Successful New-entry Prediction for Multi-Party Online Conversations via Latent Topics and Discourse Modeling

Lingzhi Wang
lzwang@se.cuhk.edu.hk
The Chinese University of Hong Kong
Hong Kong, China

Xingshan Zeng
zxshamson@gmail.com
Huawei Noah's Ark Lab
Hong Kong, China

Jing Li
jing-amelia.li@polyu.edu.hk
The Hong Kong Polytechnic University
Hong Kong, China

Kam-Fai Wong
kfwong@se.cuhk.edu.hk
The Chinese University of Hong Kong
Hong Kong, China

ABSTRACT

With the increasing popularity of social media, online interpersonal communication now plays an essential role in people's everyday information exchange. Whether and how a newcomer can better engage in the community has attracted great interest due to its application in many scenarios. Although some prior works that explore early socialization have obtained salient achievements, they are focusing on sociological surveys based on the small group. To help individuals get through the early socialization period and engage well in online conversations, we study a novel task to foresee whether a newcomer's message will be responded to by other participants in a multi-party conversation (henceforth **Successful New-entry Prediction**). The task would be an important part of the research in online assistants and social media. To further investigate the key factors indicating such engagement success, we employ an unsupervised neural network, Variational Auto-Encoder (**VAE**), to examine the topic content and discourse behavior from newcomer's chatting history and conversation's ongoing context. Furthermore, two large-scale datasets, from Reddit and Twitter, are collected to support further research on new-entries. Extensive experiments on both Twitter and Reddit datasets show that our model significantly outperforms all the baselines and popular neural models. Additional explainable and visual analyses on new-entry behavior shed light on how to better join in others' discussions.

CCS CONCEPTS

• **Human-centered computing** → **Collaborative and social computing**; • **Information systems** → **Personalization**.

KEYWORDS

response prediction, newcomer socialization, multi-party conversation, latent variable learning

1 INTRODUCTION

Online conversations are a crucial part of our daily communication — many people now turn to social media to share ideas and exchange information, especially when facing the lockdown caused by an epidemic outbreak (such as the *COVID-19*). Entrance to ongoing conversations (both online and offline [1]) requires a socialization process where newcomers seek to gain feedback, and the early socialization experiences have a long-term impact for newcomers [12]. In our everyday life, one should engage in a wide variety of conversations, ranging from online meetings advancing project collaborations to chitchats forming personal ideologies. However, not everyone is good at socializing [25] and online newcomers face special difficulties as a result of the diffuse, decentralized, and anonymous text-based interactions [2, 12]. It is rather challenging for a newcomer to engage in an online multi-party conversation.

In light of these concerns, there exists a pressing need to develop a conversation management toolkit to predict the conversation's future trajectory and advance the interpersonal communication quality [26, 54]. Therefore, here we focus on a novel task to predict **successful new-entries** — whether a newcomer's message will be replied by others and concretely contribute to the conversation's continuity. This task is inspired by some previous works, such as newcomers' socialization [12, 20, 21], social expression [16] and response prediction [3, 4], etc. The reasons why we employ "receiving a reply" to represent successful engagement can be summarized as twofold: (1) Pethe and Skiena [36] shows that replies, likes (which are used to show appreciation for a post [34]), retweets are highly correlated. (2) We conduct an interesting human evaluation (see Table 6), and the results show that posts that received replies are more successful than posts with silent responses in four indicators. Hence we decide to employ such a definition to simplify the task and relieve the burden of data collection and construction. Overall, this task can help newcomers avoid killing the conversations [26] and decrease the risk of withdrawal [2] because the silence outcomes can be interpreted as rudeness or unfriendliness [18] and decline
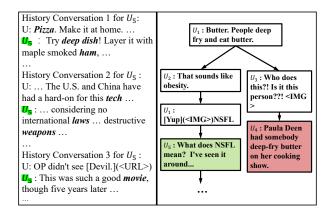
Figure 1: A Reddit conversation on the left part (U_i: **the** i**-th user).** U_3 **made a successful engagement (i.e. receiving a reply, we omit here to save space). The right column shows** U_3**'s chatting history, where the topic words are in** *bold and italic*.

the newcomers' future participation [2, 6]. More importantly, our research will potentially benefit the development of chatbots or online sales bots to understand when, what and how to say in a multi-party conversation since they should speak as few times as possible but keep the conversation going.

To solve the successful new-entry prediction task, we propose a novel framework consisting of two parts: topic and discourse modeling (TDM) and successful new-entry prediction (SNP). We examine both the conversation's ongoing contexts (henceforth **conversation contexts**) and the newcomer's chatting history (**user history**) and hypothesize that both the chatting topics (what are said) and discourse behavior (how they are said) will affect other participants' attitude to the fresh blood. To elaborate the motivation of this hypothesis, Figure 1 shows an example of a Reddit conversation snippet, where both U_4 and U_5 are newcomers. It is observed that U_5 posted a question about a new point "NSFL" (not safe for life) and hence drew future participants' attention and feedback, while U_4 made a statement via echoing "*deep-fry butter*" concerned before and didn't receive any responses. To explore the factors that affect such successful and failed new entries, we capture and distinguish topic and discourse factors with an unsupervised neural module based on variational auto-encoders (VAE) [28]. Salient words reflecting topic content (such as "pizza" and "ham" in Figure 1) derived from history conversations and discourse behavior (e.g., "what" and "?" as the question indicators) representing current trajectory are identified, serving as part of inputs for prediction module (i.e. the SNP module). Our SNP module contains a hierarchical two-layer Bi-GRU [15] to encode the conversation content for final prediction. To the best of our knowledge, we are the first to study the future trajectory of newcomers-engaged conversations and how topic and discourse factors influence their engagement success (or failure).

To summarize, the contributions of this work are as follows:

- We first formulate the task of successful new-entry prediction and contribute two large-scale datasets, Twitter and Reddit. The SNP task can benefit the development of online assistants and early socialization strategies.

