

# Improving Open-Domain Dialogue Evaluation with a Causal Inference Model

Cat P. Le, Luke Dai, Michael Johnston, Yang Liu, Marilyn Walker, Reza Ghanadan

**Abstract** Effective evaluation methods remain a significant challenge for research on open-domain conversational dialogue systems. Explicit satisfaction ratings can be elicited from users, but users often do not provide ratings when asked, and those they give can be highly subjective. Post-hoc ratings by experts are an alternative, but these can be both expensive and complex to collect. Here, we explore the creation of automated methods for predicting both expert and user ratings of open-domain dialogues. We compare four different approaches. First, we train a baseline model using an end-to-end transformer to predict ratings directly from the raw dialogue text. The other three methods are variants of a two-stage approach in which we first extract interpretable features at the turn level that capture, among other aspects, user dialogue behaviors indicating contradiction, repetition, disinterest, compliments, or criticism. We project these features to the dialogue level and train a dialogue-level MLP regression model, a dialogue-level LSTM, and a novel causal inference model called counterfactual-LSTM (CF-LSTM) to predict ratings. The proposed CF-LSTM is a sequential model over turn-level features which predicts ratings using multiple regressors depending on hypotheses derived from the turn-level features. As a causal inference model, CF-LSTM aims to learn the underlying causes of a specific event, such as a low rating. We also bin the user ratings and perform classification experiments with all four models. In evaluation experiments on conversational data from the Alexa Prize SocialBot, we show that the CF-LSTM achieves the best performance for predicting dialogue ratings and classification.

**Key words:** open-domain dialogue, user ratings, dialogue evaluation, causal inference, user satisfaction

---

Cat P. Le

Duke University, e-mail: cat.le@duke.edu.

This work was done while Cat P. Le was a research intern at Amazon Alexa AI.

Luke Dai, Michael Johnston, Yang Liu, Marilyn Walker, Reza Ghanadan

Amazon Alexa AI, e-mail: {lukedai, mjohnstn, yangliud, wamarilk.ghanadan}@amazon.com

## 1 Introduction

Evaluation has always been a complex challenge for interactive dialogue systems. For task-oriented dialogues, frameworks such as Paradise [51] model the relationship between user satisfaction, task completion, and cost factors such as dialogue length, word error rate, and dialogue behaviors. However, open-domain dialogue systems such as those built for the Alexa Prize SocialBot Grand Challenge [36, 9], where there is no clearly defined task, require new metrics and methods for evaluation that better reflect their affordances [49, 21, 13, 12, 16, 47].

In Alexa Prize SocialBot, user conversation ratings are elicited primarily for comparative evaluation of systems competing in the Grand Challenge. Numerous other use cases for dialogue ratings include selecting training data and running A/B tests to evaluate new capabilities, features, and models. Manually collected user ratings are limited in that only a fraction of users of “Alexa Let’s Chat” leave ratings, and these ratings can be highly subjective. An alternative is the post-hoc annotation of dialogues by trained annotators, which is complex, expensive, and thus not scalable. Consequently, a significant amount of dialogue data is unavailable for experimentation and modeling tasks.

To overcome this problem and increase the number of rated dialogues, we evaluate four methods for automatically predicting both user (128K) and expert (1.2K) ratings of de-identified conversations collected from interactions with an Alexa SocialBot. First, we train a baseline method using a transformer-based model [5] to predict dialogue ratings based on the raw input dialogue sequence. The other three methods are variants of a two-stage approach in which we use pre-trained models to extract interpretable features at the turn level. These models include classifiers we call ODES (Open Domain Evaluation Signals) that detect user inputs such as criticism, insults, compliments, requests to stop, and statements that indicate misunderstanding, repetition, or contradiction among other classes [49]. DialogPT [57] and DialogRPT [10] are applied to the dialogue sequence to provide measures of response relevance and specificity, and FED [30] is applied to each turn-pair to provide metrics for the interestingness, engagingness, specificity, relevance, correctness, semantic appropriateness, understandability and fluency of responses. In addition to these derived measures, we also include ASR confidence and sentiment scores based on processing the acoustics of user utterances. These features are projected to the dialogue level as averages, counts, initial, final, and penultimate values and used as input to our three models.

The first model trains an MLP regressor to predict dialogue ratings. The second model, called dialogue-level LSTM, uses an LSTM model on the sequence of turn-level features. The third approach models the rating using a shared sequential model over the turn-pair level vectors and introduces a causal structure into the prediction model. This proposed Counterfactual-LSTM (CF-LSTM) is a causal inference model that uses different models to predict the overall rating under different hypotheses derived from the ODES classifiers. In particular, a treatment (e.g., 0 or 1) is applied to each dialog based on the text classification. For instance, dialog with negative texts, such as critique, disinterest, and insult, will be assigned a treatment

of 1. On the other hand, dialogs without negative texts will be assigned a treatment of 0. Next, a causal inference neural network is trained using the dialog features and their corresponding treatments to predict the customer’s rating. This approach introduces the domain knowledge into the prediction model and helps further improve the performance. Multiple different regression functions are used based on the hypothesis. Lastly, we apply all four models to classify dialogues into bins according to their ratings (e.g., binary and 5-class classification). Our contributions are:

1. We identify and train classifiers to extract various types of dialogue features (e.g., sentiment, response relevance, and specificity) and use these features to predict dialogue ratings.
2. We observe the special conditions of the open-domain dialogue evaluation problem and make three assumptions to apply causal inference to this problem.
3. We introduce three approaches and propose a novel causal inference model, CF-LSTM, that is used to evaluate open-domain dialogues via mapping hidden dialogue features to the dialogue ratings.

