

Zero-Shot Cross-lingual Semantic Parsing

Tom Sherborne and Mirella Lapata

Institute for Language, Cognition and Computation

School of Informatics, University of Edinburgh

10 Crichton Street, Edinburgh EH8 9AB

tom.sherborne@ed.ac.uk, mlap@inf.ed.ac.uk

Abstract

Recent work in cross-lingual semantic parsing has successfully applied machine translation to localize parsers to new languages. However, these advances assume access to high-quality machine translation systems and word alignment tools. We remove these assumptions and study cross-lingual semantic parsing as a zero-shot problem, without parallel data (i.e., utterance-logical form pairs) for new languages. We propose a multi-task encoder-decoder model to transfer parsing knowledge to additional languages using only English-logical form paired data and in-domain natural language corpora in each new language. Our model encourages language-agnostic encodings by jointly optimizing for logical-form generation with auxiliary objectives designed for cross-lingual latent representation alignment. Our parser performs significantly above translation-based baselines and, in some cases, competes with the supervised upper-bound.¹

1 Introduction

Executable semantic parsing maps a natural language *utterance* to a *logical form* (LF) for execution in some *knowledge base* to return a *denotation*. The parsing task renders an utterance as a semantically identical, but machine-interpretable, expression *grounded* in a denotation. The transduction between natural and formal languages has allowed semantic parsers to become critical infrastructure in building human-computer interfaces for question answering, (Berant et al., 2013; Liang, 2016; Kollar et al., 2018), dialog systems (Artzi and Zettlemoyer, 2011), and robotics (Dukes, 2014).

Recent advances in semantic parsing have improved accuracy for neural parsers (Jia and Liang, 2016; Dong and Lapata, 2016; Wang et al., 2020a) and examined their generalization capabilities with new dataset challenges (Zhong et al., 2017; Yu

¹Our code and data are available at github.com/tomsherborne/zx-parse.

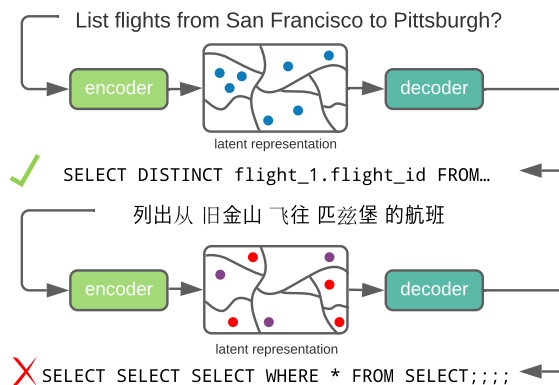


Figure 1: Accurate cross-lingual semantic parsing requires alignment of latent semantic representations across languages. The encoder generates a representation of the English utterance (blue points) to condition upon during decoding. Producing the same logical form from the equivalent Chinese utterance requires a similar encoding. However, without alignment, the representation may partially match (purple points) or not at all (red points), leading the decoder to generate an inaccurate, ill-formed query.

et al., 2018), in addition to considering languages other than English (Duong et al., 2017; *inter alia.*). Prior work largely assumes that utterance-logical form training data is parallel in all languages (Jie and Lu, 2014), or must be created with human translation (Susanto and Lu, 2017a). This entry barrier to localization for new languages has motivated the exploration of machine translation (MT) as an economical alternative (Sherborne et al., 2020; Moradshahi et al., 2020). However, MT can introduce performance-limiting artifacts and struggle to accurately model native speakers (Riley et al., 2020). Additionally, high-quality machine translation is less viable for lower resource languages, further limiting the appeal of MT-based approaches.

In this work, we propose a new approach for **zero-shot executable semantic parsing**. Our method maximizes the success of cross-lingual transfer for a parser, trained on English paired data (EN \rightarrow LF), to accurately generate logical forms

from new languages ($X \rightarrow LF$). Our goal is to parse utterances in a new language, l , without observing paired training data for this language, suitable machine translation, or bilingual dictionaries between l and English. Our critical dependencies are a pre-trained language model and utterance-logical form paired data for a source language (i.e., English). Aside from the *zero-shot* problem which is hard on its own (since paired data is not available for new languages), our semantic parsing challenge is further compounded with the difficulties inherent to *structured prediction* and the deficiency of copying strategies without gold token-level alignment (Zhu et al., 2020).

We conceptualize cross-lingual semantic parsing as a *latent representation alignment* problem. As illustrated in Figure 1, we wish to encode different languages to an overlapping latent space for the decoder to have any chance at generating accurate logical forms. To achieve this, we train a decoder, conditioned upon encodings from a source language (e.g., English), to generate logical forms and simultaneously train encodings of a new language (e.g., Chinese) to be maximally similar to English. We hypothesize that if latent representations are aligned from a language-agnostic encoder, one can generate accurate logical forms from a new language without semantic parsing training data and thus eliminate the errors outlined in Figure 1.

Our approach adopts a multi-task learning paradigm and trains a parser with auxiliary objectives, optimized to converge representations of additional new languages. We encourage language-agnostic representations by jointly optimizing for generating logical forms, reconstructing natural language, and promoting language invariance. Our intuition is that auxiliary losses can be exploited to induce similarity in a multi-lingual latent space. The effect of such alignment is that a decoder, trained only on English, can recognize an encoding from another language and generate the relevant logical form. Similar multi-task approaches have been successful in spoken-language understanding (van der Goot et al., 2021), text simplification (Mallinson et al., 2020; Zhao et al., 2020b), dependency parsing (Ahmad et al., 2019b), and machine translation (Arivazhagan et al., 2019). This work, to our knowledge is the first attempt to devise auxiliary objectives for executable semantic parsing as a zero-shot task. Our framework and hypothesis are also sufficiently flexible for application in

additional zero-shot sequence transduction tasks.

Our motivation is to improve parsing for non-English languages with maximal resource efficiency and minimal external dependencies beyond native-speaker utterances. We, therefore, induce a shared multilingual space without resorting to machine translation (Sherborne et al., 2020; Moradshahi et al., 2020) and argue that our approach is superior because it (a) nullifies the introduction of translation or word alignment errors and (b) scales to low-resource languages without reliable MT. Experimental results on Overnight (Wang et al., 2015; Sherborne et al., 2020) and a new executable version of MultiATIS++ show that our parser generates more accurate logical forms with a minimized cross-lingual transfer penalty from English to French (FR), Portuguese (PT), Spanish (ES), German (DE), Chinese (ZH), Hindi (HI), and Turkish (TR).

2 Related Work

Cross-lingual Modeling This area has recently gained increased interest across several natural language understanding settings (Zhao et al., 2020a; Nooralahzadeh et al., 2020) with benchmarks such as XGLUE (Liang et al., 2020) and XTREME (Hu et al., 2020) allowing to study classification and generation tasks for multiple languages. Cross-lingual approaches have also been developed for dependency parsing (Tiedemann et al., 2014; Schuster et al., 2019), sentence simplification (Mallinson et al., 2020), and spoken-language understanding (SLU; He et al., 2013; Upadhyay et al., 2018).

Pre-training has shown to be widely beneficial for a wide range of cross-lingual models (Devlin et al., 2019; Conneau et al., 2020a). By virtue of being trained on massive corpora, these models purportedly learn an overlapping cross-lingual latent space (Conneau et al., 2020b) but have also been identified as under-trained for some tasks (Li et al., 2021), shown poor zero-shot performance, especially for languages dissimilar to English (Pires et al., 2019), and high variance (Keung et al., 2020).

Semantic Parsing Most previous work (Lu, 2014; Susanto and Lu, 2017b,a) has focused on *multilingual* semantic parsing, i.e., learning from multiple natural languages in parallel, largely affirming the benefit of “high-resource” multilingual data and multi-language ensemble training (Jie and Lu, 2014). Shao et al. (2020) further improved cross-lingual similarity with adversarial language

identification across such ensembled training data. Code-switching in multilingual parsing has also been explored through mixed-language training datasets (Duong et al., 2017; Einolghozati et al., 2021). To adapt a parser to new languages, machine translation has been used as a reasonable proxy for in-language data (Sherborne et al., 2020; Moradshahi et al., 2020). However, machine translation, in either direction can introduce limiting artifacts (Artetxe et al., 2020) with poor generalization due to how “translationese” training data diverges from gold test utterances (Riley et al., 2020).

Zero-shot parsing has primarily focused on ‘cross-domain’ challenges to improve generalization across varying query structures and lexicons (Herzig and Berant, 2018; Givoli and Reichart, 2019) or different databases (Zhong et al., 2020; Suhr et al., 2020; Yu et al., 2018). The combination of zero-shot parsing with cross-lingual modeling has also been examined for the UCCA formalism (Hershcovich et al., 2019) and for task-oriented dialogue systems (see below).

