

# KnowMAN: Weakly Supervised Multinomial Adversarial Networks

Luisa März <sup>◇,†</sup>, Ehsaneddin Asgari <sup>†</sup>, Fabienne Braune <sup>†</sup>,  
Franziska Zimmermann<sup>†</sup> and Benjamin Roth <sup>◇</sup>

<sup>◇</sup> Digital Philology, Research Group Data Mining and Machine Learning,  
University of Vienna, Austria

<sup>†</sup> NLP Expert Center, Data:Lab, Volkswagen AG, Munich, Germany

## Abstract

The absence of labeled data for training neural models is often addressed by leveraging knowledge about the specific task, resulting in heuristic but noisy labels. The knowledge is captured in labeling functions, which detect certain regularities or patterns in the training samples and annotate corresponding labels for training. This process of weakly supervised training may result in an over-reliance on the signals captured by the labeling functions and hinder models to exploit other signals or to generalize well. We propose KnowMAN, an adversarial scheme that enables to control influence of signals associated with specific labeling functions. KnowMAN forces the network to learn representations that are invariant to those signals and to pick up other signals that are more generally associated with an output label. KnowMAN strongly improves results compared to direct weakly supervised learning with a pre-trained transformer language model and a feature-based baseline.

## 1 Introduction

Neural approaches rely on labeled data sets for training. For many tasks and languages, such data is either scarce or not available at all. Knowledge-based weak supervision tackles this problem by employing *labeling functions (LFs)*. LFs are manually specified properties, e.g. keywords, that trigger the automatic annotation of a specific label. However, these annotations contain noise and biases that need to be handled.

A recent approach for denoising weakly supervised data is Snorkel (Ratner et al., 2020). Snorkel focuses on estimating the reliability of LFs and of the resulting heuristic *labels*. However, Snorkel does not address biases on the *input side* of weakly supervised data, which might lead to learned representations that overfit the characteristics of specific LFs, hindering generalization. We address the problem of overfitting to the LFs in this paper.

Other approaches tackle such overfitting by deleting the LF signal completely from the input side of an annotated sample: For example, Go et al. (2009) strip out emoticons that were used for labeling the sentiment in tweets, and Alt et al. (2019) mask the entities used for distant supervision of relation extraction training data (Mintz et al., 2009). However, as LFs are often constructed from the most prototypical and reliable signals (e.g., keywords), deleting them entirely from the feature space might – while preventing over-reliance on them – hurt prediction quality considerably. However, we find a way to blur the signals of the LFs instead of removing them.

In this work we propose KnowMAN (Knowledge-based Weakly Supervised Multinomial Adversarial Networks), a method for controllable *soft deletion* of LF signals, allowing a trade-off between reliance and generalization. Inspired by adversarial learning for domain adaptation (Chen and Cardie, 2018a; Ganin and Lempitsky, 2015), we consider LFs as domains and aim to learn a *LF-invariant* feature extractor in our model. KnowMAN is composed of three modules: a feature extractor, a classifier, and a discriminator. Specifically, KnowMAN employs a classifier that learns the actual task and an adversarial opponent, the LF-discriminator, that learns to distinguish between the different LFs. Upstream of both is the shared feature extractor to which the gradient of the classifier and the reversed gradient of the discriminator are propagated. In our experiments, the feature extractor for encoding the input is a multi-layer perceptron on top of either a bag-of-words vector or a transformer architecture, but KnowMAN is in principle usable with any differentiable feature extractor.

