

Towards Multimodal Social Conversations with Robots: Using Vision-Language Models

Ruben Janssens and Tony Belpaeme

Abstract—Large language models have given social robots the ability to autonomously engage in open-domain conversations. However, they are still missing a fundamental social skill: making use of the multiple modalities that carry social interactions. While previous work has focused on task-oriented interactions that require referencing the environment or specific phenomena in social interactions such as dialogue breakdowns, we outline the overall needs of a multimodal system for social conversations with robots. We then argue that vision-language models are able to process this wide range of visual information in a sufficiently general manner for autonomous social robots. We describe how to adapt them to this setting, which technical challenges remain, and briefly discuss evaluation practices.

I. INTRODUCTION

Spoken dialogue is arguably the most advanced bidirectional communication method we have as humans. It enables us to exchange information quickly, but more importantly, it is the cornerstone of *social interactions*. Social interactions serve a higher purpose than merely exchanging factual information: they enable the building of bonds between social actors, and this is largely driven by *social dialogue* [1].

It is important that social robots can engage in social dialogue. Social robots are used in socially assistive scenarios, such as being a companion, where they should be able to build a bond with their user, or language tutoring, where they should be able to prepare their learner for human-human dialogue. However, social robots are currently not yet able to fully engage in social dialogue. They are still missing fundamental social skills.

After all, what makes spoken dialogue so powerful and efficient, is that it uses multiple modalities. Not only the words we speak convey information, but we also use non-verbal communication, transmitted through audible and visible signals. These signals convey information about the cognitive and affective state of the conversation partner, about the identity of the conversation partner, and about the environment—all of which allows us, crucially, to *adapt* to the conversation partner.

While large language models (LLMs) have allowed social robots to finally engage autonomously in verbal interactions that are virtually unconstrained [2], they are not yet able to make use of these multiple modalities. Vision-language models (VLMs) promise to jointly process visual input with text in dialogue. However, they are generally not designed for social conversations with an embodied agent that is




Fig. 1: Overview of important information conveyed through visual and auditory channels during social interactions.

situated in the physical world. Instead, they are designed to be assistants, which answer questions based on additional visual context, as in an image sent through a chat message.

Previous work has already explored the use of multimodal information in human-robot conversations—it is hardly a new challenge [3]. However, previous work either did not yet use language models, which inherently limited its conversational abilities [4], [5]; limited its representation of visual information to discrete categories e.g. emotions or specific visual cues [6]; or it did not focus on social conversations, but rather tackled task-oriented conversations where visual grounding referred to objects that are used in a task [7]; or focused on the start of a social conversation [8], [9].

In this position paper, we propose to explore the capabilities of VLMs to engage in *situated social conversations* for robots. We will outline what capabilities are required for this system and propose an architecture for such a system, highlighting technical challenges along the way. Finally, we discuss how such a system can be evaluated. We note that while audible signals are a very important part of non-verbal communication, we are currently not yet aware of audio description models that can represent these in a sufficiently general way, and therefore focus on visual input.

II. SYSTEM REQUIREMENTS AND CHALLENGES

We outline three categories of information that a social robot should be able to process from non-verbal signals in social dialogue, shown in Figure 1. First, it should be able to detect the *identity of the conversation partner*, possibly recognising them, but also adapting the conversation to them (such as how one will speak differently to a child than to an older adult), or making references to their behaviour or appearance where appropriate (e.g. recognising that the user is wearing a T-shirt with a particular band name, or

This research received funding from the Flemish Government (AI Research Program 2).

Authors are affiliated with IDLab-AIRO, Ghent University–imec, Ghent, Belgium. Contact: ruben.janssens@ugent.be

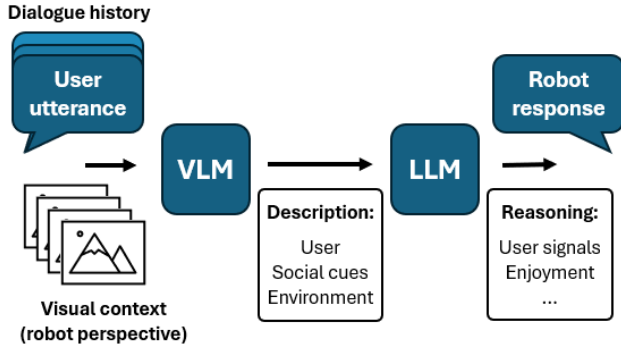


Fig. 2: Proposed architecture of a multimodal adaptive social dialogue system for a social robot using a VLM and LLM.

that the user is drinking from a mug). Second, it should be able to *recognize and interpret social cues and signals*. Using these cues, the robot can build a mental model of the user and of their emotions and cognitive processes, and adapt to them. This can help the robot to make the conversation more enjoyable for the user. Additionally, it is an essential capability for the robot to recognise the user’s feedback through these multi-modal social cues [10], potentially repairing dialogue breakdowns. In particular in educational applications such as language learning, the robot should keep track of whether the user understands its speech, in order to model and adapt to the learner’s proficiency. Finally, it should *understand the environment* the interaction is taking place in, and be able to make and understand references to that environment.

