

Good-Enough Example Extrapolation

Jason Wei

Google Research

jasonwei@google.com

Abstract

This paper asks whether extrapolating the hidden space distribution of text examples from one class onto another is a valid inductive bias for data augmentation. To operationalize this question, I propose a simple data augmentation protocol called “good-enough example extrapolation” (GE3). GE3 is lightweight and has no hyperparameters. Applied to three text classification datasets for various data imbalance scenarios, GE3 improves performance more than upsampling and other hidden-space data augmentation methods.

1 Introduction

Text classification is a fundamental task in NLP for which modern architectures have achieved high performance when training data is sufficient (Wang et al., 2019). In many applied settings where data collection and annotation is limited, however, a common challenge is data imbalance (Krawczyk, 2016), in which training data from certain categories is scarce. A classic example of such a scenario is intent classification, where developers may wish to update a conversational agent to be able to classify new intents, but the amount of training data for these new intents lags behind that of existing ones (Bapna et al., 2017; Gaddy et al., 2020).

One common method for mitigating the weaknesses of limited training data is data augmentation, a paradigm that is become increasingly seductive in the NLP landscape (see Feng et al., 2021, for a survey). While data augmentation may be of only incremental utility when training data is sufficient (Longpre et al., 2020), it can be particularly helpful for mitigating data scarcity in low-resource settings such as few-shot classification (Wei et al., 2021) or, as this paper will soon explore, data-imbalanced text classification.




Figure 1: Diagram illustrating how GE3 extrapolates the hidden-space distribution of examples in class A onto class B.

In this paper, I propose a simple data augmentation protocol called *good-enough¹ example extrapolation* (GE3) for the class-imbalanced scenario. As shown in Figure 1, GE3 extrapolates the hidden-space distribution of examples from one class onto another class. GE3 has no hyperparameters, is model-agnostic, and requires little computational overhead, making it easy-to-use. In empirical experiments, I apply GE3 to intent classification, newspaper headline classification, and relation classification, finding that in a variety of class-imbalanced scenarios, GE3 substantially outperforms upsampling and other hidden state augmentation techniques.

2 Hidden Space Extrapolation

Intuition. Representation learning aims to map inputs into a hidden space such that desired properties of inputs are easily extractable from their continuous representations (Pennington et al., 2014). For many representation learning functions, inputs with similar properties map to nearby points in hidden space, and the distances between hidden space

¹Paying good respects to “good-enough compositional augmentation” (Andreas, 2020).

representations represent meaningful relationships (Mikolov et al., 2013; Kumar et al., 2020).

So for a given classification task, inputs from the same category will have some distribution (i.e., cluster) in hidden space, where the distance between distributions represents the relationship between categories, and the distribution of points within the same category models some random variable. GE3 leverages this intuition by extrapolating the distribution of points in the same category, which models some random variable, from one category onto another. Figure 2 illustrates this intuition in a hypothetical instance from the HuffPost dataset of news headlines classification, where the hidden-space relationship between two examples in the travel category can be extrapolated to form a new example in the health category.

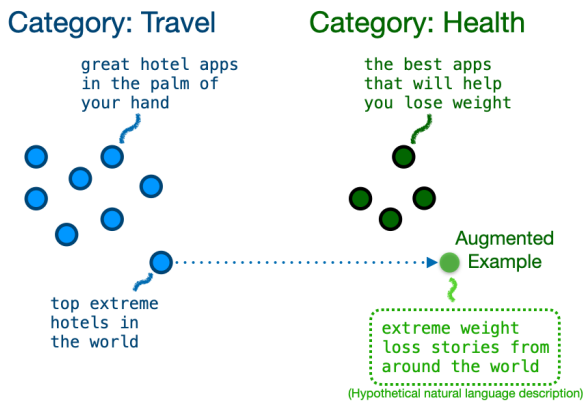


Figure 2: Intuition of extrapolating examples in hidden space from a travel category to a health category. A hypothetical natural language representation is shown for the augmented example generated via GE3.

