

# Re-entry Prediction for Online Conversations via Self-Supervised Learning

Lingzhi Wang<sup>1,2</sup>, Xingshan Zeng<sup>3</sup>, Huang Hu<sup>4</sup>, Kam-Fai Wong<sup>1,2</sup>, Daxin Jiang<sup>4</sup>

<sup>1</sup>The Chinese University of Hong Kong, Hong Kong, China

<sup>2</sup>MoE Key Laboratory of High Confidence Software Technologies, China

<sup>3</sup>Huawei Noah's Ark Lab, Hong Kong, China

<sup>4</sup>Microsoft Corporation, Beijing, China

<sup>1,2</sup>{lzwang, kfwong}@se.cuhk.edu.hk <sup>3</sup>zeng.xingshan@huawei.com

<sup>4</sup>{huahu, djiang}@microsoft.com

## Abstract

In recent years, world business in online discussions and opinion sharing on social media is booming. Re-entry prediction task is thus proposed to help people keep track of the discussions which they wish to continue. Nevertheless, existing works only focus on exploiting chatting history and context information, and ignore the potential useful learning signals underlying conversation data, such as conversation thread patterns and repeated engagement of target users, which help better understand the behavior of target users in conversations. In this paper, we propose three interesting and well-founded auxiliary tasks, namely, Spread Pattern, Repeated Target user, and Turn Authorship, as the self-supervised signals for re-entry prediction. These auxiliary tasks are trained together with the main task in a multi-task manner. Experimental results on two datasets newly collected from Twitter and Reddit show that our method outperforms the previous state-of-the-arts with fewer parameters and faster convergence. Extensive experiments and analysis show the effectiveness of our proposed models and also point out some key ideas in designing self-supervised tasks.<sup>1</sup>

## 1 Introduction

Online social media platforms are popular for individuals to discuss topics they are interested in and exchange viewpoints. However, a large number of online conversations are posted every day that hinder people from tracking the information they are interested in. As a result, there is a pressing demand for developing an automatic conversation management tool to keep track of the discussions one would like to keep engaging in.

Re-entry prediction (Zeng et al., 2019; Backstrom et al., 2013) is proposed to meet such demand. It aims to foresee whether a user (henceforth

**target user**) will come back to a conversation they once participated in. Nevertheless, the state-of-the-art work (Zeng et al., 2019) mostly focuses on rich information in users' previous chatting history and ignores the thread pattern information (Backstrom et al., 2013; Tan et al., 2019). To this end, we study in re-entry prediction by exploiting the conversation thread pattern to signal whether a user would come back since the degree of repeated engagement of users can indicate their temporary interests in the ongoing conversation.

Self-supervised learning aims to train a model on labels that are automatically derived from the data itself. Compared to previous generic self-supervised methods (e.g., Switch, Replace, and Mask), task-specific methods can achieve better performance (Jing and Tian, 2020), especially on medium-sized datasets, since task-oriented designs can better capture domain-specific features and thus achieve better performance for the target task. Therefore, we propose a prediction model (*main model*) for re-entry prediction (*main task*) with three auxiliary self-supervised tasks (*Spread Pattern Prediction*, *Repeated Target User Prediction* and *Turn Authorship Prediction*) to assist learning of main model for re-entry prediction.

Spread Pattern Prediction is inspired by *expansive* and *focused* thread in Backstrom et al. (2013), where thread pattern reflects the development of a conversation. We implement this task in a simplified but reasonable way to discriminate thread patterns based on the number of participated users. On the other hand, Zeng et al. (2019) shows that target users who contribute two or more posts in a conversation have a higher probability of coming back. Hence, we introduce a Repeated Target User Prediction task to facilitate the learning of the main model by capturing the target user's behavior, i.e., whether the target user has posted more than one message in a given context. Finally, we introduce the Turn Authorship Prediction task, in

<sup>1</sup>The code is available at <https://github.com/Lingzhi-WANG/ReEntryPrediction>

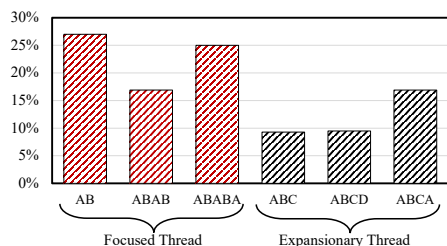


Figure 1: X-axis: thread pattern, e.g., “AB” represents thread where user A posts then B posts. Y-axis: re-entry rate, e.g., the re-entry rate for “AB” is 27%, means that 27% of the target users (user “B”) in this kind of conversations will come back.

which we step further from the Repeated Target User Prediction task to predict if each turn’s authorship is the target user. Thus, the model can track the participation of the target user and also know the thread pattern reflecting by the position of the target user who acts as a probe.

To better illustrate our motivation, Figure 1 shows the re-entry rate of six representative thread patterns on Reddit dataset. As we can see, the left three threads with user number “ $\leq 2$ ” (focused) show a higher re-entry rate than the right three threads with user number “ $> 2$ ” (expansionary). We can also see that although “ABCA” is expansionary, it has repeated target user “A”, which leads to a higher re-entry rate than the other two expansionary threads. Therefore, we can conclude that both Spread Pattern and Repeated Target User signals help predict re-entry behavior. Furthermore, since more challenging tasks get better performance (Mao, 2020), we propose Turn Authorship Prediction, where we predict whether each turn’s author is a target user or not.

