

Semantic Label Smoothing for Sequence to Sequence Problems

Michal Lukasik, Himanshu Jain, Aditya Krishna Menon, Seungyeon Kim,
Srinadh Bhojanapalli, Felix Yu, Sanjiv Kumar

Google Research
mlukasik@google.com

Abstract

Label smoothing has been shown to be an effective regularization strategy in classification, that prevents overfitting and helps in label de-noising. However, extending such methods directly to seq2seq settings, such as Machine Translation, is challenging: the large target output space of such problems makes it intractable to apply label smoothing over all possible outputs. Most existing approaches for seq2seq settings either do token level smoothing, or smooth over sequences generated by randomly substituting tokens in the target sequence. Unlike these works, in this paper, we propose a technique that smooths over *well formed* relevant sequences that not only have sufficient n-gram overlap with the target sequence, but are also *semantically similar*. Our method shows a consistent and significant improvement over the state-of-the-art techniques on different datasets.

1 Introduction

Label smoothing is a regularization technique commonly used in deep learning (Szegedy et al., 2016; Chorowski and Jaitly, 2017; Vaswani et al., 2017; Zoph et al., 2018; Real et al., 2018; Huang et al., 2019), that improves calibration (Müller et al., 2019) and helps in label de-noising (Lukasik et al., 2020a). Here, one *smooths* labels by introducing a prior in the label space (often just a uniform distribution) in order to prevent overly confident predictions and achieve better model calibration, both of which lead to better generalization.

Given these benefits, it is natural to consider whether label smoothing can be applied to *sequence-to-sequence* (seq2seq) prediction tasks in Natural Language Processing. Here, inducing a label prior involves smoothing in *sequence space*. However, this is a challenging task because the output space is exponentially large for sequences,

unlike the label space in standard classification. Previous works approached this challenge either by smoothing over individual tokens of the target sequence, or by sampling a few nearby targets according to Hamming distance or BLEU score (Norouzi et al., 2016; Elbayad et al., 2018). These techniques however do *not* guarantee that the smoothed targets lie within the space of acceptable targets (i.e., the sampled new target may no longer be grammatically correct or even preserve semantic meaning).

In this work, we propose a label smoothing approach for seq2seq problems that overcomes this limitation. Given a large-scale corpus of valid sequences, our approach selects a subset of sequences that are not only semantically similar to the target sequence, but also *well formed*. We achieve this using a pre-trained model to find semantically similar sequences from the corpus, and then use BLEU scores to rerank the closest targets. We empirically show that this approach improves over competitive baselines on multiple machine translation tasks.

2 Related Works

Token-level smoothing A popular approach used in language tasks is so called token level smoothing, where for each position’s classification loss, a prior distribution over the entire vocabulary (uniformly or with unigram probability estimates) is used for regularization (Pereyra et al., 2017; Edunov et al., 2017). This is similar to the classical label smoothing (e.g. (Szegedy et al., 2016)), as it smooths each token label independent of their context and position in the sequence. Such an approach is thus unlikely to result in semantically related targets.

Sequence-level smoothing Norouzi et al. (2016) augment the loss with a term rewarding predictions of sampled sequences. The sampling of sequences is based on their edit distance or Hamming

distance to the target. This method thus smooths the loss over to similar sequences (in terms of the edit distance) with smoothing rewards. Elbayad et al. (2018) employ a similar technique, but with a new reward function based on BLEU (Papineni et al., 2002) or CIDEr (Vedantam et al., 2015) score. Specifically, Elbayad et al. (2018) generate a smoothed version of the target sequence, wherein one replaces a token with a random token (with up-sampling of rare words). Such newly generated sequences were given a partial reward based on the cosine similarity between the two tokens in a pre-trained word-embedding space. This differs from our approach because this context-independent perturbation is limited to generating the same structure for the new sequence as that of the original sequence.

