

# Efficient Estimation of Influence of a Training Instance

Sosuke Kobayashi<sup>1,2</sup> Sho Yokoi<sup>1,3</sup> Jun Suzuki<sup>1,3</sup> Kentaro Inui<sup>1,3</sup>

Tohoku University<sup>1</sup> Preferred Networks, Inc.<sup>2</sup> RIKEN<sup>3</sup>

sosk@preferred.jp

{yokoi, jun.suzuki, inui}@ecei.tohoku.ac.jp

## Abstract

Understanding the influence of a training instance on a neural network model leads to improving interpretability. However, it is difficult and inefficient to evaluate the influence, which shows how a model’s prediction would be changed if a training instance were not used. In this paper, we propose an efficient method for estimating the influence. Our method is inspired by dropout, which zero-masks a sub-network and prevents the sub-network from learning each training instance. By switching between dropout masks, we can use sub-networks that learned or did not learn each training instance and estimate its influence. Through experiments with BERT and VGGNet on classification datasets, we demonstrate that the proposed method can capture training influences, enhance the interpretability of error predictions, and cleanse the training dataset for improving generalization.

## 1 Introduction

What is the influence of a training instance on a machine learning model? This question has attracted the attention of the community (Cook, 1977; Koh and Liang, 2017; Zhang et al., 2018; Hara et al., 2019). Evaluating the influence of a training instance leads to more interpretable models and other applications like data cleansing.

A simple evaluation is by comparing a model with another similarly trained model, whose training does not include the instance of interest. This method, however, requires computational costs of time and storage depending on the number of instances, which indicates the extreme difficulty (Table 1). While computationally cheaper estimation methods have been proposed (Koh and Liang, 2017; Hara et al., 2019), they still have computational difficulties or restrictions of model choices. The contribution of this work is to propose an estimation

method, which (i) is computationally more efficient while (ii) useful for applications (iii) without significant sacrifice of model performance.

We propose a trick for enabling a neural network without restrictions to estimate the influence, which we refer to as *turn-over dropout*. This method is computationally efficient as it requires only running two forward computations after training a single model on the entire training dataset. In addition to the efficiency, we demonstrated that it enabled BERT (Devlin et al., 2019) and VGGNet (Simonyan and Zisserman, 2015) to analyze the influences of training through various experiments, including example-based interpretation of error predictions and data cleansing to improve the accuracy on a test set with a distributional shift.

## 2 Influence of a Training Instance

### 2.1 Problem Setup

We present preliminaries on the problem setup. In this paper, we deal with the *influence* of training with an instance on prediction with another one, which has been studied in Koh and Liang (2017), Hara et al. (2019) and so on. Let  $z := (x, y)$  be an instance and represent a pair of input  $x \in X$  and its output  $y \in Y$ , and let  $D := \{z_i\}_{i=1}^N$  be a training dataset. By using an optimization method with  $D$ , we aim to find a model  $f_D: X \rightarrow Y$ . Denoting the loss function by  $L(f, z)$ , the learning problem is obtaining  $\hat{f}_D = \operatorname{argmin}_f \mathbb{E}_{z_i \in D} L(f, z_i)$ .

The *influence*,  $I(z_{\text{target}}, z_i; D)$ , is a quantitative benefit from  $z_i$  to prediction of  $z_{\text{target}}$ . Let  $f_{D \setminus \{z_i\}}$  to be a model trained on the dataset  $D$  excluding  $z_i$ , the influence is defined as

$$\begin{aligned} I(z_{\text{target}}, z_i; D) \\ := L(f_{D \setminus \{z_i\}}, z_{\text{target}}) - L(f_D, z_{\text{target}}). \end{aligned} \quad (1)$$

Intuitively, the larger this value, the more strongly a training instance  $z_i$  contributes to reduce the loss of

Method	Training	Storage	Estimation
Re-train	$O( D ^2)$	$O( \theta  D )$	$O(F D )$
Hara+	$O( D )$	$O( \theta T)$	$O(F D ) + (F+F')TB$
Koh+	$O( D )$	$O( \theta )$	$O(F D ) + (F+F')rtb$
<b>Ours</b>	$O( D )$	$O( \theta )$	$O(F D )$

