

Span-Selective Linear Attention Transformers for Effective and Robust Schema-Guided Dialogue State Tracking

Björn Bebensee Haejun Lee

Samsung Research

{b.bebensee, haejun82.lee}@samsung.com

Abstract

In schema-guided dialogue state tracking models estimate the current state of a conversation using natural language descriptions of the service schema for generalization to unseen services. Prior generative approaches which decode slot values sequentially do not generalize well to variations in schema, while discriminative approaches separately encode history and schema and fail to account for inter-slot and intent-slot dependencies. We introduce SPLAT, a novel architecture which achieves better generalization and efficiency than prior approaches by constraining outputs to a limited prediction space. At the same time, our model allows for rich attention among descriptions and history while keeping computation costs constrained by incorporating linear-time attention. We demonstrate the effectiveness of our model on the Schema-Guided Dialogue (SGD) and MultiWOZ datasets. Our approach significantly improves upon existing models achieving 85.3 JGA on the SGD dataset. Further, we show increased robustness on the SGD-X benchmark: our model outperforms the more than 30× larger D3ST-XXL model by 5.0 points.

1 Introduction

Dialogue State Tracking (DST) refers to the task of estimating and tracking the dialogue state consisting of the user’s current intent and set of slot-value pairs throughout the dialogue (Williams et al., 2013). Traditional approaches to DST assume a fixed ontology and learn a classifier for each slot (Chao and Lane, 2019). However, in real-world applications services can be added or removed requiring the model to be re-trained each time the ontology changes. Recently more flexible schema-guided approaches which take as input natural language descriptions of all available intents and slots and thus can be applied zero-shot to new services have been gaining popularity (Rastogi et al., 2020; Feng et al., 2021; Zhao et al., 2022; Gupta et al., 2022).

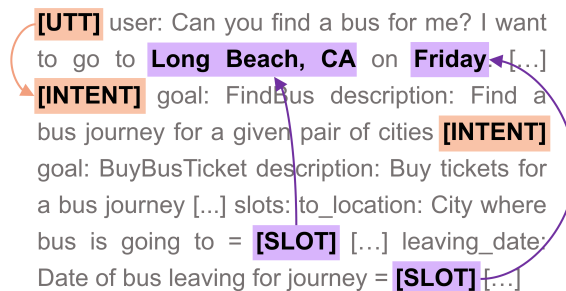


Figure 1: Span selection for schema-guided dialogue in practice. [SLOT] encodes the semantics of the natural language description of “to_location” and is matched with the span representation of “Long Beach, CA”. Similarly [UTT] encodes the semantics of the current utterance and is matched with the target [INTENT] encoding.

Discriminative DST models are based on machine reading comprehension (MRC) methods, meaning they extract and fill in non-categorical slot values directly from the user utterances (Chao and Lane, 2019; Ruan et al., 2020; Zhang et al., 2021). We use the terms discriminative and extractive interchangeably when referring to these methods. Generative DST models leverage seq2seq language models which conditioned on the dialog history and a prompt learn to sequentially generate the appropriate slot values. Prior generative methods do not generalize well to variations in schema (Lee et al., 2021, 2022; Zhao et al., 2022) whereas discriminative methods separately encode history and schema and fail to account for inter-slot and intent-slot dependencies.

In this work we introduce the **SPan-Selective Linear Attention Transformer**, short **SPLAT**, a novel architecture designed to achieve better generalization, robustness and efficiency in DST than existing approaches. SPLAT is fully extractive and, unlike prior generative approaches, constrains the output space to only those values contained in the input sequence. Figure 1 shows an example




Figure 2: An overview over the SPLAT model architecture. Intent scores are computed using the utterance representation $h_i^{[UTT]}$ and intent representations $h_k^{[INTENT]}$. A span encoder computes span representations $h_{mn}^{[SPAN]}$ for all spans x_m, \dots, x_n . The target span is selected by matching the slot query $h_q^{[SLOT]}$ to the target span $h_{ij}^{[SPAN]}$.

of the key idea behind our approach. We jointly encode the natural language schema and full dialogue history allowing for a more expressive contextualization. Spans in the input are represented by aggregating semantics of each individual span into a single representation vector. Then we take a contrastive query-based pointer network approach (Vinyals et al., 2015) to match special query tokens to the target slot value’s learned span representation in a single pass.

Our main contributions are as follows:

- We propose novel span-selective prediction layers for DST which provide better generalization and efficiency by limiting the prediction space and inferring all predictions in parallel. We achieve state-of-the-art performance on the SGD-X benchmark outperforming the 30× larger D3ST by 5.0 points.
- We adopt a Linear Attention Transformer which allows more expressive contextualization of the dialogue schema and dialogue history with constrained prediction time. We show our model already outperforms other models with similar parameter budgets even without other modules we propose in Table 1 and 5.
- We pre-train SPLAT for better span representations with a recurrent span selection objective yielding significant further span prediction performance gains of up to 1.5 points.

2 Approach

2.1 Task Formulation

For a given dialog of T turns let U describe the set of utterances in the dialog history $U = \{u_1, \dots, u_T\}$. Each u_i can represent either a user or a system utterance. The system is providing some service to the user defined by a service schema S . The service schema consists of a set of intents $I = \{i_1, \dots, i_K\}$ and their intent descriptions $D^{\text{intent}} = \{d_1^{\text{intent}}, \dots, d_K^{\text{intent}}\}$ as well as a set of slots $S = \{s_1, \dots, s_L\}$ and their slot descriptions $D^{\text{slot}} = \{d_1^{\text{slot}}, \dots, d_L^{\text{slot}}\}$.