Before the introduction of pretraining technique (Peters et al., 2018; Devlin et al., 2018; Radford et al., 2019), researchers focused on developing complex models (Lu and Ng, 2020), such as key phrase generation with neural topic model (Wang et al., 2019b) and structured models for coreference resolution (Martschat and Strube, 2015; Björkelund and Kuhn, 2014). Thus models are time-consuming in training and testing. For this reason, we propose our compact main model, which consists of three parts, turn encoder, conversation encoder and prediction layer. In addition, the chatting history information of the target user is also applied to our model by initializing the beginning hidden state of the target turn. The main model is jointly trained with the three self-supervised tasks in the manner

of multi-task training and outperforms the BERT-based (Devlin et al., 2018) model which consists of large number of parameters and is time-consuming.

In summary, our contributions are three-fold:

- Three self-supervised tasks are proposed to facilitate learning of the main model by capturing the thread pattern and participation trajectory of the target user.
- Experimental results on two newly constructed datasets, Twitter and Reddit, show that our methods outperform the previous state-of-the-arts with fewer parameters and faster convergence.
- Extensive experiments and analyses provide more insights on how our models work and how to design effective self-supervised tasks for conversational prediction task.

The remainder of this paper is organized as follows. The related work is surveyed in Section 2. Section 3 and 4 present the proposed approach, including model architecture and designed self-supervised tasks. Section 5 and 6 then present the experimental setup and results respectively. Finally, conclusions are drawn in Section 7.

## 2 Related Work

**Re-entry Prediction.** Re-entry prediction (Zeng et al., 2019; Backstrom et al., 2013; Budak and Agrawal, 2013) aims to forecast whether the users will return to a discussion they once entered and Zeng et al. (2019) achieves state-of-the-art performance by exploiting user’s history context (Flek, 2020). Re-entry prediction focuses on conversation-level response prediction (Zeng et al., 2018; Chen et al., 2011). Most of them adopt a complex framework (Zeng et al., 2018) and massive parameters (see Figure 4(b)) while our model is simple and effectively combines the current conversation and chatting history.

**Self-supervised Learning.** Self-supervised learning aims to train a network on an auxiliary task where the ground-truth label is automatically derived from the data itself (Wu et al., 2019; Lan et al., 2019; Erhan et al., 2010; Hinton et al., 2006). It has been applied to many tasks, such as text classification (Yu and Jiang, 2016), neural machine translation (Ruiter et al., 2019), multi-turn response selection (Xu et al., 2020), summarization (Chen and Wang, 2019) and dialogue learning (Wu et al., 2019). These auxiliary tasks can be categorized into word-level tasks and sentence-level tasks.

In word-level tasks, nearby word prediction (Mikolov et al., 2013) and next word prediction (Bengio et al., 2003; Wang and Gupta, 2015) are widely explored in language modeling. Masked language model (Devlin et al., 2018) is also in the line of word-level tasks.

In sentence-level tasks, Wang et al. (2019a) exploits Mask, Replace and Switch for extractive summarization. Wu et al. (2019) propose Inconsistent Order Detection for dialogue learning. Xie et al. (2020) exploit Drop, Replace, and TOV (Temporal Order Verification) for story cloze test. Xu et al. (2020) also design several self-supervised tasks to improve the performance of response selection.

Most of the previous self-supervised tasks (both in word-level and sentence-level) focus on the general domain while our work is based on task-orientated supervised methods and achieves better performance.

### 3 Re-entry Prediction Framework

This section describes our re-entry prediction framework. The left part of Figure 2 shows our overall structure. In the following, we first introduce the input and output in Section 3.1. Then in Section 3.2, we describe our prediction framework. Finally, the learning objective of the entire model will be given in Section 3.3.

#### 3.1 Input and Output

The input of our model contains two parts: the observed conversation  $c$  and chatting history  $c^h$  of target user  $u$ . The conversation  $c$  is formalized as a sequence of turns (e.g., posts or tweets)  $\langle t_1, t_2, \dots, t_m \rangle$  where  $m$  represents the length of conversation (number of turns) and  $t_m$  is posted by user  $u$ .  $t_i$  is the  $i$ -th turn of the conversation and contains words  $\langle w_{i,1}, w_{i,2}, \dots, w_{i,k_i} \rangle$ , where  $w_{ij}$  is the  $j$ -th word in  $i$ -th turn and  $k_i$  is the word length of  $i$ -th turn. The chatting history  $c^h$  is constructed by concatenating the turns (in training corpus) that are authored by the user  $u$  into a sequence following their posting time.

For output, we yield a Bernoulli distribution  $p(c, c^h)$  to indicate the estimated likelihood of whether  $u$  will re-engage in the conversation  $c$ , giving the chatting history  $c^h$  of  $u$ .

#### 3.2 Re-entry Prediction Model

Our model consists of three modules: turn encoder, conversation encoder, and prediction layer.

**Turn Encoder.** We first feed each word  $w_{i,j}$  in turn  $t_i$  into an embedding layer and get the word representation  $e_{i,j}$ . The turn representation is then modeled via a turn encoder, where a word-level bidirectional gated recurrent unit (Bi-GRU) (Cho et al., 2014) is adopted. The hidden states of Bi-GRU are defined as:

$$\vec{h}_j = f_{GRU}(e_j, \mathbf{h}_{j-1}), \overleftarrow{h}_j = f_{GRU}(e_j, \mathbf{h}_{j+1}) \quad (1)$$

The output of our turn encoder is the concatenation of the last hidden states of both directions of Bi-GRU:  $\mathbf{h} = [\vec{h}_j; \overleftarrow{h}_j]$ .