Table 1: Comparison of computational complexity for estimating the influence of all instance on another instance, with Hara et al. (2019) and Koh and Liang (2017), where  $|\theta|$  is the number of parameters,  $F$  is a forward/backward computation,  $F'$  is a double backward computation,  $T$  is the training steps,  $B$  is a training minibatch size,  $b$  is a minibatch size for stabilizing approximation,  $rt$  are the hyper-parameters; typically  $rt \approx |D|$ . See the references in detail.

prediction on another instance  $z_{\text{target}}$ . The instance of interest  $z_{\text{target}}$  is typically an instance in a test or validation dataset.

## 2.2 Related Methods

Computing the influence in Equation (1) by re-training two models for each instance is computationally expensive, and several estimation methods are proposed. Koh and Liang (2017) proposed an estimation method that assumed a strongly convex loss function and a global optimal solution<sup>1</sup>. While the method is used even with neural models (Koh and Liang, 2017; Han et al., 2020), which do not satisfy the assumption, it still requires high computational cost. Hara et al. (2019) proposed a method without these restrictions; however, it consumes large disk storage and computation time that depend on the number of optimization steps. Our proposed method is much more efficient, as shown in Table 1. For example, in a case where Koh and Liang (2017)’s method took 10 minutes to estimate the influences of 10,000 training instances on another instance with BERT (Han et al., 2020), our method only required 35 seconds<sup>2</sup>. This efficiency will expand the scope of applications of computing influence. For example, it would enable real-time interpretation of model predictions for users of the machine learning models.

<sup>1</sup> Strictly speaking, Koh and Liang (2017) studied a similar but different value from  $I$  in Equation (1). Briefly, the formulation in Koh and Liang (2017) considers convex models with the optimal parameters for  $f_{D \setminus \{z_i\}}$  and  $f_D$ . The definition in Hara et al. (2019) did not have such conditions and treated the broader problem. We follow Hara et al. (2019); therefore, the definition in Equation (1) allows any  $f_D$  and  $f_{D \setminus \{z_i\}}$ , as long as they have the same initial parameters and optimization procedures using the same mini-batches except for  $z_i$ .

<sup>2</sup> For the details, see Appendix B.

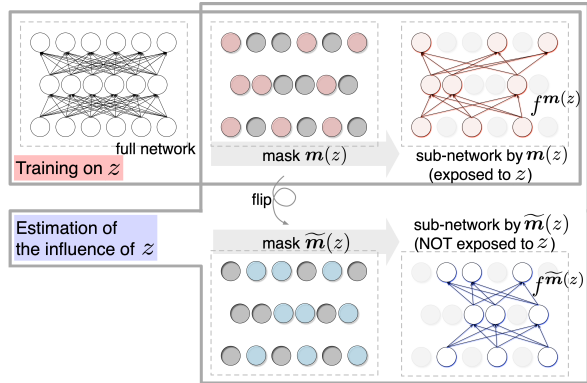


Figure 1: Dropout generates a sub-network for each training instance  $z$ , and updates its parameters (red; top) only. By contrast, the (blue; bottom) sub-network is not influenced by  $z$ . Our estimation uses the difference between the two sub-networks.

## 3 Proposed Method

### 3.1 Background: Dropout

Dropout (Hinton et al., 2012; Srivastava et al., 2014) is a popular regularization methods for deep neural networks. During training,  $d$ -dimensional random mask vector  $m$ , where  $d$  refers to the number of parameters of a layer, is sampled, and a neural network model  $f$  is transformed into a variant  $f^m$  with a parameter set multiplied with  $m$  each update<sup>3</sup>. The elements of mask  $m \in \{0, \frac{1}{p}\}^d$  are randomly sampled as follows:  $m_j := m'_j/p$ ,  $m'_j \sim \text{Bernoulli}(p)$ . Parameters masked (multiplied) with 0 are disabled in an update step like pruning. Thus, dropout randomly selects various sub-networks  $f^m$  to be updated at every step. During inference at test time, dropout is not applied. One interpretation of dropout is that it trains numerous sub-networks and uses them as ensemble (Hinton et al., 2012; Srivastava et al., 2014; Bachman et al., 2014; Baldi and Sadowski, 2014; Bul et al., 2016). In this work,  $p = 0.5$ ; approximately half of the parameters are zero-masked.

