

GECOR: An End-to-End Generative Ellipsis and Co-reference Resolution Model for Task-Oriented Dialogue

Jun Quan^{†*}, Deyi Xiong^{‡†}, Bonnie Webber[§] and Changjian Hu[¶]

[‡] School of Computer Science and Technology, Soochow University, Suzhou, China

[§] University of Edinburgh, Edinburgh, UK

[¶] Lenovo Research AI Lab, Beijing, China

terryqj0107@gmail.com, dyxiong@suda.edu.cn,

bonnie.webber@ed.ac.uk, hucj1@lenovo.com

Abstract

Ellipsis and co-reference are common and ubiquitous especially in multi-turn dialogues. In this paper, we treat the resolution of ellipsis and co-reference in dialogue as a problem of generating omitted or referred expressions from the dialogue context. We therefore propose a unified end-to-end Generative Ellipsis and CO-reference Resolution model (GECOR) in the context of dialogue. The model can generate a new pragmatically complete user utterance by alternating the generation and copy mode for each user utterance. A multi-task learning framework is further proposed to integrate the GECOR into an end-to-end task-oriented dialogue. In order to train both the GECOR and the multi-task learning framework, we manually construct a new dataset on the basis of the public dataset CamRest676 with both ellipsis and co-reference annotation. On this dataset, intrinsic evaluations on the resolution of ellipsis and co-reference show that the GECOR model significantly outperforms the sequence-to-sequence (seq2seq) baseline model in terms of EM, BLEU and F_1 while extrinsic evaluations on the downstream dialogue task demonstrate that our multi-task learning framework with GECOR achieves a higher success rate of task completion than TSCP, a state-of-the-art end-to-end task-oriented dialogue model (Lei et al., 2018).

1 Introduction

Due to the rhetorical principle of saving words and avoiding repetitions, ellipsis and co-reference occur frequently in multi-turn dialogues leaving utterances pragmatically incomplete if they are separate from context. Humans can easily understand utterances with anaphorically referenced or absent

information (e.g., Q_2 and Q_3 in Table 1) based on the dialogue context while dialogue systems often fail to understand such utterances correctly, which may result in false or incoherent responses.

If user utterances can be automatically supplemented with information that is left out or substituted by anaphora according to the dialogue context as humans do (e.g., Q_2 : *I want cheap Italian restaurants.* Q_3 : *Yes, I would like the phone number please.*), dialogue models may understand user requests correctly and would not generate wrong responses caused by ellipsis and co-reference phenomena. Especially in task-oriented dialogue systems, explicitly providing such information to the models can effectively improve the success rate of task completion.

In order to achieve this goal, we propose an end-to-end generative ellipsis and co-reference resolution model (GECOR) for task-oriented dialogue in this paper. The essential idea behind GECOR is that we treat the resolution of ellipsis and co-reference in user utterances as a generation task: transforming a user utterance with ellipsis or anaphora into a new utterance where the left-out or referred expressions are automatically generated from the dialogue context. We refer to the new utterance as the complete version of the original utterance. We use an end-to-end sequence-to-sequence model with two encoders for this transformation task, where one encoder reads the user utterance and the other the dialogue context and the decoder generates the complete utterance. Since most omitted expressions or antecedents can be found in the dialogue context, we resort to the attention and copy mechanism to detect such fragments in previous context and copy them into the generated complete utterance.

We then incorporate GECOR into an end-to-end task-oriented dialogue system in a multi-task

* Work performed during an internship at Lenovo Research AI Lab.

† Corresponding author

learning framework. The entire model contains two encoders (one for user utterance and the other for the dialogue context) and three decoders: one decoder for predicting dialogue states, the second decoder for generating complete user utterances and the third decoder for generating system responses. The three decoders are jointly trained.

In order to train GECOR with the task-oriented dialogue model, we manually annotate the public task-oriented dialogue dataset CamRest676 with omitted expressions and substitute anaphora in the dataset with corresponding antecedents. The new dataset can be used either to train a stand-alone ellipsis or co-reference resolution model or to jointly train a task-oriented dialogue model equipped with the ellipsis / co-reference resolution model.

We conduct a series of experiments and analyses, demonstrating that the proposed method can significantly outperform a strong baseline model. Our contributions are threefold:

- We propose an end-to-end generative resolution model that attempts to solve the ellipsis and co-reference resolution in a single unified framework, significantly different from previous end-to-end co-reference resolution network with two phases of detection and candidate ranking.
- To the best of our knowledge, this is the first attempt to combine the task of ellipsis and co-reference resolution with the multi-turn task-oriented dialogue. The success rate of task completion is significantly improved with the assistance of the ellipsis and co-reference resolution.
- We construct a new dataset based on CamRest676 for ellipsis and co-reference resolution in the context of task-oriented dialogue.¹

2 Related Work

Ellipsis recovery: The earliest work on ellipsis as far as we know is the PUNDIT system (Palmer et al., 1986) which discusses the communication between the syntactic, semantic and pragmatic modules that is necessary for making implicit linguistic information explicit. Dalrymple et al. (1991) and Shieber et al. (1996) establish

¹The new dataset and the code of our proposed system are available at <https://multinlp.github.io/GECOR/>

a set of linguistic theories in the ellipsis recovery of English verb phrases. Nielsen (2003) first proposes an end-to-end computable system to perform English verb phrase ellipsis recovery on the original input text. Liu et al. (2016) propose to decompose the resolution of the verb phrase ellipsis into three sub-tasks: target detection, antecedent head resolution, and antecedent boundary detection.