Zheng et al. (2018), on the other hand, constructed grammatically correct and meaning preserving sequences. However, unlike our work, their approach relies on having multiple references (target sequences per input sequence) and might not be able to generate sequences where common words or synonyms do not appear in the same order, which is a strong limitation, precluding an augmentation like: *Yesterday, he scored a 94 on his final* (original sequence), *He had 94 points in the final test yesterday* (augmented sequence).

More broadly, an important shortcoming of such approaches is that sequences deemed close can actually lack important properties such as preserving the meaning of the original sequence. In particular, swapping even a single token in a sequence may cause a drastic shift in its meaning (e.g., turning a factually correct text into a false one) even though being close in the Hamming space. We address this shortcoming by restricting augmented target sequences to the training set, and selecting sequences based on similarity obtained from a pre-trained model.

Unlike other approaches, (Bengio et al., 2015) proposed a scheduled sampling technique that does not depend on any external data source. Instead, it utilizes the self-generated sequences from the current model. Both our approach and the scheduled sampling technique bear similarity in that they aim at improving model generalization, by either providing semantically similar candidates (ours) or self-generated sequences (theirs). Indeed, these two approaches could complement each other by providing various ways of related but not exact

targets.

Hard negative mining Our work is also related to hard negative mining approaches that select a subset of confusing negatives for each input (Mikolov et al., 2013; Reddi et al., 2019; Guo et al., 2018). Different from the above, we add a soft objective function over the sampled (relevant) target sequences, rather than treating them as negatives in the classification sense.

3 Method

Sequence-to-sequence (seq2seq) learning involves learning a mapping from an input sequence \mathbf{x} (e.g., a sentence in English) to an output sequence \mathbf{y} (e.g., a sentence in French). Canonical applications include machine translation and question answering.

Formally, let \mathcal{X} denote the space of input sequences (e.g., all possible English sentences), and \mathcal{Y} the space of output sequences (e.g., all possible French sentences). We represent by $\mathbf{x} = [x_1, x_2, \dots, x_N]$ an input sequence consisting of N tokens, and similarly $\mathbf{y} = [y_1, y_2, \dots, y_{N'}]$ an output sequence with N' tokens. Our goal is to learn a function $f: \mathcal{X} \rightarrow \mathcal{Y}$ that, given an input sequence, generates a suitable target sequence.

To achieve this goal, we have a training set $\mathcal{S} \subseteq (\mathcal{X} \times \mathcal{Y})^n$ comprising pairs of input and output sequences. We then seek to minimise the objective

$$L(\theta) = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{S}} -\log p_{\theta}(\mathbf{y} | \mathbf{x}), \quad (1)$$

where $p_{\theta}(\cdot | \mathbf{x}; \theta)$ is a parametrized distribution over all possible output sequences. Given such a distribution, we choose $f(\mathbf{x}) = \operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}} p_{\theta}(\mathbf{y} | \mathbf{x})$. Observe that one may implement (1) via a token-level decomposition,

$$L(\theta) = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{S}} \sum_{i=1}^N -\log p_{\theta}(y_i | \mathbf{x}, y_1, \dots, y_{i-1}).$$

This may be understood as a maximum likelihood objective, or equivalently the *cross-entropy* between $p_{\theta}(\cdot | \mathbf{x}; \theta)$ and a one-hot distribution concentrated on \mathbf{y} .

Label smoothing meets seq2seq. Intuitively, the cross-entropy objective encourages the model to score the *observed* sequence \mathbf{y} higher than any “competing” sequence $\mathbf{y}' \neq \mathbf{y}$. While this is a sensible goal, one limitation observed from classification settings is that the loss may lead to models

that are overly confident in their predictions, which can hamper generalisation (Guo et al., 2017).

Label smoothing (Szegedy et al., 2016; Pereyra et al., 2017; Müller et al., 2019) is a simple means of correcting this in classification settings. Smoothing involves simply adding a small reward to all possible *incorrect* labels, i.e., mixing the standard one-hot label with a uniform distribution over all labels. This regularizes the training and generally leads to better predictive performance as well as probabilistic calibration (Müller et al., 2019).

