

Active Learning for Sequence Tagging with Deep Pre-trained Models and Bayesian Uncertainty Estimates

Artem Shelmanov¹, Dmitri Puzyrev^{1,2,3}, Lyubov Kupriyanova¹, Denis Belyakov², Daniil Larionov^{2,3,4}, Nikita Khromov¹, Olga Kozlova³, Ekaterina Artemova², Dmitry V. Dylov¹, Alexander Panchenko¹

¹Skolkovo Institute of Science and Technology, Russia

²Higher School of Economics University, Russia

³Mobile TeleSystems (MTS), Russia

⁴Federal Research Center “Computer Science and Control” of Russian Academy of Sciences

{a.shelmanov,dmitrii.puzyrev,lyubov.kupriyanova,
nikita.khromov,d.dylov,a.panchenko}@skoltech.ru

{dobelyakov,dslarionov}@edu.hse.ru, elartemova@hse.ru
oskozlo9@mts.ru

Abstract

Annotating training data for sequence tagging of texts is usually very time-consuming. Recent advances in transfer learning for natural language processing in conjunction with active learning open the possibility to significantly reduce the necessary annotation budget. We are the first to thoroughly investigate this powerful combination for the sequence tagging task. We conduct an extensive empirical study of various Bayesian uncertainty estimation methods and Monte Carlo dropout options for deep pre-trained models in the active learning framework and find the best combinations for different types of models. Besides, we also demonstrate that to acquire instances during active learning, a full-size Transformer can be substituted with a distilled version, which yields better computational performance and reduces obstacles for applying deep active learning in practice.

1 Introduction

In many natural language processing (NLP) tasks, such as named entity recognition (NER), obtaining gold standard labels for constructing the training dataset can be very time and labor-consuming. It makes the annotation process expensive and limits the application of supervised machine learning models. This is especially the case in such domains as biomedical or scientific text processing, where crowdsourcing is either difficult or prohibitively expensive. In these domains, highly-qualified experts are needed to annotate data correctly, which dramatically increases the annotation cost.

Active Learning (AL) is a technique that can help to reduce the amount of annotation required

to train a good model by multiple times (Settles and Craven, 2008; Settles, 2009). Opposite to exhaustive and redundant manual annotation of the entire corpus, AL drives the annotation process to focus the expensive human expert time only on the most informative objects, which contributes to a substantial increase in the model quality.

AL is an iterative process that starts from a small number of labeled seeding instances. In each iteration, an *acquisition* model is trained on the currently annotated dataset and is applied to the large pool of unannotated objects. The model predictions are used by the AL query strategy to sample the informative objects, which are then further demonstrated to the expert annotators. When the annotators provide labels for these objects, the next iteration begins. The collected data can be used for training a final *successor* model that is used in a target application.

During AL, acquisition models have to be trained on very small amounts of the labeled data, especially during the early iterations. Recently, this problem has been tackled by transfer learning with deep pre-trained models: ELMo (Peters et al., 2018), BERT (Devlin et al., 2019), ELECTRA (Clark et al., 2020), and others. Pre-trained on a large amount of unlabeled data, they are capable of demonstrating remarkable performance when only hundreds or even dozens of labeled training instances are available. This trait suits the AL framework but poses the question about the usefulness of the biased sampling provided by the AL query strategies.

In this work, we investigate AL with the aforementioned deep pre-trained models and compare the results of this combination to the outcome of

the models that do not take advantage of deep pre-training. The main contributions of this paper are the following:

- We are the first to thoroughly investigate deep pre-trained models in the AL setting for sequence tagging of natural language texts on the widely-used benchmarks in this area.
- We conduct an extensive empirical study of various AL query strategies, including Bayesian uncertainty estimation methods with multiple Monte Carlo (MC) dropout variants (Gal and Ghahramani, 2016a; Gal et al., 2017). We find the best combinations of uncertainty estimates and dropout options for different types of deep pre-trained models.
- We show that to acquire instances during AL, a full-size Transformer can be substituted with a distilled version, which yields better computational performance and reduces obstacles for applying deep AL in practice.

The remainder paper is structured as follows. Section 2 covers relevant works on AL for sequence tagging. In Section 3, we describe the sequence tagging models. Section 4 describes the AL strategies used in the experiments. In Section 5, we discuss the experimental setup and present the evaluation results. Finally, Section 6 concludes the paper.

2 Related Work

AL for sequence tagging with classical machine learning algorithms and a feature-engineering approach has a long research history, e.g. (Settles and Craven, 2008; Settles, 2009; Marcheggiani and Artières, 2014). More recently, AL in conjunction with deep learning has received much attention.

In one of the first works on this topic, Shen et al. (2018) note that practical deep learning models that can be used in AL should be computational efficient both for training and inference to reduce the delays in the annotators' work. They propose a CNN-CNN-LSTM architecture with convolutional character and word encoders and an LSTM tag decoder, which is a faster alternative to the widely adopted LSTM-CRF architecture (Lample et al., 2016) with comparable quality. They also reveal disadvantages of the standard query strategy – least confident (LC), and propose a modification, namely Maximum Normalized Log-Probability (MNLP). Sidhant and Lipton (2018) experiment with Bayesian

uncertainty estimates. They use CNN-CNN-LSTM and CNN-BiLSTM-CRF (Ma and Hovy, 2016) networks and two methods for calculating the uncertainty estimates: Bayes-by-Backprop (Blundell et al., 2015) and the MC dropout (Gal and Ghahramani, 2016a). The experiments show that the variation ratio (Freeman, 1965) has substantial improvements over MNLP. In contrast to them, we additionally experiment with the Bayesian active learning by disagreement (BALD) query strategy proposed by Hounsby et al. (2011) and perform a comparison with variation ratio.

There is a series of works that tackle AL with a trainable policy model that serves as a query strategy. For this purpose, imitation learning is used in Liu et al. (2018); Vu et al. (2019); Brantley et al. (2020), while in (Fang et al., 2017), the authors use deep reinforcement learning. Although the proposed solutions are shown to outperform other heuristic algorithms with comparably weak models (basic CRF or BERT without fine-tuning) in experiments with a small number of AL iterations, they can be not very practical due to the high computational costs of collecting training data for policy models. Other notable works on deep active learning include (Erdmann et al., 2019), which proposes an AL algorithm based on a bootstrapping approach (Jones et al., 1999) and (Lowell et al., 2019), which concerns the problem of the mismatch between a model used to construct a training dataset via AL (acquisition model) and a final model that is trained on it (successor model).

Deep pre-trained models are evaluated in the AL setting for NER by Shelmanov et al. (2019). However, they perform the evaluation only on the specific biomedical datasets and do not consider the Bayesian query strategies. Ein-Dor et al. (2020) conduct an empirical study of AL with pre-trained BERT but only on the text classification task. Brantley et al. (2020) use pre-trained BERT in experiments with NER, but they do not fine-tune it, which results in suboptimal performance. In this work, we try to fill the gap by evaluating deep pre-trained models: ELMo and various Transformers, in the AL setting with practical query strategies, including Bayesian, and on the widely-used benchmarks in this area.