Method. Formally, I describe GE3 as follows. Given a text classification task with k output classes $\{c_j\}_{j=1}^k$, denote the n_c training examples in class c as $\{X_c^i\}_{i=1}^{n_c}$ with corresponding hidden space representations $\{\mathbf{x}_c^i\}_{i=1}^{n_c}$. For each class c , let $\mu(c) = \frac{1}{n_c} \sum_{i=1}^{n_c} \mathbf{x}_c^i$ indicate the mean of all hidden space representations in that class.

GE3 generates augmented examples by extrapolating the data distribution from a source class c_s to a target class c_t in hidden space. And so for each hidden space representation $\mathbf{x}_{c_s}^i$ in the source class, I generate a corresponding augmented example $\hat{\mathbf{x}}_{c_t}$ in the target class:

$$\hat{\mathbf{x}}_{c_t} = \mathbf{x}_{c_s}^i - \mu(c_s) + \mu(c_t). \quad (1)$$

In total, for each class in the training set, I can generate a set of extrapolated points from every

other class, augmenting the size of the original training set by a factor of k . I then train the classification model on the union of original data and extrapolated examples. Notably, this extrapolation method operates without any hyperparameters, as augmented examples are generated via distributions from other classes instead of a noising function (c.f. augmentation techniques that usually have a strength parameter (Sennrich et al., 2016b; Wei and Zou, 2019)).

3 Experimental Setup

I evaluate the proposed hidden space extrapolation protocol in several data imbalance scenarios on three diverse text classification datasets.

3.1 Datasets

SNIPS. The Snips Voice Platform dataset² (Coucke et al., 2018) is an intent classification dataset that maps utterances to 7 different intents (e.g., ‘play music’, ‘get weather’, etc.). Each intent has about 1800 training examples.

HUFF. The HuffPost dataset (Misra, 2018) comprises news headlines published on HuffPost from 2012–2018. Headlines are categorized into one of 41 classes (e.g., ‘health’, ‘travel’, etc.), and I split the dataset such that the training set has 700 examples per class.

FEWREL. The few-shot relation classification dataset (Han et al., 2018) contains categorized relationships between specified tokens (e.g., ‘capital of,’ ‘birth name,’ etc). The posted training set contains 64 classes, and I perform a train-test split such that each class has 500 examples in the training set and 100 examples in the evaluation set.

For all three datasets, I create artificially imbalanced datasets via random sampling. Specifically, I randomly select half the classes to maintain the original number of examples N^{many} (i.e., $N^{\text{many}} = \{1800, 700, 500\}$ for SNIPS, HUFF, and FEWREL respectively), and for the other half of the classes, I train on only a subset of N^{few} examples. I run experiments on a range of N^{few} .

3.2 Model and Experimental Procedures

For the classification model, I use the common Sentence-BERT (Reimers and Gurevych, 2019, a modification of Devlin et al. (2019)) embedding

²I used the processed data from <https://github.com/MiuLab/SlotGated-SLU/tree/master/data/snips>.

