Simulating Human Interactions for Social Behaviour Coaching

Daniela Komenda, Livio Bürgisser, Noémie Käser, and Alexandre de Spindler
Zurich University of Applied Sciences, Winterthur, Switzerland
{komendan,buergli1,kaesenol,desa}@zhaw.ch

Abstract

Many individuals struggle with informal interactions like small-talk, which are vital in daily and professional settings. We introduce a conversational agent that combines a state-based interaction model with a social behaviour regulation (SBR) layer to provide structured coaching and real-time conversational modulation. The agent dynamically addresses issues such as oversharing or topic divergence and triggers coaching interventions based on user disengagement or inappropriateness. An exploratory study with neurodivergent-focused educators suggests the system's potential to foster socially appropriate communication. Our work shows how modular prompt orchestration can enhance both adaptability and the pedagogical value of conversational agents.

1 Introduction

Many individuals face challenges when navigating everyday verbal interactions, particularly in social or professional settings where informal conversations or spontaneous exchanges are expected. This is especially true for those with social communication difficulties, such as individuals on the autism spectrum or others who experience anxiety, cultural dissonance, or uncertainty around conversational norms.

Traditional approaches to improving these skills include self-directed learning through books, online courses, or video resources, as well as formal interventions such as coaching or workshops. However, learners lack opportunities to repeatedly engage in realistic, simulated interactions that provide immediate and contextualised feedback. Without iterative practice in conversation scenarios, the ability to generalise and apply new strategies in real-world settings may remain limited.

Conversational agents, powered by modern language models, offer an opportunity to bridge this gap. However, designing such agents for small-talk training demands more than simply role-playing a dialogue partner. For example, a training may require that the agent can realistically simulate a small-talk counterpart, regulate its conversational style to help uncover user deficiencies, detect those deficiencies in real time, deliver targeted coaching interventions, and return the user to the simulated scenario to apply what they have learnt. Achieving such complex agent behaviour necessitates multiple layers of adaptability, where the agent dynamically adjusts its behaviour based on various user cues and interaction patterns. Additionally, controlling a powerful language model across these layers must be done with precision to ensure consistent, interpretable, and reliable outputs.

We present a chatbot-based coaching system that demonstrates how two layers of adaptation can be specified independently yet operate in tandem. The system combines state-based adaptability, where structured prompts define interaction flows and transitions such as from simulated small-talk to coaching, with social behaviour regulation, which dynamically modulates in-state behaviour based on real-time conversational cues detections. By dissecting instructions into modular, minimal prompts, we retain tight control over the language model's behaviour while leveraging its natural ability to engage in open, human-like conversations. This layered and modular approach to agent adaptability enables more realistic and pedagogically effective coaching scenarios for social communication training, creating robust and responsive verbal interactions for improving small-talk skills, and serving as a building block for more immersive agentic systems such as avatars or social robots.

2 Related Work

Large language models (LLMs) have transformed conversational AI by enabling agents to engage in open-ended, context-aware dialogue. This progress has opened new opportunities for coaching applications, where agents must sustain natural conversations while guiding users toward learning goals (Aymerich-Franch and Ferrer, 2022). However, to effectively enhance learning, timely and contextualised feedback is essential (Hattie, 2008; Ajogbeje, 2023). Therefore, coaching systems must go beyond the conversational fluency LLMs inherently provide to detect user behaviours and deliver pedagogically meaningful interventions (Liu et al., 2025).

One key layer is the structuring of interaction flows to implement pedagogic coaching strategies. While fine-tuning LLMs to adopt specific coaching behaviours is possible, such methods are often resource-intensive, inflexible, and impractical when designing agents that must support a variety of coaching or training methods (Hadi et al., 2023). As a result, prompt engineering has emerged as a more feasible alternative for dynamically shaping agent behaviour (White et al., 2023). Yet, complex coaching scenarios often demand multi-step and layered prompts, which can increase the risk of inconsistent or unreliable outputs if merged into one large prompt (Long et al., 2024).

To address these challenges, the PROMISE framework (Wu et al., 2024a,b) introduced a state-based prompt orchestration approach. By decomposing complex instructions into modular and precise prompts tied to conversational states and transitions, PROMISE enhances LLM controllability (Helland et al., 2023) and supports the creation of structured, coherent coaching dialogues while leveraging the model's generative capabilities.

However, structuring the dialogue alone is insufficient for effective coaching, which also relies on real-time social adaptability. Persuasion techniques, conversational tone modulation, and user-tailored feedback must be dynamically selected based on in-the-moment user behaviour and conversational cues (Woolf et al., 2009). Such factors that cannot be fully predefined at design time.

