

MULTI³WOZ: A Multilingual, Multi-Domain, Multi-Parallel Dataset for Training and Evaluating Culturally Adapted Task-Oriented Dialog Systems

Songbo Hu^{1*} Han Zhou^{1*} Mete Hergul¹
Milan Gritta² Guchun Zhang² Ignacio Iacobacci²
Ivan Vulić^{1†} Anna Korhonen^{1†}

¹Language Technology Lab, University of Cambridge, UK

²Huawei Noah's Ark Lab, London, UK

¹{sh2091, hz416, mh2071, iv250, alk23}@cam.ac.uk

²{milan.gritta, guchun.zhang, ignacio.iacobacci}@huawei.com

Abstract

Creating high-quality annotated data for task-oriented dialog (ToD) is known to be notoriously difficult, and the challenges are amplified when the goal is to create equitable, culturally adapted, and large-scale ToD datasets for multiple languages. Therefore, the current datasets are still very scarce and suffer from limitations such as translation-based non-native dialogs with translation artifacts, small scale, or lack of cultural adaptation, among others. In this work, we first take stock of the current landscape of multilingual ToD datasets, offering a systematic overview of their properties and limitations. Aiming to reduce all the detected limitations, we then introduce **MULTI³WOZ**, a novel multilingual, multi-domain, multi-parallel ToD dataset. It is large-scale and offers culturally adapted dialogs in 4 languages to enable training and evaluation of multilingual and cross-lingual ToD systems. We describe a complex bottom-up data collection process that yielded the final dataset, and offer the first sets of baseline scores across different ToD-related tasks for future reference, also highlighting its challenging nature.

1 Introduction and Motivation

Task-oriented dialog (ToD), where a human user engages in a conversation with a system agent with the aim of completing a concrete task, is one of the central objectives, hallmarks, and applications of machine intelligence (Gupta et al., 2006; Tür et al., 2010; Young, 2010, *inter alia*). ToD technology has been proven useful across a wide

spectrum of application sectors such as hospitality industry (Henderson et al., 2014, 2019), healthcare (Laranjo et al., 2018), online shopping (Yan et al., 2017), banking (Altinok, 2018), and travel (Raux et al., 2005; El Asri et al., 2017), among others.

Wider developments in ToD have been hampered by the two conflicting requirements: **1)** large-scale in-domain datasets are crucially required in order to unlock the potential of deep learning-based ToD components and systems to handle complex dialog patterns (Budzianowski et al., 2018; Lin et al., 2021b); at the same time, **2)** data collection for ToD is known to be notoriously difficult as it is extremely time-consuming, expensive, and requires expert and domain knowledge (Shah et al., 2018; Larson and Leach, 2022). Put simply, the creation of ToD datasets for new domains and languages incurs significantly higher time and budget costs than for most other NLP tasks (Casanueva et al., 2022). Consequently, the progress in ToD until recently has been limited only to a small number of high-resource languages such as English and Chinese (Razumovskaia et al., 2022).

Recent work has recognized the need to expand the reach of multilingual ToD technology to more languages via collecting multilingual ToD data (Razumovskaia et al., 2022). Yet, as discussed in more detail later in §2, all the currently available multilingual ToD datasets suffer from one or several serious limitations: (i) the predominant reliance on translation-based data creation that introduces issues with ‘translation-ese’ and artificial performance inflation (Xu et al., 2020; Zuo et al., 2021); (ii) lack of *cultural*

*Equal contribution.

†Equal senior contribution.

Dataset (Reference)	# Langs	# Domains	# Train	# Test	No Translation?	Culturally Adapted?	Coherent?	Multi-P?
WOZ 2.0 (Mrkšić et al., 2017)	3	1	600	400	✗	✗	✓	✓
BiToD (Lin et al., 2021b)	2	5	2,894	451	✓	✓	✓	✗
AllWOZ (Zuo et al., 2021)	8	5	40	50	✗	✗	✓	✓
GlobalWOZ (Ding et al., 2022)	21	7	0 (8,437)*	500 (1,000)*	✗	✓	✗	✗
Multi ² WOZ (Hung et al., 2022)	5	7	0	1,000	✗	✗	✓	✓
MULTI³WOZ (this work)	4	7	7,440	860	✓	✓	✓	✓

Table 1: Summary of multilingual ToD datasets that support multiple languages and ToD tasks (including E2E learning), with more details concerning each dimension of comparison available in §2. For clarity, we do not show (i) monolingual ToD datasets constructed for languages other than English, we refer the reader to the survey of Razumovskaia et al. (2022) for a comprehensive overview; as well as (ii) the body of multilingual ToD datasets that focus solely on NLU for ToD (see §2). **# Langs** refers to the total number of languages in each dataset, including English. **# Train** and **# Test** refer to the average number of human-created or human-curated dialogs *per each language* in the respective portions of each dataset. **Multi-P** refers to multi-parallelism of dialogs in the dataset. (*) GlobalWOZ releases training data created automatically by an English-target NMT system, without any human curation nor post-processing, and manually curates only a portion of 500 dialogs from the target language test sets (see §2 for more details).

adaptation also results in artificial dialogs that are not localized nor adapted to real-world data and to cultural specificities of each target language and culture; (iii) small scale and lack of sufficient training data prevents truly equitable multilingual development and in-depth comparative cross-language analyses (Ding et al., 2022; Hung et al., 2022); (iv) lack of coherent and multi-parallel dialogs in all the represented languages, which are typically not created and corrected by native speakers, hinders meaningful cross-language comparisons and analyses (Ding et al., 2022); (v) some datasets focus on a single component of a full ToD system, typically Natural Language Understanding (NLU), which prevents training and evaluation of other crucial tasks such as Dialog State Tracking (DST), or Natural Language Generation (NLG) in multilingual and transfer setups.

In this work, we address all the aforementioned limitations of current multilingual ToD datasets and present a large-scale data collection process that resulted in a novel large-scale multilingual dataset for ToD: **MULTI³WOZ**. The departure point of our data collection is the established *multi-domain* English MultiWOZ dataset (Budzianowski et al., 2018), that is, its cleaned version 2.3 in particular (Han et al., 2021). MULTI³WOZ is then created via adapting a recent *bottom-up outline-based* approach of Majewska et al. (2023) which bypasses (the issues of) the translation-based design and discerns between

language-agnostic *abstract dialog schemata* (i.e., *outlines*) and adapted, language-specific *surface realizations* of the underlying schemata (i.e., the actual user and system *utterances*). We validate the usefulness and feasibility of the outline-based approach to multilingual ToD data creation for the first time on a large scale, and prove its feasibility for such large-scale endeavors: the dataset contains a total of 494,116 dialog turns created manually by human subjects.

Guided by the need to tackle the present limitations, MULTI³WOZ is the first multilingual ToD dataset with the following crucial properties; see also Table 1 for an overview. First, MULTI³WOZ is *large-scale* with the equal number of training (7,440 dialogs per language), development (860), and test dialogs (860) offered in 4 different languages: English, Arabic, French, and Turkish. It is more versatile than all prior multilingual ToD datasets as it allows for training and evaluation in monolingual, multilingual, and cross-lingual setups, and in zero-shot, few-shot, and ‘many’-shot *cross-lingual* and *cross-domain* transfer scenarios. Second, MULTI³WOZ offers *multi-parallel* dialogs, conveying comparable information over exactly the same conversational flows across all four languages. This property allows for cross-language studies and comparative analyses. Third, MULTI³WOZ enables *both* (monolingual and multilingual) training and evaluation over different constituent ToD tasks such as NLU (intent detection and slot filling), DST, NLG, as

well as full-fledged end-to-end (E2E) learning. Fourth, MULTI³WOZ is *localized* and *culturally adapted* to the actual existing entities from the cultures in which the target languages are spoken. Finally, created in a bottom-up fashion by native speakers of the target languages, hence *linguistically adapted* to the target language, it offers natural and native dialogs in all target languages, avoiding ‘translationese’ and preventing over-inflation of transfer performance (Majewska et al., 2023).

Furthermore, to guide future research, we set reference scores across different ToD tasks in all the languages of MULTI³WOZ, running a representative set of standard baselines in each relevant ToD task. The results clearly indicate the challenging nature of the dataset; we also outline the differences in performance across different languages.

2 MULTI³WOZ versus Limitations of Current Multilingual ToD Datasets

We now delve deeper into the main benefits of MULTI³WOZ, characterizing how its key properties make it a unique ToD resource. The summary and statistics of the most relevant prior work are provided in Table 1. Building upon this table, we discuss those datasets along with other related work in what follows, focusing on the five desirable properties of MULTI³WOZ and how these counteract the detected main limitations of other datasets.