3 Sequence Tagging Models

We use a tagger based on the Conditional Random Field model (Lafferty et al., 2001), two BiLSTM-

CRF taggers (Lample et al., 2016) with different word representation models, and taggers based on state-of-the-art Transformer models.

3.1 Conditional Random Field

As a baseline for comparison, we use a feature-based linear-chain Conditional Random Field (CRF) model (Lafferty et al., 2001). It is trained to maximize the conditional log-likelihood of entire tag sequences. The inference is performed using the Viterbi decoding algorithm, which maximizes the joint probability of tags of all tokens in a sequence. The features used for the CRF model are presented in Appendix B.

3.2 BiLSTM-CRF

This model encodes embedded input tokens via a bidirectional long short term memory neural network (LSTM) (Hochreiter and Schmidhuber, 1997). BiLSTM processes sequences in two passes: from left-to-right and from right-to-left producing a contextualized token vector in each pass. These vectors are concatenated and are used as features in a CRF layer that performs the final scoring of tags.

We experiment with two versions of the BiLSTM-CRF model. The first one uses GloVe (Pennington et al., 2014) word embeddings pre-trained on English Wikipedia and the 5-th edition of the Gigaword corpus, and a convolutional character encoder (Ma and Hovy, 2016), which helps to deal with out-of-vocabulary words. As in (Chiu and Nichols, 2016), the model additionally leverages the basic capitalization features, which has been shown to be useful for achieving good performance with this model. We will refer to it as CNN-BiLSTM-CRF. We consider this model as another baseline that does not exploit deep pre-training. The second version of the BiLSTM-CRF model uses pre-trained medium-size ELMo (Peters et al., 2018) to produce contextualized word representations. ELMo is a BiLSTM language model enhanced with a CNN character encoder. This model does not rely on feature-engineering at all. We will refer to it as ELMo-BiLSTM-CRF.

3.3 Transformer-based Taggers

We perform AL experiments with state-of-the-art pre-trained Transformers: BERT (Devlin et al., 2019), DistilBERT (Sanh et al., 2019), and ELECTRA (Clark et al., 2020). The sequence tagger, in this case, consists of a Transformer “body” and

a decoding classifier with one linear layer. Unlike BiLSTM that encodes text sequentially, these Transformers are designed to process the whole token sequence in parallel with the help of the self-attention mechanism (Vaswani et al., 2017). This mechanism is bi-directional since it encodes each token on multiple neural network layers taking into account all other token representations in a sequence. These models are usually faster than the recurrent counterparts and show remarkable performance on many downstream tasks (Li et al., 2020).

BERT is a masked language model (MLM). Its main learning objective is to restore randomly masked tokens, so it can be considered as a variant of a denoising autoencoder. Although this objective makes the model to learn many aspects of natural languages (Tenney et al., 2019; Rogers et al., 2020), it has multiple drawbacks, including the fact that training is performed only using a small subset of masked tokens. ELECTRA has almost the same architecture as BERT but utilizes a novel pre-training objective, called replaced token detection (RTD), which is inspired by generative adversarial networks. In this task, the model has to determine what tokens in the input are corrupted by a separate generative model, in particular, a smaller version of BERT. Therefore, the model has to classify all tokens in the sequence, which increases training efficiency compared to BERT, and the RTD task is usually harder than MLM, which makes the model learn a better understanding of a language (Clark et al., 2020).

DistilBERT is a widely-used compact version of BERT obtained via a distillation procedure (Hinton et al., 2015). The main advantages of this model are the smaller memory consumption and the higher fine-tuning and inference speed achieved by sacrificing the quality. We note that good computational performance is a must for the practical applicability of AL. Delays in the interactions between a human annotator and an AL system can be expensive. Therefore, although DistilBERT is inferior compared to other Transformers in terms of quality, it is a computationally cheaper alternative for acquiring training instances during AL that could be used for fine-tuning bigger counterparts. Lowell et al. (2019) showed that a mismatch between an acquisition model and a successor model (the model that is trained on the annotated data for the final application) could eliminate the benefits of AL. The

similarity in the architectures and the shared knowledge between the smaller distilled Transformer and its ancestor potentially can help to alleviate this problem and deliver an AL solution that is both effective and practical.

4 Active Learning Query Strategies

We experiment with four query strategies for the selection of training instances during AL.

Random sampling is the simplest query strategy possible: we just randomly select instances from the unlabeled pool for annotation. In this case, there is no active learning at all.

Uncertainty sampling (Lewis and Gale, 1994) methods select instances according to some probabilistic criteria, which indicates how uncertain the model is about the label that was given to the instance. The baseline method is Least Confident (LC): the instances are sorted in the ascending order of probabilities of the most likely tag sequence. Let y_i be a tag of a token i that can take one class c out of C values, let x_j be a representation of a token j in an input sequence of length n . Then the LC score can be formulated as follows:

$$LC = 1 - \max_{y_1, \dots, y_n} \mathbb{P}[y_1, \dots, y_n | \{\mathbf{x}_j\}]$$

This score favors longer sentences since long sentences usually have a lower probability. Maximization of probability is equivalent to maximizing the sum of log-probabilities:

$$\begin{aligned} & \max_{y_1, \dots, y_n} \mathbb{P}[y_1, \dots, y_n | \{\mathbf{x}_j\}] \Leftrightarrow \\ \Leftrightarrow & \max_{y_1, \dots, y_n} \sum_i^n \log \mathbb{P}[y_i | \{y_j\} \setminus y_i, \{\mathbf{x}_j\}] \end{aligned}$$

To make LC less biased towards longer sentences, Shen et al. (2018) propose a normalization of the log-probability sum. They call the method **Maximum Normalized Log-Probability (MNLP)**. The MNLP score can be expressed as follows:

$$MNLP = - \max_{y_1, \dots, y_n} \frac{1}{n} \sum_i^n \log \mathbb{P}[y_i | \{y_j\} \setminus y_i, \{\mathbf{x}_j\}]$$

In our experiments, we use this normalized version of the uncertainty estimate since it has been shown to be slightly better than the classical LC (Shen et al., 2018), and it is commonly applied in other works on active learning for NER.

Following Siddhant and Lipton (2018), we implement extensions for the Transformer-based

and BiLSTM-based sequence taggers applying the MC dropout technique. Gal and Ghahramani (2016a) showed that applying a dropout at the prediction time allows us to consider the model as a Bayesian neural network and calculate theoretically-grounded approximations of uncertainty estimates by analyzing its multiple stochastic predictions. Like Shen et al. (2018) and Siddhant and Lipton (2018) we experiment with **variation ratio (VR)** (Freeman, 1965): a fraction of models, which predictions differ from the majority:

$$VR_i = 1 - \frac{\text{count}(\text{mode}(\{y_i^m\}_m^M), \{y_i^m\}_m^M)}{M}$$

$$VR = \frac{1}{n} \sum_i^n VR_i,$$

where M is a number of stochastic predictions.

