

Data Augmentation for Hypernymy Detection

Thomas Kober

Rasa
Berlin, DE
t.kober@rasa.com

Julie Weeds

University of Sussex
Brighton, UK
j.e.weeds@sussex.ac.uk

Lorenzo Bertolini

University of Sussex
Brighton, UK
l.bertolini@sussex.ac.uk

David Weir

University of Sussex
Brighton, UK
d.j.weir@sussex.ac.uk

Abstract

The automatic detection of hypernymy relationships represents a challenging problem in NLP. The successful application of state-of-the-art supervised approaches using distributed representations has generally been impeded by the limited availability of high quality training data. We have developed two novel data augmentation techniques which generate new training examples from existing ones. First, we combine the linguistic principles of hypernym transitivity and intersective modifier-noun composition to generate additional pairs of vectors, such as *small dog - dog* or *small dog - animal*, for which a hypernymy relationship can be assumed. Second, we use generative adversarial networks (GANs) to generate pairs of vectors for which the hypernymy relation can also be assumed. We furthermore present two complementary strategies for *extending* an existing dataset by leveraging linguistic resources such as WordNet. Using an evaluation across 3 different datasets for hypernymy detection and 2 different vector spaces, we demonstrate that both of the proposed automatic data augmentation and dataset extension strategies substantially improve classifier performance.

1 Introduction

The detection of hypernymy relationships between terms represents a challenging commonsense inference problem and is a major component in recognising paraphrase and textual entailment in larger units of text. Consequently, it is important for Question-Answering, Text Simplification and Automatic Summarization. For example,

There are lots of cars and vans at the port today.
might be adequately summarised by

There are lots of vehicles at the port today.
as *car* and *van* both lexically entail, i.e. they are both hyponyms of the more general term *vehicle*.

Furthermore, the recognition and discovery of hyponym-hypernym relations is a foundational part of constructing taxonomies, which has a range of practical applications in a variety of domains such as Healthcare (Barisevičius et al., 2018) or Fashion¹.

While distributed representations of words are commonly used to find semantically similar words, they do not straightforwardly provide a way to distinguish more fine-grained semantic information, such as hypernymy, co-hyponymy and meronymy relationships. This deficiency has attracted substantial attention in the literature, and with regard to the task of hypernymy detection, both unsupervised approaches (Hearst, 1992; Weeds et al., 2004; Kotlerman et al., 2010; Santus et al., 2014; Rimell, 2014; Nguyen et al., 2017; Chang et al., 2018) and supervised approaches (Weeds et al., 2014; Roller et al., 2014; Roller and Erk, 2016; Shwartz et al., 2016; Vulić and Mrkšić, 2018; Rei et al., 2018; Kamath et al., 2019) have been proposed.

Supervised methods have, however, been severely hampered by a lack of adequate training data. Not only has a paucity of labelled data been an obstacle in the adoption of deep neural networks and other more complex supervised methods, but two compounding problem-specific issues have been identified. First, there is a need to avoid lexical overlap between the training and test sets in order to avoid the lexical memorisation problem (Weeds et al., 2014; Levy et al., 2015a), where a supervised method simply learns the relationships between lexemes rather than generalising to their distributional features. Second, the performance of classifiers given just the hypernym word (at training

¹See e.g. <https://www.voguebusiness.com/technology/taxonomy%2Dis%2Dthe%2Dnew%2Dfashion%2Dtech%2Dessential%2Dthe%2Dyes> or <https://fashionbrain-project.eu/fashion%2Dtaxonomy/>.

and testing) has been shown to be almost as good as performance given both words (Weeds et al., 2014; Shwartz et al., 2017). This suggests that classifiers are learning the distributional features that make something a more general term or a more specific term. Our conjecture is that in order to learn the more complex function, more complex machinery, and hence more labelled data is required.

In computer vision or speech recognition, it is common to use data augmentation to increase the size of the training set (Shrivastava et al., 2017; Park et al., 2019). The idea is that there are certain transformations of the data under which the class label remains invariant. For example, rotating an image does not change whether that image contains a face or not. By providing a supervised classifier with rotated examples, it can better generalise.

In this work, we consider the use of linguistic transformations to augment existing datasets for hypernymy detection. The challenge is to identify transformations that can be applied to the representations of two words that are known to be in a hypernym relationship, such that the entailment relation still holds between the transformed representations. We propose two ways to achieve this.

Our first augmentation technique is based on the hypothesis that lexical entailment is transitive and therefore invariant under certain compositions. For example, if A entails B and B entails C then A also entails C . Suitable candidates for A can be found by composing common intersective adjectives with the noun B . For example, if we know that *car* entails *vehicle*, then we can augment the dataset with *fast car* entails *car* and *fast car* entails *vehicle*.

Our second augmentation technique is based on the hypothesis that lexical entailment is invariant within a certain threshold of similarity. If A entails B , A' is very similar to A and B' is very similar to B then A' will also entail B' . In order to obtain vectors which are sufficiently similar to the words in the training data, we apply generative adversarial networks (GANs) to create realistic-looking synthetic vectors, from which we choose the most similar to the words in the training data.

We evaluate the proposed techniques on three hypernymy detection datasets. The first two are standard benchmark tasks in this area (Weeds et al., 2014; Baroni et al., 2012), both of which are generated from WordNet (Fellbaum, 1998). However, since many of the approaches to hypernymy classification involve vector space models which have

been specialised using the entirety of WordNet, we need to guard against the danger that evaluations are simply measuring how well WordNet has been encoded, rather than how well the general hypernymy relationship has been learned. In light of this, we introduce a new dataset (that we call **HP4K**) which does not rely on WordNet in its construction.

We evaluate our two data augmentation techniques against two methods for increasing the size of the training data which rely on finding or mining more non-synthetic examples. First, we consider the extraction of additional examples from WordNet. Second, we consider extracting examples automatically from a Wikipedia corpus using Hearst Patterns (Hearst, 1992). This provides us with what one would expect to be an upper bound on what we might reasonably expect to achieve with a similar amount of synthetic examples generated using our data augmentation techniques.

Our contributions are thus threefold. First, we have identified two novel data augmentation techniques for the task of hypernymy detection which have the potential to generate almost limitless quantities of synthetic data. Second, we show, rather surprisingly, that adding synthetic data is more effective than adding non-synthetic data in almost all cases. Third, we release a new benchmark evaluation dataset for the lexical entailment task that is not dependent on WordNet.

