

Joint Coreference Resolution and Character Linking for Multiparty Conversation

Jiabin Bai¹, Hongming Zhang¹, Yangqiu Song¹, and Kun Xu²

¹CSE, HKUST

²Tencent AI Lab

{jbai, hzhangal, yqsong}@cse.ust.hk, kxkunxu@tencent.com

Abstract

Character linking, the task of linking mentioned people in conversations to the real world, is crucial for understanding the conversations. For the efficiency of communication, humans often choose to use pronouns (e.g., “she”) or normal phrases (e.g., “that girl”) rather than named entities (e.g., “Rachel”) in the spoken language, which makes linking those mentions to real people a much more challenging than a regular entity linking task. To address this challenge, we propose to incorporate the richer context from the coreference relations among different mentions to help the linking. On the other hand, considering that finding coreference clusters itself is not a trivial task and could benefit from the global character information, we propose to jointly solve these two tasks. Specifically, we propose C², the joint learning model of Coreference resolution and Character linking. The experimental results demonstrate that C² can significantly outperform previous works on both tasks. Further analyses are conducted to analyze the contribution of all modules in the proposed model and the effect of all hyper-parameters.

1 Introduction

Understanding conversations has long been one of the ultimate goals of the natural language processing community, and a critical step towards is grounding all mentioned people to the real world. If we can achieve that, we can leverage our knowledge about these people (e.g., things that happened to them before) to better understand the conversation. On the other hand, we can also aggregate the conversation information back to our understanding about these people, which can be used for understanding future conversations that involve the same people. To simulate the real conversations and investigate the possibility for models to ground mentioned people, the character linking task was




Figure 1: The composition of the mentions in conversations for character grounding. Over 88% of the mentions are not named entities, which brings exceptional challenges when linking those to character entities.

proposed (Chen and Choi, 2016). Specifically, it uses the transcripts of TV shows (i.e., *Friends*) as the conversations and asks the models to ground all person mentions to characters.

Even though the character linking task can be viewed as a special case of the entity linking task, it is more challenging than the ordinary entity linking task for various reasons. First, the ordinary entity linking task often aims at linking named entities to external knowledge bases such as Wikipedia, where rich information (e.g., definitions) are available. However, for the character linking task, we do not have the support of such rich knowledge base and all we have are the names of these characters and simple properties (e.g., gender) about these characters. Second, the mentions in the ordinary entity linking are mostly concepts and entities, but not pronouns. However, as shown in Figure 1, 88% of the character mentions are pronouns (e.g., “he”) or personal nouns (e.g., “that guy”) while only 12% are named entities.

Considering that pronouns have relatively weak semantics by themselves, to effectively ground mentions to the correct characters, we need to fully utilize the context information of the whole conversation rather than just the local context they appear in. One potential solution is using the coreference relations among different mentions as the bridge to



Paul the Wine Guy

Monica : There's nothing to tell! **He**'s just some **guy** I work with!
Joey : C'mon, you're going out with the **guy**! There's gotta be something wrong with **him**!
Ross: All right Joey, be nice. So does **he** have a hump? A hump and a hairpiece?

Figure 2: Coreference clusters can help to connect the whole conversation to provide a richer context for each mention such that we can better link them to *Paul*. Meanwhile, the character *Paul* can also provide global information to help resolve the coreference.

connect the richer context. One example is shown in Figure 2. It is difficult to directly link the highlighted mentions to the character *Paul* based on their local context because the local context of each mention can only provide a single piece of information about its referent, e.g., “person is a male” or “the person works with Monica.” Given the coreference cluster, the mentions refer to the same person, and the pieces of information are put together to jointly determining the referent. As a result, it is easier for a model to do character linking with resolved coreference. Similar observations are also made in (Chen et al., 2017).

At the same time, we also noticed that coreference resolution, especially those involving pronouns, is also not trivial. As shown by the recent literature on the coreference resolution task (Lee et al., 2018; Kantor and Globerson, 2019), the task is still challenging for current models and the key challenge is how to utilize the global information about entities. And that is exactly what the character linking model can provide. For example, in Figure 2, it is difficult for a coreference model to correctly resolve the last mention *he* in the utterance given by *Ross* based on its local context, because another major male character (*Joey*) joins the conversation, which can distract and mislead the coreference model. However, if the model knows the mention *he* links to the character *Paul* and *Paul* works with *Monica*, it is easier to resolve *he* to some *guy* that *Monica* works with.

Motivated by these observations, we propose to jointly train the Coreference resolution and Character linking tasks and name the joint model as C^2 . C^2 adopts a transformer-based text encoder and includes a mention-level self-attention (MLSA) module that enables the model to do mention-level contextualization. Meanwhile, a joint loss function is designed and utilized so that both tasks

can be jointly optimized. The experimental results demonstrate that C^2 outperforms all previous work significantly on both tasks. Specifically, compared with the previous work (Zhou and Choi, 2018), C^2 improves the performance by 15% and 26% on the coreference resolution and character linking tasks¹ respectively comparing to the previous state-of-the-art model ACNN (Zhou and Choi, 2018). Further hyper-parameter and ablation studies testify the effectiveness of different components of C^2 and the effect of all hyper-parameters. Our code is available at <https://github.com/HKUST-KnowComp/C2>.