Processing this information is challenging. Social cues in particular are known to be subtle and fast [11], highly individual [12], and context-dependent [13]. While existing models are able to recognise or quantify facial expressions, pose, and gestures, they are not able to interpret them further than e.g. basic categories or representations of emotions. Furthermore, the vision model should be sufficiently general to work in any environment and with any user—not exhibiting harmful biases. Finally, visual recordings of users, their non-verbal expressions, and their environment, are personal data and should be processed and stored with respect for the user’s privacy, particularly in sensitive applications, e.g. with vulnerable users such as children or older adults.

The current generation of VLMs shows promising performance on a large variety of vision-language tasks [14] including integrating vision in conversations [15]. Combined with their growing context windows [16] (allowing the processing of many images at once), as well as high-resolution picture input [17], we argue they could be able to process and represent the information outlined in Figure 1. We propose to test whether they are able to do so, and whether they are able to interpret these signals to successfully adapt the conversation to these signals using the “world knowledge” they contain. While they are not trained to process visual input from an egocentric, embodied perspective, they might be able to interpret the visual information in such a scenario through a multi-step approach, leveraging the descriptive

and reasoning capabilities of VLMs and LLMs. A proposal architecture for this multi-step approach is shown in Figure 2.

Besides the interpretation of these cues, processing them still poses significant technical challenges. As social cues can be very short-lasting, such as microexpressions, visual information needs to be processed at a high framerate. This might necessitate intelligent sampling of these frames from the video, e.g. based on relevant movements. Furthermore, the visual information in the frames of the video needs to be aligned with the words that are spoken at that time. Finally, all processing and generation steps in the pipeline need to be executed in a sufficiently short delay—generally, one second is viewed as a maximum delay in order to not hinder a real-time spoken dialogue [18].

While we propose first evaluating the performance of VLMs and LLMs in this setting using zero-shot prompting or in-context learning, it could be valuable to fine-tune these models. However, the availability of data is a limiting factor, as well as the design of an appropriate loss function, as there is not one “ground truth” answer in a social conversation nor a suitable metric that can be automatically calculated [19].

III. EVALUATION

Evaluating the performance of a social dialogue system for social robots is a difficult task and there is not yet an agreed-upon standard practice. In conversations without a particular goal besides having a social interaction, the topic of the conversation could be an important influencing factor on the “quality” of the conversation, and we are essentially evaluating an open-domain conversational system. Such systems can be evaluated on large-scale datasets of conversations gathered from real-world deployments of the system [20], which is difficult to obtain with social robots. Nevertheless, a dataset of conversations with this system could be rated, either on a turn-by-turn basis or for the whole dialogue, by the user themselves [21] or an external annotator [22]. This rating could be performed for various concepts, such as the user’s experience of the conversation, e.g. their enjoyment, or the user’s perceptions of the robot. More detailed phenomena could also be rated [23], such as the presence of dialogue loops, self-disclosure by the user, dialogue breakdowns and their repairs, the accuracy of factual information, or the correctness of visual references. In particular regarding visual input, it is still unclear how to design a reproducible experiment that can evaluate the processing of visual cues in a social setting, going beyond references to visual objects or scripted conversation failures. Taking inspiration from evaluation practices for LLMs, evaluation could also be done by simple preference ratings between a pair of potential responses, and ratings could be performed by a state-of-the-art LLM acting as an “LLM-as-a-judge” [24]–[26]. Finally, objective metrics, such as the number of turns in the conversation and the length of the turns, can also provide useful information [23]. With this short discussion, we aim to prompt a discussion in the community about good practices in the unique field of multimodal social human-robot dialogue.