## 2 Related Work

Much recent work on dialogue evaluation has focused on the utterance level, measuring the goodness of a particular system response in the dialogue context. This type of evaluation can be either reference-based or referenceless. Reference-based approaches re-apply metrics from tasks such as machine translation or summarization and compare a proposed system utterance to a reference response, which may be from a pre-existing dialogue [15], or collected via crowd-sourcing [20]. For instance, BLEU [32] computes the n-gram precision of the system’s response string when compared to a reference response, while METEOR [3] considers the stems and synonyms of the reference and ROUGE [28] uses n-gram recall instead of n-gram precision. More recent neural versions of these metrics, such as BERTScore and BLEURT, use word embeddings rather than the raw strings [55, 41, 43]. However, it is well-known that these metrics have limited value for task-oriented dialogue and even less utility when evaluating the highly varied responses typical of open-domain dialogues [38, 11, 29].

Work on referenceless metrics aims to predict the quality of system response in the dialogue context without using reference utterances. These metrics evaluate the proposed system utterance in terms of coherence, interestingness, engagingness, fluency, specificity, relevance, empathy, or other measures. For example, GRADE measures coherence in terms of a graph-based representation of topic relatedness [19]. DialogRPT can be used to provide metrics that estimate the likelihood that an utterance generates many responses, directly or indirectly [10], based on training with 133M utterance pairs from Reddit, consisting of utterances and their upvotes, downvotes, and number of responses.

Another idea is to score a proposed system utterance by the probability that it elicits a particular user response type, such as disinterest or criticism [46, 31, 49, 12].

For example, the FED framework scores proposed utterances by utilizing the Di-aloGPT LM probabilities for the subsequent user utterances such as “That is interesting” [56, 31]. Predictive models can also be generated from large dialogue corpora by using the following user utterance as weak supervision [12, 46].

Other research aims to predict the overall quality of the dialogue by developing models for predicting a quality measure in terms of particular types of dialogue events. The most straightforward framework is end-to-end (E2E) approaches using the raw dialogue text as input to a transformer model [13, 5, 14]. This approach performs well with expert-rated data but does not predict user-assigned ratings well. We use this approach as our E2E baseline below.

Other work first derives several intermediate metrics at the turn or utterance level and then uses those to predict the dialogue quality, as we do here. The first framework like this was PARADISE, which included task completion as an input to the model, along with other metrics such as ASR error rates, system dialogue acts, or user requests to start over [51, 50]. Recent work on PARADISE has used other features, such as user complaints and compliments, to apply it to open-domain dialogue [49]. Many recent approaches work with embedding features, such as sentiment scores and utterance classification, before predicting the user rating [53, 14]. Sentiment analysis [21, 12] and dialogue features analysis [30, 42, 49] methods, for instance, show solid results for predicting user-assigned ratings.

Furthermore, many approaches to open-domain dialogue evaluation use large-scale corpora [4, 7, 6, 18, 44]. Since the Alexa Prize challenge was introduced in recent years, several new controlled dialogue datasets have been created [54, 37, 27, 15]. However, evaluating the Alexa Prize SocialBot remains a challenging problem due to the nature of the collected dialogues. These dialogues are de-identified conversations between thousands of users and the Alexa SocialBot covering numerous topics and genres, with widely varying lengths (e.g., the SocialBot dialogues vary in length from 3 to 230 turn-pairs). Additionally, the SocialBot is expected to have up-to-date knowledge about news, sports, movies, music, books, etc. Hence, the evaluation of the Alexa SocialBot requires a dynamic structure that takes into account domain knowledge.

In this paper, we introduce a novel causal inference structure to open-domain dialogue evaluation to improve the performance of the prediction model for the Alexa SocialBot. Causal inference [40, 39] has been extensively investigated in the neural network literature [8, 45, 1, 52, 2] because of the black box nature of neural models. Notably, the TARNet model [45], a causal inference neural network, has been applied to causal inference for various regression and classification problems. It can be applied to many machine learning models, such as image generative models [52] and Gaussian processes [1]. However, we have yet to be aware of previous work that uses causal inference for open-domain dialogue evaluation. Here we propose a novel LSTM causal structure for predicting the quality of open-domain dialogue to make the evaluation model more robust.

### 3 Proposed Methods

We propose three approaches to tackle open-domain dialogue evaluation. These approaches convert each turn pair in the dialogue into a vector representation, then process the vectors to predict the dialogue level rating. Our proposed methods consist of two main components.

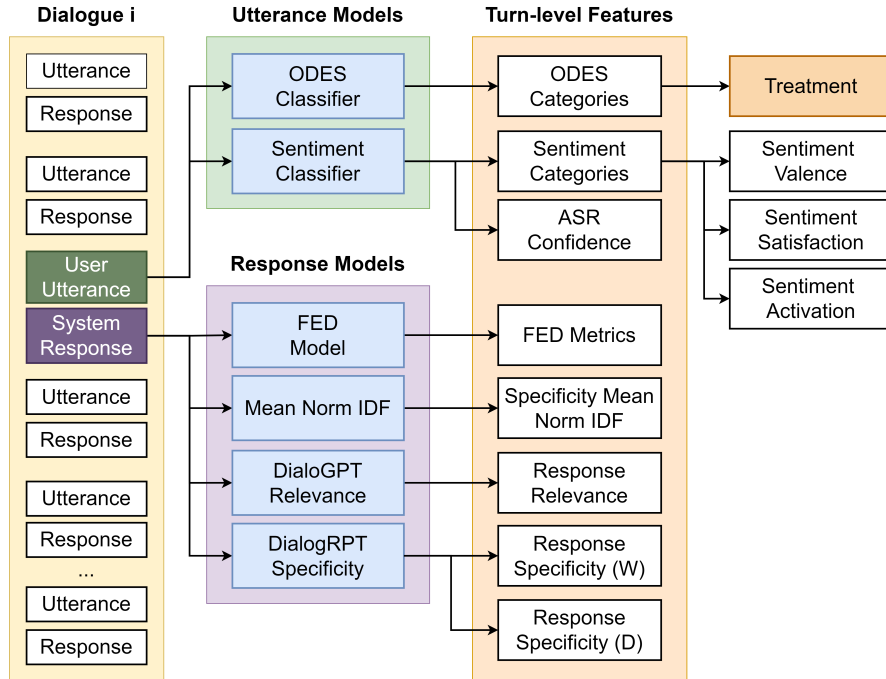
1. **Dialogue Turn-level Feature Extractors.** These extractors are applied to extract meaningful hidden features from the dialogue as in Figure 1. Subsequently, these turn-level features are converted to the dialogue-level features that are fed into the prediction models.
2. **Prediction Models.** We utilize an end-to-end transformer on the raw dialogue text to create a baseline. We then compare this baseline to three proposed models for predicting both user and expert ratings from the hidden dialogue-level features:
  - Dialogue-level MLP (Section 3.2)
  - Dialogue-level LSTM (Section 3.2)
  - Counterfactual LSTM (Section 3.3)