Siddhant and Lipton (2018) and Shen et al. (2018) refer to this method as BALD. However, the **Bayesian active learning by disagreement (BALD)** proposed by Houlsby et al. (2011) leverage mutual information between outputs y_i and model parameters θ trained on a dataset \mathcal{D} :

$$BALD_i = H(y_i | x_i, \mathcal{D}) - E_{\theta \sim p(\theta | \mathcal{D})} [H(y_i | x_i, \theta)]$$

Let p_i^{cm} be a probability of a tag c for a token i that is predicted by a m -th stochastic pass of a model with the MC dropout. Then the BALD score can be approximated according to the following expression (Gal et al., 2017):

$$\begin{aligned} BALD_i & \approx \\ & - \sum_c^C \left(\frac{1}{M} \sum_m^M p_i^{cm} \right) \log \left(\frac{1}{M} \sum_m^M p_i^{cm} \right) \\ & + \frac{1}{M} \sum_{c,m}^{C,M} p_i^{cm} \log p_i^{cm} \\ & BALD \approx \frac{1}{n} \sum_i^n BALD_i \end{aligned}$$

Although BALD can be considered similar to VR, it potentially can give better uncertainty estimates than VR since it leverages the whole probability distributions produced by the model. However, this method has not been tested in the previous works on active learning for sequence tagging.

For Transformer models, we have two variants of the Monte Carlo dropout: the MC dropout on the last layer before the classification layer (MC last), and on all layers (MC all). We note that calculating

uncertainty estimates in the case of MC all requires multiple stochastic passes, and in each of them, we have to perform inference of the whole Transformer model. However, if we replace the dropout only in the classifier, multiple recalculations are needed only for the classifier, while it is enough to perform the inference of the massive Transformer “body” only once. Therefore, in this case, the overhead of calculating the Bayesian uncertainty estimates can be less than 1% (in the case of 10 stochastic passes for ELECTRA according to the number of parameters) compared to deterministic strategies like MNL.

The BiLSTM-CRF model has two types of dropout: the word dropout that randomly drops entire words after the embedding layer and the locked dropout (Gal and Ghahramani, 2016b) that drops the same neurons in the embedding space of a recurrent layer for a whole sequence. Therefore, for the BiLSTM-CRF taggers, we have three options: replacing the locked dropout (MC locked), replacing the word dropout (MC word), and replacing both of them (MC all). We should note that obtaining Bayesian uncertainty estimates does not require the recalculation of word embeddings (in our case, ELMo).

5 Experiments and Results

5.1 Experimental Setup

We experiment with two widely-used datasets for evaluation of sequence tagging models and AL query strategies: English CoNLL-2003 (Sang and Meulder, 2003) and English OntoNotes 5.0 (Pradhan et al., 2013). The corpora statistics are presented in Appendix A. In the experiments, we use the “native” tag schemes: IOB1 for CoNLL-2003 and IOB2 for OntoNotes 5.0.

Each experiment is an *emulation* of the AL cycle: selected instances are not presented to experts for annotation but are labeled automatically according to the gold standard. Each experiment is performed for each AL query strategy and is repeated five times for CoNLL-2003 and three times for OntoNotes to report the standard deviation. A random 2% subset in tokens of the whole training set is chosen for seeding, and instances with 2% of tokens in total are selected for annotation on each iteration. Overall, 24 AL iterations are made, so in the final iteration, half of the training dataset (in tokens) is labeled. We do not use validation sets provided in the corpora but keep 25% of the

MODEL	CoNLL-2003	ONTOOTES
CRF	78.2 ± NA	79.8 ± NA
CNN-BiLSTM-CRF	88.3 ± 0.2	82.9 ± 0.3
ELMo-BiLSTM-CRF	91.2 ± 0.2	87.2 ± 0.2
DISTILBERT	89.8 ± 0.2	87.3 ± 0.1
BERT	91.1 ± 0.2	88.2 ± 0.2
ELECTRA	91.5 ± 0.2	87.6 ± 0.2

Table 1: Performance of models built on the entire training datasets without active learning

labeled corpus as the development set and train models from scratch on the rest. For OntoNotes, the AL query strategy is applied to only the 50% subsample of the unlabeled dataset. This helps to reduce the duration of the inference phase. In preliminary experiments, we note that it does not seriously affect the AL performance, but the further reduction of the subsample size has a significant negative impact.


Details of models and training procedures are presented in Appendix C. We conduct AL experiments with the pre-selected model and training hyperparameters. Tuning hyperparameters of acquisition models during AL would drastically increase the duration of an AL iteration, which makes AL impractical. The hyperparameter optimization is reasonable for successor models. However, in preliminary AL experiments, tuning hyperparameters of successor models on the development set appeared to demonstrate an insignificant difference in the performance compared to using pre-selected hyperparameters that are fixed for all AL iterations. Therefore, since tuning hyperparameters for training a successor model on each AL iteration drastically increases the amount of computation, we fix them to the values pre-selected for experiments without AL.

The evaluation is performed using the span-based F1-score (Sang and Meulder, 2003). For query strategies based on the MC dropout, we make ten stochastic predictions.


5.2 Results and Discussion

5.2.1 Training on Entire Datasets

We evaluate the performance of models trained on 75% of the available training corpus while keeping the rest 25% as the development set similarly to the experiments with AL. From Table 1, we can find that for both OntoNotes and CoNLL-2003, the model performance pattern is almost the same.

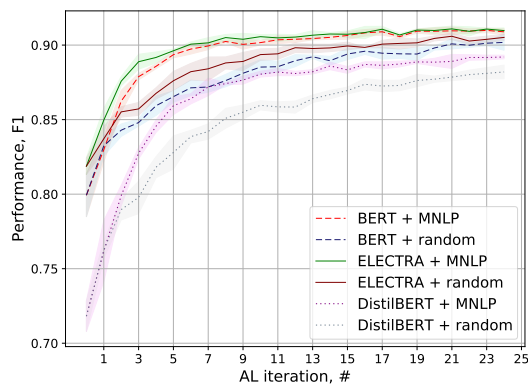


a) CoNLL-2003

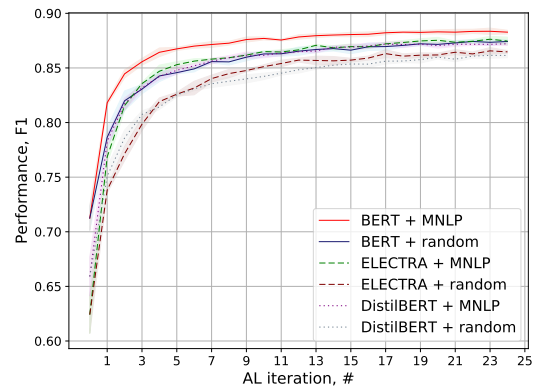


b) OntoNotes

The comparison of MNL and random query strategies for CRF and BiLSTM-based models.



c) CoNLL-2003



d) OntoNotes

The comparison of MNL and random sampling query strategies for Transformers.

