

Introducing Visual Scenes and Reasoning: A More Realistic Benchmark for Spoken Language Understanding

Di Wu^{1,2*}, Liting Jiang¹, Ruiyu Fang², Bianjing¹, Hongyan Xie³, Haoxiang Su¹, Hao Huang^{1,4†},
Zhongjiang He², Shuangyong Song², Xuelong Li^{2†}

¹School of Computer Science and Technology, Xinjiang University, Urumqi, China

²Institute of Artificial Intelligence of China Telecom (TeleAI), Beijing, China

³School of Computer, Beijing University of Aeronautics and Astronautics, Beijing, China

⁴Joint International Research Laboratory of Silk Road Multilingual Cognitive Computing, Urumqi, China

Abstract

Spoken Language Understanding (SLU) consists of two sub-tasks: intent detection (ID) and slot filling (SF). Given its broad range of real-world applications, enhancing SLU for practical deployment is increasingly critical. Profile-based SLU addresses ambiguous user utterances by incorporating context awareness (CA), user profiles (UP), and knowledge graphs (KG) to support disambiguation, thereby advancing SLU research toward real-world applicability. However, existing SLU datasets still fall short in representing real-world scenarios. Specifically, (1) CA uses one-hot vectors for representation, which is overly idealized, and (2) models typically focus solely on predicting intents and slot labels, neglecting the reasoning process that could enhance performance and interpretability. To overcome these limitations, we introduce VRSLU, a novel SLU dataset that integrates both Visual images and explicit Reasoning. For over-idealized CA, we use GPT-4o and FLUX.1-dev to generate images reflecting users' environments and statuses, followed by human verification to ensure quality. For reasoning, GPT-4o is employed to generate explanations for predicted labels, which are then refined by human annotators to ensure accuracy and coherence. Additionally, we propose an instructional template, LR-Instruct, which first predicts labels and then generates corresponding reasoning. This two-step approach helps mitigate the influence of reasoning bias on label prediction. Experimental results confirm the effectiveness of incorporating visual information and highlight the promise of explicit reasoning in advancing SLU.

Datasets — <https://github.com/Tele-AI/TeleVRSLUBench>

Introduction

Spoken Language Understanding (SLU) is a fundamental component of task-oriented dialogue (TOD) systems (Ni et al. 2023). It generally comprises two sub-tasks: intent detection (ID) and slot filling (SF) (Tur and De Mori 2011). ID aims to identify the user's underlying intent, while SF focuses on extracting entities associated with that intent. Together, the predicted intent and slot labels form the basis for

* This work was completed during the internship.

† Corresponding author.

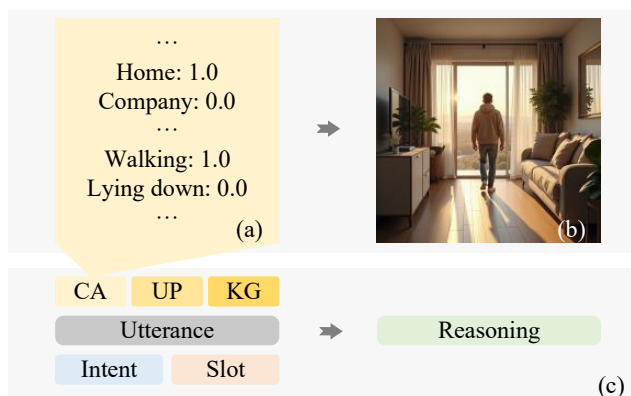


Figure 1: The construction process of scene images and reasoning in VRSLU: scene images (b) are generated based on the one-hot vectors (a) in CA, while reasoning are constructed by integrating the utterance, CA, UP, KG, intent, and slot labels (c).

generating appropriate system responses and executing subsequent actions within TOD systems. Given its wide applicability in domains such as smart speakers, virtual assistants, and smart home systems, SLU has attracted increasing research attention in recent years (Qin et al. 2021b; Weld et al. 2022; Muhammad et al. 2025).

Previous research on SLU can be broadly categorized into two stages. The first stage focuses on traditional SLU under idealized conditions, where user utterances are clear, explicit, and unambiguous. In such settings, models are able to accurately identify intent and slot information based solely on the utterance. Existing methods have already achieved promising results under these conditions (Qin et al. 2021a; Chen et al. 2022; Wu et al. 2024a). However, this scenario does not reflect the complexity of real-world applications, where user utterances are often vague or ambiguous. For example, given the utterance “Play Martial Universe by Tiancan Tudou on the smart screen”, it is difficult for the model to determine whether the user is requesting music, a video, or an audiobook. To address such challenges, Xu et al. (2022) demonstrated that traditional SLU methods

struggle with ambiguous inputs and proposed Profile-based SLU (ProSLU). This task requires the model to integrate additional sources of information to accurately infer user intent from ambiguous utterances. These sources include: Context Awareness (CA), which reflects the user’s current environment and state; User Profile (UP), which captures personal preferences and attributes; and Knowledge Graphs (KG), which provide knowledge about entities in the utterance. Although ProSLU marks an important step toward handling utterance ambiguity in real-world scenarios, significant discrepancies still remain between current research efforts and practical deployment:

- **Overly idealized representation of CA:** In ProSLU, CA is represented using one-hot vectors (e.g., “*Home: 1.0, ... Walking: 1.0, ...*”), as illustrated in Figure 1(a). However, such representations are overly simplified and fail to capture the complexity of real-world environments. In practice, CA is more realistically derived from visual inputs, where the model must infer the user’s surroundings and state from images, as shown in Figure 1(b), rather than relying on predefined categorical encodings.
- **Lack of explicit reasoning:** Traditional SLU and ProSLU systems produce only intent and slot predictions without providing the underlying reasoning process. This limitation overlooks the potential benefits of reasoning for improving both accuracy and interpretability. Explicit reasoning not only facilitates a deeper understanding of user intent but also enhances the transparency and trustworthiness of SLU systems, which is crucial for real-world deployment.