REFERENCES

- [1] V. Žegarac, “What is “phatic communication”?” in *Current Issues in Relevance Theory*, V. Rouchota and A. Jucker, Eds. John Benjamins, 1998, pp. 327–362.
- [2] M. Pinto-Bernal, M. Biondina, and T. Belpaeme, “Designing social robots with llms for engaging human interaction,” *Applied Sciences*, vol. 15, no. 11, p. 6377, 2025.
- [3] K. Nagao and A. Takeuchi, “Social interaction: Multimodal conversation with social agents,” in *AAAI*, vol. 94, 1994, pp. 22–28.
- [4] E. Fernández-Rodicio, Á. Castro-González, F. Alonso-Martín, M. Maroto-Gómez, and M. Á. Salichs, “Modelling multimodal dialogues for social robots using communicative acts,” *Sensors*, vol. 20, no. 12, p. 3440, 2020.
- [5] K. Jokinen and G. Wilcock, “Multimodal open-domain conversations with the nao robot,” in *Natural interaction with robots, knowbots and smartphones: Putting spoken dialog systems into practice*. Springer, 2013, pp. 213–224.
- [6] A. Addlesee, N. Cherakara, N. Nelson, D. H. García, N. Gunson, W. Sieińska, C. Dondrup, and O. Lemon, “Multi-party multimodal conversations between patients, their companions, and a social robot in a hospital memory clinic,” in *18th Conference of the European Chapter of the Association for Computational Linguistics 2024*. Association for Computational Linguistics, 2024, pp. 62–70.
- [7] P. Allgeuer, H. Ali, and S. Wermter, “When robots get chatty: Grounding multimodal human-robot conversation and collaboration,” in *International Conference on Artificial Neural Networks*. Springer, 2024, pp. 306–321.
- [8] R. Janssens, P. Wolfert, T. Demeester, and T. Belpaeme, “Integrating visual context into language models for situated social conversation starters,” *IEEE Transactions on Affective Computing*, 2024.
- [9] G. Pantazopoulos, J. Bruyere, M. Nikandrou, T. Boissier, S. Hemanthage, B. K. Sachish, V. Shah, C. Dondrup, and O. Lemon, “Vica: Combining visual, social, and task-oriented conversational ai in a healthcare setting,” in *Proceedings of the 2021 International Conference on Multimodal Interaction*, 2021, pp. 71–79.
- [10] A. Axelsson, H. Buschmeier, and G. Skantze, “Modeling feedback in interaction with conversational agents—a review,” *Frontiers in Computer Science*, vol. 4, p. 744574, 2022.
- [11] E. G. Krumhuber, A. Kappas, and A. S. Manstead, “Effects of dynamic aspects of facial expressions: A review,” *Emotion Review*, vol. 5, no. 1, pp. 41–46, 2013.
- [12] S. D’Mello, A. Kappas, and J. Gratch, “The affective computing approach to affect measurement,” *Emotion Review*, vol. 10, no. 2, pp. 174–183, 2018.
- [13] B. Mesquita and M. Boiger, “Emotions in context: A sociodynamic model of emotions,” *Emotion Review*, vol. 6, no. 4, pp. 298–302, 2014.
- [14] J. Zhang, J. Huang, S. Jin, and S. Lu, “Vision-language models for vision tasks: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [15] G. A. Abbo and T. Belpaeme, “I was blind but now i see: Implementing vision-enabled dialogue in social robots,” in *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2025, pp. 1176–1180.
- [16] A. Grattafiori, A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Vaughan, *et al.*, “The llama 3 herd of models,” *arXiv preprint arXiv:2407.21783*, 2024.
- [17] G. Team, “Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities,” Google DeepMind, Tech. Rep., 2025.
- [18] G. Skantze, “Turn-taking in conversational systems and human-robot interaction: a review,” *Computer Speech & Language*, vol. 67, p. 101178, 2021.
- [19] R. Janssens, P. Wolfert, T. Demeester, and T. Belpaeme, ““cool glasses, where did you get them?” generating visually grounded conversation starters for human-robot dialogue,” in *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2022, pp. 821–825.
- [20] A. Venkatesh, C. Khatri, A. Ram, F. Guo, R. Gabriel, A. Nagar, R. Prasad, M. Cheng, B. Hedayatnia, A. Metallinou, *et al.*, “On evaluating and comparing open domain dialog systems,” *arXiv preprint arXiv:1801.03625*, 2018.
- [21] G. Hoffman and K. Vanunu, “Effects of robotic companionship on music enjoyment and agent perception,” in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2013, pp. 317–324.
- [22] N. Raj Prabhu, C. Raman, and H. Hung, “Defining and quantifying conversation quality in spontaneous interactions,” in *Companion Publication of the 2020 International Conference on Multimodal Interaction*, 2020, pp. 196–205.
- [23] M. Paetzel-Prüsmann, J. F. Lehman, C. J. Gomez, and J. Kennedy, “An automatic evaluation framework for social conversations with robots,” in *Proceedings of the 2024 International Symposium on Technological Advances in Human-Robot Interaction*, 2024, pp. 56–64.
- [24] L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. Xing, *et al.*, “Judging llm-as-a-judge with mt-bench and chatbot arena,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 46 595–46 623, 2023.
- [25] A. Pereira, L. Marcinek, J. Miniota, S. Thunberg, E. Lagerstedt, J. Gustafson, G. Skantze, and B. Irfan, “Multimodal user enjoyment detection in human-robot conversation: The power of large language models,” in *Proceedings of the 26th International Conference on Multimodal Interaction*, 2024, pp. 469–478.
- [26] R. Janssens, A. Pereira, G. Skantze, B. Irfan, and T. Belpaeme, “Online prediction of user enjoyment in human-robot dialogue with llms,” in *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2025, pp. 1363–1367.