Dialog Modeling Cross-lingual transfer has been studied in the context of goal-oriented dialog for the spoken language understanding (SLU) tasks of intent classification and slot labeling (i.e., parsing an utterance into a semantic frame identifying the user’s intent and its arguments). Recently released multilingual datasets like MultiATIS++ (Xu et al., 2020) and MTOP (Li et al., 2021) have facilitated the study of zero-shot transfer through the combination of pre-training, machine translation, and word alignment (to project annotations between languages). Recent work in this setting (Zhu et al., 2020; Li et al., 2021; Krishnan et al., 2021; Nicosia et al., 2021) identifies a penalty for cross-lingual transfer that neither pre-training nor machine translation can fully overcome.

3 Problem Formulation

The primary challenge for cross-lingual parsing is learning parameters that can parse an utterance, x , from an unseen test language to an accurate logical form (LF). Typically, a parser trained on language l , or multiple languages $\{l_1, \dots, l_N\}$, is only capable for these languages and performs poorly outside this set. For a new language, prior approaches require parallel datasets and models (Jie and Lu, 2014; Haas and Riezler, 2016; Duong et al., 2017).

In our work, zero-shot parsing refers to parsing utterances in new languages *without paired data*

during training. For some language, l , there exists no pairing of x_l to a logical form, y , except for English.² This setting also excludes “silver-standard” training pairs created using machine-translation. As these models have ultimately observed some form of utterance-logical form pairs for each new language, we do not consider such approaches here and refer to Sherborne et al. (2020) as an example of using MT for this task.

It might be tempting to approach this problem as a case of fine-tuning a pre-trained (English) decoder for LF generation. Problematically, the output target is expressed in a formally defined language (e.g., SQL or λ -DCS) which models the semantics of questions very differently to natural language (e.g., without presumption or co-operation; Kaplan 1978). Formal languages (Kamp and Reyle, 1993) additionally present artifacts which render fine-tuning challenging such as unfamiliar syntax (e.g., table aliases or explicit recursion) and long output sequences. In practice, we observed fine-tuning leads to poor performance (e.g., $< 1\%$ accuracy on all languages), with the model insisting on hallucinating natural language. This is seemingly at odds with adjacent work in dialog modeling, which has found pre-trained decoders to be beneficial (Li et al., 2021). However, SLU requires learning a lightweight label vocabulary compared to the 200+ tokens required in LFs. Additionally, SLU typically maintains output sequences of similar size to natural language inputs (with tightly coupled syntactic compositionality between the two), whereas the syntactic and structural demands of LF generation are largely divorced from the input utterance.

In our solution, the model is trained to parse from utterance-logical forms pairs only in English. Other languages are incorporated using auxiliary objectives and data detailed in Section 4. We explore the hypothesis that an overlapping multi-lingual latent space can be learned through auxiliary objectives in tandem with logical form generation (see Figure 2). Our intuition is that introducing these additional losses minimizes cross-lingual variance in latent encoding space by optimizing for language-agnostic representations with high similarity to the source language (i.e., English). Our approach minimizes the cross-lingual transfer penalty such that the zero-shot parser predicts logical forms from test inputs regardless of utterance language.

²English is the “source” language for all our experiments. We refer to languages seen only at test time as “new”.




Figure 2: Our model, **ZX-Parser**, is a **Zero-shot Cross-lingual semantic Parser** which augments an encoder-decoder model with auxiliary objectives. The Encoder, E , generates a representation, z , which is input to the logical form decoder, D_{LF} , reconstruction decoder, D_{NL} , or language prediction classifier, LP . During training, English is input to all objectives and additional languages are incorporated using *only* the additional objectives $\{D_{NL}, LP\}$. Logical forms are predicted using D_{LF} for all inputs at test time.

By framing the cross-lingual parsing task as a latent representation alignment challenge, we explore a possible upper bound of parsing accuracy without errors from external dependencies. Section 6 demonstrates that our zero-shot model, using only English paired data and a small additional corpus, can generate accurate logical forms above translation baselines to compete with fully supervised in-language training.

4 Our Zero-shot Model: ZX-PARSE

We adopt a multi-task sequence-to-sequence model (Luong et al., 2016) which combines logical form generation with two auxiliary objectives. The first is a language identification discriminator and the second is a reconstruction or translation decoder. An overview of our semantic parser is given in Figure 2; we describe each component below.

Generating Logical Forms Predicting logical forms is the primary objective for our model. Given an utterance $x = (x_1, x_2, \dots, x_T)$, we wish to generate logical form $y = (y_1, y_2, \dots, y_M)$ representing the same meaning in a machine-executable language. We model this transduction task using an encoder-decoder neural network (Sutskever et al., 2014) based upon the Transformer architecture (Vaswani et al., 2017).

The sequence x is encoded to a latent representation $z = (z_1, z_2, \dots, z_T)$ through Equation (1) using a stacked self-attention Transformer encoder, E , with weights θ_E .

$$z = E(x|\theta_E) \quad (1)$$

$$p(y|x) = \prod_{i=0}^M p(y_i|y_{<i}, x) \quad (2)$$

$$p(y_i|y_{<i}, x) = \text{soft}(D_{LF}(y_{<i}|z, \theta_{D_{LF}})) \quad (3)$$

$$\mathcal{L}_{LF} = - \sum_{(x, y) \in \mathcal{S}_{LF}} \log p(y|x) \quad (4)$$

The conditional probability of the output sequence y is expressed in Equation (2) as each token y_i is autoregressively generated based upon z and prior outputs, $y_{<i}$. Equation (3) models distribution $p(y_i|y_{<i}, x)$ using a Transformer decoder for logical forms, D_{LF} , with associated weights $\theta_{D_{LF}}$ where soft is the softmax function.

We predict an output, \hat{y} , for semantic parsing dataset $\mathcal{S}_{LF} = \{x^n, y^n\}_{n=0}^N$, through the encoder and logical form decoder, $\{E, D_{LF}\}$. Equation (4) describes the loss objective minimizing the cross-entropy between y and \hat{y} .

Language Prediction Our first additional objective encourages language-agnostic representations by reducing the discriminability of the source language, l , from z . Equation (5) defines a Language Prediction (LP) network to predict l from z using a linear classifier over L training languages:

$$LP(x) = W_i x + b_i \quad (5)$$

where $W_i \in \mathbb{R}^{L \times |z|}$ and $b_i \in \mathbb{R}^L$ are a weight and bias respectively. We follow the best model from Ahmad et al. (2019b). Equation (6) describes the conditional model for the output distribution where a language label is predicted using the time-average of the input encoding z of length T :

$$p(l|x) = \text{soft}\left(LP\left(\frac{1}{T} \sum_t z_t\right)\right) \quad (6)$$

Finally, Equation (7) describes the objective function for the LP network:

$$\mathcal{L}_{LP} = - \sum_x \log p(l|x) \quad (7)$$

However, we *reverse this gradient* in the backward pass before the LP network, to encourage the encoder to produce language invariant representations (Ganin et al., 2016). The LP network is optimized to discriminate the source language from z , but the encoder is now optimized adversarially *against* this

objective. Our intuition is that discouraging language discriminability in z encourages latent representation similarity across languages, and therefore reduces the penalty for cross-lingual transfer.

Generating Natural Language The final objective acts towards both regularization and cross-lingual similarity. Motivated by *domain-adaptive pre-training* (Gururangan et al., 2020), we further adapt the encoder towards question-style utterances from native speakers of each test language lacking task-specific training data. We add an additional Transformer decoder optimized to reconstruct a noisy input from latent representation z , in Equation (1). Utterance, x , is input to the encoder, E , and a separate decoder, D_{NL} , then reconstructs x from z . We follow the *denoising* objective from Lewis et al. (2020) and replace x with noised input $\tilde{x} = N(x)$ with noising function N . The output probability of reconstruction is given in Equation (9) with each token predicted through Equation (10) using decoder, D_{NL} , with weights $\theta_{D_{\text{NL}}}$:

$$\hat{z} = E(\tilde{x}|\theta_E) \quad (8)$$

$$p(x|\tilde{x}) = \prod_{i=0}^T p(x_i|x_{<i}, \tilde{x}) \quad (9)$$

$$p(x_i|x_{<i}, \tilde{x}) = \text{soft}(D_{\text{NL}}(x_{<i}|\hat{z}, \theta_{D_{\text{NL}}})) \quad (10)$$

The auxiliary objectives are trained using both the utterances from \mathcal{S}_{LF} and monolingual data, $\mathcal{S}_{\text{NL}} = \{\{x^n\}_{n=0}^N\}_{l=0}^L$, in L languages (see Section 5). Submodel, $\{E, D_{\text{NL}}\}$, predicts the reconstruction of x from \tilde{x} with the following objective:

$$\mathcal{L}_{\text{NL}} = - \sum_x \log p(x|\tilde{x}) \quad (11)$$

In the form described above, this objective requires only unlabeled, monolingual utterances in each target language. However, we can also augment it with a translation component to exploit natural language bi-text between the new language and English (e.g., $\mathcal{S}_{\text{NL}} = \{\{x_{\text{EN}}^n, x_l^n\}_{n=0}^N\}_{l=0}^L$) to further promote cross-lingual similarity. According to some sampling factor τ , we randomly choose whether to *reconstruct* an utterance (as above) or *translate* to the parallel English utterance (i.e., replace x in Equation (11) with x_{EN}).

