Fine-tuning ERNIE for chest abnormal imaging signs extraction

Zhaoning Li, Jiangtao Ren*

School of Data and Computer Science, Guangdong Province Key Lab of Computational Science, Sun Yat-Sen University, Guangzhou, Guangdong 510006, PR China

ABSTRACT

Chest imaging reports describe the results of chest radiography procedures. Automatic extraction of abnormal imaging signs from chest imaging reports has a pivotal role in clinical research and a wide range of downstream medical tasks. However, there are few studies on information extraction from Chinese chest imaging reports. In this paper, we formulate chest abnormal imaging sign extraction as a sequence tagging and matching problem. On this basis, we propose a transferred abnormal imaging signs extractor with pretrained ERNIE as the backbone, named **EASON** (fine-tuning **E**RNIE with CRF for **A**bnormal **S**igns Extracti**ON**), which can address the problem of data insufficiency. In addition, to assign the attributes (the body part and degree) to corresponding abnormal imaging signs from the results of the sequence tagging model, we design a simple but effective **tag2relation** algorithm based on the nature of chest imaging report text. We evaluate our method on the corpus provided by a medical big data company, and the experimental results demonstrate that our method achieves significant and consistent improvement compared to other baselines.

1. Introduction

A large number of radiology reports have been accumulated for communication and documentation of diagnostic imaging since the wide use of medical information systems in China. In addition to the application of radiographs in medical image analysis [1,2], radiology reports also contain considerable meaningful knowledge to be discovered, and harnessing their potential requires efficient and automated information extraction [3]. For example, the automatic extraction of abnormal imaging signs in chest imaging reports is essential for clinical research and a wide range of downstream medical tasks: patient similarity measuring [4], diagnosis prediction [4] and automatic ICD coding [5]. However, most of the existing information extraction systems in radiology are developed for English [6–11], and too little work has been devoted to the extraction of abnormal imaging signs from Chinese chest imaging report text.

In this paper, we aim to extract structured information of abnormal imaging signs (i.e., abnormal imaging signs and their attributes "where did the abnormality occur" and "what is the degree of abnormality") from unstructured chest imaging reports. For example, "闭塞" (occluded) is an abnormal imaging sign; "右上肺" (right upper lung) and "支气管" (bronchus) are primary and secondary body parts where occlusion occurs, respectively; "部分" (partly) is the degree to which occlusion occurs, as shown in Fig. 1.

Specifically, we can divide this information extraction task into three subtasks:

1. Extracting abnormal imaging signs;
2. Extracting attributes of abnormal imaging signs;
3. Matching between abnormal imaging signs and their attributes.

To accurately and efficiently extract abnormal imaging signs and their attributes from chest imaging reports, we formulate subtasks (1) and (2) into a sequence tagging problem at the Chinese character level, which can avoid introducing errors caused by segmentation. Traditionally, researchers use machine learning methods [12–14] to perform sequence tagging tasks. Recently, with the development of deep learning, deep learning architectures based on long short-term memory networks (LSTMs) [15] or convolutional neural networks (CNNs) [16] combined with conditional random fields (CRFs) [17] have achieved state-of-the-art results in the clinical field [18–20]. In real medical situations, however, the cost of manually labeling a large training set in the medical field is too high and error-prone [21,11]. In the case of the insufficiency of high-quality annotated medical data, the data-hungry nature of deep learning limits the performance of these neural-based models. In this work, we adopt the advanced deep sequence tagging framework to address the following questions:

右上肺体积缩小，右上肺门见不规则团块状软组织密度影较前缩小，大小约3.9×2.6

（原大小约5.3cm×4.2cm），边界欠清，密度不均匀，右上肺支气管部分狭窄，部分

闭塞，右上肺见多发斑片状密影较前减少。

The **right upper lung** has a **loss of volume**, the **dense shadow of irregular mass-like soft tissue** in the **right upper hilum** is smaller than before, about 3.9 x 2.6 (original size was about 5.3 cm x 4.2 cm), not clear in the boundary and not homogeneous in the density. The **right upper lung bronchus** is **partly narrowed** and **partly occluded**. The **multiple patchy dense shadows** in the **right upper lung** have reduced than before.

t_1: {右上肺 **(right upper lung)**, 体积缩小 **(loss of volume)**}

t_2: {右上肺门 **(right upper hilum)**, 不规则团块状软组织密度影 **(dense shadow of irregular mass-like soft tissue)**}

t_3: {右上肺 **(right upper lung)**, 支气管 **(bronchus)**, 部分 **(partly)**, 狭窄 **(narrowed)**}

t_4: {右上肺 **(right upper lung)**, 支气管 **(bronchus)**, 部分 **(partly)**, 闭塞 **(occluded)**}

t_5: {右上肺 **(right upper lung)**, 多发 **(multiple)**, 斑片状密影 **(patchy dense shadows)**}

Fig. 1. A standard example of chest abnormal imaging sign extraction. In this case, t_i represents the i-th tuple in the above sentence; **cyan** denotes the abnormal imaging signs, **red** denotes the degree of the abnormal imaging sign, **olive green** and **lime green** denote the primary and secondary body part of the abnormal imaging sign, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

- How can data insufficiency be alleviated?
- How can attributes be assigned to abnormal imaging signs?

First, to alleviate the problem of data insufficiency, we propose fine-tuning ERNIE [22] (Enhanced Representation through Knowledge Integration) (Section 3.3) for our task, which is pretrained on a large corpus and has achieved ground-breaking performance across various Chinese natural language processing (NLP) tasks. Experimental results show that this transfer learning method can drastically improve the performance of our task.

Second, to assign the attributes to the corresponding abnormal imaging signs, i.e., subtask (3), we design a simple but effective **Tag2Relation** algorithm (Section 3.4) based on the nature of chest imaging report text, which can easily construct the relation between entities from the results of the sequence tagging model.

The contributions of this paper can be summarized as follows:

1. We propose **EASON** (fine-tuning **E**RNIE with CRF for **A**bnormal **S**igns Extracti**ON**), a transferred chest abnormal imaging signs extractor. To the best of our knowledge, we are the first to present such an effective method for the automatic extraction of abnormal imaging signs from chest imaging reports.
2. We design a novel tag2relation algorithm to establish the relation between abnormal imaging signs and their attributes, which can easily match the abnormal imaging sign with their attributes based on the result of the sequence tagging model.
3. We conduct extensive experiments on chest imaging reports provided by a medical big data company. Experimental results (Section 4.5) and further analysis (Section 5) show that our method achieves significant and consistent improvement compared to other baselines. We release the code and terminology to the research community for further research.[1]

2. Related work

With the integration and development of medicine and computer science technology, clinical information extraction is becoming increasingly important and attracting increasing attention. Many clinical NLP systems have been developed to extract structured information from unstructured electronic health records (EHRs). The research methods for clinical information extraction mainly include rule-based methods, machine learning-based methods, and deep learning-based methods.

