

# Lexical Entrainment for Conversational Systems

Zhengxiang Shi<sup>†</sup> Procheta Sen<sup>‡</sup> Aldo Lipani<sup>†</sup>

<sup>†</sup> University College London, United Kingdom

<sup>‡</sup> University of Liverpool, United Kingdom

{zhengxiang.shi.19,aldo.lipani}@ucl.ac.uk

procheta.sen@liverpool.ac.uk

## Abstract

Conversational agents have become ubiquitous in assisting with daily tasks, and are expected to possess human-like features. One such feature is *lexical entrainment* (LE), a phenomenon in which speakers in human-human conversations tend to naturally and subconsciously align their lexical choices with those of their interlocutors, leading to more successful and engaging conversations. As an example, if a digital assistant replies “Your appointment for Jinling Noodle Pub is at 7 pm” to the question “When is my reservation for Jinling Noodle Bar today?”, it may feel as though the assistant is trying to correct the speaker, whereas a response of “Your reservation for Jinling Noodle Bar is at 7 pm” would likely be perceived as more positive. This highlights the importance of LE in establishing a shared terminology for maximum clarity and reducing ambiguity in conversations. However, we demonstrate in this work that current response generation models do not adequately address this crucial human-like phenomenon. To address this, we propose a new dataset, named MULTIWOZ-ENTR, and a measure for LE for conversational systems. Additionally, we suggest a way to explicitly integrate LE into conversational systems with two new tasks, a LE extraction task and a LE generation task. We also present two baseline approaches for the LE extraction task, which aim to detect LE expressions from dialogue contexts.<sup>1</sup>

## 1 Introduction

There is potential for enormous variability in people’s lexical choices in dialogues, whether they are communicating with human or machine partners (Brennan, 1996b). One question that has attracted some attention from the linguistics community has been: *How do people deal with such high lexical variability when they talk to each other?* In

<sup>1</sup>The dataset is available at <https://github.com/ZhengxiangShi/LexicalEntrainment>.

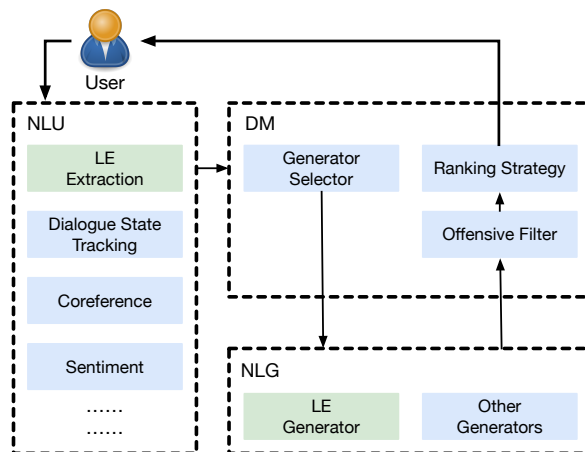


Figure 1: Illustration of a (simplified) pipeline architecture of conversational systems, which decomposes into NLU, DM, and NLG modules. LE expressions extracted by the NLU module can be leveraged by the NLG module to generate more engaging and human-like responses.

the referential communication studies (Brennan, 1996a,b), the likelihood that people in one conversation would choose the same terms to refer to the same common object as people in another conversation was only 10%. Despite the high variability across conversations, it is relatively low within a conversation (Garrod and Anderson, 1987; Brennan, 1996b). When people engage in a conversation, they naturally adapt their way of speaking to align with their conversational partner. For instance, they tend to refer to something based on how their conversational partners refer to it, using the same terms when discussing the same object repeatedly and negotiating a common description, particularly for items that may be unfamiliar to them (Brennan, 1996a). This linguistic phenomenon is known as *lexical entrainment* (LE) (Garrod and Anderson, 1987; Brennan, 1996b).

LE plays a key role in the success and naturalness of interactions in conversations. Reitter and Moore (2007) demonstrated that lexical repetition

is a reliable predictor of task success given the first five minutes in task-oriented dialogues. The degree of entrainment with respect to the most frequent words is also a distinguishing factor between dialogues rated as most natural and those rated as less natural, and is strongly correlated with task success and engaged and coordinated turn-taking behaviour (Nenkova et al., 2008). Additionally, LE is associated with a broad range of positive social behaviours and outcomes (Lopes et al., 2015; Nasir et al., 2019), such as building effective dialogues (Porzel et al., 2006) and engaging students in tutorial dialogues (Ward and Litman, 2007). In conversational systems, an explicit mechanism of lexical choice (e.g., by LE) is needed to ensure meaningful results (Pickering and Garrod, 2004), thus making conversations more engaging and successful.

Despite its importance, LE has not been systematically studied for conversational systems, and this interesting fact of language-based communication has been overlooked. There is a scarcity of datasets specifically designed to study LE, and it is not explicitly modelled in any state-of-the-art conversational systems. The only existing dataset for lexical entrainment, presented by Dušek and Jurcicek (2016), consists of 1 859 pairs of user utterances, system responses, and dialogue acts. However, it is limited to the context of a single preceding user utterance and constrained to the public transportation domain. Most importantly, this dataset lacks explicit annotations of LE, making it difficult to be used to explicitly model LE in conversational systems.

To this end, we propose a new dataset named MULTIWOZ-ENTR with annotated LE information, comprehensive dialogue context, and detailed statistical analysis. The dataset is built on the top of the MULTIWOZ (Eric et al., 2020), a task-oriented dialogue dataset with 7 distinct domains. MULTIWOZ is selected as a starting point for two reasons: (1) it is a task-oriented dataset, which tends to contain more LE expressions than other kinds of dialogue datasets (Reitter et al., 2006), such as open-domain datasets (e.g., Li et al., 2017; Wu et al., 2017); and (2) it is collected based on human-to-human conversations.

Additionally, we present a workflow of a conversational system integrated with LE modules, as illustrated in Figure 1. A standard pipeline architecture for conversational systems generally decomposes into three modules, a Natural Language

Understanding (NLU) module, a Dialogue Management (DM) module, and a Natural Language Generation (NLG) module. Each of these modules consists of several sub-modules that can be modelled independently. The integration of LE could be accomplished through two additional sub-modules (tasks), a LE extraction sub-module and a LE generator sub-module. The objective of the LE extraction sub-module is to identify lexical candidates and extract LE expressions given the dialogue context. The objective of the generator module is to make use of these extracted lexical candidates to generate more human-like responses. Additionally, we provide two baseline models for the LE extraction sub-module.

In summary, the main contributions of this paper are as follows:

- We formalize a precise and practical definition of LE (§2.1), as well as a new LE measure to evaluate the natural degree of LE in human-to-human conversations (§2.2);
- We highlight the importance of LE and provide an analysis of state-of-the-art response generation models, pointing out issues caused by the neglect of LE (§3);
- We propose a LE dataset, MULTIWOZ-ENTR, specifically designed for studying LE, and provide detailed annotations (§4) and statistical analysis (appendix B);
- We present a methodology for integrating LE into conversational systems through two novel sub-modules (tasks). Specifically, we provide two baseline models for the LE extraction sub-module (§5), providing valuable insights into the incorporation of LE into conversational systems. This approach lays the groundwork for the development of a LE generator sub-module, which is left for future research (§6).

## 2 Lexical entrainment (LE)

### 2.1 Definition of LE

According to the work of Brennan (1996b), LE is defined as “[w]hen two people repeatedly discuss the same object, they come to use the same terms.” Brennan also provides several examples of entrained expressions, which are noun phrases and also referred to as *referring expressions*. However, this definition and the examples provided are not

Table 1: An example from MULTIWOZ-ENTR. The sequence of words with the same colour represents instances of the same LE expressions. The sequence of words with the underline stands for the LE expressions entrained in that utterance, meaning that this LE expression has already been established by the current utterance.  $E_{s,j}$  indicates the number of instances of established LE expressions in the  $j^{\text{th}}$  turn from the speaker  $s$ .

