

# MULTI<sup>3</sup>NLU<sup>++</sup>: A Multilingual, Multi-Intent, Multi-Domain Dataset for Natural Language Understanding in Task-Oriented Dialogue

Nikita Moghe<sup>\*1\*</sup>, Evgeniia Razumovskaia<sup>\*2</sup>, Liane Guillou<sup>1</sup>,  
Ivan Vulić<sup>2</sup>, Anna Korhonen<sup>2</sup>, Alexandra Birch<sup>1</sup>

School of Informatics, University of Edinburgh<sup>1</sup>

Language Technology Lab, University of Cambridge<sup>2</sup>

{nikita.moghe, liane.guillou, a.birch}@ed.ac.uk, {er563, iv250, alk23}@cam.ac.uk

## Abstract

Task-oriented dialogue (ToD) systems have been widely deployed in many industries as they deliver more efficient customer support. These systems are typically constructed for a single domain or language and do not generalise well beyond this. To support work on Natural Language Understanding (NLU) in ToD across multiple languages *and* domains simultaneously, we constructed MULTI<sup>3</sup>NLU<sup>++</sup>, a **multilingual**, **multi-intent**, **multi-domain** dataset. MULTI<sup>3</sup>NLU<sup>++</sup> extends the English-only NLU<sup>++</sup> dataset to include manual translations into a range of high, medium, and low resource languages (Spanish, Marathi, Turkish and Amharic), in two domains (BANKING and HOTELS). Because of its multi-intent property, MULTI<sup>3</sup>NLU<sup>++</sup> represents complex and natural user goals, and therefore allows us to measure the realistic performance of ToD systems in a varied set of the world’s languages. We use MULTI<sup>3</sup>NLU<sup>++</sup> to benchmark state-of-the-art multilingual models for the NLU tasks of *intent detection* and *slot labelling* for ToD systems in the multilingual setting. The results demonstrate the challenging nature of the dataset, particularly in the low-resource language setting, offering ample room for future experimentation in multi-domain multilingual ToD setups.

## 1 Introduction

Task-oriented dialogue (ToD) systems (Gupta et al., 2006; Young et al., 2013), in which conversational agents assist human users to achieve their specific goals, have been used to automate telephone-based and online customer service tasks in a range of domains, including travel (Raux et al., 2003, 2005), finance and banking (Altinok, 2018), and hotel booking (Li et al., 2019).

ToD systems are often implemented as a pipeline of dedicated modules (Raux et al., 2005; Young et al., 2013). The *Natural Language Understanding (NLU) module* performs two crucial tasks:

<sup>\*</sup>Equal contribution

|     | Intent modules: debit, pin, dont_know, card, request_info, new       |
|-----|--|
| en: | I forgot my pin! Any chance you can send me a new debit card?        |
| am: | የደለፍ ቁጥሬ ረብሀት! አዲስ የድህረ ገጽ ካርድ ሊልኩልኝ የሚችሉበት እድል ይኖርይዎታል?             |
| mr: | मी माझा पिन विसरलो! तुम्ही मला नवीन डेबिट कार्ड पाठवू शकता का??      |
| tr: | PIN'imî unuttum! Bana yeni bir banka kartı gönderme şansınız var mı? |
| es: | ¡Olvidé mi PIN! ¿Me podéis enviar una nueva tarjeta de débito?       |

Figure 1: Example from the MULTI<sup>3</sup>NLU<sup>++</sup> dataset for the BANKING domain demonstrating the complex NLU tasks of multi-label intent detection across multiple languages. Intent labels consist of generic and domain-specific intents.

1) intent detection and 2) slot labelling. In the *intent detection* task the aim is to identify or classify the goal of the user’s utterance from several *pre-defined classes* (or intents) (Tur et al., 2010). These intents are then used by the *policy module* (Gašić et al., 2012; Young et al., 2013) to decide the next conversational move by the conversational agent in order to mimic the flow of a human-human dialogue. In the *slot labeling* task each token in an utterance is assigned a label describing the type of semantic information represented by the token. During this process *relevant information* e.g. named entities, times/dates, quantities etc., defining the crucial information of the user’s utterance, is identified.

Although intent detection models can reach impressive performance and have been deployed in many commercial systems (Altinok, 2018; Li et al., 2019), they are still unable to fully capture the variety and complexity of natural human interactions, and as such do not meet the requirements for deployment in more complex industry settings (Casanueva et al., 2022). This is due in part to the limitations of existing datasets for training and evaluating ToD systems. As highlighted by Casanueva et al. (2022) they are **1)** predominantly limited to detecting a single intent, **2)** focused on a single domain, and **3)** include a small set of slot types (Larson and Leach, 2022). Furthermore, the success of task-oriented dialogue is **4)** often evaluated on a small set of higher-resource languages (i.e.,

typically English) which does not test how generalisable systems are to the diverse range of the world’s languages (Razumovskaia et al., 2022a).

Arguably one of the most serious limitations, hindering their deployment to more complex conversational scenarios, is the inability to handle multiple intents. In many real-world scenarios, a user may express multiple intents in the same utterance, and TOD systems must be able to handle such scenarios (Gangadharaiah and Narayanaswamy, 2019). For example, in Figure 1, the user expresses two main intents: (i) informing that they have forgotten their pin and thus, (ii) they would like to request a new debit card instead. A single-intent detection system can detect either of the two intents (but not both), resulting in partial completion of the user’s request. Casanueva et al. (2022) recently proposed a multi-label intent detection dataset to capture such complex user requests. They further propose using intent modules as intent labels that can act as sub-intent annotations. In this example, “pin” and “don’t\_know” compose the first intent while “request\_info”, “new”, “debit”, and “card” compose the second. The use of intent modules, due to their combinatorial power, can support more complex conversational scenarios, and also allows reusability of the annotations across multiple domains. For example, “request\_info”, “new”, and “membership” can be reused for gyms, salons, etc. to request information about new memberships at the respective institutions.

Furthermore, ToD systems are typically constructed for a single language. Their extension to other languages is restricted by the lack of available training data for many of the world’s languages. Whilst the construction of multilingual TOD datasets has been given some attention (Razumovskaia et al., 2022a; Majewska et al., 2022; Xu et al., 2020), these datasets often include synthetic translations in the form of post-edited Machine Translation output (Ding et al., 2022; Zuo et al., 2021; Bellomaria et al., 2019). Post-editing may introduce undesirable effects and result in texts that are simplified, normalised, or exhibit interference from the source language as compared with manually translated texts (Toral, 2019).

To address all of the limitations discussed above, we propose MULTI<sup>3</sup>NLU<sup>++</sup>, a **multilingual**, **multi-intent**, **multi-domain** for training and evaluating TOD systems. MULTI<sup>3</sup>NLU<sup>++</sup> extends the recent monolingual English-only dataset

NLU<sup>++</sup>, which is a multi-intent, multi-domain dataset for the BANKING and HOTELS domains. MULTI<sup>3</sup>NLU<sup>++</sup> adds the element of multilinguality and thus enables simultaneous cross-domain and cross-lingual training and experimentation for TOD NLU as its unique property.

MULTI<sup>3</sup>NLU<sup>++</sup> includes expert manual translations of the 3,080 utterances in NLU<sup>++</sup> to four languages of diverse typology and data availability: Spanish, Marathi, Turkish, and Amharic. The selection of languages covers a range of language families and scripts and includes high, medium, and low-resource languages. Capturing language diversity is particularly important if we wish to design multilingual TOD systems that are robust to the variety of expressions used across languages to represent the same value or concept. Using MULTI<sup>3</sup>NLU<sup>++</sup> we demonstrate the challenges involved in extending existing state-of-the-art Machine Translation systems and multilingual language models for NLU in TOD systems. MULTI<sup>3</sup>NLU<sup>++</sup> is publicly available at <https://huggingface.co/datasets/uoel-nlp/multi3-nlu>.

## 2 Background and Related Work

NLU<sup>++</sup> (Casanueva et al., 2022), which serves as the base for MULTI<sup>3</sup>NLU<sup>++</sup>, covers two domains: BANKING and HOTELS. The intent and slot ontologies include intents and slots which are general, cross-domain types, as well as domain-specific ones. The intent ontology includes a total of 62 intents of which 23 and 14 intents are BANKING and HOTELS specific, respectively. The slot ontology includes 17 slot types, of which three and four slot types are BANKING and HOTELS specific, respectively. Such intent and slot ontology construction allows for the extension to new domains more easily. In this work, we inherit monolingual NLU<sup>++</sup>’s core benefits over previous datasets with the additional layer of multilinguality: **1**) it is based on real customer data which addresses the issue of low lexical diversity in crowd-sourced data (Larson and Leach, 2022), **2**) it supports the requirement for production systems to capture multiple intents from a single utterance, and **3**) it is also *slot-rich* – it combines a large set of fine-grained intents with a large set of fine-grained slots to facilitate more insightful evaluation of models that jointly perform the two NLU tasks (Chen et al., 2019; Gangadharaiah and Narayanaswamy, 2019).

**Multilingual NLU Datasets.** Prior work has demonstrated the importance and particular challenges posed by low-resource languages (Goyal et al., 2022; Magueresse et al., 2020; Xia et al., 2021), while an increasing number of multilingually pretrained models (Xue et al., 2021; Conneau et al., 2020; Feng et al., 2022; Liu et al., 2020) enable significant improvements in processing them. While some tasks (e.g., NER (Adelani et al., 2021) or NLI (Ebrahimi et al., 2022)) already have benchmarks to evaluate on low-resource languages, dialogue Natural Language Understanding (NLU) is still lagging behind in this respect. The reasons for this include the high cost of data collection, as well as specific challenges posed by dialogue, e.g., colloquial speech or tone used in live conversations. Additionally, while we have observed increased interest in few-shot methods for multilingual dialogue NLU (Bhathiya and Thayasivam, 2020; Moghe et al., 2021; Feng et al., 2022; Razumovskaia et al., 2022b, *among others*), none of the existing datasets (Xu et al., 2020; FitzGerald et al., 2022; van der Goot et al., 2021) allow for reproducible comparison between few-shot methods. In other words, no dataset to date has provided predefined splits for few-shot experiments.

Until recently, resources for multilingual NLU have been scarce (Razumovskaia et al., 2022a). The vast majority were built upon the English ATIS dataset (Price, 1990), which has been extended to ten target languages (Upadhyay et al., 2018; Dao et al., 2021; Xu et al., 2020). However, ATIS covers only one domain (airline booking) and has been claimed to be almost solved already, more than a decade ago (Tur et al., 2010). More recent datasets cover multiple domains and broader linguistic geography (Majewska et al., 2022; Schuster et al., 2019; FitzGerald et al., 2022), including domains such as music or alarm, which are in frequent use in production systems.

All of the existing multilingual NLU datasets label every user utterance with a single intent, although current production-level systems often rely on multi-intent labelling (Gangadharaiah and Narayanaswamy, 2019; Qin et al., 2020), allowing for faster development cycles and updates if a new, previously unseen intent is observed (Casanueva et al., 2022). MULTI<sup>3</sup>NLU<sup>++</sup> is the first multilingual, multi-intent dataset with modular intent annotations. In comparison to other multilingual NLU datasets presented in Table 1, MULTI<sup>3</sup>NLU<sup>++</sup> has

a larger intent set and is natively multi-intent/label. Unlike xSID, which is an evaluation-only dataset, it contains both training and evaluation data for all languages. Further, while MASSIVE is based on utterances specifically generated for the dataset (Bastianelli et al., 2020), MULTI<sup>3</sup>NLU<sup>++</sup> is based on *real user inputs* to a system in industrial settings (Casanueva et al., 2022). MULTI<sup>3</sup>NLU<sup>++</sup> enables systematic comparisons of dialogue NLU systems in few-shot setups for cross-lingual and cross-domain transfer for low-, medium- and high-resource languages.

### 3 Dataset Collection

Our data collection process focuses on creating datasets with natural and realistic conversations, avoiding many artefacts that arise from crowdsourcing or automatic translation. We ask professional translators to manually translate each source-language utterance into the four target languages in the dataset; this promotes equal opportunity for future research for all four languages, and enables comparative cross-language analysis.

**Choice of Base Dataset.** This work aims to create a multilingual dataset which would be useful for testing production-like systems in multiple languages. Thus, we chose NLU<sup>++</sup> (Casanueva et al., 2022) as our base dataset because i) it consists of real-world examples; and ii) every example is labelled with multiple intents which has proven useful in production settings (Qin et al., 2020; Casanueva et al., 2022).