Figure 1: The comparison of MNL and random query strategies

CRF, as the baseline model, has the lowest F1 score. Sequence taggers based on deep pre-trained models achieve substantially higher results compared to the classical CNN-BiLSTM-CRF model and CRF. BERT and ELECTRA significantly outperform ELMo-BiLSTM-CRF on OntoNotes, while on CoNLL-2003, all models have comparable scores. DistilBERT is behind larger Transformers. It also has a lower performance than ELMo-BiLSTM-CRF on CoNLL-2003 but similar scores on the OntoNotes dataset. We should note that our goal was not to show the state-of-the-art performance on each dataset but to determine reasonable hyperparameters and reference scores for experiments with AL.

5.2.2 Active Learning

The main results of experiments with AL are presented in Figures 1–3. AL shows significant improvements over the random sampling baseline

for all models and on both datasets. Performance gains are bigger for simpler models like CRF or CNN-BiLSTM-CRF without deep pre-training and for the more complex OntoNotes dataset. However, we see that both ELMo-BiLSTM-CRF and Transformers benefit from a biased sampling of AL query strategies, which magnifies their ability to be trained on extremely small amount of labeled data. For example, to get 99% of the score achieved with training on the entire CoNLL-2003 dataset, only 20% of the annotated data is required for the ELECTRA tagger accompanied with the best AL strategy. For OntoNotes and BERT, only 16% of the corpus is required. Random sampling requires more than 44% and 46% of annotated data for CoNLL-2003 and OntoNotes correspondingly.

MNL strategy. For the CoNLL-2003 corpus, the best performance in AL with the MNL query strategy is achieved by the ELECTRA model. It

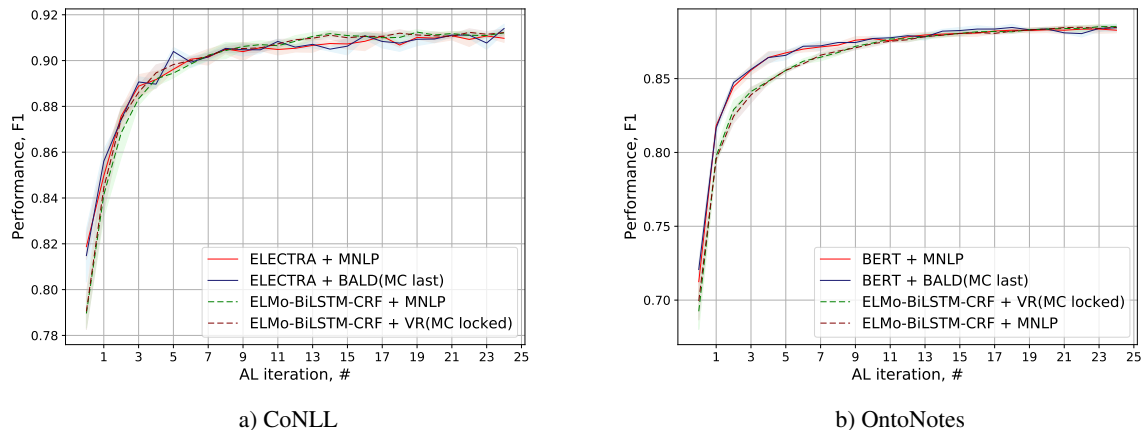


Figure 2: The comparison of the best query strategies and models overall

shows significantly better results in the beginning compared to BERT (see Figure 1c). ELECTRA also slightly outperforms the ELMo-BiLSTM-CRF tagger in the beginning and is on par with it on the rest of the AL curve (see Figure 2a). The CNN-BiLSTM-CRF model is always better than the baseline CRF but worse than the models that take advantage of deep pre-training (see Figure 1a). DistilBERT appeared to be the weakest model in the experiments on the CoNLL-2003 dataset. With random sampling, it is on par with the baseline CNN-BiLSTM-CRF model on most iterations, but with the MNLP strategy, DistilBERT significantly falls behind CNN-BiLSTM-CRF.

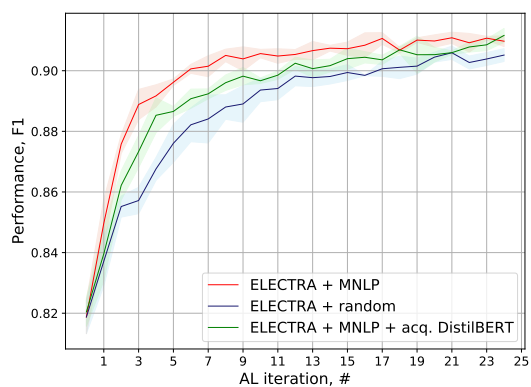
Although BERT is slightly worse than the ELECTRA and ELMo-BiLSTM-CRF taggers in the experiments on the CoNLL-2003 corpus, for OntoNotes, BERT has a significant advantage over them up to 2.3 % of the F1 score in the AL setting on early iterations. The ELMo-BiLSTM-CRF tagger falls behind the main Transformers on the OntoNotes corpus. This might be because the BiLSTM-based tagger is underfitted to the bigger corpus with only 30 training epochs. Likewise, the baseline CNN-BiLSTM-CRF model without deep pre-training significantly falls behind DistilBERT on this corpus for MNLP and random query strategies.

Bayesian active learning. Bayesian uncertainty estimates based on the MC dropout perform comparably with the deterministic MNLP strategy for Transformer-based and BiLSTM-CRF-based taggers (see Figure 2 and Tables 5, 6 in Appendix D). We consider that Bayesian uncertainty estimates do not outperform the deterministic uncertainty esti-

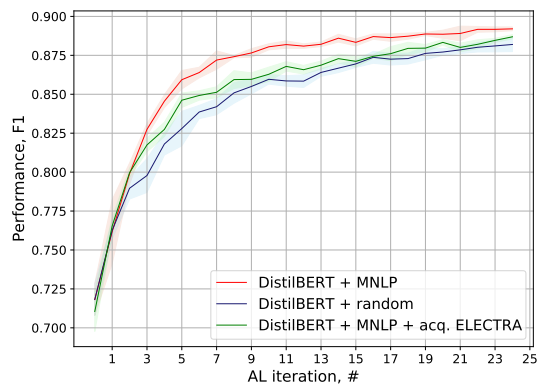
mates because the performance of the latter is very close to the maximum that can be achieved with the given amount of data.

