

Unsupervised Data Augmentation with Naive Augmentation and without Unlabeled Data

David Lowell

Northeastern University
lowell.d@husky.neu.edu

Brian E Howard

Sciome LLC
brian.howard@sciome.com

Zachary C. Lipton

Carnegie Mellon University
zlipton@cmu.edu

Byron C. Wallace

Northeastern University
b.wallace@northeastern.edu

Abstract

Unsupervised Data Augmentation (UDA) is a semi-supervised technique that applies a *consistency loss* to penalize differences between a model’s predictions on (a) observed (unlabeled) examples; and (b) corresponding ‘noised’ examples produced via data augmentation. While UDA has gained popularity for text classification, open questions linger over which of its components are important, and how to extend the method to sequence labeling tasks; this paper addresses these questions. Our main contribution is an empirical study of UDA to establish which components of the algorithm confer benefits in NLP. Notably, although prior work has emphasized use of clever augmentation techniques including back-translation, we find that enforcing consistency between predictions assigned to observed and *randomly* substituted words often yields comparable (or greater) benefits compared to these more complex perturbation models. Furthermore, we find that applying UDA’s consistency loss affords meaningful gains *without any unlabeled data at all*, i.e., in a standard supervised setting. In short, UDA need not be unsupervised to realize much of its noted benefits, and does not require complex data augmentation to be effective.

1 Introduction

While the advent of large neural models has led to rapid progress on a wide spectrum of prediction benchmarks in NLP, these methods tend to require large amounts of training data. This limitation is particularly acute in domains such as information extraction from scientific documents, where unlabeled in-domain data is plentiful but labeled data is rare and requires significant annotator experience to produce. The cost of acquiring data in such domains has spurred interest in developing models that can achieve greater extraction accuracy, even when the available labelled corpora are small (Nye et al., 2018; Maharana et al., 2018).

In this paper, we investigate Unsupervised Data Augmentation (UDA; Xie et al. 2019), a recently proposed semi-supervised learning method in which models are trained on both labeled and unlabeled in-domain data. The learning objective for the unlabeled component entails minimizing the divergence between the model’s outputs on a given example and its outputs on a perturbed version of the same example. While this combination of data augmentation with a consistency loss was previously proposed as *Invariant Representation Learning* and demonstrated utility in the context of speech recognition (Liang et al., 2018), UDA applies the method in a semi-supervised setting, in a manner similar to virtual adversarial training (Miyato et al., 2018). The original UDA paper (Xie et al., 2019) reported significant benefits on both computer vision and NLP tasks.

Producing such perturbed examples requires specifying a data augmentation pipeline. Typically, these apply one or more transformations that (the practitioner hopes) do not alter the applicable label (Goodfellow et al., 2016). In computer vision, a number of straightforward and demonstrably effective data augmentation techniques, such as horizontal flipping, cropping, rotating, small translations, and various perturbations to the color spectrum have gained widespread adoption (Huang et al., 2016; Zagoruyko and Komodakis, 2016). More recently, these methods, among others, have been successfully applied in concert with UDA to improve performance on image classification tasks (Xie et al., 2019).

By contrast, in NLP, there is less consensus about which perturbation models can be applied with confidence that they will not change the applicability of the original label. To apply UDA in NLP, researchers have primarily focused on *back-translation* (Sennrich et al., 2016; Edunov et al., 2018), generating paraphrases by applying a machine translation model to map a document into

a *pivot* language and then back into the original language. In practice, this process produces augmentations of varying quality. Another problem is that back-translation is slow, and performance may depend on arbitrary choices concerning both the translation model and pivot language. Incorporating large quantities of unlabeled data in the training process is also computationally expensive.

Given these limitations—and to better characterize why and when UDA helps—we investigate whether the benefits of UDA on NLP tasks can be achieved using less unlabeled data and/or simpler input perturbations. To this end, we investigate uniform random word replacement as an augmentation method. Random substitution for augmentation has been considered previously in the context of translation (Wang et al., 2018), and the UDA paper (Xie et al., 2019) gives preliminary results using random replacement for a single dataset, where it slightly underperforms back-translation. Here we deepen this analysis, showing that, surprisingly, random replacement is generally competitive with back-translation. Furthermore, we find that significant increases in performance can be achieved by applying consistency loss on only small labeled data sets (although large volumes of in-domain unlabeled data provide further gains).

