

Domain Adversarial Fine-Tuning as an Effective Regularizer

Giorgos Vernikos¹, Katerina Margatina², Alexandra Chronopoulou³, Ion Androutsopoulos⁴

¹ School of ECE, National Technical University of Athens, Greece

² Computer Science Department, University of Sheffield, UK

³ Center for Information and Language Processing, LMU Munich, Germany

⁴ Department of Informatics, Athens University of Economics and Business, Greece

gvernikos@mail.ntua.gr, k.margatina@sheffield.ac.uk, achron@cis.lmu.de
ion@aueb.gr

Abstract

In Natural Language Processing (NLP), pre-trained language models (LMs) that are transferred to downstream tasks have been recently shown to achieve state-of-the-art results. However, standard fine-tuning can degrade the general-domain representations captured during pretraining. To address this issue, we introduce a new regularization technique, AFTER; domain Adversarial Fine-Tuning as an Effective Regularizer. Specifically, we complement the task-specific loss used during fine-tuning with an adversarial objective. This additional loss term is related to an adversarial classifier, that aims to discriminate between *in-domain* and *out-of-domain* text representations. In-domain refers to the *labeled* dataset of the task at hand while out-of-domain refers to *unlabeled* data from a different domain. Intuitively, the adversarial classifier acts as a regularizer which prevents the model from overfitting to the task-specific domain. Empirical results on various natural language understanding tasks show that AFTER leads to improved performance compared to standard fine-tuning.

1 Introduction

Current research in NLP focuses on transferring knowledge from a language model (LM), pre-trained on large general-domain data, to a target task. The LM representations are transferred to the target task either as additional features of a task-specific model (Peters et al., 2018), or by fine-tuning (Howard and Ruder, 2018; Devlin et al., 2019; Yang et al., 2019). Standard fine-tuning involves initializing the target model with the pre-trained LM and training it with the target data.

Fine-tuning, however, can lead to catastrophic forgetting (Goodfellow et al., 2013), if the pre-trained LM representations are adjusted to such an extent to the target task, that most generic knowledge, captured during pretraining, is in effect for-

gotten (Howard and Ruder, 2018). A related problem of fine-tuning is overfitting to the target task, that often occurs when only a small number of training examples is available (Dai and Le, 2015).

Adversarial training is a method to increase robustness and regularize deep neural networks (Goodfellow et al., 2015; Miyato et al., 2017). It has been used for domain adaptation (Ganin et al., 2016) to train a model from scratch to produce representations that are invariant to different domains. Inspired by this approach, we propose a regularization technique for the fine-tuning process of a pre-trained LM, that aims to optimize knowledge transfer to the target task and avoid overfitting.

Our method, domain Adversarial Fine-Tuning as an Effective Regularizer (AFTER) extends standard fine-tuning by adding an adversarial objective to the task-specific loss. We leverage out-of-domain *unlabeled* data (i.e. from a different domain than the target task domain). The transferred LM is fine-tuned so that an adversarial classifier cannot discriminate between text representations from in-domain and out-of-domain data. This loss aims to regularize the extent to which the model representations are allowed to adapt to the target task domain. Thus, AFTER is able to preserve the general-domain knowledge acquired during the pretraining of the LM, while adapting to the target task.

Our contributions are: (1) We propose AFTER, an LM fine-tuning method that aims to avoid catastrophic forgetting of general-domain knowledge, acting as a new kind of regularizer. (2) We show that AFTER improves the performance of standard fine-tuning in four natural language understanding tasks from the GLUE benchmark (Wang et al., 2019a), with two different pre-trained LMs: BERT (Devlin et al., 2019), and XLNET (Yang et al., 2019). (3) We further conduct an ablation study to provide useful insights regarding the key factors of the proposed approach.

2 Related Work

Several approaches have been proposed for the adaptation of a model trained on a domain \mathcal{D}_S to a different domain \mathcal{D}_T , where no labeled data is available (Grauman, 2012; Tzeng et al., 2014; Sun et al., 2016). Ganin et al. (2016) were the first to propose adversarial training for *domain adaptation*. They introduced a gradient reversal layer to adversarially train a classifier that should not be able to discriminate between \mathcal{D}_S and \mathcal{D}_T , in image classification and sentiment analysis tasks.

