




Using error decay prediction to overcome practical issues of deep active learning for named entity recognition

Haw-Shiuan Chang^{1,2}  · Shankar Vembu² · Sunil Mohan² · Rheeya Uppaal¹ · Andrew McCallum¹

Received: 15 December 2019 / Revised: 16 June 2020 / Accepted: 11 July 2020 / Published online: 5 August 2020
© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2020

Abstract

Existing deep active learning algorithms achieve impressive sampling efficiency on natural language processing tasks. However, they exhibit several weaknesses in practice, including (a) inability to use uncertainty sampling with black-box models, (b) lack of robustness to labeling noise, and (c) lack of transparency. In response, we propose a transparent batch active sampling framework by estimating the error decay curves of multiple feature-defined subsets of the data. Experiments on four named entity recognition (NER) tasks demonstrate that the proposed methods significantly outperform diversification-based methods for black-box NER taggers, and can make the sampling process more robust to labeling noise when combined with uncertainty-based methods. Furthermore, the analysis of experimental results sheds light on the weaknesses of different active sampling strategies, and when traditional uncertainty-based or diversification-based methods can be expected to work well.

Keywords Active learning · Transparency · Robustness to labeling noise · Black-box models · Clustering · Named entity recognition

Editors: Ira Assent, Carlotta Domeniconi, Aristides Gionis, Eyke Hüllermeier.

✉ Andrew McCallum
mccallum@cs.umass.edu

Haw-Shiuan Chang
hschang@cs.umass.edu

Shankar Vembu
shankar@argmix.com

Sunil Mohan
smohan@chanzuckerberg.com

Rheeya Uppaal
ruppaal@cs.umass.edu

¹ University of Massachusetts Amherst, College of Information and Computer Science, Amherst, MA, USA

² Chan Zuckerberg Initiative (CZI), Redwood City, CA, USA

1 Introduction

Deep neural networks achieve state-of-the-art results on many tasks, especially when a large amount of training data is available. Their success highlights the importance of reducing the cost of collecting labels on a large scale. Active learning can be used to select data samples that will most benefit a predictor's training, thereby reducing the amount of labeled data needed without hurting the predictor's accuracy. The effectiveness of uncertainty and disagreement-based¹ active learning methods have been demonstrated on several datasets for shallow predictors (Settles and Craven 2008; Settles 2009), and more recently also for deep learning predictors (Gal et al. 2017; Shen et al. 2018; Siddhant and Lipton 2018). Nevertheless, random sampling is still the most popular method to build new datasets in several domains, including natural language processing (Tomanek and Olsson 2009). This is due to the practical issues of deploying uncertainty-based active sampling (Settles 2011; Lowell et al. 2019), including its limited *applicability*, *robustness*, and *transparency*.

Applicability Uncertainty sampling selects samples with the lowest prediction confidence of a predictor and collects their labels. However, the uncertainty in prediction may be hard to estimate for some complex models. For example, a relation extraction system is often a pipeline that includes named entity extraction, entity linking, and sentence classification. Combining the uncertainty of different components is difficult and usually relies on ad hoc approaches (Reichart et al. 2008). Another common case is that state-of-the-art commercial software packages or online services may only provide final predictions (i.e., a black-box predictor) (Wang et al. 2017a). Our first research question is whether a general active learning method can be based only on the predictions from such black-box predictors.

Robustness Uncertainty sampling assumes that the labels of uncertain examples are the most informative training data. However, the strategy may select outliers or ambiguous examples for annotators due to their low confidence scores. This selection bias often introduces substantial labeling noise into the dataset (Mussmann and Liang 2018). In fact, outlier detection methods often improve prediction performance by adopting a sample selection strategy that is exactly opposite to uncertainty sampling (Bouguelia et al. 2018). It is usually hard to distinguish an informative training example from a misleading outlier, and blindly adopting uncertainty sampling may result in much worse performance than random sampling in practice. Our second research question is whether a general active learning method can be robust to labeling noise without relying on prior knowledge.

Transparency Active sampling introduces a sampling bias, but we often know little about its effects on the performance in different aspects. This lack of transparency causes many practical issues. Not understanding why an active learning method works well for a dataset, we have few insights into whether the sampling method will work for another similar dataset well enough to compensate the effort of trying active learning, whether the deep predictors will perform worse in some crucial aspects/classes, whether the selection will emphasize undesirable biases (e.g., a person's race), and whether the sampling efficiency improvement will vanish in the long run or after switching the deep predictor for uncertainty estimation.

¹ In this work, we view disagreement scores as a kind of uncertainty estimation to simplify our discussion.

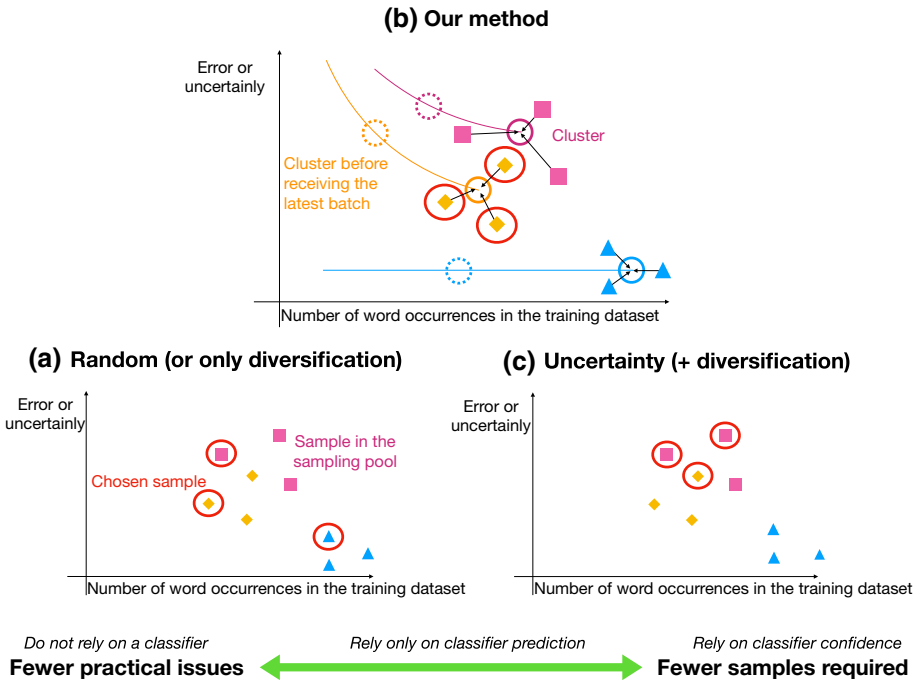


Fig. 1 Comparison of sampling strategies. Each node is a word in a NER problem and we plot its prediction error/uncertainty versus its count in the current training dataset. Random or diversification-based sampling **a** often selects samples irrelevant to the task, such as a blue triangle node. Uncertainty sampling **c** prefers to select samples with the highest error/uncertainty (e.g., pink square nodes). However, their error decay could be small because of some inherent label noise. Our method **b** is a novel balance in this spectrum which usually provides better performance than diversification while being more robust, explainable, and applicable than uncertainty sampling (+ diversification)

Furthermore, practitioners need to choose a sampling method for a new task and a given deep predictor before collecting labels (e.g., random, uncertainty or diversification sampling, or combination of the above). Existing deep active learning studies usually focus on improving prediction performance without exploring why and how much the performance gains depend on the predictor's specific ways of modeling feature interactions. The lack of such insights makes the choice extremely difficult (Lowell et al. 2019). Thus, our third research question is whether we can have a general analysis tool to provide insights into when and why an active sampling method works, and forecast the potential benefits in different aspects based on a small number of existing labels.

To answer the above research questions, we first illustrate three kinds of samples in Fig. 1. The blue triangle nodes are easy samples, which means their prediction errors are low and remain almost unchanged if we collect more such labels. The pink square nodes represent samples with noisy labels, which means their prediction errors or uncertainties are high but decrease slowly as more labels are seen. To select informative samples like the orange diamond nodes, we propose a batch active learning framework that maximizes the error reduction of each sample batch.

In this framework, we propose three sampling methods. When it's not possible to access the uncertainty of the deep predictor but labels of a validation dataset are available, we

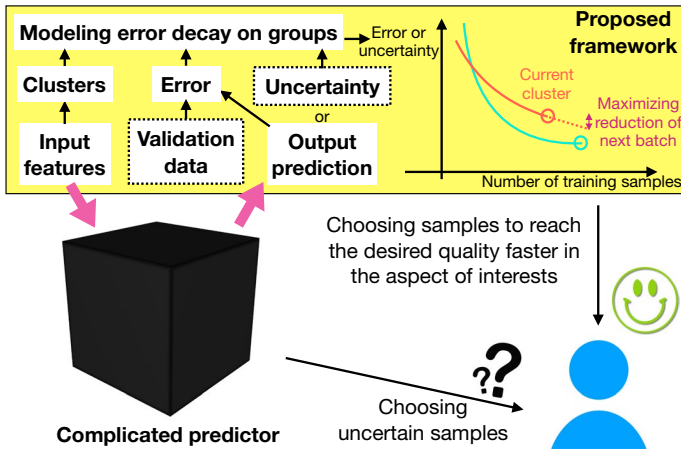


Fig. 2 The uncertainty sampling for a black-box predictor is hard to interpret. Our proposed framework predicts the reduction of error/uncertainty on clusters of samples, optionally with the help of validation data. The resulting strategies provide a more direct explanation of the choices made by batch active learning

cluster the samples and predict the validation error reduction on each cluster as shown in Fig. 2. The first method is called error decay on groups (EDG). Without a validation dataset, the second method approximates the error decay using prediction changes on groups. Finally, if uncertainty is available, we model the uncertainty reduction of each sample. In Sect. 3, we describe the first method in detail, and view the second and third methods as its extensions and denote them as EDG_ext1 and EDG_ext2, respectively.

Analyzing error decay on clusters results in several practical benefits. We can model the error decay using only model predictions; noisy samples will not be selected due to their low error decay; the strategy of our active sampling is interpretable (e.g., sampling more uppercased words in a named entity recognition problem due to their larger prediction error decay); and we can predict the performance gains from active learning in different clusters/aspects. Nevertheless, having only a few labels, we found that achieving the practical advantages sacrifices some sampling efficiency. To feasibly approximate the error reduction, we assume independence between samples in different clusters. For instance, if a sample is a word in a sentence, our selection method assumes its neighboring words do not affect the error decay of the word. In practice, violation of this assumption causes sub-optimal sampling efficiency when compared to uncertainty-based active learning.

We apply our framework to named entity recognition (NER) problems where our independence assumption is famously violated, and comprehensively evaluate the pros and cons of the proposed framework.² We summarize our results from NER experiments in Table 1, which shows that our framework provides novel ways to trading partial sampling efficiency gain for better *applicability*, *robustness*, and *transparency*.

² The proposed methods could be extended to other types of classification and sequential tagging problems given proper ways of clustering samples and a family of proper error decay functions.

