



**HAL**  
open science

## **STOP: A dataset for spoken task oriented semantic parsing**

Paden Tomasello, Akshat Shrivastava, Daniel Lazar, Po-Chun Hsu, Duc Le, Adithya Sagar, Ali Elkahky, Jade Copet, Wei-Ning Hsu, Yossi Adi, et al.

► **To cite this version:**

Paden Tomasello, Akshat Shrivastava, Daniel Lazar, Po-Chun Hsu, Duc Le, et al.. STOP: A dataset for spoken task oriented semantic parsing. SLT-2022 - IEEE Spoken Language Technology Workshop, Jan 2023, Doha, Qatar. hal-03989829

**HAL Id: hal-03989829**

**<https://inria.hal.science/hal-03989829v1>**

Submitted on 15 Feb 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# STOP: A DATASET FOR SPOKEN TASK ORIENTED SEMANTIC PARSING

Paden Tomasello, Akshat Shrivastava, Daniel Lazar, Po-Chun Hsu, Duc Le, Adithya Sagar, Ali Elkahky, Jade Copet, Wei-Ning Hsu, Yossi Adi, Robin Algayres, Tu Ahn Nguyen, Emmanuel Dupoux, Luke Zettlemoyer, Abdelrahman Mohamed

Meta AI

## ABSTRACT

End-to-end spoken language understanding (SLU) predicts intent directly from audio using a single model. It promises to improve the performance of assistant systems by leveraging acoustic information lost in the intermediate textual representation and preventing cascading errors from Automatic Speech Recognition (ASR). Further, having one unified model has efficiency advantages when deploying assistant systems on-device. However, the limited number of public audio datasets with semantic parse labels hinders the research progress in this area. In this paper, we release the Spoken Task-Oriented semantic Parsing (STOP) dataset<sup>1</sup>, the largest and most complex SLU dataset publicly available. Additionally, we define low-resource splits to establish a benchmark for improving SLU when limited labeled data is available. Furthermore, in addition to the human-recorded audio, we are releasing a TTS-generated versions to benchmark the performance for low-resource and domain adaptation of end-to-end SLU systems.

**Index Terms:** spoken language understanding, assistant, domain adaptation

## 1. INTRODUCTION

Assistant systems are becoming integrated into an increasing number of our daily devices, from cell phones and smart speakers to cars and televisions. Traditional assistant systems utilize a cascaded approach: an Automatic Speech Recognition (ASR) system transcribes the audio to text, followed by a separate Natural Language Understanding (NLU) system which predicts the user’s intentions from the decoded text. Although cascade systems enable using audio-only and text-only training resources, they are prone to error propagation when the ASR systems fail to generate the correct transcript. Further, an NLU model fed with decoded text cannot leverage paralinguistic information in the input, e.g., pausing, intonation, and stress variations. Moreover, two separate models in a cascade system increase the computational requirements for

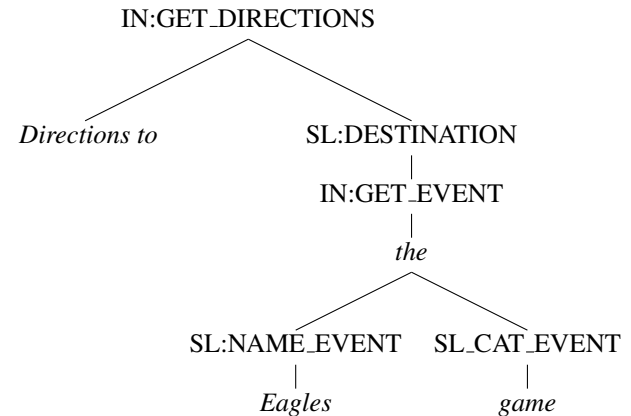
<sup>1</sup>The dataset can be downloaded at [https://github.com/facebookresearch/fairseq/tree/main/examples/audio\\_nlp/nlu](https://github.com/facebookresearch/fairseq/tree/main/examples/audio_nlp/nlu)

**Fig. 1.** Semantic Parse from TOPv2

**Utterance:** Directions to the Eagles game

**Semantic Parse:** [IN:GET DIRECTIONS Directions to [SL:DESTINATION [IN:GET\_EVENT the [SL:NAME\_EVENT Eagles ] [SL:CAT\_EVENT game ] ] ] ]

**Tree Representation:**



on-device applications, as well as duplicate efforts in building and maintaining them.

