



HAL
open science

Is one brick enough to break the wall of spoken dialogue state tracking?

Lucas Druart, Valentin Vielzeuf, Yannick Estève

► To cite this version:

Lucas Druart, Valentin Vielzeuf, Yannick Estève. Is one brick enough to break the wall of spoken dialogue state tracking?. 2024. hal-04267804v3

HAL Id: hal-04267804

<https://hal.science/hal-04267804v3>

Preprint submitted on 12 Jun 2024

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.

Is One Brick Enough to Break the Wall of Spoken Dialogue State Tracking?

Lucas Druart^{1,2}, Valentin Vielzeuf¹, Yannick Estève²

¹Orange Innovation, France

²LIA - Avignon Université, France

lucas1.druart@orange.com, first.last@univ-avignon.fr

Abstract

In Task-Oriented Dialogue (TOD) systems, correctly updating the system’s understanding of the user’s requests (*a.k.a* dialogue state tracking) is key to a smooth interaction. Traditionally, TOD systems perform this update in three steps: transcription of the user’s utterance, semantic extraction of the key concepts, and contextualization with the previously identified concepts. Such cascade approaches suffer from cascading errors and separate optimization. End-to-End approaches have been proven helpful up to the turn-level semantic extraction step. This paper goes one step further and provides (1) a novel approach for completely neural spoken DST, (2) an in depth comparison with a state of the art cascade approach and (3) avenues towards better context propagation. Our study highlights that jointly-optimized approaches are also competitive for contextually dependant tasks, such as Dialogue State Tracking (DST), especially in audio native settings. Context propagation in DST systems could benefit from training procedures accounting for the previous’ context inherent uncertainty.

Index Terms: spoken dialogue systems, context adaptation, end-to-end, dialogue state tracking

1. Introduction

Digitization enables many tasks to be automated, nevertheless users sometimes require assistance to perform complex tasks such as making a reservation at a restaurant or booking a hotel room. Task-Oriented Dialogue (TOD) systems are designed to assist such users. A common approach to implement them is to break the problem down to three iterative steps [1]: updating the system’s understanding of the users’ requests, reasoning over a database and domain knowledge to choose the next action and providing the user an answer. This paper focuses on the understanding step.

Traditionally the user’s requests update consists of three components, respectively performing the transcription of the user’s utterance, semantic extraction and contextualization of the extracted concepts [2]. Unfortunately, this method presents the inconvenience of propagating errors of a component on to the next one(s) (*i.e.* cascading errors) and of not optimizing all components on the final objective (*i.e.* separate optimization) [3], as illustrated in Figure 1. End-to-End (E2E) approaches may address these issues by designing models in which the gradient (*i.e.* error signal) can back-propagate from the output all the way to the input [4].

On the one hand, with the advent of deep-learning and textual embeddings, state of the art Dialogue State Tracking (DST) models now work directly on automatic transcriptions [5]. However, such approaches require careful, dataset specific, mechanisms to catch and correct transcription errors together

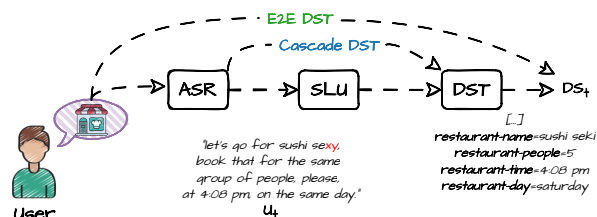


Figure 1: Spoken Dialogue State Tracking alternatives. Red characters indicate potential cascading errors.

with data augmentation to increase the downstream model’s robustness to specific upstream errors [6].

On the other hand, Spoken Language Understanding (SLU) directly from the speech signal, has successfully been applied to tasks which process utterances individually such as voice command slot filling [7] and dialogue act classification [8, 9]. Such systems often leverage transfer learning of previously trained models. This is challenging because it requires a trade-off between learning new knowledge (*e.g.* domain’s vocabulary, domain’s ontology structure) and keeping previous capabilities (*e.g.* transcription of open vocabulary concepts) on a small amount of data [3].