Table 1 Comparison of different sampling approaches

	Proposed			Classic	
	EDG	_ext1	_ext2	US (+ Div)	Rnd/Div
Interpretable Strategy	Yes	Yes	No	No	Yes
Robustness to Noise	High	Med	Med	Low	High
Output Required	Pred	Pred	Prob	Prob	None
Validation Required	Yes	No	No	No	No
Label Cost Reduction	Med	Med	High	High	Low

We compare transparency (i.e., interpretable strategy), robustness to labeling noise, required output from the predictor, whether validation data is required or not, and cost reduction of collecting clean labels. US means uncertainty sampling (including disagreement-based sampling), Rnd means random sampling, and Div means diversification-based sampling. Med, pred, and prob means medium, predictor prediction, and prediction probability, respectively. In black-box models, only predictions of the predictor are available. A more ideal property is highlighted with a darker shade in the table

1.1 Summary of main contributions

1. We propose a novel active learning method, EDG, which models validation error decay curves on clusters of samples. We extend the method by replacing validation error decay with prediction difference decay or uncertainty decay to avoid relying on validation data. This demonstrates the flexibility of EDG framework.
2. In one synthetic and three real-world NER datasets, we show that EDG significantly outperforms the diversification baseline for the black-box model with or without the help of validation data. This demonstrates the effectiveness and applicability of EDG framework.
3. Modeling the error decay on clusters can be used as an analysis tool for arbitrary active learning methods. The experimental results show that no single method always wins and the proposed analysis tool provides intuitions and guidelines on selecting a specific method. This demonstrates the transparency of EDG framework.
4. We propose a new evaluation method based on pseudo labels. The method allows us to test active learning methods on a large sampling pool and test their robustness to systematic labeling noise.
5. In the experiments on pseudo labels and synthetic dataset, we show that combining EDG with state-of-the-art uncertainty sampling methods (i.e., choosing the samples with the highest uncertainty decay rather than uncertainty) improves sampling efficiency in the presence of systematic labeling noise, random labeling noise, and model change. This demonstrates the robustness of EDG framework.

2 Related work

Several studies aim at making active learning practical (Settles 2011). For example, Phillips et al. (2018) extend the techniques of interpreting a classifier's predictions to explain the active sampling process, but still rely on the uncertainty estimation of the classifier.

Bloodgood and Vijay-Shanker (2009) survey and propose stopping criteria based on overall error decay. Instead, we propose a novel active sampling method by modeling error decay on groups of samples.

Recently, active learning on black-box models has attracted research attention due to practical needs. Wang et al. (2017a) focus on improving a black-box semantic role labeling model using another neural network with low transparency. Rubens et al. (2011) propose estimating the variance of predictions of a black-box regressor, which is computationally prohibitive unless applied to simple models such as a linear regressor.

Popular active learning methods such as uncertainty sampling are not robust to noise (Musmann and Liang 2018), and there is a body of research in active learning (referenced below) to address this issue. However, the main focus of previous studies is to identify and post-process the noisy labels given some prior knowledge of the noise-to-signal ratio. Instead, we propose active learning methods which are inherently robust to noise by avoiding the selection of difficult samples for annotators in the first place.

The methods to identify noisy labels are similar to the methods that measure the uncertainty of each sample (Sheng et al. 2008; Bouguelia et al. 2018), which could be based on disagreement between workers (Sheng et al. 2008; Zhao et al. 2011), disagreement between manually created labels and automatically generated predictions (Bouguelia et al. 2015; Khetan et al. 2018), model uncertainty (Sheng et al. 2008; Kremer et al. 2018; Bouguelia et al. 2018), or estimation of workers' quality (Zhang et al. 2015; Khetan et al. 2018). After identifying noisy labels, different methods adopt different strategies such as relabeling (Sheng et al. 2008; Bouguelia et al. 2015), excluding/down-weighting the noisy labels (Khetan et al. 2018), or both (Zhao et al. 2011; Zhang et al. 2015; Bouguelia et al. 2018), or acquiring high-quality labels (Kremer et al. 2018).

Some approaches, such as Dasgupta (2011), cluster input features and diversify samples by choosing them from different clusters, without considering information from the predictors being trained, and often showing only limited improvements in sampling efficiency. Recent approaches (Settles and Craven 2008; Wei et al. 2015; Sener and Savarese 2018; Ravi and Larochelle 2018) have combined uncertainty and diversification (e.g., by multiplying the informativeness and representativeness scores). By directly modeling the error reduction, our proposed methods naturally balance the two criteria without relying on model uncertainty estimation. This makes our methods robust to labeling noise and applicable to black-box models.

Another direction for deep active learning is to learn an error reduction predictor or a sample selector. However, the selection model is either only applicable to a simple model like naïve Bayes (Roy and McCallum 2001; Fu et al. 2018), or requires a large amount of data to train a complex predictor or selector. The training data for the non-transparent selection model usually needs to come from a similar task. For instance, they could be images with different labels for image classification, other users for recommendation, or another language for NER (Bachman et al. 2017; Fang et al. 2017; Ravi and Larochelle 2018). Although transferring the error reduction predictor between different types of datasets is possible (Konyushkova et al. 2017), it is unclear on which dataset pairs such a transfer would work (Koshorek et al. 2019).

A challenge related to active learning is to discover blind spots of the predictor (also called unknown unknowns). Lakkaraju et al. (2017) view this problem as a multi-armed bandit problem; they first cluster the samples in a pool, and select the samples from a group more often when more unknown unknowns are discovered from the group. However, it is not clear whether this strategy yields better sampling efficiency in terms of the predictor's performance.

Chen et al. (2018) define unfairness as the difference in classification accuracy between two groups and suggest additional data collection as one of the remedies for such unfairness. However, they do not study an active sampling method to efficiently reduce the classification error or the unfairness.

In the extensions of our method, we maximize the prediction change (i.e., EDG_ext1 in our experiments). The method is related to a type of active learning strategies based on maximizing model change (Settles et al. 2008; Settles and Craven 2008). However, these methods do not address the practical challenges such as transparency, applicability to black-box models, and robustness to labeling noise.

3 Method

The main goal of batch active learning is to reduce the error $E(C, D_U)$ of a classifier or tagger C on an unseen testing dataset D_U after labeling a fixed number of samples. To simplify the explanation, we first assume that testing error can be well-approximated by validation error, and the methods that do not require validation data will be described later as extensions in Sect. 4.

Based on external datasets, previous work maximizes the error reduction by modeling interactions among various types of signals such as uncertainty estimation, input features, and the state of the sampling process (Konyushkova et al. 2017; Bachman et al. 2017; Fang et al. 2017; Ravi and Larochelle 2018). In order to not rely on a large external dataset, we first cluster samples into multiple groups, and assume that the validation error of the samples in each group only depends on the number of annotated samples in the same group. Then, estimating the error reduction is decomposed into simple one-dimensional regression problems which can be done by observing only a few pairs of validation errors and its corresponding number of samples.

In the following subsections, we describe our framework in the context of solving NER problems. The framework is generally applicable to any classification or sequential tagging problems if the methods used to cluster samples (Sect. 3.2) and the error decay function (Eq. 4) are modified properly.

3.1 Error partition

Given a feature, we can derive a partition p by clustering the samples into J^p groups. For example, using the sentence embedding as our features, we can use K-means (MacQueen 1967) to cluster every sentence in the corpus into multiple groups containing sentences with similar embeddings. Then, the testing error $E(C, D_U)$ can be partitioned using the sentence groups as:

$$E(C, D_U) = \sum_{s_i \in D_U} \sum_{j=1}^{J^p} P(g_j^p | s_i) \sum_{l=1}^{|s_i|} \mathbb{1}(y_{i,l} \neq \hat{y}_{i,l}^C), \quad (1)$$

where C is the current predictor, s_i is i th sentence in the testing data D_U , $P(g_j^p | s_i)$ is the probability that the sentence s_i belongs to the j th group g_j^p , and $P(g_j^p | s_i)$ could be the indicator function $\mathbb{1}(s_i \in g_j^p)$ if a hard clustering method is used. $\sum_{l=1}^{|s_i|} \mathbb{1}(y_{i,l} \neq \hat{y}_{i,l}^C)$ is the error of the sentence s_i using predictor C , $|s_i|$ is the length of the sentence, $y_{i,l}$ is the ground-truth tag for the l th token in the sentence s_i , and $\hat{y}_{i,l}^C$ is the tag predicted by the predictor C .

By assuming that the error of sentence s_i could be approximated by the estimated average error of its groups $\hat{E}(g_j^p, C)$ (i.e., $\sum_{l=1}^{|s_i|} \mathbb{1}(y_{i,l} \neq \hat{y}_{i,l}^C) \approx \hat{E}(g_j^p, C)|s_i|$), we estimate the overall error as:

$$\hat{E}^p(C, D_U) = \sum_{j=1}^{J^p} \hat{E}(g_j^p, C)m(g_j^p, D_U), \quad (2)$$

where $m(g_j^p, D_U) = \sum_{s_i \in D_U} P(g_j^p | s_i) |s_i|$ could be viewed as the number of times group g_j^p appears in D_U .

In addition to sentence clusters, we can also rely on word features and word clusters to form a partition p . Consequently, the error becomes

$$E(C, D_U) = \sum_{s_i \in D_U} \sum_{l=1}^{|s_i|} \sum_{j=1}^{J^p} P(g_j^p | s_{i,l}) \mathbb{1}(y_{i,l} \neq \hat{y}_{i,l}^C), \quad (3)$$

where $s_{i,l}$ is l th token in the sentence s_i . Hence, $m(g_j^p, D_U) = \sum_{s_i \in D_U} \sum_{l=1}^{|s_i|} P(g_j^p | s_{i,l})$ in Eq. (2).

3.2 Clustering for NER

When we use different partitions p , we get different error estimates $\hat{E}^p(C, D_U)$. To increase the robustness of our error estimation, we adopt multiple partitions based on different features and aggregate the testing error estimates for selecting the next batch of samples to be annotated.

In our experiments on real-world NER datasets, we build four partitions using different features of sentences and words as follows:

- *Sentence* We compute sentence embeddings by averaging the word embeddings, and cluster all the sentence embeddings into 10 groups. Next, the cosine similarities between sentence embeddings and the cluster centers are passed through a softmax layer with temperature parameter 0.1 to compute $P(g_j^p | s_i)$.
- *Word* We perform a simple top-down hierarchical clustering on word embeddings, which first clusters the words into 10 groups and further partitions each group into 10 clusters. This step results in 100 clusters for words in total.
- *Word + Shape* Instead of performing clustering on the lowest layer of the hierarchy, we partition the words in each group using four different word shapes: uppercase letters, lowercase letters, first uppercase letter and rest lowercase letters, and all the shapes other than above. The same word shape features are also used in our tagger.
- *Word + Sentence* Similarly, we partition each of the 10 word groups in the lowest layer of the hierarchy. For each word, we find the sentence s_i the word belongs to, and rely on the sentence group g_j^p with highest $P(g_j^p | s_i)$ to perform the partition.