	SNIPS ($k = 7$)		HUFF ($k = 41$)		FEWREL ($k = 64$)		Average	Δ
	$N^{\text{few}}=20$	$N^{\text{few}}=50$	$N^{\text{few}}=20$	$N^{\text{few}}=50$	$N^{\text{few}}=20$	$N^{\text{few}}=50$		
Baseline (upsampling)	88.9 ± 1.1	92.1 ± 0.8	27.3 ± 0.2	30.3 ± 0.3	49.6 ± 0.6	56.7 ± 0.6	55.3	-
Interpolate	89.3 ± 1.1	92.0 ± 0.9	27.0 ± 0.1	29.9 ± 0.2	49.1 ± 0.5	56.1 ± 0.5	55.1	-0.2
Within-extrapolation	86.5 ± 1.2	90.5 ± 0.8	28.7 ± 0.2	30.8 ± 0.3	49.2 ± 0.3	56.0 ± 0.6	54.8	-0.5
Linear Delta	86.4 ± 1.0	90.7 ± 1.1	29.9 ± 0.2	32.4 ± 0.3	50.9 ± 0.6	58.2 ± 0.6	55.7	+0.4
Uniform Noise	88.8 ± 0.9	91.9 ± 0.8	30.2 ± 0.8	33.3 ± 0.2	50.9 ± 0.6	57.8 ± 0.4	56.6	+1.3
Gaussian Noise	89.1 ± 1.1	92.1 ± 0.9	31.9 ± 0.3	33.7 ± 0.2	53.0 ± 0.6	60.3 ± 0.6	58.0	+2.7
GE3 (ours)	90.6 ± 0.6	92.8 ± 0.7	32.7 ± 0.2	36.8 ± 0.1	56.3 ± 0.6	64.0 ± 0.2	59.9	+4.6

Table 1: Accuracy (%) of GE3, upsampling, and five other hidden space augmentation techniques on data-imbalanced text classification scenarios, where half of the classes are restricted to N^{few} training examples. k : number of total classes for a classification task. Δ : improvement over the upsampling baseline.

pipeline and add an additional softmax layer for classification. I implement GE3 at the final hidden layer of Sentence-BERT, which has size 768. That is, the hidden-space augmentation method only updates classifier weights after the BERT encoder. Before training, the data processing pipeline upsamples from classes with fewer examples until all classes have the same number of examples in the training set. I run all experiments for five random seeds.

3.3 Hidden Space Augmentation Baselines

As baselines for comparison, I also explore several other hidden space augmentation techniques:

Example interpolation. Given the hidden space representations of two examples \mathbf{x}_c^i and \mathbf{x}_c^j in the same class, I generate an augmented example

$$\hat{\mathbf{x}}_c = \frac{1}{2}(\mathbf{x}_c^i + \mathbf{x}_c^j). \quad (2)$$

Within-extrapolation. (DeVries and Taylor, 2017; Kumar et al., 2019). Given two examples \mathbf{x}_c^i and \mathbf{x}_c^j in the same class, I extrapolate the hidden space between the two to form an augmented example

$$\hat{\mathbf{x}}_c = \lambda \cdot (\mathbf{x}_c^i + \mathbf{x}_c^j) - \mathbf{x}_c^i. \quad (3)$$

Following Kumar et al. (2019), I use $\lambda = 0.5$.

Linear delta. (Kumar et al., 2019). The difference between two examples \mathbf{x}_c^i and \mathbf{x}_c^j in the same class can be added to a third example \mathbf{x}_c^k to form an augmented example:

$$\hat{\mathbf{x}}_c = (\mathbf{x}_c^i - \mathbf{x}_c^j) + \mathbf{x}_c^k. \quad (4)$$

Noising. Given some example \mathbf{x}_c^i , I add noise \mathbf{n} to yield an augmented example

$$\hat{\mathbf{x}}_c = \mathbf{x}_c^i + \mathbf{n}(\cdot). \quad (5)$$

For $\mathbf{n}(\cdot)$, I explore both **Uniform Noise**, where each element is uniformly sampled from $[a, b]$, where $a = -0.1$ and $b = 0.1$, as well as **Gaussian Noise**, where each element is sampled from $\mathcal{N}(\mu, \sigma)$, where $\mu = 0$ and $\sigma = 0.1$.

For these techniques, I generate augmented examples until each class has $n_{\text{aug}} \cdot N^{\text{many}}$ training examples, where $n_{\text{aug}} = 5$ (a choice which is later explored in Figure 3).