Co-reference resolution: Co-reference resolution is mainly concerned with two sub-tasks, referring expressions (i.e., mentions) detection, and entity candidate ranking. Uryupina and Moschitti (2013) propose a rule-based approach for co-reference detection which employs parse tree features with an SVM model. Peng et al. (2015) improve the performance of mention detection by applying a binary classifier on their feature set. In recent years, applying deep neural networks to the co-reference resolution has gained great success. Clark and Manning (2016) apply reinforcement learning on mention-ranking co-reference resolution. Lee et al. (2017) introduce the first end-to-end co-reference resolution model. Lee et al. (2018) present a high-order co-reference resolution model with coarse-to-fine inference.

Ellipsis and co-reference resolution in QA and Dialogue: The methods mentioned above do not generalize well to dialogues because they normally require a large amount of well-annotated contextual data with syntactic norms and candidate antecedents. In recent years, a few studies try to solve ellipsis / co-reference resolution tailored for dialogue or QA tasks. Kumar and Joshi (2016) train a semantic sequence model to learn semantic patterns and a syntactic sequence model to learn linguistic patterns to tackle with the non-sentential (incomplete) questions in a question answering system. Zheng et al. (2018) builds a seq2seq neural network model for short texts to identify and recover ellipsis. However, these methods are still limited to short texts or one-shot dialogues. Our work is the first attempt to provide both solution and dataset for ellipsis and co-reference resolution in multi-turn dialogues.

End-to-end task-oriented dialogue: Task-oriented dialogue systems have evolved from traditional modularized pipeline architectures (Rudnicky et al., 1999; Zue et al., 2000; Zue and Glass, 2000) to recent end-to-end neural frameworks (Eric and Manning, 2017a,b; Lei et al.,

Turn	Dialogue
Q ₁	I would like an Italian restaurant.
A ₁	What price range do you have in mind?
Q ₂	I want cheap ones.
A ₂	Pizza Hut Cherry Hinton serves Italian food in the south part of town. Would you like their phone number?
Q ₃	Yes, please.
User utterances after resolution	
Q ₂	I want cheap Italian restaurants.
Q ₃	Yes, I would like the phone number please.

Table 1: Examples of ellipsis and co-reference resolution

2018; Jin et al., 2018). Our work is an innovative combination of ellipsis and co-reference resolution and the end-to-end task-oriented dialogue.

3 The GECOR Model

In this section, we reformulate the ellipsis and co-reference resolution task in the context of multi-turn dialogue and detail the proposed GECOR model.

3.1 Ellipsis and Co-Reference Resolution Reformulation

Our task is to reconstruct a pragmatically complete utterance from a user utterance where the ellipsis and/or co-reference phenomena are present according to the dialogue context. Table 1 provides examples of reconstructed utterances in which the omitted information is recovered or the anaphor is substituted with referred expressions.

We attempt to solve the resolution of ellipsis and co-reference in a unified framework because in essence both ellipsis and co-reference can be understood from contextual clues. We consider these two problems in multi-turn dialogue and reformulate the resolution of them as a generation problem: generating the omitted or referred expressions. In this way, the modeling of ellipsis and co-reference is in line with response generation in dialogue modeling.

Unlike previous methods that combine detection and ranking models, our generation-based formulation is not constrained by the syntactic forms of ellipsis or co-reference in sentences. They can be either words (e.g., *noun*, *verb*) or phrases or even clauses. Furthermore, the formulation does not need to provide a set of candidate antecedents to be resolved. Previous studies usually need to

traverse the text when there are multiple ellipsis or anaphora to be resolved, which leads to a high computational complexity.

In this reformulation, we assume that the dialogue context is composed of all utterances from the beginning of the dialogue to the current user utterance. Both the context and the user utterance in question are input to the GECOR model to generate the complete version of the user utterance.

3.2 Model Structure

The GECOR model is shown in Figure 1. The model essentially contains an embedding module, a user utterance encoder, a dialogue context encoder and a decoder with either copy (Gu et al., 2016) or gated copy mechanism (modified from See et al. (2017)). Both the generation probability over the entire vocabulary and the copy probability over all words from the dialogue context are taken into account for predicting the complete user utterance.

Embedding Layer In GECOR, we first tokenize the input user utterance and the dialogue context. We then use GloVe (Pennington et al., 2014) (the pre-trained 50-dimensional word vectors) in the embedding layer to obtain word embeddings. Let $\mathbf{U} = \{u_1, \dots, u_m\}$, $\mathbf{C} = \{c_1, \dots, c_n\}$ be representations of the tokenized utterance and context sequence.