2 Related Work

Data augmentation has recently become a very popular research topic in NLP and has successfully been applied in machine translation systems (Senrich et al., 2016; Fadaee et al., 2017; Wang et al., 2018; Li et al., 2019; Tong et al., 2019; Matos Veliz et al., 2019; Xia et al., 2019; Gao et al., 2019; Li and Specia, 2019; Liu et al., 2019), but also for tasks such as relation extraction (Can et al., 2019; Yan et al., 2019), text classification (Wei and Zou, 2019), or natural language inference (Kang et al., 2018; Junghyun et al., 2020). Most similar to our usage of GANs for data augmentation is the proposal of Kang et al. (2018) who leverage a GAN-based setup together with WordNet for data augmentation for natural language inference.

While to the best of our knowledge this work represents the first application of data augmentation for lexical entailment, a number of alternative approaches have been proposed. Most proposals rely on supervised methods for injecting an exter-

nal source of knowledge into distributional representations. Starting with *retro-fitting* (Faruqui et al., 2015), vector-specialization methods modify existing representations to embed desired features (Vulić and Mrkšić, 2018; Rei et al., 2018; Kamath et al., 2019; Glavaš and Vulić, 2019).

3 Data Augmentation Strategies for Hypernymy Detection

Given a labelled dataset $\mathcal{D}_{\mathcal{X}}$ of triples $\langle x_{\text{hypo}}^{(i)}, x_{\text{hyper}}^{(i)}, y^{(i)} \rangle$, where $x_{\text{hypo}}^{(i)}, x_{\text{hyper}}^{(i)} \in \mathcal{X}$ and $y^{(i)} \in \{0, 1\}$, we define data augmentation as adding additional hyponym-hypernym triples $\langle x_{\text{hypo}}^{\prime(j)}, x_{\text{hyper}}^{\prime(j)}, y^{\prime(j)} \rangle$ coming from an automatically generated augmentation set $\mathcal{A}_{\mathcal{X}'}$, such that $x_{\text{hypo}}^{\prime(j)}, x_{\text{hyper}}^{\prime(j)} \in \mathcal{X}'$ and $y^{\prime(j)} \in \{0, 1\}$, to the existing training set of \mathcal{D} . We ensure that the data augmentation does not introduce any lexical overlap with the existing test set, i.e. $\mathcal{X} \cap \mathcal{X}' = \emptyset$.

Data augmentation strategies in NLP can roughly be divided into two categories: linguistically grounded augmentation and artificial augmentation. In the former, which has been the dominant paradigm in NLP, any additional instances that are added to a training set have an actual surface form representation, i.e. the data points correspond to actual words or sentences (Kim et al., 2019; Kumar et al., 2019; Gao et al., 2019; Andreas, 2020; Croce et al., 2020). The latter adds instances that are fully or partly artificial, meaning they do not correspond to any words or sentences. In this work we propose methods for both categories, data augmentation via distributional composition adds data points grounded in real language to a training set, and data augmentation based on GANs infers plausible points in latent space, which however, do not correspond to any real linguistic objects.

Furthermore, we distinguish between *data augmentation* and *dataset extension*, where in the former case we only leverage knowledge from the existing dataset and in the latter case we rely on expanding the training set with additionally mined hyponym-hypernym pairs. Below, we discuss two ways of augmenting and two ways of extending a training set. We make use of a cleaned October 2013 Wikipedia dump (Wilson, 2015) as reference corpus to determine word and bigram frequencies.

Distributional Composition based Augmentation. We take a modified noun as being in a hypernymy relation with the unmodified noun. For example, we treat the pairs $\langle \text{fast car}, \text{car} \rangle$ and

$\langle \text{car}, \text{vehicle} \rangle$ as expressing the same semantic relation when the modifier-noun compound is composed with an intersective composition function.

We focus on adjective-noun (AN) and noun-noun (NN) compounds, extracted from our reference corpus where each AN or NN phrase occurred at least 50 times. We filtered pairs with non-subjective adjectives using a wordlist from Nayak et al. (2014)².

We consider two strategies for automatically constructing positive hyponym-hypernym pairs: simple positive cases such as $\langle \text{small dog}, \text{dog} \rangle$ or $\langle \text{fast car}, \text{car} \rangle$; and gapped positive cases that mimic the transitivity of hypernym relations, where we pair the hypernym of an existing hyponym-hypernym pair with a compound hyponym. For example if $\langle \text{car}, \text{vehicle} \rangle$ is in the training data, we combine *car* with one of its modifiers to create the pair $\langle \text{fast car}, \text{vehicle} \rangle$.

We construct negative pairs from the simple positive cases using two strategies: creating compositional co-hyponyms such as $\langle \text{fast car}, \text{red car} \rangle$, where we keep the head noun fixed and pair it with two different modifiers; and creating perturbed simple positive examples, such as $\langle \text{small dog}, \text{cat} \rangle$ where we select the incorrect hypernym (e.g. *cat*) from the n most similar nouns to the composed hyponym (e.g. *dog*). We apply the same methodology to the perturbed gapped positive examples, replacing the correct hypernym with a noun from the top n neighbours of the compositional hyponym’s head noun. For example, given a positive pair such as $\langle \text{dog}, \text{animal} \rangle$, this would result in negative examples such as $\langle \text{small dog}, \text{vehicle} \rangle$, where the hyponym *dog* is paired with a modifier and the hypernym *animal* is replaced with one of its neighbours, in this case, *vehicle*.

In neural word embeddings, an additive composition function approximates the intersection of the corresponding feature spaces (Tian et al., 2017), hence by creating positive pairs such as $\langle \text{small dog}, \text{dog} \rangle$, we encode the distributional inclusion hypothesis (Weeds et al., 2004; Geffet and Dagan, 2005) in the augmentation set.