Various adversarial losses have been used for domain adaptation in several NLP tasks, such as question answering (Lee et al., 2019), machine reading comprehension (Wang et al., 2019b) and cross-lingual named entity recognition (Keung et al., 2019). Adversarial approaches have been also used to learn latent representations that are agnostic to different attributes of the input text, such as language (Lample et al., 2018a,b) and style (Yang et al., 2018). Contrary to previous *domain adaptation* work, we explore the addition of an adversarial loss term to serve as a *regularizer* for fine-tuning.

Other variants of LM fine-tuning include a supplementary supervised training stage in data-rich tasks (Phang et al., 2018) or multi-task learning with additional supervised tasks (Liu et al., 2019). However, such methods require additional *labeled* data. A common way to leverage *unlabeled* data during fine-tuning is through an additional stage of language modeling. For this stage, the unlabeled data can either come from the task-specific dataset (i.e. the labels are dropped and language modelling is performed on the input data) (Howard and Ruder, 2018), or additional unlabeled in-domain corpora (Sun et al., 2019; Gururangan et al., 2020). This approach adds a computationally expensive step that requires unlabeled data from a specific source. By contrast, our method leverages out-of-domain data with only a small computational overhead and minimal changes to the fine-tuning process.

Our work is compatible with the semi-supervised learning paradigm (Chapelle et al., 2010) that combines learning from both labeled and unlabeled data. In this setting, unlabeled data from the task domain is leveraged using a consistency loss which enforces invariance of the output given small perturbations of the input (Miyato et al., 2017; Clark et al., 2018). The adversarial loss term of AFTER can be interpreted as a consistency loss that ensures invariance of representations across domains.

Recently, adversarial or trust region based approaches (Zhu et al., 2020; Jiang et al., 2020; Aghajanyan et al., 2020) have been proposed as an extension to the LM fine-tuning process. These methods introduce constraints that prevent aggressive updating of the pretrained parameters or enforce smoothness during fine-tuning. However, these approaches require additional forward and backward computations while our method is more computationally efficient and can be implemented with minimal changes to the fine-tuning procedure.

3 Proposed Approach

Fig. 1 provides a high-level overview of AFTER.

Problem Definition. We tackle a `Main` task, with a labeled dataset from domain \mathcal{D}_M . We further exploit an existing *unlabeled corpus*, `Auxiliary`, that comes from a different domain \mathcal{D}_{AUX} . We label each sample with the corresponding domain label y_D , $y_D = 0$ for samples from `Main`, and $y_D = 1$ for samples from `Auxiliary`. We note that we *do not* use any real labels from `Auxiliary` (if there are any). The domain labels are used to train a classifier that discriminates between \mathcal{D}_M and \mathcal{D}_{AUX} .

Model. We initialize our model with pretrained weights from a top-performing language model, such as BERT (Devlin et al., 2019) or XLNET (Yang et al., 2019). The representation of both BERT and XLNET for the input sequence is encoded in the `[CLS]` token output embedding. We add a linear layer on top of the sequence representation (`[CLS]` output embedding) for the `Main` task, resulting in a task-specific loss L_{Main} . We also add another linear layer for the binary domain classifier (Figure 1), with a corresponding loss L_{Domain} , which has the same input.

Adversarial Fine-tuning. The domain discriminator outputs a domain label for each sample of the training set. We seek representations that are both discriminative for the `Main` task and indiscriminative for the domain classifier. Hence, we minimize L_{Main} and at the same time maximize L_{Domain} , by fine-tuning the pretrained LM with the joint loss:

$$\mathcal{L}_{AFTER} = L_{Main} - \lambda L_{Domain} \quad (1)$$

where λ ($\lambda > 0$) controls the importance of the domain loss. The parameters of the domain classifier are trained to predict the (true) domain label, while the rest of the network is trained to mislead the domain classifier, thereby developing domain-independent internal representations.




Figure 1: Illustration of the proposed approach, AFTER. The task-specific classifier leverages the labeled data from the downstream task (Main) while the domain classifier uses unlabeled data from both Main and Auxiliary datasets as well as the created domain labels.