No.	Speaker	Utterance	$E_{s,j}$
1	User	Im looking for an <u>italian restaurant</u> . I also need it to be near the <u>center of town</u> .	0
2	Agent	There are 9 <u>italian restaurants</u> in the <u>center of town</u> . Do you have a particular <u>price range</u> in mind?	2
3	User	I need a cheap <u>restaurant</u> .	1
4	Agent	There is Pizza Hut City Centre, <u>Ask</u> , and Zizzi Cambridge, which would you like to book?	0
5	User	Can you book a <u>table</u> for 7 at <u>Ask</u> ? We're looking at 14:00 on Saturday.	1
6	Agent	Sorry, there are no <u>tables</u> available at <u>Ask</u> at that time. Would you like to book another day or time, or perhaps another <u>restaurant</u> ?	3
7	User	Can you try 13:00, then?	0
8	Agent	13:00 it is. Your <u>table</u> at <u>Ask</u> will be held for 15 minutes, and your <u>reference number</u> is L2KZYL26. Can I help you with anything else?	2
9	User	I'm also looking for a <u>hotel</u> .	0
10	Agent	We have 33 different <u>hotels</u> . Did you have a <u>price range</u> or star rating in mind?	1
11	User	Something in the cheap <u>price range</u> including free parking.	1
12	Agent	What area would you prefer to stay in? This will help narrow down your options.	0
13	User	I would like to book a <u>hotel</u> close to our <u>restaurant</u> here in the <u>centre</u> , it will also need to be cheap, and have free parking.	2
14	Agent	There are two <u>hotels</u> available. Would you like me to book one for you?	1
15	User	Yes please. The our group of 7 will all be staying there for 3 nights starting <u>Saturday</u> .	0
16	Agent	I've booked you a 3 day stay at the Alexander Bed and Breakfast for 7 people. Do you need anything else?	0
17	User	Could I have the <u>reference number</u> for that reservation please?	1
18	Agent	Yes, your <u>reference number</u> is IJBO1R08.	1
19	User	Thank you so much for all of your help. That is all I needed for now. Have a nice <u>day</u> .	0
20	Agent	Thank you for choosing us to assist you on your travels. Have a good <u>day</u> .	1

specific, leaving room for multiple interpretations. This has led to several works in the field of computer science with conflicting interpretations. For instance, [Nenkova et al. \(2008\)](#); [Dušek and Jurcicek \(2016\)](#); [Hu et al. \(2016\)](#); [Dubuisson Duplessis et al. \(2017\)](#); [Duplessis et al. \(2021\)](#) have expanded upon [Brennan's](#) definition of LE to include other parts of speech beyond noun phrases, such as modal particles (e.g., “well”) and verb phrases (e.g., “work for”). This work aims to eliminate ambiguity and formalize a precise and operational definition of LE, which closely aligns with the original definition proposed by [Brennan \(1996b\)](#). Accordingly, the following definition is proposed:

**Definition 1 (Lexical Entrainment).** *Lexical entrainment refers to the natural phenomenon observed in conversations involving two or more speakers, in which equivalent referring expressions are utilised to discuss the same object. A referring expression is a noun phrase used to denote an object, which can refer to both concrete and abstract entities, and may be singular, plural, or collective.*

In the original work by [Brennan \(1996a\)](#), it was suggested that two referring expressions can be considered equivalent only if they share all the same content words. For instance, expressions such as “the red dog with its mouth open” and “the dog

with its mouth open” would not be regarded as equivalent. However, we find this definition to be excessively stringent, and thus, we propose a more relaxed criterion that allows for different modifiers in the referring expressions. As an example, consider the dialogue in Table 1, where “the reference number for that reservation” in the 17<sup>th</sup> utterance and “your reference number” in the 18<sup>th</sup> utterance would not satisfy [Brennan's](#) criterion. Despite this, the agent’s choice to use the expression “reference number” implies a lexical decision to entrain the user, given that the agent could have used alternative expressions like “confirmation number”. Therefore, we consider the two expressions to be equivalent after removing the modifier “for that reservation”. Another example of equivalent referring expressions is “red taxi” and “blue taxi”, which can be considered equivalent after disregarding the adjectives “red” and “blue”, given that the agent could have used alternative expressions like “blue cab”. The concept of equivalent referring expressions is defined as follows:

**Definition 2 (Equivalent Referring Expressions).** *Two referring expressions are deemed equivalent if, without their modifiers, their head nouns are the same, regardless of their noun forms.*

In accordance with the terminology established

by Dubuisson Duplessis et al. (2017), referring expressions that are equivalent in a conversation are referred to as a **LE expression** once they have been **established**. Each referring expression is then considered an instance of this LE expression. For example, “table” and “tables” in Table 1 belong to the same LE expression, and “table” and “tables” themselves are considered two instances of this LE expression. **An instance of LE expression** can either be free or constrained given a dialogue. A **free instance** is an instance of a LE expression that appears without being a subexpression of a larger LE expression, while a **constrained instance** is an instance of a LE expression that appears as a subexpression of a larger LE expression. For instance, in Table 1, “italian restaurant” in the 1<sup>st</sup> and 2<sup>nd</sup> utterances, as well as “restaurant” in the 3<sup>rd</sup> utterance, are all free instances, while “restaurant” in the 1<sup>st</sup> and 2<sup>nd</sup> utterances are constrained instances since they are sub-expressions of “italian restaurant”. Equivalent referring expressions are **established** if and only if the two following requirements are met: (i) they have been produced by at least two speakers, either in free or constrained form, and (ii) at least one of the instances is in free form. For example, “italian restaurant” is established in the 2<sup>nd</sup> utterance and “price range” is established in the 11<sup>th</sup> utterance. Finally, the **initiator of a LE expression** refers to the speaker who first produces an instance of the LE expression, either in a free or constrained form.

## 2.2 The proposed measure of LE

We propose a new measure, referred to as the degree of LE, to evaluate the frequency of LE expressions in dialogues after an agreement has been reached (i.e., LE expressions are established). Given a dialogue with  $n$  utterances from the speaker  $s$ , this measure is computed as follows:

$$\text{ENTR}_s = \frac{1}{n} \sum_{j=1}^n E_{s,j}, \quad (1)$$

where  $E_{s,j}$  is the number of instances of an LE expression in the  $j^{\text{th}}$  utterance from the speaker  $s$  of established LE expression.

It is worth noting that a higher value of  $\text{ENTR}_s$  does not necessarily imply higher quality of the conversation, as conversational systems should aim to find a balance between self-consistency and LE rather than simply maximising the degree of LE. Nevertheless, the degree of LE can be used to un-

derstand the natural frequency of LE in human-to-human conversations and what might be pertinent in the design of machines.

## 3 Limitations of current approaches

Current conversational systems do not sufficiently address the phenomenon of LE, leading to inconsistencies between generated responses and actual human responses. This can result in sub-optimal generated responses, which are not captured by current evaluation metrics, such as INFORM RATE, SUCCESS RATE, and BLEU SCORE. To demonstrate the discrepancies between generated responses and ground-truth responses, we evaluate four state-of-the-art response generation models (Nekvinda and Dušek, 2021) in both end-to-end fashions, that is, using only the dialogue context as input to generate responses, and policy optimization fashion, that is using the ground-truth dialogue states to generate responses. For end-to-end models, MINTL (Lin et al., 2020) and AUGPT (Kulhánek et al., 2021) are evaluated. For policy optimization models, HDSA (Chen et al., 2019) and MARCO (Wang et al., 2020) are evaluated. In the following, we present two examples of generated responses that demonstrate the consequences of neglecting LE.

In the first case study (see Table 2), the user requests a “confirmation number”, and the ground-truth (human-like) agent response entrains this expression by repeating “confirmation number”. However, all four generated responses utilise “reference number” instead. This is due to the fact that “reference number” appears more frequently in the training data or is used in their generation templates. In this scenario, using a “confirmation number” would be more appropriate and engaging for the user. In the second case study, the user requests a “cab”, which is also entrained by the ground-truth agent response. However, MARCO replies to the user with the “taxi”. Many similar situations exist where generation models struggle to choose synonyms, such as “phone number,” “contact number,” and “telephone number,” as well as “good reviews” and “good ratings.”