The dialogue-level MLP uses an MLP regressor to project the turn features to the user ratings, while the dialogue-level LSTM utilizes a standard LSTM to predict ratings from the sequence of turn-level features. The counterfactual LSTM (CF-LSTM) is a novel causal inference model. It consists of the LSTM layers augmented with the causal analysis model. This causal model consists of multiple MLP regressors, each trained to predict the ratings based on a specific hypothesis.

#### 3.1 Dialogue Turn-level Feature Extractors

To extract features for each utterance-response pair, we apply 6 models (i.e., 2 models for the utterances and 4 models for the responses), as shown in Figure 1. We apply the Open-Dialogue Evaluation Signals (ODES) Classifier and an Audio-based Sentiment Classifier for the utterance data. The ODES classifier, which is based on the T5 architecture [35], was trained on 8K annotated utterances to categorize them into the classes shown in Table 1: (i) user disinterest, (ii) user critique, (iii) user not understand, (iv) user requests topic switch, (v) user obscenity, (vi) user rejects topic switch, (vii) user requests repeat, (viii) user requests to stop, (ix) user insult, (x) user compliment, (xi) user calls out repetition, (xii) user calls out contradiction, (xiii) system not understand, and (xiv) others. This T5 classifier achieves above 95% accuracy on the testing corpora. Next, the ODES categories assign the treatment  $T$  for the CF-LSTM.

We also extract 3 different types of utterance sentiment features: sentiment valence, satisfaction, and activation, which are highly correlated with user ratings [21, 12].



**Fig. 1** The dialogue turn-level feature extractors consist of 2 utterance models (i.e., ODES classifier, sentiment classifier) and 4 response models (i.e., FED model, Mean Norm IDF, DialogPT Relevance, and DialogRPT Specificity). The treatment assignments for our causal model are extracted using the ODES classifiers.

To generate scores for the responses, we apply the FED model [30], the DialogPT model [56] for response relevance, the DialogRPT model [10] for response specificity width and depth, and Mean Norm IDF as another measure of response specificity [30, 56]. The concatenation of these extracted features forms the representation of the utterance-response pair, called turn-level features. These features are normalized to their Z-scores before being converted into dialogue-level features to train the prediction models.

### 3.2 Dialogue-level MLP and LSTM

The dialogue-level MLP is the two-level approach that uses the MLP [33] regression to project the turn pair features to the ratings. Here, we consider the dialogue features that are an aggregation of the turn-level features by choosing a few fundamental values, such as average, initial, final, and penultimate values of the sequence of turn-level features. In this case, the dialogue features (i.e., input of the

**Table 1** Examples of dialogues for each ODES category in the Alexa SocialBot dataset.

Class	ODES Name	Counts	Example
i	User disinterest	10,938	<b>SocialBot:</b> Did you know that otters sleep holding hands? <b>User:</b> I really couldn't care less.
ii	User critique	12,290	<b>SocialBot:</b> Who do you think will win Superbowl 2019? <b>User:</b> You are really a very stupid bot.
iii	User not understand	12,399	<b>SocialBot:</b> What is your favorite genre of video game? <b>User:</b> I don't know what genre means.
iv	User requests topic switch	21,278	<b>SocialBot:</b> Who is your favorite Batman actor? <b>User:</b> Can we talk about something else?
v	User obscenity	66,532	<b>SocialBot:</b> Do you have a favorite movie? <b>User:</b> B*** me.
vi	User rejects topic switch	4,389	<b>SocialBot:</b> I love country music too. But I was wondering, do you have a favorite sport? <b>User:</b> No, I want to keep talking about music.
vii	User requests to repeat	26,779	<b>SocialBot:</b> Would you rather climb the Eiffel Tower or the Empire State Building? <b>User:</b> Could you say that again?
viii	User requests to stop	78,504	<b>SocialBot:</b> I just love talking about music. What is your favorite kind of music? <b>User:</b> Please stop I need to go to bed.
ix	User insult	12,173	<b>SocialBot:</b> Wow, Taylor Swift has 97 albums. That's a lot! <b>User:</b> You are so full of sh**.
x	User compliment	57,539	<b>SocialBot:</b> Did you know LeBron James Jr has a college football scholarship, and he's only 12 years old? <b>User:</b> That's really interesting.
xi	User calls out repetition	15,052	<b>SocialBot:</b> What's your favourite football team? <b>User:</b> You already asked me that question twice.
xii	User calls out contradiction	3,147	<b>SocialBot:</b> I don't have any pets. <b>User:</b> You just said you had a cat.
xiii	System not understand	12,534	<b>User:</b> Can we talk about Elle King? <b>SocialBot:</b> I like BB King too. <b>User:</b> That's not what I said.
xiv	Others	2,320,515	<b>SocialBot:</b> Do you like K-Pop music? <b>User:</b> Yes, I often listen to Blackpink and BTS.

Dialogue-level MLP) have the same length and can be fed into the MLP regressor for prediction.