In TOD systems the semantic extraction also depends on the dialogue’s current context. While E2E SLU has been efficiently designed for single independent utterance processing [7, 10, 11], contextually dependant E2E SLU remains unexplored to the best of our knowledge. Indeed, dialogue history integration to guide the current turn’s prediction (*e.g.* better spelling of technical vocabulary) has already been implemented [8, 9, 12, 13]. Yet, the tasks described in these studies can be achieved without resorting to the context which is impossible for DST (*e.g.* processing cross-turn reference resolution).

Producing high quality annotated dialogue datasets is expensive because of the cognitive load required to analyse the context and adapt the annotation. Such datasets exist for chat based dialogue understanding [14, 15] but lack for spoken dialogues, explaining the gap between E2E SLU and DST. Recently, two datasets have been introduced in an attempt to fill this gap: Spoken MultiWOZ [16] and SpokenWOZ [17].

This paper lies at the intersection of both directions. We focus on contextually dependant semantic extraction, such as DST, in which the previous dialogue context is mandatory to correctly process the current one (*e.g.* cross-turn references resolution). In fact, the spoken DST models presented in this paper output a summary of the user’s requests since the beginning of the dialogue. State of the art DST systems use cascade ap-

proaches which create a textual bottleneck both in terms of data and model inference. E2E approaches do not require ground-truth transcriptions and can be jointly optimized.

This paper paves the path towards E2E spoken DST with: (1) a novel completely neural DST approach, (2) a detailed comparison with a state-of-the-art cascade approach and (3) avenues towards better context propagation.

2. Method

2.1. Task-Oriented Dialogues

In TODs users require assistance from an agent to complete a task such as making a reservation at a restaurant or booking a hotel room. More formally, let us define a TOD as a sequence of t dialogue turns $U_1, A_2, \dots, A_{t-1}, U_t$ where A_{t-1} and U_t respectively correspond to agent’s turn $t - 1$ and user’s turn t . The goal of DST is to keep up to date a condensed representation of the user’s requests. In this paper, users requests are represented as Dialogue States (DS) and correspond to a list of n slot-value pairs flattened as $\text{slot}_1=\text{value}_1; \dots; \text{slot}_n=\text{value}_n$. At a given turn t , a DST system is thus inputted the previous context DS_{t-2} and both agent’s and user’s most recent turns A_{t-1} and U_t from which it should output the updated user’s requests DS_t^1 .

2.2. Context Propagation

As the dialogue unfolds, the user might refer to previously mentioned entities. In order to design a contextually dependant SLU model, we need to propagate the context of the previous turns DS_{t-2} to inform the prediction \hat{DS}_t associated to the current user turn U_t . This paper compares two alternatives of spoken DST models detailed in the following sections, and illustrated in Figure 2. For a fairer comparison, we use the same pre-trained components for both approaches and train the model(s) with the same consumer grade 24Gb GPU².

2.2.1. Traditional Cascade Approach

The cascade approach consists of an Automatic Speech Recognition (ASR) model which transcribes the agent and user’s turns, respectively A_{t-1} and U_t , concatenates them to the previous’ turns context DS_{t-2} , and uses it as input of a Natural Language Understanding (NLU) model which predicts the updated Dialogue State \hat{DS}_t .

In our experiments, we consider two ASR models: a WavLM model [18] fine-tuned (with two additional linear layers outputting tokens’ probabilities and CTC loss) on the dataset’s transcriptions and an off-the-shelf Whisper [19] model. Regarding the NLU component, we focus on a T5 Encoder-Decoder [20] model. Note that the NLU model is trained with user turns transcriptions of the ASR model in order to be as close as possible to its inference regime.

2.2.2. Completely Neural Approach

The completely neural approach leverages the same pre-trained components. It removes the textual bottleneck by fusing the current dialogue turns with the context later on, in an embedding high dimensional space. This enriched context is then used to condition T5’s decoder generation. The goal of this approach is to enable joint optimization of all components.