In practice we prepend each u_i with the speaker name (user or system) and a special utterance query token [UTT] which will serve as the encoding of the system-user utterance pair.

Each d_i^{slot} consists of the slot name, a natural language description of the semantics of the slot and for categorical values an enumeration of all possible values this slot can assume. We also append a special slot query embedding token [SLOT] which serves as the slot encoding.

Some slot values are shared across all slots and their representation can be modeled jointly. Unless denoted otherwise these shared target values T are special tokens [NONE] and [DONTCARE] which correspond to the "none" and "dontcare" slot values in SGD and MultiWOZ.

2.2 Joint Encoding with Linear Attention

Linear Attention Transformers. In order to better capture the semantics of the input and to al-

low for a longer context as well as all the relevant schema descriptions to be encoded jointly we use a Transformer (Vaswani et al., 2017) with linear-time attention. Instead of computing the full attention matrix as the original Transformer does, its linear attention variants compute either an approximation of it (Choromanski et al., 2021) or only compute full attention for a fixed context window of size w around the current token and additional n_{global} global tokens, thus lowering the complexity of the attention computation from $\mathcal{O}(n^2)$ for a sequence of length n to $\mathcal{O}(w + n_{\text{global}})$ (Beltagy et al., 2020; Zaheer et al., 2020).

We focus on the windowed variant and incorporate it to DST. We denote the Linear Attention Transformer with selective global attention parametrized by θ with input sequence \mathcal{I} and its subset of global input tokens $\mathcal{G} \subseteq \mathcal{I}$, i.e. inputs corresponding to tokens at positions that are attended using the global attention mechanism, as $\text{LAT}(\mathcal{I}; \mathcal{G}; \theta)$. While we choose the Longformer (Beltagy et al., 2020) for our implementation, in practice any variants with windowed and global attention can be used instead.

Joint encoding. The full input sequence of length N is given as the concatenation of its components. We define the set of globally-attended tokens as the union of sets of tokens corresponding to the intent descriptions D^{intent} , the slot descriptions D^{slot} , and the shared target values T . Then, the joint encoding of N hidden states is obtained as the output of the last Transformer layer as

$$\begin{aligned} \mathcal{I} &= [\text{CLS}] U [\text{SEP}] T D^{\text{intent}} D^{\text{slot}} [\text{SEP}] \\ \mathcal{G} &= T \cup D^{\text{intent}} \cup D^{\text{slot}} \\ E &= \text{LAT}(\mathcal{I}; \mathcal{G}; \theta). \end{aligned} \quad (1)$$

2.3 Intent Classification

Let $\mathbf{x}_i^{[\text{UTT}]}$ denote the representation of the encoded [UTT] token corresponding to the i -th turn. Given the encoded sequence E , we obtain the final utterance representations by feeding $\mathbf{x}_i^{[\text{UTT}]}$ into the utterance encoder. Similarly for each intent $I = \{i_1, \dots, i_t\}$ and its respective [INTENT] token, we obtain final intent representations using the intent encoder:

$$\begin{aligned} \mathbf{h}_i^{[\text{UTT}]} &= \text{LN}(\text{FFN}(\mathbf{x}_i^{[\text{UTT}]})) \\ \mathbf{h}_j^{[\text{INTENT}]} &= \text{LN}(\text{FFN}(\mathbf{x}_j^{[\text{INTENT}]})) \end{aligned} \quad (2)$$

Here LN refers to a LayerNorm and FFN to a feed-forward network.

We maximize the dot product similarity between each utterance representation and the ground truth active intent’s representation via cross-entropy:

$$\begin{aligned} \text{score}_{i \rightarrow j} &= \text{sim}(\mathbf{h}_i^{[\text{UTT}]}, \mathbf{h}_j^{[\text{INTENT}]}) \\ \mathcal{L}_{\text{intent}} &= -\frac{1}{T} \sum_{i=1}^T \log \frac{\exp(\text{score}_{i \rightarrow j})}{\sum_{k=1}^K \exp(\text{score}_{i \rightarrow k})} \cdot \mathbb{1}_{\text{GT}} \end{aligned} \quad (3)$$

where K is the number of intents and $\mathbb{1}_{\text{GT}}$ is an indicator function which equals 1 if and only if j is the ground truth matching i .

2.4 Span Pointer Module

We introduce a novel Span Pointer Module which computes span representations via a span encoder and extracts slot values by matching slot queries via a similarity-based span pointing mechanism (Vinyals et al., 2015).