2 Problem Formulations and Notations

We first introduce the coreference resolution and character linking tasks as well as used notations. Given a conversation, which contains multiple utterances and n character mentions c_1, c_2, \dots, c_n , and a pre-defined character set \mathcal{Z} , which contains m characters z_1, z_2, \dots, z_m . The coreference resolution task is grouping all mentions to clusters such that all mentions in the same cluster refer to the same character. The character linking task is linking each mention to its corresponding character.

3 Model

In this section, we introduce the proposed C^2 framework, which is illustrated in Figure 3. With the conversation and all mentions as input, we first encode them with a shared mention representation encoder module, which includes a pre-trained transformer text encoder and a mention-level self-attention (MLSA) module. After that, we make predictions for both tasks via two separate modules. In the end, a joint loss function is devised so that the model can be effectively trained on both tasks simultaneously. Details are as follows.

3.1 Mention Representation

We use pre-trained language models (Devlin et al., 2018; Joshi et al., 2019a) to obtain the contextualized representations for mentions. As speaker information is critical for the conversation understanding, we also include that information by appending speaker embeddings to each mention. As

¹The performance on the coreference resolution is evaluated based on the average F1 score of B3, CEAF _{ϕ} 4, and BLANC. The performance on the character linking task is evaluated by the average F1 score of the micro and macro F1.

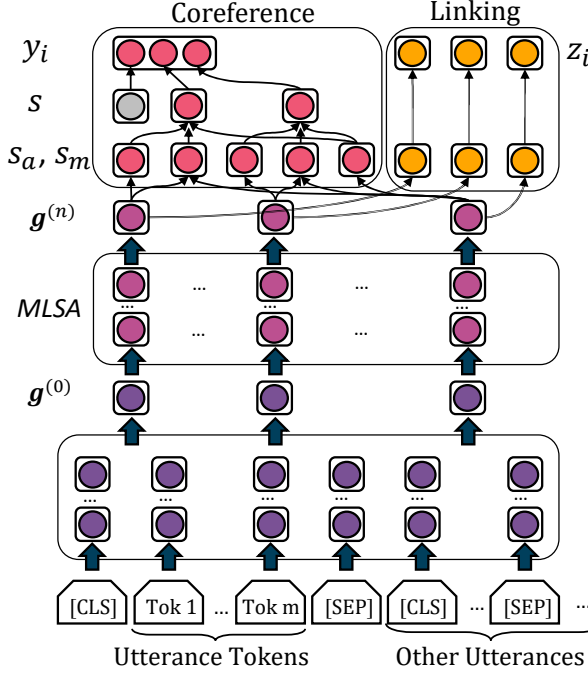


Figure 3: The coreference module and the linking module share the same mention representation $g^{(n)}$ as inputs. The mention representation $g^{(i)}$ are iteratively refined through the mention-level self-attention layers. The initial mention representations $g^{(0)}$ are the sum of text span representations from a pre-trained text encoder and corresponding speaker embeddings.

as a result, the initial representation of mention i is:

$$g_i^{(0)} = t_{start_i} + t_{end_i} + e_{speaker_i}, \quad (1)$$

where t_{start_i} and t_{end_i} are the contextualized representation of the beginning and the end tokens of mention i , and the $e_{speaker_i}$ is the speaker embedding for the current speaker. Here, we omit the embeddings of inner tokens because their semantics has been effectively encoded via the language model. The speaker embeddings are randomly initialized before training.

Sometimes the local context of a mention is not enough to make reasonable predictions, and it is observed that the co-occurred mentions can provide document-level context information. To refine the mention representations given the presence of other mentions in the document, we introduce the Mention-Level Self-Attention (MLSA) layer, which has n layers of transformer encoder structure (Vaswani et al., 2017) and is denoted as T . Formally, this iterative mention refinement process can be described by

$$g_1^{(i+1)}, \dots, g_k^{(i+1)} = T(g_1^{(i)}, \dots, g_k^{(i)}), \quad (2)$$

where k indicates the number of mentions in a document, and the $g^{(i)}$ means the mention representation from the i -th layer of MLSA.

