

# Elastic weight consolidation for better bias inoculation

**James Thorne**

Department of Computer Science  
University of Cambridge  
jt719@cam.ac.uk

**Andreas Vlachos**

Department of Computer Science  
University of Cambridge  
av308@cam.ac.uk

## Abstract

The biases present in training datasets have been shown to affect models for sentence pair classification tasks such as natural language inference (NLI) and fact verification. While fine-tuning models on additional data has been used to mitigate them, a common issue is that of catastrophic forgetting of the original training dataset. In this paper, we show that elastic weight consolidation (EWC) allows fine-tuning of models to mitigate biases while being less susceptible to catastrophic forgetting. In our evaluation on fact verification and NLI stress tests, we show that fine-tuning with EWC dominates standard fine-tuning, yielding models with lower levels of forgetting on the original (biased) dataset for equivalent gains in accuracy on the fine-tuning (unbiased) dataset.

## 1 Introduction

A number of recent works have illustrated shortcomings in sentence-pair classification models that are used for Natural Language Inference (NLI). These arise from limited or biased training data and the lack of suitable inductive bias in models. Naik et al. (2018) demonstrated that phenomena such as the presence of negation or a high degree of lexical overlap induce misclassifications on models trained on the MultiNLI dataset (Williams et al., 2018). Poliak et al. (2018) and Gururangan et al. (2018) identified biases introduced during dataset construction that were exploited by models to learn associations between the label and only one of the two input sentences without considering the other – known as *hypothesis-only* bias.

Such biases also affect fact verification (Schuster et al., 2019), typically modeled as a text-pair classification between a claim and evidence retrieved from a trusted source (Pomerleau and Rao, 2017; Thorne et al., 2018).

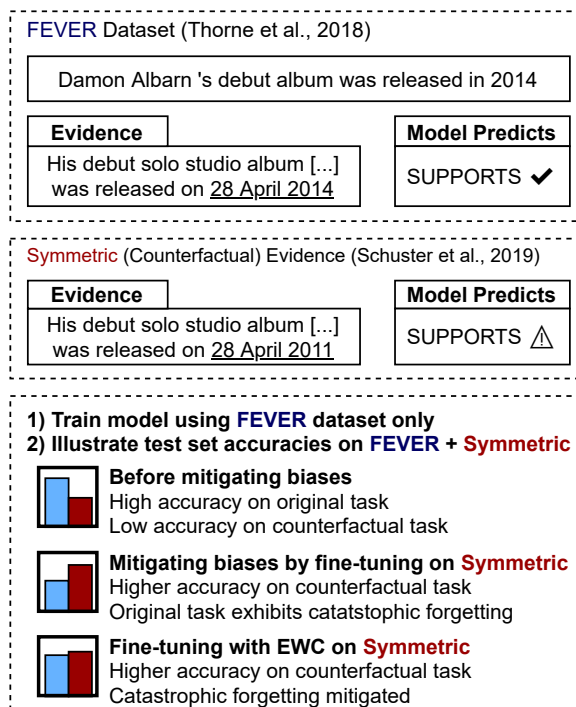


Figure 1: Hypothesis only bias in FEVER contributes to low accuracy when testing against counterfactual evidence. This is mitigated by fine-tuning on counterfactual evidence. Catastrophic forgetting from fine-tuning is reduced when using elastic weight consolidation (EWC), preserving the original task accuracy.

To mitigate these undesirable behaviors Liu et al. (2019a) fine-tune models with a small targeted number of labeled instances to “inoculate” the models. This can be contrasted to methods such as Debaised Focal Loss (DFL) and Product of Experts (POE) (Mahabadi et al., 2020) which require architectural changes to separately encode and penalize biases. In related work, Sunwal et al. (2019) delexicalize instances – replacing tokens with placeholders, preventing classifiers from exploiting mutual information between domain-specific noun-phrases and class labels. Finally, Schuster et al. (2019) re-weight the loss of in-

stances during training guided by the mutual information between tokens and the instance labels. Each of these methods corresponds to a multi-objective optimization problem where the original dataset accuracy is often sacrificed in favor of accuracy on a different evaluation set.

