U}_\Lambda \mathbf{H}_{\text{text}}),$$

where \mathbf{W}_Λ and \mathbf{U}_Λ are model parameters. Note that this gating mechanism has been a building block for many recent MMT systems (Zhang et al., 2020; Lin et al., 2020; Yin et al., 2020). We are, however, the first to focus on its interpretability. Finally, we generate the output vector \mathbf{H} by:

$$\mathbf{H} = \mathbf{H}_{\text{text}} + \Lambda \text{Embed}_{\text{image}}(z). \quad (3)$$

\mathbf{H} is then fed into the decoder directly for translation as in vanilla Transformer.

3.2 Retrieval-Augmented MMT (RMMT)

RMMT consists of two sequential components: (1) an image retriever $p(z|x)$ that takes x as input and returns Top- K most relevant images from an image database; (2) a multi-modal translator $p(y|x, \mathcal{Z}) = \prod_i^N p_\theta(y_i | x, \mathcal{Z}, y_{<i})$ that generates each y_i conditioned on the input sentence x , the image set \mathcal{Z} returned by the retriever, and the previously generated tokens $y_{<i}$.

Image Retriever Based on the TF-IDF model, searching in existing retrieval-based MMT (Zhang et al., 2020) ignores the context information of a given query, which could lead to poor performance. To improve the recall of our image retriever, we compute the similarity between a sentence x and an image z with inner product:

$$\text{sim}(x, z) = \text{Embed}_{\text{text}}(x)^\top \text{Embed}_{\text{image}}(z),$$

where $\text{Embed}_{\text{text}}(x)$ and $\text{Embed}_{\text{image}}(z)$ are d -dimensional representations of x and z , respectively. We then retrieve top- K images that are closest to x . For $\text{Embed}_{\text{image}}(z)$, we compute it by Eq. 2. For $\text{Embed}_{\text{text}}(x)$, we implement it using BERT (Devlin et al., 2019):

$$\text{Embed}_{\text{text}}(x) = \mathbf{W}_{\text{text}} \text{BERT}_{\text{CLS}}(x). \quad (4)$$

Following standard practices, we use a pre-trained BERT model¹ to obtain the ‘‘pooled’’ representation of the sequence (denoted as $\text{BERT}_{\text{CLS}}(x)$). Here, \mathbf{W}_{text} is a projection matrix.

Multimodal Translator Different from Gated Fusion, $p(y|x, \mathcal{Z})$ now is conditioning on a set of images rather than one single image. For each z in \mathcal{Z} , we represent it using $\text{Embed}_{\text{image}}(z) \in \mathbb{R}^d$ as in Equation 2. The image set \mathcal{Z} then forms a feature matrix $\text{Embed}_{\text{image}}(\mathcal{Z}) \in \mathbb{R}^{K \times d}$, where $K = |\mathcal{Z}|$ and each row corresponds to the feature vector of an image. We use a transformation layer $f_\theta(*)$ to extract salient features from $\text{Embed}_{\text{image}}(\mathcal{Z})$ and obtain a compressed representation \mathbb{R}^d of \mathcal{Z} . After the transformation, ideally, we can implement $p(y|x, \mathcal{Z})$ using any existing MMT models. For interpretability, we follow the Gated Fusion model to fuse the textual and visual representations with a learnable gating matrix Λ :

$$\mathbf{H} = \mathbf{H}_{\text{text}} + \Lambda f_\theta(\text{Embed}_{\text{image}}(\mathcal{Z})). \quad (5)$$

Here, $f_\theta(*)$ denotes a max-pooling layer with window size $K \times 1$.

4 Experiment

In this section, we evaluate our models on the Multi30k and VaTex benchmark.

4.1 Dataset

We perform experiments on the widely-used MMT datasets: Multi30k. We follow a standard split

¹Here we use bert-base-uncased version.

of 29,000 instances for training, 1,014 for validation and 1,000 for testing (Test2016). We also report results on the 2017 test set (Test2017) with extra 1,000 instances and the MSCOCO test set that includes 461 more challenging out-of-domain instances with ambiguous verbs. We merge the source and target sentences in the officially pre-processed version of Multi30k² to build a joint vocabulary. We then apply the byte pair encoding (BPE) algorithm (Sennrich et al., 2016) with 10,000 merging operations to segment words into subwords, which generates a vocabulary of 9,712 (9,544) tokens for En-De (En-Fr).

Retriever pre-training. We pre-train the retriever on a subset of the Flickr30k dataset (Plummer et al., 2015) that has overlapping instances with Multi30k removed. We use Multi30k’s validation set to evaluate the retriever. We measure the performance by recall-at- K ($R@K$), which is defined as the fraction of queries whose closest K images retrieved contain the correct images. The pre-trained retriever achieves $R@1$ of 22.8% and $R@5$ of 39.6%.