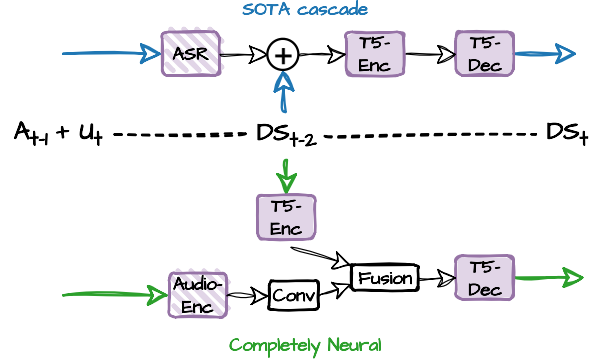


Figure 2: Two approaches for context propagation in spoken DST: SOTA cascade (top) and completely neural models (bottom). The inputs are displayed in the middle: agent previous turn A_{t-1} , user current turn U_t and previous dialogue state DS_{t-2} . The output is the current dialogue state DS_t . Hatched components are speech-related while solid ones are text-related. Colored blocks are fine-tuned while white ones are trained from scratch.

An audio encoder E_{audio} (e.g. WavLM or Whisper’s encoder) and a textual encoder E_{text} (e.g. T5’s encoder) respectively encode the current agent and user dialogue turns (audio) and the dialogue’s context in the form of the previous dialogue state DS_{t-2} (textual). Given that both models do not have the same processing windows, two convolution layers (stride 3, kernel size 9) are added to down-sample the audio encoder’s outputs. The fusion layer is a self-attention layer over the concatenation, noted \parallel , of both encoder’s outputs. The goal is to enable the model to select and mix the information from both encoders. Finally, a textual decoder (e.g. T5’s decoder) predicts \hat{DS}_t conditioned on the fusion of both encoders’ outputs. More formally, we have:

$$\begin{aligned} h_{state} &= E_{text}(DS_{t-2}) \\ h_{turns} &= \text{Conv}(E_{audio}(A_{t-1} + U_t)) \\ h &= \text{Self-Attention}(h_{state} \parallel h_{turns}) \\ \hat{DS}_t &= w_1 \dots w_n \\ \text{with } w_i &= \text{argmax}_w p(w | w_{i-1} \dots, w_1, h) \end{aligned}$$

3. Results

3.1. Datasets

MultiWOZ is a human-human chat-based English Task-Oriented Dialogue (TOD) dataset commonly used for training and evaluating dialogue systems [15]. A spoken version with vocalized user turns was published in the context of the *Speech Aware Dialogue Systems* track of the 11th edition of the Dialogue System Technology Challenge³ (DSTC11) [16]. Given that only the user turns are vocalized, the agent turns are concatenated as context with the previous dialogue state in our models.

The user utterances in the training set are available as synthetic speech, whereas the dev and test sets (**Dev|Test**) include both synthetic and human speech versions (**TTS|Human**). The dataset contains close to 10,000 dialogues with a 80/10/10 train-

¹For $t = 0$, both the context and agent turns are empty.

²Code will be made available upon publication.

³<https://dstc11.dstc.community/>

	Dev		Test	
	71.4		71.2	
	[70.3, 72.6]		[70.1, 72.4]	
	TTS	Human	TTS	Human
Cascade (g.t. text)				
dstc11 baseline [16]	38.4	31.8	n/a	
dstc11 best [5]	47.2	43.2	44.0	39.5
Cascade (WavLM)	58.2	55.0	57.2	53.5
	[57.2, 59.3]	[53.9, 56.2]	[56.0, 58.3]	[52.3, 54.7]
Cascade (Whisper)	63.7	63.6	64.4	62.3
	[62.5, 64.8]	[62.4, 64.8]	[63.3, 65.6]	[61.1, 63.5]
E2E (WavLM)	56.4	54.0	53.4	53.0
	[55.3, 57.4]	[52.9, 55.1]	[52.3, 54.5]	[51.8, 54.2]
E2E (Whisper)	59.0	56.9	58.3	56.6
	[58.0, 60.2]	[55.7, 58.0]	[57.2, 59.4]	[55.5, 57.7]

