

A Dataset for Building Code-Mixed Goal Oriented Conversation Systems

Suman Banerjee, Nikita Moghe, Siddhartha Arora and Mitesh M. Khapra
Indian Institute of Technology Madras, India

{suman, nikitavam, sidarora, miteshk}@cse.iitm.ac.in

Abstract

There is an increasing demand for goal-oriented conversation systems which can assist users in various day-to-day activities such as booking tickets, restaurant reservations, shopping, *etc.* Most of the existing datasets for building such conversation systems focus on monolingual conversations and there is hardly any work on multilingual and/or code-mixed conversations. Such datasets and systems thus do not cater to the multilingual regions of the world, such as India, where it is very common for people to speak more than one language and seamlessly switch between them resulting in code-mixed conversations. For example, a Hindi speaking user looking to book a restaurant would typically ask, “Kya tum is *restaurant* mein ek *table book* karne mein meri *help* karoge?” (“Can you help me in booking a table at this restaurant?”). To facilitate the development of such code-mixed conversation models, we build a goal-oriented dialog dataset containing code-mixed conversations. Specifically, we take the text from the DSTC2 restaurant reservation dataset and create code-mixed versions of it in Hindi-English, Bengali-English, Gujarati-English and Tamil-English. We also establish initial baselines on this dataset using existing state of the art models. This dataset along with our baseline implementations is made publicly available for research purposes.

1 Introduction

Over the past few years, there has been an increasing demand for virtual assistants which can help users in a wide variety of tasks in several domains such as entertainment, finance, healthcare, e-commerce, *etc.* To cater to this demand, several commercial conversation systems such as Siri, Cortana, Allo have been developed. While these systems are still far from general purpose open domain chat, they perform reasonably well for certain goal-oriented tasks such as setting alarms/reminders, booking appointments, checking movie show timings, finding directions for navigation, *etc.* Apart from these commercial systems, there has also been significant academic research to advance the state of the art in conversation systems (Shang et al., 2015; Vinyals and Le, 2015; Yao et al., 2015; Li et al., 2016a; Li et al., 2016b; Serban et al., 2017). Most of this academic research is driven by publicly available datasets such as Twitter conversation dataset (Ritter et al., 2010), Ubuntu dialog dataset (Lowe et al., 2015), Movie subtitles dataset (Lison and Tiedemann, 2016) and DSTC2 restaurant reservation dataset (Henderson et al., 2014a). In this work, we focus on goal-oriented conversations such as the ones contained in the DSTC2 dataset.

Most of the datasets and state of the art systems mentioned above are monolingual. Specifically, all the utterances and responses in the conversations are in one language (typically, English) and there are no multilingual and/or code-mixed utterances/responses. However, in several multilingual regions of the world, such as India, it is natural for speakers to produce utterances and responses which are multilingual and code-mixed. For example, Table 1 shows real examples of how bilingual speakers from India talk when requesting someone to help them reserve a restaurant or book movie tickets. As it can

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

Languages	Utterances
English	Speaker 1: Hi, Can you help me in booking a table at this restaurant? Speaker 2: Sure, would you like something in cheap, moderate or expensive price range?
Hindi-English	Speaker 1: <i>Hi</i> , kya tum is <i>restaurant</i> mein ek <i>table book</i> karne mein meri <i>help</i> karoge? Speaker 2: <i>Sure</i> , kya aap <i>cheap</i> , <i>moderate</i> ya <i>expensive price range</i> mein kuch <i>like</i> karenge?
Bengali-English	Speaker 1: <i>Hi</i> , tumi ki ei <i>restaurant</i> ey ekta <i>table book</i> korte <i>help</i> korbe amake? Speaker 2: <i>Sure</i> , aapni ki <i>cheap</i> , <i>moderate</i> na <i>expensive price range</i> ey kichu <i>like</i> korben ?
English	Speaker 1: Hello, can you tell me about the show timings of “Black Panther”? Speaker 2: Sure, would you like to book tickets for today or any other day?
Gujarati-English	Speaker 1: <i>Hello</i> , mane Black Panther na <i>show timings</i> janavo. Speaker 2: <i>Sure</i> , shu tame aaj ni ke koi anya divas ni <i>ticket book</i> karva mango cho?
Tamil-English	Speaker 1: <i>Hello</i> , “Black Panther” <i>show timings</i> eppo epponu solla mudiuma Speaker 2: <i>Kandipa</i> , <i>tickets</i> innaiku <i>book</i> pannanuma illana vera ennaikku?

Table 1: Example code-mixed utterances in the specified languages.

be seen, when engaging in such informal conversations it is very natural for such speakers to use code-mixed utterances, mixing their native language with English. Apart from India, such code-mixing is also prevalent in other multilingual regions of the world, for example, Spanglish (Spanish-English), Frenglish (French-English), Porglish (Portuguese-English) and so on. To cater to such users, it is essential to create datasets containing code-mixed conversations and thus facilitate the development of code-mixed conversation systems.

With the above motivation, we build a dataset containing code-mixed goal-oriented conversations for four Indian languages. Specifically, we take every utterance from the DSTC2 restaurant reservation dataset and ask a mix of in-house and crowdsourced workers to create a corresponding code-mixed utterance involving their native language and English. We simply instructed the workers to (i) assume that they were chatting with a friend who spoke the same native language as them in addition to English, (ii) not try very hard to translate the sentence completely to their native language but feel free to switch to English whenever they wanted (just as they would in a normal conversation with a friend) and (iii) use Romanized text instead of the native language’s script. The resulting dataset contains utterances of the type shown in Table 1. We found that 87.73% of the created utterances were code-mixed, 7.18% had only English words and 5.09% had only native language words. The four languages that we chose were Hindi, Bengali, Tamil and Gujarati which have 422M, 83M, 60M and 46M native speakers respectively.

Apart from reporting various statistics about this data (such as CM-index (Gambäck and Das, 2016) and I-index (Guzmán et al., 2016)), we also report some initial baselines by evaluating some state of the art approaches on the proposed dataset. Specifically, we evaluate a standard sequence-to-sequence model with an attention mechanism (Bahdanau et al., 2015) and a hierarchical recurrent encoder-decoder model (Serban et al., 2016). Our code implementing these models along with the dataset is available freely for research purposes¹. To the best of our knowledge, this is the first conversation dataset containing code-mixed conversations and will hopefully enable further research in this area. In particular, since the data is 5-way parallel (English, Bengali, Hindi, Tamil, Gujarati) it would be useful for building jointly trained code-mixed models.