For fine-tuning specifically, the reduction in accuracy on the original task called catastrophic forgetting (French, 1999) as parameters for the original task are overridden during fine-tuning. In this paper, we show that regularizing fine-tuning with Elastic Weight Consolidation (Kirkpatrick et al., 2017, EWC) minimizes catastrophic forgetting by penalizing weight updates to parameters crucial to modeling the original dataset. EWC has previously been applied to machine translation, where Thompson et al. (2019) and Saunders et al. (2019) minimized catastrophic forgetting for models undergoing domain adaption. Extending this line of research further, we demonstrate that EWC can be used to mitigate biases present in two sentence-pair classification tasks without architectural changes to the models. We evaluate a number of popular model architectures, trained on MultiNLI and FEVER, and demonstrate that fine-tuning with EWC mitigates model bias while yielding lower reductions in accuracy on the original dataset compared to standard fine-tuning.

On all experiments on the FEVER dataset, fine-tuning with symmetric counterfactual data (Schuster et al., 2019) mitigated hypothesis-only bias, increasing the absolute accuracy of a BERT model by approximately 10%. Without EWC, accuracy on the original dataset was reduced from 87% to 79%, whereas with EWC catastrophic forgetting was mitigated and accuracy was 82%. In all experiments with EWC, the original task accuracy was significantly higher than fine-tuning without regularization and fine-tuning with L2 regularization. These gains were attained while maintaining similar performance on the fine-tuning data. Plotting the Pareto frontier, we show that equivalent gains in accuracy can be made with less forgetting of the original dataset. Furthermore, we demonstrate that fine-tuning methods can be combined with POE and DFL, yielding improvements on both the original task as well as fine-tuning data used for debiasing. Similar patterns were observed when fine-tuning MultiNLI models with lexical bias evaluation datasets (Naik et al., 2018).

## 2 Mitigating biases with fine-tuning

Fine-tuning broadly refers to approaches where a model is initially trained on one dataset and then further improved by training on another. We refer to these datasets as fine-tuning training and test data as *FT-train* and *FT-test* respectively. This technique is commonly used to mitigate model biases (Liu et al., 2019a), where the original data, while useful in model training, often contain biases, which are addressed by further training the model on a small set of instances targeting these biases. Fine-tuning to mitigate bias, however, can result in model parameters over-adjusting to the instances targeting it, reducing the accuracy on the original task, referred to as catastrophic forgetting (French, 1999). To ameliorate this issue, one can regularize the parameter updates so that they do not deviate too much from the original training, similar to the intuition behind multi-task training approaches (Ruder, 2017).

Elastic Weight Consolidation (Kirkpatrick et al., 2017, EWC) penalizes parameter updates according to the model’s sensitivity to changes in these parameters. The model sensitivity is estimated by the Fisher information matrix, which describes the model’s expected sensitivity to a change in parameters, and near the (local) minimum of the loss function used for training is equivalent to the second-order derivative  $F = \mathbb{E}_{(x,y) \sim \mathcal{D}_{original}} [\nabla^2 \log p(y|x; \theta)]$ . When fine-tuning with EWC (which we refer to as FT+EWC), the Fisher information is used to elastically scale the cost of updating parameters  $\theta_i$  from the original value  $\theta_i^*$ , controlled by the  $\lambda$  hyper-parameter, as follows:

$$\mathcal{L}(\theta) = \mathcal{L}_{FT}(\theta) + \sum_i \frac{\lambda}{2} F_{i,i} (\theta_i - \theta_i^*)^2 \quad (1)$$

For efficiency, we use the *empirical Fisher* (Martens, 2014): diagonal elements are approximated through squaring first-order gradients from a sample of instances, recomputed before each epoch. If the Fisher information is not used (i.e.  $F_{i,i} = 1$ ), Eq. 1 is equivalent to L2 regularization (which we refer to as FT+L2).