The above-mentioned issues stem from neglecting the phenomenon of LE, which occurs naturally and unconsciously in human-to-human conversations. We consolidate our observation with a quantifiable analysis in §6.1.



Table 2: Case study for LE in existing conversational systems. The sequence of words with the underline stands for the LE expressions entrained in that utterance, meaning that this LE expression has already been established by the current utterance. In case study 1, when the user asks for a “confirmation number”, response generation models consistently provide a response using the term “reference number” instead. In case study 2, the user requests the agent to arrange a “cab” but the response generation model provides a response using the term “taxi”.

Model	Utterance (Case Study 1)
USER	Yes, please book it for 2 people on Saturday for 4 nights. Will you give me the <u>confirmation number</u> when that is complete?
AGENT	Sure, booking was successful. Your <u>confirmation number</u> is 7MG032V9.
AUGPT	Booking was successful. <u>Reference number</u> is: 00000016.
MINTL	I have booked you at the Express by Holiday Inn Cambridge for 4 nights starting Saturday. Your <u>reference number</u> is D77JCGJ2. Is there anything else I can help?
HDSA	Booking was successful. <u>Reference number</u> is 7MG032V9. Is there anything else I can help you with?
MARCO	I have booked you at the [hotel_name]. Your <u>reference number</u> is 7MG032V9. Is there anything else I can help you with?
Model	Utterance (Case Study 2)
USER	I need a <u>cab</u> too please, arriving by my <u>booked time</u> , car type and contact number would be helpful too please.
AGENT	You need a <u>cab</u> to get you to the restaurant before the <u>booked time</u> and where do you need <u>cab</u> to pick you up at?
AUGPT	I have booked a blue Skoda for you. The contact number is 85534112262.
MINTL	I have booked you a BMW. The contact number is 07338019809. Where will you be departing from?
HDSA	Where will you be departing from and where are you going?
MARCO	I would be happy to book a <u>taxi</u> for you. where are you departing from and what is your destination?

## 4 MULTIWOZ-ENTR

To facilitate LE research for conversational systems, we introduce a new dataset named MULTIWOZ-ENTR based on the MULTIWOZ 2.1 (Eric et al., 2020).

### 4.1 MULTIWOZ.

MULTIWOZ 2.1 contains 8 438, 1 000, and 1 000 samples for training, validation, and test sets respectively. It spans over seven domains, including Taxi, Hotel, Attraction, Train, Restaurant, Police, and Hospital. The Police and Hospital domains are not included in the validation and testing sets, but the training set contains 245 and 287 samples for these domains, respectively. Detailed statistics for the other five domains are summarised in Table 3.

Table 3: Statistics of MultiWOZ 2.1: the number of dialogues in different domains.

Domain	Taxi	Hotel	Attraction	Train	Restaurant
Train	1 655	3 387	2 718	3 117	3 817
Valid	207	416	401	484	438
Test	195	394	396	495	437

### 4.2 MULTIWOZ-ENTR.

We use the following steps to generate annotations of LE expressions in MULTIWOZ. Please refer to appendix B for a detailed statistical analysis of the proposed dataset.

**Step 1: Preprocessing.** The goal of this step is to ensure that all the LE expressions can be easily captured from the given dialogues. In previous works (e.g., Dubuisson Duplessis et al., 2017; Duplessis et al., 2021), two text sequences are regarded equal only if they are exactly the same. However, even if they contain LE expressions, two sentences may have minor grammatical differences, such as different forms (e.g., “table” and “tables”), British and American English spelling (e.g., “centre” and “center”), and numerical characters (e.g., “4 star hotel” and “four star hotel”). To tackle these issues, we first convert all British English to American English using an open-source toolkit, thereby converting “centre” to “center”, for example. We also standardize numerical characters. We then use a stemming algorithm with the NLTK toolkit (Loper and Bird, 2002) for each token in a dialogue. For example, “italian restaurant” and “italian restaurants” in Table 1 will both be converted to “italian restaur”, which ensures that they can be recognized in the next step. Meanwhile, we replace all punctuation (except in-word punctuation such as “pre-trained”) with pseudo-random numbers with  $k$  random bits to prevent them from being identified as parts of LE expressions in the subsequent step. In our work, LE expressions are punctuation-free and not case-sensitive, which are factors that were not considered in previous works (e.g., Dubuisson Duplessis et al., 2017; Duplessis et al., 2021).

**Step 2: Maximize recall.** The target of this step is to enhance the recall rate while extracting potential candidates of LE expressions from the given dialogues. To achieve this, we utilise an open-source toolkit<sup>2</sup> (Duplessis et al., 2021), which is specifically designed for LE in dyadic conversations. The toolkit retrieves all sequences of tokens that share the same surface text form as that used by the two interlocutors. Nonetheless, this approach may introduce noise by including irrelevant phrases such as “in the”, “for 7”, and “like to book”.

**Step 3: Maximize precision.** The purpose of this step is to eliminate noise from the pool of candidates for LE expressions. We post-process the candidates of LE expressions by removing stopwords (Loper and Bird, 2002) from the beginning or at the end of LE expression candidates. For example, “the postcode” will be converted to “postcode” and “in the center” will be converted to “center”. However, stopwords are allowed to appear in the middle of LE expression candidates, for instance, “cheap hotel with free park” will not be post-processed since “with” is a stopword in the middle of this expression. Additionally, we only consider noun phrases as per the definition of LE, introduced in §2.1. To identify noun phrases, we manually create dictionaries for verbs and adjectives/adverbs in the MULTIWOZ dataset, along with a dictionary for undesired phrases such as particles, modal verbs, and prepositions. This helps filter out negative examples from the candidate pool. However, some special cases require rule-based methods. For example, some words such as “help” and “booking” can function as both nouns and verbs, making it difficult to accurately distinguish them. In these cases, we examine the special cases and apply specific rules. Since MULTIWOZ is a task-oriented dataset with a limited lexicon size, our rule-based methods are feasible and appropriate.

**Step 4: Human validation** This step aims to ensure the quality of the MULTIWOZ-ENTR dataset by conducting a human validation process. The authors of the dataset serve as validators to assess the annotations made in the previous steps. At least two human validators review every annotated LE expression, and any disagreements among validators are resolved through discussion. To perform this validation, a batch of 50 dialogues from all 10 438 dialogues without replacement is randomly

selected, and all LE expressions within each dialogue are manually annotated by the validators as the ground-truth labels. For each dialogue, the automatic annotations obtained in the previous three steps are compared with the ground-truth labels to compute the recall and precision rates, and the  $F_1$  Score is computed. If the  $F_1$  Score is not 100%, additional rule-based methods are applied in the previous steps. This annotation process is repeated until successive 100%  $F_1$  scores are achieved. This human validation step took 24 runs in total. This validation process is effective because MULTIWOZ is a task-oriented conversational dataset with a limited range of variations in sentence patterns and lexicons.

## 5 LE extraction

As depicted in Figure 1, to incorporate the LE phenomenon into conversational systems, it is necessary first to identify lexical choices that can potentially be used as LE expressions through an extraction process. The resulting candidates can then be fed into LE generators to produce appropriate responses. In this section, we introduce a novel task of extracting LE expressions from the preceding dialogue context and present two baseline models: an end-to-end approach and a pipeline approach.

### 5.1 LE extraction task definition

In this section, we introduce the task of extracting LE expressions. The task involves identifying all instances of LE expressions established in or before a given utterance, denoted by underlined text in Table 1. For example, in Table 1, if we aim to extract LE expressions for the 2<sup>nd</sup> utterance, the model must be able to identify all occurrences of “italian restaurant” and “center of town” that have already been established at the 2<sup>nd</sup> utterance (thus they are underlined) from the preceding dialogue context. However, “price range” will not be considered until the 11<sup>th</sup> utterance is established. For each sample, the model receives the preceding dialogue context as input, and its goal is to identify all instances of established LE expressions in the target utterance.

