

LIUM-CVC Submissions for WMT18 Multimodal Translation Task

Ozan Caglayan, Adrien Bardet,
Fethi Bougares, Loïc Barrault

LIUM, Le Mans University

FirstName.LastName@univ-lemans.fr

Kai Wang, Marc Masana, Luis Herranz and Joost van de Weijer

CVC, Universitat Autònoma de Barcelona

{kwang,mmasana,lherranz,joost}@cvc.uab.es

Abstract

This paper describes the multimodal Neural Machine Translation systems developed by LIUM and CVC for WMT18 Shared Task on Multimodal Translation. This year we propose several modifications to our previous multimodal attention architecture in order to better integrate convolutional features and refine them using encoder-side information. Our final constrained submissions ranked first for English→French and second for English→German language pairs among the constrained submissions according to the automatic evaluation metric METEOR.

1 Introduction

In this paper, we present the neural machine translation (NMT) and multimodal NMT (MMT) systems developed by LIUM and CVC for the third edition of the shared task. Several lines of work have been conducted since the introduction of the shared task on MMT in 2016 (Specia et al., 2016). The majority of last year submissions including ours (Caglayan et al., 2017a) were based on the integration of global visual features into various parts of the NMT architecture (Elliott et al., 2017). Apart from these, hierarchical multimodal attention (Helcl and Libovický, 2017) and multi-task learning (Elliott and Kádár, 2017) were also explored by the participants.

This year we decided to revisit the multimodal attention (Caglayan et al., 2016) since our previous observations about qualitative analysis of the visual attention was not satisfying. In order to improve the multimodal attention both qualitatively and quantitatively, we experiment with several refinements to it: first, we try to use different input image sizes prior to feature extraction and second we normalize the final convolutional feature maps to assess its impact on the final MMT performance. In terms of architecture, we propose to

refine the visual features by learning an encoder-guided early spatial attention. In overall, we find that normalizing feature maps is crucial for the multimodal attention to obtain a comparable performance to monomodal NMT while the impact of the input image size remains unclear. Finally, with the help of the refined attention, we obtain modest improvements in terms of BLEU (Papineni et al., 2002) and METEOR (Lavie and Agarwal, 2007).

The paper is organized as follows: data preprocessing, model details and training hyperparameters are detailed respectively in section 2 and section 3. The results based on automatic evaluation metrics are reported in section 4. Finally the paper ends with a conclusion in section 5.

2 Data

We use Multi30k (Elliott et al., 2016) dataset provided by the organizers which contains 29000, 1014, 1000 and 1000 English→{German,French} sentence pairs respectively for `train`, `dev`, `test2016` and `test2017`. A new training split of 30014 pairs is formed by concatenating the `train` and `dev` splits. Early-stopping is performed based on METEOR computed over the `test2016` set and the final model selection is done over `test2017`.

Punctuation normalization, lowercasing and aggressive hyphen splitting were applied to all sentences prior to training. A Byte Pair Encoding (BPE) model (Sennrich et al., 2016) with 10K merge operations is jointly learned on English-German and English-French resulting in vocabularies of 5189-7090 and 5830-6608 subwords respectively.

2.1 Visual Features

Since Multi30k images involve much more complex region-level relationships and scene compositions compared to ImageNet (Russakovsky et al.,




Figure 1: Filtered attention (FA): the convolutional feature maps are dynamically masked using an attention conditioned on the source sentence representation.

2015) object classification task, we explore different input image sizes to quantify its impact in the context of MMT since rescaling the input image has a direct effect on the size of the receptive fields of the CNN. After normalizing the images using ImageNet mean and standard deviation, we resize and crop the images to 224x224 and 448x448. Features are then extracted from the final convolutional layer (`res5c_relu`) of a pre-trained ResNet50 (He et al., 2016) CNN.¹ This led to feature maps $V \in \mathbb{R}^{2048 \times w \times w}$ where the spatial dimensionality w is 7 or 14.

2.1.1 Feature Normalization

We conjecture that transferring ReLU features from a CNN into a model that only makes use of bounded non-linearities like *sigmoid* and *tanh*, can saturate the non-linear neurons in the very early stages of training if their weights are not carefully initialized. Instead of tuning the initialization, we experiment with L_2 normalization over the channel dimension so that each feature vector ($\in \mathbb{R}^{2048}$) has an L_2 norm of 1.