P1. Supporting Multiple Languages and ToD Tasks. There has been a surge of interest in the creation of multilingual ToD datasets, aiming to mitigate the language resource gap in multilingual NLP (Ponti et al., 2019; Joshi et al., 2020b). Despite the effort, the gap is still much more pronounced for dialog tasks and data than for some other NLP tasks such as NLI (Conneau et al., 2018; Ebrahimi et al., 2022) or NER (Adelani et al., 2021), also due to its increased time demands and cost of annotation.¹ Further, the ma-

¹For instance, the creation of the validation and test sets of the XCOPA dataset requires a total time ranging from 12 to 20 hours per language (Ponti et al., 2020). In contrast, the creation of the validation and testing sets for each individual language in MULTI³WOZ requires over 300 hours of effort. Even when considering the annotation cost per sentence (utterance), which amounts to approximately

majority of multilingual ToD datasets focused only on two standard NLU tasks (i.e., intent detection and slot labeling), again due to the high cost and specific challenges posed by collecting full dialog data (Budzianowski et al., 2018). The first wave of such NLU datasets were built upon the single-domain English ATIS dataset (Hemphill et al., 1990), extending it to 10 languages via human translation (Upadhyay et al., 2018; Xu et al., 2020; Dao et al., 2021). More recent NLU datasets cover multiple domains and wider linguistic typology and geography (Schuster et al., 2019; FitzGerald et al., 2022; Moghe et al., 2023; Majewska et al., 2023). However, current NLU datasets (i) still support only the two NLU tasks, and (ii) provide utterances ‘in isolation’ (i.e., out of the context of the full dialog which facilitates their multilingual construction). Further, (iii) some datasets do not provide any training data and are useful only for evaluation of (zero-shot) cross-lingual transfer; (iv) all the datasets except that of Majewska et al. (2023) and the concurrent work of Goel et al. (2023) were constructed via translation from the source English datasets.

Monolingual ‘end-to-end’ ToD datasets, which support NLU as well as other ToD tasks (i.e., modeling and evaluation of the full ToD pipeline), have been created only for particular high-resource languages. MultiWOZ (Budzianowski et al., 2018) and Taskmaster (Byrne et al., 2019) are two large-scale multi-domain English datasets spanning 7 and 6 domains, respectively, containing both single-domain and multi-domain dialogs. Inspired by MultiWOZ, monolingual RisaWOZ (Quan et al., 2020) and CrossWOZ (Zhu et al., 2020) datasets have been created for Chinese. Crucially, multilingual multi-domain ToD datasets that support full ToD modeling are still scarce, see Table 1, and they all come with some core limitations, as discussed next.

P2. Avoiding Translation-Based Design. The majority of datasets have been obtained via *manual or semi-automatic translation* (e.g., via post-editing MT output [PEMT]) of an English source dataset (Zuo et al., 2021; Ding et al., 2022; Hung et al., 2022). The translation-based approach is cost-efficient and can natively yield data which

\$0.17 per utterance, the cost is notably higher than the per sentence annotation cost for NER (\$0.06 as reported by Bontcheva et al. [2017]) and NLI (\$0.01015 per instance as reported by Marelli et al. [2014]).

is comparable across languages, but results in (i) undesired ‘translationese’ effects (Artetxe et al., 2020), (ii) lacks dialog naturalness (Ding et al., 2022), and (iii) typically leads to overinflated and thus misleading performance of ToD systems. For instance, Majewska et al. (2023) empirically validate that cross-lingual transfer performance substantially increases when exactly the same dialogs are obtained via automatic or manual translation rather than via a bottom-up approach relying on native speakers of the target languages.

Unlike prior work (i.e., all datasets from Table 1 except BiToD), the honed outline-based construction of MULTI³WOZ (see §3 later) avoids all the negative implications of translation, while maintaining cost efficiency (and thus enabling its large scale), supporting cultural adaptation, and enabling coherence and multi-parallelism.

P3. Dataset Scale and Large-Scale Training.

MULTI³WOZ offers a substantially larger number of dialogs for training than any previous multilingual ‘full ToD’ dataset, and it treats the four supported languages in an equitable way: i.e., it provides the same set of manually (bottom-up) constructed dialogs for training, development, and testing in each language. Previous work (Multi²WOZ, AllWOZ, GlobalWOZ) targeted the creation of test data only, for evaluating cross-lingual transfer scenarios. These datasets come (i) without providing any training data at all (Multi²WOZ), or (ii) with a very small set of post edited MT-obtained dialogs (AllWOZ),² or (iii) with automatically created MT-based training data only (GlobalWOZ). The only exception is BiToD (Lin et al., 2021b), but it spans only two, highest-resourced languages, a smaller number of domains, and has approximately three times fewer training data than MULTI³WOZ. For instance, MULTI³WOZ contains almost 124,000 turns *per each represented language* (~98,000/12,500/12,500), with a total of 494,116 turns; for comparison, the *total* number of turns in BiToD is 115,638, while it is 143,048 in the original English-only MultiWOZ.

P4. (Improved) Cultural Adaptation. A large number of datasets for multilingual NLP ignores the fact that the data should also be adapted to

²The tiny size of AllWOZ is even more problematic at the level of single domains, e.g., it contains only 13 dialogs for the *Taxi* domain, hindering any generalizable evaluations.

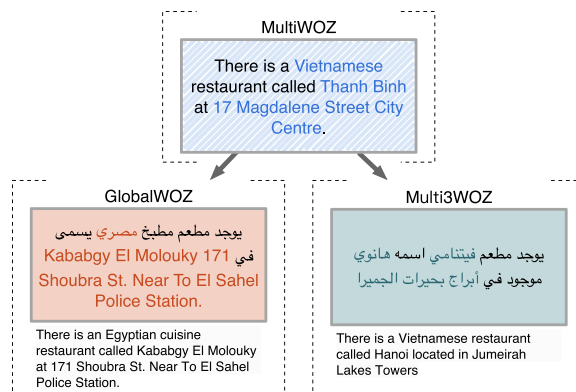


Figure 1: An example of dialog turns from culturally adapted GlobalWOZ versus MULTI³WOZ, with culturally specific entities highlighted and English translations provided below each text box. In general, due to its design, a proportion of GlobalWOZ dialogs exhibit inconsistent similar code-switched and script-switched utterances (e.g., also with phone and reference numbers); GlobalWOZ comes with other design-triggered dialog-level inconsistencies, not shown for brevity.

the target cultures and concepts (Ponti et al., 2020; Hershovich et al., 2022). Besides (i) propagating the source language bias towards possible *conversational concepts* (e.g., the US-tied concept of *tailgating* or conversations about *baseball*) (Ponti et al., 2020), the lack of the so-called *cultural adaptation* also (ii) creates peculiar or more unlikely *conversational contexts* (e.g., a user speaking to a Turkish ToD system about restaurants in Cambridge) (Ding et al., 2022), or (iii) even ignores specificities of a particular culture (e.g., postcodes are not used in Arabic-speaking countries). The only two datasets that try to incorporate the notion of cultural adaptation into their design are BiToD and GlobalWOZ (see Table 1). However, BiToD’s adaptation is based on a very specific bilingual region of the world (Hong Kong), while GlobalWOZ’s automatic cultural adaptation approach results in a large number of incoherent dialogs and annotation errors, e.g., see Figure 1. We thus adopt a new and improved cultural adaptation approach that ensures high-quality, coherent and multi-parallel dialogs across languages while respecting the underlying cultural traits, see §3 later.

P5. Dialog Coherence and ‘Multi-Parallelism’.

Finally, due to their design properties and oversimplifying assumptions, some datasets break

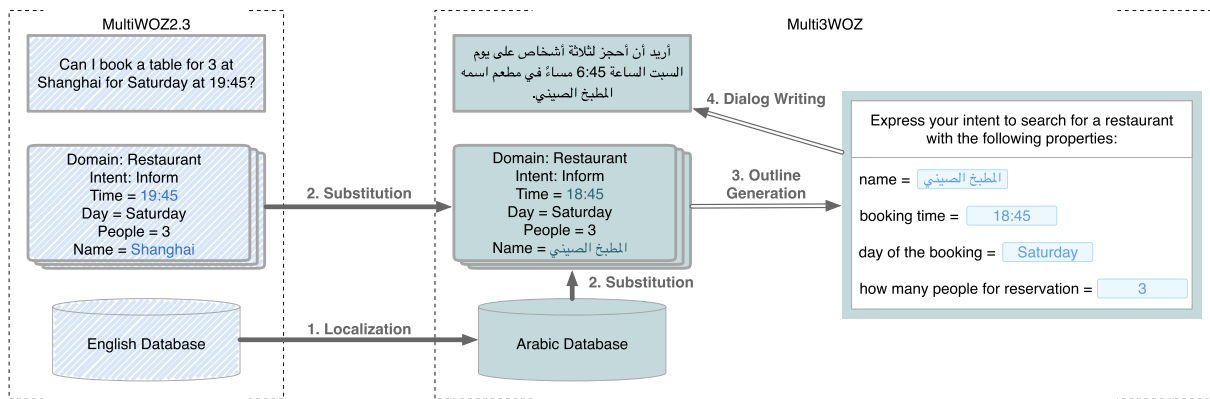


Figure 2: Overview of the full data collection pipeline for **MULTI³WOZ**. It is derived from the **MultiWOZ** dataset v2.3, with two phases: (i) *cultural adaptation* and (ii) *outline-based generation*. Cultural adaptation \rightarrow spans two subtasks *localization* and *value substitution*, and it adapts dialogs and contextualizes them to the actual existing entities from the cultures in which the target languages are spoken. Outline-based generation \Rightarrow is a bottom-up dialog collection method to collect language-specific and linguistically adapted surface forms from the target language native speakers based on language-agnostic abstract dialog schemata. In both datasets, each utterance is annotated with task-specific meaning representations. In the above figure, a rectangle \square denotes an utterance and stacked rectangles \square denote its corresponding dialog act. Further, each dialog is conditioned on a culture-adapted ontology database cylinder as an extra-linguistic context, and it must be coherent with the database content.