### 3.2 Proposed Method: Turn-over Dropout

In the standard dropout method, dropout masks are sampled independently at every update. In our proposed method, however, we use *instance-specific dropout masks*  $m(z)$ , which are also random vectors but *deterministically* generated and

<sup>3</sup> Typically, dropout is applied to the layers of the neural network rather than its parameter matrices. In this case, each instance in a minibatch drops different column-wise parameters of matrices at once.

tied with each instance  $z$ . Thus, when the network is trained with an instance  $z$ , only a deterministic subset of its parameters is updated, as shown in Figure 1. In other words, the sub-network  $f^{\mathbf{m}(z)}$  is updated; however the corresponding counterpart of the network  $f^{\widetilde{\mathbf{m}}(z)}$  is not at all affected by  $z$ , where  $\widetilde{\mathbf{m}}(z)$  is the *flipped* mask of  $\mathbf{m}(z)$ , i.e.,  $\widetilde{\mathbf{m}}(z) := \frac{1}{p} - \mathbf{m}(z)$ . Both sub-networks,  $f^{\mathbf{m}(z)}$  and  $f^{\widetilde{\mathbf{m}}(z)}$ , can be used by applying the individual masks to  $f$ . These sub-networks are analogously comprehended as two different networks trained on a dataset with or without an instance, respectively,  $f_D$  and  $f_{D \setminus \{z_i\}}$ <sup>4</sup>. From this analogy, the influence of a training instance can be evaluated by considering these two sub-networks. The influence  $I(z_{\text{target}}, z_i; D) = L(f_{D \setminus \{z_i\}}, z_{\text{target}}) - L(f_D, z_{\text{target}})$  is estimated as

$$\begin{aligned} \hat{I}(z_{\text{target}}, z_i; D) \\ := L(f_D^{\widetilde{\mathbf{m}}(z_i)}, z_{\text{target}}) - L(f_D^{\mathbf{m}(z_i)}, z_{\text{target}}), \quad (2) \end{aligned}$$

which corresponds to the gain when using  $f_D^{\mathbf{m}(z_i)}$ , instead of  $f_D^{\widetilde{\mathbf{m}}(z_i)}$  for a prediction on  $z_{\text{target}}$ . We call this estimation method *turn-over dropout*. Its summarized advantages are as follows:

- **Lower computation time:** The method only requires running forward procedure two times.
- **No snapshot or re-training:** A single model can be used for all training instances.
- **Easy to implement:** The model modification and estimation procedure are very simple.

### 3.3 Memory-efficient Instance-specific Masks

One may think that using instance-specific masks require a large space, depending on the dataset size and the number of parameters to be masked. However, this cost is drastically reduced to a constant  $O(1)$ , using a trick. As the masks are not updated, we do not have to save them directly. Instead, we can deterministically generate the random masks with a fixed random seed number anytime. Thus, models can avoid storing masks and generate masks when using them. We call this trick as *volatile mask generation*<sup>5</sup>.

<sup>4</sup> In this paper, we associate  $f^{\mathbf{m}(z)}$  and  $f^{\widetilde{\mathbf{m}}(z)}$  with  $f_D$  and  $f_{D \setminus \{z_i\}}$ , respectively. However, while  $f_D$  does not focus on any instance in  $D$  so much, its substitute  $f^{\mathbf{m}(z)}$  may be a little biased to some characteristic of  $z$ . For ignoring bias, we can use  $f_D$  itself (i.e., full network) instead of  $f^{\mathbf{m}(z)}$ , while the representation powers of  $f_D$  and  $f_{D \setminus \{z_i\}}$  are different. We tested the alternative but did not find large improvements. Further exploration is an interesting future work.

<sup>5</sup> The volatile mask generation method solved storage and memory issues in our experiments. However, the memory

## 4 Experiments

The computational efficiency of our method is discussed in Section 2. Moreover, we answer a question: even if it is efficient, does it work well on applications? To demonstrate the applicability, we conducted experiments using different models and datasets.