Performing clustering on the concatenation of multiple feature spaces is less interpretable, so we choose to model the feature interdependency by hierarchical clustering (i.e., concatenating the clustering results). For example, in the third partition (i.e., Word + Shape), a cluster contains all words that have the same shape feature and belong to one of the 10 word embedding clusters.

Among the four partitions, the first one (i.e., Sentence) uses soft clustering because a sentence might contain multiple aspects that belong to different groups. We perform hard clustering on word features because it achieves similar performance when compared to soft clustering, and speeds up updating the cluster size when a new sample is added. For efficiency, all the clustering is done by mini-batch K-means (Sculley 2010) in D_A , the union of training data, sampling pool, and validation data.

To simplify the method and to have a better control on experimental settings, we use the same method to form groups based on the tagger's input features for all datasets, and use the same groups when selecting all batches in a dataset. Nevertheless, we note that the framework allows us to model error decay on more fine-grained clusters as more training data are collected, or use other external features (e.g., the journal where the sentence is published) that might not be easily incorporated into the tagger or uncertainty sampling.

3.3 Error decay modeling

Within each group, we assume that the error depends only on the number of samples in the group being observed in the training data. This is to avoid a complicated and uninterpretable error decay model built by many pairs of training data subset and validation error. We model the error of predictor C_{T_i} on j th group $\hat{E}(g_j^p, C_{T_i})$ in Eq. (2) using a one-dimensional function $e(n)$, where $n = m(g_j^p, T_i)$ is the size of group g_j^p in the training data T_i after i th batch is collected, and further constrain the class of decay functions $e(n)$ using prior knowledge of the tasks.

The decay function of prediction error $e(n)$ depends on the task (Si et al. 1992; Hestness et al. 2017). The error decay rate of many tasks has been shown to be $1/n^k$, both theoretically and empirically (Hestness et al. 2017), and k is typically between 0.5 and 2 (Si et al. 1992).

In sequence tagging tasks, the error decay rate depends on the importance of context. To intuitively explain how the importance of context affects the error decay rate, we discuss the form of error decay functions in one case where context does not affect the label and in another case where context matters.

Case 1 (context does not matter) Assuming we are classifying each token in a sentence into two classes and its label does not depend on context (like predicting the outcome of a coin toss), we only make reducible errors when we observe the less-likely label more times than the other label. Applying Chernoff bounds (Mitzenmacher and Upfal 2017), we can show that the error decay rate is as fast as an exponential function.

Without loss of generality, we assume the probability q of observing positive class (i.e., head) in the i th token is smaller than 0.5. Let n be the number of coin tosses be n and the random variable $X_j^i = \mathbb{1}(\text{jth toss on } i\text{th coin is head})$. In order to classify the testing tokens optimally (i.e., predict tail whenever seeing the i th token), we would like to observe $X^i = \sum_{j=1}^n X_j^i < \frac{n}{2}$. Therefore, the error rate is $P(X^i \geq \frac{n}{2})(1 - q) + q(1 - P(X^i \geq \frac{n}{2})) = P(X^i \geq \frac{n}{2})(1 - 2q) + q$.

Since all X_j^i are assumed to be independent, we can use Chernoff bounds to model the decay of $P(X^i \geq \frac{n}{2}) = P(X^i \geq (1 + \delta)\mu)$ as n increases, where $\mu = q \cdot n$, and $\delta = \frac{0.5 - q}{q}$. Chernoff bounds tell us that $P(X^i \geq \frac{n}{2}) \leq \exp(-n \cdot h(q))$, where $h(q)$ is an error decay speed function that depends on q . Different versions of Chernoff bounds lead to different $h(\cdot)$, but all $h(\cdot)$ increase as q decreases. That is, when coins are more biased, error rate decays faster.

Case 2 (context matters) If we assume that the influence of each word in the context to the label is independent, we need to estimate the probability of having the label given a word in the context to predict the label accurately. For example, we want to know how likely the label is a person name when we observe “Dr” in the context, so that we can estimate how likely the Pepper in “Dr Pepper” should be labeled as a person. The error of the probability estimation decays with rate $1/\sqrt{n}$ in the long run according to Chernoff bounds or the central limit theorem, so the error decay function is likely to be as slow as $1/\sqrt{n}$ when the words in context affect the label.

The error decay rate of most of the NER tasks should lie between the decay rates in the above two cases because the taggers will gradually learn to utilize longer contexts. Thus, we model the error decay $\hat{E}(g_j^p, C_{T_i}) = e(m(g_j^p, T_i))$ for NER by a fractional polynomial:

$$e(n) = c_j + b_j \left(\frac{a_{0.5}}{(a_0 \cdot n)^{0.5}} + \sum_{k=1}^3 \frac{a_k}{(a_0 \cdot n)^k} \right), \tag{4}$$

where $a_{0.5}$, a_{0-3} , b_j and c_j are parameters to be optimized, and we constrain these parameters to be non-negative. c_j is an estimate of irreducible error in this model, b_j tries to predict the initial error (when the first batch is collected), $a_{0.5}$ and a_{1-3} weight the curves with different decay rates, and a_0 scales the number of training samples. If $e(n)$ is proportional to $1/\sqrt{n}$, the estimated a_{1-3} would be close to 0. If $e(n)$ is proportional to an exponential function, a_{1-3} would become the coefficients in its Taylor expansion.

The parameters $a_{0.5}$, a_{0-3} , b_j , and c_j are estimated by solving

$$\arg \min_{\substack{\{a_{0-3}, a_{0.5}, \\ b_j, c_j\} \in \mathbf{R}_{\geq 0}^M}} \sum_{j=1}^{J^p} w_j \sum_{i=1}^{t_m} v_{ij} \left(\hat{E}(g_j^p, C_{T_i}) - \frac{E_j^p(C_{T_i}, D_V)}{m(g_j^p, D_V)} \right)^2, \tag{5}$$

where t_m is the number of annotated batches and the estimated error $\hat{E}(g_j^p, C_{T_i}) = e(m(g_j^p, T_i))$. The average error of j th group in validation dataset D_V , $E_j^p(C_{T_i}, D_V) = \sum_{s_i \in D_V} P(g_j^p | s_i) \sum_l \mathbb{1}(y_{i,l} \neq \hat{y}_{i,l}^{C_{T_i}})$ for a partition using sentence clusters and $E_j^p(C_{T_i}, D_V) = \sum_{s_i \in D_V} \sum_l P(g_j^p | s_{i,l}) \mathbb{1}(y_{i,l} \neq \hat{y}_{i,l}^{C_{T_i}})$ for a partition using word clusters. w_j and v_{ij} are constant weights,³ and $M = 2J^p + 5$ is the number of parameters. Due to the small number of parameters, error decay curves could be modeled by retraining deep neural networks only a few times (we set $t_m = 5$ when selecting the first batch in our experiments).

3.4 Query batch selection

Modeling the error decay on each cluster based on different features could be used as an analysis tool to increase the transparency of existing active learning methods. Such an analysis reveals the weaknesses (i.e., the groups of samples with high validation error) of the current tagger and allows us to estimate the number of samples that needs to be collected to reach a desirable error rate.

³ See “Appendix 3” for details on setting the weights.

We propose a novel active learning method to actively address the fixable weaknesses of the tagger discovered by the analysis tool. When a single partition p is used, we select the next batch B by maximizing

$$H^p(B \cup T) = - \sum_j e\left(m(g_j^p, B \cup T)\right)m(g_j^p, D_A), \tag{6}$$

where T is the collected training data, and D_A is the union of the pools of candidate samples, training data T , and validation data D_V , which are used to approximate group occurrence statistics in the testing data D_U . Note that we use $e(m(g_j^p, B \cup T))$ to approximate $\hat{E}(g_j^p, C_{B \cup T})$, so we prevent retraining the predictor C within each batch selection.

Proposition 1 *Suppose that $\hat{E}(g_j^p, C_{T_i})$ is a twice differentiable, non-increasing and convex function with respect to $m(g_j^p, T_i)$ for all j , then $H^p(T)$ is non-decreasing and submodular.*

The convexity of $\hat{E}(g_j^p, C_{T_i})$ is a reasonable assumption because the error usually decays at a slower rate as more samples are collected. Since selecting more samples only decreases the value of adding other samples, $H^p(T)$ is submodular (see ‘‘Appendix 2’’ for a rigorous proof).

Finding the optimal B in Eq. (6) is NP-complete because the set cover problem can be reduced to this optimization problem (Guillory and Bilmes 2010), but the submodularity implies that a greedy algorithm could achieve $1 - 1/e$ approximation, which is the best possible approximation for a polynomial time algorithm (up to a constant factor) (Lund and Yannakakis 1994; Guillory and Bilmes 2010).

Algorithm 1: Error decay on groups (EDG) selection algorithm

Input : Sampling pool without labels, validation data with labels D_V , a predictor C , number of clusters J^p for every partition p , burn-in batch number t_b , total batch number t_{max}
Output: Labels for selected batches
foreach partition p **do**
 | Cluster samples into groups g_j^p , where $1 \leq j \leq J^p$
end
for $t_m \leftarrow 1$ **to** t_{max} **do**
 | **if** $t_m \leq t_b$ **then**
 | Randomly sample batch B
 | **else**
 | Model error decay by solving (5)
 | Select batch B by solving (7)
 | **end**
 | Add the batch B to training data T_{t_m} , remove the batch B from sampling pool
 | Train the predictor $C_{T_{t_m}}$
 | Test the predictor on validation data to compute $E_j^p(C_{T_{t_m}}, D_V), \forall j, p$
end

When having multiple partitions p based on different features, we select the next sentence in the batch according to:

$$\arg \max_{s_i} \left(\prod_p \left(\frac{H^p(S_i \cup T) - H^p(T)}{|s_i|} + \epsilon \right) \right)^{\frac{1}{F}}, \tag{7}$$

where ϵ is a small smoothness term. If the partition p is a set of sentence clusters, then $S_i = \{s_i\}$. If it is a set of word clusters, then $S_i = \{s_{i,l}\}_{l=1}^{|s_i|}$, where $s_{i,l}$ is l th word in i th sentence. We normalize the error reduction $H^p(S_i \cup T) - H^p(T)$ by the sentence length $|s_i|$

to avoid the bias of selecting longer sentences as done in previous work (Settles and Craven 2008; Shen et al. 2018). After annotators label the whole batch, we retrain the tagger model and update the error decay prediction by solving (5) before selecting the next batch.

The selection process is summarized in Algorithm 1. In the first few batches, we perform random sampling to collect pairs of every cluster size and the prediction error in the cluster for solving the one-dimensional regression problems. After the number of collected batches t_m is larger than the number of burn-in epochs t_b , we have sufficient size and error pairs required to model the error decay in (5). Then, we can predict the future error $H^p(S_i \cup T)$ and select the samples that minimize the error using (7).

Note that (7) naturally balances informativeness and representativeness. From the informativeness perspective, the sample without error decay won't be selected. From the representativeness and diversification perspectives, we will decrease the value of choosing a sample in a batch after the samples in the same clusters are selected.

