

Towards More Fine-grained and Reliable NLP Performance Prediction

Zihuiwen Ye, Pengfei Liu, Jinlan Fu[†], Graham Neubig

Carnegie Mellon University, [†]Fudan University

{zihuiwey, pliu3, gneubig}@andrew.cmu.edu,

fujl16@fudan.edu.cn

Abstract

Performance prediction, the task of estimating a system’s performance without performing experiments, allows us to reduce the experimental burden caused by the combinatorial explosion of different datasets, languages, tasks, and models. In this paper, we make two contributions to improving performance prediction for NLP tasks. First, we examine performance predictors not only for holistic measures of accuracy like F1 or BLEU, but also *fine-grained* performance measures such as accuracy over individual classes of examples. Second, we propose methods to understand the *reliability* of a performance prediction model from two angles: confidence intervals and calibration. We perform an analysis of four types of NLP tasks, and both demonstrate the feasibility of fine-grained performance prediction and the necessity to perform reliability analysis for performance prediction methods in the future.

We make our code publicly available: <https://github.com/neulab/Reliable-NLPPP>

1 Introduction

Performance prediction (P²) aims to predict a machine learning system’s performance based on features of the underlying problem, dataset, or learning algorithm. While this topic is still relatively unexplored in the NLP context, there are a few examples of predicting performance as: (i) a function of training or model parameters for determining the number of training iterations (Kolachina et al., 2012) or value of hyperparameters (Rosenfeld et al., 2019) and identifying and terminating bad training runs (Domhan et al., 2015). (ii) a function of dataset characteristics to illustrate which factors are significant predictors of system performance (Birch et al., 2008; Turchi et al., 2008), or find a subset of representative experiments to run in order to obtain plausible predictions (Xia et al., 2020). In this paper, we ask two

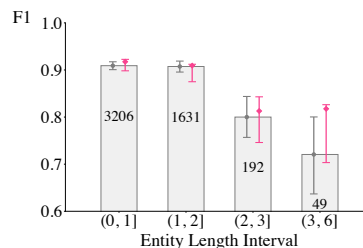


Figure 1: Breakdown of performance over different entity lengths of an NER system. *Actual F1* (gray point) is calculated from actual results while *predicted F1* (red point) is estimated by a performance prediction model. Gray and red lines represent corresponding confidence intervals. Numbers in each bar indicate the number of test samples in each length bucket.

research questions with respect to performance prediction: can we predict performance on a more *fine-grained* level, and can we quantify the *reliability* of performance predictions?

With respect to the first contribution, previous P² methods have almost entirely focused on predicting *holistic* measures of accuracy such as entity F1 (Ratinov and Roth, 2009) or BLEU score (Papineni et al., 2002) over the entire dataset (§2.2). However, from a perspective of understanding the workings of our models, work on model analysis has demonstrated the need for more *fine-grained* analysis over a wide variety of tasks (Kummerfeld et al., 2012; Kummerfeld and Klein, 2013; Karpathy et al., 2015; Neubig et al., 2019; Fu et al., 2020a,b,c). These methods calculate separate accuracy scores for different types of examples (e.g. accuracies for entity recognition by entity length). Our first contribution is to examine experimental settings where we predict these fine-grained evaluation scores (§2.3), and also propose performance prediction methods particularly suited to this fine-grained evaluation setting (§3).

Our second contribution is the development of methods for estimating the reliability of performance predictions. While allowing estimation

of experimental results without actually having to run the corresponding experiments may improve efficiency, if the performance predictor is wrong it may lead to missing results of a potentially important experiment. This particularly becomes an issue when developing methods for fine-grained performance prediction, as the number of data points which can be used to predict each performance number decreases as we subdivide datasets into finer-grained categories. Thus, we make methodological steps towards answering two specific questions: (i) how can we define and calculate a confidence interval over performance predictions? (ii) how well does the confidence interval of prediction performance calibrate with the true probability of an experimental result? Fig. 1 is an example of performance prediction and reliability analysis over fine-grained performance estimates (F1 scores over different entity length buckets) of an NER system are obtained in two ways: (i) calculated based on results from the NER system itself (in gray); (ii) estimated based on a performance prediction model, without running an actual experiment (in red). We can observe that: (1) with fewer test samples (e.g. 49), confidence intervals of both actual and predicted F1 become much wider, suggesting larger uncertainty. (2) in the last bucket, the predicted F1 (red point) is far from the actual F1 (gray point), but with a confidence interval of predicted performance (red bar), the actual F1 still falls within it, indicating the importance of knowing the level of confidence.

In experiments, we investigate the efficacy of different performance prediction models on four typical NLP tasks under both holistic and fine-grained settings, then explore methods for the reliability analysis of these performance prediction models. Major experimental results show: 1) there is no one-size-fits-all model: best-scoring performance prediction systems in different scenarios are diverse. In particular, one of our proposed models achieved the best results on the Part-of-Speech task (§6.1). 2) a better performance prediction model doesn't imply better calibration (§6.2). 3) all four performance prediction models (including previous top-scoring ones) produce confidence intervals over-confidently (§6.2).

2 Performance Prediction: Formulation and Applicable Scenarios

In this section, we will mathematically define performance prediction and its application in the

holistic and fine-grained evaluation.

2.1 Formulation

Given a machine learning model \mathcal{M} , which is trained over a training set \mathcal{D}^{tr} based on a specific training strategy \mathcal{S} , we then test the dataset \mathcal{D}^{ts} under evaluation setting \mathcal{E} and the test result y can be formulated as a function of the following inputs:

$$y = f(\mathcal{M}, \mathcal{D}^{tr}, \mathcal{D}^{ts}, \mathcal{S}, \mathcal{E}), \quad (1)$$

This we will refer to as the *actual* performance (e.g., F1 score), which requires us to run an actual experiment.