First, for any given span of token representations $\mathbf{x}_i, \dots, \mathbf{x}_j$ in the joint encoding E we obtain the span representation $\mathbf{h}_{ij}^{\text{SPAN}}$ by concatenating the span’s first and last token representation and feeding them into a 2-layer feed-forward span encoder (Joshi et al., 2020):

$$\begin{aligned} \mathbf{y}_{ij} &= [\mathbf{x}_i; \mathbf{x}_j] \\ \mathbf{h}_{ij}^{\text{SPAN}} &= \text{LN}(\text{FFN}_{\text{GeLU}}(\mathbf{y}_{ij})) \times \text{n_layers} \end{aligned} \quad (4)$$

Similarly, for each slot token representation $\mathbf{x}^{[\text{SLOT}]}$ in E we compute a slot query representation $\mathbf{h}^{[\text{SLOT}]}$ with a 2-layer feed-forward slot encoder:

$$\mathbf{h}^{[\text{SLOT}]} = \text{LN}(\text{FFN}_{\text{GeLU}}(\mathbf{x}^{[\text{SLOT}]})) \times \text{n_layers} \quad (5)$$

Given slots $S = \{s_1, \dots, s_L\}$ and corresponding slot query representations $\mathbf{h}_1^{[\text{SLOT}]}, \dots, \mathbf{h}_L^{[\text{SLOT}]}$ we score candidate target spans by dot product similarity of the slot queries with their span representations. That is, for each slot query q with ground truth target span x_i, \dots, x_j we maximize $\text{sim}(\mathbf{h}_q^{[\text{SLOT}]}, \mathbf{h}_{ij}^{\text{SPAN}})$ by cross-entropy. The loss function is given by

$$\begin{aligned} \text{score}_{q \rightarrow ij} &= \text{sim}(\mathbf{h}_q^{[\text{SLOT}]}, \mathbf{h}_{ij}^{\text{SPAN}}) \\ \mathcal{L}_{\text{slot}} &= -\frac{1}{L} \sum_{q=1}^L \log \frac{\exp(\text{score}_{q \rightarrow ij})}{\sum_{k=1}^K \exp(\text{score}_{q \rightarrow k})} \cdot \mathbb{1}_{\text{GT}} \end{aligned} \quad (6)$$

where L is the number of slots and K is the number of spans. $\text{sim}(\mathbf{h}_q^{[\text{SLOT}]}, \mathbf{h}_{ij}^{\text{SPAN}})$ denotes the

similarity between the q -th slot query representation and the span representation of its ground truth slot value.

It is computationally too expensive to compute span representations for all possible spans. In practice however the length of slot values rarely exceeds some L_{ans} . Thus, we limit the maximum span length to L_{ans} and do not compute scores for spans longer than this threshold. This gives us a total number of $N \cdot L_{\text{ans}}$ candidate spans.

Joint optimization. We optimize the intent and slot losses jointly via the following objective:

$$\mathcal{L} = \frac{\mathcal{L}_{\text{slot}} + \mathcal{L}_{\text{intent}}}{2} \quad (7)$$

2.5 Pre-Training via Recurrent Span Selection

Since the span pointer module relies on span embedding similarity for slot classification we believe it is crucial to learn good and robust span representations. In order to improve span representations for down-stream applications to DST we pre-train SPLAT in a self-supervised manner using a modified recurrent span selection objective (Ram et al., 2021).

Given an input text \mathcal{I} let $\mathcal{R} = \{\mathcal{R}_1, \dots, \mathcal{R}_a\}$ be the clusters of identical spans that occur more than once. Following Ram et al. (2021) we randomly select a subset $\mathcal{M} \subseteq \mathcal{R}$ of J recurring spans such that the number of their occurrences sums up to a maximum of 30 occurrences. Then, for each selected cluster of recurring spans \mathcal{M}_j we randomly replace all but one occurrence with the query token [SLOT].

The slot query tokens act as the queries while the respective unmasked span occurrences act as the targets. Unlike the original recurrent span selection objective we do not use separate start and end pointers for the target spans but instead use our Span Pointer Module to learn a single representation for each target span.

We pre-train SPLAT to maximize the dot product similarity between the query token and the unmasked target span representation. The loss for the j -th cluster of identical masked spans is given by Equation (6) and the total loss is given as the sum of losses of over all clusters.

Effectively each sentence containing a masked occurrence of the span acts as the span description while the target span acts as the span value. This can be seen as analogous to slot descriptions and slot values in DST.

3 Experimental Setup

We describe our experimental setup including datasets used for pre-training and evaluation, implementation details, baselines and evaluation metrics in detail below.

3.1 Benchmark Datasets

We conduct experiments on the Schema-Guided Dialogue (SGD) (Rastogi et al., 2020), SGD-X (Lee et al., 2022) and MultiWOZ 2.2 (Zang et al., 2020) datasets.

Schema-Guided Dialogue. Unlike other task-oriented dialogue datasets which assume a single, fixed ontology at training and test time the SGD dataset includes new and unseen slots and services in the test set. This allows us to not only measure DST performance but also zero-shot generalization to unseen services. The dataset includes natural language descriptions for all intents and slots in its schema. We follow the standard evaluation setting and data split suggested by the authors.

SGD-X. The SGD-X benchmark is an extension of the SGD dataset which provides five additional schema variants of different linguistic styles which increasingly diverge in style from the original schema with v_1 being most similar and v_5 least similar. We can evaluate our model’s robustness to variations in schema descriptions by training our model on SGD and comparing evaluation results using the different included schema variants.