(a) Spoken MultiWOZ: ground-truth previous state DS_{t-2}

	Dev	Test
Cascade (WavLM)	82.3	63.0
	[81.3, 83.3]	[61.7, 64.3]
Cascade (Whisper)	80.7	64.2
	[79.5, 81.8]	[62.8, 65.5]
E2E (WavLM)	70.7	61.8
	[69.4, 72.0]	[60.7, 63.0]
E2E (Whisper)	81.6	80.5
	[80.4, 82.8]	[79.6, 81.3]

(c) SpokenWOZ: ground-truth previous state DS_{t-2}

	Dev		Test	
	32.0		30.3	
	[30.3, 33.7]		[28.5, 31.9]	
	TTS	Human	TTS	Human
Cascade (g.t. text)				
Cascade (WavLM)	19.5	16.2	17.6	15.3
	[18.4, 20.7]	[15.1, 17.2]	[17.2, 19.3]	[15.2, 17.4]
Cascade (Whisper)	24.0	21.9	23.1	21.3
	[22.7, 25.3]	[20.6, 23.2]	[21.9, 24.4]	[20.0, 22.6]
E2E (WavLM)	15.1	14.4	13.7	14.6
	[14.1, 16.0]	[13.4, 15.4]	[12.8, 14.6]	[13.6, 15.6]
E2E (Whisper)	19.1	17.6	18.5	16.6
	[18.0, 20.2]	[16.5, 18.7]	[17.4, 19.5]	[15.7, 17.7]

(b) Spoken MultiWOZ: predicted previous state \hat{DS}_{t-2}

	Dev	Test
Cascade (WavLM)	24.6	23.4
	[23.0, 26.3]	[22.4, 24.6]
Cascade (Whisper)	24.3	23.5
	[22.8, 25.8]	[22.5, 24.6]
E2E (WavLM)	22.2	20.3
	[20.7, 23.7]	[19.3, 21.3]
E2E (Whisper)	26.5	24.1
	[24.7, 28.5]	[23.1, 25.2]

(d) SpokenWOZ: predicted previous state \hat{DS}_{t-2}

Table 1: JGA \uparrow with bootstrapped 95% confidence intervals. **Cascade (g.t. text)** shows upper-bound performance on the ground-truth transcriptions. Note that [16] and [5] use the complete dialogue history as input to the DST model which is not possible for audio native E2E approaches [13].

dev-test split and an average of 13.3 turns per dialogue. Among the pre-defined slots, we can distinguish 3 groups: categorical slots with a closed set of values ($\sim 60\%$), non-categorical slots with an open set of values ($\sim 30\%$) and time slots ($\sim 10\%$). Note that, in order to reduce the value overlap across sets, non-categorical slots were replaced and time slots offset in the **Dev** and **Test** sets.

SpokenWOZ [17] is a human-human multi-domain spoken English TOD dataset. It extends the MultiWOZ’s set of slots with cross-turn and reasoning slots. It contains 5,700 dialogue recordings with a 4200/500/1000 train/dev/test split which corresponds to a total of 203,074 dialogue turns and 249 hours of audio. Given the native audio nature of the dataset, no ground-truth transcriptions are available. We thus consider the dataset’s provided ASR transcriptions instead of the outputs of a fine-tuned WavLM.

Given the low quantity of data at our disposal, we use 100% of the training sets for both dataset.

3.2. Evaluation

We evaluate all approaches with a turn-level exact match metric known as Joint-Goal Accuracy (JGA \uparrow) [21]. This metric requires to post-process the comma separated slot-value output format to convert it into a valid dictionary which does not take into account the order of the slot-value pairs. We present the results of all three approaches in two scenarios: with ground truth previous context DS_{t-2} , and with the previously predicted context \hat{DS}_{t-2} in Table 1. We further analyse the performance per slot group in Figure 3 and per dialogue turn in Figure 4.