Gradient Reversal Layer. We use a Gradient Reversal Layer (GRL) (Ganin et al., 2016) between the [CLS] output embedding and the domain discriminator layer, as shown in Figure 1, to maximize L_{Domain} . During the forward pass, GRL acts as an identity transform, but during backpropagation, GRL *reverses* the gradients. In effect, the pretrained LM parameters are updated towards the opposite direction of the gradient of L_{Main} and, adversarially, towards the direction of the gradient of L_{Domain} .

4 Experiments

Datasets. We experiment with four Main datasets from the GLUE benchmark (Wang et al., 2019a). The chosen datasets represent the broad variety of natural language understanding tasks, such as linguistic acceptability (CoLA) (Warstadt et al., 2019), sentiment analysis (SST-2) (Socher et al., 2013), paraphrase detection (MRPC) (Dolan and Brockett, 2005) and textual entailment (RTE) (Dagan et al., 2005; Bar-Haim et al., 2006; Giampiccolo et al., 2007; Bentivogli et al., 2009). The datasets used represent both high (SST-2) and low-resource (RTE, CoLA, MRPC) tasks, as well as single-sentence (CoLA, SST-2) and sentence-pair (MRPC, RTE) tasks. For Auxiliary data we select corpora from various domains. For the NEWS domain we use the AG NEWS dataset (Zhang et al., 2015) and for the REVIEWS domain we use a part of the Electronics reviews of He and McAuley (2016). For the LEGAL domain we use the English part of EUROPARL (Koehn, 2004) and for the MEDICAL domain we use papers from PubMed, provided by Cohan et al. (2018). We also use math questions from the dataset of Saxton et al. (2019) for the MATH domain. Table 1 summarizes all datasets. More details regarding the selection and processing of the datasets can be found in Appendix A.1.

Baselines. We compare our approach (AFTER) with the standard fine-tuning (SFT) scheme of the

DATASET	DOMAIN	N_{train}
Main		
CoLA	Miscellaneous	8.5K
SST-2	Movie Reviews	67K
MRPC	News	3.7K
RTE	News, Wikipedia	2.5K
Auxiliary		
AG NEWS	Agricultural News (NEWS)	120K
EUROPART	Legal Documents (LEGAL)	120K
AMAZON	Electronics Reviews (REVIEWS)	120K
PUBMED	Medical Papers (MEDICAL)	120K
MATHEMATICS	Mathematics Questions (MATH)	120K

Table 1: Datasets used. N_{train} denotes the number of training examples. The indicator (DOMAIN) summarizes the domain of each Auxiliary dataset.

pretrained LMs. As our baselines we fine-tune two pretrained LMs (BERT-BASE and XLNET-BASE), using the suggested hyperparameters from Devlin et al. (2019) and Yang et al. (2019) respectively.

Implementation Details. We base our implementation on Hugging Face’s Transformers library (Wolf et al., 2019) in PyTorch (Paszke et al., 2019). We make our code publicly available¹. We tune the λ hyperparameter of Eq. 1 on the validation set for each experiment, finding that most values of λ improve over the baseline. We fine-tune each model for 4 epochs and evaluate the model 5 times per epoch, as suggested by Dodge et al. (2020). We select the best model based on the validation loss. For more implementations details see Appendix A.2.

5 Results

Table 2 shows the results on the validation sets of the four GLUE datasets for the two pretrained LMs. We compare the two baselines with AFTER using the Auxiliary data from Table 1. We do not report results on CoLA with XLNET (–) because the model demonstrated degenerate performance with the available resources for the batch size (see Appendix A.2 for more details).