KnowMAN consistently outperforms our baselines by 2 to 30% depending on the dataset. By setting a hyperparameter  $\lambda$  that controls the influence of the adversarial part we can control the degree of

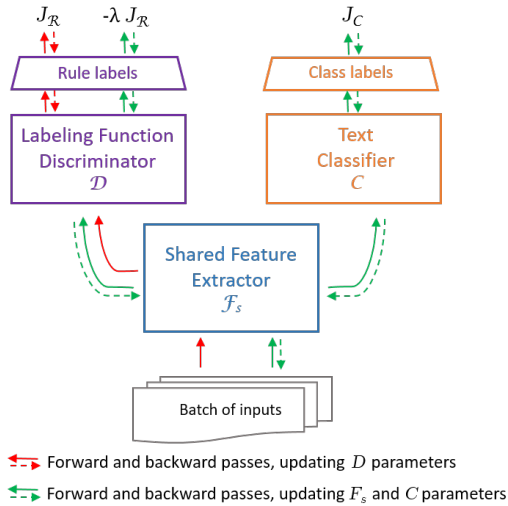


Figure 1: KnowMAN architecture. The figure depicts one iteration over a batch of inputs. The parameters of  $\mathcal{C}$  and  $\mathcal{F}_s$  are updated together, following the green arrows. The LF discriminator  $\mathcal{D}$  is updated following the red arrows. Solid lines indicate forward, dashed lines the backward pass.

discarding the information of LF-specific signals. The optimal  $\lambda$  value depends on the dataset and its properties.

The contributions of this work are i) proposing an adversarial architecture for controlling the influence of signals associated with specific LFs, ii) consistent improvements over weakly supervised baselines, iii) release of our code<sup>1</sup>. To our knowledge, we are the first that apply adversarial learning to overcome the noisiness of labels in weak supervision.

## 2 Method

Our approach is composed of three interacting modules i) the shared feature extractor  $\mathcal{F}_s$ , ii) the classifier  $\mathcal{C}$  and iii) the LF discriminator  $\mathcal{D}$ . The loss function of  $\mathcal{C}$  rewards the classifier  $\mathcal{C}$  for predicting the correct label for the instance, and the gradient is used for optimizing the shared feature extractor and classifier modules towards that goal. At the same time, the loss function for the LF-discriminator  $\mathcal{D}$  rewards predicting which LF was responsible for labeling an instance. However, in adversarial optimization, KnowMAN backpropagates the *reversed* gradient for the LF-discriminator, hence the information indicative for distinguishing between specific LFs is weakened throughout the network. The hyperparameter  $\lambda$  is used to control the level

<sup>1</sup><https://github.com/LuisaMaerz/KnowMAN>

of weakening the signals - the higher we choose the value the more influence is assigned to the discriminator information that goes into  $\mathcal{D}$ . The result of the interplay between classifier and LF-discriminator is a shared feature representation that is good at predicting the labels while reducing the influence of LF-specific signals, encouraging the shared feature extractor to take other information (correlated with all LFs for a class) into account.

In Figure 1, the arrows illustrate the training flow of the three modules. Due to the adversarial nature of the LF discriminator  $\mathcal{D}$ , it has to be trained with a separate optimizer (red arrows), while the rest of the network is updated with the main optimizer (green arrows). When  $\mathcal{D}$  is trained the parameters of  $\mathcal{C}$  and  $\mathcal{F}_s$  are frozen and vice versa.

To calculate the losses we utilize canonical negative log-likelihood loss (NLL) and use it for both, the classifier and the LF discriminator. The classification NLL can be formalized as:

$$\mathcal{L}_C(\hat{y}_i, y_i) = -\log P(\hat{y}_i = y_i) \quad (1)$$

where  $y_i$  is the (weakly supervised) annotated label and  $\hat{y}_i$  is the prediction of the classifier module  $\mathcal{C}$ , for a training sample  $i$ . Analogously, we can define the NLL for the LF discriminator:

$$\mathcal{L}_D(\hat{l}f_i, lf_i) = -\log P(\hat{l}f_i = lf_i) \quad (2)$$

where  $lf_i$  is the actual LF used for annotating sample  $i$  and  $\hat{l}f_i$  is the predicted LF by the discriminator  $\mathcal{D}$ . Accordingly, we minimize two different objectives within KnowMAN:

$$J_C = \sum_{i=1}^N \mathcal{L}_C(\mathcal{C}(\mathcal{F}_s(x_i); y_i)) \quad (3)$$

$$J_D = \sum_{i=1}^N \mathcal{L}_D(\mathcal{D}(\mathcal{F}_s(x_i); lf_i)) \quad (4)$$