3.2 Coreference Resolution

Following the previous work (Joshi et al., 2019a), we model the coreference resolution task as an antecedent finding problem. For each mention, we aim at finding one of the previous mentions that refer to the same person. If no such previous mention exists, it should be linked to the dummy mention ε . Thus the goal of a coreference model is to learn a distribution, $P(y_i)$ over each antecedent for each mention i :

$$P(y_i) = \frac{e^{s(i, y_i)}}{\sum_{y' \in \mathcal{Y}(i)} e^{s(i, y')}} \quad (3)$$

where $s(i, j)$ is the score for the antecedent assignment of mention i to j . The score $s(i, j)$ contains two parts: (1) the plausibility score of the mentions $s_a(i, j)$; (2) the mention score measuring the plausibility of being a proper mention $s_m(i)$. Formally, the $s(i, j)$ can be expressed by

$$s(i, j) = s_m(i) + s_m(j) + s_a(i, j), \quad (4)$$

$$s_m(i) = FFNN_m(g_i^{(n)}), \quad (5)$$

$$s_a(i, j) = FFNN_a([g_i^{(n)}, g_j^{(n)}]), \quad (6)$$

where $g^{(n)}$ stands for the last layer mention representation resulted from the MLSA and $FFNN$ indicates the feed-forward neural network.

3.3 Character Linking

The character linking is formulated as a multi-class classification problem, following previous work (Zhou and Choi, 2018). Given the mention representations $g^{(n)}$, the linking can be done with a simple feed-forward network, denoted as $FFNN(\cdot)$. Specifically, the probability of character entity z_i is linked with a given mention i can be calculated by:

$$Q(z_i) = \text{Softmax}(FFNN_l(g_i^{(n)}))_{z_i}, \quad (7)$$

where the notation $(\cdot)_z$ represents the z -th composition of a given vector.

3.4 Joint Learning

To jointly optimize both coreference resolution and entity linking, we design a joint loss of both tasks. For coreference resolution, given the gold clusters, we minimize the negative log-likelihood of

DATASET	EPISODES	SCENES	UTTERANCES	SPEAKERS	MENTIONS	ENTITIES
TRN	76	987	18,789	265	36,385	628
DEV	8	122	2,142	48	3,932	102
TST	13	192	3,597	91	7,050	165
TOTAL	97	1,301	24,528	331	47,367	781

Table 1: The detailed information about the datasets. For each season, the episode 1 to 19 are used for training, the episode 20 to 21 for development, and the remaining for testing.

the possibility that each mention is linked to a gold antecedent. Then the coreference loss L_c becomes

$$L_c = - \sum_{i=1}^N \log \sum_{y \in \mathcal{Y}(i) \cap GOLD(i)} P(y), \quad (8)$$

where the $GOLD(i)$ denotes the gold coreference cluster that mention i belongs to. Similarly, for character linking, we minimize the negative log-likelihood of the joint probability for each mention being linked to the correct referent character:

$$L_l = - \sum_{i=1}^N \log Q(z_i). \quad (9)$$

Finally, the joint loss can be the arithmetic average of the coreference loss and linking loss:

$$L = \frac{1}{2} (L_l + L_c). \quad (10)$$

4 Experiments

In this section, we introduce the experimental details to demonstrate the effectiveness of C^2 .

4.1 Data Description

We use the latest released character identification V2.0² as the experimental dataset, and we follow the standard training, developing, and testing separation provided by the dataset. In the dataset, all mentions are annotated with their referent global entities. For example, in Figure 4, the mention *I* is assigned to *ROSS*, and the mentions *mom* and *dad* are assigned to *JUDY* and *JACK* respectively in the first utterance given by *Ross*. The gold coreference clusters are derived by grouping the mentions assigned to the same character entity. Statistically, the dataset includes four seasons of the TV show *Friends*, which contain 97 episodes, 1,301 scenes, and 24,528 utterances. In total, there are 47,367 mentions, which are assigned to 781 unique characters. The detailed statistics are shown in Table 1.

²<https://github.com/emorynlp/character-identification>

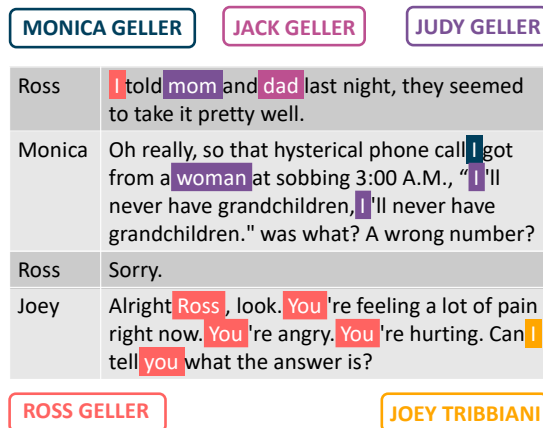


Figure 4: The example annotations for character identification. The arrows in the figure are pointing from the character mentions to their referent character entities.

4.2 Baseline Methods

The effectiveness of the joint learning model is evaluated on both the coreference resolution and character linking tasks. To fairly compare with existing models, only the singular mentions are used following the singular-only setting (S-only) in the previous work (Zhou and Choi, 2018).

For the **coreference resolution** task, we compare with the following methods.