### 5.2 End-to-end approach

The end-to-end model comprises three major components: *dialogue context encoder*, *contingent fusion module*, and *linear layer*. At first, the model uses the dialogue context encoder (Devlin et al., 2019) to encode the dialogue context  $C$  (sepa-

<sup>2</sup><https://github.com/GuillaumeDD/dialign>

rated by ‘‘Human’’ and ‘‘Agent’’ annotations) into a sequence of token embeddings  $\{u_j\}_{j=1}^s$ , where  $u_j \in \mathbb{R}^{d_w}$  is the embedding for the  $j^{\text{th}}$  token and  $s$  is the sequence length of the dialogue context. Next, the fusion module, stacked by  $l$  sub-layers where each sub-layer consists of one attention layer and one feed-forward layer, is employed to update the token embeddings if and only if dialogue acts of agent utterances are inputted. For each turn, the dialogue act, whose tokens of slot and its value are concatenated together, is represented as a vector  $x \in \mathbb{R}^{d_w}$  via the dialogue context encoder to update token embeddings  $\{u_j\}_{j=1}^s$  via the attention mechanism (Bahdanau et al., 2015):  $u_j^l = \sum_j a_j u_j^{l-1}$ , where  $a_j = \frac{\exp(s_j^{l-1})}{\sum_{j=1}^K \exp(s_j^{l-1})}$  and  $s_j^{l-1}$  is the matching score computed based on the query vector  $x$  and each context vector  $u_j^{l-1}$ . Finally, each token embedding  $u_j$  is then mapped into a scalar score using a trainable linear layer to minimize the cross-entropy loss.

Given the ground-truth token labels for each token in the dialogue context, denoted as  $y = \{y_1, y_2, \dots, y_N\}$ , we minimize the cross entropy loss to predict whether the token belongs to the entrainment expression:  $\mathcal{L} = -\sum_j [\alpha \cdot y_j \log \hat{y}_j + (1 - y_j) \log(1 - \hat{y}_j)]$ , where the hyperparameter  $\alpha$  can be increased to improve the recall rate or reduced to increase the precision rate.

### 5.3 Pipeline approach

The pipeline approach for LE expression extraction involves two distinct steps: *named entity recognition* and *classification*. In the first step, the dialogue context is provided to a named entity recognition model, which extracts all the noun phrases from the context. In the next step, these extracted phrases are concatenated with the dialogue context and fed into a binary classifier model to determine whether each phrase represents a LE expression. For the named entity recognition model, we employ NLTK (Loper and Bird, 2002), and for the binary classifier model, we use BERT (Devlin et al., 2019) with an additional trainable linear layer. The fusion module described in §5.2 is utilized if system acts are included as input.

## 6 Experiments and Results

In this section, we first implement quantitative analysis using our new proposed LE measure. Then we evaluate two baseline models on the MULTIWOZ-




Figure 2:  $ENTR_{agent}$  wrt the number of turns. Solid lines represent mean values of  $ENTR_{agent}$  with a 1-D Gaussian filter. Shaded regions stand for the standard deviation.

ENTR for the LE extraction task. Additionally, we study the effect of different domains on the model performance.

Table 4: Average  $ENTR_{agent}$  on the test set. † indicates  $p < 0.001$  statistical significance.

Human	AUGPT	MINTL	HDSA	MARCO
0.798	0.676†	0.558†	0.654†	0.719†

### 6.1 Quantitative Analysis

To evaluate the discrepancy in LE between generation models and the ground-truth responses, we compute the average value of  $ENTR_{agent}$  for each response generation model in the test set. Table 4 presents the results, which show that the ground-truth agent (Human) tends to entrain the user about 0.8 times on average, while the four response generation models entrain the user with a lower degree of frequency. Specifically, in Table 1,  $ENTR_{user}$  is 0.6 and  $ENTR_{agent}$  is 1.1, which is an average value calculated across utterances. The t-test results reveal a significant difference in the means of  $ENTR_{agent}$  between the generation model and ground truth responses, indicating that there is considerable scope for improvement in state-of-the-art response generation models in terms of incorporating LE.

We also compute the  $ENTR_{agent}$  against the number of turns in dialogues. As shown in Figure 2, the degree of LE tends higher as the number of turns increases, in line with expectations that more LE happens when people speak more. The curves of response generation models are mainly under their counterparts for humans, indicating discrepan-

Table 5: The test results for the Lexical Entrainment Extraction Task under different settings, with colored arrows indicating the corresponding changes in performance relative to the baseline results. The best-performing results are highlighted in light orange and blue highlights the second place.

	Dialogue History Length = 1			Dialogue History Length = 2			Full Dialogue History Length		
	Recall	Precision	F <sub>1</sub> Score	Recall	Precision	F <sub>1</sub> Score	Recall	Precision	F <sub>1</sub> Score
<i>*End-to-End Approach</i>									
Baseline	30.7	58.3	40.3	40.5	49.7	44.7	70.0	49.3	57.8
+ user samples	33.9(▲10.4%)	52.0(▼10.8%)	41.0(▲1.7%)	40.8(▲0.7%)	47.9(▼3.6%)	44.1(▼1.3%)	72.8(▲4%)	45.9(▼6.9%)	56.3(▼2.6%)
+ agent acts	28.2(▼8.1%)	65.5(▲12.3%)	39.4(▼2.2%)	38.7(▼4.4%)	54.6(▲9.9%)	45.3(▲1.3%)	77.7(▲11%)	39.9(▼19.1%)	52.7(▼8.8%)
<i>*Pipeline Approach</i>									
Baseline	34.9	39.4	37.0	40.8	33.3	36.7	53.8	21.4	30.7
+ user samples	35.5(▲1.7%)	40.6(▲3.0%)	37.9(▲2.4%)	41.5(▲1.7%)	35.1(▲5.4%)	38.1(▲3.8%)	53.5(▲0.6%)	22.7(▲0.6%)	31.8(▲0.6%)
+ agent acts	36.9(▲5.7%)	21.1(▼46.4%)	26.9(▼27.3%)	44.6(▲9.3%)	12.2(▼63.4%)	19.1(▼48.0%)	55.7(▲4.1%)	7.8(▼53.6%)	13.7(▼55.4%)

cies between generated responses and human-level responses.

## 6.2 LE Extraction Task

We first compare the performance of two different baseline approaches and then evaluate the effect of different domains on the model performance.

**Experimental setup.** We conduct training for our models using three distinct input settings. Firstly, as our primary objective is to model a human-like agent rather than replicate the speaking style of users, we exclusively train our models on agent samples. This yields 31 436, 4 093, and 4 211 samples for the training, validation, and testing sets. Secondly, we investigate the impact of user samples by adding them to the training set, resulting in a total of 50 130 training samples. Thirdly, we consider the potential usefulness of incorporating dialogue act information, as dialogue acts can provide valuable information that may influence the next utterance. Furthermore, Table 8 reveals that the median and mode of the priming distance are both 1, indicating that speakers are more likely to use a LE expression immediately after it appears. Thus, we vary the dialogue history input range, including the last one, last two, or full history, denoted as 1, 2, full.

**Comparison between end2end and pipeline approaches.** In Table 5, we observe that introducing user samples or agent acts into the training data does not improve the performance of the baseline models. This is consistent with our finding in Section B that there are few overlaps between LE expressions and dialogue acts. Moreover, we present a breakdown of the model’s performance based on the range of dialogue contexts for each input setting. Experimental results for the end-to-end

approach reveal that decreasing the length of dialogue context has a negative impact on performance by significantly reducing the recall rate. Although the model’s attention is focused on a smaller set of dialogue contexts (precision rate increases), it is insufficient to offset the loss of missing some instances in the earlier dialogue context (recall rate decreases). For the pipeline approach, the inclusion of false positives and the low recall rate at the first named entity steps lead to a poorer performance compared to the end-to-end approach.

Table 6: Test results for five different domains using the end-to-end approach. We train the model using the agents’ utterances only. The best performance for each column is highlighted in bold.