Rule-based methods are the earliest attempt to extract information from EHRs [23,6,7,24–26,9]. For example, Friedman et al. [6] proposed MEDLEE (medical language extraction and encoding system) to extract information from textual patient reports with controlled vocabulary and grammatical rules. Johnson et al. [7] designed RADA (radiology analysis tool) to extract and structure key medical concepts and their attributes contained in radiology reports through predefined rules. Harkema et al. [26] proposed the ConText algorithm for determining whether clinical conditions mentioned in clinical reports are negated, hypothetical, historical, or experienced by someone other than the patient. ConText is based on the simple approach used by NegEx [27] (a regular expression algorithm) for finding negated conditions in text. These rule-based methods require formulating rules that consume significant time and effort, and their limited coverage and generalizability are the main drawbacks.

The traditional machine learning-based methods include handcrafted features for hidden Markov models (HMMs) [13,28], maximum entropy Markov models (MEMMs) [12,29], CRFs [14,30] and support vector machines (SVMs) [31,32]. In other related work targeting radiology reports, Esuli et al. [8] performed information extraction from free-text radiology reports with the CRF-based method. Hassanpour and Langlotz [10] used the conditional Markov model (CMM) and CRFs to extract radiological observations from reports. These machine learning-based methods heavily rely on a large number of feature engineering and thus require considerable human effort and time on feature engineering.

In recent years, deep learning has ushered in incredible advances in NLP tasks. Different from shallow machine learning methods, deep neural networks rely on powerful representation learning ability to automatically discover features, which significantly reduces feature engineering and saves human resources and time. For example, Gupta et al. [11] used an unsupervised model to extract relations and their associated entities from radiology reports using automated clustering of similar relations in narrative mammography radiology reports. The proposed approach based on distributional semantics (neural representation) and clustering to find similar relations outperforms other approaches. In addition, sequence tagging methods based on LSTMs or CNNs combined with a CRF [17] layer achieve state-of-the-art

---

[1] https://github.com/Das-Boot/eason.

performance in the clinical field and outperform traditional statistical methods [18–20]. However, the data-hungry nature of deep learning limits the performance of these neural-based methods for medical tasks with small datasets. The recent development of language representation models [33–35,22,36,37] trained on a large corpus demonstrate the possibility of transfer learning for sequence tagging.

3. Methods

3.1. Problem definition

To better illustrate our method, we introduce some terminologies as follows:

- **Abnormal imaging sign (*Abn*):** An entity refers to abnormal results of chest radiographs, CT, MR, etc.
- **Body part (*P*):** An entity refers to the specific organ or tissue structure where the abnormal imaging sign occurs, which serves as an attribute of abnormal imaging signs, including the primary body part (*PP*) and the secondary body part (*SP*).
- **Degree (*D*):** An entity refers to the scope (e.g., "弥漫" (diffusely)), severity (e.g., "轻度" (slightly)), frequency (e.g., "多发" (multiple)), and quantity (e.g., "单个" (single)) of abnormal imaging sign occurrence, which also serves as an attribute of abnormal imaging signs.[2]

Then, the structured information of an abnormal imaging sign can be formally defined as a quadruple: {*PP*, *SP*, *D*, *Abn*}, where *PP*, *SP*, *D* are attributes of the corresponding *Abn*; note that the attribute may be null, as shown in Fig. 1.

Thus, in this work, our task is to extract all the quadruples in a given chest imaging report, which includes abnormal imaging sign *Abn* identification, attribute *PP/SP/D* identification, and matching between *Abn* and *PP/SP/D*.

3.2. Tagging scheme

We use the "BIO" (begin, inside, other) and "Abn, P, D" signs to represent the position information and the semantic roles of the Chinese characters, respectively. Note that we only label the *P* or *D* serving as an attribute of *Abn*. Both the primary and secondary body parts are marked as different body parts; for example, "右上肺支气管" (right upper lung bronchi) is labeled as "右上肺" (right upper lung) and "支气管" (the bronchi), respectively.

Fig. 2 shows an example of such a tagging scheme for the sentence "右上肺见多发斑片状密影较前减少。" (The multiple patchy dense shadows in the right upper lung have reduced than before). Based on our tagging scheme, we can label the abnormal imaging sign: "斑片状密影" (patchy dense shadows), body part: "右上肺" (right upper lung) and degree: "多发" (multiple) separately with our unique tags. Specifically, tag "O" represents the "other", which means that the corresponding character is irrelevant in any entity components. Tag "B-P" represents the "body part begin", tag "I-P" represents the "body part inside", tag "B-D" represents the "degree begin", tag "I-D" represents the "degree inside", tag "B-Abn" represents the "abnormal imaging sign begin" and tag "I-Abn" represents the "abnormal imaging sign inside".

3.3. Extracting abnormal imaging signs and attributes with EASON

Fig. 3 shows the main structure of our EASON model. We take the input sentence $S = \{x_i\}_{i=1}^n$ and its corresponding label sequence $y = \{y_i\}_{i=1}^n$ as an example to introduce each component of EASON from bottom to top as follows, where n is the length of the S.

3.3.1. Encoding sentences with ERNIE

It is difficult to train a superior deep learning model without any prior knowledge in the case of data insufficiency. In this paper, we use transfer learning to alleviate the problem of data insufficiency. Specifically, we propose to fine-tune ERNIE [22] with CRF for our task, where ERNIE is a novel language representation model based on a multilayer transformer [38] and a multistage knowledge masking strategy.

To encode sentence S with ERNIE, we first construct input representations $\boldsymbol{H}_0$ by summing the corresponding token embeddings ($\boldsymbol{E}_{token}$), segment embeddings ($\boldsymbol{E}_{seg}$), and position embeddings ($\boldsymbol{E}_{pos}$) as follows:

$$\boldsymbol{H}_0 = \boldsymbol{E}_{token} + \boldsymbol{E}_{seg} + \boldsymbol{E}_{pos} \tag{1}$$

Then, N transformer layers are applied to calculate the context-dependent representations:

$$\boldsymbol{H}_\alpha = Transformer(\boldsymbol{H}_{\alpha-1}), \alpha \in [1, N], \tag{2}$$

where $\boldsymbol{H}_\alpha$ represents the hidden representations of the input sentence at the α-th layer. Note that in addition to fine-tuning, there is another paradigm for transfer learning: feature extraction. In feature extraction, the weights of ERNIE are "frozen," and the pretrained representations (can also be obtained from the above equation) are used in a downstream model. Section 5.2 shows the relative performance of fine-tuning vs. feature extraction.

