



Prompting Whisper for QA-driven Zero-shot End-to-end Spoken Language Understanding

Mohan Li, Simon Keizer, Rama Doddipatla

Cambridge Research Laboratory, Toshiba Europe Ltd, Cambridge, UK

{mohan.li, simon.keizer, rama.doddipatla}@toshiba.eu

Abstract

Zero-shot spoken language understanding (SLU) enables systems to comprehend user utterances in new domains without prior exposure to training data. Recent studies often rely on large language models (LLMs), leading to excessive footprints and complexity. This paper proposes the use of Whisper, a standalone speech processing model, for zero-shot end-to-end (E2E) SLU. To handle unseen semantic labels, SLU tasks are integrated into a question-answering (QA) framework, which prompts the Whisper decoder for semantics deduction. The system is efficiently trained with prefix-tuning, optimising a minimal set of parameters rather than the entire Whisper model. We show that the proposed system achieves a 40.7% absolute gain for slot filling (SLU-F1) on SLURP compared to a recently introduced zero-shot benchmark. Furthermore, it performs comparably to a Whisper-GPT-2 modular system under both in-corpora and cross-corpora evaluation settings, but with a relative 34.8% reduction in model parameters.

Index Terms: zero-shot learning, spoken language understanding, Whisper, question-answering, prefix-tuning

1. Introduction

Spoken Language Understanding (SLU) is a fundamental technology in conversational AI, dedicated to converting spoken utterances into semantic elements such as user intents, entities and emotions [1]. The training of SLU models involves expensive data collection, requiring crowdsourced speech recording and expert annotation. This has led to increased research targeting SLU within zero-shot learning scenarios, which aims to develop a generalised system capable of predicting unseen semantic labels without additional training on annotated data.

Previous studies primarily focused on text-based zero-shot natural language understanding (NLU) [2, 3, 4, 5], which processes transcripts produced by an automatic speech recognition (ASR) model to create a modular solution to zero-shot SLU. Among these studies, the prompt-based question-answering (QA) framework [6, 7, 8] has gained popularity, driven by the recent advancements in generative large language models (LLMs) [9, 10, 11]. This approach involves crafting a descriptive question for each semantic label (e.g. “*What is the name of the game?*” for the slot type `game_name`) and prompting an LLM to produce answers (slot values) that fulfill the SLU task. Unlike traditional classification or sequence-tagging methods [12, 13] that require predefined output classes, the QA framework offers flexibility in accommodating unseen semantics, relying solely on generating relevant questions for the new labels.

Over the years, SLU has evolved from modular systems to more efficient end-to-end (E2E) models [14, 15, 16]. However, research on zero-shot E2E SLU remains limited. A recent study

introduces a zero-shot audio-to-intent classification framework [17], identifying unseen intents by comparing the speech embedding similarity between the test utterance and enrollment samples. Furthermore, [18] presents a knowledge-aware audio-grounded (KA2G) generative framework for zero-shot slot filling. Similar to modular systems, KA2G incorporates ASR and NLU components to implement the QA-style slot value generation. In this model, these components are aligned in terms of decoder states to facilitate E2E training and inference.

In tandem with the rise of LLMs, advanced speech-to-text (STT) models are emerging, represented by OpenAI’s Whisper [19]. Developed using a substantial volume of audio-transcript pairs from the internet, Whisper demonstrates outstanding performance in ASR, speech translation (ST), and has been extended to other SLU tasks [20, 21, 22] under supervised learning settings. Particularly, these implementations do not depend on external language models. In this paper, we aim to unveil Whisper’s capabilities in simultaneously addressing two zero-shot SLU challenges: intent classification and slot filling. Our contributions are summarised as follows:

- We investigate the application of the Whisper model for zero-shot E2E SLU, and show that it significantly outperforms the existing baseline on the SLURP benchmark [23].
- To enable zero-shot functionalities, we reformulate the SLU tasks as QA problems and accordingly adapt the Whisper model through decoder-prompting and prefix-tuning.
- The proposed system achieves competitive performance compared to a Whisper-GPT-2 pipeline structure, while maintaining a substantially smaller footprint.
- We conduct cross-corpora zero-shot evaluations to validate the generalisation of the proposed system, trained on SLURP and tested on FSC [14] and SmartLight [24] datasets.

2. Proposed Zero-shot Whisper SLU

In this section, we provide an overview of the Whisper model and introduce its integration into the prompt-based QA framework. We further elaborate the use of prefix-tuning on Whisper for multitasking ASR and zero-shot E2E SLU. A novel semantic question generation method is also described.