We therefore extend PROMISE with a Social Behaviour Regulation (SBR) layer, which enables conversational agents to detect and respond to verbal cues such as oversharing, awkwardness, or deep talk divergence. This two-layered approach separates dialogue management via PROMISE from fine-grained behavioural modulation via SBR, allowing for orthogonal and layered adaptability. Together, these mechanisms enhance the agent's ability to deliver personalised and impactful coaching interventions, while maintaining reliable control over the language model's behaviour.

3 Use Case

We present a conversational agent designed to support users in practising small-talk while receiving adaptive, real-time coaching as exemplified in Fig. 1. The agent applies two layers of adaptability, a state-based interaction model that transitions the user between small-talk (light green) and coaching (dark green), and a SBR layer that modulates the simulated colleague's behaviour within the small-talk state.

In the first user utterance, the agent detects that the user introduces a philosophical topic, shifting the conversation into deep talk. In response, the agent's SBR adapts the behaviour by redirecting the conversation to light and casual topics. As the user exhibits signs of introversion by providing a mismatched and disengaged response in their second utterance, the agent transitions to a coaching state. As a result, the user receives feedback, suggestions for alternative responses, and the option to obtain more advice or return to the conversation.

While this scenario is highly simplified to convey the key idea within limited space, it demonstrates how in-state behaviour modulation and state-based transitions complement each other. The SBR layer helps surface conversational challenges through subtle adjustments, while the state model delivers targeted coaching interventions when necessary. Together, these layers create a responsive and iterative learning environment.

4 Approach

The agent architecture integrates two complementary layers of adaptability: state-based interaction management and SBR. Both layers rely on the PROMISE framework to structure and control language model prompts through modular and precise instructions.

4.1 State-Based Adaptability

PROMISE coordinates both state prompts for response generation and transition prompts for managing flow between states. At the core of an agent system, the state-machine-based interaction model depicted in Fig. 2 orchestrates the overall conversation flow. This model defines an outer state (grey) that ensures the agent consistently alternates be-

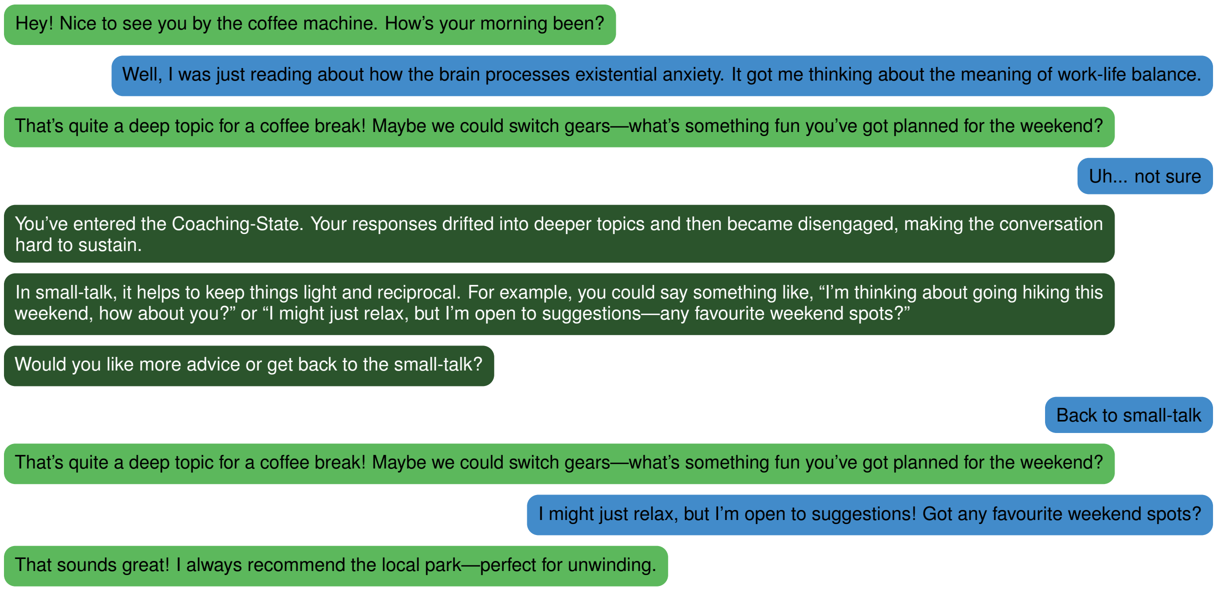

Figure 1: Example of how SBR and state transitions combine to redirect and coach the user during small-talk.

tween two roles: simulating a small-talking colleague or delivering coaching interventions.