4.2 Setup

We experiment with different model sizes (*Base*, *Small*, and *Tiny*, see Appendix A for details). *Base* is a widely-used model configuration for Transformer in both text-only translation (Vaswani et al., 2017) and MMT (Grönroos et al., 2018b; Ive et al., 2019). However, for small datasets like Multi30k, training such a large model (about 50 million parameters) could cause overfitting. In our preliminary study, we found that even a *Small* configuration, which is commonly used for low-resourced translation (Zhu et al., 2019), can still overfit on Multi30k. We therefore perform grid search on the En→De validation set in Multi30k and obtain a *Tiny* configuration that works surprisingly well.

We use Adam with $\beta_1 = 0.9$, $\beta_2 = 0.98$ for model optimization. We start training with a warm-up phase (2,000 steps) where we linearly increase the learning rate from 10^{-7} to 0.005. Thereafter we decay the learning rate proportional to the number of updates. Each training batch contains at most 4,096 source/target tokens. We set label smoothing weight to 0.1, dropout to 0.3. We follow (Zhang et al., 2020) to early-stop the training if validation loss does not improve for ten epochs. We average the last ten checkpoints for inference as in (Vaswani et al., 2017) and (Wu et al., 2018). We perform

²<https://github.com/multi30k/dataset>

beam search with beam size set to 5. We report 4-gram BLEU and METEOR scores for all test sets. All models are trained and evaluated on one single machine with two Titan P100 GPUs.

4.3 Baselines

Our baselines can be categorized into three types:

- The text-only **Transformer**;
- The conventional MMT models: **Doubly-ATT** and **Imagination**;
- The retrieval-based MMT models: **UVR-NMT**.

Details of these methods can be found in Section 2. For fairness, all the baselines are implemented by ourselves based on FairSeq (Ott et al., 2019). We use top-5 retrieved images for both UVR-NMT and our RMMT. We also consider two more recent state-of-the-art conventional methods for reference: **GMNMT** (Yin et al., 2020) and **DCCN** (Lin et al., 2020), whose results are reported as in their papers.

Note that most MMT methods are difficult (or even impossible) to interpret. While there exist some interpretable methods (e.g., UVR-NMT) that contain gated fusion layers similar to ours, they perform sophisticated transformations on visual representation before fusion, which lowers the interpretability of the gating matrix. For example, in the gated fusion layer of UVR-NMT, we observe that the visual vector is order-of-magnitude smaller than the textual vector. As a result, interpreting gating weight is meaningless because visual vector has negligible influence on the fused vector.

4.4 Results

Table 1 shows the BLEU scores of these methods on the Multi30k dataset. From the table, we see that although we can replicate similar BLEU scores of Transformer-Base as reported in (Grönroos et al., 2018b; Ive et al., 2019), these scores (Row 1) are significantly outperformed by Transformer-Small and Transformer-Tiny, which have fewer parameters. This shows that Transformer-Base could overfit the Multi30k dataset. Transformer-Tiny, whose number of parameters is about 20 times smaller than that of Transformer-Base, is more robust and efficient in our test cases. We therefore use it as the base model for all our MMT systems in the following discussion.

Based on the Transformer-tiny model, both our proposed models (Gated Fusion and RMMT) and baseline MMT models (Doubly-ATT, Imagination and UVR-NMT) significantly outperform the

#	Model	#Params	En→De			#Params	En→Fr		
			Test2016	Test2017	MSCOCO		Test2016	Test2017	MSCOCO
<i>Text-only Transformer</i>									
1	Transformer-Base	49.1M	38.33	31.36	27.54	49.0M	60.60	53.16	42.83
2	Transformer-Small	36.5M	39.68	32.99	28.50	36.4M	61.31	53.85	44.03
3	Transformer-Tiny	2.6M	41.02	33.36	29.88	2.6M	61.80	53.46	44.52
<i>Existing MMT Systems</i>									
4	GMNMT [♠]	4.0M	39.8	32.2	28.7	-	60.9	53.9	-
5	DCCN [♠]	17.1M	39.7	31.0	26.7	16.9M	61.2	54.3	45.4
6	Doubly-ATT [♠]	3.2M	41.45	33.95	29.63	3.2M	61.99	53.72	45.16
7	Imagination [♠]	7.0M	41.31	32.89	29.90	6.9M	61.90	54.07	44.81
8	UVR-NMT [◇]	2.9M	40.79	32.16	29.02	2.9M	61.00	53.20	43.71
<i>Our MMT Systems</i>									
9	Gated Fusion [♠]	2.9M	41.96	33.59	29.04	2.8M	61.69	54.85	44.86
10	RMMT [◇]	2.9M	41.45	32.94	30.01	2.9M	62.12	54.39	44.52

Table 1: BLEU scores on Multi30k. Results in row 4 and 5 are taken from the original papers. ♠ indicates conventional MMT models, while ◇ refer to retrieval-based models. Without further specified, all our implementations are based on the *Tiny* configuration.

state-of-the-arts (GMNMT and DCCN) on En→De translation. However, the improvement of all these methods (Rows 4-10) over the base Transformer-Tiny model (Row 3) is very marginal. This shows that visual context might not be as important as we expected for translation, at least on datasets we explored.