Table 1 highlights that the joint optimization does indeed seem to robustify the WavLM audio encoder in the sense that it reduces the performance gap between TTS and Human for

spoken MultiWOZ. Moreover, the completely neural approach leveraging a robust audio encoder such as Whisper’s performs on par with cascade approaches and even slightly better in an audio native setting such as with SpokenWOZ.

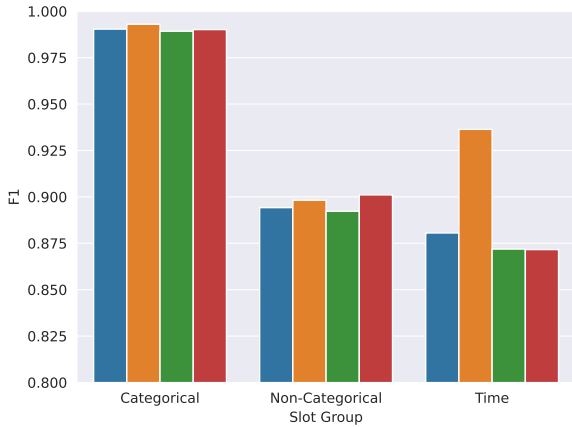
3.2.1. Slot group analysis

In order to get a more precise understanding of the differences between these approaches, we further evaluate each slot’s F1-measure. We present each slot group’s average F1-measure on the spoken MultiWOZ **Test-Human** and SpokenWOZ **Test** sets in Figure 3. Categorical slots present little difficulty while non-categorical and time slots are more challenging especially when an effort is made to have less overlap between those slots’ values such as in spoken MultiWOZ. In such a setting, Whisper’s transcription formatting seems appreciated for time slots values. In a native audio setting, such as SpokenWOZ, the end-to-end approaches perform better than cascade approaches which underlines the advantage of joint-optimization in such settings. In addition, Whisper’s encoder seems to improve time formatting capabilities suggesting that part of the formatting information might be already present in its encoder.

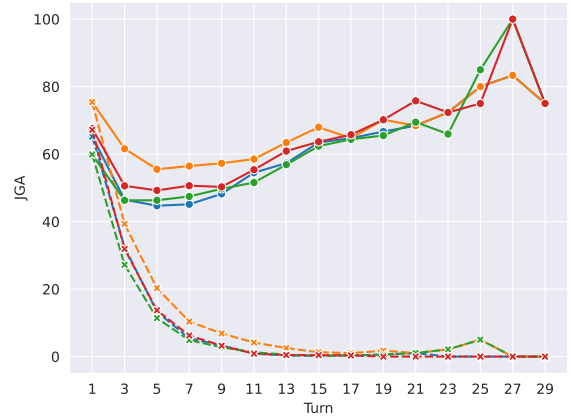
3.2.2. Dialogue turn analysis

When comparing the per-turn performance, we observe that, as illustrated in Fig. 4, the robustness of the audio encoder prevents a too fast collapse of the turn accuracy as the dialogue unfolds. However, in a more realistic scenario where we base our next prediction on the previous one \hat{DS}_{t-2} ⁴, all approaches have trouble following the course of the dialogue. This suggests

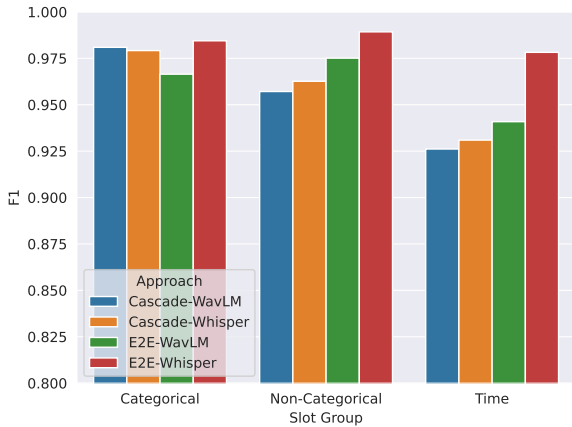
⁴Note that A_{t-1} and U_t remain unchanged which might lead to some incoherences.