Domains	Recall	Precision	F <sub>1</sub> Score
Taxi	63.53	44.17	52.11
Hotel	68.79	42.16	52.28
Attraction	62.56	49.80	55.01
Train	64.68	45.85	53.66
Restaurant	63.77	48.11	54.84
All	<b>70.0</b>	<b>49.3</b>	<b>57.8</b>

**The effect of different domains.** In Section B, we performed an ANOVA test which provided strong evidence of differences in the mean of lexicon sizes for the Taxi, Hotel, Attraction, Train, and Restaurant domains. Thus, we proceeded to evaluate the proposed end-to-end model with respect to these five domains while excluding the hospital and police domains, which do not have a validation and testing set, following previous works on task-oriented conversation systems (Wu et al., 2019; Kim et al., 2020; Zhu et al., 2020). Furthermore, our analysis was restricted to agent samples only.

Table 6 presents the performance of the baseline model on the MULTIWOZ-ENTR dataset across the five domains. The experimental results indi-



cate that the  $F_1$  scores for the five domains fall within a narrow range, with no clear variation in the model’s performance. However, training the model on a single domain resulted in lower performance compared to training it on all domains jointly, indicating that the model may benefit from more training data.

## 7 Related work

**Conversational systems.** Conversational AI has been a long-standing area of exploration in the natural language processing (NLP) research community (Gao et al., 2018; Fu et al., 2022; Hendriksen et al., 2020; Shi et al., 2022b; Rahmani et al., 2022b; Ramos et al., 2022; Shi and Lipani, 2023b,a; Hendriksen et al., 2022; Rahmani et al., 2023). There are broadly two categories of conversational systems, open-domain (Wu et al., 2017; Huang et al., 2019) and task-oriented (Hosseini-Asl et al., 2020; Fu and Lipani, 2023; Hendriksen et al., 2021; Sen et al., 2023). Open-domain conversational systems mainly focus on open-ended conversations such as chit-chat, while task-oriented conversational systems target accomplishing a goal described by users in natural language such as hotel booking, movie recommendation, collaborative building (Budzianowski et al., 2018; Heck et al., 2020; Shi et al., 2022a, 2023; Rahmani et al., 2022a, 2020; Veldkamp et al., 2023). A number of corpora with various domains and annotations have been collected for task-oriented conversational systems. Some datasets (Henderson et al., 2014; El Asri et al., 2017) contain conversations with a single domain, while others (Budzianowski et al., 2018; Byrne et al., 2019; Eric et al., 2020) contain multi-domains conversations. MULTIWOZ dataset (Budzianowski et al., 2018; Eric et al., 2020) spans 7 distinct domains with over 10 000 dialogues.

**Lexical Entrainment (LE) for conversational systems.** Some previous works (Hirschberg, 2008; Brockmann et al., 2005; Bakshi et al., 2021; de Jong et al., 2008; Campano et al., 2015; Dušek and Jurčiček, 2016) attempted to integrate LE in the conversationAL system. de Jong et al. (2008) presented an entrainment model for aligning the language use of a virtual agent to the level of politeness and formality displayed by the user’s utterances. Campano et al. (2015) proposed an embodied conversational agent with the ability to choose when to share appreciation with a human partner in a museum setting. Lopes et al. (2015) used rule-

based and statistical models for the integration of LE of an information-providing task in the public transport domain for spoken dialogue systems. Dušek (2017) implemented a context-aware generator to model to entrain the user’s way of speaking by learning implicitly from data. Hu et al. (2016) designed a rule-based natural language generator to produce utterances entrained to a range of utterance features used in prior utterances by a human user, including referring expressions, syntactic template selection, and tense/modal choices.

## 8 Conclusion

In this paper, we highlight the importance of incorporating LE into conversational systems and propose a new dataset, named MULTIWOZ-ENTR, designed to facilitate the study of LE. We conduct a comprehensive statistical analysis of the MULTIWOZ-ENTR dataset, providing valuable insights into its properties. Our investigation of current state-of-the-art response generation models highlights the issue of neglecting LE, and our proposed measure serves as evidence for the lack of LE in these models. To address this deficiency, we introduce two new tasks, namely, LE extraction and generation, and propose two baseline models for the LE extraction task under various settings. Our work lays the foundation for further research on the LE generation task.

## Limitations

The discussion above is based on the assumption that researchers and designers should pursue the most human-like agents. However, there is a risk of experiencing "an eerie sensation" —Mori et al. (2012) noted that some specific human-likeness designs may lead to an agent becoming repulsive, that is, falling into the uncanny valley. For example, people could be startled because of the prosthetic hand’s limp boneless grip together with its texture and coldness, although the prosthetic hand has achieved a degree of resemblance to the human form. In this paper, we are not arguing for the necessity to precisely mimic human beings. Rather we agree with Amershi et al. (2019) that Artificial Intelligence (AI) systems should still follow social norms rather than “pretend” to be human beings. We believe, as Thomas et al. (2021) suggested, that the study of natural phenomena in human-human conversations could give us inspiration for the design of machines, which is the goal of this work.

## References

- Saleema Amershi, Daniel S. Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi T. Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. [Guidelines for human-ai interaction](#). In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Shreyan Bakshi, Soumya Batra, Peyman Heidari, Ankit Arun, Shashank Jain, and Michael White. 2021. [Structure-to-text generation with self-training, acceptability classifiers and context-conditioning for the GEM shared task](#). In *Proceedings of the 1st Workshop on Natural Language Generation, Evaluation, and Metrics (GEM 2021)*, pages 136–147, Online. Association for Computational Linguistics.
- ES Brennan. 1996a. [Lexical choice and conceptual pacts in conversation](#). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22:1482–1493.
- Susan E Brennan. 1996b. [Lexical entrainment in spontaneous dialog](#). *Proceedings of ISSD*, 96:41–44.
- Carsten Brockmann, Amy Isard, Jon Oberlander, and Michael White. 2005. [Modelling alignment for affective dialogue](#). In *Workshop on adapting the interaction style to affective factors at the 10th international conference on user modeling (UM-05)*.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. [MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.
- Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Ben Goodrich, Daniel Duckworth, Semih Yavuz, Amit Dubey, Kyu-Young Kim, and Andy Cedilnik. 2019. [Taskmaster-1: Toward a realistic and diverse dialog dataset](#). In *EMNLP-IJCNLP*, Hong Kong, China. Association for Computational Linguistics.
- Sabrina Campano, Chloé Clavel, and Catherine Pelachaud. 2015. ["i like this painting too": When an eca shares appreciations to engage users](#). In *AAMAS*.
- Wenhu Chen, Jianshu Chen, Pengda Qin, Xifeng Yan, and William Yang Wang. 2019. [Semantically conditioned dialog response generation via hierarchical disentangled self-attention](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3696–3709, Florence, Italy. Association for Computational Linguistics.
- Markus de Jong, Mariët Theune, and Dennis Hofs. 2008. [Politeness and alignment in dialogues with a virtual guide](#). In *AAMAS 2008*. International Foundation for Autonomous Agents and Multiagent Systems.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *NAACL*, pages 4171–4186. Association for Computational Linguistics.
- Guillaume Dubuisson Duplessis, Chloé Clavel, and Frédéric Landragin. 2017. [Automatic measures to characterise verbal alignment in human-agent interaction](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 71–81, Saarbrücken, Germany. Association for Computational Linguistics.
- Guillaume Dubuisson Duplessis, Caroline Langlet, Chloé Clavel, and Frédéric Landragin. 2021. [Towards alignment strategies in human-agent interactions based on measures of lexical repetitions](#). *Language Resources and Evaluation*, 55(2).
- Ondřej Dušek. 2017. [Novel methods for natural language generation in spoken dialogue systems](#). *PhD Thesis*.
- Ondrej Dušek and Filip Jurčiček. 2016. [A context-aware natural language generation dataset for dialogue systems](#). In *Workshop on collecting and generating resources for chatbots and conversational agents-development and evaluation*.
- Ondřej Dušek and Filip Jurčiček. 2016. [A context-aware natural language generator for dialogue systems](#). In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 185–190, Los Angeles. Association for Computational Linguistics.
- Layla El Asri, Hannes Schulz, Shikhar Sharma, Jeremie Zumer, Justin Harris, Emery Fine, Rahul Mehrotra, and Kaheer Suleman. 2017. [Frames: a corpus for adding memory to goal-oriented dialogue systems](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 207–219, Saarbrücken, Germany. Association for Computational Linguistics.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, Adarsh Kumar, Anuj Goyal, Peter Ku, and Dilek Hakkani-Tur. 2020. [MultiWOZ 2.1: A consolidated multi-domain dialogue dataset with state corrections and state tracking baselines](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 422–428, Marseille, France. European Language Resources Association.