2.1. Whisper model

Whisper is an open-sourced speech processing model developed by OpenAI, trained on 68000 hours of multilingual and multi-task supervised data sourced from the web. This diverse dataset contributes to enhanced robustness, enabling Whisper to effectively handle variations in accents, background noise, and technical language [19]. Whisper follows a simple E2E architecture that builds upon an encoder-decoder Transformer. The en-




Figure 1: The workflow of the proposed ZS-Whisper-SLU system. The SLU tasks are performed in three stages: ASR (left), intent classification and intent-restricted slot filling (right). The Whisper model is frozen and adapted with prefix-tuning.

coder processes the log-Mel spectrogram of a 30-second audio chunk as input. Depending on the resulting speech embeddings and user-provided special tokens, the decoder executes tasks including speech transcription, language identification, and \mathbf{X} -to-English translation. The decoding prompt may even incorporate transcription history and user instructions to influence the model output [25]. With its versatility, the Whisper decoder is recognised as a robust audio-conditioned language model, making it a suitable choice for zero-shot SLU.

2.2. Prompt-based QA framework for SLU

Given a spoken utterance \mathbf{X} , our goal is to: (1) identify its intent from a set of candidates \mathbf{I} and (2) find all the entities in \mathbf{X} corresponding to a set of slot types \mathbf{S} . In conventional supervised learning methods, these tasks are usually treated as utterance or word-level classification problems [26]. However, this approach lacks flexibility in accommodating new elements introduced to \mathbf{I} and \mathbf{S} without further training the classifier, a scenario typically known as zero-shot learning.

In this work, we propose ZS-Whisper-SLU, a generative question-answering (QA) framework based on Whisper, to conduct zero-shot intent classification and slot filling (see Figure 1). Descriptive questions are created for each element in \mathbf{I} and \mathbf{S} to construct question sets \mathbf{Q}_I and \mathbf{Q}_S . The tasks are executed in three stages as follows:

1. ASR: We adhere to Whisper’s standard transcription procedure and reserve the ASR decoder states in a cache. These states encompass both acoustic features from the encoder and textual dependencies captured at the decoder side.
2. Intent classification: The 1-best ASR transcript from stage 1 is duplicated and combined with all questions in \mathbf{Q}_I . The resulting sequences (intent prompts) are organised into a mini-batch and directly input to the Whisper decoder to generate binary answers, “Yes” or “No”, in parallel. Both answers, in plain texts, are tokenised using Whisper’s own vocabulary, so that no additional verbaliser (binary classifier) is needed. As the intent questions in these text prompts do not have paired speech embeddings, the cross-attention module at each de-

coder layer is skipped [27]. Apart from the ASR transcript in the prompt, the answer generation also relies on the cached ASR states for the purpose of E2E modelling.

3. Intent-restricted slot filling: The decoding process is similar to intent classification. However, instead of feeding all questions in \mathbf{Q}_S to the Whisper decoder, we select those relevant to the intent identified in stage 2. This not only reduces computation costs but also eliminates distracting slot types (e.g. the slot type *artist* may not be pertinent to the intent *takeaway_order*). The slot answers are open-ended and correspond to the entities mentioned in \mathbf{X} . If an entity does not exist for a certain slot type, the system is supposed to produce an *[EOT]* (end of transcript) symbol alone.

Given the potential size of \mathbf{Q}_I and \mathbf{Q}_S , we employ contrastive learning on each utterance for efficient training. This involves creating a mini-batch by combining positive questions representing the ground-truth intent and slots, along with N randomly-sampled negative questions for both semantics [18]. The system is optimised using cross-entropy loss, with prepared answers serving as labels. During inference, the entire \mathbf{Q}_I set is prompted to Whisper, and the question with the highest likelihood of “Yes” answer is picked as the predicted intent. Non-empty answers to the chosen slot questions are collected as entity outputs. In cases where multiple slot types yield the same answer, the one with the highest probability is selected. Within such a framework, one can easily introduce new intents and slot types by adding their questions to \mathbf{Q}_I and \mathbf{Q}_S , respectively.

2.3. Prefix-tuning

Prefix-tuning [28, 29] is a widely used parameter-efficient fine-tuning (PEFT) technique for adapting deep neural networks such as LLMs and computer vision (CV) models. Recent works have expanded its application to speech-related tasks and yielded promising outcomes [30]. Despite optimising a minimal amount of additional parameters (typically less than 1%) to the existing model, prefix-tuning could achieve superior adaptation effects compared to the traditional full fine-tuning strategies.