Alternatively, to calculate y , instead of performing a full training and evaluation cycle, one can directly estimate it by extracting features of \mathcal{M} , \mathcal{D}^{tr} , \mathcal{D}^{ts} , \mathcal{S} , and running them through a prediction function

$$\hat{y} = g(\Phi_{\mathcal{M}}, \Phi_{\mathcal{D}^{tr}}, \Phi_{\mathcal{D}^{ts}}, \Phi_{\mathcal{S}}, \mathcal{E}; \Theta), \quad (2)$$

where $\Phi(\cdot)$ represents features of the input, and Θ denotes learnable parameters. We will refer to this as our *predicted* performance. As long as Eq. 2 is fast to calculate and a relatively accurate approximation of Eq. 1, it allows us to get a reasonable idea of expected experimental results much more efficiently than if we had to actually experiment.

In a real scenario, not all inputs in Eq. 2 need to be taken into account, and researchers can adopt different inputs for a particular use. For example, Domhan et al. (2015) define \hat{y} as a function of training strategy \mathcal{S} (e.g., different hyper-parameter settings) so that they can know which training setting can lead to bad performance without running. Dodge et al. (2020) estimate validation performance as a function of computation budget to conduct more robust model comparisons.

Why Performance Prediction matters for NLP tasks

Firstly, for some NLP tasks with few resources, it is challenging to build and test systems for all languages or domains. For example, the task of Machine Translation (MT) for low resource languages is hard due to the lack of the large parallel corpora, preventing us from measuring system performance in these scenarios (Xia et al., 2019, 2020). Therefore, performance prediction is useful in that it can efficiently and comprehensively give insights about the workings of models over a wide variety of task settings. Secondly, performance prediction can be used to alleviate the data

sparsity problem in fine-grained evaluation, which plays an important role in current NLP task evaluation (Fu et al., 2020a).

In this paper, we consider two performance prediction scenarios, a holistic evaluation setting that most previous works have explored, and a novel setting of predicting fine-grained evaluation metrics. Below, we briefly describe them.

2.2 Holistic Evaluation

Performance prediction in holistic evaluation aims to estimate an overall score (e.g., BLEU) based on dataset characteristics, specifically,

$$\hat{y} = g^{\text{holistic}}(\Phi_{\mathcal{D}^{tr}}, \Phi_{\mathcal{D}^{ts}}; \Theta), \quad (3)$$

where $\Phi(\cdot)$ represents features of input and Θ denotes learnable parameters.

Featurization In practice, we choose a machine translation (MT) task and a Part-of-Speech task (POS) task in this setting. We use the same set of dataset features as (Xia et al., 2020), including the language features and the source and the target, or transfer language.

2.3 Fine-grained Evaluation

In contrast, fine-grained evaluation aims to break down the overall score into different interpretable parts, allowing us to identify the strengths and weakness of learning systems. For example, the accuracy of an NER system with an overall F1 score 90 (%) can be partitioned into four buckets based on different entity lengths l (e.g., [$l = 1, 1 < l \leq 3, 3 < l \leq 5, l > 5$]) of test entities, thereby obtaining fine-grained F1 scores: [93, 91, 89, 75], identifying that the model struggles on longer entities ($l > 5$).

Although fine-grained evaluation is advantageous in interpreting systems’ performance, it frequently suffers from the data sparsity problem—a few or no test samples may be included within certain buckets. For example, in the above case it’s difficult to calculate the F1 score for entities whose lengths satisfy $l > 7$ since few entities can be found in the whole test set.

With the above dilemma in mind, we define a performance prediction problem in fine-grained evaluation where the paucity of test samples in some buckets leads to an inability to compute performance accurately.

$$\hat{y} = g^{\text{fine}}(\Phi_M, \Phi_{\mathcal{D}^{tr}}, \Phi_{\mathcal{D}^{ts}}; \Theta), \quad (4)$$

where $\Phi(\cdot)$ represents features of input and Θ denotes learnable parameters.

Featurization Performing fine-grained evaluation involves two major steps: (i) partition the test set into different *buckets* based on a certain *aspect* (e.g., entity length), (ii) and calculate performance (e.g., F1 score) for each bucket. Therefore, data-wise ($\Phi_{\mathcal{D}^{ts}}$), the input of performance prediction function in Eq. 4 ($g^{\text{fine}}(\cdot)$) can be featurized as different types of (i) buckets (ii) aspects (iii) datasets. Additionally, we take (iv) different types of models as input. We present brief descriptions of the above four types of features.

1. *Models*: We choose 12 models for the NER task and 8 models for the Chinese Word Segmentation (CWS) task. The models are built by choosing the different character encoder (e.g., ELMo (Peters et al., 2018) and Flair (Akbi et al., 2018; Akbi et al.)), word embedding (e.g., GloVe (Pennington et al., 2014) and Word2Vec (Mikolov et al., 2013b)), sentence-level encoder (e.g., LSTM (Hochreiter and Schmidhuber, 1997) and CNN (Kalchbrenner et al., 2014)), and decoder (e.g., MLP and CRF (Lample et al., 2016; Collobert et al., 2011)).

2. *Datasets*: We consider 6 (5) datasets for the NER (CWS) task, detailed in appendix.

3. *Attributes*: We consider the interpretable evaluation aspects proposed in works (Fu et al., 2020a). We consider 9 attributes for the NER task and 8 attributes for the CWS task in this paper (e.g, entity length and sentence length).

4. *Buckets*: The test entities (words) of the NER (CWS) task are partitioned into four buckets according to their attribute value. We compute the F1 score for the entities.

3 Parameterized Regression Functions

The performance prediction model takes in a set of features that characterize an experiment’s peculiarities and predict performances based on different parameterized regressors $g(\cdot)$ in Eq. 2. We first describe methods explored by previous works and then present a tensor regression-based approach that is particularly well-suited for fine-grained performance prediction.