## 3 Experimental setup

We assess the application of EWC to minimize catastrophic forgetting when mitigating model biases in the context of two sentence-pair classi-

fication tasks: fact verification and natural language inference. We compare the untreated model (original), fine-tuning (FT), FT+EWC, FT+L2, and merging instances from the FT-train dataset when training (Merged). Each model is first trained using the original dataset and splits from the respective task, using the AllenNLP implementations (Gardner et al., 2017) with default hyper-parameters and tokenize with SpaCy or pre-trained transformer tokenizers. We train five random initializations of each model, reporting the mean accuracy, standard deviation, and  $p$ -value with an unpaired  $t$ -test. For fine-tuning, the learning rate, regularization strength  $\lambda$ , and number of epochs are selected through 5-fold cross-validation on the FT-train data, selecting the model with the highest FT-train accuracy. 30 hyper parameter choices were evaluated with grid search over 10 choices for regularization strength between  $10^6$  and  $10^8$  and 3 choices of learning rates in  $\{2 \cdot 10^{-6}, 4 \cdot 10^{-6}, 6 \cdot 10^{-6}\}$  for transformer models and  $\{2 \cdot 10^{-4}, 4 \cdot 10^{-4}, 6 \cdot 10^{-4}\}$  for ESIM models. We trained the models for a max of 8 epochs. For the transformer-based models, the highest cross validation accuracy on the FT-train dataset was achieved with  $LR = 4 \cdot 10^{-6}$ ,  $\lambda = 10^7$  and 6 epochs. For the ESIM-based models, the highest FT-train accuracy was achieved with  $LR = 2 \cdot 10^{-4}$ ,  $\lambda = 10^7$  and 5 epochs. Full hyper-parameter choices are in Appendix D.

**Mitigating hypothesis only bias in fact verification:** The FEVER<sup>1</sup> task (Thorne et al., 2018) is to predict whether short factoid sentences called claims are Supported or Refuted against evidence (in the form of sentences from Wikipedia) or whether there is not enough info (NEI). When training the models, the NEI instances by definition don’t have evidence: we sample negative instances for these with random sentences from the Wikipedia page closest to the claim using TF-IDF. This preprocessing is the same as Thorne et al. (2018). Schuster et al. (2019) identified a bias where the label for some claims can be predicted without the need for evidence. To evaluate this bias, they released<sup>2</sup> a set of 1420 symmetric counterfactual instances where each claim is supported by one Wikipedia passage and refuted by another (approximately 1% of the FEVER dataset).

<sup>1</sup><https://fever.ai/>

<sup>2</sup><https://github.com/TalSchuster/FeverSymmetric/>

This is mitigating the claim-only bias by reducing the mutual information between claims and labels. The availability of counterfactual data meant that it is possible to experiment with fine-tuning as a mitigation strategy, using the published dev and test data as FT-train and FT-test respectively. Following Schuster et al. (2019)’s evaluation, we train two ESIM (Enhanced LSTM) variants (Chen et al., 2016), and a BERT (Devlin et al., 2019) transformer. We also evaluate RoBERTa (Liu et al., 2019b), as it has been shown to be more robust to adversarial testing (Bartolo et al., 2020).

**Mitigating model limitations in NLI stress tests:** The MultiNLI<sup>3</sup> task (Williams et al., 2018) requires systems to predict whether a hypothesis is entailed by a premise. Naik et al. (2018) identify limitations of models trained on this dataset where 6 phenomena such as *lexical overlap*, *numerical reasoning* and presence of *antonyms* were evaluated with ‘stress-tests’. In this paper, we report on *antonyms* and *numerical reasoning* as these stress-tests exhibited catastrophic forgetting when used to fine-tune models (Liu et al., 2019a). To this end, we do not evaluate on HANS (McCoy et al., 2020), as high accuracies can be attained without forgetting.

Like Liu et al. (2019a), we mitigate these biases through fine-tuning both an ESIM (Chen et al., 2016) model on stress-test data<sup>4</sup>. Each contains a small number of procedurally generated instances (between 1500-9800) that specifically target one of these phenomena. We evaluate FT and FT+EWC using the same methodology, controlling the number of instances, sampled at random, in FT-train and report the change in accuracy on the FT-test and original test sets.

## 4 Results

### 4.1 Fact verification

Fine-tuning the models, rather than merging datasets, yielded the greatest improvements in accuracy on FT-test. All improvements from the untreated model were significant ( $p < 0.05$ , denoted #). Without L2 or EWC, catastrophic forgetting occurs due to the shift in label distribution between the FEVER and FT-train dataset, which only contains 2 of the original 3 label classes.