As illustrated in Figure 1, to steer Whisper towards pro-

ducing reliable ASR transcripts and desired SLU answers, we concatenate a sequence of trainable prefix vectors with the speech/text embeddings at each encoder/decoder layer. Specifically, these vectors are injected to the key \mathbf{K} and value \mathbf{V} elements involved in Transformer’s self-attention module:

$$\mathbf{K} := \text{Concat}(\mathbf{p}^{\mathbf{K}}, \mathbf{K}), \mathbf{V} := \text{Concat}(\mathbf{p}^{\mathbf{V}}, \mathbf{V}), \quad (1)$$

where $\mathbf{p}^{\mathbf{K}}$ and $\mathbf{p}^{\mathbf{V}}$ denote the prefix vectors generated by a prefix encoder (implemented as an embedding layer), serving as the soft contextual information for the ongoing encoding/decoding process. Note that the prefix vectors are shared across ASR and SLU tasks at each decoder layer, facilitating the exchange of shared knowledge such as entity names.

In our proposed system, the prefix vectors play three key roles: (1) adapting Whisper to the acoustic characteristics of the SLU corpus, (2) enhancing language modelling for task-specific domains, and addressing text normalisation issues like capitalisation and punctuation in ASR transcripts, and (3) guiding the Whisper decoder to produce intent and slot answers conditioned on the ASR decoder states. Moreover, the employment of prefix-tuning fully reserves the capabilities of the original Whisper model, ensuring flexibility for other use cases.

2.4. SLU question generation

High-quality questions are essential for accurately extracting semantic arguments within QA-based NLU/SLU frameworks. Common question design methods in prior research involve: i) adopting handcrafted templates [31] and ii) deriving questions from raw label names [8, 18]. However, these methods are prone to human biases and suffer from poor generalisation. For instance, the SLURP corpus includes an intent class labeled `convert_datetime`, where users seek to convert date or time between different time zones. It is obvious that the true intent is not explicitly conveyed in the label name. Transforming such names directly into semantic questions may produce misleading prompts, adversely impacting SLU performance.

Consequently, we leverage LLMs such as GPT-3.5 [32] to generate semantic questions, providing both the label name and a few user utterance examples to mitigate naming biases. We first instruct the LLM to produce a concise description for each semantic label, using the prompt templates presented in Figure 2. Then the output description is translated into a question of a fixed format. We observe enhanced stability in adopting such a two-step question generation process, compared to prompting the LLM for a single-shot question production.

3. Experiments

3.1. Datasets

We evaluate the proposed system on three datasets:

SLURP [23]: The official SLURP dataset includes 65 intents and 56 slots within the domain of in-home personal robot assistant. Following [18], a new data split is created for a zero-shot evaluation on 5 randomly selected slot types `{podcast_name, artist_name, audio-book_name, business_name, radio_name}`. This adapted dataset is utilised for our in-corpus evaluation.

FSC [14]: A widely used voice assistant corpus, featuring spoken commands labelled with `{action, object, location}` arguments. The argument values are drawn from a predefined set and constitute a total of 31 intents. We reorganise them into 15 intents and 2 slots (`{language, location}`,

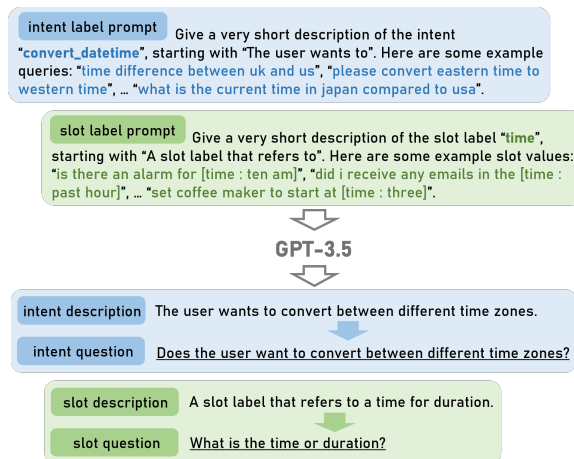


Figure 2: An example of LLM-based question generation process for intent and slot labels.

in order to accommodate the `[None]` value), and transform the results back to the original format for evaluation. This dataset is used for our cross-corpus experiments.