coherence and multi-parallelism of dialogs. GlobalWOZ, while performing a form of cultural adaptation, (i) creates erroneous slot value annotations that are inconsistent with the dialog ontology and database in the particular language, and (ii) even induces inconsistent annotations within an individual dialog. Another problem with GlobalWOZ is that the authors select a subset of 500 test set dialogs for human PEMT work based on a simple heuristic: they opt for dialogs for which the sum of corpus-level frequencies of their constitutive 4-grams, normalized by dialog length, is the largest. This selection, not motivated in the original paper and performed independently for each language, entails that different portions of the original English MultiWOZ are included into the final language-specific test sets. This design choice, besides (i) artificially decreasing linguistic diversity of dialogs chosen for the test set in each language,³ also (ii) breaks the desired multi-parallel nature of the test set. As a consequence, GlobalWOZ overestimates downstream ToD performance for target languages, and cannot be used for any direct comparison of ToD task performance across different languages since test sets

³The selection heuristic favors dialogs that contain the same most frequent 4-grams globally.

per language contain different dialogs, as also pointed out by Hung et al. (2022).

MULTI³WOZ is the only dataset which performs cultural adaptation and avoids confounding factors such as GlobalWOZ’s selection heuristics, while maintaining the desired properties of dialog coherence and multi-parallelism.

3 **MULTI³WOZ**

MULTI³WOZ comprises linguistically and culturally adapted task-oriented dialogs in four languages: Arabic (ARA; Afro-Asiatic), English (ENG; Indo-European), French (FRA; Indo-European), and Turkish (TUR; Turkic). A total of 27,480 ($3 \times 9,160$) dialogs is collected for ARA, FRA, TUR, while the dataset also includes a subset of 9,160 normalized and corrected MultiWOZ v2.3 dialogs.⁴

In what follows, we describe its creation, as depicted in Figure 2. Our approach involves three key steps: (i) *normalizing annotations* from the original MultiWOZ v2.3 with canonical values; (ii) *cultural adaptation* by contextualizing

⁴We select 9,160 out of MultiWOZ’s full set of 10,438 dialogs by filtering out erroneous dialogs identified during the normalization and cultural adaptation process; problematic dialogs were also recorded by our annotators during the dialog generation and quality control phases (see later in §3).

dialogs to entities from the relevant cultures; and (iii) *collecting linguistically adapted dialogs* from target language native speakers using a bottom-up outlined-based method.

Preliminaries and Notation. In ToD, the domains of a dataset (e.g., MultiWOZ) and the systems built upon it are typically defined by an *ontology*, which provides a structured representation of an underlying *database*. The ontology specifies slots that encompass all entity attributes and their corresponding values (Budzianowski et al., 2018). MULTI³WOZ is designed to be fully compatible with the original English MultiWOZ’s ontology and data format, but now with culturally adapted database entries (see Figure 2).

MULTI³WOZ \mathbb{D} contains four multi-parallel sets of dialogs, namely \mathbb{D}^{ARA} , \mathbb{D}^{ENG} , \mathbb{D}^{FRA} , and \mathbb{D}^{TUR} , along with their corresponding *cultural-specific databases* denoted as \mathbb{E}^{ARA} , \mathbb{E}^{ENG} , \mathbb{E}^{FRA} , and \mathbb{E}^{TUR} .⁵ Each database entry, $\mathcal{E} \in \mathbb{E}$, contains a set of slot-value pairs, such that $\mathcal{E} = \{(s_1, v_1), (s_2, v_2), \dots, (s_n, v_n)\}$.⁶ Each dialog in the dataset is represented as a list of natural language utterances, with alternating turns between the user and system initiated by the user. Each turn is annotated with its corresponding sentence-level meaning representation. Namely, for $\mathcal{D} \in \mathbb{D}$, $\mathcal{D} = [(\mathbf{u}_1, \mathbf{a}_1), \dots, (\mathbf{u}_j, \mathbf{a}_j)]$, where \mathbf{u} is a surface form (user or system) utterance; \mathbf{a} is a dialog act representation; j is the length of the dialog \mathcal{D} .

A dialog act \mathbf{a} is then defined as a set of tuples $\mathbf{a} = \{(d_1, i_1, s_1, v_1), \dots, (d_k, i_k, s_k, v_k)\}$, where each tuple consists of domain d , intent i , slot s , and slot value v .

Slot-Value Normalization. In the English MultiWOZ dataset, slot values are annotated as text spans within the corresponding utterances. This annotation scheme allows for more flexible and natural language expressions of the canonical value v_{truth} described in the ontology and database (e.g., *13:00*), resulting in various surface forms

$v^{(1)}, \dots, v^{(l)}$ (e.g., *1 pm*, *1:00 pm*, *one*). However, this flexibility can create a discrepancy between the expected canonical value required by the backend API and the predicted value by the model.⁷

Moreover, the absence of a 1-to-1 mapping between the canonical value in the database and the annotations in MultiWOZ, coupled with erroneous or misspelled entries, hinders the consistent and systematic adaptation of culture-dependent entities to the target language. To address this, we *manually* created a normalization dictionary and assigned canonical values to all slot values across the English MultiWOZ dataset. For example, we created a normalization dictionary for the *restaurant-name* slot, mapping 544 distinct surface forms to 110 canonical names. These canonical names correspond exactly to the entities in the English *restaurants* domain’s database, enabling a one-to-one mapping between the entities described in dialogs and those in the database. Besides facilitating cultural adaptation through the creation of surface form agnostic outlines, we believe that this time-consuming yet crucial normalization process will also enable consistent evaluations of models built on MULTI³WOZ. Henceforth, any mention of a slot value v assumes that it is in its canonical form.⁸

Cultural Adaptation. While English MultiWOZ contains only dialogs describing entities in the Cambridge (UK) area, MULTI³WOZ expands the scope to three additional languages targeting three cities where the target languages are considered native: Dubai for Arabic, Paris for French, and Ankara for Turkish.⁹ To ensure

⁵In order to simplify our notation, we represent a backend database as a set of data entries, where each entry corresponds to a real-world entity within the target culture.

⁶We denote each attribute of an entity as a slot and consider the domain of an entity as an inherent attribute. For example, $\{(domain, police), (name, parkside\ police\ station), (address, Parkside, Cambridge), (phone, 01223358966), (postcode, cb11jg)\}$ is a database entry in \mathbb{E}^{ENG} .

⁷The query sent to the backend API is formulated using a formal language that lacks the flexibility of natural language. This issue can significantly affect the performance of extractive models, such as extractive DST models (Heck et al., 2020; Zhou et al., 2023).

⁸The introduction of slot values in canonical forms offers supplementary information to the original MultiWOZ annotation. The original format can be automatically derived, enabling backward compatibility with previous models.

⁹We fully acknowledge that here we use the term ‘culture’ (imprecisely) as a proxy for the limited set of properties, customs, and entities to be expected or common at the target location. We also acknowledge that language-culture mappings are typically many-to-many, with the possibility of multiple languages being native to the same culture, and one language spreading over more than one culture or subculture (Hershcovich et al., 2022). Our (simplified) choice is primarily driven by pragmatic considerations and feasibility requirements.

that our dataset respects and reflects the cultural traits pertaining to each target city and language, we propose a systematic approach for cultural adaptation, which ensures dialog coherence and multi-parallelism across all languages, and includes the following steps: 1. *slot-value localization/redistribution* with cultural awareness, 2. *controlled entity replacement* with one-to-one entity mappings, 3. *slot-value randomization* to avoid verbatim memorization.

We perform *slot-value redistribution* to adjust the original slot and value to align with the target ‘culture’. These modifications are based on the feedback from native speakers of the target language with expertise in the corresponding cultural context. To better fit the target culture, we remove ENG-specific slots and values that are irrelevant to the culture. For example, we obliterate the *postcode* slot in the Arabic dataset \mathbb{D}^{ARA} due to its limited relevance in the associated culture.¹⁰

The main objective of our proposed cultural adaptation method is to perform *controlled entity replacement* using a 1-to-1 entity mapping. As a prerequisite, we first construct a localized database (e.g., \mathbb{E}^{ARA} for Arabic) for each target language. This database aims to reflect real-world entities and properties, and has been constructed by human participants in our project, native speakers of the target languages, who referred to a variety of public knowledge sources on the Internet, including the Google Places API and TripAdvisor API.¹¹

In order to construct such a 1-to-1 mapping, an English entity \mathcal{E}^{ENG} and a target entity (e.g., \mathcal{E}^{ARA}) can be mapped to each other only if all categorical slot values attributed to each entity are identical.¹² Namely, the following condition holds:

¹⁰We also consider religious factors: e.g., to respect local culture, we replace the ‘gastropub’ *restaurant type* with the value ‘Arab’, or ‘nightclub’ with ‘waterpark’ for the *attractions* slot. Moreover, we address the issue of unbalanced entity distribution in the original MultiWOZ, which is heavily skewed towards Cambridge (UK) and contains a disproportionate number of mentions of ‘colleges’ and ‘guest houses’. To mitigate this bias, we swap certain types of entities; e.g., we exchange the very specific term ‘college’ with ‘architecture’ and ‘guest house’ with ‘hotel’ to offer a better localization of the entity distribution for the target location.