3.1 Gradient Boosting Methods

Previous work on performance prediction has used gradient boosted decision tree models (Ganjisaf-




Figure 2: Illustration of performance tensor in the fine-grained evaluation scenario. Colored entries represent missing performances that would be predicted.

far et al., 2011; Chen and Guestrin, 2016), which demonstrate robust performance on the relatively low-data scenarios we often encounter in performance prediction tasks. We specifically explore the following two models:

XGBoost (Chen and Guestrin, 2016) is a tree boosting system widely used to solve problems such as ranking, classification, and regression. We use the same experimental setting as described in (Xia et al., 2020).

LightGBM (Ke et al., 2017) is a gradient boosting framework. Compared with XGBoost, which utilizes a level-wise tree growth in the decision tree, LightGBM uses a leaf-wise splitting method.

3.2 Tensor Regression

Besides gradient boosted trees, we also present tensor regression-based performance prediction models. Tensors are multidimensional arrays that can concisely depict the structure of the data. The order of a tensor is its number of dimensions. For example, in the NER task, the four feature dimensions of a tensor are model, dataset, attribute and bucket, with each slice representing one underlying relationship between the two dimensions. Applying tensor factorization algorithms in the performance prediction setting allows us to determine the interdependencies between multiple aspects of the tasks simultaneously.

Performance Prediction as Tensor Completion

To formulate the performance prediction task as a tensor regression problem: (i) we first define a *performance tensor* that each entry stores a performance value under a specific setting determined by input features (described in §2.3); (ii) missing entries in performance tensor can be predicted based on different tensor completion techniques.

Specifically, taking fine-grained evaluation for example, we define a fine-grained performance tensor as $\mathbf{Y} \in \mathbb{R}^{I_1 \times I_2 \times I_3 \times I_4}$, where \mathbf{Y}_{ijkl} denotes

the performance (e.g. F1 score) of the i -th model (e.g. BERT-based Tagger) on the j -th bucket (e.g. 2nd) that is obtained by partitioning the k -th dataset (e.g. CoNLL03) based on the t -th attribute (e.g. entity length). I_1, I_2, I_3, I_4 denote the number of models, buckets, datasets, attributes. Fig. 2 elaborates on this, in which three dimensions (buckets, datasets, and attributes) are considered for the sake of presentation.

CP Decomposition The CP decomposition (Hitchcock, 1927) expresses a tensor \mathbf{Y} as a sum of lower rank tensors. For example, an order 4 tensor can be decomposed as the sum of R rank-1 tensors, each being the outer product of four vectors in each dimension.

Robust PCA Robust PCA is a modification of principal component analysis (PCA) (Candès et al., 2009). If a tensor can be conceived as a superposition of low-rank components and a sparse component, Robust PCA attempts to recover the low-rank and sparse components. The sparse components can be considered as the gross, but sparse noise in the dataset.

4 Statistical Preliminaries

Before going into our second contribution to establishing reliability of performance prediction, we describe two relevant concepts from statistics.

4.1 Confidence Interval (CI)

The confidence interval (CI) is a range of possible values for an unknown parameter associated with a *confidence level* of γ (Nakagawa and Cuthill, 2007; Dror et al., 2018) that the actual parameter can fall into the suggested range. Specifically, suppose that we are interested in estimating the underlying true parameter of ω . Given an observed parameter estimate of $\hat{\omega}$, obtained from the data, we aim to compute an interval with a confidence level γ that ω lies in an interval CI.

Commonly, there are two approaches to calculate confidence intervals, depending on our knowledge about the distribution of the statistics of interest. When an analytical form exists and we have reasonable assumptions on the distribution, we can employ the normal theory or use Student’s t-distribution to construct a confidence interval.

Regarding data drawn from a completely unknown distribution, a CI can be calculated by a bootstrapping method (Efron, 1992; Johnson, 2001). The main idea behind the bootstrapping

method is to simulate the real distribution by sampling with replacement from a distribution that approximates it, thereby allowing us to make inferences about the statistics of interest and construct confidence intervals. Common methods to construct the CI with bootstrap include the percentile method, where after specifying a confidence level γ , we take the range of points that cover the middle γ proportion of bootstrap sampling distribution \hat{Y} as the desired confidence interval, represented by $(Q_{\hat{Y}((1-\gamma)/2)}, Q_{\hat{Y}((1+\gamma)/2)})$, where Q denotes the quantile. Works on establishing confidence for results in NLP tasks using this bootstrap method include Koehn (2004) and Li et al. (2017).

4.2 Model Calibration (MC)

Calibration (Gleser, 1996), also known as reliability, refers to the ability of a model to make good probabilistic predictions. For a discrete distribution over events, a model is said to be well-calibrated if for those events that the model assigns a probability of p , the long-run proportion that the event actually occurs turns out to be p . For example, if a weather forecast model predicts that there is a 0.1 probability of rain at 7 a.m., then when observed on a large number of random trials at 7 a.m., the model is well-calibrated if 0.1 of them actually do result in rain. Similarly, for a classification model matching the probability a model assigns to a predicted label (i.e., *confidence*) and the correctness measure of the prediction (i.e., *accuracy*) (Wang et al., 2020) is desired.

Nonetheless, it is common that a model could have a high predictive accuracy, but poor calibration if the model systematically over- or under-estimates its confidence in the predictions it makes. One way to quantify miscalibration is to use *Expected Calibration Error* (ECE; Nae (2015)), which aims to quantitatively characterize the difference in expectation between confidence and accuracy. To calculate ECE, the predictions should first be partitioned into M buckets based on the confidence of the predictions, where N represents the total number of prediction samples and $|B_m|$ is the number of samples in the m -th bucket. Given these buckets, ECE can be defined as,

$$\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{N} |\text{acc}(B_m) - \text{conf}(B_m)|, \quad (5)$$

where $\text{acc}(B_m)$ denotes the accuracy of B_m ,

$$\text{acc}(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} \mathbf{1}(\hat{y}_i = y_i), \quad (6)$$

where \hat{y} and y represent predicted and ground truth labels respectively. $\text{conf}(B_m)$ represents the average confidence of bucket B_m ,

$$\text{conf}(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} \hat{p}_i. \quad (7)$$

where \hat{p} represents the prediction confidence of sample i .