- ACNN: A CNN-based model (Zhou and Choi, 2018) coreference resolution model that can also produce the mention and mention-cluster embeddings at the same time.
- C2F: The end-to-end coarse-to-fine coreference model (Joshi et al., 2019b) with BERT (Devlin et al., 2018) or SpanBERT (Joshi et al., 2019a) as the encoder.
- CorefQA: An approach that reformulates the coreference resolution problem as a question answering problem (Wu et al., 2020) and being able to be benefited from fine-tuned question-answer text encoders.

MODEL	B3			CEAF ϕ_4			BLANC			AVE.F1
	PREC.	REC.	F1	PREC.	REC.	F1	PREC.	REC.	F1	
ACNN	84.30	71.90	77.60	54.50	71.80	62.00	84.30	80.40	82.10	73.96 (0.97)
COREFQA (SPANBERT-LARGE)	73.72	75.55	74.62	65.82	72.38	68.94	86.82	84.69	85.75	76.44 (0.20)
C2F (BERT-BASE)	69.62	76.11	72.72	66.44	60.92	63.56	79.38	86.05	82.38	72.88 (0.23)
C2F (BERT-LARGE)	71.72	80.25	75.75	69.97	62.61	66.08	81.65	88.23	84.63	75.49 (0.18)
C2F (SPANBERT-BASE)	72.49	77.88	75.08	66.00	64.23	65.10	81.60	87.43	84.27	74.81 (0.19)
C2F (SPANBERT-LARGE)	81.93	84.38	82.57	78.04	71.99	74.89	88.15	91.09	89.56	82.34 (0.17)
C ² (BERT-BASE)	78.10	81.56	79.79	72.48	69.87	71.15	86.14	89.49	87.74	80.14 (0.21)
C ² (BERT-LARGE)	78.49	81.90	80.16	73.81	71.15	72.46	86.20	89.93	87.97	80.17 (0.23)
C ² (SPANBERT-BASE)	81.18	83.59	82.36	73.64	73.09	73.36	88.06	91.04	89.49	81.74 (0.19)
C ² (SPANBERT-LARGE)	85.83	85.27	85.55	77.13	77.84	77.48	92.31	92.03	92.17	85.06 (0.16)

Table 2: Experimental results on the coreference resolution task. The results are presented in a 2-digit decimal following previous work. Standard deviations of the average F1 scores are shown in brackets.

For the **character linking** task, we also include ACNN as a baseline method. Considering existing general entity linking models (Kolitsas et al., 2018; van Hulst et al., 2020; Raiman and Raiman, 2018; Onando Mulang et al., 2020) cannot be applied to the character linking problem because they are not designed to handle pronouns, we propose another text-span classification model with transformer encoder as another strong baseline for the character linking task.

- ACNN: A model that uses the mention and mention-cluster embeddings as input to do character linking (Zhou and Choi, 2018).
- BERT/SpanBERT: A text-span classification model consists of a transformer text encoder followed by a feed-forward network.

4.3 Evaluation Metrics

We follow the previous work (Zhou and Choi, 2018) for the evaluation metrics. Specifically, for coreference resolution, three evaluation metrics, B³, CEAF ϕ_4 , and BLANC, are used. The metrics are all proposed by the CoNLL’12 shared task (Pradhan et al., 2012) to evaluate the output coreference cluster against the gold clusters. We follow Zhou and Choi (2018) to use BLANC (Recasens and Hovy, 2011) to replace MUC (Vilain et al., 1995) because BLANC takes singletons into consideration but MUC does not. As for the character linking task, we use the Micro and Macro F1 scores to evaluate the multi-class classification performance.

4.4 Implementation Details

In our experiments, we consider four different pre-trained language encoders: BERT-Base, BERT-Large, SpanBERT-Base, and SpanBERT-Large,

and we use $n = 2$ layers of the mention-level self-attention (MLSA). The feed-forward networks are implemented by two fully connected layers with ReLU activations. Following the previous work, (Zhou and Choi, 2018), the scene-level setting is used, where, each scene is regarded as a document for coreference resolution and linking. During the training, each mini-batch consists of segments obtained from a single document. The joint learning model is optimized with the Adam optimizer (Kingma and Ba, 2015) with an initial learning rate of 3e-5, and a warming-up rate of 10%. The model is set to be trained for 100 epochs with an early stop. All the experiments are repeated three times, and the average results are reported.

5 Results and Analysis

In this section, we discuss the experimental results and present a detailed analysis.

5.1 Coreference Resolution Results

The performances of coreference resolution models are shown in Table 2. C² with SpanBERT-large achieves the best performance on all evaluation metrics. Comparing to the baseline ACNN model, which uses hand-crafted features, C² uses a transformer to better encode the contextual information. Besides that, even though ACNN formulates the coreference resolution and character linking tasks in a pipe-line and uses the coreference resolution result to help character linking, the character linking result cannot be used to help to resolve coreference clusters. As a comparison, we treat both tasks jointly such that they can help each other.