- We propose a novel framework combining unsupervised and supervised neural networks. VAE and RNN-based modules are incorporated for the personalized user engagement prediction via learning latent topic and discourse representation.

- Experimental results on both Twitter and Reddit show that the proposed model significantly outperforms the baselines. For example, we achieve 34.6 F1 on Reddit compared with 32.5 achieved by a BERT-based method [19].

- Extensive analytical experiments are conducted to show the effectiveness of the modules. We probe into the learned topics and discourse and make the results explainable. Some successful entry strategies for early socialization are given based on the analysis on top of the differences between successful and failed new-entries.

2 RELATED WORK

Our work is in line with newcomer socialization, response prediction, and conversation modeling on social media platforms.

2.1 Newcomer Socialization

Newcomer socialization [12, 20, 21], which analyzes the process that newcomers make the transition from being organizational outsiders to being insiders [5], is an essential research for supporting, socializing and integrating members to virtual environments [24]. A range of user behaviors have been investigated, including information process [1], socialization [20], and social expression [16].

Most of the previous researches are based on sociological surveys in a small group while our work is based on large-scale datasets. Despite focusing on early socialization, our work also benefits the development of online assistants, e.g., chatbots.

2.2 Response Prediction

Research on response prediction aims to predict whether a given online content will receive desired responses by analyzing the propagation patterns of social media content, such as the prediction of user responses [3, 7], thread-ending turns [26], and re-entry behavior [46, 54], etc. Some of them tackle post-level responses [3, 7, 39, 40], which focus on modeling single post content; and others concern users' future behavior prediction enhanced with context modeling [26, 52, 54, 56], to achieve conversational level trajectory prediction. For the methodology, earlier works depend on the extraction of handcrafted features related to their objectives, e.g., social and sentiment features [3], arrival patterns and timing features [4], group information [11], etc. More recent achievements obtained from probabilistic graphical methods [8], neural models like RNN-based models [26, 54] and graph nerual networks [55], due to their better ability in leveraging abundant context information.

However, none of them considers newcomers in a multi-party conversation and how to help them better engage in, which is extensively investigated in this work.

2.3 Conversation Modeling

We are also inspired by the existing methods for conversation modeling, which develop models to encode conversation context from rich information, such as user interactions [47, 50, 55], temporal orders of turns [14, 26, 44], latent topics and discourse [45, 51], etc. Among them, Jiao et al. [26] and Tanaka et al. [44] only consider temporal
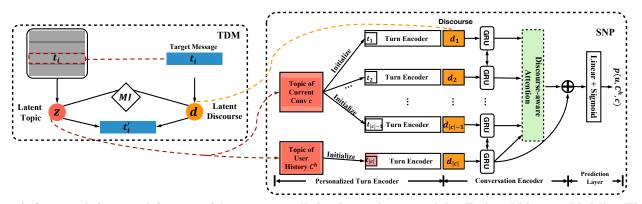
Figure 2: Our generic framework for successful new-entry prediction. It contains two modules: Topic and Discourse Modeling (TDM) and Successful New-entry Prediction (SNP). SNP consists of three parts: turn encoder, conversation encoder and prediction layer.

orders of conversation turns and use RNN-based methods to learn temporal and content information, while Wei et al. [49] and Zeng et al. [55] formulate multi-party conversations as tree structure and apply Graph Convolutional Networks (GCN) [29] to extract context-aware representations. Ritter et al. [37] and Zeng et al. [51] explore unsupervised methods to discover discourse-level information for online conversations, while the latter further leverages neural models to jointly learn distinguished topics and discourse.

In our model, the way we learn latent topic and discourse representations is built upon the success of neural latent variable models [32] for unsupervised conversation understanding [51]. Compared with the original design, we also consider user chatting history to leverage newcomers' personal interests and investigate how topics and discourse factors affect newcomers' involvement in multi-party conversations, which is beyond the capability of Zeng et al. [51].

3 OUR MODEL

Here we present our successful new-entry prediction framework. Figure 2 shows our overall framework. In the following, we first introduce the input and output in Section 3.1. Then in Section 3.2, we discuss how the TDM module works. The SNP module will be later described in Section 3.3. The learning objective of the entire framework will be given at last in Section 3.4.

3.1 Input and Output

The input for our model can be divided into two parts: the observed conversation c and the history conversation set $C^h = \{c_1^h, c_2^h, ..., c_k^h\}$ of the newcomer u, where k is the number of history conversations obtained from training set. The conversation c is formalized as a sequence of turns (e.g., posts or tweets) $\{t_1, t_2, ..., t_{|c|}\}$, and the $|c|^{th}$ turn is posted by the newcomer u (we predict whether u can get others' response afterwards). The conversations in user's history conversation set C^h are organized similarly into the sequences of turns. For output, we yield a Bernoulli distribution $p(u, C^h, c)$ to indicate the estimated likelihood of whether u gets responses from other participants (successful new-entries).

3.2 Topic and Discourse Modeling (TDM)

Inspired by Zeng et al. [51], we learn distributional word clusters that reflect the latent topic z in conversation c, and discourse behaviors $d = \langle d_1, d_2, ..., d_{|c|} \rangle$ for each turn in c. At corpus level, we assume there are K topics, each is represented by a word distribution ϕ_k^T ($k = 1, 2, ..., K$). It also contains D discourse behaviors represented with word distributions ϕ_d^D ($d = 1, 2, ..., D$).

To learn the topic and discourse representations, each turn t_i in one conversation, referred to as a **target message**, is fed into the TDM sequentially in the form of bag-of-words (BoW). For factor modeling, TDM employs the extended framework of variational auto-encoders (VAE) [28] to resemble the data generative process via two steps [32, 41]. First, given the target turn t_i and its conversation context c (i.e. the other turns in the same conversation), TDM converts them into two latent variables: topic variable z and discourse variable d_i. Then, we reconstruct the target turn t_i with the intermediate representations captured by z and d_i. In the following, we first describe the encoder followed by the decoder in detail.