We further evaluate all the methods on the ME-TEOR scores (see Appendix C). We also run experiments on the VaTex dataset (see Appendix B). Similar results are observed as Table 1. Although various MMT systems have been proposed recently, a well-tuned model that uses text only remain competitive. This motivates us to revisit the importance of visual context for translation in MMT models.

5 Model Analysis

Taking a closer look at the results given in the previous section, we are surprised by the observation that our models learn to *ignore* visual context when translating (Sec 5.1). This motivates us to revisit the contribution of visual context in MMT systems (Sec 5.2). Our adversarial evaluation shows that adding model regularization achieves comparable results as incorporating visual context. Finally, we discuss when visual context is needed (Sec 5.3) and how these findings could benefit future research.

5.1 Probe the need for visual context in MMT

To explore the need for visual context in our models, we focus on the interpretable component: the gated fusion layer (see Equation 3 and 5). Intuitively, a larger gating weight Λ_{ij} indicates the model learns to depend more on vi-

Multi30k	Gated Fusion	RMMT
En→De		
Test2016	4.5e-21	8.6e-13
Test2017	7.0e-17	4.0e-13
MSCOCO	9.7e-21	3.5e-14
En→Fr		
Test2016	1.6e-18	1.1e-11
Test2017	7.2e-15	5.0e-12
MSCOCO	2.3e-18	5.3e-13

Table 2: Micro-averaged gating weight $\bar{\Lambda}$ on Multi30k.

sual context to perform better translation. We quantify the degree to which visual context is used by the micro-averaged gating weight $\bar{\Lambda} = \sum_{m=1}^M \text{sum}(\Lambda^m) / (d \times V)$. Here M , V are the total number of sentences and words in the corpus, respectively. $\text{sum}(\cdot)$ add up all elements in a given matrix, and $\bar{\Lambda}$ is a scalar value ranges from 0 to 1. A larger $\bar{\Lambda}$ implies more usage of the visual context.

We first study models’ behavior after convergence. From Table 2, we observe that $\bar{\Lambda}$ is negligibly small, suggesting that both models learn to discard visual context. In other words, visual context may not be as important for translation as previously thought. Since $\bar{\Lambda}$ is insensitive to outliers (e.g., large gating weight at few dimensions), we further compute $p(\Lambda_{ij} > 1e-10)$: percentage of gating weight entries in Λ that are larger than $1e-10$. With no surprise, we find that on all test splits $p(\Lambda_{ij} > 1e-10)$ are always zero, which again shows that visual input is not used by the model in inference.

The Gated Fusion’s training process also shed




Figure 1: Training dynamic of Multi30k En→de and En→Fr translation, from Epoch 1.

some light on how the model accommodates the visual information during training. Figure 1 (a) and (b) shows how $\bar{\lambda}$ changes during training, from the first epoch. We find that, Gated Fusion starts with a relatively high $\bar{\lambda}$ (>0.5), but quickly decreases to ≈ 0.48 after the first epoch. As the training continues, $\bar{\lambda}$ gradually decreases to roughly zero. In the early stages, the model relies heavily on images, possibly because they could provide meaningful features extracted from a pre-trained ResNet-50 CNN, while the textual encoder is randomly initialized. Compared with text-only NMT, utilizing visual features lowers MMT models’ trust in the hidden representations generated from the textual encoders. As the training continues, the textual encoder learns to represent source text better and the importance of visual context gradually decreases. In the end, the textual encoder carries sufficient context for translation and supersedes the contributions from the visual features. Nevertheless, this doesn’t explain the superior performance of the multimodal systems (Table 1). We speculate that visual context is acting as regularization that helps model training in the early stages. We further explore this hypothesis in the next section.

5.2 Revisit need for visual context in MMT

In the previous section, we hypothesize that the gains of MMT systems come from some regularization effects. To verify our hypothesis, we conduct experiments based on two widely used regularization techniques: random noise injection (Bishop,

1995) and weight decay (Hanson and Pratt, 1989). The former simulates the effects of assumably uninformative visual representations and the later is a more principled way of regularization that does not get enough attention in the current hyperparameter tuning stage. Inspecting the results, we find that applying these regularization techniques achieves similar gains over the text-only baseline as incorporating multimodal information does.