We compare the performance of Bayesian AL strategies, when different dropout layers are replaced with the MC dropout. For ELMo-BiLSTM-CRF, we compare three options: replacing the dropout that follows word embeddings (embeddings acquired from ELMo), locked dropout in the recurrent layer, and both. Replacing the dropout that follows the embedding layer degrades the performance of AL significantly, especially for BALD. Replacing both yields the same performance as replacing only the locked dropout. We consider that the latter option is the best for AL since it requires fewer changes to the architecture. Overall, for both datasets, variation ratio has a slight advantage over the BALD strategy for the ELMo-BiLSTM-CRF model for all MC dropout options.

For Transformers, we compare two options: replacing the dropout only on the last classification layer and all dropouts in the model. When the variation ratio strategy is used, replacing only the last dropout layer with the MC dropout degrades the performance compared to MNLP, while replacing all dropout layers shows comparable results with it. However, for the BALD strategy, we see the inverse situation: replacing the last dropout layer leads to the significantly better performance than replacing all layers. This pattern can be noted for both ELECTRA and DistilBERT on the CoNLL-2003 corpus and for both BERT and DistilBERT on the OntoNotes corpus. Therefore, for Transformers, BALD with the MC dropout on the last layer is the best Bayesian query strategy since it

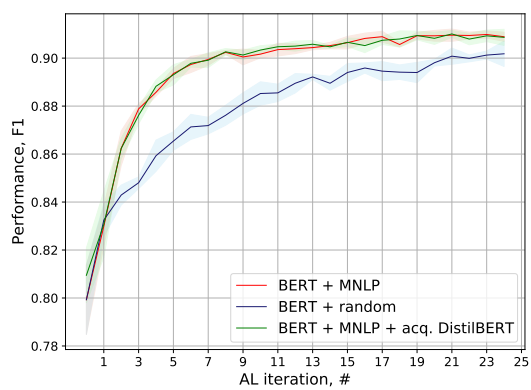


a) DistilBERT is an acquisition model




b) ELECTRA is an acquisition model

DistilBERT/ELECTRA as acquisition/successor models.



c) DistilBERT is an acquisition model



d) BERT is an acquisition model

DistilBERT/BERT as acquisition/successor models.

Figure 3: AL experiments on the CoNLL-2003 dataset, in which a successor model does not match an acquisition model.

Table 2: Duration of active learning iteration phases (seconds). The presented values correspond to the 5-th iteration of AL on the CoNLL-2003 dataset with the MNL query strategy. The duration is averaged over three runs.

MODEL	ACQ. MODEL TRAINING	QUERYING INST.	TOTAL
CNN-BiLSTM-CRF	1052	68	1120
ELMo-BiLSTM-CRF	339	100	439
ELECTRA	100	50	150
DISTILBERT	61	27	88

provides both good quality and low computational overhead.

Duration of AL iterations. The valuable benefit of deep pre-trained models is also their high inference and training speed, which helps to significantly reduce the duration of an AL iteration

and makes it feasible to implement text annotation tools empowered with interactive deep active learning. We measure the duration of the acquisition model training phase and the duration of the instance querying phase, which includes model inference on the whole unlabeled dataset (see Table 2). The experiments were conducted using the Nvidia V100 GPU and Xeon E5-2698 CPU.

According to obtained results, the large ELECTRA model can be trained more than ten times faster than the basic CNN-BiLSTM-CRF model because of a hardware-optimized architecture, a smaller number of the necessary training epochs, and the absence of the validation on the development set. ELECTRA also helps to reduce the duration of the query selection by more than 26%. The lightweight DistilBERT model is expectedly even faster. We also should note that using BiLSTM-

CRF with the pre-trained contextualized word representation model ELMo also helps to make the duration of AL iterations shorter compared to CNN-BiLSTM-CRF due to no need for training of the CNN word representation subnetwork (ELMo usually is not fine-tuned) and caching of the forward passes of ELMo during training across different epochs.

Mismatch between a successor model and an acquisition model. Since AL creates substantial computational overhead caused by training and inference of an acquisition model, for practical usage, it is reasonable to keep an acquisition model as lightweight as possible, while retaining the successor model complex and full-fledged. However, [Lowell et al. \(2019\)](#) demonstrate that the mismatch between an acquisition model and a successor model can diminish the benefits of applying the AL. They show that using AL can even harm the performance of the successor model compared to random sampling. We investigate the situation when an acquisition model does not match a successor model. Figure 3 shows the results of BERT, DistilBERT, and ELECTRA in this setting.

From Figures 3a and 3b, we can find that, when DistilBERT is used as an acquisition model for ELECTRA and vice versa, improvements achieved due to AL over the random sampling baseline are substantially lower compared to the “native” acquisition model. The same negative effect can be seen due to a mismatch between ELECTRA and BERT models in Figure 4 in Appendix E. These results support the findings presented in ([Lowell et al., 2019](#)). However, in our case, AL still gives notable improvements over the random sampling baseline. From Figure 3d, we can see that using BERT as an acquisition model for DistilBERT results only in a slight reduction of performance compared to using the “native” acquisition model. Moreover, in a reversed experiment, when DistilBERT is used as an acquisition model (see Figure 3c), there is no performance drop in AL at all. Such results for BERT-DistilBERT can be explained by the relationship between the distilled model and its ancestor resulting in similar uncertainty estimates for unlabeled instances in the annotation pool. This finding reveals the possibility of replacing a big acquisition model such as BERT with a distilled version that is faster and requires much less amount of memory. This can help to alleviate practical obstacles of deploying AL in real-world scenarios.

6 Conclusion

In this work, we investigated the combination of AL with sequence taggers that take advantage of deep pre-trained models. In the AL setting, these sequence taggers substantially outperform the models that do not use deep pre-training. We show that AL and transfer learning is a very powerful combination that can help to produce remarkably performing models with just a small fraction of the annotated data. For the CoNLL-2003 corpus, the combination of the best performing pre-trained model and AL strategy achieves 99% of the score that can be obtained with training on the full corpus, while using only 20% of the annotated data. For the OntoNotes corpus, one needs just 16%.

We performed a large empirical study of AL query strategies based on the Monte Carlo dropout in conjunction with deep pre-trained models and are the first to apply Bayesian active learning by disagreement to sequence tagging tasks. Bayesian active learning by disagreement achieves better results than the variation ratio for Transformers. However, we find that the variation ratio is slightly better for the ELMo-BiLSTM-CRF model. It is reasonable to use both MC dropout-based query strategies when only the last dropout layer works in a stochastic mode during the inference. This makes this type of query strategies suitable for practical usage due to little computational overhead. Finally, we demonstrate that it is possible to reduce the computational overhead of AL with deep pre-trained models by using a smaller distilled version of a Transformer model for acquiring instances.