To address the two challenges discussed above and advance SLU research toward real-world applicability, we construct a new SLU dataset that incorporates both **Visual** scenes and **Reasoning**, VRSLU. To mitigate the issue of overly idealized CA representation, we first empirically demonstrate that CA significantly enhance the model’s understanding of user utterances, underscoring their importance. We further argue for the use of images instead of one-hot vectors. Then, we use GPT-4o to convert the one-hot encoded CA vectors into descriptive paragraphs, and generate corresponding images using FLUX.1-dev. All generated images undergo manual verification to ensure semantic alignment and visual quality. To address the lack of explicit reasoning, we input the utterance, CA, UP, KG, and the corresponding intent and slot labels into GPT-4o, prompting it to generate explanatory reasoning for each label, as illustrated in Figure 1(c). All reasoning outputs are manually reviewed and refined to ensure both accuracy and clarity. Furthermore, we propose LR-Instruct, an instructional method that first predicts labels and subsequently generates reasoning. This two-step approach effectively mitigates the impact of reasoning deviations on label prediction in multimodal large language models (LLMs). Experimental results demonstrate the effectiveness of our approach across multiple LLMs. Finally, we show that both visual images and explicit reasoning play critical roles in accurately understanding user requests.

The contributions of this work are as follows:

1. To extend SLU to more practical scenarios, we construct a novel dataset, VRSLU, which is the first to include both visual scene images and explicit reasoning processes.

2. We propose LR-Instruct, an instruction-based method that enhances performance across multiple MLLMs on VRSLU. Additionally, we demonstrate the critical role of visual context and the potential of explicit reasoning in accurately identifying intent and slot labels.

Related Work

Traditional SLU Dataset Progress

The Airline Travel Information Systems (ATIS) dataset has gained significant popularity in the NLP community since its release (Hemphill, Godfrey, and Doddington 1990), and has been extended to various languages (Upadhyay et al. 2018; Xu, Haider, and Mansour 2020). However, due to its single-domain focus on aviation, Tur and De Mori (2011) argues that it is difficult to learn transferable knowledge applicable to other domains from this dataset. Similarly, CAIS, the first dataset primarily targeting the music domain, suffers from comparable limitations (Liu et al. 2019). The SNIPS benchmark, collected via crowdsourcing by Snips (Coucke et al. 2018), contains utterances from voice assistants and smart home devices across seven domains. Its multi-domain nature has contributed to its widespread adoption in SLU research (Song et al. 2025). Furthermore, Saade et al. (2019) expanded the scope of SLU by introducing the Spoken Language Understanding Resource Package (SLURP) (Bastianelli et al. 2020), which covers an even broader range of domains.

ProSLU Dataset Progress

Xu et al. (2022) demonstrated that many studies based on the aforementioned datasets fail to effectively handle ambiguous utterances (Li et al. 2022a; Zhou et al. 2021). To address this issue, they proposed ProSLU, a dataset that incorporates not only user utterances as input to the model, but also integrates UP, CA, and KG information to enhance the model’s ability to accurately interpret user requests (Pham and Nguyen 2024; Teng et al. 2024). Compared with traditional SLU datasets, ProSLU is more representative of real-world SLU scenarios.

However, a significant disconnect remains between ProSLU and real-world SLU scenarios, primarily due to overly idealized CA and a lack of explicit reasoning. To address these limitations, we introduce VRSLU, a new dataset designed to overcome the shortcomings of existing SLU benchmarks. Key differences are summarized in Table 1.

Advances in LLMs for SLU

LLMs have demonstrated potential across a wide range of NLP tasks (He et al. 2024; Li et al. 2024b; Wang et al. 2024c; Li et al. 2024a; Wang et al. 2025), including SLU. Previous SLU approaches (Wu et al. 2024b; Aimaiti et al. 2024) typically formulate ID as a text classification task (Li et al. 2022b; Xiong et al. 2024), and SF as a sequence labeling task (Qu et al. 2023). In contrast, most LLM-based methods treat SLU as a generative task. Pan et al. (2023) explores the

Model	Ambiguous?	Visual?	Reasoning?
ATIS	×	×	×
CAIS	×	×	×
SNIPS	×	×	×
SLURP	×	×	×
ProSLU	✓	×	×
VRSLU	✓	✓	✓

Table 1: Differences between the VRSLU and previous SLU datasets. Here, × denotes the absence of a particular characteristic in the dataset, while ✓ indicates its presence.

zero-shot performance of ChatGPT on SLU. Subsequently, Zhu et al. (2024) proposes GPT-SLU, which models the relationship between ID and SF through prompting. Similarly, Teng et al. (2024) applied GPT to the ProSLU dataset.

Due to the lack of reasoning annotations, existing models generate only intent and slot labels, limiting the generative potential of LLMs. Moreover, applying MLLMs to SLU remains underexplored because few SLU datasets include visual information.

Visual Scene Construction

Significance of CA

Users’ requests are often influenced by their surrounding environment and current state. To further validate the importance of CA, we conduct experiments on the ProSLU dataset using both General SLU (Xu et al. 2022) and ProHAN (Teng et al. 2024) under two settings: with CA information (*w/ CA*) and without CA information (*w/o CA*). The results are presented in Figure 2.

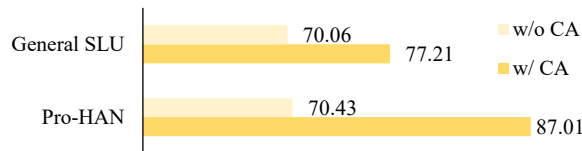


Figure 2: Overall accuracy of General SLU and ProHAN under *w/ CA* and *w/o CA* Settings.

As shown in the figure, both models exhibit a decrease in overall accuracy under the *w/o CA* setting compared to the *w/ CA* setting. This indicates that incorporating CA information enables the model to better complete ID and SF.

Necessity of Visual Representation

As shown in Figure 1(a), CA in ProSLU is represented as a one-hot vector. However, in real-world scenarios, conveying the user’s surrounding environment and state in this manner is impractical. A more realistic setting involves providing the model with an image that captures this context, from which it can infer the relevant CA information. Learning CA from images helps narrow the divide between research and practical deployment.