Prior work on English multi-intent classification often uses MixSNIPS and MixATIS as benchmarks (Qin et al., 2020). However, the multi-intent examples in these datasets are synthetic, i.e., they are obtained through a simple concatenation of two single-intent examples, e.g., “Play this song and book a restaurant”. This process leads to repetitive content and unnatural examples in the datasets. In contrast, NLU<sup>++</sup> consists of intrinsically multi-intent examples such as “I cannot login to my account because I don’t remember my pin.”, and is much more semantically variable.

**Languages.** We selected Spanish, a widely-spoken Romance language, as our high-resource language. Marathi, an Indo-Aryan language predominantly spoken in the Indian state of Maharashtra, is our medium-resource language. Turkish, an agglutinative Turkic language, may be regarded as a low-resource language from a Machine Translation per-

| Dataset                                     | # Languages | Domains | Size                                    | # intents | # slots | Multi-Intent? |
|---|-------------|---------|---|-----------|---------|---------------|
| Multilingual TOP<br>(Schuster et al., 2019) | 3           | 3       | 43,323 [en]<br>8,643 [es]<br>5,082 [th] | 12        | 11      | ×             |
| MultiATIS++<br>(Xu et al., 2020)            | 8           | 1       | 5,871                                   | 18        | 84      | ×             |
| xSID<br>(van der Goot et al., 2021)         | 13          | 6       | 800                                     | 16        | 33      | ×             |
| MASSIVE<br>(FitzGerald et al., 2022)        | 51          | 18      | 19,521                                  | 60        | 55      | ×             |
| MULTI <sup>3</sup> NLU <sup>++</sup>        | 4           | 2       | 3,080                                   | 62        | 17      | ✓             |

Table 1: Statistics of representative multilingual dialogue NLU datasets. MultiATIS++ contains some utterances with more than one intent label but the vast majority are single-label.

spective, or as a medium-resource language based on the amount of training data in XLM-R (Conneau et al., 2020).<sup>1</sup> Amharic, an Ethiopian Semitic language belonging to the Afro-Asiatic language family, is our low-resource language. Spanish and Turkish are written in Latin script, Marathi is written in Devanagari, and Amharic in Ge’ez script.

**Manual Translation.** The use of crowd-sourcing agencies to collect multilingual datasets has often resulted in crowd workers using a Machine Translation (MT) API (plus post-editing) to complete the task, or even simply transliterating the given sentence (Goyal et al., 2022). In the case of post-editing MT output, translations may exhibit *post-editese* – post-edited sentences are often simplified, normalised, or exhibit a higher degree of interference from the source language than manual translations (Toral, 2019). In our case, we wish to preserve the register of the original utterances, in particular with respect to the colloquial nature of many of the utterances in the original English NLU++ dataset. Furthermore, we wish to collect high-quality, natural translations. We, therefore, opted to recruit professional translators to perform manual translation, via two online platforms: Proz.com<sup>2</sup> and Blend Express.<sup>3</sup> We instructed our translators to treat the task as a *creative writing task* (Ponti et al., 2020) and maintain the *colloquial nature of the utterances*. We also provided instructions to annotate the spans in the generated translations. We provide the instructions given to our translators in Appendix B. We recruited three translators per

language; Spanish translators were recruited via Proz.com, and translators for the remaining languages were assigned to us by Blend Express.

We first conducted a pilot task in which we asked professional translators to translate and annotate 50 sentences per domain. We conducted an in-house evaluation by native speakers of the respective language to verify that the translations were colloquial, that named entities were appropriately translated, and that the translation was of high quality. These evaluations were communicated to the translators.<sup>4</sup>

After the pilot, we asked the same translators to complete the translation of the remaining utterances in the dataset. We ran an automatic checker to ensure that the slot values marked by the translators were present within their translated sentences. Further corrections such as incorrect annotations were also communicated to the translators.

**Duration and Cost.** Data collection was carried out over five months and involved (i) selecting translation agencies, (ii) running the pilot task, (iii) providing feedback to the translators, and (iv) the full-fledged data collection phase. The professional translators were compensated at £0.06/word for Spanish and £0.07/word for the remaining languages. The total cost of MULTI<sup>3</sup>NLU<sup>++</sup> is £7,624; see Appendix C for further details.

**Slot Span Verification.** Initially, the translators performed slot labelling simultaneously with translation. To ensure the quality of the annotation, the translated data was revised and annotated by three native speakers of the target language. Similar to Majewska et al. (2022), we used an automated inter-annotator reliability method to automatically

<sup>1</sup>The amount of Turkish data is higher than, e.g., Marathi (which we consider to be a medium-resource language).

<sup>2</sup>[www.proz.com/](http://www.proz.com/); Proz.com is a platform for recruiting freelance translators who self-quote their remuneration.

<sup>3</sup>[www.getblend.com/online-translation/](http://www.getblend.com/online-translation/)

<sup>4</sup>As translators are generally not used to annotating spans in text, a significant amount of time was spent on training the translators for this new task.

verify the annotation quality. We conducted slot span labelling for 200 examples in Spanish for both BANKING and HOTELS domains from three native Spanish annotators. The accuracy score<sup>5</sup> for a given sub-sample was at 86.8% revealing high agreement between the annotators.

## 4 Baseline Experiments

We benchmark several state-of-the-art approaches on MULTI<sup>3</sup>NLU<sup>++</sup> to **1**) provide reference points for future dataset use and **2**) demonstrate various aspects of multilingual dialogue NLU systems which can be evaluated using the dataset, such as cross-lingual and cross-domain generalisation. We provide baseline numbers for intent detection and slot labelling and analyse the performance across languages and methods, and for different sizes of training data. We follow the main experimental setups from Casanueva et al. (2022) where possible, but we extend them to multilingual contexts.

**Training Data Setups.** We follow an N-fold cross-validation setup following the setup in Casanueva et al. (2022). The experiments were run in three setups: *low*, *mid* and *large*. The *low* data setup corresponds to 20-fold cross-validation, where the model is trained on  $\frac{1}{20}$ th of the dataset and tested on the remaining 19 folds. The *mid* and *large* setups correspond to 10-fold cross-validation, where in *mid* setup the model is trained on  $\frac{1}{10}$ th of the data and tested on the other nine folds, and vice versa for the *large* setup.

**Domain Setups.** MULTI<sup>3</sup>NLU<sup>++</sup> contains training and evaluation data for two domains: BANKING and HOTELS. It thus enables evaluation of NLU systems in three domain setups: *in-domain*, *cross-domain* and *all-domain*. In the *in-domain* setup, the model is trained and evaluated on the same domain, that is, we are testing how well the model can generalise to unseen user utterances while operating in the same intent space as the training data and without any domain distribution shift. In the *cross-domain* setup, the model is trained on one domain and tested on the other domain. We evaluate on the union of label sets of two domains rather than on the intersection as done by Casanueva et al. (2022). In this setup, we are testing how well a model can generalise to a new, unseen domain including intents unseen in training. In the *all-domain* setup,

<sup>5</sup>The accuracy is the ratio between the number of exactly matching spans between the annotators to the total number of slot values annotated.

we train and evaluate the models on data from both domains. In this setup, we are testing how models perform on the larger label set (where some labels are shared between the domains) when examples are provided for all classes.

**Language Setups.** MULTI<sup>3</sup>NLU<sup>++</sup>, offering comparable sets of annotated data points across languages, allows for systematic comparisons of multilingual dialogue NLU systems on languages with different amounts of resources and diverse typological properties. We evaluate NLU systems in two setups: *in-language* and *cross-lingually*. *Cross-lingual* benchmarking is conducted with the established approaches: (i) direct transfer using multilingually pretrained large language models (e.g., XLM-R; Conneau et al. (2020)), and (ii) transfer via translation, i.e., when either the test utterances are translated into the source language (*Translate-Test*). We source our translations from the M2M100 translation model (Fan et al., 2021).

### 4.1 Classification-Based Methods

**Experimental Setup and Hyperparameters.** We evaluate two standard classification approaches to intent detection: (i) MLP-based with a fixed encoder; and (ii) full-model fine-tuning. Prior work has demonstrated that strong intent detection results can be attained without fine-tuning the full encoder model both in monolingual (Casanueva et al., 2020) and multilingual setups (Gerz et al., 2021). The idea is to use a fixed efficient sentence encoder to encode sentences and train only the multi-layer perceptron (MLP) classifier on top of that to identify the intents. As we are dealing with multi-label classification, a sigmoid layer is stacked on top of the classifier. Intent classes for which the probability scores are higher than a (predefined) threshold are considered active. As in (Casanueva et al., 2022), we use the threshold of 0.3. In the experiments we evaluate two state-of-the-art multilingual sentence encoders: 1) *mpnet*, a multilingual sentence encoder trained using multilingual knowledge distillation (Reimers and Gurevych (2020)), and 2) *LaBSE*, a language-agnostic BERT sentence encoder (Feng et al., 2022) which was trained using dual-encoder training. LaBSE was especially tailored to produce improved sentence encodings in low-resource languages. The models were loaded from the sentence-transformers library (Reimers and Gurevych, 2019).

In the full fine-tuning setup, which is the current

|                         | ENGLISH     |             | AMHARIC     |             | MARATHI     |             | SPANISH     |             | TURKISH     |             | AVG         |             |
|-------------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
|                         | H           | B           | H           | B           | H           | B           | H           | B           | H           | B           | H           | B           |
| Full-Fine-tuning: XLM-R | 50.0        | 61.7        | 43.7        | 51.3        | 46.4        | 57.9        | 48.3        | 59.3        | 49.9        | 56.9        | 47.6        | 55.8        |
| MLP-based: LaBSE        | 55.4        | 64.3        | <b>48.8</b> | <b>56.6</b> | <b>53.5</b> | <b>63.5</b> | 56.6        | 64.0        | 55.9        | <b>61.6</b> | 54.1        | <b>60.7</b> |
| MLP-based: mpnet        | <b>61.8</b> | <b>68.8</b> | 39.9        | 44.1        | 53.1        | 58.8        | <b>61.1</b> | <b>65.2</b> | <b>59.6</b> | 60.3        | <b>55.1</b> | 58.7        |

Table 2: Comparison between *Full fine-tuning* and *MLP-based* setups for intent detection ( $F_1 \times 100$ ). The experiments are run for the **10-Fold** in-language in-domain setup. H is HOTELS domain and B is BANKING domain.

| Model   | BANKING     |             |             | HOTELS      |             |             | ALL         |             |             |
|---------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
|         | 20-Fold     | 10-Fold     | Large       | 20-Fold     | 10-Fold     | Large       | 20-Fold     | 10-Fold     | Large       |
| AMHARIC |             |             |             |             |             |             |             |             |             |
| LaBSE   | <b>46.2</b> | <b>56.6</b> | <b>77.3</b> | <b>38.6</b> | <b>48.8</b> | <b>71.9</b> | <b>46.4</b> | <b>56.1</b> | <b>76.1</b> |
| mpnet   | 37.0        | 44.1        | 64.7        | 31.9        | 39.9        | 60.6        | 36.2        | 43.1        | 63.5        |
| SPANISH |             |             |             |             |             |             |             |             |             |
| LaBSE   | 51.4        | 64.0        | <b>85.2</b> | 46.4        | 56.6        | 79.5        | 52.2        | 63.5        | <b>83.8</b> |
| mpnet   | <b>55.2</b> | <b>65.2</b> | 84.4        | <b>50.9</b> | <b>61.1</b> | <b>80.3</b> | <b>55.4</b> | <b>65.1</b> | 83.5        |

Table 3: In-language in-domain intent detection performance for Amharic and Spanish ( $F_1 \times 100$ ). Results for other languages are provided in Appendix D, Table 12.

standard in multilingual NLP (Xu et al., 2020), a multilingually pretrained encoder is fine-tuned on task-specific data. In our case, XLM-R (Base) (Conneau et al., 2020) is fine-tuned for multi-label intent detection.

All MLP-based models were trained with the same hyperparameters, following the suggested values from Casanueva et al. (2022). The MLP classifier comprises one 512-dimensional hidden layer and *tanh* as non-linear activation. The learning rate is 0.003 for MLP-based models and  $2e-5$  for full model fine-tuning, with linear weight decay in both setups. For all setups, the models were trained with AdamW (Loshchilov and Hutter, 2019) for 500 and 100 epochs for intent detection and slot labelling, respectively. We used a batch size of 32. The evaluation metric is micro- $F_1$ .

