

Multimodal Human-Autonomous Agents Interaction Using Pre-Trained Language and Visual Foundation Models

Linus Nwankwo

linus.nwankwo@unileoben.ac.at
Chair of Cyber-Physical Systems, Montanuniversität
Leoben, Austria

Elmar Rueckert

Chair of Cyber-Physical Systems, Montanuniversität
Leoben, Austria

ABSTRACT

In this paper, we extended the method proposed in [17] to enable humans to interact naturally with autonomous agents through vocal and textual conversations. Our extended method exploits the inherent capabilities of pre-trained large language models (LLMs), multimodal visual language models (VLMs), and speech recognition (SR) models to decode the high-level natural language conversations and semantic understanding of the robot’s task environment, and abstract them to the robot’s actionable commands or queries. We performed a quantitative evaluation of our framework’s natural vocal conversation understanding with participants from different racial backgrounds and English language accents. The participants interacted with the robot using both spoken and textual instructional commands. Based on the logged interaction data, our framework achieved 87.55% vocal commands decoding accuracy, 86.27% commands execution success, and an average latency of 0.89 seconds from receiving the participants’ vocal chat commands to initiating the robot’s actual physical action. The video demonstrations of this paper can be found at <https://linusnep.github.io/MTCC-IRoNL/>.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI); Interaction paradigms; Natural language;**

KEYWORDS

Human-robot interaction, LLMs, VLMs, ROS, foundation models, natural language interaction, vocal conversation.

ACM Reference Format:

Linus Nwankwo and Elmar Rueckert. 2024. Multimodal Human-Autonomous Agents Interaction Using Pre-Trained Language and Visual Foundation Models. In *Workshop of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24 Workshop)*, March 11–14, 2024, Boulder, CO, USA. ACM, New York, NY, USA, 3 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Existing approaches for interacting with autonomous robots in the real world have been dominated by complex teleoperation controllers [16] and rigid command protocols [18], where the robots

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HRI '24 Workshop, March 11–14, 2024, Boulder, CO, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/24/06

<https://doi.org/XXXXXXXX.XXXXXXX>

execute predefined tasks based on specialized programming languages. As the challenges we present to these robots become more intricate and the environments they operate in grow more unpredictable [21], there arises an unmistakable need for more natural and intuitive interaction mechanisms.

The last few years have witnessed tremendous advancement in generative AI and natural language processing (NLP) [13], [19]. These advancements driven primarily by foundation models, specifically transformer-based large language models (LLMs) like OpenAI GPT-3 [7], Google BERT [1], Meta AI LLaMA [9], and multimodal visual language models (VLMs) e.g., CLIP [5], DALL-E [12], and their successors, has opened new avenues for human-robot interaction (HRI) [23]. The inherent abilities of these models to understand language patterns, and structure and generate human-like responses as well as visual observations have led to several interesting applications, such as [3], [6] and [2].

In this paper, we exploit the inherent natural language capabilities of the pre-trained foundation models, as well as a speech recognition (SR) model to enable humans to interact naturally with autonomous agents through both spoken and textual dialogues. As demonstrated in the video at the project website¹, our framework aims to realize a new approach to human-robot interactions—one where the vocal or textual conversation is the command (refer to Section 3 for more details). Therefore, our contributions are twofold:

- we introduced a dual-modality framework that can leverage independent pre-trained LLMs, VLMs, and SR models to enable humans to interact with real-world autonomous robots or other entities through spoken or textual conversations.
- we performed real-world experiments with our developed framework to ensure that the robot’s actions are always aligned with the user’s spoken or textual instructions.

2 RELATED WORK

Prior works such as [14], [10], and [15] have explored the incorporation of vocal instructions into robotic systems. However, while these works are exceptional, they have relied primarily on a direct speech-to-action (STA) approach, where the robot’s actions are dependent upon the accurate transcription of the vocal commands by the SR model employed in the respective works. In most noisy real-world scenarios, their approach may introduce stochastic behaviour in the robot’s actions due to vulnerability to acoustic distortions present in real-world environments.

Further, L. Nwankwo et al. [17] incorporated text-based interaction techniques into autonomous robots leveraging LLMs and VLMs. Nonetheless, the framework lacks complete naturalness due to the absence of a mechanism to understand vocal instructions. In

¹<https://linusnep.github.io/MTCC-IRoNL/>

this work, we build upon the foundation provided in [17]. Instead of relying on the accuracy of the SR model to plan the robot’s actions or depend on the text-based approach as a standalone, we propose a dual-modality approach that synergizes both the textual and vocal modalities. We leveraged the LLMs and the SR models’ abilities to maintain robustness in diverse environments. In environments where ambient noise levels may compromise the accuracy of the SR model’s vocal instructions decoding, our framework provides the flexibility to revert to the text-based interaction method. Conversely, in a quieter environment, the user can leverage the vocal modality pipeline for more natural and seamless interaction.

