

Outlier-Aware Training for Improving Group Accuracy Disparities

Li-Kuang Chen^{1*} Canasai Kruengkrai² Junichi Yamagishi²

¹National Tsing Hua University, Taiwan

lkchen@nlpplab.cc

²National Institute of Informatics, Japan

{canasai, jyamagishi}@nii.ac.jp

Abstract

Methods addressing spurious correlations such as Just Train Twice (JTT, Liu et al. 2021) involve reweighting a subset of the training set to maximize the worst-group accuracy. However, the reweighted set of examples may potentially contain unlearnable examples that hamper the model’s learning. We propose mitigating this by detecting outliers to the training set and removing them before reweighting. Our experiments show that our method achieves competitive or better accuracy compared with JTT and can detect and remove annotation errors in the subset being reweighted in JTT.¹

1 Introduction

Machine learning models trained with empirical risk minimization (ERM, Vapnik 1992) can achieve a high average accuracy by minimizing the overall loss during training. Despite this, ERM models are also known to perform poorly on certain minority groups of examples. When specific attributes in a dataset frequently co-occur with a class label, ERM models often learn to correlate the co-occurring attributes and the label, using the attributes as “shortcuts” for classifying examples. These “shortcuts” are also called *spurious correlations*, because model performance can significantly decrease when the model encounters examples that belong to a minority group where the correlations between the attributes and class label do not hold.

More specifically, each class in a dataset can be divided by whether their examples contain such spurious attributes. Each set of examples with a class-attribute combination is called a “group”. The worst group is characterized by having the poorest ERM model performance among other groups. As an example, Figure 1 shows accuracy disparities among groups in the FEVER dataset. The

* This work was conducted during the author’s internship under National Institute of Informatics, Japan.

¹Our code is available at <https://github.com/nii-yamagishilab/jtt-m>.




Figure 1: Results for the FEVER test set (Thorne et al., 2018; Schuster et al., 2021). The data are divided into six groups in accordance with *class-attribute* combinations, where class = {REFUTES (REF), SUPPORTS (SUP), NOT ENOUGH INFO (NEI)} and attribute = {no neg, neg}, indicating the presence of a negation word in the claim. Both methods perform well on groups with strong spurious correlations (e.g., [REF, neg]). Our proposed method (JTT-m) helps improve accuracies for groups where such spurious correlations do not appear (e.g., [SUP, neg] and [NEI, neg]).

ERM-trained model can achieve close to perfect accuracy on the group with a spurious correlation (the REFUTES class with negation), but only half the accuracy on the worst group (the SUPPORTS class with negation).

Improving the worst-group performance of ERM models while maintaining the overall accuracy is an active topic of research that has applications in fair machine learning classifiers or robustness against adversarial examples (Słowiak and Bottou, 2022). Methods aiming to maximize worst-group accuracy can be roughly categorized into two categories: those that utilize group information and those that do not. Group Distributionally Robust Optimization (Group DRO, Sagawa et al. 2020) uses attribute (and thus group) information during training to dynamically minimize the loss of each group. While Group DRO achieves a high worst-group and overall accuracy, it requires annotation

on group information during training, which can be expensive to obtain and unavailable for less popular datasets. On the other hand, methods such as DRO with Conditional Value-at-Risk (CVaR DRO, [Duchi et al. 2019](#); [Levy et al. 2020](#)), Learning from Failure (LfF, [Nam et al. 2020](#)), Predict then Interpolate (PI, [Bao et al. 2021](#)), Spectral Decoupling (SD, [Pezeshki et al. 2021](#)), Just Train Twice (JTT, [Liu et al. 2021](#)), and RWY and SUBY from ([Idrissi et al., 2022](#)) all aim to minimize worst-group loss without group information.

CVaR DRO minimizes worst-case loss over all subpopulations of a specific size and requires computing the worst-case loss at each step. LfF trains an intentionally biased model and upweights the minority examples. PI interpolates distributions of correct and incorrect predictions and can minimize worst-case loss over all interpolations. SD replaces the L_2 weight decay in the cross entropy loss function with logits. RWY reweights sampling probabilities so that mini-batches are class-balanced. SUBY subsamples large classes so that every class is the same size as the smallest class. JTT simply obtains misclassified examples (the *error set*) from the training set once and upweights the fixed set of erroneous examples. We focus on JTT due to its simplicity and relative effectiveness and because it does not require group information for improving worst-group accuracy. While [Idrissi et al. \(2022\)](#)’s SUBY and RWY also follow JTT in improving worst-group accuracies, their methods target only datasets with imbalanced classes, and are not applicable to class-balanced datasets such as MultiNLI ([Williams et al., 2018](#)).