To incorporate the information of the user’s chatting history, we also use the same procedure described above to encode each history turn  $t_i^h$  in  $c^h$ . We then apply another Bi-GRU layer to capture the temporal features among these history turns and derive the final representation of chatting history  $\mathbf{h}^h$  for target user  $u$ . Finally, we use  $\mathbf{h}^h$  to initialize the hidden states of the last conversation turn (posted by target user  $u$ )  $t_m$ ’s turn encoder, and the initialization mechanism is proven to be helpful in Wang et al. (2020). The following equation describes the initialization:

$$\vec{h}_{m,0} = \overleftarrow{h}_{m,k_m} = \eta(\mathbf{W}_0 \mathbf{h}^h + b_0) \quad (2)$$

where  $\vec{h}_{m,0}$  and  $\overleftarrow{h}_{m,k_m}$  are the initial states of both directions, and  $\eta$  is a Tanh activated function.  $\mathbf{W}_0$  and  $b_0$  are learnable parameters.

With the initialization process, we produce more informative representation of the final turn  $\mathbf{h}_m$ , since it can encode information from both current conversation  $c$  and target user  $u$ ’s chatting history.

**Conversation Encoder.** To learn the conversational structure representations for  $c$ , we apply a third Bi-GRU, to capture the interactions between adjacent context turns:

$$\vec{r}_j = f_{GRU}(\mathbf{h}_j, \mathbf{r}_{j-1}), \overleftarrow{r}_j = f_{GRU}(\mathbf{h}_j, \mathbf{r}_{j+1}) \quad (3)$$

We concatenate the outputs of both directions and get the turn representations of  $c$ :  $\mathbf{r}_j = [\vec{r}_j; \overleftarrow{r}_j]$ .

Since different turns might play different roles in predicting target user  $u$ ’s re-entry behavior (e.g. the turns that  $u$  directly replied before should be more important than other turns), we apply an attention mechanism here to force our model to pay more attention to important turns. Concretely, the final representation of conversation  $c$  is defined as:

$$\mathbf{r} = \sum_{j=0}^m a_j \mathbf{h}_j, a_j = \text{softmax}(\mathbf{W}_{att} \mathbf{h}_j + b_{att}) \quad (4)$$

where  $\mathbf{W}_{att}$  and  $b_{att}$  are learnable parameters.

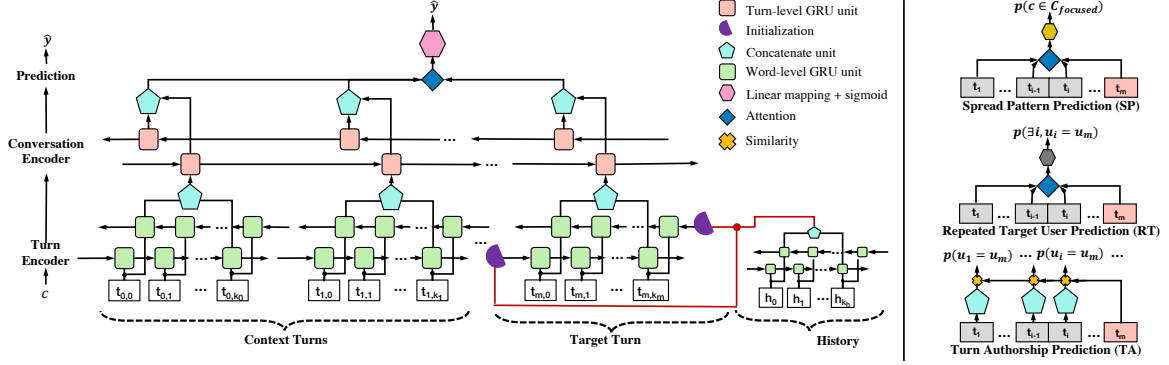


Figure 2: Our main model (left part) and three self-supervised tasks (right part) for re-entry prediction.

**Prediction Layer.** We predict the final output  $\hat{y} \in [0, 1]$ , which signals how likely  $u$  will re-engage in  $c$ , with the following prediction layer:

$$\hat{y} = \sigma(\mathbf{v}^T[\mathbf{r}, \mathbf{h}_m] + b) \quad (5)$$

where  $\mathbf{v}$  and  $b$  are learnable parameters and  $\sigma(\cdot)$  is the sigmoid function. Here we concatenate the conversation representation  $\mathbf{r}$  and hidden state of final turn  $\mathbf{h}_m$  as input, to emphasize the role of final turn posted by target user  $u$ .

### 3.3 Learning Objective

Following Zeng et al. (2019), we use *binary cross-entropy loss* as our learning objective. Also, to deal with the imbalance of positive and negative instances in training corpus, we weigh differently for their losses. The equation for our main task is defined as follows:

$$\mathcal{L}_{main} = - \sum_{i \in \mathcal{T}} [\lambda \cdot y_i \log(\hat{y}_i) + \mu(1 - y_i) \log(1 - \hat{y}_i)] \quad (6)$$

where  $\mathcal{T}$  is the training corpus,  $y_i$  denotes the ground-truth label for  $i$ -th instance in the training corpus (label is 1 if target user re-engages later, otherwise 0), and  $\lambda$  and  $\mu$  are hyper-parameters to trade off the weights between positive and negative instances. Generally, the values for  $\lambda$  and  $\mu$  can be tuned based on the ratios of positive and negative examples in the training corpus.

## 4 Self-Supervised Tasks

This section describes the proposed self-supervised tasks that guide the re-entry prediction model to better capture user behaviors in online conversations. The right part of Figure 2 illustrates our three self-supervised tasks.