As an additional contribution, we adapt UDA to sequence tagging tasks, which are common in NLP. Back-translation is ill-suited to such tasks, because we lack alignment between spans of interest in the original text and the back-translated paraphrase. For these problems, we propose and evaluate word replacement augmentation strategies for sequence tagging. Interestingly, we observe that augmentation via uniform random word replacement yields improvements, but that it is more effective to employ a masked language model to predict *reasonable* candidate replacements.

2 Unsupervised Data Augmentation

UDA is a semi-supervised method in which a model is trained to make similar predictions for an observed example and a corresponding perturbed instance produced via some data augmentation technique (in addition to satisfying the standard objectives over labeled data). Applying UDA requires specifying both (i) a consistency loss to be applied on an (original, augmented) example pair; and (ii) a data augmentation technique to produce the perturbed examples in the first place.

2.1 Consistency Loss

As originally proposed by Xie et al. (2019), UDA’s loss function is a sum over a supervised component and an unsupervised component. Assuming cross-entropy loss for the former, the loss function is:

$$\sum_{x,y \in L} [-\log(\hat{p}(y|x))] + \lambda \sum_{x \in \mathcal{U}} [\mathcal{L}(x)], \quad (1)$$

where L is the set of labeled data, \mathcal{U} is the set of unlabeled data, and λ weights the relative contribution of the unlabeled term to the total loss. The consistency loss \mathcal{L} is defined as the KL-Divergence between model predictions for the original and augmented examples:

$$\mathcal{L}(x) = \mathcal{D}_{\text{KL}}\{\hat{p}(y|x) || \hat{p}(y|q(x))\}, \quad (2)$$

where q is a data perturbation operation and $\hat{p}(y|x)$ is the probability distribution over labels output by the model given input x .

For sequence tagging tasks where examples correspond to multiple labels, we define the consistency loss as the average KL-Divergence between per-word model predictions for the original and augmented examples. Specifically, we replace the consistency loss above with

$$\mathcal{L}(x) = \frac{\sum_{j=1}^n \mathcal{D}_{\text{KL}}(\hat{p}(y_j|x) || \hat{p}(y_j|q(x)))}{n}. \quad (3)$$

Here, n denotes sequence length and $\hat{p}(y_j|x)$ are the predicted probabilities assigned by our model to labels corresponding to word j of sentence x .

2.2 Data Augmentation Strategies for Text

As defined above, consistency loss (for both classification and sequence tagging) requires specifying a data perturbation operation q that can be applied to observed instances, yielding a new, but similar, example. If we assume that q transforms x such that $q(x)$ and x share the same true label, then it seems reasonable to desire that a model would make similar predictions for x and $q(x)$, as encouraged by consistency loss.

However, there is a trade-off between the diversity of instances produced by q for x and the likelihood that these will share the same label as x . For example, consider the strategy of paraphrasing via backtranslation. A valid paraphrase is one that, with high probability, shares the ground-truth label of the input. Xie et al. (2019) observed that

diversity among the generated paraphrases is more important than validity when applying UDA to text classification. This suggests that it may be possible to effectively use an alternative augmentation strategy that prioritizes diversity and simplicity at the expense of validity.

Uniform Random Word Replacement We propose a variant of q that performs a simple uniform random word replacement operation. Specifically, we define $q(x)$ such that most of the time it copies directly from x , but with some probability p it replaces each x_j in x with some other word x'_j drawn at random from the vocabulary of words that appear in $\mathcal{U} \cup L$. Formally:

$$q(x)_j = \begin{cases} x'_j & \text{with probability } p \\ x_j & \text{with probability } 1 - p. \end{cases} \quad (4)$$

This method is simple (and does not require a learned language model), but naive. It produces output which is diverse, but not necessarily valid or even grammatical. We compare this technique to two more clever model-based data augmentation techniques: One for text classification (proposed in prior work) and one suitable for sequence tagging (which we introduce here).