Combined Model The combined model uses a single encoder, E , and the three objective decoders $\{D_{\text{LF}}, D_{\text{NL}}, \text{LP}\}$ (see Figure 2). During

training, an English query is encoded and input to all three objectives to express output loss as $\mathcal{L}_{\text{LF}} + \mathcal{L}_{\text{NL}} + \mathcal{L}_{\text{LP}}$. For new languages without (x, y) pairs, the utterance is encoded and input only to the auxiliary objectives for a combined loss as $\mathcal{L}_{\text{NL}} + \mathcal{L}_{\text{LP}}$. During inference, an utterance is encoded and *always* input to D_{LF} to predict a logical form, \hat{y} , regardless of test language, l . During the backward pass, each output loss back-propagates the gradient signal from the respective objective function. For the encoder, these signals are combined as:

$$\frac{\partial \mathcal{L}}{\partial \theta_E} = \frac{\partial \mathcal{L}_{\text{LF}}}{\partial \theta_E} - \lambda \alpha_{\text{LP}} \frac{\partial \mathcal{L}_{\text{LP}}}{\partial \theta_E} + \alpha_{\text{NL}} \frac{\partial \mathcal{L}_{\text{NL}}}{\partial \theta_E} \quad (12)$$

$$\lambda = \frac{2}{1 + e^{-\gamma p}} - 1 \quad (13)$$

where $\alpha_{\{\text{LP}, \text{NL}\}}$ are loss weightings for auxiliary objectives and λ is the reversed gradient scheduling parameter from Ganin et al. (2016). The λ value increments with training progress p , scaled by γ , according to Equation (13), to limit the impact of noisy predictions during early training.

We expect that the parser will adapt and recognize an encoding from an unfamiliar language through our joint training process, and successfully connect new language representations to the logical-form decoder at test time. This sequence-to-sequence approach is highly flexible and may be useful for zero-shot approaches to additional generation tasks (e.g., paraphrasing).

5 Experimental Setup

Semantic Parsing Datasets Our experiments examine whether our zero-shot approach generalizes across languages and domains. We evaluate performance on a new version of the **ATIS** dataset of travel queries (Hemphill et al., 1990; Dahl et al., 1994). We align existing English utterances and SQL logical forms from Iyer et al. (2017) to the multi-lingual utterances from the MultiATIS++ dataset for spoken language understanding (Xu et al., 2020). This alignment adds executable SQL queries to utterances in Chinese (ZH), German (DE), French (FR), Spanish (ES), and Portuguese (PT). We use the same 4,473/493/448 dataset split for training/validation/test as Kwiatkowski et al. (2011). We also add to the test set Hindi (HI) and Turkish (TR) utterances from Upadhyay et al.

(2018).³ We can now predict SQL from the ATIS test questions in eight natural languages. The Multi-ATIS++ Japanese set was excluded as the utterance alignment to this language was not recoverable.

We also examine **Overnight** (Wang et al., 2015), an eight-domain dataset covering *Basketball*, *Blocks*, *Calendar*, *Housing*, *Publications*, *Recipes*, *Restaurants*, and *Social Network* domains. Overnight comprises 13,682 English utterances paired with λ -DCS logical forms, executable in SEMPRES (Berant et al., 2013), split into 8,754/2,188/2,740 for training/validation/test respectively. This training data exists only in EN and we use the ZH and DE test data from Sherborne et al. (2020) for multilingual evaluation. Given the varying linguistic phenomena across domains (e.g. relative spatial reasoning in *Blocks* or temporal arithmetic in *Calendar*), this dataset presents a harder challenge for cross-lingual transfer.

We measure performance with *denotation accuracy* as all inferred logical forms are executable in some knowledge base. This metric compares the retrieved denotation from the prediction, \hat{y} , to that from executing the gold-standard logical form. Dataset sizes are outlined in Appendix A.

Natural Language Data For the reconstruction objective, we used the MKQA corpus (Longpre et al., 2020), a multi-lingual translation of 10,000 samples from NaturalQuestions (Kwiatkowski et al., 2019). This is suitable for our auxiliary objective as the utterances are native-speaker question surface forms, matching our test set while varying in subject. MKQA is also balanced across new languages to limit overexposure bias to one new language. For bi-text, we use the original English and the professionally translated question as a pair.

We also report experiments using a sample of crawled data from ParaCrawl 7.1 (Bañón et al., 2020). The sample comprises 10,000 web scraped sentences paired with equivalent English to form bi-text. Note that these samples are mostly declarative sentences and as such do not match the surface form of our test inputs (i.e., questions) and are also not parallel between sampled languages. We contrast this to MKQA to examine how the *style* of natural language data influences performance.

For ATIS experiments, we use 60,000 utterances from each source in languages with training data (EN, FR, PT, ES, DE, ZH). For Overnight, we use

³Misalignment between ATIS versions result in the test sets containing 442 and 381 utterances for HI and TR respectively.

30,000 utterances in EN, DE, and ZH.

Model Configuration The implementation of ZX-PARSE (see Section 4) largely follows parameter settings from Liu et al. (2020) for Transformer encoder and decoder layers (see Appendix A for details on model configuration). ZX-PARSE requires an encoder model to generate multi-lingual latent representations for all objectives. Our main results use only the encoder component of *mBART50* (Tang et al., 2020) and we present experiments using other pre-trained models in Appendix B. We use all pre-trained encoder layers and append one additional learnable layer. All decoders are randomly initialized six-layer stacks. Early experiments found this approach superior to any pre-trained decoder initialization.

The language predictor follows from Ahmad et al. (2019b) as a single linear classification layer mapping from 1,024 inputs to L output languages. Earlier findings supported that if the LP network is larger, then the reversed gradient signal is too strong and therefore less useful as the LP network can memorize the language.

Comparison Models We primarily compare to a “Translate-Test” back-translation baseline wherein the new language test set is translated to English using Google Translate (Wu et al., 2016) and input to a reference sequence-to-sequence model trained on English. We also compare to “Translate-Train”, where we use MT from English to generate a proxy dataset in each new language (e.g., French, Portuguese, Spanish, German, Chinese, Hindi and Turkish) to train a monolingual parser. We consider improving upon these “minimum effort” baselines as a lower bound for justifying our approach.

Additionally, we compare to an upper-bound monolingual model trained on professional translations of the new languages. We report results on MultiATIS++ for FR, PT, ES, DE, and ZH (professional translations are not available for Overnight training data). This is the “maximum effort” strategy that we desire to avoid. Parameters for these reference systems match those outlined above e.g., *mBART50* encoder to logical form decoder.

6 Results

Our results are outlined to answer four core questions, with additional ablations in Appendix B. Our findings support the hypothesis that we can minimize the cross-lingual transfer penalty by im-

	ATIS								Overnight		
	EN	FR	PT	ES	DE	ZH	HI	TR	EN	DE	ZH
Monolingual Training	77.2	67.8	66.1	64.1	66.6	64.9	—	—	80.5	—	—
Translate-Train	—	55.9	56.1	57.1	60.1	56.1	56.3	45.4	—	62.2	59.4
Translate-Test	—	58.2	57.3	57.9	56.9	51.4	52.6	52.7	—	60.1	48.1
ZX-PARSE	76.9	70.2	63.4	59.7	69.3	60.2	54.9	48.3	81.9	66.2	60.0

Table 1: Denotation accuracy for ATIS (Dahl et al., 1994) and Overnight (eight-domain average; Wang et al., 2015) for supervised monolingual upper-bound, Translate-Test, and our best ZX-PARSE model. Results for English (EN), French (FR), Portuguese (PT), Spanish (ES), German (DE), Chinese (ZH), Hindi (HI) and Turkish (TR) ranked by similarity to English (Ahmad et al., 2019a). Best results compared to baselines are bolded.