Smartlight (close-field) [24]: A subset of the SNIPS SLU benchmark that contains instructions for controlling smart lights. The dataset comes with 6 intents and 3 slots to cover the functions and attributes of lights. In addition to FSC, we incorporate SmartLight into the cross-corpus evaluations.

3.2. Experimental setups

We choose `Whisper-large-v2` (with 1.5 billion parameters) as the foundation ASR model for zero-shot SLU adaptation. The proposed system is compared against three baselines:

- A Modular system in [18]: Comprising a Conformer-LSTM-based ASR module and a GPT-2-based [10] NLU module.
- KA2G [18]: The E2E counterpart to the above modular structure, with the ASR and NLU components connected through a neural aligner at the decoder state level.
- A Modular system developed in-house: Consisting of a Whisper ASR module and a GPT-2 NLU module. For a fair comparison, the ASR module also utilises `Whisper-large-v2`, and the NLU module employs GPT-2-large (with 0.8 billion parameters), which shares a similar architecture with the ASR decoder. Both modules are prefixed-tuned using the same hyper-parameters as for the proposed system.

We do not compare our ZS-Whisper-SLU system with the audio-to-intent approach in [17] due to undisclosed data splits and enrollment signals not utilised by our system.

For the in-corpus experiments on SLURP, both ZS-Whisper-SLU and the GPT-2 module in the modular system are trained using $N = 10$ negative intent and slot examples for each utterance. The prefix length is set to 10 per task, resulting in 10 prefix vectors at each encoder layer (for ASR) and 30 prefix vectors at each decoder layer (for joint ASR, intent classification, and slot filling). All models undergo 10 epochs of training using the AdamW optimiser with a mini-batch size of 12. A linear learning rate scheduler, starting from 0.002 without warmup, is employed for weight decay. The SLURP-developed models are directly applied to the cross-corpus evaluations on FSC and SmartLight. However, the encoder prefix vectors are disabled to prevent transfers of acoustic knowledge. We utilise the official test set of FSC and all utterances in SmartLight, since it

Table 1: **In-corporus evaluation.** Word error rate (WER %) and slot filling SLU-F1 (%) on the zero-shot test set of SLURP. The system architecture, trainable and total number of parameters of the systems are also provided. † indicates components in the E2E system.

Model	Architecture	#Trainable Param	#Total Param	WER	SLU-F1
Modular [18]	ASR: Conformer-LSTM NLU: GPT-2-small	≈220M	≈220M	-	10.1
KA2G [18]	ASR†: Conformer-LSTM NLU†: GPT-2-small	≈220M	≈220M	18.0	9.3
Modular	ASR: Whisper-large-v2 NLU: GPT-2-large	3.3M	2.3B	8.5	43.5
ZS-Whisper-SLU (proposed)				8.3	50.0
- ASR transcript	Whisper-large-v2	3.3M	1.5B	8.4	45.8
- ASR states				8.5	36.3

Table 2: **Cross-corpus evaluation.** WER (%) and intent classification accuracy (Acc. %) on the test set of FSC. Systems with * are supervised trained with labelled data.

Model	WER	Acc.
Whisper-SLU* [21]	-	99.7
Modular	0.8	95.1
ZS-Whisper-SLU	0.8	95.0

Table 3: **Cross-corpus evaluation.** WER (%), intent classification accuracy (Acc. %), slot filling SLU-F1 (%) and perfect-parsing (PP %) on the full set of SmartLight. Systems with * are supervised trained with labelled data.

Model	WER	Acc.	SLU-F1	PP
Finstreder* [33]	6.1	-	-	88.0
Whisper-SLU* [21]	-	95.4	-	-
Modular	2.8	91.9	90.8	82.0
ZS-Whisper-SLU	2.7	91.6	90.9	82.5

lacks a data split. For fast inference, all transcripts and semantic answers are generated with greedy-search decoding.

3.3. Experimental results

Table 1 presents the in-corporus ASR and slot filling results on the zero-shot test set of SLURP. It is observed that the proposed ZS-Whisper-SLU system outperforms both the modular and E2E KA2G systems from [18] in both metrics. Thanks to the ASR power of Whisper, our system achieves a substantial 9.7% reduction in word-error-rate (WER) compared to KA2G. This further leads to an absolute 40.7% improvement in SLU-F1 score [23] on the unseen slots. ZS-Whisper-SLU also demonstrates superiority over a much stronger baseline, our own modular system that incorporates a GPT-2-large LLM for NLU, with a relative SLU-F1 gain of 14.9%. In addition to enhanced performance, ZS-Whisper-SLU utilises significantly fewer parameters (34.8% relative reduction) than the Whisper-GPT-2 modular system, highlighting the robust language understanding capabilities of the Whisper decoder and the efficacy of E2E modelling. From a training efficiency perspective, the zero-shot SLU functionalities in the proposed system are achieved by optimising a prefix encoder with only 3.3 million parameters, which constitutes just 0.2% of the original Whisper model.