Back-Translation We use the back-translation machinery described by Xie et al. (2019). Specifically, this entails use of WMT’14 English-French (Bojar et al., 2014) translation models in both directions with random sampling with a tunable temperature in place of beam search for generation. We set our temperature to 0.9, one of the recommended settings in prior work (Xie et al., 2019).

Masked Language Model Back-translation is not suitable for use in UDA for sequence tagging tasks, as in such tasks labels apply to *tokens* and it is not obvious how to align tokens in a given paraphrase with those in the original text. More specifically, as we have defined it for sequence tagging (Equation 3), consistency loss penalizes dissimilarity between model predictions $p(y_j|x)$ and $p(y_j|q(x))$ for all indices j . However, when q is defined as a back-translation process, there is no expectation that x ’s and $q(x)$ ’s ground-truth labeling will be aligned. They may not even be of the same length. Therefore, we instead consider word replacement strategies (at each index j), including (i.i.d.) random replacement, and a model-based word replacement strategy that ensures alignment

between x and $q(x)$. Both of these involve individual word substitutions, so x and $q(x)$ will have the same length.

For the model-based replacement strategy, we again define q such that it replaces a given word in x with probability p (otherwise copying from x). However, here we select x'_j using a masked language model. Specifically, we mask x_j and use BERT (Devlin et al., 2018) to induce a probability distribution over all possible words (in its vocabulary) that might appear at position j . We then draw x'_j from the ten most probable words (excluding the original word x_j) with probabilities proportional to the likelihood assigned to these words by BERT. We hypothesize that this method will provide a substantially greater expectation of validity than random replacement, on the assumption that BERT is sensitive enough to context that it is likely to replace words of one category with other words of the same category.

In Table 1 we show examples of random and BERT-based replacement of entities, in randomly selected sentence from CoNLL. As expected, random selection does not respect grammaticality or entity classes. BERT performs better, albeit imperfectly. For example, “aitken” is replaced with the first token of a surname, “mc”, rather than a full surname and “antoine” is replaced with “he”. While the latter substitution is grammatical and semantically similar to the original, “he” is not considered a named entity in CoNLL.

3 Experimental Setup

We evaluate our proposed training method on four English language text classification and three English language sequence tagging datasets. Of these, the classification datasets include three benchmark sentiment sets (IMDB, Yelp, and Amazon), and one scientific classification task (evidence inference). The sequence tagging datasets include one standard NER benchmark dataset (CoNLL-2003) and two scientific sequence labeling tasks (EBM-NLP and TAC). The scientific tasks are of particular interest for this work because it is expensive to collect annotations in these specialized domains.

IMDB (Maas et al., 2011) is a sentiment classification dataset consisting of movie reviews (25,000 examples in both the train and test sets) drawn from the IMDB website. Reviews with a score $\leq 4/10$ are considered negative, those with scores $\geq 7/10$ are considered positive. Neutral reviews are excluded.

	ORG	PER	PER		PER
	aberdeen	manager	roy	aitken	said: "it's unfortunate for us that antoine cannot play..."
BERT	rangers		glen	mc	he
Rand	delegates		cancer	peripheral	51,000

Table 1: Example replacements selected for entity tokens in a randomly selected sentence from CoNLL-2003 using different selection methods. In training, all words are equally likely to be selected for replacement. We focus on named entities in this example for illustrative purposes.

Yelp (Zhang et al., 2015) is a sentiment classification dataset comprising reviews drawn from Yelp (560,000 in the train set, 38,000 in the test set). One and two star reviews are considered negative. Three and four star reviews are considered positive.

Amazon (Zhang et al., 2015; McAuley and Leskovec, 2013) is a sentiment classification dataset consisting of Amazon reviews (3,600,000 in the train set, 400,000 in the test set). One and two star reviews are considered negative. Four and five star reviews are considered positive. Three star reviews are not included.