For random noise injection, we keep all hyperparameters unchanged but replace visual features extracted using ResNet with randomly initialized vectors, which are noise drawn from a standard Gaussian distribution. A MMT model equipped with ResNet features is denoted as a *ResNet-based model*, while the same model with random initialization is denoted as a *noise-based model*. We run each experiment three times and report the averaged results. Note that values in parentheses indicate the performance gap between the ResNet-based model and its noise-based adversary.

Table 3 shows BLEU scores on the Multi30k dataset. Each column in the table corresponds to a test set “contest”. From the table, we observe that, among 18 (3 methods \times 3 test sets \times 2 tasks) contests with the Transformer model (row 1), noise-based models (rows 2-4) achieve better performance 13 times, while ResNet-based models win 14 cases. This shows that noise-based models perform comparably with ResNet-based models. A further comparison between noise-based models and ResNet-based models shows that they are compatible after 18 contests, in which the former wins 8 times and the latter wins 10 times.

We observe similar results when repeating above evaluation using METEOR (Table 9) and on VaTex (Table 7). These observations deduce that random noise could function as visual context. In MMT systems, adding random noise or visual context can help reduce overfitting (Bishop et al., 1995) when translating sentences in Multi30k, which are short and repetitive (Caglayan et al., 2019). Moreover, we find that the ℓ_2 norm of model weights in ResNet-based Gated Fusion and noise-based Gated Fusion are only 97.7% and 95.2% of that in Transformer on En→De, respectively. This further verifies our speculation that, as random noise injection (An, 1996), visual context can help weight smoothing and improve model generalization.

Further, we regularize the models with weight decay. We consider three models: the text-only Trans-

#	Model	En→De			En→Fr		
		Test2016	Test2017	MSCOCO	Test2016	Test2017	MSCOCO
1	Transformer	41.02	33.36	29.88	61.80	53.46	44.52
2	Doubly-ATT	41.53(+0.08)	33.90(-0.05)	29.76(+0.15)	61.85(-0.35)	54.61(+0.46)	44.85(-0.80)
3	Imagination	41.20(-0.11)	33.32(+0.42)	29.92(+0.02)	61.28(-0.62)	53.74(-0.33)	44.89(+0.08)
4	Gated Fusion	41.53(-0.45)	33.52(-0.07)	29.87(+0.83)	61.58(-0.11)	54.21(-0.64)	44.88(+0.02)

Table 3: BLEU scores on Multi30k with randomly initialized visual representation. Numbers in parentheses indicate the relative improvement/deterioration compared with the same model with ResNet feature initialization.

		BLEU	METEOR	$\bar{\Lambda}$
1	Transformer	11.39	35.53	-
2	+weight decay 0.1	11.66	35.95	-
w. ResNet features				
3	Gated Fusion	14.79	40.41	0.047
4	RMMT	16.67	43.62	0.011
w. random noise				
5	Gated Fusion	11.40	35.44	0.032
6	RMMT	12.08	37.60	0.010

Table 4: Adversarial evaluation with limited textual context on Multi30k En-De Test2016.

former, the representative existing MMT method Doubly-ATT, and our Gated Fusion method. Figure 2 and 3 (in Appendix C) show the BLEU and METEOR scores of these methods on En→De translation as weight decay rate changes, respectively. We see that the best results of the text-only Transformer model with fine-tuned weight decay are comparable or even better than that of the MMT models Doubly-ATT and Gated Fusion that utilize visual context. This again shows that visual context is not as useful as we expected and it essentially plays the role of regularization.

5.3 When is visual context needed in MMT

Despite the less importance of visual information we showed in previous sections, there also exist works that support its usefulness. For example, Caglayan et al. (2019) experimentally show that, with limited textual context (e.g., masking some input tokens), MMT models will utilize the visual input for translation. This further motivates us to investigate when visual context is needed in MMT models. We conduct experiment with a new masking strategy that does not need any entity linking annotations as in Caglayan et al. (2019). Specifically, we follow Tan and Bansal (2020) to collect a list of *visually grounded tokens*. A visually grounded token is the one that has more than 30 occurrences in the Multi30k dataset with stop words removed.

Masking all visually grounded tokens will affect around 45% of tokens in Multi30k.