Visual Scene Construction Process

Due to the high costs associated with capturing real user scenarios and states, which involve diverse environments, camera angles, and recruiting users of various ages and genders, we use FLUX.1-dev to directly generate images without specifying scene configurations, camera perspectives, or user attributes. This approach enables us to obtain a wide variety of scene images while keeping costs manageable.

Initially, we attempt to generate images directly from CA one-hot vectors using FLUX.1-dev; however, the results were unsatisfactory. Consequently, we adopt an alternative approach: first, GPT-4o generates descriptive paragraphs based on the CA information, and then FLUX.1-dev produces images from these descriptions. The overall process is illustrated in Figure 3. (Note: The GPT-4o version used in this work is dated 2023-05-15, with a temperature setting of 0.7. The FLUX.1-dev hyperparameters are as follows: guidance scale is set to 3.5 and the number of inference steps is 50. The scene images are generated on a single NVIDIA A6000 GPU with 48GB of memory.)

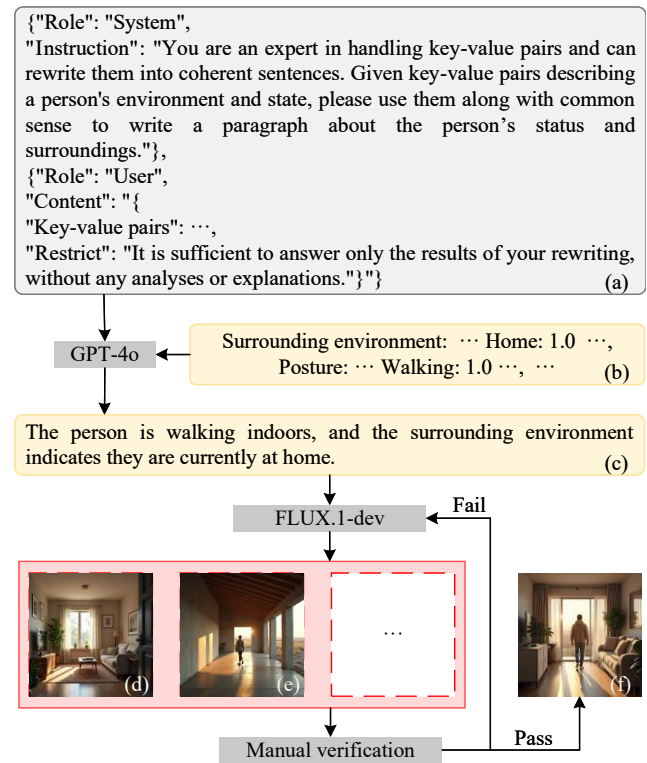


Figure 3: The process of image construction. Where, (a) shows the prompt, (b) represents CA one-hot vectors, (c) is the corresponding descriptive paragraph, (d) and (e) are invalid images, highlighted in red, and (f) is a valid image.

Due to the high cost of using GPT-4o, we initially sample 100 instances from the dataset to iteratively refine the prompt. We start by designing the initial prompt and using the API to generate descriptive paragraphs for these instances. Subsequently, the paragraphs are used to generate corresponding images with FLUX.1-dev.

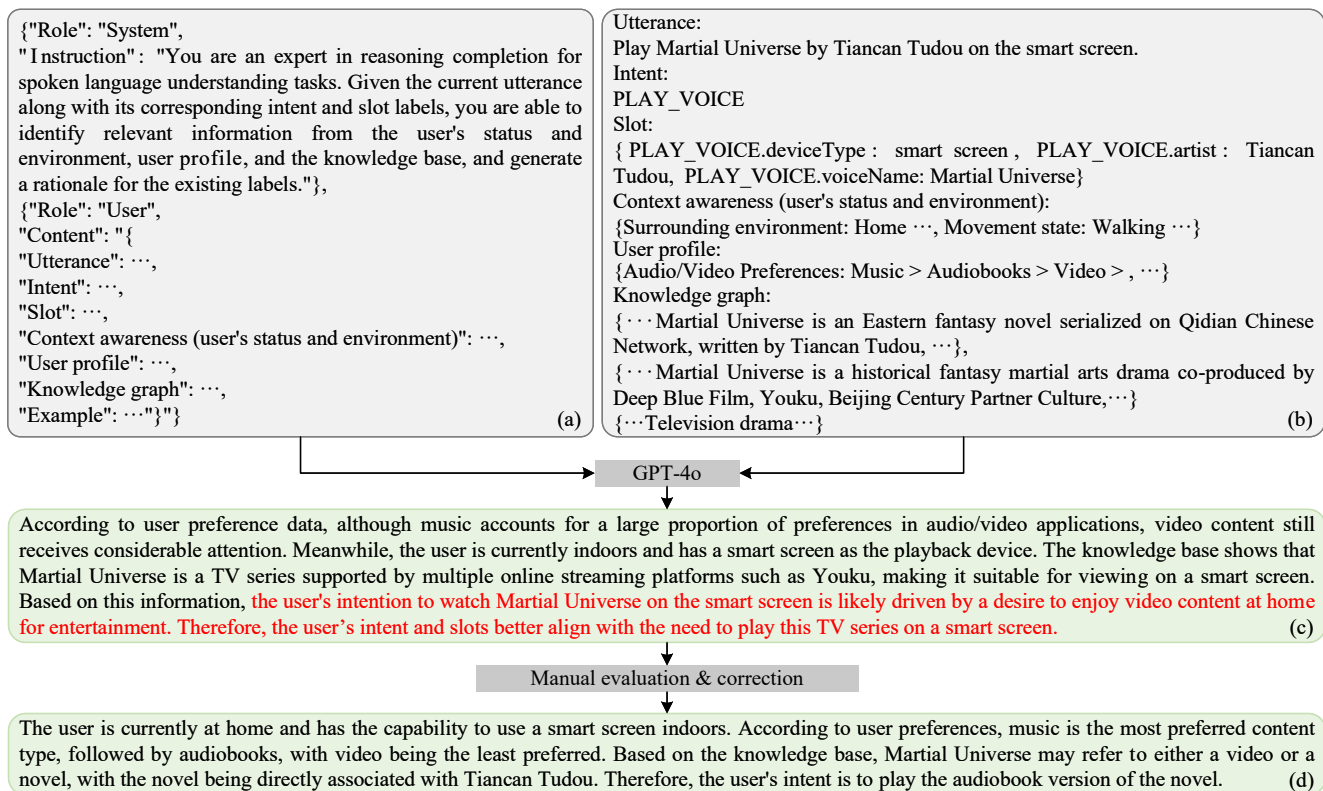


Figure 4: The process of reasoning construction. Where, (a) is the prompt, (b) is an example (information irrelevant to the user’s request has been omitted for clarity), (c) is the original reasoning generated by GPT-4o with errors highlighted in red, and (d) is the manually corrected result.