Finally, we take the final hidden representations $\boldsymbol{H}_N \in \mathbb{R}^{n \times d}$,[3] and then project $\boldsymbol{H}_N$ with a linear projection to matrix $\boldsymbol{H}_N \boldsymbol{W}$, where d is the size of the transformer layer, weight matrix $\boldsymbol{W} \in \mathbb{R}^{d \times k}$ is the parameter of the model to be learned in training, and k is the number of distinct tags.

3.3.2. Conditional random field

The conditional random field (CRF) [17] can obtain a globally optimal chain of labels for a given sequence considering the correlations between adjacent tags. In a sequence tagging task, there are usually strong dependencies between the output labels. Therefore, instead of only fine-tuning ERNIE to model tagging decisions separately, we stack the CRF layer on top of the ERNIE outputs to jointly decode labels for the whole sentence.

We use $\boldsymbol{P} \in \mathbb{R}^{n \times k}$ as the matrix of scores output by the linear layer, where $\boldsymbol{P}_{ij}$ represents the score of the j^{th} label of the i^{th} character within a sentence. For the sentence $S = \{x_i\}_{i=1}^n$ along with a path of tags $y = \{y_i\}_{i=1}^n$, CRF obtains a real-valued score as follows

$$score(S, y) = \sum_{i=0}^n \boldsymbol{A}_{y_i, y_{i+1}} + \sum_{i=1}^n \boldsymbol{P}_{i, y_i}, \tag{3}$$

where $\boldsymbol{A}$ is the transition matrix, and $\boldsymbol{A}_{i,j}$ denotes the score of a transition from tag i to tag j. y_0 and y_n are the special tags at the beginning and the end of a sentence, so $\boldsymbol{A}$ is a square matrix of size $k + 2$. Therefore, the probability for the label sequence y given a sentence S is:

$$p(y|S) = \frac{e^{score(S,y)}}{\sum_{\tilde{y} \in Y_S} e^{score(S,\tilde{y})}} \tag{4}$$

We now maximize the log-likelihood of the correct tag sequence:

---

[2] We only focus on the degree entities in the description of the current chest imaging report. We do not care about the degree entities related to historical condition, so we do not annotate "较前减少" (reduced than before) as degree in the sentence "右上肺见多发斑片状密影较前减少。" (The multiple patchy dense shadows in the right upper lung have reduced than before).

[3] Since we focus on intrasentence sequence tagging in this work, we ignore the special classification token CLS and separation token SEP.

Character Sentence: 右 上 肺 见 多 发 斑 片 状 密 影 较 前 减 少 。

Tag Sequence: B-P I-P I-P O B-D I-D B-Abn I-Abn I-Abn I-Abn I-Abn O O O O O

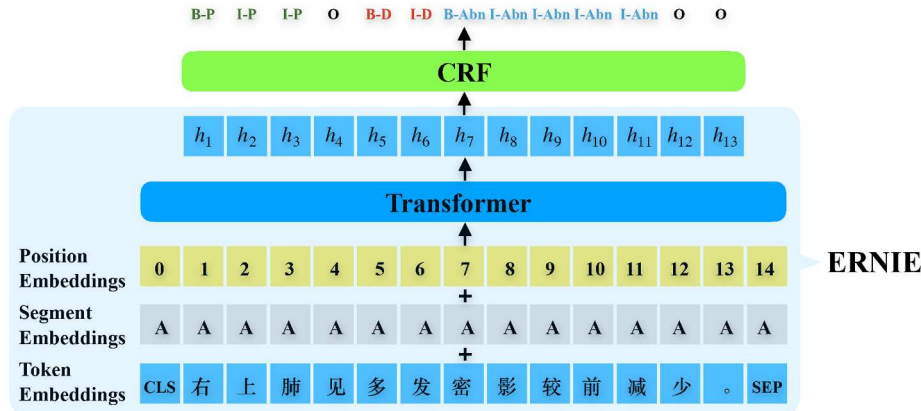**Fig. 2.** A standard annotation for the example sentence based on our tagging scheme.

Fig. 3. The overview architecture of the EASON model.

[右上肺体积缩小，]$_{C_1}$ [右上肺门见不规则团块状软组织密度影较前缩小，大小约

3.9×2.6（原大小约5.3cm×4.2cm），边界欠清，密度不均匀，]$_{C_2}$ [右上肺支气管部分

狭窄，部分闭塞，]$_{C_3}$ [右上肺见多发斑片状密影较前减少。]$_{C_4}$

C_1: {右上肺 (right upper lung), **P2Abn**, 体积缩小 (loss of volume)}

C_2: {右上肺门 (right upper hilum), **P2Abn**, 不规则团块状软组织密度影
   (dense shadow of irregular mass-like soft tissue)}

C_3: {右上肺 (right upper lung), **P2P**, 支气管 (bronchus)} {支气管 (bronchus), **P2Abn**, 狭窄 (narrowed)}
   {支气管 (bronchus), **P2Abn**, 闭塞 (occluded)} {部分 (partly), **D2Abn**, 狭窄 (narrowed)}
   {部分 (partly), **D2Abn**, 闭塞 (occluded)}

C_4: {右上肺 (right upper lung), **P2Abn**, 斑片状密影 (patchy dense shadows)}
   {多发 (multiple), **D2Abn**, 斑片状密影 (patchy dense shadows)}

Fig. 4. The primary body parts divide the example sentence into four chunks. In this case, C_i represents the i-th chunk in the above sentence; **cyan** denotes the abnormal imaging signs, **red** denotes the degree of the abnormal imaging sign, **olive green** and **lime green** denote the primary and secondary body part of the abnormal imaging sign, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

$$\log(p(y|S)) = score(S, y) - \log\left(\sum_{\widetilde{y} \in Y_S} e^{score(S,\widetilde{y})}\right), \qquad (5)$$

where Y_S represents all possible tag sequences for an input sentence S. From the formulation above, we can obtain a valid output sequence. When decoding, the sequence with the maximum score is output by:

$$y^* = \underset{\widetilde{y} \in Y_S}{arg\max}\, score(S, \widetilde{y}) \qquad (6)$$

In general, we can use the Viterbi algorithm [39] to decode the optimal label sequence. Note that the CRF layer is jointly fine-tuned with ERNIE.