Utterance and Context Encoder We use a single-layer bidirectional GRU to construct both encoders. The forward and backward hidden states over the input embeddings from the embedding layer are concatenated to form the hidden states of the two encoders.

Decoder The decoder is a single-layer unidirectional GRU. In the decoder, the attention distribution a^t is calculated as in Bahdanau et al. (2015):

$$e_i^t = v^T \tanh(W_h h_i + W_s s_{t-1} + b_{attn}) \quad (1)$$

$$a^t = \text{softmax}(e^t) \quad (2)$$

where v , W_h , W_s and b_{attn} are learnable parameters, h_i is the hidden state for word u_i from the sequence produced by the utterance encoder. The attention distribution a^t is used to produce a weighted sum of the encoder hidden states, known as the context vector h_t^* :

$$h_t^* = \sum_i a_i^t h_i \quad (3)$$

It is fed into the single-layer unidirectional GRU together with the previous decoder state s_t and the word embedding y_{t-1} of the previously generated

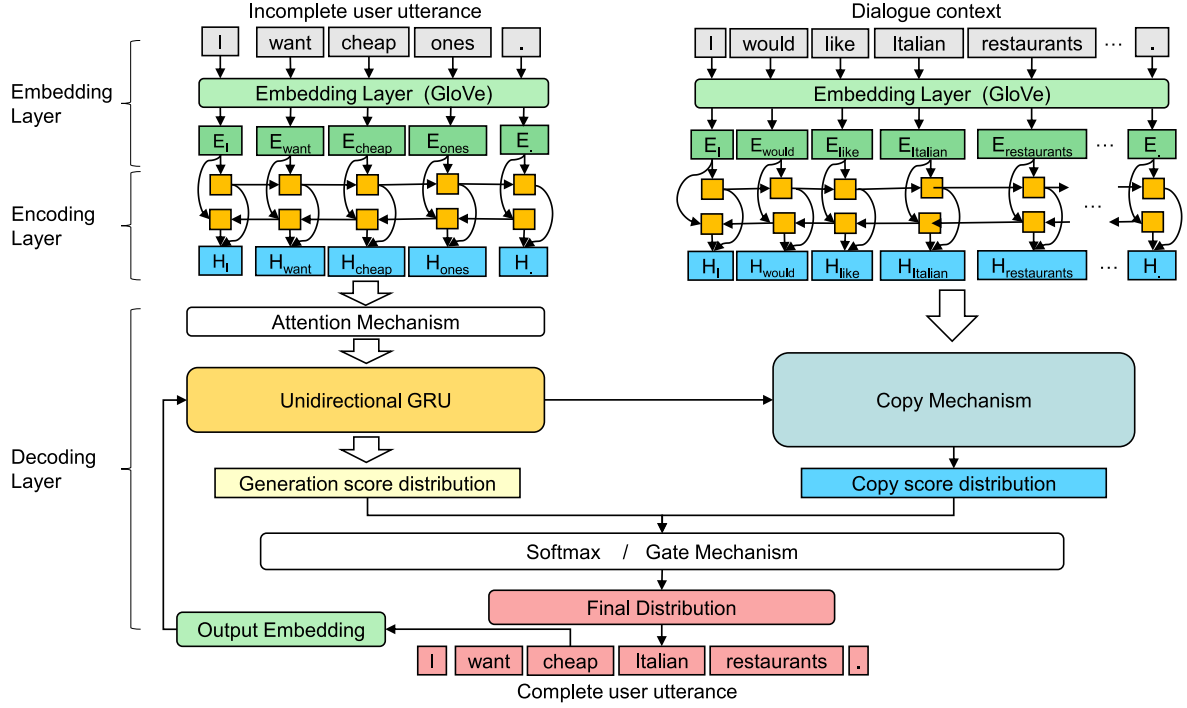


Figure 1: The end-to-end generative model for ellipsis and co-reference resolution (GECOR).

word to obtain the decoder state s_t . The updated s_{t-1} is then concatenated with the context vector h_t^* to produce the generation probability distribution over the vocabulary \mathbf{V} as follows:

$$P^g(y_t) = \frac{1}{Z} e^{\psi_g(y_t)}, \quad y_t \in \mathbf{V} \quad (4)$$

$$\psi_g(y_t = v_i) = \mathbf{v}_i^T (W_g^h h_t^* + W_g^s s_t + b_g) \quad (5)$$

$$s_t = GRU([y_{t-1}; h_t^*], s_{t-1}) \quad (6)$$

where W_g^h , W_g^s and b_g are learnable parameters and \mathbf{v}_i is the one-hot indicator vector for word $v_i \in \mathbf{V}$. ψ_g is the score function for the generation-mode and Z is the normalization term shared by the generation-mode and copy-mode.

Copy Network The copy network (Gu et al., 2016) is used to calculate the probabilities for words copied from the dialogue context. These words are parts of the omitted or referred expressions to be predicted. We build the copy network on the top of the context encoder. The probability for copying each word from the dialogue context is computed as follows:

$$P^c(y_t) = \frac{1}{Z} \sum_{i:c_i=y_t}^{|C|} e^{\psi_c(c_i)}, \quad y_t \in \mathbf{C} \quad (7)$$

$$\psi_c(y_t = c_i) = \sigma(W_c h_i^c + b_c) s_t \quad (8)$$

where W_c and b_c are learnable parameters, h_i^c is the output for word c_i from the context encoder,

and σ is a non-linear activation function. ψ_c is the score function for the copy-mode and Z is the normalization term shared by equation (4) and (7).