For scene image verification, we recruit two professional annotators with prior experience in image annotation to assess the quality of the 100 generated images. The annotators are asked to identify errors such as logical inconsistencies or failures to accurately reflect the user’s environment and state. After multiple rounds of prompt refinement, the final prompt used for image generation is presented in Figure 3(a). Images generated with this prompt achieve accuracies of 85.00% and 87.00% according to the two annotators, respectively. The inter-annotator agreement, measured by Gwet’s AC1, is 0.9170, indicating high reliability.

After finalizing the prompt for GPT-4o, we follow the workflow illustrated in Figure 3 to generate a corresponding visual scene for each CA entry and conduct manual verification on all generated images. An image is accepted only if both annotators agree that it meets the predefined quality criteria; otherwise, it is regenerated. As shown in Figure 3(d), there is no person in the image, although the scene is a home. In Figure 3(e), the character is depicted as being neither at home nor in a walking state. Both images fail to accurately reflect CA provided in Figure 3(b), they are considered invalid. Invalid images are regenerated using FLUX.1-dev until a valid version is produced. As illustrated in Figure 3(f), a valid image is semantically consistent with CA and visually reasonable, thereby satisfying the annotation criteria.

Reasoning Construction

Potential Benefits of Reasoning

Reasoning is the process of thinking in a logical and systematic manner, leveraging evidence and prior experience to reach conclusions or make decisions (Valmeekam et al. 2022). It enables models to provide valid justifications for their predictions, thereby improving interpretability and transparency (Huang and Chang 2023). Moreover, incorporating appropriate reasoning processes has been shown to enhance model performance across a range of tasks (Zhang et al. 2023; Wang et al. 2024a). However, existing SLU research has largely focused on intent and slot label prediction, while neglecting the generation of reasoning processes. This oversight limits both the accuracy and interpretability of models, thereby hindering the deployment of SLU systems in real-world scenarios. A major contributing factor is the lack of annotated reasoning processes in current SLU datasets, which deprives models of clear learning targets. To address this issue, we construct reasoning explanations for each sample using the following approach.

Reasoning Construction Process

The overall reasoning construction process is illustrated in Figure 4. Specifically, for each sample, we first use GPT-4o to generate an explanation of the user’s request, which is

then refined through manual evaluation and correction.

To optimize the prompting strategy, we randomly select 100 instances and iteratively refine the prompt used for reasoning generation within our computational constraints. After several rounds of refinement, we determine that the optimal input should include the utterance, intent label, slot labels, UP, KG, and CA (in textual form). We exclude previously generated images, as they tend to reduce the quality of the reasoning, likely due to GPT-4o needing to interpret the image in order to extract CA. In contrast, providing CA directly in text form allows for more accurate reasoning. To further enhance quality, we manually construct a reasoning example for a randomly selected instance and include it in the prompt. The final version of the prompt is shown in Figure 4(a).

Although we iteratively refine the prompt based on 100 sampled instances, the generated reasoning is not always coherent. As shown in Figure 4(b), the intent label of the user utterance is “PLAY.VOICE”, and the corresponding slot labels also indicate a voice-related request. While the KG provides relevant information about the novel, GPT-4o still misinterprets the user’s intent as playing a video, as illustrated in Figure 4(c). This may be due to the ambiguity of the utterance itself, as well as the presence of irrelevant information in UP, KG, and CA. Furthermore, GPT-4o occasionally hallucinates and fails to correctly interpret the intent and slot labels, resulting in incorrect reasoning in the output.

To address these issues, we invite two annotators with experience in SLU-related labeling tasks (different from those who annotated the images) and two doctoral students specializing in TOD research to manually evaluate and revise the generated reasoning. The procedure is as follows:

1. **Define Annotation Guidelines:** Annotators compare the reasoning generated by GPT-4o with the intent and slot to assess its clarity, relevance, and accuracy. If revision is needed, they must provide justifications, such as vague explanations or weak alignment with the labels. This feedback then informs the discussion and refinement process in 5.

2. **Group Assignment:** Groups are formed by drawing lots. Each group consists of one annotator with experience in SLU and one doctoral student with a background in TOD system research.

3. **First-round Review:** Each group independently reviews and revises the entire dataset. Within each group, the two members collaborate to evaluate and edit the reasoning.

4. **Cross Review:** After the initial review, the two groups exchange their annotated samples. Each group then jointly evaluates the annotations made by the other group.

5. **Result Integration:** Each sample produces two revised versions of the reasoning. If the versions are consistent, they are merged. In case of discrepancies, all four participants engage in a joint discussion to finalize the reasoning.

These steps complete the construction of the reasoning component in the dataset. A revised reasoning example is shown in Figure 4(d). Compared to the original reasoning generated by GPT-4o, the revised version emphasizes that the user’s current environment is suitable for using a smart screen, removes the video-related content, clarifies that “*Martial Universe*” refers to the novel by the author

“*Tiancan Tudou*”, and concludes that the user intends to play the audiobook version of the novel.