3.4. Matching between abnormal imaging signs and attributes: tag2relation algorithm

After extracting abnormal imaging signs and attributes, we design a simple but effective matching algorithm **tag2relation** to assign attributes to the corresponding abnormal imaging signs automatically. Based on the nature of chest imaging text, we find that the secondary body parts of the reports are enumerable, so we asked professional medical practitioners to develop a dictionary of secondary body parts, which was constructed according to the information of secondary body parts in chest imaging reports as well as some medical literature such as 《医学影像学诊断图谱和报告（第一版）》 (*Medical Imaging Diagnostic Atlas and Reports [First Edition]*). To better illustrate this algorithm, we define the semantic unit chunk as follows:

- **Chunk:** A chunk refers to the textual content between two primary body parts in the sentence.

Furthermore, the relations between attributes and corresponding abnormal imaging signs can be defined as follows:

- **P2Abn:** A relation "P2Abn" indicates that a primary or secondary body part serves as an attribute of a corresponding abnormal imaging sign
- **D2Abn:** A relation "D2Abn" indicates that a degree serves as an attribute of a corresponding abnormal imaging sign
- **P2P:** A relation "P2P" indicates that a secondary body part is a subdivision of a primary body part.

Algorithm 1. Tag2triplet($\mathscr{S}$, $\mathscr{Y}$, $\mathscr{D}$)

> **input** : A tag sequence $\mathcal{Y}$ corresponding to sentence $\mathcal{S}$, the dictionary $\mathcal{D}$ of secondary body parts
>
> **output:** The relations set $\mathcal{T}$ in sentence $\mathcal{S}$
>
> 1   *find the primary body parts $\mathcal{P}$ in $\mathcal{S}$ with $\mathcal{Y}$ and $\mathcal{D}$;*
>
> 2   *find the chunk set $\mathcal{C}$ in $\mathcal{S}$ with $\mathcal{P}$;*
>
> 3 $\mathcal{T} \leftarrow \emptyset$;
>
> 4   **for** $C \in \mathcal{C}$ **do**
>
> 5   *find entities E in C with $\mathcal{Y}$;*
>
> 6 $R \leftarrow$ *perform cartesian products over E;*
>
> 7 $\mathcal{T} \leftarrow \mathcal{T} \cup$ Filter$(R, \mathcal{S}, \mathcal{Y})$
>
> 8   **end**
>
> 9   **return** $\mathcal{T}$

The tag2relation algorithm is described in Algorithm 1. We elaborate on this algorithm by taking the sentence *S* in Fig. 4 as an example (the English translation of example sentence is shown in Fig. 1).

First, we use the predefined dictionary to identify the primary body parts (i.e., all the body parts that are not in the dictionary) in the example sentence and then identify each chunk through these parts. As shown in Fig. 4, we can divide the example sentence into four chunks. In each chunk, we apply a cartesian product over the entity tags to obtain the candidates of the relation. Finally, we **filter** the candidates by selecting the relations with the shortest distance between attribute and *Abn* to obtain the final matching results.

4. Experiments

4.1. Dataset

The experimental dataset consists of chest imaging reports provided by a medical big data company, which is a Chinese high-tech enterprise focusing on the construction of a big data management cloud platform for respiratory disease. We asked two annotators with the medical background to manually annotate abnormal imaging signs and corresponding attributes in reports, and the disagreements between two annotators were resolved by a senior medical practitioner. Specifically, our annotation task consists of two subtasks: (1) entity annotation: choosing nonoverlapping entity spans and (2) semantic relation annotation: building a directed graph on top of the entity spans. Based on the annotation results of the two annotators (i.e., A and B), we use F1-score (F1) for consistency evaluation to ensure the quality of the data annotation. The F1 value can be calculated by the following formulas:

$$P = \frac{\#\text{identical samples}}{\#\text{total samples labeled by A}}, \tag{7}$$

$$R = \frac{\#\text{identical samples}}{\#\text{total samples labeled by B}}, \tag{8}$$

$$F_1 = 2\frac{P \cdot R}{P + R}, \tag{9}$$

The annotation took place in two preannotation rounds and one official annotation round. The purpose of preannotation is to let the annotators

Table 1
Annotation statistics.

Annotation Process	Entity		Relation	
	Total	*F1*	Total	*F1*
First-Round	1398	66.25	1079	60.09
Second-Round	1199	85.63	885	81.84
Official Round	3831	93.35	3004	88.01

Table 2
Statistics of abnormal imaging signs and corresponding attributes.

Entity Type	Training Set	Test Set
Abnormal Imaging Sign	2362	428
Body Part	2154	396
Degree	926	162
Sum	5442	986

Table 3
Statistics of different types of relations.

Relation Type	Training Set	Test Set
P2Abn	2721	466
D2Abn	1126	195
P2P	388	72
Sum	4235	733

fully understand and familiarize themselves with the annotation guidelines, while the senior medical practitioner will further refine the guidelines based on disagreements.[4] Table 1 shows the total number of entities, semantic relations, and consistency F1 values for each round.

---

[4] The annotation guidelines are available at https://github.com/Das-Boot/eason.

Table 4
Comparative results of our EASON model and baseline models on the test set.

Model	*Abn* Identification			Attributes Identification			Matching		
	P	*R*	*F1*	*P*	*R*	*F1*	*P*	*R*	*F1*
IDCNN	93.93	93.93	93.93	90.58	89.61	90.09	83.36	84.38	83.87
RDCNN-CRF	93.36	95.33	94.34	91.70	91.04	91.37	85.14	86.30	85.71
IDCNN-CRF	95.08	94.86	94.97	92.55	91.22	91.88	86.30	86.30	86.30
BiLSTM	94.12	93.46	93.79	91.97	90.32	91.14	84.75	84.52	84.64
BiLSTM-CRF	93.72	94.16	93.94	93.65	92.47	93.06	86.10	87.40	86.74
BERT	94.23	95.33	94.77	93.15	92.65	92.90	84.51	88.22	86.33
BERT-wwm	94.21	95.09	94.65	92.83	92.83	92.83	84.94	88.08	86.48
ERNIE	95.29	94.63	94.96	94.04	93.37	93.71	85.64	87.40	86.51
EASON	**96.46**	**95.56**	**96.01**	**94.24**	**93.91**	**94.08**	**87.89**	**89.45**	**88.66**

The F1 values for entities and semantic relation annotation in the second annotation round are 85.63 and 81.84, respectively. When the F1 value is greater than 80%, we believe that the annotators are already familiar with the annotation guidelines [40]. Then we started the official annotation round. After annotation, we randomly divided the dataset into the training set (253 reports, 2596 sentences) and test set (45 reports, 458 sentences) according to the ratio 0.85:0.15. Table 2 shows the statistics of abnormal imaging signs and their attributes, and Table 3 shows the statistics of relations between abnormal imaging signs and corresponding attributes.