3 Models

In this section we will describe our baseline NMT and multimodal NMT systems. All models use 128 dimensional embeddings and GRU (Cho et al., 2014) layers with 256 hidden states. Dropout (Srivastava et al., 2014) is applied over source embeddings x_s , encoder states H^{enc} and pre-softmax activations o_t . We also apply L_2 regularization with a factor of $1e-5$ on all parameters except biases. The parameters are initialized using the method proposed by He et al. (2015) and optimized with Adam (Kingma and Ba, 2014). The total gradient norm is clipped to 1 (Pascanu et al., 2013). We use batches of size 64 and an initial learning rate of $4e-4$. All systems are im-

¹We use torchvision for feature extraction.

plemented using the PyTorch version of *nmtpy*² (Caglayan et al., 2017b).

3.1 Baseline NMT

Let us denote the length of the source sentence $\{x_1, \dots, x_S\}$ and the target sentence $\{y_1, \dots, y_T\}$ by S and T respectively. The source sentence is first encoded with a 2-layer bidirectional GRU to obtain the set of hidden states:

$$H^{enc} \leftarrow Enc(\{x_1, \dots, x_S\}), H^{enc} \in \mathbb{R}^{S \times 512}$$

The decoder is a 2-layer conditional GRU (CGRU) (Sennrich et al., 2017) with tied embeddings (Press and Wolf, 2016). CGRU is a stacked 2-layer recurrence block with the attention mechanism in the middle. We use feed-forward attention (Bahdanau et al., 2014) which encapsulates a learnable layer. The first decoder (which is initialized with a zero vector) receives the previous target embeddings as inputs (equation 1). At each timestep of the decoding stage, the attention mechanism produces a context vector c_t^{txt} (equation 2) that becomes the input to the second GRU (equation 3). Finally, the probability over the target vocabulary is conditioned over a transformation of the final hidden state h_t^{dec2} (equation 4, 5).

$$h_t^{dec1} = DEC_1(y_{t-1}, h_{t-1}^{dec2}) \quad (1)$$

$$c_t^{txt} = ATT_{txt}(H^{enc}, h_t^{dec1}) \quad (2)$$

$$h_t^{dec2} = DEC_2(c_t^{txt}, h_t^{dec1}) \quad (3)$$

$$o_t = \tanh(\mathbf{W}_o h_t^{dec2} + b_o) \quad (4)$$

$$P(y_t) = softmax(\mathbf{W}_v o_t) \quad (5)$$

3.2 Multimodal Attention (MA)

Our baseline multimodal attention (MA) system (Caglayan et al., 2016) applies a spatial attention mechanism (Xu et al., 2015) over the visual features. At each timestep t of the decoding stage,

²github.com/lium-1st/nmtpytorch

a multimodal context vector c_t is computed and given as input to the second decoder (equation 3):

$$c_t = \mathbf{W}_f [c_t^{txt}; \mathbf{W}_{vis} c_t^{vis}] \quad (6)$$

$$c_t^{vis} = \text{ATT}_{vis}(V, h_t^{dec1}) \quad (7)$$

Previous analysis showed that the attention over the visual features is inconsistent and weak. We argue that this is because of the diluted relevant visual information, and the competition with the far more relevant source text information.

3.3 Filtered Attention (FA)

In order to enhance the visual attention, we propose an extension to the multimodal attention where the objective is to filter the convolutional feature maps using the last hidden state of the source language encoder (Figure 1). We conjecture that a learnable masking operation over the convolutional feature maps can help the decoder-side visual attention mechanism by filtering out regions irrelevant to translation and focus on the most important part of the visual input. The filtered convolutional feature map \tilde{V} is computed as follows:

$$\beta^{pre} = \text{ConvAtt}([\text{Tile}(h_S^{enc}); V]) \quad (8)$$

$$\tilde{V} = \beta^{pre} \odot V, \beta^{pre} \in \mathbb{R}^{1 \times w \times w} \quad (9)$$

ConvAtt block is inspired from previous works in visual question answering (VQA) (Yang et al., 2016; Kazemi and Elqursh, 2017). It basically computes a spatial attention distribution β^{pre} which we further use to mask the actual convolutional features V . The filtered \tilde{V} replaces V in the equation 7 instead of being pooled into a single visual embedding in contrast to VQA models.