- Xiao Fu and Aldo Lipani. 2023. [Priming and actions: An analysis in conversational search systems](#). In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '23, page 2277–2281, New York, NY, USA. Association for Computing Machinery.
- Xiao Fu, Emine Yilmaz, and Aldo Lipani. 2022. [Evaluating the cranfield paradigm for conversational search systems](#). In *Proceedings of the 2022 ACM SIGIR International Conference on Theory of Information Retrieval*, ICTIR '22, page 275–280, New York, NY, USA. Association for Computing Machinery.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. [Neural approaches to conversational AI](#). In *SIGIR*, pages 1371–1374. ACM.
- Simon Garrod and Anthony Anderson. 1987. [Saying what you mean in dialogue: A study in conceptual and semantic co-ordination](#). *Cognition*, 27(2).
- Michael Heck, Carel van Niekerk, Nurul Lubis, Christian Geishauser, Hsien-Chin Lin, Marco Moresi, and Milica Gasic. 2020. [TripPy: A triple copy strategy for value independent neural dialog state tracking](#). In *SIGDIAL*.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014. [The second dialog state tracking challenge](#). In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*.
- Mariya Hendriksen, Maurits Bleeker, Svitlana Vakulenko, Nanne van Noord, Ernst Kuiper, and Maarten de Rijke. 2022. [Extending clip for category-to-image retrieval in e-commerce](#). In *Advances in Information Retrieval: 44th European Conference on IR Research, ECIR 2022, Stavanger, Norway, April 10–14, 2022, Proceedings, Part I*, Berlin, Heidelberg.
- Mariya Hendriksen, Ernst Kuiper, Pim Nauts, Sebastian Schelter, and Maarten de Rijke. 2020. [Analyzing and predicting purchase intent in e-commerce: anonymous vs. identified customers](#). *ArXiv preprint*, abs/2012.08777.
- Mariya Hendriksen, Artuur Leeuwenberg, and Marie-Francine Moens. 2021. [Lstm for dialogue breakdown detection: exploration of different model types and word embeddings](#). In *Increasing Naturalness and Flexibility in Spoken Dialogue Interaction: 10th International Workshop on Spoken Dialogue Systems*, pages 443–453. Springer.
- Julia Hirschberg. 2008. [Speaking more like you: Lexical, acoustic/prosodic, and discourse entrainment in spoken dialogue systems](#). In *Proceedings of the 9th SIGdial Workshop on Discourse and Dialogue*, Columbus, Ohio.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. [A simple language model for task-oriented dialogue](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Zhichao Hu, Gabrielle Halberg, Carolyn R Jimenez, and Marilyn A Walker. 2016. [Entrainment in pedestrian direction giving: How many kinds of entrainment?](#) In *Situated dialog in speech-based human-computer interaction*. Springer.
- Minlie Huang, Xiaoyan Zhu, and Jianfeng Gao. 2019. [Challenges in building intelligent open-domain dialog systems](#). *ArXiv preprint*, abs/1905.05709.
- Sungdong Kim, Sohee Yang, Gyuwan Kim, and Sang-Woo Lee. 2020. [Efficient dialogue state tracking by selectively overwriting memory](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Jonáš Kulhánek, Vojtěch Hudeček, Tomáš Nekvinda, and Ondřej Dušek. 2021. [AuGPT: Auxiliary tasks and data augmentation for end-to-end dialogue with pre-trained language models](#). In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. [DailyDialog: A manually labelled multi-turn dialogue dataset](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, and Pascale Fung. 2020. [MinTL: Minimalist transfer learning for task-oriented dialogue systems](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3391–3405, Online. Association for Computational Linguistics.
- 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*, Philadelphia, Pennsylvania, USA.
- José Lopes, Maxine Eskenazi, and Isabel Trancoso. 2015. [From rule-based to data-driven lexical entrainment models in spoken dialog systems](#). *Computer Speech & Language*, 31(1).
- Masahiro Mori, Karl F MacDorman, and Norri Kageki. 2012. [The uncanny valley \[from the field\]](#). *IEEE Robotics & automation magazine*, 19(2):98–100.
- Md. Nasir, Sandeep Nallan Chakravarthula, Brian R. W. Baucom, David C. Atkins, Panayiotis G. Georgiou, and Shrikanth Narayanan. 2019. [Modeling interpersonal linguistic coordination in conversations using word mover's distance](#). In *Interspeech*. ISCA.



- Tomáš Nekvinda and Ondřej Dušek. 2021. [Shades of BLEU, flavours of success: The case of MultiWOZ](#). In *Proceedings of the 1st Workshop on Natural Language Generation, Evaluation, and Metrics (GEM 2021)*. Association for Computational Linguistics.
- Ani Nenkova, Agustín Gravano, and Julia Hirschberg. 2008. [High frequency word entrainment in spoken dialogue](#). In *Proceedings of ACL-08: HLT, Short Papers*. Association for Computational Linguistics.
- Martin J Pickering and Simon Garrod. 2004. [Toward a mechanistic psychology of dialogue](#). *Behavioral and brain sciences*, 27(2).
- Robert Porzel, Annika Scheffler, and Rainer Malaka. 2006. [How entrainment increases dialogical effectiveness](#). In *Proceedings of the IUI*. Citeseer.
- Hossein A. Rahmani, Mohammad Aliannejadi, Mitra Baratchi, and Fabio Crestani. 2020. Joint geographical and temporal modeling based on matrix factorization for point-of-interest recommendation. In *Advances in Information Retrieval*, pages 205–219, Cham. Springer International Publishing.
- Hossein A. Rahmani, Mohammadmehdi Naghiaei, Mahdi Dehghan, and Mohammad Aliannejadi. 2022a. [Experiments on generalizability of user-oriented fairness in recommender systems](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*. Association for Computing Machinery.
- Hossein A. Rahmani, Mohammadmehdi Naghiaei, Ali Tourani, and Yashar Deldjoo. 2022b. [Exploring the impact of temporal bias in point-of-interest recommendation](#). In *Proceedings of the 16th ACM Conference on Recommender Systems, RecSys '22*. Association for Computing Machinery.
- Hossein A. Rahmani, Xi Wang, Yue Feng, Qiang Zhang, Emine Yilmaz, and Aldo Lipani. 2023. [A survey on asking clarification questions datasets in conversational systems](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Jerome Ramos, To Eun Kim, Zhengxiang Shi, Xiao Fu, Fanghua Ye, Yue Feng, and Aldo Lipani. 2022. [Condit: A state machine like architecture for multi-modal task bots](#). *Alexa Prize TaskBot Challenge Proceedings*.
- David Reitter, Frank Keller, and Johanna D. Moore. 2006. [Computational modelling of structural priming in dialogue](#). In *NAACL: Short Papers*.
- David Reitter and Johanna D. Moore. 2007. [Predicting success in dialogue](#). In *ACL*. Association for Computational Linguistics.
- Procheta Sen, Xi Wang, Ruiqing Xu, and Emine Yilmaz. 2023. [Task2kb: A public task-oriented knowledge base](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(13):16482–16484.
- Zhengxiang Shi, Yue Feng, and Aldo Lipani. 2022a. [Learning to execute actions or ask clarification questions](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2060–2070, Seattle, United States. Association for Computational Linguistics.
- Zhengxiang Shi and Aldo Lipani. 2023a. [Dept: Decomposed prompt tuning for parameter-efficient fine-tuning](#). *ArXiv preprint*, abs/2309.05173.
- Zhengxiang Shi and Aldo Lipani. 2023b. [Don't stop pretraining? make prompt-based fine-tuning powerful learner](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Zhengxiang Shi, Francesco Tonolini, Nikolaos Aletras, Emine Yilmaz, Gabriella Kazai, and Yunlong Jiao. 2023. [Rethinking semi-supervised learning with language models](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5614–5634, Toronto, Canada. Association for Computational Linguistics.
- Zhengxiang Shi, Qiang Zhang, and Aldo Lipani. 2022b. [Stepgame: A new benchmark for robust multi-hop spatial reasoning in texts](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11321–11329.
- Paul Thomas, Mary Czerwinski, Daniel McDuff, and Nick Craswell. 2021. [Theories of conversation for conversational IR](#). *ACM Transactions on Information Systems (TOIS)*, 39(4):1–23.
- Karel Veldkamp, Mariya Hendriksen, Zoltán Szilávik, and Alexander Keijser. 2023. [Towards contrastive learning in music video domain](#). *ArXiv preprint*, abs/2309.00347.
- Kai Wang, Junfeng Tian, Rui Wang, Xiaojun Quan, and Jianxing Yu. 2020. [Multi-domain dialogue acts and response co-generation](#). In *ACL*, pages 7125–7134, Online. Association for Computational Linguistics.
- Arthur Ward and Diane Litman. 2007. [Measuring convergence and priming in tutorial dialog](#). *University of Pittsburgh*.
- Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019. [Transferable multi-domain state generator for task-oriented dialogue systems](#). In *ACL*. Association for Computational Linguistics.
- Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. [Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots](#). In *ACL*, pages 496–505, Vancouver, Canada. Association for Computational Linguistics.
- Su Zhu, Jieyu Li, Lu Chen, and Kai Yu. 2020. [Efficient context and schema fusion networks for multi-domain dialogue state tracking](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 766–781, Online. Association for Computational Linguistics.