<sup>3</sup><https://cims.nyu.edu/~sbowman/multinli/>

<sup>4</sup>[https://abhilasharavichander.github.io/NLI\\_StressTest/](https://abhilasharavichander.github.io/NLI_StressTest/)

Model	FEVER Dataset (Original Task) Accuracy (%)				
	Original	Merged	FineTune	FineTune+L2	FineTune+EWC
ESIM+GloVe	79.94 ± 0.4	79.57 ± 0.4	70.78 ± 1.1	73.29 ± 0.4*	74.64 ± 0.7*†
ESIM+ELMo	80.15 ± 0.2	80.33 ± 0.8	76.45 ± 0.8	73.72 ± 0.6*	78.09 ± 0.4*†
BERT Base	86.88 ± 0.5	86.87 ± 0.5	78.82 ± 0.9	79.90 ± 1.4*	82.23 ± 1.1*†
RoBERTa Base	88.12 ± 0.3	88.11 ± 0.1	82.51 ± 1.5	83.14 ± 1.4*	85.12 ± 1.1*†
Symmetric Dataset (Fine-tuning Task) Accuracy (%)					
ESIM+GloVe	68.37 ± 1.0	69.35 ± 0.5	74.21 ± 1.3#	73.34 ± 1.2#◇	73.20 ± 1.4#◇♡
ESIM+ELMo	64.04 ± 0.7	66.46 ± 1.3#	68.68 ± 0.7#	70.31 ± 0.5#	69.16 ± 0.7#
BERT Base	74.77 ± 1.4	79.24 ± 0.7#	87.07 ± 0.6#	86.66 ± 0.4#◇	85.11 ± 0.4#
RoBERTa Base	78.34 ± 0.2	87.03 ± 2.3#	91.01 ± 0.6#	90.98 ± 0.5#◇	89.63 ± 1.3#◇♡

Table 1: Bias mitigation for FEVER classifiers comparing no treatment (original), against merging from instances from the FT-train with the original task training dataset (Merged) and FineTuning (with EWC and L2). Improvements  $p < 0.05$  are marked with the following symbols: \* against FT, † against FT+L2, # against original. Deteriorations  $p > 0.05$  on the symmetric dataset are marked with ◇ against FT and ♡ against FT+L2

Both L2 and EWC reduced catastrophic forgetting. Improvements on the original task are significant (denoted \*) compared to FT. However, EWC regularization retained more of the original task accuracy than L2 for all models; this was also significant (denoted †). In all cases, there is a trade-off between original and fine-tuning task accuracies. With regularization, the FT accuracy was higher than FT+L2 and FT+EWC (with the exception of ESIM+ELMo). However, the deterioration from FT when regularizing was not significant ( $p > 0.05$ , denoted ◇). Furthermore, for the highest performing model (RoBERTa base), the deterioration of using FT+EWC against FT+L2 was also not significant (denoted ♡).

Training a model where the FEVER training and FT-train were merged yielded modest improvements on the FT-test without harming the original FEVER task accuracy. We attribute this to the impact of these 700 instances being diluted by the large number of training instances in FEVER (FT-train is <1% the size of FEVER).

## 4.2 Combining FT and bias modeling

Fine-tuning can be applied to any task using a small amount of bias-mitigating labeled data, whereas explicit modeling of hypothesis-only biases (Mahabadi et al., 2020) requires architectural changes that are specific to the task and model. We consider two techniques from Mahabadi et al. (2020): Product of Experts (PoE), which multiplies the sentence-pair label probabilities with hypothesis-only label probabilities, and Debaised

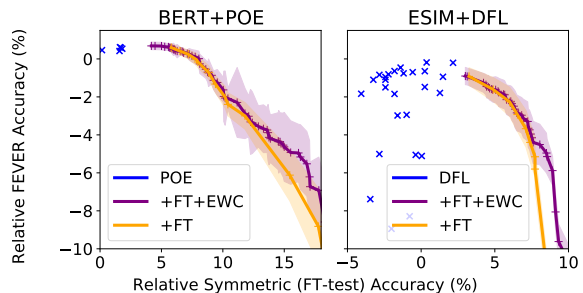


Figure 2: Pareto frontiers of fine-tuning BERT and ESIM models showing FT+EWC dominates FT.

Focal Loss (DFL) which explicitly modulates the loss of instances according to the accuracy of a hypothesis-only classifier. For both ESIM and BERT, accuracy, when trained with PoE, was stable across different choices of hyper-parameters, whereas some hyper-parameter choices for DFL resulted in lower accuracy on both tasks. We report results for BERT+PoE and ESIM+DFL as these were best.