We propose further enhancing JTT by removing outliers from the error set before upweighting it. The outliers might be examples that are difficult to learn, such as annotation errors. Keeping them from being upweighted allows the model to train on a cleaner error set and thus better show the intended effect of the original JTT. We focus on worst-group performance caused by the spurious correlations of negation words and evaluate on datasets susceptible to spurious correlations of this type. Our experiments on the FEVER and MultiNLI datasets show that our method can outperform JTT in terms of either the average or the worst-group accuracy while maintaining the same level of performance for the other groups.

Our contributions are as follows. We devise a method for improving worst-group accuracy with-

out group information during training based on JTT (Section 3). We show that by removing outliers from the error set being upweighted, we can achieve similar or better overall and worst-group performance (Section 4.2). Our examination of the outliers being removed also suggests that the improvement may come from removing annotation errors in the upweighted error set (Section 4.3).

2 Background

Spurious correlations and minority groups

We investigate the spurious correlations occurring in two natural-language datasets: FEVER ([Thorne et al., 2018](#)) and MultiNLI ([Williams et al., 2018](#)). The task for FEVER involves retrieving documents related to a given *claim*, finding sentences to form *evidence* against the claim, and then classifying the claim on the basis of the evidence into three classes: SUPPORTS (SUP), REFUTES (REF), or NOT ENOUGH INFORMATION (NEI). We focus on improving the worst-group classification performance for the final part of the task. The task for MultiNLI is to classify whether the *hypothesis* is entailed by, neutral with, or contradicted by the *premise*. We use [Schuster et al. \(2021\)](#)’s preprocessing of both datasets, containing 178,059/11,620/11,710 training/dev/test examples for FEVER and 392,702/9,832 training/test examples for MultiNLI.

Attributes known to cause spurious correlations for these datasets are negation words ([Gururangan et al., 2018](#)) and verbs that suggest negating actions ([Schuster et al., 2019](#)). We merge these two sources of negation words into a single set: {*no, never, nothing, nobody, not, yet, refuse, refuses, refused, fail, fails, failed, only, incapable, unable, neither, none*}. Each class can be split into two groups based on whether each claim/hypothesis contains a spurious attribute (i.e., the negation words listed above). Models tend to perform well on groups where the attributes are highly correlated with the label. Groups where the correlation between the label and the attribute does not hold are called *minority groups* or *worst groups*, since models often fail to classify their examples correctly. For example, the claim “*Luis Fonsi does **not** go by his given name on stage.*”, labeled SUPPORTS, belongs to the worst group [SUP, neg].

Table 1(a) shows that most claims containing negation are from the class REFUTES. The relatively small amount of examples from the groups

	No negation	Negation	
REF	27,575 (17.1%)	14,275 (86.3%)	41,850 (23.5%)
SUP	99,303 (61.5%)	1,267 (7.7%)	100,570 (56.5%)
NEI	34,633 (21.4%)	1,006 (6.0%)	35,639 (20.0%)

(a) FEVER

	No negation	Negation	
Contr	88,180 (27.3%)	42,723 (61.2%)	130,903 (33.3%)
Ent	118,554 (36.7%)	12,345 (17.7%)	130,899 (33.3%)
Neut	116,185 (36.0%)	14,715 (21.1%)	130,900 (33.3%)

(b) MultiNLI

Table 1: Class and group distributions for (a) FEVER and (b) MultiNLI training sets. Both datasets show a high spurious correlation between the REF (Contr) class and the attribute neg. Minority groups where the spurious correlation does not hold are [SUP (Ent), neg] and [NEI (Neut), neg].

(SUP, negation) and (NEI, negation) form the minority groups, where the ERM model performance fails. A similar trend can be seen in Table 1(b).

