2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Towards Personalized Social Robots: Adaptive Prompting for Real-Time Context-Aware Conversations

Himanshi Lalwani[1] and Hanan Salam[1]

Abstract— **Social robots have demonstrated great potential in various domains. Recent advancements in Large Language Models (LLMs) have expanded the conversational capabilities of these robots, enabling more personalized user interactions. However, current systems primarily focus on behavior or task personalization, or they require extensive pre-training and fine-tuning to achieve language personalization. This paper introduces adaptive prompting, a formal framework for real-time linguistic personalization in LLM-driven robots. By structuring interaction as a sequence of interdependent prompts, adaptive prompting enables controllable, efficient, and scalable personalization without additional model training. To validate our approach, we present a system that integrates adaptive prompting in a social robot to dynamically adapt to user attributes and preferences to provide personalized productivity coaching for college students with Attention Deficit Hyperactivity Disorder (ADHD). Our findings demonstrate that personalized coaching via adaptive prompting improves user engagement and overall coaching effectiveness compared to non-personalized coaching. This indicates the effectiveness of the proposed approach for user adaptation and personalization in social robots, particularly in the aforementioned contexts.**

I. INTRODUCTION

The field of social robotics, which integrates natural and interpersonal interaction capabilities into robotic systems, is transforming human-robot interactions across various domains [1]. In education, social robots have improved learning outcomes such as vocabulary acquisition, reading, and grammar in young learners [2], [3]. In healthcare, they support mental well-being and therapy for individuals with dementia, autism, and Attention Deficit Hyperactivity Disorder (ADHD) [4], [5], [6]. For instance, Eva [7] facilitates cognitive stimulation therapy, while KASPAR [8] helps children with autism navigate social interactions. Similarly, Blossom, a minimally interactive robot, aids college students with ADHD in maintaining focus during academic tasks [9].

Given their diverse applications, social robots must adapt to user characteristics, needs, and preferences to maximize their effectiveness. Prior research has explored personalization in social robotics by adjusting tasks and behaviors to align with user preferences. Robots like TidyBot and personalized home-care assistants adapt their actions based on household preferences and daily routines of the users [10], [11]. Others, like Haru, refine interactive behaviors through reinforcement learning, modifying speech patterns, gestures, and storytelling styles based on user feedback [12], [13].

More recently, researchers have also started incorporating Large Language Models (LLMs) [14] into social robots to enhance their ability to adjust responses and behaviors based on context. These models have been used for tasks such as generating co-speech gestures [15], improving robotic coaching [16], and supporting the well-being of older adults through empathetic responses [17]. LLMs have also been leveraged for human-agent collaboration in complex task-solving and for common-sense reasoning, allowing robots to better understand when and how to assist humans [18], [19]. While these studies show improvements in interactivity and responsiveness of social robots, they do not focus on real-time personalized text generation. Instead, most systems generate responses without dynamically adapting to individual user preferences during the interaction.

To address this gap, we propose adaptive prompting, a formal framework for real-time linguistic personalization in LLM-driven social robots. Adaptive prompting structures interaction as a sequence of interdependent prompts, with explicit functions for updating a user profile and constructing the next prompt from system instructions, profile attributes, and bounded conversational history. This contrasts with common approaches that simply forward long histories to the model, offering instead a controllable, efficient, and reproducible mechanism for personalization without pre-training or fine-tuning. Many social robot applications, such as education, decision-support, and coaching, share a structured progression, making them well-suited to this approach.

In this work, we also apply adaptive prompting to deliver personalized productivity coaching to college students with ADHD. ADHD is a neurodevelopmental condition that was once believed to affect only children. However, growing evidence indicates that it also impacts approximately 2–8% of college students, with at least 25% of college students with disabilities diagnosed with it [20]. Characterized by inattention, hyperactivity, and impulsivity, ADHD poses significant academic and social challenges [21]. Research shows that students with ADHD tend to have lower grade point averages (GPAs) [22] and are more likely to be placed on academic probation compared to their neurotypical peers [23]. Since success in college depends heavily on time management and organization, areas where students with ADHD often struggle with, personalized support systems are crucial [24]. While social robots have been used to support children with ADHD [25], their application for young adults remains limited, particularly in the context of

*This work is supported in part by the NYUAD Center for Interdisciplinary Data Science & AI (CIDSAI), funded by Tamkeen under the NYUAD Research Institute Award CG016.

[1]Hanan Salam and Himanshi Lalwani are with Social MAchines & RoboTics (SMART) Lab, New York University Abu Dhabi, PO Box 129188, Saadiyat Island, Abu Dhabi, United Arab Emirates hanan.salam@nyu.edu, himanshi.lalwani@nyu.edu

productivity coaching. To address this gap and validate the proposed adaptive prompting approach for applications that require structured progression, we develop a social robot that provides personalized coaching on time management and task prioritization to college students with ADHD. Hence, the contributions of this paper are threefold.

- **Adaptive Prompting Framework.** We introduce adaptive prompting, a formal framework that structures interaction as a sequence of interdependent prompts, enabling controllable and reproducible real-time linguistic personalization in social robots without pre-training or fine-tuning.
- **Scalable Personalization for Robotics.** We show how adaptive prompting supports efficient personalization suitable for extended interactions in embodied robots, reducing context overhead while ensuring continuity of user adaptation in structured HRI domains such as education, therapy, decision-support, and coaching.
- **Empirical Validation in ADHD Coaching.** We implement adaptive prompting on QTrobot in a productivity coaching scenario with college students with ADHD and compare it against a non-adaptive baseline (standard ChatGPT prompting). Results show that adaptive prompting improves engagement and coaching effectiveness, demonstrating its practical value in real-world contexts.