End-to-end spoken language understanding (SLU) has recently shown great promise by directly predicting intent and slots values from audio using a single system [1] [2]. It resolves the ambiguity of similarly sounding inputs that would be competing hypotheses of an ASR system. SLU has the advantage of leveraging acoustic information lost in the cascaded approach, like prosody and rhythm. One of the challenges for SLU research today is a lack of large and challenging audio datasets with semantic parse labels. In this paper, we release the Spoken Task-Oriented Semantic Parsing (STOP) dataset, which contains 1) three times as many audio samples as existing datasets 2) five times as many speakers as existing datasets 3) and compositional queries to support the ever-increasing demands of supporting real-world assis-

tant systems.

We recognize that collecting audio and intent pairs is a challenging and resource intensive task. Speech audio could contain sensitive information and proper precaution needs to be taken during collection. These privacy considerations pose a challenge to large-scale dataset collection. Accordingly, we aim to establish the first-ever low-resource splits for an end-to-end SLU dataset to support research on training with limited training data. Our baselines show the importance of audio and ASR pretraining for obtaining competitive performance versus cascaded systems.

We summarize our contributions as follows:

- Release of the STOP dataset, the largest and most semantically complex end-to-end spoken language understanding dataset publicly available,
- Defining low-resource splits as a benchmark for training with limited natural speech data samples,
- TTS-generated versions of the dataset for benchmarking low-resource training
- Experimental analysis on training end-to-end spoken language understanding models, with insights into speech pretraining, performance in the low-resource setting, domain adaptation, and how TTS audio can be incorporated into these settings.
- Open source release of our models and code for reproduction at [https://github.com/facebookresearch/fairseq/tree/main/examples/audio\\_nlp/nlu](https://github.com/facebookresearch/fairseq/tree/main/examples/audio_nlp/nlu)

We conclude with promising future directions of research.

## 2. RELATED WORK

In the 1990s, the first audio dataset containing semantic labels was first released in the Air Travel Information System (ATIS) corpus [3], followed by the Switchboard-DAMSL Labeling Project [4]. The successful application of sequence-to-sequence neural models to machine translation systems [5] and speech recognition [6] ignited the interest in end-to-end systems for SLU, which were introduced in [1] and [2]. Due to the scarcity of audio to intent data for the ever-increasing number of domains in practical assistant scenarios, transfer learning has been an area of active research explored in [7], [8] as well as ASR Pretraining in [9].

More recently, there have been several efforts around newer, larger datasets for SLU, namely, the Snips benchmark [10], Fluent Speech Command (FSC) corpus [9], and the SLURP dataset [11]. In this paper, we introduce the STOP dataset, which contains three times as many audio files and five times as many speakers compared to the largest of these datasets, SLURP. Further, our dataset includes compositional

queries with nested intents, which no previous publicly available SLU dataset included. Compositional or nested queries for text-based systems were first introduced in [12] with the first version of the TOP dataset with 2 domains. TOPv2 was release in [13], which adds 6 additional domains for a total of 137k more samples across 8 domains.

## 3. DATASETS

### 3.1. STOP Data Collection and Verification

The STOP dataset builds upon the TOPv2 dataset [13] which provides text-only inputs and target semantic parse trees. For every utterance-semantic parse pair, we collect an audio recording of the input utterance by requesting workers through Amazon’s Mechanical Turk to record themselves. We ran the recordings through an ASR system to control the quality of the recorded audio. In our case, it is the fine-tuned wav2vec 2.0 [14] ASR model in the HuggingFace [15] library. Recordings are automatically added to the dataset only if their character error rate is below 50% from the expected input text. Given that the wav2vec 2.0 model is trained on English audiobooks data, it performed poorly for non-native speakers, so any sample below the threshold was reviewed manually. In the event that the recording does not match the utterance it is discarded and recollected by another Mechanical Turk worker.

### 3.2. Normalizing text to spoken form

The TOPv2 dataset contains utterances in written text form, which may include digits or abbreviations like “4” and “CA.” However, ASR systems are typically trained on spoken text forms like “four” and “California,” respectively. To make our dataset more aligned with ASR system, we used an in-house spoken text normalization tool to enumerate possible pronunciations for each utterance and an in-house ASR model to predict the audio transcription. We then computed the edit distances across the possible normalizations and chose the normalization with the lowest distance. We include both the normalized and unnormalized form for each sample.