Empirical Risk Minimization (ERM)

Let $x \in \mathcal{X}$ be a training example and $y \in \mathcal{Y}$ be its label. Given a dataset $D = \{(x_i, y_i)\}_{i=1}^N$, ERM aims to minimize the average loss (“empirical risk”), defined as:

$$J_{\text{ERM}}(\theta) = \frac{1}{N} \sum_{(x,y) \in D} \ell(g_\theta(x), y), \quad (1)$$

where N is the number of training examples, $g_\theta(\cdot)$ is the model, and θ represents model parameters. We use cross-entropy loss as the loss function:

$$\ell(g_\theta(x), y) = - \sum_{y \in \mathcal{Y}} \mathbb{1}\{y = \hat{y}\} \log(p_\theta(\hat{y}|x)), \quad (2)$$

where $\mathbb{1}\{\cdot\}$ is the indicator function, x represents the input sentence pair (s_1, s_2) , and $y \in \mathcal{Y} = \{\text{SUP}, \text{REF}, \text{NEI}\}$ ($\{\text{Ent}, \text{Contr}, \text{Neut}\}$ for MultiNLI). We first encode the input sentence pairs with BERT (Devlin et al., 2019) and feed the resulting embedding \mathbf{e} into a multi-layer perceptron (MLP) followed by a softmax function for classification:

$$\begin{aligned} p_\theta(\hat{y}|x) &= \text{softmax}(\text{MLP}(\mathbf{e})), \\ \mathbf{e} &= \text{BERT}(s_1, s_2). \end{aligned} \quad (3)$$

Just Train Twice (JTT)

Liu et al. (2021) propose improving worst-group performance by simply training with an up-

weighted error set. During the first round of training, the set of incorrectly classified training examples is identified via an ERM model. The training error set E is then upweighted with a real and positive upweight factor $\lambda_{\text{up}} \in \mathbb{R}^+$, and a final model is trained on the reweighted objective:

$$\begin{aligned} J_{\text{up}}(\theta, E) &= \frac{1}{N_{\text{up}}} \left(\lambda_{\text{up}} \sum_{\substack{(x,y) \\ \in E}} \ell(g_\theta(x), y) + \sum_{\substack{(x,y) \\ \notin E}} \ell(g_\theta(x), y) \right), \end{aligned} \quad (4)$$

where λ_{up} is a hyperparameter, and N_{up} is the size of the training set after upweighting.

3 Proposed method

Even though the upweighted ERM error set can improve worst-group accuracy, it is possible that the error set contains *unlearnable* or *out-of-distribution* (OOD) examples, e.g., annotation errors. When upweighting the entire error set, these examples will get amplified along with the rest of the error set, lessening the overall benefits of upweighting and retraining.

We propose modifying the JTT algorithm by removing outliers in the ERM error set before training the second time. We adopt a similar approach from Lee et al. (2018) for detecting outliers. Let \mathbf{x} be the output of the penultimate layer (i.e., the last layer before the logits) and belong to class y . First, we calculate the Mahalanobis distance for each \mathbf{x} from the mean of each class y :

$$M(\mathbf{x}) = \sqrt{(\mathbf{x} - \boldsymbol{\mu}_y)^\top \boldsymbol{\Sigma}_y^{-1} (\mathbf{x} - \boldsymbol{\mu}_y)}, \quad (5)$$

where $\boldsymbol{\mu}_y$ and $\boldsymbol{\Sigma}_y$ are the class mean and covariance.² The greater the distance of \mathbf{x} is from $\boldsymbol{\mu}_y$, the likelier it is to be an OOD example.

Then, we filter OOD examples by comparing the calculated Mahalanobis distance against a chi-squared distribution with a critical value α of 0.001 and a degree of freedom df :³

$$x_i \in \begin{cases} S_{\text{in}} & \text{if } p_i < \alpha, \\ S_{\text{out}} & \text{if } p_i \geq \alpha, \end{cases} \quad (6)$$

²We compute $\boldsymbol{\Sigma}_y$ using the standard covariance maximum likelihood estimate (MLE) implemented in scikit-learn.

³We select a value of df that yields the best worst-group accuracy on the dev set.