Both probabilities from the two modes contribute to the final probability distribution over the extended vocabulary (the vocabulary plus the words from the dialogue context) which is calculated as follows:

$$P(y_t) = P^g(y_t) + P^c(y_t), \quad y_t \in \mathbf{V} \cup \mathbf{C} \quad (9)$$

which is used to predict the final output word.

Gated Copy An alternative to the copy network is the gated copy mechanism that use a gate to regulate the contributions of the generation and copy mode to the final prediction. The gate p_{gen} is calculated as follows:

$$p_{gen} = \sigma(W_h h_t^* + W_s s_t + W_y y_{t-1} + b_t) \quad (10)$$

$$P(y_t) = p_{gen} P^g(y_t) + (1 - p_{gen}) P^c(y_t) \quad (11)$$

where W_h , W_s , W_y and b_t are learnable parameters and σ is the sigmoid function.

Training The standard cross-entropy loss is adopted as the loss function to train the GECOR model.

4 Task-Oriented Dialogue with GECOR

We integrate the proposed GECOR into an end-to-end task-oriented dialogue system TSCP proposed

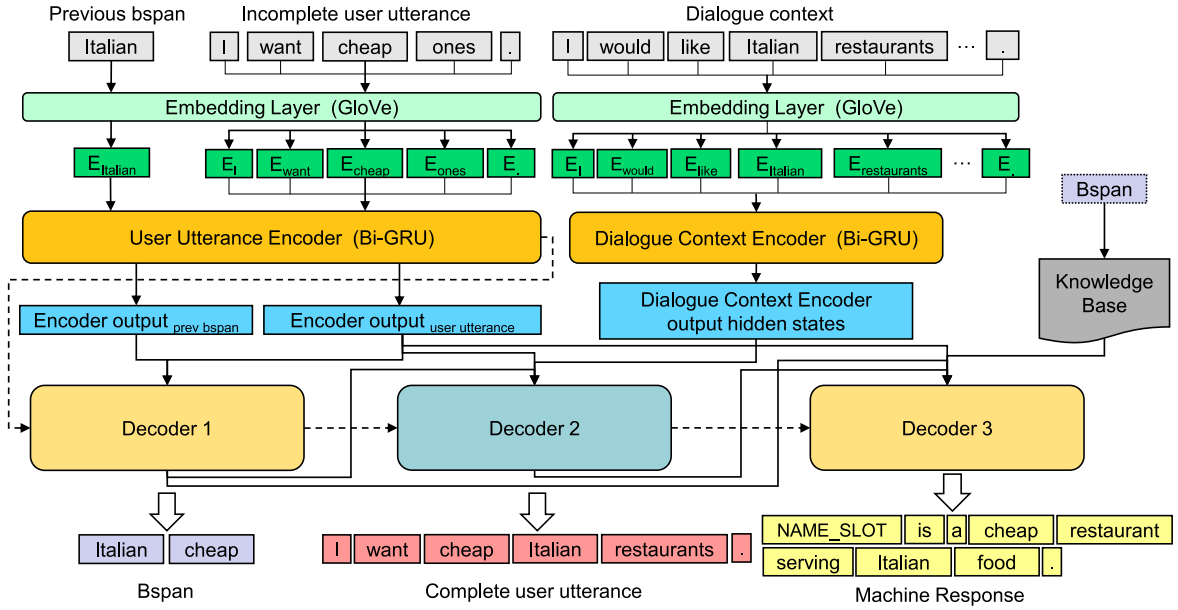


Figure 2: The architecture of the end-to-end task-oriented dialogue enhanced with the GECOR model. Decoder 1: BSpan decoder. Decoder 2: completed user utterance decoder. Decoder 3: machine response decoder.

by Lei et al. (2018) in a multi-task learning framework, which is shown in Figure 2. The GECOR-equipped TSCP model contains the embedding layer, the utterance and context encoders, and three decoders: decoder 1 for generating belief spans (BSpan) defined in (Lei et al., 2018) which are text spans for tracking dialogue states (e.g., $\langle inf \rangle Italian, cheap \langle /inf \rangle$; $\langle req \rangle phone \langle /req \rangle$), decoder 2 for complete user utterances and decoder 3 for machine responses. The embedding layer and encoders are the same as described in section 3.

BSpan Decoder Unlike Lei et al. (2018), we do not concatenate current user utterance with previously generated machine response. At each turn of dialogue, the user utterance and the previous BSpan (the dialogue states updated to the previous turn) are used as the inputs to the user utterance encoder. The outputs of this encoder are then fed into the BSpan decoder for predicting the new BSpan for the current turn and a cross-entropy loss L_1 is calculated. The user utterance encoder hidden states, the last hidden state and the output of the BSpan decoder are input into the other two decoders.