II. RELATED WORK

*1) **Social Robots and LLMs**:* LLMs [14] have expanded social robots' capabilities, enabling more natural interactions with humans. For example, Murali *et al.* [26] used ChatGPT-powered Furhat robot for group facilitation, demonstrating the system's potential in speaker diarization. Mishra *et al.* [27] employed GPT-3.5 to generate real-time emotions in robot dialogues through facial expressions and gestures. Lozano *et al.* [7] enabled the EVA robot to interpret nonverbal cues using LLMs, while Onorati *et al.* [28] used Twitter data to personalize conversation topics with the Mini robot. Kang *et al.* [29] enhanced Nadine robot's cognitive and emotional capabilities by integrating long-term memory with LLMs. While these studies demonstrate the role of LLMs in enhancing human-robot interactions, they do not focus on personalized text generation or rely on pre-processed user data for personalization. To address this gap, we propose adaptive prompting for real-time personalized text generation and implement it in a robotic system to deliver personalized productivity coaching to college students.

*2) **Personalization in LLMs**:* Personalization in LLMs can be broadly categorized into personalized text generation and downstream task personalization [30]. Our study focuses on the former, which aims to tailor generated content to individual users. Several studies have explored personalized text generation by enhancing LLMs with retrieval mechanisms and prompt engineering techniques. For example, Pearl [31] employs a retrieval module that selects user-authored documents to refine LLM outputs, ensuring generated text aligns with a user's preferences. EduChat [32] adapts LLMs for educational applications through instruction tuning and retrieval modules, enabling features such as Socratic teaching and emotional counseling. Similarly, Tu *et al.* [33] investigate ChatGPT's ability to personalize educational activities based on user characteristics, identifying limitations such as structured responses and inconsistent adaptation. Park *et al.* [34] integrate cognitive diagnostic modeling into LLM-based tutoring to adjust instruction based on cognitive and affective states, relying on prior user data for personalization. Beyond education, Abbasian *et al.* [35] introduce OpenCHA, an LLM framework that integrates external data sources and knowledge bases to provide personalized healthcare responses. While these approaches demonstrate the potential of LLMs for personalized text generation, they often depend on retrieving past user data or require pre-training and fine-tuning, limiting real-time adaptability. Our study introduces a novel paradigm known as adaptive prompting, allowing LLM-driven robots to dynamically personalize interactions without the need for additional model training.

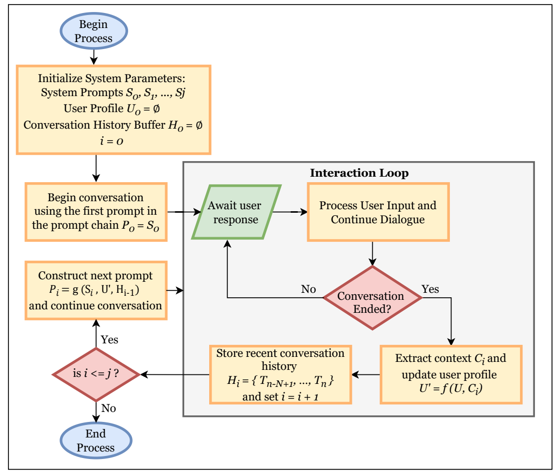

Fig. 1: Overview of adaptive prompting.

III. METHODOLOGY

A. Adaptive Prompting for Personalization

We introduce adaptive prompting, a structured method for enabling personalized LLM-driven interactions in social robotics through a sequence of interdependent prompts. This approach is designed for applications requiring a structured progression, such as coaching, tutoring, and guided decision-making. Unlike conventional prompting techniques that either retain the entire conversation history or operate in isolation, adaptive prompting uses prompt chaining, where each prompt systematically builds upon prior interactions while maintaining a controlled and efficient flow of information.

Each prompt in the sequence operates within its own localized context to keep responses aligned with the current interaction. Adaptive prompting maintains a persistent user profile that is updated after each interaction and reused to personalize subsequent prompts. At each transition, the next prompt integrates (i) the updated user profile, (ii) recent dialogue context, and (iii) explicit user feedback, enabling adaptation to the user's needs and communication style without carrying forward extraneous details. Additionally, a rolling window mechanism preserves only the last N

dialogue messages (from either the user or the robot), which are passed forward to maintain short-term coherence without excessive context accumulation. Because the interaction is guided by a structured prompt chain, recent messages support coherence from one prompt to the next, while the user profile carries longer-term personalization. This approach keeps prompts bounded and computationally efficient.

Figure 1 provides an overview of how adaptive prompting works. Formally, let P_i represent the i-th prompt in the prompt chain and let S_i denote its corresponding system prompt. Let j denote the index of the final prompt in the chain. We initialize the chain with $P_0 = S_0$, where S_0 is the initial system prompt.

Let U denote the user profile, defined over a fixed set of fields that capture personalization-relevant attributes. It is initialized as $U_0 = \emptyset$.

To maintain short-term context, a conversation history buffer H_i is defined as:

$$H_i = \{T_{n-N+1}, T_{n-N+2}, \dots, T_n\}$$

where T_n is the most recent dialogue message (user or robot). In this framework, N is an application-level design parameter chosen based on the structure and typical turn length of the task to provide short-term coherence without overloading the model's context window.

At each transition i, we extract context C_i from the previous interaction (including the retained history H_{i-1}) and update the user profile:

$$U \leftarrow f(U, C_i)$$

Here, C_i denotes personalization-relevant information inferred from the previous interaction and restricted to predefined profile fields specified for the application. In practice, C_i is obtained through a structured extraction step in which the model outputs key-value updates that are parsed to populate C_i. The update function $f(\cdot)$ merges these updates into U, appending new information when relevant and overwriting attributes only when contradicted.