Given the success of label smoothing in classification settings, it is natural to explore its value in seq2seq problems. However, standard label smoothing is clearly inadmissible: it would require smoothing over all possible outputs $\mathbf{y}' \in \mathcal{Y}$, which is typically an intractably large set. Nonetheless, we may follow the basic intuition of smoothing by adding a subset of *related* targets to the observed sequence \mathbf{y} , yielding a *smoothed loss*

$$-\log p_{\theta}(\mathbf{y} | \mathbf{x}) + \frac{\alpha}{|\mathcal{R}(\mathbf{y})|} \cdot \sum_{\mathbf{y}' \in \mathcal{R}(\mathbf{y})} -\log p_{\theta}(\mathbf{y}' | \mathbf{x}). \quad (2)$$

Here, $\mathcal{R}(\mathbf{y})$ is a set of *related sequences* that are similar to the ground truth \mathbf{y} , and $\alpha > 0$ is a tuning parameter that controls how much we rely on the observed versus related sequences.

The quality of $\mathcal{R}(\mathbf{y})$ is important for our task. Ideally, we would like an $\mathcal{R}(\mathbf{y})$ that: (i) is efficient to compute, and (ii) comprises sequences which meaningfully align with \mathbf{x} (e.g., are plausible alternate translations). We now assess several options for constructing $\mathcal{R}(\mathbf{y})$ in light of the above.

Random sequences. One simple choice is to choose a random subset of output sequences from the training set. In the common setting where f is learned by minibatch SGD on randomly drawn minibatches $\mathcal{B} = \{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}$, one may simply pick $\mathcal{R}(\mathbf{y})$ to be all output sequences in \mathcal{B} .

Such random sequences contain general target language understanding (e.g., French grammar for an English to French translation task). However, these sequences are unlikely to have any semantic correlation with the true label.

Token-level smoothing. To ensure greater semantic correlation between the selected sequences and the original \mathbf{y} , one idea is to perform *token-level smoothing*. For example, Vaswani et al. (2017) proposed to smooth uniformly over all tokens from the vocabulary. Elbayad et al. (2018) proposed to construct sequences $\mathbf{y}' = [y'_1, y'_2, \dots, y'_{N'}]$ where for a randomly selected subset of tokens $j \in [N']$,

Algorithm 1 Sampling of related sequences.

Input: example (\mathbf{x}, \mathbf{y}) ; sequences \mathcal{Y}_{ref}

Output: related sequences $\mathcal{R}(\mathbf{y})$

- 1: Embed reference sequences, e.g., using BERT
 - 2: $\mathcal{N}(\mathbf{y}) \leftarrow k$ closest sequences to \mathbf{y} from \mathcal{Y}_{ref} in the embedding space.
 - 3: Sort elements of $\mathcal{N}(\mathbf{y})$ by BLEU score to \mathbf{y} .
 - 4: $\mathcal{R}(\mathbf{y}) \leftarrow$ top k' elements from \mathcal{N} .
-

Orig:	Yesterday, he scored a 94 on his final.
1st:	He had 94 points in the final test yesterday.
2nd:	But the child just scored 9 points on the Apgar test.
Orig:	Exchange of experience and good practices.
1st:	Exchange of best practices.
2nd:	Exchange of information and best practices.
Orig:	Nothing else I can do?
1st:	Is there anything else I can do for you, sir?
2nd:	Can I do something for you?

Table 1: English translations of top two augmentations from BERT+BLEU4 on examples from EN-CS.

y'_i is some *related* token in the minibatch; for other tokens, $y'_i = y_i$. These related tokens are chosen so as to maximise the BLEU score between \mathbf{y} and \mathbf{y}' .

While this approach increases the semantic similarity to \mathbf{y} , operating on a token level is limiting. For example, one may change the meaning of a factual sentence by changing even a few words. Further, operating at a per-token level limits the *diversity* of $\mathcal{R}(\mathbf{y})$, since, e.g., all sequences have the same number of tokens and structure as \mathbf{y} .

