

End-to-end Neural Information Status Classification

Yufang Hou
IBM Research Europe
Dublin, Ireland
yhou@ie.ibm.com

Abstract

Most previous studies on information status (IS) classification and bridging anaphora recognition assume that the gold mention or syntactic tree information is given (Hou et al., 2013a; Hou, 2016; Rösiger et al., 2018; Hou, 2020b; Yu and Poesio, 2020). In this paper, we propose an end-to-end neural approach for information status classification. Our approach consists of a mention extraction component and an information status assignment component. During the inference time, our system takes a raw text as the input and generates mentions together with their information status. On the ISNotes corpus (Markert et al., 2012), we show that our information status assignment component achieves new state-of-the-art results on fine-grained IS classification based on gold mentions. Furthermore, our system performs significantly better than other baselines for both mention extraction and fine-grained IS classification in the end-to-end setting. Finally, we apply our system on BASHI (Rösiger, 2018) and SciCorp (Rösiger, 2016) to recognize referential bridging anaphora. We find that our end-to-end system trained on ISNotes achieves competitive results on bridging anaphora recognition compared to the previous state-of-the-art system that relies on syntactic information and is trained on the in-domain datasets (Yu and Poesio, 2020).

1 Introduction

Information status (IS henceforth) studies how discourse entities are referred to in a text (Halliday, 1967; Prince, 1981; Nissim et al., 2004; Markert et al., 2012) and is related to a wide range of discourse phenomena, such as *coreference* (Ng, 2010), *bridging* (Clark, 1975), and *comparative anaphora* (Modjeska et al., 2003). In general, IS reflects the accessibility of a discourse entity based on the evolving discourse context and the speaker’s assumption about the hearer’s knowledge and beliefs.

Based on Prince (1992) and Nissim et al. (2004), Markert et al. (2012) proposed an IS scheme for written English which consists of eight IS categories. According to Markert et al. (2012), a *mention* is a discourse entity that carries information status. An entity is *old* if it refers to the same entity that has been mentioned previously. *Mediated* entities are discourse-new and hearer-old. They have not been introduced in the discourse directly, but are inferrable from previously mentioned entities, or generally known to the hearer. Finally, an entity is *new* if it is introduced into the discourse for the first time and the hearer cannot infer it from previously mentioned entities. Figure 1 explains each IS category with examples from a short text.

Unlike coreference resolution, most previous work on IS classification and bridging resolution assumes that the gold mention information is given (Rahman and Ng, 2012; Hou et al., 2013a; Hou, 2016; Rösiger et al., 2018; Hou, 2020b; Yu and Poesio, 2020). On a few corpora where gold mentions are not annotated (e.g., BASHI), researchers rely on syntactic parsers to extract mentions (Hou, 2018; Yu and Poesio, 2020). In this paper, we demonstrate that mentions can be reliably extracted without relying on syntactic information. Note that we extract all *singleton* as well as *non-singleton* mentions and assign information status to them. This is different from most previous work on coreference resolution (Lee et al., 2011; Clark and Manning, 2016; Lee et al., 2017) where the mention extraction component is focused on identifying non-singleton mentions only.

Our system consists of two models. The first model (*MenExt*) uses context-dependent boundary representations to extract mentions of discourse entities including *singleton* mentions. Unlike Lee et al. (2017), our model does not have the constraint of the maximum mention length, and it reasons over the space of all spans during the inference time. The second model (*ISAssign*) assigns infor-

S1: [The Bakersfield Supermarket]_{new} went bankrupt [last May]_{m/worldKnow}.

S2: [The business located in [northern Manhattan]_{m/worldKnow}]_{old} closed when [[its]_{old} owner]_{m/syntactic} was murdered.

S3: [Friends]_{m/bridging} expressed [outrage]_{new} at [the murder]_{old}.

S4: Unfortunately, [such crimes]_{m/comparative} are not unusual.

S5: In [[Brooklyn]_{m/worldKnow} and [the Bronx]_{m/worldKnow}]_{m/aggregate}, [one out of four commercial firms]_{new} is burglarized [each year]_{new}.

S6: Since [1984]_{m/worldKnow}, [[New York City]_{m/worldKnow}'s crime rate]_{m/syntactic} has increased by [62%]_{m/function}.

	Description	Example
old	coreferent with an already introduced entity	The business located in northern Manhattan, its, the murder
m/syntactic	syntactically linked to other old or mediated mentions	its owner, New York City's crime rate
m/worldKnow.	generally known to the hearer	last May, northern Manhattan, Brooklyn, the Bronx, New York City, 1984
m/bridging	associative anaphors which link to previously introduced related entities/events	Friends
m/comparative	usually contain a premodifier to indicate that this entity is compared to another entity	such crimes
m/aggregate	coordinated NPs where at least one element is old or mediated	Brooklyn and the Bronx
m/function	refer to a value of a previously explicitly mentioned rise/fall function	62%
new	introduced into the discourse for the first time and not known to the hearer before	The Bakersfield Supermarket, outrage, one out of four commercial firms, each year

Figure 1: Information status categories and examples. IS definitions are from Hou (2020b). Examples are adapted from a news article in ISNotes.

mation status to the extracted mentions. We use a mention's boundary representations in its corresponding context to predict its information status. The context information is adapted from Hou (2020b) which is syntactic-agnostic.

For information status classification based on gold mentions, our second model *ISAssign* significantly outperforms the existing state-of-the-art model (Hou, 2020b) on the ISNotes benchmark based on the same context information. In the end-to-end setting, our system achieves strong results for both mention extraction and IS classification compared to other baselines on ISNotes. We further demonstrate the usefulness of our system by applying it to identify referential bridging anaphora on BASHI (Rösiger, 2018) and SciCorp (Rösiger, 2016). We find that our end-to-end system trained on ISNotes achieves competitive results for bridging anaphora recognition compared to a recent approach (Yu and Poesio, 2020) which uses syntactic tree information for mention extraction and is trained on the in-domain datasets. We will release the code and the trained models at <https://github.com/IBM/bridging-resolution/>.