2 Related Work

Serban et al. (2015) report an excellent (and up-to-date) survey of existing dialog datasets. For brevity, we only mention some of the important points from their survey and refer the reader to the original paper for more details. To begin with, we note that existing dialog datasets can be categorized along 3 main dimensions. The first dimension is the modality of the dataset, *i.e.*, whether the dataset contains spoken conversations (Godfrey et al., 1992; Kim et al., 2016) or text conversations (Forsyth and Martell, 2007; Ritter et al., 2010; Lowe et al., 2015). The second dimension is whether the dataset contains goal-oriented conversations or open-ended conversations. A goal-oriented conversation typically involves chatting for the sake of completing a task such as the Dialog State Tracking Challenge (DSTC) datasets which involve tasks for reserving a restaurant (Henderson et al., 2014a), checking bus

¹<https://github.com/sumanbanerjeel/Code-Mixed-Dialog>

User	Bot
1. <SILENCE>	Hello, welcome to the Cambridge restaurant system. You can ask for restaurants by area, price range or food type. How may I help you?
2. i'm looking for a moderately priced restaurant.	What kind of food would you like?
3. it doesn't matter.	What part of town do you have in mind?
4. the south part of town.	api_call R_cuisine south moderate
	<p>pizza_hut_cherry_hinton R_post_code pizza_hut_cherry_hinton_post_code</p> <p>pizza_hut_cherry_hinton R_cuisine italian</p> <p>pizza_hut_cherry_hinton R_location south</p> <p>pizza_hut_cherry_hinton R_phone pizza_hut_cherry_hinton_phone</p> <p>pizza_hut_cherry_hinton R_address pizza_hut_cherry_hinton_address</p> <p>pizza_hut_cherry_hinton R_price moderate</p> <p>pizza_hut_cherry_hinton R_rating 3</p> <p>restaurant_alimentum R_post_code restaurant_alimentum_post_code</p> <p>restaurant_alimentum R_cuisine european</p> <p>restaurant_alimentum R_location south</p> <p>restaurant_alimentum R_phone restaurant_alimentum_phone</p> <p>restaurant_alimentum R_address restaurant_alimentum_address</p> <p>restaurant_alimentum R_price moderate</p> <p>restaurant_alimentum R_rating 10</p>
5. <SILENCE>	restaurant_alimentum is a nice restaurant in the south of town serving modern_european food.
6. can i get the address?	Sure, restaurant_alimentum is on restaurant_alimentum_address.
7. and whats the phone number for that?	The phone number of restaurant_alimentum is restaurant_alimentum_phone.
8. okay, thank you, good bye.	you are welcome.

Table 2: An example chat from the English version of DSTC2 dataset (Bordes and Weston, 2017).

schedules (Williams et al., 2013), collecting tourist information (Henderson et al., 2014b) and so on. Such datasets are also typically domain-specific. Open-ended conversations on the other hand involve general chat on any topic and there is no specific end task. Some popular examples of datasets containing such open-ended conversations are the Ritel Corpus (Rosset and Petel, 2006), NPS Chat Corpus (Forsythand and Martell, 2007), Twitter Corpus (Ritter et al., 2010), *etc.* The third dimension is whether the dataset contains human-human conversations or human-bot conversations. As the name suggests, human-bot conversation datasets contain conversations between humans and an existing conversation system (typically a domain-specific goal-oriented bot) (Williams et al., 2013; Henderson et al., 2014a; Henderson et al., 2014b). Human-human conversations, on the other hand, can contain spontaneous conversations between humans, as are typically observed in discussion forums (Walker et al., 2012), chat rooms (Forsythand and Martell, 2007), SMS messages (Chen and Kan, 2013) and so on. Human-human conversations can also contain scripted dialogs such as scripts of movies (Banchs, 2012), TV shows (Roy et al., 2014), *etc.* It is surprising that of the 63 conversation datasets developed in the past (Serban et al., 2015), none contain multilingual conversations. In particular, none of them contain code-mixed conversations from multilingual regions of the world. There is clearly a need to fill this gap and we believe that the dataset developed as a part of this work is a small step in that direction.

In general, the research community has been interested in developing datasets, tools and approaches for code-mixed content. This interest is largely triggered by the abundance of code-mixed content found in chats, emails, social media platforms, *etc.* In the context of such code-mixed content, existing works have looked at the problems of language identification (Nguyen and Dogruöz, 2013; Solorio et al., 2014; Barman et al., 2014; Molina et al., 2016), part-of-speech tagging (Barman et al., 2016; Ghosh et al., 2016; AlGhamdi et al., 2016), user profiling (Khapra et al., 2013), topic modeling (Rosner and Farrugia, 2007), information retrieval (Chakma and Das, 2016) and language modeling (Adel et al., 2013a; Adel et al., 2013b; Adel et al., 2015). However, to the best of our knowledge, ours is the first work on developing code-mixed conversation systems for goal-oriented dialogs.

# of Utterances	49167
# of Unique utterances	6733
Average # of utterances per dialog	15.19
Average # of words per utterance	7.71
Average # of words per dialog	120.33
Average # of KB triples per dialog	38.24
# of Train Dialogs	1168
# of Validation Dialogs	500
# of Test Dialogs	1117
Vocabulary size	1229

Table 3: Statistics of the English version of DSTC2 dataset

3 Background: DSTC2 Restaurant Reservation Dataset

We build on top of the goal-oriented restaurant reservation dialog dataset which was released as part of the second Dialog State Tracking Challenge (DSTC2) (Henderson et al., 2014a). This dataset contains conversations between crowdsourced workers and existing dialog systems (bots). Specifically, the workers were asked to book a table at a restaurant with the help of a bot. These dialog systems consisted of modules like automatic speech recognizer, natural language interpreter, dialog manager, response generator and a speech synthesizer (Young, 2000). The dialog manager used policies which were either hand-crafted or learned by formulating the problem as a partially observable Markov decision process (POMDP) (Williams and Young, 2007). The speech input from the user was first converted to text and then fed to the dialog system. For this, the authors used two Automatic Speech Recognition (ASR) modules out of which one was artificially degraded in order to simulate noisy environments. The workers could request for restaurants based on 3 slots: *area* (5 possible values), *cuisine* (91 possible values) and *price range* (3 possible values). The workers were also instructed to change their goals and look for alternative *areas*, *cuisines* and *price ranges* in the middle of the dialog. This was done to account for the unpredictability in natural conversations. The conversations were then transcribed and the utterances were labeled with different dialog states. For example, each utterance was labeled with its semantic intent representation (*request[area]*, *inform[area = north]*) and the dialog turns were labeled with annotations such as constraints on the slots (*cuisine = italian*), requested slots (*requested = {phone, address}*) and the method of search (*by_constraints*, *by_alternatives*). Such annotations are useful for domain-specific slot-filling based dialog systems.