The dialogue-level LSTM is a method that uses a standard LSTM [17] on the dialogue-level features (i.e., a sequence of turn-level features) as in Table 3.1. This model is similar to the one used for speech sentiment analysis [21], in which the LSTM model is used to map the sequence of turn-level sentiment features (i.e., sentiment valence, sentiment satisfaction, sentiment activation sentiment scores) to the ratings. Here, the input of the model is the concatenated turn-level features. Since this model is based on LSTM layers, it can handle the varied-length input (e.g., dialogue-level feature vectors).

### 3.3 Counterfactual LSTM

In recent natural language processing (NLP) tasks, the assumption that training and test data are identically distributed is often not satisfied. Bias and variance from the data are often introduced into the standard NLP approaches (e.g., non-causal models), which makes the training procedure difficult, and leads to poor results in real-world test data. We propose using a causal structure containing domain knowledge to overcome this issue. This structure helps introduce inductive biases and leads to more robust predictors [8]. To apply causal inference to our problem, three common conditions must be satisfied: Ignorability, Positivity, and Consistency [8].

- **Ignorability.** In causal inference and statistics, the ignorability requires that the data collection method be independent of the treatment assignment. In other words, the outcomes (i.e., factual and counterfactual) do not depend on the treatment assignment.

$$T \perp\!\!\!\perp Y(T = a), \quad \forall a \in \{0, 1\} \quad (1)$$

- **Positivity.** The probability of a sample being assigned the treatment  $T = a$  for  $a \in \{0, 1\}$  is bounded between 0 and 1.

$$0 < \mathbb{P}(T = a | X = x) < 1, \quad \forall a \in \{0, 1\} \quad (2)$$

- **Consistency.** A sample under a treatment  $T = a$ , for  $a \in \{0, 1\}$  has a unique outcome. The outcome of a sample under treatment  $T = a$  is identical to the outcome if that sample is again assigned to  $T = a$ .


$$T = a \iff Y(a) = Y, \quad \forall a \in \{0, 1\} \quad (3)$$

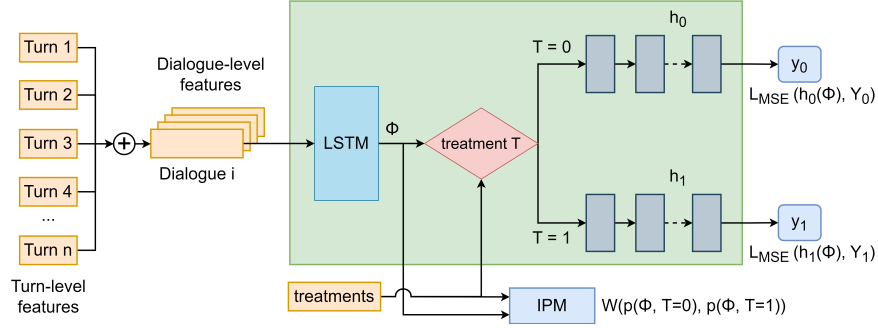
Here, the positivity and consistency conditions are satisfied. To satisfy the ignorability condition, we must assume that the extracted dialogue features, described in Section 3.1, capture everything relevant about the dialogue to predict the ratings. In other words, we assume that all confounding variables  $X$  are observed and that there is no hidden confounder. This property, called *conditional ignorability*, is described as follows:

$$(T \perp\!\!\!\perp (Y(a)) | X), \quad \forall a \in \{0, 1\} \quad (4)$$

We attempt to satisfy the ignorability condition by using many features in our models that were proposed in previous work, as well as introducing our own novel ODES classifiers, but we acknowledge that future work may derive additional fea-

**Fig. 2** The causal relationship between outcomes, treatments, and input features. The outcome  $Y$  (e.g., customer ratings) is a function of the dialogue features  $X$ , and the treatments  $T$ .





**Fig. 3** The diagram of the proposed counterfactual LSTM model. The model consists of the LSTM layers and the causal structure  $h_0, h_1$ . It learns to map the dialogue-level features to the user ratings based on the treatment  $T$ .

tures from the dialogue events, such as dialogue acts or sub-dialogue level features [50, 34].

Next, we propose a novel causal inference model called the Counterfactual LSTM, which maps the dialogue-level features to the ratings based on a specific hypothesis. This paper considers a simple case with only 2 hypotheses. First, the turn-level features are stacked into a dialogue-level feature vector for each dialogue. Since each dialogue has a different number of turn pairs, the dialogue-level features are varied-length vectors. These dialogue-level features represent the dialogue textual data fed into the CF-LSTM model to predict the dialogue ratings. As shown in Figure 3, the proposed CF-LSTM model consists of the LSTM layers and 2 individual MLP regressors, each used to train on a specific hypothesis or treatment. The inputs of the LSTM layers are the dialogue-level feature vectors, which are sequences of turn-level features. Since the LSTM can take varied-length inputs, this property also holds for CF-LSTM and can take varied-length dialogue-level feature vectors. Next, the output features  $\Phi$  of the LSTM are fed to the specific MLP ( $h_0$  or  $h_1$ ) based on the hypothesis and trained to predict dialogue ratings.

### 3.3.1 Definition of Treatments

We utilize the extracted ODES features from Section 3.1 as the treatment groups. Let  $T$  denote the treatment assignment. The treatment  $T = 0$  is the hypothesis group whose dialogues consist of only the ODES category (xiv), i.e., other. The treatment  $T = 1$  is the hypothesis group whose dialogues consist of one or more ODES categories (i) to (xii). Since the CF-LSTM has a dedicated MLP for each dialogue hypothesis, it can predict dialogue ratings for similar dialogue with a different hypothesis. For instance, two dialogues with different dialogue ratings have the same negative sentiment features. One dialogue is categorized into ODES group (i), i.e., user disinterest, and is assigned  $T = 1$ . On the other hand, the sentiment of the other dialogue is classified as unfavorable, and the user appears to be satisfied with the

conversation, e.g.,  $T = 0$ . Non-causal regressors can quickly group these dialogues since they contain the same sentiment features. Still, the CF-LSTM with a causal structure can identify the difference between these two dialogues and better predict the rating.