Currently, CorefQA is the best-performing general coreference resolution model on the OntoNotes dataset (Pradhan et al., 2012). However, its performance is limited on the conversation dataset due to

MODEL	Ro	RA	Ch	Mo	Jo	PH	EM	Ri	MICRO	MACRO
ACNN	78.3	86.5	78.8	81.7	78.3	88.8	69.2	83.9	73.7 (0.6)	59.6 (2.3)
BERT-BASE	87.4	89.9	86.6	88.2	87.1	91.1	94.3	62.4	84.0 (0.1)	77.3 (0.2)
BERT-LARGE	88.2	89.9	87.9	88.8	87.7	93.1	93.5	68.0	84.8 (0.2)	79.1 (0.2)
SPANBERT-BASE	87.6	91.8	86.7	88.2	86.8	92.6	94.6	73.3	84.2 (0.1)	77.3 (0.2)
SPANBERT-LARGE	90.9	92.8	88.3	90.3	90.2	94.3	94.6	71.7	85.5 (0.1)	79.8 (0.2)
C^2 (BERT-BASE)	86.5	87.8	85.6	86.8	88.1	92.4	93.0	66.0	84.0 (0.1)	78.6 (0.2)
C^2 (BERT-LARGE)	85.9	90.0	87.3	86.9	87.2	93.0	96.1	66.0	84.9 (0.1)	79.5 (0.2)
C^2 (SPANBERT-BASE)	89.8	91.3	90.5	90.9	87.8	93.2	93.4	71.3	85.7 (0.1)	81.0 (0.1)
C^2 (SPANBERT-LARGE)	91.2	94.1	91.1	92.5	90.4	94.4	89.2	77.1	87.0 (0.1)	81.1 (0.1)

Table 3: Experimental results per character on the character linking. The results are presented in a 1-digit decimal following previous work. Standard deviations of the Micro and Macro F1 scores are shown in brackets. The names in the table are written in two-letter acronyms. Ro: Ross, Ra: Rachel, Ch: Chandler, Mo: Monica, Jo: Joey, Ph: Phoebe, Em: Emily, Ri: Richard

two reasons. First, different from the experimental setting of OntoNotes, the mentions in our experiment setting are gold mentions. Consequently, the flexible span predicting strategy of CorefQA loses its advantages because of the absence of the mention proposal stage. Second, the CorefQA leverages the fine-tuning on other question answering (QA) datasets and it is possible that the used QA dataset (i.e., SQuAD-2.0 (Rajpurkar et al., 2018)) is more similar to OntoNotes rather than the used multiparty conversation dataset, which is typically much more informal. As a result, the effect of such fine-tuning process only works on OntoNotes.

The coarse-to-fine (C2F) model (Joshi et al., 2019b) with a transformer encoder was the previous state-of-the-art model on OntoNotes. Referring to Table 2, given the same text encoder, the proposed C^2 model can constantly outperform the C2F model. These results further demonstrate that with the help of the proposed joint learning framework, the out-of-context character information can help achieve better mention representations so that the coreference models can resolve them more easily.

5.2 Character Linking Results

As shown in Table 3, the proposed joint learning model also achieves the best performance on the character linking task and there are mainly two reasons for that. First, the contextualized mention representations obtained from pre-trained language encoders can better encode the context information than those representations used in ACNN. Second, with the help of coreference clusters, richer context about the whole conversation is encoded for each mention. For example, when using the same pre-trained language model as the encoder, C^2 can always outperform the baseline classification model. These empirical results confirm that, though the

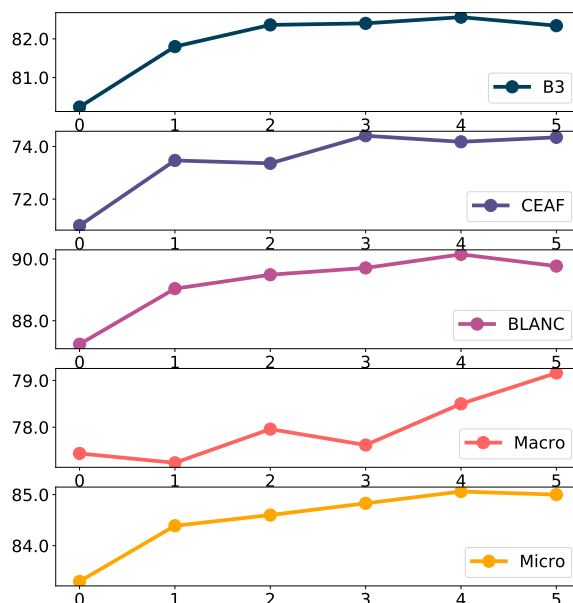


Figure 5: The x-axis is the number of MLSA layers used in the C^2 . The y-axes are the F1 scores on each metric for their corresponding tasks. The curves have general trends of going up, which indicates that the model performs better when there are more layers.