**Setup** First, we used the Stanford Sentiment TreeBank (SST-2) (Socher et al., 2013) binary sentiment classification task. Five thousand instances were sampled from the training set, and 872 instances in the development set were used. We trained BERT-base classifiers (Wolf et al., 2019) with the adapter modules (Houlsby et al., 2019), which froze the pre-trained BERT parameters but newly trained branch networks in addition to the output layers. We applied the turn-over dropout on the adapter modules and output layers.

In addition, we used the CIFAR-10 (Krizhevsky, 2009) 10-class image classification task, with the 50,000 training instances and 10,000 validation instances. We trained the VGGNet19 classifier (Simonyan and Zisserman, 2015) with the turn-over dropout.

Models were trained with the cross-entropy loss. Further details of the setup are shown in Appendix A.

### 4.1 Side Effect on Model Performance

Note that turn-over dropout is not for improving the accuracy of models. It gives the models the method of efficiently estimating the influence of each training instance. A possible side effect is a deterioration of accuracy due to introducing instance-specific dropout with  $p = 0.5$ <sup>6</sup>. Thus, we first explored the change of classification accuracy when using the turn-over dropout.

For BERT with the adapter modules on SST-2, if we use a small dataset ( $N=5,000$ ), the accuracy slightly decreased from the baseline model, from 90.0% to 88.3%. If we use a larger dataset ( $N=20,000$ ), the change is negligible; 90.5% and 90.2%. Thus, in a case with large datasets, where

issue could occur even with the method, depending on implementations. For such a particular case and another solution for it, see Appendix C in detail.

<sup>6</sup>Dropout with  $p = 0.5$  is often used in various neural networks, especially on linear layers of them, and improves the accuracy. However, dropout on all layers could damage. It is also unclear how dropout with “static” masks effect because the idea is novel.

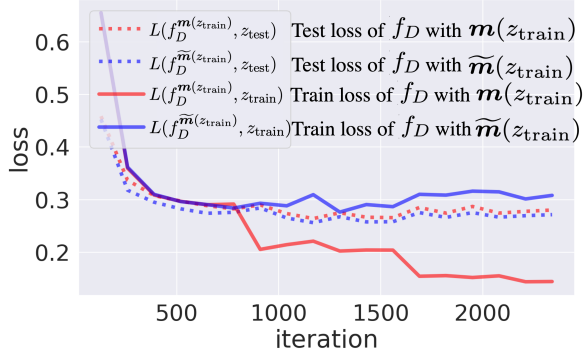


Figure 2: Loss curves of BERT on SST-2.

we typically want to use turn-over dropout for efficiency, applying the turn-over dropout does not decrease the validation accuracy compared with the baseline. However, when we use turn-over dropout on all layers of BERT *without* the adapter modules using makes training unstable. Furthermore the same is true for VGGNet on CIFAR-10. Instead, we first applied the turn-over dropout only for all layers after the 11th layer, although this means early layers can learn all instances in the training dataset and make the turn-over dropout leaky<sup>7</sup>. We found that VGGNet with turn-over dropout can overfit more than the baseline does; their accuracies are 86.2% and 92.0%, respectively. If we add regularization using the original dropout, the accuracy is recovered to 91.3%. Thus, in some cases, we have to care about the decrease of model performances when using turn-over dropout. While we experimented with the successful architectures only, exploring the side effect in various architectures and its remedy is important future work.

## 4.2 Sanity Check: Learning Curves

We first observed an interesting property of the turn-over dropout from the loss curves during training, as shown in Figure 2. The solid red line of training loss using  $m(z_{\text{train}})$ ,  $L(f_D^{m(z_{\text{train}})}, z_{\text{train}})$ , showed a typical tendency of training loss. However, the solid blue line of training loss using  $\tilde{m}(z_{\text{train}})$ ,  $L(f_D^{\tilde{m}(z_{\text{train}})}, z_{\text{train}})$ , indicated loss values close to the test losses (in dotted lines), without overfitting. This fact agrees with the idea behind the turn-over dropout; the sub-network  $f^{\tilde{m}(z_{\text{train}})}$  us-

<sup>7</sup> Yuki M. Asano (2020) demonstrated that early layers of CNN contained limited information about the statistics of images, and such low-level statistics can be learned even through a single image. Based on the finding, we assumed that early layers did not fit each instance so much, and the effect of leakage was small.