Encode step. We learn the parameter μ, σ and π from the input c_{bow} and $t_{i_{bow}}$ (the BoW form of the conversation c and target turn t_i) following the formula below:

$$\mu = f_\mu(f_e(c_{bow})), \ log\ \sigma = f_\sigma(f_e(c_{bow}))$$
$$\pi = \text{softmax}(f_\pi(t_{i_{bow}})) \tag{1}$$

where $f_*(\cdot)$ is neural perceptrons performing linear transformations activated with an ReLU function [33].

Decode step. In general, the decoder learns to reconstruct the words in target message t_i. The following is the procedure:

- Draw latent topic $\mathbf{z} \sim \mathcal{N}(\mu, \sigma^2)$.
- Topic mixture $\boldsymbol{\theta} = \text{softmax}(f_\theta(\mathbf{z}))$.
- Draw the latent discourse $\boldsymbol{d} \sim \text{Multi}(\pi)$.
- For the n-th word in the conversation:
  - $\beta_n = \text{softmax}(f_{\phi^T}(\boldsymbol{\theta}) + f_{\phi^D}(\boldsymbol{d}))$
  - Draw the word $w_n \sim \text{Multi}(\beta_n)$.

In particular, the weight matrix of $f_{\phi^T}(\cdot)$ (after the softmax normalization) is considered as the topic-word distribution ϕ^T. We can also get the discourse-word distribution ϕ^D in a similar way.

We use TDM to encode both contexts of the target conversation and chatting history. For the target conversation c, we model and denote the topic variable as e_c, and the discourse behaviors of its turns as $d = \langle d_1, ..., d_{|c|} \rangle$. For the chatting history conversation set C^h of the newcomer u, we learn the topic variables for all the conversations. Then they are averaged as u's representation, denoted as e_u. This can be regarded as a kind of user embedding, which reflects their preferences and interests learned from the user history.

3.3 Successful New-entry Prediction (SNP)

This section describes how we encode conversation and predict successful new-entry via leveraging topic and discourse variables learned by TDM (described in Section 3.2). It mainly contains three parts: personalized turn encoder, discourse-aware conversation encoder, and the final prediction layer.

Personalized Turn Encoder. For each turn t_i in conversation c, we learn turn representations that are aware of personalized topic vector e_u. To that end, we first feed each word w_j in t_i into an embedding layer and get the word representation r_{ij}. Then a bidirectional gated recurrent unit (Bi-GRU) [15] is used to encode the word vector sequence of turn t_i, denoted as $\langle r_{i1}, r_{i2}, ..., r_{i,n_i} \rangle$. n_i is the number of words in t_i.

We divide the observed conversation turns into context turns (turns before the last turn) and query turn (last turn, posted by newcomer u). For query turn, we use u's topic representation e_u (produced by TDM module in Section 3.2) to initialize the aforementioned Bi-GRU. For the context turns, the topic representation e_c for conversation c is similarly utilized for initialization. Concretely, the initial states for both directions are $\overrightarrow{h_{i,0}} = \overleftarrow{h_{i,n_i}} = W^P e + b^P$, where e is e_u or e_c. For all turns, the hidden states of Bi-GRU are defined as:

$$\overrightarrow{h_{ij}} = f_{GRU}(r_{ij}, h_{i,j-1}), \overleftarrow{h_{ij}} = f_{GRU}(r_{ij}, h_{i,j+1}) \tag{2}$$

The representation of turn t_i is the concatenation of last hidden states of both directions of Bi-GRU: $h_i = [\overrightarrow{h_{i,n_i}}; \overleftarrow{h_{i,0}}]$. Finally, we get the turn-level representations of conversation c: $\langle h_1, h_2, .., h_{|c|} \rangle$.

Discourse-aware Conversation Encoder. We incorporate latent discourse behavior to model the turn interactions in the conversations, which allows better understanding of how users interact with each other in the discourse. We first concatenate the turn-level representations with the discourse variable d learned in TDM, as the input for the second Bi-GRU layer. This Bi-GRU layer is used to model conversation structure and defined as:

$$\overrightarrow{h_j^d} = f_{GRU}(s_j, h_{j-1}^d), \overleftarrow{h_j^d} = f_{GRU}(s_j, h_{j+1}^d) \tag{3}$$

where $s_j = [h_j; d_j]$ and the representation of each turn after GRU is $h_j^d = [\overrightarrow{h_j^d}; \overleftarrow{h_j^d}]$.

Then, we design an attention mechanism (henceforth discourse-aware attention) to identify discourse behavior in contexts that contribute more to signal successful engagement. Our intuition is that different discourse behaviors represent different functions, and therefore should be distinguished in the weights to make predictions. For example, a turn raising a question might be more important than a simple agreement response. Therefore, we assign different attention weights to the turns, based on their discourse behaviors:

$$a_j = f_d(\text{argmax}(d_j)) \tag{4}$$

where $\text{argmax}(d_j)$ means the learned discourse behavior to turn j, and $f_d(\cdot)$ maps the discourse behaviors to different weight values.

Finally, to produce the whole conversation representation, we concatenate the hidden state of $|c|^{th}$ turn (i.e. the query turn) $h_{|c|}^d$ and the weighted sum of all conversation turns:

$$h^c = [h_{|c|}^d; \sum_j \text{softmax}(a_j) h_j^d] \tag{5}$$

Prediction Layer. For prediction, we employ a linear projection function, activated by a sigmoid activation function, to predict how likely the newcomer can successfully chip in the conversation:

$$p(u, C^h, c) = \sigma(w^T h^c + b) \tag{6}$$

where w^T and b are trainable, and $\sigma(\cdot)$ is the sigmoid activation function.