Table 4 shows the adversarial study with visually grounded tokens masked. In particular, we select Transformer, Gated Fusion and RMMT as representative methods. From the table, we see that random noise injection (row 5,6) and weight decay (row 2) can only bring marginal improvement over the text-only Transformer model. However, ResNet-based models that utilize visual context significantly improve the translation results. For example, RMMT achieves almost 50% gain over the Transformer on the BLEU score. Moreover, both Gated Fusion and RMMT using ResNet features lead to a larger $\bar{\Lambda}$ value than that when textual context is sufficient as shown in Table 2. Those results further suggest that visual context is needed when textual context is insufficient. In addition to token masking, sentences with incorrect, ambiguous and gender-neutral words (Frank et al., 2018) might also need visual context to help translation. Therefore, to fully exert the power of MMT systems, we emphasize the need for a new MMT benchmark, in which visual context is deemed necessary to generate correct translation.

Interestingly, even with ResNet features, we observe a significant drop in both BLEU and METEOR scores compared with those in Table 1 and 8, similar to that reported in (Chowdhury and Elliott, 2019). The reason could be two-fold. On the one hand, there are many words that can not be visualized. For example, in Table 5 (a), although Gated Fusion can successfully identify the main objects in the image (“little boys pose with a puppy”), it fails to generate the more abstract concept “family picture”. On the other hand, when translating different words, it is difficult to capture correct regions in images. For example, in Table 5 (b), we see that Gated Fusion incorrectly generates the word *frauen* (*women*) because it captures the woman at the top-right corner of the image.




Figure 2: BLEU score curves on En→De translation with different weight decay rate.



<p>(a) </p>	<p>SRC: two <u>young boys</u> pose with a puppy for a <u>family picture</u> NMT: zwei braune hunde spielen mit einem <u>spielzeug</u> für einen <u>tennisball</u> <i>(two brown dogs play with a toy for a tennis ball)</i> MMT: zwei kleine jungen posieren mit einem welpen für ein foto <i>(two little boys pose with a puppy for a photo)</i> REF: zwei kleine jungen posieren mit einem welpen für eine familienfoto <i>(two little boys pose with a puppy for a family photo)</i></p>
<p>(b) </p>	<p>SRC: two men sitting in a <u>restaurant</u> NMT: zwei <u>kinder</u> spielen in einem <u>springbrunnen</u> <i>(two children are playing in a fountain)</i> MMT: zwei frauen sitzen in einem restaurant <i>(two women are sitting in a restaurant)</i> REF: zwei männer sitzen in einem restaurant <i>(two men are sitting in a restaurant)</i></p>

Table 5: Case studies under limited textual input. We use underline to denote masked tokens, and ~~strikethrough~~ (**bold**) font to denote ~~incorrect~~ (**correct**) lexical choices. We use Gated Fusion for analysis.

5.4 Discussion

Finally, we discuss how our findings might benefit future MMT research. First, a benchmark that requires more visual information than Multi30k to solve is desired. As shown in Section 5.2, sentences in Multi30k are rather simple and easy-to-understand. Thus textual context could provide sufficient information for correct translation, making visual modules relatively redundant in these systems. While the MSCOCO test set in Multi30k contains ambiguous verbs and encourages models to use image sources for disambiguation, we still lack a corresponding training set.

Second, our methods can serve as a verification tool to investigate whether visual grounding is needed in translation for a new benchmark.

Third, we find that visual feature selection is also critical for MMT’s performance. While most methods employ the attention mechanism to learn to attend relevant regions in an image, the shortage of annotated data could impair the attention module (see Table 5 (b)). Some recent efforts (Yin

et al., 2020; Lin et al., 2020; Caglayan et al., 2020) address the issue by feeding models with pre-extracted visual objects instead of the whole image. However, these methods are easily affected by the quality of the extracted objects. Therefore, a more effective end-to-end visual feature selection technique is needed, which can be further integrated into MMT systems to improve performance.

6 Conclusion

In this paper we devise two interpretable models that exhibit state-of-the-art performance on the widely adopted MMT datasets — Multi30k and the new video-based dataset — VaTex. Our analysis on the proposed models, as well as on other existing MMT systems, suggests that visual context helps MMT in the similar vein as regularization methods (e.g., weight decay), under sufficient textual context. Those empirical findings, however, should not be understood as us downplaying the importance existing datasets and models; we believe that sophisticated MMT models are necessary

for effective grounding of visual context into translation. Our goal, rather, is to (1) provide additional clarity on the *remaining* shortcomings of current dataset and stress the need for new datasets to move the field forward; (2) emphasise the importance of interpretability in MMT research.

7 Acknowledgement

Zhiyong Wu is partially supported by a research grant from the HKU-TCL Joint Research Centre for Artificial Intelligence.