4.2. Evaluation metrics

The standard and widely used performance measures [41,42], such as precision (P), recall (R) and F1, are used as evaluation metrics in the following experiments, which can be calculated by the following formulas:

$$P = \frac{\#correct\ predicted\ samples}{\#predicted\ samples}, \tag{10}$$

$$R = \frac{\#correct\ predicted\ samples}{\#total\ samples\ in\ D}, \tag{11}$$

$$F_1 = 2\frac{P \cdot R}{P + R}, \tag{12}$$

where D is the set of all the sentences in the dataset. In the tasks of abnormal imaging sign (*Abn*) identification and attribute identification (*P* and *D*), a predicted sample is regarded as correct if and only if it precisely matches an annotated entity. In the task of matching between *Abn* and attributes, a predicted sample is considered to be correct when its relation type and two corresponding entities are both correct.

4.3. Hyperparameters

The model was implemented by using Keras[5] version 2.2.4 and the "ERNIE 1.0 Base for Chinese"[6] version of ERNIE in which it uses 12 transformer encoder layers, 768 hidden units, and 12 attention heads. The optimization method of the fine-tuning process was Adam [43] with $\beta_1 = 0.9$, $\beta_2 = 0.999$. The learning rate reached 5e-5 in the first epoch, decayed to 1e-5 in the second epoch, and maintained this learning rate until the end of the training. We let the mini-batch size be 16. In the experiments, we performed the grid search and 10-fold cross-validation on the training set to find the optimal hyperparameters. On the test set, we selected the optimal model among all 200 epochs with

the highest validation F1-score.

4.4. Baselines

For a comprehensive comparison, we compare our method against several classical sequence tagging models, which can be divided into two categories: CNN-based models and BiLSTM-based models.

For the CNN-based models, the baselines are as follows:

- **IDCNN** [44]: This model uses a deep iterated dilated CNN (IDCNN) architecture to aggregate context from the entire text, which has better capacity than traditional CNN and faster computation speed than LSTMs and then maps the output of IDCNNs to predict each label independently through a softmax classifier.
- **IDCNN-CRF** [44]: This model uses CRF to maximize the label probability of the complete sentence based on IDCNNs. Compared to the softmax classifier, the CRF classifier is more appropriate for tasks with strong output label dependency.
- **RDCNN-CRF** [20]: This model is an extension of IDCNN-CRF that uses residual connection [45] between IDCNN layers (RDCNNs) to ease the training of networks and then sums the output of standard CNNs and RDCNNs as the input of CRF.

The baselines for the BiLSTM-based models are listed as follows:

- **BiLSTM** [46]: The model consists of two parts: a BiLSTM encoder and a softmax classifier.
- **BiLSTM-CRF** [47]: A classic and popular choice for sequence tagging tasks, which consists of a BiLSTM encoder and a CRF classifier.

To further analyze the performance of the fine-tuning pretrained Chinese language representation model on our task, we also fine-tune several advanced pretrained models as the experimental baselines:

- **BERT** [35]: BERT (bidirectional encoder representations from the transformer) is the first language representation model based on the bidirectional transformer and masking strategy, and it has shown marvelous improvements across various NLP tasks.
- **BERT-wwm** [36]: Bert-wwm is an upgraded the version of BERT in which they adapted the whole word masking (WWM) strategy in Chinese text for the language model pretraining task.
- **ERNIE** [22]: ERNIE is an upgraded version of BERT in which they use phrase-level and entity-level masking strategy in addition to basic masking strategy. Note that ERNIE was trained on not only Chinese Wikipedia data but also Baidu Baike (similar to Wikipedia), Baidu news and Baidu Tieba (similar to Reddit). The numbers of sentences are 21 M, 51 M, 47 M, 54 M, respectively.
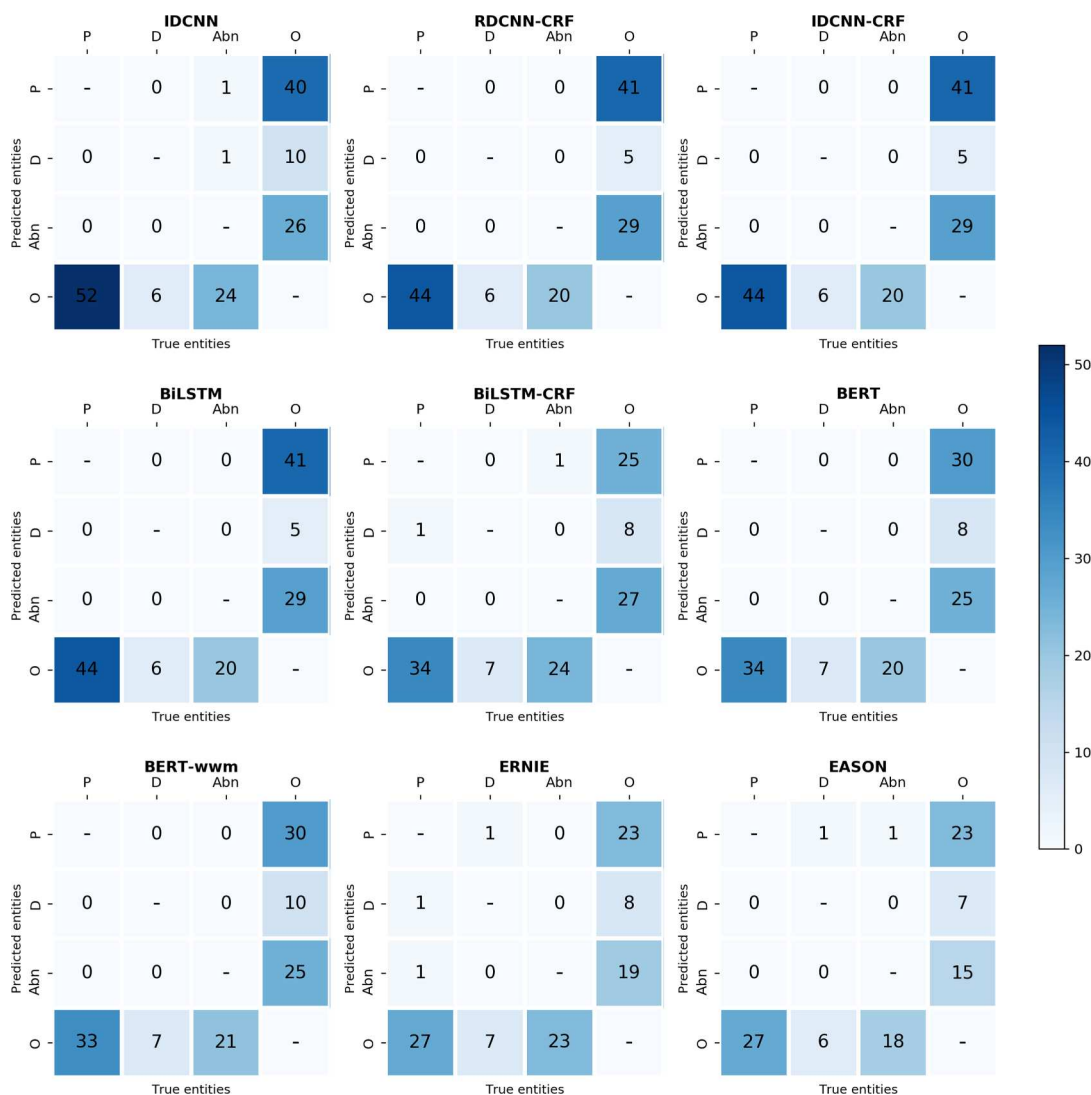
---