Test input: if s ##te ##ven so ##der ##berg ##h ' s ` solar ##is ' is a failure it is a glorious failure .  
(if steven soderbergh ' s ` solaris ' is a failure it is a glorious failure .)  
[positive → negative]  
Influential: un ##fo ##cus ##ed , ex ##c ##ru ##ciating ##ly te ##dious  
(unfocused , excruciatingly tedious)

Figure 3: A misclassified text in the test set and the text with the highest influence with the error label in the training set for BERT on SST-2.

Test input:  
- why do some people assume they know who the ask ##er is , based on the question he asks ? from reading back the answers that i get on my questions , i am now officially a m ##us ##lim - ch ##rist ##ian - b ##udd ##his ##t , a straight ... it ok with you that i still don ' t have an identity - crisis ?  
[Society & Culture → Business & Finance]  
Influentials:  
- is ch ##rist ##mas eve in the evening or is it just the day before ch ##rist ##mas ? my sister thinks that ch ##rist ##mas eve ...  
- how i can be ch ##rist ##ain ? i want to be real ch ##rist ##ain , how i can be ch ##rist ##ain ask ch ##rist into ...

Figure 4: A misclassified text in the test set and the texts with the highest influence with the error label in the training set for BERT on Yahoo Answers.

ing the flipped mask does not learn each training instance  $z_{\text{train}}$ .

## 4.3 Interpretation of Error of Predictions

Neural network models are notorious for their black-box prediction, which harms the trust and usability (Ribeiro et al., 2016). The influence estimation can mitigate this problem by suggesting possible reasons for a wrong model prediction by identifying influential training instances.

To verify this benefit, we collected the misclassified instances of the validation or test set and searched for the training instances that most influenced the wrong predictions. Figure 3 indicates a text example from the results. Rare words of named entities were divided into many subwords (Schuster and Nakajima, 2012; Sennrich et al., 2016; Wu et al., 2016) and requiring more complex processing. A guess is that BERT might fail to understand the input due to the cluttered subwords, and predict a wrong label, which depended on a training instance similarly with many subwords. Additionally, we conducted the same experiment on Yahoo Answers 10-label question classification dataset (Zhang et al., 2015)<sup>8</sup>, which is more complex than sentiment analysis. Figure 4 shows the results on Yahoo Answers. The misclassified text shares the phrase “ch ##rist” with the two influen-

<sup>8</sup>We used 5,000 training instances as well as SST-2.



Figure 5: Misclassified images in the validation set (upper row) and images with the highest influence in the training set (lower row) for VGGNet on CIFAR-10.

	Accuracy (%)	Loss
1% Random Removal	$76.8 \pm 1.1$	$0.521 \pm 0.030$
No Cleansing	$77.0 \pm 0.9$	$0.536 \pm 0.063$
<b>1% Cleansing</b>	<b><math>78.3 \pm 0.2</math></b>	<b><math>0.484 \pm 0.008</math></b>

Table 2: The results of data cleansing. Loss is the cross entropy loss. The averages and standard deviations from four difference runs are shown.

tial instances. Such a low-level cue is not critical in the test. However, it seemed that the model focused on the phrase and predicted the label of training instances containing the phrase.

In addition, more intuitively, image results are shown in Figure 5. The two leftmost instances with the “bird” label were wrongly predicted as “airplane.” The training instances of airplane with the highest influence on the error predictions are shown in the row below. The corresponding images had similar visual features, such as shape, layout, or color, which probably led to the wrong predictions.

#### 4.4 Data Cleansing

Another possible application of the influence estimation is to eliminate harmful instances from the training dataset. If the mean influence of a training instance on unseen instances is negative, the instance can be harmful for generalization. We experimented with data cleansing in a case of domain shift, where the training dataset is of SST-2 (movie review); however, the validation and test dataset are of the ‘electronics’ subset in Multi-Domain Sentiment Dataset (Blitzer et al., 2007) (Elec). We split the Elec dataset into 200 instances for validation and 1,800 instances for the test. Note that we do not use Elec dataset as a training dataset for studying only the effect of data cleansing.