3.4 Learning Objective

For parameter learning in our model, we design the objective function as follows, to allow the joint learning of TDM and SNP modules:

$$\mathcal{L} = \mathcal{L}_{TDM} + \mathcal{L}_{SNP} \tag{7}$$

Objective Function of TDM. Following Zeng et al. [51], $\mathcal{L}_{TDM}$ is defined as follows:

$$\mathcal{L}_{TDM} = \mathcal{L}_z + \mathcal{L}_d + \mathcal{L}_t - \lambda \mathcal{L}_{MI} \tag{8}$$

where $\mathcal{L}_z$ and $\mathcal{L}_d$ are objectives about learning topics and discourse, $\mathcal{L}_t$ is the loss for target message reconstruction, and $\mathcal{L}_{MI}$ ensures that topics and discourse learn differently.

To learn the latent topics and discourse, TDM employs the variational inference [9] to approximate posterior distribution over the latent topic z and the latent discourse d given all the training data. $\mathcal{L}_z$ and $\mathcal{L}_d$ is defined as follows:

$$\mathcal{L}_z = \mathbb{E}_{q(z|c)}[p(c \mid z)] - D_{KL}(q(z \mid c)||p(z)) \tag{9}$$

$$\mathcal{L}_d = \mathbb{E}_{q(d|t)}[p(t \mid d)] - D_{KL}(q(d \mid t)||p(d)) \tag{10}$$

where $q(z|c)$ and $q(d|t)$ are approximated posterior probabilities describing how the latent topic z and the latent discourse d are generated from the conversations and message turns. $p(c \mid z)$ and $p(t \mid d)$ represent the corpus likelihoods conditioned on the latent variables. $p(z)$ follows the standard normal prior $\mathcal{N}(0, I)$ and $p(d)$ is the uniform distribution $Unif(0, 1)$. D_{KL} refers to the Kullback-Leibler divergence that ensures the approximated posteriors to be close to the true ones.

For $\mathcal{L}_t$ that ensures the learned latent topics and discourse can reconstruct target turn t, it is defined as below:

$$\mathcal{L}_t = \mathbb{E}_{q(z|c)q(d|t)}[\log p(t \mid z, d)] \tag{11}$$

We also leverage on $\mathcal{L}_{MI}$ to guide the model to separate word distributions that represent topics and discourse. $\mathcal{L}_{MI}$ is defined as:

$$\mathcal{L}_{MI} = \mathbb{E}_{q(z)} D_{KL}(p(d \mid z)||p(d)) \tag{12}$$

Table 1: Statistics of Twitter and Reddit datasets.

	Twitter	Reddit
# of users	53,488	96,001
# of convs	37,339	69,428
# of conv turns	179,265	236,764
# of successful entries	29,340	12,199
# of failed entries	7,999	57,229
Avg turn number per conv	4.8	3.4
Avg token number per turn	20.5	20.7
Ratio of newcomer with history	0.59	0.62
Avg # of history for newcomers	2.5	6.3

Objective Function of SNP. The objective function of SNP is designed to be binary cross-entropy loss as following:

$$\mathcal{L}_{SNP} = -\sum_i \mu\, y_i log(\hat{y}_i) + (1 - y_i) log(1 - \hat{y}_i) \tag{13}$$

where $\hat{y}_i$ denotes the probability estimated from $p(u, C^h, c)$ for the i-th instance, and y_i is the corresponding binary ground truth label (1 for successful entries and 0 for the opposite). To take the potential data imbalance into account, we also adopt a trade-off weight μ to give more weight to the minority class. μ is set based on the proportion of positive and negative instances in the training set.
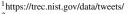
4 EXPERIMENTAL SETUP

4.1 Datasets.

We construct two new conversation datasets from **Twitter** and **Reddit**. The raw data for the **Twitter** dataset is released by Zeng et al. [53], containing Twitter conversations formed based on the TREC 2011 Microblog Track[1]. The raw data for the **Reddit** dataset contains posts and comments from Jan to May 2015, which is obtained from a publicly available Reddit corpus[2]. For both datasets, we follow the common practice to form conversations with in-reply-to relations [48, 53], where a post or a reply (comment) is considered as a conversation turn.

Our work focuses on the multi-party conversations with new users engaging in later. To that end, we remove the conversations with < 4 turns and those with < 3 participants. Finally, the datasets are randomly divided into 80%, 10%, and 10%, for training, validation, and test.

The statistics of the two datasets are shown in Table 1. As can be seen, Twitter users tend to respond to newcomers while new-entries in Reddit are more likely to be failed, probably because Twitter users are more open to public discussions compared with Reddit. We can also see that about 60% newcomers has user chatting history, which means that 60% newcomers in the test set are involved in other discussions in training data. We further study the newcomer's distribution for the number of history conversations in Figure 3. Most of the newcomers engaged in less than 5 conversations before. The sparsity in user history might pose challenges to learn their interests.

[1]https://trec.nist.gov/data/tweets/
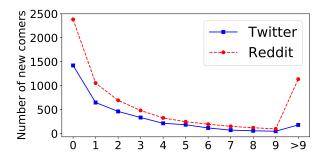[2]https://files.pushshift.io/reddit/comments/

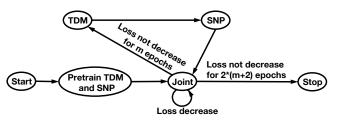**Figure 3: The distribution over the number of history conversations (X-axis). Y-axis: number of newcomers.**

Figure 4: State transition diagram of training process.

4.2 Preprocessing.

For Twitter dataset, we applied Glove tweet preprocessing toolkit [35]. As for Reddit, we first utilized the open-source natural language toolkit (NLTK) [31] for word tokenization. We then removed all the non-alphabetic tokens and replaced links with the generic tag "URL". For both datasets, a vocabulary was built and maintained with all the remaining tokens, including emoticons and punctuation. For the TDM module of our model, all stopwords were removed for topic modeling following common practice [10].

4.3 Parameter Setting.

The parameters in the TDM module are set up following Zeng et al. [51]. For the parameters in SNP module, we first initialize the embedding layer with 200-dimensional Glove embedding [35], whose Twitter version is used for Twitter dataset and Common Crawl version is applied for Reddit[3]. For the Bi-GRU layer, we set the size of its hidden states for each direction to 100 (200 for final output). The batch size is set to 64. In model training, we employ Adam optimizer [27] with initial learning rate selected among {1e-3, 1e-4, 1e-5} and early stop adoption [13]. Dropout strategy [42] is used to alleviate overfitting. All the hyper-parameters are tuned on the validation set by grid search.