(a) Spoken MultiWOZ Test-Human



(a) Spoken MultiWOZ Test-Human



(b) SpokenWOZ Test

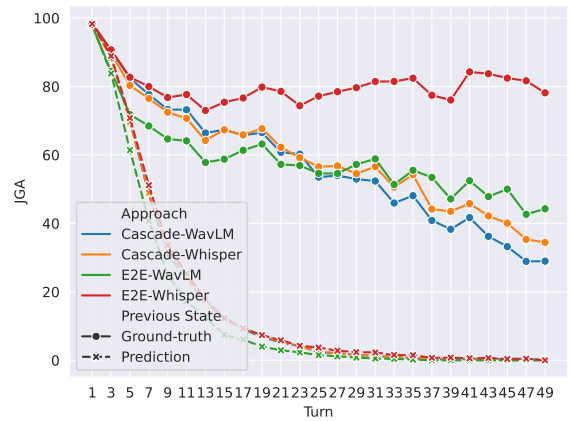
Figure 3: Slot group average F1

that additional training mechanisms such as done for other sequence prediction tasks [22] might be required to compensate this training-inference discrepancy.

When comparing datasets we also notice that the native audio setting of SpokenWOZ underlines the advantage of joint-optimization since the completely neural approaches perform better on further dialogue turns. Note that only the completely neural approach with Whisper’s encoder seems to perform roughly equally among dialogue turns. This suggests that this encoder captures better the details necessary to update the users’ requests.

4. Discussion

This paper focuses on spoken DST, for which, to the best of our knowledge, only two datasets are available. While E2E approaches often require more data to reach the same level of performance [23], this paper compares cascade and E2E approaches at the same resource level. Meaning that we use 100% of each dataset’s training set and the same pre-trained backbone models for both approaches. In addition, both datasets assume the dialogue turns to be perfectly separable. Future datasets will



(b) SpokenWOZ Test

Figure 4: Turn accuracy with and without ground-truth previous state for each approach. Note that there are fewer and fewer dialogues as the number of turns increases.

enable to study E2E approaches in higher resource and more realistic settings. Given the exact match nature of JGA, a more fine-grained evaluation to assess which errors are low-impact errors (*e.g.* rectified with the help of a database, with no impact on the dialogue trajectory) and an adapted post processing of the non categorical slot values is left as future work.

5. Conclusion

In order to pave the path towards E2E spoken DST, we analyze the differences between a state of the art cascade approach and a completely neural approach. Our study highlights that although the cascade approach remains the most accurate approach, completely neural approaches are competitive especially in audio native settings such as SpokenWOZ. However, context propagation in completely neural approaches remains an open challenge. Integrating the previous context’s uncertainty into the training process such as done for sequence prediction tasks [22] seems an interesting step in this direction.