### 3.3. Low resource splits

Collecting audio recordings and annotating them with their semantic parses is laborious. We, therefore, believe a crucial research area for end-to-end NLU systems is low resource training, where the amount of training data is limited. Therefore, for our STOP dataset, in addition to the complete training set, we are also releasing low resource splits, which can be used as baselines for low-resource end-to-end NLU training and target datasets for domain adaptation. TOPv2 [13] defined **SPIS**, or *samples per intent and slot*, which indicates the number of training samples available for each slot and in-

tent and establishes a benchmark for training with only 25 SPIS.

### 3.4. TTS TOPv2

Going deeper in simulating real-world pre-release conditions in the STOP benchmark, we are releasing two different TTS versions of TOPv2. First, the utterances in TOPv2 are converted to spoken form as detailed in Section 3.2. We then run the VITS [16] text-to-speech (TTS) system trained on LibriTTS [17] and the CSTR VCTK Corpus [18]. In the low-resource and domain adaptation parts of Section 4, the performance of NLU models trained with these machine-generated datasets is compared to ones trained on natural human-spoken audio. Although it is a common practice in the industry to use TTS-generated speech for enabling pre-release NLU domains, STOP is the first publicly-available benchmark to offer the opportunity to academic institutions to conduct research in this area.

### 3.5. Dataset statistics and comparisons

We collected two audio recordings for every utterance in the validation and test sets of TOPv2 and one for each of the utterances in the train set. Speakers may only be in one split to ensure no data pollution exists between different partitions. Additionally, we limit the number of recordings any given speaker can submit to 3000. We ask speakers to self-identify their gender with options, Male, Female, or prefer not to say, or non-binary. Table 1 shows dataset statistics across different genders. Also, we ask speakers to self-identify whether they are English-native speakers, non-native speakers, or prefer not to say. The distribution can be seen in Table 2

**Table 1.** Gender Statistics

Female	Male	Non-binary	Unanswered
44.55%	54.62%	0.79%	0.04%

**Table 2.** Native English Statistics

Native	Non-native	Prefer not to say
95.18%	4.09%	0.73%

**Table 3.** Spoken dataset comparison

	FSC	SNIPS	SLURP	STOP
Speakers	97	67	177	<b>885</b>
Audio Files	30043	5886	72277	<b>236477</b>
Duration[hrs]	19	5.5	58	<b>218</b>

Overall, our STOP dataset is three times as large as previous datasets and contains five times as many speakers. A complete comparison with other Spoken NLU datasets is in Table 3.

## 4. EXPERIMENTS

### 4.1. Metrics

We use two metrics when evaluating our models. The first is the Exact Match (EM) accuracy, which reports the model’s accuracy in perfectly reproducing the semantic parse tree with the correct intent and slot tags. The second is the Exact Match Tree (EM-Tree), which reports the accuracy of producing the parse trees’ intent and slot labels but does not evaluate the tags’ correctness. In End-to-End NLU, EM-Tree is a good proxy if the model understands the intent and the structure of the user interaction but cannot transcribe the audio perfectly, e.g., ASR errors in rare entity names. In contrast, EM represents understanding the intent and being able to transcribe the audio perfectly.

### 4.2. Models

Our experiments contain baseline results using cascaded and end-to-end NLU systems. The cascaded system first transcribes the audio using an ASR system, which is a pre-trained wav2vec2.0 model fine-tuned on either 960h of Librispeech [19] or the full STOP training set. The text is then fed into an NLU model that is a pretrained BART-base model [20] containing a 6 layer transformer encoder and 6 layer transformer decoder, augmented with a copy-generate mechanism [21].

The end-to-end baseline system is an encoder-decoder Transformer architecture with an output vocabulary containing letters and all the intents and parses in STOP. For training, we use an Adam optimizer [22] with the betas set to 0.9 and 0.98. We use a tri-stage learning rate scheduler [23], where the learning rate is warmed up for the first 10% of updates to a maximum of  $1e-4$ , held constant for the next 40% and then linearly decayed for the remainder updates to  $5e-6$ . We trained all our models with a maximum number of updates of 320,000 and chose the best checkpoint based on the validation set.

