

Learning to Recognize Discontiguous Entities

Aldrian Obaja Muis Wei Lu

Singapore University of Technology and Design
{aldrian_muis, luwei}@sutd.edu.sg

Abstract

This paper focuses on the study of recognizing discontiguous entities. Motivated by a previous work, we propose to use a novel hypergraph representation to jointly encode discontiguous entities of unbounded length, which can overlap with one another. To compare with existing approaches, we first formally introduce the notion of *model ambiguity*, which defines the difficulty level of interpreting the outputs of a model, and then formally analyze the theoretical advantages of our model over previous existing approaches based on linear-chain CRFs. Our empirical results also show that our model is able to achieve significantly better results when evaluated on standard data with many discontiguous entities.

1 Introduction

Building effective automatic named entity recognition (NER) systems that is capable of extracting useful semantic shallow information from texts has been one of the most important tasks in the field of natural language processing. An effective NER system can typically play an important role in certain downstream NLP tasks such as relation extraction, event extraction, and knowledge base construction (Hasegawa et al., 2004; Al-Rfou and Skiena, 2012).

Most traditional NER systems are capable of extracting entities¹ as short spans of texts. Two basic assumptions are typically made when extract-

¹Or sometimes *mentions* are considered, which can be named, nominal or pronominal references to entities (Florian et al., 2004). In this paper we use “mentions” and “entities” interchangeably.

EGD showed [*hiatal hernia*]₁ and vertical [*laceration*]₂ in distal [*esophagus*]₂ with [*blood in [stomach]*]₄₃ and overlying [*lac*]₄.

Figure 1: Discontiguous entities in a medical domain. Words annotated with the same index are part of the same entity. Note that entity 3 and entity 4 overlap with one another.

ing entities: 1) entities do not overlap with one another, and 2) each entity consists of a contiguous sequence of words. These assumptions allow the task to be modeled as a sequence labeling task, for which many existing models are readily available, such as linear-chain CRFs (McCallum and Li, 2003).

While the above two assumptions are valid for most cases, they are not always true. For example, in the entity *University of New Hampshire* of type *ORG* there exists another entity *New Hampshire* of type *LOC*. This violates the first assumption above, yet it is crucial to extract both entities for subsequent tasks such as relation extraction and knowledge base construction. Researchers therefore have proposed to tackle the above issues in NER using more sophisticated models (Finkel and Manning, 2009; Lu and Roth, 2015). Such efforts still largely rely on the second assumption.

Unfortunately, the second assumption is also not always true in practice. There are also cases where the entities are composed of multiple discontiguous sequences of words, such as in disorder mention recognition in clinical texts (Pradhan et al., 2014b), where the entities (disorder mentions in this case) may be discontiguous. Consider the example shown in Figure 1. In this example there are four enti-

ties, the first one, *hiatal hernia*, is a conventional contiguous entity. The second one, *laceration ... esophagus*, is a discontinuous entity, consisting of two parts. The third and fourth ones, *blood in stomach* and *stomach ... lac* (for *stomach laceration*), are overlapping with each other, with the fourth being discontinuous at the same time.

For such discontinuous entities which can potentially overlap with other entities in complex manners, existing approaches such as those based on simple sequence tagging models have difficulties handling them accurately. This stems from the fact that there is a very large number of possible entity combinations in a sentence when the entities can be discontinuous and overlapping.

Motivated by this, in this paper we propose a novel model that can better represent both contiguous and discontinuous entities which can overlap with one another. Our major contributions can be summarized as follows:

- We propose a novel model that is able to represent both contiguous and discontinuous entities.
- Theoretically, we introduce the notion of *model ambiguity* for quantifying the ambiguity of different NER models that can handle discontinuous entities. We present a study and make comparisons about different models' ambiguity under this theoretical framework.
- Empirically, we demonstrate that our model can significantly outperform conventional approaches designed for handling discontinuous entities on data which contains many discontinuous entities.

2 Related Work

Learning to recognize named entities is a popular task in the field of natural language processing. A survey by Nadeau (2007) lists several approaches in NER, including Hidden Markov Models (HMM) (Bikel et al., 1997), Decision Trees (Sekine, 1998), Maximum Entropy Models (Borthwick and Sterling, 1998), Support Vector Machines (SVM) (Asahara and Matsumoto, 2003), and also semi-supervised and unsupervised approaches. Ratnov (2009) utilizes averaged perceptron to solve this problem and also focused on four key design decisions, achieving state-of-the-art in MUC-7 dataset. These ap-

proaches work on standard texts, such as news articles, and the entities to be recognized are defined to be contiguous and non-overlapping.

Noticing that many named entities contain other named entities inside them, Finkel and Manning (2009) proposed a model that is capable of extracting nested named entities by representing the sentence as a constituency parse tree, with named entities as phrases. As a parsing-based model, the approach has a time complexity that is cubic in the number of words in the sentence.

Recently, Lu and Roth (2015) proposed a model that can represent overlapping entities. In addition to supporting nested entities, theoretically this model can also represent overlapping entities where neither is nested in another. The model represents each sentence as a hypergraph with nodes indicating entity types and boundaries. Compared to the previous model, this model has a lower time complexity, which is linear in the number of words in the sentence.