2 Related Work

Information Status Classification. Previous work on information status classification all assumes that the gold mention information is available (Markert et al., 2012; Rahman and Ng, 2012; Cahill and Riester, 2012; Hou et al., 2013a; Hou, 2016, 2020b). In this work, we remove this constraint and propose a system to tackle the task in an end-to-end manner.

Bridging Anaphora Recognition. Most previous work on bridging focuses on selecting the correct antecedents for gold bridging anaphors (Poesio et al., 2004; Hou et al., 2013b; Hou, 2018, 2020a). As a sub-task of bridging resolution (Hou et al., 2018), bridging anaphora recognition requires a system to identify all bridging anaphors in a text. Some previous work models this task as part of IS classification problem (Hou et al., 2013a; Hou, 2020b). Yu and Poesio (2020) propose a multi-task framework for bridging resolution and report results for bridging anaphora recognition on a wide range of bridging corpora that follow different bridging definitions. All these studies are mainly

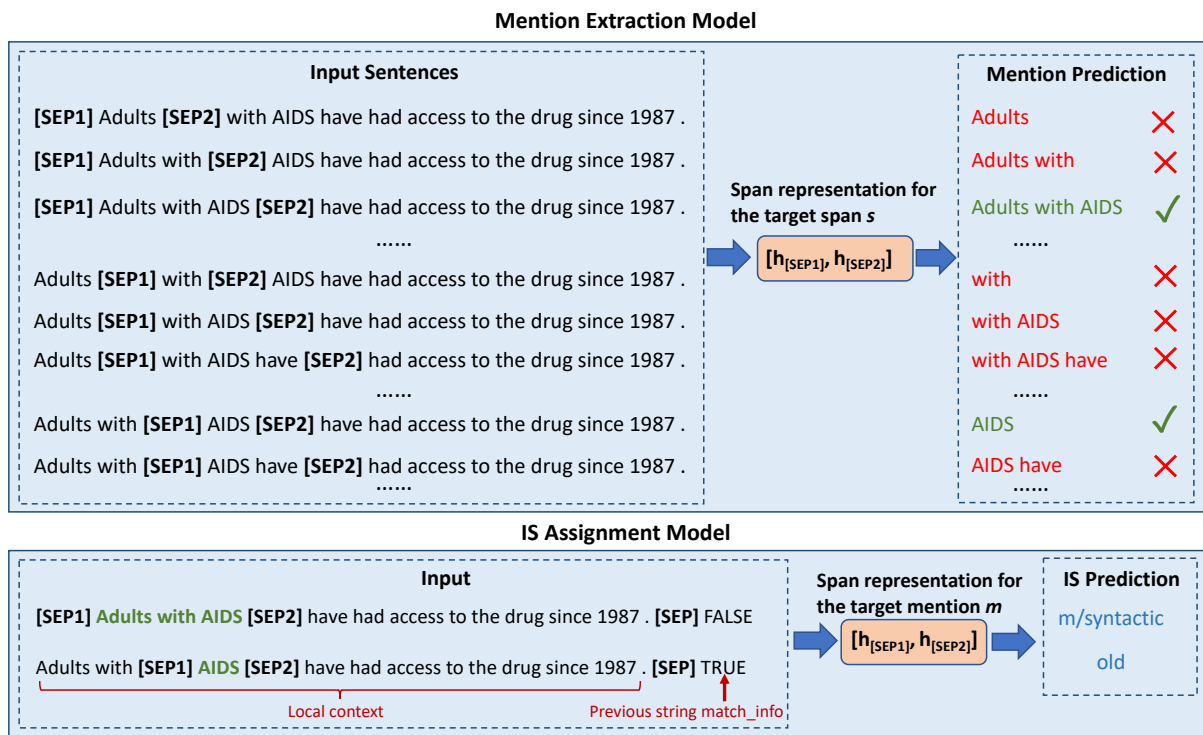


Figure 2: System architecture of our proposed approach for end-to-end IS classification.

based on gold mentions. On a few corpora (BASHI and SciCorp) in which mentions are not annotated, gold or predicted syntactic information is explored to extract “NPs” as the mentions. In this work, we focus on identifying bridging anaphors on the three corpora (i.e., ISNotes, BASHI, and SciCorp) that contain anaphoric referential bridging annotations.¹ Our system does not rely on any gold mention or syntactic information and works well across different domains.

Mention Extraction. Mention extraction is an important task for a lot of down-streaming tasks, such as *relation extraction* (Mintz et al., 2009), *entity linking* (Kolitsas et al., 2018), and *coreference resolution* (Lee et al., 2011). Note that the meaning of “mention” varies across different tasks. Our definition for mentions is close to the one used in coreference resolution. However, mention detection in coreference resolution usually focuses on non-singletons only. In our work, mentions are referential discourse entities which include both singleton and non-singleton entities. Furthermore, IS classification is also related to “anaphorosity de-

¹Rösiger et al. (2018) call this type of bridging as “referential bridging” in which bridging anaphors are truly anaphoric. This is different from bridging annotations in ARRAU (Poesio et al., 2018) where most bridging links are “non-anaphoric referential bridging pairs”, such as *Europe–Spain*.

tection” in coreference resolution (Wiseman et al., 2015) as some *old* mentions are anaphors in coreference.

3 Method

In this section, we describe our system for end-to-end information status classification in detail. Figure 2 illustrates how the system extracts mentions from the input sentences (*Mention Extraction Model*) and assigns information status for the predicted mentions (*IS Assignment Model*).