### 4.3. In-domain vs Out-of-domain ASR

Given the surge of publicly-available ASR models trained on public data, we sought to answer how well these models perform for the assistant domain. Including out-of-domain experiments in our benchmark opens the space for research work in transfer learning for end-to-end SLU models. We used publicly-available Wav2Vec2.0[14] and HuBERT[24] pre-trained models, fine-tuned using a CTC loss[25] on Librispeech 960h and evaluated them on the STOP eval and test set in addition to the Libripeech test and dev sets. We then

took the same models and fine-tuned them using the STOP training set to perform the same evaluation.

#### 4.4. Speech pre-training for NLU

Similar to [9], we demonstrate that speech pre-training is a critical component for training end-to-end NLU systems. This work evaluates two speech pre-training methods, namely Wav2Vec 2.0 [14] and HuBERT [24].

#### 4.5. ASR pre-training

In cascaded NLU systems, the semantic parse model can benefit significantly from employing a pre-training like BART [20]. Currently audio pre-training is encoder only, leaving the decoder randomly initialized when starting the training for a downstream task in a seq2seq manner. This work also benchmarks the benefit of supervised pre-training of the decoder component. As a secondary pre-training task, we train a seq2seq ASR model that uses the pre-trained encoder. We present systems with ASR pre-training utilizing the STOP and the larger Librispeech datasets.

#### 4.6. Low resource end-to-end NLU

Since collecting audio samples with the corresponding parse labels can be impractical due to a variety of concerns, we investigate how much audio intent pairs are needed to build an end-to-end system for a single domain. To this end, we train a HuBERT encoder on the Librispeech training set, then fine-tune it on the held-in domain (either Weather or Reminder). We leverage the low resource splits in 3.3, with SPIS in {10, 25, 50, 100, 250, 500, 1000}, and analyze the error rates with progressively more data.

#### 4.7. Domain adaptation for end-to-end NLU models

One way of improving the performance of models trained in the low-resource setting is to perform source training on held in domains, then fine-tune on the low-resource target domain. Similar to the original TOPv2 paper, we study how end-to-end NLU models adapt to new domains.

We study this method by first taking the HuBERT encoder trained on Librispeech then source training on all data samples from “held-in” TOPv2 domains – all domains aside from Weather and Reminder as described in Section 3.3. We then fine-tune our model using the low-resource splits with varying amounts of data for both the Weather and the Reminder domain. We compare our results to the results of 4.6 to see how our models benefits from having previously been train on semantic parse data from different domains.

#### 4.8. Leveraging TTS for zero resource and low resource end-to-end NLU

Another way to overcome the cost of collecting large natural datasets is to supplement or replace the training data with synthetic high resource data, which is much cheaper to generate than natural speech. Accordingly, we explore the feasibility of building end-to-end NLU systems using only TTS data, or leveraging mostly TTS with a small amount of audio and intent pairs from our low-resource splits. To do so, we use the TTS versions of TOPv2 details in Section 3.4, optionally combined with the low-resource training split detailed in section 3.3. In this analysis, all data is from the held-out domain considered (Weather or Reminder).

A comparison of our these ablation studies can be seen in Table 4

Experiment	Parse Data?	TTS Data?
Low-resource Baseline	No	No
Domain Adaptation	Yes	No
TTS - Low-resource	No	Yes

**Table 4.** A comparison of low-resource and domain adaptation ablation studies for the Weather and Reminder domains.

## 5. RESULTS AND ANALYSIS

### 5.1. ASR Results for In-domain vs Out-of-domain data

Table 5 shows the results of training ASR models on Librispeech, STOP and a combined dataset. As can be seen, training purely on LibriSpeech performs quite poorly on the STOP eval and test sets, with WER above 25%. By fine-tuning on the STOP training set however, we are able to achieve the best WERs on the STOP eval and test sets of 4.29 and 4.26 respectively using a HuBERT encoder model. We hope this result demonstrates the significance of our dataset outside of end-to-end systems assistant systems, since it is the largest in-domain assistant ASR dataset as far as we know.

### 5.2. Speech pre-training for NLU

As can be seen in Table 6, speech pre-training is a critical component for end-to-end NLU. Compared to using a randomly initialized model, Wav2Vec2.0 and HuBERT have almost twice the EM accuracy for both valid and test sets. For Tree accuracy, Wav2Vec2.0 and HuBERT improve upon the randomly initialized model by roughly 25% points.