Complete User Utterance Decoder The basic structure of this decoder is the same as the decoder described in the last section. We pass the last hidden state of the BSpan decoder to the initial state of this decoder. In addition to the inputs from the user utterance encoder and the dialogue context

encoder, we also input the output of the BSpan decoder into this decoder. The generation probability P_t^g , copy probability P_t^{c1} for copying tokens in BSpan, and copy probability P_t^{c2} for copying words in the dialogue context are calculated with a shared normalization term and combined for the final probability computation:

$$P_t = P_t^g + P_t^{c1} + P_t^{c2} \quad (12)$$

P_t is then used to decode the words in the complete user utterance. For this decoder, the second cross-entropy loss L_2 is calculated.

Machine Response Decoder Similar to the previous two decoders, the machine response decoder is also a single-layer unidirectional GRU, the initial state of which is set to the last hidden state of the complete user utterance decoder. In this decoder, we compute three context vectors for each decoder state s_t . The first context vector h_{t1}^* is calculated over the user utterance encoder hidden states while the other two context vectors h_{t2}^* , h_{t3}^* are calculated over the BSpan decoder hidden states and the complete user utterance decoder hidden states, respectively. The concatenation of s_t , h_{t1}^* , h_{t2}^* , h_{t3}^* and the Knowledge Base (KB) matching vector K_t (a one-hot representation of the retrieval results in KB according to the constraints in the corresponding BSpan) is used to generate the output and update the decoder state. The generated output is then concatenated with the three context vectors to feed into a layer to produce the gener-

Turn	Dialogue
Q ₁	I would like a traditional food restaurant.
A ₁	What price range do you have in mind?
Q ₂	I don't care.
Q ₂ (Complete)	I don't care about the price range.
Q ₂ (Ellipsis)	I don't care.
Q ₂ (Co-reference)	I don't care about it.

Table 2: An example of the ellipsis / co-reference annotation

ation probability distribution over the vocabulary. Similar to the complete user utterance decoder, we also use the copy mechanism in the machine response decoder. The third cross-entropy loss L_3 is then calculated.

Training The final loss for the multi-task learning framework is estimated as follows:

$$L = L_1 + L_2 + L_3 \quad (13)$$

We learn parameters to minimize the final loss.

5 Data Annotation for Ellipsis and Co-Reference Resolution in Dialogue

Since there are no publicly available labeled data for the resolution of ellipsis and co-reference in dialogue, we manually annotate such a new dataset based on the public dataset CamRest676 (Wen et al., 2016a,b) from the restaurant domain.

Annotation Specification Annotation cases for user utterances can be summarized into the following three conventions:

- As shown in Table 2, if a user utterance contains an ellipsis or anaphor, we manually resolve the ambiguity of ellipsis or anaphor and supplement the user utterance with a correct expression by checking the dialogue context. In doing so, we create a pragmatically *complete* version for the utterance. If the utterance only contains an ellipsis and the ellipsis can be replaced with an anaphor, we create a *co-reference* version for it. Similarly, if the utterance only contains an anaphor and the anaphor can be omitted, we create an *ellipsis* version for the utterance.
- If the user utterance itself is pragmatically complete, without any ellipsis or anaphora, we create an anaphor and ellipsis version for it if such a creation is appropriate.
- If the utterance itself is complete and it is not suitable to create an ellipsis or anaphor version, we just do nothing.

With the annotation convention described above, for each user utterance in the dataset, we can label it as $l \in \{\textit{ellipsis}, \textit{co-reference}, \textit{complete}\}$ or create two other versions for it if appropriate. Please notice that these labels are used only for dataset statistics or for designing experiments, not for training our models.

Dataset statistics The CamRest676 dataset contains 676 dialogues, with 2,744 user utterances. After annotation, 1,174 *ellipsis* versions and 1,209 *co-reference* versions are created from the 2,744 user utterances. 1,331 incomplete utterances are created that they are either *ellipsis* or *co-reference* version. 1,413 of the 2,744 user utterances are complete and not amenable to change. No new versions are created from these 1,413 utterances.

Dataset Split for Experiments We split the new dataset into a training set (accounting for 80%) and validation set (accounting for 20%) which can be used for the stand-alone ellipsis/co-reference resolution task and the multi-task learning of both the ellipsis/co-reference resolution and end-to-end task-oriented dialogue.

6 Experiments

In this section we conducted experiments on the new dataset to examine the generative ellipsis/co-reference resolution model and its integration into the end-to-end task-oriented dialogue.

6.1 Evaluation Metrics

As far as we know, there is no end-to-end generative ellipsis and co-reference resolution model applied to multi-turn dialogues. Therefore there are no off-the-shelf metrics tailored to this evaluation. Since we deal with two tasks: the task of ellipsis/co-reference resolution (resolution task for short) and the task-oriented dialogue with integrated ellipsis/co-reference resolution (hereafter dialogue task), we use two sets of evaluation metrics. For the resolution task, we use the exact match rate (EM) that measures whether the generated utterances exactly match the gold utterances.