4.4 Joint Train Procedure.

How to jointly train the TDM and SNP modules is critical for the performance of our model. We design to first train the TDM module for N_1 epochs (e.g., $N_1 = 100$) and train SNP module for N_2 epochs (e.g., $N_2 = 5$) respectively (while the other one is fixed), as pre-training procedure. Then we jointly train the whole model alternatively. Figure 4 is a state transition diagram showing the process of our joint training method.

[3]https://nlp.stanford.edu/projects/glove/

Table 2: Comparison results on Twitter and Reddit datasets (in %). Higher scores indicate better performance. The best results in each column are in bold. Our model gets significantly better scores than all other comparisons for all metrics ($p<0.01$, paired t-test).

Models	Twitter				Reddit			
	AUC	Accuracy	Precision	F1	AUC	Accuracy	Precision	F1
Simple Baselines								
RANDOM	50.1	49.5	78.0	60.1	49.9	50.2	15.4	24.1
HISTORY	44.1	41.2	73.2	50.6	53.6	46.5	18.1	27.5
Comparisons								
SVM	51.5	56.3	75.7	74.2	54.3	50.1	18.9	29.1
BiLSTM	52.5	77.4	78.2	87.2	59.3	54.1	22.2	31.7
BERT	70.5	80.2	80.4	89.0	63.2	51.1	21.6	32.5
CONVERNET	73.6	79.2	78.9	88.2	60.6	55.3	21.6	31.2
JECUH	75.2	80.1	80.3	88.4	60.7	57.6	22.6	31.9
Our Model	**83.2**	**82.9**	**84.7**	**90.2**	**64.8**	**62.7**	**24.9**	**34.6**

4.5 Comparisons and Evaluation Metrics

Comparisons. We first compare with two simple baselines: RANDOM, which randomly selects label 0 or 1 for prediction; HISTORY, which predicts based on the ratio of successful entries in newcomers' history conversations (random prediction is adopted for newcomers without user history).

We further compare our model with 5 different models.

1) **SVM**: SVM-based binary classifier [17] with features (e.g., TF-IDF, topic distribution, post length, thread length, etc.) gathering from Jiao et al. [26], Suh et al. [43], and Hong et al. [23],

2) **BiLSTM**: hierarchical BiLSTM used to encode conversation contexts — a BiLSTM as turn encoder and another BiLSTM to model the turn sequence, and an MLP layer works for the predictions (like our model).

3) **BERT**: a pretrained BERT is adopted to learn the turn representations and another BiLSTM to model the whole conversation, an MLP layer works for the predictions (like our model).

4) **CONVERNET**: the state-of-the-art model to predict conversation killers [26], where a few features are ignored (sentiment, background, etc.) because they are unavailable in our datasets.

5) **JECUH**: the state-of-the-art model to predict conversation re-entries [54], where the implementation is based on their code.

Evaluation Metrics. To measure the performance, we adopt popular evaluation metrics for binary classification and consider area under the ROC Curve (AUC), accuracy, precision, and F1 scores.

5 EXPERIMENTAL RESULTS

We first compare model performance in Section 5.1 including an ablation study. Then in Section 5.2, we analyze the effects of topics and discourse over successful new-entries. Section 5.3 shows the differences between successful and failed cases, and verifies with a human evaluation. At last, Section 5.4 gives more discussions on our model output with quantitative and qualitative analyses.

5.1 Main Comparison Results

Table 2 reports the main results on the two datasets, where our model significantly outperforms all comparison models. Here are more observations:

Table 3: Comparison results with ablations(in %). Higher scores indicate better performance.

Models	Twitter		Reddit	
	Accuracy	F1	Accuracy	F1
W/O TOPIC INIT	79.5	88.1	60.5	32.7
W/O DISC CONCAT	80.8	88.3	61.4	33.2
W/O DISC ATT	81.1	88.6	60.2	33.5
Full Model	**82.9**	**90.2**	**62.7**	**34.6**

• *All models perform better on Twitter than Reddit.* This suggests models' sensitivity to label imbalance, where on Reddit, we observe more sparse positive samples to learn successful new-entries compared with the negative (see Table 1).

• *Successful New-entry Prediction (SNP) is challenging.* Simple baselines such as RANDOM and HISTORY perform poorly. This indicates that SNP is challenging and impossible to be well tackled relying on simple strategies. We also find that neural models exhibit better performance than non-neural, probably benefited from their ability to learn deep semantic features from complex online interactions, which is beyond the capability of shallow features crafted manually.

• *Newcomers' chatting history is useful.* By leveraging user history, JECUH and our model perform better than both BiLSTM and ConverNet. It might be attributed to the continuity of user interests, which allows the use of more context from user history to better understand the new-entry and how it is related to the conversation contexts.

• *The coupled effects of topics and discourse are helpful to SNP.* It is seen that the joint modeling of topics and discourse helps our model to obtain the best performance. More performance gain is observed on Reddit dataset, suggesting our ability to alleviate the overfitting caused by sparse positive samples, possibly because richer features can be learned from latent topic and discourse clusters.

Ablation Study. We also examine the contributions of some components in our framework with an ablation study presented in Table 3. We compare our full model with its variants without using topics for turn encoder initialization (W/O TOPIC INIT), without concatenating

Table 4: C_v scores for Top 5 and 10 words of leaned topics. The values range from 0.0 to 1.0, and higher scores indicate better topic coherence.

Models	Twitter		Reddit	
	5	10	5	10
LDA	0.498	0.393	0.483	0.377
NTM	0.499	0.425	0.492	0.397
Our	**0.504**	**0.431**	**0.495**	**0.412**

Table 5: 5 sample latent discourse behaviors discovered from Reddit (The top 5 terms by likelihood are shown here). Names in the first column are our interpretation of the discourse behaviors according to the learned clusters. Discourse words indicating the behavior are highlighted in *blue and italic*.