Here the shared feature extractor has two different objectives: i) help  $\mathcal{C}$  to achieve better classification performance and ii) make the feature distribution invariant to the signals from the LFs. This is captured by the shared objective:

$$J_{\mathcal{F}_s} = J_C + \lambda \cdot (-J_D) \quad (5)$$

where  $\lambda$  is the parameter that controls the adversarial influence i.e. the degree of LF signal blur.  $-J_D$  is the reversed loss of the LF discriminator  $\mathcal{D}$  that represents  $\mathcal{C}$ 's adversarial opponent. In general, the

exact implementation or architecture of the individual modules is interchangeable and can be set up as required. This makes KnowMAN a universally applicable and easily customizable architecture.

### 3 Experiments

#### 3.1 Data

For our experiments we use three standard datasets for weak supervision.

**Spam.** Based on the YouTube comments dataset (Alberto et al., 2015) there is a smaller Spam dataset from Snorkel (Ratner et al., 2020) where the task is to classify if a text is relevant to a certain YouTube video or contains spam. This dataset is very small and does consist of a train and a test set only. For the 10 LFs keywords and regular expressions are used.

**Spouse.** This dataset for extracting the *spouse* relation has also been created by Snorkel, it is based on the Signal Media One-Million News Articles Dataset (Corney et al., 2016). The 9 LFs use information from a knowledge base, keywords and patterns. One peculiarity of this dataset is that over 90% of the instances do not hold a spouse relation.

**IMDb.** The IMDb dataset contains movie reviews that should be classified in terms of their sentiment (binary, positive or negative sentiment). The LFs used for this dataset are occurrences of positive and negative keywords from (Hu and Liu, 2004). A particular characteristic of this data set is the large amount of 6800 LFs, which constitutes a particular challenge to the Snorkel denoising framework. As a result Snorkel fails to calculate its generative model, since its memory consumption exceeds the available limit of 32GB RAM.

#### 3.2 Experimental setup

For the experiments we use two different methods for encoding the input: i) TF-IDF encoding and ii) a DistilBERT transformer. For TF-IDF encoding, we vectorize<sup>2</sup> the input sentences and feed them to a simple MLP. In the transformer setting, the sequences of words are encoded using a pre-trained DistilBERT. Similar to BERT (Devlin et al., 2019), DistilBERT is a masked transformer language model, which is a smaller, lighter, and faster version leveraging knowledge distillation while retaining 97% of BERT’s language understanding

<sup>2</sup>[https://scikit-learn.org/stable/modules/generated/sklearn.feature\\_extraction.text.TfidfVectorizer.html](https://scikit-learn.org/stable/modules/generated/sklearn.feature_extraction.text.TfidfVectorizer.html)

capabilities (Sanh et al., 2019).

Our encoder takes the representation of the CLS token from a frozen DistilBERT and learns a non-linear transformation with a drop-out layer to avoid overfitting (Srivastava et al., 2014):

$$h_i = \text{DistilBERT}(\text{Sentence}_i)_{[CLS]}$$

$$F_{s_i} = \text{Dropout}(\text{ReLU}(f(h_i)))$$

where  $\text{DistilBERT}(\cdot)_{[CLS]}$  generates the hidden state of the BERT’s classifier token (CLS) and the function  $f$  represents a linear transformation for the  $i^{\text{th}}$  sentence.

The classifier and discriminator networks following the feature extractor are in line with the implementation of Chen and Cardie (2018a) for domain-adversarial learning. Both are simple sequential models with dropout, batch normalization, *ReLU* activation and softmax as the last layer. Please see our code for implementation details. In the TF-IDF setup we use Adam (Kingma and Ba, 2014) for both optimizers. When using transformer encoding the  $\mathcal{D}$  optimizer again is Adam and the  $\mathcal{C}$  optimizer is AdamW (Loshchilov and Hutter, 2018), as this yielded more stable results.