5 On Reliability of P^2 Models

Now we discuss our methodology for predicting the reliability of performance prediction models through confidence intervals and the calibration of those confidence intervals.

5.1 CIs of Predicted Performance

We refer to $y \sim Y$ as an actual observed performance as in Eq. 1 for a specific task (e.g., NER). y is the output of an NLP system learned on a dataset $\mathcal{D} = (\mathcal{D}^{tr}, \mathcal{D}^{ts})$. We refer to $\hat{y} \sim \hat{Y}$ as a predicted performance estimated as Eq. 2. \hat{y} is the output of a performance prediction model learned from a dataset $\Phi(\mathcal{D}) = (\Phi(\mathcal{D}^{tr}), \Phi(\mathcal{D}^{ts}))$, where $\Phi(\cdot)$ represents the input dataset features. Our goal is to compute a confidence interval w.r.t a predicted performance \hat{y} , to make inference about Y .

Bootstrap for CI of Predicted Performance

One potential challenge is that we cannot make plausible assumptions about the distribution of predicted performances \hat{Y} , which prevents us from using popular parametric methods (as mentioned in § 4.1) to calculate the confidence interval. Instead, we resort to a bootstrap resampling method as adopted in (Efron, 1992), to simulate \hat{Y} .

To achieve this, we first (i) sample different training sets for the performance prediction model $\Phi(\mathcal{D})_1^{tr}, \Phi(\mathcal{D})_2^{tr}, \dots, \Phi(\mathcal{D})_K^{tr} \sim \Phi(\mathcal{D})^{tr}$, and then (ii) train K performance prediction models using Eq. 2 on each of the K partitions, and (iii) evaluate K models on $\Phi(\mathcal{D})^{ts}$, thereby obtaining a prediction distribution \hat{Y} . From this resampling distribution, we use the percentile method, taking the top $(1 - \gamma)/2$ and the bottom $(1 + \gamma)/2$ of the distribution as higher and lower bounds for the confidence interval.

5.2 Calibration of CI

Because we calculate confidence intervals of the predicted performance \hat{y} , drawn from the distribution of \hat{Y} , rather than the actual y from Y , it’s still unclear if our predicted CI is reliable enough to cover the actual performance. In other words, “from an infinite number of independent trials, does the true value *actually* lie within the intervals approximately 95% of the times?”

To answer this question, we establish a method to measure calibration for the confidence interval of predicted performance. To check (i) if y could be generally contained in the prediction intervals reasonably well, and (ii) if a prediction model produces predictions that are not over or under-confident, we empirically examine the prediction distributions and establish the reliability of the confidence intervals.

To this end, we extend the definition of calibration in classification setting to our regression problem. Specifically, we formulate confidence level γ as prediction confidence conf defined in Eq. 7, and then the original definition of different M buckets can be instantiated as different confidence levels here: $\gamma_1, \dots, \gamma_M$. The *accuracy* at each confidence level γ_b defined as follows:

$$\text{acc}(\gamma_b) = \frac{\sum_{i=1}^N \mathbb{1}(A < y_i < B)}{N}, \quad (8)$$

where $i \in [1, N]$, $b \in [1, M]$. N represents the number of test samples. y_i denotes the actual performance for the test sample i . $A = (Q_{\hat{Y}((1-\gamma_b)/2)})$ and $B = (Q_{\hat{Y}((1+\gamma_b)/2)})$.

Intuitively, $\text{acc}(\gamma_b)$ represents the relative frequency of the actual value y falling into the predicted confidence interval w.r.t. \hat{y} . Fig. 3 illustrates how $\text{acc}(\gamma_b)$ is calculated: given three samples whose performances are to be predicted, the denominator of $\text{acc}(\gamma_b = 0.8)$ is 3 while the numerator tallies how many times (2 in this case) the actual performances (i.e., y_1, y_2, y_3) of three samples fall into the confidence interval (with $\gamma_b = 0.8$) of corresponding bootstrapped distributions.

Based on Eq. 8, we can re-write a calibration error CE as:

$$\begin{aligned} \text{CE} &= \sum_{b=1}^M | \text{acc}(\gamma_b) - \gamma_b | \\ &= \sum_{b=1}^M \left| \frac{\sum_{i=1}^N \mathbb{1}(A < y_i < B)}{N} - \gamma_b \right| \end{aligned} \quad (9)$$




Figure 3: Illustration of calibration on confidence intervals ($\gamma_b = 0.8$) w.r.t. the predicted performance. Solid lines represent actual performances (y) while dashed lines denote resampled predicted performances using bootstrap. Boundaries of shaded areas indicate confidence interval $(Q_{\hat{Y}((1-\gamma_b)/2)}, Q_{\hat{Y}((1+\gamma_b)/2)})$. Intuitively, y_1 and y_2 fall into confidence intervals of corresponding bootstrapped distributions \hat{Y}_1 and \hat{Y}_2 .

6 Experiments

In this section, we break down our experimental results into answering two research questions sections: (1) how well do our underlying performance predictors work, particularly the newly proposed tensor-based predictors and on the newly proposed task of fine-grained performance prediction? (2) how well can we estimate the reliability of our performance predictions?