BLEU (Papineni et al., 2002) and F_1 score (a balance between word-level precision and recall) are also used for the resolution task to evaluate the quality of generated utterances at the n-gram and word level. We use the success F_1 which is defined as the F_1 score of requested slots correctly answered in dialogues to evaluate task comple-

Data	Model	Resolution Task							
		EM(%)	EM 1(%)	EM 2(%)	BLEU(%)	F ₁ (%)	Prec.(%)	Rec.(%)	Reso.F ₁ (%)
Ellipsis	Baseline	49.99	68.88	27.31	73.26	90.89	92.14	89.67	44.47
	GECOR 1	67.56	92.07	37.18	83.69	96.25	98.28	94.30	70.48
	GECOR 2	67.75	91.38	38.46	82.94	96.58	98.48	94.76	70.85
Co-reference	Baseline	55.64	76.03	33.60	78.12	92.58	93.28	91.89	44.24
	GECOR 1	71.35	91.67	47.68	85.89	96.49	98.19	94.86	64.93
	GECOR 2	71.18	93.80	44.92	85.93	97.09	98.46	95.76	71.26
Mixed	Baseline	50.38	70.89	28.57	74.11	90.93	91.72	90.15	44.10
	GECOR 1	68.52	92.03	42.04	83.91	95.88	98.12	93.74	66.06
	GECOR 2	66.22	91.64	37.45	82.98	96.47	98.41	94.60	66.16

Table 3: Results of the resolution task on the dataset. GECOR 1/2: the GECOR model with the copy/gated copy mechanism. EM 1 and EM 2 respectively indicate the situation that the input utterance is complete or incomplete while EM is the comprehensive evaluation of the two situations. Reso.F₁: Resolution.F₁

tion rate for the dialogue task, similar to Lei et al. (2018).

6.2 Parameter Settings

For all our models, both the size of hidden states and word embeddings were set to 50. The vocabulary size $|\mathbf{V}|$ was set to 800 and the batch size was set to 32. We trained our models via the Adam optimizer (Kingma and Ba, 2015), with a learning rate of 0.003 and a decay parameter of 0.5. Early stopping and dropout were used to prevent overfitting, and the dropout rate was set to 0.5.

6.3 Baselines and Comparisons

For the resolution task, we compared our GECOR model with the baseline model proposed by Zheng et al. (2018) which is a seq2seq neural network model that identifies and recovers ellipsis for short texts.

For the dialogue task, we compared our multi-task learning framework with the baseline model TSCP proposed by Lei et al. (2018) which is a seq2seq model enhanced with reinforcement learning. We ran the source code² on our dataset to get the baseline results for comparison.

For the resolution task, we also performed a comparison study to examine the impacts of the gate mechanism incorporated into the copy network on the GECOR model and on the multi-task learning dialogue model.

6.4 The GECOR Model

Our generative resolution model was trained on three types of data: the ellipsis data where only ellipsis version utterances from the annotated dataset were used, the co-reference data where

only co-reference version utterances from the annotated dataset were used, and the mixed data where we randomly selected a version for each user utterance from $\{\textit{ellipsis}, \textit{co-reference}, \textit{complete}\}$. In the mixed data, 633 turns are with ellipsis user utterances, 698 turns are with co-reference user utterances, and the rest are with complete user utterances. The experimental results of the GECOR and baseline model (Zheng et al., 2018) on the three different datasets are shown in Table 3.

Overall results From the third column of the table, we find that the GECOR model with the copy mechanism (GECOR 1) improves the exact match rate (EM) by more than 17 points on the ellipsis version data, more than 15 points on the co-reference data, and more than 18 points on the mixed data. We further define a metric we term as **Resolution.F₁** that is an F₁ score calculated by comparing machine-generated words with ground truth words for only the ellipsis / co-reference part of user utterances. The GECOR model achieves consistent and significant improvements over the baseline in terms of BLEU, F₁ and Resolution.F₁ in addition to the EM metric. The major difference between the GECOR and the baseline is that the former tries to copy words from the dialogue context. The improvements, especially the improvements on the ellipsis resolution (higher than those on the co-reference resolution) indicate that the copy mechanism is crucial for the recovery of ellipsis and co-reference.

Effect of the two copy mechanisms Comparing the GECOR 1 against the GECOR 2 (with the gated copy mechanism), we can find that the gating between copy and generation is helpful in terms of the word-level quality (F₁ and Resolution.F₁ score) but not in terms of the fragment

²<https://github.com/WING-NUS/seqquicity>

Data	Model	Resolution Task					Dialogue Task		
		EM(%)	BLEU(%)	F1(%)	Prec.(%)	Rec.(%)	Succ.F1(%)	Prec.(%)	Rec.(%)
Complete	TSCP	-	-	-	-	-	86.30	89.60	83.23
Ellipsis	TSCP	-	-	-	-	-	84.56	87.25	82.02
	Our Model	60.83	78.89	95.64	97.79	93.58	85.33	88.69	82.21
Co-reference	TSCP	-	-	-	-	-	82.17	88.91	76.38
	Our Model	68.56	83.98	96.61	98.09	95.18	86.00	90.46	81.95
Mixed	TSCP	-	-	-	-	-	83.25	86.91	79.89
	Our Model	66.47	83.63	96.26	98.16	94.44	85.97	87.98	84.05

Table 4: Results of the multi-task learning model. This table is split into two parts: performance of resolution for the integrated GECOR on the left side and performance of dialogue task on the right side.

or sequence-based metrics (i.e., BLEU and EM). Therefore, we only integrate the GECOR model with the copy mechanism into the dialogue system.