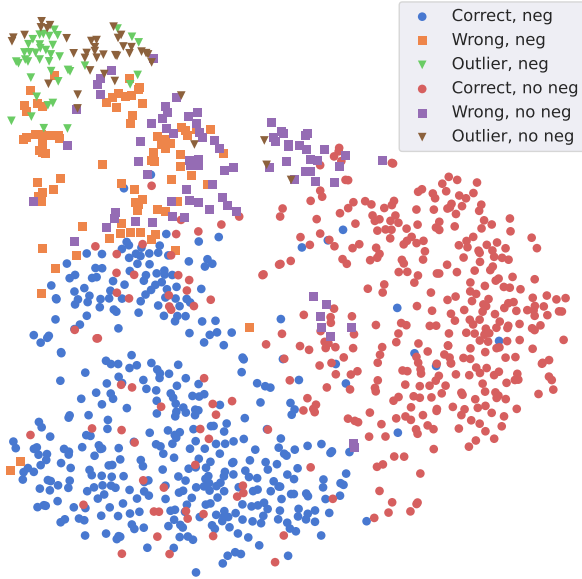


Figure 2: T-SNE visualization of samples from the class Entailment of the MultiNLI training set. Correct predictions of groups with and without negation (blue and red) are quite well separated. Wrong predictions lie at the top left, and outliers lie further away. Outliers are detected by their Mahalanobis scores.

where S_{in} and S_{out} are the sets of in-distribution and OOD training examples, and p_i is the p -value of the i -th example. We show the T-SNE visualization in Figure 2.

Once the OOD examples are identified, we remove the subset of misclassified OOD examples from the error set E , forming a new error set E_{in} :

$$E_{in} = \{(x_i, y_i) \text{ s.t. } \hat{y}_i \neq y_i \wedge x_i \notin S_{out}\}, \quad (7)$$

which is then upweighted as per JTT:

$$\begin{aligned} J_{up-in}(\theta, E_{in}) &= \frac{1}{N_{up}} \left(\lambda_{up} \sum_{\substack{(x,y) \\ \in E_{in}}} \ell(g_\theta(x), y) + \sum_{\substack{(x,y) \\ \notin E_{in}}} \ell(g_\theta(x), y) \right), \end{aligned} \quad (8)$$

4 Experiments

4.1 Training details

We follow Sagawa et al. (2020); Liu et al. (2021); Idrissi et al. (2022) in using different optimization settings for different training methods to maximize the validation accuracy. For ERM, we used the AdamW optimizer (Loshchilov and Hutter, 2019), linear learning rate decay, and a gradient clipping of 1. For the first training of JTT, we used the SGD

optimizer without gradient clipping. The second training used the same settings as those of ERM.

We used HuggingFace’s implementation (Wolf et al., 2020) of BERT-base with default parameter settings. For all methods, we used a batch size of 32, initial learning rate of $2e-5$, and we trained them for 2 epochs. We tried $df \in \{4, 5, 6\}$ and $\lambda_{up} \in \{1, 2, 3, 4\}$ and selected the values yielding the best worst-group accuracy on the dev set. Since no dev set is provided for MultiNLI, we tuned the hyperparameters on FEVER and applied them to MultiNLI.

4.2 Results

We compared our proposed method (referred to as **JTT-m**, Eq. (8)) against two baselines: **ERM** (Eq. (1)) and **JTT** (Eq. (4)). Table 2 shows the results for the average and worst-group performances of various approaches.

As expected, **ERM** had the best average accuracy but performed poorly on the worst group across the two datasets. **JTT** and **JTT-m** had improved performance on the worst group with slightly decreased average accuracies on both datasets compared with **ERM**. On FEVER, **JTT-m** outperformed **JTT** in average accuracy while maintaining the same worst-group [SUP, neg] accuracy. On MultiNLI, **JTT-m** performed significantly better on the worst group [Neut, neg] and maintained the same average accuracy as **JTT**.

We also observed larger variations in the results for FEVER. This is likely due to the smaller group sizes in FEVER. The worst group of MultiNLI [Neut, neg] accounted for around 3.5% of the test set, while FEVER’s [SUP, neg] was only 0.5% of the test set and was about 5 times lower than the smallest group in MultiNLI in absolute numbers. For the same reason, another minority group of FEVER, [NEI, neg], also displayed a higher variation.

In addition, **JTT-m** slightly reduced training time due to the smaller training set. Our Mahalanobis distance method detected 2,077 and 1,821 examples as outliers in the FEVER and MultiNLI error sets. By eliminating these examples, we could reduce the training time while achieving results similar to or better than **JTT**.