BERT and SpanBERT can produce very good vector representation for the mentions based on the local context, the coreference clusters can still provide useful document-level contextual information for linking them to a global character entity.

5.3 The Number of MLSA Layers

Another contribution of the proposed C^2 model is the proposed mention-level self-attention (MLSA) module, which helps iteratively refine the mention representations according to the other mentions co-occurred within the same document. In this section, to show its effect and the influence of iteration layers, we tried different layers and show their per-

1	Chandler	Oh, Danielle! I wasn't expecting the machine... Give me a call when you get a chance. Bye-bye. Oh God!	DANIELLE
2	Monica	That's what you've been working on for the past two hours?!	
3	Chandler	Hey, I've been honing!	
4	Ross	What was with the dishes?	CHANDLER
5	Chandler	Oh, uh.. I want her to think I might be in a restaurant.. y'know? I might have some kind of life, like I haven't been sitting around here honing for the past few hours.	
6	Monica	Look look! He's doing it again, the guy with the telescope!	MAN 1
7	Phoebe	Oh my God! Go away! Stop looking in here!	
8	Monica	"Great, now he's waving back."	
9	Joey	Man, we gotta do something about that guy. This morning, I caught him looking into our apartment. It creeps me out! I feel like I can't do stuff!	MONICA
10	Monica	What kinda stuff?	
11	Joey	Will you grow up? I'm not talking about sexy stuff, but, like, when I'm cooking naked.	
12	Phoebe	You cook naked?	JOEY
13	Joey	Yeah, toast, oatmeal... nothing that spatters.	
14	Chandler	What are you looking at me for? I didn't know that.	

Figure 6: Case study. All mentions that are linked to the same character and in the same coreference cluster are highlighted with the same color. The misclassified mention is marked with the red cross.

MODEL	COREFERENCE F1			LINKING F1	
	B3	CEAF ϕ 4	BLANC	MICRO	MACRO
C ²	85.54	77.48	92.17	87.05	81.09
- MLSA	83.57	75.32	90.51	86.26	80.32
- LINKING	83.50	76.10	90.08	-	-
- COREF.	-	-	-	86.94	79.58

Table 4: Three ablation studies are conducted concerning the MLSA layers, the coreference resolution module, and the character linking module.

performances on the test set in Figure 5. We conducted the experiments with the SpanBERT-Base encoder and all other hyper-parameters are the same. The x-axis is the number of layers, and the y-axes are F1 scores of B3, CEAF, and BLANC for coreference resolution, the Macro and Micro F1 scores for character linking. From the results, we can see that with the increase of layer number from zero to five, the F1 scores on both tasks gradually increase. This trend demonstrates that the model can perform better on both tasks when there are more layers. Meanwhile, the marginal performance improvement of the MLSA layer is decreasing. This indicates that adding too many layers of MLSA may not further help improve the performance because enough context has been included. Considering the balance between performance and computational efficiency, we chose the iteration layers to be two in our current model based on similar observations made on the development set.

5.4 Ablation Study

In this section, we present the ablation study to clearly show the effect of different modules in the proposed framework C² in Table 4. First, we try to remove the mention-level self-attention (MLSA) from our joint learning model and a clear performance drop is observed on both tasks. Specifically, the performance on coreference resolution is reduced by 1.21 on the average F1, and meanwhile, the macro-F1 and micro-F1 scores on character linking decreased by 0.77 and 0.79 respectively. The reduction reveals that the MLSA indeed helps achieve better mention representations with the help from co-occurred mentions. Second, we try to remove the coreference resolution and character linking modules. When the character linking module is removed, it is observed that the performance on coreference resolution decreased by 1.94 on the averaged F1 score. When the coreference module is removed, the performance of C² on character linking dropped by 0.83 on the average of Micro and Macro F1 scores. These results prove that the modeling of coreference resolution and character linking can indeed help each other and improve the performance significantly, and the proposed joint learning framework can help to achieve that goal.

5.5 Case Study

Besides the quantitative evaluation, in this section, we present the case study to qualitatively evaluate

the strengths and weaknesses of the proposed C^2 model. As shown in Figure 6, we randomly select an example from the development set to show the prediction results of the proposed model on both tasks. To illustrate the coreference resolution and character linking results from the C^2 model, the mentions from the same coreference cluster are highlighted with the same color. Also, we use the same color to indicate to which character the mentions are referring. Meanwhile, the falsely predicted result is marked with a red cross.

5.5.1 Strengths

For this example, the results on both tasks are consistent. The mentions that are linked to the same character entity are in the same coreference group and vice versa. Based on this observation and previous experimental results, it is more convincing that the proposed model can effectively solve the two problems at the same time. Besides that, we also notice that the model does not overfit the popular characters. It can correctly solve all the mentions referring to not only main characters, and also for the characters that only appear several times such as *MAN I*. Last but not least, the proposed model can correctly resolve the mention to the correct antecedent even though there is a long distance between them in the conversation. For example, the mention *me* in utterance 14 can be correctly assigned to the mention *you* in utterance 2, though there are 11 utterances in between. It shows that by putting two tasks together, the proposed model can better utilize the whole conversation context. The only error made by the model is incorrectly classifying a mention and at the same time putting it into a wrong coreference cluster.