MultiWOZ. The MultiWOZ dataset is set of human-human dialogues collected in the Wizard-of-OZ setup. Unlike in SGD the ontology is fixed and there are no unseen services at test time. There are multiple updated versions of the original MultiWOZ dataset (Budzianowski et al., 2018): MultiWOZ 2.1 (Eric et al., 2020) and MultiWOZ 2.2 (Zang et al., 2020) fix annotation errors of previous versions, MultiWOZ 2.3 (Han et al., 2021) is based on version 2.1 and adds co-reference annotations, MultiWOZ 2.4 (Ye et al., 2022) is also based on version 2.1 and includes test set corrections. However, MultiWOZ 2.2 is the only version of the dataset which includes a fully defined schema matching the ontology. We therefore choose the MultiWOZ 2.2 dataset for our experiments. We follow the standard evaluation setting and data split.

3.2 Evaluation Metrics

In line with prior work (Rastogi et al., 2020) we evaluate our approach according to the following two metrics.

Intent Accuracy: For intent detection the intent accuracy describes the fraction of turns for which the active intent has been correctly inferred.

Joint Goal Accuracy (JGA): For slot prediction JGA describes the fraction of turns for which all slot values have been predicted correctly. Following the evaluation setting from each dataset we use a fuzzy matching score for slot values in SGD and exact match in MultiWOZ.

3.3 Implementation Details

We base our implementation on the Longformer code included in the HuggingFace Transformers library (Wolf et al., 2020) and continue training from the base model (110M parameters) and large model (340M parameters) checkpoints. We keep the default Longformer hyperparameters in place, in particular we keep the attention window size set to 512. The maximum sequence length is 4096. During pre-training we train the base model for a total of 850k training steps and the large model for 800k training steps. During fine-tuning we train all models for a single run of 10 epochs and choose the model with the highest joint goal accuracy on the development set. We use the Adam optimizer (Kingma and Ba, 2014) with a maximum learning rate of 10^{-5} which is warmed up for the first 10% of steps and subsequently decays linearly. We set the batch size to 32 for base models and to 16 for large models. We pre-train SPLAT on English Wikipedia. Specifically we use the KILT Wikipedia snapshot¹ from 2019 (Petroni et al., 2021) as provided by the HuggingFace Datasets library (Lhoest et al., 2021).

For both SGD and MultiWOZ we set the shared target values T as the [NONE] and [DONTCARE] tokens and include a special intent with the name "NONE" for each service which is used as the target intent when no other intent is active. We set the maximum answer length L_{ans} to 30 tokens.

All experiments are conducted on a machine with eight A100 80GB GPUs. A single training run takes around 12 hours for the base model and 1.5 days for the large model.

¹https://huggingface.co/datasets/kilt_wikipedia

4 Evaluation

We evaluate the effectiveness of our model through a series of experiments designed to answer the following questions: 1) How effective is the proposed model architecture at DST in general? 2) Does the model generalize well to unseen services? 3) Is the model robust to changes in schema such as different slot names and descriptions? 4) Which parts of the model contribute most to its performance?

4.1 Baselines

We compare our model to various discriminative and generative baseline approaches. Note that not all of them are directly comparable due to differences in their experimental setups.

Extractive baselines. SGD baseline (Rastogi et al., 2020) is a simple extractive BERT-based model which encodes the schema and last utterance separately and uses the embeddings in downstream classifiers to predict relative slot updates for the current turn. SGP-DST (Ruan et al., 2020) and DS-DST (Zhang et al., 2020) are similar but jointly encode utterance and slot schema. Multi-Task BERT (Kapelonis et al., 2022) is also similar but uses system action annotations which include annotations of slots offered or requested by the system (e.g. "[ACTION] Offer [SLOT] location [VALUE] Fremont"). paDST (Ma et al., 2019) combines an extractive component for non-categorical slots with a classifier that uses 83 hand-crafted features (including system action annotations) for categorical slots. Additionally it augments training data via back-translation achieving strong results but making a direct comparison difficult. LUNA (Wang et al., 2022) separately encodes dialogue history, slots and slot values and learns to first predict the correct utterance to condition the slot value prediction on.

Generative baselines. Seq2Seq-DU (Feng et al., 2021) first separately encodes utterance and schema and then conditions the decoder on the cross-attended utterance and schema embeddings. The decoder generates a state representation consisting of pointers to schema elements and utterance tokens. AG-DST (Tian et al., 2021) takes as input the previous state and the current turn and learns to generate the new state in a first pass and correcting mistakes in a second generation pass. AG-DST does not condition generation on the schema and slot semantics are learned implic-

Model	Pretrained Model	Single-Pass	Intent	JGA
<i>With system action annotations</i>				
MT-BERT (Kapelonis et al., 2022)	BERT-base (110M)	✗	94.7	82.7
paDST (Ma et al., 2019)	XLNet-large (340M)	✗	94.8	86.5
<i>No additional data</i>				
SGD baseline (Rastogi et al., 2020)	BERT-base (110M)	✗	90.6	25.4
MT-BERT (Kapelonis et al., 2022)	BERT-base (110M)	✗	-	71.9
DaP (ind) (Lee et al., 2021)	T5-base (220M)	✗	90.2	71.8
SGP-DST (Ruan et al., 2020)	T5-base (220M)	✗	91.8	72.2
D3ST (Base) (Zhao et al., 2022)	T5-base (220M)	✓	97.2	72.9
D3ST (Large) (Zhao et al., 2022)	T5-large (770M)	✓	97.1	80.0
D3ST (XXL) (Zhao et al., 2022)	T5-XXL (11B)	✓	98.8	86.4
SPLAT (Base)	Longformer-base (110M)	✓	96.7	80.1
SPLAT (Large)	Longformer-large (340M)	✓	97.6	85.3

Table 1: Results on the SGD test set.