4.3 Discussion

The improvements for the MultiNLI worst group agree with our hypothesis: removing outliers from

Dataset	FEVER		MultiNLI	
	Avg. (%)	Worst (%)	Avg. (%)	Worst (%)
ERM	87.8 \pm 0.2	48.6 \pm 0.7	84.9 \pm 0.1	72.0 \pm 1.0
JTT	86.8 \pm 0.2	50.5 \pm 3.5	83.0 \pm 0.2	75.5 \pm 1.5
JTT-m	87.4 \pm 0.1*	50.2 \pm 2.8	83.0 \pm 0.3	77.3 \pm 0.4*

Table 2: Average and worst-group test accuracies for all methods. The ‘‘Worst’’ column indicates the worst-group accuracies on [SUP, neg] and [Neutr, neg] for FEVER and MultiNLI, respectively. We report mean and standard deviation computed across five runs using different random seeds. ‘‘*’’ indicates the statistical significance compared with JTT (paired t-test, $p < 0.05$).

Group	JTT	JTT-m
[REF, no neg]	79.9 \pm 0.5	80.7 \pm 0.3
[REF, neg]	93.8 \pm 0.6	96.2 \pm 0.6*
[SUP, no neg]	94.7 \pm 0.2	94.5 \pm 0.1
[SUP, neg]	50.5 \pm 3.5	50.2 \pm 2.8
[NEI, no neg]	82.5 \pm 0.5	83.0 \pm 0.3
[NEI, neg]	71.5 \pm 0.9	72.1 \pm 3.3

(a) FEVER

Group	JTT	JTT-m
[Contr, no neg]	82.8 \pm 0.7	82.8 \pm 1.0
[Contr, neg]	91.9 \pm 0.1	91.8 \pm 0.6
[Ent, no neg]	82.6 \pm 0.2	82.2 \pm 1.1
[Ent, neg]	79.5 \pm 0.5	78.9 \pm 1.9
[Neut, no neg]	81.2 \pm 0.6	81.7 \pm 0.8
[Neut, neg]	75.5 \pm 1.5	77.3 \pm 0.4*

(b) MultiNLI

Table 3: Accuracies and standard deviations for each group on (a) FEVER and (b) MultiNLI. ‘‘*’’ indicates statistical significance (paired t-test, $p < 0.05$).

the upweighted error set improves model performance. As seen in Table 3, all other groups of MultiNLI were either not affected by the removal of outliers or showed insignificant changes. On the other hand, removing outliers from the FEVER error set seemed to have a larger effect on groups other than the worst group [SUP, neg], especially on [REF, neg] and [NEI, neg].

We examined the group-wise percentage of the error-set OOD examples (i.e., the ones removed in **JTT-m**) to see how each group may be affected by the removal of their OOD examples (Figure 3). Despite the improvements in groups [REF, neg] and [Neut, neg], few to no examples from these groups were regarded as outliers by the Mahalanobis dis-

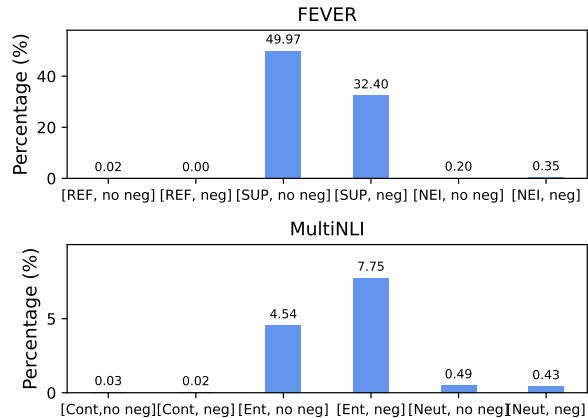


Figure 3: Percentage of OOD examples in the error set of each group. A large percentage of examples from classes SUP and Ent are regarded as outliers. FEVER’s SUP has a much higher percentage removed compared with MultiNLI’s Ent. All other groups contain only less than 1% of examples regarded as outliers.

tance method. Instead, groups of classes SUP and Ent, whose performance does not improve when outliers are removed, contained the highest percentage of OOD examples. This suggests that these outliers can affect the model’s decision boundaries among classes.

To investigate the properties of the OOD examples detected, we randomly sampled 100 examples from S_{in} and S_{out} for both FEVER and MultiNLI. For FEVER, we found 24 annotation errors in S_{out} , much higher than the 1 annotation error in S_{in} . For MultiNLI, S_{out} contained 10 annotation errors, whereas S_{in} contained 4. We show a sample of the annotation errors found in Table 4. This suggests that (1) the Mahalanobis distance method can detect at least a subset of annotation errors as outliers, and (2) the improvements in either the group or the overall performance may be partially due to the removal of these annotation errors.