The cross-corpus zero-shot evaluation results for FSC and SmartLight are displayed in Table 2 and 3, respectively. To illustrate the robustness of ZS-Whisper-SLU, we also compared it with the state-of-the-art supervised training systems from the literature on both datasets. The results indicate that the proposed system consistently achieves competitive performance with the modular structure. Notably, it performs comparably well to the supervised models, showing a slight difference of 4.7% on FSC and 5.5% (PP) on SmartLight, despite not being exposed to any in-corporus data during training. This underscores the ability of ZS-Whisper-SLU to effectively transfer knowledge from seen domains to unseen domains.

3.4. Ablation study

The ablation study results on SLURP are presented in the final row of Table 1. To assess the key factor influencing ZS-Whisper-SLU’s zero-shot performance, we selectively prune certain components during the QA decoding process:

- Removal of ASR transcript in the question prompt: In this case, decoding relies solely on the intent/slot question and cached ASR states generated during transcription.
- Exclusion of ASR states: This aligns with the modular approach, where the Whisper decoder functions as an independent NLU module, processing text-based ASR transcripts and semantic questions in the prompts.

A significant performance drop of 13.7% in SLU-F1 is observed when excluding ASR states from the decoding history in the proposed system. This operation undermines the E2E decoding property of the Whisper model, thus leading to inferior behaviour. Omitting ASR transcript from the prompt results in a moderate SLU-F1 decrease (4.2%), which suggests that including the text stabilises semantics extraction by mitigating acoustic variance embedded in the ASR states. Therefore, employing joint modelling on speech and text modalities achieves optimal zero-shot performance for ZS-Whisper-SLU.

4. Conclusions

This paper presents the use of OpenAI’s Whisper model for zero-shot E2E SLU. We address intent classification and slot filling tasks within a QA framework to handle unseen labels. The intent and slot values are derived by prompting questions to the Whisper decoder through prefix-tuning. Our experiments show that ZS-Whisper-SLU achieves competitive performance in both in-corporus and cross-corpus settings compared to a modular system, while using notably fewer model parameters.