Incomplete vs. complete utterances In multi-turn dialogues, user utterances may be incomplete or complete. A robust resolution model needs to be able to accurately identify whether the input utterance is complete or not. The model needs to keep it unchanged when it is complete and to predict the corresponding complete version when it is incomplete. For these cases, we tested our models and made statistical analysis on the three versions of data as shown in column 3, 4 and 5 of Table 3 (EM, EM 1, EM 2). We can find that the GECOR model beats the baseline model in all respects. However, the GECOR model needs further improvement when the input utterances are incomplete, compared with its good performance on complete utterances.

Analysis on GECOR results for complete utterances We then analyzed the experimental results of the GECOR 1 on the mixed data in detail. When the input user utterances are complete, the GECOR model can amazingly generate 92.03% utterances that exactly match the input utterances. Only 7.97% do not match perfectly. Most unmatched cases, as we found, are with: **(1) missed words** (e.g., User: *Can I get a Korean restaurant in the town centre?* GECOR: *Can I get a Korean restaurant in the town?*) **(2) Repetition** (e.g., User: *OK, thank you. That is all for today then.* GECOR: *OK, thank you. That is all for today for today then.*)

Analysis on GECOR results for incomplete utterances For incomplete input user utterances, GECOR can generate 42.04% exactly matched cases. Among the 57.96% cases that do not exactly match ground truth utterances, only 6.3% are not complete, which still contains unresolved el-

lipsis or co-reference, while 93.7% of these cases are complete with GECOR-generated words that do not match ground truth words. An in-depth analysis on these show that they can be clustered into 4 classes. **(1) Paraphrases.** We found that the majority of the unmatched complete utterances generated by GECOR are actually paraphrases to the ground truth complete utterances (e.g., User: *Any will be fine.* GECOR: *Any food type will be fine.* Reference: *Any type of restaurant will be fine.*) This is also confirmed by the high scores of the word-level evaluation metrics in Table 3. **(2) Partial resolution.** When a pronoun refers to more than one items, GECOR sometimes generate a partial resolution for the pronoun (e.g., User: *I do not care about them.* GECOR: *I do not care about the price range.* Reference: *I do not care about the price range or location.*) **(3) Minor errors.** In a few cases, the resolution part is correct while there are some errors elsewhere. (e.g., User: *How about Chinese food?* Prediction: *How about international food on the south side of town?* Reference: *How about Chinese food on the south side of town?.*) **(4) Repetition.** Some cases contain repeatedly generated words.

We think that although not exactly matched, paraphrased complete utterances generated by GECOR are acceptable. These utterances are helpful for the downstream dialogue task. For other errors, such as partial resolution or repetition, it may be necessary to enhance the attention or copy mechanism further in GECOR.

6.5 The Multi-Task Learning Model

We further conducted experiments to extrinsically evaluate the GECOR model in task-oriented dialogue with the success F_1 metric. This is also to evaluate our multi-task learning framework in integrating the GECOR model into the end-to-end dialogue model. In addition to training the base-

line TSCP model on the ellipsis, co-reference and mixed dataset, we also trained it on the dataset with only complete user utterances. This is to examine the ability of the baseline model in using correct contextual information presented in user utterances. The experimental results are shown in Table 4.

Overall results In comparison to the baseline, we can see that our model improves the success F_1 score by nearly 4 points on the co-reference dataset, which is close to the score obtained by the baseline trained with the complete user utterances. On the mixed and ellipsis dataset, our model also achieves 2.7 points and 0.8 points of success F_1 score improvements, respectively.

Resolution performance of the integrated GECOR We also provide the performance of the integrated GECOR on the resolution task in Table 4. The performance is slightly lower than when the GECOR is trained independently as a stand-alone system. This suggests that the GECOR is able to perform well when integrated into a dialogue system. The overall results demonstrate that the proposed multi-task learning framework for the end-to-end dialogue is able to improve the task completion rate by incorporating an auxiliary ellipsis/co-reference resolution task.

Since the BSpan decoder is also used in the baseline system to capture contextual information and track dialogue states, we believe that our multi-task learning model with the integrated GECOR will play a more important role in end-to-end dialogue models that do not use state tracking modules, e.g., neural open-domain conversation models (Vinyals and Le, 2015; Li et al., 2016).

7 Conclusion and Future Work

In this paper, we have extensively investigated the ellipsis and co-reference resolution in the context of multi-turn task-oriented dialogues. We have presented the GECOR, a unified end-to-end generative model for both ellipsis and co-reference resolution in multi-turn dialogues. A multi-task learning framework is further proposed to integrate the GECOR into the end-to-end task-oriented dialogue. In order to train and test the proposed model and framework, we manually created a new dataset with annotated ellipsis and co-reference information based on the publicly available CamRest676 dataset. Experiments on the resolution task show that the GECOR is able to significantly

improve the performance in terms of the exact match rate, BLEU and word-level F_1 score. Experiments on the dialogue task demonstrate that the task completion rate of the task-oriented dialogue system is significantly improved with the aid of ellipsis and co-reference resolution.