Evidence Inference (Lehman et al., 2019; DeYoung et al., 2020) We construct a classification dataset derived from the evidence inference dataset, a biomedical corpus in which the task is to infer the effect of an intervention on an outcome from an article describing a randomized controlled trial. The classes correspond to the intervention leading to a *significant increase*, *significant decrease*, or *no significant change* in outcome. In the original task, the model must first extract relevant evidence sentences from the full text article, and then make a prediction based on this. We evaluate in the ‘oracle’ setting, in which the model must only classify *given* relevant evidence sentences ($\sim 17,000$ train examples, and $\sim 2,000$ instances in the test set).

CoNLL-2003 (Tjong Kim Sang and De Meulder, 2003) is an NER dataset consisting of annotated Reuters news articles ($\sim 14,000$ sentences in the training set, $\sim 3,000$ in the test set), labeled with entity categories *person*, *organization*, *location*, and *miscellaneous*.

TAC (Schmitt et al., 2018) comprises annotated “materials and methods” sections from PubMed Central articles ($\sim 5,500$ sentences in the training set, $\sim 6,500$ in the test set). Labels are available for 24 entity classes, of which we consider the two best represented: *end point* and *test article*.

EBMNLP (Nye et al., 2018) is a corpus of annotated abstracts drawn from medical articles describ-

ing medical randomized controlled trials ($\sim 28,000$ sentences in the training set, $\sim 2,000$ in the test set). Spans are tagged as describing the *patient population*, the *intervention* studied, and the *outcome* measured in the trial being described.

For each dataset we simulate a pool of unlabeled data by hiding the annotations of the training set. We then create five distinct sets of labeled data (ten for sequence tagging tasks) by revealing the annotations for a random subset of the pool, forming a labeled training set L and an unlabeled training set \mathcal{U} . For classification tasks, we sample ten examples per class to form L , while for CoNLL and EBMNLP, we sample two hundred examples. For the smaller TAC dataset, we use one hundred.

Training details For each L , we then train a model both using only standard supervised learning over L and with an additional consistency loss, using either uniform random word replacement or more complex data augmentation technique (back-translation for classification; BERT-based replacement for sequence tagging). When training with consistency loss, we evaluate variants in which we apply this to both L and \mathcal{U} , and where we apply it only over L . The latter corresponds to a standard supervised setting with an additional loss term.

BERT’s pretraining task already incorporates unsupervised data (Devlin et al., 2018). We therefore also repeat the above experiments using finetuned weights instead of off-the-shelf pretrained weights. The finetuned weights are produced by training on BERT’s masked language model task with $L \cup \mathcal{U}$ as the training data.

In our classification experiments, we use a linear model on top of BERT (Devlin et al., 2018) as a classifier. For sequence tagging, we follow architecture and hyperparameter choices in prior work (Beltagy et al., 2019), adding a conditional random field (Lafferty et al., 2001) on top of BERT representations. We train all models using Adam (Kingma and Ba, 2015) with a learning rate of $2e-5$ for classification and $1e-3$ for sequence tagging.

In exploratory experiments we observed that

model performance is relatively robust to the choice of λ (the weight assigned to the consistency loss term) when large quantities of unlabeled data are available. We therefore set λ to 1 in all of our semi-supervised experiments. In our supervised learning only experiments, we found it necessary to compensate for the lack of unlabeled examples and the corresponding change in the relative weightings of standard and consistency losses. This can be effectively done by repeating labeled examples, imposing only the consistency loss for them, as though they were unlabeled. We use a ratio of 20 “unlabeled” examples to 1 labeled example in our supervised experiments.

In our augmentation procedure, we set p to 30%. For biomedical tasks (Evidence Inference, EBMNLP, TAC) model weights are initialized using SciBERT, a model pretrained over scientific papers (Beltagy et al., 2019). For all other tasks, we initialized parameters to the pretrained BERT_{BASE} weights (Devlin et al., 2018). BERT_{BASE} weights are used for CoNLL, and SciBERT weights are used for EBMNLP and TAC when using BERT as a masked-language model for data augmentation.

4 Results

4.1 Classification

Figure 1 presents the results of our experiments for the classification tasks. Both back-translation and random replacement perform well on the IMDB, Yelp, and Amazon datasets. Notably, *random replacement consistently achieves results equivalent to or better than those attained with the more computationally complex back-translation method*. This is particularly noteworthy, as Xie et al. (2019)’s theoretical analysis of UDA assumes augmentations which are both label-preserving and which produce in-domain data. Random replacement, however, will likely produce augmentations which are not grammatical and may produce outputs that are not label preserving.