GAN based Augmentation. We create an augmentation set using Generative Adversarial Networks (Goodfellow et al., 2014). GANs consist of two model components — the generator and the discriminator — which are typically implemented

²In preliminary experiments we did not find that filtering non-subjective adjectives had much of an effect, but decided to move forward with the filtered data nonetheless.

as neural networks. The generator’s task is to create data that mimics the distribution of the original data, while the discriminator’s task is to distinguish between data coming from the real distribution and synthetic data coming from the generator. Both components are trained jointly until the generator succeeds in creating realistic data. Using GANs for data augmentation has been shown to be a successful strategy for a number of computer vision tasks (Shrivastava et al., 2017; Frid-Adar et al., 2018; Neff, 2018). Our goal is to create synthetic hyponym-hypernym pairs that are similar to real examples. Unlike most other scenarios involving GANs for NLP tasks, our generated vectors do not need to correspond to actual words.

For our model - *GANDALF*³ - we used a list of $\approx 40K$ nouns for which we had vector representations as the “real” data input to *GANDALF*, and sampled the synthetic vectors from a Gaussian distribution, optimising a binary cross-entropy error criterion for the generator and the discriminator, which are both simple feedforward networks with a single hidden layer. We provide *GANDALF*’s full model details in Appendix A. As an additional quality check for the generated vectors, we tested whether a logistic regression classifier could distinguish the synthetic and non-synthetic vectors. Typically, the accuracy of the classifier was between 0.55-0.65, meaning the classifier is barely able to distinguish between “real” vectors and generated ones.

Figure 1 illustrates the training loop of *GANDALF* as well as the selection process for constructing an augmented training set. Essentially, once *GANDALF* has been trained, the generator is used to create a large collection⁴ of synthetic noun vectors. To augment a dataset, $\mathcal{D}_{\mathcal{X}}$, for each triple, $\langle x_{hypo}, x_{hyper}, y \rangle \in \mathcal{D}_{\mathcal{X}}$ we find the n synthetic vectors most similar to x_{hypo} and the n synthetic vectors most similar to x_{hyper} and for each of the n^2 synthetic vector pairs, $\langle x'_{hypo}, x'_{hyper} \rangle$, we create the triple $\langle x'_{hypo}, x'_{hyper}, y \rangle$. The augmented training set is formed by randomly sub-sampling this set of triples.

WordNet based Extension. WordNet (Fellbaum, 1998) is a large manually curated lexical resource, covering a wide range of lexical relations between words, where groups of semantically similar words form “synsets”⁵. For each synset we

extract all hypernyms and hyponyms of a given lexeme, and add it as a positive hyponym-hypernym pair if the original lexeme and any extracted hypernym/hyponym occurs at least 30 times in our reference corpus.

We construct negative pairs based on distributional similarity, where we calculate the pairwise cosine similarities between all lexemes in the positive set. Subsequently we use all antecedent (LHS) lexemes from the extracted positive pairs and select the top n most similar words for each antecedent as negative examples⁶.

Pattern based Extension. Hearst Patterns (Hearst, 1992) are textual patterns such as *a car is-a vehicle* and can be automatically mined from text corpora in an unsupervised way. This has recently been shown to deliver strong performance on the hypernymy detection task (Roller et al., 2018). In this work, we leverage Hearst Patterns to mine additional hyponym-hypernym pairs in order to extend a training set. We treat any extracted noun pairs as additional positive examples and create the negative pairs in the same way as for the WordNet-based approach above.

4 Experiments

We evaluate our models on the datasets **Weeds** (Weeds et al., 2014) and **LEDS** (Baroni et al., 2012): well-studied and frequently used benchmarks for the hypernymy detection task (Roller et al., 2014; Vilnis and McCallum, 2014; Roller and Erk, 2016; Carmona and Riedel, 2017; Shwartz et al., 2017). Since both datasets use WordNet during construction, this can give rise to a bias in favour of those models that also make use of WordNet. To address this concern, we have created a new entailment dataset, **HP4K**, that makes use of Hearst Patterns, and is manually annotated, thereby avoiding the use of WordNet.

Weeds: The dataset is based on nouns sampled from WordNet where each noun had to occur at least 100 times in Wikipedia, and its predominant sense had to account for more than 50% of the occurrences in SemCor (Miller et al., 1993). We use the predefined split of Weeds et al. (2014), that avoids any lexical overlap between the training and evaluation sets. The split contains 2012 examples in the training set, evenly balanced between positive and negative hyponym-hypernym pairs, and

³GAN-based Data Augmentation for Lexical Inference.

⁴We would typically create half a million synthetic nouns.

⁵We use the API provided by NLTK (Loper and Bird,

2002), using WordNet 3.0.

⁶Ensuring we don’t accidentally add any real positive pairs.

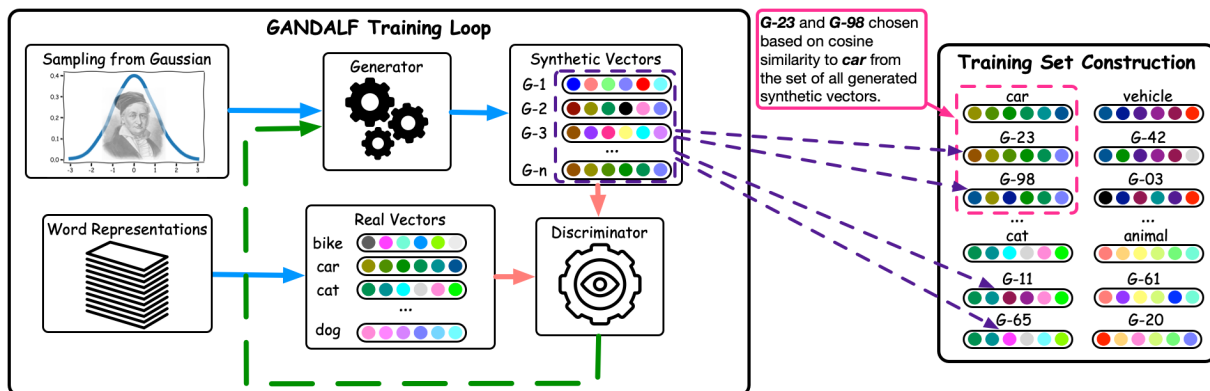


Figure 1: The left panel shows the training loop for *GANDALF* where synthetic noun vectors are generated from Gaussian noise. The discriminator’s task is to distinguish between the generated synthetic vectors and real word representations from word2vec or HyperVec. The dashed line indicates error backpropagation, on the basis of the discriminator’s current performance, through to the generator. After a final set of synthetic vectors has been generated with *GANDALF*, we choose the top n most similar synthetic vectors to a real representation and add these with the same label as the original hyponym-hypernym pair from an existing training set (e.g. $\langle \text{cat}, \text{animal} \rangle$) to the augmented training set (right panel).