**Results and Discussion.** Table 2 presents the comparison between the *full fine-tuning* based on XLM-R and *MLP-based* classification approaches. The results demonstrate that for the in-domain in-language 10-fold setup, the *MLP-based* approach works consistently better than *full fine-tuning* across domains and languages. We assume that the reason is that the *MLP-based* approach is more parameter-efficient than *full fine-tuning*, making it more suited for such low-data setups. Due to their computational efficiency combined with competitive performance, we focus on MLP-based models for the remainder of this section.

The main MLP-based intent detection results are presented in Tables 3 and 4 for in-domain intent

detection for the in-language and zero-shot cross-lingual setups, respectively. When we compare the performance on low- and high-resource languages, although the sizes and content of the training data are the same across languages, we observe a large gap between the performance of the same models on Spanish and Amharic. In addition, the results in Table 3 reveal the properties of the multilingual sentence encoders with respect to the intent detection task. While *mpnet* performs consistently better on Spanish (our high-resource language), *LaBSE* is a much stronger encoder for low-resource Amharic. The differences are especially pronounced in the low-data setups (**20-Fold** and **10-Fold**). While for Spanish this difference can be recovered with more training data (cf. **Large** setup), for Amharic these differences persist and are amplified across large training data setups.

We consider zero-shot transfer from two source languages: English (the most commonly used high-resource source language) and Amharic (our low-resource option). Surprisingly, the results in Table 4 suggest that using English as a default source language, as typically done in work on cross-lingual transfer, might not be optimal. In fact, using Amharic as a source language leads to stronger transfer results across languages. One trend we observed is that the lower resource the source language is, the stronger the performance on diverse target languages.<sup>6</sup> We speculate that by training on lower-resource language data we might ‘unearth’

<sup>6</sup>This trend is observed across other source languages and both sentence encoders, as shown in Appendix D.

| SRC | DOMAIN  | TGT  |      |      |      |      | SRC | DOMAIN  | TGT  |      |      |      |      |
|-----|---------|------|------|------|------|------|-----|---------|------|------|------|------|------|
|     |         | EN   | AM   | MR   | ES   | TR   |     |         | AM   | EN   | MR   | ES   | TR   |
| EN  | BANKING | 51.6 | 38.9 | 34.1 | 36.2 | 29.5 | AM  | BANKING | 46.2 | 48.4 | 38.9 | 40.2 | 33.8 |
|     | HOTELS  | 45.1 | 29.8 | 29.2 | 34.4 | 27.5 |     | HOTELS  | 38.6 | 43.5 | 33.5 | 37.4 | 31.3 |
|     | ALL     | 51.6 | 37.7 | 32.9 | 35.9 | 28.8 |     | ALL     | 46.4 | 48.4 | 39.1 | 41.5 | 34.6 |

Table 4: Cross-lingual in-domain results for intent detection with MLP-based and LaBSE as the fixed sentence encoder ( $F_1 \times 100$ ). The results are presented for transfer from English and Amharic for the **20-Fold** setup, while the results for other source languages and training data setups are available in Appendix D, Table 15.

| SRC | DOMAIN  | TGT  |      |      |      |      | SRC | DOMAIN  | TGT  |      |      |      |      |
|-----|---------|------|------|------|------|------|-----|---------|------|------|------|------|------|
|     |         | EN   | AM   | MR   | ES   | TR   |     |         | AM   | EN   | MR   | ES   | TR   |
| EN  | BANKING | 37.3 | 15.3 | 19.7 | 22.9 | 18.8 | AM  | BANKING | 27.3 | 19.9 | 22.8 | 15.5 | 20.4 |
|     | HOTELS  | 24.3 | 6.1  | 10.8 | 14.3 | 14.3 |     | HOTELS  | 10.4 | 7.9  | 6.4  | 5.6  | 6.8  |
|     | ALL     | 43.7 | 16.4 | 22.0 | 28.4 | 26.3 |     | ALL     | 28.6 | 20.0 | 21.9 | 15.4 | 21.6 |

Table 5: Cross-lingual in-domain results for slot labelling with full fine-tuning of XLM-R ( $F_1 \times 100$ ). The results are presented for transfer from English and Amharic for the **20-Fold** setup, while the results for other source languages and training data setups are available in Appendix G, Table 24.

the sentence encoders’ multi-lingual capabilities.

The main results for slot labelling are presented in Table 5 for the in-domain in-language and zero-shot cross-lingual setups. The comparison of in-language results for high-resource (en) and low-resource (am) demonstrate a similar trend to intent classification: in-language performance on the high-resource language is stronger than that on the low-resource language. However, unlike for intent detection, the high-resource language serves as a stronger source for cross-lingual transfer than the low-resource language.

## 4.2 QA-Based Methods

**Experimental Setup and Hyperparameters.** Reformulating TOD tasks as question-answering (QA) problems has achieved state-of-the-art performance (Namazifar et al., 2021). As these methods were the best performing for the NLU++ dataset, we now investigate this approach in the multilingual setting for both NLU tasks. To formulate intent detection as an extractive QA task, the utterance is appended with “yes. no. [UTTERANCE]” and acts as the context. Intent labels are converted into questions as “Is the intent to ask about [INTENT\_DESCRIPTION]” where INTENT\_DESCRIPTION is the free-form description of the intent. The QA model must learn to predict the span as “yes” across all the questions corresponding to the specific intents in the utterance and “no” otherwise. During the evaluation of transfer performance to other languages, the utterance is in the target language while the question

is in the source language.<sup>7</sup> For slot labelling, the template is “none. [UTTERANCE]”. We follow the strict evaluation approach for slot labelling - a span is marked as correct only if it exactly matches the ground truth span (Namazifar et al., 2021).

To build an extractive QA model for these tasks, we first fine-tune a multilingual language model mDeBERTa (He et al., 2021) with a general-purpose QA dataset such as SQuADv2 (Rajpurkar et al., 2018) and then with the respective languages from our dataset. We fine-tune for 10 epochs (5 for the Large setting) with a learning rate of  $1e - 5$ , weight decay is 0, and batch size is 4.

**Results and Discussion.** We report the results for intent detection in the 20-Fold setup using English and Amharic in Table 6, and the remaining results are in Appendix F. We also compare cross-lingual transfer with translation-based methods. We find that mDeBERTa-based QA models have comparable results for English multi-intent detection even with the monolingual models in Casanueva et al. (2022). We notice a drastic drop in the zero-shot transfer performance across all languages and domains. This is in line with recent findings that zero-shot transfer is harder for dialogue tasks (Ding et al., 2022; Hung et al., 2022; Majewska et al., 2022) as opposed to other cross-lingual tasks (Hu et al., 2020). Unlike the observation in §4.1, using a higher resource language is beneficial for transfer learning in the QA setup. We find that improvement in transfer performance depends more on the matching of the script of the source and tar-

<sup>7</sup>We found this setup to be empirically better than posing the question and utterance in the same language.

| Method                                  | Domain  | EN         | AM         | MR         | ES         | TR         |
|---|---------|------------|------------|------------|------------|------------|
| <b>Cross-lingual zero-shot transfer</b> |         |            |            |            |            |            |
| EN                                      | BANKING | 80.8 ± 1.3 | 42.7 ± 4.5 | 58.4 ± 3.2 | 66.2 ± 3.1 | 56.5 ± 3.6 |
|   | HOTELS  | 66.9 ± 3.0 | 37.1 ± 3.3 | 51.1 ± 2.9 | 62.2 ± 2.5 | 59.6 ± 2.8 |
|   | ALL     | 79.6 ± 1.3 | 40.7 ± 4.8 | 56.5 ± 3.7 | 67.9 ± 3.5 | 58.2 ± 4.3 |
| AM                                      | BANKING | 56.1 ± 2.2 | 30.1 ± 4.7 | 31.3 ± 4.5 | 27.9 ± 4.9 | 28.5 ± 4.7 |
|   | HOTELS  | 30.5 ± 8.4 | 51.1 ± 2.1 | 27.5 ± 7.6 | 31.5 ± 8.3 | 33.0 ± 8.5 |
|   | ALL     | 29.4 ± 4.9 | 57.2 ± 2.6 | 32.3 ± 4.1 | 27.7 ± 5.2 | 30.7 ± 4.5 |
| <b>Translate-Test</b>                   |         |            |            |            |            |            |
| EN                                      | ALL     | -          | 16.2 ± 3.0 | 63.3 ± 2.1 | 67.8 ± 1.8 | 46.3 ± 5.0 |
| AM                                      | ALL     | 16.2 ± 4.1 | -          | 15.2 ± 4.0 | 11.7 ± 2.5 | 11.7 ± 2.5 |

Table 6:  $F_1 \times 100$  for the QA-based intent detection models in the **20-Fold** setup using EN and AM for training. The remaining results are in Appendix F.

| Domain  | EN   | AM   | MR   | ES   | TR   |
|---------|------|------|------|------|------|
| BANKING | 59.7 | 44.2 | 48.7 | 42.9 | 52.3 |
| HOTELS  | 61.6 | 47.6 | 42.0 | 49.0 | 50.9 |
| ALL     | 63.0 | 46.8 | 46.9 | 47.7 | 53.6 |

Table 7:  $F_1 \times 100$  for cross-lingual zero-shot results of QA-based slot labelling models when trained with English data for the **20-Fold** setup. The results for other languages and data setups are in Appendix G.

get language followed by the amount of training data present per language during the pre-training of the mDeBERTa model. The exception is Amharic as the source language, where transfer performance is poor across the board. The scores also indicate QA models are better at this task than MLP methods, cf. Table 5.

While comparing zero-shot cross-lingual transfer with Translate-Test, we find that the latter is poorer (except in the case of EN-MR), opposing the findings in Majewska et al. (2022). We attribute this to the poor quality of the translations as repetition/nonsensical generation is rampant in them. Further, translation-based methods have an additional overhead of translating the sentences.

We find that the standard deviation is quite variable across both tasks. We check for standard deviations across different data setups and find that the trend is consistent irrespective of the underlying training data. There is also no visible trend between the language similarity or the amount of training data present during pre-training of mDeBERTa.

Table 7 reports the slot labelling results. Unlike intent detection results, slot labelling results are comparable across the languages in the same data settings. Similar to the intent detection results, using mDeBERTa QA models yields better results than those reported in Table G for high resource languages. Overall, we find that QA-based methods are indeed a promising research avenue (Liu et al.,

2021b), even in the multilingual setting (Zhao and Schütze, 2021).<sup>8</sup>

## 5 Discussion and Future Work

### High-Resource and Low-Resource Languages.

Language models such as mDeBERTa and XLM-R are pretrained on  $\sim 100$  languages. However, their representational power is uneven for high- and low-resource languages (Lauscher et al., 2020; Ebrahimi et al., 2022; Wu et al., 2022). MULTI<sup>3</sup>NLU<sup>++</sup> includes the same training and evaluation data for all languages, allowing us to systematically analyse model performance on dialogue NLU for both high- and low-resource languages. We compare the performance for different languages in the in-domain setup. As seen in Figure 2a, the overall trend in performance is the same across languages: with more training data, we gain higher performance overall. Interestingly, the absolute numbers are indicative of the resources available in pre-training for a given language. For instance, Amharic (am) has the lowest performance while Spanish (es) has the highest performance.<sup>9</sup>

**Cross-Domain Generalisation.** We now consider the intent detection task in the cross-domain in-language setting. The results in Figure 2b corroborate the findings from the in-domain experiments: the lower-resource the language is, the lower the performance on the task is. It is noticeable that the performance in the cross-domain setup is much lower than for the in-domain setup, additionally exposing the complexity of MULTI<sup>3</sup>NLU<sup>++</sup>. Additionally, in Figure 2b we observe that for the

<sup>8</sup>The efficiency overhead of the QA-based approach is that every utterance is accompanied by  $X$  questions where  $X$  is the number of intent/slot labels.

<sup>9</sup>This trend holds across domains, sentence encoders, and the zero-shot QA setup. These additional results are provided in Appendix E and Appendix F.