3.1 Mention Extraction Model

We formulate the task of mention extraction as a binary classification problem for every possible span in a sentence *sent*. Specifically, for a span s containing one or multiple words, we insert two special tokens “[SEP1]” and “[SEP2]” immediately before and after s , respectively. This gives us a new sentence *sent'*. We then use a transformer encoder to encode the new sentence *sent'* and concatenate the hidden states of “[SEP1]” and “[SEP2]” in the last layer to generate the contextual representation for the span s (i.e., $[h_{[SEP1]}, h_{[SEP2]}]$). The probability of s being a valid mention is calculated by a softmax function over the output layer based on the span representation “[$h_{[SEP1]}, h_{[SEP2]}$]”.

Training. Training data for mention extraction is generated by considering spans up to L words for all sentences. For a sentence containing N words, we generate the corresponding training instances by collecting spans that start from each word up to L words (illustrated in “Input Sentences” in Figure 2). The size of the training instances generated for each sentence will be $\mathcal{O}(NL)$ if $L < N$. To make training more efficient, following Lee et al. (2017), we set $L = 10$.²

Inference. During the inference stage, for a given sentence s that contains N words, we apply the learned model to all spans ($\mathcal{O}(N^2)$) without any pruning. The spans that belong to the positive class are extracted as the predicted mentions from this sentence.

It is worth noting that previous work on coreference resolution usually uses aggressive pruning strategies to extract non-singleton mentions (Lee et al., 2017; Wu et al., 2020). In our work, the mention extraction model is trained using all spans up to ten words which include both singleton and non-singleton mentions as the positive instances. Although this filters out a small portion of long mentions, we think the fact that our training data contains the signals of mention/non-mentions for all spans up to ten words is sufficient for the model to learn the differences between mentions and non-mentions. Therefore, we do not apply any pruning during inference.

3.2 IS Assignment Model

We model IS assignment as a multi-class classification problem for every predicted mention in a sentence *sent*. More specifically, for a predicted mention m in *sent*, we insert two special tokens “[SEP1]” and “[SEP2]” immediately before and after m , respectively. We treat this new sentence as the local context to predict the IS category for m . Furthermore, we add an additional token which indicates whether the current mention m has the same string with any previous mentions in the text. Hou (2020b) found that this additional previous context information is important for predicting *old* mentions. As illustrated in Figure 2 (IS Assignment Model), we generate the input for the model

² Lee et al. (2017) reported that mentions that exceed the maximum mention width of ten words only account for less than 2% of the training mentions in OntoNotes. Note that these mentions are all non-singletons. In ISNotes, we found that 10% of the mentions contain more than ten words and most of them (80%) are singletons.

by concatenating the local context and the additional token using the special token “[SEP]”. Similar to the mention extraction model architecture, we use another transformer encoder to encode the input and generate the contextual representation ($[h_{[SEP1]}, h_{[SEP2]}]$) for the mention m by concatenating the hidden states of “[SEP1]” and “[SEP2]” in the last layer. Finally, we use a softmax function to predict the IS category for m based on its contextual representation “[$h_{[SEP1]}, h_{[SEP2]}$]”.

Note that our model is different from the one proposed in Hou (2020b). Hou (2020b) treats IS assignment as a sentence classification task by designing a special sequence for each mention. The author applies BERT to encode the sequence and uses the hidden state of the “[CLS]” token for prediction. Although this model achieves a great improvement compared to previous approaches on IS classification on ISNotes, we think that our mention representation based on “[SEP1]” and “[SEP2]” is more accurate compared to the “[CLS]” token because the latter does not capture a mention’s position information in its local context.³

Training and Inference. The training dataset is generated based on the gold mentions. During inference, we applied the trained model to assign information status to the predicted mentions or the gold mentions in the testing dataset.

3.3 Model Parameters

Our system is developed based on Huggingface Transformers (Wolf et al., 2020). For mention extraction, we train the model for one epoch with a learning rate of $1e - 5$ ⁴ and a batch size of 32. For IS assignment, the model is trained for three epochs with a learning rate of $3e - 5$ and a batch size of 32. Both models are initialized using pre-trained RoBERTa_{LARGE} contextual embeddings (Liu et al., 2019). During training and testing, the maximum text length is set to 128 tokens.⁵

³For both mention extraction and IS assignment, we have tried to extract representation for a mention based on its first token and last token. In our experiments, we found that using the special tokens ($[SEP1]/[SEP2]$) leads to better results.

⁴In practice, we found that a higher learning rate leads to unstable models for mention extraction. This corresponds to the observations from Mosbach et al. (2020) that small learning rates can avoid vanishing gradients early in training when fine-tuning transformer-based masked language models.

⁵The maximum text length is set to 256 tokens when applying the system to SciCorp to recognize bridging anaphors. This is because SciCorp contains a few long sentences due to the noise from sentence splitting.

4 Experiments

4.1 Datasets

We use three datasets for experiments. The first dataset ISNotes (Markert et al., 2012) contains 50 texts with 10,980 mentions from the World Street Journal (WSJ) portion of the OntoNotes corpus. Each mention is assigned to one of the eight IS categories described in Figure 1. Table 1 shows the IS distribution in ISNotes.

For mention extraction and IS assignment, we mainly test our approach on ISNotes since it contains mention annotations for all mentions including both singletons and non-singletons.

Additionally, we apply our model trained on ISNotes to identify bridging anaphors in BASHI (Rösiger, 2018) and SciCorp (Rösiger, 2016) in which bridging anaphors are truly anaphoric⁶. In general, bridging anaphora recognition is a relatively less explored research area, especially in the end-to-end setting.