¹¹However, we note that, for database completeness, a portion of the entity information has been synthetically generated due to missing information on the Web, e.g., when a restaurant does not provide a phone number on its website.

¹²A categorical slot is defined by the ontology such that the possible values for this slot are a closed set. For example,

$\forall (s^{\text{ENG}}, v^{\text{ENG}}) \in \mathcal{E}^{\text{ENG}}, \exists (s^{\text{ARA}}, v^{\text{ARA}}) \in \mathcal{E}^{\text{ARA}} : v^{\text{ENG}} = v^{\text{ARA}}$ if $\text{is_categorical}(s^{\text{ENG}})$. This strategy guarantees a consistent distribution of entities with respect to each categorical property as MultiWOZ. It further facilitates the coherent and multi-parallel creation of dialogs, particularly when the user requests a certain property of a desired entity along the progress of dialogs (e.g., ‘an *expensive* restaurant’). This stands in contrast to the random sampling cultural adaptation solution of GlobalWOZ, which results in frequently mismatched entities being returned in response to the user request, and often results in dialog incoherence.

The original MultiWOZ contains a substantial number of randomized slot values, such as *time*, *reference*, and *taxi-phone*. To prevent verbatim memorization and undesired data artefacts, we perform *slot-value randomization* independently in each target dialog subset in MULTI³WOZ. For *time*-related slot values in MULTI³WOZ, we apply the randomization by adding a 1-hour random offset drawn from a uniform distribution $[-1, 1]$ to the original value, as also illustrated in Figure 2. We ensure that all *time* relevant slots (e.g., *leaving time* and *arriving time*) in a dialog are equivalently shifted by the same randomized offset. For *reference* numbers, we employ the 1-to-1 randomly generated reference mapping. Regarding *taxi-phone* values, we first adhere to the target culture’s specific phone pattern followed by a 1-to-1 randomly generated phone mapping. In general, this procedure mitigates the risk of exploiting annotation artifacts and consequent overfitting when conducting cross-lingual transfer learning experiments.

Outline-Based Dialog Generation. By adopting the outline-based dialog generation process we simultaneously enable cultural adaptation while eliminating the impact of syntactic and lexical grounding in the source language (i.e., the so-called “translation artifacts”), while keeping the annotation protocol feasible (Majewska et al., 2023). The outline-based method can be decomposed into two steps: *outline creation* (i.e., creating dialog schemata) and *dialog writing* (i.e.,

the slot ‘price range’ can only have the values of ‘cheap’, ‘moderate’, and ‘expensive’. In contrast, the value for a *hotel name* is an open set and not categorical, as it can be any string.

creating the actual surface realizations, utterances, from the dialog schemata).

Following Majewska et al. (2023), *outline creation* involves creating minimal but comprehensive instructions for the so-called *dialog creators* (termed DCs henceforth) to generate dialogs that fully convey specific intents and slots while avoiding the imposition of predefined syntactic structures or linguistic expressions. As depicted in Figure 2, we convert a culturally adapted (termed *CA-ed* henceforth) dialog act (e.g., using ARA as an example language, \mathbf{a}^{ARA}) into a human-interpretable outline based on a set of manually defined templates, where different sets of templates are used for the user and system utterances. Given a tuple $(d, i, s, v^{\text{ARA}}) \in \mathbf{a}^{\text{ARA}}$, we transform a domain-intent pair $d-i$ into a natural language instruction, e.g., *Restaurant-Inform* \Rightarrow “Express your intent to search for a restaurant with the following properties:”. In addition, the slot s is mapped to a predefined natural language description, and it is presented along with the CA-ed slot value v^{ARA} (e.g., *booking time = 18:45*). As illustrated in Figure 2, in cases where there are multiple tuples with the same pair $d-i$, we group them together and present within a “card”. We note that a target language utterance (e.g., \mathbf{u}^{ARA}) can be constructed based on multiple cards, with each card corresponding to a unique domain-intent pair $d-i$.¹³ Additionally, each card may contain multiple slot-value pairs, where each slot value is shown as a CA-ed value (e.g., v^{ARA}). To take full advantage of our outline-based framework, we have developed a Web-based annotation toolkit along with detailed annotation guidelines; the latter is made publicly available.

Dialog writing is then carried out by bilingual speakers as DCs. They are (i) native in the target language and (ii) fluent in English: following the results from our pilots, we opted for keeping the English templates as it facilitated the quality control of templates and cards while it did not have any detrimental effect on the quality of finally generated target language dialogs. The DCs were instructed to write natural-sounding exchanges in their native language between a hypothetical user and an assistant, based on the outlines derived from the CA-ed dialog act (e.g., \mathbf{a}^{ARA}) and a set of user goals that the hypothetical user

¹³*Restaurant-Inform* is the domain-intent pair for the utterance *There will be 5 of us and 19:45 would be great.*

wants to achieve (e.g., *You are looking for a place to stay.*). For each utterance \mathbf{u} from the source ENG dataset, the tasks of the DCs were then as follows: **1)** writing a native dialog utterance from the card(s) that covers all the slot values from the cards; **2)** annotating character-level span indices for each slot value v^{ARA} ; **3)** indicating with a binary flag for each domain-intent pair $d-i$ whether this dialog act retains coherence of the full dialog, this way also signaling and capturing errors still present in the English MultiWOZ v2.3 dataset.

Duration, Cost, Dialog Creators, Quality Control. The logistically and technically complex data collection process spanned 14 months, starting in January 2022. The full cost of data collection was \sim \\$64,500, equally distributed across the three target languages. The recruited DCs are (i) professional translators and (ii) college students, recruited via the ProZ platform (www.proz.com) or from universities worldwide. A total of 133 native Arabic speakers, 112 native French speakers, and 75 native Turkish speakers contributed to the dataset.

We applied a number of quality control mechanisms throughout the data collection process. First, to ensure that the DCs have fully understood the instructions and all (sub)tasks, they were required to complete a qualification round before creating any actually deployed data. Second, our annotation platform features a real-time automatic check for all submissions, providing feedback and highlighting issues for the collected dialogs. Finally, we also ran two rounds of *post-collection dialog editing*: We invited a carefully selected small group of dialog creators, who had consistently produced exceptional high-quality dialogs, to review and, if necessary, edit all the dialogs in the validation and test sets of all three target languages.

Ethical and Responsible Data Creation and Use. Following the principles from Rogers et al. (2021), the project has placed a high priority on ethical and responsible data creation and use. It underwent the full Ethics Approval process at University of Cambridge, and we describe other ethics-related aspects here.

Terms of Use: MULTI³WOZ is released under the same MIT License as the original MultiWOZ.

Privacy: To comply with the EU General Data Protection Regulation (GDPR), we have acted as

English	Arabic	French	Turkish
I would like to eat Indian food.	أريد أن أجد مطعم هندي	Je souhaiterais avoir celui qui propose les plats indiens.	Türk yemekleri sunan bir restoran arıyorum.
There are 5 Indian restaurants available in the centre. What day are you looking to dine?	هناك خمسة اختيارات لي مطاعم هندية في وسط المدينة. ما هو يوم الحجز الذي تريده؟	Je souhaiterais avoir celui qui propose les plats indiens.	merkez bölgesinde türk yemekleri sunan beş farklı seçenek bulunuyor. Hangi güne rezervasyon yaptırmak istersiniz?
I need a reservation for Thursday. There will be 1 person at 11:15. Any of the 5 restaurants would be fine.	أريد أن أحجز للشخص واحد اليوم الخميس في الساعة 11:15.	Je réserve pour 1 personne à 10 h 45 le jeudi.	perşembe günü saat 11:15 için 1 kişilik rezervasyon yaptırmak istiyorum.
I was able to get you a table at The Golden Curry. Your reference number is 7DD4DK3N and they'll hold your table for 15 minutes. Can I help with anything else?	إن اسم المطعم هو سانجاء ورقم حجزك هو SVRJYBEW. هل تحتاج أي مساعدة أخرى؟	Bien réservé! Numéro de référence UXQXMAA8. Nom de l'emplacement: Indien 6. Autres choses?	Göksu Lokantaları restoranına yapılan A1IA291R referans kodlu rezervasyonunuz onaylanmıştır.
Thank you and goodbye.	شكرا للمساعدة	Non, je vous remercie.	Teşekkürler.
You're welcome. Goodbye.	عفوا وبجاءك الله يا عزيزي	Bonne journée à vous.	Rica ederim. İyi günler.