4 Method extensions

For some applications, a validation set is not large enough to be used to model the error decay curves, and our independent assumption may be too strong. To address this concern, we also test two extensions of our method.

4.1 Prediction difference decay

We replace the ground truth labels in (5) with the prediction $\hat{y}_{i,l}^{C_{T_m}}$ based on the current training data. That is, our sampling method computes the difference between the current prediction and the previous predictions $C_{T_1}, \dots, C_{T_{m-1}}$ in each group, and models the decay of the difference to maximize the convergence rate of predictions. We denote this method as EDG_ext1 in our experiments.

4.2 Uncertainty or disagreement decay

When the uncertainty or disagreement information is available, we can model their decay and choose the sentences with highest uncertainty decay rather than highest uncertainty. To avoid making the independence assumption, we skip the clustering step and assume that future uncertainty decay is proportional to the previous uncertainty decay, and set the score of i th sentence to be

$$\min(\max(u_i^{t_f} - u_i^{t_m}, 0), u_i^{t_m}), \quad (8)$$

where $u_i^{t_m}$ is the current uncertainty of i th sentence, and $u_i^{t_f}$ is its previous uncertainty. Note that we take the minimum between the difference and $u_i^{t_m}$ to ensure that the predicted future uncertainty is always non-negative. This method is denoted as EDG_ext2 in our experiments.

5 Experimental setup

NER problems are often used as benchmark to evaluate (deep) active learning methods (Settles and Craven 2008; Shen et al. 2018; Siddhant and Lipton 2018) because they are the foundation for many information extraction tasks, and acquiring tags for each token requires a large amount of human effort. We follow Strubell et al. (2017) and use phrase-level micro-F1 as the performance metric for NER tasks. Precision and recall are computed by counting the number of correct boundary and type predictions. Unless otherwise stated, we use a four-layer convolutional neural network (CNN) as our tagger (Strubell et al. 2017).⁴

5.1 Simulation on gold labels

This is one of the most widely used setups to evaluate active learning methods. We compare the performance of NER tagger trained on different data subsets chosen by different methods. In the “Appendix 1”, we also compare the performance of applying different active learning methods to BiLSTM-CRF models.

5.1.1 Synthetic dataset

We synthesize a dataset with 100 words; each word could be tagged as one of four entity types or none (not an entity). There are three categories of words. The first category consists of half of the words which are always tagged as none. This setup reflects the fact that a substantial amount of words such as verbs are almost always tagged as none in NER tasks.

One-fourth of the words belong to the second category where every word mention has equal probability of being tagged as one of the entity types or none. In real-word NER tasks, the noisy label assignment may be due to inherently ambiguous or difficult words.

The remaining 25 words are in the third category where the labels are predictable and depend on the other context words. The likelihood of words in the third category being tagged as one of the four entity types is sampled from a Dirichlet distribution with $\alpha_{1-4} = 1$, while the likelihood of being none is zero. Whenever one of these words w appear in the sentence, we check two of its preceding and succeeding words that are also in the third category, average their likelihoods of entity types, and assign the type with the highest likelihood to the word w .

When generating a sentence, the first word is picked randomly. The transition probability within each category is 0.9. Inside the first and second categories, the transition probability is uniformly distributed, while the probability of transition to each w inside the third category is proportional to a predetermined random number between 0.1 and 1. The sentence length is between 5 and 50 and there is a 0.1 probability of ending the sentence after generating a word within the range.

In this dataset, each word is a group in our method and no clustering is performed. The i th word has a word embedding vector $[\mathbb{1}(k = i)]_{k=1}^{100}$. When modeling error decay, we start from 1000 tokens and use a batch size of 500. When evaluating the sampling methods, we start with 3000 using a batch size of 1000. That is, after the first batch is selected, we

⁴ We choose to test the methods on CNN because Strubell et al. (2017) showed that CNN can achieve performance which is close to state of the art while being much more efficient than BiLSTM-CRF.

Table 2 The size of datasets for the simulation on gold labels

	Synthetic data		CoNLL 2003	
	Token	Sentence	Token	Sentence
Train	99,956	6726	204,567	14,987
Val	10,045	679	51,578	3466
Test	10,004	677	46,666	3684
	NCBI disease		MedMentions ST19	
	Token	Sentence	Token	Sentence
Train	135,900	5725	758,449	28,227
Val	23,836	941	254,539	9303
Test	24,255	970	253,737	9383

update the error decay curves based on the prediction of taggers trained on 1000, 1500, 2000, 2500, 3000 and 4000 tokens.

5.1.2 Real-world datasets

We test the active sampling methods on CoNLL 2003 English NER (Tjong Kim Sang and De Meulder 2003), NCBI disease (Doğan et al. 2014), and MedMentions (Murty et al. 2018) datasets. The size of these datasets are presented in Table 2. CoNLL 2003 dataset has four entity types: people name (PER), organization name (ORG), location name (LOC), and other entities (MISC). NCBI disease dataset has only one type (disease name). For MedMentions, we only consider semantic types that are at level 3 or 4 (higher means more specific) in UMLS (Bodenreider 2004). Any concept mapping to more abstract semantic types is removed as was done by Greenberg et al. (2018) and this subset is called MedMentions ST19. The 19 concept types in MedMentions ST19 are *virus, bacterium, anatomical structure, body substance, injury or poisoning, biologic function, health care activity, research activity, medical device, spatial concept, biomedical occupation or discipline, organization, professional or occupational group, population group, chemical, food, intellectual product, clinical attribute, and Eukaryote*.

In all the three datasets, the first 30,000 tokens are from randomly sampled sentences. To model error decay, we start from 10,000 tokens and retrain the tagger whenever 5000 new tokens are added. When evaluating the sampling methods, we start from 30,000 tokens, using a batch size of 10,000.⁵

⁵ Selecting only part of a sentence is not reasonable in NER, so each chosen batch may be slightly larger than the desired batch size. The difference is smaller than the length of the last chosen sentence. Since the desired batch size is set to be large (10,000 words in real-world datasets), the difference in batch size between sampling methods is negligible.



Fig. 3 Simulation on pseudo labels compares active learning methods on a large pool with noisy labels. In addition to the original validation and testing set, we also use a testing set with pseudo labels for evaluation

5.2 Simulation on pseudo labels

In practice, we often observe systematic noise from annotators. The noise could come from some inherently difficult or ambiguous cases in the task or from incapable workers in the crowdsourcing platforms. Thus, we propose a novel evaluation method to test the robustness of different sampling methods in the presence of such noise.

As shown in Fig. 3, we first train a high-quality tagger using all the training data and use it to tag a large sampling pool. Then, different active learning methods are used to optimize the tagger trained on these pseudo labels. The micro-F1 is measured by comparing the tagger's prediction with pseudo labels or gold labels on unseen sentences. The evaluation method also allows us to perform sampling on a much larger sampling pool,⁶ which is usually used in the actual deployment of active learning methods. We randomly select 100,000 abstracts from PubMed as our new sampling pool, which is around 180 times larger than the pool we used in the simulation on gold labels. Precisely, the sampling pool consists of 24,572,575 words and 921,327 sentences. The testing data with pseudo labels have 2,447,607 tokens and 91,591 sentences.

We also evaluate the sampling methods on two practical variations of the above setting. We use the data collected for optimizing a CNN to train BiLSTM-CRF models, which can be used to test the robustness of active learning methods after switching tagger models (Lowell et al. 2019). In addition, when collecting gold labels for biomedical NER, annotators often tag the whole abstract at a time, which can only be tested using a large sampling pool.

For all sampling methods, the sampling score of an abstract is the average of the sampling scores of the sentences in the abstract weighted by the sentence length. That is, the selection criteria view an abstract as a bag of sentences similar to how a sentence is considered as a bag of words when clustering is performed on words. We greedily select the abstract with the highest sampling score.

5.3 Sampling strategies

We compare the following sampling methods:

- **Random (RND):** We select sentences randomly with uniform probability.

⁶ Running sampling methods on a large pool is time-consuming, so we only compare the methods after the first batch is collected, i.e., when the size of the training dataset reaches 40,000.

- **Error Decay on Groups (EDG)**: This is our method where we optimize (7) using validation data.
- **EDG_ext1 (w/o Val)**: As described in Sect. 4.1, we replace the validation error in EDG with the prediction difference.
- **Maximum Normalized Log-Probability (US)**: We use the least confidence sampling (Culotta and McCallum 2005). This variant of uncertainty sampling has been shown to be very effective in NER tasks (Shen et al. 2018). When applying maximum normalized log-probability to the CNN model, we select the sentences via $\arg \min_{s_i} (1/|s_i|) \sum_{l=1}^{|s_i|} \max_{y_{il}} \log P(y_{il})$.
- **Maximum Normalized Log-Probability with Diversification (US + Div)**: We diversify uncertain samples based on sentence embeddings (i.e., the average embedding of its words) (Wei et al. 2015; Shen et al. 2018). We implement the US + Div, also called filtered active submodular selection (FASS), described in Shen et al. (2018). We use cosine similarity to measure the similarity between sentence embeddings. The number of candidate sentences is the batch size times $t = 100$.
- **Diversification (Div)**: We use the same algorithm as US + Div, except that all samples are equally uncertain.
- **US + Div + EDG_ext2**: This is the same algorithm as US + Div, but with uncertainty scores replaced with their difference in Eq. (8).
- **Bayesian Active Learning by Disagreement (BALD)**: We select samples based on the disagreement among forward passes with different dropouts (Gal et al. 2017). The prediction disagreement of l th token in i th sentence is computed by $\frac{\sum_{k=1}^K \mathbb{1}(y_{il}^k \neq \text{mode}_k(y_{il}^k))}{K}$. The number of forward passes K is set to 10 in our experiments. The sentence disagreement is the average of tokens' disagreement. We use the default hyperparameter values for the dropouts as in Strubell et al. (2017).
- **BALD + EDG_ext2**: Here, disagreement scores are replaced with their difference in Eq. (8).

5.4 NER tagger details

We use the published hyperparameters⁷ for real-world datasets, and simplify the tagger for the synthetic dataset to decrease the standard deviation of micro-F1 scores. In the synthetic dataset, we reduce the number of layers of CNN to two because the label depends only on the left two words and right two words. Furthermore, we change the learning rate from 5×10^{-4} to 10^{-4} , batch size from 128 to 32, and max epochs from 250 to 1000 to make the performance more stable. When training BiLSTM-CRF, we also use the implementation and its default hyperparameters from Strubell et al. (2017). In all the experiments, the number of epochs is chosen using validation data.

The word embeddings for CoNLL 2003 are vectors with 50 dimensions from SENNA (Collobert et al. 2011). The word embeddings for NCBI disease and MedMentions ST19 are word2vec (Mikolov et al. 2013) with 50 dimensions trained on randomly sampled 10% of all PubMed text. Before clustering, we normalize all the word embedding vectors such that the square of the ℓ_2 distance between two words is twice their cosine distance.

⁷ <https://github.com/iesl/dilated-cnn-ner>.