Method

Task Description

Following previous research on LLM-based SLU (Wu et al. 2022; Xing et al. 2024), ID and SF are typically formulated as question-answering tasks. Notably, in VRSLU, the model input includes not only the user utterance U , but also three types of knowledge: UP T_{up} and KG T_{kg} in textual form, and CA I_{ca} as visual scenes. The model generates the reasoning process R that underlies the user’s request, as well as the intent label L_I and slot labels L_S .

$$L_I, L_S, R = \text{MLLM}(U, T_{up}, T_{kg}, I_{ca}). \quad (1)$$

Instructional Template

VRSLU requires models to predict both intent and slot labels, as well as to generate the corresponding reasoning process. However, deviations in the reasoning process of LLMs can lead to incorrect label predictions (Chua et al. 2024). To mitigate this, inspired by Hase and Bansal (2021); Lyu, Apidianaki, and Callison-Burch (2024), we propose LR-Instruct, an instructional template framework that first performs label prediction followed by reasoning generation.

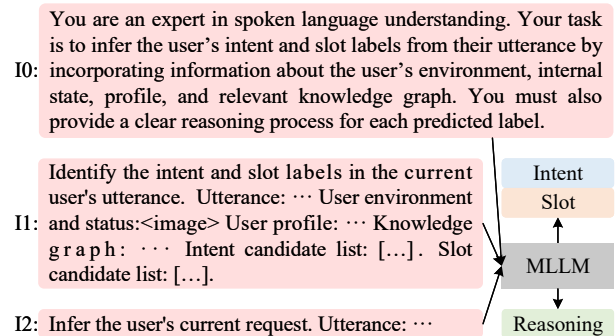


Figure 5: The proposed LR-Instruct for VRSLU. Due to space limitations, we omit the case details as well as the candidate lists for intent and slot labels.

As illustrated in Figure 5, we construct a multi-turn question-answering prompt consisting of three parts. The first part (I0) instructs the model to role-play as an SLU expert and introduces the task. The remaining two parts correspond to the model’s expected outputs. In I1, the model is required to select appropriate intent and slot labels from the candidate list based on the utterance, the provided CA image, UP, and KG. In I2, the model is asked to provide the reasoning behind the label selection. All three components (I0–I2) are used during both training and inference.

Experiments

Multimodal Large Language Model

Since VRSLU includes images, we adopt the following MLLMs as baselines in our experiments: **Yi-VL-6B** (AI

Model	Intent (Acc)	Slot (F1)	Overall (Acc)	Reasoning (Bleu-1)	Reasoning (Bleu-4)	Reasoning (Sim)
Qwen2-VL-2B	74.01	74.42	69.11	44.93	12.57	74.79
+ Ours	78.96	76.34	72.76	45.19	12.89	76.33
Qwen2.5-VL-3B	84.37	85.10	79.66	46.97	13.90	77.43
+ Ours	87.19	87.37	81.54	47.21	14.02	77.56
Yi-VL-6B	80.98	82.14	77.78	47.20	14.15	77.81
+ Ours	86.93	87.14	82.86	47.66	14.73	78.66
Qwen2-VL-7B	84.56	86.60	81.54	47.17	14.56	78.12
+ Ours	91.64	92.12	85.08	48.01	14.89	78.45
Qwen2.5-VL-7B	86.44	87.91	82.86	47.91	14.72	78.76
+ Ours	92.66	92.62	86.63	48.61	14.98	78.94

Table 2: Main experiment results.

et al. 2024). The model employs a Vision Transformer for image encoding, while the projection module facilitates the alignment between image and text. **Qwen2-VL** (Wang et al. 2024b). It is a multimodal extension of Qwen2-7B that can process visual and language inputs simultaneously. We use two parameter scales for this model: 2B and 7B. **Qwen2.5-VL** (Bai et al. 2025). It is an upgraded version of Qwen2-VL. We use two parameter scales for this mode: 3B and 7B. These MLLMs are required to perform reasoning before predicting the intent and slot labels.

Experiment Setup

For fairness, both the proposed LR-Instruct and all baselines are fine-tuned using LoRA (Hu et al. 2022). We use the AdamW optimizer (Loshchilov and Hutter 2017) with a learning rate of $5e-5$ and train each model for 3 epochs, with a per-device batch size of 1. The LoRA rank is set to 8 and the dropout rate to 0.1. To ensure reproducibility, we use a fixed random seed of 42 across all MLLMs. All experiments are conducted on eight NVIDIA A100 GPUs (40GB each).

Evaluation Metrics

Consistent with prior SLU work (Wu et al. 2025), F1-score (F1) is used to evaluate the performance of SF, accuracy (Acc) to evaluate ID performance, and overall accuracy to evaluate sentence-level semantic frame parsing performance. Since VRSLU requires reasoning in the output, we use BLEU-1 and BLEU-4 scores to evaluate reasoning quality. We also compute the cosine similarity between the generated and reference reasoning texts using Qwen3-Embedding-0.6B (Zhang et al. 2025), following the official Sentence Transformers guidelines.

Main Result

The experimental results of various MLLMs combined with the proposed LR-Instruct are presented in Table 2.

1. When combined with various MLLMs, the proposed LR-Instruct consistently improves intent Acc, slot F1, and overall Acc performance. It also achieves notable gains in reasoning metrics, including BLEU-1, BLEU-4, and cosine similarity. A key factor contributing to these improvements is that LR-Instruct requires MLLMs to predict labels prior to

generating reasoning, thereby mitigating the negative effects of reasoning errors. Moreover, the improvements achieved by our method are robust across variations in MLLM parameters and architectures, demonstrating its consistent effectiveness.

2. The quality of reasoning is positively correlated with intent and slot prediction accuracy when using the same underlying MLLM. Higher reasoning performance reflects a deeper understanding of the user’s request, which in turn enhances the model’s ability to accurately identify intent and slot labels.

Analysis

We construct VRSLU, the first SLU dataset that integrates both visual information and reasoning processes. This raises three key questions: Q1: Can the constructed scene images effectively reflect the user’s current environment and state, thereby enhancing the model’s understanding of the user’s utterances? Q2: What are the potential benefits of requiring the model to output not only intent and slot labels but also the corresponding reasoning process? Q3: What risks might arise from explicitly generating the reasoning process as part of the model’s output?

To investigate the above questions, we carry out the following analyses.