All the above models focus on NER in conventional texts, where the assumption of contiguous entities is valid. In the past few years, there is a growing body of works on recognizing disorder mentions in clinical text. These disorder mentions may be discontinuous and also overlapping. To tackle such an issue, a research group from University of Texas Health Science Center at Houston (Tang et al., 2013; Zhang et al., 2014; Xu et al., 2015) first utilized a conventional linear-chain CRF to recognize disorder mention parts by extending the standard BIO (*Begin, Inside, Outside*) format, and next did some post-processing to combine different components. Though effective, as we will see later, such a model comes with some drawbacks. Nevertheless, their work motivated us to perform further analysis on this issue and propose a novel model specifically designed for discontinuous entity extraction.

3 Models

3.1 Linear-chain CRF Model

Before we present our approach, we would like to spend some time to discuss a simple approach based on linear-chain CRFs (Lafferty et al., 2001). This approach is primarily based on the system by Tang et al. (2013), and this will be the baseline system

EGD showed <i>hiatal</i> _[B] <i>hernia</i> _[I] and vertical <i>laceration</i> _[BD] in distal <i>esophagus</i> _[BD] with <i>blood</i> _[B] <i>in</i> _[I] <i>stomach</i> _[BH] and overlying <i>lac</i> _[BD] .
<i>Infarctions</i> _[BH] either <i>water</i> _[BD] <i>shed</i> _[ID] or <i>embolic</i> _[BD]

Figure 2: Entity encoding in the linear-chain model. **Top:** for the example in Fig 1. **Bottom:** for the second example in Fig 4. The **O** labels are not shown.

that we will make comparison with in later sections.

The problem is regarded as a sequence prediction task, where each word is assigned a label similar to BIO format often used for NER. We used the encoding used by Tang et al. (2013), which uses 7 tags to handle entities that can be discontinuous and overlapping. Specifically, we used **B**, **I**, **O**, **BD**, **ID**, **BH**, and **IH** to denote **B**eginning of entity, **I**nside entity, **O**utside of entity, **B**eginning of **D**iscontiguous entity, **I**nside of **D**iscontiguous entity, **B**eginning of **H**ead, and **I**nside of **H**ead. To encode a sentence in this format, first we identify the contiguous word sequences which are parts of multiple entities. We call these *head components* and we label each word inside each component with **BH** (for the first word in *each component*) or **IH**. Then we find contiguous word sequences which are parts of a discontinuous entity, which we call the *body components*. Words inside those components which have not been labeled are labeled with **BD** (for the first word in *each component*) or **ID**. Finally, words that are parts of a contiguous entity are called *contiguous component*, and, if they have not been labeled, are labeled as **B** (for the first word in *each component*) or **I**.

This encoding is lossy, since the information on which parts constitute the same entity is lost. The top example in Figure 2 is the encoding of the example shown in Figure 1. During decoding, based on the labels only it is not entirely clear whether “*laceration*” should be combined with “*esophagus*” or with “*stomach*” to form a single mention. For the bottom example, we cannot deduce that “*Infarctions*” alone is a mention, since there is no difference in the encoding of a sentence with only two mentions {“*Infarctions ... water shed*”, “*Infarctions ... embolic*”} or having three mentions with “*Infarctions*” as another mention, since in both cases, the word “*Infarctions*” is labeled with **BH**.

Also, it should be noted that some of the label sequences are not valid. For example, a sentence in which there is only one word labeled as **BD** is invalid, since a discontinuous entity requires at least two words to be labeled as **BD** or **BH**. This is, however, a possible output from the linear CRF model, due to the Markov assumption inherent in linear CRF models. Later we see that our models do not have this problem.

3.2 Our Model

Linear-chain CRF models are limited in their representational power when handling complex entities, especially when they can be discontinuous and can overlap with one another. While recent models have been proposed to effectively handle overlapping entities, how to effectively handle discontinuous entities remains a research question to be answered. Motivated by previous efforts on handling overlapping entities (Lu and Roth, 2015), in this work we propose a model based on hypergraphs that can better represent entities that can be discontinuous and at the same time be overlapping with each other.

Unlike the previous work (Lu and Roth, 2015), we establish a novel theoretical framework to formally quantify the ambiguity of our hypergraph-based models and justify their effectiveness by making comparisons with the linear-chain CRF approach.

Now let us introduce our novel hypergraph representation. A hypergraph can be used to represent entities of different types and their combinations in a given sentence. Specifically, a hypergraph is constructed as follows. For the word at position k , we have the following nodes:

- \mathbf{A}^k : this node represents all entities that begin with the current or a future word (to the right of the current word).
- \mathbf{E}^k : this node represents all entities that begin with the current word.
- \mathbf{T}_t^k : this node represents entities of certain specific type t that begin with the current word. There is one \mathbf{T}_t^k for each different type.
- $\mathbf{B}_{t,i}^k$: this node indicates that the current word is part of the i -th component of an entity of type t .
- $\mathbf{O}_{t,i}^k$: this node indicates that the current word appears in between $(i-1)$ -th and i -th components of an entity of type t .