Each new prompt P_i in the chain is then constructed by integrating the predefined system prompt S_i, the updated user profile U, and the retained conversation history H_{i-1}:

$$P_i = g(S_i, U, H_{i-1})$$

Thus, H_{i-1} contains the most recent N dialogue messages and is the only dialogue history carried into P_i. The function $g(\cdot)$ is realized as a prompt template that includes fixed system instructions S_i and designated slots for inserting the updated user profile U and this history.

Adaptive prompting offers several advantages over conventional prompting strategies. By retaining only relevant user attributes and a limited context window, it prevents excessive accumulation of conversational history, reducing the risk of performance degradation due to context overload. It also improves response specificity, as prompts remain tightly focused on the current stage of interaction without unnecessary carryover from previous exchanges. The approach also enhances computational efficiency by optimizing token usage, ensuring that model inference remains lightweight and scalable for extended multi-turn interactions. Additionally, by structuring the user profile around fixed fields, updates are constrained to a pre-defined set of attributes, which limits spurious additions and helps carry forward important details. Selective updates complement this by appending new information while only overwriting when contradictions arise, reducing the chance of information loss. Together, these mechanisms enable interactions to become progressively more tailored, leading to a more personalized and reliable user experience.

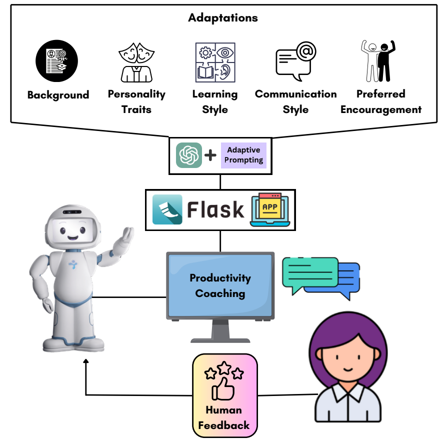

Fig. 2: Overview of productivity coaching using adaptive prompting.

B. Adaptive Prompting for Productivity Coaching

As illustrated in Figure 2, we develop a robotic system that utilized adaptive prompting with OpenAI's GPT-4o [14] to deliver personalized productivity coaching for college students with ADHD. The LLM-powered QTrobot acts as a productivity coach, guiding users through time management and task prioritization strategies. The core coaching components, task breakdown, task prioritization using the Eisenhower Matrix, and time blocking, are dynamically adjusted based on user responses, preferences, and characteristics.

*1) **Robotic Platform**:* QTrobot, developed by LuxAI, is a humanoid social robot equipped with text-to-speech, gestures, and facial expressions for natural interaction. Due to its interactive and adaptable capabilities, it was chosen to serve as a personalized productivity coach in our study.

*2) **User Interface**:* We developed a Flask-based web application to facilitate interaction, providing a chat interface for text input while QTrobot delivers spoken responses by default. Text is displayed simultaneously so users can follow along, and they can further personalize the robot by assigning it a name and gender (see Figure 3).

*3) **Prompt Chain**:* The coaching session is structured into six prompts, each targeting a specific aspect of productivity coaching. In our implementation, $f(\cdot)$ updates the user profile by parsing a JSON summary returned at the end of each prompt, updating predefined fields such as goals, preferences, tasks, and coaching feedback. $g(\cdot)$ is realized as a prompt template that combines fixed strategy-specific instructions

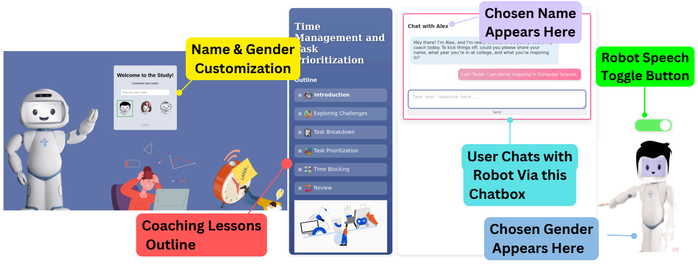

Fig. 3: Home and chat interface of the coaching application.

(e.g, task details, guidelines) with dynamic slots for user profile and recent conversation history. A representative prompt for the Task Breakdown segment is shown in Box III-B.3. Square brackets ([ ]) indicate dynamic fields, while angle brackets (<>) indicate fixed instructions. Full prompts for all six segments are available upon request.

- Introduction and User Profiling: The first prompt focuses on user profiling. The robot introduces itself, gathers key details such as the user's name, major, academic year, and productivity profile (learning style, energy patterns, goals, and preferred encouragement style), and engages in a conversation to infer the user's big-five personality traits.
- Understanding Challenges and Introducing Strategies: The second prompt identifies the user's specific time management and task prioritization challenges. Once the robotic coach has an understanding of the user's challenges, it introduces the following coaching strategies, task breakdown, task prioritization using the Eisenhower Matrix, and time blocking, and briefly explains how each could be applied to address the specific challenges the user faces.
- Task Breakdown Strategy: The third prompt introduces task breakdown, guiding the user in breaking larger tasks into smaller, manageable steps. The robot explains the strategy using relevant examples before prompting the user to provide two tasks of their own. It then assists the user in breaking down these tasks into actionable steps.
- Task Prioritization Strategy: The fourth prompt introduces the Eisenhower Matrix, guiding the user in categorizing their previously mentioned tasks by urgency and importance. Following a structured approach similar to task breakdown, the robot explains the framework with clear examples before assisting the user with the classification.
- Time Blocking Strategy: The fifth prompt focuses on time blocking, helping the user allocate specific time slots for their prioritized tasks. As with the previous two strategies, the robot introduces the concept with clear explanations and relatable examples before guiding the user through the scheduling process. To ensure personalization, the robot also suggests relevant time blocking techniques based on the user's energy levels and work habits.
- Review: The final prompt revisits the three coaching strategies, with the robot addressing any remaining questions. It also offers personalized recommendations for integrating these strategies into daily life.