References

- Guozhong An. 1996. The effects of adding noise during backpropagation training on a generalization performance. *Neural computation*, 8(3):643–674.
- Hasan Sait Arslan, Mark Fishel, and Gholamreza Anbarjafari. 2018. Doubly attentive transformer machine translation. *arXiv preprint arXiv:1807.11605*.
- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- Loïc Barrault, Fethi Bougares, Lucia Specia, Chiraag Lala, Desmond Elliott, and Stella Frank. 2018. Findings of the third shared task on multimodal machine translation. In *WMT18*.
- Chris M Bishop. 1995. Training with noise is equivalent to tikhonov regularization. *Neural computation*, 7(1):108–116.
- Christopher M Bishop et al. 1995. *Neural networks for pattern recognition*. Oxford university press.
- Ozan Caglayan, Julia Ive, Veneta Haralampieva, Pranava Madhyastha, Loïc Barrault, and Lucia Specia. 2020. Simultaneous machine translation with visual context. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2350–2361, Online. Association for Computational Linguistics.
- Ozan Caglayan, Pranava Madhyastha, Lucia Specia, and Loïc Barrault. 2019. Probing the need for visual context in multimodal machine translation. In *NAACL-HLT (1)*.
- Iacer Calixto, Desmond Elliott, and Stella Frank. 2016. Dcu-uva multimodal mt system report. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 634–638.
- Iacer Calixto and Qun Liu. 2017. Sentence-level multilingual multi-modal embedding for natural language processing. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 139–148, Varna, Bulgaria. INCOMA Ltd.
- Iacer Calixto, Qun Liu, and Nick Campbell. 2017. Doubly-attentive decoder for multi-modal neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1913–1924, Vancouver, Canada. Association for Computational Linguistics.
- Koel Dutta Chowdhury and Desmond Elliott. 2019. Understanding the effect of textual adversaries in multimodal machine translation. In *Proceedings of the Beyond Vision and LAnGuage: inTEgrating Real-world kNowledge (LANTERN)*, pages 35–40.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Jesse Dodge, Suchin Gururangan, Dallas Card, Roy Schwartz, and Noah A. Smith. 2019. Show your work: Improved reporting of experimental results. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2185–2194, Hong Kong, China. Association for Computational Linguistics.
- Desmond Elliott. 2018. Adversarial evaluation of multimodal machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2974–2978.
- 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.
- Desmond Elliott, Stella Frank, and Eva Hasler. 2015. Multilingual image description with neural sequence models. *arXiv preprint arXiv:1510.04709*.
- 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.
- 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.

- Stella Frank, Desmond Elliott, and Lucia Specia. 2018. Assessing multilingual multimodal image description: Studies of native speaker preferences and translator choices. *Natural Language Engineering*, 24(3):393–413.
- Stig-Arne Grönroos, Benoit Huet, Mikko Kurimo, Jorma Laaksonen, Bernard Merialdo, Phu Pham, Mats Sjöberg, Umut Sulubacak, Jörg Tiedemann, Raphael Troncy, and Raúl Vázquez. 2018a. The MeMAD submission to the WMT18 multimodal translation task. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 603–611, Belgium, Brussels. 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. 2018b. The memad submission to the wmt18 multimodal translation task. *arXiv*.
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor OK Li. 2018. Search engine guided neural machine translation. In *AAAI*, pages 5133–5140.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. Realm: Retrieval-augmented language model pre-training. *ICML*.
- Stephen José Hanson and Lorien Y Pratt. 1989. Comparing biases for minimal network construction with back-propagation. In *Advances in neural information processing systems*, pages 177–185.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Jindřich Helcl, Jindřich Libovický, and Dušan Variš. 2018. Cuni system for the wmt18 multimodal translation task. *arXiv preprint arXiv:1811.04697*.
- Jack Hessel and Lillian Lee. 2020. Does my multimodal model learn cross-modal interactions? it’s harder to tell than you might think! In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 861–877.
- Po-Yao Huang, Frederick Liu, Sz-Rung Shiang, Jean Oh, and Chris Dyer. 2016. Attention-based multimodal neural machine translation. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 639–645, Berlin, Germany. Association for Computational Linguistics.
- Julia Ive, Pranava Madhyastha, and Lucia Specia. 2019. Distilling translations with visual awareness. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6525–6538, Florence, Italy. Association for Computational Linguistics.
- Allan Jabri, Armand Joulin, and Laurens Van Der Maaten. 2016. Revisiting visual question answering baselines. In *European conference on computer vision*, pages 727–739. Springer.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2019. Generalization through memorization: Nearest neighbor language models. In *International Conference on Learning Representations*.
- Jan Kukačka, Vladimir Golkov, and Daniel Cremers. 2017. Regularization for deep learning: A taxonomy. *arXiv preprint arXiv:1710.10686*.
- Patrick Lewis, Ethan Perez, Aleksandara Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*.
- Jindřich Libovický and Jindřich Helcl. 2017. Attention strategies for multi-source sequence-to-sequence learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 196–202, Vancouver, Canada. Association for Computational Linguistics.
- Jindřich Libovický, Jindřich Helcl, and David Mareček. 2018. Input combination strategies for multi-source transformer decoder. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 253–260, Belgium, Brussels. Association for Computational Linguistics.
- Huan Lin, Fandong Meng, Jinsong Su, Yongjing Yin, Zhengyuan Yang, Yubin Ge, Jie Zhou, and Jiebo Luo. 2020. Dynamic context-guided capsule network for multimodal machine translation. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1320–1329.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649.
- Justin Salamon and Juan Pablo Bello. 2017. Deep convolutional neural networks and data augmentation for environmental sound classification. *IEEE Signal Processing Letters*, 24(3):279–283.

- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725.
- 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.
- Hao Tan and Mohit Bansal. 2020. Vokenization: Improving language understanding via contextualized, visually-grounded supervision. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2066–2080.
- 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*, pages 5998–6008.
- Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuanfang Wang, and William Yang Wang. 2019. VateX: A large-scale, high-quality multilingual dataset for video-and-language research. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Jason Weston, Emily Dinan, and Alexander Miller. 2018. Retrieve and refine: Improved sequence generation models for dialogue. In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*, pages 87–92.
- Felix Wu, Angela Fan, Alexei Baevski, Yann Dauphin, and Michael Auli. 2018. Pay less attention with lightweight and dynamic convolutions. In *International Conference on Learning Representations*.
- 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.
- Yongjing Yin, Fandong Meng, Jinsong Su, Chulun Zhou, Zhengyuan Yang, Jie Zhou, and Jiebo Luo. 2020. A novel graph-based multi-modal fusion encoder for neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3025–3035.
- Zhuosheng Zhang, Kehai Chen, Rui Wang, Masao Utiyama, Eiichiro Sumita, Zuchao Li, and Hai Zhao. 2020. Neural machine translation with universal visual representation. In *International Conference on Learning Representations*.
- Mingyang Zhou, Runxiang Cheng, Yong Jae Lee, and Zhou Yu. 2018. A visual attention grounding neural model for multimodal machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3643–3653.
- Jinhua Zhu, Yingce Xia, Lijun Wu, Di He, Tao Qin, Wengang Zhou, Houqiang Li, and Tiejun Liu. 2019. Incorporating bert into neural machine translation. In *International Conference on Learning Representations*.

A Training Settings

Table 6 shows the configuration of different model sizes.

Model component	Base	Small	Tiny
Number of encoder/decoder layers	6	6	4
Input/Output layer dimension	512	512	128
Inner feed-forward layer dimension	2048	1024	256
Number of attention heads	8	4	4

Table 6: Model configurations for *Base*, *Small*, and *Tiny*.

B Results on VaTex

VaTex is a video-based MMT corpus that contains 129,955 English-Chinese sentence pairs for training, 15,000 sentence pairs for validation, and 30,000 sentence pairs for testing. Each pair of sentences is associated with a video clip. Since the testing set is not publicly available, we use half of the validation set for validating and the other half for testing. We apply the byte pair encoding algorithm on the lower-cased English sentences and split Chinese sentences into sequences of characters, resulting in a vocabulary of 17,216 English tokens and 3,384 Chinese tokens. We use the video features provided along with the VaTex dataset, in which each video is represented as \mathbb{R}^{k*1024} , where k is the number of segments. Since some MMT systems take a “global” visual feature as input, we use 3D-Max-Pooling to extract the pooled representation \mathbb{R}^{1024} for each video.

Model	BLEU	METEOR
Transformer	35.82	59.02
+weight decay 0.1	36.32	59.38
+weight decay 0.01	36.07	59.14
+weight decay 0.001	35.92	59.22
Doubly-ATT	36.05 (35.46)	59.26 (58.84)
Imagination	36.25 (36.10)	59.26 (59.15)
Gated Fusion	36.06 (36.01)	59.34 (59.33)
RMMT	36.35 (36.43)	59.44 (59.57)

Table 7: Results on VaTex En-Zh translation. Numbers in parentheses are the performance of the same model with random noise initialization.

The results are shown in Table 7. We observe that although most MMT systems show improvement over the Transformer baseline, the gains are quite marginal. Indicating that although image-based MMT models can be directly applied to

video-based MMT, there is still room for improvement due to the challenge of video understanding. We also note that (a) regularize the text-only Transformer with weight decay demonstrates similar gains as injecting video information into the models; (b) replacing video features with random noise replicate comparable performance, which further supports our findings in Section 5.2.