¹<https://github.com/GeorgeVern/AFTERV1.0>

	CoLA	SST-2	MRPC	RTE
	<i>Matthews corr.</i>	<i>Accuracy</i>	<i>Accuracy / F1</i>	<i>Accuracy</i>
BERT SFT	55.5 ± 3.2	92.0 ± 0.5	85.4 ± 1.1 / 89.6 ± 0.6	64.3 ± 3.1
AFTER w/ NEWS	57.3 ± 1.5	92.5 ± 0.4	87.5 ± 1.7 / 91.1 ± 1.2	64.7 ± 1.9
AFTER w/ REVIEWS	<u>57.1</u> ± 1.2	92.4 ± 0.3	86.4 ± 0.3 / 90.1 ± 0.4	64.6 ± 0.8
AFTER w/ LEGAL	55.0 ± 1.5	<u>92.4</u> ± 0.3	86.6 ± 0.6 / 90.3 ± 0.5	64.8 ± 1.9
AFTER w/ MEDICAL	<u>55.9</u> ± 2.9	92.6 ± 0.3	86.9 ± 1.3 / 90.7 ± 1.0	62.6 ± 3.4
AFTER w/ MATH	<u>56.1</u> ± 2.8	92.3 ± 0.8	87.3 ± 0.9 / 90.8 ± 0.7	62.5 ± 1.3
XLNET SFT	–	93.0 ± 0.7	86.4 ± 0.7 / 90.1 ± 0.5	64.7 ± 4.4
AFTER w/ NEWS	–	93.9 ± 0.3	87.3 ± 0.7 / 91.0 ± 0.5	63.9 ± 2.3
AFTER w/ REVIEWS	–	93.5 ± 0.3	86.9 ± 0.6 / 90.5 ± 0.5	65.1 ± 2.8
AFTER w/ LEGAL	–	<u>93.6</u> ± 0.5	87.5 ± 1.6 / 90.9 ± 1.2	64.8 ± 1.6
AFTER w/ MEDICAL	–	<u>93.3</u> ± 0.5	87.0 ± 1.1 / 90.5 ± 0.7	64.5 ± 2.1
AFTER w/ MATH	–	93.9 ± 0.4	87.3 ± 1.2 / 90.8 ± 0.9	66.1 ± 1.9

Table 2: Comparison of standard of fine-tuning (SFT) and AFTER for BERT (Top) and XLNET (Bottom). Underlined scores outperform the baseline. Best scores for each pretrained LM are shown in **bold**. We report the mean and standard deviation across five runs on the validation set.

BERT. We observe that the proposed approach (AFTER) outperforms the first baseline (BERT SFT) in all four tasks. For most of these tasks, AFTER results in improved performance with every *Auxiliary* dataset, demonstrating the robustness of our approach across domains.

Specifically, in CoLA, we observe that fine-tuning with the adversarial loss substantially outperforms standard fine-tuning. Specifically, using an *Auxiliary* dataset from the NEWS domain improves the baseline by 1.8 points. In SST-2, we notice that although standard fine-tuning achieves high accuracy, the use of AFTER still results in slight performance gains ($\sim 0.4\%$). Similar to CoLA, these improvements are consistent across *Auxiliary* datasets and often come with reduced variance, compared to SFT. In MRPC, we observe gains of 1.5 points on average in accuracy and 1.0 in F1 over SFT. Using NEWS data as *Auxiliary*, AFTER outperforms the baseline by 2.1 points in accuracy and 1.5 in F1. In RTE, the proposed approach improves upon the baseline from 64.3% to 64.8% in accuracy, using data from the LEGAL domain. However, we also observe deteriorated performance with the use of some *Auxiliary* datasets (e.g. MEDICAL, MATH). We attribute this result to the similarity between the domain of RTE (Wikipedia) and the domain of the pretraining corpus of BERT (Wikipedia and Books). We test this hypothesis in section 6.

XLNET. We observe in Table 2 that AFTER consistently outperforms standard fine-tuning for an even higher-performing LM (XLNET SFT).

Specifically, in SST-2 AFTER improves the accuracy of standard fine-tuning (SFT) by 0.6% on average and reduces variance, as well. For instance,

with the use of *Auxiliary* data from NEWS or MATH domains, AFTER results in 0.9% improvement in accuracy. In MRPC, the performance boost is also consistent across *Auxiliary* data. In particular, the use of LEGAL data leads in absolute improvement of 1.1% in accuracy and 0.8% in F1. In RTE, adversarial fine-tuning outperforms the baseline by 1.4% in accuracy. However, similar to BERT, we observe lower performance when using AFTER with some *Auxiliary* data (e.g. NEWS, MEDICAL). We attribute this performance degradation to the same reason as BERT, the similarity between the pretraining corpus domain and the target task domain (both LMs have similar pretraining corpora).