Model	Pretrained Model	Single-Pass	Intent	JGA
DS-DST [†] (Zhang et al., 2020)	BERT-base (110M)	✗	-	51.7
Seq2Seq-DU (Feng et al., 2021)	BERT-base (110M)	✓	90.9	54.4
LUNA (Wang et al., 2022)	BERT-base (110M)	✗	-	56.1
AG-DST (Tian et al., 2021)	GPT-2 (117M)	✗ [‡]	-	56.1
AG-DST (Tian et al., 2021)	PLATO-2 (310M)	✗ [‡]	-	57.3
DaP (seq) (Lee et al., 2021)	T5-base (220M)	✓	-	51.2
DaP (ind) (Lee et al., 2021)	T5-base (220M)	✗	-	57.5
D3ST (Base) (Zhao et al., 2022)	T5-base (220M)	✓	-	56.1
D3ST (Large) (Zhao et al., 2022)	T5-large (770M)	✓	-	54.2
D3ST (XXL) (Zhao et al., 2022)	T5-XXL (11B)	✓	-	58.7
SPLAT (Base)	Longformer-base (110M)	✓	91.4	56.6
SPLAT (Large)	Longformer-large (340M)	✓	91.5	57.4

Table 2: Results on the MultiWOZ 2.2 test set. Results denoted by [†] were reported in the original MultiWOZ 2.2 paper (Zang et al., 2020). [‡]: AG-DST uses a fixed two-pass generation procedure.

itly so it is unclear how well AG-DST transfers to new services. DaP (Lee et al., 2021) comes in two variants which we denote as DaP (seq) and DaP (ind). DaP (ind) takes as input the entire dialogue history and an individual slot description and decodes the inferred slot value directly but requires one inference pass for each slot in the schema. DaP (seq) instead takes as input the dialogue history and the sequence of all slot descriptions and decodes all inferred slot values in a single pass. D3ST (Zhao et al., 2022) takes a similar approach and decodes the entire dialogue state including the active intent in a single pass. Categorical slot values are predicted via an index-picking mechanism.

4.2 Main Results

Schema-Guided Dialogue. Table 1 shows results on the SGD test set. We report results for intent accuracy and JGA. We find that our model significantly outperforms models of comparable size in terms of JGA. In particular our 110M parameter SPLAT base model outperforms the 220M model D3ST base model by 7.2 JGA points and even achieves comparable performance to the much larger D3ST large model. Going from SPLAT base to SPLAT large we observe a significant performance improvement. In particular SPLAT large outperforms the D3ST large model by 5.3 JGA and nearly achieves comparable performance to the

Model	Params.	Orig.	Avg. (v_1-v_5)	Avg. Δ	Max Δ
DaP (ind) (Lee et al., 2021)	220M	71.8	64.0	-7.8	-
SGP-DST (Ruan et al., 2020)	220M	72.2 / 60.5*	49.9*	-10.6	-
D3ST (Large) (Zhao et al., 2022)	770M	80.0	75.3	-4.7	-10.9
D3ST (XXL) (Zhao et al., 2022)	11B	86.4	77.8	-8.6	-17.5
SPLAT (Base)	110M	80.1	76.0	-4.1	-7.8
SPLAT (Large)	340M	85.3	82.8	-2.5	-5.3

Table 3: Joint goal accuracy on the five different SGD-X schema variants. Results denoted by * are based on a reimplementation in the SGD-X paper which could not reproduce the original results.

Model	Params.	Seen	Unseen	Overall
SGP-DST ¹	220M	88.0	67.0	72.2
D3ST (Base) ²	220M	92.5	66.4	72.9
D3ST (Large) ²	770M	93.8	75.4	80.0
D3ST (XXL) ²	11B	95.8	83.3	86.4
SPLAT (Base)	110M	94.5	75.2	80.1
SPLAT (Large)	340M	94.6	82.2	85.3

Table 4: Joint goal accuracy on the SGD test set on seen and unseen services. Baseline results are reported by ¹Ruan et al. (2020) and ²Zhao et al. (2022) respectively.

more than 30 \times larger D3ST XXL model.

We note that although paDST achieves the best performance of all baseline models in terms of JGA, it is not directly comparable because it is trained with hand-crafted features and additional back-translation data for training which has been shown to significantly improve robustness and generalization to unseen descriptions in schema-guided DST (Lee et al., 2022). Similarly, although Multi-Task BERT achieves good performance this can mostly be attributed to the use of system action annotation as Kapelonis et al. (2022) themselves demonstrate. Without system action annotations its performance drops to 71.9 JGA.

In terms of intent accuracy SPLAT base slightly underperforms D3ST base and D3ST large by 0.5 and 0.4 JGA while SPLAT large achieves better performance and slightly improves upon the D3ST large performance. Overall, SPLAT achieves strong performance on SGD.

MultiWOZ. Table 2 shows results on the MultiWOZ 2.2 test set. As the majority of papers does not report intent accuracy on MultiWOZ 2.2 we focus our analysis on JGA. We find that SPLAT base outperforms most similarly-sized models including D3ST base and large and that SPLAT large

performs better than all models aside from the more than 30 \times larger D3ST XXL. The notable exceptions to this are AG-DST and DaP (ind). AG-DST large achieves performance that is similar to SPLAT large using a generative approach but it performs two decoding passes, employs a negative sampling strategy to focus on more difficult examples and is trained for a fixed schema. DaP (ind) also achieves similar performance but needs one inference pass for every slot at every turn of the dialogue. This is much slower and simply not realistic in real-world scenarios with a large number of available services and slots. The sequential variant DaP (seq) which instead outputs the full state in a single pass performs much worse.