502 examples in the evaluation set.

LEDS: The dataset consists of 2770 examples, evenly balanced between positive and negative hyponym-hypernym pairs. The positive examples are based on direct hyponym-hypernym relationships from WordNet and the negative examples are based on a random permutation of the hypernyms of the positive pairs. As there is no predefined training/evaluation split, we make use of the 20-fold cross-validation methodology of Roller and Erk (2016) that avoids any lexical overlap between training and evaluation sets.

HP4K: We extracted Hearst Patterns from our reference Wikipedia corpus and randomly selected 4500 unigram pairs. Subsequently, we manually annotated each pair according to whether it constitutes a correct hyponymy-hypernymy relation or not. The labelling was carried out by 4 experienced annotators — all domain experts, familiar with the problem of hypernymy detection. We then split up the annotators in two teams, with each team annotating one half of the dataset. The initial round of annotations resulted in a Cohen’s κ score of 0.714, indicating substantial agreement (Viera and Garrett, 2005). Conflicts were resolved on a cross-team basis such that team A would resolve team B’s annotation conflicts and vice-versa.

During annotation we noticed that positive pairs typically fall into one of two categories. Either they were “true” subtype-supertype relations, such as $\langle \text{dog}, \text{animal} \rangle$, or they were individual-class relationships where the hyponym is typically a named entity and represents a *specific* instance of the more general class, as for example in $\langle \text{Nirvana}, \text{band} \rangle$.

Negative pairs were of a more diverse nature and included a range of different relations, such as co-hyponyms, meronyms or reverse hyponym-hypernyms. Negative pairs can also be comprised of two random nouns or two nouns without any semantic relation due to some amount of noise in extracting candidates solely on the basis of Hearst Patterns. Table 1 shows positive and negative examples from the dataset.

Pair	Relationship	Label
$\langle \text{dog}, \text{animal} \rangle$	hyponymy-hypernymy (Subtype)	True
$\langle \text{Nirvana}, \text{band} \rangle$	hyponymy-hypernymy (Instance)	True
$\langle \text{beef}, \text{stew} \rangle$	meronymy	False
$\langle \text{pedestrian}, \text{road} \rangle$	topical relatedness	False
$\langle \text{chemical}, \text{adenosine} \rangle$	reverse hyponymy-hypernymy	False
$\langle \text{cherry}, \text{plum} \rangle$	co-hyponymy	False
$\langle \text{form}, \text{situation} \rangle$	none	False

Table 1: Examples from our proposed HP4K dataset.

HP4K consists of 4369 pairs with a class distribution of 45:55 (positive : negative). Subsequently we split the dataset into a training and evaluation set, ensuring that there is no lexical overlap between the two sets. This resulted in a training set of size 3426 and an evaluation set of size 943⁷.

4.1 Models

We conduct experiments with two distributional vector space models, word2vec (Mikolov et al., 2013) and HyperVec (Nguyen et al., 2017). HyperVec is based on word2vec’s skip-gram architecture and leverages WordNet to optimise the word representations for the hypernym detection task. Hierarchical information is encoded in the norm of

⁷All resources are available from <https://github.com/tttthomasssss/le-augmentation>.

the learned vectors, such that lexemes higher up in the hypernymy hierarchy have larger norms than lexemes in lower parts.

For word2vec we use the 300-dimensional pre-trained Google News vectors⁸ and for HyperVec we trained 100-dimensional embeddings on a October 2013 Wikipedia dump (Wilson, 2015), using the recommended settings of Nguyen et al. (2017), as our augmentation sets contained many words that were OOV in the pre-trained HyperVec vectors⁹.

In our experiments, we consider a supervised scenario where a classifier predicts a hyponym-hypernym relation between two given word embeddings. We use two different models as classifier a logistic regression classifier (LR), and a 3-layer feedforward neural network (FF). In both cases, the classifier takes the aggregated hypothesised hyponym-hypernym pair as input and predicts whether the pair is in a hyponym-hypernym relation. We report a detailed overview of the model parameterisation in Appendix A.

The two models share the same procedure for aggregating the word embeddings of the hypothesised hyponym-hypernym pair. For data augmentation based on distributional composition, we use vector averaging as composition function, which gave substantially better performance than addition in preliminary experiments.

4.2 Results

For the FF network, we performed 10-fold cross-validation on the **Weeds** and **HP4K** training sets. As our evaluation for **LEDS** is based on a 20-fold cross-validation split, rather than a pre-defined training/evaluation split as for **Weeds** and **HP4K**, the same procedure for hyperparameter tuning is not straightforwardly applicable without exposing the model to some of the evaluation data. However, we found that the top parameterisations for **Weeds** and **HP4K** were quite similar and therefore applied hyperparameters to the FF model for **LEDS** that performed well in 10-fold cross-validation on **Weeds** and **HP4K**. For data augmentation and dataset extension, we consider the following amounts of additional data: {0.2K, 1K, 2K, 4K, 10K, 20K, 40K}. All augmentation sets are balanced between positive and negative pairs.

Figure 2 shows the increase in absolute points

⁸Available from: <https://code.google.com/archive/p/word2vec/>.

⁹We used the HyperVec code from www.ims.uni-stuttgart.de/data/hypervec.

of accuracy for the LR and FF model, as well as both vector spaces, averaged across all datasets. While in total data augmentation as well as dataset extension has a positive impact, the gains are larger for the FF model, suggesting that a higher capacity model is necessary to more effectively leverage the additional information from the augmentation source. Furthermore, before starting our experiments we expected that extending an existing dataset with WordNet represents an upper bound on performance, given that WordNet is manually annotated and curated. However in our experiments we found that data augmentation by either distributional composition or by using *GANDALF* remarkably *surpassed* performance of the WordNet-based extension technique regularly.

The effect of data augmentation and dataset extension in absolute points of accuracy on each dataset individually for the FF model is shown in Figure 3. It highlights consistent improvements across the board with only a single performance degradation in the case of extending the **LEDS** dataset with Hearst Patterns when using HyperVec-based word representations. The results per dataset for the LR model are presented in Appendix A and show that the LR model is less effective in leveraging the augmented data, causing more frequent performance drops. This suggests that models with more capacity are able to make more efficient use of additional data and are more robust in the presence of noise which is inevitably introduced by automatic methods.