Summary. The experiments of this section reveal that AFTER can boost target task performance and reduce variance compared to standard fine-tuning across different pretrained LMs. We can therefore attribute the effectiveness of AFTER to regularization itself and not to the model architecture. We can also observe in Table 2 that the target task performance of our approach (BERT AFTER) is on par (RTE) or higher (MRPC) than using standard fine-tuning with a higher-performing pretrained LM (XLNET SFT). This finding demonstrates the effectiveness of the proposed approach and motivates the need for more effective fine-tuning schemes as a way to improve the target of pretrained LMs on downstream tasks.

6 Ablation Study

We investigate the effect of some key factors of AFTER such as the relation of the target task domain and the domain of the pretraining corpus of the

LM, the selection of Auxiliary data and the emergence of domain-invariant characteristics. For the experiments of this section we used BERT, unless otherwise stated.

LM pretraining and Task Domains. To explore why AFTER fails to improve upon the baseline on RTE, we examine if the pretrained representations are already well suited for the task (i.e. no regularization is needed). We calculate the average masked LM (MLM) loss of BERT for each Main dataset. We observe in Table 3 that SST-2 produces the largest loss which can be partially attributed to the dataset format (it contains short sentences that make the MLM task very challenging). RTE produces the lowest loss confirming our intuition regarding the similarity of the pretraining corpus of BERT and RTE. In this case, general-domain and domain-specific representations are close, rendering domain-adversarial regularization undesirable. This is also confirmed by the the vocabulary overlap between RTE and a Wikipedia corpus (Table 3). The more distant the pretraining domain of BERT is to the specific task (measured by vocabulary overlap and MLM loss), the more benefits AFTER demonstrates, confirming our intuition regarding domain-adversarial regularization.

	RTE	MRPC	CoLA	SST-2
MLM Loss	2.17	2.37	2.53	3.39
Overlap with WIKI (%)	38.3	34.0	24.0	26.1
AFTER Improvement (%)	0.8	2.5	3.2	0.7

Table 3: Masked LM loss of BERT (the lower the better), vocabulary overlap with the Wikipedia domain (WIKI) and improvement of AFTER (best) for each task.

Domain Distance. We measure the domain distance for all Main-Auxiliary pairs to evaluate how the choice of the latter affects the performance of AFTER. We represent the word distribution of each dataset using term distributions $t \in \mathcal{R}^{|V|}$ where t_i is the probability of the i -th word in the joint vocabulary V (see Appendix A.4) and calculate Jensen-Shannon (JS) divergence (Plank and van Noord, 2011). Combining the results of Table 2 and Fig. 2, no clear pattern emerges demonstrating, perhaps, our method’s robustness to domain distance. We leave a further investigation of selection criteria for the Auxiliary data for future work.

Domain-invariant vs. Domain-specific Features. To investigate if the benefits of AFTER can be attributed only to data augmentation we compare adversarial ($\lambda > 0$ in Eq. 1) and multi-task ($\lambda < 0$) fine-tuning. We experiment with MRPC and CoLA

	NEWS	REVIEWS	LEGAL	MEDICAL	MATH
CoLA	0.58	0.59	0.55	0.46	0.33
SST-2	0.62	0.66	0.62	0.55	0.34
MRPC	0.85	0.62	0.72	0.63	0.41
RTE	0.84	0.62	0.75	0.63	0.38

Figure 2: JS divergence between all dataset pairs.




Figure 3: AFTER ($\lambda > 0$) vs. MULTI-TASK ($\lambda < 0$).

for both settings (tuning each λ separately). We observe that during multi-task fine-tuning (Fig. 3), L_{Domain} is close to zero (even in the first epoch). This implies that domain classification is an easy auxiliary task, confirming our intuition that a non-adversarial fine-tuning setting favors domain-specific features. Although the multi-task approach leverages the same unlabeled data, its performance is worse than AFTER (Table 4), which highlights the need for an adversarial domain discriminator.

	CoLA	MRPC
AFTER w/ NEWS	57.3	87.5/91.1
MULTI-TASK w/ NEWS	56.5	86.7/90.5

Table 4: Comparison of AFTER vs. MULTI-TASK.