### 5.3. ASR pre-training

In Table 7, we compare the results of performing ASR pre-training across Librispeech and STOP for both HuBERT and Wav2Vec2.0 encoders. As can be seen, leveraging ASR

**Table 5.** ASR Results

Encoder	Training Set	dev_other	dev_clean	test_clean	test_other	stop_eval	stop_test
HuBERT	Librispeech	<b>8.47</b>	<b>2.99</b>	<b>3.25</b>	<b>8.06</b>	25.68	26.19
Wav2Vec2.0	Librispeech	9.02	3.31	3.42	8.81	26.53	26.93
HuBERT	STOP	46.31	31.30	31.52	47.16	<b>4.29</b>	<b>4.26</b>
Wav2Vec2.0	STOP	48.67	32.51	33.04	49.71	4.43	4.45

**Table 6.** Speech Pre-training

Encoder	Valid (EM/Tree)	Test (EM/Tree)
End-to-end		
Random	35.28 / 57.73	36.54 / 57.01
Wav2Vec2.0	67.71 / 83.35	68.05 / 82.53
HuBERT	67.86 / 83.43	68.40 / 82.85
Cascaded		
STOP ASR	72.43 / 86.58	72.36 / 85.77
Librispeech ASR	24.14 / 49.32	31.32 / 60.70
Text GT	85.65 / 88.43	85.25 / 87.85

**Table 7.** ASR Pretraining

	Valid (EM/Tree)	Test EM/Tree)
HuBERT		
None	67.86 / 83.43	68.40 / 82.85
Librispeech	68.83 / 83.45	<b>69.36</b> / 82.84
STOP	<b>69.12</b> / <b>83.89</b>	69.23 / <b>82.87</b>
wav2vec2.0		
None	67.71 / 83.35	68.05 / 82.53
Librispeech	68.24/83.55	68.47/82.49
STOP	68.52 / 83.83	68.70 / 82.78
Cascaded		
STOP ASR	72.43 / 86.58	72.36 / 85.77
LibriSpeech ASR	24.14 / 49.32	31.32 / 60.70
GT Text	85.65 / 88.43	85.25 / 87.85

pre-training improves in all scenarios by about 1%. Looking at HuBERT, the difference between Librispeech ASR pre-training and STOP ASR pretraining is small, with Librispeech performing slightly better on Test EM accuracy and STOP performing slightly better on Test Tree Accuracy. Similarly for Wav2Vec2.0, we see STOP ASR pre-training perform only slightly better than Librispeech pre-training for both Test EM accuracy and Tree Accuracy. We conclude that for decoder initialization, having some intelligent initialization obtained from ASR pre-training is adequate and does not require in-domain pre-training for decent performance gains.

#### 5.4. Low resource Baseline

As shown in Figure 2, we find that scaling up training data greatly improves training model accuracy. EM and EM-tree trend together, though in Reminder, EM benefits more from scaling up training data than EM-tree. This may be a result of the compositional nature of the Reminder domain, as opposed to Weather, which has a more flat semantic structures.

#### 5.5. Domain Adaptation

We find that source training significantly increases downstream performance on the target task, especially for the very low resource regime (SPIS < 100), and especially when Weather is the target domain. Interestingly, in the higher resource setting for Reminder, sometimes not performing source training performs better than source training. We speculate that some domains transfer well, while others inter-

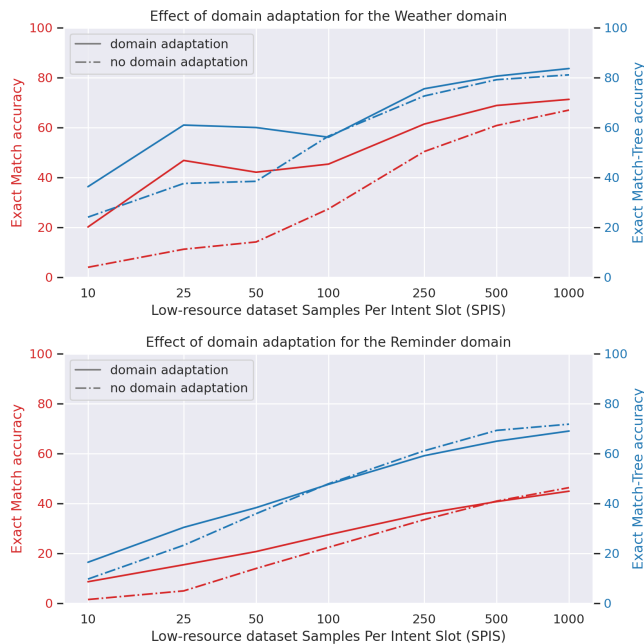
fere with downstream performance. Further study is needed to understand and mitigate this effect.