5. References

- [1] G. Tur and R. De Mori, *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons, 2011.
- [2] C. Xia, C. Zhang, X. Yan, Y. Chang, and P. Yu, “Zero-shot user intent detection via capsule neural networks,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018, pp. 3090–3099.
- [3] A. Bapna, G. Tur, D. Hakkani-Tur, and L. Heck, “Towards zero-shot frame semantic parsing for domain scaling,” in *Proceeding of the 18th Annual Conference of the International Speech Communication Association (Interspeech)*, 2017.
- [4] D. Shah, R. Gupta, A. Fayazi, and D. Hakkani-Tur, “Robust zero-shot cross-domain slot filling with example values,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 5484–5490.
- [5] S. Lee and R. Jha, “Zero-shot adaptive transfer for conversational language understanding,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 6642–6649.
- [6] X. Du, L. He, Q. Li, D. Yu, P. Pasupat, and Y. Zhang, “QA-driven zero-shot slot filling with weak supervision pretraining,” in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, 2021, pp. 654–664.
- [7] S. Mehri and M. Eskenazi, “Gensf: Simultaneous adaptation of generative pre-trained models and slot filling,” in *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 2021, pp. 489–498.
- [8] X. Li, L. Wang, G. Dong, K. He, J. Zhao, H. Lei, J. Liu, and W. Xu, “Generative zero-shot prompt learning for cross-domain slot filling with inverse prompting,” in *Findings of the Association for Computational Linguistics: ACL 2023*, 2023, pp. 825–834.
- [9] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019, pp. 4171–4186.
- [10] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, “Language models are unsupervised multitask learners,” *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [11] 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,” *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 5485–5551, 2020.
- [12] G. Mesnil, Y. Dauphin, K. Yao, Y. Bengio, L. Deng, D. Hakkani-Tur, X. He, L. Heck, G. Tur, D. Yu *et al.*, “Using recurrent neural networks for slot filling in spoken language understanding,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 3, pp. 530–539, 2014.
- [13] S. Louvan and B. Magnini, “Recent neural methods on slot filling and intent classification for task-oriented dialogue systems: A survey,” in *Proceedings of the 28th International Conference on Computational Linguistics*, 2020, pp. 480–496.
- [14] L. Lugosch, M. Ravanelli, P. Ignoto, V. S. Tomar, and Y. Bengio, “Speech model pre-training for end-to-end spoken language understanding,” *arXiv preprint arXiv:1904.03670*, 2019.
- [15] D. Serdyuk, Y. Wang, C. Fuegen, A. Kumar, B. Liu, and Y. Bengio, “Towards end-to-end spoken language understanding,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5754–5758.
- [16] 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,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 720–726.
- [17] V. R. Elluru, D. Kulshreshtha, R. Paturi, S. Bodapati, and S. Ronanki, “Generalized zero-shot audio-to-intent classification,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.
- [18] G. Sun, C. Zhang, I. Vulić, P. Budzianowski, and P. C. Woodland, “Knowledge-aware audio-grounded generative slot filling for limited annotated data,” *arXiv preprint arXiv:2307.01764*, 2023.
- [19] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *International Conference on Machine Learning*. PMLR, 2023, pp. 28 492–28 518.
- [20] M. Wang, Y. Li, J. Guo, X. Qiao, Z. Li, H. Shang, D. Wei, S. Tao, M. Zhang, and H. Yang, “WhiSLU: End-to-end spoken language understanding with Whisper,” in *Interspeech*, 2023, pp. 770–774.
- [21] Q. Meeus, M.-F. Moens, and H. Van Hamme, “Whisper-SLU: Extending a pretrained speech-to-text transformer for low resource spoken language understanding,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–6.
- [22] D. Porjazovski, Y. Getman, T. Grósz, and M. Kurimo, “Advancing audio emotion and intent recognition with large pre-trained models and bayesian inference,” in *Proceedings of the 31st ACM International Conference on Multimedia*, 2023, pp. 9477–9481.
- [23] E. Bastianelli, A. Vanzo, P. Swietojanski, and V. Rieser, “SLURP: A spoken language understanding resource package,” in *2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020, pp. 7252–7262.
- [24] A. Saade, J. Dureau, D. Leroy, F. Caltagirone, A. Coucke, A. Ball, C. Doumouro, T. Lavril, A. Caulier, T. Bluche *et al.*, “Spoken language understanding on the edge,” in *2019 Fifth Workshop on Energy Efficient Machine Learning and Cognitive Computing-NeurIPS Edition (EMC2-NIPS)*. IEEE, 2019, pp. 57–61.
- [25] P. Peng, B. Yan, S. Watanabe, and D. Harwath, “Prompting the hidden talent of web-scale speech models for zero-shot task generalization,” *arXiv preprint arXiv:2305.11095*, 2023.
- [26] Q. Chen, Z. Zhuo, and W. Wang, “BERT for joint intent classification and slot filling,” *arXiv preprint arXiv:1902.10909*, 2019.
- [27] M. Li, C. Zorilá, C.-T. Do, and R. Doddipatla, “Towards a unified end-to-end language understanding system for speech and text inputs,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.
- [28] X. L. Li and P. Liang, “Prefix-tuning: Optimizing continuous prompts for generation,” 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)*, 2021, pp. 4582–4597.
- [29] X. Liu, K. Ji, Y. Fu, W. L. Tam, Z. Du, Z. Yang, and J. Tang, “P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks,” *arXiv preprint arXiv:2110.07602*, 2021.
- [30] K.-W. Chang, Y.-K. Wang, H. Shen, I.-t. Kang, W.-C. Tseng, S.-W. Li, and H.-y. Lee, “SpeechPrompt v2: Prompt tuning for speech classification tasks,” *arXiv preprint arXiv:2303.00733*, 2023.
- [31] D. Lamanov, P. Burnyshev, E. Artemova, V. Malykh, A. Bout, and I. Piontkovskaya, “Template-based approach to zero-shot intent recognition,” in *Proceedings of the 15th International Conference on Natural Language Generation*, 2022, pp. 15–28.
- [32] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, “Language models are few-shot learners,” *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [33] D. Bermuth, A. Poeppel, and W. Reif, “Finstredet: Simple and fast spoken language understanding with finite state transducers using modern speech-to-text models,” *arXiv preprint arXiv:2206.14589*, 2022.