7 Conclusions and Future Work

We propose AFTER, a domain adversarial method to regularize the fine-tuning process of a pretrained LM. Empirical results demonstrate that our method can lead to improved performance over standard fine-tuning. AFTER can be widely applied to any transfer learning setting and model architecture, with minimal changes to the fine-tuning process, without requiring any additional labeled data. We aim to further explore the effect of Auxiliary data on the final performance and the use of multiple Auxiliary datasets. We also aim to extend the proposed approach as a way to fine-tune a pretrained LM to a different language, in order to produce language-invariant representations.

Acknowledgements

We thank Andrei Popescu-Belis for his valuable comments and help with an initial draft of this paper.

References

- Armen Aghajanyan, Akshat Shrivastava, Anchit Gupta, Naman Goyal, Luke Zettlemoyer, and Sonal Gupta. 2020. [Better fine-tuning by reducing representational collapse](#). *CoRR*, abs/2008.03156.
- Roy Bar-Haim, Ido Dagan, Bill Dolan, Lisa Ferro, Danilo Giampiccolo, and Bernardo Magnini. 2006. The second pascal recognising textual entailment challenge.
- Luisa Bentivogli, Ido Dagan, Hoa Trang Dang, Danilo Giampiccolo, and Bernardo Magnini. 2009. The fifth pascal recognizing textual entailment challenge. In *Proceedings of the Text Analysis Conference (TAC)*.
- Joachim Bingel and Anders Søgaard. 2017. [Identifying beneficial task relations for multi-task learning in deep neural networks](#). In *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics*, pages 164–169.
- Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien. 2010. *Semi-Supervised Learning*, 1st edition.
- Kevin Clark, Minh-Thang Luong, Christopher D. Manning, and Quoc Le. 2018. [Semi-supervised sequence modeling with cross-view training](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1914–1925.
- Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. [A discourse-aware attention model for abstractive summarization of long documents](#). 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 615–621.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. [The pascal recognising textual entailment challenge](#). In *Proceedings of the First International Conference on Machine Learning Challenges: Evaluating Predictive Uncertainty Visual Object Classification, and Recognizing Textual Entailment, MLCW’05*, page 177–190.
- Andrew M Dai and Quoc V Le. 2015. [Semi-supervised sequence learning](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 3079–3087.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *In Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186.
- Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah A. Smith. 2020. [Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping](#). *ArXiv*, abs/2002.06305.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. [Domain-adversarial training of neural networks](#). *Journal of Machine Learning Research*, 17(1):2096–2030.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. [The third PASCAL recognizing textual entailment challenge](#). In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 1–9.
- Ian J. Goodfellow, Mehdi Mirza, Xia Da, Aaron C. Courville, and Yoshua Bengio. 2013. [An empirical investigation of catastrophic forgetting in gradient-based neural networks](#). *CoRR*, abs/1312.6211.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. [Explaining and harnessing adversarial examples](#). In *Proceedings of the International Conference on Learning Representations*.
- Kristen Grauman. 2012. [Geodesic flow kernel for unsupervised domain adaptation](#). In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 2066–2073.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). *ArXiv*, abs/2004.10964.
- Ruining He and Julian McAuley. 2016. [Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering](#). In *Proceedings of the International Conference on World Wide Web, WWW ’16*, pages 507–517.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 328–339.

- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2020. [SMART: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2177–2190.
- Phillip Keung, Yichao Lu, and Vikas Bhardwaj. 2019. [Adversarial learning with contextual embeddings for zero-resource cross-lingual classification and NER](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing*, pages 1355–1360.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Philipp Koehn. 2004. Europarl: A parallel corpus for statistical machine translation. 5.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018a. [Unsupervised machine translation using monolingual corpora only](#). In *Proceedings of the International Conference on Learning Representations*.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018b. [Word translation without parallel data](#). In *Proceedings of the International Conference on Learning Representations*.
- Seanie Lee, Donggyu Kim, and Jangwon Park. 2019. [Domain-agnostic question-answering with adversarial training](#). In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 196–202.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019. [Multi-task deep neural networks for natural language understanding](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4487–4496.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). In *5th International Conference on Learning Representations, ICLR 2017, Conference Track Proceedings*.
- Takeru Miyato, Shin-Ichi Maeda, Masanori Koyama, and Shin Ishii. 2017. [Virtual adversarial training: A regularization method for supervised and Semi-Supervised learning](#).
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’ Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2227–2237.
- Jason Phang, Thibault Févry, and Samuel R. Bowman. 2018. [Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks](#). *ArXiv*, abs/1811.01088.
- Barbara Plank and Gertjan van Noord. 2011. [Effective measures of domain similarity for parsing](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1566–1576.
- David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. 2019. [Analysing mathematical reasoning abilities of neural models](#). *ArXiv*, abs/1904.01557.
- 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 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642.
- Baochen Sun, Jiashi Feng, and Kate Saenko. 2016. [Return of frustratingly easy domain adaptation](#). In *Proceedings of the Association for the Advancement of Artificial Intelligence Conference on Artificial Intelligence*, pages 2058–2065.
- Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. [How to fine-tune bert for text classification?](#) *ArXiv*, abs/1905.05583.
- Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. 2014. [Deep domain confusion: Maximizing for domain invariance](#). *CoRR*, abs/1412.3474.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019a. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *International Conference on Learning Representations*.
- Huazheng Wang, Zhe Gan, Xiaodong Liu, Jingjing Liu, Jianfeng Gao, and Hongning Wang. 2019b. [Adversarial domain adaptation for machine reading comprehension](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing*, pages 2510–2520.