BASHI contains 459 bridging anaphors from 50 WSJ texts which are different from the ones in ISNotes. Note that BASHI does not contain mention annotations and its bridging anaphors also include comparative anaphors which are treated as a separate IS category in ISNotes.

SciCorp contains 1,366 bridging anaphors from 14 full English scientific papers from two disciplines (i.e., genetics and computational linguistics). Note that bridging anaphors in this corpus are restricted to definite noun phrases only. Following Rösiger et al. (2018), we filter out bridging anaphors that are “containing inferrable” (Prince, 1981), such as “their interest” or “our aim”. These mentions are “mediated/syntactic” in ISNotes.

4.2 Experimental Setup

On ISNotes, all experiments are performed using 10-fold cross-validation on documents. For mention extraction, we report *Recall*, *Precision* and *F-score* based on exact match on gold mentions. For IS assignment, we report *Recall*, *Precision* and *F-score* per IS category. In the end-to-end setting, a prediction is counted as correct if it matches with the boundaries of a gold mention and has the same IS type as the gold mention.

⁶Rösiger et al. (2018) call this type of bridging as “referential bridging”.

Mentions	10,980
old	3,237
mediated	3,708
syntactic	1,592
world knowledge	924
bridging	663
comparative	253
aggregate	211
func	65
new	4,035

Table 1: IS distribution in ISNotes.

	R	P	F
Baselines			
<i>syntactic NPs</i>	78.0	67.1	72.1
<i>Lee et al. (2011)</i>	74.2	70.8	72.5
<i>Yu et al. (2020)</i>	88.7	86.6	87.7
Our model			
<i>our model (test L = 10)</i>	83.5	92.4	87.7
<i>our model</i>	92.8	91.5	92.1

Table 2: Results of our mention extraction model compared to the baselines on ISNotes. The differences between our model and the baselines are statistically significant at $p < 0.01$ using randomization test.

4.3 Results for Mention Extraction on ISNotes

For mention extraction, we compare our method with three baselines. The first baseline (*syntactic NPs*) is simply to extract all NPs from the automatically parsed syntactic trees as mentions. Yu and Poesio (2020) use this strategy to predict mentions for BASHI and SciCorp because gold mentions are not annotated in these two datasets. We use the Stanford CoreNLP toolkit (Manning et al., 2014) to obtain the syntactic trees. The second baseline (*Lee et al. (2011)*) is the mention detection component from Lee et al. (2011) which is widely used in various coreference resolution systems. It is a rule-based system based on automatically parsed trees. The third baseline (*Yu et al. (2020)*) is the best neural mention detection approach proposed in Yu et al. (2020) which uses biaffine attention over a bi-directional LSTM to predict mentions.

Table 2 shows the results of our model on mention extraction compared to the three baselines. Our model outperforms the three baselines in all metrics. In particular, our model achieves an F-score of 92.1, obtaining a 4.4 absolute improvement in

	baseline self-attention with <i>BERT</i> _{LARGE}			this work					
	R	P	F	baseline with <i>RoBERTa</i> _{LARGE}			our model <i>RoBERTa</i> _{LARGE}		
	R	P	F	R	P	F	R	P	F
old	88.4	90.0	89.2	89.0	92.0	90.5	88.8	91.8	90.3
m/worldKnow.	77.7	79.5	78.6	78.0	80.9	79.4	79.2	79.4	79.3
m/syntactic	83.7	81.1	82.4	85.7	81.8	83.7	84.7	83.5	84.1
m/aggregate	80.1	79.3	79.7	77.7	75.9	76.8	76.8	77.5	77.1
m/function	73.4	85.5	79.0	71.9	80.7	76.0	90.6	81.7	85.9
m/comparative	90.5	86.7	88.6	78.7	85.4	81.9	87.7	88.1	87.9
m/bridging	51.0	54.5	52.7	47.8	53.5	50.5	54.1	59.9	56.9
new	86.6	85.2	85.9	88.2	84.9	86.5	88.8	85.8	87.3
acc	83.7			84.3			85.1		

Table 3: Results of the IS assignment model compared to the previous state-of-the-art approach on ISNotes. Bolded scores indicate the best performance for each IS class. The difference between *our model with RoBERTa*_{LARGE} and the baseline is statistically significant at $p < 0.01$ using randomization test.

F-score compared to the baseline *Yu et al. (2020)*.⁷

Table 2 also reports the results of our model using $L = 10$ as the pruning strategy during inference. That is, we filter out predicted mentions which contain more than ten words. This results in a 4.4 F-score drop in performance, which verifies our assumption in Section 3.1 that our model can learn the differences between mentions and non-mentions well from the signals of all spans up to ten words and we do not need to apply pruning during inference.

Inference Efficiency. During the inference stage for mention extraction, previous studies on coreference resolution normally set the maximum mention length to 10 tokens (*Lee et al., 2017*). *Yu et al. (2020)* set this number to 30. We implemented both options as our baselines (see Table 2). For a sentence containing k words, our inference algorithm will have to classify $k(k+1)/2$ spans, which leads to a time complexity of $O(k^2)$. The other two options will have a time complexity of $O(k * 10)$ and $O(k * 30)$, respectively. The average sentence length on ISNotes is 24.3 tokens. Among all sentences, 10.5% have less than 10 tokens, and 73% have less than 30 tokens. Normally k would not be a large number. On a sentence containing 100 tokens, our algorithm is 5 times slower than the pruning with $L = 10$, and 1.7 times slower than the pruning with $L = 30$. In practice, we have tested our algorithm on scientific papers using a

⁷*Yu et al. (2020)* report 87.9 (recall), 89.7 (precision) and 88.8 (F-score) for mention detection on the CRAC testing dataset in the “High F1” setting. Our mention extraction model trained on the CRAC training dataset achieves better results on the same test set: 91.1 (recall), 90.4 (precision) and 90.8 (F-score).

simple heuristic to filter out non-mentions during inference. More specifically, we do not perform inference for spans starting with verbs, punctuation or prepositions using a dictionary. This greatly helps us to speed up the inference process. In addition, one main goal of our work is to test the performance of this algorithm without adding any pruning constraints. As a result, it is interesting to see that the model can effectively filter out long non-mentions by itself.