The simulation state (green) models casual workplace interactions, allowing users to engage in small-talk with a fictitious co-worker. From this state, transitions occur when conversational issues are detected, leading into issue-specific coaching states that provide targeted feedback. Offensiveness coaching (blue) is activated when user utterances are perceived as offensive or inappropriate. Introversion coaching (purple) is activated when users display minimal engagement or withdrawal. After receiving feedback in a coaching state, users return to the small-talk simulation to apply newly acquired strategies.

4.2 Social Behaviour Regulation (SBR)

The SBR layer runs in parallel with the state machine and continuously analyses conversational cues during small-talk simulations. It performs real-time detections of conversational aspects.

For each detection, the system applies behaviour adaptations by appending specific prompt instructions to the original state prompt, as illustrated in Fig. 3. This helps to modulate the baseline behaviour of the simulated person as illustrated with the following examples.

Deep topic divergence	$\rightarrow$ Redirect to light, casual topics
Awkwardness	$\rightarrow$ Select familiar, casual topics
Oversharing	$\rightarrow$ Limit to safe and neutral topics

These adaptations help surface conversational deficiencies and prepare the ground for targeted coaching interventions.

4.3 Integrated Adaptation

The dual-layer architecture of our system integrates a state-based interaction model (Baseline Behaviour) with either a smalltalk state or a coaching state, layered on top of an SBR module. The baseline behaviour, implemented via the PROMISE framework, manages the overall conversation flow through a state machine. Upon receiving each user utterance, PROMISE appends it to the conversation history, determines the current state, and checks transition conditions to decide whether to move into a new state. Transition conditions are encoded as prompts used to instruct the underlying language model to analyse the conversation history and make a decision. The active state points at the prompt to be used to generate the agent's response to the user. As a result, different prompts are used depending on the state currently active.

Running in parallel with the state machine, the SBR layer processes the same interaction history to detect conversational issues using dedicated prompts. Based on these detections, it selects behaviour regulation strategies, encoded as prompt elements, and appends them to the active state prompt provided by the baseline behaviour. This combined prompt enables the agent to produce socially adaptive and context-sensitive responses. Thus, while the state model governs high-level conversation flow and role-switching between simulation and coaching, the SBR layer fine-tunes in-state

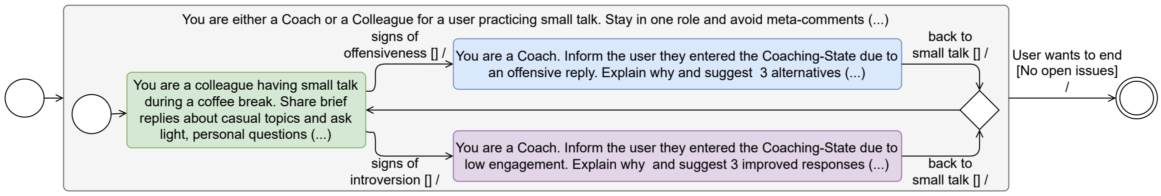

Figure 2: State model with transitions between small-talk and coaching based on detected conversational issues.

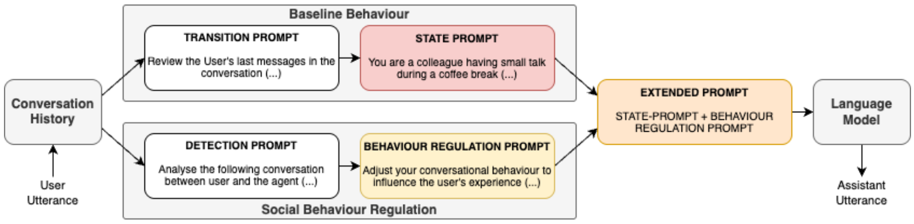

Figure 3: Social behaviour regulation flow combining state prompt and behaviour regulation prompt

responses, ensuring that conversations remain both structured and dynamically responsive to evolving user inputs.

5 Validation

To assess the chatbot's ability to support users in improving small-talk skills, we conducted exploratory user testing with eight special needs teachers specialised in supporting neurodivergent learners. The objective was to evaluate whether the agent could effectively simulate workplace small-talk scenarios and deliver timely coaching interventions.

Participants engaged in five to ten minute sessions with the chatbot, alternating between two predefined colleague personas (male, female). The agent successfully maintained realistic small-talk while triggering coaching interventions in cases of conversational disengagement (introversion) or inappropriate responses (offensiveness). Notably, only three users entered a coaching state, all cases activating the introversion coaching, while no occurrences of offensiveness were observed.

User evaluations were given by responses to an adapted chatbot usability questionnaire (Holmes et al., 2019). The responses summarised in Fig. 4 highlight positive perceptions of the chatbot's conversational realism, appropriateness, and perceived understanding. The three users who entered coaching states also rated the provided feedback as helpful and relevant to small-talk scenarios.