That non-label preserving augmentations can lead to performance gain is intriguing and perhaps counter-intuitive. We speculate that that random replacement acts as a form of regularization by inserting noise into input sequences. This may discourage over-reliance on individual tokens.

On the evidence inference task, UDA with back-translation under-performs the supervised baseline, with a loss of 7.5 F1 when the full unlabeled dataset is used. This is perhaps unsurprising, given that the

models used for back-translation were not trained on scientific text. These results suggest that the effectiveness of back-translation is contingent on the domain similarity of the back-translation model’s training data and that of the downstream task. By contrast, UDA with random replacement produces a meaningful gain over the supervised baseline: 2.5 F1 with only the labeled data and 6 F1 with the full unlabeled dataset.

Across all classification experiments—excepting evidence inference using back-translation—applying consistency loss to only the labeled data yields improvements over the supervised baseline, albeit less than what is achieved using the full amount of unsupervised data. Further, these gains are disproportionate to the quantity of data used to attain them. In the worst case (Yelp with back-translation), using only the supervised data results in only 20% of the potential performance improvement that could be attained using the full set of unlabeled data. In the best case (Amazon with random replacement), 48% of the potential gain can be achieved without using any unsupervised data at all, despite the fact that the labeled Amazon dataset represents less than 0.001% of the full set.

4.2 Sequence Tagging

Figure 2 presents results from sequence tagging experiments. BERT-based replacement provides a meaningful advantage over the supervised baseline on the CoNLL and TAC datasets. Random replacement also offers gains, but these are smaller and less consistent. Without access to the full unlabeled dataset, random replacement results in a small decrease in performance. With access to unlabeled data, it produces only a small benefit on CoNLL. The gain on TAC is larger, but still smaller than that achieved using BERT-based replacement.

BERT-based replacement is more effective than random replacement for sequence tagging. But that random replacement provides any benefits at all for such tasks is perhaps counter-intuitive, given that predictions are made at the word level for these tasks. It is therefore likely that random replacement will lead to a change in ground truth labeling for any replaced entities. We hypothesize that this encourages the model to place greater weight on the context in which words appear. This may render models more robust at recognizing unfamiliar entities based on their contexts (Agarwal et al., 2021).

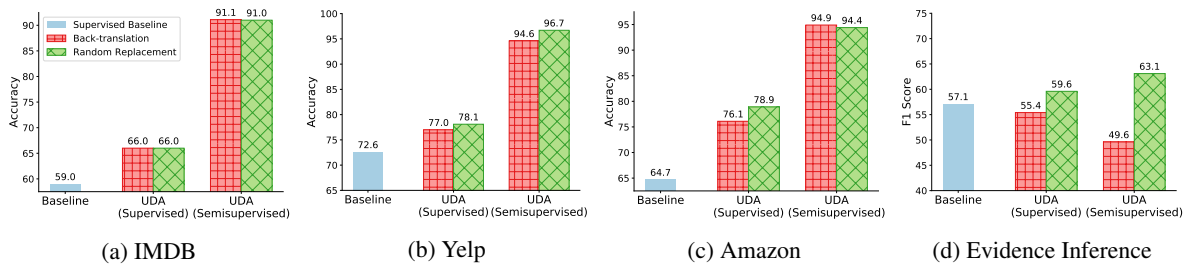


Figure 1: Comparison of performance achieved on classification tasks using different variants of UDA. Each bar represents the average performance across five sets of labeled data (labeled data quantity is noted parenthetically). The supervised baseline represents standard ML on the supervised data set only, without any consistency loss. Supervised with consistency loss represents use of consistency loss, but only over the labeled data, with the unlabeled data discarded. Semisupervised with consistency loss represents use of consistency loss over the entire dataset, both labeled and unlabeled.