Discourse	Top 5 representative terms
Disagreement	*but* have ask *different* see
Explanation	*because* stil when of that
Opinion	*think* my here ! *never*
Doubt	*n't* always want like *why*
Question	*?* For ! *where what*

discourse as turn representation (W/O DISC CONCAT) and without discourse-aware attention (W/O DISC ATT). The results indicate that topic factors might contribute more than discourse (W/O TOPIC INIT performs the worst), while their joint effects help exhibit the best results. We can also find that newcomers' user history is crucial to the model, as it is used to learn their personal interests and infer new-entries' topics.

5.2 Effects of Topic and Discourse

Topic Analysis. To analyze the topics extracted from our TDM model, we adopt C_v scores measured via the open-source Palmetto toolkit [4] to evaluate the topic coherence. C_v scores assume the top N words in a coherent topic (ranked by likelihood) tend to co-occur in the same document and have shown comparable evaluation results to human judgments [38]. We present C_v scores for top 5 and 10 words of the learned topics in Table 4 and compare them with LDA [22] and NTM [32]. The comparison results show that our model can learn a better topic representation.

To further analyze the topics learned from user history, we set the topic number $K = 2$, sample three users, and visualize their history conversations' topic mixtures in Figure 5. It is seen that users tend to be drawn by discussions with similar topics, which interprets why incorporating chatting history allows better understanding of the new-entries' topics.

Discourse Analysis. Compared with latent topics, discourse behavior is harder to be understood. So to interpret what is learned for discourse, Table 5 shows 5 example discourse behaviors from Reddit with the top 5 terms by likelihood, where meaningful discourse words are found to represent different user discourse behaviors.
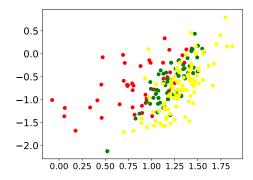
---

[4] https://github.com/dice-group/Palmetto

Figure 5: Topic mixture visualization (before softmax normalization) of three users' history conversations. A point refers to a conversation while different users are in distinguished colors.
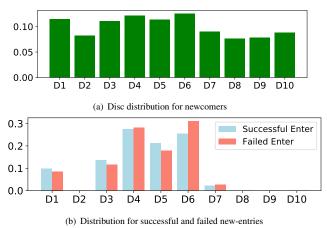

(a) Disc distribution for newcomers

(b) Distribution for successful and failed new-entries

Figure 6: 6(a) is the distribution over discourse behaviors used in new-entries. 6(b) is the distribution of discourse behaviors for successful and failed new-entries. For both, X-axis: the 10 discourse behavior learned by our model; Y-axis: corresponding probability.

To further examine newcomers' discourse behavior, Figure 6(a) shows the distribution of the latent discourse learned for new-entries. The results show that all discourse behaviors have similar frequency to be used, while some are relatively more popular, such as D4 and D6. We then compare the distribution of discourse behavior over successful and failed new-entries to Reddit conversations in Figure 6(b) (the top 2 discourse predictions are considered). As can be seen, some discourse behaviors tend to result in successful outcomes (e.g., D3 and D5), while some may easily lead to the opposite (e.g., D4 and D6). This indicates the importance for newcomers to use the right manner so as to better chip in others' discussions.

5.3 Successful and Failed New-entry

Differences in Topic Similarity. We carry out another study to examine how topics affect new-entries. We measure the similarity between the newcomers' topic distribution in their chatting history and the target conversations' contexts for successful cases (SN) and the failed ones (FN). It turns out that the similarity scores for

Table 6: Human evaluation results (%). The overall inter-rater agreement achieved Krippendorff's α of 0.74, which indicates reliable results [30].

Models	Twitter				Reddit			
	OT	AQ	CL	CS	OT	AQ	CL	CS
SUCCESSFUL	100	28	0	54	98	48	14	38
FAILED	96	4	0	6	88	20	22	20

(a) F1 with Varying History #
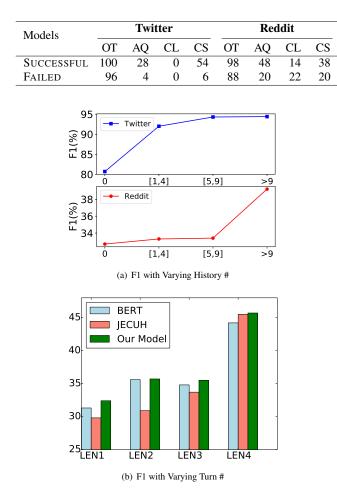

(b) F1 with Varying Turn #

Figure 7: Y-axis: F1 score. In 7(a), X-axis: user history conversation numbers. In 7(b), LENi in X-axis: the i-th quantile by turn numbers (smaller i, shorter length).

SN and FN are 38.7 and 42.7 in Twitter, while in Reddit they are 20.0 and 22.8 respectively. Interestingly, lower similarity results in more chances for success — newcomers mentioning points that are discussed less in prior history (probably cutting in with a question) are more likely to be responded to by others.

Human Evaluation. We conduct a human evaluation to explore the differences between successful and failed new-entry. Four postgraduate volunteers proficient in English are invited, and 50 conversations (25 conversations are successful cases and the others are failed) are randomly sampled from Twitter and Reddit datasets, respectively. Inspired by Arguello et al. [2] and Burke et al. [12], we propose four evaluation indicators, on-topics (OT), asking questions (AQ), complex language (CL) and controversial statement (CS). Volunteers give 0 (not match) or 1 (match) for each conversation according to the four indicators. From the results showed in Table 6,

the posts of successful newcomers are more on-topic, asking questions more often, in simple language, and usually more controversial, which is consistent with Arguello et al. [2].