We finetuned BERT models (with turn-over dropout) on SST-2 dataset and calculated the mean influences considering Elec’s validation set. Af-

ter that, we re-trained models without turn-over dropout on datasets that removed training instances with 1% of the most negative influences. Finally, the model performances on Elec’s test dataset are compared, as shown in Table 2. The models trained on the cleansed datasets achieved better accuracy and lower loss than those trained on the original dataset. This result demonstrated that our estimation of the influence could also be used for data cleansing.

## 5 Conclusion

This paper proposed a method that required a low computational cost for estimating the influence of a training instance. The method alters dropout with instance-specific masks and, for estimation, uses sub-networks that are not trained with each instance. The experiments demonstrated that this method could be applied even for complex models.

## Acknowledgments

We appreciate the helpful comments from the anonymous reviewers. We thank Sho Takase, Hiroshi Noji, Hitomi Yanaka, Koki Washio, Saku Sugawara, Benjamin Heinzerling, and Kazuaki Hanawa for constructive comments. This work was supported by JSPS KAKENHI Grant Number JP19H04162.

## References

- Philip Bachman, Ouais Alsharif, and Doina Precup. 2014. [Learning with pseudo-ensembles](#). In *Advances in Neural Information Processing Systems* 27, pages 3365–3373.
- Pierre Baldi and Peter Sadowski. 2014. [The dropout learning algorithm](#). *Artificial Intelligence*, 210(C):78–122.
- John Blitzer, Mark Dredze, and Fernando Pereira. 2007. [Biographies, Bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 440–447, Prague, Czech Republic. Association for Computational Linguistics.
- Samuel Rota Bul, Lorenzo Porzi, and Peter Kotschieder. 2016. [Dropout distillation](#). In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, pages 99–107.
- R Dennis Cook. 1977. Detection of influential observation in linear regression. *Technometrics*, pages 15–18.

- 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.
- Xiaochuang Han, Byron C. Wallace, and Yulia Tsvetkov. 2020. [Explaining black box predictions and unveiling data artifacts through influence functions](#). In *Proceedings of the 2020 Annual Conference of the Association for Computational Linguistics (to appear)*.
- Satoshi Hara, Atsushi Nitanda, and Takanori Maehara. 2019. [Data cleansing for models trained with sgd](#). In *Advances in Neural Information Processing Systems 32*, pages 4215–4224.
- Geoffrey E. Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R. Salakhutdinov. 2012. [Improving neural networks by preventing co-adaptation of feature detectors](#). *CoRR*, abs/1207.0580.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799, Long Beach, California, USA. PMLR.
- Pang Wei Koh and Percy Liang. 2017. [Understanding black-box predictions via influence functions](#). In *Proceedings of the 34th International Conference on Machine Learning*, pages 1885–1894.
- Alex Krizhevsky. 2009. [Learning multiple layers of features from tiny images](#). Technical report.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A Python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.
- Marco Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. [“why should I trust you?”: Explaining the predictions of any classifier](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 97–101, San Diego, California. Association for Computational Linguistics.
- Mike Schuster and Kaisuke Nakajima. 2012. Japanese and korean voice search. In *International Conference on Acoustics, Speech and Signal Processing*, pages 5149–5152.
- 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.
- Karen Simonyan and Andrew Zisserman. 2015. [Very deep convolutional networks for large-scale image recognition](#). In *International Conference on Learning Representations*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *Journal of Machine Learning Research*, 15:1929–1958.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *arXiv*, abs/1910.03771.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#). *CoRR*, abs/1609.08144.
- Andrea Vedaldi Yuki M. Asano, Christian Rupprecht. 2020. [A critical analysis of self-supervision, or what we can learn from a single image](#). In *International Conference on Learning Representations (ICLR)*.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In C. Cortes, N. D. Lawrence, D. D. Lee,

M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 649–657. Curran Associates, Inc.

Xuezhou Zhang, Xiaojin Zhu, and Stephen Wright. 2018. [Training set debugging using trusted items](#). In *AAAI Conference on Artificial Intelligence*.