Comparison. While DaP (ind) shows strong performance that matches SPLAT on MultiWOZ, SPLAT fares much better than DaP (ind) on the SGD dataset. This can be seen to be indicative of a stronger generalization ability as MultiWOZ uses the same schema at training and test time whereas SGD includes new, unseen services at test time and thus requires the model to generalize and understand the natural language schema descriptions.

4.3 Robustness

DST models which take natural language descriptions of intents and slots as input naturally may be sensitive to changes in these descriptions. In order to evaluate the robustness of our model to such linguistic variations we perform experiments on the SGD-X benchmark. The SGD-X benchmark comes with five crowd-sourced schema variants v_1 to v_5 which increasingly diverge in style from the original schema. We train SPLAT on SGD and evaluate it on the test set using all five different schema variants.

As shown in Table 3, our model is considerably more robust to linguistic variations than all of the

Model	Params.	SGD		MultiWOZ	
		Intent	JGA	Intent	JGA
Longformer (extr.)	110M	95.9	78.5	91.4	55.5
+ SPM	110M	97.0	79.0	91.4	56.1
+ SPM + RSS-PT	110M	96.7	80.1	91.4	56.6
Longformer (extr.)	340M	97.5	83.5	91.4	56.3
+ SPM	340M	98.2	83.8	91.4	57.8
+ SPM + RSS-PT	340M	97.6	85.3	91.5	57.4

Table 5: Ablation results on the SGD and MultiWOZ test sets. Longformer (extr.) refers to an extractive model with no span representations and simple start and end pointers for answer prediction, SPM refers to the Span Pointer Module and RSS-PT to pre-training with the Recurrent Span Selection objective.

baseline models. On average SPLAT base loses around 4.1 points and SPLAT large loses around 2.5 points joint goal accuracy when compared to the results on the original schema. When considering the mean performance across all unseen schema variants SPLAT large significantly outperforms the more than $30\times$ larger D3ST XXL by 5.0 points. These observations also hold for the base model: the 110M parameter SPLAT base even outperforms the 11B parameter D3ST XXL on the least similar schema variant v_5 further highlighting the superior robustness of our model.

4.4 Generalization to unseen domains

In real-world scenarios virtual assistants cover a wide range of services that can change over time as new services get added or removed requiring dialogue models to be re-trained. One of our goals is to improve generalization to unseen services thus minimizing the need for expensive data collection and frequent re-training. As the MultiWOZ dataset does not include any new and unseen services in its test set our analysis primarily focuses on the SGD dataset. Table 4 shows results on SGD with a separate evaluation for dialogues in seen and unseen domains. We find that SPLAT achieves better generalization and improves upon the baselines with a particularly large margin on unseen domains where SPLAT base outperforms D3ST base by 8.8 points and SPLAT base outperforms D3ST large by 6.8 points.

4.5 Ablation Study

We conduct an ablation study to identify the contribution of the different components to model performance. Results can be seen in Table 5. We

compare a variant of our model which does not use span representations (referred to as “Longformer (extractive)”) but instead has two pointers [SL0T] and [/SL0T] which are used to select the start and end of the answer span. We find that using the Span Pointer Module to directly select the span improves performance across both model sizes and datasets. Furthermore, we find pre-training our model for better span representations via the recurrent span selection task to be crucial giving further significant performance gains for all sizes and datasets except the 340M parameter model on the MultiWOZ dataset where JGA slightly deteriorates. Across both model sizes gains from RSS pre-training are larger on the SGD dataset. We hypothesize that this may be attributed to better span representations learned through RSS pre-training which in turn generalize better to unseen domains.

5 Related Work

Extractive DST. Following the traditional extractive setting [Chao and Lane \(2019\)](#) propose a machine reading comprehension (MRC) approach which decodes slot values turn-by-turn using a different learned classifier for each slot. As a classifier has to be learned for each new slot this approach cannot easily be transferred to new slots.

Schema-guided approaches address this by explicitly conditioning predictions on a variable schema which describes intents and slots in natural language ([Rastogi et al., 2020](#)). Both [Ruan et al. \(2020\)](#) and [Zhang et al. \(2021\)](#) introduce schema-guided models but predict slots independently from one another requiring multiple encoder passes for each turn and failing to model intent-slot and inter-slot dependencies. [Ma et al. \(2019\)](#) use MRC for non-categorical and handcrafted features for categorical slots.

Generative DST. In an attempt to address the lack of ability to generalize to new domains and ontologies, [Wu et al. \(2019\)](#) propose incorporating a generative component into DST. Based on the dialog history and a domain-slot pair a state generator decodes a value for each slot. However as each slot is decoded independently the approach cannot model slot interdependencies. [Feng et al. \(2021\)](#) instead generate the entire state as a single sequence of pointers to the dialogue history and input schema but separately encode history and schema. [Zhao et al. \(2021\)](#) model DST fully as a text-to-text problem and directly generate the entire current state as

a string. Lin et al. (2021) transfer a language model fine-tuned for seq2seq question answering to DST zero-shot using the dialog history as context and simply asking the model for the slot values. By also including a natural language schema in the input, Zhao et al. (2022) show that full joint modeling and rich attention between history and schema lead to better results in DST. Furthermore, they demonstrate the flexibility of this fully language driven paradigm by leveraging strong pre-trained language models for cross-domain zero-shot transfer to unseen domains. Gupta et al. (2022) show the effectiveness of using demonstrations of slots being used in practice instead of a natural language descriptions in the prompt.