- Alex Warstadt, Amanpreet Singh, and Samuel Bowman. 2019. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7(0):625–641.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#).
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’ Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 5753–5763.
- Zichao Yang, Zhiting Hu, Chris Dyer, Eric P Xing, and Taylor Berg-Kirkpatrick. 2018. [Unsupervised text style transfer using language models as discriminators](#). In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 7287–7298.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Proceedings of the Conference on Neural Information Processing Systems*, pages 649–657.
- Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu. 2020. [FreeLB: Enhanced adversarial training for natural language understanding](#). In *International Conference on Learning Representations*.

A Appendices

In this supplementary material, we provide additional information for producing the results in the paper, and results that could not fit into the main body of the paper.

A.1 Dataset Details

Main datasets. We use only four datasets of the GLUE benchmark as `Main` for our experiments, due to resources constraints. All `Main` datasets are open source and can be found in <https://gluebenchmark.com/tasks>.

Auxiliary datasets. We choose `Auxiliary` datasets that are larger than `Main`, which we consider as the most realistic scenario, given the availability of unlabeled compared to labeled data. We under-sample the `Auxiliary` dataset to ensure that the two domains are equally represented, motivated by the observation of [Bingel and Søgaard \(2017\)](#) that balanced datasets tend to be better in auxiliary tasks. For each mini-batch, we sample equally from the `Main` and `Auxiliary` datasets.

The `Auxiliary` datasets are a mixed of labeled and unlabeled datasets from different domains. The labeled `Auxiliary` datasets (e.g. AG NEWS) are handled as unlabeled corpora, by dropping the task-specific labels and using only the domain labels. Although some domains might seem similar to those of the `Main` datasets, e.g. Electronics Reviews vs. Movies reviews and Agricultural News vs. News this is not the case as can be seen in Figure 6.

The maximum sequence length for all datasets was 128, so all samples were truncated to 128 tokens and lower-cased. For EUROPARL, which contains parallel corpora in multiple languages, only the English part is used. We therefore sample 120K sentences from the English corpus. For PUBMED we use 120K abstracts from medical papers, from the dataset of [Cohan et al. \(2018\)](#). For MATH we use 120K questions of medium difficulty from the dataset of [Saxton et al. \(2019\)](#). We note that all corpora used are in English.

A.2 Hyperparameters and Model details

For BERT we use the `bert-base-uncased` pretrained model and we fine-tune it with the following hyperparameters: dropout 0.1, batch size 28 and a maximum length of 128 tokens. For the optimization we use the Adam optimizer ([Kingma and Ba, 2015](#)) with a learning rate of $2e-5$, adam ep-

silon $1e-6$ and weight decay 0.01. We use a linear warmup schedule with 0.1 warmup proportion.

For XLNET we use the `xlnet-base-cased`. We use the last hidden state output embedding, as the input sequence representation. We fine-tune XLNET with the following hyperparameters: 26 batch size and the same learning rate ($2e-5$) and sequence length (128) as BERT. We do not use weight decay or warmup. In order to replicate the results of [Yang et al. \(2019\)](#) in CoLA, the authors suggested using a considerably larger batch size ($\times 4$), which was not possible in our case, due to resources constraints².