There is also a special leaf node, **X**-node, which indicates the end (*i.e.*, right boundary) of an entity.

The nodes are connected by directed hyperedges, which for the purpose of explaining our models are defined as those edges that connect one node, called the parent node, to one or more child nodes. For ease of notation, in the rest of this paper we use *edge* to refer to directed hyperedge.

The edges Each \mathbf{A}^k is a parent to \mathbf{E}^k and \mathbf{A}^{k+1} , encoding the fact that the set of all entities at position k is the union of the set of entities starting exactly at current position (\mathbf{E}^k) with the set of entities starting at or after position $k + 1$ (\mathbf{A}^{k+1}).

Each \mathbf{E}^k is a parent to $\mathbf{T}_1^k, \dots, \mathbf{T}_T^k$, where T is the total number of possible types that we consider. Each \mathbf{T}_t^k has two edges where it serves as a parent, within one it is parent to $\mathbf{B}_{t,0}^k$ and within another it is to \mathbf{X} . These edges encode the fact that at position k , either there is an entity of type t that begins with the current word (to $\mathbf{B}_{t,0}^k$), or there is no entity of type t that begins with the current word (to \mathbf{X}).

In the full hypergraph, each $\mathbf{B}_{t,i}^k$ is a parent to $\mathbf{B}_{t,i}^{k+1}$ (encoding the fact that the next word also belongs to the same component of the same entity), to $\mathbf{O}_{t,i+1}^{k+1}$ (encoding the fact that this word is part of a discontinuous entity, and the next word is the first word separating current component and the next component), and to \mathbf{X} (representing that the entity ends at this word). Also there are edges with all possible combinations of $\mathbf{B}_{t,i}^{k+1}$, $\mathbf{O}_{t,i+1}^{k+1}$, and \mathbf{X} as the child nodes, representing overlapping entities. For example, the edge $\mathbf{B}_{t,i}^k \rightarrow (\mathbf{B}_{t,i}^{k+1}, \mathbf{X})$ denotes that there is an entity which continues to the next word (the edge to $\mathbf{B}_{t,i}^{k+1}$), while there is another entity ending at k -th word (the edge to \mathbf{X}). In total there are 7 edges in which $\mathbf{B}_{t,i}^k$ is a parent, which are:

- $\mathbf{B}_{t,i}^k \rightarrow (\mathbf{X})$
- $\mathbf{B}_{t,i}^k \rightarrow (\mathbf{O}_{t,i+1}^{k+1})$
- $\mathbf{B}_{t,i}^k \rightarrow (\mathbf{O}_{t,i+1}^{k+1}, \mathbf{X})$
- $\mathbf{B}_{t,i}^k \rightarrow (\mathbf{B}_{t,i}^{k+1})$
- $\mathbf{B}_{t,i}^k \rightarrow (\mathbf{B}_{t,i}^{k+1}, \mathbf{X})$
- $\mathbf{B}_{t,i}^k \rightarrow (\mathbf{B}_{t,i}^{k+1}, \mathbf{O}_{t,i+1}^{k+1})$
- $\mathbf{B}_{t,i}^k \rightarrow (\mathbf{B}_{t,i}^{k+1}, \mathbf{O}_{t,i+1}^{k+1}, \mathbf{X})$

Analogously, $\mathbf{O}_{t,i}^k$ has three edges that connect to

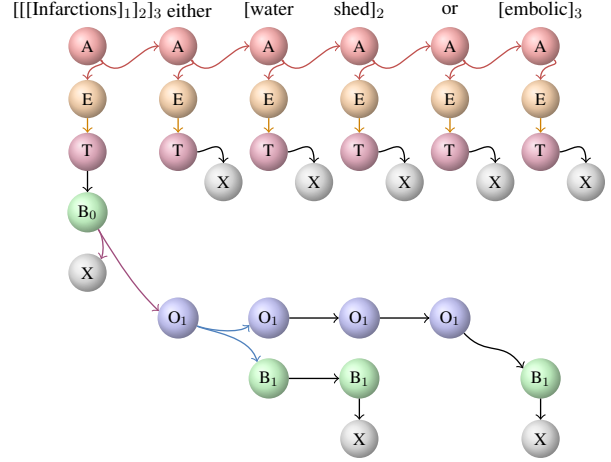


Figure 3: The hypergraph for SHARED model for the second example in Figure 4. The type information in **T**, **B**, and **O**-nodes is not shown. The **X**-node is drawn multiple times for better visualization.

$\mathbf{O}_{t,i}^{k+1}$, $\mathbf{B}_{t,i+1}^{k+1}$, and both. Note that $\mathbf{O}_{t,i}^k$ is not a parent to \mathbf{X} by definition.

During testing, the model will predict a **subgraph** which will result in the predicted entities after decoding. We call this subgraph representing certain entity combination **entity-encoded hypergraph**.

For example, Figure 3 shows the entity-encoded hypergraph of our model encoding the three mentions in the second example in Figure 4. The edge from the **T**-node for the first word to the **B**-node for the first word shows that there is at least one entity starting with this word. The three places where an **X**-node is connected to a **B**-node show the end of the three entities. Note that this hypergraph clearly shows the presence of the three mentions without ambiguity, unlike a linear-chain encoding of this example where it cannot be inferred that “*Infarctions*” alone is a mention, as discussed previously. In this paper, we set the maximum number of components to be 3 since the dataset does not contain any mention with more than 3 components.