6 Conclusion

In this work we introduced SPLAT, a novel architecture for schema-guided dialogue state tracking which learns to infer slots by learning to select target spans based on natural language descriptions of slot semantics, and further showed how to pre-train SPLAT via a recurrent span selection objective for better span representations and a stronger slot prediction performance. We find that our proposed architecture yields significant improvements over existing models and achieving 85.3 JGA on the SGD dataset and 57.4 JGA on the MultiWOZ dataset. In schema-guided DST the ability to generalize to new schemas and robustness to changes in schema descriptions is of particular interest. We demonstrated that our model is much more robust to such changes in experiments on the SGD-X benchmark where SPLAT outperforms the more than $30\times$ larger D3ST-XXL model by 5.0 points.

Limitations

One trade-off of limiting the prediction space using an extractive pointer module is that it does not support prediction of multiple slot values which is necessary for some dialogues in the MultiWOZ 2.3 and 2.4 datasets. To keep the architecture simple we do not consider cases in which slots take multiple values in this work, but we can effectively adapt our model for this setting by introducing sequential query tokens for each slot. Another limitation is that the span representation requires a computation of $O(N \cdot L_{\text{ans}})$ complexity where N and L_{ans} represent the length of context and answer span, respectively. For very long answers this might occur significant computational costs compared to exist-

ing span prediction approaches which have $O(N)$ complexity. However, this can be alleviated by adding a simple sampling and filtering step during training and prediction. We plan to further study and address these limitations in future work.

Ethics Statement

We introduced a novel model architecture for schema-guided dialogue state tracking which leverages a natural language schema and a span pointer module to achieve higher accuracy in dialogue state tracking. All experiments were conducted on publicly available datasets which are commonly used in research on dialogue systems.

References

- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.
- Guan-Lin Chao and Ian Lane. 2019. BERT-DST: Scalable End-to-End Dialogue State Tracking with Bidirectional Encoder Representations from Transformer. In *Proc. Interspeech 2019*, pages 1468–1472.
- Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, David Benjamin Beller, Lucy J Colwell, and Adrian Weller. 2021. Rethinking attention with performers. In *International Conference on Learning Representations*.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, Adarsh Kumar, Anuj Goyal, Peter Ku, and Dilek Hakkani-Tur. 2020. MultiWOZ 2.1: A consolidated multi-domain dialogue dataset with state corrections and state tracking baselines. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 422–428, Marseille, France. European Language Resources Association.
- Yue Feng, Yang Wang, and Hang Li. 2021. A sequence-to-sequence approach to dialogue state tracking. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1714–1725, Online. Association for Computational Linguistics.

- Raghav Gupta, Harrison Lee, Jeffrey Zhao, Yuan Cao, Abhinav Rastogi, and Yonghui Wu. 2022. [Show, don't tell: Demonstrations outperform descriptions for schema-guided task-oriented dialogue](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4541–4549, Seattle, United States. Association for Computational Linguistics.
- Ting Han, Ximing Liu, Ryuichi Takanabu, Yixin Lian, Chongxuan Huang, Dazhen Wan, Wei Peng, and Minlie Huang. 2021. Multiwoz 2.3: A multi-domain task-oriented dialogue dataset enhanced with annotation corrections and co-reference annotation. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 206–218. Springer.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. 2020. Spanbert: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Eleftherios Kapelonis, Efthymios Georgiou, and Alexandros Potamianos. 2022. A multi-task bert model for schema-guided dialogue state tracking. *arXiv preprint arXiv:2207.00828*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Chia-Hsuan Lee, Hao Cheng, and Mari Ostendorf. 2021. [Dialogue state tracking with a language model using schema-driven prompting](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4937–4949, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Harrison Lee, Raghav Gupta, Abhinav Rastogi, Yuan Cao, Bin Zhang, and Yonghui Wu. 2022. [Sgd-x: A benchmark for robust generalization in schema-guided dialogue systems](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):10938–10946.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. [Datasets: A community library for natural language processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zhaojiang Lin, Bing Liu, Andrea Madotto, Seungwhan Moon, Paul Crook, Zhenpeng Zhou, Zhiguang Wang, Zhou Yu, Eunjoon Cho, Rajen Subba, et al. 2021. Zero-shot dialogue state tracking via cross-task transfer. *arXiv preprint arXiv:2109.04655*.
- Yue Ma, Zengfeng Zeng, Dawei Zhu, Xuan Li, Yiyang Yang, Xiaoyuan Yao, Kaijie Zhou, and Jianping Shen. 2019. An end-to-end dialogue state tracking system with machine reading comprehension and wide & deep classification. *arXiv preprint arXiv:1912.09297*.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim Rocktäschel, and Sebastian Riedel. 2021. [KILT: a benchmark for knowledge intensive language tasks](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2523–2544, Online. Association for Computational Linguistics.
- Ori Ram, Yuval Kirstain, Jonathan Berant, Amir Globerson, and Omer Levy. 2021. [Few-shot question answering by pretraining span selection](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3066–3079, Online. Association for Computational Linguistics.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8689–8696.
- Yu-Ping Ruan, Zhen-Hua Ling, Jia-Chen Gu, and Quan Liu. 2020. Fine-tuning bert for schema-guided zero-shot dialogue state tracking. *arXiv preprint arXiv:2002.00181*.
- Xin Tian, Liankai Huang, Yingzhan Lin, Siqi Bao, Huang He, Yunyi Yang, Hua Wu, Fan Wang, and Shuqi Sun. 2021. [Amendable generation for dialogue state tracking](#). In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*, pages 80–92, Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. *Advances in neural information processing systems*, 28.
- Yifan Wang, Jing Zhao, Junwei Bao, Chaoqun Duan, Youzheng Wu, and Xiaodong He. 2022. [LUNA:](#)