EN→DE test2017	BLEU	METEOR
Baseline NMT	31.0 ± 0.3	52.1 ± 0.4
MA ₄₄₈	28.6 ± 0.8	50.1 ± 0.3
MA ₄₄₈ + L ₂ -norm	30.8 ± 0.5	52.0 ± 0.2

Table 1: Impact of L₂ normalization on the performance of multimodal attention.

4 Results

We train each model 4 times using different seeds and report mean and standard deviation for the final results using *multeval* (Clark et al., 2011)

Feature Normalization We can see from Table 1 that without L₂ normalization, multimodal attention is not able to reach the performance of baseline NMT. Applying the normalization consistently improves the results for all input sizes by around ~ 2 points in BLEU and METEOR. From now on, we only present systems trained with normalized features.

EN→DE test2017	BLEU	METEOR
MA ₂₂₄	30.6 ± 0.4	51.8 ± 0.2
MA ₄₄₈	30.8 ± 0.5	52.0 ± 0.2
FA ₂₂₄	31.5 ± 0.5	52.2 ± 0.5
FA ₄₄₈	31.6 ± 0.5	52.5 ± 0.4

Table 2: Impact of input image width on the performance of multimodal attention variants.

Image Size Although the impact of doubling the image width and height at the input seems marginal (Table 2), we switch to 448x448 images to benefit from the slight gains which are consistent across both attention variants.

4.1 Monomodal vs Multimodal Comparison

We first present the mean and standard deviation of BLEU and METEOR over 4 runs on the internal test set test2017 (Table 3). With the help of L₂ normalization, MA system almost reaches the monomodal system but fails to improve over it. On the contrary, the filtered attention (FA) mechanism improves over the baseline and produces hypotheses that are statistically different than the baseline with $p \leq 0.02$.

The improvements obtained for EN→DE language pair are not reflected on the EN→FR performance. One should note that the hyperparameters from EN→DE task were transferred to EN→FR without any other tuning.

The automatic evaluation of our final submissions (which are ensembles of 4 runs) on the official test set test2018 is presented in Table 5. In addition to our submissions, we also provide the best constrained and unconstrained systems³ in terms of METEOR. However, it should be noted that the submitted systems will be primarily evaluated using human direct assessment.

On EN→DE, our constrained FA system is comparable to the constrained UMONS submission. On EN→FR, our submission obtained the

³www.statmt.org/wmt18/multimodal-task.html

English→German	# Params	test2017 ($\mu \pm \sigma$)		
		BLEU	METEOR	TER
Baseline NMT	4.6M	31.0 \pm 0.3	52.1 \pm 0.4	51.2 \pm 0.5
Multimodal Attention (MA)	10.0M	30.8 \pm 0.5	52.0 \pm 0.2	51.1 \pm 0.7
Filtered Attention (FA)	11.3M	31.6 \pm 0.5	52.5 \pm 0.4	50.5 \pm 0.5

Table 3: EN→DE results: Filtered attention is statistically different than the NMT ($p \leq 0.02$).

English→French	# Params	test2017 ($\mu \pm \sigma$)		
		BLEU	METEOR	TER
Baseline NMT	4.6M	53.1 \pm 0.3	69.9 \pm 0.2	31.9 \pm 0.8
Multimodal Attention (MA)	10.0M	52.6 \pm 0.3	69.6 \pm 0.3	31.9 \pm 0.4
Filtered Attention (FA)	11.3M	52.8 \pm 0.2	69.6 \pm 0.1	31.9 \pm 0.1

Table 4: EN→FR results: multimodal systems are not able to improve over NMT in terms of automatic metrics.

EN→DE	BLEU	METEOR	TER
MeMAD†	38.5	56.6	44.5
UMONS*	31.1	51.6	53.4
LIUMCVC-FA*	31.4	51.4	52.1
LIUMCVC-NMT*	31.1	51.5	52.6
EN→FR			
CUNI†	40.4	60.7	40.7
LIUMCVC-FA*	39.5	59.9	41.7
LIUMCVC-NMT*	39.1	59.8	41.9

Table 5: Official test2018 results (†: Unconstrained, *: Constrained.)

highest automatic evaluation scores among the constrained submissions and is slightly worse than the unconstrained CUNI system.