Table 2 compares our FF model using word2vec embeddings with all proposed techniques for augmenting or extending a dataset. Our techniques are able to outperform a non-augmented model by 4-6 points in accuracy, representing a relative error reduction of 14%-26%. While the primary objective in this work is to improve an existing model setup with data augmentation, our augmented models compare favourably with previously published results.¹⁰ In general, data augmentation by distributional composition or by *GANDALF* overcomes two key weaknesses of simply extending a dataset with more data from WordNet or Hearst Patterns. First, many of the hyponym-hypernym pairs we mined from WordNet contain low-frequency words,

¹⁰We note that due to the use of different performance metrics and cross-validation splits, direct model-to-model comparisons are difficult on the **LEDS** and **Weeds** datasets. Thus we only compare to approaches that use the same evaluation protocol as we do.




Figure 2: Effect of data augmentation and dataset extension in absolute points of accuracy averaged across all datasets over the same model without augmentation or extension. The 2 heatmaps on the left are based on data augmentation (DC=Distributional Composition, GAN=*GANDALF*). The 2 heatmaps on the right use dataset extension (WN=WordNet, HP=Hearst Patterns).




Figure 3: Effect of data augmentation and dataset extension in absolute points of accuracy on all datasets for the FF model.

Model	Weeds	LEDS	HP4K
No Augmentation	0.72	0.77	0.67
Distributional Composition <i>GANDALF</i>	0.76	0.83	0.70
WordNet Extension	0.75	0.83	0.69
Hearst Patterns Extension	0.74	0.81	0.68
<i>Weeds et al. (2014)</i>	0.75	-	-
<i>Carmona and Riedel (2017)</i>	0.63	0.81	-

Table 2: Accuracy scores for the data augmentation and the two dataset extension strategies in comparison to the same FF model without any augmentation or extension.

which may have poor representations in our vector space models. Second, while using Hearst Patterns typically returned higher frequency words, the retrieved candidates frequently did not represent hyponymy-hypernymy relationships.

5 Analysis

The concrete amount of data augmentation, i.e. the number of additional hyponym-hypernym pairs that are added to the training set, represents a tuneable parameter. Figure 4 shows the effect of varying amounts of data augmentation for the FF model, using word2vec representations, across all datasets. We note that all amounts of additional augmentation data share the same quality, i.e. it is not the case that a smaller augmentation set consists of “better data” or contains less noise than a larger set. For the **Weeds** and **LEDS** datasets, peak performance is typically achieved with smaller amounts of additional data, whereas for the **HP4K**

dataset optimal performance is achieved with larger amounts of augmentation data. One explanation for the different augmentation characteristics of the **HP4K** dataset in comparison to the other two datasets is its independence of WordNet during the development of the dataset.

5.1 Data Augmentation in Space

In order to visualise what area of the vector space the *GANDALF* vectors and the composed vectors inhabit, we created a t-SNE (van der Maaten and Hinton, 2008) projection of the vector spaces in Figure 5. For the visualisation we produced the nearest neighbours of standard word2vec embeddings and augmentation embeddings for 5 exemplary words and project them into the same space. Figure 5 shows that the generated augmentation points, marked with an “x”, fit in with the real neighbours and do not deviate much from the “natural” neighbourhood of a given word. *GANDALF* vectors typically inhabit the edges of a neighbour cluster, whereas composed vectors are frequently closer to the cluster centroid. Table 3 lists the nearest neighbours for the example words. For word2vec and the composed augmentation vectors, we simply list the nearest neighbours of each query word. For *GANDALF* we list the nearest neighbours of the generated vector that correspond to actual words. For example, if the vector *GANDALF-234* is closest to the representations of *sugar*, *GANDALF-451* and *mountain*, we only list *sugar* and *mountain* as




Figure 4: Impact of different amounts of data augmentation.




Figure 5: t-SNE visualisation of the data augmentation spaces. Data points marked with “x” denote the representation as coming from the data augmentation set.

neighbours of *GANDALF-234*.

The composed neighbours for each word are typically closely related to the original query, e.g. *raw sugar* for *sugar*, or *zoo animal* for *animal*. The *GANDALF* neighbours on the other hand have a much weaker association with the query word, but are frequently still related to it *somehow* as in the example of *akpeteshie*, which is a spirit on sugar cane basis, as a neighbour for *sugar*.

5.2 Data Augmentation as Regularisation

In the past, a prominent criticism of distributional methods for hypernymy detection was that such models were found to frequently identify features of a prototypical hypernym in the distributional representations, rather than being able to dynamically focus on the relevant features that are indicative of a hypernymy relation for a specific pair of words (Weeds et al., 2014; Levy et al., 2015b). We therefore briefly investigate whether data augmentation can be used as a regularisation mechanism that helps prevent models from overfitting on prototypical hypernym features.

Table 4 shows the results on the **Weeds** dataset using a hypernym-only FF model with word2vec representations, in comparison to the same model variant that makes use of the hyponym and the hypernym. Ideally, we would hope to see weak performance for the hypernym-only and strong performance on the full model. This would indicate that the classifier does not rely on prototypical features

in the hypernym, but is able to focus on specific features in a given hyponym-hypernym pair. For data augmentation by distributional composition there appears to be a correlation between the performance of the hypernym-only and the full model, i.e. a stronger model on the whole dataset also results in better performance for the hypernym-only model. Hence augmentation by distributional composition might not be effective in helping the model to generalise in its current form. For augmentation with *GANDALF* however, performance for the full model improves, while performance of the hypernym-only model slightly drops, suggesting that the evoked *GANDALF* representations have a regularisation effect, while also improving generalisation. Hence, a fruitful avenue for future work will be further leveraging data augmentation for regularisation.

6 Conclusion

In NLP, in contrast to computer vision, data augmentation has not been applied as standard due to the apparent lack of universal rules for label-invariant language transformations.