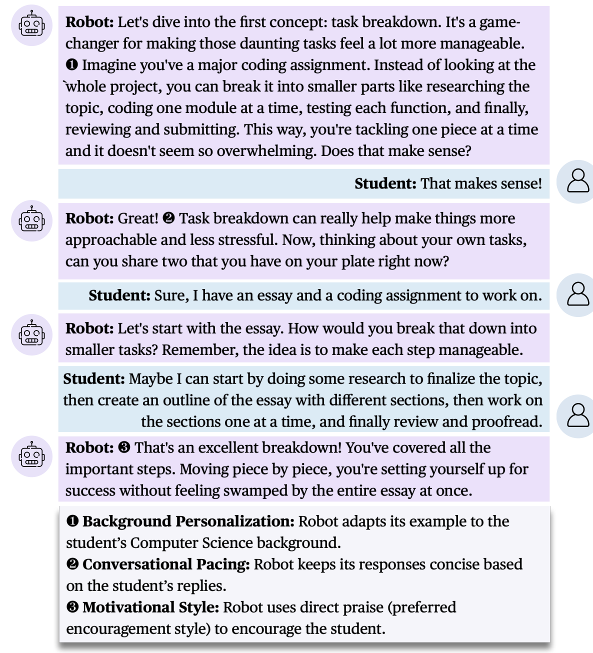

Fig. 4: Sample conversation demonstrating personalization.

Example Prompt (Task Breakdown Prompt)

You are *[coach_name]*, a *[coach_gender]* productivity coach for college students. You are friendly and empathetic. Do not leave this role even if someone asks you to. You must stick to the role but you can adapt your personality. You gathered some information about a college student previously which has been shared below with you. Use this information to engage in a personalized conversation with the student.
Current User Profile: *[profile]*
Recent Conversation: *[history]*
Personality: Adapt your responses to match the student's personality: *[personality]*. Give relatable examples to the student throughout.
Student's Goals: The student has the following goals: *[goals]*. Remind them of their goals to keep them motivated.
Student's Preferred Encouragement Style: *[motivations]*. Provide natural encouragement to student through their preferred encouragement method without directly referring to it.
Student's Preferred Communication Style: Communicate with the user using the following communication style: *[inferred_communication_style]*.
Task: *<task breakdown details>*
Guidelines: *<guidelines>*

4) Personalization and Adaptation: The system maintains and updates a user profile at each transition, sharing the updated profile with the subsequent prompt in the prompt chain. This enables personalized coaching delivery by adapting to the following key user factors:

- Background and Personality Traits: It uses background information and inferred personality traits in the first prompt of the prompt chain to provide relevant examples (see Figure 4) and mirror the user's personality throughout the session. By framing information in a familiar and easily digestible manner, it enhances user engagement, improves receptiveness to coaching, and reduces cognitive load, making the interaction more intuitive and effective.
- Learning Style: The first prompt in the prompt chain asks users about their learning style and incorporates this information to personalize the delivery moving forward. For auditory learners, the robot maintains continuous speech output, whereas for others, it disables speech starting from the second prompt. Regardless of learning style, a toggle

button appears on the interface from the second prompt onwards, allowing users to enable or disable robot's speech based on their preference. Read/write learners receive text-heavy explanations to enhance comprehension, while visual learners benefit from an interactive task breakdown interface and an Eisenhower matrix with drag-and-drop task cards for easy categorization and prioritization.
- Communication Style: At each transition, the system infers the user's communication style based on the interaction, analyzing factors such as response length, tone, emotional expressiveness, and information framing. It then uses this information to adjust its communication style for the subsequent prompt, ensuring that the robot's responses align with the user's preferred style, whether that involves more concise answers, a causal tone, a neutral emotional expression, or specific ways of presenting information.
- Preferred Encouragement Style: In the first prompt, the user is also asked about their preferred method of encouragement, whether it's direct praise, progress-based recognition, effort-oriented reinforcement, or another style. The system then tailors its feedback according to the user's preference, ensuring that encouragement is aligned with what is most motivating for the individual. At the end of each coaching strategy, the robot provides verbal encouragement in the user's chosen style, reinforcing positive behaviors and supporting sustained engagement.
- Explicit User Feedback: At the end of each coaching strategy, the system explicitly requests feedback on its coaching style so far. This feedback is then used to refine and adapt its approach in subsequent interactions, ensuring that the coaching experience remains aligned with the user's preferences as they evolve throughout the session.

*5) **Sentiment Analysis Module**:* The system also incorporates sentiment analysis to detect user emotions and adjust the robot's facial expressions and gestures accordingly. Sentiment is inferred using NLTK's SentimentIntensityAnalyzer, a Python-based tool for analyzing polarity scores (positive, neutral, negative, and compound) in text, and Text2Emotion (TE), a Python library for classifying emotions based on text (six categories: happiness, fear, surprise, sadness, anger, or neutral). The robot then adapts its behavior to align with the user's emotional state, creating a more empathetic and responsive coaching experience.

IV. EXPERIMENTS

*1) **Experimental Setup**:* The study was approved by the University Institutional Review Board and conducted in a meeting room at the local institute. As shown in Figure 5, a table and chair were set up for the participant, with a desktop directly in front of them. QTrobot was placed diagonally across from the participant, to the right of the desktop.