5.5.2 Weaknesses

By analyzing the error case, it is noticed that the model may have trouble in handling the mentions that require common sense knowledge. Humans can successfully resolve the mention *her* to Danielle because they know Danielle is on the other side of the telephone, but Monica is in the house. As a result, Chandler can only deceive Danielle but not Monica. But the current model, which only relies on the context, cannot tell the difference.

5.6 Error Analysis

We use the example in Figure 6 to emphasize the error analysis that compares the performance of our model and the baseline models. The details are as

follows. In this example, the only mistake made by our model is related to common-sense knowledge, and the baseline models are also not able to make a correct prediction.

For coreference resolution, 3 out of 25 mentions are put into a wrong cluster by the c2f baseline model. The baseline model failed to do long-distance antecedent assignments (e.g., the “me” in utterance 14). Meanwhile, our model is better in this case because it successfully predicts the antecedent of the mention “me”, even though its corresponding antecedent is far away in utterance 2. This example demonstrates the advantage that our joint model can use global information obtained from character linking to better resolve the co-referents that are far away from each other.

For character linking, 2 out of 25 mentions are linked to the wrong characters by the baseline model. It is observed that the baseline model cannot consistently make correct linking predictions to less-appeared characters, for example, the “He” in utterance 6. In this case, our model performs better mainly because it can use the information gathered from the nearby co-referents to adjust its linking prediction, as its nearby co-referents are correctly linked to corresponding entities.

6 Related Works

Coreference resolution is the task of grouping mentions to clusters such that all the mentions in the same cluster refer to the same real-world entity (Pradhan et al., 2012; Zhang et al., 2019a,b; Yu et al., 2019). With the help of higher-order coreference resolution mechanism (Lee et al., 2018) and strong pre-trained language models (e.g., SpanBERT (Joshi et al., 2019b)), the end-to-end based coreference resolution systems have been achieving impressive performance on the standard evaluation dataset (Pradhan et al., 2012). Recently, motivated by the success of the transfer learning, Wu et al. (2020) propose to model the coreference resolution task as a question answering problem. Through the careful fine-tuning on a high-quality QA dataset (i.e., SQUAD-2.0 (Rajpurkar et al., 2018)), it achieves the state-of-the-art performance on the standard evaluation benchmark. However, as disclosed by Zhang et al. (2020), current systems are still not perfect. For example, they still cannot effectively handle pronouns, especially those in informal language usage scenarios like conversations. In this paper, we propose to leverage the out-of-

context character information to help resolve the coreference relations with a joint learning model, which has been proven effective in the experiments.

As a traditional NLP task, entity linking (Mihalcea and Csomai, 2007; Ji et al., 2015; Kolitsas et al., 2018; Raiman and Raiman, 2018; Onando Mulang et al., 2020; van Hulst et al., 2020) aims at linking mentions in context to entities in the real world (typically in the format of knowledge graph). Typically, the mentions are named entities and the main challenge is the disambiguation. However, as a special case of the entity linking, the character linking task has its challenge that the majority of the mentions are pronouns. In the experiments, we have demonstrated that when the local context is not enough, the richer context information provided by the coreference clusters could be very helpful for linking mentions to the correct characters.

In the NLP community, people have long been thinking that the coreference resolution task and entity linking should be able to help each other. For example, Ratnoff and Roth (2012) show how to use knowledge from named-entity linking to improve the coreference resolution, but do not consider doing it in a joint learning approach. After that, Hajishirzi et al. (2013) demonstrate that the coreference resolution and entity linking are complementary in terms of reducing the errors in both tasks. Motivated by these observations, a joint model for coreference, typing, and linking is proposed (Durrett and Klein, 2014) to improve the performance on three tasks at the same time. Compared with previous works, the main contributions of this paper are two-fold: (1) we tackle the challenging character linking problem; (2) we design a novel mention representation encoding method, which has been shown effective on both the coreference resolution and character linking tasks.

7 Conclusion

In this paper, we propose to solve the coreference resolution and character linking tasks jointly. The experimental results show that the proposed model C^2 performs better than all previous models on both tasks. Detailed analysis is also conducted to show the contribution of different modules and the effect of the hyper-parameter.

8 Acknowledgements

This paper was supported by the NSFC Grant U20B2053 from China, the Early Career Scheme

(ECS, No. 26206717), the General Research Fund (GRF, No. 16211520), and the Research Impact Fund (RIF, No. R6020-19) from the Research Grants Council (RGC) of Hong Kong, with special thanks to the Tencent AI Lab Rhino-Bird Focused Research Program.