Proposal: semantic smoothing. To overcome the limitations of token-level smoothing, we would ideally like to directly smooth over related *sequences*. Our basic idea is to seek sequences

$$\mathcal{R}(\mathbf{y}) = \{\mathbf{y}' : s_{\text{sem}}(\mathbf{y}, \mathbf{y}') \wedge s_{\text{bleu}}(\mathbf{y}, \mathbf{y}') > 1 - \epsilon\},$$

where s_{sem} is a score of semantic similarity, and s_{bleu} is the BLEU score. Intuitively, our relevant sequences comprise those that are both semantically similar to \mathbf{y} , and have sufficient unigram overlap.

A key challenge is efficiently identifying semantically similar sequences to \mathbf{y} . To achieve this in a tractable manner, we propose the following procedure (see Algorithm 1). First, we assume the existence of an *embedding space* for output sequences. For example, this could be the result of BERT (Devlin et al., 2019), which embeds each sequence into a fixed vector representation. Given such an embedding space and a corpus \mathcal{Y}_{ref} of reference sequences, we may now efficiently compute

the neighbors of y , $\mathcal{N}(y)$, comprising the top- k closest sequences in \mathcal{Y}_{ref} for the given y (Indyk and Motwani, 1998).¹

The elements of $\mathcal{N}(y)$ can be expected to have high semantic similarity with y , which is desirable. However, such sequences may *not* meaningfully align with the original input x (e.g., may not be sufficiently close translations). To account for this, we prune the elements from $\mathcal{N}(y)$ based on the BLEU score. Intuitively, this pruning retains sequences that are both semantically similar *and* have non-trivial token overlap with y .

We use \mathcal{Y}_{ref} as all output sequences in the training set. In practice, one may however use any set of sequences that are valid for the domain in question. We find $k = 100$ closest sequences in this space, and smooth over $k' = 5$ pruned sequences with the highest BLEU score to y . In Table 1 we show example augmentations. Notice both the diversity of augmentations, as well as relatedness to the original targets.

4 Experiments

Setup. We use the Transformer model for our experiments, and follow the experimental setup and hyperparameters from Vaswani et al. (2017). We experiment on three popular machine translation tasks: English-German (EN-DE), English-Czech (EN-CS) and English-French (EN-FR), using the WMT training datasets, and on the tensor2tensor framework (Vaswani et al., 2018).² We evaluate on the Newstest 2015 for EN-DE and EN-CS, and WMT 2014 for EN-FR.

Baselines. We use the seq2seq model results by Vaswani et al. (2017) as a baseline. We compare our approach with the following alternate smoothing methods: i) smoothing is done over all possible tokens from the vocabulary at each next token prediction (Szegedy et al., 2016), ii) smoothing is conducted over random targets from within batch (Guo et al., 2018), and iii) smoothing is done over artificially generated targets that are close to the actual target sequence according to BLEU score (Elbayad et al., 2018). For all these methods we experiment with values of α in $\{0.1, 0.001, 0.0001, 0.00001\}$, and report the best results in each case. For the

(Elbayad et al., 2018) baseline, we follow the reported best performing variant, randomly swapping tokens with others from the target sequence.

Main results. In Table 2 we report results from our method (BERT+BLEU) and the different state-of-the-art methods mentioned above. Our most direct comparison is against (Elbayad et al., 2018), as both the methods smooth over sequences that have high BLEU score. However, instead of generating sequences by randomly replacing tokens, we retrieve them from a corpus of well formed text sequences. In particular, we use BERT-base multilingual model to embed all the training target sequences into 768 dim fixed vector representation (corresponding to CLS token) and then identify top-100 nearest neighbors for each of the target sequence. Consequently, our method outperforms (Elbayad et al., 2018) by a large margin on all three benchmarks. This demonstrates the importance of smoothing over sequences that not only have significant n-gram overlap with the ground truth target sequence but are also well formed and are semantically similar to the ground truth. In Table 3 we report the comparison between our model and the strongest baseline on EN-CS across multiple metrics, confirming the improvement we report in Table 2 for BLEU score.