6. References

- [1] G. Tur and R. De Mori, *SLU in Commercial and Research Spoken Dialogue Systems*. Wiley Telecom, 2011.
- [2] J. D. Williams, A. Raux, and M. Henderson, “The Dialog State Tracking Challenge Series: A Review,” *Dialogue & Discourse*, 2016.
- [3] S. Mdhaffar, V. Pelloin, A. Caubrière, G. Laperrière, S. Ghanay, B. Jabaian, N. Camelin, and Y. Estève, “Impact Analysis of the Use of Speech and Language Models Pretrained by Self-Supervision for Spoken Language Understanding,” in *Conference on Language Resources and Evaluation (LREC)*, 2022.
- [4] D. Serdyuk, Y. Wang, C. Fuegen, A. Kumar, B. Liu, and Y. Bengio, “Towards end-to-end spoken language understanding,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [5] L. Jacqmin, L. Druart, V. Vielzeuf, L. M. Rojas-Barahona, Y. Estève, and B. Favre, “OLISIA: a Cascade System for Spoken Dialogue State Tracking,” in *Proceedings of The Eleventh Dialog System Technology Challenge*. Association for Computational Linguistics, 2023.
- [6] M. Faruqi and D. Hakkani-Tür, “Revisiting the boundary between ASR and NLU in the age of conversational dialog systems,” *Computational Linguistics*, 2022. [Online]. Available: <https://aclanthology.org/2022.cl-1.8>
- [7] S. Arora, H. Futami, S.-L. Wu, J. Huynh, Y. Peng, Y. Kashiwagi, E. Tsunoo, B. Yan, and S. Watanabe, “A study on the integration of pipeline and e2e slu systems for spoken semantic parsing toward stop quality challenge,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [8] V. Sunder, S. Thomas, H.-K. J. Kuo, J. Ganhotra, B. Kingsbury, and E. Fosler-Lussier, “Towards end-to-end integration of dialog history for improved spoken language understanding,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
- [9] J. Ganhotra, S. Thomas, H.-K. J. Kuo, S. Joshi, G. Saon, Z. Tüske, and B. Kingsbury, “Integrating Dialog History into End-to-End Spoken Language Understanding Systems,” in *Interspeech*, 2021.
- [10] G. Sun, C. Zhang, and P. C. Woodland, “End-to-end spoken language understanding with tree-constrained pointer generator,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [11] V. Pelloin, N. Camelin, A. Laurent, R. De Mori, A. Caubrière, Y. Estève, and S. Meignier, “End2end acoustic to semantic transduction,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.
- [12] N. Tomashenko, C. Raymond, A. Caubrière, R. D. Mori, and Y. Estève, “Dialogue history integration into end-to-end signal-to-concept spoken language understanding systems,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [13] V. Sunder, E. Fosler-Lussier, S. Thomas, H.-K. J. Kuo, and B. Kingsbury, “ConvKT: Conversation-Level Knowledge Transfer for Context Aware End-to-End Spoken Language Understanding,” in *Interspeech*, 2023.
- [14] X. Hu, J. Dai, H. Yan, Y. Zhang, Q. Guo, X. Qiu, and Z. Zhang, “Dialogue meaning representation for task-oriented dialogue systems,” in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2022. [Online]. Available: <https://aclanthology.org/2022.findings-emnlp.17>
- [15] P. Budzianowski, T.-H. Wen, B.-H. Tseng, I. Casanueva, S. Ultes, O. Ramadan, and M. Gasic, “Multiwoz - a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling,” in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018. [Online]. Available: <https://api.semanticscholar.org/CorpusID:52897360>
- [16] H. Soltau, I. Shafran, M. Wang, A. Rastogi, W. Han, and Y. Cao, “DSTC-11: Speech aware task-oriented dialog modeling track,” in *Proceedings of The Eleventh Dialog System Technology Challenge*. Association for Computational Linguistics, 2023.
- [17] S. Si, W. Ma, H. Gao, Y. Wu, T.-E. Lin, Y. Dai, H. Li, R. Yan, F. Huang, and Y. Li, “SpokenWOZ: A Large-Scale Speech-Text Benchmark for Spoken Task-Oriented Dialogue Agents,” in *NeurIPS Datasets and Benchmarks Track*, 2023. [Online]. Available: <https://openreview.net/forum?id=viktK3nO5b>
- [18] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, M. Zeng, and F. Wei, “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing (JSTSP)*, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:239885872>
- [19] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” *ArXiv*, vol. abs/2212.04356, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:252923993>
- [20] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, “Exploring the limits of transfer learning with a unified text-to-text transformer,” *Journal Machine Learning Research (JMLR)*, 2020.
- [21] V. Zhong, C. Xiong, and R. Socher, “Global-locally self-attentive encoder for dialogue state tracking,” in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 2018. [Online]. Available: <https://aclanthology.org/P18-1135>
- [22] S. Bengio, O. Vinyals, N. Jaitly, and N. Shazeer, “Scheduled sampling for sequence prediction with recurrent neural networks,” *Advances in neural information processing systems*, vol. 28, 2015.
- [23] 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.