4.4 Results for IS Classification on ISNotes Based on Gold Mentions

In this section, we compare our IS assignment model to the previous state-of-the-art approach (*Hou, 2020b*) using gold mentions as input. Note that the baseline model proposed by *Hou (2020b)* treats IS classification as a sentence classification task based on carefully designed “pseudo sentences”. It is fine-tuned on *BERT*_{LARGE} and uses the hidden state of the “[CLS]” token for prediction.⁸

We find that changing *BERT*_{LARGE} to *RoBERTa*_{LARGE} improves the overall accuracy by 0.6% (see *baseline with RoBERTa*_{LARGE} in Table 3). Our model further improves the accuracy by 0.8% (85.1 vs. 84.3 in Table 3), which demonstrates that the mention representation based on the boundary embeddings in our model is more accurate than the baseline model. In addition, it seems that our mention representation is especially effective for detecting bridging anaphors and *mediated/function* mentions.

⁸Please check Appendix A.4 for more results of IS classification on ISNotes from previous approaches.

	<i>NPs from predicted syntactic trees</i>			<i>baselines system mentions Lee et al. (2011)</i>			<i>system mentions Yu et al. (2020)</i>			<i>this work our mention extraction model</i>		
	R	P	F	R	P	F	R	P	F	R	P	F
old	66.4	77.3	71.4	79.8	81.4	80.6	83.2	86.6	84.9	85.0	89.1	87.0
m/worldKnow.	49.8	58.9	54.0	65.0	61.3	63.1	72.4	69.3	70.8	74.9	70.8	72.8
m/syntactic	67.0	53.9	59.8	57.0	60.8	58.9	77.4	70.2	73.6	80.4	76.3	78.3
m/aggregate	60.7	61.5	61.1	28.0	31.9	29.8	61.1	63.2	62.2	69.7	72.1	70.8
m/function	76.6	63.6	69.5	70.3	50.6	58.8	73.4	73.4	73.4	81.2	81.2	81.2
m/comparative	73.5	56.9	64.1	52.2	55.2	53.7	74.3	72.3	73.3	82.2	79.7	80.9
m/bridging	50.2	42.1	45.8	42.2	44.5	43.3	51.6	53.1	52.3	51.3	53.5	52.4
new	70.8	49.1	58.0	58.2	49.8	53.7	74.1	70.6	72.3	79.0	75.4	77.1
overall	65.8	56.7	60.9	63.4	60.5	61.9	75.5	73.8	74.6	78.9	77.8	78.3

Table 4: Results of IS classification on ISNotes in the end-to-end setting using different mention extraction approaches. Bolded scores indicate the best performance for each IS class. The differences between our system and other baselines are statistically significant at $p < 0.01$ using randomization test.

4.5 Results for IS Classification on ISNotes Based on Extracted Mentions

Previous studies on IS classification only focus on gold mentions (Markert et al., 2012; Rahman and Ng, 2012; Cahill and Riester, 2012; Hou et al., 2013a; Hou, 2020b). It is unclear how these approaches will perform in realistic scenarios where gold mentions are often not available. In this section, we evaluate our system for IS classification in the end-to-end setting. More specifically, we apply our IS assignment model described in Section 3.2 to four sets of system mentions separately. The system mentions are predicted by four different approaches that we have compared in Section 4.3.

Table 4 reports the results of IS classification on ISNotes in the end-to-end setting. It is interesting to notice that although the first two baselines achieve similar results for mention extraction in terms of F-score (*Syntactic NP* vs. *Lee et al. (2011)* in table 2), their results on IS classification are quite different on each IS category. More specifically, the mention detection component from Lee et al. (2011) performs much better on identifying *old* and *mediated/WorldKnowledge* mentions compared to the other simple baseline that predicts all NPs as mentions. This is in line with the fact that system mentions from Lee et al. (2011) are optimized to identify non-singletons for coreference resolution. In ISNotes, among all mentions that are the first mentions in coreference chains, 46.2% are *new* mentions, 23.2% are *m/worldKnowledge* mentions, and 18.0% are *m/syntactic* mentions.

In general, our proposed system outperforms the three baselines in all metrics for most IS categories by a large margin. We also notice that within our

system, the IS classification results in the end-to-end setting degrade for all IS categories compared to the setting with gold mentions (Table 4 vs. Table 3). Particularly, the *old* and *mediated/bridging* categories have less performance degradation compared to other IS classes.

4.6 Results for Bridging Anaphora Recognition on BASHI and SciCorp

To test the effectiveness of our system in a more realistic scenario, we apply our models trained on the whole ISNotes corpus to BASHI and SciCorp to recognize bridging anaphors. Note that BASHI also includes comparative anaphors as bridging anaphors. Therefore, we merge the predictions for *mediated/bridging* and *mediated/comparative* and treat them as bridging anaphors for BASHI. For SciCorp, we use a small list of determiners⁹ to choose definite bridging anaphors from all predicted *mediated/bridging* mentions since only definite noun phrases are annotated in SciCorp.

Table 5 shows the results of our system for bridging anaphora recognition on BASHI and SciCorp. We also list the results from a recent approach (Yu and Poesio, 2020) for these two corpora. Yu and Poesio (2020) propose a multi-task learning framework to identify bridging links and report results on BASHI and SciCorp using NPs as the predicted mentions. Note that the results from Yu and Poesio (2020) are based on the models trained on the in-domain datasets using 10-fold cross-validation¹⁰.