ZX-PARSE	ATIS								Overnight		
	EN	FR	PT	ES	DE	ZH	HI	TR	EN	DE	ZH
(a) D_{LF} only	77.2	61.3	42.5	46.5	50.2	38.5	40.4	37.3	80.5	58.4	48.0
(b) $D_{LF} + D_{NL}$	77.7	62.7	54.9	58.2	61.1	51.2	49.5	44.7	81.3	62.7	49.5
(c) $D_{LF} + LP$	76.3	57.2	53.7	51.8	58.6	44.1	39.8	38.8	80.6	60.6	49.4
(d) $D_{LF} + LP + D_{NL}$	76.9	70.2	63.4	59.7	69.3	60.2	54.9	48.3	81.9	66.2	60.0

Table 2: Denotation accuracy for ATIS and Overnight (eight-domain average) comparing ablations of ZX-PARSE: (a) no auxiliary objectives, (b) logical form (LF) generation and reconstruction, (c) LF generation and language prediction, (d) all objectives. Best results compared to baselines are bolded.

proving latent alignment with auxiliary objectives. We also examine the latent space directly and find ZX-PARSE learns more similar representations between languages. Our parser achieves state-of-the-art zero-shot results for all non-English languages in the MultiATIS++ and Overnight benchmarks.

Better than Translation? We compare between ZX-PARSE and the upper- and lower-bounds in Table 1. Our multi-task approach significantly improves upon “Translate-Test” for all languages included within the auxiliary objectives ($p < 0.01$). For ATIS, we find that “Translate-Train” performs below “Translate-Test” for languages similar to English (FR, ES, PT) but worse for more distant languages (DE, ZH). ZX-PARSE performance improves on “Translate-Train” for all languages included in reconstruction (EN, FR, PT, ES, DE, ZH), however, the general cross-lingual improvement insufficiently extends to additional languages (HI, TR) to perform above baselines.

Within ZX-PARSE, French and German demonstrate the strongest zero-shot accuracy — +2.4% and +2.7% above the monolingual upper bound for ATIS. We do not observe similar improvement for Portuguese or Spanish despite their similarity to English. This may be a result of German and French dominating the pre-training corpora com-

pared to other new languages. (Tang et al., 2020, their Table 6).

Our model demonstrates similar significant improvement for Overnight ($p < 0.01$), however, we find lesser gain compared to ATIS. This may be a consequence of the compounded challenge of evaluating eight varied domains of complex linguistic constructs. Here, we find that “Translate-Train” is a stronger approach than “Translate-Test”, which may be a consequence of machine-translation direction. Our best approach on German still improves above “Translate-Train” (+4.0%), however, we find performance on Chinese to be only marginally improved by comparison (+0.6%). We also observe some contrast in ZX-PARSE performance related to orthographic similarity to English. Parsing accuracy on Overnight in German is +6.2% above Chinese, with a similar +9.1% gap between these same languages for ATIS.

Which Objective Matters? Ablations to the model are shown in Table 2, identifying the contributions of different objectives. Model (a) shows that without auxiliary objectives, performance in new languages is generally below Translate-Test. This is unsurprising, as this approach uses only pre-trained cross-lingual information without additional effort to improve similarity. Such efforts

are incorporated in Model (b) using the additional reconstruction decoder. Even without the LP loss, domain targeted adaptation (with translation) improves cross-lingual parsing by an average across new languages of +9.3% for ATIS and +2.9% for Overnight. Notably, we identified an optimal ratio of translation to reconstruction of 50% (i.e., $\tau = 0.5$). This suggests that both monolingual utterances (for domain-adaptive tuning) and bi-text (for translation) contribute to the utility of our method beyond reliance on one technique.

Evaluating the LP objective within Model (c) and (d), we find the reversed gradient successfully reduces language discriminability. For Model (d), language prediction accuracy during training peaks at 93% after 2% progress and subsequently decreases to <8% beyond 10% of training. Language prediction accuracy for the test set is 7.2%. We observe a similar trend for Model (c). Comparing individual objectives, we find the addition of the language predictor alone less helpful than the reconstruction decoder. Comparing Model (a) and (c), we observe a smaller average improvement on new languages of +4.3% for ATIS and +1.8% for Overnight. This suggests adaptation towards specific surface form patterns can be more effective here than modeling languages as discrete labels.

Considering the combination of objectives in Model (d), we identify cumulative benefit to parsing with both objectives. Compared to Model (a), the full model improves by an average of +16.3% for ATIS and +9.9% for Overnight across new languages. Our findings support our claim that latent cross-lingual similarity can be improved using auxiliary objectives and we specifically identify that a combination of approaches yields superior parsing. We suggest that this combination benefits from constructive interference, as the language prediction loss promotes invariance in tandem with multilingual generation tasks adapting the encoder to improve modeling the surface form (e.g., questions from native speakers) of the new language test data.

Additional objectives also improve parsing for Hindi and Turkish despite neither being included within auxiliary training data (see HI and TR columns in Table 3). By adapting our latent representation to encourage similarity, we improve parsing accuracy for two typologically diverse languages without explicit guidance. To further examine this, we visualize the MultiATIS++ test set in Figure 3 and observe *less discriminable* encodings

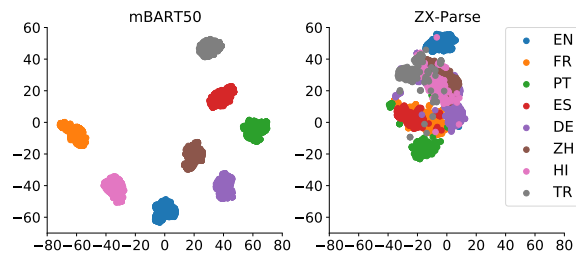


Figure 3: t-SNE comparison using mBART50 and ZX-PARSE encoders (MultiATIS++ test set). Our approach improves the latent alignment across languages.

from ZX-PARSE compared to *mBART50*. Quantitatively, we find the average cosine distance between the sentence-mean of parallel utterances reduces from 0.58 to 0.47. Similarly, the average token-level symmetric Hausdorff distance (Taha and Hanbury, 2015) between languages reduces from 0.72 to 0.41. This further supports that we learn more similar representations and our method has wider utility beyond explicitly targeted languages.

Does Language Style Matter? In Table 3 we examine whether our auxiliary objectives are influenced by the style of natural language corpora for reconstruction. We find the use of questions positively improves performance compared to crawled sentences. Using questions either as monolingual utterances (i.e., no translation in D_{NL}) or with as a bi-text sample (i.e., reconstruction and translation in D_{NL}) improves above the Translate-Test baseline. We observe modest improvements with ParaCrawl, especially when introducing bi-text into D_{NL} , but this is less consistent across languages. Overall, our results suggest that ZX-PARSE is robust even when question-style data is unavailable but can be particularly effective when adapting towards both new languages *and* domains. We also examined the influence of language family on performance (see Appendix B) and found that best performance utilizes a linguistically varied ensemble of languages. Omitting either Romance (ES/FR/PT) or Sino-Tibetan (ZH) languages in reconstruction negatively impacts performance.

Where Does Improvement Come from? Comparing to Translate-Test, on ATIS, our best model generates 32% fewer ill-formed SQL requests and 24% fewer extraneous queries accessing unrelated tables in the database. Translation can fail when entities are mishandled and our model generates 36% fewer queries with erroneous named entities. For Overnight, gains are strongly related to improved numeracy in the model. Between our full

Baselines	ATIS								Overnight		
	EN	FR	PT	ES	DE	ZH	HI	TR	EN	DE	ZH
Translate-Train	—	55.9	56.1	57.1	60.1	56.1	56.3	45.4	—	62.2	59.4
Translate-Test	—	58.2	57.3	57.9	56.9	51.4	52.6	52.7	—	60.1	48.1
ZX-PARSE											
MKQA $\tau = 0.0$	76.3	67.1	60.5	58.2	68.3	59.2	54.1	47.1	81.3	64.3	52.7
MKQA $\tau = 0.5$	76.9	70.2	63.4	59.7	69.3	60.2	54.9	48.3	81.9	66.2	60.0
ParaCrawl $\tau = 0.0$	72.7	63.4	58.0	54.1	62.0	50.9	46.9	39.9	78.4	62.4	51.1
ParaCrawl $\tau = 0.5$	76.5	64.6	60.3	59.2	63.1	52.6	47.8	45.8	81.2	63.2	52.9

Table 3: Denotation accuracy for ATIS and Overnight (eight-domain average) comparing between data sources: MKQA (questions) or sampled web data from ParaCrawl. We additionally contrast between modeling corpora as monolingual text ($\tau = 0$) or partially as bi-text ($\tau = 0.5$). Best results compared to baselines are bolded.

model and simplest approach (Model (a) in Table 2), we find more well-formed logical forms account for the largest improvement (32.5% fewer ill-formed SQL queries for ATIS and 35.2% fewer ill-formed λ -DCS queries for Overnight). This supports our notion in Figure 1 that better latent alignment can minimize cross-lingual penalty. However, improved structure prediction is insufficient to solve this task on its own; 58.7% of remaining errors in the best model are due to mishandled entities with the highest entity errors for Chinese (60.2%) and lowest for French (36.7%). This suggests that aligning entities across languages might be necessary for further improvement.