We have considered the problem of hypernymy detection, and proposed two novel techniques for data augmentation. These techniques rely on semantic rules rather than an external knowledge source, and have the potential to generate almost limitless synthetic data for this task. We demonstrate that these techniques perform better than

Word	word2vec Neighbours	GANDALF Neighbours	Composed Neighbours
sugar	refined sugar, cane sugar, turbinado, cocoa, sugars	cmh rawalpindi, prescribed antipsychotic medication, sugar, mumtaz bhutto, akpeteshie	raw sugar, white sugar, brown sugar, sugar price, sugar industry
dog	dogs, puppy, pit bull, pooch, cat	ellis burks, sniffing glue, microchip implants, liz klekamp, cf rocco baldelli	pet cat, dog fighting, cat breed, rat terrier terrier breed
ocean	sea, oceans, pacific ocean, atlantic ocean, oceanic	pacific ocean, heavily vegetated, alaska aleutian, seagrasses, plutoid	ocean basin, pacific ocean, shallow sea, sea fish, sea mammal
animal	animals, animal welfare, dog, pet, cats	animals, pet, hallway feeds, poop scooping, panhandle animal welfare	first animal, adult animal, zoo animal, different animal, animal organization
energy	renewable energy, enery, electricity, enegy, fossil fuel	radial velocity measurements, stopped, renewable energy, bicycle advisory, steinkuehler	energy efficiency, solar energy, state energy, food energy, energy company

Table 3: Nearest neighbours for word2vec, GAN vectors and composed vectors.

Augmentation	Hypernym-Only	Full
No Data Augmentation	0.59	0.72
Distributional Composition (size=100)	0.60	0.74
Distributional Composition (size=500)	0.57	0.71
GANDALF (size=500)	0.58	0.75
GANDALF (size=1000)	0.60	0.73

Table 4: Accuracy for the hypernym-only and full models on the **Weeds** dataset with no, DC or GAN augmentation.

extending the training set with additional non-synthetic data, drawn from an external knowledge source in most cases. Our results are consistent across evaluation benchmarks, word vector spaces and classification architectures. We have also shown that our approach is effective even when the word vector space model has already been specialised for hypernymy detection.

Since WordNet is widely used as a source of information about semantic relations, we have proposed a new evaluation benchmark that is independent of WordNet. Whilst results are lower across the board on this dataset, suggesting that it is more difficult than the others, we see the same pattern of increasing performance with a more complex classifier and the use of data augmentation.

Future work includes leveraging data augmentation for more complex models and the extension of our approach to a multilingual setup as well as domains with a more specialised vocabulary such as Healthcare or Fashion.

Acknowledgements

Part of this research was supported by EPSRC grant no. 2129720: *Composition and Entailment in Distributed Word Representations*. We thank Mark Steedman for valuable feedback and insightful discussions that had substantial impact on the quality of the paper. We also thank Nicola De Cao, without whose help *GANDALF* wouldn't have been able to defeat the Balrog of mediocre performance. Lastly, we thank several mini-batches of reviewers for their ~~stochastic gradients~~ helpful comments that iteratively improved the paper, rejection after rejection,

until it *finally* converged towards acceptance.

References

- Jacob Andreas. 2020. [Good-enough compositional data augmentation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7556–7566, Online. Association for Computational Linguistics.
- Gintaras Barisevičius, Martin Coste, David Geleta, Damir Juric, Mohammad Khodadadi, Giorgos Stoulos, and Ilya Zaihrayeu. 2018. Supporting digital healthcare services using semantic web technologies. In *The Semantic Web – ISWC 2018*, pages 291–306. Springer International Publishing.
- Marco Baroni, Raffaella Bernardi, Ngoc-Quynh Do, and Chung-chieh Shan. 2012. Entailment above the word level in distributional semantics. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 23–32, Avignon, France. Association for Computational Linguistics.
- Duy-Cat Can, Hoang-Quynh Le, Quang-Thuy Ha, and Nigel Collier. 2019. [A richer-but-smarter shortest dependency path with attentive augmentation for relation extraction](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2902–2912, Minneapolis, Minnesota. Association for Computational Linguistics.
- Vicente Ivan Sanchez Carmona and Sebastian Riedel. 2017. How well can we predict hypernyms from word embeddings? a dataset-centric analysis. In *Proceedings of the European Chapter of the Association for Computational Linguistics*.
- Haw-Shiuan Chang, Ziyun Wang, Luke Vilnis, and Andrew McCallum. 2018. [Distributional inclusion vector embedding for unsupervised hypernymy detection](#). 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 485–495, New Orleans, Louisiana. Association for Computational Linguistics.

- Danilo Croce, Giuseppe Castellucci, and Roberto Basili. 2020. [GAN-BERT: Generative adversarial learning for robust text classification with a bunch of labeled examples](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2114–2119, Online. Association for Computational Linguistics.
- Marzieh Fadaee, Arianna Bisazza, and Christof Monz. 2017. [Data augmentation for low-resource neural machine translation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 567–573, Vancouver, Canada. Association for Computational Linguistics.
- Manaal Faruqui, Jesse Dodge, Sujay Kumar Jauhar, Chris Dyer, Eduard Hovy, and Noah A. Smith. 2015. [Retrofitting word vectors to semantic lexicons](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1606–1615, Denver, Colorado. Association for Computational Linguistics.
- Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. Bradford Books.
- M. Frid-Adar, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan. 2018. [Synthetic data augmentation using gan for improved liver lesion classification](#). In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 289–293.
- Fei Gao, Jinhua Zhu, Lijun Wu, Yingce Xia, Tao Qin, Xueqi Cheng, Wengang Zhou, and Tie-Yan Liu. 2019. [Soft contextual data augmentation for neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5539–5544, Florence, Italy. Association for Computational Linguistics.
- M. Geffet and I. Dagan. 2005. [The distributional inclusion hypotheses and lexical entailment](#). In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 107–114. Association for Computational Linguistics.
- Goran Glavaš and Ivan Vulić. 2019. [Generalized tuning of distributional word vectors for monolingual and cross-lingual lexical entailment](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4824–4830, Florence, Italy. Association for Computational Linguistics.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. [Generative adversarial nets](#). In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc.
- Marti Hearst. 1992. [Automatic acquisition of hyponyms from large text corpora](#). In *Proceedings of the International Conference on Computational Linguistics*.
- Sergey Ioffe and Christian Szegedy. 2015. [Batch normalization: Accelerating deep network training by reducing internal covariate shift](#). In *Proceedings of the 32Nd International Conference on International Conference on Machine Learning - Volume 37, ICML’15*, pages 448–456. JMLR.org.
- Min Junghyun, McCoy R. Thomas, Dipanjan Das, Pitler Emily, and Linzen Tal. 2020. [Syntactic data augmentation increases robustness to inference heuristics](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Seattle, Washington. Association for Computational Linguistics.
- Aishwarya Kamath, Jonas Pfeiffer, Edoardo Maria Ponti, Goran Glavaš, and Ivan Vulić. 2019. [Specializing distributional vectors of all words for lexical entailment](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 72–83, Florence, Italy. Association for Computational Linguistics.
- Dongyeop Kang, Tushar Khot, Ashish Sabharwal, and Eduard Hovy. 2018. [AdvEntuRe: Adversarial training for textual entailment with knowledge-guided examples](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2418–2428, Melbourne, Australia. Association for Computational Linguistics.
- Hwa-Yeon Kim, Yoon-Hyung Roh, and Young-Kil Kim. 2019. [Data augmentation by data noising for open-vocabulary slots in spoken language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 97–102, Minneapolis, Minnesota. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Lili Kotlerman, Ido Dagan, Idan Szpektor, and Maayan Zhitomirsky-Geffet. 2010. [Directional distributional similarity for lexical inference](#). *Special Issue of Natural Language Engineering on Distributional Lexical Semantics*, 4(16):359–389.
- Varun Kumar, Hadrien Glaude, Cyprien de Lichy, and William Campbell. 2019. [A closer look at feature space data augmentation for few-shot intent classification](#). In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 1–10, Hong Kong, China. Association for Computational Linguistics.