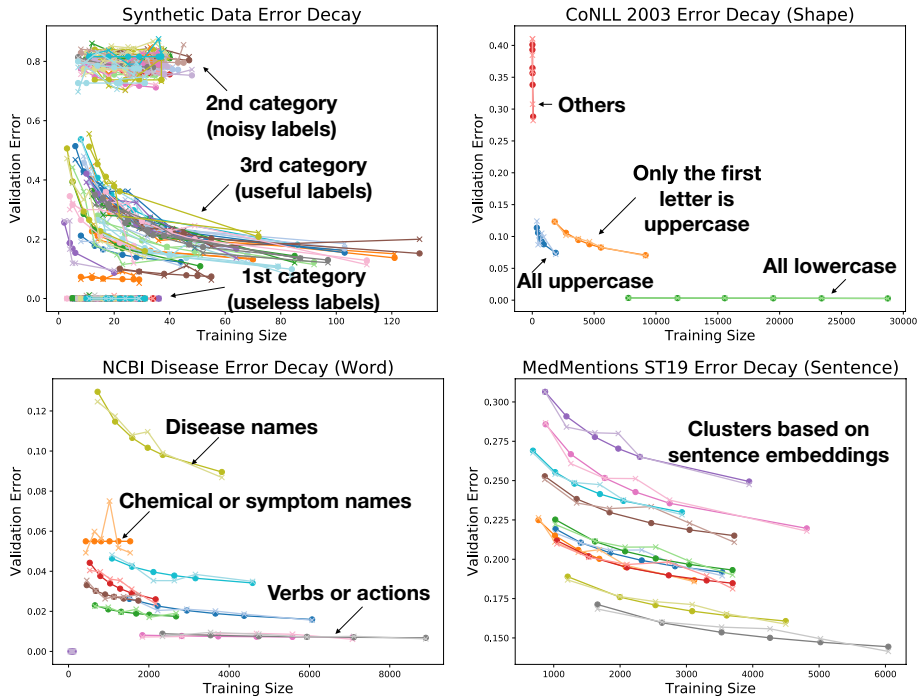


Fig. 4 Error decay on groups that are modeled after the first batch is collected by EDG. The x markers on the curves are the real error and • means prediction from the fitting curve. The groups shown in the figure for NCBI disease and MedMentions ST19 are formed by clustering word and sentence embeddings, respectively

5.5 Visualization of the error decay model

In addition to qualitatively evaluating our methods, we also visualize our error decay models. The visualization examines whether our error decay function in Eq. (4) can accurately model the empirical error decay in NER datasets, and whether our sample selection strategies are transparent and interpretable.

Given a dataset and a partition p , the empirical and predicted errors in each group g_j^p are plotted as two curves in Fig. 4. We compare the y value of empirical error $E_j^p(C_{T_t}, D_V)/m(g_j^p, D_V)$ and predicted error $\hat{E}(g_j^p, C_{T_t}) = e(m(g_j^p, T_t))$ given different x values $m(g_j^p, T_t)$, the cluster size of group g_j^p in the training set. Each curve connects six points corresponding to $t = 1 \dots 6$, and the total number of words in training set $|T_t|$ are 10,000, 15,000, 20,000, 25,000, 30,000, and 40,000 in real world datasets, respectively. The first 30,000 words are selected randomly and the last 10,000 words are selected by EDG.

In different datasets, we visualize the different partitions derived from different features. In synthetic data, a group is a word. In CoNLL 2003, we plot each group that contains all the words with the same shape. In NCBI disease, we plot each group that contains words with similar word embeddings. In MedMentions ST19, we plot each group that contains sentences with similar sentence embeddings. Note that in the quantitative experiments, we actually use word + shape in CoNLL 2003 and 100 clusters in NCBI disease, but we illustrate only shape in CoNLL 2003 and 10 clusters in NCBI disease to simplify the figure.

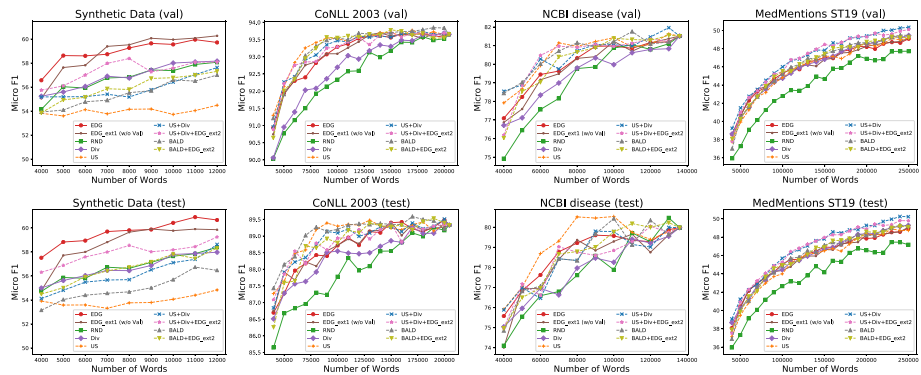


Fig. 5 Comparison of different sampling methods on the four NER tasks. The validation (first row) and testing scores (second row) are averaged from the micro-F1 (%) of three CNNs trained with different random initializations. The performance of methods which cannot be applied to black-box taggers is plotted using dotted curves

6 Results and analysis

The results of error decay visualization, the simulation on gold labels, and pseudo labels are shown in Figs. 4, 5 and Table 3, respectively.

As shown in Table 3, the scores on the gold validation set, the gold testing set, and the testing set using pseudo labels follow a similar trend and most active learning methods do better than random regardless of which test set is used. The observation indicates that the taggers do not overfit the label noise in the training data severely and justifies our pseudo label experiments.

We first qualitatively analyze the error decay modeling in Sect. 6.1. Next, we quantitatively compare different methods in Sects. 6.2, 6.3, 6.4, and 6.5.

6.1 Error decay visualization

In Fig. 4, the predicted error decay curves usually fit the empirical values well. Some deviations of empirical error come from the randomness of training CNNs. For example, the fourth point in MedMentions ST19 has higher empirical error in almost all clusters because the parameters of CNN trained on the set with 25,000 words happen to converge to a worse state.

In the figures, we can see that the length between the fifth and sixth points in each curve varies because the last 10,000 words in the training set are actively selected. The clusters that might have a larger error decay (e.g., the orange curve in CoNLL 2003) would get more training instances in the last sample batch. The figures demonstrate that the one-dimensional regression problem for each cluster could be solved well even though the sampling process is not random and only six training pairs for each curve are observed.

We can interpret the sampling strategy of EDG from the different number of samples selected in different groups. For example, EDG improves the sampling efficiency when compared to random sampling by selecting more words whose first letter is uppercase in CoNLL 2003, and by selecting more disease name candidates than verbs or actions in NCBI disease.

Table 3 Simulation on pseudo labels for NCBI disease dataset

	Whole abstract						Sentence						Avg
	CNN for CNN			CNN for BiLSTM-CRF			CNN for CNN			CNN for BiLSTM-CRF			
	Val	Test	Pseudo	Val	Test	Pseudo	Val	Test	Pseudo	Val	Test	Pseudo	
EDG ^a	61.0	59.2	54.9	66.5	66.8	58.7	60.1	58.2	54.7	68.2	66.5	59.1	61.2
EDG_ext1 (w/o Val)	59.6	58.5	54.9	65.7	65.5	57.7	60.2	58.4	52.5	67.7	67.8	58.5	60.6
RND	56.0	56.3	50.4	62.6	62.6	55.0	57.7	56.1	52.9	63.9	64.1	56.2	57.8
Div	55.2	55.6	52.8	63.2	63.9	56.1	57.7	56.2	52.5	63.9	64.2	56.6	58.2
US	59.7	59.0	54.1	67.2	67.4	57.8	58.3	57.7	53.1	68.5	68.0	58.1	60.7
US + Div	60.1	59.5	56.0	64.1	65.7	57.9	60.1	56.8	54.5	68.7	67.4	58.5	60.8
US + Div + EDG_ext2	61.2	58.6	56.0	69.5	65.4	58.7	62.2	59.4	55.9	69.8	66.6	58.5	61.8
BALD	59.8	59.4	55.5	67.4	67.7	58.3	59.9	58.8	54.9	69.7	67.3	59.5	61.5
BALD + EDG_ext2	60.7	58.2	56.0	68.4	65.8	58.9	63.7	60.0	54.9	71.7	67.6	59.2	62.1

After selecting the first batch using different sampling methods, the micro-F1 (%) is computed by averaging across five neural networks trained using different random initializations. Whole abstract and sentence mean we sample one abstract and sentence at a time, respectively. CNN for BiLSTM-CRF means that we report the F1 of BiLSTM-CRF that is trained on the data selected for optimizing CNN. The highest score for black-box and uncertainty-based models are highlighted, and the last column shows the unweighted average of all values in each row. The average F1 difference of EDG_ext1 versus Div, US + Div + EDG_ext2 versus US + Div, BALD + EDG_ext2 versus BALD are significant ($p < 0.01$) according to two-sample t-test

^aIndicates the method uses ground-truth labels in the validation set to collect samples in the training set

The transparency of EDG explains some important empirical observations in previous work. For example, Lowell et al. (2019) observed that the benefit of active learning is usually more significant in NER than in sentence classification, and the improvement in NER is robust against the change of predictor model. Shen et al. (2018) observed that in some NER datasets (e.g., CoNLL 2003), we can train a neural tagger that reaches a similar performance using only a selected small portion of the training set compared to using all the data. Figure 4 indicates that one of the main reasons is that there are more lowercase words in CoNLL 2003 than uppercase words, and lowercase words are almost never tagged as names of people, organizations, or locations. Therefore, the active learning methods could easily achieve high sampling efficiency by selecting more uppercase words, and the selection tendency can benefit various kinds of predictor models.

In real-world datasets, the error decay rate usually follows the function $1/\sqrt{n}$ when n is large for most of the groups regardless of the feature being used. For example, $a_{0.5}$ in Fig. 4 is at least two times larger than the $a_1 + a_2 + a_3$ in CoNLL 2003, NCBI disease, and MedMentions ST19. The small difference between empirical and predicted error also justifies our assumption that the weight parameters of terms are shared across all the groups (i.e., a_{0-3} and $a_{0.5}$ do not depend on the cluster index j).

6.2 EDG versus Div and RND

Among the methods we compared against, only random (RND) and diversification (Div) can be applied to black-box taggers. Our method (EDG) significantly outperforms Div, and Div outperforms RND in synthetic, CoNLL 2003, and NCBI disease datasets, which demonstrates the effectiveness of EDG. This also justifies our assumptions and indicates that the error decay curves are modeled well enough for the purpose of active sampling.

6.3 US versus US + Div

Shen et al. (2018) found that diversification is surprisingly not helpful in batch active learning. However, our results suggest that this finding might be valid only when the sampling pool size is small and/or some groups of frequent words/sentences are clearly not helpful. When the pool is sufficiently large and the task is to jointly extract many different types of entities like MedMentions ST19, sampling almost all kinds of sentences can be helpful to the task as all sentence clusters have similar decay on the far right of Fig. 4. Then, the diversification approach (Div) can be as effective as US and EDG, while US + Div(+ EDG_ext2) provides the best result.