Also note that this model supports discontinuous and overlapping mentions of different types since each type has its own set of **O**-nodes and **B**-nodes, unlike the linear-chain model, which supports only overlapping mentions of the same type.

We also experimented with a variant of this model, where we split the **T**-nodes, **B**-nodes, and **O**-nodes further according to the number of components. We split $\mathbf{B}_{t,i}^k$ into $\mathbf{B}_{t,i,j}^k$, $i = 1 \dots j, j =$

1 . . . 3 which represents that the word is part of the i -th component of a mention with total j components. Similarly we split $\mathbf{O}_{t,i}^k$ into $\mathbf{O}_{t,i,j}^k$ and \mathbf{T}_t^k into $\mathbf{T}_{t,j}^k$. We call the original version SHARED model, and this variant SPLIT model. The motivation for this variant is that the majority of overlaps in the data are between discontinuous and contiguous entities, and so splitting the two cases – one component (contiguous) and more (discontinuous) – will reduce ambiguity for those cases.

These models are still ambiguous to some degree, for example when an **O**-node has two child nodes and two parents, we cannot decide which of the parent node is paired with which child node. However, in this paper we argue that:

- This model is less ambiguous compared to the linear-chain model, as we will show later theoretically and empirically.
- Every output of our model is a valid prediction, unlike the linear-chain model since this model will always produce a valid path from **T**-nodes to the **X**-nodes representing some entities.

We will also show through experiments that our models can encode the entities more accurately.

3.3 Interpreting Output Structures

Both the linear-chain CRF model and our models are still ambiguous to some degree, so we need to handle the ambiguity in interpreting the output structures into entities. For all models, we define two general heuristics: ENOUGH and ALL. The ENOUGH heuristic handles ambiguity by trying to produce a minimal set of entities which encodes to the one produced by the model, while ALL heuristic handles ambiguity by producing the union of all possible entity combinations that encode to the one produced by the model. For more details on how these heuristics are implemented for each model, please refer to the supplementary material.

3.4 Training

For both models, the training follows a log-linear formulation, by maximizing the loglikelihood of the training data \mathcal{D} :

$$\mathcal{L}(\mathcal{D}) = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \left[\sum_{e \in \mathcal{E}(\mathbf{x}, \mathbf{y})} [\mathbf{w}^T \mathbf{f}(e)] - \log Z_{\mathbf{w}}(\mathbf{x}) \right] - \lambda \|\mathbf{w}\|^2$$

Here (\mathbf{x}, \mathbf{y}) is a training instance consisting of the sentence \mathbf{x} and the entity-encoded hypergraph $\mathbf{y} \in \mathcal{Y}$ where \mathcal{Y} is the set of all possible mention-encoded hypergraphs. The vector \mathbf{w} consists of feature weights, which are the model parameters to be learned. The set $\mathcal{E}(\mathbf{x}, \mathbf{y})$ consists of all edges present in the entity-encoded hypergraph \mathbf{y} for input \mathbf{x} . The function $\mathbf{f}(e)$ returns the features defined over the edge e , $Z_{\mathbf{w}}(\mathbf{x})$ is the normalization term which gives the sum of scores over all possible entity-encoded hypergraphs in \mathcal{Y} that is relevant to the input \mathbf{x} , and finally λ is the ℓ_2 -regularization parameter.

4 Model Ambiguity

The main aim of this paper is to assess how well each model can represent the discontinuous entities, even in the presence of overlapping entities.

In this section, we will theoretically compare the models’ ambiguity, which is defined as the average number of mention combinations that map to the same encoding in a model. Now, to compare two models, instead of calculating the ambiguity directly, we can calculate the *relative ambiguity* between the two models directly by comparing the number of *canonical encodings* in the two models.

A canonical encoding is a fixed, selected representation of a particular set of mentions in a sentence, among (possibly) several alternative representations. Several alternatives may be present due to the ambiguity of the encoding-decoding process and also since the output of the model is not restricted to a specific rule. For example, for the text “John Smith”, a model trained in BIO format might output “B-PER I-PER” or “I-PER I-PER”, and both will still mean that “John Smith” is a person, although the “correct” encoding would of course be “B-PER I-PER”, which is selected as the canonical encoding. Intuitively, a canonical encoding is a formal way to say that we only consider the “correct” encodings.

A model with larger number of canonical encodings will, on average, have less ambiguity compared to the one with smaller number of canonical encodings. Subsequently, a model with less ambiguity will be more precise in predicting entities.