The trade-off between accuracy on the original and FT-test datasets is visualized in Figure 2 indicating the change in accuracy for both bias modeling techniques in isolation, as well as in combination with fine-tuning.<sup>5</sup> This further shows that FT+EWC Pareto dominates FT for both the ESIM and BERT model. With EWC, equivalent gains on the symmetric FT-test can be attained with a lower

<sup>5</sup>All Pareto frontier lines are the result of sweeping the following: for POE,  $\beta$ ; for DFL,  $\beta$  and  $\gamma$ ; for FT, learning rate; and for FT+EWC, learning rate + EWC strength ( $\lambda$  in Equation 1) in Appendix D.8.



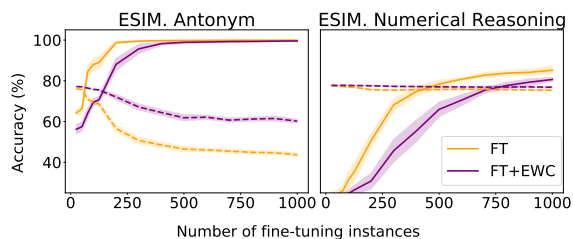


Figure 3: Training curves fine tuning MultiNLI with stress-test data. Solid lines indicate challenge dataset accuracy. Dashed lines indicate MultiNLI accuracy.

reduction in accuracy on the original FEVER task.

### 4.3 Natural Language Inference (NLI)

In a separate experiment, we apply EWC to a different domain. We inoculate biases on the ESIM model for natural language inference reported in Liu et al. (2019a). For both *Antonym* and *Numerical Reasoning* challenges, MultiNLI accuracy was higher with FT+EWC compared to FT (dashed lines in Figure 3).

**Antonym challenge:** The ESIM model was sensitive to fine-tuning, attaining near perfect accuracy (top row of Figure 3) on the FT-test data. The antonym stress-test only contains instances labeled *contradiction*: a change in label distribution that causes catastrophic forgetting. Without EWC, accuracy on MultiNLI fell to just above chance levels as the model learned only to predict contradiction (yellow dashed line). However, using an appropriate EWC penalty attained near-perfect accuracy with a smaller reduction in MultiNLI accuracy (purple dashed line).

**Numerical reasoning challenge:** The ESIM model was sensitive to fine-tuning to introduce numerical reasoning behaviours to the model. As the difference in label distribution in the inoculation dataset was less severe than the Antonym dataset, the catastrophic forgetting was less severe. Nevertheless, FT+EWC minimized catastrophic forgetting at the expense of reducing sample efficiency: accuracy on MultiNLI fell from 77.9% to 75.4% without EWC and 76.8% with EWC.

## 5 Conclusions

Fine-tuning can be used to mitigate model bias but has the risk that the model catastrophically forgets the data it was originally trained on. Incorporating elastic weight consolidation (EWC) when fine-tuning minimizes catastrophic forgetting, yielding

higher accuracy on the original task. We show this holds for both the NLI stress-tests, as well debiasing fact-verification systems (Schuster et al., 2019) through fine-tuning.

## Acknowledgements

James Thorne is funded by an Amazon Alexa Graduate Fellowship. Andreas Vlachos is supported by the ERC grant AVeriTeC. Thank you to Danielle Saunders for the insightful discussions.

## References

- Tariq Alhindi, Savvas Petridis, and Smaranda Muresan. 2018. [Where is Your Evidence: Improving Fact-checking by Justification Modeling](#). In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 85–90, Brussels, Belgium. Association for Computational Linguistics.
- Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. [MultiFC: A Real-World Multi-Domain Dataset for Evidence-Based Fact Checking of Claims](#). In *Proceedings of EMNLP-IJCNLP*, Hong Kong, China. Association for Computational Linguistics.
- Max Bartolo, Alastair Roberts, Johannes Welbl, Sebastian Riedel, and Pontus Stenetorp. 2020. [Beat the AI: Investigating Adversarial Human Annotations for Reading Comprehension](#).
- Qian Chen, Xiaodan Zhu, Zhenhua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2016. [Enhanced LSTM for Natural Language Inference](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, page 1657–1668, Vancouver, Canada.
- 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)*, Minneapolis, Minnesota.
- Robert M French. 1999. Catastrophic forgetting in connectionists networks. *Trends in Cognitive Sciences*, 3(4):128–135.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2017. AllenNLP: A Deep Semantic Natural Language Processing Platform.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R Bowman, and Noah A Smith. 2018. [Annotation Artifacts in Natural Language Inference Data](#).

- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2017. [Overcoming catastrophic forgetting in neural networks](#). *Proceedings of the National Academy of Sciences of the United States of America*, 114(13):3521–3526.
- Nelson F. Liu, Roy Schwartz, and Noah A. Smith. 2019a. [Inoculation by Fine-Tuning: A Method for Analyzing Challenge Datasets](#). pages 2171–2179.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#).
- Rabeeh Karimi Mahabadi, Yonatan Belinkov, and James Henderson. 2020. [End-to-End Bias Mitigation by Modelling Biases in Corpora](#). In *Proceedings of the 58th meeting of the Association for Computational Linguistics*, Seattle, WA. Association for Computational Linguistics.
- James Martens. 2014. [New insights and perspectives on the natural gradient method](#).
- Thomas R. McCoy, Ellie Pavlick, and Tal Linzen. 2020. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *ACL 2019 - 57th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, 2*, pages 3428–3448.
- Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. [Stress Test Evaluation for Natural Language Inference](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353, Santa Fe.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. [Hypothesis Only Baselines in Natural Language Inference](#). In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.
- Dean Pomerleau and Delip Rao. 2017. [Fake News Challenge](#). <http://fakenewschallenge.org/>.
- Sebastian Ruder. 2017. [An Overview of Multi-Task Learning in Deep Neural Networks](#). May.
- Danielle Saunders, Felix Stahlberg, Adrià de Gispert, and Bill Byrne. 2019. [Domain Adaptive Inference for Neural Machine Translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 222–228, Florence, Italy. Association for Computational Linguistics.
- Tal Schuster, Darsh J Shah, Yun Jie Serene Yeo, Daniel Filizzola, Enrico Santus, and Regina Barzilay. 2019. [Towards Debiasing Fact Verification Models](#). In *Proceedings of EMNLP-IJCNLP*, Hong Kong, China.
- Sandeep Sunawal, Mithun Paul, Rebecca Sharp, and Mihai Surdeanu. 2019. [On the Importance of Delexicalization for Fact Verification](#). pages 3411–3416.
- Brian Thompson, Jeremy Gwinnup, Huda Khayrallah, Kevin Duh, and Philipp Koehn. 2019. [Overcoming Catastrophic Forgetting During Domain Adaptation of Neural Machine Translation](#). (2018):2062–2068.
- 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.
- 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.

## A Hypothesis-only bias in other fact-verification datasets

We use the FEVER dataset in this paper due to the availability of the symmetric counterfactual data released by [Schuster et al. \(2019\)](#). Other datasets for fake news detection and fact verification also have the same task signature (sentence-pair classification with a claim and another sentence), but second sentence is not evidential. All tasks exhibit a hypothesis-only bias where information from the claim can be used to predict the label without considering the second sentence in the sentence-pair. If there was no mutual information between claims (without evidence) and labels, this should be 33%.

For FEVER, this bias is introduced through synthetic generation of the claims and is more problematic than the biases that occur in the datasets consisting of naturally occurring claims. In Liar and MultiFC, the claims arise from real-world events and the biases in the data reflect political viewpoints, rather than cognitive shortcuts taken by the FEVER annotators.

The RoBERTa model, trained on only the claims outperforms the sentence-pair setup for

both the MultiFC (Augenstein et al., 2019) and Liar-Plus tasks (Alhindi et al., 2018) whereas for the the FEVER data the sentence-pair accuracy is higher. As FEVER is the only task that requires the use of evidence for classification, this is expected.

Dataset	Accuracy (%)	
	Claim Only	Sentence Pair
Liar-Plus	28.74	20.48
Liar-Plus (binary)	72.59	70.48
MultiFC	46.02	44.83
FEVER	61.50	88.93
FEVER (2-way)	79.09	92.24

Table 2: Validation accuracy for claim-only vs sentence pair classification for fact verification datasets trained on RoBERTa. For 2-way fever we discard instances labelled NOTENOUGHINFO. For binary Liar-Plus, we map all positive labels to true and all negative labels to false and discard neutral instances.