## Appendix Overview

The appendix is structured as follows:

**Appendix §A.** introduces several measurements for LE.

**Appendix §B.** presents statistical analysis of the proposed MULTIWOZ-ENTR datasets.

**Appendix §C.** provides three examples of the MULTIWOZ-ENTR dataset.

**Appendix §D.** provides examples of rule-based filters used in the dataset construction.

**Appendix §E.** provides implementation details and hyperparameters used in our experiments.

### A Other Measures of LE

In this section, we introduce several additional measures for LE, adapted from [Nenkova et al. \(2008\)](#); [Dubuisson Duplessis et al. \(2017\)](#):

**Expression Lexicon Size (ELS).** Let  $\mathcal{E}$  be the set of all established LE expressions given a dialogue. Then the expression lexicon size is the number of established LE expressions given a dialogue, that is,  $|\mathcal{E}|$ , where  $|\cdot|$  represents the cardinality of the set. The ELS is 10 for the dialogue in Table 1.

**Initial Expression Ratio ( $IER_s$ ).** The initial expression ratio is the percentage of established LE expressions initiated by the speaker  $s$ , and is computed as:

$$IER_s = \frac{|\{e \in \mathcal{E} : e \text{ initiated by } s\}|}{|\mathcal{E}|}, \quad (2)$$

where  $IER_{human}$  is 40% and  $IER_{agent}$  is 60% in Table 1.

**Expression Repetition Ratio ( $ERR_s$ ).** Let  $\mathcal{T}_E$  be the set of all tokens from all instances of established LE expressions from a dialogue, and  $\mathcal{T}$  be the set of all tokens from this dialogue, where  $\mathcal{T}_E$  is a subset of  $\mathcal{T}$ . Then the expression repetition ratio is the number of tokens from LE expressions in the speaker  $s$ 's utterances divided by the total number of tokens in a dialogue:

$$ERR_s = \frac{|\{t \in \mathcal{T}_E : t \text{ from speaker } s\}|}{|\mathcal{T}|} \quad (3)$$

For example, there are total 302 tokens for the dialogue in Table 1. Among 302 tokens, 17 are from instances uttered by the user and 21 are from instances uttered by the agent. Thus, in this case,  $ERR_{user}$  is 5.6% and  $ERR_{agent}$  is 7.0%.

**Frequency.** Frequency refers to the number of utterances in which the established LE expression appears. For example, in Table 1, the frequency of “reference number” is 3.

**Size.** Size refers to the number of tokens which the established LE expression is made up of. For example, in Table 1, the size of “restaurant” is 1, and the size of “reference number” is 2.

**Span.** Span refers to the number of utterances in which the first appearance of the established LE expression comes across to its last appearance, where both the first and last utterances count. For example, in Table 1, the span of “reference number” is 11, and the span of “day” is 2.

**Density.** Density refers to the ratio between the established LE expression’s frequency and its span. For example, in Table 1, the density of “ask” is 0.8, and the density of “price range” is 0.3.

**Priming.** Priming refers to the number of repetitions of the established LE expression by the initiator before being used by the other interlocutor. In Table 1, the priming of “price range” is 2.

**Priming Distance.** Priming Distance refers to the number of utterances in which the first appearance of the established LE expression comes across to its first appearance from another speaker. For example, in Table 1, the priming distance of “italian restaurant”.

### B Statistical Analysis

This process results in the MULTIWOZ-ENTR dataset, consisting of 12 038, 1 000, and 1 000 dialogues for the training, validation, and testing sets with 62 961 LE expressions. We present a thorough statistical analysis of the MULTIWOZ-ENTR dataset below.

**Measures’ statistics.** The first step of our statistical analysis is to compute the expression lexicon size (ELS) for each of the 10, 438 dialogues in the MULTIWOZ-ENTR dataset, as defined in appendix A. As depicted in Figure 3, the ELS varies across the dialogues, with an average of 6.03. Approximately 15% of the dialogues have more than 10 LE expressions. In contrast, 288 of 10 438 dialogues have zero LE expression. The maximum ELS value for a single dialogue is 30.

Table 7 presents the statistical analysis of two dialogue-level measures, initial expression ratio




Figure 3: Distribution of Expression Lexicon Size (ELS).

Table 7: Statistics with respect to speakers: *a* represents agent and *u* represents human.

	IER <sub>u</sub>	IER <sub>a</sub>	ERR <sub>u</sub>	ERR <sub>a</sub>
<b>Mean.</b>	75.02	24.98	5.29	5.76
<b>Std.</b>	21.39	21.39	2.59	2.94

(IER) and expression repetition ratio (ERR). The results indicate that the ERR values for the agent are higher than those for the human, suggesting that the agent’s role players are more likely to entrain than the human’s role players. This finding aligns with our observation that about 75% of the LE expressions are initiated by the human speaker. It is reasonable to assume that the agent’s utterances are more likely to be entrained when the human speaker initiates more LE expressions since the human speakers are those initiating the conversations and making requests. As the MULTIWOZ dataset involves human-to-human interactions, this implies that users are more likely to initiate a LE expression, and a human-like agent’s responses need to incorporate a higher degree of LE to make conversations more engaging and successful. This underscores the importance of integrating LE expressions into agent utterances to enhance their effectiveness.

The statistical analysis of expression-level measures is presented in Table 8. The median and mode of the priming distance are both 1, which means that a large proportion of the LE expressions (63.79%) occur within two turns of the conversation. This finding suggests that the majority of LE expressions are produced immediately after the first use of the expression. To investigate the effects

Table 8: Statistics with respect to 62 961 expressions: Distance stands for the priming distance.

	Mean.	Std.	Min.	Max.	Median.	Mode.
<b>Frequency</b>	2.62	0.98	2	17	2	2
<b>Size</b>	1.47	0.88	1	17	1	1
<b>Span</b>	5.02	4.05	2	40	4	2
<b>Priming</b>	1.11	0.34	1	6	1	1
<b>Distance</b>	2.55	2.90	1	35	1	1
<b>Density</b>	0.72	0.30	0.06	2.00	0.75	1.00

of this phenomenon on the model’s performance, we limit the model’s attention to the last two utterances of the preceding dialogue context in the entrainment extraction task in §6.2.