Let $\mathcal{M}_{LI}(n)$, $\mathcal{M}_{SH}(n)$, $\mathcal{M}_{SP}(n)$ denote the number of canonical encodings of the linear-chain, SHARED, and SPLIT model, respectively, for a sen-

tence with n words. Then we formally define the relative ambiguity of model M_1 over model M_2 , $\mathcal{A}_r(M_1, M_2)$, as follows:

$$\mathcal{A}_r(M_1, M_2) = \lim_{n \rightarrow \infty} \frac{\log \sum_{i=1}^n \mathcal{M}_{M_2}(i)}{\log \sum_{i=1}^n \mathcal{M}_{M_1}(i)} \quad (1)$$

$\mathcal{A}_r(M_1, M_2) > 1$ means model M_1 is more ambiguous than M_2 . Now, we claim the following:

Theorem 4.1. $\mathcal{A}_r(\text{LI}, \text{SH}) > 1$

We provide a proof sketch below. Due to space limitation, we cannot provide the full dynamic programming calculation. We refer the reader to the supplementary material for the details.

Proof Sketch The number of canonical encodings in the linear-chain model is less than 7^n since there are 7 possible tags for each of the n words and not all of the 7^n tag sequences are canonical encodings. So we have $\mathcal{M}_{\text{LI}}(n) < 7^n$ and thus we can derive $\log \sum_{i=1}^n \mathcal{M}_{\text{LI}}(i) < 3n \log 2$.

For our models, by employing some dynamic programming adapted from the inside algorithm (Baker, 1979), we can calculate the growth order of the number of canonical encodings for SHARED model to arrive at a conclusion that $\forall n > n_0, \sum_{i=1}^n \mathcal{M}_{\text{SH}}(i) > C \cdot 2^{10n}$ for some constants n_0, C . Then we have:

$$\mathcal{A}_r(\text{LI}, \text{SH}) \geq \lim_{n \rightarrow \infty} \frac{\log C + 10n \log 2}{3n \log 2} = \frac{10}{3} > 1 \quad \square$$

Theorem 4.1 says that the linear-chain model is more ambiguous compared to our SHARED model. Similarly, we can also establish $\mathcal{A}_r(\text{SH}, \text{SP}) > 1$. Later we also see this empirically from experiments.

5 Experiments

5.1 Data

To allow us to conduct experiments to empirically assess different models’ capability in handling entities that can be discontinuous and can potentially overlap with one another, we need a text corpus annotated with entities which can be discontinuous and overlapping with other entities. We found the largest of such corpus to be the dataset from the task to recognize disorder mentions in clinical text, initially organized by ShARe/CLEF eHealth Evaluation Lab (SHEL) in 2013 (Suominen et al., 2013) and continued in SemEval-2014 (Pradhan et al., 2014a).

The definition of the task is to recognize mentions of concepts that belong to the Unified Medical Language System (UMLS) semantic group *disorders* from a set of clinical texts. Each text has been annotated with a list of disorder mentions by two professional coders trained for this task, followed by an open adjudication step (Suominen et al., 2013).

Unfortunately, even in this dataset, only 8.95% of the mentions are discontinuous. Working directly on such data would prevent us from understanding the true effectiveness of different models when handling entities which can be discontinuous and overlapping. In order to truly understand how different models behave on data with discontinuous entities, we consider a subset of the data where we consider those sentences which contain at least one discontinuous entity. We call the resulting subset the “Discontinuous” subset of the “Original” dataset. Later we will also still use the training data of the “Original” dataset in the experiments.

Note that this “Discontinuous” subset still contains contiguous entities since a sentence usually contains more than one entity. The subset is a balanced dataset with 53.61% of the entities being discontinuous and the rest contiguous. We then split this dataset into training, development, and test set, according to the split given in SemEval 2014 setting (henceforth LARGE dataset). To see the impact of dataset size, we also experiment on a subset of the LARGE dataset, following the SHEL 2013 setting, with the development set in the LARGE dataset used as test set (henceforth SMALL dataset). The training and development set of the SMALL dataset comes from a random 80% (Tr80) and 20% (Tr20) split of the training set in LARGE dataset.

The statistics of the datasets, including the number of overlaps between the entities in the “All” column, are shown in Table 1.

We note that this dataset only contains one type of entity. In later experiments, in order to evaluate the models on multiple types, we create another dataset where we split the entities based on the entity-level semantic category. This information is available for some entities through the Concept Unique Identifier (CUI) annotation in the data. In total we have three types: two types (type A and B) based on the semantic category, and one type (type N) for those entities

Split	#Sentences	Number of mentions				#Overlaps	
		1 part	2 parts	3 parts	Total	All	Diff
Train	534	544	607	44	1,195	205	58
- Tr80	416	448	476	33	957	164	48
- Tr20	118	96	131	11	238	41	10
Dev	303	357	421	18	796	240	28
Test	430	584	610	16	1,210	327	61

Table 1: The statistics of the data. Tr80 and Tr20 refers to the 80% and 20% partitions of the full training data.

having no semantic category information². See the supplementary material for more details. The number of overlaps between different types is shown in the “Diff” column in Table 1. Except for a handful overlaps in development set, all overlaps involve at least one discontinuous entity. Our main result will still be based on the dataset with one type of entity.