⁹The whole list of determiners we used to detect definite NPs is: {*the, this, that, these, those*}.

¹⁰The results on SciCorp from Yu and Poesio (2020) may not be directly comparable to ours since it is unclear whether they filter out *containing inferrable* mentions from bridging

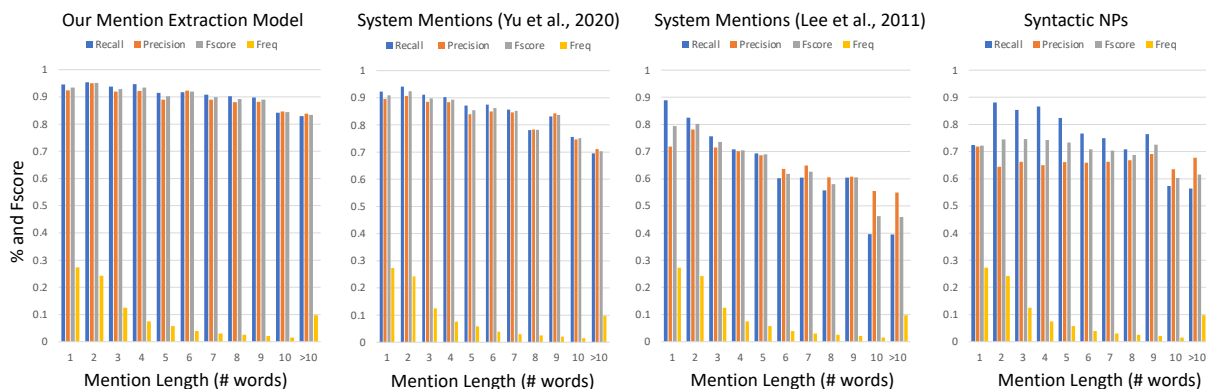


Figure 3: Mention extraction performance for mentions of various lengths in different systems. “Freq” is the percentage of gold mentions with each length.

	R	P	F
BASHI			
<i>Yu and Poesio (2020)</i>	34.2	34.4	34.3
<i>our system</i>	59.7	25.5	35.8
SciCorp			
<i>Yu and Poesio (2020)</i>	35.7	45.0	39.8
<i>our system</i>	47.6	36.0	41.0

Table 5: Results of our system for bridging anaphora recognition on BASHI and SciCorp. “*Yu and Poesio (2020)*” is the results from [Yu and Poesio \(2020\)](#) which include gold coreferent anaphors for evaluation.

This is because SciCorp contains texts from a very different domain compared to ISNotes and BASHI.

In general, our system trained on ISNotes achieves competitive results to identify bridging anaphors on BASHI and SciCorp compared to [Yu and Poesio \(2020\)](#). We notice that our system achieves relatively high recall scores on both corpora (59.7% on BASHI and 47.6% on SciCorp).

More interestingly, although our system is trained on a dataset consisting of news articles, it still can identify certain domain specific bridging anaphors in genetics and computational linguistic scientific papers from SciCorp, such as “*the underlying siRNA*”, “*the missing two Nucleotides*”, “*the target mRNA*”, “*the previous optimization*”, “*the objective function*”, “*the next clause*”, and “*the most predictive features*”. [Hou \(2020b\)](#) states that some bridging can be indicated by referential patterns without world knowledge about the anaphor/antecedent NPs. It seems that our IS assignment model can capture some of these patterns and generalize them into different domains.

anaphors annotated in SciCorp.

5 Analysis

To better understand the strengths and weaknesses of our system, we provide both quantitative and qualitative analyses. In the following discussion, we use predictions from ISNotes.

Mention Extraction Performance. The training dataset we generated offers strong signals for spans that correspond to entity mentions, including singletons and non-singletons. In Figure 3, we show recall, precision, and F-score for mentions of various lengths. In general, our mention extraction model (Figure 3: left) performs significantly better than the three baseline models for all length groups. Although there is some performance degradation for longer mentions containing more than ten words, the results of our model for this group are still quite strong with an F-score of 83.4.

It is also worth noting that the results using the mention detection component from [Lee et al. \(2011\)](#) degrade gradually with the increase of the mention length (Figure 3: middle). We think this is because the system is optimized to identify non-singleton mentions for coreference resolution and most coreferent mentions are short. In ISNotes, 88.6% of non-singleton mentions contain less than six words.

Common Errors on Identifying Bridging. Although our system achieves substantial improvements for bridging anaphora recognition on ISNotes compared to previous work ([Hou et al., 2013a](#); [Yu and Poesio, 2020](#); [Hou, 2020b](#)), detecting bridging anaphors still remains a challenging problem compared to other IS categories. We examine the confusion matrix of our IS assignment model based on the gold mentions. We find that

the model mostly confuses bridging anaphors with *new* mentions. The highest portion of recall errors is due to the fact that the model misclassifies 202 bridging anaphors as *new* mentions, while most precision errors are 142 *new* mentions being misclassified as *mediated/bridging*. It seems that our model struggles to distinguish bridging anaphors from some generic *new* mentions that are relational nouns (Löbner, 1985) (see Example 1 and Example 2).

(1) **Rescue crews**, however, gave up hope that others would be found alive under the collapsed roadway.

(2) Lang is cutting **costs** and will attempt to operate the magazine with only subscription revenue.

6 Conclusions

In this paper, we propose a system that addresses information status classification in the end-to-end setting for the first time. Our mention extraction model does not require any pruning process and can still achieve strong performance for longer spans. We show that our system outperforms other baselines for mention extraction and IS assignment in different settings. We further demonstrate that our system trained on ISNotes can be applied to identify bridging anaphors in different domains. In the future, we plan to further improve our algorithm for bridging anaphora recognition. Applying our system to detect mentions and anaphoricity for coreference resolution is another interesting avenue for future work.