[5] https://github.com/keras-team/keras.
[6] https://github.com/PaddlePaddle/ERNIE.

Fig. 5. The confusion matrix of our EASON model and other baseline models for entity errors. *X*-Axis: true entities; *y*-axis: predicted entities; P, D, Abn represent body part, degree, abnormal imaging sign entities, respectively, and O is the entities unrelated to the task (i.e., the corresponding character of O is irrelevant in P, D, and Abn).

4.5. Experimental Results

As mentioned in Section 3.1, our task includes abnormal imaging sign (*Abn*) identification, attribute (*P* and *D*) identification, and matching between *Abn* and attributes. Thus, in this work, we compare our model with baselines in these three subtasks, as shown in Table 4.

First, we observe that our EASON model outperforms all other models with 96.01% in *Abn* identification, 94.08% in attributes identification, and 88.66% in matching in terms of F1-score. This demonstrates the effectiveness of our proposed method.

Second, compared with superiors of the classical sequence tagging models in F1-score, we can see EASON achieves an improvement of 1.04 points (compared with IDCNN-CRF) in *Abn* identification, 1.02 points (compared with BiLSTM-CRF) in attributes identification, and 1.92 points (compared with BiLSTM-CRF) in matching, which verifies our assumption that the current annotated dataset is not large enough to train a deep learning model sufficiently. With the help of transferred prior knowledge from the pretrained language representation model, we can obtain better performance in all three subtasks.

Moreover, it shows that the performance of the ERNIE (the output layer is a softmax classifier) improved after jointly fine-tuning with the CRF. The benefits in F1-score brought by jointly fine-tuning with CRF

are 1.05, 0.37, 2.15 points in all subtasks because CRF models use the whole label sequence instead of independent label classification and thus can avoid some invalid label sequences.

5. Analysis and discussion

5.1. Error analysis

5.1.1. Comparison on confusion matrix

In this paper, we focus on extracting all the quadruples (Section 3.1) from chest imaging reports, which includes abnormal imaging sign identification, attribute identification, and matching between them. Accurate identification of *Abn*, *P*, and *D* plays a vital role in our task. To visually compare how many errors each model makes at the entity level, we present a confusion matrix for entities *Abn*, *P*, and *D* shown in Fig. 5. We can see that our model EASON (shown in the lower right corner of Fig. 5) can better identify *Abn* and *P* compared with other baselines. Next, we elaborate on the errors produced by EASON.

5.1.2. Error analysis of EASON

To perform error analysis of EASON, we divide the errors under the entity level strict evaluation into four categories according to [48,49]:

Table 5
Statistics of different types of errors produced by EASON.

	Count	% of Errors
TYPE	2	2.5%
EXTENT	30	37.5%
SPURIOUS	20	25.0%
MISSING	28	35.0%

Table 6
The distribution of **type error**, missing error, and spurious error.

Gold Standard	EASON Output				
	P	*D*	*Abn*	Missing	Total
P	369			9 (2.4%)	378
D	**1**	155		5 (3.1%)	161
Abn	**1**		409	14 (3.3%)	424
Spurious	9 (2.4%)	3 (1.9%)	8 (1.9%)		20
Total	380	158	417	28	983

- **Type error**: The entity was identified by a correct span but the incorrect label type.
- **Extent error**: The span of the entity overlapped with that of a gold-standard entity but did not match it exactly.
- **Spurious error**: The span of the entity had no overlap with any gold-standard entity.
- **Missing error**: The span of the entity in the gold standard had no overlap with that of any entity in the system output.

Table 5 shows the statistics of all errors produced by EASON. Next, we analyze each type of error in detail. Table 6 shows the distribution of type, missing, and spurious errors produced by EASON. Correctly recognized entities can be seen reading down the diagonal in Table 6. Type errors appear in **bold**. Type errors are rare, with only two occurrences, and they are all mistagged as body parts. For example, in the sentence "食管全程扩张，局部较前增著" (The esophagus expands throughout the area, increasing locally), the gold standard has an instance of "全程" (throughout the area) tagged as degree, while EASON output is body part.

Spurious errors appear in blue. From the row of spurious errors, we see that spurious errors occur most often for body part (2.4% or 9 out of 380) because ordinary body parts had lexical or semantic features similar to those of body parts. However, body parts serve as attributes to the abnormal imaging signs, and ordinary body parts do not. For example, in the sentence "两肺膨胀良好" (Both lungs are well distended), EASON identifies "两肺" (both lungs) as body part, though it is not. Missing errors appear in red and occur most often for the abnormal imaging signal (14 out of 424 tags or 3.3% missed). For example, EASON does not identify "食糜及液体潴留" (chyme and fluid retention) and "胃腔突破食管裂孔" (gastric cavity breaks through esophageal hiatus) as an abnormal imaging signal. The reason may be that EASON requires slightly more training data to learn these less common abnormal imaging signals.

Table 7 shows statistics for three types of extent errors. Numbers in the short row indicate instances where the span of the entity produced

Table 7
Statistics of extent error in EASON output.