7 Conclusion

We presented a multi-task model for zero-shot cross-lingual semantic parsing which combines logical form generation with auxiliary objectives that require only modest natural language corpora for localization. Through aligning latent representations, ZX-PARSE minimizes the error from cross-lingual transfer and improves accuracy across languages unseen during training.

Although we focused exclusively on executable semantic parsing, our approach is general and potentially relevant for linguistically motivated frameworks such as Abstract Meaning Representation (Banarescu et al., 2013; Damonte and Cohen, 2018) or Discourse Representation Theory (Kamp and Reyle, 1993; Evang and Bos, 2016). In the future, we will investigate a few-shot scenario and study sample efficient cross-lingual transfer by explicitly promoting generalization using techniques such as meta-learning (Finn et al., 2017).

Ethics Statement

A key limitation of our work is the limited coverage of eight higher-resource languages. As such, we are unable to test our approach in a **genuinely** low-resource scenario. We must also consider the risk of over-generalization to dominant dialects within each language as we lack an evaluation of additional dialects (e.g. our English dataset is representative of American English but not Indian English). We hope that such issues can be addressed with additional data collection.

Our training requirements are detailed in Appendix A. We hope our work contributes to further usage and development of singular multilingual models as opposed to learning N monolingual models for N languages.

Acknowledgements

We thank the anonymous reviewers for their feedback and Bailin Wang, Kate McCurdy, and Rui Zhang for insightful discussion. The authors gratefully acknowledge the support of the UK Engineering and Physical Sciences Research Council (grant EP/L016427/1; Sherborne) and the European Research Council (award number 681760; Lapata).

References

Wasi Ahmad, Zhisong Zhang, Xuezhe Ma, Eduard Hovy, Kai-Wei Chang, and Nanyun Peng. 2019a. On difficulties of cross-lingual transfer with order differences: A case study on dependency parsing. 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

- 2440–2452, Minneapolis, Minnesota. Association for Computational Linguistics.
- Wasi Uddin Ahmad, Zhisong Zhang, Xuezhe Ma, Kai-Wei Chang, and Nanyun Peng. 2019b. [Cross-lingual dependency parsing with unlabeled auxiliary languages](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 372–382, Hong Kong, China. Association for Computational Linguistics.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Roei Aharoni, Melvin Johnson, and Wolfgang Macherey. 2019. [The missing ingredient in zero-shot neural machine translation](#). *CoRR*, abs/1903.07091.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2020. [Translation artifacts in cross-lingual transfer learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7674–7684, Online. Association for Computational Linguistics.
- Yoav Artzi and Luke Zettlemoyer. 2011. [Bootstrapping semantic parsers from conversations](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 421–432, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. [Abstract Meaning Representation for sembanking](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.
- Marta Bañón, Pinzhen Chen, Barry Haddow, Kenneth Heafield, Hieu Hoang, Miquel Esplà-Gomis, Mikel L. Forcada, Amir Kamran, Faheem Kirefu, Philipp Koehn, Sergio Ortiz Rojas, Leopoldo Pla Sempere, Gema Ramírez-Sánchez, Elsa Sarrías, Marek Strelec, Brian Thompson, William Waites, Dion Wiggins, and Jaume Zaragoza. 2020. [ParaCrawl: Web-scale acquisition of parallel corpora](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4555–4567, Online. Association for Computational Linguistics.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on Freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, Seattle, Washington, USA. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020a. [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.
- Alexis Conneau, Shijie Wu, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2020b. [Emerging cross-lingual structure in pretrained language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6022–6034, Online. Association for Computational Linguistics.
- Deborah A. Dahl, Madeleine Bates, Michael Brown, William Fisher, Kate Hunicke-Smith, David Pallett, Christine Pao, Alexander Rudnicky, and Elizabeth Shriberg. 1994. [Expanding the scope of the ATIS task: The ATIS-3 corpus](#). In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Marco Damonte and Shay B. Cohen. 2018. [Cross-lingual Abstract Meaning Representation parsing](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1146–1155, New Orleans, Louisiana. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Li Dong and Mirella Lapata. 2016. [Language to logical form with neural attention](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 33–43, Berlin, Germany. Association for Computational Linguistics.
- Matthew S. Dryer and Martin Haspelmath, editors. 2013. *WALS Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig.
- Kais Dukes. 2014. [SemEval-2014 task 6: Supervised semantic parsing of robotic spatial commands](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 45–53, Dublin, Ireland. Association for Computational Linguistics.
- Long Duong, Hadi Afshar, Dominique Estival, Glen Pink, Philip Cohen, and Mark Johnson. 2017. [Multilingual semantic parsing and code-switching](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 379–389, Vancouver, Canada. Association for Computational Linguistics.

- Arash Einolghozati, Abhinav Arora, Lorena Sainz-Maza Lecanda, Anuj Kumar, and Sonal Gupta. 2021. [El volumen louder por favor: Code-switching in task-oriented semantic parsing](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1009–1021, Online. Association for Computational Linguistics.
- Kilian Evang and Johan Bos. 2016. [Cross-lingual learning of an open-domain semantic parser](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 579–588, Osaka, Japan. The COLING 2016 Organizing Committee.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. [Model-agnostic meta-learning for fast adaptation of deep networks](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135. PMLR.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(1):2096–2030.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. [AllenNLP: A deep semantic natural language processing platform](#). In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia. Association for Computational Linguistics.
- Ofer Givoli and Roi Reichart. 2019. [Zero-shot semantic parsing for instructions](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4454–4464, Florence, Italy. Association for Computational Linguistics.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Carolin Haas and Stefan Riezler. 2016. [A corpus and semantic parser for multilingual natural language querying of OpenStreetMap](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 740–750, San Diego, California. Association for Computational Linguistics.
- Xiaodong He, Li Deng, Dilek Hakkani-Tur, and Gokhan Tur. 2013. [Multi-style adaptive training for robust cross-lingual spoken language understanding](#). In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 8342–8346.
- Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. [The ATIS spoken language systems pilot corpus](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- Daniel Hershcovich, Zohar Aizenbud, Leshem Choshen, Elior Sulem, Ari Rappoport, and Omri Abend. 2019. [SemEval-2019 task 1: Cross-lingual semantic parsing with UCCA](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 1–10, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Jonathan Herzig and Jonathan Berant. 2018. [Decoupling structure and lexicon for zero-shot semantic parsing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1619–1629, Brussels, Belgium. Association for Computational Linguistics.
- 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 generalization](#). *CoRR*, abs/2003.11080.
- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, Jayant Krishnamurthy, and Luke Zettlemoyer. 2017. [Learning a neural semantic parser from user feedback](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 963–973, Vancouver, Canada. Association for Computational Linguistics.
- Robin Jia and Percy Liang. 2016. [Data recombination for neural semantic parsing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12–22, Berlin, Germany. Association for Computational Linguistics.
- Zhanming Jie and Wei Lu. 2014. [Multilingual semantic parsing : Parsing multiple languages into semantic representations](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1291–1301, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Hans Kamp and Uwe Reyle. 1993. *From Discourse to Logic; An Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer, Dordrecht.
- S. Jerrold Kaplan. 1978. [On the difference between natural language and high level query languages](#). In *Proceedings of the 1978 Annual Conference, ACM ’78*, page 27–38, New York, NY, USA. Association for Computing Machinery.