Acknowledgments

The author thanks the anonymous reviewers for their the valuable feedback.

References

Aoife Cahill and Arndt Riester. 2012. Automatically acquiring fine-grained information status distinctions in German. In *Proceedings of the SIGdial 2012 Conference: The 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, Seoul, Korea, 5–6 July 2012, pages 232–236.

Herbert H. Clark. 1975. Bridging. In *Proceedings of the Conference on Theoretical Issues in Natural Language Processing*, Cambridge, Mass., June 1975, pages 169–174.

Kevin Clark and Christopher D. Manning. 2016. Improving coreference resolution by learning entity-level distributed representations. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 643–653, Berlin, Germany. Association for Computational Linguistics.

M. A. K. Halliday. 1967. Notes on transitivity and theme in English, Part 2. *Journal of Linguistics*, 3:199–244.

Yufang Hou. 2016. Incremental fine-grained information status classification using attention-based LSTMs. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pages 1880–1890.

Yufang Hou. 2018. A deterministic algorithm for bridging anaphora resolution. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium, 31 October–4 November 2018, pages 1938–1948.

Yufang Hou. 2020a. Bridging anaphora resolution as question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Seattle, Wash., 5–10 July 2020, pages 1428–1438.

Yufang Hou. 2020b. Fine-grained information status classification using discourse context-aware BERT. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6101–6112, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Yufang Hou, Katja Markert, and Michael Strube. 2013a. Cascading collective classification for bridging anaphora recognition using a rich linguistic feature set. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Wash., 18–21 October 2013, pages 814–820.

Yufang Hou, Katja Markert, and Michael Strube. 2013b. Global inference for bridging anaphora resolution. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Atlanta, Georgia, 9–14 June 2013, pages 907–917.

Yufang Hou, Katja Markert, and Michael Strube. 2018. Unrestricted bridging resolution. *Computational Linguistics*, 44(2):237–284.

Nikolaos Kolitsas, Octavian-Eugen Ganea, and Thomas Hofmann. 2018. End-to-end neural entity linking. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 519–529, Brussels, Belgium. Association for Computational Linguistics.

- Heeyoung Lee, Yves Peirsman, Angel Chang, Nathanael Chambers, Mihai Surdeanu, and Dan Jurafsky. 2011. Stanford’s multi-pass sieve coreference resolution system at the CoNLL-2011 shared task. In *Proceedings of the Shared Task of the 15th Conference on Computational Natural Language Learning*, Portland, Oreg., 23–24 June 2011, pages 28–34.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. [End-to-end neural coreference resolution](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 188–197, Copenhagen, Denmark. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Sebastian Löbner. 1985. Definites. *Journal of Semantics*, 4:279–326.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](#). In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Baltimore, Maryland. Association for Computational Linguistics.
- Katja Markert, Yufang Hou, and Michael Strube. 2012. [Collective classification for fine-grained information status](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*, Jeju Island, Korea, 8–14 July 2012, pages 795–804.
- Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. [Distant supervision for relation extraction without labeled data](#). In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 1003–1011, Suntec, Singapore. Association for Computational Linguistics.
- Natalia N. Modjeska, Katja Markert, and Malvina Nissim. 2003. [Using the web in machine learning for other-anaphora resolution](#). In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*, pages 176–183.
- Marius Mosbach, Maksym Andriushchenko, and Dietrich Klakow. 2020. [On the stability of fine-tuning bert: Misconceptions, explanations, and strong baselines](#).
- Vincent Ng. 2010. [Supervised noun phrase coreference research: The first fifteen years](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1396–1411, Uppsala, Sweden. Association for Computational Linguistics.
- Malvina Nissim, Shipara Dingare, Jean Carletta, and Mark Steedman. 2004. An annotation scheme for information status in dialogue. In *Proceedings of the 4th International Conference on Language Resources and Evaluation*, Lisbon, Portugal, 26–28 May 2004, pages 1023–1026.
- Massimo Poesio, Yulia Grishina, Varada Kolhatkar, Nafise Sadat Moosavi, Ina Rösiger, Adam Roussel, Fabian Simonjetz, Alexandra Uma, Olga Uryupina, Juntao Yu, and Heike Zinsmeister. 2018. Anaphora resolution with the ARRAU corpus. In *Proceedings of the Workshop on Computational Models of Reference, Anaphora and Coreference*. New Orleans, Louisiana, June 6, 2018, pages 11–22.
- Massimo Poesio, Rahul Mehta, Axel Maroudas, and Janet Hitzeman. 2004. Learning to resolve bridging references. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics*, Barcelona, Spain, 21–26 July 2004, pages 143–150.
- Ellen F. Prince. 1981. Towards a taxonomy of given-new information. In P. Cole, editor, *Radical Pragmatics*, pages 223–255. Academic Press, New York, N.Y.
- Ellen F. Prince. 1992. The ZPG letter: Subjects, definiteness, and information-status. In W.C. Mann and S.A. Thompson, editors, *Discourse Description. Diverse Linguistic Analyses of a Fund-Raising Text*, pages 295–325. John Benjamins, Amsterdam.
- Altaf Rahman and Vincent Ng. 2012. Learning the fine-grained information status of discourse entities. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, Avignon, France, 23–27 April 2012, pages 798–807.
- Ina Rösiger. 2016. [SciCorp: A corpus of English scientific articles annotated for information status analysis](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 1743–1749, Portorož, Slovenia. European Language Resources Association (ELRA).
- Ina Rösiger. 2018. [BASHI: A corpus of wall street journal articles annotated with bridging links](#). In *Proceedings of the 11th International Conference on Language Resources and Evaluation*, Miyazaki, Japan, 7–12 May 2018, pages 382–388.
- Ina Rösiger, Arndt Riester, and Jonas Kuhn. 2018. [Bridging resolution: Task definition, corpus resources and rule-based experiments](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, Santa Fe, New-Mexico, USA, 20–26 August 2018, pages 3516–3528.
- Sam Wiseman, Alexander M. Rush, Stuart Shieber, and Jason Weston. 2015. [Learning anaphoricity and antecedent ranking features for coreference resolution](#). In *Proceedings of the 53rd Annual Meeting of the*

Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 1416–1426, Beijing, China. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. **Transformers: State-of-the-art natural language processing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Wei Wu, Fei Wang, Arianna Yuan, Fei Wu, and Jiwei Li. 2020. **CorefQA: Coreference resolution as query-based span prediction**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6953–6963, Online. Association for Computational Linguistics.

Juntao Yu, Bernd Bohnet, and Massimo Poesio. 2020. **Neural mention detection**. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1–10, Marseille, France. European Language Resources Association.

Juntao Yu and Massimo Poesio. 2020. **Multitask learning-based neural bridging reference resolution**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3534–3546, Barcelona, Spain (Online). International Committee on Computational Linguistics.

A Appendices

A.1 Datasets

The ISNotes corpus can be downloaded from <https://www.h-its.org/software/isnotes-corpus/>. The BASHI corpus can be downloaded from <https://www.ims.uni-stuttgart.de/en/research/resources/corpora/bashi/>. The SciCorp dataset can be downloaded from <https://www.ims.uni-stuttgart.de/en/research/resources/corpora/scicorp/>. Note that ISNotes needs to be integrated with the original OntoNotes ONF files, and BASHI needs to be merged with the OntoNotes CoNLL format annotations.

A.2 Model Parameters

Our system is developed based on Huggingface Transformers¹¹. The mention extraction model is trained for one epoch with a learning rate of $1e - 5$ and a batch size of 32. The IS assignment model is trained for three epochs with a learning rate of $3e - 5$ and a batch size of 32. Both models are initialized using pre-trained RoBERTa_{LARGE} contextual embeddings. They have 24 transformer blocks, 1024 hidden units, 16 self-attention heads, and around 355M parameters.

During training and testing, the maximum text length is set to 128 tokens. The maximum text length is set to 256 tokens when we apply the system to SciCorp to recognize bridging anaphors. This is because SciCorp contains a few long sentences due to the noise from sentence splitting.

A.3 Computing Infrastructure and Running Time

We carried out experiments in a computing cluster environment. We use a node with a v100 GPU and 32 GB RAM for all experiments. We perform 10-fold cross-validation on ISNotes. It takes around 120 minutes to train and test the mention extraction model for one fold. And it takes around 25 minutes to train and test the IS assignment model for one fold without any pruning.

A.4 Results of previous studies

In this section we compare our IS assignment model to previous approaches which use gold mentions as input. Table 6 shows the results of fine-

¹¹<https://github.com/huggingface/transformers>

	<i>baselines</i>												<i>this work</i>		
	Hou et al.(2013) <i>collective</i>			Hou et al.(2013) <i>cascade collective</i>			Hou (2016) <i>incremental LSTM</i>			Hou (2020) <i>self-attention with</i> <i>BERT_{LARGE}</i>					
	R	P	F	R	P	F	R	P	F	R	P	F	R	P	F
old	84.4	86.0	85.2	82.2	87.2	84.7	85.4	84.9	85.2	88.4	90.0	89.2	88.8	91.8	90.3
m/worldKnow.	67.4	77.3	72.0	67.2	77.2	71.9	67.1	74.5	70.6	77.7	79.5	78.6	79.2	79.4	79.3
m/syntactic	82.2	81.9	82.0	81.6	82.5	82.0	80.8	81.9	81.4	83.7	81.1	82.4	84.7	83.5	84.1
m/aggregate	64.5	79.5	71.2	63.5	77.9	70.0	67.8	84.6	75.3	80.1	79.3	79.7	76.8	77.5	77.1
m/function	67.7	72.1	69.8	67.7	72.1	69.8	64.6	76.4	70.0	73.4	85.5	79.0	90.6	81.7	85.9
m/comparative	81.8	82.1	82.0	86.6	78.2	82.2	77.9	83.1	80.4	90.5	86.7	88.6	87.7	88.1	87.9
m/bridging	19.3	39.0	25.8	44.9	39.8	42.2	15.7	32.3	21.1	51.0	54.5	52.7	54.1	59.9	56.9
new	86.5	76.1	81.0	83.0	78.1	80.5	87.2	74.8	80.5	86.6	85.2	85.9	88.8	85.8	87.3
acc	78.9			78.6			78.6			83.7			85.1		

Table 6: Results of the IS assignment model compared to previous approaches on ISNotes based on gold mentions.

grained IS classification of different models on IS-Notes. Below is a short description of each model.

collective. Hou et al. (2013a) applied collective classification to account for the linguistic relations among IS categories. They explored a wide range of features (34 in total), including a large number of lexico-semantic features (for recognizing bridging) as well as a couple of surface features and syntactic features.

cascaded collective. This is the cascading minority preference system for bridging anaphora recognition from Hou et al. (2013a).

incremental LSTM. This is the attention-based LSTM model proposed by Hou (2016). The model uses one-hot vectors to encode IS classes and predicts information status for all mentions of a document from left to right incrementally.

self-attention with BERT_{LARGE}. Hou (2020b) treats IS classification as a sentence classification task based on carefully designed “pseudo sentences”. It is fine-tuned on BERT_{LARGE} and uses the hidden state of the “[CLS]” token for prediction.

this work. The model described in Section 3.2.