With our proposed dual-modality approach, we aim to provide the user with the autonomy to select the mode of interaction most suited to the prevailing conditions. Specifically, our framework mitigates the risk of misinterpretations and erroneous robotic actions that may have arisen due to sole dependence on the STA method, thereby ensuring consistent and reliable HRI [23] in the real world.

3 METHOD

Figure 1 shows the architectural overview of our proposed framework. The proposed model contains five main components: the

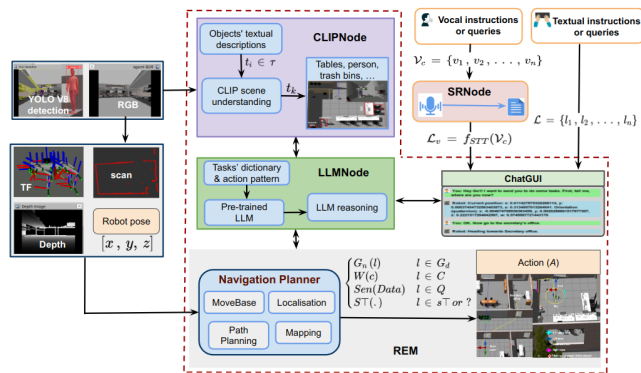


Figure 1: Overview of our framework’s architecture. The area enclosed with the red dotted line decodes the textual-based natural language conversations and visual understanding. In the SRNode, we employed Google’s SR model [8] to decode the vocal conversation from humans and abstract them to the textual representations required by the ChatGUI to interact with the LLMNode.

LLMNode to decode the high-level textual-based conversations from humans, the CLIPNode to provide a visual and semantic understanding of the robot’s task environment, the REM node to abstract the high-level understanding from the LLMNode to actual robot actions, the ChatGUI to serve as the user’s primary textual-based interaction point, and the SRNode to provide vocal or auditory commands understanding.

In this section, we provide details on the incorporation of the vocal conversation understanding pipeline. For details about the implementation of the first four sections (the area enclosed in red dotted line in Figure 1) and how the pre-trained LLMs and VLMs are prompted to generate the actions used by the REM node, we refer the reader to [17].

3.1 Vocal Conversation Decoding

In order to decode the vocal natural language conversation and abstract them to the robot’s actions, we developed the SRNode. The SRNode employs Google’s SR model [8] to capture the high-level auditory input from a microphone device, transcribing the auditory inputs to textual representations. The textual representation is subsequently used by the ChatGUI to establish communication between the LLMNode and the rest of the interfaces of Figure 1 within ROS ecosystem [11]. Formally, given a vocal command \mathcal{V}_c e.g., task descriptions, queries captured by the microphone device, we developed a function f_{STT} employing the Google’s SR model [8] such that $\mathcal{V}_c = \{v_1, v_2, \dots, v_n\}$ is transcribed to a textual natural language representation \mathcal{L} as depicted in Eq. 1. The elements $v_i, i \in n$ represent distinct vocal commands or instructions given to the SR model through speech, e.g., “Hello robot, can you move forward?”, “What is your current location?”, “Navigate to the kitchen area”, etc.

$$\mathcal{L}_v = f_{STT}(\mathcal{V}_c) = \{l_1, l_2, \dots, l_n\}, \quad l_i \in \mathcal{L} \quad (1)$$

where l_i denotes the transcribed natural language command from the set \mathcal{L}_v . We sent the resulting output from Eq.1 as input to the ChatGUI. We then used the LLMNode to handle the incoming natural language inputs from the ChatGUI by first passing them through the pre-trained LLM [4]. The resultant output from the LLMNode is then mapped to the robot’s actionable commands or information request by the robot’s execution mechanism (REM) node, consisting of the ROS[11] navigation planner packages shown at the lower bottom-left of Figure 1.

4 EXPERIMENTS

We conducted both real-world and simulated experiments to validate the performance of our framework. In simulation, we utilised the Unitree Go1 ROS & Gazebo packages² and a ROS-based open-source mobile robot adapted from [20]. We ran all the simulations with a ground station PC with Nvidia Geforce RTX 3060 Ti GPU, 8GB memory running Ubuntu 20.04, ROS Noetic distribution.