4 Results

Table 1 shows results for GE3 on the three datasets for $N^{\text{few}} = 20$ and $N^{\text{few}} = 50$. GE3 outperforms the upsampling baseline by an average of 4.6%, with strongest improvements on HUFF and FEWREL. Of the other augmentation techniques, Gaussian noising and uniform noising had the best performance, with an average improvement of 2.7% and 1.3%, respectively. Whereas these techniques only enforce smoothness around the distribution of points in a single class, I hypothesize that GE3 improved performance more because it injects a stronger inductive bias that the distribution of examples of the same class around their mean can be extrapolated to other classes.

Moreover, as Table 1 only shows results for $N^{\text{few}} \in \{20, 50\}$, in Figure 4 I compare GE3 with upsampling, as well as Gaussian and uniform noise (the strongest baselines), for $N^{\text{few}} \in \{10, 20, 40, 60, 100, 200, 300, 400, 500\}$. GE3 improves performance across a wide variety of N^{few} values, with improvements over the baselines slightly diminishing when training data is more balanced (as expected).

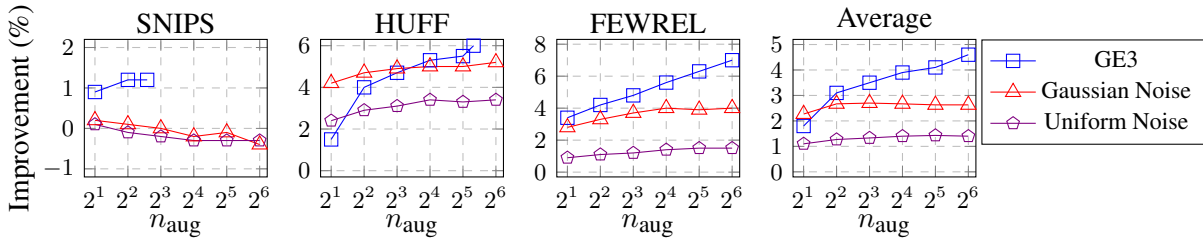


Figure 3: Improvements over the upsampling baseline from augmentation methods based on how many duplicates of training data were made (n_{aug}). GE3 extrapolates examples from one class to another, so it generates at most $n_{aug} = 6$ for SNIPS, $n_{aug} = 40$ for HUFF and $n_{aug} = 63$ for FEWREL.

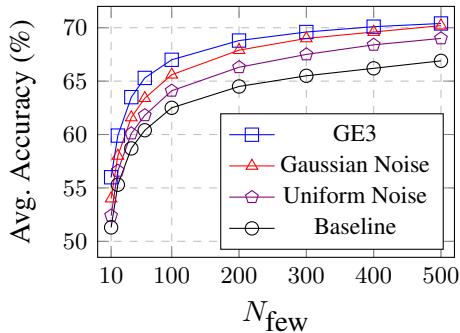


Figure 4: Comparing GE3 with the upsampling baseline, gaussian noise, and uniform noise for a range of N_{few} (the number of training examples in the imbalanced classes). The performance shown is averaged over three datasets.

Finally, one of the reasons that GE3 improves performance more than other techniques could be that each class gets extrapolated examples from every other class, and if classes have unique distributions, then these extrapolated examples are valuable as additional training data. Therefore, I perform an ablation study using a variable n_{aug} , which restricts the number of other classes a given class can receive extrapolated examples from. For instance, if $n_{aug} = 2$, then any given class may only receive extrapolated examples from two other random classes, even if there are 63 other classes (as is the case in FEWREL). I also perform a similar ablation for Gaussian and uniform noise, in which I generate augmented examples until each class has $n_{aug} \cdot N^{many}$ training examples.

Figure 3 shows these results. For uniform and Gaussian noise, additional augmented examples did not further improve performance after around $n_{aug} = 16$. For GE3, on the other hand, improvement continued to increase as n_{aug} increased (although the marginal improvement decreases for each n_{aug}). This result confirms the intuition that extrapolations from more classes provided additional value during training.

5 Related Work

Text data augmentation. Data augmentation methods for NLP have garnered increased interest in recent years. Many common techniques modify data using either token perturbations (Zhang et al., 2015; Sennrich et al., 2016a) or language models (Sennrich et al., 2016b; Kobayashi, 2018; Liu et al., 2020; Ross et al., 2021). These techniques occur at the input-level, where all augmented data is represented by discrete tokens in natural language.