A1: As shown in Figure 2, the CA one-hot vector helps the model understand user requests. If constructed images are truly effective, removing them or replacing them with low-quality alternatives should degrade model performance. To verify this, Table 3 presents the performance of our LR-Instruct method (Qwen2.5-VL-7B) under two conditions: without the CA image (*w/o* I_{ca}) and with CA images that have not been manually verified (*w/ unverified* I_{ca}). Compared to the original setting, overall accuracy drops by 3.39% and 1.18%, respectively. These results suggest that **manually verified images more accurately represent the user’s environment and state, thereby confirming the effectiveness of constructed scene images in enhancing the model’s understanding of user utterances.**

A2: Requiring the model to generate a reasoning process offers two main advantages:

- (1) Intuitively, **requiring the model to generate a rea-**

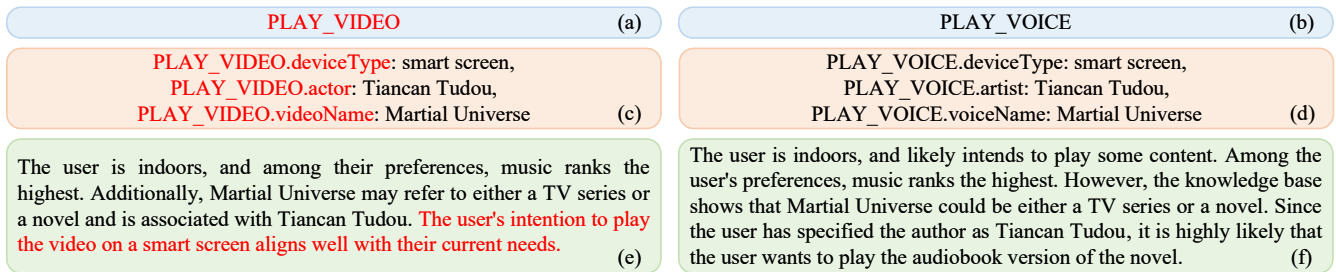


Figure 6: Predictions and reasoning from Qwen2.5-VL-7B and Qwen2.5-VL-7B+LR-Instruct. (a), (c), and (e) display the intent, slot, and reasoning of Qwen2.5-VL-7B; (b), (d), and (f) show the corresponding outputs from Qwen2.5-VL-7B+LR-Instruct.

Setting	Intent (Acc)	Slot (F1)	Overall (Acc)
Ours	92.66	92.62	86.63
w/o I_{ca}	89.08	89.78	83.24
w/ unverified I_{ca}	90.45	91.24	85.45
w/o R	90.40	91.04	85.15
w/ uncorrected R	89.45	90.01	83.86
w/o $I_{ca}&R$	87.76	87.81	80.98
w/ automated $I_{ca}&R$	89.45	89.24	83.05

Table 3: Impact of CA image and reasoning on results.

soning process enhances interpretability. As shown in Figure 6, we compare the outputs of Qwen2.5-VL-7B and Qwen2.5-VL-7B+LR-Instruct for the case in Figure 4(b). While the former correctly identifies “Tiancan Tudou” as a novel or TV series, it ultimately misinterprets the user’s intent as requesting a video. In contrast, the proposed LR-Instruct retrieves the author “Tiancan Tudou” from the knowledge base and grounds this information in the user’s utterance. This enables the model to infer that the user intends to play the audiobook version of the novel, thereby accurately predicting the intent and slot labels. Regardless of correctness, generating an explicit reasoning process could improve model transparency and facilitate interpretability.

(2) Moreover, **requiring the model to generate reasoning helps unlock its potential in ID and SF, thereby improving overall performance.** We evaluate two settings: removing the reasoning process (*w/o R*) and using uncorrected automatically generated reasoning (*w/ uncorrected R*). As shown in Table 3, *w/o R* leads to drops in intent Acc, slot F1, and overall Acc compared to the vanilla setting. This indicates that requiring reasoning helps the model better understand the utterance and make more accurate predictions. The performance decline in *w/ uncorrected R* also suggests that GPT-4o’s generated reasoning still benefits from human refinement. Additionally, we conduct experiments by removing both CA images and reasoning (*w/o $I_{ca}&R$*), as well as using unverified CA images and uncorrected reasoning (*w/ automated $I_{ca}&R$*). The results show consistent drops across all metrics compared to the vanilla setting, confirming the usefulness of both CA (image) and reasoning, and highlighting the importance of human verification.

A3: As illustrated in Figure 6, a key risk of explicitly

generating the reasoning process is that flawed reasoning can negatively impact the accuracy of label predictions. Supporting this observation, Tables 2 and 3 show that Qwen2.5-VL-7B, which generates reasoning before making predictions, performs worse than its counterpart without reasoning (*w/o R*). This suggests that if not well controlled, reasoning generation may hinder model performance. The proposed LR-Instruct mitigates this issue by first predicting labels and then generating reasoning, which enhances the model’s ability to discriminate between intent and slot categories. As shown in Figure 7, after applying LR-Instruct to the Qwen2.5-VL-7B, its recognition performance on most intent and slot categories surpasses that of the baseline, intuitively demonstrating the effectiveness of our method.

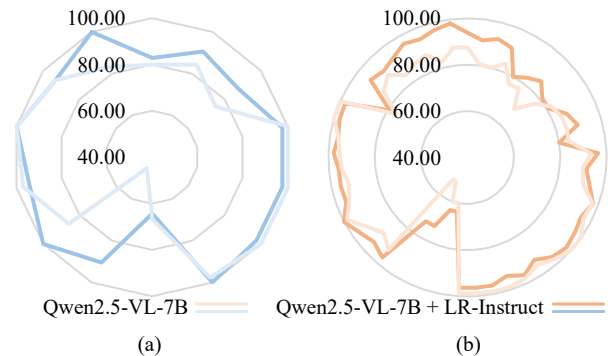


Figure 7: Comparison of results for different categories of intent (a) and slot (b).