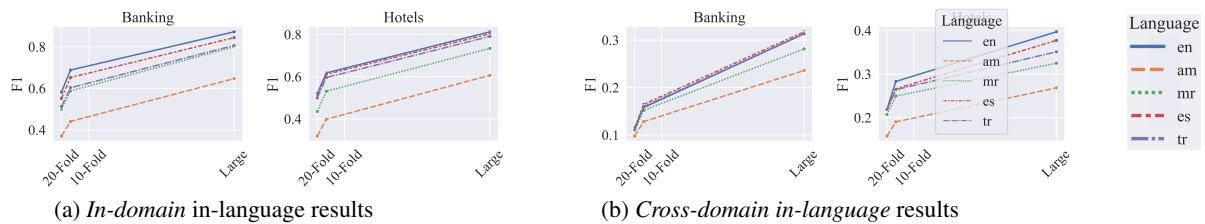


Figure 2: Comparison of in-language (a) in-domain and (b) cross-domain results for intent detection ( $F_1$ ). The model is the *MLP-based* baseline with mpnet as the underlying encoder.

cross-domain setup, high-resource languages benefit more from the increase in training data size than low-resource languages. This shows that the gap in performance on low- and high-resource languages is rooted not only in the amount of in-task training data available but also in the representational power of multilingual models for low-resource languages.

**Future Directions.** Our work focuses on collecting high-quality parallel data through expert translators. With recent advances in MT (Kocmi et al., 2022), it would be worth investigating if the quality of datasets collected with machine translation + human post-editing (Hung et al., 2022; Ding et al., 2022) is on par with the human translators, especially for the higher resource language pairs.

Another possible direction for future work would be to further diversify the dataset by including additional languages, i.e. with a focus on increasing coverage of language families, branches, and/or scripts, or properties that pose particular challenges in multilingual settings (e.g. free word order in Machine Translation, etc.).

We believe our resource will have an impact beyond multilingual multi-intent multi-domain systems. We hope the community addresses interesting questions in data augmentation (generating paraphrases with multiple intents), analysing representation learning in multilingual models, translation studies, and MT evaluation.

## 6 Conclusion

We collected MULTI<sup>3</sup>NLU<sup>++</sup>, a dataset that facilitates multilingual, multi-label, multi-domain NLU for task-oriented dialogue. Our dataset incorporates core properties from its predecessor NLU<sup>++</sup> (Casanueva et al., 2022): a multi-intent and slot-rich ontology, a mixture of generic and domain-specific intents for reusability, and utterances that are based on complex scenarios. We investigated these properties in a multilingual setting for Spanish, Marathi, Turkish, and Amharic.

We implemented MLP-based and QA-based baselines for intent detection and slot labelling across different data setups, transfer learning setups, and multilingual models. From a wide set of observations, we highlight that (i) there is a significant drop in performance across all languages as compared to NLU<sup>++</sup> with performance drops increasing as we progress from high- to low-resource languages; (ii) zero-shot performance improves when the source language has lower resources in the MLP setup; (iii) cross-lingual transfer in the QA-based intent detection is dependent on matching the script of the source language and amount of data during pretraining setup.

We hope that the community finds this dataset valuable while working on advancing research in multilingual NLU for task-oriented dialogue.

## 7 Limitations

MULTI<sup>3</sup>NLU<sup>++</sup>, like NLU<sup>++</sup> on which it was based, comprises utterances extracted from real dialogues between users and conversational agents as well as synthetic human-authored utterances constructed with the aim of introducing additional combinations of intents and slots. The utterances, therefore, lack the wider context that would be present in a complete dialogue. As such the dataset cannot be used to evaluate systems with respect to discourse-level phenomena present in dialogue.

Our source dataset includes scenarios and names that are Anglo-centric and will not capture the nuances of intent detection at the regional and cultural level. Future efforts in collecting multilingual datasets should focus on appropriate localisation of the mentioned domains, intents, and scenarios (Liu et al., 2021a; Majewska et al., 2022).

## 8 Ethics Statement

Our dataset builds on the NLU<sup>++</sup> dataset (Casanueva et al., 2022) which was collected by removing any personal information. All the names

in the dataset are also collected by randomly combining names and surnames from the list of the top 10K names from the US registry.

Our data collection process was thoroughly reviewed by the School of Informatics, University of Edinburgh under the number 2019/59295. Our translators are on legal contracts with the respective translation agencies (See Appendix C) for details.

Although we have carefully vetted our datasets to exclude problematic examples, the larger problem of unethical uses and unfairness in conversational systems cannot be neglected (Dinan et al., 2021). Our work also uses multilingual language models that are shown to harm marginalised populations (Kaneko et al., 2022). Our dataset is publicly available under Creative Commons Attribution 4.0 International (CC-BY-4.0).

## Acknowledgements

We thank our translators on Proz.com and Blend Express for helping us with the translations in this dataset. We thank Blend Express for offering us 10% discount on their services. We thank Nina Markl, Stephanie Droop, Laurie Burchell, Ronald Cardenas, Resul Tugay, and Hallelujah Kebede for helping us with our internal evaluation. We are also grateful to Iñigo Casanueva for his feedback and advice while working on the data collection. We thank the anonymous reviewers for their helpful suggestions.

This work was conducted within the EU project GoURMET H2020–825299. This work was supported in part by the UKRI Centre for Doctoral Training in Natural Language Processing, funded by the UKRI (grant EP/S022481/1) and the University of Edinburgh (Moghe) and by the ERC H2020 Advanced Fellowship GA 742137 SEMANTAX (Guillou). We also thank Huawei for their support (Moghe). The work was also supported in part by a personal Royal Society University Research Fellowship (no 221137; 2022-) awarded to Ivan Vulić, as well as by a Huawei research donation to the Language Technology Lab at the University of Cambridge.

## 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 Rabiou Gwadabe, Ignatius Ezeani, Rubungo Andre Niyongabo, Jonathan Mukiibi, Verah 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.

Duygu Altinok. 2018. An ontology-based dialogue management system for banking and finance dialogue systems. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Paris, France. European Language Resources Association (ELRA).

Emanuele Bastianelli, Andrea Vanzo, Pawel Swietojanski, and Verena Rieser. 2020. [SLURP: A spoken language understanding resource package](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7252–7262, Online. Association for Computational Linguistics.

Valentina Bellomaria, Giuseppe Castellucci, Andrea Favalli, and Raniero Romagnoli. 2019. [Almawave-slu: A new dataset for SLU in italian](#). In *Proceedings of the Sixth Italian Conference on Computational Linguistics, Bari, Italy, November 13-15, 2019*, volume 2481 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Hemanthage S Bhatthiya and Uthayasanker Thayasivam. 2020. [Meta learning for few-shot joint intent detection and slot-filling](#). In *Proceedings of the 2020 5th International Conference on Machine Learning Technologies*, pages 86–92.

Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. [Efficient intent detection with dual sentence encoders](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 38–45, Online. Association for Computational Linguistics.

Inigo Casanueva, Ivan Vulić, Georgios Spithourakis, and Pawel 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.
- Qian Chen, Zhu Zhuo, and Wen Wang. 2019. [BERT for joint intent classification and slot filling](#). *ArXiv preprint*, abs/1902.10909.
- 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.
- Mai Hoang Dao, Thanh 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.
- Emily Dinan, Gavin Abercrombie, A. Stevie Bergman, Shannon L. Spruit, Dirk Hovy, Y-Lan Boureau, and Verena Rieser. 2021. [Anticipating safety issues in E2E conversational AI: framework and tooling](#). *ArXiv preprint*, abs/2107.03451.
- Bosheng Ding, Junjie Hu, Lidong Bing, Mahani Aljunied, Shafiq Joty, Luo Si, and Chunyan Miao. 2022. [GlobalWoZ: Globalizing MultiWoZ 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.
- 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 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.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. [Beyond english-centric multilingual machine translation](#). *J. Mach. Learn. Res.*, 22:107:1–107:48.
- 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.
- Jack FitzGerald, Christopher Hench, Charith Peris, Scott Mackie, Kay Rottmann, Ana Sanchez, Aaron Nash, Liam Urbach, Vishesh Kakarala, Richa Singh, et al. 2022. [Massive: A 1m-example multilingual natural language understanding dataset with 51 typologically-diverse languages](#). *ArXiv preprint*, abs/2204.08582.
- Rashmi Gangadharaiah and Balakrishnan Narayanaswamy. 2019. [Joint multiple intent detection and slot labeling for goal-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 564–569, Minneapolis, Minnesota. Association for Computational Linguistics.
- M. Gašić, M. Henderson, B. Thomson, P. Tsiakoulis, and S. Young. 2012. [Policy optimisation of pomdp-based dialogue systems without state space compression](#). In *2012 IEEE Spoken Language Technology Workshop (SLT)*, pages 31–36.
- Daniela Gerz, Pei-Hao Su, Razvan Kusztos, Avishek Mondal, Michał Lis, Eshan Singhal, Nikola Mrkšić, Tsung-Hsien Wen, and Ivan Vulić. 2021. [Multilingual and cross-lingual intent detection from spoken data](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7468–7475, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’ Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022. [The Flores-101 evaluation benchmark for low-resource and multilingual machine translation](#). *Transactions of the Association for Computational Linguistics*, 10:522–538.
- N. Gupta, G. Tur, D. Hakkani-Tur, S. Bangalore, G. Ricciardi, and M. Gilbert. 2006. [The AT&T spoken language understanding system](#). *IEEE Transactions on Audio, Speech, and Language Processing*, 14(1):213–222.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: decoding-enhanced bert with disentangled attention](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- 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 pre-training 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.
- Masahiro Kaneko, Aizhan Imankulova, Danushka Bollegala, and Naoaki Okazaki. 2022. [Gender bias in masked language models for multiple languages](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2740–2750, Seattle, United States. Association for Computational Linguistics.
- Tom Kocmi, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Thamme Gowda, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Rebecca Knowles, Philipp Koehn, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Michal Novák, Martin Popel, and Maja Popović. 2022. [Findings of the 2022 conference on machine translation \(WMT22\)](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 1–45, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Stefan Larson and Kevin Leach. 2022. [A survey of intent classification and slot-filling datasets for task-oriented dialog](#). *ArXiv preprint*, abs/2207.13211.
- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. [From zero to hero: On the limitations of zero-shot language transfer with multilingual Transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499, Online. Association for Computational Linguistics.
- Bai Li, Nanyi Jiang, Joey Sham, Henry Shi, and Hussein Fazal. 2019. [Real-world conversational ai for hotel bookings](#). In *2019 Second International Conference on Artificial Intelligence for Industries (AI4I)*, pages 58–62.
- Fangyu Liu, Emanuele Bugliarello, Edoardo Maria Ponti, Siva Reddy, Nigel Collier, and Desmond Elliott. 2021a. [Visually grounded reasoning across languages and cultures](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10467–10485, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021b. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *ArXiv preprint*, abs/2107.13586.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Alexandre Magueresse, Vincent Carles, and Evan Heetderks. 2020. [Low-resource languages: A review of past work and future challenges](#). *ArXiv preprint*, abs/2006.07264.
- Olga Majewska, Evgeniia Razumovskaia, Edoardo Maria Ponti, Ivan Vulić, and Anna Korhonen. 2022. [Cross-lingual dialogue dataset creation via outline-based generation](#). *ArXiv preprint*, abs/2201.13405.
- Nikita Moghe, Mark Steedman, and Alexandra Birch. 2021. [Cross-lingual intermediate fine-tuning improves dialogue state tracking](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1137–1150, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mahdi Namazifar, Alexandros Papangelis, Gökhan Tür, and Dilek Hakkani-Tür. 2021. [Language model is all you need: Natural language understanding as question answering](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2021, Toronto, ON, Canada, June 6-11, 2021*, pages 7803–7807. IEEE.
- 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.
- P. J. Price. 1990. [Evaluation of spoken language systems: the ATIS domain](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- Libo Qin, Xiao Xu, Wanxiang Che, and Ting Liu. 2020. [AGIF: An adaptive graph-interactive framework for joint multiple intent detection and slot filling](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1807–1816, Online. Association for Computational Linguistics.

- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don't know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Antoine Raux, Brian Langner, Alan W. Black, and Maxine Eskénazi. 2003. [Let's go: improving spoken dialog systems for the elderly and non-natives](#). *8th European Conference on Speech Communication and Technology (Eurospeech 2003)*.
- Antoine Raux, Brian Langner, Dan Bohus, Alan W. Black, and Maxine Eskenazi. 2005. [Let's go public! taking a spoken dialog system to the real world](#). In *9th European Conference on Speech Communication and Technology, Lisbon, Portugal*, pages 432–437.
- Evgeniia Razumovskaia, Goran Glavas, Olga Majewska, Edoardo M Ponti, Anna Korhonen, and Ivan Vulic. 2022a. [Crossing the conversational chasm: A primer on natural language processing for multilingual task-oriented dialogue systems](#). *Journal of Artificial Intelligence Research*, 74:1351–1402.
- Evgeniia Razumovskaia, Ivan Vulić, and Anna Korhonen. 2022b. [Data augmentation and learned layer aggregation for improved multilingual language understanding in dialogue](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2017–2033, Dublin, Ireland. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). 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 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2020. [Making monolingual sentence embeddings multilingual using knowledge distillation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4512–4525, Online. Association for Computational Linguistics.
- 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.
- Antonio Toral. 2019. [Post-editeese: an exacerbated translationese](#). In *Proceedings of Machine Translation Summit XVII: Research Track*, pages 273–281, Dublin, Ireland. European Association for Machine Translation.
- Gokhan Tur, Dilek Hakkani-Tür, and Larry Heck. 2010. [What is left to be understood in ATIS?](#) In *2010 IEEE Spoken Language Technology Workshop*, pages 19–24. IEEE.
- 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.
- Rob van der Goot, Ibrahim Sharaf, Aizhan Imankulova, Ahmet Üstün, Marija Stepanović, Alan Ramponi, Siti Oryza Khairunnisa, Mamoru Komachi, and Barbara Plank. 2021. [From masked language modeling to translation: Non-English auxiliary tasks improve zero-shot spoken language understanding](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2479–2497, Online. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface's transformers: State-of-the-art natural language processing](#). *ArXiv preprint*, abs/1910.03771.
- Shijie Wu, Benjamin Van Durme, and Mark Dredze. 2022. [Zero-shot cross-lingual transfer is underspecified optimization](#). In *Proceedings of the 7th Workshop on Representation Learning for NLP*, pages 236–248, Dublin, Ireland. Association for Computational Linguistics.
- Mengzhou Xia, Guoqing Zheng, Subhabrata Mukherjee, Milad Shokouhi, Graham Neubig, and Ahmed Hassan Awadallah. 2021. [MetaXL: Meta representation transformation for low-resource cross-lingual learning](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 499–511, Online. Association for Computational Linguistics.
- Weijia Xu, Batool 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.
- 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.

Steve J. Young, Milica Gasic, Blaise Thomson, and Jason D. Williams. 2013. [POMDP-Based Statistical Spoken Dialog Systems: A Review](#). *Proceedings of the IEEE*, 101(5):1160–1179.

Mengjie Zhao and Hinrich Schütze. 2021. [Discrete and soft prompting for multilingual models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8547–8555, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Lei Zuo, Kun Qian, Bowen Yang, and Zhou Yu. 2021. [Allwoz: Towards multilingual task-oriented dialog systems for all](#). *ArXiv preprint*, abs/2112.08333.

## A Language Codes

Language codes which are used in the paper are provided in Table 8.

| Language code | Language |
|---------------|----------|
| EN            | English  |
| AM            | Amharic  |
| MR            | Marathi  |
| ES            | Spanish  |
| TR            | Turkish  |

Table 8: Language codes used in the paper

## B Translation Guidelines

We are an academic team interested in evaluating modern automatic machine translation systems for building better multilingual chatbots. The sentences that you will create should be colloquial. We do not want exact translations but rather what would have been your utterance if you were to say the given content in Spanish. Please do not use any form of machine translation during the process.

The sentences in this document are spoken to a customer service bot that provides banking services (e.g., making transfers, depositing cheques, reporting lost cards, requesting mortgage information) and hotel ‘bell desk’ reception tasks (e.g., booking rooms, asking about pools or gyms, requesting room service).

Please translate the sentences under ‘source\_text’. Please write the translation under ‘target\_text’. In the translation please:

- maintain the meaning and style as close to English text as possible;  
Example: “Exactly, it was declined”  
It is important that the colloquial word “Exactly” is reflected in the translation.

- if the example includes pronouns (e.g., “this one”), maintain the pronouns in the translation as well.
- you are encouraged to translate the proper names and time values in the most natural form for the target language.  
Example: “cancel the one at 2:35 p.m. on Dec 8th”  
The “2:35 p.m.” can be translated in any way time is usually expressed in the target language, e.g., “25 minutes to 3” or “2:35 in the day”, if “p.m.” is non-existent in the target language.
- If there is no exact translation for a concept or the concept is absent from the culture, feel free to substitute it with a description of the concept or a similar concept familiar to the population.  
Example: “book a hotel via Booking.com”  
If there is no access to Booking com, feel free to substitute it with “book a hotel by phone”.

You may observe that some of the sentences have some spans under slot\_1, slot\_2, slot\_3 and so on. After you have written the Spanish sentence in the target\_text, please replace the values under this column with the corresponding spans in your Spanish sentence in the respective slot columns

1. For example, in the sentence “how much more have I spent on take out since last week?”, there is “last week” under slot\_1 and “take out” under slot\_2. We would expect you to copy the corresponding phrases (i.e: the translations for last week and take out respectively) from your written Spanish sentence in the ‘slot\_1’ and ‘slot\_2’. If there is no exact phrase that matches the span from English, copy its equivalent in the columns.
2. Please do not change the order of values while writing the corresponding value columns in the target sentence.

## C Dataset Collection Details

Our annotators for Spanish are based in Spain, Amharic are based in Ethiopia, Marathi are based in India, and Turkish are based in Turkey. The rates for translation were fixed by the translators or the translation agencies, ensuring fair pay for the translators. Our internal annotators were compensated

| Type                | Platform | Language | Cost (£) |
|---------------------|----------|----------|----------|
| Translation         | Proz.com | Spanish  | 1,500    |
| Translation         | BLEND    | Marathi  | 2,031    |
| Translation         | BLEND    | Amharic  | 2,036    |
| Translation         | BLEND    | Turkish  | 1,954    |
| Internal Evaluation | N/A      | All      | 103      |
| Total               |          |          | 7,624    |

Table 9: Dataset collection cost breakdown

with £15/hour for their work. We provide details on the dataset collection costs in Table 9.

## D Full Results for MLP-Based Intent Detection Baselines

We provide full results for the *MLP-based* baseline in tables 12 to 17. Our implementation uses the Transformers library (Wolf et al., 2019) and the SentenceBERT library (Reimers and Gurevych, 2019). The models were trained on NVIDIA Titan xP GPUs. Approximate training times for every fold are provided in Table 10.

| Domain  | 20  | 10   | LARGE |
|---------|-----|------|-------|
| BANKING | 1.5 | 5    | 25    |
| HOTELS  | 1   | 2    | 17    |
| ALL     | 3   | 10.5 | 40    |

Table 10: Approximate GPU training times for *MLP-based* setup for every fold (in mins).

| Domain  | 20   | 10   | LARGE |
|---------|------|------|-------|
| BANKING | 1    | 2    | 8     |
| HOTELS  | 0.33 | 0.75 | 33    |
| ALL     | 2    | 4    | 15    |

Table 11: Approximate GPU training times for *QA-based* setup for every fold (in hours).

## E In-Domain Cross-Lingual Results

We compare the in-domain cross-lingual results using mpnet and LaBSE as an underlying encoder for all domains in Figures 3 and 4, respectively.

## F Full Results for QA-Based Intent Detection

We provide full results for the QA-based baseline in tables 18 to 22.

For both slot labelling and intent detection, we performed the hyperparameter search over fold 0 of the HOTELS domain in the 20-fold setup. The hyperparameters varied include learning rate [1e-5, 2e-5, 3e-5], batch size [4,8,16], epochs [5,10,15]

and multilingual models of XLM-R and mDeBERTa. The other hyperparameters are the same as (Casanueva et al., 2022). Approximate training times for every fold on NVIDIA GeForce RTX 2080 are provided in Table 11. Our implementation of the QA model uses the QA fine-tuning scripts from the Transformers Library (Wolf et al., 2019).

## G Full Results for Slot Labelling Baselines

We provide full results for slot labelling for *token classification* setup in Tables 23 and 24. The results for the *QA-based* setup are provided in Tables 25 and 26.

| Model   | BANKING     |              |             | HOTELS      |             |             | ALL         |             |             |
|---------|-------------|--------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
|         | 20-Fold     | 10-Fold      | Large       | 20-Fold     | 10-Fold     | Large       | 20-Fold     | 10-Fold     | Large       |
| ENGLISH |             |              |             |             |             |             |             |             |             |
| LaBSE   | 51.5        | 64.3         | 87.1        | 45.1        | 55.4        | 80.4        | 51.8        | 63.4        | 85.2        |
| mpnet   | <b>58.1</b> | <b>68.8</b>  | <b>87.2</b> | <b>51.9</b> | <b>61.8</b> | <b>81.2</b> | <b>57.7</b> | <b>68.1</b> | <b>86.1</b> |
| AMHARIC |             |              |             |             |             |             |             |             |             |
| LaBSE   | <b>46.2</b> | <b>56.60</b> | <b>77.3</b> | <b>38.6</b> | <b>48.8</b> | <b>71.9</b> | <b>46.4</b> | <b>56.1</b> | <b>76.1</b> |
| mpnet   | 37.0        | 44.1         | 64.7        | 31.9        | 39.9        | 60.6        | 36.2        | 43.1        | 63.5        |
| MARATHI |             |              |             |             |             |             |             |             |             |
| LaBSE   | <b>50.9</b> | <b>63.5</b>  | <b>84.1</b> | 43.5        | <b>53.5</b> | <b>77.0</b> | <b>50.8</b> | <b>61.9</b> | <b>81.9</b> |
| mpnet   | 50.0        | 58.8         | 80.0        | <b>43.5</b> | 53.1        | 73.4        | 49.2        | 57.8        | 78.0        |
| SPANISH |             |              |             |             |             |             |             |             |             |
| LaBSE   | 51.4        | 64.0         | <b>85.2</b> | 46.4        | 56.6        | 79.5        | 52.2        | 63.5        | <b>83.8</b> |
| mpnet   | <b>55.2</b> | <b>65.2</b>  | 84.4        | <b>50.9</b> | <b>61.1</b> | <b>80.3</b> | <b>55.4</b> | <b>65.1</b> | 83.5        |
| TURKISH |             |              |             |             |             |             |             |             |             |
| LaBSE   | 49.2        | <b>61.6</b>  | <b>82.8</b> | 45.2        | 55.9        | <b>80.1</b> | 50.1        | 61.3        | <b>82.0</b> |
| mpnet   | <b>51.2</b> | 60.3         | 80.7        | <b>49.9</b> | <b>59.6</b> | 79.1        | <b>52.3</b> | <b>61.3</b> | 80.9        |

Table 12: In-language in-domain results for intent detection with MLP-based setup ( $F_1 \times 100$ ).

| SRC | TGT | BANKING |         |       | HOTELS  |         |       | ALL     |         |       |
|-----|-----|---------|---------|-------|---------|---------|-------|---------|---------|-------|
|     |     | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| en  | am  | 38.9    | 49.6    | 71.0  | 29.8    | 38.3    | 65.1  | 37.7    | 47.5    | 70.2  |
|     | mr  | 34.1    | 47.4    | 77.6  | 29.2    | 38.5    | 69.1  | 32.8    | 45.1    | 75.9  |
|     | es  | 36.2    | 50.1    | 80.8  | 34.3    | 44.2    | 74.8  | 35.9    | 49.2    | 80.3  |
|     | tr  | 29.4    | 42.0    | 72.9  | 27.4    | 36.2    | 72.0  | 28.7    | 41.1    | 74.2  |
| am  | en  | 48.3    | 57.7    | 78.8  | 43.5    | 51.2    | 73.7  | 48.4    | 56.9    | 77.8  |
|     | mr  | 38.9    | 51.7    | 76.0  | 33.4    | 44.0    | 70.1  | 39.1    | 50.8    | 74.4  |
|     | es  | 40.1    | 52.6    | 76.5  | 37.4    | 48.1    | 73.3  | 41.5    | 53.0    | 76.2  |
|     | tr  | 33.8    | 46.5    | 73.2  | 31.2    | 42.2    | 70.8  | 34.5    | 46.4    | 73.0  |
| mr  | en  | 51.3    | 59.8    | 79.4  | 47.2    | 54.0    | 74.6  | 49.6    | 58.0    | 78.7  |
|     | am  | 48.5    | 57.1    | 71.3  | 41.7    | 50.2    | 69.8  | 47.8    | 55.8    | 70.3  |
|     | es  | 50.0    | 61.4    | 79.3  | 45.5    | 55.1    | 76.6  | 50.8    | 60.6    | 78.2  |
|     | tr  | 43.8    | 55.4    | 74.9  | 41.5    | 50.8    | 75.0  | 45.2    | 55.4    | 75.1  |
| es  | en  | 52.2    | 60.6    | 80.7  | 48.7    | 55.4    | 74.3  | 50.9    | 59.0    | 80.2  |
|     | am  | 47.7    | 56.5    | 70.7  | 40.8    | 49.2    | 68.8  | 47.3    | 55.7    | 70.1  |
|     | mr  | 47.1    | 59.0    | 78.9  | 40.7    | 50.7    | 73.3  | 47.3    | 58.0    | 77.2  |
|     | tr  | 43.1    | 55.6    | 76.5  | 39.9    | 50.3    | 76.2  | 44.1    | 55.8    | 76.7  |
| tr  | en  | 49.7    | 56.8    | 75.7  | 47.3    | 53.7    | 73.9  | 48.0    | 55.5    | 76.5  |
|     | am  | 48.2    | 56.1    | 70.6  | 42.4    | 50.7    | 66.8  | 47.4    | 54.9    | 69.5  |
|     | mr  | 49.1    | 59.4    | 76.6  | 44.4    | 53.9    | 74.1  | 49.3    | 58.5    | 75.4  |
|     | es  | 51.3    | 61.5    | 78.5  | 48.0    | 57.6    | 78.1  | 52.3    | 61.5    | 78.7  |

Table 13: Cross-lingual in-domain results for intent detection with MLP-based and LaBSE as the fixed sentence encoder ( $F_1 \times 100$ ).