### 3.3.2 Loss Function of CF-LSTM

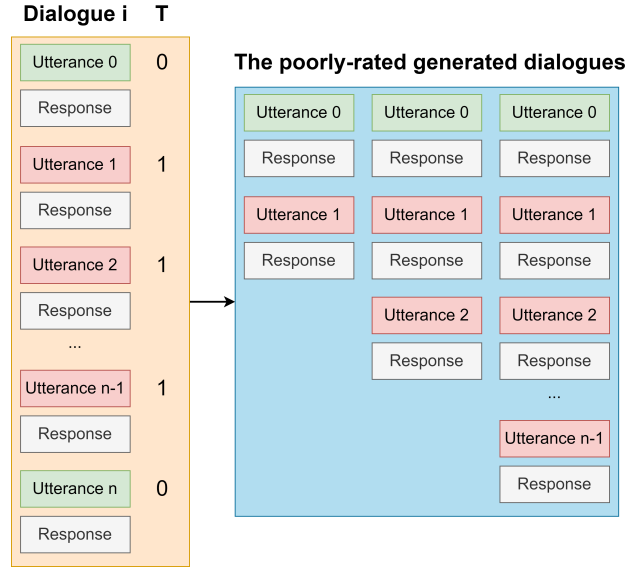
Here, we define the loss function for our proposed CF-LSTM model. First, most dialogues with the treatment  $T = 1$  have relatively low ratings. For example, the dialogues in group (viii), i.e., user insult, will often receive ratings of 2 or less. Additionally, the ratings are given non-uniformly among different groups of users. These are the source of bias and variance in the dataset, and a large proportion of users choose not to leave a rating. In order to avoid introducing variance and bias into the model, the Integral probability metric (IPM) is established in [45]. This metric is used as the upper bound for this source of variance. The IPM measures the distance between two distributions:  $p(\Phi, T = 0), p(\Phi, T = 1)$ . In this paper, we use Wasserstein distance [48] as the IPM measurement in our CF-LSTM model. Let  $Y_0$  and  $Y_1$  denote the ground truth ratings for the treatment  $T = 0$  and  $T = 1$ , respectively. The loss function for the CF-LSTM model is defined as follows:

$$L = L_{MSE}(h_0(\Phi), Y_0 | T = 0) + L_{MSE}(h_1(\Phi), Y_1 | T = 1) + W(p(\Phi, T = 0), p(\Phi, T = 1)) \quad (5)$$

where  $L$  denotes the loss function of the CF-LSTM model,  $L_{MSE}$  is the mean square error (MSE) loss,  $W$  is the Wasserstein distance,  $h_0()$  and  $h_1()$  are the outputs of the MLP regressors  $h_0$  and  $h_1$ , respectively. This loss function consists of two main parts, the MSE loss, and the IPM loss. The MSE loss is the error between the predicted and ground truth ratings under every hypothesis (e.g.,  $T = 0$  and  $T = 1$ ). The IPM loss is defined as the Wasserstein distance between 2 distributions of ratings described above.

### 3.3.3 Data Generation for Classification Tasks

We consider the classification task of detecting low-quality dialogues at run-time. The model needs to be trained with sufficient data samples to dynamically identify the low-rated dialogues with high accuracy. It is often not satisfied in practice since most collected data samples have good ratings. As a result, we want to generate more dialogues with low ratings so that the model can be trained to identify these types of data quickly. It is also essential since it determines whether the model can be applied to evaluate dialogues at run-time. This online-evaluation mechanism can also improve the overall user experience since the system could navigate to a new topic if it realizes that the current conversation is rated poorly. Here, we first select the poorly-rated dialogue  $i$  and augment these data by masking utterances. We can



**Fig. 4** The procedure of generating new dialogues by masking utterances. Let dialogue  $i$  be the low-rated dialogue with  $n + 1$  utterances. Each utterance is given the treatment of 1 and 0. At turn  $m$ , if the treatment for utterance  $m$  is 1, we define a new poorly-rated dialogue with the utterances 0 to  $m$  and masking the remaining utterances, i.e., utterances  $(m + 1)$  to  $n$ .

assign a treatment to each utterance in the dialogue from the extracted turn-level features. At utterance  $m$ , if the treatment is 1, we define a new dialogue, labeled as 0 (i.e., poorly rated dialogue), that consists of the utterance 0 to the utterance  $m$ . In other words, we mask all of the utterances after  $m$  from the original dialogue, as illustrated in Figure 4. Lastly, this newly generated data is used to fine-tune the classification model. The resulting model can be applied online to dynamically improve the user experience at run-time or used as a data selection mechanism that identifies high-quality dialogues from low-quality dialogues.

### 3.3.4 Transfer Learning with CF-LSTM

As described above, the CF-LSTM is designed to predict the rating based on a fixed number of hypotheses (e.g., 2) in the dataset. However, CF-LSTM can handle the increase or decrease in the number of hypotheses using transfer learning [23, 25, 22, 24, 26]. In other words, we can transfer the prior knowledge from the pre-trained CF-LSTM to a new causal NLP task using transfer learning techniques [23, 22]. For example, the new data samples arrive with added treatment types. Here, we modify the trained model by adding new MLP regressors (e.g., representing different hypotheses) to the existing structure. We initialize the weights randomly for the added regressors while freezing the other weights of the pre-trained

model. Lastly, we update the entire model by fine-tuning it using the new data samples. In the next section, we discuss the experimental results of our proposed methods when applied to our real-world dataset for both regression and classification problems.