Ablating BLEU pruning. Table 4 reveals it is useful to use a sufficiently restrictive criterion in BLEU pruning; however, excess pruning (BLEU5) is harmful. Thus, we seek to retrieve semantically related targets which do not necessarily have highest scoring n-gram overlap to the original target. This is intuitive: enforcing too high n-gram overlap may cause all augmented targets to be too lexically similar, limiting the benefit of seeing new targets in training. We also experimented with not reranking neighbors using BLEU pruning, which resulted in no improvement over the baseline. In other words, it was essential to use this kind of postprocessing for obtaining improvements.

Ablating the number of neighbors. We experimented with how the number of neighbors influences the results. For EN-CS, we obtained the following BLEU4 scores correspondingly for 10, 5 and 3 neighbors: 21.86, 22.82, 22.23. Overall, we find that too few or too many neighbors harm the performance compared to the 5 neighbors we used in other experiments. At the same time, the time complexity increases linearly as number of neighbors increases.

¹Alternatively, one could consider selecting highest scoring augmentations based on a pre-trained seq2seq model. However, the resulting quadratic computational complexity renders such an approach impractical.

²Data available at <https://tensorflow.github.io/tensor2tensor/>.

Method	α	EN-DE	EN-CS	EN-FR
Base setup (Vaswani et al., 2017)	—	28.03	21.19	39.66
Token LS (Vaswani et al., 2017; Szegedy et al., 2016)	0.1	28.72	21.47	39.87
Within batch sequence LS (Guo et al., 2018)	0.001	28.81	21.26	39.21
Sampled augmentations BLEU4 (Elbayad et al., 2018)	0.01	29.19	20.94	40.19
BERT+BLEU4	0.1	29.99	22.82	39.84
BERT+BLEU4	0.01	29.51	22.30	40.82

Table 2: BLEU4 evaluation scores on translation tasks from different label smoothing methods. We ran a bootstrap test (Koehn, 2004) for estimating the significance of improvement over the strongest baseline and found that on all three datasets the improvement is statistically significant, $p < 0.05$.

	BLEU3	BLEU4	BLEU5	METEOR	ROUGE	CIDER
(Elbayad et al., 2018)	27.9	20.94	15.93	24.92	50.98	211.49
BERT+BLEU4	29.8	22.82	17.73	26.03	52.29	228.26

Table 3: Comparison of our model against the strongest baseline (Elbayad et al., 2018) as reported in Table 2 on EN-CS across multiple metrics.

	BLEU3	BLEU4	BLEU5
BERT+BLEU3	29.12	22.03	16.89
BERT+BLEU4	29.80	22.82	17.73
BERT+BLEU5	29.41	22.38	17.26

Table 4: Results on EN-CS from targets smoothing with varying n-gram overlap enforced for the final selection of top 5 augmented targets. Enforcing higher overlap to the original target worsens the performance.

5 Conclusion

We propose a novel label smoothing approach for sequence to sequence problems that selects a subset of sequences that are not only semantically similar to the target sequences, but are also *well formed*. We achieve this by using a pre-trained model to find semantically similar sequences from the training corpus, and then we use BLEU score to rerank the closest targets. Our method shows a consistent and significant improvement over state-of-the-art techniques across different datasets.

In future work, we plan to apply our semantic label smoothing technique to various sequence to sequence problems, including Text Summarization (Zhang et al., 2019) and Text Segmentation (Lukasik et al., 2020b). We also plan to study the relation between pretraining and data augmentation techniques.

References

Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for se-

quence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 1171–1179.

Jan Chorowski and Navdeep Jaitly. 2017. Towards better decoding and language model integration in sequence to sequence models. In *Interspeech 2017, 18th Annual Conference of the International Speech Communication Association, Stockholm, Sweden, August 20-24, 2017*, pages 523–527.