**Baselines** For each input encoding we implemented several baselines. Weakly supervised TF-IDF (*WS TF-IDF*) and Weakly supervised DistilBERT (*WS DistilBERT*). Both calculate the labels for each instance in the train set based on their matching LFs. *WS TF-IDF* directly applies a logistic regression classifier to the input and the calculated labels. *WS DistilBERT* directly uses the DistilBERT uncased model for English (Sanh et al., 2019) as a prediction model. The second baseline (*Feature TF-IDF*, *Feature DistilBERT*) uses feature extractor and classifier layers of KnowMAN without taking the information of  $\mathcal{D}$  into account (this is equal to setting  $\lambda$  to zero). We also fine-tuned the pure language model (*Fine-tuned DistilBERT*) without further transformations and without integrating the KnowMAN architecture.

We also compare with training *TF-IDF* and *DistilBERT* models on labels denoised by Snorkel (*Snorkel TF-IDF*, *Snorkel DistilBERT*). However, Snorkel denoising failed for the IMDb data set due to the large amount of LFs.

**KnowMAN** We refer to the KnowMAN architecture as *TF-IDF KnowMAN* and *DistilBERT KnowMAN*. Depending on the dataset we choose different  $\lambda$  values. We also implemented two ways

	Spam	Spouse			IMDb
	Acc	P	R	F1	Acc
WS TF-IDF	0.87	0.12	<b>0.83</b>	0.20*	0.65*
Feature TF-IDF	0.91	0.12	0.76	0.21*	0.75*
Snorkel TF-IDF	0.81	0.18	0.63	0.28*	0.50*
KnowMAN TF-IDF	<b>0.94</b>	<b>0.16</b>	0.72	<b>0.35</b>	<b>0.77</b>
Fine-tuned DistilBERT	<b>0.92</b>	0.14	0.78	0.24	0.70
WS DistilBERT	0.87	0.09	<b>0.90</b>	0.17*	0.67*
Feature DistilBERT	0.86	0.18	0.80	0.29*	0.74
Snorkel DistilBERT	0.88	0.13	0.70	0.23*	0.49*
KnowMAN DistilBERT	0.90	<b>0.27</b>	0.67	<b>0.39</b>	<b>0.76</b>

Table 1: Results on the test sets. The \* indicates that KnowMAN performs significantly better than the marked model. For the Spouse data set we do report significance for the F1 scores only.

of evaluation and best model saving during training: i) evaluate after each batch and save the best model, ii) evaluate after a certain number of steps in between the batches and save the best model.

**Hyperparameters** We perform hyperparameter tuning using Bayesian optimization (Snoek et al., 2012) for the IMDb and Spouse datasets. For Spam, hyperparameters are not optimized, as no validation set is available. Sampling history and resulting hyperparameters are reported in the Appendix, Figures 2, 3 as well as hyperparameters chosen for the Spam data set.

**Evaluation** For the evaluation of the IMDb and the Spam datasets we use accuracy, for the Spouse dataset we use the macro F1 score of the positive class. To check statistical significance we use randomized testing (Yeh, 2000). Results are considered significant if  $\rho < 0.05$ .

### 3.3 Results

The results of the experiments are shown in Table 1. For the TF-IDF setup *KnowMAN TF-IDF* outperforms the baselines across all datasets. We find the optimal  $\lambda$  values as follows: Spam/Spouse/IMDb = 2/5/4.9. Using the additional feature extractor layer (*Feature TF-IDF*) is beneficial compared to direct logistic regression for all datasets. *Snorkel TF-IDF* can outperform the other two baselines for the Spouse dataset only.