### 4.1 Spread Pattern Prediction

Backstrom et al. (2013) shows that *expansive* (engagement among many users) and *focused* (repeated engagement among few users) are two kinds of spread patterns in online multi-party conversations. Distinguishing spread patterns of conversations is helpful in predicting the future trajectory of the conversation. Therefore, we propose the *Spread Pattern Prediction* task (SP task in Figure 2) which is a simplified form of the work of Backstrom et al. (2013).

We divide conversations into two types – *focused conversation* ( $C_{focused}$ ) and *expansive conversation* ( $C_{exp}$ ). Focused conversations are composed of repeated discussions between only two active users while expansive conversations contain more than two users. We then make binary prediction between focused (label  $y^{sp} = 1$ ) and expansive ( $y^{sp} = 0$ ) conversation with another prediction layer (the reason for assigning label 1 to focused conversation can be found in Section 6.3):

$$\hat{y}^{sp} = p(c \in C_{focused}) = \sigma(\mathbf{v}_{sp}^T[\mathbf{r}, \mathbf{h}_m] + b_{sp}) \quad (7)$$

where  $\mathbf{r}$  and  $\mathbf{h}_m$  are the same as Eq. 5,  $\mathbf{v}_{sp}$  and  $b_{sp}$  are learnable parameters.

We still apply weighted binary cross entropy introduced in Eq. 6 as our learning objective. To simplify the hyper parameter tuning, we force the trade off weight to be the ratio between positive and negative instances. Below describes the equation:

$$\mathcal{L}_{SP} = - \sum_{i \in \mathcal{T}} [\lambda_{sp} \cdot y_i^{sp} \log(\hat{y}_i^{sp}) + (1 - y_i^{sp}) \log(1 - \hat{y}_i^{sp})] \quad (8)$$

where  $\lambda_{sp}$  equals to the number of positive instances divided by that of negative ones in training corpus.  $\hat{y}_i^{sp}$  is the output of the  $i$ -th instance.

## 4.2 Repeated Target Prediction

Zeng et al. (2019) shows that their model achieves better performance in second or third re-entry prediction (i.e. the target user has already contributed two or three turns) than first re-entry prediction. It might be attributed to the fact that users who participated in the conversation twice or more are more likely to return to this conversation (see statistic in Table 1). Therefore, we design *Repeated Target Prediction* task (refer to **RT** task in Figure 2). We label those conversations containing repeated target users with 1 ( $y^{rt} = 1$ ) and other conversations with 0 ( $y^{rt} = 0$ ) and carry out binary prediction:

$$\hat{y}^{rt} = p(\exists i, u_i = u_m) = \sigma(\mathbf{v}_{rt}^T[\mathbf{r}, \mathbf{h}_m] + b_{rt}) \quad (9)$$

where  $\mathbf{r}$  and  $\mathbf{h}_m$  are the same as Eq. 5,  $\mathbf{v}_{rt}$  and  $b_{rt}$  are learnable parameters.

The learning objective for this task can be summarized as below, following similar idea in SP:

$$\mathcal{L}_{RT} = - \sum_{i \in \mathcal{T}} [\lambda_{rt} \cdot y_i^{rt} \log(\hat{y}_i^{rt}) + (1 - y_i^{rt}) \log(1 - \hat{y}_i^{rt})] \quad (10)$$

## 4.3 Turn Authorship Prediction

By combining the intention of the SP and RT tasks, we further design *Turn Authorship Prediction* (henceforth **TA**) task. The TA task aims to predict whether the turn’s author is the target user and we label "yes" with 1 and "no" with 0. This task benefits the main task by signaling both the conversation spread pattern and repeated user pattern. Specifically, this is a turn-level authorship prediction and can help learn meaningful turn representations, which are essential for conversation modeling.

Formally, we output each turn’s score as the similarity between hidden states of current turn ( $\mathbf{h}_i$ ,  $i = 1, 2, \dots, m - 1$ ) and target turn ( $\mathbf{h}_m$ ), followed by a sigmoid activated function:

$$\hat{y}_j^{ta} = p(u_j = u_m) = \sigma(\mathbf{h}_j \cdot \mathbf{h}_m) \quad (11)$$

which reflects the probability of turn  $j$  being authored by the target user. Mean Square Error (MSE) loss is applied for TA task:

$$\mathcal{L}_{TA} = \sum_{i \in \mathcal{T}} \sum_{j=1}^{m_i-1} (y_{ij}^{ta} - \hat{y}_{ij}^{ta})^2 \quad (12)$$

where  $m_i$  is the turn number of conversation  $i$ .

	Twitter	Reddit
# of convs	45,111	16,340
# of turns	229,435	58,189
Avg # of turns per conv	5.09	3.56
Avg len of turn per conv	20.3	42.9
% with repeated target	63.2	39.7
% of positive instances	48.9	21.3

Table 1: Statistics of Twitter and Reddit datasets. ‘‘Avg # of turns’’ means the average turn number. ‘‘len’’ refers to the number of tokens. ‘‘% with repeated target’’ represents the ratio of conversations that target users have appeared at least twice in context. ‘‘Positive instances’’ are the conversations which target users re-engage later.

## 4.4 Training Procedure

All three auxiliary tasks are trained on parameters shared with the main task except for the final prediction layer (Section 3.2) in a multi-task learning manner. The final total loss is:

$$\mathcal{L}_{final} = \mathcal{L}_{main} + \alpha_{sp} \mathcal{L}_{SP} + \alpha_{rt} \mathcal{L}_{RT} + \alpha_{ta} \mathcal{L}_{TA} \quad (13)$$

where  $\alpha_{sp}$ ,  $\alpha_{rt}$ ,  $\alpha_{ta}$  are hyper-parameters.