In the future work, we are seeking to extend the empirical investigation of deep pre-trained models in active learning to query strategies that aim at better handling the batch selection of instances for annotation.

Acknowledgments

We thank the reviewers for their valuable feedback. This work was done in the framework of the joint MTS-Skoltech lab. The development of the experimental setup for the study of active learning methods and its application to sequence tagging tasks (Section 4) was supported by the Russian Science Foundation grant 20-71-10135. The Zhores supercomputer ([Zacharov et al., 2019](#)) was used for computations. We are very grateful to Sergey Ustyantsev for help with the implementation and the execution of the experiments.

References

- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. [Contextual string embeddings for sequence labeling](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1638–1649, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. 2015. [Weight uncertainty in neural network](#). In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 1613–1622. JMLR.org.
- Kianté Brantley, Hal Daumé III, and Amr Sharaf. 2020. [Active imitation learning with noisy guidance](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2093–2105, Online. Association for Computational Linguistics.
- Jason PC Chiu and Eric Nichols. 2016. [Named entity recognition with bidirectional LSTM-CNNs](#). *Transactions of the Association for Computational Linguistics*, 4:357–370.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: pre-training text encoders as discriminators rather than generators](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- 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.
- Liat Ein-Dor, Alon Halfon, Ariel Gera, Eyal Shnarch, Lena Dankin, Leshem Choshen, Marina Danilevsky, Ranit Aharonov, Yoav Katz, and Noam Slonim. 2020. [Active Learning for BERT: An Empirical Study](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7949–7962, Online. Association for Computational Linguistics.
- Alexander Erdmann, David Joseph Wisley, Benjamin Allen, Christopher Brown, Sophie Cohen-Bodénès, Micha Elsner, Yukun Feng, Brian Joseph, Béatrice Joyeux-Prunel, and Marie-Catherine de Marneffe. 2019. [Practical, efficient, and customizable active learning for named entity recognition in the digital humanities](#). 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 2223–2234, Minneapolis, Minnesota. Association for Computational Linguistics.
- Meng Fang, Yuan Li, and Trevor Cohn. 2017. [Learning how to active learn: A deep reinforcement learning approach](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 595–605, Copenhagen, Denmark. Association for Computational Linguistics.
- Linton C Freeman. 1965. *Elementary applied statistics: for students in behavioral science*. John Wiley & Sons.
- Yarin Gal and Zoubin Ghahramani. 2016a. [Dropout as a Bayesian approximation: Representing model uncertainty in deep learning](#). In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1050–1059.
- Yarin Gal and Zoubin Ghahramani. 2016b. [A theoretically grounded application of dropout in recurrent neural networks](#). In *Advances in Neural Information Processing Systems*, volume 29, pages 1019–1027. Curran Associates, Inc.
- Yarin Gal, Riashat Islam, and Zoubin Ghahramani. 2017. [Deep Bayesian active learning with image data](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1183–1192. PMLR.

- Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. [Distilling the knowledge in a neural network](#). *CoRR*, abs/1503.02531.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- Neil Houlsby, Ferenc Huszar, Zoubin Ghahramani, and Máté Lengyel. 2011. [Bayesian active learning for classification and preference learning](#). *CoRR*, abs/1112.5745.
- Rosie Jones, Andrew McCallum, Kamal Nigam, and Ellen Riloff. 1999. [Bootstrapping for text learning tasks](#). In *IJCAI-99 Workshop on Text Mining: Foundations, Techniques and Applications*, volume 1.
- 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 2001)*, Williams College, Williamstown, MA, USA, June 28 - July 1, 2001, pages 282–289. Morgan Kaufmann.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. [Neural architectures for named entity recognition](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California. Association for Computational Linguistics.
- David D. Lewis and William A. Gale. 1994. [A sequential algorithm for training text classifiers](#). In *Proceedings of the 17th Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval. Dublin, Ireland, 3-6 July 1994 (Special Issue of the SIGIR Forum)*, pages 3–12. ACM/Springer.
- Jing Li, Aixin Sun, Jianglei Han, and Chenliang Li. 2020. [A survey on deep learning for named entity recognition](#). *IEEE Transactions on Knowledge and Data Engineering*.
- Ming Liu, Wray Buntine, and Gholamreza Haffari. 2018. [Learning how to actively learn: A deep imitation learning approach](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1874–1883, Melbourne, Australia. Association for Computational Linguistics.
- David Lowell, Zachary C. Lipton, and Byron C. Wallace. 2019. [Practical obstacles to deploying active learning](#). 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 21–30, Hong Kong, China. Association for Computational Linguistics.
- Xuezhe Ma and Eduard Hovy. 2016. [End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1064–1074, Berlin, Germany. Association for Computational Linguistics.
- Diego Marcheggiani and Thierry Artières. 2014. [An experimental comparison of active learning strategies for partially labeled sequences](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 898–906, Doha, Qatar. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). 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 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Hwee Tou Ng, Anders Björkelund, Olga Uryupina, Yuchen Zhang, and Zhi Zhong. 2013. [Towards robust linguistic analysis using](#)

- [OntoNotes](#). In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 143–152, Sofia, Bulgaria. Association for Computational Linguistics.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A primer in bertology: What we know about how BERT works](#). *CoRR*, abs/2002.12327.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). pages 142–147.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter](#). *arXiv preprint arXiv:1910.01108*.
- Burr Settles. 2009. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison.
- Burr Settles and Mark Craven. 2008. [An analysis of active learning strategies for sequence labeling tasks](#). In *2008 Conference on Empirical Methods in Natural Language Processing, EMNLP 2008, Proceedings of the Conference, 25-27 October 2008, Honolulu, Hawaii, USA, A meeting of SIG-DAT, a Special Interest Group of the ACL*, pages 1070–1079. Association for Natural Language Processing.
- Artem Shelmanov, Vadim Liventsev, Danil Kireev, Nikita Khromov, Alexander Panchenko, Irina Fedulova, and Dmitry V. Dylov. 2019. [Active learning with deep pre-trained models for sequence tagging of clinical and biomedical texts](#). In *2019 IEEE International Conference on Bioinformatics and Biomedicine, BIBM 2019, San Diego, CA, USA, November 18-21, 2019*, pages 482–489. IEEE.
- Yanyao Shen, Hyokun Yun, Zachary C. Lipton, Yakov Kronrod, and Animashree Anandkumar. 2018. [Deep active learning for named entity recognition](#). In *International Conference on Learning Representations*.
- Aditya Siddhant and Zachary C. Lipton. 2018. [Deep Bayesian active learning for natural language processing: Results of a large-scale empirical study](#). In *Proceedings of the 2018 Con-*
- ference on Empirical Methods in Natural Language Processing*, pages 2904–2909, Brussels, Belgium. Association for Computational Linguistics.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008.
- Thuy-Trang Vu, Ming Liu, Dinh Phung, and Gholamreza Haffari. 2019. [Learning how to active learn by dreaming](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4091–4101, Florence, Italy. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Hugging-face’s Transformers: State-of-the-art natural language processing](#). *CoRR*, abs/1910.03771.
- Igor Zacharov, Rinat Arslanov, Maksim Gunin, Daniil Stefonishin, Andrey Bykov, Sergey Pavlov, Oleg Panarin, Anton Maliutin, Sergey Rykovanov, and Maxim Fedorov. 2019. [“Zhores” — petaflops supercomputer for data-driven modeling, machine learning and artificial intelligence installed in Skolkovo Institute of Science and Technology](#). *Open Engineering*, 9(1):512 – 520.