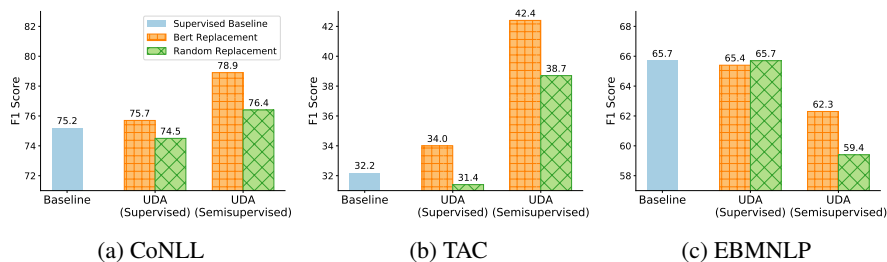


Figure 2: Comparison of performance achieved on sequence tagging tasks using different variants of UDA. Each bar represents the average performance across ten sets of labeled data (labeled data quantity is noted in parenthesis). The supervised baseline represents standard ML on the supervised data set only, without any consistency loss. Supervised with consistency loss represents use of consistency loss, but only over the labeled data, with the unlabeled data discarded. Semisupervised with consistency loss represents use of consistency loss over the entire dataset, both labeled and unlabeled.

UDA does not offer performance gains on the EBMNLP dataset using either augmentation strategy. When consistency loss is applied to only the labeled data, the performance is largely unchanged. However, when unlabeled data is incorporated, performance decreases by 3.4 F1 when using BERT-based replacement, and 6.3 F1 when using random replacement.

Crowdsourced (lay) workers annotated the training data in EBMNLP, while doctors annotated the test data (Nye et al., 2018), and we speculate that this may play a role in the difference in observed performance, as we observed similar gains to those attained on CoNLL and TAC when performing exploratory studies on a development set. It may be that UDA performs poorly on EBMNLP relative to the supervised learning baseline because relying more heavily on context is harmful when the training set annotations are noisy. We note that the lay training set annotators consistently included more words in their labeled entity spans than the test set annotators (see Table 2). This may indicate

	P	I	O
Train	8.2	3.9	4.8
Test	6.5	1.8	3.7

Table 2: Average length of PIO spans in words for EBMNLP’s train and test sets

that the training set annotators included context words which do not truly belong to an entity class in their spans. Encouraging the model to infer the implications of context based on these misidentified context words may be compounding that error.

Our results indicate that unlabeled data is more critical for UDA in sequence tagging than in classification. Here, in the best case (TAC with Replacement), we see UDA without unlabeled data achieves only 18% of the performance increase that may ultimately be achieved by including unlabeled data. This is lower than the worst case observed in the classification task.




Figure 3: Comparison of performance achieved using varying quantities of unlabeled data. Curves are averaged across 5 experiments for Yelp, and 10 for CoNLL. Each labeled dataset consists of 20 labeled examples for Yelp and 100 examples for CoNLL. Random replacement is used as the augmentation method for Yelp and BERT-based replacement is used as the augmentation method for CoNLL.




Figure 4: Analysis of performance achieved using varying quantities of labeled data on the CoNLL set. The blue solid line represents the case that UDA is used and all unlabeled data is incorporated in training. The orange dashed line represents the case that UDA is used, but with only the labeled data. The green dotted line represents training without any consistency loss. Each curve is averaged across ten experiments. BERT-based replacement is used as the augmentation method.

4.3 Varying Quantities of Labeled and Unlabeled Data

We next examine the question: How much unlabeled data is necessary? Incorporating additional unlabeled data extends the training process and we hypothesize that, at some point, we will observe diminishing returns. To this end, we run experiments varying the quantity of unlabeled data used when training on the Yelp and CoNLL datasets. These results are presented in Figure 3. For Yelp, we begin to observe diminishing returns as we approach use of the full unlabeled set. Interestingly, on CoNLL, too high a quantity of unlabeled data appears to actually degrade performance.