Bordes and Weston (2017) argued that for various domains collecting such explicit annotations for every state in the dialog is tedious and expensive. Instead, they emphasized on building end-to-end dialog systems (as opposed to slot-filling based systems) by adapting this dataset and treating it as a simple sequence of utterance-response pairs (without any explicit dialog states associated with the utterances). In addition, the authors also created API calls which can be issued to an underlying Knowledge Base (KB) and appended the resultant KB triples to each dialog. Table 2 shows one small sample dialog from this adapted dataset along with the API calls. Notice that the API call uses the information of all the constraints specified by the user so far and then receives all triples from the restaurant KB which match the user’s requirements. This dataset facilitated the development of models (Bordes and Weston, 2017; Seo et al., 2017; Williams et al., 2017; Eric and Manning, 2017) which just predict the bot utterances and API calls without explicitly tracking the slots. Table 3 reports the statistics of this dataset. In this work, we create code-mixed versions of this dataset in 4 different languages as described below.

4 Code-Mixed Dialog Dataset

In this section, we describe the process used for creating a new dataset containing code-mixed conversations. Specifically, we describe (i) the process used for extracting unique utterance templates from the original DSTC2 dataset, (ii) the process of creating code-mixed translations of these utterances with the help of in-house and crowdsourced workers and (iii) the process used for evaluating the collected conversations. Finally, we report some statistics about the dataset.

4.1 Extracting Unique Utterance Templates

We found that many utterances in the original English version of DSTC2 dataset (henceforth referred to as En-DSTC2) have the same sentence structure but only differ in the values of the *areas*, *cuisines*, *price ranges* and entities such as *restaurant names*, *addresses*, *phone numbers* and *post codes*. For example, consider these two sentences which only differ in the *area* and *cuisine*: (i) “Sorry, there is no *chinese* restaurant in the *north* part of town.” and (ii) “Sorry, there is no *italian* restaurant in the *west* part of town”. Both these sentences can be thought of as instantiations of the generic template: “Sorry, there is no [CUISINE] restaurant in the [AREA] part of town.” wherein the placeholders [AREA] and [CUISINE] get replaced by different values. We used the KB provided by Bordes and Weston (2017) to find all the entities appearing in all the utterances and replaced them by placeholders such as: [AREA], [CUISINE], [PRICE], [RESTAURANT], [ADDRESS], [PHONE] and [POST_CODE]. Further, since the authors had mentioned that the KB provided was not perfect/complete, we did some manual inspection to find all such entities and came up with a list of 536 such entity words. After replacing these words with their respective placeholders we obtained 3590 unique English utterances.

4.2 Creating Code-Mixed Translations

According to Myers-Scotton (1993) code-mixing involves a native language which provides the morphosyntactic frame and a foreign language whose *linguistic units* such as *phrases*, *words* and *morphemes* are inserted into this morphosyntactic frame. The native language is called the *Matrix* while the foreign language is called the *Embedding*. Our work focuses on creating a conversation dataset wherein 4 different Indian languages, *viz.*, Hindi, Bengali, Gujarati and Tamil serve as the *Matrix* and English serves as the *Embedding*. We used a mix of in-house and crowdsourced workers to create a code-mixed version of the original DSTC2 dataset. For example, for Hindi and Gujarati, we did not have enough in-house speakers so we completely relied on crowdsourcing for creating the data but then used in-house workers to verify the collected data. For Bengali, all the data was created by in-house annotators who were native Bengali speakers and proficient in English. Lastly, for Tamil, roughly 40% of the data was created with the help of crowdsourced workers and the rest with the help of in-house workers. Irrespective of whether the workers were crowdsourced or in-house we used the same set of instructions as described below.

We instructed the annotators to assume that they were chatting with a friend who is a native speaker of Hindi (or Gujarati, Bengali, Tamil) but also speaks English well (typically, because English was the language in which the friend did most of his/her education). To explain the idea of code-mixing, we showed them example utterances where it was natural for the user to mix English words while chatting in the native language. They were then shown an English utterance from the DSTC2 dataset and asked to create its code-mixed translation in the native language keeping the above code-mixed examples in mind. They were asked to use Roman script irrespective of whether the word being used belongs to English or the native language (in particular, they were clearly instructed to not use the native language’s script). As expected, we observed that while translating, the annotators tend to retain some difficult-to-translate and colloquially relevant English words as it is. The annotators were also clearly instructed to refrain from producing pure translations (*i.e.*, they were asked to not try hard to translate English words which they would typically not translate in an informal conversation). Also, the annotators were instructed to retain the placeholder words ([AREA], [CUISINE], *etc.*) as it is and not translate them.

We used Amazon Mechanical Turk (AMT) as the platform for crowdsourcing. Each Human Intelligence Task (HIT) required the user to give code-mixed translations of 5 utterances and was priced at \$0.2. Once we collected the code-mixed translations of all the utterance templates that were extracted using the procedure described in the previous subsection, we then instantiated them into proper sentences by replacing the placeholders ([AREA], [CUISINE], *etc.*) with the corresponding entities as present in the original DSTC2 dataset. For every dialog in the original DSTC2 dataset, every utterance was then replaced by its code-mixed translation resulting in an end-to-end code-mixed conversation.

		In-house Workers	Evaluators
Avg. Age		25.2	24.6
Gender	Female	33.3%	25.0%
	Male	66.7%	75.0%
Highest Education	Undergraduate	25.0%	33.3%
	Graduate	41.7%	33.3%
	Postgraduate	33.3%	33.3%
English Medium Schooling	Yes	100%	100%
	No	0%	0%
Frequency of English usage	Frequently	75.0%	91.7%
	Occasionally	25.0%	8.3%
	Rarely	0%	0%
Frequency of native language usage	Frequently	100%	91.7%
	Occasionally	0%	8.3%
	Rarely	0%	0%

Table 4: Demographic details of the in-house workers and the human evaluators.

4.3 Evaluating the Collected Dataset

We did evaluations at two levels. The first evaluation was at the level of utterances wherein if the code-mixed translation of an utterance was obtained via crowdsourcing, then we got this translation verified by in-house evaluators. The evaluators were asked to check if (i) the translation was faithful to the source sentence, (ii) the code-mixing was natural and not forced and (iii) all translations used Roman script and not the native language’s script. Any utterance which was flagged as erroneous by the evaluator was again crowdsourced and a new translation was solicited from AMT workers. If a worker’s utterances were flagged erroneous often then we barred him/her from doing any more tasks.

As mentioned in the previous section, once we collected such verified translations for all the utterance templates, we instantiated them and created complete end-to-end dialogs containing code-mixed utterances. Once the entire dialog was constructed, we conducted a separate human evaluation wherein we asked 12 in-house evaluators (3 evaluators per language) to read 100 code-mixed dialogs (entire dialogs as opposed to just some utterances) from each language and rate them on three metrics namely colloquialism, intelligibility and coherence on a scale of 1 (very poor) to 5 (very good) as defined below.