Conclusion

In this work, we propose a novel benchmark, VRSLU, which is the first SLU dataset that incorporates user environment and statue scene images as part of the input, and requires models to generate reasoning processes. We also propose LR-Instruct, a method that first prompts the model to predict labels before generating the reasoning process, thereby effectively mitigating the influence of reasoning errors on prediction outcomes. Finally, we highlight the contributions of scene images and the potential benefits of explicit reasoning. In future work, we will explore extending VRSLU to AI networks (Shao and Li 2025) to make SLU research more closely aligned with real-world applications.

Acknowledgments

The paper was supported by the following projects: National Natural Science Foundation of China (62466055, 61663044), Research Project of the State Language Commission (ZD1145-96), and the Excellent Doctoral Student Research Innovation Project of Xinjiang University (No. XJU2022BS077).

References

- AI, .; Young, A.; Chen, B.; Li, C.; Huang, C.; Zhang, G.; Zhang, G.; Li, H.; Zhu, J.; Chen, J.; Chang, J.; Yu, K.; Liu, P.; Liu, Q.; Yue, S.; Yang, S.; Yang, S.; Yu, T.; Xie, W.; Huang, W.; Hu, X.; Ren, X.; Niu, X.; Nie, P.; Xu, Y.; Liu, Y.; Wang, Y.; Cai, Y.; Gu, Z.; Liu, Z.; and Dai, Z. 2024. Yi: Open Foundation Models by 01.AI. *arXiv:2403.04652*.
- Aimaiti, A.; Wu, D.; Jiang, L.; Abudouwaili, G.; Huang, H.; and Silamu, W. 2024. An Uyghur Extension to the MASSIVE Multi-lingual Spoken Language Understanding Corpus with Comprehensive Evaluations. In *Proc. Interspeech 2024*, 3525–3529.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-v1 technical report. *arXiv preprint arXiv:2502.13923*.
- Bastianelli, E.; Vanzo, A.; Swietojanski, P.; and Rieser, V. 2020. SLURP: A Spoken Language Understanding Resource Package. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 7252–7262.
- Chen, D.; Huang, Z.; Wu, X.; Ge, S.; and Zou, Y. 2022. Towards Joint Intent Detection and Slot Filling via Higher-order Attention. In *IJCAI*, 4072–4078.
- Chua, J.; Rees, E.; Batra, H.; Bowman, S. R.; Michael, J.; Perez, E.; and Turpin, M. 2024. Bias-augmented consistency training reduces biased reasoning in chain-of-thought. *arXiv preprint arXiv:2403.05518*.
- Coucke, A.; Saade, A.; Ball, A.; Bluche, T.; Caulier, A.; Leroy, D.; Doumouro, C.; Gisselbrecht, T.; Caltagirone, F.; Lavril, T.; et al. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *arXiv preprint arXiv:1805.10190*.
- Hase, P.; and Bansal, M. 2021. When can models learn from explanations? a formal framework for understanding the roles of explanation data. *arXiv preprint arXiv:2102.02201*.
- He, Z.; Wang, Z.; Liu, X.; Liu, S.; Yao, Y.; Huang, Y.; Li, X.; Li, Y.; Che, Z.; Zhang, Z.; et al. 2024. Telechat technical report. *arXiv preprint arXiv:2401.03804*.
- Hemphill, C. T.; Godfrey, J. J.; and Doddington, G. R. 1990. The ATIS spoken language systems pilot corpus. In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Huang, J.; and Chang, K. C.-C. 2023. Towards Reasoning in Large Language Models: A Survey. In *Findings of the Association for Computational Linguistics: ACL 2023*, 1049–1065.
- Li, C.; Zhou, Y.; Chao, G.; and Chu, D. 2022a. Understanding users’ requirements precisely: a double Bi-LSTM-CRF joint model for detecting user’s intentions and slot tags. *Neural Computing and Applications*, 34(16): 13639–13648.
- Li, Q.; Peng, H.; Li, J.; Xia, C.; Yang, R.; Sun, L.; Yu, P. S.; and He, L. 2022b. A survey on text classification: From traditional to deep learning. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 13(2): 1–41.
- Li, X.; Yao, Y.; Jiang, X.; Fang, X.; Wang, C.; Liu, X.; Wang, Z.; Zhao, Y.; Wang, X.; Huang, Y.; et al. 2024a. 52b to 1t: Lessons learned via tele-flm series. *arXiv preprint arXiv:2407.02783*.
- Li, X.; Yao, Y.; Jiang, X.; Fang, X.; Wang, C.; Liu, X.; Wang, Z.; Zhao, Y.; Wang, X.; Huang, Y.; et al. 2024b. Tele-flm technical report. *arXiv preprint arXiv:2404.16645*.
- Liu, Y.; Meng, F.; Zhang, J.; Zhou, J.; Chen, Y.; and Xu, J. 2019. CM-Net: A Novel Collaborative Memory Network for Spoken Language Understanding. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 1051–1060.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lyu, Q.; Apidianaki, M.; and Callison-Burch, C. 2024. Towards faithful model explanation in nlp: A survey. *Computational Linguistics*, 50(2): 657–723.
- Muhammad, Y. I.; Salim, N.; Zainal, A.; and Suhaili, S. M. 2025. A joint learning classification for intent detection and slot filling from classical to deep learning: a review. *Neural Computing and Applications*, 1–45.
- Ni, J.; Young, T.; Pandelea, V.; Xue, F.; and Cambria, E. 2023. Recent advances in deep learning based dialogue systems: A systematic survey. *Artificial intelligence review*, 56(4): 3055–3155.
- Pan, W.; Chen, Q.; Xu, X.; Che, W.; and Qin, L. 2023. A preliminary evaluation of chatgpt for zero-shot dialogue understanding. *arXiv preprint arXiv:2304.04256*.
- Pham, T.; and Nguyen, D. Q. 2024. Jpis: A joint model for profile-based intent detection and slot filling with slot-to-intent attention. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 10446–10450. IEEE.
- Qin, L.; Liu, T.; Che, W.; Kang, B.; Zhao, S.; and Liu, T. 2021a. A co-interactive transformer for joint slot filling and intent detection. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8193–8197. IEEE.
- Qin, L.; Xie, T.; Che, W.; and Liu, T. 2021b. A survey on spoken language understanding: Recent advances and new frontiers. *arXiv preprint arXiv:2103.03095*.
- Qu, X.; Gu, Y.; Xia, Q.; Li, Z.; Wang, Z.; and Huai, B. 2023. A survey on arabic named entity recognition: Past, recent advances, and future trends. *IEEE Transactions on Knowledge and Data Engineering*, 36(3): 943–959.