C Results on METEOR

We also report our results based on METEOR (Banerjee and Lavie, 2005), which consistently demonstrates higher correlation with human judgments than BLEU does in independent evaluations such as in EMNLP WMT 2011³. From Table 8, we can see that on En-Fr translation, MMT systems demonstrate similar improvements over text-only baselines in both METEOR and BLEU (see Table 1). On En-De translation, however, MMT systems are mostly on-par with Transformer-tiny on METEOR and do not show consistent gains as BLEU. We hypothesis the reason being that En-De sets are created in a *image-blind* fashion, in which the crowd-sourcing workers produce translations without seeing the images (Frank et al., 2018). Such that source sentence can already provide sufficient context for translation. When creating the En-Fr corpus, the image-blind issue is fixed (Elliott et al., 2017), thus images are perceived as “needed” in the translation for whatever reason. Although BLEU is unable to elicit this difference, evaluation based on METEOR captured it and confirmed previous research. We also compute METEOR scores for our experiments that regularize models with random noise (see Table 9) and weight decay (see Figure 3). The results are consistent with those evaluated using BLEU and further complement our early findings.

D Results on IWSLT’14

We also evaluate the retrieval-based model RMMT on text-only corpus — IWSLT’14. The IWSLT’14 dataset contains 160k bilingual sentence pairs for En-De translation task. Following the common practice, we lowercase all words, split 7k sentence pairs from the training dataset for validation and concatenate *dev2010*, *dev2012*, *tst2010*, *tst2011*, *tst2012* as the test set. The number of BPE operations is set to 20,000. We use the *Small* configuration in all our experiments. The dropout and label

³<http://statmt.org/wmt11/papers.html>

#	Model	#Params	En→De			#Params	En→Fr		
			Test2016	Test2017	MSCOCO		Test2016	Test2017	MSCOCO
<i>Text-only Transformer</i>									
1	Transformer-Base	49.1M	65.92	60.02	54.73	49.0M	80.09	74.93	68.57
2	Transformer-Small	36.5M	66.01	60.80	55.95	36.4M	80.71	75.74	69.10
3	Transformer-Tiny	2.6M	68.22	62.05	56.64	2.6M	81.02	75.62	69.43
<i>Existing MMT Systems</i>									
4	GMNMT [♣]	4.0M	57.6	51.9	47.6	-	74.9	69.3	-
5	DCCN [♣]	17.1M	56.8	49.9	45.7	16.9M	76.4	70.3	65.0
6	Doubly-ATT [♣]	3.2M	68.04	61.83	56.21	3.2M	81.12	75.71	70.25
7	Imagination [♣]	7.0M	68.06	61.29	56.57	6.9M	81.2	76.03	70.35
<i>Our MMT Systems</i>									
9	Gated Fusion [♣]	2.9M	67.84	61.94	56.15	2.8M	80.97	76.34	70.51
10	RMMT [◇]	2.9M	67.97	61.71	56.33	2.9M	81.29	76.09	70.24

Table 8: METEOR scores on Multi30k. Results in row 4 and 5 are taken from the original papers. ♣ indicates conventional MMT models, while ◇ refers to retrieval-based models. Without further specification, all our implementations are based on the *Tiny* configuration.

#	Model	En→De			En→Fr		
		Test2016	Test2017	MSCOCO	Test2016	Test2017	MSCOCO
1	Transformer	68.22	62.05	56.64	81.02	75.62	69.43
2	Doubly-ATT	68.39(+0.35)	61.83(+0.0)	56.46(+0.25)	81.27(+0.15)	76.22(+0.51)	70.21(-0.04)
3	Imagination	67.93(-0.13)	61.84(+0.55)	56.49(-0.08)	80.75(-0.45)	76.57(+0.54)	69.88(-0.47)
4	Gated Fusion	68.25(+0.41)	61.5(-0.44)	55.93(-0.22)	81.22(+0.25)	76.01(-0.33)	70.33(-0.18)

Table 9: METEOR scores on Multi30k with randomly initialized visual representation. Numbers in parentheses indicate the relative improvement/deterioration compared with the original model with ResNet features.




Figure 3: METEOR score curves on En→De translation with different weight decay rate.

Model	BLEU
Transformer- <i>Small</i>	28.62
+weight decay 0.0001	29.14
RMMT- <i>Small</i>	29.03

Table 10: BLEU score on IWSLT’14 EN→DE translation.

smoothing rate are set to 0.3 and 0.1, respectively. Since there is no images associated with IWSLT, we follow (Zhang et al., 2020) and retrieve top-5 images from Multi30K corpus.

From Table 10, we see that Transformer without weight decay is marginally outperformed by RMMT, but achieves slightly higher BLEU scores when trained with a 0.0001 weight decay. Our discussion in Section 5.2 sheds light on why visual context is helpful on non-grounded low-resourced datasets like IWSLT’14 — for low-resourced dataset like IWSLT’14, injecting visual context help regularize model training and avoid overfitting.