5.4 Further Discussion

Effect of History Number. As discussed in Section 5.1 and 5.3, user history is essential to learn topic factors. We further analyze the change of prediction results over varying user history lengths. Test set is then divided into four subsets with the number of history conversations (x) involving a newcomer, where $x = 0$, $1 - 4$, $5 - 9$, and > 9. Our F1 scores over them are displayed in Figure 7(a). Better F1 is achieved on newcomers with longer history (engaged in more conversations), as sparse history provides limited contexts to learn topics, which will further affect SNP performance.

Effect of Conversation Length. After showing how our model performs over varying sparsity degree of user history (in Figure 7(a)), here we are interested in the model sensitivity over varying lengths of conversation contexts. Figure 7(b) shows the F1 scores over varying turn numbers in conversation contexts in Reddit. Better F1 scores are observed for conversations with more turns, as longer contexts can benefit feature learning. Also, more performance gain is observed from our model for LEN1 (very short contexts), signaling our ability to cope with data sparsity.

Qualitative Analysis. To provide more insights, we use the example in Figure 1 to conduct a qualitative analysis. It is found that our model assigns more attention weights (a_j in Eq.4) to the turn posted before U_3 as it concerns the topic "NSFL" leading to successful new-entry. We also notice that U_3 has a wide range of interests (shown in user history) and open to engaging in different types of discussions. This might be another reason why our model predicts positive outcomes, following our previous findings (Section 5.2).

Advises for Newcomers. As we showed in Section 5.3 that successful and failed newcomers show differences in topic similarity and four evaluation indicators, we give two suggestions as follows: 1) Contribute new and interesting information to the community, even it's controversial. 2) Use simple language and ask on-topic questions more for better communication.

6 CONCLUSION

This paper first formulate the task of successful new-entry prediction and collect two large-scale datasets, Twitter and Reddit. A novel model is proposed to predict successful new-entries via modeling latent topics and discourse in conversation contexts and user chatting history. This paper also explores the roles of topic and discourse played in newcomers' engagement to multi-party conversations. Extensive experiments have shown that the proposed model achieves significantly better performance than baselines and the model has learned meaningful topics and discourse representations, which are able to further signal how to make successful new-entries.

ACKNOWLEDGEMENTS

REFERENCES

[1] Manju K Ahuja and John E Galvin. 2003. Socialization in virtual groups. *Journal of Management* 29, 2 (2003), 161–185.

[2] Jaime Arguello, Brian S Butler, Elisabeth Joyce, Robert Kraut, Kimberly S Ling, Carolyn Rosé, and Xiaoqing Wang. 2006. Talk to me: foundations for successful individual-group interactions in online communities. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. 959–968.

[3] Yoav Artzi, Patrick Pantel, and Michael Gamon. 2012. Predicting responses to microblog posts. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, 602–606.

[4] 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 WSDM*. 13–22.

[5] Talya N Bauer, Todd Bodner, Berrin Erdogan, Donald M Truxillo, and Jennifer S Tucker. 2007. Newcomer adjustment during organizational socialization: a meta-analytic review of antecedents, outcomes, and methods. *Journal of applied psychology* 92, 3 (2007), 707.

[6] Nancy K Baym. 1993. Interpreting soap operas and creating community: Inside a computer-mediated fan culture. *Journal of folklore research* (1993), 143–176.

[7] Bin Bi and Junghoo Cho. 2016. Modeling a Retweet Network via an Adaptive Bayesian Approach. In *Proceedings of the 25th International Conference on World Wide Web* (Montréal, Québec, Canada) *(WWW '16)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 459–469.

[8] Bin Bi and Junghoo Cho. 2016. Modeling a retweet network via an adaptive bayesian approach. In *proceedings of the 25th WWW*. 459–469.

[9] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. 2017. Variational inference: A review for statisticians. *Journal of the American statistical Association* 112, 518 (2017), 859–877.

[10] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* 3 (March 2003), 993–1022.

[11] Ceren Budak and Rakesh Agrawal. 2013. On participation in group chats on twitter. In *Proceedings of the 22nd WWW*. 165–176.

[12] Moira Burke, Robert Kraut, and Elisabeth Joyce. 2010. Membership claims and requests: Conversation-level newcomer socialization strategies in online groups. *Small group research* 41, 1 (2010), 4–40.

[13] 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*. 402–408.

[14] Hao Cheng, Hao Fang, and Mari Ostendorf. 2017. A factored neural network model for characterizing online discussions in vector space. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. 2296–2306.

[15] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* (2014).

[16] Alton YK Chua and Radhika Shenoy Balkunje. 2013. Beyond knowledge sharing: interactions in online discussion communities. *International Journal of Web Based Communities* 9, 1 (2013), 67–82.

[17] Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning* 20, 3 (1995), 273–297.

[18] Catherine Durnell Cramton. 2002. Attribution in distributed work groups. *Distributed work* (2002), 191–212.

[19] 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* (2018).

[20] Rosta Farzan, Robert Kraut, Aditya Pal, and Joseph Konstan. 2012. Socializing volunteers in an online community: a field experiment. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. 325–334.

[21] John E Galvin and Manju K Ahuja. 2001. Am I doing what's expected? New member socialization in virtual groups. In *Our virtual world: The transformation of work, play and life via technology*. IGI Global, 40–55.

[22] Thomas L Griffiths, Michael I Jordan, Joshua B Tenenbaum, and David M Blei. 2004. Hierarchical topic models and the nested chinese restaurant process. In *Advances in neural information processing systems*. 17–24.

[23] Liangjie Hong, Ovidiu Dan, and Brian D Davison. 2011. Predicting popular messages in twitter. In *Proceedings of the 20th international conference companion on World wide web*. 57–58.

[24] Ian Horrocks. 2002. *The Semantic Web-ISWC 2002: First International Semantic Web Conference, Sardinia, Italy, June 9-12, 2002, Proceedings*. Vol. 1. Springer Science & Business Media.

[25] Yeh Hsueh, Jun Hao, and Hui Zhang. 2015. The needs and difficulties in socializing the young in contemporary China: Early childhood education experts' perspectives. *Policy Futures in Education* 14, 1 (2015), 123–128.