5 Conclusion

MMT task consists of translating a source sentence into a target language with the help of an image representing the source sentence. The different level of relevance of both input modalities makes it a difficult task where the image should be used with parsimony. With the aim of improving the attention over visual input, we introduced a filtering technique to allow the network to ignore irrelevant parts of the image that should not be considered during decoding. This is done by using an attention-like mechanism between the source sentence and the convolutional feature maps. Results show that this mechanism significantly improves the results for English→German on one of the test sets. In the future, we plan to qualitatively analyze

the spatial attention and try to improve it further.

Acknowledgments

This work was supported by the French National Research Agency (ANR) through the CHIST-ERA M2CR project⁴, under the contract number ANR-15-CHR2-0006-01 and by MINECO through APCIN 2015 under the contract number PCIN-2015-251.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473.
- Ozan Caglayan, Walid Aransa, Adrien Bardet, Mercedes García-Martínez, Fethi Bougares, Loïc Barrault, Marc Masana, Luis Herranz, and Joost van de Weijer. 2017a. Lium-cvc submissions for wmt17 multimodal translation task. In *Proceedings of the Second Conference on Machine Translation, Volume 2: Shared Task Papers*, pages 432–439, Copenhagen, Denmark. Association for Computational Linguistics.
- Ozan Caglayan, Loïc Barrault, and Fethi Bougares. 2016. Multimodal attention for neural machine translation. *CoRR*, abs/1609.03976.
- Ozan Caglayan, Mercedes García-Martínez, Adrien Bardet, Walid Aransa, Fethi Bougares, and Loïc Barrault. 2017b. Nmtpy: A flexible toolkit for advanced neural machine translation systems. *Prague Bull. Math. Linguistics*, 109:15–28.

⁴<http://m2cr.univ-lemans.fr>

- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Jonathan H. Clark, Chris Dyer, Alon Lavie, and Noah A. Smith. 2011. Better hypothesis testing for statistical machine translation: Controlling for optimizer instability. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Short Papers - Volume 2*, HLT '11, pages 176–181, Stroudsburg, PA, USA. Association for Computational Linguistics.
- 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*, Copenhagen, Denmark.
- 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.
- Desmond Elliott and Ákos Kádár. 2017. Imagination improves multimodal translation. *CoRR*, abs/1705.04350.
- K. He, X. Zhang, S. Ren, and J. Sun. 2016. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Jindřich Helcl and Jindřich Libovický. 2017. Cuni system for the wmt17 multimodal translation task. In *Proceedings of the Second Conference on Machine Translation, Volume 2: Shared Task Papers*, pages 450–457, Copenhagen, Denmark. Association for Computational Linguistics.
- Vahid Kazemi and Ali Elqursh. 2017. Show, ask, attend, and answer: A strong baseline for visual question answering.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Alon Lavie and Abhaya Agarwal. 2007. Meteor: An automatic metric for mt evaluation with high levels of correlation with human judgments. In *Proceedings of the Second Workshop on Statistical Machine Translation, StatMT '07*, pages 228–231, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL '02*, pages 311–318, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28, ICML'13*, pages III–1310–III–1318. JMLR.org.
- Ofir Press and Lior Wolf. 2016. Using the output embedding to improve language models. *arXiv preprint arXiv:1608.05859*.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252.
- Rico Sennrich, Orhan Firat, Kyunghyun Cho, Alexandra Birch-Mayne, Barry Haddow, Julian Hitschler, Marcin Junczys-Dowmunt, Samuel Lubli, Antonio Miceli Barone, Jozef Mokry, and Maria Nadejde. 2017. Nematus: a toolkit for neural machine translation. In *Proceedings of the EACL 2017 Software Demonstrations*, pages 65–68. Association for Computational Linguistics (ACL).
- 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, Berlin, Germany. Association for Computational Linguistics.
- 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*, pages 543–553, Berlin, Germany. Association for Computational Linguistics.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958.

Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 2048–2057. JMLR Workshop and Conference Proceedings.

Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alexander J. Smola. 2016. Stacked attention networks for image question answering. In *CVPR*, pages 21–29. IEEE Computer Society.