Fine tuning of DistilBERT can not outperform our best KnowMAN. However, for the Spam dataset *Fine-tuned DistilBERT* gives better results than *KnowMAN DistilBERT* but still is worse than *KnowMAN TF-IDF*. Using *WS DistilBERT* gives the same results for the Spam dataset and slightly better results for IMDb, when compared to *WS*

*TF-IDF*, for Spouse the performance decreases. *Snorkel DistilBERT* can outperform the other two baselines for the Spam dataset only. The low performance of Snorkel on IMDb (for both DistilBERT and TF-IDF) might be explained by the very large amount of LF for this dataset. The *KnowMAN DistilBERT* results across datasets are in line with the TF-IDF setup - KnowMAN can outperform all baselines for the Spouse and IMDb dataset. We observe that  $\lambda = 5$  for Spouse and  $\lambda = 1$  for IMDb is most beneficial when using DistilBERT. For the Spam dataset we observe that KnowMAN (with  $\lambda = 2$ ) outperforms all the baselines, except for the fine-tuned DistilBERT model.

**Discussion** The performance drop we observe with DistilBERT for KnowMAN compared to the tf-idf setup of the IMDb dataset could be explained by implementation details. Due to memory issues we have to truncate the input when using DistilBERT. Since the movie reviews from IMDb are rather long this could harm performance. Since the Spam dataset is very small a single wrongly classified instance can have great impact on the results. This could explain why *KnowMAN TF-IDF* outperforms *KnowMAN DistilBERT* here as well. In general we could not perform hyperparameter optimization for the DistilBERT experiments due to memory issues. Therefore the results for that experiments might not have reached their optimum. However, the results show the value of using KnowMAN though. Overall our results confirm the assumption that KnowMAN enables a focus shift of the shared feature extractor from the signals of the LFs towards signals of other valuable information. KnowMAN consistently improves over the other experiments significantly - except

for the Spam dataset. We assume that the dataset size is too small to see significant changes in the results. Compared to the implementation of [Chen and Cardie \(2018a\)](#) we could not use the specialized domain feature extractor for our datasets in the experiments. This is due to the fact that our test sets do not contain information about LF matches. However, we will address this issue by integrating a mixture of experts module for the specialized feature extractor as recommended by [Chen et al. \(2019\)](#).

## 4 Related Work

Adversarial neural networks have been used to reduce the divergence between distributions, such as [Goodfellow et al. \(2014\)](#), [Chen et al. \(2018\)](#) and [Ganin and Lempitsky \(2015\)](#). The latter proposed an architecture for gradient reversal and a shared feature extractor. Unlike us, they focused on a binary domain discriminator. Similarly, ([Chen and Cardie, 2018a](#)) use an adversarial approach in a multinomial scenario for domain adaptation.

Some works on adversarial learning in the context of weak supervision focus on different aspects and only share similarity in name with our approach: [Wu et al. \(2017\)](#) use *virtual adversarial training* ([Miyato et al., 2017](#)) for perturbing input representations, which can be viewed as a general regularization technique not specific to weakly supervised learning. [Qin et al. \(2018\)](#); [Zeng et al. \(2018\)](#) use generative adversarial mechanisms for selecting *negative* training instances that are difficult to discriminate from heuristically annotated ones for a classifier.

Several approaches have focused on denoising the labels for weakly supervised learning ([Takamatsu et al., 2012](#); [Manning et al., 2014](#); [Lin et al., 2016](#)). Snorkel ([Ratner et al., 2020](#)) is one of the most general approaches in this line of work. However, Snorkel only models biases and correlations of LFs, and does not consider problems of weak supervision that may stem from biases in the features and learned representations.

A recent approach that focuses on denoising weakly supervised data is ([Sedova et al., 2021](#)). Knodle is a framework for comparison of different methods that improve weakly supervised learning. We use some of their datasets for our approach but denoise the signals of the LFs during training.