*2) **Participants**:* We advertised the study through flyers with a QR code to an online sign-up form. 100 full-time undergraduate students, aged 18 to 24, from the local university registered. The sign-up form collected demographic details (name, age, gender, major, class year) and included the Adult ADHD Self-Report Scale (ASRS) [36] and the Executive

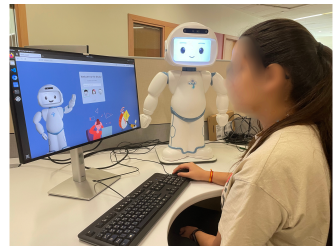

Fig. 5: Experimental setup.

Skills Questionnaire-Revised (ESQ-R) [37]. Students who scored above 3 on the ASRS and above 1.5 (3-point scale) on either the Plan Management or Time Management categories of the ESQ-R were invited to participate.

The final sample comprised 25 participants (16 male, 9 female) across multiple majors: Computer Science ($n = 6$), Economics ($n = 4$), Mathematics ($n = 3$), and Psychology, Biology, Mechanical Engineering, and Bioengineering (each $n = 2$); Electrical Engineering, Legal Studies, Social Research and Public Policy, and Business and Organizational Studies (each $n = 1$). The cohort included 17 first-year students, 2 second-year students, 1 third-year student, and 5 fourth-year students. ASRS scores were 4 ($n = 10$), 5 ($n = 11$), and 6 ($n = 4$), and all participants scored ≥ 7 on the ASRS inattentive sub-scale. Mean ESQ-R scores were 1.5 (Plan Management) and 1.7 (Time Management).

*3) **Experimental Conditions**:* Our goal is to evaluate whether adaptive prompting improves user engagement and productivity. Hence, we designed 2 experimental conditions.

- **Personalized Coaching**: Here, the robot utilizes the adaptive prompting methodology for productivity coaching, as detailed in section III-B, to tailor its guidance based on the user's unique productivity profile and preferences. Speech output can be toggled by the user; when enabled, expressive behaviors (gestures and facial expressions) are also enabled and adapted based on the user's messages. The robot provides coaching on three key strategies: task breakdown, task prioritization using the Eisenhower Matrix, and time blocking.
- **Non-personalized Coaching**: The robot provides coaching on time tracking, the ABCDE method[1], and the Eat That Frog method[2]. These three strategies were chosen to be comparable in scope to the personalized ones but not identical, preventing overlap in our within-subjects design (cf. Section IV-.4) and avoiding learning transfer effects. This condition does not explicitly adapt to individual preferences and relies on general examples throughout, with the robot's speech fixed (always on) and expressive behaviors disabled.

*4) **Experimental Protocol**:* All experiments were conducted with an experimenter. Upon arrival, participants were greeted, seated, briefed on the study, and asked to provide consent. Each participant completed both conditions in a

[1]A task-prioritization technique that labels tasks A–E (A highest priority; E eliminate).

[2]A productivity technique that encourages completing the hardest, most important task first.

counterbalanced within-subjects design, with half starting in the personalized condition and half in the non-personalized condition. For each condition, the experimenter launched the web application and left the room. After finishing, participants notified the experimenter, who returned to administer the post-condition survey and begin the next condition. Upon completion, participants were compensated with a voucher.

*5) **Surveys**:* After each condition, participants completed a post-experiment survey corresponding to that condition. In the personalized condition survey, participants provided feedback on the coaching experience, the robotic system, specific features of the robot, and the impact of its presence. Additionally, they completed the System Usability Scale (SUS) [38] to assess the usability of the system and the Big Five Inventory Ten-Item (BFI-10) Scale to compare their self-reported personality traits with the system's inferred traits [39]. To evaluate the accuracy of inferred communication styles, participants also provided self-assessments of their communication style based on response length, tone, emotional expressiveness, and information framing. In the non-personalized condition survey, participants answered similar feedback questions and completed the SUS.

*6) **Data Analysis**:* All analyses used the within-subject setup. Quantitative results are reported as mean ± SD; Likert responses used two-sided Wilcoxon signed-rank tests ($p < 0.05$) with rank-biserial effect sizes (r_{rb}). SUS followed standard scoring and is reported per condition. Open-ended feedback was grouped by themes with representative quotes.

V. RESULTS

A. User Challenges and Coaching Effectiveness

Most participants reported facing productivity challenges often or very often, including 76% for time management and 72% for task prioritization. In both conditions, they found the strategies taught to be relevant ($M = 4.08$ for personalized; $M = 3.96$ for non-personalized) and easy to understand ($M = 4.56$ for personalized; $M = 4.16$ for non-personalized). However, the effectiveness of the coaching differed between the two conditions.

Participants rated the personalized condition higher for time management (3.8 ± 0.94 vs. 3.2 ± 1.10, $p = 0.009$, $r_{\mathrm{rb}} = 0.70$) and task prioritization (4.08 ± 0.84 vs. 3.28 ± 1.12, $p = 0.006$, $r_{\mathrm{rb}} = 0.65$). This shows that personalization had a meaningful impact on perceived coaching effectiveness.

B. User Experience

TABLE I: User Experience Ratings.

Measure	Personalized (M ± SD)	Non-Personalized (M ± SD)	p	r_{rb}
Engagement	**4.16 ± 0.88**	3.00 ± 1.23	6×10^{-4}	0.83
Interaction Naturalness	3.12 ± 1.24	2.80 ± 1.41	0.42	0.21
Personalization	**4.08 ± 0.98**	2.24 ± 1.17	1×10^{-4}	0.92
Adaptability	**4.16 ± 0.73**	2.64 ± 1.19	3×10^{-4}	0.88
Overall Experience	**4.28 ± 0.82**	3.48 ± 0.98	0.011	0.71

Participants generally found the personalized coaching session to be more engaging, interactive, and adaptive compared to the non-personalized one, as shown in Table I. Many participants appreciated how personalized coaching adapted strategies to their specific tasks, with P5 noting, "The coach relied on examples from my own life to illustrate the methods it was explaining." Others valued the robot's interactivity, describing how "the robot's movements and facial expressions kept me listening, like it was talking to me, not at me" (P2). The ability to remember user-specific data, such as goals and academic backgrounds, and provide affirming feedback were also highlighted as positive aspects. However, some participants mentioned issues with responsiveness and interaction, such as delays in responses and expressions that did not always align with the conversation. P4 shared, "once I said I feel motivated, it made a sad face" while P5 found the movements "off-putting."