- Phillip Keung, Yichao Lu, Julian Salazar, and Vikas Bhardwaj. 2020. [Don't use English dev: On the zero-shot cross-lingual evaluation of contextual embeddings](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 549–554, Online. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Thomas Kollar, Danielle Berry, Lauren Stuart, Karolina Owczarzak, Tagyoung Chung, Lambert Mathias, Michael Kayser, Bradford Snow, and Spyros Matsoukas. 2018. [The Alexa meaning representation language](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, pages 177–184, New Orleans - Louisiana. Association for Computational Linguistics.
- Jitin Krishnan, Antonios Anastasopoulos, Hemant Purohit, and H. Rangwala. 2021. [Multilingual code-switching for zero-shot cross-lingual intent prediction and slot filling](#). *ArXiv*, abs/2103.07792.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Tom Kwiatkowski, Luke Zettlemoyer, Sharon Goldwater, and Mark Steedman. 2011. [Lexical generalization in CCG grammar induction for semantic parsing](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1512–1523, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Haoran Li, Abhinav Arora, Shuohui Chen, Anshit Gupta, Sonal Gupta, and Yashar Mehdad. 2021. [MTOP: A comprehensive multilingual task-oriented semantic parsing benchmark](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2950–2962, Online. Association for Computational Linguistics.
- Percy Liang. 2016. [Learning executable semantic parsers for natural language understanding](#). *Commun. ACM*, 59(9):68–76.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fefei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- 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.
- Shayne Longpre, Yi Lu, and Joachim Daiber. 2020. [MKQA: A linguistically diverse benchmark for multilingual open domain question answering](#). *CoRR*, abs/2007.15207.
- Wei Lu. 2014. [Semantic parsing with relaxed hybrid trees](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1308–1318, Doha, Qatar. Association for Computational Linguistics.
- Minh-Thang Luong, Quoc V. Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2016. [Multi-task sequence to sequence learning](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Jonathan Mallinson, Rico Sennrich, and Mirella Lapata. 2020. [Zero-shot crosslingual sentence simplification](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5109–5126, Online. Association for Computational Linguistics.
- Mehrad Moradshahi, Giovanni Campagna, Sina Semnani, Silei Xu, and Monica Lam. 2020. [Localizing open-ontology QA semantic parsers in a day using machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5970–5983, Online. Association for Computational Linguistics.
- Massimo Nicosia, Zhongdi Qu, and Yasemin Altun. 2021. [Translate & Fill: Improving zero-shot multilingual semantic parsing with synthetic data](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3272–3284, Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Farhad Nooralahzadeh, Giannis Bekoulis, Johannes Bjerva, and Isabelle Augenstein. 2020. [Zero-shot cross-lingual transfer with meta learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4547–4562, Online. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 8024–8035.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Parker Riley, Isaac Caswell, Markus Freitag, and David Grangier. 2020. [Translationese as a language in “multilingual” NMT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7737–7746, Online. Association for Computational Linguistics.
- Tal Schuster, Ori Ram, Regina Barzilay, and Amir Globerson. 2019. [Cross-lingual alignment of contextual word embeddings, with applications to zero-shot dependency parsing](#). 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 1599–1613, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bo Shao, Yeyun Gong, Weizhen Qi, Nan Duan, and Xiaola Lin. 2020. [Multi-level alignment pretraining for multi-lingual semantic parsing](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3246–3256, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Tom Sherborne, Yumo Xu, and Mirella Lapata. 2020. [Bootstrapping a crosslingual semantic parser](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 499–517, Online. Association for Computational Linguistics.
- Alane Suhr, Ming-Wei Chang, Peter Shaw, and Kenton Lee. 2020. [Exploring unexplored generalization challenges for cross-database semantic parsing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8372–8388, Online. Association for Computational Linguistics.
- Raymond Hendy Susanto and Wei Lu. 2017a. [Neural architectures for multilingual semantic parsing](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 38–44, Vancouver, Canada. Association for Computational Linguistics.
- Raymond Hendy Susanto and Wei Lu. 2017b. [Semantic parsing with neural hybrid trees](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 3309–3315. AAAI Press.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3104–3112.
- Abdel Aziz Taha and Allan Hanbury. 2015. [An efficient algorithm for calculating the exact hausdorff distance](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(11):2153–2163.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. [Multilingual translation with extensible multilingual pretraining and finetuning](#). *CoRR*, abs/2008.00401.
- Jörg Tiedemann, Željko Agić, and Joakim Nivre. 2014. [Treebank translation for cross-lingual parser induction](#). In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning*, pages 130–140, Ann Arbor, Michigan. Association for Computational Linguistics.
- 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.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz

- Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Bailin Wang, Richard Shin, Xiaodong Liu, Oleksandr Polozov, and Matthew Richardson. 2020a. [RAT-SQL: Relation-aware schema encoding and linking for text-to-SQL parsers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7567–7578, Online. Association for Computational Linguistics.
- Changhan Wang, Kyunghyun Cho, and Jiatao Gu. 2020b. [Neural machine translation with byte-level subwords](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9154–9160.
- Yushi Wang, Jonathan Berant, and Percy Liang. 2015. [Building a semantic parser overnight](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1332–1342, Beijing, China. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *ArXiv*, abs/1910.03771.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#). *CoRR*, abs/1609.08144.
- 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.
- Tao Yu, Rui Zhang, Kai Yang, Michihito Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir Radev. 2018. [Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3911–3921, Brussels, Belgium. Association for Computational Linguistics.
- Mengjie Zhao, Yi Zhu, Ehsan Shareghi, Roi Reichart, Anna Korhonen, and Hinrich Schütze. 2020a. [A closer look at few-shot crosslingual transfer: Variance, benchmarks and baselines](#). *CoRR*, abs/2012.15682.
- Yanbin Zhao, Lu Chen, Zhi Chen, and Kai Yu. 2020b. [Semi-supervised text simplification with back-translation and asymmetric denoising autoencoders](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9668–9675.
- Victor Zhong, Mike Lewis, Sida I. Wang, and Luke Zettlemoyer. 2020. [Grounded adaptation for zero-shot executable semantic parsing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6869–6882, Online. Association for Computational Linguistics.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. [Seq2SQL: Generating structured queries from natural language using reinforcement learning](#). *CoRR*, abs/1709.00103.
- Qile Zhu, Haidar Khan, Saleh Soltan, Stephen Rawls, and Wael Hamza. 2020. [Don’t parse, insert: Multi-lingual semantic parsing with insertion based decoding](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 496–506, Online. Association for Computational Linguistics.

A Experimental Setup

	Train	Validation	Test	Total
ATIS				
EN	4,473	497	448	5,418
FR	4,473	497	448	5,418
PT	4,473	497	448	5,418
ES	4,473	497	448	5,418
DE	4,473	497	448	5,418
ZH	4,473	497	448	5,418
HI	—	—	442	442
TR	—	—	381	381
Overnight				
EN	8,754	2,188	2,740	13,682
DE	—	2,188	2,740	4,928
ZH	—	2,188	2,740	4,928

Table 4: Semantic parsing dataset partitions per language for ATIS (Dahl et al., 1994; Upadhyay et al., 2018; Xu et al., 2020) and Overnight (Wang et al., 2015; Sherborne et al., 2020). Each example is an utterance paired with a logical form.

Model	# Layers	# Parameters	# Vocabulary	Tokenization	# Languages
mBART-large (encoder)	12	408M	250,027	bBPE	25
XLM-R-large	24	550M	250,002	bBPE	100
mBART50-large (encoder)	12	408M	250,054	bBPE	50
ZX-PARSE	6 (decoder)	208M	593 (ATIS) 226 (Overnight)	Whitespace	8 (ATIS) 3 (Overnight)

Table 5: Pretrained model configurations and configuration for the trainable components of ZX-PARSE (e.g., the objectives). All models use a hidden dimension of 1,024, a feed-forward hidden projection of 4,096 and 16 heads per multi-head attention layer. For natural language, all models use byte-level BPE tokenization (Wang et al., 2020b) and logical forms are tokenized using whitespace.

Zero-shot Model Configuration The encoder, E , decoders, $\{D_{LF}, D_{NL}\}$, and embedding matrices all use a dimension size of 1,024 with the self-attention projection of 4,096 and 16 heads per layer. Both decoders are 6-layer stacks. Weights were initialized by sampling from normal distribution $\mathcal{N}(0, 0.02)$. The language prediction network is a two-layer feed-forward network projecting from z to 1,024 hidden units then to $|L|$ for L languages. L is six for experiments on ATIS and three for experiments on Overnight.

Configurations for models used in this work are reported in Table 5 with similar details for the objective components of ZX-PARSE. Initial experiments examined *XLM-R-base*, which is 12 layers opposed to 24, however, performance was significantly worse and, therefore, this model was not considered further. Experiments reported in Section 6 all use *mBART50* as the pre-trained encoder as all other pre-trained models performed significantly worse (see Appendix B). In all our experiments, we found that a randomly initialized decoder was superior to using pre-trained weights.

A complete outline of dataset partitions per language is shown in Table 4 for both datasets. ZX-PARSE uses only English training and validation data and tests on all additional languages. We did not use multi-lingual validation data as recommended in Keung et al. (2020) as this approach did not prove critically beneficial in early experiments and doing so would explode the data requirements for a multi-lingual system.

Experimental Setting The system was trained using the Adam optimizer (Kingma and Ba, 2015) with a learning rate of 1×10^{-4} , and a weight decay factor of 0.1. We use a “Noam” schedule for the learning rate (Vaswani et al., 2017) with a warmup of 5,000 steps. For pre-trained components, we fine-tune XLM-R and mBART

encoders with learning rate 1×10^{-5} but freeze the encoder when using mBART50. Loss weighting values for $\alpha_{\{LP, NL\}}$ were empirically optimized to $\{0.33, 0.1\}$ respectively from a range $\{0.5, 0.33, 0.1, 0.05, 0.01, 0.005, 0.001\}$. Batches during training were size 50 and homogeneously sampled from either \mathcal{S}_{LF} or \mathcal{S}_{NL} , with an epoch consuming one pass over both. Models were trained for a maximum of 100 epochs with early stopping. Model selection and hyper-parameters were tuned on the \mathcal{S}_{LF} validation set in English e.g., validation only evaluates performance on logical-form generation and not additional objectives. Test predictions were generated using beam search with 5 hypotheses.