## 4 Experimental Results

In the following experiments, we apply our proposed methods to a dataset of dialogues with user ratings collected from users interacting with the Alexa SocialBot [36]. This dataset consists of 128K de-identified conversational dialogues. On average, each dialogue has 20 utterances, and the data is limited to dialogue consisting of at least three turn-pairs. For each utterance in a dialogue, we assign treatments based on the ODES categories (Please see Section 3.3.1). Table 1 indicates the number of utterances of each ODES category in the Alexa SocialBot dataset. When considering this dataset, 70% of dialogues are assigned  $T = 0$ , and 30% of dialogues are assigned  $T = 1$ . Next, we compare our methods to the baseline approach, the End-to-end (E2E) Transformer. It directly uses a transformer-based model [5] to map the dialogue textual data into the user ratings.

First, we consider the regression problem of predicting the user ratings for the given dialogue. Let L1d denote the daily average prediction and L7d denote the 7-day rolling daily average prediction. Table 2 shows the three methods’ performances when compared with the E2E baseline in terms of the Pearson correlation between the predicted ratings and the ground truth. As shown in Table 2, CF-LSTM achieves higher correlations in the individual prediction, the daily average prediction, and the 7-day rolling daily average prediction. It also shows a performance boost compared to the Dialogue-level MLP approach, which suffers from information loss when converting turn-level to dialogue-level features.

**Table 2** The comparisons between open-dialogue evaluation methods for the regression problem in terms of Pearson correlation (e.g., individual, L1d, L7d predictions) and the classification problems (e.g., binary, 5-class) in terms of prediction accuracy, on SocialBot conversations

METHODS	INDIVIDUAL PREDICTION	L1D* PREDICTION	L7D† PREDICTION	BINARY CLASS.	5-CLASS CLASS.
E2E TRANSFORMER	0.22	0.30	0.47	54.1%	32.8%
DIALOGUE-LEVEL LSTM	0.30	0.41	0.59	64.6%	43.5%
DIALOGUE-LEVEL MLP	0.31	0.40	0.66	62.5%	46.1%
CF-LSTM	<b>0.34</b>	<b>0.46</b>	<b>0.68</b>	<b>67.8%</b>	<b>48.2%</b>

\* DAILY AVERAGE PREDICTION, † 7-DAY ROLLING AVERAGE PREDICTION

In addition, to understand the impact of assigning the treatment (i.e., ODES categories) to the dialogues, we consider the situation where the treatments of the col-

**Table 3** The correlations of CF-LSTM when the treatment assignments for all dialogues are inverted.

PEARSON CORRELATION	ORIGINAL TREATMENTS	INVERTED TREATMENTS
INDIVIDUAL PREDICTION	0.34	0.26
L1D PREDICTION	0.46	0.35
L7D PREDICTION	0.68	0.52

\* DAILY AVERAGE, † 7-DAY ROLLING AVERAGE

lected dialogues are inverted (e.g., treatment  $T = 0$  becomes  $T = 1$ , and vice versa). For the same dialogue feature vector, its treatment is inverted. This setting allows us to investigate the impact of the ODES Classifier when assigning the incorrect treatments and whether the models can overcome this misinformation. We feed the modified data samples to the well-trained CF-LSTM capable of predicting the counterfactual outcomes for these dialogues. Table 3 shows the correlations between the predicted counterfactual outcomes and the original ground truth. Here, we want to investigate whether the model can predict correctly even if the treatment assignment is inverted (e.g., the ODES classifier malfunctions). As shown in Table 3, the individual prediction and daily average prediction (L1d) of the inverted treatments scenario still have a relatively high correlation with the original ground truth (i.e., when the treatments are not inverted). The mean square error between the factual outcomes  $Y_1$  and the counterfactual outcomes  $Y_0$  is 0.3034. The average treatment effect (ATE) is defined as follows for our problem:

$$ATE = E[Y_{1i} - Y_{0i}] = -0.7809$$

The ATE shows that, on average, a dialogue being assigned treatment  $T = 1$  (e.g., consists of disinterest, insult) will have a lower user rating than the dialogue with  $T = 0$  by 0.7809.

We randomly select a set of 1213 dialogues and ask an expert to rate the conversational quality. Let  $R_{expert}$  be the rating given by the expert, and  $R_{user}$  be the ratings given by users. The expert ratings have a low Pearson correlation with the user ratings, demonstrating the challenge in predicting user ratings in open-domain dialogue since they are only loosely correlated with an expert evaluation of dialogue quality.

$$\rho = corr(R_{expert}, R_{user}) = 0.2056$$

Next, we use the CF-LSTM and the Dialogue-level MLP to predict the expert’s rating in the given dialogues. Table 4 indicates the correlations obtained from CF-LSTM and Dialogue-level MLP on expert-rated dialogues from Socialbot conversations. Due to the simplicity of the network architecture, the Dialogue-level MLP achieves competitive results in individual and daily average predictions. However, it performs poorly in 7-day rolling average predictions because of the sparsity across time of the small dataset of 1213 dialogues. On the other hand, CF-LSTM performs well on all types of predictions and shows flexibility in training for both large and small datasets.

**Table 4** The correlations of prediction models on 1213 expert-rated dialogues from Socialbot conversations.

PEARSON CORRELATION	DIALOGUE-LEVEL MLP	CF-LSTM (OURS)
INDIVIDUAL PREDICTION	0.35	0.26
L1D PREDICTION	0.42	0.36
L7D PREDICTION	0.12	0.51

\* DAILY AVERAGE, † 7-DAY ROLLING AVERAGE

Lastly, we consider the problem of classifying the dialogues into 2-class (i.e., binary) and 5-class classification. For binary classification, we label any dialogue with a rating less than 3 as 0 (e.g., low-rated dialogues). Hence, the remaining dialogues with ratings greater or equal to 3 are labeled as 1 (i.e., highly-rated dialogues). For the 5-class classification, we round up or down (i.e., round half-up) the ratings into five groups. For example, a dialogue with a rating of 4.5 is classified into group 5, and a dialogue with a rating of 3.4 is classified into group 3. In this dataset, we observe that most dialogues are often rated above 3 out of 5. As a result, we apply the data generation method described in Section 3.3.3 to generate more poorly-rated dialogues. This newly generated data is used to fine-tune the binary classification model. The resulting model can be applied online to improve the user experience at run-time. Table 2 indicates the performance of our model and other methods, in terms of classification accuracy, in binary and 5-class classification problems. Our proposed method CF-LSTM outperforms other approaches in both binary classification and 5-class classification.