- **Colloquialism:** To check if the code-mixing was colloquial throughout the dialog and not forced.
- **Intelligibility:** To check if the entire dialog could be easily understood by a bilingual speaker who could speak the native language as well as English.
- **Coherence:** To check if the entire dialog looked coherent even though it was constructed by stitching together utterances which were independently translated and code-mixed (*i.e.*, while translating an utterance annotators did not know what their preceding and following utterances were).

These 100 dialogs were chosen randomly from across the entire dataset for each language. The evaluators used for this were different from the in-house annotators used to create the original translations in order to reduce the bias in evaluations. The average ratings given by the evaluators for each of the languages are shown in Table 5 and are encouraging. The demographic details of the in-house workers and evaluators are shown in Table 4.

4.4 Dataset Statistics and Analysis

For every word in the code-mixed corpus, we were able to identify whether it was a word from the native language or English or language agnostic (named entities). It was easy to do this because we had the vocabulary of the original English DSTC2 corpus as well as named entities (so any word which was not in the original DSTC2 vocabulary or a named entity was a word from the native language). We also manually verified this list of words marked as native words and corrected discrepancies if any (*i.e.*, we ensured that all the words which were marked as native words were actually native words). Note that the

Datasets	Colloquialism	Intelligibility	Coherent
Hi-DSTC2	4.20	4.06	4.21
Be-DSTC2	4.07	4.05	4.11
Gu-DSTC2	3.66	3.60	3.76
Ta-DSTC2	4.17	3.96	3.93

Table 5: Average human ratings for different metrics.

	Hindi	Bengali	Gujarati	Tamil
Vocabulary Size	1676	1372	1858	2185
Code-Mixed English Vocabulary	386	360	387	424
Native Language Vocabulary	739	477	912	1214
Others Vocabulary	551	535	559	547
Unique Utterances	6549	6274	6417	6666
Utterances with code-mixed words	5750	5703	5643	5632
Pure Native Language utterances	348	210	340	420
Pure English utterances	451	361	434	614
Average length of utterances	8.16	7.74	8.04	6.78
Average # of code-mixed utterances per dialog	12.11	14.28	11.80	12.96

Table 6: Statistics of the code-mixed dataset

cuisine names such as *Australian*, *Italian*, *etc.* have their own dedicated words in the native language. Table 6 summarizes various statistics about the dataset such as total vocabulary size, native language vocabulary size, *etc.* We refer to the original English dataset as En-DSTC2 and the Hindi, Bengali, Tamil and Gujarati code-mixed datasets created as a part of this work as Hi-DSTC2, Be-DSTC2, Ta-DSTC2 and Gu-DSTC2 respectively. Below, we make a few observations from Table 6.

The percentage of code-mixed English words in the vocabulary of Hi-DSTC2, Be-DSTC2, Gu-DSTC2 and Ta-DSTC2 are 23.03%, 26.24%, 20.83% and 19.40% respectively. From these English words, the most high frequency words across all the four versions of the dataset were *restaurant*, *food*, *town* and *serve*. Although these words have their own dedicated counterparts in all the other four languages, people colloquially use these code-mixed English words very often when talking about restaurants in their native language. The percentage of code-mixed utterances out of all the unique utterances in Hi-DSTC2, Be-DSTC2, Gu-DSTC2 and Ta-DSTC2 are 87.80%, 90.90%, 87.94% and 84.49% respectively (from Table 6). This shows that a significant portion of the dataset contains code-mixed utterances and very few utterances are in pure native languages or in pure English. This fact is also evident from the average number of code-mixed utterances per dialog in Table 6 compared to the average number of utterances per dialog in Table 3. We also calculated the number of utterances which contain k non-native words and then plotted a histogram (Figure 1) where the x-axis shows the number of k non-native words and the y-axis shows the number of utterances which had k non-native words. These histograms show a similar trend across all the languages. Apart from such intra-utterance code-mixing, we also noticed some intra-word code-mixing in mostly Bengali (*restauranter*, *townner*) and Tamil (*addressum*, *numberum*, *addressah*) versions of the dataset.

4.5 Quantitative Measures of Code-Mixing

Gambäck and Das (2016) introduced a measure to quantify the amount of code-mixing in a sentence as:

$$C_u(x) = \begin{cases} \frac{N(x) - \max_{L_i \in \mathcal{L}} \{t_{L_i}\}}{N(x)} & : N(x) > 0 \\ 0 & : N(x) = 0 \end{cases} \quad (1)$$

Here, \mathcal{L} is the set of all languages in the corpus, t_{L_i} is the number of tokens of language L_i in the given sentence x , $\max_{L_i \in \mathcal{L}} \{t_{L_i}\}$ is the maximum number of tokens of a language L_i in the sentence x and $N(x)$ is the number of language-specific tokens in the sentence ($N(x)$ does not include named entities as they are language agnostic). The authors make a crucial assumption that $\arg \max_{L_i \in \mathcal{L}} \{t_{L_i}\}$ is the *Matrix* language and hence the numerator of Equation 1 gives the number of foreign language tokens in x . This measure does not take into account the number of language switch points in a sentence (denoted by $P(x)$) and so the authors modify it further:

$$C_u(x) = 100 \cdot \frac{N(x) - \max_{L_i \in \mathcal{L}} \{t_{L_i}\} + P(x)}{2N(x)} : (if N(x) > 0) \quad (2)$$