	P	*D*	*Abn*	Total
SHORT	5	4	7	16
LONG	12	0	0	12
S&L	2	0	0	2
Total	19	4	7	30

Table 8
Test set performance of our EASON model and feature-based models.

Model	F1		
	Abn	Attributes	Matching
First Layer (Embeddings)	94.55	92.46	86.39
Second-to-Last Hidden	94.68	**94.28**	86.84
Last Hidden	94.83	94.03	87.14
Sum Last Four Hidden	94.50	94.07	87.48
Concat Last Four Hidden	94.91	93.48	85.98
Sum All 12 Layers	95.25	93.92	87.02
EASON	**96.01**	94.08	**88.66**

by EASON fell within that of a gold standard. In contrast, numbers in the long row indicate instances where the span of the entity covered that of a gold standard. Moreover, numbers in the short and long row indicate the span of the entity neither fell into nor covered that of a gold standard. Nearly two-thirds of the extent errors occurred on body part, with many instances of long errors. The occurrence of extent error indicates that EASON sometimes cannot detect the boundaries of entities correctly. There are two main reasons for these errors. The first cause is tokenization. For example, the gold standard has an instance of "肝" (liver), while EASON tagged "肝、" as body part (long extent error) because of the failure of tokenizing "、" from "肝." Another cause is that EASON tags one entity as two (short extent error) or tags two entities as one (long extent error). For example, the body part "食管下端贲门区" (the lower esophagus from the cardia) identified by EASON is, in fact, two entities in the gold standard: "食管下端" (lower esophagus) and "贲门区" (cardia area).

5.2. Feature extraction or fine-tuning?

To further investigate the performance of feature extraction (where the weights of the pretrained model are frozen) and fine-tuning the pretrained model in our task, we compare our model with feature-based models that use one or more layers of ERNIE as input to a one-layer 512-dimensional BiLSTM before the CRF layer. The results are illustrated in Table 8. From the table, we see that EASON achieves No. 1 in *Abn* identification and matching and No. 2 in attributes identification in terms of F1-score. This indicates that our proposed fine-tuning model is more suitable for the task of chest abnormal imaging sign extraction than feature-based models.

6. Conclusion

In this paper, we formulate chest abnormal imaging sign extraction as a sequence tagging and matching problem and deliver an effective solution for this task. In particular, we propose EASON to extract abnormal imaging signs and their attributes in Chinese chest imaging reports. To alleviate the problem of data insufficiency, we fine-tune ERNIE trained from the large corpus with CRF to perform sequence tagging in our task. In addition, we design a tag2relation algorithm to assign the attributes to corresponding abnormal imaging signs from the results of the sequence tagging model. Experimental results on the corpus provided by a medical big data company show that our proposed EASON model achieves superior performance compared to other baseline models, i.e., reaching the F1-score of 96.01%, 94.08%, 88.66% in *Abn* identification, attributes identification and matching, respectively.

Declaration of Competing Interest

None.

References

[1] G.J.S. Litjens, T. Kooi, B.E. Bejnordi, A.A.A. Setio, F. Ciompi, M. Ghafoorian, J.A.W.M. van der Laak, B. van Ginneken, C.I. Sánchez, A survey on deep learning in medical image analysis, Med. Image Anal. 42 (2017) 60–88.

[2] A.S. Lundervold, A. Lundervold, An overview of deep learning in medical imaging focusing on MRI, Z Med. Phys. 29 (2) (2019) 102–127.

[3] E. Pons, L.M.M. Braun, M.G.M. Hunink, J.A. Kors, Natural language processing in radiology: a systematic review, Radiology 279 (2) (2016) 329–343.

[4] J. Ni, J. Liu, C. Zhang, D. Ye, Z. Ma, Fine-grained patient similarity measuring using deep metric learning, Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM, 2017, pp. 1189–1198.

[5] J. Mullenbach, S. Wiegreffe, J. Duke, J. Sun, J. Eisenstein, Explainable prediction of medical codes from clinical text, in: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, 1101–1111, 2018.

[6] C. Friedman, G. Hripcsak, W. DuMouchel, S.B. Johnson, P.D. Clayton, Natural language processing in an operational clinical information system, Nat. Lang. Eng. 1 (1) (1995) 83–108.

[7] D.B. Johnson, R.K. Taira, A.F. Cardenas, D.R. Aberle, Extracting information from free text radiology reports, Int. J. Digit. Libr. 1 (3) (1997) 297–308.

[8] A. Esuli, D. Marcheggiani, F. Sebastiani, An enhanced CRFs-based system for information extraction from radiology reports, J. Biomed. Inform. 46 (3) (2013) 425–435.

[9] S. Bozkurt, J.A. Lipson, U. Senol, D.L. Rubin, Automatic abstraction of imaging observations with their characteristics from mammography reports, JAMIA 22 (e1) (2015) e81–e92.

[10] S. Hassanpour, C.P. Langlotz, Information extraction from multi-institutional radiology reports, Artif. Intell. Med. 66 (2016) 29–39.

[11] A. Gupta, I. Banerjee, D.L. Rubin, Automatic information extraction from unstructured mammography reports using distributed semantics, J. Biomed. Inform. 78 (2018) 78–86.

[12] A. McCallum, D. Freitag, F.C.N. Pereira, Maximum entropy Markov models for information extraction and segmentation, in: Proceedings of the Seventeenth International Conference on Machine Learning, ICML, 591–598, 2000.

[13] G. Zhou, J. Su, Named entity recognition using an HMM-based chunk tagger, in: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, ACL, 473–480, 2002.

[14] A. McCallum, W. Li, Early results for named entity recognition with conditional random fields, feature induction and web-enhanced lexicons, in: Proceedings of the Seventh Conference on Natural Language Learning, CoNLL, 188–191, 2003.

[15] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural Comput. 9 (8) (1997) 1735–1780.

[16] Y. LeCun, B.E. Boser, J.S. Denker, D. Henderson, R.E. Howard, W.E. Hubbard, L.D. Jackel, Backpropagation applied to handwritten zip code recognition, Neural Comput. 1 (4) (1989) 541–551.

[17] J.D. Lafferty, A. McCallum, F.C.N. Pereira, Conditional random fields: probabilistic models for segmenting and labeling sequence data, in: Proceedings of the Eighteenth International Conference on Machine Learning, ICML, 282–289, 2001.

[18] M. Habibi, L. Weber, M.L. Neves, D.L. Wiegandt, U. Leser, Deep learning with word embeddings improves biomedical named entity recognition, Bioinformatics 33 (14) (2017) i37–i48.