In the real-world experiments, we used a Lenovo ThinkBook Intel Core i7 with Intel iRIS Graphics running Ubuntu 20.04, ROS Noetic distribution. Segway RMP Lite 220 mobile robot was used. The robot is equipped with an RGB-D camera and a 2D RPLiDAR for both visual and spatial observations of the task environment. We used our PC’s inbuilt microphone and a plug-and-play AmazonBasics Pro Gaming Headset with a microphone function in all the experiments. We performed all the real-world experiments in our laboratory office (11 rooms) and outside corridor environment, measuring approximately 18×20 m and 6×120 m respectively.

We experimented with OpenAI GPT-2 [4], Google BERT [1], and Meta AI LLaMA [9]. OpenAI GPT-3 [7] and GPT-4 [22] are also adaptable to our framework. However, due to their API access limitations, we mostly utilised the open-access and free versions of the LLMs (GPT-2 [4] specifically) in our experiments.

4.1 Preliminary Results

We invited 5 participants (average age of 27 (± 3) and gender distribution, 80% male and 20% female) with different English language accents to interact with the robot via natural vocal conversation.

²https://github.com/unitreerobotics/unitree_guide

We logged the interaction data i.e., the SR models' transcription of the participants' spoken words, the LLMNode predicted labels, the true action labels, etc. We used the logged interaction data to quantitatively evaluate the performance of our framework. We defined vocal commands understanding accuracy (VCUA) metric in addition to the navigation success rate (NSR), object identification accuracy (OIA), and average response time (ART) metrics utilised in [17] to assess our framework's performance. With the VCUA, we assess how accurately the LLMNode predicts the commands based on the transcribed vocal instructions from the SRNode. We computed the accuracy as the percentage proportion of the correctly transcribed instruction to the generated instructions (from LLMNode) fed to the REM node for the actual robot's execution.

Figure 2 presents the statistical results that we obtained from the interaction data analysis. The top-left figure shows the VCUA, NSR, and OIA metrics.

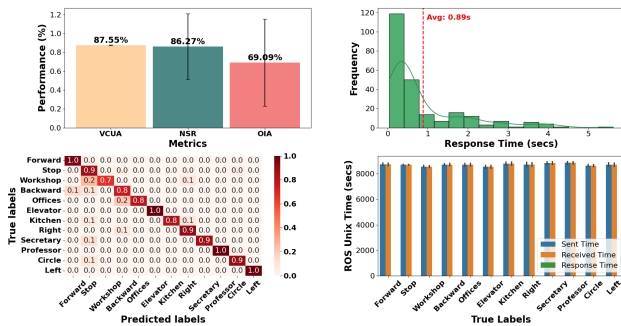


Figure 2: Quantitative evaluation results illustrating VCUA, NSR, OIA, and ART based on the logged interaction data.

and OIA metrics for selected labels. We achieved 87.55% VCUA and 86.27% NSR, which indicates a good level of accuracy in the vocal commands decoding. In comparison to the results obtained in the textual-based method [17], we observed a slight difference in the command recognition accuracy, with VCUA achieving about 12% less than the textual-based CRA (99.13%) as well as a 11.69% reduction in the NSR. This is expected because of the ambient environmental noise and variation in the participant's accents, which affect the vocal transcription from the SRNode, as could be seen in the confusion matrix (bottom-left of Figure 2).

Further, the ART (right column of Figure 2) across all the selected commands is approximately 0.89 seconds. This indicates that, on average, the robot takes less than a second from receiving a vocal chat command to initiating the robot's actual physical action, which suggests a relatively quick response time for our framework.

5 CONCLUSION AND FUTURE WORK

We introduced a framework that leverages the inherent capabilities of LLMs, VLMs, and SR models to enhance human-robot interaction through natural vocal and textual conversations. Our evaluation from logged human interaction data achieved high vocal command understanding accuracy and effective task execution. This shows that our framework can enhance the intuitiveness and naturalness of human-robot interaction in the real world. In our future work, we aim to refine our framework to resist the impact of environmental noise. We intend to incorporate adaptive noise-cancellation

algorithms and context-aware speech recognition techniques to mitigate the impacts of random noise.