<p>The patient had blood in his mouth and on his tongue, pupils were pinpoint and reactive.</p> <ul style="list-style-type: none"> - <i>blood in his mouth</i> - <i>blood ... on his tongue</i> - <i>pupils ... pinpoint</i>
<p>Infarctions either water shed or embolic</p> <ul style="list-style-type: none"> - <i>Infarctions</i> - <i>Infarctions ... water shed</i> - <i>Infarctions ... embolic</i>
<p>You see blood or dark/black material when you vomit or have a bowel movement.</p> <ul style="list-style-type: none"> - <i>blood ... vomit</i> - <i>blood ... bowel movement</i> - <i>dark ... material ... vomit</i> - <i>dark ... bowel movement</i> - <i>black material ... vomit</i> - <i>black material ... bowel movement</i>

Figure 4: Examples of discontinuous and overlapping mentions, taken from the dataset.

Figure 4 shows some examples of the mentions. The first example shows two discontinuous mentions that do not overlap. The second example shows a typical discontinuous and overlapping case. The last example shows a very hard case of overlapping

²It is tempting to just ignore these entities since the N type does not convey any specific information about the entities in it. However, due to the dataset size, excluding this type will lead to very small number of interactions between types. So we decided to keep this type

and discontinuous mentions, as each of the components in $\{\mathit{blood}, \mathit{dark}, \mathit{black\ material}\}$ is paired with each of the word in $\{\mathit{vomit}, \mathit{bowel\ movement}\}$, resulting in six mentions in total, with one having three components (*dark ... material ... vomit*).

5.2 Features

Motivated by the features used by Zhang et al. (2014), for both the linear-chain CRF model and our models we use the following features: neighbouring words with relative position information (we consider previous and next k words, where $k=1, 2, 3$), neighbouring words with relative position information paired with the current word, word n -grams containing the current word ($n=2,3$), POS tag for the current word, POS tag n -grams containing the current word ($n=2,3$), orthographic features (prefix, suffix, capitalization, lemma), note type (discharge summary, echo report, radiology, and ECG report), section name (e.g. Medications, Past Medical History)³, Brown cluster, and word-level semantic category information⁴. We used Stanford POS tagger (Toutanova et al., 2003) for POS tagging, and NLP4J package⁵ for lemmatization. For Brown cluster features, following Tang et al. (2013), we used 1,000 clusters from the combination of training, development, and test set, and used all the subpaths of the cluster IDs as features.

5.3 Experimental Setup

We evaluated the three models on the SMALL dataset and the LARGE dataset.

Note that in both the SMALL and LARGE dataset, about half of all mentions are discontinuous, both in training and test set. We also want to see whether training on a set where the majority of the mentions are contiguous will affect the performance on recognizing discontinuous mentions. So we also performed another experiment where we trained each model on the original training set where the majority of the entities are contiguous. We refer to this original dataset as “Train-Orig” (it contains 10,405 sentences, including those with no entities) and the

³Section names were determined by some heuristics, refer to the supplementary material for more information

⁴This is standard information that can be extracted from UMLS. See (Zhang et al., 2014) for more details.

⁵<http://www.github.com/emorynlp/nlp4j/>

	SMALL						LARGE					
	Train-Disc			Train-Orig			Train-Disc			Train-Orig		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
LI-ENH	59.7	39.8	47.8	71.0	45.8	55.7	54.7	41.2	47.0	64.1	46.5	53.9
LI-ALL	16.6	43.5	24.1	55.5	49.2	52.2	15.2	44.9	22.7	52.8	49.4	51.1
SH-ENH	85.9	39.7	54.3	82.2	48.0	60.6	76.9	40.1	52.7	73.9	49.1	59.0
SH-ALL	85.9	39.7	54.3	82.2	48.0	60.6	76.0	40.5	52.8	73.4	49.5	59.1
SP-ENH	86.7	37.8	52.7	82.5	48.0	60.7	79.4	38.6	52.0	75.3	48.8	59.2
SP-ALL	86.7	37.8	52.7	82.5	48.0	60.7	79.4	38.6	52.0	75.3	48.8	59.2

Table 2: Results on the two datasets and two different training data after optimizing regularization hyperparameter λ in development set. The -ENH and -ALL suffixes refer to the ENOUGH and ALL heuristics. The best result in each column is put in boldface.

earlier one as “Train-Disc”.

First we trained each model on the training set, varying the regularization hyperparameter λ ,⁶ then the λ with best result in the development set using the respective ENOUGH heuristic for each model is chosen for final result in the test set.

For each experiment setting, we show precision (P), recall (R) and F1 measure. Precision is the percentage of the mentions predicted by the model which are correct, recall is the percentage of mentions in the dataset correctly discovered by the model, and F1 measure is the harmonic mean of precision and recall.

5.4 Results and Discussions

The full results are recorded in Table 2.

We see that in general our models have higher precision compared to the linear-chain baseline. This is expected, since our models have less ambiguity, which means that from a given output structure it is easier in our model to get the correct interpretation. We will explore this more in Section 5.5.

The ALL heuristic, as expected, results in higher recall, and this is more pronounced in the linear-chain model, with up to 4% increase from the ENOUGH heuristic, achieving the highest recall in three out of four settings. The high recall of the ALL heuristic in the linear-chain model can be explained by the high level of ambiguity the model has. Since it has more ambiguity compared to our models, one label sequence predicted by the model produces a lot of entities, and so it is more likely to overlap with the gold entities. But this has the drawback of very low precision as we can see in the result.