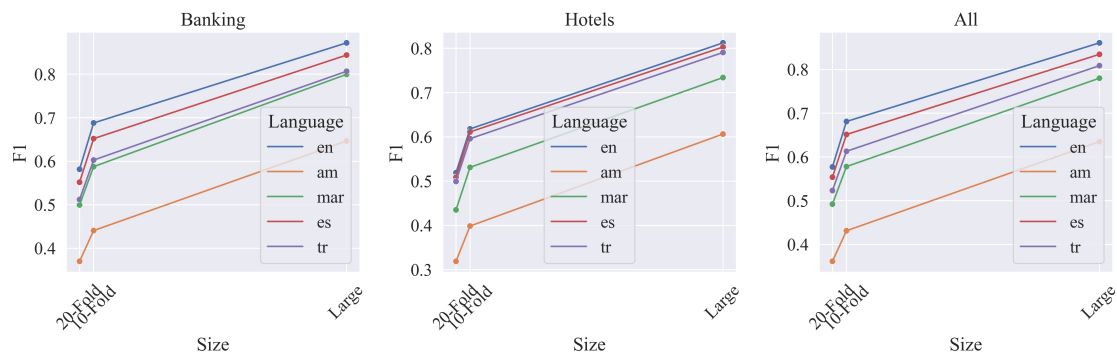


Figure 3: Comparison of in-domain in-lingual results. The results are for *MLP-based* baseline with mpnet as the underlying encoder. English is used as the source language in the experiments.



| SRC | TGT | BANKING |         |       | HOTELS  |         |       | ALL     |         |       |
|-----|-----|---------|---------|-------|---------|---------|-------|---------|---------|-------|
|     |     | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| en  | am  | 24.8    | 28.6    | 32.9  | 12.9    | 17.5    | 23.0  | 22.0    | 25.6    | 30.2  |
|     | mr  | 42.0    | 49.6    | 63.7  | 37.5    | 46.1    | 62.2  | 41.0    | 48.3    | 61.4  |
|     | es  | 50.1    | 59.1    | 75.3  | 45.8    | 55.5    | 75.3  | 49.2    | 58.5    | 75.1  |
|     | tr  | 43.9    | 52.1    | 66.7  | 44.8    | 54.1    | 73.5  | 44.5    | 52.8    | 68.0  |
| am  | en  | 37.2    | 41.3    | 51.2  | 35.6    | 41.5    | 57.4  | 36.3    | 39.8    | 50.8  |
|     | mr  | 37.4    | 42.4    | 54.2  | 34.7    | 40.6    | 57.1  | 36.4    | 40.8    | 53.5  |
|     | es  | 38.9    | 44.5    | 56.9  | 36.0    | 42.7    | 60.4  | 38.3    | 43.1    | 55.9  |
|     | tr  | 38.1    | 43.5    | 55.4  | 35.8    | 42.3    | 60.6  | 37.7    | 42.4    | 55.3  |
| mr  | en  | 51.2    | 58.8    | 75.5  | 48.5    | 57.4    | 75.4  | 51.2    | 58.9    | 75.9  |
|     | am  | 32.3    | 37.4    | 47.8  | 20.5    | 28.1    | 43.1  | 30.6    | 35.2    | 46.2  |
|     | es  | 50.3    | 58.0    | 75.5  | 46.7    | 56.2    | 76.5  | 50.3    | 58.3    | 75.8  |
|     | tr  | 47.2    | 54.8    | 72.0  | 46.5    | 55.2    | 74.9  | 48.2    | 55.7    | 73.4  |
| es  | en  | 55.9    | 64.7    | 82.6  | 52.4    | 61.5    | 78.7  | 56.2    | 64.6    | 81.6  |
|     | am  | 30.8    | 36.4    | 45.5  | 16.8    | 23.6    | 37.5  | 28.7    | 34.0    | 43.1  |
|     | mr  | 46.8    | 54.9    | 70.4  | 41.8    | 51.3    | 68.2  | 46.5    | 54.4    | 69.4  |
|     | tr  | 49.4    | 58.1    | 74.8  | 48.4    | 58.2    | 77.9  | 50.4    | 59.1    | 75.6  |
| tr  | en  | 52.8    | 60.6    | 77.8  | 52.3    | 61.3    | 77.6  | 53.6    | 61.3    | 78.1  |
|     | am  | 32.5    | 37.9    | 48.1  | 18.1    | 24.8    | 39.5  | 30.4    | 35.7    | 44.6  |
|     | mr  | 46.8    | 54.6    | 70.9  | 42.5    | 51.7    | 68.1  | 46.5    | 54.4    | 70.1  |
|     | es  | 52.3    | 60.8    | 78.7  | 49.8    | 59.5    | 78.3  | 53.0    | 61.4    | 79.3  |

Table 14: Cross-lingual in-domain results for intent detection with MLP-based and mpnet as the fixed sentence encoder ( $F_1 \times 100$ ).

| Model   | BANKING     |             |             | HOTELS      |             |             |
|---------|-------------|-------------|-------------|-------------|-------------|-------------|
|         | 20-Fold     | 10-Fold     | Large       | 20-Fold     | 10-Fold     | Large       |
| ENGLISH |             |             |             |             |             |             |
| LaBSE   | 10.6        | 15.8        | <b>35.2</b> | <b>23.0</b> | <b>29.9</b> | <b>43.4</b> |
| mpnet   | <b>11.4</b> | <b>15.8</b> | 31.4        | 22.0        | 28.3        | 39.7        |
| AMHARIC |             |             |             |             |             |             |
| LaBSE   | 8.6         | <b>13.6</b> | <b>32.1</b> | <b>22.4</b> | <b>28.5</b> | <b>38.6</b> |
| mpnet   | <b>9.8</b>  | 12.8        | 23.6        | 15.8        | 19.0        | 26.8        |
| MARATHI |             |             |             |             |             |             |
| LaBSE   | 7.2         | 12.1        | <b>30.1</b> | <b>21.8</b> | <b>28.6</b> | <b>40.5</b> |
| mpnet   | <b>10.9</b> | <b>15.2</b> | 28.1        | 20.7        | 25.0        | 32.5        |
| SPANISH |             |             |             |             |             |             |
| LaBSE   | 8.2         | 13.7        | <b>33.9</b> | <b>22.4</b> | <b>29.7</b> | <b>42.0</b> |
| mpnet   | <b>11.5</b> | <b>16.5</b> | 31.7        | 21.8        | 26.6        | 37.7        |
| TURKISH |             |             |             |             |             |             |
| LaBSE   | 7.6         | 12.6        | <b>32.8</b> | <b>22.1</b> | <b>28.8</b> | <b>42.2</b> |
| mpnet   | <b>11.1</b> | <b>16.0</b> | 31.3        | 21.9        | 26.4        | 35.1        |

Table 15: In-language cross-domain results for intent detection with MLP-based setup ( $F_1 \times 100$ ).

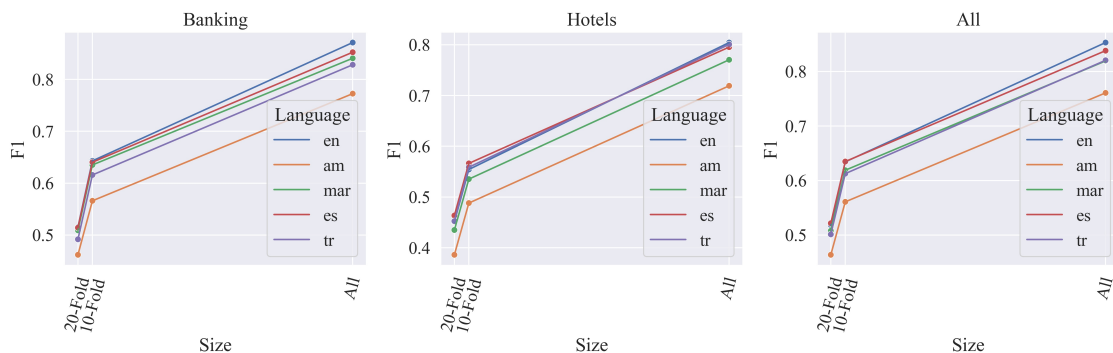


Figure 4: Comparison of in-domain in-lingual results. The results are for *MLP-based* baseline with LaBSE as the underlying encoder. English is used as the source language in the experiments.

| SRC | TGT | BANKING |         |       | HOTELS  |         |       |
|-----|-----|---------|---------|-------|---------|---------|-------|
|     |     | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| en  | am  | 6.8     | 10.4    | 29.0  | 19.3    | 25.9    | 40.8  |
|     | mr  | 3.8     | 6.1     | 23.9  | 15.1    | 20.8    | 39.1  |
|     | es  | 4.3     | 7.2     | 26.4  | 16.0    | 22.5    | 41.0  |
|     | tr  | 2.7     | 4.5     | 22.7  | 13.3    | 19.2    | 40.8  |
| am  | en  | 12.4    | 18.0    | 33.5  | 22.6    | 27.2    | 37.5  |
|     | mr  | 4.7     | 8.92    | 27.14 | 18.54   | 25.04   | 36.05 |
|     | es  | 5.5     | 10.4    | 29.0  | 19.3    | 26.0    | 38.1  |
|     | tr  | 3.4     | 6.9     | 25.2  | 16.1    | 23.3    | 38.8  |
| mr  | en  | 14.1    | 19.3    | 33.1  | 24.4    | 29.0    | 38.2  |
|     | am  | 11.0    | 16.3    | 31.5  | 24.1    | 29.7    | 39.3  |
|     | es  | 8.1     | 13.5    | 30.5  | 22.9    | 29.3    | 39.9  |
|     | tr  | 5.5     | 9.9     | 27.9  | 19.4    | 26.6    | 40.8  |
| es  | en  | 14.6    | 19.6    | 32.8  | 25.0    | 29.3    | 38.6  |
|     | am  | 11.0    | 16.4    | 32.8  | 24.0    | 30.0    | 40.3  |
|     | mr  | 7.0     | 11.7    | 32.1  | 20.6    | 27.4    | 40.7  |
|     | tr  | 5.3     | 9.6     | 29.6  | 18.9    | 26.7    | 41.3  |
| tr  | en  | 15.9    | 20.5    | 34.0  | 24.0    | 27.8    | 36.8  |
|     | am  | 13.3    | 18.6    | 33.8  | 24.4    | 29.2    | 38.7  |
|     | mr  | 9.3     | 15.0    | 34.1  | 22.4    | 28.3    | 38.7  |
|     | es  | 10.5    | 16.7    | 35.4  | 23.9    | 30.0    | 40.4  |