While the small sample limits generalisability, the results suggest that the dual-layered adaptability approach creates a plausible and responsive environment for practising small-talk.

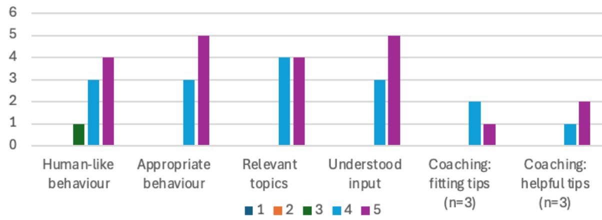

Figure 4: Chatbot usability ratings (1 = low, 5 = high)

6 Conclusion

We introduced a conversational coaching agent designed to help users improve their small-talk skills through a layered approach to adaptability. By combining a state-based interaction model with SBR, the system delivers both structured conversation flows and real-time behavioural adjustments. This dual mechanism enables the agent to simulate realistic social interactions, detect conversational challenges, and deliver targeted coaching interventions. The modularity of the approach ensures precise control over the language model's behaviour, while maintaining the flexibility needed to support dynamic, user-centred training scenarios.

Initial user testing suggests that the agent can provide a usable and engaging environment for practising small-talk. However, due to the small number of test participants, further studies are necessary to validate these findings.

In future work, we aim to extend this framework to support richer interaction modalities, including non-verbal cues. Additionally, the layered design can be adapted to other socially complex domains such as customer support, negotiation training, or conflict resolution.

References

Oke James Ajogbeje. 2023. Enhancing classroom learning outcomes: The power of immediate feedback strategy. *International Journal of Disabilities Sports and Health Sciences*, 6(3):453–465.

Laura Aymerich-Franch and Iliana Ferrer. 2022. Investigating the use of speech-based conversational agents for life coaching. *International Journal of Human-Computer Studies*, 159:102745.

Muhammad Usman Hadi, Rizwan Qureshi, Abbas Shah, Muhammad Irfan, Anas Zafar, Muhammad Bilal Shaikh, Naveed Akhtar, Jia Wu, Seyedali Mirjalili, et al. 2023. Large language models: a comprehensive survey of its applications, challenges, limitations, and future prospects. *Authorea Preprints*.

John Hattie. 2008. *Visible learning: A synthesis of over 800 meta-analyses relating to achievement*. routledge.

Solveig Helland, Elena Gavagnin, and Alexandre de Spindler. 2023. Divide et impera: Multi-transformer architectures for complex NLP-tasks. In *Proceedings of the 8th edition of the Swiss Text Analytics Conference*, pages 70–75, Neuchatel, Switzerland. Association for Computational Linguistics.

Samuel Holmes, Anne Moorhead, Raymond Bond, Huiru Zheng, Vivien Coates, and Michael Mctear. 2019. Usability testing of a healthcare chatbot: Can we use conventional methods to assess conversational user interfaces? In *Proceedings of the 31st European Conference on Cognitive Ergonomics (ECCE 2019)*, pages 207–214, New York, NY, USA. ACM.

Vincent Liu, Ehsan Latif, and Xiaoming Zhai. 2025. Advancing education through tutoring systems: A systematic literature review. *arXiv preprint arXiv:2503.09748*.

Do Xuan Long, Duong Ngoc Yen, Anh Tuan Luu, Kenji Kawaguchi, Min-Yen Kan, and Nancy F Chen. 2024. Multi-expert prompting improves reliability, safety, and usefulness of large language models. *arXiv preprint arXiv:2411.00492*.

Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382*.

Beverly Woolf, Winslow Burleson, Ivon Arroyo, Toby Dragon, David Cooper, and Rosalind Picard. 2009. Affect-aware tutors: recognising and responding to student affect. *International Journal of Learning Technology*, 4(3/4):129–164.

Wenyuan Wu, Jasmin Heierli, Max Meisterhans, Adrian Moser, Andri Färber, Mateusz Dolata, Elena Gavagnin, Alexandre de Spindler, and Gerhard Schwabe. 2024a. Promise: A framework for model-driven stateful prompt orchestration. In *Intelligent Information Systems: CAiSE Forum 2024, Limassol, Cyprus, June 3–7, 2024, Proceedings*, volume 520 of *Lecture Notes in Business Information Processing*. Springer International Publishing.

Wenyuan Wu, Jasmin Heierli, Max Meisterhans, Adrian Moser, Andri Färber, Mateusz Dolata, Elena Gavagnin, Alexandre de Spindler, and Gerhard Schwabe. 2024b. Promise: Model-driven stateful prompt orchestration for persuasive conversational interactions. In *Proceedings of the 9th edition of the Swiss Text Analytics Conference*, pages 185–185, Chur, Switzerland. Association for Computational Linguistics.