6.4 EDG versus uncertainty-based methods

As shown in Wang et al. (2017a), it is difficult to perform better in terms of sampling efficiency when comparing a black-box active learning method with uncertainty sampling (US). In real-world datasets, EDG achieves part of the performance gain from US that is easily explainable (e.g., coming from ignoring those easy words), controllable by humans,⁸ and

⁸ In the supplementary material A.5, we show that EDG can be easily modified for situations where each label class has a different penalty.

does not involve the specifics of the tagger to model the interaction between each word and its context.

Furthermore, US is not robust to labeling noise or ambiguous samples (Mussmann and Liang 2018), which have high errors but low error decay. For instance, US almost always selects difficult words with high irreducible errors in the synthetic data. In real-world dataset, we could also observe ambiguous or difficult words. For example, *insulin* is a chemical, but *insulin resistance* could be a disease or a symptom in NCBI disease dataset. In Fig. 4, we can see that EDG does not select many words in the group. Our error curves in the supplementary material A.4 show that US selects many such words with incorrect pseudo labels. The vulnerability makes EDG outperform US in synthetic data and Table 3 on average.

US + Div and BALD are more robust to labeling noise than US, but still suffers from a similar problem. Thus, the sampling strategies that choose more samples with reducible uncertainty (i.e., US + Div + EDG_ext2 and BALD + EDG_ext2) could significantly improve the accuracy of taggers in noisy datasets like our synthetic data and NCBI disease dataset with pseudo labels, while having comparable performance on the other clean datasets with gold labels.

6.5 EDG_ext1 (w/o Val) versus EDG

In all datasets, modeling the error decay using pseudo labels (EDG_ext1) achieves similar performance when compared to using gold validation data (EDG), and also outperforms Div. In addition, the micro-F1 scores of EDG on validation and testing data roughly show a similar trend, which suggests that our method does not overfit the validation data even though it has access to its gold labels during sampling.

7 Conclusions

We proposed a general active learning framework which is based only on the predictions from black-box predictors, is robust to labeling noise without relying on prior knowledge, and forecasts the potential error reduction in different aspects based on a small number of existing labels.

Our experimental results suggest that no single batch active learning method wins in all the cases and every method has its own weaknesses. We recommend practitioners to analyze the error decay on groups in order to choose a proper sampling algorithm. If the sampling pool is small and the error decay analysis shows that many samples could be easily tagged, then uncertainty sampling methods are expected to perform well. Otherwise, diversification should be considered or combined with uncertainty sampling. Finally, error decay on groups (EDG) or its extensions should be adopted if there are practical deployment challenges such as issues of *applicability* (e.g., only a black-box predictor is available), *robustness* (e.g., labels are inherently noisy), or *transparency* (e.g., an interpretable sampling process or an error reduction estimation is desired).

8 Future work

In our experiments, we demonstrated that our methods are transparent and robust to labeling noise. However, we have not yet applied them to tasks other than NER. For example, when we annotate a corpus for relation extraction, we usually want to select a document which is informative for the named entity recognizer, entity linker, and sentence classifier. This challenge is also called multi-task active learning (Reichart et al. 2008; Settles 2011). Compared to heuristically combining uncertainty from different models (Reichart et al. 2008), our methods provide more flexibility because it allows us to assign weights on the error reduction of each task and select the next batch by considering all tasks jointly.

In addition to the above pipeline system, question answering (QA) is another example where uncertainty is difficult to estimate. Many reading comprehension models such as pointer networks predict the start and end positions of the answer in a paragraph (Wang et al. 2017b). However, higher uncertainty on the position prediction does not necessarily mean the model is uncertain about the answer. It is possible that the correct answer appears in many places in the paragraph and the network points to all the right places with similar low probability. By modeling the error decay directly, our methods avoid the issue.

Finally, we have not compared EDG with active learning methods that are designed for a specific task to solve a specific practical issue. For example, the active sampling methods proposed by Wang et al. (2017a) are designed for semantic role labeling and focus on the applicability issue (i.e., black-box setting). Due to the difficulty of adapting their methods to NER and making fair comparisons, we leave such comparisons for future work.

Acknowledgements We thank Akshay Krishnamurthy for many helpful discussions. We also thank the anonymous reviewers for their constructive feedback.

Funding This work was supported in part by the Center for Data Science and the Center for Intelligent Information Retrieval, in part by the Chan Zuckerberg Initiative under the project “Scientific Knowledge Base Construction, in part using high performance computing equipment obtained under a grant from the Collaborative R&D Fund managed by the Massachusetts Technology Collaborative, in part by the National Science Foundation (NSF) grant numbers DMR-1534431 and IIS-1514053.

Compliance with ethical standards

Conflict of interest Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

Appendix A: Additional experimental results

A.1 CNN for BiLSTM-CRF using gold labels

Baldrige and Osborne (2004) and Lowell et al. (2019) demonstrate that samples collected to optimize one model might not be helpful to another model. To test the robustness of different active learning methods to such model switch, we train BiLSTM-CRF models using the batches collected based on CNN (i.e., CNN for BiLSTM-CRF in Table 3). When training BiLSTM-CRF on the synthetic dataset, we increase the max epochs from 250 to 1000.

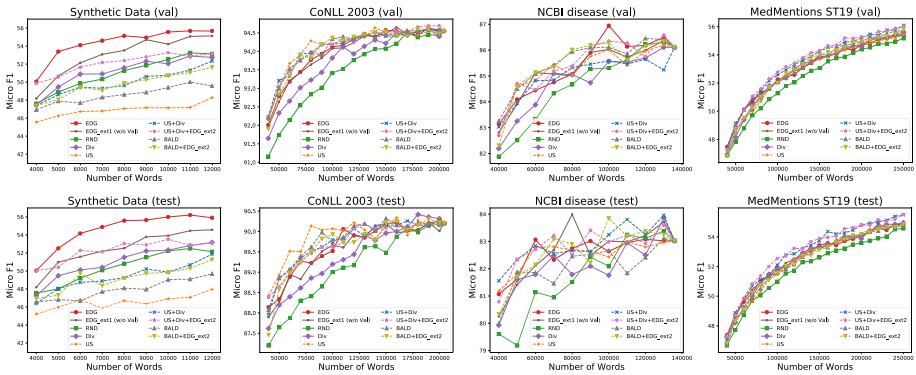


Fig. 6 Performance of BiLSTM-CRF models trained on training sets in Fig. 5 (i.e., CNN for BiLSTM-CRF setting). The performance metrics are the average micro-F1 (%) of three BiLSTM-CRF trained with different random initializations

The results are presented in Fig. 6. Nearly all the observations from Fig. 5 also hold in Fig. 6. One difference is that EDG and its extensions perform slightly better. For example, in Fig. 5 (CNN for CNN setting), BALD + EDG_ext2 sometimes performs slightly worse than BALD (e.g., in the testing data of CoNLL 2003 and NCBI disease). However, BALD + EDG_ext2 seems to always perform similar or better (e.g., in MedMentions ST19 and testing data of NCBI disease) than BALD after we switch the tagger model.

Another difference is that the performance gain between active sampling and random sampling is smaller in MedMentions ST19 than the gap in Fig. 5, while we do not observe a similar reduction in CoNLL 2003 and NCBI disease. We hypothesize that active learning methods will skip the groups of unhelpful words in CoNLL 2003 (like lowercase words) and NCBI disease, and those words are usually also unhelpful to other models. Thus, the performance gains in such datasets are more transferable and less dependant on the model choice in the first place.

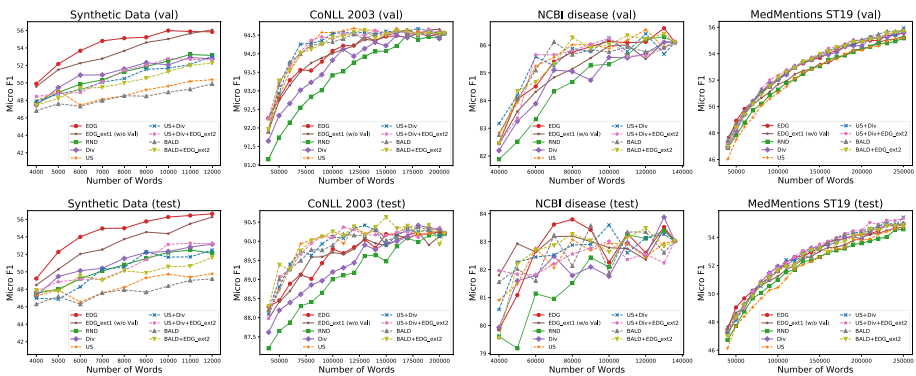


Fig. 7 Comparison of applying different sampling methods to BiLSTM-CRF models. The performance metrics are the average micro-F1 (%) of three BiLSTM-CRF trained with different random initializations

A.2 BiLSTM-CRF for BiLSTM-CRF using gold labels

Following Shen et al. (2018), we set the uncertainty of each sentence as the negative log likelihood of the predicted label sequence. The results are presented in Fig. 7. Most of the observations from Fig. 5 also hold in Fig. 7. For example, EDG and its extensions significantly improve the performances in our synthetic dataset, and are significantly better than diversification methods in CoNLL 2003 and NCBI disease datasets. There are some minor differences. For instance, the performance gap between EDG and uncertainty-based methods is larger in CoNLL 2003, which implies that much of the sampling efficiency improvement of uncertainty-based methods comes from the specific way of how BiLSTM-CRF models dependency between words. In addition, the US and EDG without validation data do not perform well in MedMentions ST19. We hypothesize that this is because the transition probabilities play an important role in this dataset, but sampling many uncertain label sequences does not help the BiLSTM model.

A.3 Sensitivity to number of clusters

We choose our clustering approach and all other hyperparameters based on the validation performance after collecting the first batch in CoNLL 2003 and MedMentions ST19, and we find the performance is not sensitive to the choices unless some crucial information is missing (e.g., not considering the word shape in CoNLL 2003). We report the performances with varying number of clusters in each layer of our hierarchical clustering in Table 4. We see that the micro-F1 scores are very close to each other, except in the testing set of NCBI disease dataset. We suspect the score variation in NCBI disease mainly comes from the randomness in the training process of neural networks because we conduct only one trial of experiments when filling Table 4. The results suggest that number of clusters is a trade-off in EDG. Increasing the number of clusters decreases the bias, but increases the variance in the error decay estimation.