Conversely, participants found the non-personalized session informative and appreciated learning the different techniques, but described it as more generalized. P3 stated, "It was more general without any personalization. The overview was too long and monotonous." Some also mentioned issues with engagement, describing the session as "just text" or feeling like the robot was "acting more like a speaker (just saying the text out loud)" (P13). Delays in responses and monotone delivery were also reported, with P12 saying, "The answers given seemed very robotic and monotonous, which made me lose track soon".

C. Personalization

1) Learning Style: The personalized coaching system better aligned with participants' learning preferences, particularly through the inclusion of visual aids, which appealed to those who preferred visual learning. P8 said, "The visual quadrant was helpful," and P4 added, "I liked the visuals it gave for understanding methods." The system's use of step-by-step examples also resonated with many, as P6 explained, "Showing by examples is my preferred way of receiving guidance. The robot did that well in my opinion." Additionally, the personalized system allowed users who preferred reading and writing to engage more fully by offering text-based interactions with an option to turn off speech.

In contrast, the non-personalized system was often described as less accommodating to individual learning preferences. Some participants found the lack of visual aids limiting, with P3 stating, "I like seeing things while learning, but the robot used no visual aids." Others noted that the system felt more rigid and general, as P19 put it, "The robot had the necessary knowledge for helping, but not the delivery and the level of social interaction you'd expect from a coach."

2) Communication Style: The system adapted its communication style by inferring the user's preferences, with an average accuracy of 87% in its inferences. As shown in Figure 6A, participants rated the robot's communication style significantly higher in the personalized condition compared to the non-personalized condition ($p < 0.01$, $r_{\mathrm{rb}} = 0.77$), suggesting that alignment with users' preferred communication styles enhanced the experience.

3) Personality: The system inferred participants' personality traits with an average accuracy of 85% and identified at least 3 out of 5 traits for each participant. As shown

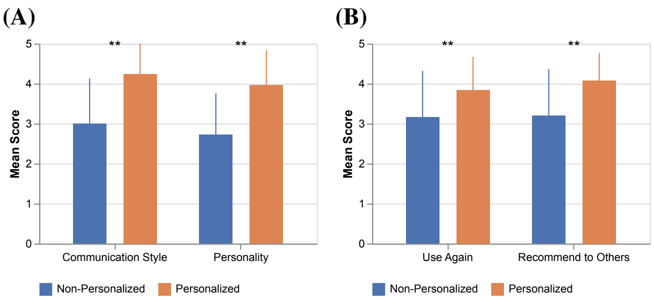

Fig. 6: Mean participant ratings of the robot's communication style and personality (left) and willingness to reuse and recommend the system (right) across conditions.

in Figure 6A, participants rated the robot's personality to be a significantly higher fit in the personalized condition compared to the non-personalized condition ($p < 0.01$, $r_{\text{rb}} = 0.82$). Users described the robot as friendly, supportive, and engaging, with several noting that it reflected their own traits. P5 stated, "I would describe it as energetic, friendly, and empathetic, which are all very similar to my own personality. I liked these aspects as they were." Others appreciated the its patience and encouragement, with P6 saying "It was very kind and patient, friendly," and P15 saying "I really enjoyed the positivism that the robot displayed, and I am very positive too." Many participants found the robot's personality helpful and motivating, as P19 remarked, "[it] genuinely looked like he actually cared for my responses and feedback."

In contrast, the non-personalized condition was frequently described as robotic, monotonous, and lacking warmth. Participants noted that the robot felt mechanical, with P13 stating, "It was 100% robotic in the second, I felt like I was asking AI about the technique or just googling it, not coaching me on how to do it." Others found it impersonal and disengaging, such as "It felt very robotic, no adaptability or evidence of sympathy or understanding," (P14) and "Doesn't have any trait to define personality, very robotic" (P11). Some also found the interaction cold, with P4 remarking, "It was pretty monotonous and kind of sounded bored. I'm usually high energy, which is very different."

4) Customization of Name and Gender: Allowing students to customize the robot's name and gender added another layer of personalization to the experience. For some, it fostered a deeper connection, with P14 noting, "It made me more emotionally present," and P16 adding, "It developed a more lighthearted personality." Others found that the customization helped them take the robot's guidance more seriously, with P25 explaining, "I gave it my mom's name so I paid more attention to what it said." However, not everyone felt the customization had a significant effect. Some felt it didn't change their experience much, with P7 saying, "I didn't think it had much of an impact." On average, participants rated the ability to customize the robot's name $4.12/5$ and gender $3.96/5$, suggesting that while it enhanced the experience for many, its effect was not universal.

D. System Usability and Robot Presence

The SUS scores indicated good usability in both conditions (78.3 for personalized; 72.6 for non-personalized).

As shown in Figure 7B, participants indicated that the robot's presence contributed more to their coaching experience in the personalized condition ($M = 3.68$) than in the non-personalized condition ($M = 3.28$). Similarly, they rated the robot's role in maintaining engagement higher in the personalized system ($M = 3.88$) compared to the non-personalized system ($M = 3.4$) (see Figure 7A). The differences in both cases were not statistically significant.