Figure 3: An example set of parallel dialogs in four languages: English, Arabic, French, and Turkish, extracted from the MULTI³WOZ dataset. The dialogs illustrate different aspects of cultural adaptation, including slot-value redistribution, slot-value randomization, and controlled entity replacement, which are highlighted with distinct colors. Due to space limitations, we only show a set of single-domain short dialogs. However, it is important to note that the MULTI³WOZ dataset contains multi-domain dialogs with diverse dialog patterns and linguistic constructions. The dialog ID for this specific example is SSNG0101.

a data controller and collected the minimum of personal data required for this project. All participants provided informed consent by signing a *Participant Consent Form* before any data collection occurred. To adhere to the principle of data minimization, we collected only the participants' email addresses as individually identifiable information for the sole purpose of processing payments. Our dataset consists solely of hypothetical dialogs in which the domains and content have been restricted and predefined, minimizing the risk of personal data being present in MULTI³WOZ.

Compensation. The DCs were compensated based on the number of dialogs they contributed to the dataset, with a payment rate of approximately \$12/h. As stated in our consent form, they were able to withdraw from the study at any time.

Data Structure and Statistics. Figure 3 presents an example of multi-parallel dialogs from MULTI³WOZ. All dialogs in MULTI³WOZ consist of parallel surface form utterances in multiple languages and retain the same annotations as the original MultiWOZ. Precisely, each dialog \mathcal{D} is annotated with a CA-ed user goal, as well as for each utterance u in the dialog: a CA-ed dialog act, a CA-ed dialog state. In addition, MULTI³WOZ offers (i) annotations for character-level textual spans for all the slot values in the dialog act to steer span extraction-based solutions to slot labeling (Joshi et al., 2020a), and (ii) a binary coherence indicator. The dataset is released in three




Figure 4: Utterance length in MULTI³WOZ.

standard formats: (i) json files following the structure of MultiWOZ (Budzianowski et al., 2018); (ii) a format compatible with the Hugging-face repository (Wolf et al., 2020; Lhoest et al., 2021); (iii) ConvLab-3-compatible format (Zhu et al., 2022).

MULTI³WOZ's language-independent features, e.g., the frequency of dialog acts and average dialog length, closely resemble those of the original MultiWOZ; we thus focus on the statistics pertaining to language and cultural adaptation. Figure 4 presents the distribution of the number of tokens per turn, with white spaces as the token delimiter. Note that each language exhibits variance in its morphosyntactic properties (e.g., Turkish is an agglutinative language), which naturally impacts the expected utterance length. Further, we find that 13.3% of the slot values in the dialog acts are normalized with canonical values, while 38.7% of the dialog acts' slot values are provided with CA-ed values. The type-to-token ratio (TTR) varies across languages, with English

having a lower TTR value (0.010) compared to Arabic (0.032), French (0.023), and Turkish (0.035). In comparison to the GlobalWOZ dataset, which is an MT-based dataset without CA, our dataset (MULTI³WOZ) achieves an increased TTR for Arabic (\uparrow 0.013), French (\uparrow 0.006), and Turkish (\uparrow 0.014).¹⁴ This outcome highlights that MULTI³WOZ’s bottom-up design sparked higher semantic variability and naturalness in the target languages (Majewska et al., 2023). We further highlight the higher semantic diversity of utterances in MULTI³WOZ in comparison to PEMT-based methods such as the one used by Multi²WOZ. We select a subset of 1,586 Arabic dialogs of flows shared between the two datasets and calculate the average pairwise cosine similarity between utterances in each data subset and their corresponding utterances in the English MultiWOZ, relying on LaBSE (Feng et al., 2022) as a state-of-the-art multilingual sentence encoder. The scores of 0.54 (MULTI³WOZ) and 0.91 (Multi²WOZ) suggest the higher semantic variability created through the outline-based approach with cultural adaptation.

4 MULTI³WOZ as a ToD Benchmark

MULTI³WOZ establishes a multilingual and cross-lingual benchmark for ToD systems and their sub-modules. We now present a first ‘benchmarking study’ on the dataset, evaluating representative models for NLU, DST, NLG, and E2E tasks in ToD, merely scratching the surface of possible experimental work now enabled by MULTI³WOZ.

Natural Language Understanding. NLU is typically decomposed into two established tasks: intent detection (ID) and slot labeling (SL). ID can be cast as a multi-class classification task that identifies the presence of a domain-intent pair $d-i$ (e.g., *Restaurant-Inform*) from the user’s utterance, where the set of intents is predefined in the ontology. SL is a sequence tagging task that identifies the presence of a value v and its corresponding slot s within the utterance.

¹⁴For this comparison, we utilize the ‘F&E’ proportion of the GlobalWOZ dataset. In this dataset, English utterances are translated into the target language using Google Translate, while preserving the slot values associated with English entities. The calculation of the TTR is limited to the dialogs that are included in both the GlobalWOZ dataset and our dataset.

We evaluate ID and SL methods backed by XLM-R_{base} (Conneau et al., 2020). Precisely, at each dialog turn t , the model encodes the concatenation of the previous two utterances (\mathbf{u}_{t-2} and \mathbf{u}_{t-1}) along with the current utterance (\mathbf{u}_t) to provide embedding vectors at both the sequence and token levels. To implement the intent detector, for each domain-intent pair $d-i$ defined by the ontology, the representation of the ‘< s >’ token is subsequently projected down to two logits and passed through a Sigmoid layer to form a Bernoulli distribution indicating if $d-i$ appears in the \mathbf{u}_t . Performance is evaluated by measuring its accuracy in identifying the exact presence of all domain-intent pairs in a dialog act, as well as its F1 score. For SL, we adopt the widely used BIO labeling scheme to annotate each token in the user’s utterance.^{15,16}

In Table 2, we observe that the fully supervised ID model achieves similarly high accuracy across all languages, and we also observe a large cross-lingual transfer gap (Hu et al., 2020) for both tasks. Further, there is a substantial decrease in performance for Arabic SL. Note that in MULTI³WOZ the slot-value spans are annotated at the character level, and we only consider a span to be correctly identified if there is an exact match. At the same time, Rust et al. (2021) observed that the sub-optimal performance of the tokenizers for the multilingual models may yield degraded downstream performance. To investigate the limitations of tokenization, we then aligned the slot boundaries with the token boundaries. Specifically, we defined the slot span as the minimal token span that covered the entire slot in the utterance. With this approach, the identical model achieved F1 of 78.44 (\uparrow 30.00) for Arabic SL, confirming that the suboptimal XLM-R’s tokenization was the primary contributor to the original performance degradation in Arabic.

Dialog State Tracking. For DST, we follow the standard MultiWOZ preprocessing and evaluation

¹⁵Specifically, each token is labeled with either B- $d-i-s$ (e.g., *B-Restaurant-Inform-Food*), denoting the beginning of a slot-value pair with the corresponding slot name, I- $d-i-s$ indicating it is inside the slot-value, or O indicating that the token is not associated with any slot-value pair.

¹⁶We conducted all NLU experiments on a single RTX 24 GiB GPU with a batch size of 64 and a learning rate of $2e - 5$. We trained each model for 10 epochs and selected the model with the best F1 score on the validation set as the final model.

Language	Intent Detection		Slot Labeling			Dialog State Tracking		
	Accuracy	F1	Precision	Recall	F1	JGA	Turn Acc.	F1
Fully Supervised (Monolingual)								
ENG	92.71	95.77	95.92	94.08	94.99	59.90	97.87	93.67
ARA	92.20	94.59	48.44	42.47	45.26	47.72	96.85	89.26
FRA	88.92	92.93	79.57	77.76	78.65	49.77	97.02	89.93
TUR	91.50	94.52	87.25	86.72	86.94	53.59	97.28	91.04
AVG.	91.33	94.45	77.80	75.26	76.46	52.74	97.26	90.97
Zero-Shot Cross-lingual Transfer (from English)								
ARA	60.28	71.06	17.56	27.74	21.47	1.47	80.73	5.80
FRA	72.88	81.71	48.53	60.52	53.86	3.66	85.08	32.83
TUR	69.52	79.09	48.47	66.80	56.18	1.30	82.05	15.22
AVG.	67.56	77.29	38.19	51.69	43.84	2.14	82.62	17.95

Table 2: Fully supervised and zero-shot cross-lingual transfer from English (\mathbb{D}^{ENG} as the source) for ID, SL, and DST tasks on MULTI³WOZ. AVG. shows the mean average of the evaluation scores across all four languages. The reported scores are averaged over 3 random runs.

setups (Wu et al., 2019), excluding the ‘hospital’ and ‘police’ domains due to the absence of test dialogs in these domains. We report the Joint Goal Accuracy (JGA), Turn Accuracy, and Joint F1.