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, Minneapolis, Minnesota. Association for Computational Linguistics.

Sergey Edunov, Myle Ott, Michael Auli, David Grangier, and Marc’Aurelio Ranzato. 2017. Classical structured prediction losses for sequence to sequence learning. *arXiv preprint arXiv:1711.04956*.

Maha Elbayad, Laurent Besacier, and Jakob Verbeek. 2018. **Token-level and sequence-level loss smoothing for RNN language models**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2094–2103, Melbourne, Australia. Association for Computational Linguistics.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 1321–1330.

- Mandy Guo, Qinlan Shen, Yinfei Yang, Heming Ge, Daniel Cer, Gustavo Hernandez Abrego, Keith Stevens, Noah Constant, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. [Effective parallel corpus mining using bilingual sentence embeddings](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 165–176, Brussels, Belgium. Association for Computational Linguistics.
- Yanping Huang, Youlong Cheng, Ankur Bapna, Orhan Firat, Dehao Chen, Mia Chen, HyounJoong Lee, Jiquan Ngiam, Quoc V Le, Yonghui Wu, and zhifeng Chen. 2019. Gpipe: Efficient training of giant neural networks using pipeline parallelism. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Álché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 103–112. Curran Associates, Inc.
- Piotr Indyk and Rajeev Motwani. 1998. [Approximate nearest neighbors: Towards removing the curse of dimensionality](#). In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing, STOC '98*, page 604–613, New York, NY, USA. Association for Computing Machinery.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Michal Lukasik, Srinadh Bhojanapalli, Aditya Krishna Menon, and Sanjiv Kumar. 2020a. Does label smoothing mitigate label noise? *arXiv preprint arXiv:2003.02819*.
- Michal Lukasik, Boris Dachev, Gonçalo Simões, and Kishore Papineni. 2020b. Text segmentation by cross segment attention. *arXiv preprint arXiv:2004.14535*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Rafael Müller, Simon Kornblith, and Geoffrey E. Hinton. 2019. When does label smoothing help? In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 4696–4705.
- Mohammad Norouzi, Samy Bengio, zhifeng Chen, Navdeep Jaitly, Mike Schuster, Yonghui Wu, and Dale Schuurmans. 2016. [Reward augmented maximum likelihood for neural structured prediction](#). In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 1723–1731. Curran Associates, Inc.
- 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*, pages 311–318. Association for Computational Linguistics.
- Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. 2017. Regularizing neural networks by penalizing confident output distributions. In *International Conference on Learning Representations Workshop*.
- Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V. Le. 2018. [Regularized evolution for image classifier architecture search](#). *arXiv preprint arXiv:1802.01548*.
- Sashank J Reddi, Satyen Kale, Felix Yu, Daniel Holtmann-Rice, Jiecao Chen, and Sanjiv Kumar. 2019. Stochastic negative mining for learning with large output spaces. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1940–1949.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2818–2826.
- Ashish Vaswani, Samy Bengio, Eugene Brevdo, François Chollet, Aidan Gomez, Stephan Gouws, Llion Jones, Łukasz Kaiser, Nal Kalchbrenner, Niki Parmar, Ryan Sepassi, Noam Shazeer, and Jakob Uszkoreit. 2018. [Tensor2Tensor for neural machine translation](#). In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Papers)*, pages 193–199, Boston, MA. Association for Machine Translation in the Americas.
- 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 *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, pages 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575.
- Xingxing Zhang, Furu Wei, and Ming Zhou. 2019. [HiBERT: Document level pre-training of hierarchical bidirectional transformers for document summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5059–5069.

- Renjie Zheng, Mingbo Ma, and Liang Huang. 2018. Multi-reference training with pseudo-references for neural translation and text generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3188–3197.
- B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le. 2018. Learning transferable architectures for scalable image recognition. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8697–8710.