- [Learning slot-turn alignment for dialogue state tracking](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3319–3328, Seattle, United States. Association for Computational Linguistics.
- Jason Williams, Antoine Raux, Deepak Ramachandran, and Alan Black. 2013. [The dialog state tracking challenge](#). In *Proceedings of the SIGDIAL 2013 Conference*, pages 404–413, Metz, France. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019. [Transferable multi-domain state generator for task-oriented dialogue systems](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 808–819, Florence, Italy. Association for Computational Linguistics.
- Fanghua Ye, Jarana Manotumruksa, and Emine Yilmaz. 2022. [MultiWOZ 2.4: A multi-domain task-oriented dialogue dataset with essential annotation corrections to improve state tracking evaluation](#). In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 351–360, Edinburgh, UK. Association for Computational Linguistics.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. 2020. Big bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems*, 33:17283–17297.
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020. [MultiWOZ 2.2 : A dialogue dataset with additional annotation corrections and state tracking baselines](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 109–117, Online. Association for Computational Linguistics.
- Jianguo Zhang, Kazuma Hashimoto, Chien-Sheng Wu, Yao Wang, Philip Yu, Richard Socher, and Caiming Xiong. 2020. [Find or classify? dual strategy for slot-value predictions on multi-domain dialog state tracking](#). In *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*, pages 154–167, Barcelona, Spain (Online). Association for Computational Linguistics.
- Yang Zhang, Vahid Noroozi, Evelina Bakhturina, and Boris Ginsburg. 2021. Sgd-qa: Fast schema-guided dialogue state tracking for unseen services. *arXiv preprint arXiv:2105.08049*.
- Jeffrey Zhao, Raghav Gupta, Yuan Cao, Dian Yu, Mingqiu Wang, Harrison Lee, Abhinav Rastogi, Izhak Shafran, and Yonghui Wu. 2022. Description-driven task-oriented dialog modeling. *arXiv preprint arXiv:2201.08904*.
- Jeffrey Zhao, Mahdis Mahdieh, Ye Zhang, Yuan Cao, and Yonghui Wu. 2021. [Effective sequence-to-sequence dialogue state tracking](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7486–7493, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

A Appendix

Symbol	Definition
LAT	Linear Attention Transformer
\mathcal{I}	Input sequence
\mathcal{G}	Global inputs
\mathcal{M}	Set of masked recurring span clusters
\mathcal{R}	Set of all recurring span clusters
D^{intent}	Intent descriptions
D^{slot}	Intent descriptions
E	Joint encoding obtained from LAT
I	Intents
S	Slots
T	Shared target tokens
U	Utterances
$\mathbf{h}^{\text{[INTENT]}}$	Intent embedding
$\mathbf{h}^{\text{[SLOT]}}$	Slot embedding
$\mathbf{h}^{\text{[UTT]}}$	Utterance embedding
$\mathbf{h}_{ij}^{\text{SPAN}}$	Span embedding from position i to j
\mathbf{x}_i	Token representation at position i
θ	Model parameters

Table 6: Glossary of symbols

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
We discussed the limitations of our work in the unnumbered limitations section.
- A2. Did you discuss any potential risks of your work?
We only used publically available datasets that are commonly used in research on dialogue systems. We believe there are no significant risks associated with our work.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract and section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Discussed in section 3.1 and 3.3

- B1. Did you cite the creators of artifacts you used?
Section 3.1 and 3.3
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
We only used publically available data and adhere to the creator’s license terms. The SGD dataset is freely available under the CC-BY-SA 4.0 and the MultiWOZ dataset is freely available under the MIT license.
- 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)?
We only used publically available data and adhere to the creator’s license terms and their intended use. The SGD dataset is freely available under the CC-BY-SA 4.0 and the MultiWOZ dataset is freely available under the MIT license.
- 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?
We only used publically available data that is commonly used in dialogue systems research and which does not uniquely identify people and which does not contain any personal data or offensive content.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
We did not create artifacts. Documentation of the artifacts used is provided in section 3.1 and 3.3.
- 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.
Section 3.1

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

C Did you run computational experiments?

Section 3 and Section 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?
Section 3.3
- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 3.3
- 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?
Section 3.3
- 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.)?
Section 3.3

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- 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.?
Not applicable. Left blank.
- 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)?
Not applicable. Left blank.
- 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?
Not applicable. Left blank.
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Not applicable. Left blank.
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
Not applicable. Left blank.