We adapt T5DST (Lin et al., 2021a) as a strong baseline that reformulates the DST as a QA task with slot descriptions. The DST model is back-boned with mT5_{small} (Xue et al., 2021) (as very similar scores were obtained with mT5_{base}). Regarding the model and training details, readers are referred to the original work (Lin et al., 2021a).¹⁷

Fully supervised DST scores provide a strong benchmark with the multilingual T5DST model over all languages in MULTI³WOZ. We observe the highest performance in English (59.9% JGA), followed by Turkish, French, and Arabic, indicating the levels of difficulty of DST for each language. Table 2 presents the zero-shot cross-lingual transfer-from-English results, revealing poor transferability of the DST models across languages (all below 4% JGA). This indicates the limitations of current multilingual models in zero-shot setups and the challenge of transfer learning for culturally adapted dialogs in MULTI³WOZ.

Natural Language Generation. We approach the NLG task as a sequence-to-sequence prob-

lem, again supported by mT5_{small}. Specifically, at each dialog turn t , the model takes the input of its dialog context, and generates a system response \mathbf{u}_t . Traditionally, NLG in ToD systems is defined as the task of converting a dialog act into a natural language utterance (Williams and Young, 2007). In our study, we evaluate NLG performance in both a traditional setup, where the goal is to realize the surface form of the dialog act, and an end-to-end LM setup, where we model response generation as a transduction problem from the dialog history to a natural response. Third, we consider the setup where both the dialog history and the ‘oracle’ dialog act are available, serving as a performance upper bound. For the *surface realization* setup, we convert the dialog act \mathbf{a}_t into a flattened string format (e.g., `[inform][restaurant]([price range][expensive], [area][center])`) to serve as the input. For the *language modeling* setup, the model generates a response \mathbf{u}_t solely based on the preceding dialog history \mathbf{u}_{t-2} and \mathbf{u}_{t-1} . In this setup, the generation model does not have any knowledge about the system’s ontology and database. In the *language modeling with oracle* setup, the model takes the concatenation of the two preceding utterances \mathbf{u}_{t-2} and \mathbf{u}_{t-1} , as well as \mathbf{a}_t as input.

Following MultiWOZ, we evaluate with the corpus BLEU score (Papineni et al., 2002); we evaluate lexicalized utterances without performing delexicalization. We also report ROUGE-L

¹⁷The experiments were run on a single RTX 24 GiB GPU; batch size of 4, a learning rate of $1e-4$; 5 epochs.

Language	Surface Realization			Language Modeling			Language Modeling with Oracle		
	BLEU	ROUGE	METEOR	BLEU	ROUGE	METEOR	BLEU	ROUGE	METEOR
ENG	20.67	47.76	44.16	8.66	27.95	25.18	21.20	48.52	44.31
ARA	9.57	14.04	21.92	7.22	20.77	18.11	17.56	15.99	35.22
FRA	9.96	35.31	29.17	6.19	24.47	19.78	13.61	40.69	34.87
TUR	13.59	39.29	33.99	9.87	30.07	26.84	24.23	53.76	48.49
AVG.	13.45	34.10	32.31	7.98	21.14	22.48	19.15	39.74	40.72

Table 3: Fully supervised NLG performance for mT5_{small}. AVG. shows the mean average of the evaluation scores across all four languages. The reported scores are averaged over 3 random runs.

(Lin, 2004) and METEOR (Banerjee and Lavie, 2005).¹⁸

The results are summarized in Table 3. We observe that the performance of English is significantly higher than other languages in the first setup. This disparity can be attributed to the fact that dialog acts are considered a formal language for the system to process internally and, except for culturally adapted values, they are provided in English. Therefore, it is more challenging for a model to learn how to generate natural language utterances in other languages. Furthermore, by incorporating the dialog history and the oracle dialog act, the performance of all three languages improved significantly, indicating that modeling the dialog history contributes to more coherent responses. Lastly, in the absence of database information, the performance for all languages is considerably low. This highlights the challenge of modeling ToD, and underlines the necessity of incorporating databases into the ToD models in future work.

End-to-End Modeling. Finally, E2E modeling performance serves as an even more comprehensive, challenging and arguably more important indicator for assessing the progress of ToD research, garnering intensified research attention (Hosseini-Asl et al., 2020; Lin et al., 2020; Peng et al., 2021; Su et al., 2022; Wu et al., 2023, *inter alia*). Developing an E2E system offers several advantages over focusing on individual sub-components like NLU modules or dialog state trackers. The E2E approach achieves increased applicability, enabling the development of practical real-world applications. Moreover, it reduces

vulnerability to error propagation across sub-components and offers a simpler system design compared to the traditional pipelined approaches.

To the best of our knowledge, no previous publicly available implementation of a multilingual E2E ToD system exists that would be compatible with the MultiWOZ dataset and its derivatives. Other available multilingual ToD benchmarks either lack E2E results (Hung et al., 2022; Ding et al., 2022), or do not release their implementation (Zuo et al., 2021). The only exception is BiToD (Lin et al., 2021b); however, the BiToD dataset and the associated system use a different annotation schema, which is incompatible with MultiWOZ. Therefore, we present the first publicly available implementation of a multilingual E2E system compatible with the MultiWOZ-related datasets. We release this implementation as a baseline for further research and experimentation on MULTI³WOZ.

Our system is composed of three key components: a Dialog State Tracking (DST) model, a Database (DB) Interface component, and a Response Generation (RG) model. First, the DST model is a sequence-to-sequence model, which takes the concatenated lexicalized form of all the historical utterances as input and generates a linearized dialog state (e.g., *hotel price range = cheap; type = hotel*). Then, the DB Interface transforms the predicted dialog state into an SQL query. This query is executed, resulting in a list of entities that satisfy the specified constraints, which are then returned to the system. Finally, the RG model, also implemented as a seq2seq model, takes as input the concatenation of historical utterances, predicted dialog state, and a database summary that indicates the number of entities returned for each active domain (e.g., *restaurant more than five*). It generates a delexicalized response, which can be further lexicalized using

¹⁸All NLG experiments were run on a single A100 80 GiB GPU; batch size of 32, a learning rate of $1e-3$; 10 epochs.

Language	End-to-End Modeling		
	Inform	Success	BLEU
ENG	67.9	39.0	15.7
ARA	66.8	36.7	14.0
FRA	47.9	22.2	12.0
TUR	45.9	21.2	16.7
AVG.	57.1	29.8	14.6

Table 4: Fully supervised E2E performance for mT5_{large}. AVG. shows the mean average of the evaluation scores across all four languages. The reported scores are averaged over 3 random runs.

the values in the predicted dialog state and the returned entities from the database.

In our implementation, we utilize two separate mT5_{large} models as the backbone for the DST model and the RG model. As discussed later, we opt for the large model because it demonstrates a substantial performance advantage over its smaller counterpart. The data preprocessing, including the linearization of dialog state annotations for training, and the evaluation protocol are based on the established implementation of the SOLOIST system (Peng et al., 2021). To ensure up-to-date functionality, our implementation is based on the most recent version 4.30 of the HuggingFace transformers repository. Our system is designed to prioritize simplicity and efficiency, with the primary goal of minimizing the complexity and effort required for training, evaluation, and future development. We report the standard evaluation metrics for the E2E task, including the Inform Rate, Success Rate, and the delexicalized corpus BLEU score.¹⁹

Table 4 presents the results of the fully supervised E2E experiments. As anticipated, we observe noticeable performance disparities across languages, particularly in comparison to English. Furthermore, we find that the size of the pre-trained language model significantly impacts system performance. Specifically, the mT5_{large} model exhibits a substantial (mean average) performance improvement of 16.4 Inform Rate, 17.2 Success Rate, and 4.6 BLEU points, compared to mT5_{small}.

¹⁹All E2E experiments were run on a single A100 80 GiB GPU; batch size of 4, learning rate of $5e - 5$; 5 epochs.

5 Conclusion

We have introduced a large-scale, culturally adapted, multilingual, and multi-parallel training and evaluation framework for ToD, which covers $\sim 495,000$ dialog turns over 4 languages. The dataset was motivated by the limitations of current ToD datasets in multilingual setups, which we systematically analyzed as one contribution of this work. Owing to its unique set of properties and scale, beyond initial analyses and experiments conducted in this work, we hope that MULTI³WOZ will inspire a wide array of further developments in modeling, analysis, and interpretability of multilingual and cross-lingual multi-domain ToD.

For instance, future work could replicate the data collection process to expand the dataset to even more languages (including low-resource ones). Further, one could analyze the performance disparities observed in Tables 2–4 within each language-specific ToD system, as well as explore methods to mitigate such disparities, e.g., through the utilization of cross-lingual transfer techniques. Future work could also explore evaluation metrics beyond the ones explored in this work, e.g., it would be interesting to explore the correlation between the increase in evaluation scores in multilingual ToD systems and the resulting performance gain in terms of factors such as utility, user experience, and user satisfaction. Additionally, it would be important to investigate how ToD systems should, ideally, be constructed and evaluated across different languages to ensure their inclusiveness and robustness in diverse linguistic contexts.

Code and Data. We release the dataset and code at github.com/cambridgeltl/multi3woz.

Acknowledgments

Songbo Hu is supported by Cambridge International Scholarship. Ivan Vulić acknowledges the support of a personal Royal Society University Research Fellowship (no 221137; 2022–).