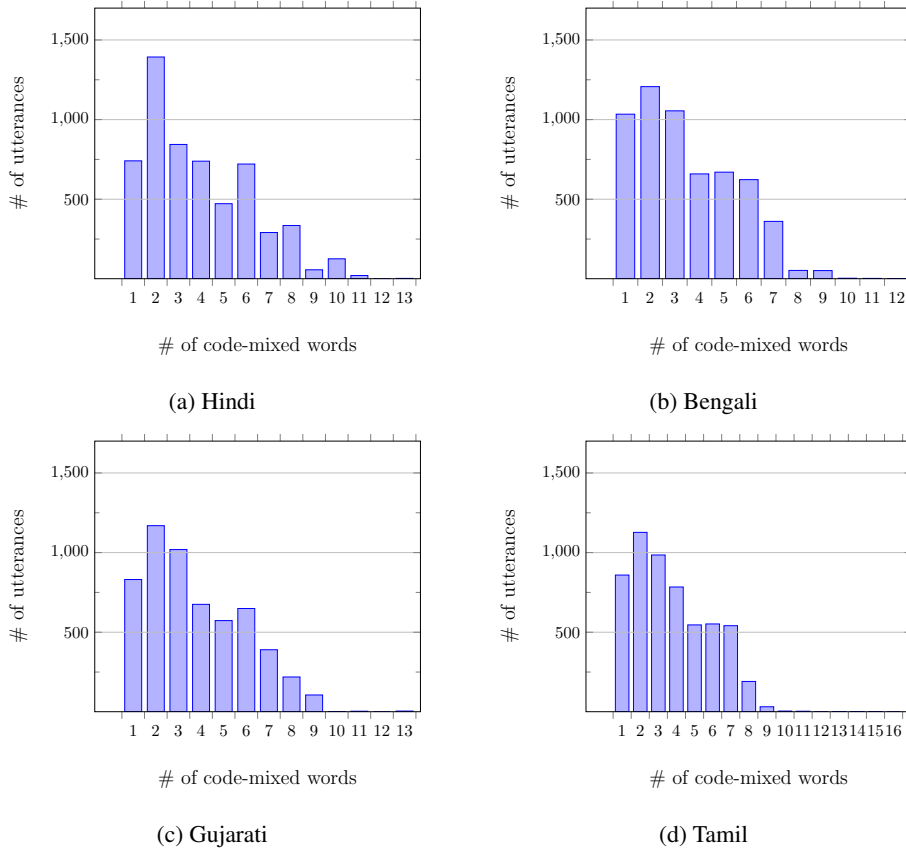


Figure 1: Histogram of the number of code mixed words in all the unique utterances for each language.

The code-mixing in the entire corpus can then be quantified by taking an average of the above measure across all sentences in the corpus:

$$C_{avg} = \frac{1}{U} \sum_{i=1}^U C_u(x_i) \quad (3)$$

where U is the number of sentences in the corpus. However, their main assumption that the language which has the maximum number of tokens in a sentence is the *Matrix* language, may not always hold. Consider a counter example: “Prezzo ek accha *restaurant* hain *in the north part of town* jo *tasty chinese food* serve karta hain.” Here the word ‘Prezzo’ is a named entity and hence treated as a language independent token. The most frequent language (*italicized*) is English but the *Matrix* language is essentially Hindi. So we propose a small modification to their measure and replace $\max_{L_i \in \mathcal{L}} \{t_{L_i}\}$ by the following:

$$native(x) = \begin{cases} t_{L_n} & : t_{L_n} > 0 \\ N(x) & : t_{L_n} = 0 \end{cases} \quad (4)$$

where L_n is the native (*Matrix*) language of the utterance and t_{L_n} is the number of tokens of the native language in the utterance. Note that we know the native (*Matrix*) language of every utterance beforehand because of the manner in which the dataset was created. Gambäck and Das (2016) also pointed out that C_{avg} does not take the inter-utterance code-mixing and frequency of code-mixed utterances into account. To overcome this they proposed to use a term $\delta(x_i)$ which assigns a score of 1 if the *Matrix* language of x_i is different from that of x_{i-1} or a score of 0 if they are same or $i = 1$. Note that in our case $\delta(x_i)$ would mostly be 0 except for cases where x_i is a pure English utterance. The authors also used a term for the fraction of code-mixed utterances ($\frac{S}{U}$) in the corpus, where S is the total number of code-mixed utterances. We use a modified version of their final Code-Mixing index² by replacing the maximum

²We refer the reader to Gambäck and Das (2016) for the detailed derivation.

Language-pair	En-Be	En-Hi	En-Hi	En-Hi	En-Hi	En-Hi	En-Be	En-Gu	En-Ta
Dataset	TW (Jamatia)	TW (Jamatia)	FB (Jamatia)	FB+TW (Jamatia)	Vyas	Hi-DSTC2	Be-DSTC2	Gu-DSTC2	Ta-DSTC2
I-index	-	-	-	-	-	0.04	0.04	0.03	0.03
C_{avg}	8.34	21.19	3.92	11.82	11.44	32.12	31.80	31.66	29.54
δ	22.09	30.99	6.70	17.81	53.50	26.38	29.06	24.50	38.32
C_c	25.14	64.38	16.76	38.53	31.31	73.31	76.27	71.63	80.49

Table 7: Comparison of the quantitative measures of code-mixing in the dataset.

function by $native(x)$:

$$C_c = \frac{100}{U} \left[\frac{1}{2} \sum_{i=1}^U \left(1 - \frac{native(x) + P(x)}{N(x)} + \delta(x) \right) + \frac{5}{6} S \right] \quad (5)$$

Similarly, Guzmán et al. (2016) introduced the I-index measure to quantify the integration of different languages in a corpus. This metric is much simpler and simply computes the number of switch points in the corpus. For example, if a corpus contains n words and there are k positions at which the language of $word_i$ is not the same as the language of $word_j$ then the I-index is given by $\frac{k}{n-1}$. We compute the I-index for every utterance in a dialog, then compute the average over all utterances in a dialog and finally report the average across all dialogs in the code-mixed corpus. These measures of our dataset are shown in Table 7 and are compared with that of the existing datasets (Jamatia et al., 2016; Vyas et al., 2014). Jamatia et al. (2016) collected the code-mixed text from Twitter (TW) and Facebook (FB) posts whereas Vyas et al. (2014) collected their dataset only from Facebook forums. Although the dataset of Vyas et al. (2014) show the highest inter-utterance code-mixing (δ), Hi-DSTC2 and Ta-DSTC2 show the highest level of overall code-mixing at the utterance level (C_{avg}) and the corpus level (C_c) respectively.

5 Baseline Models

We establish some initial baseline results on this code-mixed dataset by evaluating two different generation based models: (i) sequence-to-sequence with attention (Bahdanau et al., 2015) and (ii) Hierarchical Recurrent Encoder-Decoder (HRED) model (Serban et al., 2016). Due to lack of space we don’t describe these popular models here but refer the reader to the original papers. Apart from the above models, models which fetch the correct response from a set of candidate responses such as Query Reduction Networks (Seo et al., 2017), Memory Networks (Bordes and Weston, 2017) and Hybrid Code Networks (Williams et al., 2017) have also been evaluated on En-DSTC2. However, it is difficult to get candidate responses for every domain in practice and hence we stick to generation based models.

5.1 Experimental Setup

We use the train, validation and test splits of Bordes and Weston (2017) mentioned in Table 3. We create training instances from the dialogs by creating pairs of $\{context, response\}$ where $response$ is every even numbered utterance and $context$ contains all the previous utterances. Thus, if a dialog has 10 utterances, we create 5 training instances from it. Similarly at the test time the model is given the $context$ and it has to generate the $response$. For both the models, we used Adam optimizer (Kingma and Ba, 2015) to train the network with a mini batch size of 32. We used dropouts (Srivastava et al., 2014) of 0.25 and 0.35, initial learning rate of 0.0004 and Gated Recurrent Units (GRU) (Cho et al., 2014) with hidden dimensions of size 350. We used word embeddings of size 300 with Glorot initialization (Glorot and Bengio, 2010). We also clipped the gradients at a maximum norm of 10 to avoid exploding gradients.