Claim:	Nice & Slow was released in 1968.
Evidence:	"Nice & Slow" is a 1998 single from Usher's second album My Way.
Annotated label:	SUPPORTS
Predicted label:	REFUTES

(a) FEVER

Premise:	So far, no promising treatments exist according to Larry Gentilello.
Hypothesis:	Larry Gentilello asserted that effective treatments already exist, not just treatments that hold promise.
Annotated label:	Entailment
Predicted label:	Contradiction

(b) MultiNLI

Table 4: Example of annotation errors from (a) FEVER and (b) MultiNLI.

5 Conclusion

We have shown that the JTT algorithm can benefit from pruning the error set before upweighting and training a second time, improving worst-group accuracy or overall accuracy on two popular datasets. We also showed that annotation errors may occur in the error set, hampering JTT's effectiveness. These annotation errors can be mitigated by detecting and removing them with our Mahalanobis distance method. Investigating the effects of using other OOD-detection methods and finding a more effective way to tune the additional hyperparameters are directions for our future work.

Acknowledgments

This work is supported by JST CREST Grants (JP-MJCR18A6 and JPMJCR20D3) and MEXT KAKENHI Grants (21H04906), Japan.

References

- Yujia Bao, Shiyu Chang, and Regina Barzilay. 2021. [Predict then interpolate: A simple algorithm to learn stable classifiers](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 640–650. PMLR.
- 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. Association for Computational Linguistics.
- John Duchi, Tatsunori Hashimoto, and Hongseok Namkoong. 2019. [Distributionally robust losses against mixture covariate shifts](https://web.stanford.edu/~hnamk/papers/DuchiHaNa19.pdf). <https://web.stanford.edu/~hnamk/papers/DuchiHaNa19.pdf>.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). 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 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Badr Youbi Idrissi, Martin Arjovsky, Mohammad Pezeshki, and David Lopez-Paz. 2022. [Simple data balancing achieves competitive worst-group-accuracy](#). In *Proceedings of the First Conference on Causal Learning and Reasoning*, volume 177 of *Proceedings of Machine Learning Research*, pages 336–351. PMLR.
- Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. 2018. [A simple unified framework for detecting out-of-distribution samples and adversarial attacks](#). In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Daniel Levy, Yair Carmon, John C Duchi, and Aaron Sidford. 2020. [Large-scale methods for distributionally robust optimization](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 8847–8860. Curran Associates, Inc.
- Evan Z Liu, Behzad Haghgoo, Annie S Chen, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa, Percy Liang, and Chelsea Finn. 2021. [Just train twice: Improving group robustness without training group information](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 6781–6792. PMLR.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations (ICLR)*.
- Junhyun Nam, Hyuntak Cha, Sungsoo Ahn, Jaeho Lee, and Jinwoo Shin. 2020. [Learning from failure: De-biasing classifier from biased classifier](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 20673–20684. Curran Associates, Inc.
- Mohammad Pezeshki, Oumar Kaba, Yoshua Bengio, Aaron C Courville, Doina Precup, and Guillaume Lajoie. 2021. [Gradient starvation: A learning proclivity in neural networks](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 1256–1272. Curran Associates, Inc.

- Shiori Sagawa, Pang Wei Koh, Tatsunori B. Hashimoto, and Percy Liang. 2020. [Distributionally robust neural networks](#). In *International Conference on Learning Representations (ICLR)*.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. [Get your vitamin C! robust fact verification with contrastive evidence](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 624–643, Online. Association for Computational Linguistics.
- Tal Schuster, Darsh Shah, Yun Jie Serene Yeo, Daniel Roberto Filizzola Ortiz, Enrico Santus, and Regina Barzilay. 2019. [Towards debiasing fact verification models](#). 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 3419–3425. Association for Computational Linguistics.
- Agnieszka Słowik and Leon Bottou. 2022. [On distributionally robust optimization and data rebalancing](#). In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 1283–1297. PMLR.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: a large-scale dataset for fact extraction and VERification](#). 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 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- Vladimir Vapnik. 1992. [Principles of risk minimization for learning theory](#). In *Advances in Neural Information Processing Systems*, volume 4. Morgan-Kaufmann.
- 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. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.