## 5 Conclusion and Future Work

We propose a novel causal inference structure for open-domain dialogue evaluation. This structure utilizes turn-level features, such as sentiment analysis, textual categories, response relevance, and response specificity, and uses them to predict user ratings. Our method performs competitively in an evaluation with rated conversations with the Alexa SocialBot, in regression and classification problems. As a causal inference model, this CF-LSTM model is more robust in learning complex representations and capable of predicting the ratings for the dialogues under different hypotheses. In addition to offline uses, this model can also be applied at run-time to dynamically identify low-quality dialogues and trigger the introduction of different topics to help improve the user experience. This paper considers only 2 treatments, which are all the ODES categories, versus others. In future works, we can further boost the flexibility of our model by increasing the number of treatments corresponding to the different ODES groups. Additionally, other essential dialogue features (e.g., text and sentiment embeddings) can be included in the input features to help improve the overall performance of the prediction model.

## References

1. A. M. Alaa and M. Van Der Schaar. Bayesian inference of individualized treatment effects using multi-task gaussian processes. *Advances in neural information processing systems*, 30, 2017.
2. A. Aloui, J. Dong, C. P. Le, and V. Tarokh. Causal knowledge transfer from task affinity. *arXiv preprint arXiv:2210.00380*, 2022.
3. S. Banerjee and A. Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005.
4. F. Charras, G. D. Duplessis, V. Letard, A.-L. Ligozat, and S. Rosset. Comparing system-response retrieval models for open-domain and casual conversational agent. In *Second Workshop on Chatbots and Conversational Agent Technologies (WOCHAT@ IVA2016)*, 2016.
5. J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proc. of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, 2019.
6. J. Dodge, A. Gane, X. Zhang, A. Bordes, S. Chopra, A. Miller, A. Szlam, and J. Weston. Evaluating prerequisite qualities for learning end-to-end dialog systems. *arXiv preprint arXiv:1511.06931*, 2015.
7. G. D. Duplessis, V. Letard, A.-L. Ligozat, and S. Rosset. Purely corpus-based automatic conversation authoring. In *10th edition of the Language Resources and Evaluation Conference (LREC)*, pages ISBN–978, 2016.
8. A. Feder, K. A. Keith, E. Manzoor, R. Pryzant, D. Sridhar, Z. Wood-Doughty, J. Eisenstein, J. Grimmer, R. Reichart, M. E. Roberts, et al. Causal inference in natural language processing: Estimation, prediction, interpretation and beyond. *arXiv preprint arXiv:2109.00725*, 2021.
9. R. Gabriel, Y. Liu, A. Gottardi, M. Eric, A. Khatri, A. Chadha, Q. Chen, B. Hedayatnia, P. Rajan, A. Binici, et al. Further advances in open domain dialog systems in the third alexa prize socialbot grand challenge. *Alexa Prize Proceedings*, 2020.
10. X. Gao, Y. Zhang, M. Galley, C. Brockett, and B. Dolan. Dialogue response ranking training with large-scale human feedback data. In *Proc. of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 386–395, Online, Nov. 2020. Association for Computational Linguistics.
11. A. Gatt and E. Kraemer. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*, 61:65–170, 2018.
12. S. Ghazarian, B. Hedayatnia, A. Papangelis, Y. Liu, and D. Hakkani-Tur. What is wrong with you?: Leveraging user sentiment for automatic dialog evaluation. *arXiv preprint arXiv:2203.13927*, 2022.
13. S. Ghazarian, J. Wei, A. Galstyan, and N. Peng. Better automatic evaluation of open-domain dialogue systems with contextualized embeddings. In *Proc. of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*, pages 82–89, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
14. S. Ghazarian, R. Weischedel, A. Galstyan, and N. Peng. Predictive engagement: An efficient metric for automatic evaluation of open-domain dialogue systems. In *Proc. of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7789–7796, 2020.
15. K. Gopalakrishnan, B. Hedayatnia, Q. Chen, A. Gottardi, S. Kwatra, A. Venkatesh, R. Gabriel, D. Hakkani-Tür, and A. A. AI. Topical-chat: Towards knowledge-grounded open-domain conversations. *Proc. Interspeech 2019*, pages 1891–1895, 2019.
16. R. Higashinaka, M. Araki, H. Tsukahara, and M. Mizukami. Improving taxonomy of errors in chat-oriented dialogue systems. In *9th International Workshop on Spoken Dialogue System Technology*, page 331. Springer, 2019.
17. S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, nov 1997.