References

- Henry Y. Chen, Ethan Zhou, and Jinho D. Choi. 2017. Robust coreference resolution and entity linking on dialogues: Character identification on TV show transcripts. In *Proceedings of CoNLL*, pages 216–225.
- Yu-Hsin Chen and Jinho D. Choi. 2016. Character identification on multiparty conversation: Identifying mentions of characters in TV shows. In *Proceedings of SIGDIAL*, pages 90–100.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Greg Durrett and Dan Klein. 2014. A joint model for entity analysis: Coreference, typing, and linking. *Transactions of the association for computational linguistics*, 2:477–490.
- Hannaneh Hajishirzi, Leila Zilles, Daniel S. Weld, and Luke Zettlemoyer. 2013. Joint coreference resolution and named-entity linking with multi-pass sieves. In *Proceedings of the EMNLP 2013*, pages 289–299.
- Johannes M. van Hulst, Faegheh Hasibi, Koen Dercksen, Krisztian Balog, and Arjen P. de Vries. 2020. Rel: An entity linker standing on the shoulders of giants. *Proceedings of SIGIR 2020*.
- Heng Ji, Joel Nothman, Ben Hachey, and Radu Florian. 2015. Overview of tac-kbp2015 tri-lingual entity discovery and linking. *Theory and Applications of Categories*.
- Mandar Joshi, Danqi Chen, Y. Liu, Daniel S. Weld, L. Zettlemoyer, and Omer Levy. 2019a. Spanbert: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Mandar Joshi, Omer Levy, Daniel S. Weld, and Luke Zettlemoyer. 2019b. BERT for coreference resolution: Baselines and analysis. In *Proceedings of EMNLP*.
- Ben Kantor and Amir Globerson. 2019. Coreference resolution with entity equalization. In *Proceedings of ACL 2019*, pages 673–677.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of ICLR 2015*.

- Nikolaos Kolitsas, Octavian-Eugen Ganea, and Thomas Hofmann. 2018. End-to-end neural entity linking. In *Proceedings of CoNLL 2018*, pages 519–529.
- Kenton Lee, Luheng He, and Luke Zettlemoyer. 2018. Higher-order coreference resolution with coarse-to-fine inference. In *Proceedings of NAACL (Short Papers)*, pages 687–692.
- Rada Mihalcea and Andras Csomai. 2007. Wikify! linking documents to encyclopedic knowledge. *CIKM '07*, page 233–242.
- Isaiah Onando Mulang, Kuldeep Singh, Chaitali Prabh, Abhishek Nadgeri, Johannes Hoffart, and Jens Lehmann. 2020. Evaluating the impact of knowledge graph context on entity disambiguation models. *arXiv e-prints*, pages arXiv–2008.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. Conll-2012 shared task: Modeling multilingual unrestricted coreference in ontonotes. pages 1–40.
- Jonathan Raiman and Olivier Raiman. 2018. Deep-type: multilingual entity linking by neural type system evolution. *arXiv preprint arXiv:1802.01021*.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don't know: Unanswerable questions for squad. In *Proceedings of ACL 2018*, pages 784–789.
- Lev Ratinov and Dan Roth. 2012. Learning-based multi-sieve co-reference resolution with knowledge. In *Proceedings of EMNLP 2012*, pages 1234–1244.
- Marta Recasens and Eduard Hovy. 2011. Blanc: Implementing the rand index for coreference evaluation. *Natural Language Engineering*, 17:485 – 510.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, L. Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *ArXiv*, abs/1706.03762.
- Marc Vilain, John Burger, John Aberdeen, Dennis Connolly, and Lynette Hirschman. 1995. A model-theoretic coreference scoring scheme. *Proceedings of the 6th conference on message understanding*, pages 45–52.
- Wei Wu, Fei Wang, Arianna Yuan, Fei Wu, and Jiwei Li. 2020. CorefQA: Coreference resolution as query-based span prediction. In *Proceedings ACL 2020*, pages 6953–6963.
- Xintong Yu, Hongming Zhang, Yangqiu Song, Yan Song, and Changshui Zhang. 2019. What you see is what you get: Visual pronoun coreference resolution in dialogues. In *Proceedings of EMNLP-IJCNLP 2019*, pages 5122–5131.
- Hongming Zhang, Yan Song, and Yangqiu Song. 2019a. Incorporating context and external knowledge for pronoun coreference resolution. In *Proceedings of NAACL-HLT 2019*, pages 872–881.
- Hongming Zhang, Yan Song, Yangqiu Song, and Dong Yu. 2019b. Knowledge-aware pronoun coreference resolution. In *Proceedings of ACL 2019*, pages 867–876.
- Hongming Zhang, Xinran Zhao, and Yangqiu Song. 2020. A brief survey and comparative study of recent development of pronoun coreference resolution. *CoRR*, abs/2009.12721.
- Ethan Zhou and Jinho D. Choi. 2018. They exist! introducing plural mentions to coreference resolution and entity linking. In *Proceedings of ICCL*, pages 24–34.