Table 4 Performance sensitivity to the number of clusters (J)

	CoNLL 2003		NCBI disease		MedMentions ST19	
	Val	Test	Val	Test	Val	Test
EDG ($J = 5$)	93.0	88.6	79.7	78.5	45.4	44.9
EDG ($J = 10$)	93.1	88.7	80.1	78.4	46.0	45.7
EDG ($J = 15$)	93.1	88.6	79.8	78.3	46.0	45.7
EDG_ext1 (w/o Val) ($J = 5$)	93.0	88.5	79.7	78.5	45.3	44.7
EDG_ext1 (w/o Val) ($J = 10$)	93.1	88.6	80.0	78.1	46.0	45.7
EDG_ext1 (w/o Val) ($J = 15$)	93.0	88.6	79.7	78.7	45.9	45.7

Notice that the number of total clusters for the word and word + sentence feature is J^2 (e.g., 225 for $J = 15$). The micro-F1 scores (%) are the average over all the training set sizes in Fig. 5. The range of training set sizes are 40,000–200,000, 40,000–130,000, 40,000–250,000 for CoNLL 2003, NCBI disease, and MedMentions ST19, respectively. The highest F1 scores using different numbers of clusters for each sampling method are highlighted

A.4 Error decay of uncertainty sampling

The error curve modeling can be used not only to select the next batch but also to analyze the existing sampling strategy. For instance, Fig. 8 shows that the last points in the word group 3, 8, and 9 are farther away from the fifth points in x-axis compared to the corresponding distances in other groups. This implies that uncertainty sampling tends to select samples with high errors when it chooses the 30,000th to 40,000th tokens as shown in Fig. 1. However, the high errors do not necessarily lead to high error reduction in this dataset. This explains why US only achieves 58.3 validation micro-F1 in Table 3, which is significantly worse than other methods like EDG or US + Div.

A.5 Weighted class evaluation

To test the controllability of our methods, we set up a variation of MedMentions ST19 where we give different penalties to different classes. Specifically, we focus on four related types: *research activity*, *health care activity*, *population group*, and *spatial concept*, and assign 0.9 weights to these classes. Other 15 classes are assigned 0.1 weights. That is, when we compute micro-F1, we weight the total number of correct predictions, the total number of predictions, and the total number of entities according to the importance of classes.

To incorporate the weights into our method, we modify $E_j^p(C_{T_i}, D_V)$ in (5) as

$$E_j^p(C_{T_i}, D_V) = \sum_{s_i \in D_V} P(g_j^p | s_i) \sum_l \left(\frac{r(y_{il}) + r(\hat{y}_{il}^{C_{T_i}})}{2} \right) \mathbb{1}(y_{il} \neq \hat{y}_{il}^{C_{T_i}}), \tag{9}$$

where $r(y_{il})$ and $r(\hat{y}_{il}^{C_{T_i}})$ are the weights of the ground truth class and the predicted class, respectively. After this simple modification, our method boosts the weighted micro-F1 in the testing data from 34.3 (using same $r(y)$ for all the classes) to 35.5 (using different $r(y)$ for different classes) after collecting the first batch in MedMentions ST19 where the score of random sampling is 31.9. The scores come from averaging the results of five randomly initialized CNNs.

Fig. 8 Error decay curves of taggers trained using pseudo labels in NCBI disease dataset. The six points in each curve come from the taggers trained by 10,000, 15,000, 20,000, 25,000, 30,000, and 40,000 words. The first 30,000 words are selected randomly and uncertainty sampling (US) selects 30,000–40,000 words. As in Fig. 4, the x markers on the curves are the real error and • means prediction from the fitting curve. The groups are formed by clustering their word embeddings, and the index of each group is presented

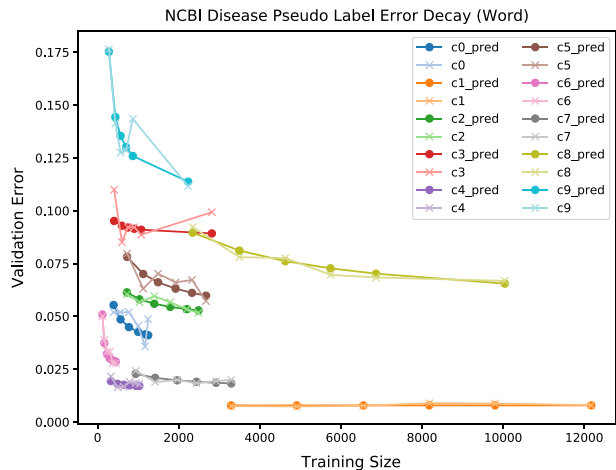


Table 5 The experimental setup is the same as Table 3 except that the performances are the maximal micro-F1 (%) of five neural networks rather than their average

	Whole abstract												Avg
	Sentence						Sentence						
	CNN for CNN			CNN for BiLSTM-CRF			CNN for CNN			CNN for BiLSTM-CRF			
	Val	Test	Pseudo	Val	Test	Pseudo	Val	Test	Pseudo	Val	Test	Pseudo	
EDG	61.6	61.9	55.7	68.7	68.2	59.2	62.1	60.2	56.1	70.0	68.2	59.8	62.4
EDG_ext1 (w/o Val)	60.4	61.5	55.4	67.4	67.0	58.6	61.9	60.0	55.2	69.0	69.2	59.1	62.0
RND	56.5	57.9	51.3	64.9	64.1	55.7	58.8	57.2	54.1	65.4	65.2	56.5	59.0
Div	57.9	57.8	54.2	64.4	64.6	57.2	59.3	58.4	53.8	66.2	65.7	57.0	59.7
US	60.9	60.9	54.9	68.1	68.5	58.3	60.0	59.6	54.0	69.7	68.5	59.9	61.9
US + Div	61.1	60.4	56.4	65.4	66.8	58.3	60.5	58.9	55.4	70.1	68.1	59.3	61.7
US + Div + EDG_ext2	61.9	60.1	56.8	71.7	66.1	59.1	63.1	60.8	56.8	70.6	68.6	60.1	63.0
BALD	61.0	61.4	56.4	68.9	69.4	58.8	61.0	59.7	55.1	70.9	67.8	60.1	62.5
BALD + EDG_ext2	61.3	60.0	56.9	69.4	66.5	59.5	65.0	62.2	55.9	72.4	68.7	59.7	63.1

A.6 Results statistics

The micro-F1 in Fig. 5 is the average of three trials. Different trials use different random initializations for the CNN. The performance variance among different trials is usually small. The average standard error in the validation set across all batches and all the methods is 0.13 for synthetic data, 0.06 for CoNLL 2003, 0.28 for NCBI disease, and 0.16 for MedMentions ST19. The average standard error in the testing set is 0.18 for synthetic data, 0.14 for CoNLL 2003, 0.49 for NCBI disease, and 0.17 for MedMentions ST19.

In Table 3, we show the average micro-F1 of five CNNs with different initializations. Sometimes, we care more about the maximal performance, so we also report the highest F1 score out of five runs in Table 5; the results show a similar trend. When applying two-sample t-test to the comparison of average performance in Table 3, we assume that every micro-F1 score is a true hidden value plus Gaussian noise, and the variance of the noise is the same given a sampling method. Based on the assumption, the one-tailed two-sample t-test gives us $p < 0.00003$ for the difference between BALD + EDG_ext2 and BALD, between US + Div + EDG_ext2 and US + Div, and between EDG_ext1 and Div.

Appendix B: Proof of Proposition 1

We would like to prove that $H^p(T) = -\sum_j \hat{E}(g_j^p, C_T)m(g_j^p, D_A)$ is submodular and non-decreasing by assuming $\frac{d\hat{E}(g_j^p, C_T)}{dm(g_j^p, T)} \leq 0$ and $\frac{d^2\hat{E}(g_j^p, C_T)}{d^2m(g_j^p, T)} \geq 0$ for every group g_j^p in partition p .

First, we prove that $-\hat{E}(g_j^p, C_T)$ is submodular and non-decreasing. Assuming we have two subsets X, Z , and one sample s_i such that $X \subseteq Z \subseteq D_A$ and $s_i \in D_A \setminus Z$. Based on the assumption, we get $m(g_j^p, X) \leq m(g_j^p, Z)$ and $m(g_j^p, X \cup \{s_i\}) \leq m(g_j^p, Z \cup \{s_i\})$ for all g_j^p . Since $\frac{d\hat{E}(g_j^p, C_T)}{dm(g_j^p, T)} \leq 0$, $-\hat{E}(g_j^p, C_T)$ is non-decreasing. In order to consider the case that $m(g_j^p, X \cup \{s_i\}) \geq m(g_j^p, Z)$, we first decompose

$$\begin{aligned} & (-\hat{E}(g_j^p, C_{X \cup \{s_i\}})) - (-\hat{E}(g_j^p, C_X)) \\ &= (-\hat{E}(g_j^p, C_{X \cup \{s_i\}})) - (-\max(\hat{E}(g_j^p, C_{X \cup \{s_i\}}, \hat{E}(g_j^p, C_Z))) \\ & \quad + (-\max(\hat{E}(g_j^p, C_{X \cup \{s_i\}}, \hat{E}(g_j^p, C_Z))) - (-\hat{E}(g_j^p, C_X)), \end{aligned} \quad (10)$$

and

$$\begin{aligned} & (-\hat{E}(g_j^p, C_{Z \cup \{s_i\}})) - (-\hat{E}(g_j^p, C_Z)) \\ &= (-\hat{E}(g_j^p, C_{Z \cup \{s_i\}})) - (-\min(\hat{E}(g_j^p, C_{X \cup \{s_i\}}, \hat{E}(g_j^p, C_Z))) \\ & \quad + (-\min(\hat{E}(g_j^p, C_{X \cup \{s_i\}}, \hat{E}(g_j^p, C_Z))) - (-\hat{E}(g_j^p, C_Z))). \end{aligned} \quad (11)$$

Based on mean value theorem and

$$\begin{aligned} & (-\hat{E}(g_j^p, C_{X \cup \{s_i\}})) - (-\max(\hat{E}(g_j^p, C_{X \cup \{s_i\}}, \hat{E}(g_j^p, C_Z))) \\ &= (-\min(\hat{E}(g_j^p, C_{X \cup \{s_i\}}, \hat{E}(g_j^p, C_Z))) - (-\hat{E}(g_j^p, C_Z))), \end{aligned} \quad (12)$$

we get

$$\begin{aligned}
 & (-\hat{E}(g_j^p, C_{X \cup \{s_i\}})) - (-\hat{E}(g_j^p, C_X)) - (-\hat{E}(g_j^p, C_{Z \cup \{s_i\}})) + (-\hat{E}(g_j^p, C_Z)) \\
 &= (-\max(\hat{E}(g_j^p, C_{X \cup \{s_i\}}, \hat{E}(g_j^p, C_Z))) - (-\hat{E}(g_j^p, C_X)) - (-\hat{E}(g_j^p, C_{Z \cup \{s_i\}}))) \\
 &+ (-\min(\hat{E}(g_j^p, C_{X \cup \{s_i\}}, \hat{E}(g_j^p, C_Z))) = \left(-\frac{d\hat{E}(g_j^p, C_X)}{dm(g_j^p, X)}\Big|_{m(g_j^p, X)=x_1} + \frac{d\hat{E}(g_j^p, C_X)}{dm(g_j^p, X)}\Big|_{m(g_j^p, X)=x_2}\right) \cdot \\
 &\min(m(g_j^p, Z) - m(g_j^p, X), |s_i|),
 \end{aligned} \tag{13}$$

where $m(g_j^p, X) < x_1 < \min(m(g_j^p, X \cup \{s_i\}), m(g_j^p, Z))$ and $\max(m(g_j^p, X \cup \{s_i\}), m(g_j^p, Z)) < x_2 < m(g_j^p, Z \cup \{s_i\})$.