## 5 Experimental Setup

**Datasets.** For experiments, we construct two new datasets from **Twitter** and **Reddit**. The raw Twitter and Reddit data is released by Zeng et al. (2018, 2019) and both in English. For both Twitter and Reddit, we form the conversations with postings and replies (all the comments and replies also viewed as a single turn) following the practice in Li et al. (2015) and Zeng et al. (2018).

In our main experiment, different from Zeng et al. (2019), we do not focus on predicting first re-entries (i.e. only giving the context until the target user’s first participation), we generalize the setting into re-entry prediction regardless of the number of user’s past participation. In this way, our model can learn more general and applicable features for re-entry prediction in diverse scenarios.

The statistics of the two datasets are shown in Table 1. As can be seen, Twitter dataset is much larger than Reddit dataset, with longer conversations (derived from the average number of turns) and shorter turns (observed from the average length of turns). Besides, it contains more conversations with repeated target users, which means that we are more likely to predict the second or third re-entries. At last, Reddit dataset is severely imbalanced in the ratio of positive and negative samples. This indicates that users in Reddit usually do not come back to the conversation they once participated in.

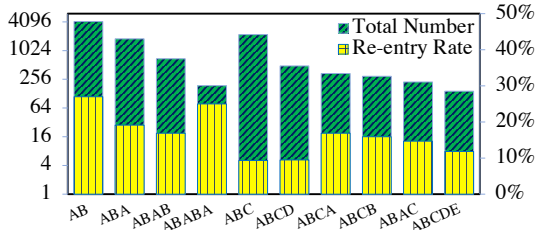


Figure 3: X-axis: thread patterns (the same meaning as in Fig. 1). Left Y-axis: the number of user patterns; Right Y-axis: re-entry rate for each user pattern.

We also present the distribution of thread patterns with their re-entry rate for Reddit in Figure 3. It can be seen that “AB”, “ABA” and “ABC” are the most frequent patterns. And re-entry rate for focused conversations (i.e. only two users participate, such as “AB” and “ABAB”) is generally higher than expansionary conversations, since prior contributions in one conversation may result in continued participation. Such a phenomenon verifies our motivation to design self-supervised tasks.

**Preprocessing.** We applied the Glove tweet preprocessing toolkit (Pennington et al., 2014) to the Twitter dataset. As for the Reddit dataset, we first tokenized the words with the open-source natural language toolkit (NLTK) (Loper and Bird, 2002). We then removed all the non-alphabetic tokens and replaced links with the generic tag “URL”. For both datasets, a vocabulary was maintained with all the remaining tokens, including emoticons and punctuation marks.

**Parameter Setting.** For the parameters in the main model, we first initialize the embedding layer with 200-dimensional Glove embedding (Pennington et al., 2014), whose Twitter version is used for the Twitter dataset and Common Crawl version is applied to Reddit<sup>2</sup>. For our BiGRU layers, we set the size of hidden states for each direction to 200. We employ Adam optimizer (Kingma and Ba, 2015) with initial learning rate  $1e-4$  and early stop adoption (Caruana et al., 2001) in training. The batch size is set to 32. Dropout strategy (Srivastava et al., 2014) and  $L_2$  regularization are used to alleviate overfitting. And the tradeoff parameters  $\alpha_{sp}$ ,  $\alpha_{rt}$ ,  $\alpha_{ta}$  are all set to 0.2. All the hyper-parameters above are tuned on the validation set by grid search.

<sup>2</sup><https://nlp.stanford.edu/projects/glove/>

**Evaluation Metrics.** We use area under the Curve of ROC (AUC), accuracy (ACC), precision (Pre), and F1-scores (F1) to evaluate baselines and our method. Note that, to save spaces, we do not include Recall scores since it can be calculated with Pre and F1.

**Baselines and Comparisons.** We first compare four baselines. The first method is a weak baseline RANDOM that randomly predicts “yes-or-no” labels. The second model, referred to as CCCT, is from an earlier work (Backstrom et al., 2013) that trains a bagged decision tree with manually-crafted features including arrival patterns, time effects, and most related terms, etc. The third compared model, BiLSTM+BIA (Zeng et al., 2019), yields state-of-the-art results with a BiLSTM modeling turn information and a bi-attention mechanism extracting the interaction effects between context and history. We also compare BERT+BiLSTM, where the turn representations are extracted with BERT (Devlin et al., 2018), and a BiLSTM is applied for modeling the conversation structure. For our proposed method, we further compare different self-supervised tasks (SP, RT, TA).

## 6 Experimental Results

In this section, we first introduce the main comparison results in Section 6.1. Then the effects of our methods and how our methods make effects are given in Section 6.2 and Section 6.3 respectively. Finally, Section 6.4 yields further discussion on user history and error analysis.