Models Besides the four performance prediction models (CP, PCA, XGBoost, LGBM) that we have introduced in §3, following Xia et al. (2020), we additionally use a simple mean value baseline model which predicts an average of scores s from the training folds for all test entries in the left-out evaluation fold:

$$\hat{s}_{\text{mean}}^{(i)} = \frac{1}{|\mathcal{D} \setminus \mathcal{D}^{(i)}|} \sum_{s \in \mathcal{D} \setminus \mathcal{D}^{(i)}} s; i \in 1, \dots, k, \quad (11)$$

where $\mathcal{D}^{(i)}$ is the left-out data used to evaluate the model performance.

Hyper-parameters Detailed information about the hyper-parameters used in training the performance prediction models in various tasks is provided in the appendix.

Tasks We explore performance prediction on four tasks: (1) Machine Translation (MT) (Schwenk et al., 2019), (2) Part-of-Speech tagging (POS), (3) Named Entity Recognition (NER), (4) Chinese Word Segmentation (CWS). To compare the performance of tensor-based models and gradient boosting models on the same dataset, we convert the datasets used in different prediction tasks

to tensors. Statistics of the tensor data are shown in the appendix.

6.1 Evaluation of Performance Prediction

Setup To investigate the effectiveness of the performance prediction models across different tasks, we conduct k -fold cross-validation for evaluation. Specifically, we randomly partition the entire experimental data \mathcal{D} into $k = 5$ folds, use 4 folds for training, and test the model’s performance on the remaining fold. To evaluate the result, we calculate the average root mean square error (RMSE) between the predicted scores and the true scores.

Results The RMSE scores of different performance tasks are shown in Tab. 1. Notably, RMSE scores across different tasks should not be compared directly, because the scales of the evaluation metrics are different. We observed that:

(1) Overall, all four models we investigated outperform the baseline by a large margin, indicating their effectiveness on these four performance prediction tasks. (2) Comparing two tensor-based models, PCA consistently outperforms CP. Notably, our proposed tensor regression model (PCA) has surpassed the previous best-performing system (XGBoost (Xia et al., 2020)) on the POS dataset and achieved comparable result on the MT dataset despite the relatively high sparsity of the tensor (0.346). (3) We observe that CP achieves much worse performance on the POS dataset. One potential reason is that: CP is sensitive to datasets (like POS) that exhibit large variance along some feature dimensions, which can not be alleviated by feature scaling. (4) There is no one-size-fits-all model: on different datasets, the corresponding best-scoring performance prediction models are diverse, suggesting that we should take dataset’s characteristics into account when selecting a model for a specific performance prediction scenario.

Prediction Error Analysis In §1 and Fig. 1, we reveal how entities with different lengths influence the performance prediction, a result of the underlying paucity of data. Here we perform a more detailed error analysis to understand the factors that influence the performance of performance prediction models. Specifically, we perform a case study on the NER task using XGBoost and look for feature combinations on which performance predictions show poor results. We use XGBoost to predict F1 scores on all possible combinations of four

Model	Fine-grained		Holistic	
	NER	CWS	MT	POS
Baseline	0.209	0.137	6.388	29.09
XGBoost	0.055	0.021	2.463	7.319
LGBM	0.059	0.041	2.389	7.673
CP	0.068	0.043	4.065	24.70
PCA	0.057	0.029	2.920	5.860

Table 1: Results (RMSE, lower scores indicate better performances) of different performance prediction models on four tasks. The lowest value of each column is bold.

feature dimensions (models, datasets, attributes, and buckets) to obtain \hat{y}_{ijkl} using the combined test sets from 5-fold cross-validation. For each prediction, we calculate a square residual $(\hat{y} - y)^2$. Then, we group the square residuals by 2 of the 4 dimensions¹ and take their mean value aggregated over the other 2 dimensions. Fig. 4 shows the aggregated mean square residual (MSR) fixed on the model and dataset dimensions, and Fig. 5 shows the result fixed on the attribute and bucket dimensions. In both figures, a high MSR (dark grid) means a poor performance prediction. In Fig. 4, we notice that (1) dataset-wise: WB and WNUT, and (2) model-wise: *CnnWgloveLstmMlp* and *CnoneWrandLstmCrf* show poor results. We observe that (1) WB is generated from weblogs and WNUT is generated from Twitter, both of which are noisy. (2) *CnnWgloveLstmMlp* does not use a CRF-decoder, and *CnoneWrandLstmCrf* does not encode character-level features, both of which are important characteristics in building an NER model. It is plausible that the systems have an unstable performance in those experimental settings and thus make them harder to predict. In Fig. 5, we notice that (1) a lower bucket value along the attributes *entity consistency*, *token consistency*, and *entity density*, (2) a higher bucket value along the attributes *token frequency* or *entity length* lead to poor performance prediction results. In other words, the performance prediction model finds it hard to predict when there is a low label consistency of token or entity, a low entity density, and when token frequency is high and entity is long.

6.2 Evaluation of Reliability

Setup As described in §5.1, we use non-parametric bootstrap to produce confidence intervals for \hat{y} . For the holistic evaluation setting, we

¹Readers can refer to this work (Fu et al., 2020a) to get more details.




Figure 4: Each grid in the heatmap denotes the mean square residual fixed on the corresponding model (y-axis) and dataset (x-axis) aggregated over all attributes and buckets. The colorbar on the right denotes the value of the mean square residual. Readers can refer to this work (Fu et al., 2020a) to get more details about the information of models and attributes.




Figure 5: Each grid denotes the mean square residual fixed on the corresponding attribute (y-axis) and bucket (x-axis) aggregated over all models and datasets. Buckets 1 to 4 are ordered in increasing attribute value.

do not include tensor-based models since the property, “with replacement”, of the bootstrap makes it difficult to construct resampled tensors in the holistic evaluation setting.