We see switching from one heuristic to the other

⁶Taken from the set {0.125, 0.25, 0.5, 1.0, 2.0}

does not affect the results of our models much. Looking at the output of our models, they tend to produce output structures with less ambiguity, which causes little difference in the two heuristics.

One example where the baseline made a mistake is the sentence: “Ethanol Intoxication and withdrawal”. The gold mentions are “Ethanol Intoxication” and “Ethanol withdrawal”. But the linear-chain model labeled it as “[Ethanol]_[B] [Intoxication]_[I] and [withdrawal]_[BD]”, which is inconsistent since there is only one discontinuous component. Our models do not have this issue because in our models every subgraph that may be predicted translates to valid mention combinations, as discussed in Section 3.2.

In the “Train-Orig” column, we see that all models can recognize discontinuous entities better when given more data, even though the majority of the entities in “Train-Orig” are contiguous.

5.5 Experiments on Ambiguity

To see the ambiguity of each model empirically, we run the decoding process for each model given the gold output structure, which is the true label sequence for the linear-chain model and the true mention-encoded hypergraph for our models.

We used the entities from the training and development sets for this experiment, and we compare the “Original” datasets with the “Discontiguous” subset to see that the ambiguity is more pronounced when there are more discontinuous entities. Then we show the precision and recall errors (defined as $1 - P$ and $1 - R$, respectively) in Table 3.

Since the ALL heuristics generates all possible mentions from the given encoding, theoretically it should give perfect recall. However, due to errors in the training data, there are mentions which can-

	Discontiguous		Original	
	Prec Err	Rec Err	Prec Err	Rec Err
LI-ALL	63.66%	0.30%	23.81%	0.17%
SH-ALL	1.73%	0.30%	0.35%	0.17%
SP-ALL	1.05%	0.30%	0.22%	0.17%
LI-ENH	2.74%	3.82%	0.52%	0.90%
SH-ENH	1.21%	1.46%	0.25%	0.38%
SP-ENH	0.75%	0.90%	0.17%	0.28%

Table 3: Precision and recall errors (%) of each model in the “Discontiguous” and “Original” datasets when given the gold output structure (label sequence in linear-chain model, hypergraph in our models). Lower numbers are better.

Type	#Ent	Linear-chain			SHARED			SPLIT		
		P	R	F	P	R	F	P	R	F
A	289	69.8	59.9	64.4	79.4	56.1	65.7	81.0	56.1	66.3
B	418	50.0	34.0	40.5	56.8	29.0	38.4	58.2	28.0	37.8
N	503	62.1	37.8	47.0	84.8	43.3	57.4	84.9	42.4	56.5
Total	1210	60.3	41.7	49.3	74.3	41.4	53.2	75.5	40.7	52.9

Table 4: Results on the LARGE dataset when entities are split into three types: A, B, and N. #Ent is the number of entities

not be properly encoded in the models⁷. Removing these errors results in perfect recall (0% recall error). This means that all models are complete: they can encode any mention combinations.

We see however, a very huge difference on the precision error between the linear-chain model and our models, even more when most of the entities are discontiguous. For the discontiguous subset with the ALL heuristic, the linear-chain model produced 5,463 entities, while the SHARED and SPLIT model produced 2,020 and 2,006 entities, respectively. The total number of gold entities is 1,991. This means one encoding in the linear-chain model produces much more distinct mention combinations compared to our model, which again shows that the linear-chain model has more ambiguity. Similarly, we can deduce that the SHARED model has slightly more ambiguity compared to the SPLIT model. This confirms our theoretical result presented previously.

It is also worth noting that in the ENOUGH heuristic our models have smaller errors compared to the linear-chain model, showing that when both models can predict the true output structure (the correct

⁷There are 19 errors in the original dataset, and 6 in the discontiguous subset, which include duplicate mentions and mentions with incorrect boundaries

label sequence for the baseline model and mention-encoded hypergraph for our models), it is easier in our models to get the desired mention combinations.

5.6 Experiments on Multiple Entity Types

We used the LARGE dataset with the multiple-type entities for this experiment. We ran our two models and the linear-chain CRF model with the ENOUGH heuristic on this multi-type dataset, in the same setting as Train-Orig in previous experiments, and the result is shown in Table 4. We used the best lambda from the main experiment for this experiment.

There is a performance drop compared to the LARGE-Train-Orig results in Table 2, which is expected since the presence of multiple types make the task harder. But in general we still see that our models are still better than the baseline, especially the SPLIT model, which shows that in the presence of multiple types, our models can still work better than the baseline model.

6 Conclusions and Future Work

In this paper we proposed new models that can better represent discontiguous entities that can be overlapping at the same time. We validated our claims through theoretical analysis and empirical analysis on the models’ ambiguity, as well as their performances on the task of recognizing disorder mentions on datasets with a substantial number of discontiguous entities. When the true output structure is given, which is still ambiguous in all models, our models show that it is easier to produce the desired mention combinations compared to the linear-chain CRF model with reasonable heuristics. We note that an extension similar to semi-Markov or weak semi-Markov (Muis and Lu, 2016) is possible for our models. We leave this for future investigations.