Hidden space augmentation. A growing direction in data augmentation has proposed to augment data in hidden space instead of at the input-level. In computer vision, DeVries and Taylor (2017) explored noising, interpolation, and extrapolation, and MIXUP (Zhang et al., 2018) combines pairs of examples. These methods have since been adopted to NLP—Chen et al. (2020) modify MIXUP to improve semi-supervised text-classification, and Kumar et al. (2019) explore various hidden space augmentation techniques for few-shot intent classification, which I evaluated as baselines. Whereas the extrapolation technique used by DeVries and Taylor (2017) and Kumar et al. (2019) (which I call “within-extrapolation” in this paper) extrapolates the hidden space between a pair of points in the hidden space, GE3 extrapolates the hidden space distribution of one class onto another class.

Example extrapolation. In vision, Schwartz et al. (2018) used a modified auto-encoder to synthesize new examples from category after seeing a few examples from it, improving few-shot object recognition. Perhaps most similar to this work, Lee et al. (2021) train T5 (Raffel et al., 2020) to, given some examples of a class as an input sequence, generate additional examples. Because GE3 operates in hidden space, it is simpler and more computationally accessible compared with fine-tuning T5 for each classification task.

6 Discussion

The motivation for this work emerged from a mixture of failed experiments (I tried to devise an algorithm to select better augmented sentences in hidden space) in addition to an admiration for the elegance of the Ex2 (Lee et al., 2021). In hindsight, it would have been helpful to compare the performances of these two techniques in the same setting (notably, whereas I artificially restrict the sample size for certain classes in this paper, Ex2 uses the original data distributions of the datasets, which is a harder setting to show improvements from data augmentation).

I would be remiss not to mention at least one weakness that I see in my own work. There has been an influx of recent work proposing various augmentation techniques for different NLP tasks, and due to the lack of standardized evaluation datasets and models, many papers³ do not perform a full comparison with respect to relevant baselines. This paper circumvents comparing with many data augmentation baselines (e.g., Chawla et al. (2002)) by focusing on the question of whether hidden-space example extrapolation is a valid inductive bias (and not whether it is the best augmentation technique). Hence, although I find example extrapolation to be a nice idea, I should concede that the particular GE3 operationalization of example extrapolation should undergo more comprehensive comparison with baselines before I can recommend it as a go-to augmentation technique.

In summary, I have proposed a data augmentation protocol called GE3, which extrapolates the hidden space distribution of one class onto another. The empirical experiments in this paper suggests that example extrapolation in hidden space is a valid inductive bias for data augmentation. Moreover, GE3 is appealing because it has no hyperparameters, is model agnostic, and is lightweight. If example extrapolation is an idea deserving of further exploration by our field, I hope this paper adds a leaf to the tree of knowledge in that space.

Acknowledgements

Thanks Dan Garrette and Hyung Won Chung for providing feedback on the manuscript, and Barret Zoph and Ekin Dogus Cubuk for general feedback.

³I'll offer up my own paper (Wei and Zou, 2019) as an example. That paper should have compared with backtranslation and contextual augmentation at least.