When calculating a calibration error as defined in Eq. 10, we set $M = 20$, choosing a range of 20 increasing confidence levels, (i.e. $\gamma_1 = 0.05, \gamma_2 = 0.10 \dots, \gamma_{20} = 1.00$), to evaluate the correctness of the confidence intervals given by prediction models. Besides using a reliability diagram and a calibration error, to compare the calibration performances of different models more comprehensively, we additionally use the following quantitative metrics: (1) *average width* is the mean range of all the prediction distributions, formally, the difference between the maximum and minimum: $\frac{1}{N} \sum_{i \in [1, N]} \max(\hat{Y}_i) - \min(\hat{Y}_i)$. (2) *coverage* is the value of acc_b evaluated at $\gamma = 1$. (i.e. proportion of $y \hat{y}$ that fall into the distributions \hat{Y} , out of all the N prediction entries).

Results The reliability diagram of different models and their corresponding metrics on four tasks are illustrated in Fig. 6 and Tab. 2 respec-

Model	CE	Wid.	Cov.	CE	Wid.	Cov.
	NER			CWS		
Baseline	10.47	0.006	0.01	10.50	0.003	0.01
XGBoost	4.60	0.093	0.75	4.38	0.045	0.77
LGBM	6.78	0.093	0.50	7.41	0.029	0.46
CP	6.33	0.110	0.68	5.22	0.099	0.85
PCA	8.76	0.051	0.31	9.87	0.015	0.12

Model	MT			POS		
	Baseline	9.98	1.46	0.08	10.30	3.96
XGBoost	3.75	5.55	0.81	3.96	17.61	0.82
LGBM	7.23	3.01	0.44	8.02	12.90	0.34

Table 2: Calibration errors (CE), average width (Wid.), coverage (Cov.) of different models over four tasks. (NER, CWS for fine-grained evaluation setting and MT, POS for holistic evaluation setting)

tively. Intuitively, the smaller the CE (calibration error) value is, the more closely the black dotted line becomes diagonal. Ideally, a perfectly calibrated model should have a CE of 0.

From these two tables, we see that: (1) Overall, in both holistic (MT and POS) and fine-grained settings (NER and CWS), we see that XGBoost achieves the lowest calibration error together with a higher coverage, especially in the holistic setting. (2) We observe that all of the plots indicate that the intervals produced by the models are over-confident, as the dots lie under the identity function. In other words, given a confidence level γ , the actual accuracy is lower than γ . (3) In Tab. 1, we find that LGBM achieves the lowest RMSE (2.389) in task MT, but its calibration error (7.23) is worse than XGBoost (3.75), implying that a model that predicts accurately is not necessarily well calibrated. This could be explained by the observation that the predicted distribution \hat{Y} of LGBM has a narrower width (3.01). Given a large number of trials predicted by LGBM, we cannot be confident that the true y is contained in the range of values predicted.

Case Analysis To get a better understanding of how calibration analysis is conducted on different performance models, we perform a case study on NER task. Fig. 7(a-b) illustrates two plots that artificially simulate two common relations following Diebold et al. (1997) between actual and predicted distribution: (i) *Bias* (ii) *Over-confidence*. From Tab. 2, we see that XGBoost is better calibrated than LGBM in NER task. To interpret this gap, we (i) first randomly select test samples from NER dataset and then (ii) use two performance prediction models XGBoost and LGBM to produce blue

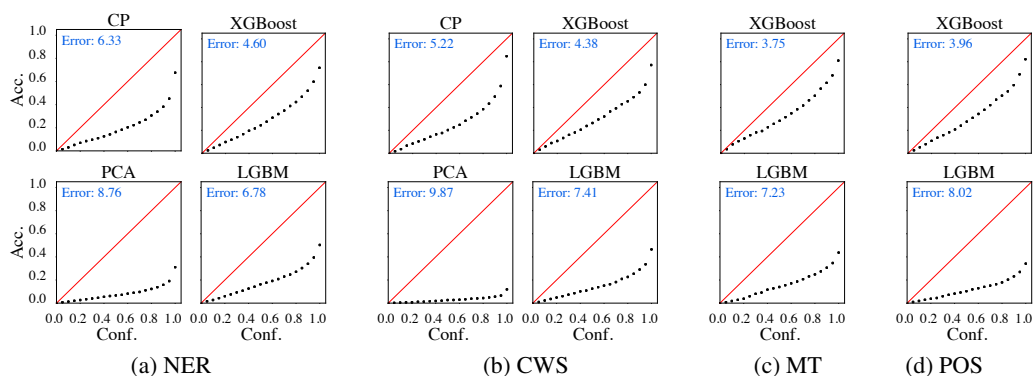


Figure 6: Calibration of different performance prediction models on four tasks (NER, CWS for fine-grained evaluation setting and MT, POS for holistic evaluation setting).

distributions in Fig. 7(c-d) using the bootstrap (as §5.1). A perfectly calibrated model will show a histogram shape that resembles the actual one. We can see that the histogram shape of (d) signifies an over-confidence problem, in which the predicted distribution (in blue) is covered by the actual distribution (in red). By contrast, in (c) the histogram of XGBoost in blue shifts to the left compared with the actual observed distribution, indicating that the prediction on this bucket is biased.

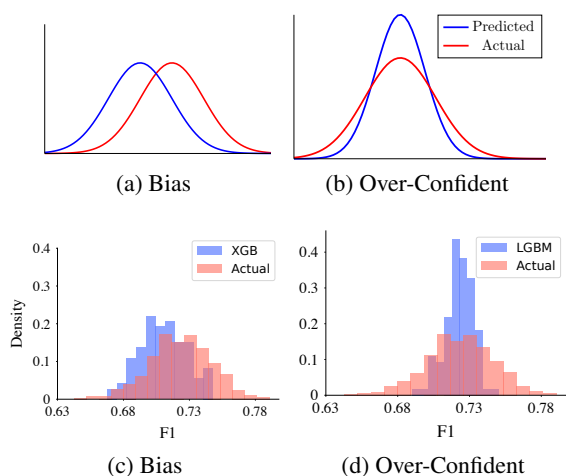


Figure 7: The first row of two plots (a,b) artificially simulate two typical relations between actual and predicted distributions. The second row of two plots (c,d) show two real-world distributions of predicted performance w.r.t one test sample from NER task against corresponding actual distributions.