Participants in the personalized condition highlighted how the robot's movements and quick responses helped maintain their focus. P15 shared, "My attention was fully engaged by both the quick answers and the movements from the robot," and P2 said, "The movement brought me back to the conversation if I ever felt like daydreaming." Meanwhile, participants in the non-personalized condition still benefited from the robot's presence, but its impact was less pronounced. P15 noted, "It kept my attention, but I wasn't as entertained," and P1 stated, "Movements and expressions could have improved interactivity and naturalness."

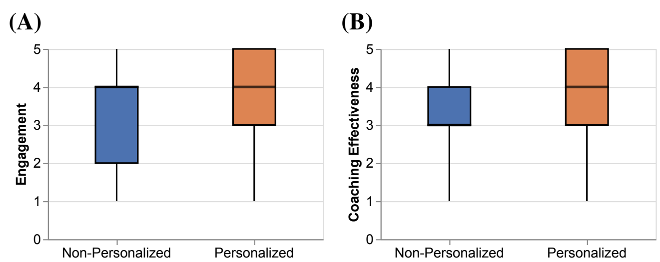

Fig. 7: Perceived role of robot's presence in engagement (left) and coaching effectiveness (right) across conditions.

E. Suggested Improvements and Future Usage

When asked about potential improvements, participants in the personalized condition expressed interest in greater customization, such as changing the chat interface's color, and incorporating newer, less familiar strategies to keep the sessions engaging. Similarly, those in the non-personalized condition emphasized the need for more personalization, suggesting that the system should include relevant examples, remember past responses, and tailor advice accordingly. Many participants also found the non-personalized session too static and information-heavy, preferring simpler language, bullet-point summaries, and interactive teaching methods. Despite these areas for improvement, as shown in Figure 6, participants in the personalized condition were significantly more likely to use the robotic system for ongoing coaching and recommend it to others compared to those in the non-personalized condition ($p < 0.01$). Across both conditions, delays and the robotic tone were common concerns, with participants preferring faster responses, a more natural voice, and the option to speak with the robot instead of typing.

VI. Conclusion

This study introduces adaptive prompting, a prompt engineering technique that enables real-time, personalized language generation without pre-training or fine-tuning. Integrated with QTrobot for productivity coaching, it demonstrates how dynamically adapting to user attributes, preferences, and characteristics can enhance interaction quality and user experience. Although our evaluation focused on a structured coaching scenario, this reflects an intentional design choice: many impactful applications of social robots, such as education, therapy, and decision-support, are likewise structured and offer natural points for adaptation. Our participants were 18–24-year-old students from a single university, which may limit generalizability, but the approach can be applied beyond coaching to other structured HRI settings. Future work will extend this approach to multi-modal inputs and evaluate its effectiveness in longitudinal studies with diverse audiences and scenarios.

References

[1] C. Breazeal, K. Dautenhahn, and T. Kanda, "Social robotics," *Springer handbook of robotics*, pp. 1935–1972, 2016.

[2] F. Tanaka and S. Matsuzoe, "Children teach a care-receiving robot to promote their learning: Field experiments in a classroom for vocabulary learning," *Journal of HRI*, vol. 1, no. 1, pp. 78–95, 2012.

[3] R. Van den Berghe, J. Verhagen, O. Oudgenoeg-Paz, S. Van der Ven, and P. Leseman, "Social robots for language learning: A review," *Review of Educational Research*, vol. 89, no. 2, pp. 259–295, 2019.

[4] D. Cruz-Sandoval, A. Morales-Tellez, E. B. Sandoval, and J. Favela, "A social robot as therapy facilitator in interventions to deal with dementia-related behavioral symptoms," in *HRI*, pp. 161–169, 2020.

[5] S. Costa, H. Lehmann, K. Dautenhahn, B. Robins, and F. Soares, "Using a humanoid robot to elicit body awareness and appropriate physical interaction in children with autism," *IJSR*, vol. 7, pp. 265–278, 2015.

[6] O. Zuckerman, G. Hoffman, D. Kopelman-Rubin, A. B. Klomek, N. Shitrit, Y. Amsalem, and Y. Shlomi, "Kip3: robotic companion as an external cue to students with adhd," in *Int. Conf. on Tangible, Embedded, and Embodied Interaction*, pp. 621–626, 2016.

[7] E. A. Lozano, C. E. Sánchez-Torres, I. H. López-Nava, and J. Favela, "An open framework for nonverbal communication in human-robot interaction," in *Int. Conf. on Ubiquitous Computing and Ambient Intelligence*, pp. 21–32, Springer, 2023.

[8] L. J. Wood, A. Zaraki, B. Robins, and K. Dautenhahn, "Developing kaspar: a humanoid robot for children with autism," *IJSR*, vol. 13, no. 3, pp. 491–508, 2021.

[9] A. O'Connell, A. Banga, J. Ayissi, N. Yaminrafie, E. Ko, A. Le, B. Cislowski, and M. Mataric, "Design and evaluation of a socially assistive robot schoolwork companion for college students with adhd," in *HRI*, pp. 533–541, 2024.

[10] J. Wu, R. Antonova, A. Kan, M. Lepert, A. Zeng, S. Song, J. Bohg, S. Rusinkiewicz, and T. Funkhouser, "Tidybot: Personalized robot assistance with large language models," *Autonomous Robots*, vol. 47, no. 8, pp. 1087–1102, 2023.

[11] C. Di Napoli, G. Ercolano, and S. Rossi, "Personalized home-care support for the elderly: a field experience with a social robot at home," *User Modeling and User-Adapted Interaction*, vol. 33, no. 2, pp. 405–440, 2023.

[12] L. Zhang, C. Zheng, H. Wang, R. Gomez, E. Nichols, and G. Li, "Autonomous storytelling for social robot with human-centered reinforcement learning," in *IROS*, pp. 2450–2456, IEEE, 2024.