## B Compute infrastructure

All experiments were performed on a single workstation with a single Xeon E5-2630 CPU, 64GB RAM and an Nvidia 1080Ti GPU.

## C Average run-time for fine-tuning

Estimating the Fisher matrix diagonal takes approximately 2 seconds for the ESIM model using 2000 instances sampled from the original dataset. Average training duration (excluding estimating Fisher information) was 11 seconds in total with 700 instances.

Estimating the Fisher matrix diagonal takes approximately 25 seconds for the BERT and RoBERTa models using 2000 instances sampled from the original dataset. Average training duration (excluding estimating Fisher information) was 2 minutes 20 seconds in total with 700 instances.

## D Hyper-parameter configurations

### D.1 Base models

For the base-models, the default hyperparameters in AllenNLP are used for the ESIM, BERT and RoBERTa models.

#### ESIM

- Embedding dimension: 300, bidirectional

- Dropout: 0.5
- Optimizer: Adam
- Gradient Norm: 10.0
- Batch Size: 64
- Learning Rate: 0.0004
- Learning Rate Schedule: reduce on plateau, patience 0, factor 0.5
- Number of Epochs: 75
- Early Stopping: Patience 10

#### BERT+RoBERTa

- Embedding dimension: 768
- Optimizer: AdamW
- Gradient Norm: 10.0
- Batch Size: 8
- Learning Rate: 0.0004
- Learning Rate Schedule: slanted triangular, cut frac 0.06
- Number of Epochs: 5
- Early Stopping: Patience 0

## D.2 Fine-tuning without EWC

### ESIM

- FT Learning Rate: 0.0002
- FT Epochs: 8

### BERT

- FT Learning Rate: 0.000004
- FT Epochs: 6

### RoBERTa

- FT Learning Rate: 0.000004
- FT Epochs: 7

## D.3 Fine-tuning using EWC

### ESIM

- FT Learning Rate: 0.0002
- FT Epochs: 5
- EWC: 10000000

### BERT+RoBERTa

- FT Learning Rate: 0.000004
- FT Epochs: 6
- EWC: 10000000

## D.4 Unsupervised bias mitigation

Hyperparameters:  $\beta$  controls the weight update for the hypothesis-only model and  $\gamma$  controls the modulation of hypothesis-only model in the loss function.

- POE (BERT)  $\beta = 0.4$
- POE (ESIM)  $\beta = 0.05$
- DFL (BERT)  $\beta = 0.4, \gamma = 0.6$
- DFL (ESIM)  $\beta = 0.05, \gamma = 2.0$

## D.5 Search bounds for fine-tuning

Approximately 30 configurations were considered (Cartesian product of LR+EWC). The best performing system was selected through max accuracy on the FT cross validation dataset through 5 fold cross validation.

- EWC  $\{10^6, 2 \cdot 10^6, 4 \cdot 10^6, 8 \cdot 10^6, 10^7, 2 \cdot 10^7, 4 \cdot 10^7, 6 \cdot 10^7, 8 \cdot 10^7, 10^8\}$
- Fine-tuning learning rate (ESIMs):  $\{0.0002, 0.0004, 0.0006\}$
- Fine-tuning learning rate (Transformer):  $\{0.000002, 0.000004, 0.000006\}$
- Epochs: Up to 10 epochs cross valudating on the FT-training dataset

## D.6 Search bounds for end2end bias mitigation

We used the same range of values published by [Mahabadi et al. \(2020\)](#). For DFL, we performed a grid search of every value totalling 30 trials.

$$\gamma \in \{0.02, 0.05, 0.1, 0.6, 2.0, 4.0, 5.0\}$$
$$\beta \in \{0.05, 0.2, 0.4, 0.8, 1.0\}$$

## D.7 Stress test sizes

Following the evaluation of ([Liu et al., 2019a](#)) we vary the number of instances from the stress test (between 1500-9800) instances. To plot Figure 3, we use 25, 50, 75, 100, 250, 400, 500, 600, 700, 800, 900, 1000 instances in our evaluation.

## D.8 Pareto frontier sweeps

The Pareto frontier sweeps in Figure 2 were generated by plotting all hyperparameters from Appendix D.5. The end2end (blue) crosses are plotted by plotting the choices from Appendix D.6 in combination with the best FT model.