7 Implications and Future Directions

In this work, we not only widen the applicability of performance prediction, extending it to fine-grained evaluation scenarios, but also establish a set of reliability analysis mechanisms to improve its practicality. In closing, we highlight some potential future directions:

Confidence over confidence: Our work provides an idea for reliability analysis of the predicted confidence interval, which could also be explored on other scenarios, e.g., density forecasting (Diebold et al., 1997). Another potentially valuable research topic is to build connections with the probability integral transform (Angus, 1994), which is a typical method of calibration evaluation in financial risks, and our proposed calibration method.

Calibration for automated evaluation metrics: From a broader point of view, the role of existing learnable automatic evaluation metrics for text generation, such as BLEURT (Sellam et al., 2020) and COMET (Rei et al., 2020), is similar to a performance prediction model (i.e., both take features of input data as input and then output an evaluation score). Reliability analysis of these metrics is also an important topic since they determine the direction of model optimization.

Acknowledgements

We sincerely thank all reviewers for their insightful comments and suggestions. We also thank Mengzhou Xia for discussions on details of performance prediction for different NLP tasks. This work was supported in part by a grant under the Northrop Grumman SOTERIA project and the Air Force Research Laboratory under agreement number FA8750-19-2-0200. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory or the U.S. Government.

References

2015. Obtaining well calibrated probabilities using Bayesian Binning. *Proceedings of the National Conference on Artificial Intelligence*, 4:2901–2907.
- Alan Akbik, Tanja Bergmann, and Roland Vollgraf. Pooled contextualized embeddings for named entity recognition.
- 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.
- John E Angus. 1994. The probability integral transform and related results. *SIAM review*, 36(4):652–654.
- Alexandra Birch, Miles Osborne, and Philipp Koehn. 2008. [Predicting success in machine translation](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 745–754, Honolulu, Hawaii. Association for Computational Linguistics.
- Emmanuel J. Candès, Xiaodong Li, Yi Ma, and John Wright. 2009. [Robust principal component analysis?](#) *CoRR*, abs/0912.3599.
- Hui Chen, Zijia Lin, Guiguang Ding, Jianguang Lou, Yusen Zhang, and Borje Karlsson. 2019. Grn: Gated relation network to enhance convolutional neural network for named entity recognition. 33(01):6236–6243.
- Tianqi Chen and Carlos Guestrin. 2016. [Xgboost: A scalable tree boosting system](#). *CoRR*, abs/1603.02754.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *The Journal of Machine Learning Research*, 12:2493–2537.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Francis X Diebold, Todd A Gunther, and Anthony S Tay. 1997. Evaluating Density Forecasts. *International Economic Review*, 39:863–883.
- Jesse Dodge, Suchin Gururangan, Dallas Card, Roy Schwartz, and Noah A. Smith. 2020. [Show your work: Improved reporting of experimental results](#). *EMNLP-IJCNLP 2019 - 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing, Proceedings of the Conference*, (2):2185–2194.
- Tobias Domhan, Jost Tobias Springenberg, and Frank Hutter. 2015. Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- Rotem Dror, Gili Baumer, Segev Shlomov, and Roi Reichart. 2018. [The hitchhiker’s guide to testing statistical significance in natural language processing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1383–1392, Melbourne, Australia. Association for Computational Linguistics.
- Bradley Efron. 1992. Bootstrap methods: another look at the jackknife. In *Breakthroughs in statistics*, pages 569–593. Springer.
- Jinlan Fu, Pengfei Liu, and Graham Neubig. 2020a. Interpretable multi-dataset evaluation for named entity recognition. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Jinlan Fu, Pengfei Liu, Qi Zhang, and Xuan-Jing Huang. 2020b. Is chinese word segmentation a solved task? rethinking neural chinese word segmentation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5676–5686.
- Jinlan Fu, Pengfei Liu, Qi Zhang, and Xuanjing Huang. 2020c. Rethinking generalization of neural models: A named entity recognition case study. In *AAAI*, pages 7732–7739.
- Yasser Ganjisaffar, Rich Caruana, and Cristina Videira Lopes. 2011. Bagging gradient-boosted trees for high precision, low variance ranking models. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, pages 85–94.
- Leon Jay Gleser. 1996. Measurement, regression, and calibration.
- Frank L. Hitchcock. 1927. [The expression of a tensor or a polyadic as a sum of products](#). *Journal of Mathematics and Physics*, 6(1-4):164–189.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Roger W Johnson. 2001. An introduction to the bootstrap. *Teaching statistics*, 23(2):49–54.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *Proceedings of ACL*.
- Andrej Karpathy, Justin Johnson, and Li Fei-Fei. 2015. Visualizing and understanding recurrent networks. *arXiv preprint arXiv:1506.02078*.

- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie Yan Liu. 2017. LightGBM: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems*, 2017-December(Nips):3147–3155.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 4:388–395.
- Prasanth Kolachina, Nicola Cancedda, Marc Dymetman, and Sriram Venkatapathy. 2012. [Prediction of learning curves in machine translation](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 22–30, Jeju Island, Korea. Association for Computational Linguistics.
- Jonathan K. Kummerfeld, David Hall, James R. Curran, and Dan Klein. 2012. [Parser showdown at the wall street corral: An empirical investigation of error types in parser output](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1048–1059, Jeju Island, Korea. Association for Computational Linguistics.
- Jonathan K. Kummerfeld and Dan Klein. 2013. [Error-driven analysis of challenges in coreference resolution](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 265–277, Seattle, Washington, USA. Association for Computational Linguistics.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *Proceedings of NAACL-HLT*, pages 260–270.
- Junhui Li, Deyi Xiong, Zhaopeng Tu, Muhua Zhu, Min Zhang, and Guodong Zhou. 2017. [Modeling source syntax for neural machine translation](#).
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Shinichi Nakagawa and Innes C Cuthill. 2007. Effect size, confidence interval and statistical significance: a practical guide for biologists. *Biological reviews*, 82(4):591–605.
- Graham Neubig, Zi-Yi Dou, Junjie Hu, Paul Michel, Danish Pruthi, and Xinyi Wang. 2019. [compare-mt: A tool for holistic comparison of language generation systems](#). *arXiv preprint arXiv:1903.07926*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. 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.
- 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 NAACL*, volume 1, pages 2227–2237.
- Lev Ratinov and Dan Roth. 2009. Design challenges and misconceptions in named entity recognition. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL-2009)*, pages 147–155.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. Comet: A neural framework for mt evaluation. *arXiv preprint arXiv:2009.09025*.
- Jonathan S Rosenfeld, Amir Rosenfeld, Yonatan Belinkov, and Nir Shavit. 2019. A constructive prediction of the generalization error across scales. *arXiv preprint arXiv:1909.12673*.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2019. [Wikimatrix: Mining 135m parallel sentences in 1620 language pairs from wikipedia](#). *CoRR*, abs/1907.05791.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Marco Turchi, Tijl De Bie, and Nello Cristianini. 2008. Learning performance of a machine translation system: a statistical and computational analysis. In *Proceedings of the Third Workshop on Statistical Machine Translation*, pages 35–43.
- Shuo Wang, Zhaopeng Tu, Shuming Shi, and Yang Liu. 2020. [On the inference calibration of neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3070–3079, Online. Association for Computational Linguistics.
- Mengzhou Xia, Antonios Anastasopoulos, Ruochen Xu, Yiming Yang, and Graham Neubig. 2020. [Predicting performance for natural language processing tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*,

pages 8625–8646, Online. Association for Computational Linguistics.

Mengzhou Xia, Xiang Kong, Antonios Anastasopoulos, and Graham Neubig. 2019. [Generalized data augmentation for low-resource translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5786–5796, Florence, Italy. Association for Computational Linguistics.

A Datasets

Named Entity Recognition (NER) We choose two well-established benchmark datasets: CoNLL-2003² and OntoNotes 5.0.³ CoNLL-2003 is drawn from Reuters news. OntoNotes 5.0 is collected from newsgroups (NW), broadcast news (BN), broadcast conversation (BC), weblogs (WB), magazine genre (MZ), and telephone speech (TC), in which the first five genres of text are used in this paper.

Chinese Word Segmentation (CWS) We consider five mainstream datasets: CKIP, CTB, MSR, NCC, and SXU, from SIGHAN2005⁴ and SIGHAN2008⁵. The traditional Chinese characters in CKIP are mapped to simplified Chinese characters in our experiment.

B Models

Our NER (CWS) models can be decomposed into four aspects: 1) character/subword encoders; 2) word (bigram) embeddings; 3) sentence-level encoders; 4) decoders.

The four aspects of NER can be summarized as: 1) character/subword encoder: ELMo (Peters et al., 2018), Flair (Akbik et al., 2018; Akbik et al.), BERT⁶ (Peters et al., 2018; Devlin et al., 2018); 2) additional word embeddings: GloVe (Pennington et al., 2014); 3) sentence-level encoders: LSTM (Hochreiter and Schmidhuber, 1997), CNN (Kalchbrenner et al., 2014; Chen et al., 2019); 4) decoders: MLP or CRF (Lample et al., 2016; Collobert et al., 2011).

The four aspects’ setting of CWS: 1) character/subword encoder: ELMo, BERT; 2) bigram embeddings: Word2Vec (Mikolov et al., 2013a), averaging the embedding of two contiguous charac-

Feature	NER	CWS	MT	POS
Sparisity	0.0	0.0	0.346	0.019
Shape	(11,6,9,4)	(5,8,8,4)	(39,39,22)	(26,60,14)

Table 3: Statistics of performance tensors for four tasks. *Sparisity* denotes the percentage of missing values in the tensor

ters; the settings of 3) the sentence-level encoders and 4) decoders are equal to NER.

We can also do bootstrap for predictions using regression models. If we consider recovering missing data with CP decomposition as a prediction method, we can construct a CI on the predicted values too.

C Hyper-parameters

For XGBoost, we use squared error as the objective function for regression and set the learning rate as 0.1. We allow the maximum tree depth to be 10, the number of trees to be 100, and use the default regularization terms to prevent the model from overfitting. For LGBM, we set the objective as regression for LGBMRegressor, the number of boosted trees and maximum tree leaves to be 100, adopt a learning rate of 0.1, and use the default regularization terms. For the Robust PCA model, we scale all the datasets, adopt the default regularization parameter of 1 for both the low rank and the sparse tensor, and set the learning rate as 1.1. For CP Decomposition, we do not standardize the features in CWS and NER, but do so for WMT and POS. We adopt a rank $r = 5$ in training and performance prediction, expressing the recovered tensor used for prediction to be a sum of 5 rank-1 tensors.

Statistics of Tensor Tab. 3, where *sparisity* denotes the percentage of missing values in the tensor.

²<https://www.clips.uantwerpen.be/conll2003/ner/>

³<https://catalog.ldc.upenn.edu/LDC2013T19>

⁴<http://sighan.cs.uchicago.edu/bakeoff2005/>

⁵<https://www.aclweb.org/mirror/ijcnlp08/sighan6/chinesebakeoff.htm>

⁶BERT is grouped into the subword encoder because we use it to obtain the representation of subwords.