### 6.1 Main Comparison Results

Table 2 reports the main results on the two datasets. Several interesting observations can be drawn:

- *History and attention mechanism are useful.* Compared to BiGRU, both BiGRU+HISTORY and BiGRU+ATT achieve better performance. The integration of them, i.e., FULL MAIN, brings greater improvement, which means both user’s past behaviors and the key turns of current context are important to signal the user’s re-entry behavior.
- *Self-supervised tasks are all beneficial.* The main model trained with any one of the three self-supervised tasks outperforms the main model itself. Specifically, TA task achieves the best performance on AUC and F1 on both datasets.
- *Self-supervised methods perform better than BERT-based model.* Compared to BERT+BiLSTM, FULL MAIN trained with SP

Models	Twitter				Reddit			
	AUC	Acc	Pre	F1	AUC	Acc	Pre	F1
<b>Baselines</b>								
RANDOM	50.3±0.56	50.1±0.54	51.2±0.55	50.5±0.55	49.3±1.38	49.5±1.44	22.0±1.24	30.6±1.73
CCCT(Backstrom et al., 2013)	62.4	60.3	57.9	64.9	60.1	57.2	29.5	36.1
BiLSTM+BiA(Zeng et al., 2019)	57.3±1.18	51.8±0.45	52.3±0.21	67.9±0.19	60.9±2.75	55.8±3.81	28.1±1.70	38.3±2.08
BERT+BiLSTM(Devlin et al., 2018)	67.8±0.26	60.8±0.33	57.2±0.22	69.7±0.19	62.5±0.11	55.3±2.00	28.2±1.06	39.5±0.33
<b>W/O Self-supervised Task</b>								
BiGRU	65.4±0.69	58.1±1.59	54.9±1.55	67.8±0.35	58.6±3.00	48.1±3.85	26.1±1.87	38.1±1.24
BiGRU+HISTORY	65.1±0.64	58.2±1.21	55.2±1.13	68.6±0.53	61.8±3.15	52.4±2.42	27.4±1.63	39.1±1.46
BiGRU+ATT	66.5±0.79	59.3±1.11	56.2±1.14	68.7±0.40	59.3±3.95	51.7±3.39	27.1±2.50	37.5±1.35
BiGRU+His+ATT(FULL MAIN)	67.3±0.62	59.9±1.39	56.7±1.15	69.4±0.83	61.6±3.93	53.4±4.11	29.1±2.93	39.4±1.51
<b>With Self-supervised Task(s)</b>								
FULL MAIN+SP	67.1±0.47	59.9±1.10	57.1±1.04	69.9±0.30	62.8±0.82	58.1±2.18	29.6±1.42	40.0±0.25
FULL MAIN+RT	67.4±0.41	60.0±1.06	57.1±1.02	69.3±0.16	63.2±1.40	<b>59.6±1.86</b>	<b>30.1±1.14</b>	39.9±0.98
FULL MAIN+TA	<b>68.6±0.86</b>	<b>61.0±0.90</b>	<b>58.4±0.91</b>	<b>70.5±0.19</b>	<b>64.6±0.91</b>	57.7±2.12	29.1±1.81	<b>40.6±0.27</b>

Table 2: Main comparison results displayed with average scores (in %) and their standard deviations over the results with 5 sets of random initialization seeds. The best results in each column are in **bold**. Our model yields better scores than all comparisons for all metrics.

Tasks	Twitter				Reddit			
	AUC	Acc	Pre	F1	AUC	Acc	Pre	F1
REPLACE	65.6	58.2	55.5	69.3	62.3	54.9	28.8	39.2
SWITCH	65.8	58.0	56.8	69.1	61.1	52.8	27.2	38.9
MASK	64.3	57.5	55.3	68.5	60.7	53.0	27.6	38.7
TA (OUR)	<b>68.6</b>	<b>61.0</b>	<b>58.4</b>	<b>70.5</b>	<b>64.6</b>	<b>57.7</b>	<b>29.1</b>	<b>40.6</b>

Table 3: Comparison results (in %) of different self-supervised tasks. TA task yields better performance than Replace, Switch and Mask on all metrics.

or RT task achieves comparable performance on Twitter and better performance on Reddit. Besides, FULL MAIN trained with TA task consistently outperforms BERT+BiLSTM on both datasets. The reason might be that the TA task can better capture the user’s re-entry behaviors and thus leads to better performance of the main model.

- *Self-supervised learning can make the performance more stable.* We can see that all models with auxiliary self-supervised tasks have a smaller standard deviation, which means self-supervised learning can reduce the impact of the model’s parameter initialization and make the performance more stable.

## 6.2 Effects of Our Self-Supervised Tasks

To further validate the effects of our self-supervised tasks, we compare them with three generic self-supervised tasks. Also, we investigate the training efficiency of our main model.

**Compare with other self-supervised tasks.** We compare our best task TA with three popular self-supervised tasks: **Replace**, **Switch** and **Mask**. We follow the turn-level setting in Wang et al. (2019a) and implement them as follows. **Replace**: ran-

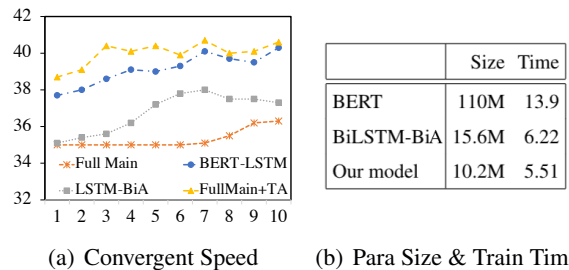


Figure 4: For Fig. 4(a), X-axis: epoch index, Y-axis: F1 score. For Tab. 4(b), “Size” means the parameter size, “Time” refers to the time that one epoch needed.

domly replaces some turns in a conversation with random turns from other conversations, then predict which turns are replaced (each turn has one label, while 1 means replaced, 0 otherwise). **Switch**: randomly switches some turns of the conversation, then predict which turns are not in the original positions (1 means not in the original position, otherwise 0). **Mask**: randomly masks the representations of some turns, then predicts them from a candidate list. Refer to Table 3, our self-supervised task outperforms other generic tasks on the both datasets. This is probably because our tasks can capture more useful information (e.g., thread pattern and user trajectory) which are vital to re-entry prediction.