REFERENCES

- [1] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *North American Chapter of the Association for Computational Linguistics*. <https://api.semanticscholar.org/CorpusID:52967399>
- [2] Anthony Brohan et al. 2023. RT-1: Robotics Transformer for Real-World Control at Scale. arXiv:2212.06817 [cs.RO]
- [3] Anthony Brohan et al. 2023. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control. arXiv:2307.15818 [cs.RO]
- [4] Alec Radford et al. 2019. Language Models are Unsupervised Multitask Learners. <https://api.semanticscholar.org/CorpusID:160025533>
- [5] Alec Radford et al. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *International Conference on Machine Learning*. <https://api.semanticscholar.org/CorpusID:231591445>
- [6] Brohan Anthony et al. 2023. Do as i can, not as i say: Grounding language in robotic affordances. In *Proceedings of The 6th Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 205)*, Karen Liu, Dana Kulic, and Jeff Ichnowski (Eds.). PMLR, 287–318. <https://proceedings.mlr.press/v205/ichter23a.html>
- [7] Brown Tom et al. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 1877–1901.
- [8] Chiu Chung-Cheng et al. 2018. State-of-the-Art Speech Recognition with Sequence-to-Sequence Models. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (Calgary, AB, Canada). IEEE Press, 4774–4778. <https://doi.org/10.1109/ICASSP.2018.8462105>
- [9] Hugo Touvron et al. 2023. LLaMA: Open and Efficient Foundation Language Models. arXiv:2302.13971 [cs.CL]
- [10] Matthew Marge et al. 2022. Spoken language interaction with robots: Recommendations for future research. *Computer Speech & Language* 71 (2022), 101255. <https://doi.org/10.1016/j.csl.2021.101255>
- [11] Quigley Morgan et al. 2009. ROS: an open-source Robot Operating System. *ICRA Workshop on Open Source Software* 3.
- [12] Ramesh Aditya et al. 2021. Zero-Shot Text-to-Image Generation. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 8821–8831. <https://proceedings.mlr.press/v139/ramesh21a.html>
- [13] Staphord Bengesi et al. 2023. Advancements in Generative AI: A Comprehensive Review of GANs, GPT, Autoencoders, Diffusion Model, and Transformers. arXiv:2311.10242 [cs.LG]
- [14] Davide Ferrari, Filippo Alberi, and Cristian Secchi. 2023. Facilitating Human-Robot Collaboration through Natural Vocal Conversations. arXiv:2311.13973 [cs.RO]
- [15] Krishna Kodur, Manizheh Zand, and Maria Kyrarini. 2023. Towards Robot Learning from Spoken Language. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction* (Stockholm, Sweden) (HRI '23). Association for Computing Machinery, New York, NY, USA, 112–116. <https://doi.org/10.1145/3568294.3580053>
- [16] Tsung-Chi Lin, Achyuthan Unni Krishnan, and Zhi Li. 2023. Perception-Motion Coupling in Active Telepresence: Human Behavior and Teleoperation Interface Design. *J. Hum.-Robot Interact.* 12, 3, Article 31 (mar 2023), 24 pages. <https://doi.org/10.1145/3571599>
- [17] Nwankwo Linus and Rueckert Elmar. 2024. The Conversation is the Command: Interacting with Real-World Autonomous Robot Through Natural Language. arXiv:2401.11838 [cs.RO]
- [18] Debasmita Mukherjee, Kashish Gupta, and Homayoun Najjaran. 2022. A Critical Analysis of Industrial Human-Robot Communication and Its Quest for Naturalness Through the Lens of Complexity Theory. *Frontiers in Robotics and AI* 9 (2022). <https://doi.org/10.3389/frobt.2022.870477>
- [19] San Murugesan and Aswani Kumar Cherukuri. 2023. The Rise of Generative Artificial Intelligence and Its Impact on Education: The Promises and Perils. *Computer* 56, 5 (2023), 116–121. <https://doi.org/10.1109/MC.2023.3253292>
- [20] Linus Nwankwo, Clemens Fritze, Konrad Bartsch, and Elmar Ruckert. 2023. ROMR: A ROS-based open-source mobile robot. *Hardware X* 14 (2023), e00426. <https://doi.org/10.1016/j.ohx.2023.e00426>
- [21] Linus Nwankwo and Elmar Ruckert. 2023. Understanding Why SLAM Algorithms Fail in Modern Indoor Environments. In *Advances in Service and Industrial Robotics*, Tadej Petrič, Aleš Ude, and Leon Žlajpah (Eds.). Springer Nature Switzerland, Cham, 186–194. https://doi.org/10.1007/978-3-031-32606-6_22
- [22] OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774 [cs.CL]
- [23] Hang Su, Wen Qi, Jiahao Chen, Chenguang Yang, Juan Sandoval, and Med Amine Laribi. 2023. Recent advancements in multimodal human-robot interaction. *Frontiers in Neurobotics* 17 (2023). <https://doi.org/10.3389/fnbot.2023.1084000>