Our work could be extended to end-to-end open-domain multi-turn dialogue. We will further improve our model by incorporating syntactic and location information. We would also like to adapt the proposed methods to document-level neural machine translation in the future.

Acknowledgments

The present research was supported by the National Natural Science Foundation of China (Grant No.61861130364) and the Royal Society (London) (NAF\R1\180122). We would like to thank the anonymous reviewers for their insightful comments.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015)*.
- Kevin Clark and Christopher D Manning. 2016. Deep reinforcement learning for mention-ranking coreference models. *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2256–2262.
- Mary Dalrymple, Stuart M Shieber, and Fernando CN Pereira. 1991. Ellipsis and higher-order unification. *Linguistics and philosophy*, 14(4):399–452.
- Mihail Eric and Christopher D Manning. 2017a. A copy-augmented sequence-to-sequence architecture gives good performance on task-oriented dialogue. *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 468–473.
- Mihail Eric and Christopher D Manning. 2017b. Key-value retrieval networks for task-oriented dialogue. *Proceedings of the SIGDIAL 2017 Conference*, pages 37–49.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor OK Li. 2016. Incorporating copying mechanism in sequence-to-sequence learning. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1631–1640.
- Xisen Jin, Wenqiang Lei, Zhaochun Ren, Hongshen Chen, Shangsong Liang, Yihong Zhao, and Dawei

- Yin. 2018. Explicit state tracking with semi-supervision for neural dialogue generation. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 1403–1412. ACM.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015)*.
- Vineet Kumar and Sachindra Joshi. 2016. Non-sentential question resolution using sequence to sequence learning. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2022–2031.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. End-to-end neural coreference resolution. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 188–197.
- Kenton Lee, Luheng He, and Luke Zettlemoyer. 2018. Higher-order coreference resolution with coarse-to-fine inference. *Proceedings of NAACL-HLT 2018*, pages 687–692.
- Wenqiang Lei, Xisen Jin, Min-Yen Kan, Zhaochun Ren, Xiangnan He, and Dawei Yin. 2018. Sequicity: Simplifying task-oriented dialogue systems with single sequence-to-sequence architectures. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1437–1447.
- Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016. Deep reinforcement learning for dialogue generation. *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202.
- Zhengzhong Liu, Edgar González Pellicer, and Daniel Gillick. 2016. Exploring the steps of verb phrase ellipsis. In *Proceedings of the Workshop on Coreference Resolution Beyond OntoNotes (CORBON 2016)*, pages 32–40.
- Leif Arda Nielsen. 2003. A corpus-based study of verb phrase ellipsis. In *Proceedings of the 6th Annual cluk Research Colloquium*, pages 109–115.
- Martha S. Palmer, Deborah A. Dahl, Rebecca J. Schiffman, Lynette Hirschman, Marcia Linebarger, and John Dowding. 1986. Recovering implicit information. In *Proceedings of the Workshop on Strategic Computing Natural Language, HLT '86*, pages 96–113, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Haoruo Peng, Kai-Wei Chang, and Dan Roth. 2015. A joint framework for coreference resolution and mention head detection. In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*, pages 12–21.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Alexander I Rudnicky, Eric Thayer, Paul Constantinides, Chris Tchou, R Shern, Kevin Lenzo, Wei Xu, and Alice Oh. 1999. Creating natural dialogs in the carnegie mellon communicator system. In *Sixth European Conference on Speech Communication and Technology*.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Stuart M Shieber, Fernando CN Pereira, and Mary Dalrymple. 1996. Interactions of scope and ellipsis. *Linguistics and philosophy*, 19(5):527–552.
- Olga Uryupina and Alessandro Moschitti. 2013. Multilingual mention detection for coreference resolution. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 100–108.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *Proceedings of the International Conference on Machine Learning, Deep Learning Workshop (2015)*.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrkšić, Lina M Rojas Barahona, Pei-Hao Su, Stefan Ultes, David Vandyke, and Steve Young. 2016a. Conditional generation and snapshot learning in neural dialogue systems. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2153–2162.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2016b. A network-based end-to-end trainable task-oriented dialogue system. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 438–449.
- Jie Zheng, Fang Kong, and Guodong Zhou. 2018. Sequence to sequence model to ellipsis recovery for chinese short text. *Journal of Chinese Information Processing*, 32(12):92–99.
- Victor Zue, Stephanie Seneff, James R Glass, Joseph Polifroni, Christine Pao, Timothy J Hazen, and Lee Hetherington. 2000. Juplter: a telephone-based

conversational interface for weather information. *IEEE Transactions on speech and audio processing*, 8(1):85–96.

Victor W Zue and James R Glass. 2000. Conversational interfaces: Advances and challenges. *Proceedings of the IEEE*, 88(8):1166–1180.