5.2 Evaluation

We evaluate the performance of the above models using BLEU-4 (Papineni et al., 2002), ROUGE-1, ROUGE-2 and ROUGE-L (Lin, 2004) which are widely used to evaluate the performance of Natural Language Generation systems. We also compute the per utterance accuracy (exact match) by comparing

Metrics	SEQ2SEQ WITH ATTENTION					HRED				
	English	Hindi	Bengali	Gujarati	Tamil	English	Hindi	Bengali	Gujarati	Tamil
BLEU-4	56.6	54.0	56.8	53.8	62.1	57.8	54.1	56.7	54.1	60.7
ROUGE-1	67.2	62.9	67.4	64.7	67.8	67.9	63.3	67.1	65.3	67.1
ROUGE-2	55.9	52.4	57.5	54.8	56.3	57.5	52.6	56.9	55.2	55.6
ROUGE-L	64.8	61.0	65.1	62.6	65.6	65.7	61.5	64.8	63.2	65.1
Per response acc.	46.0	48.0	50.4	47.6	49.3	48.8	47.2	47.7	47.9	47.8
Per dialog acc.	1.4	1.2	1.5	1.5	1.3	1.4	1.5	1.6	1.6	1.0

Table 8: Performance of the baseline models on all the languages

the generated response with the ground truth response. The generated response is considered to be accurate only if it exactly matches the ground truth response. This is obviously a more strict metric for generation based models (Eric and Manning, 2017). We also compute the per dialog accuracy by matching all the generated responses in a dialog with all the ground truth responses for that dialog. This metric measures whether the model was able to produce the entire dialog correctly end-to-end and hence complete the goal. We summarize the performance of the two models in Table 8. We observe that the performance of these models is very similar across all the languages. We observe that the models are still far from 100% accuracy and there is clearly scope for further improvement.

6 Conclusion

Code-mixing is an emerging trend of communication in the multilingual regions. The community has already addressed this phenomenon by introducing challenges on POS-Tagging, Language Identification, Language Modeling, *etc* on the code-mixed corpora. However, the approaches to development of dialog systems still rely on monolingual conversation datasets. To alleviate this problem we introduced a goal-oriented code-mixed dialog dataset for four languages (Hindi-English, Bengali-English, Gujarati-English and Tamil-English respectively). The dataset was created using a mix of in-house and crowdsourced workers. All the utterances in the dataset were evaluated by in-house evaluators and the overall dialogs were also evaluated for colloquialism, intelligibility and coherence. On all these measures, the dialogs in our dataset received a high score. To facilitate further research on these datasets, we provide the implementation of two popular neural dialog models *viz.* sequence-to-sequence and HRED. The evaluation of these models suggest that there is a clear scope for development of new architectures which can understand and converse in code-mixed languages.

Acknowledgements

We would like to thank Accenture Technology Labs, India for supporting this work through their generous academic research grant.

References

- Heike Adel, Ngoc Thang Vu, Franziska Kraus, Tim Schlippe, Haizhou Li, and Tanja Schultz. 2013a. Recurrent neural network language modeling for code switching conversational speech. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2013, Vancouver, BC, Canada, May 26-31, 2013*, pages 8411–8415.
- Heike Adel, Ngoc Thang Vu, and Tanja Schultz. 2013b. Combination of recurrent neural networks and factored language models for code-switching language modeling. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, ACL 2013, 4-9 August 2013, Sofia, Bulgaria, Volume 2: Short Papers*, pages 206–211.
- Heike Adel, Ngoc Thang Vu, Katrin Kirchhoff, Dominic Telaar, and Tanja Schultz. 2015. Syntactic and semantic features for code-switching factored language models. *Trans. Audio, Speech and Lang. Proc.*, 23(3):431–440, March.

- Fahad AlGhamdi, Giovanni Molina, Mona Diab, Thamar Solorio, Abdelati Hawwari, Victor Soto, and Julia Hirschberg. 2016. Part of speech tagging for code switched data. In *Proceedings of the Second Workshop on Computational Approaches to Code Switching*, pages 98–107. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *International Conference on Learning Representations*.
- Rafael E. Banchs. 2012. Movie-dic: A movie dialogue corpus for research and development. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers - Volume 2*, ACL '12, pages 203–207, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Utsab Barman, Amitava Das, Joachim Wagner, and Jennifer Foster. 2014. Code mixing: A challenge for language identification in the language of social media. In *Proceedings of the First Workshop on Computational Approaches to Code Switching@EMNLP 2014, Doha, Qatar, October 25, 2014*, pages 13–23.
- Utsab Barman, Joachim Wagner, and Jennifer Foster. 2016. Part-of-speech tagging of code-mixed social media content: Pipeline, stacking and joint modelling. In *Proceedings of the Second Workshop on Computational Approaches to Code Switching*, pages 30–39. Association for Computational Linguistics.
- Antoine Bordes and Jason Weston. 2017. Learning end-to-end goal-oriented dialog. *International Conference on Learning Representations*.
- Kunal Chakma and Amitava Das. 2016. CMIR: A corpus for evaluation of code mixed information retrieval of hindi-english tweets. *Computación y Sistemas*, 20(3):425–434.
- Tao Chen and Min-Yen Kan. 2013. Creating a live, public short message service corpus: the NUS SMS corpus. *Language Resources and Evaluation*, 47(2):299–335.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734. Association for Computational Linguistics.
- Mihail Eric and Christopher Manning. 2017. A copy-augmented sequence-to-sequence architecture gives good performance on task-oriented dialogue. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 468–473. Association for Computational Linguistics.
- Eric N. Forsyth and Craig H. Martell. 2007. Lexical and discourse analysis of online chat dialog. In *Proceedings of the First IEEE International Conference on Semantic Computing (ICSC 2007), September 17-19, 2007, Irvine, California, USA*, pages 19–26.
- Björn Gambäck and Amitava Das. 2016. Comparing the level of code-switching in corpora. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portorož, Slovenia, May 23-28, 2016*.
- Souvick Ghosh, Satanu Ghosh, and Dipankar Das. 2016. Part-of-speech tagging of code-mixed social media text. In *Proceedings of the Second Workshop on Computational Approaches to Code Switching*, pages 90–97. Association for Computational Linguistics.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010*, pages 249–256.
- John J. Godfrey, Edward C. Holliman, and Jane McDaniel. 1992. Switchboard: Telephone speech corpus for research and development. In *Proceedings of the 1992 IEEE International Conference on Acoustics, Speech and Signal Processing - Volume 1, ICASSP'92*, pages 517–520, Washington, DC, USA. IEEE Computer Society.
- Gualberto A. Guzmán, Jacqueline Serigos, Barbara E. Bullock, and Almeida Jacqueline Toribio. 2016. Simple tools for exploring variation in code-switching for linguists. In *Proceedings of the Second Workshop on Computational Approaches to Code Switching@EMNLP 2016, Austin, Texas, USA, November 1, 2016*, pages 12–20.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014a. The second dialog state tracking challenge. In *Proceedings of the SIGDIAL 2014 Conference, The 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 18-20 June 2014, Philadelphia, PA, USA*, pages 263–272.

- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014b. The third dialog state tracking challenge. In *2014 IEEE Spoken Language Technology Workshop, SLT 2014, South Lake Tahoe, NV, USA, December 7-10, 2014*, pages 324–329.
- Anupam Jamatia, Björn Gambäck, and Amitava Das. 2016. Collecting and annotating indian social media code-mixed corpora. In *Computational Linguistics and Intelligent Text Processing - 17th International Conference, CICLing 2016, Konya, Turkey, April 3-9, 2016, Revised Selected Papers, Part II*, pages 406–417.
- Mitesh M. Khapra, Salil Joshi, Ananthakrishnan Ramanathan, and Karthik Visweswariah. 2013. Offering language based services on social media by identifying user’s preferred language(s) from romanized text. In *22nd International World Wide Web Conference, WWW ’13, Rio de Janeiro, Brazil, May 13-17, 2013, Companion Volume*, pages 71–72.
- Seokhwan Kim, Luis Fernando D’Haro, Rafael E. Banchs, Jason D. Williams, Matthew Henderson, and Koichiro Yoshino. 2016. The fifth dialog state tracking challenge. In *2016 IEEE Spoken Language Technology Workshop, SLT 2016, San Diego, CA, USA, December 13-16, 2016*, pages 511–517.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *International Conference on Learning Representations*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 110–119.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios P. Spithourakis, Jianfeng Gao, and William B. Dolan. 2016b. A persona-based neural conversation model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*.
- Pierre Lison and Jörg Tiedemann. 2016. Opensubtitles2016: Extracting large parallel corpora from movie and TV subtitles. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portorož, Slovenia, May 23-28, 2016*.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. In *Proceedings of the SIGDIAL 2015 Conference, The 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 2-4 September 2015, Prague, Czech Republic*, pages 285–294.
- Giovanni Molina, Fahad AlGhamdi, Mahmoud Ghoneim, Abdelati Hawwari, Nicolas Rey-Villamizar, Mona Diab, and Thamar Solorio. 2016. Overview for the second shared task on language identification in code-switched data. In *Proceedings of the Second Workshop on Computational Approaches to Code Switching*, pages 40–49. Association for Computational Linguistics.
- Carol Myers-Scotton. 1993. *Duelling languages: Grammatical structure in codeswitching*. Oxford University Press.
- Dong Nguyen and A. Seza Dogruöz. 2013. Word level language identification in online multilingual communication. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 857–862.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA.*, pages 311–318.
- Alan Ritter, Colin Cherry, and Bill Dolan. 2010. Unsupervised modeling of twitter conversations. In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 2-4, 2010, Los Angeles, California, USA*, pages 172–180.
- Mike Rosner and Paulseph-John Farrugia. 2007. A tagging algorithm for mixed language identification in a noisy domain. In *INTERSPEECH 2007, 8th Annual Conference of the International Speech Communication Association, Antwerp, Belgium, August 27-31, 2007*, pages 190–193.

- Sophie Rosset and Sandra Petel. 2006. The ritel corpus - an annotated human-machine open-domain question answering spoken dialog corpus. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation, LREC 2006, Genoa, Italy, May 22-28, 2006.*, pages 1640–1643.
- Anindya Roy, Camille Guinaudeau, Hervé Bredin, and Claude Barras. 2014. Tvd: A reproducible and multiply aligned tv series dataset. In *LREC*, pages 418–425.
- Minjoon Seo, Sewon Min, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Query-reduction networks for question answering. *International Conference on Learning Representations*.
- Iulian Vlad Serban, Ryan Lowe, Peter Henderson, Laurent Charlin, and Joelle Pineau. 2015. A survey of available corpora for building data-driven dialogue systems. *CoRR*, abs/1512.05742.
- Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 3776–3784.
- Iulian Vlad Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron C. Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3295–3301.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. Neural responding machine for short-text conversation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 1577–1586.
- Tamar Solorio, Elizabeth Blair, Suraj Maharjan, Steven Bethard, Mona Diab, Mahmoud Ghoneim, Abdelati Hawwari, Fahad AlGhamdi, Julia Hirschberg, Alison Chang, and Pascale Fung. 2014. Overview for the first shared task on language identification in code-switched data. In *Proceedings of the First Workshop on Computational Approaches to Code Switching*, pages 62–72, Doha, Qatar, October. Association for Computational Linguistics.
- Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.
- Oriol Vinyals and Quoc V. Le. 2015. A neural conversational model. In *ICML Deep Learning Workshop*.
- Yogarshi Vyas, Spandana Gella, Jatin Sharma, Kalika Bali, and Monojit Choudhury. 2014. Pos tagging of english-hindi code-mixed social media content. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 974–979. Association for Computational Linguistics.
- Marilyn A. Walker, Jean E. Fox Tree, Pranav Anand, Rob Abbott, and Joseph King. 2012. A corpus for research on deliberation and debate. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC 2012, Istanbul, Turkey, May 23-25, 2012*, pages 812–817.
- Jason D. Williams and Steve J. Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422.
- Jason D. Williams, Antoine Raux, Deepak Ramachandran, and Alan W. Black. 2013. The dialog state tracking challenge. In *Proceedings of the SIGDIAL 2013 Conference, The 14th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 22-24 August 2013, SUPELEC, Metz, France*, pages 404–413.
- Jason D. Williams, Kavosh Asadi, and Geoffrey Zweig. 2017. Hybrid code networks: practical and efficient end-to-end dialog control with supervised and reinforcement learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 665–677.
- Kaisheng Yao, Geoffrey Zweig, and Baolin Peng. 2015. Attention with intention for a neural network conversation model. *CoRR*, abs/1510.08565.
- Steve J. Young. 2000. Probabilistic methods in spoken-dialogue systems. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 358(1769):1389–1402.

Appendix A. Instructions to Crowdsourced Workers

Instructions

Attempt this HIT only if you know HINDI.