References

- Jacob Andreas. 2020. [Good-enough compositional data augmentation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7556–7566, Online. Association for Computational Linguistics.
- Ankur Bapna, Gokhan Tür, Dilek Hakkani-Tür, and Larry Heck. 2017. [Towards zero-shot frame semantic parsing for domain scaling](#). In *Proc. Interspeech 2017*, pages 2476–2480.
- Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357.
- Jiaao Chen, Zichao Yang, and Diyi Yang. 2020. [Mix-Text: Linguistically-informed interpolation of hidden space for semi-supervised text classification](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2147–2157, Online. Association for Computational Linguistics.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, Thibaut Lavril, et al. 2018. [Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces](#).
- 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.
- Terrance DeVries and Graham W Taylor. 2017. Dataset augmentation in feature space. *arXiv preprint arXiv:1702.05538*.
- Steven Y Feng, Varun Gangal, Jason Wei, Sarath Chandar, Soroush Vosoughi, Teruko Mitamura, and Eduard Hovy. 2021. [A survey of data augmentation approaches for NLP](#). *Findings of the Association of Computational Linguistics*.
- David Gaddy, Alex Kouzemtchenko, Pavan Kumar Reddy, Prateek Kolhar, and Rushin Shah. 2020. Overcoming conflicting data for model updates. *arXiv preprint arXiv:2010.12675*.
- Xu Han, Hao Zhu, Pengfei Yu, Ziyun Wang, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2018. [FewRel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4803–4809, Brussels, Belgium. Association for Computational Linguistics.

- Sosuke Kobayashi. 2018. [Contextual augmentation: Data augmentation by words with paradigmatic relations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 452–457, New Orleans, Louisiana. Association for Computational Linguistics.
- Bartosz Krawczyk. 2016. Learning from imbalanced data: open challenges and future directions. *Progress in Artificial Intelligence*, 5(4):221–232.
- Vaibhav Kumar, Tenzin Singhay Bhotia, Vaibhav Kumar, and Tanmoy Chakraborty. 2020. [Nurse is Closer to Woman than Surgeon? Mitigating Gender-Biased Proximities in Word Embeddings](#). *Transactions of the Association for Computational Linguistics*, 8:486–503.
- Varun Kumar, Hadrien Glaude, Cyprien de Lichy, and William Campbell. 2019. [A closer look at feature space data augmentation for few-shot intent classification](#). In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 1–10, Hong Kong, China. Association for Computational Linguistics.
- Kenton Lee, Kelvin Guu, Luheng He, Tim Dozat, and Hyung Won Chung. 2021. [Neural data augmentation via example extrapolation](#). *arXiv preprint*.
- Ruibo Liu, Guangxuan Xu, Chenyan Jia, Weicheng Ma, Lili Wang, and Soroush Vosoughi. 2020. [Data boost: Text data augmentation through reinforcement learning guided conditional generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9031–9041, Online. Association for Computational Linguistics.
- Shayne Longpre, Yu Wang, and Chris DuBois. 2020. [How effective is task-agnostic data augmentation for pretrained transformers?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4401–4411, Online. Association for Computational Linguistics.
- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013. [Linguistic regularities in continuous space word representations](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 746–751, Atlanta, Georgia. Association for Computational Linguistics.
- Rishabh Misra. 2018. [HuffPost news category dataset](#).
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). 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 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Alexis Ross, Tongshuang Wu, Hao Peng, Matthew E Peters, and Matt Gardner. 2021. [Tailor: Generating and perturbing text with semantic controls](#). *arXiv preprint arXiv:2107.07150*.
- Eli Schwartz, Leonid Karlinsky, Joseph Shtok, Sivan Harary, Mattias Marder, Abhishek Kumar, Rogerio Feris, Raja Giryes, and Alex Bronstein. 2018. [Delta-encoder: an effective sample synthesis method for few-shot object recognition](#). In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [Edinburgh neural machine translation systems for WMT 16](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 371–376, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019. [Glue: A multi-task benchmark and analysis platform for natural language understanding](#). *Proceedings of the International Conference on Learning Representations*.
- Jason Wei, Chengyu Huang, Soroush Vosoughi, Yu Cheng, and Shiqi Xu. 2021. [Few-shot text classification with triplet networks, data augmentation, and curriculum learning](#). *arXiv preprint*.
- Jason Wei and Kai Zou. 2019. [EDA: Easy data augmentation techniques for boosting performance on text classification tasks](#). 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 6382–6388, Hong Kong, China. Association for Computational Linguistics.

Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. 2018. [mixup: Beyond empirical risk minimization](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.

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.