Table 16: Cross-lingual cross-domain results for intent detection with MLP-based and LaBSE as the fixed sentence encoder ( $F_1 \times 100$ ).

| SRC | TGT | BANKING |         |       | HOTELS  |         |       |
|-----|-----|---------|---------|-------|---------|---------|-------|
|     |     | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| en  | am  | 6.5     | 8.7     | 16.2  | 11.2    | 14.4    | 18.5  |
|     | mr  | 9.5     | 13.1    | 25.9  | 18.9    | 24.1    | 31.9  |
|     | es  | 9.9     | 13.6    | 28.4  | 20.1    | 25.4    | 35.9  |
|     | tr  | 9.2     | 12.9    | 25.9  | 19.7    | 25.3    | 33.5  |
| am  | en  | 11.9    | 14.3    | 23.1  | 15.7    | 17.2    | 18.9  |
|     | mr  | 11.4    | 14.4    | 24.2  | 17.0    | 18.9    | 22.2  |
|     | es  | 11.7    | 14.3    | 24.8  | 16.7    | 18.5    | 21.5  |
|     | tr  | 11.3    | 14.0    | 25.0  | 16.9    | 19.0    | 22.3  |
| mr  | en  | 11.9    | 16.1    | 27.4  | 20.7    | 24.0    | 30.3  |
|     | am  | 7.9     | 10.9    | 21.5  | 15.1    | 18.9    | 25.7  |
|     | es  | 11.0    | 15.4    | 28.6  | 20.8    | 24.8    | 32.3  |
|     | tr  | 10.6    | 15.0    | 27.6  | 20.9    | 24.8    | 32.3  |
| es  | en  | 12.3    | 17.1    | 30.9  | 22.0    | 26.0    | 36.0  |
|     | am  | 7.4     | 11.0    | 23.2  | 15.0    | 19.7    | 28.3  |
|     | mr  | 10.8    | 15.6    | 30.4  | 21.2    | 26.1    | 35.8  |
|     | tr  | 10.7    | 15.7    | 30.3  | 21.8    | 26.7    | 36.8  |
| tr  | en  | 12.6    | 17.3    | 30.7  | 21.5    | 25.4    | 32.2  |
|     | am  | 8.1     | 11.6    | 24.3  | 15.6    | 19.7    | 27.3  |
|     | mr  | 11.1    | 15.8    | 30.9  | 21.2    | 25.4    | 33.8  |
|     | es  | 11.6    | 16.3    | 31.7  | 21.8    | 26.1    | 33.7  |

Table 17: Cross-lingual cross-domain results for intent detection with MLP-based and mpnet as the fixed sentence encoder ( $F_1 \times 100$ ).

| Model          | BANKING |         |       | HOTELS  |         |       | ALL     |         |       |
|----------------|---------|---------|-------|---------|---------|-------|---------|---------|-------|
|                | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| <b>ENGLISH</b> |         |         |       |         |         |       |         |         |       |
| mDeBERTa       | 80.8    | 85.0    | 93.2  | 66.9    | 73.2    | 87.3  | 79.6    | 83.6    | 91.7  |
| <b>AMHARIC</b> |         |         |       |         |         |       |         |         |       |
| mDeBERTa       | 56.1    | 64.8    | 83.4  | 51.1    | 57.3    | 77.9  | 57.2    | 64.2    | 81.8  |
| <b>MARATHI</b> |         |         |       |         |         |       |         |         |       |
| mDeBERTa       | 70.4    | 77.4    | 88.6  | 53.2    | 63.2    | 81.7  | 69.1    | 75.5    | 86.8  |
| <b>SPANISH</b> |         |         |       |         |         |       |         |         |       |
| mDeBERTa       | 78.8    | 83.5    | 91.6  | 65.9    | 72.7    | 86.0  | 77.7    | 82.1    | 89.9  |
| <b>TURKISH</b> |         |         |       |         |         |       |         |         |       |
| mDeBERTa       | 75.8    | 80.2    | 90.1  | 65.6    | 72.8    | 86.7  | 75.0    | 79.5    | 89.2  |

Table 18: In-language in-domain results for intent detection with QA-based setup and mDeBERTa as the base model ( $F_1 \times 100$ ).

| SRC | TGT | BANKING |         |       | HOTELS  |         |       | ALL     |         |       |
|-----|-----|---------|---------|-------|---------|---------|-------|---------|---------|-------|
|     |     | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| en  | am  | 42.7    | 46.7    | 52.1  | 37.1    | 38.0    | 45.2  | 40.7    | 41.2    | 49.5  |
|     | mr  | 58.4    | 63.2    | 71.0  | 51.1    | 55.6    | 66.0  | 56.5    | 59.1    | 66.6  |
|     | es  | 66.2    | 71.3    | 82.4  | 62.2    | 69.1    | 84.1  | 67.9    | 71.9    | 81.2  |
|     | tr  | 56.5    | 62.2    | 73.0  | 59.6    | 65.3    | 79.6  | 58.2    | 62.1    | 73.6  |
| am  | en  | 30.1    | 37.9    | 55.0  | 30.5    | 38.6    | 60.0  | 29.4    | 32.4    | 53.1  |
|     | mr  | 31.3    | 40.6    | 53.7  | 27.5    | 34.2    | 55.3  | 32.3    | 35.4    | 47.8  |
|     | es  | 27.9    | 35.5    | 53.1  | 31.5    | 39.7    | 61.2  | 27.7    | 31.2    | 52.2  |
|     | tr  | 28.5    | 38.2    | 54.3  | 33.0    | 41.3    | 62.5  | 30.7    | 35.0    | 52.4  |
| mr  | en  | 49.7    | 55.0    | 69.2  | 44.3    | 52.4    | 68.0  | 47.9    | 51.3    | 66.6  |
|     | am  | 32.4    | 35.6    | 48.4  | 27.1    | 31.0    | 44.4  | 30.4    | 34.0    | 43.5  |
|     | es  | 46.4    | 50.7    | 65.8  | 43.8    | 51.5    | 66.2  | 45.6    | 49.4    | 63.7  |
|     | tr  | 44.9    | 50.1    | 62.9  | 44.3    | 52.2    | 67.6  | 45.4    | 49.5    | 61.5  |
| es  | en  | 67.6    | 74.1    | 84.6  | 62.9    | 69.1    | 84.6  | 67.9    | 73.9    | 84.2  |
|     | am  | 41.6    | 45.4    | 53.7  | 35.7    | 38.1    | 47.5  | 40.1    | 43.6    | 51.2  |
|     | mr  | 54.1    | 60.2    | 66.7  | 49.6    | 53.8    | 62.6  | 52.8    | 56.6    | 62.7  |
|     | tr  | 52.3    | 58.9    | 70.5  | 58.1    | 63.5    | 78.5  | 55.7    | 61.4    | 71.7  |
| tr  | en  | 52.78   | 60.57   | 77.78 | 52.27   | 61.25   | 77.63 | 53.61   | 61.25   | 78.13 |
|     | am  | 32.46   | 37.88   | 48.12 | 18.05   | 24.75   | 39.54 | 30.37   | 35.72   | 44.58 |
|     | mr  | 46.83   | 54.59   | 70.90 | 42.46   | 51.65   | 68.06 | 46.5    | 54.35   | 70.11 |
|     | es  | 52.30   | 60.83   | 78.73 | 49.76   | 59.48   | 78.25 | 52.98   | 61.44   | 79.27 |

Table 19: Cross-lingual in-domain results for intent detection with QA-based setup and mDeBERTa as the base model ( $F_1 \times 100$ ).

| SRC | TGT | BANKING |         |       | HOTELS  |         |       | ALL     |         |       |
|-----|-----|---------|---------|-------|---------|---------|-------|---------|---------|-------|
|     |     | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| en  | am  | 17.6    | 21.4    | 25.4  | 11.5    | 13.7    | 20.2  | 16.2    | 18.9    | 23.9  |
|     | mr  | 63.5    | 68.1    | 75.8  | 57.2    | 62.7    | 76.8  | 63.3    | 67.8    | 75.4  |
|     | es  | 67.7    | 72.3    | 81.1  | 59.9    | 66.5    | 80.9  | 67.8    | 71.6    | 79.8  |
|     | tr  | 45.3    | 52.3    | 63.8  | 42.6    | 47.4    | 67.3  | 46.3    | 52.1    | 64.1  |
| am  | en  | 20.2    | 24.0    | 32.9  | 12.2    | 13.2    | 21.2  | 16.2    | 19.3    | 28.1  |
|     | mr  | 18.8    | 22.2    | 31.2  | 11.7    | 12.4    | 18.2  | 15.2    | 18.1    | 27.7  |
|     | es  | 14.2    | 16.6    | 24.1  | 11.4    | 12.3    | 21.5  | 11.7    | 13.7    | 22.3  |
|     | tr  | 13.8    | 15.3    | 20.7  | 11.8    | 12.5    | 21.3  | 11.7    | 13.3    | 19.5  |

Table 20: Translate-test in-domain results for intent detection with QA-based setup and mDeBERTa as the base model ( $F_1 \times 100$ ).

| Model    | BANKING |         |       | HOTELS  |         |       |
|----------|---------|---------|-------|---------|---------|-------|
|          | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| ENGLISH  |         |         |       |         |         |       |
| mDeBERTa | 58.3    | 63.6    | 72.8  | 55.9    | 59.0    | 61.8  |
| AMHARIC  |         |         |       |         |         |       |
| mDeBERTa | 41.8    | 47.0    | 56.7  | 35.6    | 39.5    | 49.0  |
| MARATHI  |         |         |       |         |         |       |
| mDeBERTa | 47.0    | 51.1    | 59.3  | 42.3    | 46.5    | 58.0  |
| SPANISH  |         |         |       |         |         |       |
| mDeBERTa | 62.1    | 66.6    | 72.0  | 51.3    | 56.6    | 60.1  |
| TURKISH  |         |         |       |         |         |       |
| mDeBERTa | 61.6    | 64.9    | 70.9  | 47.6    | 52.6    | 58.7  |

Table 21: In-language cross-domain results for intent detection with QA-based setup and mDeBERTa as the base model ( $F_1 \times 100$ ).

| SRC | TGT | BANKING |         |       | HOTELS  |         |       |
|-----|-----|---------|---------|-------|---------|---------|-------|
|     |     | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| en  | am  | 29.0    | 33.7    | 40.3  | 30.5    | 33.9    | 42.6  |
|     | mar | 45.0    | 50.0    | 56.6  | 40.4    | 44.4    | 52.7  |
|     | sp  | 54.2    | 61.5    | 71.6  | 49.0    | 52.0    | 59.2  |
|     | tr  | 52.1    | 59.0    | 68.1  | 42.8    | 46.1    | 53.8  |
| am  | en  | 19.7    | 27.7    | 36.3  | 19.4    | 22.1    | 33.1  |
|     | mr  | 19.6    | 27.7    | 35.5  | 19.6    | 23.6    | 35.6  |
|     | sp  | 18.9    | 26.5    | 34.6  | 17.8    | 19.0    | 30.5  |
|     | tr  | 20.8    | 31.5    | 38.2  | 16.9    | 19.5    | 32.7  |
| mr  | en  | 36.3    | 40.7    | 51.7  | 35.5    | 38.5    | 46.9  |
|     | am  | 20.5    | 24.8    | 35.6  | 21.7    | 25.4    | 34.6  |
|     | sp  | 37.5    | 40.5    | 51.9  | 34.5    | 37.5    | 46.9  |
|     | tr  | 36.5    | 41.0    | 50.0  | 31.6    | 35.8    | 43.4  |
| es  | en  | 60.4    | 65.4    | 71.1  | 46.3    | 50.9    | 53.6  |
|     | am  | 33.3    | 35.5    | 41.1  | 27.0    | 29.4    | 36.8  |
|     | mr  | 47.5    | 50.2    | 53.4  | 34.7    | 39.2    | 45.6  |
|     | tr  | 55.9    | 60.5    | 66.7  | 36.3    | 40.7    | 46.9  |
| tr  | en  | 57.0    | 61.0    | 71.0  | 42.6    | 47.5    | 55.6  |
|     | am  | 33.2    | 36.0    | 43.8  | 25.7    | 29.8    | 38.9  |
|     | mr  | 46.9    | 50.2    | 57.2  | 34.1    | 39.9    | 48.9  |
|     | es  | 56.3    | 60.6    | 69.0  | 38.2    | 42.6    | 53.0  |