Since $x_2 > x_1$ and $\frac{d^2\hat{E}(g_j^p, C_X)}{d^2m(g_j^p, X)} \geq 0$, we know

$\frac{d\hat{E}(g_j^p, C_X)}{dm(g_j^p, X)}\Big|_{m(g_j^p, X)=x_1} \leq \frac{d\hat{E}(g_j^p, C_X)}{dm(g_j^p, X)}\Big|_{m(g_j^p, X)=x_2}$, which leads to $(-\hat{E}(g_j^p, C_{X \cup \{s_i\}})) - (-\hat{E}(g_j^p, C_X)) \geq (-\hat{E}(g_j^p, C_{Z \cup \{s_i\}})) - (-\hat{E}(g_j^p, C_Z))$, so $-\hat{E}(g_j^p, C_T)$ is submodular and non-decreasing.

Finally, $H^p(T) = -\sum_j \hat{E}(g_j^p, C_T)m(g_j^p, D_A)$ is submodular and non-decreasing because $m(g_j^p, D_A)$ does not change over X , and the linear combination of submodular and non-decreasing functions with non-negative weights is still submodular and non-decreasing.

Appendix C: Implementation details

In Eq. (8), the difference between the current training data size t_m and previous training data size t_p is always 20,000 words for real-world datasets, and 2000 words for the synthetic dataset in our experiments. When choosing samples using Eq. (8), we need to be careful about the starvation problem. That is, some types of samples are not selected in the recent history, and the samples would have low uncertainty changes which further prevents them from being selected in the future. To mitigate this issue, we alternate between using the scores in Eq. (8) and the current uncertainty $u_i^{t_m}$ to choose the next batch. For example, when plotting the performance of BALD + EDG_ext2 in Fig. 5, we select the first batch using BALD + EDG_ext2 (i.e., Eq. (8)) and the second batch using only BALD (i.e., $u_i^{t_m}$), the third batch using BALD + EDG_ext2, and so on. The same strategy is applied to US + Div + EDG_ext2 as well.

In (5), we find that w_j and v_{ij} could be set as 1 in most of the cases. However, we observe some predicted error decay curves collapse into a flat line (i.e., $b_j = 0$ in Eq. (4)) due to the unstable performance in validation set. To increase robustness, we set $w_j = \min(100, m(g_j^p, D_V))$, and $v_{ij} = 3$ if $t = \arg \min_x E_j^p(C_{T_x}, D_V)$ (i.e., lowest error for j th group across t) and $v_{ij} = 1$ otherwise in our experiments, and optimize (5) using Newton Conjugate-Gradient (Nash 1984).

In (7), we use geometric mean to combine multiple features because we usually want a sample that has large error reduction in all the groups it belongs to. Our preliminary experiments indicate that using geometric mean is better than arithmetic mean. The smoothness constant ϵ in (7) should be proportional to the size of the dataset D_A because larger error reduction could be made in a larger dataset. In our experiments, we set ϵ to be 0.01 for MedMentions ST19, 0.001 for NCBI disease and CoNLL 2003 dataset.

References

- Amari, S., Fujita, N., & Shinomoto, S. (1992). Four types of learning curves. *Neural Computation*, 4(4), 605–618.
- Bachman, P., Sordoni, A., & Trischler, A. (2017). Learning algorithms for active learning. In *ICML*.
- Baldrige, J., & Osborne, M. (2004). Active learning and the total cost of annotation. In *NIPS*.
- Bloodgood, M., & Vijay-Shanker, K. (2009). A method for stopping active learning based on stabilizing predictions and the need for user-adjustable stopping. In *CoNLL*.
- Bodenreider, O. (2004). The unified medical language system (umls): Integrating biomedical terminology. *Nucleic Acids Research*, 32(suppl 1), D267–D270.
- Bouguelia, M. R., Belaïd, Y., & Belaïd, A. (2015). Stream-based active learning in the presence of label noise. In *4th International conference on pattern recognition applications and methods-ICPRAM 2015*.
- Bouguelia, M. R., Nowaczyk, S., Santosh, K., & Verikas, A. (2018). Agreeing to disagree: Active learning with noisy labels without crowdsourcing. *International Journal of Machine Learning and Cybernetics*, 9(8), 1307–1319.
- Chen, I., Johansson, F. D., & Sontag, D. (2018). Why is my classifier discriminatory? In *Advances in neural information processing systems* (pp. 3543–3554).
- Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., & Kuksa, P. (2011). Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12(Aug), 2493–2537.
- Culotta, A., & McCallum, A. (2005). Reducing labeling effort for structured prediction tasks. In *AAAI*.
- Dasgupta, S. (2011). Two faces of active learning. *Theoretical Computer Science*, 412(19), 1767–1781.
- Doğan, R. I., Leaman, R., & Lu, Z. (2014). NCBI disease corpus: A resource for disease name recognition and concept normalization. *Journal of Biomedical Informatics*, 47, 1–10.
- Fang, M., Li, Y., & Cohn, T. (2017). Learning how to active learn: A deep reinforcement learning approach. In *EMNLP*.
- Fu, W., Wang, M., Hao, S., & Wu, X. (2018). Scalable active learning by approximated error reduction. In *SIGKDD*.
- Gal, Y., Islam, R., & Ghahramani, Z. (2017). Deep bayesian active learning with image data. In *ICML*.
- Greenberg, N., Bansal, T., Verga, P., & McCallum, A. (2018). Marginal likelihood training of bilstm-crf for biomedical named entity recognition from disjoint label sets. In *EMNLP*.
- Guillory, A., & Bilmes, J. (2010). *Interactive submodular set cover*. arXiv preprint [arXiv:10023345](https://arxiv.org/abs/10023345).
- Hestness, J., Narang, S., Ardalani, N., Diamos, G., Jun, H., Kianinejad, H., et al. (2017). *Deep learning scaling is predictable, empirically*. arXiv preprint [arXiv:171200409](https://arxiv.org/abs/171200409).
- Khetan, A., Lipton, Z. C., & Anandkumar, A. (2018). Learning from noisy singly-labeled data. In *ICLR*.
- Konyushkova, K., Sznitman, R., & Fua, P. (2017). Learning active learning from data. In *NIPS*.
- Koshorek, O., Stanovsky, G., Zhou, Y., Srikumar, V., & Berant, J. (2019). On the limits of learning to actively learn semantic representations. In *CoNLL*.
- Kremer, J., Sha, F., & Igel, C. (2018). Robust active label correction. In *International conference on artificial intelligence and statistics*.
- Lakkaraju, H., Kamar, E., Caruana, R., & Horvitz, E. (2017). Identifying unknown unknowns in the open world: Representations and policies for guided exploration. In *AAAI*.
- Lowell, D., Lipton, Z. C., & Wallace, BC. (2019). Practical obstacles to deploying active learning. In *EMNLP*.
- Lund, C., & Yannakakis, M. (1994). On the hardness of approximating minimization problems. *Journal of the ACM (JACM)*, 41(5), 960–981.
- MacQueen, J., et al. (1967). Some methods for classification and analysis of multivariate observations. *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability, Oakland, CA, USA, 1*, 281–297.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *NIPS*.
- Mitzenmacher, M., & Upfal, E. (2017). *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Oxford: Cambridge University Press.
- Murty, S., Verga, P., Vilnis, L., Radovanovic, I., & McCallum, A. (2018). Hierarchical losses and new resources for fine-grained entity typing and linking. In *ACL*.
- Musmann, S., & Liang, P. (2018). On the relationship between data efficiency and error for uncertainty sampling. In *ICML*.
- Nash, S. G. (1984). Newton-type minimization via the Lanczos method. *SIAM Journal on Numerical Analysis*, 21(4), 770–788.

- Phillips, R., Chang, K. H., & Friedler, S. A. (2018). Interpretable active learning. In *Conference on fairness, accountability and transparency*.
- Ravi, S., & Larochelle, H. (2018). Meta-learning for batch mode active learning. In *ICLR workshop*.
- Reichart, R., Tomanek, K., Hahn, U., & Rappoport, A. (2008). Multi-task active learning for linguistic annotations. In *ACL*.
- Roy, N., & McCallum, A. (2001). Toward optimal active learning through Monte Carlo estimation of error reduction. In *ICML*.
- Rubens, N., Sheinman, V., Tomioka, R., & Sugiyama, M. (2011). Active learning in black-box settings. *Austrian Journal of Statistics*, 40(1&2), 125–135.
- Sculley, D. (2010). Web-scale k-means clustering. In *WWW*.
- Sener, O., & Savarese, S. (2018). Active learning for convolutional neural networks: Acore-set approach. In *ICLR*.
- Settles, B. (2009). Active learning literature survey. *Computer Sciences Technical Report 1648*, University of Wisconsin-Madison.
- Settles, B. (2011). From theories to queries: Active learning in practice. In *Active Learning and Experimental Design workshop In conjunction with AISTATS 2010*.
- Settles, B., & Craven, M. (2008). An analysis of active learning strategies for sequence labeling tasks. In *EMNLP*.
- Settles, B., Craven, M., & Ray, S. (2008). Multiple-instance active learning. In *NIPS*.
- Shen, Y., Yun, H., Lipton, Z. C., Kronrod, Y., & Anandkumar, A. (2018). Deep active learning for named entity recognition. In *ICLR*.
- Sheng, V. S., Provost, F., & Ipeirotis, P. G. (2008). Get another label? Improving data quality and data mining using multiple, noisy labelers. In *SIGKDD*.
- Siddhant, A., & Lipton, Z. C. (2018). Deep bayesian active learning for natural language processing: Results of a large-scale empirical study. In *EMNLP*.
- Strubell, E., Verga, P., Belanger, D., & McCallum, A. (2017). Fast and accurate entity recognition with iterated dilated convolutions. In *EMNLP*.
- Sang, E. F. T. K., & De Meulder, F. (2003). Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *NAACL*.
- Tomanek, K., & Olsson, F. (2009). A web survey on the use of active learning to support annotation of text data. In *NAACL HLT 2009 workshop on active learning for natural language processing*.
- Wang, C., Chiticariu, L., & Li, Y. (2017a). Active learning for black-box semantic role labeling with neural factors. In *IJCAI*.
- Wang, W., Yang, N., Wei, F., Chang, B., & Zhou, M. (2017b). Gated self-matching networks for reading comprehension and question answering. In *ACL*.
- Wei, K., Iyer, R., & Bilmes, J. (2015). Submodularity in data subset selection and active learning. In *ICML*.
- Zhang, J., Wu, X., & Sheng, V. S. (2015). Active learning with imbalanced multiple noisy labeling. *IEEE Transactions on Cybernetics*, 45(5), 1095–1107.
- Zhao, L., Sukthankar, G., & Sukthankar, R. (2011). Incremental relabeling for active learning with noisy crowdsourced annotations. In *SocialCom/PASSAT*.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.