- Omer Levy, Yoav Goldberg, and Ido Dagan. 2015a. Improving distributional similarity with lessons learned from word embeddings. *Transactions of the Association for Computational Linguistics*, 3:211–225.
- Omer Levy, Steffen Remus, Chris Biemann, and Ido Dagan. 2015b. Do supervised distributional methods really learn lexical inference relations? In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 970–976, Denver, Colorado.
- Guanlin Li, Lemaou Liu, Guoping Huang, Conghui Zhu, and Tiejun Zhao. 2019. Understanding data augmentation in neural machine translation: Two perspectives towards generalization. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5689–5695, Hong Kong, China. Association for Computational Linguistics.
- Zhenhao Li and Lucia Specia. 2019. Improving neural machine translation robustness via data augmentation: Beyond back-translation. In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 328–336, Hong Kong, China. Association for Computational Linguistics.
- Angli Liu, Jingfei Du, and Veselin Stoyanov. 2019. Knowledge-augmented language model and its application to unsupervised named-entity recognition. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1142–1150, Minneapolis, Minnesota. Association for Computational Linguistics.
- Edward Loper and Steven Bird. 2002. Nltk: The natural language toolkit. In *Proceedings of the Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics - Volume 1*, pages 63–70. Association for Computational Linguistics.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605.
- Claudia Matos Veliz, Orphee De Clercq, and Veronique Hoste. 2019. Benefits of data augmentation for NMT-based text normalization of user-generated content. In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 275–285, Hong Kong, China. Association for Computational Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc.
- George A. Miller, Claudia Leacock, Randeep Tenji, and Ross T. Bunker. 1993. A semantic concordance. In *Proceedings of the Arpa Workshop on Human Language Technology*, pages 303–308.
- Neha Nayak, Mark Kowarsky, Gabor Angeli, and Christopher D. Manning. 2014. A dictionary of non-subjective adjectives. Technical Report CSTR 2014-04, Department of Computer Science, Stanford University.
- Thomas Neff. 2018. Data augmentation in deep learning using generative adversarial networks. Master’s thesis, Graz University of Technology.
- Kim Anh Nguyen, Maximilian Köper, Sabine Schulte im Walde, and Ngoc Thang Vu. 2017. Hierarchical embeddings for hypernymy detection and directionality. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 233–243, Copenhagen, Denmark. Association for Computational Linguistics.
- Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le. 2019. SpecAugment: A simple data augmentation method for automatic speech recognition. *CoRR*, abs/1904.08779.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. In *NIPS-W*.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12:2825–2830.
- Marek Rei, Daniela Gerz, and Ivan Vulic. 2018. Scoring lexical entailment with a supervised directional similarity network. In *ACL*.
- Laura Rimell. 2014. Distributional lexical entailment by topic coherence. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 511–519, Gothenburg, Sweden. Association for Computational Linguistics.
- Stephan Roller, Katrin Erk, and Gemma Boleda. 2014. Inclusive yet selective: Supervised distributional hypernymy detection. In *Proceedings of the International Conference on Computational Linguistics*.
- Stephan Roller and Katrin Erk. 2016. Relations such as hypernymy: Identifying and exploiting hearst patterns in distributional vectors for lexical entailment. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2163–2172, Austin, Texas. Association for Computational Linguistics.