A Dataset Characteristics

Tables 3 and 4 present the characteristics of the datasets used in experiments.

Table 3: Characteristics of the OntoNotes 5.0 corpus (without the PT section)

	ENG.TRAIN	ENG.TEST
# OF TOKENS	1,088,503	152,728
# OF SENTENCES	59,924	8,262
ENTITY TYPES:		
PERSON	15,429	1,988
GPE	15,405	2,240
ORG	12,820	1,795
DATE	10,922	1,602
CARDINAL	7,367	935
NORP	6,870	841
MONEY	2,434	314
PERCENT	1,763	349
ORDINAL	1,640	195
LOC	1,514	179
TIME	1,233	212
WORK_OF_ART	974	166
FAC	860	135
EVENT	748	63
QUANTITY	657	105
PRODUCT	606	76
LANGUAGE	304	22
LAW	282	40
TOTAL ENTITIES:	81,828	11,257

Table 4: Characteristics of the CoNLL-2003 corpus

	ENG.TRAIN	ENG.TESTB
# OF TOKENS	203,621	46,435
# OF SENTENCES	14,041	3,453
ENTITY TYPES:		
LOC	7,140	1,668
PER	6,600	1,617
ORG	6,321	1,661
MISC	3,438	702
TOTAL ENTITIES:	23,499	5,648

B Features Used by the CRF Model

1. A lowercased word form.
2. Trigram and bigram suffixes of words.
3. Capitalization features.
4. An indicator that shows whether a word is a digit.
5. A part-of-speech tag of a word with specific info (plurality, verb tense, etc.)
6. A generalized part-of-speech.

7. An indicator whether a word is at the beginning or ending of a sentence.
8. The aforementioned characteristics for the next word and previous word except suffixes.

C Model and Training Details

C.1 CRF

We set CRF L1 and L2 regularization terms equal to 0.1, and limit the number of iterations by 100.

C.2 BiLSTM-CRF Taggers

We implement the BiLSTM-CRF sequence tagger on the basis of the `Flair` package¹ (Akbik et al., 2018). We use the same parameters for both types of BiLSTM-CRF models. The recurrent network has one layer with 128 neurons. During training, we anneal the learning rate by half, when the performance of the model stops improving on the development set for 3 epochs. After annealing, we restore the model from the epoch with the best validation score. The starting learning rate is 0.1. The maximal number of epochs is 30, and the batch size is 32. For optimization, we use the standard SGD algorithm.

C.3 Transformer-based Taggers

The implementation of Transformer-based taggers is based on the Hugging Face Transformers (Wolf et al., 2019)² library. We use the following pre-trained versions of BERT, ELECTRA, and DistilBERT accordingly: ‘*bert-base-cased*’, ‘*google/electra-base-discriminator*’, and ‘*distilbert-base-cased*’. The corrected version of Adam (AdamW from the Transformers library) is used for optimization with the base learning rate of 5e-5. The linear decay of the learning rate is applied following the (Devlin et al., 2019). The number of epochs is 4 and the batch size is 16. As in (Shen et al., 2018), we see that it is critical to adjust the batch size on early AL iterations, when only small amount of labeled data is available. We reduce the batch size to keep the number of iterations per epoch over 50, but limit the minimal batch size to 4.

¹<https://github.com/flairNLP/flair>

²<https://huggingface.co/transformers/>

D Comparison of Various MC Dropout Options

Table 5: Results of AL with various MC dropout options on the CoNLL-2003 dataset

		1	5	10	15	20	24
ELMo-BiLSTM-CRF	MNLP	84.1 ± 1.0	89.5 ± 0.2	90.7 ± 0.2	91.1 ± 0.1	91.1 ± 0.1	91.2 ± 0.1
	Random	82.5 ± 0.5	86.8 ± 0.4	88.7 ± 0.4	89.4 ± 0.2	90.1 ± 0.2	90.3 ± 0.2
	VR(MC word)	83.3 ± 1.0	89.4 ± 0.1	90.5 ± 0.1	90.8 ± 0.3	91.0 ± 0.2	91.1 ± 0.1
	VR(MC all)	84.7 ± 0.7	89.7 ± 0.2	90.7 ± 0.1	90.9 ± 0.1	90.9 ± 0.1	91.2 ± 0.2
	VR(MC locked)	84.4 ± 1.0	89.8 ± 0.2	90.6 ± 0.2	91.0 ± 0.3	91.1 ± 0.2	91.2 ± 0.2
	BALD(MC word)	83.5 ± 1.0	88.8 ± 0.5	90.3 ± 0.2	90.5 ± 0.2	90.9 ± 0.2	91.1 ± 0.2
	BALD(MC locked)	84.3 ± 0.4	89.5 ± 0.4	90.5 ± 0.0	90.8 ± 0.1	91.1 ± 0.1	91.3 ± 0.1
	BALD(MC all)	84.2 ± 0.7	89.6 ± 0.2	90.5 ± 0.2	90.7 ± 0.1	91.1 ± 0.1	91.1 ± 0.1
DistilBERT	MNLP	76.2 ± 2.0	85.9 ± 0.6	88.0 ± 0.2	88.3 ± 0.2	88.9 ± 0.3	89.2 ± 0.1
	Random	76.3 ± 0.6	82.8 ± 1.1	86.0 ± 0.3	87.0 ± 0.2	87.7 ± 0.2	88.2 ± 0.5
	VR (MC last)	78.0 ± 0.2	85.3 ± 0.2	87.7 ± 0.5	88.1 ± 0.4	88.6 ± 0.5	89.0 ± 0.3
	VR (MC all)	78.3 ± 1.4	86.3 ± 0.6	88.3 ± 0.1	88.8 ± 0.3	88.8 ± 0.3	89.0 ± 0.2
	BALD (MC last)	78.0 ± 0.9	86.0 ± 0.2	87.7 ± 0.3	88.3 ± 0.3	89.1 ± 0.0	89.1 ± 0.1
	BALD (MC all)	77.3 ± 1.0	85.6 ± 0.4	87.9 ± 0.1	88.6 ± 0.3	89.0 ± 0.1	89.0 ± 0.1
	ELECTRA	MNLP	85.0 ± 0.8	89.6 ± 0.1	90.6 ± 0.1	90.7 ± 0.2	91.0 ± 0.2
Random	83.7 ± 0.8	87.6 ± 0.6	89.4 ± 0.4	89.9 ± 0.3	90.4 ± 0.2	90.5 ± 0.2	
VR (MC last)	85.9 ± 0.7	88.7 ± 0.2	90.3 ± 0.2	90.6 ± 0.5	91.0 ± 0.3	91.0 ± 0.1	
VR (MC all)	84.3 ± 1.0	89.8 ± 0.2	90.4 ± 0.3	90.9 ± 0.2	91.0 ± 0.1	91.2 ± 0.3	
BALD (MC last)	85.6 ± 0.5	90.4 ± 0.2	90.5 ± 0.2	90.6 ± 0.4	90.9 ± 0.2	91.4 ± 0.2	
BALD (MC all)	84.0 ± 0.9	89.5 ± 0.4	90.6 ± 0.2	90.7 ± 0.3	91.1 ± 0.2	90.8 ± 0.3	