[19] Q. Wang, Y. Zhou, T. Ruan, D. Gao, Y. Xia, P. He, Incorporating dictionaries into deep neural networks for the Chinese clinical named entity recognition, Journal of Biomedical Informatics 92.

[20] J. Qiu, Q. Wang, Y. Zhou, T. Ruan, J. Gao, Fast and accurate recognition of Chinese clinical named entities with residual dilated convolutions, in: IEEE International Conference on Bioinformatics and Biomedicine, BIBM, 2018, pp. 935–942.

[21] S. Zheng, F. Wang, H. Bao, Y. Hao, P. Zhou, B. Xu, Joint Extraction of Entities and Relations Based on a Novel Tagging Scheme, Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL, 2017, pp. 1227–1236.

[22] Y. Sun, S. Wang, Y. Li, S. Feng, X. Chen, H. Zhang, X. Tian, D. Zhu, H. Tian, H. Wu, ERNIE: Enhanced representation through knowledge integration, arXiv preprint arXiv:1904.09223.

[23] C. Friedman, P.O. Alderson, J.H.M. Austin, J.J. Cimino, S.B. Johnson, Research paper: a general natural-language text processor for clinical radiology, JAMIA 1 (2) (1994) 161–174.

[24] Q.T. Zeng, S. Goryachev, S.T. Weiss, M. Sordo, S.N. Murphy, R. Lazarus, Extracting principal diagnosis, co-morbidity and smoking status for asthma research: evaluation of a natural language processing system, BMC Med. Inf. & Decisi. Making 6 (2006) 30.

[25] A. Coden, G.K. Savova, I.L. Sominsky, M.A. Tanenblatt, J.J. Masanz, K. Schuler, J.W. Cooper, W. Guan, P.C. de Groen, Automatically extracting cancer disease characteristics from pathology reports into a Disease Knowledge Representation Model, J. Biomed. Inform. 42 (5) (2009) 937–949.

[26] H. Harkema, J.N. Dowling, T. Thornblade, W.W. Chapman, ConText: An algorithm for determining negation, experiencer, and temporal status from clinical reports, J. Biomed. Inform. 42 (5) (2009) 839–851.

[27] W.W. Chapman, W. Bridewell, P. Hanbury, G.F. Cooper, B.G. Buchanan, A Simple Algorithm for Identifying Negated Findings and Diseases in Discharge Summaries, J. Biomed. Inform. 34 (5) (2001) 301–310.

[28] M. Song, H. Yu, W. Han, Developing a hybrid dictionary-based bio-entity recognition technique, BMC Med. Inf. & Decision Making 15 (S-1) (2015) S9.

[29] J.R. Finkel, S. Dingare, H. Nguyen, M. Nissim, C.D. Manning, G. Sinclair, Exploiting context for biomedical entity recognition: from syntax to the web, Proceedings of the International Joint Workshop on Natural Language Processing in Biomedicine and its Applications, NLPBA/BioNLP, 2004.

[30] M. Skeppstedt, M. Kvist, G.H. Nilsson, H. Dalianis, Automatic recognition of disorders, findings, pharmaceuticals and body structures from clinical text: An annotation and machine learning study, J. Biomed. Inform. 49 (2014) 148–158.

[31] Y. Wu, T. Fan, Y. Lee, S. Yen, Extracting named entities using support vector machines, in: Knowledge Discovery in Life Science Literature, PAKDD 2006 International Workshop, KDLL, 91–103, 2006.

[32] Z. Ju, J. Wang, F. Zhu, Named entity recognition from biomedical text using SVM, in: International Conference on Bioinformatics and Biomedical Engineering, 1–4, 2011.

[33] M.E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, L. Zettlemoyer, Deep contextualized word representations, in: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, 2018, pp. 2227–2237.

[34] A. Akbik, D. Blythe, R. Vollgraf, Contextual string embeddings for sequence labeling, Proceedings of the 27th International Conference on Computational Linguistics, COLING, 2018, pp. 1638–1649.

[35] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, 4171–4186, 2019.

[36] Y. Cui, W. Che, T. Liu, B. Qin, Z. Yang, S. Wang, G. Hu, Pre-training with whole word masking for Chinese BERT, arXiv preprint arXiv:1906.08101.

[37] Y. Sun, S. Wang, Y. Li, S. Feng, H. Tian, H. Wu, H. Wang, Ernie 2.0: A continual pre-training framework for language understanding, arXiv preprint arXiv:1907.12412.

[38] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, in: Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems, NIPS, 5998–6008, 2017.

[39] A.J. Viterbi, Error bounds for convolutional codes and an asymptotically optimum decoding algorithm, IEEE Trans. Inform. Theory 13 (2) (1967) 260–269.

[40] R. Artstein, M. Poesio, Inter-coder agreement for computational linguistics, Comput. Linguist. 34 (4) (2008) 555–596.

[41] Y. Liu, Y. Zhou, S. Wen, C. Tang, A strategy on selecting performance metrics for classifier evaluation, IJMCMC 6 (4) (2014) 20–35.

[42] Y. Zhou, Y. Liu, Correlation analysis of performance metrics for classifier, in: Decision Making and Soft Computing: Proceedings of the 11th International FLINS Conference, World Scientific, 2014, pp. 487–492.

[43] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: 3rd International Conference on Learning Representations, ICLR, 2015.

[44] E. Strubell, P. Verga, D. Belanger, A. McCallum, Fast and accurate entity recognition with iterated dilated convolutions, Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP, 2017, pp. 2670–2680.

[45] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2016, pp. 770–778.

[46] P. Wang, Y. Qian, F.K. Soong, L. He, H. Zhao, Part-of-speech tagging with bidirectional long short-term memory recurrent neural network, arXiv preprint arXiv:1510.06168.

[47] Z. Huang, W. Xu, K. Yu, Bidirectional LSTM-CRF models for sequence tagging, arXiv preprint arXiv:1508.01991.

[48] B. Wellner, M. Huyck, S.A. Mardis, J.S. Aberdeen, A.A. Morgan, L. Peshkin, A.S. Yeh, J. Hitzeman, L. Hirschman, Research Paper: Rapidly Retargetable Approaches to De-identification in Medical Records, JAMIA 14 (5) (2007) 564–573.

[49] Z. Jiang, C. Zhao, B. He, Y. Guan, J. Jiang, De-identification of medical records using conditional random fields and long short-term memory networks, J. Biomed. Inform. 75 (2017) S43–S53.