We also analyze the change in performance when varying the amount of labeled data used for se-




Figure 5: Performance on IMDB using off-the-shelf pretrained BERT weights compared to BERT weights finetuned to IMDB (i.e., after continuing pretraining BERT on IMDB)

quence labeling tasks. To investigate this, we train using UDA with 10, 100, 1000 or 10000 labeled sentences drawn from CoNLL. Results from these experiments are presented in Figure 4. We observe consistent, modest gains from using UDA in a semi-supervised fashion, excepting the extreme ends of the curve, where almost all or almost none of the data is labeled. In these two cases, we see performance with UDA and without UDA converge.

4.4 Finetuning BERT Weights

BERT’s pretraining tasks (masked-language modeling and next sentence prediction) already provide a method for incorporating unlabeled data (Devlin et al., 2018). Given this, we finetune by training BERT’s pretraining tasks on the full unlabeled datasets. We then investigate the resulting performance on the downstream tasks, to determine whether there is still a benefit to using UDA with the full unlabeled dataset after that data has been incorporated into the model via finetuning.

Figure 5 presents results from this experiment for the IMDB dataset. Our results show that, when BERT’s weights have already been finetuned on the unlabeled data, incorporating that data again when training with UDA is less valuable. Applying UDA using only the labeled data and random replacement allows us to realize 74% of the possible performance increase when using the full unlabeled data. This is compared to only 22% when BERT has not been finetuned.

However, we still observe that performance gains do continue to accrue when unlabeled data is incorporated into UDA training, even when the




Figure 6: Performance ranges on the Amazon dataset with spans indicating the minimum-to-maximum performance over 5 independent samples (of the labeled subset). Triangles indicate means.

BERT weights have been finetuned. Since training with UDA is comparatively computationally inexpensive to robustly finetuning BERT, it is likely practical and advantageous to use both in concert.

4.5 Variability of Results

Throughout our experiments we observed that performance varies greatly as a function of the particular labeled set used. Figure 6 illustrates the range of observed results on the Amazon dataset. The difference between the maximum performance and the minimum performance in the supervised baseline is 17.8 points of accuracy. The delta for UDA using only labeled data and random word replacement is even higher, at 21.4 points of accuracy. This has important implications for a practitioner: While one might reasonably have an expectation of achieving high performance on average, in practice only a single labeled dataset will be constructed and used for training. Our results show that, in a low resource classification setting, such a practitioner might actually achieve significantly lower or higher performance than expected.

This illustrates a further advantage of UDA. When exploiting a large quantity of unlabeled data, performance not only improves, but becomes more consistent across labeled dataset choices as well. We observe similar trends across all classification datasets, with the exception of Evidence Inference with random replacement, for which the variability of results remains relatively high, even when all unlabeled data is employed.

By contrast, we do not observe such trends in the sequence tagging tasks, where UDA variant

choice does not consistently affect the variability of performances across labeled set choices.

5 Conclusions

In this paper we evaluated and extended UDA in the context of NLP tasks. We proposed and evaluated new techniques for applying UDA to classification tasks and extended it to sequence tagging tasks by imposing a consistency loss over word label distributions. We showed that naive data augmentation methods may often be just as effective as the complex, model based augmentations that are currently fashionable, and that performance improvements may still be attained even in the absence of any unlabeled data.

We proposed a simple, effective augmentation method: randomly replace words with other words. The replacement word may be selected either uniformly at random, or by using BERT (Devlin et al., 2018) as a masked language model to induce a probability distribution over tokens. We found that the former method is effective for classification tasks, and the latter for sequence tagging. We further investigated the practicality of using UDA without unlabeled data, applying a consistency loss only to a small labeled dataset. We experimentally evaluated various augmentation strategies and settings on four classification datasets and three sequence tagging datasets.

We found reliable performance increases on all four classification datasets. For classification, we found that random word replacement is as effective as—and sometimes more effective than—back-translation, which is what has been proposed in prior work Xie et al. (2019). In particular, random word-replacement is effective on our scientific classification task, where back-translation is ineffective, perhaps due to its reliance on machine translation models not trained with scientific literature.

Our experiments show that both random replacement and BERT-based replacement are effective on two out of three sequence tagging tasks, with BERT-based replacement that we have proposed consistently outperforming random replacement. On the third sequence tagging dataset, we observed a degradation of performance when using any variety of UDA, which we hypothesize owes to the noisy annotation of this training set.