Table 6: Results of AL with various MC dropout options on the OntoNotes 5.0 dataset

Model	Query strat.	1	5	10	15	20	24
ELMo-BiLSTM-CRF	MNLP	79.5 ± 0.1	85.6 ± 0.2	87.4 ± 0.2	88.1 ± 0.1	88.4 ± 0.1	88.5 ± 0.1
	Random	75.1 ± 1.2	81.5 ± 0.2	84.6 ± 0.2	86.0 ± 0.2	86.9 ± 0.2	87.3 ± 0.2
	VR(MC word)	79.4 ± 0.6	85.1 ± 0.3	87.0 ± 0.1	87.7 ± 0.1	88.0 ± 0.2	88.2 ± 0.2
	VR(MC all)	79.6 ± 1.0	85.6 ± 0.1	87.4 ± 0.2	88.0 ± 0.3	88.3 ± 0.1	88.4 ± 0.1
	VR(MC locked)	79.8 ± 0.5	85.6 ± 0.1	87.4 ± 0.1	88.1 ± 0.2	88.4 ± 0.1	88.5 ± 0.2
	BALD(MC word)	78.6 ± 0.8	84.4 ± 0.3	86.3 ± 0.4	87.4 ± 0.3	87.9 ± 0.2	88.2 ± 0.0
	BALD(MC locked)	78.8 ± 1.5	85.1 ± 0.5	87.0 ± 0.3	87.9 ± 0.2	88.2 ± 0.2	88.4 ± 0.1
	BALD(MC all)	78.9 ± 1.3	85.0 ± 0.5	87.0 ± 0.6	87.9 ± 0.3	88.1 ± 0.2	88.3 ± 0.2
DistilBERT	MNLP	78.3 ± 0.5	84.8 ± 0.4	86.2 ± 0.2	86.9 ± 0.2	87.0 ± 0.2	87.2 ± 0.1
	Random	75.1 ± 0.9	82.5 ± 0.2	84.2 ± 0.3	85.4 ± 0.2	86.0 ± 0.2	86.1 ± 0.2
	VR (MC last)	77.8 ± 0.3	84.4 ± 0.3	85.8 ± 0.0	86.6 ± 0.3	87.0 ± 0.3	87.0 ± 0.3
	VR (MC all)	78.5 ± 0.2	84.6 ± 0.3	86.1 ± 0.0	86.7 ± 0.2	87.2 ± 0.1	87.2 ± 0.1
	BALD (MC last)	78.3 ± 0.3	84.9 ± 0.1	86.3 ± 0.0	86.9 ± 0.0	87.1 ± 0.2	87.0 ± 0.1
	BALD (MC all)	78.3 ± 0.1	84.8 ± 0.1	86.4 ± 0.2	87.2 ± 0.1	87.1 ± 0.1	87.2 ± 0.2
BERT	MNLP	81.8 ± 0.2	86.7 ± 0.1	87.7 ± 0.1	88.1 ± 0.2	88.3 ± 0.2	88.3 ± 0.2
	Random	78.7 ± 0.8	84.6 ± 0.1	86.3 ± 0.3	86.6 ± 0.3	87.2 ± 0.2	87.4 ± 0.1
	VR (MC last)	81.6 ± 0.8	86.4 ± 0.3	87.4 ± 0.2	87.7 ± 0.2	88.0 ± 0.4	88.1 ± 0.1
	VR (MC all)	82.2 ± 0.5	86.8 ± 0.3	87.7 ± 0.2	88.0 ± 0.2	88.5 ± 0.2	88.5 ± 0.2
	BALD (MC last)	81.7 ± 0.5	86.6 ± 0.3	87.7 ± 0.2	88.3 ± 0.2	88.4 ± 0.1	88.4 ± 0.2
	BALD (MC all)	82.3 ± 0.7	86.7 ± 0.0	87.8 ± 0.2	88.3 ± 0.2	88.4 ± 0.1	88.6 ± 0.3

E Experiments with a Mismatch between a Successor Model and an Acquisition Model

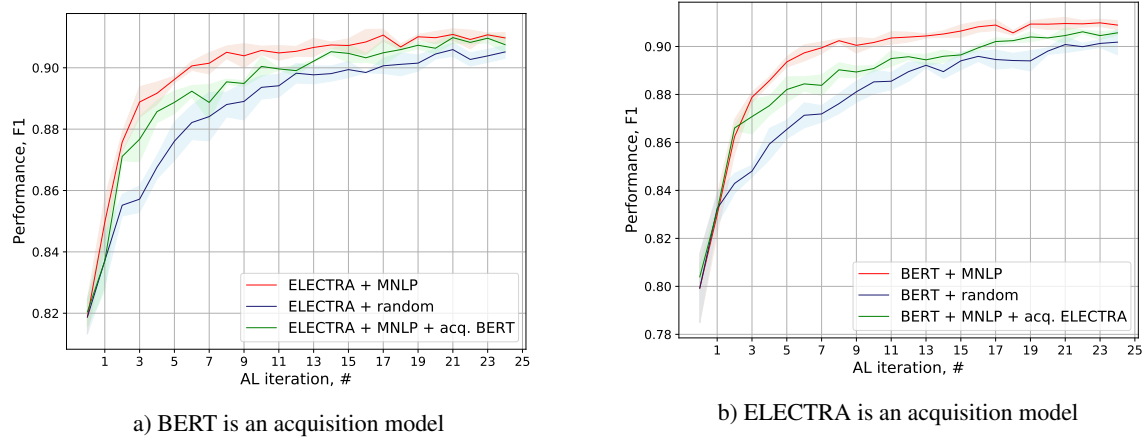


Figure 4: AL experiments on the CoNLL-2003 dataset, in which a successor model does not match an acquisition model (for BERT and ELECTRA).