We would like to thank our internship students, Bassil Alaeddin (for the work on the Arabic portion of the dataset) and Max Letellier (for French), for their contributions and dedication to this project. We are grateful to a large number of

our diligent annotators for their significant efforts and contributions to this work. Furthermore, we would like to express our gratitude to the ACL editors and anonymous reviewers for their insightful feedback, which greatly improved the quality of this paper.

References

- David Ifeoluwa Adelani, Jade Abbott, Graham Neubig, Daniel D'souza, Julia Kreutzer, Constantine Lignos, Chester Palen-Michel, Happy Buzaaba, Shruti Rijhwani, Sebastian Ruder, Stephen Mayhew, Israel Abebe Azime, Shamsuddeen H. Muhammad, Chris Chinenye Emezue, Joyce Nakatumba-Nabende, Perez Ogayo, Aremu Anuoluwapo, Catherine Gitau, Derguene Mbaye, Jesujoba Alabi, Seid Muhie Yimam, Tajuddeen Rabi'u Gwadabe, Ignatius Ezeani, Rubungo Andre Niyongabo, Jonathan Mukiibi, Verrah Otiende, Iroro Orife, Davis David, Samba Ngom, Tosin Adewumi, Paul Rayson, Mofetoluwa Adeyemi, Gerald Muriuki, Emmanuel Anebi, Chiamaka Chukwunke, Nkiruka Odu, Eric Peter Wairagala, Samuel Oyerinde, Clemencia Siro, Tobius Saul Bateesa, Temilola Oloyede, Yvonne Wambui, Victor Akinode, Deborah Nabagereka, Maurice Katusiime, Ayodele Awokoya, Mouhamadane MBOUP, Dibora Gebreyohannes, Henok Tilaye, Kelechi Nwaike, Degaga Wolde, Abdoulaye Faye, Blessing Sibanda, Orevaoghene Ahia, Bonaventure F. P. Dossou, Kelechi Ogueji, Thierno Ibrahima DIOP, Abdoulaye Diallo, Adewale Akinfaderin, Tendai Marengereke, and Salomey Osei. 2021. MasakhaNER: Named entity recognition for African languages. *Transactions of the Association for Computational Linguistics*, 9:1116–1131. <https://doi.org/10.1162/tacl.a.00416>
- Duygu Altinok. 2018. An ontology-based dialogue management system for banking and finance dialogue systems. *CoRR*, abs/1804.04838. Version 1. <https://doi.org/10.48550/arXiv.1804.04838>
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2020. Translation artifacts in cross-lingual transfer learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7674–7684, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.618>
- Satanjeev Banerjee and Alon Lavie. 2005. ME-THEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Kalina Bontcheva, Leon Derczynski, and Ian Roberts. 2017. *Crowdsourcing Named Entity Recognition and Entity Linking Corpora*. Springer Netherlands, Dordrecht. https://doi.org/10.1007/978-94-024-0881-2_32
- 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. <https://doi.org/10.18653/v1/D18-1547>
- 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 *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4516–4525, Hong Kong, China. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-1459>
- Inigo Casanueva, Ivan Vulić, Georgios Spithourakis, and Paweł Budzianowski. 2022. NLU++: A multi-label, slot-rich, generalisable dataset for natural language understanding in task-oriented dialogue. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1998–2013, Seattle, United States. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.findings-naacl.154>

- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.747>
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. XNLI: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D18-1269>
- Mai Hoang Dao, Thinh Hung Truong, and Dat Quoc Nguyen. 2021. Intent detection and slot filling for Vietnamese. In *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association, Brno, Czechia, 30 August – 3 September 2021*, pages 4698–4702. ISCA.
- Bosheng Ding, Junjie Hu, Lidong Bing, Mahani Aljunied, Shafiq Joty, Luo Si, and Chunyan Miao. 2022. GlobalWoZ: Globalizing Multi-WoZ to develop multilingual task-oriented dialogue systems. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1639–1657, Dublin, Ireland. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.acl-long.115>
- Abteen Ebrahimi, Manuel Mager, Arturo Oncevay, Vishrav Chaudhary, Luis Chiruzzo, Angela Fan, John Ortega, Ricardo Ramos, Annette Rios, Ivan Vladimir Meza Ruiz, Gustavo Giménez-Lugo, Elisabeth Mager, Graham Neubig, Alexis Palmer, Rolando Coto-Solano, Thang Vu, and Katharina Kann. 2022. AmericasNLI: Evaluating zero-shot natural language understanding of pretrained multilingual models in truly low-resource languages. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6279–6299, Dublin, Ireland. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.acl-long.435>
- 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. <https://doi.org/10.18653/v1/W17-5526>
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2022. Language-agnostic BERT sentence embedding. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.acl-long.62>
- Jack FitzGerald, Christopher Hench, Charith Peris, Scott Mackie, Kay Rottmann, Ana Sanchez, Aaron Nash, Liam Urbach, Vishesh Kakarala, Richa Singh, Swetha Ranganath, Laurie Crist, Misha Britan, Wouter Leeuwis, Gökhan Tür, and Prem Natarajan. 2022. MASSIVE: A 1m-example multilingual natural language understanding dataset with 51 typologically-diverse languages. *CoRR*, abs/2204.08582. Version 2. <https://doi.org/10.18653/v1/2023.acl-long.235>
- Rahul Goel, Waleed Ammar, Aditya Gupta, Siddharth Vashishtha, Motoki Sano, Faiz Surani, Max Chang, HyunJeong Choe, David Greene, Kyle He, Rattima Nitisaroj, Anna Trukhina, Shachi Paul, Pararth Shah, Rushin Shah, and Zhou Yu. 2023. PRESTO: A multilingual dataset for parsing realistic task-oriented dialogs. *CoRR*, abs/2303.08954. Version 2. <https://doi.org/10.48550/arXiv.2303.08954>
- Narendra K. Gupta, Gökhan Tür, Dilek Hakkani-Tür, Srinivas Bangalore, Giuseppe Riccardi, and Mazin Gilbert. 2006. The AT&T spoken language understanding system. *IEEE*

- Transactions on Speech and Audio Processing*, 14(1):213–222. <https://doi.org/10.1109/TSA.2005.854085>
- Ting Han, Ximing Liu, Ryuichi Takanobu, Yixin Lian, Chongxuan Huang, Dazhen Wan, Wei Peng, and Minlie Huang. 2021. Multiwoz 2.3: A multi-domain task-oriented dialogue dataset enhanced with annotation corrections and co-reference annotation. In *Natural Language Processing and Chinese Computing - 10th CCF International Conference, NLPCC 2021, Qingdao, China, October 13–17, 2021, Proceedings, Part II*, volume 13029 of *Lecture Notes in Computer Science*, pages 206–218. Springer. https://doi.org/10.1007/978-3-030-88483-3_16
- Michael Heck, Carel van Niekerk, Nurul Lubis, Christian Geishauer, Hsien-Chin Lin, Marco Moresi, and Milica Gasic. 2020. TripPy: A triple copy strategy for value independent neural dialog state tracking. In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 35–44, 1st virtual meeting. Association for Computational Linguistics.
- Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. The ATIS spoken language systems pilot corpus. In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24–27, 1990*. <https://doi.org/10.3115/116580.116613>
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014. 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. IEEE. <https://doi.org/10.1109/SLT.2014.7078595>
- Matthew Henderson, Ivan Vulić, Iñigo Casanueva, Pawel Budzianowski, Daniela Gerz, Sam Coope, Georgios Spithourakis, Tsung-Hsien Wen, Nikola Mrkšić, and Pei-Hao Su. 2019. Polyresponse: A rank-based approach to task-oriented dialogue with application in restaurant search and booking. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3–7, 2019 - System Demonstrations*, pages 181–186. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-3031>
- Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. 2022. Challenges and strategies in cross-cultural NLP. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6997–7013, Dublin, Ireland. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.acl-long.482>
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. A simple language model for task-oriented dialogue. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS’20, Red Hook, NY, USA. Curran Associates Inc.*
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.
- Chia-Chien Hung, Anne Lauscher, Ivan Vulić, Simone Ponzetto, and Goran Glavaš. 2022. Multi2WOZ: A robust multilingual dataset and conversational pretraining for task-oriented dialog. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3687–3703, Seattle, United States. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.naacl-main.270>
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020a. SpanBERT: Improving pre-training by representing and predicting spans. *Transactions*