## 5 Conclusion

We propose KnowMAN - an adversarial neural network for training models with noisy weakly supervised data. By integrating a shared feature extractor that learns labeling function invariant features, KnowMAN can improve results on weakly supervised data drastically across all experiments and datasets in our setup. The experiments also show that the adverse effect of labeling function-specific signals is highly dependent on the datasets and their properties. Therefore, it is crucial to fine-tune the  $\lambda$  parameter on a validation set to find the optimal degree of blurring the labeling function signals. Since the modules in the KnowMAN architecture are easily exchangeable, KnowMAN can be applied to any architecture and dataset labeled with heuristic labeling functions.

## Acknowledgements

This research was funded by the WWTF through the project "Knowledge-infused Deep Learning for Natural Language Processing" (WWTF Vienna Research Group VRG19-008), by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) -RO 5127/2-1.

## References

- Tulio Alberto, Johannes Lochter, and Tiago Almeida. 2015. [Tubespam: Comment spam filtering on youtube](#). pages 138–143.
- Christoph Alt, Marc Hübner, and Leonhard Hennig. 2019. [Improving relation extraction by pre-trained language representations](#). In *Automated Knowledge Base Construction (AKBC)*.
- Xilun Chen, Ahmed Hassan Awadallah, Hany Hassan, Wei Wang, and Claire Cardie. 2019. [Multi-source cross-lingual model transfer: Learning what to share](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3098–3112, Florence, Italy. Association for Computational Linguistics.
- Xilun Chen and Claire Cardie. 2018a. [Multinomial adversarial networks for multi-domain text classification](#). 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 1226–1240, New Orleans, Louisiana. Association for Computational Linguistics.
- Xilun Chen and Claire Cardie. 2018b. [Multinomial adversarial networks for multi-domain text classification](#). 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 1226–1240, New Orleans, Louisiana. Association for Computational Linguistics.
- Xilun Chen, Yu Sun, Ben Athiwaratkun, Claire Cardie, and Kilian Weinberger. 2018. Adversarial deep averaging networks for cross-lingual sentiment classification. volume 6, pages 557–570.
- D. Corney, M. Albakour, Miguel Martinez-Alvarez, and Samir Moussa. 2016. What do a million news articles look like? In *NewsIR@ECIR*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yaroslav Ganin and Victor Lempitsky. 2015. Unsupervised domain adaptation by backpropagation. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML'15*, page 1180–1189. JMLR.org.
- Alec Go, Richa Bhayani, and Lei Huang. 2009. Twitter sentiment classification using distant supervision. *CS224N project report, Stanford*, 1(12):2009.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. In *NIPS*.
- Minqing Hu and Bing Liu. 2004. **Mining and summarizing customer reviews**. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '04*, page 168–177, New York, NY, USA. Association for Computing Machinery.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *International Conference on Learning Representations*.
- Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. Neural relation extraction with selective attention over instances. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2124–2133.
- Ilya Loshchilov and Frank Hutter. 2018. **Fixing weight decay regularization in adam**.
- Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. 2014. The stanford corenlp natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, pages 55–60.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 1003–1011.
- Takeru Miyato, Andrew M. Dai, and Ian Goodfellow. 2017. **Adversarial training methods for semi-supervised text classification**.
- Pengda Qin, Weiran Xu, and William Yang Wang. 2018. **DSGAN: Generative adversarial training for distant supervision relation extraction**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 496–505, Melbourne, Australia. Association for Computational Linguistics.
- Alexander Ratner, Stephen H. Bach, Henry R. Ehrenberg, Jason A. Fries, Sen Wu, and Christopher Ré. 2020. **Snorkel: rapid training data creation with weak supervision**. *VLDB J.*, 29(2-3):709–730.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. **Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter**. *CoRR*, abs/1910.01108.
- Anastasiya Sedova, Andreas Stephan, Marina Speranskaya, and Benjamin Roth. 2021. **Knodle: Modular weakly supervised learning with pytorch**. *CoRR*, abs/2104.11557.
- Jasper Snoek, Hugo Larochelle, and Ryan Adams. 2012. Practical bayesian optimization of machine learning algorithms. In *Proc. NIPS*.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Shingo Takamatsu, Issei Sato, and Hiroshi Nakagawa. 2012. Reducing wrong labels in distant supervision for relation extraction. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 721–729.
- Yi Wu, David Bamman, and Stuart Russell. 2017. Adversarial training for relation extraction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1778–1783.
- Alexander Yeh. 2000. **More accurate tests for the statistical significance of result differences**. In *COLING 2000 Volume 2: The 18th International Conference on Computational Linguistics*.
- Daojian Zeng, Yuan Dai, Feng Li, R Simon Sherratt, and Jin Wang. 2018. Adversarial learning for distant supervised relation extraction. *Computers, Materials & Continua*, 55(1):121–136.