Please Note that there are some changes in Instructions.

We have given you 5 sentences in ENGLISH. Your job is to

- Convert the sentences into HINDI by translating only the most common words in the sentence to HINDI words, using **only English alphabets(a,b,c,d,...)**. (See the examples below)
- The words in the sentence which are **difficult to translate** (eg. Restaurant, table, sorry) should be kept as it is in ENGLISH.
- The goal is to **NOT totally translate** the ENGLISH sentence to HINDI, but **translate partially** while keeping the difficult-to-translate words in ENGLISH. (The way in which people generally chat over instant messengers.) This is called **CODE-MIXING**.
- The sentences will sometimes have some **placeholders** denoted by **square brackets([,])** and **capital letters**. While translating you need to **keep the placeholders as it is**.
- All possible **placeholders** and their possible values are:
 - [CUISINE] = chinese, spanish, vietnamese, japanese, british,...
 - [RESTAURANT] = the_golden_curry, the_hotpot, maharajah_tandoori_restaurant,...
 - [ADDRESS] = the_golden_curry_address, the_hotpot_address, maharajah_tandoori_restaurant_address,.....
 - [PHONE] = the_golden_curry_phone, the_hotpot_phone, maharajah_tandoori_restaurant_phone,.....
 - [LOCATION] = north, south, east, west, centre
 - [PRICE]= cheap, expensive, moderate...
- A word with underscores("the_golden_curry") is a single word and you need to keep it as it is.(SEE EXAMPLE NUMBER 5)
- For the **ease of translation**, you can **imagine a placeholder to be its value**. But you have to **keep the original placeholder in the translated sentence**. (SEE EXAMPLES 2 and 4 BELOW)
- You need to convert the sentences and write them in the boxes below for each given sentence.
- **DO NOT USE HINDI ALPHABETS, USE ONLY ENGLISH ALPHABETS FOR THE TRANSLATIONS.** (Examples: main, kaun, pasand, e.t.c)

Example 1:

Given ENGLISH sentence: I am sorry but there is no other spanish restaurant in the moderate price range.

HINDI sentence: Sorry, but moderate price range mein aur koi spanish restaurant nahi hain.

Example 2:

While doing such a sentence you can **imagine** [PRICE] to be 'moderate' and [CUISINE] to be 'spanish' , but while writing the HINDI sentence you have to **keep** [PRICE], [CUISINE] as it is.

Given ENGLISH sentence with placeholders: I am sorry but there is no other [CUISINE] restaurant in the [PRICE] price range.

HINDI sentence: Sorry, but [PRICE] price range mein aur koi [CUISINE] restaurant nahi hain.

Example 3:

Given ENGLISH sentence: Let me confirm, You are looking for a restaurant in the expensive price range right ?

HINDI sentence: Ek bar confirm kar leta hu, aap expensive price range mein ek restuarant dhund rahe hain, right ?

Example 4:

While doing such a sentence you can **imagine** [PRICE] to be 'cheap' and [LOCATION] to be 'north' , but while writing the HINDI sentence you have to **keep** [PRICE], [LOCATION] as it is.

Given ENGLISH sentence with placeholders: i want [PRICE] food in the [LOCATION] part of town.

HINDI sentence: Mujhko [LOCATION] part of town mein [PRICE] food chahiye.

Example 5:

While doing such a sentence you can **imagine** [RESTAURANT] to be "the_golden_curry" , but while writing the HINDI sentence you have to **keep** [RESTAURANT] as it is. You also need to keep "the_golden_curry_post_code" as a single word.

Given ENGLISH sentence with placeholders: The post code of [RESTAURANT] is the_golden_curry_post_code.

HINDI sentence: [RESTAURANT] ka post code hain the_golden_curry_post_code.

Figure 2: Screenshot of the instructions and examples provided to the crowdsourced workers

Appendix B. Sample Code-Mixed Chats from the Dataset

User	Bot
1. <SILENCE>	Hello , Cambridge restaurant system mein aapka swagat hai. Aap pa sakte hai restaurant by area , price range ya food type. Aap ko kaise help kar sakta hu main ?
2. main north part of town mein cheap restaurant dhund raha hoon	Aap kaunse type ka food lena chahege?
3. mujhe parvah nahi hain.	api_call R_cuisine north cheap da_vinci_pizzeria R_post_code da_vinci_pizzeria_post_code da_vinci_pizzeria R_cuisine italian da_vinci_pizzeria R_location north da_vinci_pizzeria R_phone da_vinci_pizzeria_phone da_vinci_pizzeria R_address da_vinci_pizzeria_address da_vinci_pizzeria R_price cheap da_vinci_pizzeria R_rating 10 royal_spice R_post_code royal_spice_post_code royal_spice R_cuisine indian royal_spice R_location north royal_spice R_phone royal_spice_phone royal_spice R_address royal_spice_address royal_spice R_price cheap royal_spice R_rating 8
5. <SILENCE>	north of town me da_vinci_pizzeria ek achaa jagah hai aur prices cheap hai
6. aur kuch	north of town me royal_spice ek achaa jagah hai aur prices cheap hai
7. kyaa mujhe phone number mil sakta hai	royal_spice ka phone number hai royal_spice_phone
8. thank you goodbye	Welcome , alvida

Table 9: An example chat from the Hindi version of DSTC2 dataset

User	Bot
1. <SILENCE>	Hello, Cambridge restaurant system ey apnake swagato janai. Apni restaurants khujte paren food type, price ar area diye. Apnake ki bhabe help korte pari ami ?
2. Ami ekta cheap restaurant khujchi town er north part e	Kon type er khabar apni pochondo korben ?
3. ami care kori na	api_call R_cuisine north cheap da_vinci_pizzeria R_post_code da_vinci_pizzeria_post_code da_vinci_pizzeria R_cuisine italian da_vinci_pizzeria R_location north da_vinci_pizzeria R_phone da_vinci_pizzeria_phone da_vinci_pizzeria R_address da_vinci_pizzeria_address da_vinci_pizzeria R_price cheap da_vinci_pizzeria R_rating 10 royal_spice R_post_code royal_spice_post_code royal_spice R_cuisine indian royal_spice R_location north royal_spice R_phone royal_spice_phone royal_spice R_address royal_spice_address royal_spice R_price cheap royal_spice R_rating 8
5. <SILENCE>	da_vinci_pizzeria ekta nice jayega town er north ey jekhane prices cheap
6. ar onno kichu ache	royal_spice ekta nice jayega town er north ey jekhane prices cheap
7. ami ki phone number ta pete pari	royal_spice er phone number holo royal_spice_phone
8. thank you goodbye	you are welcome

Table 10: An example chat from the Bengali version of DSTC2 dataset