**Compare with baselines.** As discussed in Section 6.1, models with self-supervised learning show more stable performance. We further explore their differences with respect to convergence speed, parameter size, and training time. We present F1 scores (in validation set) of the first 10 epochs for the four models, BERT+BiLSTM, BiLSTM-BiA (Zeng et al., 2019), our main model

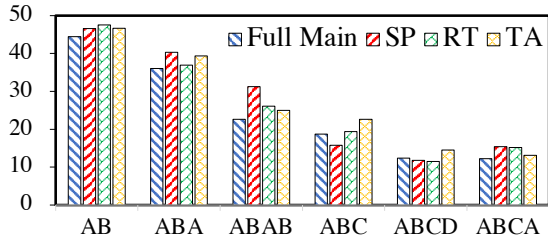


Figure 5: X-axis: thread pattern (the same meaning as in Fig. 3). Y-axis: F1 score (in %).

without self-supervised learning (FULL MAIN) and our main model with self-supervised learning (FULLMAIN+TA) in Figure 4(a). As we can see, FULLMAIN+TA achieves the highest F1 scores from the first epoch. This is due to the benefits of our efficient pattern-guided self-supervised learning. On the other hand, FULLMAIN+TA converges early at around the third epoch while the other models are trained slowly and converge later. We also present the parameter size and train time of one epoch for BERT-BiLSTM (BERT), BiLSTM-BIA and FULLMAIN+TA (Our model) in Table 4(b). It can be seen that FULL MAIN+TA has fewer parameters and faster training speed.

### 6.3 How Do Our Methods Work?

We also explore the inherent properties of our methods and show how they work. In this way, we would like to point out some key ideas in designing task-oriented self-supervised tasks.

**What types of conversations are benefited?** To understand how our self-supervised tasks work, we explore the performance of six different kinds of conversations categorized by their thread patterns. Three of them (“AB”, “ABA” and “ABAB”) are focused conversations, the others (“ABC”, “ABCD” and “ABCA”) are expansionary conversations. The results together with our main model without self-supervised (FULL MAIN) are displayed in Figure 5. It can be seen that SP brings larger gains for focused conversations than expansionary ones; and RT improves the cases with repeated target users (“ABA”, “ABAB” and “ABCA”) most. Such results show that SP and RT tasks benefit the main task by improving performance on their positive instances. This raises a suggestion on designing task-oriented self-supervised tasks, i.e., *choosing tasks related to the instances that your current model is not good at*. Also, different self-supervised tasks can be proposed for different purposes in a real sys-

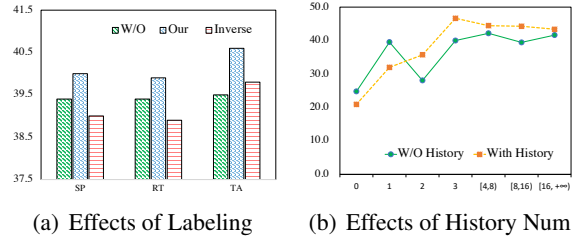


Figure 6: Fig.6(a) displays the F1 scores (in %, Y-axis) for SP, RT and TA in three different scenarios: without these tasks (W/O), the labels for these tasks are the same as main results (Our) and the labels are inverted (Inverse). For Fig.6(b), X-axis: the number of history turns that target user has, Y-axis: F1 score (in %).

tem. On the other hand, TA performs consistently better in all six cases, because the turn-level labeling emphasizes the model capability of tackling all kinds of conversations. This raises another suggestion, i.e., *designing tasks that can reflect model’s ability in different dimensions*.

**Will our methods still work if the labels are inverted?** In general, when we evaluate the performance of a task, positive instances count more than negative instances, since we care more about true positives in calculating precision and recall. Therefore, we wonder whether the labeling strategies will affect the results. To this end, we invert the labels of our self-supervised tasks by changing the label 1 to 0 and label 0 to 1. For example, we used to label focused conversations as 1 and expansionary conversations as 0 in SP task. Now we label them with opposite labels to explore how the methods work. From the results shown in Figure 6(a), the performance of inverted SP and inverted RT is even poorer than the model without self-supervised tasks (W/O). Inverted TA shows better performance than W/O, but the F1 score is still lower than the original labeled TA. We attribute such performance drop to the inconsistent labeling between auxiliary task and main task. This means that the positive label in the auxiliary task should be related to that in the main task so as to enhance learning. Therefore, we turn to the final finding in our experiments, i.e., *labeling strategies make a difference for the designed self-supervised tasks*.

### 6.4 Further Discussion

**Effects of user history.** To understand how user history affects prediction, we present F1 scores for the model with history and without history in Fig-



ure 6(b). Our model with history performs better for users having more than 1 history conversation and perform worse in the cases of only 0 and 1. This is because our model needs sufficient information to capture personalized features.

**Error analysis.** We have tried the joint training of all three auxiliary tasks and find that performance is similar to training only with TA task. This might be attributed to the difficulty of balancing among so many tasks during joint training. Another reason is that TA task has already covered the information in SP and RT task as its idea comes from the combination of the previous two tasks. On the other hand, our model performs worse in predicting the expansionary conversations (Figure 5), since most users in such conversations tend to not return, and the reasons for that might be diverse, e.g., too busy to reply. We leave how to enhance the performance in such cases as our future work.

## 7 Conclusion

We present a basic model with three novel self-supervised tasks for re-entry prediction. Experiments on two newly constructed conversation datasets, Twitter and Reddit, show that our model outperforms the previous models with fewer parameters and faster convergence. Further discussions provide more insights on how our model works and how to design task-oriented self-supervised tasks.