- of the Association for Computational Linguistics, 8:64–77. <https://doi.org/10.1162/tacl.a.00300>
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020b. The state and fate of linguistic diversity and inclusion in the NLP world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.560>
- Liliana Laranjo, Adam G. Dunn, Huong Ly Tong, Ahmet Baki Kocaballi, Jessica A. Chen, Rabia Bashir, Didi Surian, Blanca Gallego, Farah Magrabi, Annie Y. S. Lau, and Enrico W. Coiera. 2018. Conversational agents in healthcare: A systematic review. *Journal of the American Medical Informatics Association*, 25(9):1248–1258. <https://doi.org/10.1093/jamia/ocy072>, PubMed: 30010941
- Stefan Larson and Kevin Leach. 2022. A survey of intent classification and slot-filling datasets for task-oriented dialog. *CoRR*, abs/2207.13211. Version 1. <https://doi.org/10.48550/arXiv.2207.13211>
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. Datasets: A community library for natural language processing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.emnlp-demo.21>
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Zhaojiang Lin, Bing Liu, Seungwhan Moon, Paul Crook, Zhenpeng Zhou, Zhiguang Wang, Zhou Yu, Andrea Madotto, Eunjoon Cho, and Rajen Subba. 2021a. Leveraging slot descriptions for zero-shot cross-domain dialogue StateTracking. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5640–5648, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.448>
- 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. <https://doi.org/10.18653/v1/2020.emnlp-main.273>
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, Peng Xu, Feijun Jiang, Yuxiang Hu, Chen Shi, and Pascale Fung. 2021b. BiToD: A bilingual multi-domain dataset for task-oriented dialogue modeling. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- Olga Majewska, Evgeniia Razumovskaia, Edoardo Maria Ponti, Ivan Vulić, and Anna Korhonen. 2023. Cross-lingual dialogue dataset creation via outline-based generation. *Transactions of the Association for Computational Linguistics*, 11:139–156. <https://doi.org/10.1162/tacl.a.00539>
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. A SICK cure for the evaluation of compositional distributional semantic models. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 216–223, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Nikita Moghe, Evgeniia Razumovskaia, Liane Guillou, Ivan Vulić, Anna Korhonen, and

- Alexandra Birch. 2023. Multi3NLU++: A multilingual, multi-intent, multi-domain dataset for natural language understanding in task-oriented dialogue. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3732–3755, Toronto, Canada. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-acl.230>
- Nikola Mrkšić, Ivan Vulić, Diarmuid Ó. Séaghdha, Ira Leviant, Roi Reichart, Milica Gašić, Anna Korhonen, and Steve Young. 2017. Semantic specialization of distributional word vector spaces using monolingual and cross-lingual constraints. *Transactions of the Association for Computational Linguistics*, 5:309–324. https://doi.org/10.1162/tacl_a_00063
- 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*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics. <https://doi.org/10.3115/1073083.1073135>
- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayandeh, Lars Liden, and Jianfeng Gao. 2021. Soloist: Building task bots at scale with transfer learning and machine teaching. *Transactions of the Association for Computational Linguistics*, 9:807–824. https://doi.org/10.1162/tacl_a_00399
- Edoardo Maria Ponti, Goran Glavaš, Olga Majewska, Qianchu Liu, Ivan Vulić, and Anna Korhonen. 2020. XCOPA: A multilingual dataset for causal commonsense reasoning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2362–2376, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.185>
- Edoardo Maria Ponti, Helen O’Horan, Yevgeni Berzak, Ivan Vulić, Roi Reichart, Thierry Poibeau, Ekaterina Shutova, and Anna Korhonen. 2019. Modeling language variation and universals: A survey on typological linguistics for natural language processing. *Computational Linguistics*, 45(3):559–601. https://doi.org/10.1162/colila_00357
- Jun Quan, Shian Zhang, Qian Cao, Zizhong Li, and Deyi Xiong. 2020. RiSAWOZ: A large-scale multi-domain Wizard-of-Oz dataset with rich semantic annotations for task-oriented dialogue modeling. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 930–940, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.67>
- Antoine Raux, Brian Langner, Dan Bohus, Alan W. Black, and Maxine Eskénazi. 2005. Let’s go public! Taking a spoken dialog system to the real world. In *INTERSPEECH 2005 - Eurospeech, 9th European Conference on Speech Communication and Technology, Lisbon, Portugal, September 4–8, 2005*, pages 885–888. ISCA. <https://doi.org/10.21437/Interspeech.2005-399>
- Evgeniia Razumovskaia, Goran Glavaš, Olga Majewska, Edoardo Maria Ponti, Anna Korhonen, and Ivan Vulić. 2022. Crossing the conversational chasm: A primer on natural language processing for multilingual task-oriented dialogue systems. *Journal of Artificial Intelligence Research*, 74:1351–1402. <https://doi.org/10.1613/jair.1.13083>
- Anna Rogers, Timothy Baldwin, and Kobi Leins. 2021. ‘Just what do you think you’re doing, dave?’ A checklist for responsible data use in NLP. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4821–4833, Punta Cana, Dominican Republic. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.findings-emnlp.414>
- Phillip Rust, Jonas Pfeiffer, Ivan Vulić, Sebastian Ruder, and Iryna Gurevych. 2021. How good is your tokenizer? On the monolingual performance of multilingual language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3118–3135, Online. Association for Computational Linguistics.

<https://doi.org/10.18653/v1/2021.acl-long.243>

- Sebastian Schuster, Sonal Gupta, Rushin Shah, and Mike Lewis. 2019. Cross-lingual transfer learning for multilingual task oriented dialog. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3795–3805, Minneapolis, Minnesota. Association for Computational Linguistics. <https://doi.org/10.18653/v1/N19-1380>
- Pararth Shah, Dilek Hakkani-Tür, Bing Liu, and Gökhan Tür. 2018. Bootstrapping a neural conversational agent with dialogue self-play, crowdsourcing and on-line reinforcement learning. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1–6, 2018, Volume 3 (Industry Papers)*, pages 41–51. Association for Computational Linguistics. <https://doi.org/10.18653/v1/N18-3006>
- Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai, and Yi Zhang. 2022. Multi-task pre-training for plug-and-play task-oriented dialogue system. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4661–4676, Dublin, Ireland. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.acl-long.319>
- Gökhan Tür, Andreas Stolcke, L. Lynn Voss, Stanley Peters, Dilek Hakkani-Tür, John Dowding, Benoît Favre, Raquel Fernández, Matthew Frampton, Michael W. Frandsen, Clint Frederickson, Martin Graciarena, Donald Kintzing, Kyle Leveque, Shane Mason, John Niekrasz, Matthew Purver, Korbinian Riedhammer, Elizabeth Shriberg, Jing Tien, Dimitra Vergyri, and Fan Yang. 2010. The CALO meeting assistant system. *IEEE Transactions on Speech Audio Processing*, 18(6):1601–1611. <https://doi.org/10.1109/TASL.2009.2038810>
- Shyam Upadhyay, Manaal Faruqui, Gökhan Tür, Dilek Hakkani-Tür, and Larry P. Heck. 2018. (almost) zero-shot cross-lingual spoken language understanding. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2018, Calgary, AB, Canada, April 15–20, 2018*, pages 6034–6038. IEEE. <https://doi.org/10.1109/ICASSP.2018.8461905>
- Jason D. Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422. <https://doi.org/10.1016/j.csl.2006.06.008>
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-demos.6>
- 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 *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 808–819, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1078>
- Qingyang Wu, Deema Alnuhait, Derek Chen, and Zhou Yu. 2023. Using textual interface to align external knowledge for end-to-end task-oriented dialogue systems. *CoRR*, abs/2305.13710. Version 1. <https://doi.org/10.48550/arXiv.2305.13710>
- Weijia Xu, Batoool Haider, and Saab Mansour. 2020. End-to-end slot alignment and recognition for cross-lingual NLU. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5052–5063, Online. Association for

- Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.410>
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.41>
- Zhao Yan, Nan Duan, Peng Chen, Ming Zhou, Jianshe Zhou, and Zhoujun Li. 2017. Building task-oriented dialogue systems for online shopping. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4–9, 2017, San Francisco, California, USA*, pages 4618–4626. AAAI Press. <https://doi.org/10.1609/aaai.v31i1.11182>
- Steve J. Young. 2010. Cognitive user interfaces. *IEEE Signal Processing Magazine*, 27(3):128–140. <https://doi.org/10.1109/MSP.2010.935874>
- Han Zhou, Ignacio Iacobacci, and Pasquale Minervini. 2023. XQA-DST: Multi-domain and multi-lingual dialogue state tracking. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 969–979, Dubrovnik, Croatia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-eacl.73>
- Qi Zhu, Christian Geishauser, Hsien-Chin Lin, Carel van Niekerk, Baolin Peng, Zheng Zhang, Michael Heck, Nurul Lubis, Dazhen Wan, Xiaochen Zhu, Jianfeng Gao, Milica Gasic, and Minlie Huang. 2022. Convlab-3: A flexible dialogue system toolkit based on a unified data format. *CoRR*, abs/2211.17148. Version 1. <https://doi.org/10.48550/arXiv.2211.17148>
- Qi Zhu, Kaili Huang, Zheng Zhang, Xiaoyan Zhu, and Minlie Huang. 2020. CrossWOZ: A large-scale Chinese cross-domain task-oriented dialogue dataset. *Transactions of the Association for Computational Linguistics*, 8:281–295. <https://doi.org/10.1162/tacl.a.00314>
- Lei Zuo, Kun Qian, Bowen Yang, and Zhou Yu. 2021. AllWOZ: Towards multilingual task-oriented dialog systems for all. *CoRR*, abs/2112.08333. Version 1. <https://doi.org/10.48550/arXiv.2112.08333>