For the reconstruction noising function, we use token masking to randomly replace u tokens in x with “<mask>” where u is sampled from $U(0, v)$. We found $v = 3$ as the empirically optimal maximum tokens to mask in an input. Similarly, we found $\gamma = 40$ optimal for the language prediction loss and $\tau = 0.5$ as the optimal sampling factor for translation versus reconstruction. This value of τ corresponds to using half the reconstruction data as mono-lingual utterances and half as bi-text paired with English.

Reproducibility All models were implemented using AllenNLP (Gardner et al., 2018) and PyTorch (Paszke et al., 2019), using pre-trained models from HuggingFace (Wolf et al., 2019). Each model is trained on 1 NVIDIA RTX3090 GPU in a cluster configuration, with no model requiring over 24 hours to complete training. Hyper-parameters were chosen by training a reference model for parsing English utterances and selecting the system with minimum validation loss. Our optimization grid-search explored: $\{6, 9, 12\}$ decoder layers; freezing or unfreezing the pre-trained encoder; $\{0, 1, 2\}$ additional encoder layers ap-

pended to the pre-trained encoder; learning rates of $1 \times 10^{\{-3, -4, -5\}}$ and a weight decay factor of $\{0, 0.1, 0.01\}$. Optimal parameters in these early tests were carried through for all additional models.

Additionally, we optimized hyper-parameters for auxiliary objectives through linear search with all other factors fixed. The upper limit, v , for the number of tokens to mask during reconstruction, $U(0, v)$, was optimized from integers 1-6. The MKQA dataset used for auxiliary tasks contains shorter sentences than prior work using masking, such as Lewis et al. (2020), and we observed that high levels of masking ultimately destroys the input sentence and handicaps the overall task. τ was optimized between values of 0.0 (e.g. ignore bi-text) to 1.0 (e.g. all data is used as bi-text) in increments of 0.1. Finally, we optimize the γ parameter within Equation 7 between $\{0, 5, 10, 20, 40, 50, 100\}$ on an approximately logarithmic scale. The optimal value of $\gamma = 40$ results in loss \mathcal{L}_{LF} reaching 99% of the maximum value at approximately 13.6% of training progress.

B Additional Results

We extend the results in Section 6 to include additional ablations for all pre-trained models. Table 6 details all results for ATIS across eight test languages. Additionally, complete results across all domains in Overnight are reported for English in Table 7, German in Table 8, and Chinese in Table 9. Table 3, comparing between reconstruction data sources, is expanded on for Overnight in Table 10. Finally, we present additional ablations to our model considering reconstruction language families in Table 11 and 12.

Which Pre-trained Encoder? Our full results using three different pre-trained encoders are outlined in Tables 6–9. Our experiments identify *mBART* as the weakest pre-trained model, reporting the lowest accuracies for all ATIS test languages. ZX-PARSE using *XLM-R* generally improved upon *mBART* for ATIS but proved worse for Overnight. As *XLM-R* is not pre-trained for sequence-to-sequence tasks, this result suggests this model could be poorer at representing input content in more complex queries. Despite being half the size of *XLM-R*, *mBART50* is the only pre-trained encoder able to perform competitively across all languages. Despite lower performance with different pre-trained models, we identify that introducing additional objectives yields improved

accuracy in most cases. Similar to our results using *mBART50*, we find that combining tasks is the optimal strategy for ZX-PARSE using either *XLM-R* or *mBART* as an encoder.

We additionally explored if pre-training is required for our approach by training a comparable model from scratch. While performance on English was similar to our best results, we found that cross-lingual transfer was extremely poor and these results are omitted due to negligible accuracies ($< 2\%$) for non-English languages. Overall, this suggests that our methodology is optimal when *aligning an existing multi-lingual latent space* rather than *inducing a multi-lingual latent space from scratch*.

Ablations of Reconstruction Language Data

We present ablations to our main experiments examining the influence of language similarity in reconstruction data for ATIS in Table 11 and for Overnight in Table 12. Similar to our results for Hindi and Turkish in Table 2, we find that using our auxiliary objectives in our model improves overall cross-lingual alignment in languages that we did not intentionally target with reconstruction data.

In our first case, we consider omitting the Romance genus languages (French, Spanish, Portuguese) from the reconstruction corpus for experiments on ATIS. The observed reduction in performance across all languages is likely a consequence of reduced training data leading to weaker cross-lingual alignment. Notably, this drop is largest for French (-11.2%) and Spanish (-7.1%). In contrast, the smallest reduction is for Chinese (-3.9%) and English (-2.8%). We additionally examine the effect of omitting the only Sino-Tibetan language (Chinese) from experiments on both ATIS and Overnight. While we observe a similar overall reduction in performance here – our notable finding is a larger reduction in parsing accuracy for Chinese across both ATIS (-17.0%) and Overnight (-11.1%). Without a similar language to Chinese (in the same family or with a similar orthography) in this experiment, we suggest there is little to “support” better cross-lingual alignment for Chinese relative to others. This contrasts with the performance drop for Romance languages, which are still relatively similar to English and German.

Overall, these ablations support that both *variety* and *similarity* are important for considering language data for auxiliary objectives. Performance on omitted languages can improve from a base-

	EN	FR	PT	ES	DE	ZH	HI	TR
Monolingual Training	77.2	67.8	66.1	64.1	66.6	64.9	—	—
Translate-Train	—	55.9	56.1	57.1	60.1	56.1	56.3	45.4
Translate-Test	—	58.2	57.3	57.9	56.9	51.4	52.6	52.7
ZX-PARSE using mBART								
D_{LF} only	74.6	35.4	18.3	55.6	35.9	10.8	10.7	21.4
$D_{LF} + D_{NL}$	77.7	27.9	17.6	54.5	34.5	10.3	12.6	21.1
$D_{LF} + LP$	75.3	32.3	15.3	49.3	32.3	7.4	10.7	19.2
$D_{LF} + LP + D_{NL}$	77.0	39.4	24.6	56.3	37.6	12.9	11.0	32.6
ZX-PARSE using XLM-R								
D_{LF} only	76.5	44.6	47.2	57.1	41.8	16.2	11.2	14.5
$D_{LF} + D_{NL}$	78.6	36.9	43.4	57.0	46.0	11.7	11.2	12.6
$D_{LF} + LP$	74.9	30.8	32.6	56.6	30.8	11.5	11.4	21.4
$D_{LF} + LP + D_{NL}$	78.2	48.1	46.0	60.8	55.4	18.3	18.3	35.8
ZX-PARSE using mBART50								
D_{LF} only	77.2	61.3	42.5	46.5	50.2	38.5	40.4	37.3
$D_{LF} + D_{NL}$	77.7	62.7	54.9	58.2	61.1	51.2	49.5	44.7
$D_{LF} + LP$	76.3	57.2	53.7	51.8	58.6	44.1	39.8	38.8
$D_{LF} + LP + D_{NL}$	76.9	70.2	63.4	59.7	69.3	60.2	54.9	48.3

Table 6: Complete denotation accuracy results for ATIS across all languages: English (EN), French (FR), Portuguese (PT), Spanish (ES), German (DE), Chinese (ZH), Hindi (HI) and Turkish (TR). Models shown use (i) no auxiliary objectives, (ii) logical form generation and reconstruction, (iii) logical form generation and language prediction and (iv) finally all objectives.

line, but better localization is achievable using a linguistically varied ensemble of languages closely modeling the desired languages to parse.

ZX-PARSE using mBART	Avg.	Ba.	Bl.	Ca.	Ho.	Pu.	Rec.	Res.	So.
D_{LF} only	81.0	89.0	64.7	81.5	77.8	77.6	88.0	86.1	83.6
$D_{LF} + D_{NL}$	81.7	89.0	65.7	86.3	77.2	81.4	86.4	85.2	82.5
$D_{LF} + LP$	79.8	85.9	65.9	82.1	74.1	77.0	88.0	83.3	81.9
$D_{LF} + D_{NL} + LP$	80.5	86.7	65.7	84.5	75.1	82.6	85.8	83.3	80.3
ZX-PARSE using XLM-R	Avg.	Ba.	Bl.	Ca.	Ho.	Pu.	Rec.	Res.	So.
D_{LF} only	82.3	89.3	67.7	85.1	75.7	85.7	89.8	81.0	84.4
$D_{LF} + D_{NL}$	81.7	89.3	61.4	84.5	77.2	82.6	88.3	86.6	83.8
$D_{LF} + LP$	76.7	77.0	61.7	74.4	75.7	74.5	86.1	85.2	79.4
$D_{LF} + D_{NL} + LP$	82.2	87.7	64.9	86.3	77.2	82.6	89.2	85.6	84.2
ZX-PARSE using mBART50	Avg.	Ba.	Bl.	Ca.	Ho.	Pu.	Rec.	Res.	So.
D_{LF} only	80.5	90.0	66.7	82.7	76.7	75.8	87.7	83.3	80.9
$D_{LF} + D_{NL}$	81.3	88.5	60.9	83.3	78.8	83.9	88.6	83.8	82.6
$D_{LF} + LP$	80.6	90.0	63.4	78.6	76.2	81.4	84.7	86.4	83.8
$D_{LF} + D_{NL} + LP$	81.9	87.7	65.4	84.5	77.8	81.4	88.0	87.0	83.1

Table 7: Denotation accuracy for the Overnight dataset (Wang et al., 2015) from **English** utterances. Domains are *Basketball, Blocks, Calendar, Housing, Publications, Recipes, Restaurants* and *Social Network*.