## Acknowledgements

The research described in this paper is partially supported by HK GRF #14204118 and HK RSFS #3133237. We thank the three anonymous reviewers for the insightful suggestions on various aspects of this work.

## References

- Lars Backstrom, Jon Kleinberg, Lillian Lee, and Cristian Danescu-Niculescu-Mizil. 2013. Characterizing and curating conversation threads: expansion, focus, volume, re-entry. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 13–22.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. 2003. A neural probabilistic language model. *Journal of machine learning research*, 3(Feb):1137–1155.
- Anders Björkelund and Jonas Kuhn. 2014. Learning structured perceptrons for coreference resolution

with latent antecedents and non-local features. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 47–57.

- Ceren Budak and Rakesh Agrawal. 2013. On participation in group chats on twitter. In *Proceedings of the 22nd international conference on World Wide Web*, pages 165–176.
- Rich Caruana, Steve Lawrence, and C Lee Giles. 2001. Overfitting in neural nets: Backpropagation, conjugate gradient, and early stopping. In *Advances in neural information processing systems*, pages 402–408.
- Jilin Chen, Rowan Nairn, and Ed Chi. 2011. Speak little and well: recommending conversations in online social streams. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 217–226.
- Qian Chen and Wen Wang. 2019. Sequential matching model for end-to-end multi-turn response selection. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7350–7354. IEEE.
- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1724–1734.
- 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*.
- Dumitru Erhan, Aaron Courville, Yoshua Bengio, and Pascal Vincent. 2010. Why does unsupervised pre-training help deep learning? In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 201–208.
- Lucie Flek. 2020. Returning the n to nlp: Towards contextually personalized classification models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7828–7838.
- Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. 2006. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554.
- Longlong Jing and Yingli Tian. 2020. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

- Diederick P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Jing Li, Wei Gao, Zhongyu Wei, Baolin Peng, and Kam-Fai Wong. 2015. Using content-level structures for summarizing microblog repost trees. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 2168–2178.
- Edward Loper and Steven Bird. 2002. Nltk: The natural language toolkit. In *Proceedings of the ACL-02 Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics - Volume 1, ETMTNLP '02*, page 63–70, USA. Association for Computational Linguistics.
- Jing Lu and Vincent Ng. 2020. Conundrums in entity reference resolution. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6620–6631.
- Huanru Henry Mao. 2020. A survey on self-supervised pre-training for sequential transfer learning in neural networks. *arXiv preprint arXiv:2007.00800*.
- Sebastian Martschat and Michael Strube. 2015. Latent structures for coreference resolution. *Transactions of the Association for Computational Linguistics*, 3:405–418.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Jeffrey Pennington, Richard Socher, and Christopher D 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.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Dana Ruitter, Cristina España-Bonet, and Josef van Genabith. 2019. Self-supervised neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1828–1834, Florence, Italy. Association for Computational Linguistics.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958.
- Ming Tan, Dakuo Wang, Yupeng Gao, Haoyu Wang, Saloni Potdar, Xiaoxiao Guo, Shiyu Chang, and Mo Yu. 2019. Context-aware conversation thread detection in multi-party chat. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6456–6461, Hong Kong, China. Association for Computational Linguistics.
- Hong Wang, Xin Wang, Wenhan Xiong, Mo Yu, Xiaoxiao Guo, Shiyu Chang, and William Yang Wang. 2019a. Self-supervised learning for contextualized extractive summarization. *arXiv preprint arXiv:1906.04466*.
- Lingzhi Wang, Jing Li, Xingshan Zeng, Haisong Zhang, and Kam-Fai Wong. 2020. Continuity of topic, interaction, and query: Learning to quote in online conversations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6640–6650, Online. Association for Computational Linguistics.
- Xiaolong Wang and Abhinav Gupta. 2015. Unsupervised learning of visual representations using videos. In *Proceedings of the IEEE international conference on computer vision*, pages 2794–2802.
- Yue Wang, Jing Li, Hou Pong Chan, Irwin King, Michael R Lyu, and Shuming Shi. 2019b. Topic-aware neural keyphrase generation for social media language. *arXiv preprint arXiv:1906.03889*.
- Jiawei Wu, Xin Wang, and William Yang Wang. 2019. Self-supervised dialogue learning. *arXiv preprint arXiv:1907.00448*.
- Yuqiang Xie, Yue Hu, Luxi Xing, Chunhui Wang, Yong Hu, Xiangpeng Wei, and Yajing Sun. 2020. Enhancing pre-trained language models by self-supervised learning for story cloze test. In *Knowledge Science, Engineering and Management*, pages 271–279, Cham. Springer International Publishing.
- Ruijian Xu, Chongyang Tao, Daxin Jiang, Xueliang Zhao, Dongyan Zhao, and Rui Yan. 2020. Learning an effective context-response matching model with self-supervised tasks for retrieval-based dialogues. *arXiv preprint arXiv:2009.06265*.
- Jianfei Yu and Jing Jiang. 2016. Learning sentence embeddings with auxiliary tasks for cross-domain sentiment classification. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016*. Association for Computational Linguistics.

Xingshan Zeng, Jing Li, Lu Wang, Nicholas Beauchamp, Sarah Shugars, and Kam-Fai Wong. 2018. Microblog conversation recommendation via joint modeling of topics and discourse. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 375–385.

Xingshan Zeng, Jing Li, Lu Wang, and Kam-Fai Wong. 2019. Joint effects of context and user history for predicting online conversation re-entries. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2809–2818, Florence, Italy. Association for Computational Linguistics.