## A Appendix

Please find our implementation at <https://github.com/LuisaMaerz/KnowMAN>.

### A.1 Dataset statistics

The datasets used for the KnowMAN experiments have different properties. Especially the number of labeling functions and the dataset sizes varies a lot.

dataset	classes	train/test samples	lfs
Spam	2	1586/250	10
Spouse	2	22254/2701	9
IMDb	2	40000/5000	6786

Table 2: Dataset statistics for KnowMAN experiments. Lfs are labeling functions.

### A.2 Hyperparameter optimization

We perform hyperparameter tuning using Bayesian optimization (Snoek et al., 2012). Bayesian Optimization is an approach that uses the Bayes Theorem to direct the search in order to find the minimum or maximum of a black-box objective function. In comparison with random search and grid search, it tends to obtain better hyperparameters in fewer steps by making a proper balance between exploration and exploitation steps. Our hyperparameter space includes batch size, dropout, number of iterations over  $\mathcal{D}$ , the shared hidden size of the models, learning rate for  $\mathcal{D}$  and  $\mathcal{F}_s, \mathcal{C}$  and the number of layers of  $\mathcal{C}, \mathcal{D}$  and  $\mathcal{F}_s$ . We implemented two ways of evaluation and best model saving during training: i) evaluate after each batch and save the best model, ii) evaluate after a certain number of steps in between the batches and save the best model. We also optimized the number of steps if logging in between a batch.

We evaluated the models for IMDb and Spouse on the respective validation set. For the Spam dataset, there is no development set available and we used the following hyperparameters for *KnowMAN TF-IDF* following the parameters used in Chen and Cardie (2018b): **Batch size:** 32, **dropout:** 0.4, **n critic:** 5, **lambda:** 2.0, **shared hidden size:** 700, **learning rate C & F:** 0.0001, **learning rate D:** 0.0001, **number of F layers:** 1, **number of C layers:** 1, **number of D layers:** 1.

### A.3 Experimental details

We ran our experiments on a DGX-1 server with one V100 GPU per experiment. The runtime of one model depends on the dataset: 0.25 hours for the Spam dataset, 0.25 hours for the Spouse dataset, and 8 hours for the IMDb dataset.




Figure 2: Sampled hyperparameters for KnowMAN TF-IDF on IMDb. Optimal hyperparameters are indicated in red.

**Batch size: 895, dropout: 0.275, n critic: 50, lambda: 4.9, shared hidden size: 585, learning rate C & F: 0.0001, learning rate D: 0.0001, number of F layers: 1, number of C layers: 1, number of D layers: 10.**

Histograms on the diagonal show how, for each hyperparameter, how many samples have been drawn during optimization.






Figure 3: Sampled hyperparameters for KnowMAN DistilBERT on Spouse. Optimal hyperparameters are indicated in red.

**Batch size: 16, dropout: 0.379, n critic: 1, lambda: 5.0, shared hidden size: 988, learning rate C & F: 0.0005, learning rate D: 0.001, number of F layers: 5, number of C layers: 10, number of D layers: 1.**

Histograms on the diagonal show how, for each hyperparameter, how many samples have been drawn during optimization.