18. C. Hori and T. Hori. End-to-end conversation modeling track in dstc6. *arXiv preprint arXiv:1706.07440*, 2017.
19. L. Huang, Z. Ye, J. Qin, L. Lin, and X. Liang. GRADE: Automatic graph-enhanced coherence metric for evaluating open-domain dialogue systems. In *Proc. of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9230–9240, Online, Nov. 2020. Association for Computational Linguistics.
20. J. Juraska, K. K. Bowden, and M. Walker. ViGGO: A video game corpus for data-to-text generation in open-domain conversation. In *Proceedings of the 12th International Conference on Natural Language Generation*, 2019.
21. Y. Kim, J. Levy, and Y. Liu. Speech sentiment and customer satisfaction estimation in social-bot conversations. *arXiv preprint arXiv:2008.12376*, 2020.
22. C. P. Le, J. Dong, M. Soltani, and V. Tarokh. Task affinity with maximum bipartite matching in few-shot learning. In *International Conference on Learning Representations*, 2021.
23. C. P. Le, M. Soltani, J. Dong, and V. Tarokh. Fisher task distance and its application in neural architecture search. *IEEE Access*, 10:47235–47249, 2022.
24. C. P. Le, M. Soltani, R. Ravier, and V. Tarokh. Improved automated machine learning from transfer learning. *arXiv e-prints*, pages arXiv–2103, 2021.
25. C. P. Le, M. Soltani, R. Ravier, and V. Tarokh. Task-aware neural architecture search. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4090–4094. IEEE, 2021.
26. C. P. Le, Y. Zhou, J. Ding, and V. Tarokh. Supervised encoding for discrete representation learning. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3447–3451, 2020.
27. Y. Li, H. Su, X. Shen, W. Li, Z. Cao, and S. Niu. Dailydialog: A manually labelled multi-turn dialogue dataset. *arXiv preprint arXiv:1710.03957*, 2017.
28. C.-Y. Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.
29. C.-W. Liu, R. Lowe, I. V. Serban, M. Noseworthy, L. Charlin, and J. Pineau. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, 2016.
30. S. Mehri and M. Eskenazi. Unsupervised evaluation of interactive dialog with dialogpt. In *Proc. of SIGDIAL*, 2020.
31. S. Mehri and M. Eskenazi. Unsupervised evaluation of interactive dialog with dialogpt. In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 225–235, 2020.
32. K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proc. of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002.
33. M.-C. Popescu, V. E. Balas, L. Perescu-Popescu, and N. Mastorakis. Multilayer perceptron and neural networks. *WSEAS Trans. Cir. and Sys.*, 8(7):579–588, jul 2009.
34. R. Prasad and M. Walker. Training a dialogue act tagger for human-human and human-computer travel dialogues. In *Proceedings of the Third SIGdial Workshop on Discourse and Dialogue*, pages 162–173, 2002.
35. C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67, 2020.
36. A. Ram, R. Prasad, C. Khatri, A. Venkatesh, R. Gabriel, Q. Liu, J. Nunn, B. Hedayatnia, M. Cheng, A. Nagar, et al. Conversational AI: the science behind the alexa prize. *arXiv preprint arXiv:1801.03604*, 2018.
37. H. Rashkin, E. M. Smith, M. Li, and Y.-L. Boureau. Towards empathetic open-domain conversation models: A new benchmark and dataset. *arXiv preprint arXiv:1811.00207*, 2018.
38. E. Reiter and R. Dale. Building applied natural language generation systems. *Natural Language Engineering*, 3(1):57–87, 1997.

39. P. R. Rosenbaum and D. B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
40. D. B. Rubin. Bayesian inference for causal effects: The role of randomization. *The Annals of statistics*, pages 34–58, 1978.
41. T. Scialom, P.-A. Dray, S. Lamprier, B. Piwowarski, J. Staiano, A. Wang, and P. Gallinari. QuestEval: Summarization asks for fact-based evaluation. In *Proc. of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6594–6604, Online and Punta Cana, Dominican Republic, Nov. 2021. Association for Computational Linguistics.
42. A. See, S. Roller, D. Kiela, and J. Weston. What makes a good conversation? how controllable attributes affect human judgments. *arXiv preprint arXiv:1902.08654*, 2019.
43. T. Sellam, D. Das, and A. Parikh. BLEURT: Learning robust metrics for text generation. In *Proc. of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online, July 2020. Association for Computational Linguistics.
44. I. Serban, A. Sordoni, Y. Bengio, A. Courville, and J. Pineau. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proc. of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
45. U. Shalit, F. D. Johansson, and D. Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning*, pages 3076–3085. PMLR, 2017.
46. I. Shalyminov, O. Dušek, and O. Lemon. Neural response ranking for social conversation: A data-efficient approach. In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*, pages 1–8, 2018.
47. K. Sinha, P. Parthasarathi, J. Wang, R. Lowe, W. L. Hamilton, and J. Pineau. Learning an unreferenced metric for online dialogue evaluation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2430–2441, Online, July 2020. Association for Computational Linguistics.
48. C. Villani. *The Wasserstein distances*, pages 93–111. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
49. M. Walker, C. Harmon, J. Graupera, D. Harrison, and S. Whittaker. Modeling performance in open-domain dialogue with paradise. *International Workshop on Spoken Dialogue Systems, IWSDS*, 2021.
50. M. Walker and R. J. Passonneau. Date: A dialogue act tagging scheme for evaluation of spoken dialogue systems. In *Proceedings of the First International Conference on Human Language Technology Research*, 2001.
51. M. A. Walker, D. J. Litman, C. A. Kamm, and A. Abella. PARADISE: A framework for evaluating spoken dialogue agents. In *35th Annual Meeting of the Association for Computational Linguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics*, pages 271–280, Madrid, Spain, July 1997. Association for Computational Linguistics.
52. J. Yoon, J. Jordon, and M. van der Schaar. GANITE: Estimation of individualized treatment effects using generative adversarial nets. In *International Conference on Learning Representations*, 2018.
53. C. Zhang, Y. Chen, L. F. D’Haro, Y. Zhang, T. Friedrichs, G. Lee, and H. Li. Dynaeval: Unifying turn and dialogue level evaluation. *arXiv preprint arXiv:2106.01112*, 2021.
54. S. Zhang, E. Dinan, J. Urbanek, A. Szlam, D. Kiela, and J. Weston. Personalizing dialogue agents: I have a dog, do you have pets too? *arXiv preprint arXiv:1801.07243*, 2018.
55. T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019.
56. Y. Zhang, S. Sun, M. Galley, Y.-C. Chen, C. Brockett, X. Gao, J. Gao, J. Liu, and B. Dolan. Dialogpt: Large-scale generative pre-training for conversational response generation. *arXiv preprint arXiv:1911.00536*, 2019.
57. Y. Zhang, S. Sun, M. Galley, Y.-C. Chen, C. Brockett, X. Gao, J. Gao, J. Liu, and B. Dolan. Dialogpt: Large-scale generative pre-training for conversational response generation, 2020.