**The effect of different domains.** To investigate the penitential effects of different domains on LE, we compute the expression lexicon size with respect to different domains. We observe that the Police and Hospital domains have relatively smaller lexicon sizes compared to the other five domains. However, it should be noted that these two domains have fewer training samples and no validation or testing samples. To determine if there are statistically significant differences between the domains, we conducted a one-way ANOVA test with the null hypothesis  $H_0: \mu_{Taxi} = \mu_{Hotel} = \mu_{Attraction} = \mu_{Train} = \mu_{Restaurant}$ , where  $\mu_d$  represents the mean of the expression lexicon size for domain  $d$ . Our experimental results indicate that the null hypothesis can be rejected with a p-value of  $2.59e^{-43}$ , suggesting strong evidence of differences in the mean of lexicon sizes among the five domains with large sample sizes. Therefore, we further examine the effects of different domains on the entrainment extraction task in Section 6.2.

**The effect of dialogue acts.** In the response generation process, the entrainment extraction module’s output is used in conjunction with the outputs from other modules, such as dialogue state tracking and intent detection, as shown in Figure 1. It is essential to understand the extent of overlap among these outputs. To address this issue, we examine the lexicon overlap between LE expressions and dialogue acts. We stem each token with the NLTK toolkit and obtain 17 594 unique tokens for dialogue acts and 3 477 for LE expressions, with only 874 tokens belonging to both sets. The analysis shows that the lexicon of LE expressions shares only a small fraction of overlap with the lexicon of dialogue acts, indicating that the degree of overlap with original forms and contextual meaning is also

likely to be low. This suggests that the incorporation of dialogue acts may have a limited effect on the LE extraction task. We evaluate the impact of incorporating dialogue acts on the entrainment extraction task empirically in §6.2.

## C Case Studies

In Table 10, 11 and 12, we provide three examples as a supplementary of §3.

## D Rules Used in Dataset Construction

We provide two examples of rules defined in §4 to check whether "help" and "booking" are verbs. For example, we find all cases where "booking" is used as a verb. We add one space before "in booking" to avoid "booking" in "train booking" being counted as a verb. There are many similar words such as "contact" and "cost". For the sake of the space, we will not list all these rules. The code will be released after this paper is accepted.

```

1 def special_case_booking(utterance: str)
2     -> bool:
3     """
4     Check whether booking is a verb
5     given an utterance
6     """
7     special_cases = [
8         'are you booking', ' before
9         booking',
10        'booking your ', ' i booking',
11        'booking the ', 'booking that',
12        ' in booking', ' be booking'
13    ]
14    for sc in special_cases:
15        if sc in utterance:
16            return True
17    return False
18
19 def special_case_help(utterance: str) ->
20     bool:
21     """
22     Check whether help is a verb given
23     an utterance
24     """
25     special_cases = [
26         'can you help', "can I help",
27         'could you help', 'help you',
28         "could help", 'would help',
29         'may help', 'might help',
30         'help me', 'will help',
31         "to help", "can help",
32         "can you please help",
33    ]
34    for sc in special_cases:
35        if sc in utterance:
36            return True
37    return False

```

Table 9: Hyperparameters for the LE extraction task.

Hyperparameter	Assignment
number of steps	100 epochs or 50K steps
batch size	24
maximum learning rate	5e-5
learning rate optimizer	AdamW
max sequence length	512
Adam epsilon	1e-6
Adam beta weights	0.9, 0.98
learning rate scheduler	Warmup linear
Weight decay	0.01
Warmup proportion	0.06
learning rate decay	linear
fusion layers	2

## E Training Details

During the tokenization process of BERT, any previously unseen words are segmented into sub-words, which results in the creation of multiple tokens for a single word. To obtain a score for the entire word, we use the score assigned to the first sub-word and exclude the remaining sub-words when calculating the cross-entropy loss. Additionally, we exclude the special tokens [CLS], [SEP], and [PAD] when calculating the loss. In scenarios where the dialogue history is incomplete, some instances of LE expressions may not be detected. To ensure unbiased evaluation, we consider instances of these expressions outside the dialogue history as false negative examples. When the length of the dialogue history is not complete, it is inevitable that certain instances of LE expressions will not be captured. To account for this, instances of these expressions outside the dialogue history are automatically counted as false negative examples to ensure a fair evaluation when compared to scenarios where the dialogue history is complete. To evaluate the end-to-end approach, we assess its performance at the noun phrase level, calculating the recall rate, precision rate, and  $F_1$  score. The positive weight  $\alpha$  for the cross-entropy loss is varied between 4 and 12.

Table 10: Example with two agent responses. Current response generation models could partially entrain the user's way of speaking such as "*contact number*". However, it is not optimal for these generation models to use "*taxi*" when users are using "*cab*".

Model	Utterance
USER	I need a <b>cab</b> too please, arriving by my <b>booked time</b> , car type and <b>contact number</b> would be helpful too please.
Agent	You need a <b>cab</b> to get you to the restaurant before the <b>booked time</b> and where do you need <b>cab</b> to pick you up at?
USER	From the Cambridge Punter, please.
AGENT	A black tesla is booked, with <b>contact number</b> 07774975314.
<b>Generated responses for the first agent response</b>	
AUGPT	I have booked a blue skoda for you. The contact number is 85534112262.
MINTL	I have booked you a BMW . The contact number is 07338019809 . Where will you be departing from ?
HDSA	Where will you be departing from and where are you going ?
MARCO	I would be happy to book a <b>taxi</b> for you . where are you departing from and what is your destination ?
<b>Generated responses for the second agent response</b>	
AUGPT	I have booked you a black honda and the <b>contact number</b> is 91788856987. Is there anything else I can help you with?
MINTL	I have booked your taxi . Be expecting a BMW . If you need to reach them , please call 07338019809 .
HDSA	I have booked you a [taxi_type] . the <b>contact number</b> is [taxi_phone] . is there anything else i can help you with ?
MARCO	I have booked a [taxi_type] for you . the <b>contact number</b> is [taxi_phone] . is there anything else i can help you with ?

Table 11: Example with one agent response. In the ground-truth agent response, lexical entrainment with the user's requirement including "*Monday*" and "*5 nights*" makes the response more informative and less ambiguous.

Model	Utterance
USER	Yes, I would like to book a stay for 1 person, starting <b>Monday</b> and for <b>5 nights</b> .
AGENT	Your room has been booked at the Bridge Guest House for <b>5 nights</b> starting <b>Monday</b> . Your reference code is 7VM9FIPC. Is there anything else I can do to help you?
<b>Generated responses for the first agent response</b>	
AUGPT	I have made those reservations and your reference number is 00000011.
MINTL	I have booked your room. Your reference number is QJTIXK8C.
HDSA	Booking was successful. Reference number is 7vm9fipc. Is there anything else i can help you with?
MARCO	I was able to book you at the [hotel_name] for [value_count] nights. Your reference number is 7vm9fipc. Is there anything else i can help you with?

Table 12: Example with one agent response. In the ground-truth agent response, lexical entrainment with the user's requirement including "*8 people*" and "*4 nights*" makes the response more informative and less ambiguous.

Model	Utterance
USER	I need it for <b>8 people</b> , and for <b>4 nights</b> starting Saturday.
AGENT	Great, I have booked your stay for <b>8 people</b> for <b>4 nights</b> . Your reference number is PB41WUOT. May I help with anything else?
<b>Generated responses for the first agent response</b>	
AUGPT	Booking was successful. Reference number is: 00000003.
MINTL	I have booked your room at Allenbell for saturday for <b>4 nights</b> for <b>8 people</b> . Your reference number is 6KNJSQXY. Is there anything else i can help you with?
HDSA	Booking was successful. Reference number is pb41wuot. Is there anything else i can help you with?
MARCO	I was able to book you at the [hotel_name] for [value_count] nights. your reference number is pb41wuot. Is there anything else i can help you with?