- Stephen Roller, Douwe Kiela, and Maximilian Nickel. 2018. [Hearst patterns revisited: Automatic hypernym detection from large text corpora](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 358–363. Association for Computational Linguistics.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. 2016. [Improved techniques for training gans](#). In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 2234–2242. Curran Associates, Inc.
- Enrico Santus, Alessandro Lenci, and Qin Lu. 2014. [Chasing hypernyms in vector space with entropy](#). In *Proceedings of the European Chapter of the Association for Computational Linguistics*.
- 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.
- A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb. 2017. [Learning from simulated and unsupervised images through adversarial training](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2242–2251.
- Vered Shwartz, Yoav Goldberg, and Ido Dagan. 2016. [Improving hypernymy detection with an integrated path-based and distributional method](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2389–2398, Berlin, Germany.
- Vered Shwartz, Enrico Santus, and Dominik Schlechtweg. 2017. [Hypernyms under siege: Linguistically-motivated artillery for hypernymy detection](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 65–75, Valencia, Spain. Association for Computational Linguistics.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *Journal of Machine Learning Research*, 15:1929–1958.
- Ran Tian, Naoaki Okazaki, and Kentaro Inui. 2017. [The mechanism of additive composition](#). *Machine Learning*, 106(7):1083–1130.
- Yixuan Tong, Liang Liang, Boyan Liu, Shanshan Jiang, and Bin Dong. 2019. [Supervised neural machine translation based on data augmentation and improved training & inference process](#). In *Proceedings of the 6th Workshop on Asian Translation*, pages 147–151, Hong Kong, China. Association for Computational Linguistics.
- Anthony J. Viera and Joanne M. Garrett. 2005. [Understanding interobserver agreement: The kappa statistic](#). *Family medicine*, 37:360–3.
- Luke Vilnis and Andrew McCallum. 2014. [Word representations via gaussian embedding](#). *CoRR*, abs/1412.6623.
- Ivan Vulić and Nikola Mrkšić. 2018. [Specialising word vectors for lexical entailment](#). 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 1134–1145. Association for Computational Linguistics.
- 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.
- Julie Weeds, Daoud Clarke, Jeremy Reffin, David Weir, and Bill Keller. 2014. [Learning to distinguish hypernyms and co-hyponyms](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2249–2259, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Julie Weeds, David Weir, and Diana McCarthy. 2004. [Characterising measures of lexical distributional similarity](#). In *Proceedings of the 20th International Conference on Computational Linguistics*, pages 1015–1021, Geneva, Switzerland.
- Jason Wei and Kai Zou. 2019. [EDA: Easy data augmentation techniques for boosting performance on text classification tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6382–6388, Hong Kong, China. Association for Computational Linguistics.
- Benjamin Wilson. 2015. [The unknown perils of mining wikipedia](#). <https://blog.lateral.io/2015/06/the-unknown-perils-of-mining-wikipedia/>.
- Mengzhou Xia, Xiang Kong, Antonios Anastasopoulos, and Graham Neubig. 2019. [Generalized data augmentation for low-resource translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5786–5796, Florence, Italy. Association for Computational Linguistics.
- Jianhao Yan, Lin He, Ruqin Huang, Jian Li, and Ying Liu. 2019. [Relation extraction with temporal reasoning based on memory augmented distant supervision](#). In *Proceedings of the 2019 Conference of*

the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pages 1019–1030, Minneapolis, Minnesota. Association for Computational Linguistics.

A Supplemental Material

A.1 GANDALF Model Details

The generator and discriminator in *GANDALF* are single layer feedforward networks, with tanh activations and a dropout ratio (Srivastava et al., 2014) of 0.3. We used ADAM (Kingma and Ba, 2014) to optimise a binary cross entropy error criterion with a learning rate of 0.0002 and β values of 0.5 and 0.999. We found that *GANDALF* required quite a bit of wizardry to achieve strong performance and we found the website <https://github.com/soumith/ganhacks> very helpful. For example we applied label noise and soft labels (Salimans et al., 2016) and used a batch normalisation layer (Ioffe and Szegedy, 2015), which had the largest impact on model performance. *GANDALF* is implemented in PyTorch (Paszke et al., 2017) and we release our code on <https://github.com/ttthomassss/le-augmentation>.

A.2 Model Details

For our linear model we use the logistic regression classifier implemented in scikit-learn (Pedregosa et al., 2011). Our neural network model is 3-layer feedforward model implemented in PyTorch (Paszke et al., 2017).

We tuned the parameters of the Feedforward neural network by 10-fold cross-validation on the respective training sets, except for **LEDS**, where we chose the parameters on the basis of a model that performed well on the **Weeds** and **HP4K**. Our parameter grid consisted of activation function: {tanh, relu}, dropout: { 0.0, 0.1, 0.3 } and hidden layer sizes, where we considered { 200-200-200, 200-100-50, 200-50-30 } for Hypervec and { 600-600-600, 600-400-200, 600-300-100, 600-200-50} for word2vec. We furthermore considered 3 different aggregation functions: `diff` (Weeds et al., 2014), which simply takes the elementwise difference of the embedding pair; `asym` (Roller et al., 2014) which is the concatenation of the difference and the squared difference of the embedding pair; and `concat-asym` (Roller and Erk, 2016), which is the concatenation of the embedding pair, their difference, and their squared difference. We trained all models for 30 epochs with early stopping and used ADAM with a learning rate of 0.01 to optimise a cross entropy error criterion.

A.3 Results

Figure 6 below shows the effect of data augmentation in terms of points of Accuracy for the logistic regression classifier per vector space model and dataset. Unlike for the higher-capacity feedforward model, data augmentation frequently causes performance to go down for the simpler linear model. This suggests that more complex models are required to fully leverage the additional information from the augmentation sets.

Table 5 below gives an overview over the complete results for both classifier and vector space models, across all datasets. It shows the consistent positive effect of data augmentation on the more complex feedforward model in comparison to the logistic regression classifier, which is less robust to the small amounts of noise that is inevitably introduced by the automatic augmentation algorithm.




Figure 6: Effect of data augmentation and dataset extension in absolute points of accuracy on all datasets for the LR model.

Model	Weeds					LEDS					HP4K					Vector Space
	None	DC	GAN	WN	HP	None	DC	GAN	WN	HP	None	DC	GAN	WN	HP	
LR	0.69	<u>0.73</u>	<u>0.72</u>	<u>0.74</u>	<u>0.70</u>	0.81	0.80	0.78	0.80	0.80	0.64	0.61	0.63	0.63	0.63	word2vec
FF	0.72	0.76	<u>0.75</u>	<u>0.75</u>	<u>0.74</u>	0.77	0.83	0.80	0.83	0.81	0.67	<u>0.70</u>	0.71	<u>0.69</u>	<u>0.68</u>	
LR	0.70	<u>0.71</u>	<u>0.74</u>	<u>0.73</u>	<u>0.70</u>	0.79	0.76	0.78	0.79	0.78	0.63	<u>0.69</u>	<u>0.69</u>	<u>0.68</u>	<u>0.67</u>	hypervec
FF	0.71	<u>0.74</u>	0.75	0.75	<u>0.74</u>	0.79	<u>0.80</u>	0.81	<u>0.80</u>	0.75	0.66	<u>0.70</u>	0.72	<u>0.70</u>	<u>0.70</u>	

Table 5: Accuracy scores for the data augmentation strategies (DC and GAN), and the two dataset extension strategies (WN and HP), and the baseline that neither uses augmentation nor extension (None). Boldfaced results denote top performance per vector space and dataset, underlined results denote improved performance in comparison to the baseline without data augmentation.