Table 22: Cross-lingual cross-domain results for intent detection with QA-based setup and mDeBERTa as the base model ( $F_1 \times 100$ ).

| Model          | BANKING |         |       | HOTELS  |         |       | ALL     |         |       |
|----------------|---------|---------|-------|---------|---------|-------|---------|---------|-------|
|                | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| <b>ENGLISH</b> |         |         |       |         |         |       |         |         |       |
| XLM-R          | 37.2    | 59.7    | 85.2  | 24.2    | 54.5    | 79.0  | 43.7    | 62.9    | 83.2  |
| <b>AMHARIC</b> |         |         |       |         |         |       |         |         |       |
| XLM-R          | 27.2    | 44.9    | 69.3  | 20.0    | 34.6    | 65.2  | 33.3    | 44.4    | 68.2  |
| <b>MARATHI</b> |         |         |       |         |         |       |         |         |       |
| XLM-R          | 35.1    | 46.4    | 71.0  | 24.9    | 37.9    | 62.2  | 31.6    | 46.1    | 67.8  |
| <b>SPANISH</b> |         |         |       |         |         |       |         |         |       |
| XLM-R          | 36.5    | 47.7    | 69.9  | 29.5    | 44.5    | 65.0  | 36.6    | 51.72   | 68.9  |
| <b>TURKISH</b> |         |         |       |         |         |       |         |         |       |
| XLM-R          | 32.2    | 52.0    | 82.1  | 19.3    | 49.0    | 71.2  | 38.2    | 55.7    | 76.6  |

Table 23: In-domain in-language slot labelling performance using XLM-Roberta ( $F_1 \times 100$ )

| SRC | TGT | BANKING |         |       | HOTELS  |         |       | ALL     |         |       |
|-----|-----|---------|---------|-------|---------|---------|-------|---------|---------|-------|
|     |     | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| en  | en  | 37.2    | 59.7    | 85.2  | 24.2    | 54.5    | 79.0  | 43.6    | 62.9    | 83.2  |
|     | am  | 15.2    | 27.1    | 41.2  | 6.1     | 21.2    | 31.2  | 16.4    | 26.8    | 37.0  |
|     | mr  | 19.7    | 33.4    | 50.5  | 10.8    | 29.5    | 40.9  | 21.9    | 34.0    | 46.8  |
|     | es  | 22.8    | 29.6    | 58.5  | 14.2    | 39.4    | 56.1  | 28.3    | 43.2    | 58.8  |
|     | tr  | 18.7    | 35.4    | 56.9  | 14.3    | 37.8    | 49.8  | 26.3    | 41.0    | 54.7  |
| am  | am  | 27.2    | 44.8    | 69.2  | 20.0    | 34.5    | 65.2  | 33.2    | 44.3    | 68.2  |
|     | en  | 19.8    | 30.4    | 44.5  | 7.9     | 25.0    | 39.4  | 19.9    | 28.3    | 43.8  |
|     | mr  | 22.8    | 38.3    | 60.5  | 6.3     | 26.1    | 45.1  | 21.8    | 34.0    | 53.0  |
|     | es  | 15.5    | 25.6    | 36.5  | 5.5     | 19.5    | 33.9  | 15.4    | 23.6    | 36.6  |
|     | tr  | 20.3    | 35.7    | 50.6  | 6.7     | 27.3    | 43.3  | 21.5    | 31.8    | 47.7  |
| mr  | mr  | 35.1    | 46.4    | 71.0  | 24.9    | 37.9    | 62.2  | 31.6    | 46.1    | 67.7  |
|     | en  | 19.5    | 34.7    | 46.9  | 12.4    | 29.5    | 46.3  | 24.7    | 35.3    | 48.4  |
|     | am  | 18.6    | 30.5    | 60.0  | 6.7     | 20.9    | 42.8  | 19.2    | 31.3    | 47.5  |
|     | es  | 13.8    | 25.9    | 37.0  | 08.1    | 21.9    | 36.5  | 17.0    | 26.2    | 37.1  |
|     | tr  | 20.3    | 35.2    | 50.9  | 11.6    | 32.5    | 47.0  | 25.7    | 37.1    | 50.5  |
| es  | es  | 36.5    | 47.6    | 69.8  | 29.5    | 44.4    | 65.0  | 36.6    | 51.7    | 68.8  |
|     | en  | 33.5    | 51.0    | 73.0  | 21.9    | 46.3    | 60.3  | 37.4    | 53.3    | 69.4  |
|     | am  | 16.6    | 28.6    | 43.2  | 7.8     | 24.0    | 31.7  | 19.1    | 30.4    | 37.5  |
|     | mr  | 21.1    | 32.3    | 51.0  | 12.7    | 31.7    | 41.4  | 24.3    | 35.3    | 45.1  |
|     | tr  | 20.0    | 33.9    | 55.4  | 15.6    | 36.7    | 45.6  | 27.1    | 40.5    | 50.9  |
| tr  | tr  | 32.2    | 52.0    | 82.0  | 19.3    | 49.0    | 71.2  | 38.2    | 55.6    | 76.6  |
|     | en  | 27.4    | 43.9    | 63.9  | 15.0    | 36.4    | 55.5  | 29.3    | 43.5    | 58.8  |
|     | am  | 21.6    | 34.1    | 48.0  | 07.1    | 23.1    | 38.6  | 21.7    | 33.4    | 45.0  |
|     | mr  | 26.5    | 41.1    | 58.6  | 11.6    | 33.6    | 51.2  | 26.9    | 40.2    | 56.0  |
|     | es  | 19.6    | 31.8    | 48.8  | 09.0    | 17.6    | 40.8  | 20.8    | 31.7    | 45.6  |

Table 24: Cross-lingual in-domain slot labelling results using XLM-R ( $F_1 \times 100$ )

| Model          | BANKING |         |       | HOTELS  |         |       | ALL     |         |       |
|----------------|---------|---------|-------|---------|---------|-------|---------|---------|-------|
|                | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| <b>ENGLISH</b> |         |         |       |         |         |       |         |         |       |
| mDeBERTa       | 59.7    | 66.5    | 76.0  | 61.6    | 67.3    | 77.3  | 63.0    | 68.5    | 78.0  |
| <b>AMHARIC</b> |         |         |       |         |         |       |         |         |       |
| mDeBERTa       | 03.2    | 03.8    | 05.0  | 05.8    | 06.8    | 09.9  | 03.5    | 03.9    | 05.4  |
| <b>MARATHI</b> |         |         |       |         |         |       |         |         |       |
| mDeBERTa       | 03.2    | 03.7    | 04.9  | 05.5    | 06.3    | 08.9  | 03.3    | 03.8    | 05.1  |
| <b>SPANISH</b> |         |         |       |         |         |       |         |         |       |
| mDeBERTa       | 46.0    | 52.4    | 60.1  | 52.6    | 58.7    | 68.8  | 52.5    | 57.9    | 65.7  |
| <b>TURKISH</b> |         |         |       |         |         |       |         |         |       |
| mDeBERTa       | 59.4    | 67.1    | 84.1  | 58.7    | 65.4    | 77.7  | 62.4    | 68.1    | 79.1  |

Table 25: In-domain in-language slot labelling with QA models ( $F_1 \times 100$ )

| SRC | TGT | BANKING |         |       | HOTELS  |         |       | ALL     |         |       |
|-----|-----|---------|---------|-------|---------|---------|-------|---------|---------|-------|
|     |     | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large | 20-Fold | 10-Fold | Large |
| en  | en  | 59.7    | 66.5    | 76.0  | 61.6    | 67.3    | 77.3  | 63.0    | 68.5    | 78.0  |
|     | am  | 44.2    | 47.9    | 53.8  | 47.6    | 49.7    | 55.8  | 46.8    | 48.4    | 52.9  |
|     | mr  | 48.7    | 54.6    | 62.1  | 42.0    | 46.5    | 54.4  | 46.9    | 50.9    | 57.3  |
|     | es  | 42.9    | 48.0    | 55.1  | 49.0    | 54.1    | 60.7  | 47.7    | 51.7    | 58.1  |
|     | tr  | 52.3    | 58.4    | 69.0  | 50.9    | 55.8    | 63.0  | 53.6    | 58.4    | 65.2  |
| am  | en  | 2.5     | 2.8     | 3.6   | 4.1     | 4.3     | 6.1   | 2.4     | 2.6     | 3.4   |
|     | am  | 3.2     | 3.8     | 5.0   | 5.8     | 6.8     | 9.9   | 3.5     | 3.9     | 5.4   |
|     | mr  | 2.5     | 2.9     | 3.6   | 2.6     | 3.1     | 5.0   | 2.2     | 2.4     | 3.7   |
|     | es  | 1.8     | 2.0     | 2.7   | 3.1     | 3.2     | 4.9   | 1.8     | 2.0     | 2.5   |
|     | tr  | 2.9     | 3.4     | 4.4   | 3.6     | 4.1     | 5.8   | 2.6     | 2.9     | 4.3   |
| mr  | en  | 2.7     | 2.9     | 3.5   | 5.1     | 5.6     | 7.2   | 2.6     | 2.9     | 3.6   |
|     | am  | 2.6     | 2.9     | 3.8   | 4.2     | 4.9     | 6.3   | 2.6     | 2.9     | 3.7   |
|     | mr  | 3.2     | 3.7     | 4.9   | 5.5     | 6.3     | 8.9   | 3.3     | 3.8     | 5.1   |
|     | es  | 1.8     | 2.0     | 2.5   | 3.8     | 4.1     | 5.6   | 1.9     | 2.1     | 2.6   |
|     | tr  | 2.9     | 3.3     | 4.4   | 4.4     | 5.1     | 7.5   | 2.8     | 3.2     | 4.5   |
| es  | en  | 54.5    | 61.7    | 69.9  | 54.8    | 59.3    | 66.9  | 55.5    | 60.2    | 68.0  |
|     | am  | 43.2    | 50.2    | 54.2  | 47.3    | 51.1    | 52.6  | 44.2    | 48.2    | 52.9  |
|     | mr  | 44.4    | 53.4    | 58.6  | 38.8    | 42.9    | 48.3  | 42.4    | 46.0    | 50.3  |
|     | es  | 46.0    | 52.4    | 60.1  | 52.6    | 58.7    | 68.8  | 52.5    | 57.9    | 65.7  |
|     | tr  | 50.0    | 58.0    | 66.0  | 48.4    | 53.2    | 59.0  | 50.6    | 54.5    | 62.1  |
| tr  | en  | 49.3    | 54.5    | 64.0  | 52.8    | 57.1    | 61.9  | 52.6    | 55.5    | 58.8  |
|     | am  | 44.4    | 51.7    | 56.5  | 45.8    | 49.8    | 55.1  | 45.5    | 49.9    | 51.8  |
|     | mr  | 47.9    | 55.9    | 68.0  | 38.5    | 43.1    | 51.5  | 45.2    | 50.5    | 58.9  |
|     | es  | 37.0    | 40.8    | 49.5  | 43.2    | 46.7    | 52.2  | 40.4    | 43.8    | 48.6  |
|     | tr  | 59.4    | 67.1    | 84.1  | 58.7    | 65.4    | 77.7  | 62.4    | 68.1    | 79.1  |

Table 26: Cross-lingual in-domain slot labelling results using XLM-R ( $F_1 \times 100$ )

## ACL 2023 Responsible NLP Checklist

---

### A For every submission:

- A1. Did you describe the limitations of your work?  
7
- A2. Did you discuss any potential risks of your work?  
8
- A3. Do the abstract and introduction summarize the paper’s main claims?  
1
- A4. Have you used AI writing assistants when working on this paper?  
*Left blank.*

### B Did you use or create scientific artifacts?

3

- B1. Did you cite the creators of artifacts you used?  
3
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?  
1
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?  
3
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?  
8
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?  
3
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.  
3

### C Did you run computational experiments?

4

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?  
4, Appendix

---

*The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.*

- ✓ C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?  
*4, Appendix*
  - ✓ C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?  
*4*
  - ✓ C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?  
*4*
- D** ✓ **Did you use human annotators (e.g., crowdworkers) or research with human participants?**  
*3*
- ✓ D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?  
*Appendix B*
  - ✓ D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?  
*3, 8*
  - ✓ D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?  
*8, Appendix B*
  - ✓ D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?  
*Appendix C*
  - ✓ D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?  
*Appendix C*