#### 5.6. TTS-Augmented Low Resource NLU

We compare three models: (1) trained on low-resource natural speech only, (2) trained on a combination of low-resource natural speech and high-resource synthetic speech data, and (3) trained only on high-resource synthetic speech. We find that even with no natural speech included, (3) performs better than (1). Adding natural speech further improves performance, and (2) always performs better than (1).

We again observe a difference between the Weather and Reminder domains: the gap between (1) and (3) is larger in Weather than Reminder, whereas adding natural speech to TTS improves Reminder more than Weather.

We also compare the performance of the two different TTS methods described in Section 3.3 – for convenience abbreviated as Librispeech and VCTK. We compare models trained on low-resource natural speech supplemented with high-resource Librispeech, VCTK, and the union of the two TTS datasets. We find that including both Librispeech and VCTK generally improves model performance beyond models trained on either of the two individually. This suggests that adding diversity to TTS data transfers to improved per-




**Fig. 2.** Low resource training in the Weather and Reminder domains, compared to low-resource training after source training on high resource held-in domains (domain adaptation). Source training generally improves downstream target performance.


formance on natural speech test samples.

## 6. CONCLUSION

In this work, we: release STOP, the largest SLU dataset for end-to-end semantic parsing, provide initial baselines using both an end-to-end as well as a cascaded system, and show the importance of in-domain ASR training for cascaded systems. Additionally, we find ASR pre-training on top of audio pretraining can lead to improvements in E2E systems, closing the gap to cascaded systems. We believe there are three particularly promising directions to explore so we may one day reach parity and eventually out-perform cascaded systems. First, is pseudolabeling or model distillation, where we can use a cascaded model to train an end-to-end model on large amounts of unlabeled audio. Second, is building more semantically meaningful representations. We hypothesize cascaded systems surpass E2E systems due to large text pre-training datasets. Third, is building a larger TTS dataset on-top of synthetic utterance-parse pairs. Our current results show impressive performance using only TTS data, and our current TTS data is built only on-top of the TOPv2 utterances, but could be expanded further easily. Finally, while this work did not explore on-device models, we believe end-to-end NLU is especially promising in compute and memory limited environments.



**Fig. 3.** Comparison of low-resource training with only low-resource natural speech, low-resource natural speech and full-resource TTS, and full-resource TTS only. TTS generally greatly improves model accuracy, and adding natural speech improves it further in almost all cases.