	Avg.	Ba.	Bl.	Ca.	Ho.	Pu.	Rec.	Res.	So.
Translate-Train	62.2	73.5	45.4	68.0	49.3	64.0	67.5	59.8	70.2
Translate-Test	60.1	75.7	50.9	61.0	55.6	50.4	69.9	46.3	71.4
ZX-PARSE using mBART									
D_{LF} only	40.8	60.1	42.6	32.1	42.9	23.6	39.8	35.6	49.9
$D_{LF} + D_{NL}$	39.0	55.8	41.9	39.3	35.4	20.5	38.9	32.9	47.2
$D_{LF} + LP$	38.6	61.1	45.6	28.0	45.5	14.9	38.9	18.1	57.1
$D_{LF} + D_{NL} + LP$	52.6	69.8	47.1	47.6	51.3	51.6	56.6	38.9	57.5
ZX-PARSE using XLM-R									
D_{LF} only	38.6	45.0	43.4	21.4	45.5	32.3	37.7	39.8	44.1
$D_{LF} + D_{NL}$	45.8	58.6	48.4	33.3	38.1	39.1	44.3	51.9	52.9
$D_{LF} + LP$	41.1	64.2	41.4	26.2	31.2	38.5	39.2	44.4	43.8
$D_{LF} + D_{NL} + LP$	49.0	68.3	48.6	48.2	42.3	46.6	48.2	34.7	55.4
ZX-PARSE using mBART50									
D_{LF} only	58.4	70.3	51.1	61.9	54.0	49.7	65.4	42.1	73.1
$D_{LF} + D_{NL}$	62.7	73.1	56.1	66.1	58.7	49.7	70.2	57.9	70.1
$D_{LF} + LP$	60.6	76.5	57.9	68.5	52.9	52.8	36.6	66.6	73.2
$D_{LF} + D_{NL} + LP$	66.2	79.8	60.4	72.6	60.3	62.1	45.8	74.4	73.9

Table 8: Denotation accuracy for the Overnight dataset (Wang et al., 2015) using the **German** test set from Sherborne et al. (2020). Domains are *Basketball, Blocks, Calendar, Housing, Publications, Recipes, Restaurants* and *Social Network*.

	Avg.	Ba.	Bl.	Ca.	Ho.	Pu.	Rec.	Res.	So.
Translate-Train	59.4	75.4	46.5	50.5	57.8	56.7	62.1	60.1	66.1
Translate-Test	48.1	62.3	39.6	49.8	43.1	48.3	51.4	29.2	61.2
ZX-PARSE using mBART									
D_{LF} only	19.4	11.5	23.6	23.8	33.9	6.8	21.1	12.5	22.4
$D_{LF} + D_{NL}$	16.7	1.0	28.1	20.8	32.8	9.3	15.7	7.9	18.1
$D_{LF} + LP$	17.0	15.9	26.8	7.7	28.0	2.5	24.4	5.6	25.5
$D_{LF} + D_{NL} + LP$	36.1	26.1	30.6	42.9	49.2	28.6	44.9	23.1	43.2
ZX-PARSE using XLM-R									
D_{LF} only	17.6	6.1	24.8	20.2	21.7	14.9	22.6	20.4	10.2
$D_{LF} + D_{NL}$	18.0	17.6	13.3	8.9	33.9	13.7	21.4	10.6	24.4
$D_{LF} + LP$	18.5	20.5	4.3	13.1	39.7	5.0	19.3	13.4	32.5
$D_{LF} + D_{NL} + LP$	22.7	18.4	30.3	15.5	37.0	13.0	19.6	7.9	39.5
ZX-PARSE using mBART50									
D_{LF} only	48.0	53.7	49.6	53.0	50.8	36.0	52.1	23.1	65.3
$D_{LF} + D_{NL}$	49.5	56.6	49.4	55.4	56.1	35.4	54.2	24.9	64.1
$D_{LF} + LP$	49.4	55.5	54.9	73.8	53.4	19.3	21.3	52.7	63.9
$D_{LF} + D_{NL} + LP$	60.0	59.4	57.4	74.4	62.2	41.1	59.3	57.9	68.4

Table 9: Denotation accuracy for the Overnight dataset (Wang et al., 2015) using the **Chinese** test set from Sherborne et al. (2020). Domains are *Basketball*, *Blocks*, *Calendar*, *Housing*, *Publications*, *Recipes*, *Restaurants* and *Social Network*.

	Avg.	Ba.	Bl.	Ca.	Ho.	Pu.	Rec.	Res.	So.
EN									
MKQA $\tau = 0.0$	81.3	89.3	63.9	82.7	79.4	82.6	83.3	87.7	81.9
MKQA $\tau = 0.5$	81.9	87.7	65.4	84.5	77.8	81.4	88.0	87.0	83.1
ParaCrawl $\tau = 0.0$	78.4	87.5	58.9	77.4	72.0	78.3	84.7	86.4	81.9
ParaCrawl $\tau = 0.5$	81.2	87.2	65.2	84.5	74.6	80.7	86.6	87.0	83.4
DE									
MKQA $\tau = 0.0$	64.3	78.8	56.6	68.5	58.7	46.6	70.8	59.3	75.5
MKQA $\tau = 0.5$	66.2	79.8	60.4	72.6	60.3	62.1	45.8	74.4	73.9
ParaCrawl $\tau = 0.0$	62.4	76.3	53.7	64.4	57.6	51.3	50.1	72.3	73.6
ParaCrawl $\tau = 0.5$	63.2	78.0	55.9	69.0	57.1	56.5	44.0	70.2	74.4
ZH									
MKQA $\tau = 0.0$	52.7	59.1	50.1	67.9	54.5	41.6	59.6	20.8	67.6
MKQA $\tau = 0.5$	60.0	59.4	57.4	74.4	62.2	41.1	59.3	57.9	68.4
ParaCrawl $\tau = 0.0$	51.1	56.3	55.5	64.8	54.5	26.6	27.0	53.8	70.5
ParaCrawl $\tau = 0.5$	52.9	58.8	55.9	71.4	60.8	30.4	30.1	46.8	68.6

Table 10: Denotation accuracy for the Overnight dataset (Wang et al., 2015) compared across reconstruction data usage for English, German and Chinese. We compare between MKQA (Longpre et al., 2020) and ParaCrawl (Bañón et al., 2020) with additional contrast between using reconstruction data as monolingual utterances (e.g. $\tau = 0.0$) or with some proportion as bi-text where the target sequence is replaced with the parallel English utterance (e.g. $\tau = 0.5$). Domains are *Basketball*, *Blocks*, *Calendar*, *Housing*, *Publications*, *Recipes*, *Restaurants* and *Social Network*. Best results for each language are bolded.

	EN	FR	PT	ES	DE	ZH	HI	TR
Omit Romance Genus (FR, ES, PT)	74.1	59.0	58.0	52.6	64.4	56.3	45.1	39.7
Omit Sino-Tibetan Family (ZH)	74.1	65.1	58.3	55.8	65.1	43.2	42.7	38.7

Table 11: Denotation accuracy for the ATIS dataset compared across reconstruction language ablations. We experiment with omitting the Romance genus (French, Spanish, Portuguese) and the Sino-Tibetan family (ZH only). Language groupings are sourced from [Dryer and Haspelmath \(2013\)](#).

	Avg.	Ba.	Bl.	Ca.	Ho.	Pu.	Rec.	Res.	So.
Omit Sino-Tibetan Family (ZH)									
EN	80.7	87.7	63.9	81.5	77.8	82.0	85.6	84.6	82.2
DE	61.8	75.7	53.6	63.7	56.6	50.3	49.5	71.7	73.5
ZH	48.9	54.7	52.9	63.1	51.9	23.6	25.9	51.5	67.5

Table 12: Denotation accuracy for the Overnight dataset compared across reconstruction language ablations. We report results for English, German and Chinese when omitting the Sino-Tibetan family (ZH only). Language groupings are sourced from [Dryer and Haspelmath \(2013\)](#) and domains are *Basketball, Blocks, Calendar, Housing, Publications, Recipes, Restaurants* and *Social Network*.