- Saade, A.; Dureau, J.; Leroy, D.; Caltagirone, F.; Coucke, A.; Ball, A.; Doumouro, C.; Lavril, T.; Caulier, A.; Bluche, T.; et al. 2019. Spoken language understanding on the edge. In *2019 Fifth Workshop on Energy Efficient Machine Learning and Cognitive Computing-NeurIPS Edition (EMC2-NIPS)*, 57–61. IEEE.
- Shao, J.; and Li, X. 2025. Ai flow at the network edge. *IEEE Network*.
- Song, M.; Liu, T.; Li, Q.; Ma, D.; Sun, M.; and Tian, L. 2025. Zero-Shot Cross-Domain Slot Filling with Retrieval Augmented In-Context Learning. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Teng, D.; Lu, C.; Xu, X.; Che, W.; and Qin, L. 2024. Prohan: A heterogeneous graph attention network for profile-based spoken language understanding. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 10726–10730. IEEE.
- Tur, G.; and De Mori, R. 2011. *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons.
- Upadhyay, S.; Faruqui, M.; Tür, G.; Dilek, H.-T.; and Heck, L. 2018. (Almost) zero-shot cross-lingual spoken language understanding. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 6034–6038. IEEE.
- Valmeekam, K.; Olmo, A.; Sreedharan, S.; and Kambhampati, S. 2022. Large language models still can't plan (a benchmark for LLMs on planning and reasoning about change). In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.
- Wang, L.; Hu, Y.; He, J.; Xu, X.; Liu, N.; Liu, H.; and Shen, H. T. 2024a. T-sciq: Teaching multimodal chain-of-thought reasoning via large language model signals for science question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 19162–19170.
- Wang, P.; Bai, S.; Tan, S.; Wang, S.; Fan, Z.; Bai, J.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; et al. 2024b. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Wang, Z.; Liu, X.; Yao, Y.; Wang, C.; Zhao, Y.; Yang, Z.; Deng, W.; Jia, K.; Peng, J.; Huang, Y.; et al. 2025. Technical Report of TeleChat2, TeleChat2. 5 and T1. *arXiv preprint arXiv:2507.18013*.
- Wang, Z.; Yao, Y.; Mengxiang, L.; He, Z.; Wang, C.; Song, S.; et al. 2024c. TeleChat: An Open-source Bilingual Large Language Model. In *Proceedings of the 10th SIGHAN Workshop on Chinese Language Processing (SIGHAN-10)*, 10–20.
- Weld, H.; Huang, X.; Long, S.; Poon, J.; and Han, S. C. 2022. A survey of joint intent detection and slot filling models in natural language understanding. *ACM Computing Surveys*, 55(8): 1–38.
- Wu, D.; Jiang, L.; Mao, B.; Xie, H.; Su, H.; He, Z.; Fang, R.; Song, S.; Huang, H.; and Li, X. 2025. INT: Establishing Information Transfer for Multilingual Intent Detection and Slot Filling. In *Findings of the Association for Computational Linguistics: ACL 2025*, 15120–15142.
- Wu, D.; Jiang, L.; Yin, L.; Li, Z.; and Huang, H. 2024a. Ceagnet: a co-interactive external attention network for joint intent detection and slot filling. *Neural Computing and Applications*, 36(22): 13513–13525.
- Wu, D.; Jiang, L.; Yin, L.; Wang, K.; Su, H.; Li, Z.; and Huang, H. 2024b. Dual level intent-slot interaction for improved multi-intent spoken language understanding. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 12301–12305. IEEE.
- Wu, Y.; Wang, H.; Zhang, D.; Chen, G.; and Zhang, H. 2022. Incorporating instructional prompts into a unified generative framework for joint multiple intent detection and slot filling. In *Proceedings of the 29th International Conference on Computational Linguistics*, 7203–7208.
- Xing, B.; Liao, L.; Huang, M.; and Tsang, I. 2024. Dc-instruct: An effective framework for generative multi-intent spoken language understanding. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 14520–14534.
- Xiong, S.; Zhao, Y.; Zhang, J.; Mengxiang, L.; He, Z.; Li, X.; and Song, S. 2024. Dual prompt tuning based contrastive learning for hierarchical text classification. In *Findings of the association for computational linguistics ACL 2024*, 12146–12158.
- Xu, W.; Haider, B.; and Mansour, S. 2020. End-to-end slot alignment and recognition for cross-lingual NLU. *arXiv preprint arXiv:2004.14353*.
- Xu, X.; Qin, L.; Chen, K.; Wu, G.; Li, L.; and Che, W. 2022. Text is no more enough! a benchmark for profile-based spoken language understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 11575–11585.
- Zhang, Y.; Li, M.; Long, D.; Zhang, X.; Lin, H.; Yang, B.; Xie, P.; Yang, A.; Liu, D.; Lin, J.; et al. 2025. Qwen3 Embedding: Advancing Text Embedding and Reranking Through Foundation Models. *arXiv preprint arXiv:2506.05176*.
- Zhang, Z.; Zhang, A.; Li, M.; Zhao, H.; Karypis, G.; and Smola, A. 2023. Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*.
- Zhou, P.; Huang, Z.; Liu, F.; and Zou, Y. 2021. PIN: A novel parallel interactive network for spoken language understanding. In *2020 25th International Conference on Pattern Recognition (ICPR)*, 2950–2957. IEEE.
- Zhu, Z.; Cheng, X.; An, H.; Wang, Z.; Chen, D.; and Huang, Z. 2024. Zero-shot spoken language understanding via large language models: A preliminary study. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 17877–17883.