The supplementary material and our implementations for the models are available at:

<http://statnlp.org/research/ie>

Acknowledgments

We would like to thank the anonymous reviewers for their helpful feedback, and also the ShARE/CLEF eHealth Evaluation Lab for providing us the dataset. This work is supported by MOE Tier 1 grant SUTDT12015008.

References

- Rami Al-Rfou and Steven Skiena. 2012. SpeedRead: A Fast Named Entity Recognition Pipeline. *Proceedings of COLING 2012*, pages 51–66.
- Masayuki Asahara and Yuji Matsumoto. 2003. Japanese Named Entity Extraction with Redundant Morphological Analysis. In *Proceedings of HLT-NAACL '03*, volume 1, pages 8–15.
- James K Baker. 1979. Trainable Grammars for Speech Recognition. *Journal of the Acoustical Society of America*, 65(S1):S132.
- Daniel M. Bikel, Scott Miller, Richard M. Schwartz, and Ralph Weischedel. 1997. Nymble: a high-performance learning name-finder. *Proceedings of the fifth conference on Applied Natural Language Processing (ANLP '97)*, pages 194–201.
- Andrew Borthwick and John Sterling. 1998. NYU: Description of the MENE named entity system as used in MUC-7. In *Proceedings of the 7th Message Understanding Conference (MUC-7)*.
- Jenny Rose Finkel and Christopher D. Manning. 2009. Nested named entity recognition. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing (EMNLP '09)*, volume 1, pages 141–150.
- Radu Florian, Hany Hassan, Abraham Ittycheriah, Hongyan Jing, Nanda Kambhatla, Xiaoqiang Luo, H Nicolov, and Salim Roukos. 2004. A statistical model for multilingual entity detection and tracking. In *Proceedings of HLT-NAACL '04*, pages 1–8.
- Takaaki Hasegawa, Satoshi Sekine, and Ralph Grishman. 2004. Discovering Relations Among Named Entities from Large Corpora. *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, pages 415–422.
- John Lafferty, Andrew McCallum, and Fernando CN Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of International Conference on Machine Learning (ICML '01)*, pages 282–289.
- Wei Lu and Dan Roth. 2015. Joint Mention Extraction and Classification with Mention Hypergraphs. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP '15)*, pages 857–867.
- Andrew McCallum and Wei Li. 2003. Early results for named entity recognition with conditional random fields, feature induction and web-enhanced lexicons. In *Proceedings of HLT-NAACL '03*, volume 4, pages 188–191.
- Aldrian Obaja Muis and Wei Lu. 2016. Weak Semi-Markov CRFs for Noun Phrase Chunking in Informal Text. In *Proceedings of HLT-NAACL '16*, pages 714–719.
- David Nadeau and Satoshi Sekine. 2007. A survey of named entity recognition and classification. *Linguisticae Investigationes*, 30(1):3–26.
- Sameer Pradhan, Noémie Elhadad, Wendy W. Chapman, Suresh Manandhar, and Guergana Savova. 2014a. SemEval-2014 Task 7: Analysis of Clinical Text. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 54–62.
- Sameer Pradhan, Noémie Elhadad, Brett R. South, David Martinez, Lee Christensen, Amy Vogel, Hanna Suominen, Wendy W. Chapman, and Guergana Savova. 2014b. Evaluating the state of the art in disorder recognition and normalization of the clinical narrative. *Journal of the American Medical Informatics Association : JAMIA*, 22(1):143–54.
- 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 '09)*, pages 147–155.
- Satoshi Sekine. 1998. NYU: Description of the Japanese NE system used for MET-2. In *Proceedings of the 7th Message Understanding Conference (MUC-7)*.
- Hanna Suominen, Sanna Salanterä, Sumithra Velupillai, Wendy W. Chapman, Guergana Savova, Noemie Elhadad, Sameer Pradhan, Brett R. South, Danielle L. Mowery, Gareth J. F. Jones, Johannes Leveling, Liadh Kelly, Lorraine Goeuriot, David Martinez, and Guido Zuccon, 2013. *Overview of the ShARE/CLEF eHealth Evaluation Lab 2013*, pages 212–231. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Buzhou Tang, Hongxin Cao, Yonghui Wu, Min Jiang, and Hua Xu. 2013. Recognizing clinical entities in hospital discharge summaries using Structural Support Vector Machines with word representation features. *BMC medical informatics and decision making*, 13 Suppl 1(Suppl 1):S1.
- Kristina Toutanova, Dan Klein, and Christopher D Manning. 2003. Feature-rich part-of-speech tagging with a cyclic dependency network. In *Proceedings of HLT-NAACL '03*, volume 1, pages 252–259.
- Jun Xu, Yaoyun Zhang, Jingqi Wang, Yonghui Wu, and Min Jiang. 2015. UTH-CCB : The Participation of the SemEval 2015 Challenge Task 14. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 311–314.
- Yaoyun Zhang, Jingqi Wang, Buzhou Tang, Yonghui Wu, Min Jiang, Yukun Chen, and Hua Xu. 2014. UTH-CCB: A report for SemEval 2014 – Task 7 Analysis of Clinical Text. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 802–806.