**Fig. 4.** Comparison of TTS methods for aiding low-resource training.

## 7. REFERENCES

- [1] D. Serdyuk, Y. Wang, C. Fuegen, A. Kumar, B. Liu, and Y. Bengio, “Towards end-to-end spoken language understanding,” *CoRR*, vol. abs/1802.08395, 2018. [Online]. Available: <http://arxiv.org/abs/1802.08395>
- [2] P. Haghani, A. Narayanan, M. Bacchiani, G. Chuang, N. Gaur, P. Moreno, R. Prabhavalkar, Z. Qu, and A. Waters, “From audio to semantics: Approaches to end-to-end spoken language understanding,” 2018.
- [3] C. T. Hemphill, J. J. Godfrey, and G. R. Doddington, “The atis spoken language systems pilot corpus,” in *Proceedings of the Workshop on Speech and Natural Language*, ser. HLT '90. USA: Association for Computational Linguistics, 1990, p. 96–101. [Online]. Available: <https://doi.org/10.3115/116580.116613>
- [4] D. Jurafsky and E. Shriberg, “Switchboard-damsl labeling project coder”’s manual.” 1997.
- [5] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to sequence learning with neural networks,” *CoRR*, vol. abs/1409.3215, 2014. [Online]. Available: <http://arxiv.org/abs/1409.3215>
- [6] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, “Listen, attend and spell,” *CoRR*, vol. abs/1508.01211, 2015. [Online]. Available: <http://arxiv.org/abs/1508.01211>
- [7] S. Schuster, S. Gupta, R. Shah, and M. Lewis, “Cross-lingual transfer learning for multilingual task oriented dialog,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019, pp. 3795–3805. [Online]. Available: <https://aclanthology.org/N19-1380>
- [8] N. Tomashenko, A. Caubrière, and Y. Estève, “Investigating Adaptation and Transfer Learning for End-to-End Spoken Language Understanding from Speech,” in *Proc. Interspeech 2019*, 2019, pp. 824–828.
- [9] L. Lugosch, M. Ravanelli, P. Ignoto, V. S. Tomar, and Y. Bengio, “Speech Model Pre-Training for End-to-End Spoken Language Understanding,” in *Proc. Interspeech 2019*, 2019, pp. 814–818.
- [10] A. Coucke, A. Saade, A. Ball, T. Bluche, A. Caulier, D. Leroy, C. Doumouro, T. Gisselbrecht, F. Caltagirone, T. Lavril, M. Primet, and J. Dureau, “Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces,” *CoRR*, vol. abs/1805.10190, 2018. [Online]. Available: <http://arxiv.org/abs/1805.10190>
- [11] E. Bastianelli, A. Vanzo, P. Swietojanski, and V. Rieser, “SLURP: A spoken language understanding resource package,” *CoRR*, vol. abs/2011.13205, 2020. [Online]. Available: <https://arxiv.org/abs/2011.13205>
- [12] S. Gupta, R. Shah, M. Mohit, A. Kumar, and M. Lewis, “Semantic parsing for task oriented dialog using hierarchical representations,” *CoRR*, vol. abs/1810.07942, 2018. [Online]. Available: <http://arxiv.org/abs/1810.07942>
- [13] X. Chen, A. Ghoshal, Y. Mehdad, L. Zettlemoyer, and S. Gupta, “Low-resource domain adaptation for compositional task-oriented semantic parsing,” *CoRR*, vol. abs/2010.03546, 2020. [Online]. Available: <https://arxiv.org/abs/2010.03546>
- [14] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *CoRR*, vol. abs/2006.11477, 2020. [Online]. Available: <https://arxiv.org/abs/2006.11477>
- [15] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, and J. Brew, “Huggingface’s transformers: State-of-the-art natural language processing,” *CoRR*, vol. abs/1910.03771, 2019. [Online]. Available: <http://arxiv.org/abs/1910.03771>
- [16] J. Kim, J. Kong, and J. Son, “Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech,” *CoRR*, vol. abs/2106.06103, 2021. [Online]. Available: <https://arxiv.org/abs/2106.06103>
- [17] H. Zen, R. Clark, R. J. Weiss, V. Dang, Y. Jia, Y. Wu, Y. Zhang, and Z. Chen, “Libritts: A corpus derived from librispeech for text-to-speech,” in *Interspeech*, 2019. [Online]. Available: <https://arxiv.org/abs/1904.02882>
- [18] J. Yamagishi, C. Veaux, and K. MacDonald, “CSTR VCTK Corpus: English multi-speaker corpus for CSTR voice cloning toolkit (version 0.92),” 2019.
- [19] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An asr corpus based on public domain audio books,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [20] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, “BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension,” in *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.



- [21] A. Aghajanyan, J. Maillard, A. Shrivastava, K. Diedrick, M. Haeger, H. Li, Y. Mehdad, V. Stoyanov, A. Kumar, M. Lewis, and S. Gupta, “Conversational Semantic Parsing,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2020.
- [22] D. P. Kingma and J. Ba, “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*, Y. Bengio and Y. LeCun, Eds., 2015. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [23] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Interspeech 2019*. ISCA, sep 2019. [Online]. Available: <https://doi.org/10.21437/Interspeech.2019-2680>
- [24] W. Hsu, B. Bolte, Y. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *CoRR*, vol. abs/2106.07447, 2021. [Online]. Available: <https://arxiv.org/abs/2106.07447>
- [25] A. Graves, S. Fernandez, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural nets,” in *ICML '06: Proceedings of the International Conference on Machine Learning*, 2006.