When we combine AFTER with either BERT or XLNET we use the same hyperparameters as above. We note that both models have approximately 110M parameters and this is (almost) the same using AFTER, as well. Our approach only introduces a binary domain discriminator in the form of a linear layer.

For all experiments we used a 6G GeForce GTX 1080. The duration of the experiments depended on the datasets. For SST-2, which is the largest dataset, the experiments for the baseline (BERT, XLNET) had a runtime of approximately 100mins (for all 4 epochs) and 200mins for AFTER, due to the implicit dataset augmentation. Smaller datasets such as MRPC and CoLA had an approximate runtime of 30mins with standard fine-tuning and 60mins with AFTER.

A.3 Tuning the λ hyperparameter

We tune λ on each development set, choosing from $\{0.1, 0.01, 0.001, 0.0001\}$. In Figure 4 we compare the performance of BERT and AFTER for different `Main-Auxiliary` combinations, as we vary the value of λ .

We observe that the various values of λ can have different effect on the performance and variance of AFTER. We observe that most values of λ significantly improve the performance of the baseline, BERT and an exhaustive search is not required. Table 5 presents the values of λ that were used for the results reported in Table 2. Best values of λ were chosen based on the task-specific metric (e.g. Accuracy, Matthews correlation).

²The authors' response regarding the hyperparameters : Clarification on reported dev numbers on GLUE tasks. Similar problems have been reported for other BERT-based models on CoLA, as well: [Xlnet](#), [Alberta](#), [Roberta](#) are not finetuned for CoLA task.

	BERT				XLNET		
	CoLA	SST-2	MRPC	RTE	SST-2	MRPC	RTE
NEWS	0.1	0.1	0.1	0.001	0.01	0.1	0.0001
REVIEWS	0.1	0.01	0.1	0.001	0.001	0.0001	0.01
LEGAL	0.01	0.1	0.01	0.1	0.001	0.01	0.0001
MEDICAL	0.01	0.1	0.1	0.0001	0.1	0.1	0.01
MATH	0.001	0.1	0.001	0.001	0.0001	0.01	0.01

Table 5: Best λ value of AFTER for each experiment.




Figure 4: Performance of standard BERT fine-tuning vs. AFTER for different λ values on SST-2 (Top) and CoLA (Bottom). The errorbars correspond to one standard deviation.

A.4 More Domain Distance Results

In order to create a common vocabulary for all data for Figure 2 we find the 5k most frequent words in each dataset and we then take the union of these sub-vocabularies which results in 23k words. We also calculate the vocabulary overlap, by creating each domain (or task) vocabulary with the 10k most frequent words in each dataset (in case a dataset contains less words we use all the words in the dataset).

We then calculate the vocabulary overlap between domains (Figure 5) and between each task and all domains (Figure 6). For the latter, we also include the WIKI domain to account for the pre-training domain of BERT and XLNET. For the vocabulary of WIKI we use the WikiText-2 corpus from Merity et al. (2017). We observe in Figure 5 that most domains are dissimilar, with the exception of NEWS and LEGAL domains, that have 36.6% vocabulary overlap. In Figure 6, we observe that RTE has the most overlap in vocabulary with WIKI which is a possible cause for the deteriorated performance of AFTER, since the model has already

been pretrained in this domain and does not require further regularization, as described in Section 6.

NEWS	100.0	30.4	36.6	21.6	1.8
REVIEWS	30.4	100.0	28.0	23.2	2.3
LEGAL	36.6	28.0	100.0	27.0	1.7
MEDICAL	21.6	23.2	27.0	100.0	2.7
MATH	1.8	2.3	1.7	2.7	100.0
	NEWS	REVIEWS	LEGAL	MEDICAL	MATH

Figure 5: Vocabulary overlap (%) between domains.

	NEWS	REVIEWS	LEGAL	MEDICAL	MATH	Wiki
CoLA	22.7	20.8	20.1	13.1	0.6	24.0
SST-2	24.1	24.4	24.6	16.1	0.9	26.1
MRPC	40.7	24.6	31.3	20.3	2.7	34.0
RTE	40.6	23.3	32.6	20.1	2.5	38.3

Figure 6: Vocabulary overlap (%) between tasks and domains.