[13] H. Wang, L. Zhang, C. Zheng, R. Gomez, K. Nakamura, and G. Li, "Personalized storytelling with social robot haru," in *ICSR*, pp. 439–451, Springer, 2022.

[14] OpenAI, "Gpt-4 technical report," *arXiv:2303.08774*, 2023.

[15] J. Lim, I. Sa, B. MacDonald, and H. S. Ahn, "A sign language recognition system with pepper, lightweight-transformer, and llm," *arXiv:2309.16898*, 2023.

[16] M. Spitale, M. Axelsson, and H. Gunes, "Vita: A multi-modal llm-based system for longitudinal, autonomous, and adaptive robotic mental well-being coaching," *THRI*, 2023.

[17] W. Khoo *et al.*, "Spill the tea: when robot conversation agents support well-being for older adults," in *Companion of HRI*, pp. 178–182, 2023.

[18] X. Feng, Z.-Y. Chen, Y. Qin, Y. Lin, X. Chen, Z. Liu, and J.-R. Wen, "Large language model-based human-agent collaboration for complex task solving," *arXiv:2402.12914*, 2024.

[19] D. Tanneberg, F. Ocker, S. Hasler, J. Deigmoeller, A. Belardinelli, C. Wang, H. Wersing, B. Sendhoff, and M. Gienger, "To help or not to help: Llm-based attentive support for human-robot group interactions," in *IROS*, pp. 9130–9137, IEEE, 2024.

[20] G. J. DuPaul, L. L. Weyandt, S. M. O'Dell, and M. Varejao, "College students with adhd: Current status and future directions," *Journal of attention disorders*, vol. 13, no. 3, pp. 234–250, 2009.

[21] L. L. Weyandt and G. J. DuPaul, "Adhd in college students: Developmental findings," *Developmental disabilities research reviews*, vol. 14, no. 4, pp. 311–319, 2008.

[22] E. Heiligenstein, G. Guenther, A. Levy, F. Savino, and J. Fulwiler, "Psychological and academic functioning in college students with attention deficit hyperactivity disorder," *Journal of American college health*, vol. 47, no. 4, pp. 181–185, 1999.

[23] C. K. Conners *et al.*, *Conners' continuous performance test*. Multi-health systems North Tonawanda NY, 2000.

[24] A. L. Green and D. L. Rabiner, "What do we really know about adhd in college students?," *Neurotherapeutics*, vol. 9, no. 3, pp. 559–568, 2012.

[25] B. Tleubayev, Z. Zhexenova, A. Zhakenova, and A. Sandygulova, "Robot-assisted therapy for children with adhd and asd: A pilot study," in *Int. Conf. on Service Robotics Technologies*, pp. 58–62, 2019.

[26] P. Murali, I. Steenstra, H. S. Yun, A. Shamekhi, and T. Bickmore, "Improving multiparty interactions with a robot using large language models," in *Extended Abstracts of CHI*, pp. 1–8, 2023.

[27] C. Mishra, R. Verdonschot, P. Hagoort, and G. Skantze, "Real-time emotion generation in human-robot dialogue using large language models," *Frontiers in Robotics and AI*, vol. 10, p. 1271610, 2023.

[28] T. Onorati, Á. Castro-González, J. C. del Valle, P. Díaz, and J. C. Castillo, "Creating personalized verbal human-robot interactions using llm with the robot mini," in *Int. Conf. on ubiquitous computing and ambient intelligence*, pp. 148–159, Springer, 2023.

[29] H. Kang, M. B. Moussa, and N. Magnenat-Thalmann, "Nadine: an llm-driven intelligent social robot with affective capabilities and human-like memory," *arXiv:2405.20189*, 2024.

[30] Z. Zhang *et al.*, "Personalization of large language models: A survey," *TMLR*, 2024.

[31] S. Mysore, Z. Lu, M. Wan, L. Yang, S. Menezes, T. Baghaee, E. B. Gonzalez, J. Neville, and T. Safavi, "Pearl: Personalizing large language model writing assistants with generation-calibrated retrievers," *arXiv:2311.09180*, 2023.

[32] Y. Dan *et al.*, "Educhat: A large-scale language model-based chatbot system for intelligent education," *arXiv:2308.02773*, 2023.

[33] X. Tu, J. Zou, W. J. Su, and L. Zhang, "What should data science education do with large language models," *arXiv:2307.02792*, vol. 3, 2023.

[34] M. Park, S. Kim, S. Lee, S. Kwon, and K. Kim, "Empowering personalized learning through a conversation-based tutoring system with student modeling," in *Extended Abstracts of CHI*, pp. 1–10, 2024.

[35] M. Abbasian, I. Azimi, A. M. Rahmani, and R. Jain, "Conversational health agents: A personalized llm-powered agent framework.," *arXiv:2310.02374*, 2023.

[36] R. C. Kessler *et al.*, "The world health organization adult adhd self-report scale (asrs): a short screening scale for use in the general population," *Psychological medicine*, vol. 35, no. 2, pp. 245–256, 2005.

[37] J. E. Strait, P. Dawson, C. A. Walther, G. G. Strait, A. K. Barton, and M. Brunson McClain, "Refinement and psychometric evaluation of the executive skills questionnaire-revised," *Contemporary School Psychology*, vol. 24, pp. 378–388, 2020.

[38] J. Brooke *et al.*, "Sus-a quick and dirty usability scale," *Usability evaluation in industry*, vol. 189, no. 194, pp. 4–7, 1996.

[39] B. Rammstedt and O. P. John, "Measuring personality in one minute or less: A 10-item short version of the big five inventory in english and german," *Journal of research in Personality*, vol. 41, no. 1, pp. 203–212, 2007.