[26] Yunhao Jiao, Cheng Li, Fei Wu, and Qiaozhu Mei. 2018. Find the conversation killers: A predictive study of thread-ending posts. In *Proceedings of the 2018 World Wide Web Conference*. 1145–1154.

[27] Diederick P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*.

[28] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).

[29] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.

[30] Klaus Krippendorff. 2004. *Content analysis: An introduction to its methodology*.

[31] 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* (Philadelphia, Pennsylvania) *(ETMTNLP '02)*. Association for Computational Linguistics, USA, 63–70.

[32] Yishu Miao, Edward Grefenstette, and Phil Blunsom. 2017. Discovering Discrete Latent Topics with Neural Variational Inference. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 70)*, Doina Precup and Yee Whye Teh (Eds.). PMLR, International Convention Centre, Sydney, Australia, 2410–2419.

[33] Vinod Nair and Geoffrey E. Hinton. 2010. Rectified Linear Units Improve Restricted Boltzmann Machines. In *Proceedings of the 27th International Conference on International Conference on Machine Learning* (Haifa, Israel) *(ICML'10)*. Omnipress, USA, 807–814.

[34] Bridianne O'dea, Melinda R Achilles, Mark E Larsen, Philip J Batterham, Alison L Calear, and Helen Christensen. 2018. The rate of reply and nature of responses to suicide-related posts on Twitter. *Internet interventions* 13 (2018), 105–107.

[35] 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)*. 1532–1543.

[36] Charuta Pethe and Steven Skiena. 2019. The Trumpiest Trump? Identifying a Subject's Most Characteristic Tweets. *arXiv preprint arXiv:1909.04002* (2019).

[37] Alan Ritter, Colin Cherry, and Bill Dolan. 2010. Unsupervised modeling of twitter conversations. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 172–180.

[38] Michael Röder, Andreas Both, and Alexander Hinneburg. 2015. Exploring the space of topic coherence measures. In *Proceedings of the eighth ACM international conference on Web search and data mining*. 399–408.

[39] Matthew Rowe and Harith Alani. 2014. Mining and comparing engagement dynamics across multiple social media platforms. In *Proceedings of the 2014 ACM conference on Web science*. 229–238.

[40] Matthew Rowe, Sofia Angeletou, and Harith Alani. 2011. Predicting discussions on the social semantic web. In *Extended Semantic Web Conference*. Springer, 405–420.

[41] Akash Srivastava and Charles Sutton. 2017. Autoencoding variational inference for topic models. *arXiv preprint arXiv:1703.01488* (2017).

[42] 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 (Jan. 2014), 1929–1958.

[43] Bongwon Suh, Lichan Hong, Peter Pirolli, and Ed H Chi. 2010. Want to be retweeted? large scale analytics on factors impacting retweet in twitter network. In *2010 IEEE Second International Conference on Social Computing*. IEEE, 177–184.

[44] Koji Tanaka, Junya Takayama, and Yuki Arase. 2019. Dialogue-Act Prediction of Future Responses Based on Conversation History. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28 - August 2, 2019, Volume 2: Student Research Workshop*. 197–202.

[45] Lingzhi Wang, Jing Li, Xingshan Zeng, Haisong Zhang, and Kam-Fai Wong. [n. d.]. Continuity of Topic, Interaction, and Query: Learning to Quote in Online Conversations. In *Proceedings of the EMNLP2020*. 6640–6650.

[46] Lingzhi Wang, Xingshan Zeng, Huang Hu, Kam-Fai Wong, and Daxin Jiang. 2021. Re-entry Prediction for Online Conversations via Self-Supervised Learning. In *Findings of the Association for Computational Linguistics: EMNLP 2021*. Association for Computational Linguistics, Punta Cana, Dominican Republic, 2127–2137. https://doi.org/10.18653/v1/2021.findings-emnlp.183

[47] Lingzhi Wang, Xingshan Zeng, and Kam-Fai Wong. 2021. Quotation Recommendation and Interpretation Based on Transformation from Queries to Quotations. In *Proceedings of the 59th ACL*. 754–758.

[48] Yue Wang, Jing Li, Hou Pong Chan, Irwin King, Michael R. Lyu, and Shuming Shi. 2019. Topic-Aware Neural Keyphrase Generation for Social Media Language. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.

[49] Penghui Wei, Nan Xu, and Wenji Mao. 2019. Modeling Conversation Structure and Temporal Dynamics for Jointly Predicting Rumor Stance and Veracity. 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 2019, Hong Kong, China, November 3-7, 2019. 4786–4797.

[50] Victoria Zayats and Mari Ostendorf. 2018. Conversation Modeling on Reddit Using a Graph-Structured LSTM. *TACL* 6 (2018), 121–132.

[51] Jichuan Zeng, Jing Li, Yulan He, Cuiyun Gao, Michael R. Lyu, and Irwin King. 2019. What You Say and How You Say it: Joint Modeling of Topics and Discourse in Microblog Conversations. *Transactions of the Association for Computational Linguistics* 7 (March 2019), 267–281.

[52] Jichuan Zeng, Jing Li, Yulan He, Cuiyun Gao, Michael R. Lyu, and Irwin King. 2020. What Changed Your Mind: The Roles of Dynamic Topics and Discourse in Argumentation Process. In *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, Yennun Huang, Irwin King, Tie-Yan Liu, and Maarten van Steen (Eds.). ACM / IW3C2, 1502–1513.

[53] 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). 375–385.

[54] 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*. Association for Computational Linguistics, Florence, Italy, 2809–2818.

[55] Xingshan Zeng, Jing Li, Lu Wang, and Kam-Fai Wong. 2019. Neural Conversation Recommendation with Online Interaction Modeling. 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)*. Hong Kong, China, 4625–4635.

[56] Xingshan Zeng, Jing Li, Lingzhi Wang, and Kam-Fai Wong. 2021. Modeling Global and Local Interactions for Online Conversation Recommendation. *ACM Trans. Inf. Syst.* 40, 3, Article 49 (nov 2021), 33 pages. https://doi.org/10.1145/3473970