Finally, we found that UDA may produce meaningful increases in performance even when unlabeled data is unavailable, particularly if the weights

have already been finetuned to the task. The magnitude of this increase depends upon the task, and may be relatively large (as in the case of the Amazon dataset) or relatively small (as in the CoNLL dataset). In general this approach is more effective for classification tasks than sequence tagging tasks.

To summarize our findings: UDA is effective in low-supervision natural language tasks, even when used with naive augmentation methods and without unlabeled data.

Ethical Considerations

We are not aware of any social harm which our research might cause. We note that one advantage of using our proposed naive, random word replacement augmentation over more complex model driven augmentations (such as back-translation) is a reduction in required compute. This is important as emissions and energy costs of compute heavy model training continue to come under scrutiny (Strubell et al., 2019). We hope that by presenting a simple but effective training method, our work may serve the public good by helping to address this rising challenge.

Acknowledgements

This work was funded in part by the National Institutes of Health (NIH) grant R43ES029901, and by the Army Research Office (ARO), W911NF1810328.

References

- Oshin Agarwal, Yinfei Yang, Byron C Wallace, and Ani Nenkova. 2021. Interpretability analysis for named entity recognition to understand system predictions and how they can improve. *Computational Linguistics*, 47(1):117–140.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. **Scibert: Pretrained language model for scientific text**. In *EMNLP*.
- Ondrej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. **Findings of the 2014 workshop on statistical machine translation**. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 12–58, Baltimore, Maryland, USA. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Jay DeYoung, Eric Lehman, Ben Nye, Iain J. Marshall, and Byron C. Wallace. 2020. **Evidence inference 2.0: More data, better models**.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. **Understanding back-translation at scale**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500, Brussels, Belgium. Association for Computational Linguistics.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. The MIT Press.
- Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. 2016. **Densely connected convolutional networks**.
- 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*.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. **Conditional random fields: Probabilistic models for segmenting and labeling sequence data**. In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01*, pages 282–289, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Eric Lehman, Jay DeYoung, Regina Barzilay, and Byron C Wallace. 2019. Inferring which medical treatments work from reports of clinical trials. *arXiv preprint arXiv:1904.01606*.
- Davis Liang, Zhiheng Huang, and Zachary C Lipton. 2018. Learning noise-invariant representations for robust speech recognition. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 56–63. IEEE.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. **Learning word vectors for sentiment analysis**. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Adyasha Maharana, Arpit Tandon, Eric Wimberley, Mihir Shah, Ruchir Shah, and Brian E Howard. 2018. A pragmatic approach to information extraction for systematic review.
- Julian McAuley and Jure Leskovec. 2013. **Hidden factors and hidden topics: Understanding rating dimensions with review text**. In *Proceedings of the 7th ACM Conference on Recommender Systems, RecSys '13*, page 165–172, New York, NY, USA. Association for Computing Machinery.

- Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993.
- Benjamin Nye, Junyi Jessy Li, Roma Patel, Yinfei Yang, Iain James Marshall, Ani Nenkova, and Byron C. Wallace. 2018. [A corpus with multi-level annotations of patients, interventions and outcomes to support language processing for medical literature](#). *CoRR*, abs/1806.04185.
- Charles Schmitt, Vickie Walker, Ashley Williams, Arun Varghese, Yousuf Ahmad, Andy Rooney, and Mary Wolfe. 2018. Overview of the tac 2018 systematic review information extraction track.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. [Energy and policy considerations for deep learning in NLP](#). *CoRR*, abs/1906.02243.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In *Proceedings of CoNLL-2003*, pages 142–147. Edmonton, Canada.
- Xinyi Wang, Hieu Pham, Zihang Dai, and Graham Neubig. 2018. [SwitchOut: an efficient data augmentation algorithm for neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 856–861, Brussels, Belgium. Association for Computational Linguistics.
- Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. 2019. Unsupervised data augmentation for consistency training. *arXiv preprint arXiv:1904.12848*.
- Sergey Zagoruyko and Nikos Komodakis. 2016. [Wide residual networks](#). *CoRR*, abs/1605.07146.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#).