Everyone Deserves a Teacher They Like: Designing Personality-Matched Multimodal Educational Agents in Virtual Reality

Yuchao Zhuo
Division of Integrative Systems and Design
The Hong Kong University of Science and Technology
Hong Kong, Hong Kong
yzhuo332@connect.ust.hk

Yukun Zhao
Department of Computer Science and Engineering
The Hong Kong University of Science and Technology
Hong Kong, Hong Kong
yzhaoeg@connect.ust.hk

Xuanyu Wang
State Key Laboratory for Manufacturing Systems Engineering
Xi'an Jiaotong University
Xi'an, China
XJTU-POLIMI Joint School
Xi'an Jiaotong University
Xi'an, China
xuanyuwang@xjtu.edu.cn

Tristan Braud*
Division of Integrative Systems and Design
The Hong Kong University of Science and Technology
ClearWater Bay, Kowloon, Hong Kong
braudt@ust.hk

Figure 1: We present an automatically personality-matched multimodal educational agent that adapts its verbal and nonverbal behaviors to individual learners. By expressing varying interaction styles based on the learners' personalities, the system aims to improve learners' motivation, engagement, and performance through personalized personality expression during instruction.

*Corresponding author.

CHI EA '26, Barcelona, Spain

ACM ISBN 979-8-4007-2281-3/26/04
https://doi.org/10.1145/3772363.3798875

Abstract

Virtual Reality (VR) provides immersive and interactive learning environments in which virtual humans can act as educational agents, supporting instruction through embodied interaction. Prior research suggests that learners respond to such agents as social actors and that personality plays an important role in shaping motivation and learning experiences. However, existing studies on personality-adaptive educational agents often focus on isolated personality traits of the Five-Factor Model (FFM), rely on single-modality expression, and lack automatic mechanisms for matching agent personalities to individual learners. To address these limitations, we propose a multimodal educational agent system that automatically matches agent personalities to individual learners based on empirical data from user studies. The system expresses different personalities through verbal and nonverbal behaviors, enabling more engaging interactions in learning environments. Our work explores the potential of personality-aware recommendation and multimodal expression to enhance perceived engagement and support personalized learning experiences in virtual educational environments.

CCS Concepts

• **Applied computing** → *Computer-assisted instruction.*

Keywords

Large Language Model (LLM), Five-factor personality, Educational Agent, Virtual Reality (VR), Conversational Agent, Animation

ACM Reference Format:
Yuchao Zhuo, Yukun Zhao, Xuanyu Wang, and Tristan Braud. 2026. Everyone Deserves a Teacher They Like: Designing Personality-Matched Multimodal Educational Agents in Virtual Reality. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3772363.3798875

1 Introduction

Virtual Reality (VR) has attracted growing attention in educational research due to its ability to support immersive and interactive learning experiences that extend beyond traditional screen-based instruction [13]. Virtual reality can enhance learners' motivation and interest as well [11]. Virtual humans represent an important form of the educational agent, which are widely employed to guide learners, deliver instructional content, and scaffold learning activities [1, 3]. From a social and cognitive perspective, studies on computers have indicated that learners tend to perceive and respond to virtual instructors as social actors, with agent behavior shaping users' perceptions and interactions [9, 10].

Through coordinated verbal and nonverbal behaviors, virtual humans can enact diverse roles that support learners' understanding, reasoning, and higher-order thinking. In psychology, personality is strongly associated with how agents are perceived and evaluated during interactions [5]. Previous studies have also demonstrated that the personalities of both students and teachers would influence the learning process and outcomes for the student [2, 8]. Personality matching between students and teachers is associated with learners' motivation, learning strategies, and willingness to communicate. It means that different students will have different preferences for the agent [14]. When Large Language Models (LLMs) are used as tutors, users with different demographic attributes and personality traits have different needs and expectations for the interaction experience [12]. Therefore, a match between learners' and teachers' personalities is an important predictor of improved learning performance.

Existing studies on personality adaptation in educational agents have primarily examined single personality dimensions (e.g., extraversion), rather than evaluating the combined effects of multiple traits within the Five-Factor Model (FFM) [7]. In addition, prior work has largely focused on personality expression through a single modality—most commonly, text-based interaction. While the role of multimodal personality expression in educational agents remains insufficiently explored [6]. Furthermore, even when multimodal educational agents are considered, existing systems typically rely on manually predefined personality settings. There remains a lack of automatic recommendation mechanisms for identifying agent personalities that are well matched with individual learners.

Motivated by these gaps, we introduce a multimodal educational agent system designed to support personality-aware instructional interactions. The system incorporates a recommendation-based approach that leverages empirical evidence from user studies. The approach automatically selects educational agent personalities that are well matched with individual learners, with the aim of supporting effective and personalized learning experiences. The system enables the expression of virtual human personality through synchronized verbal and nonverbal behaviors, allowing personality traits to be consistently conveyed during instructional interactions in immersive environments.

2 System Design

2.1 Data Collection and Experimental Procedure

Establishing such a mapping mechanism for personality matching is crucial for enabling reliable and scalable personalization in educational agent systems to enhance learning outcomes. We adopt a Multilayer Perceptron (MLP) as the core model for personality matching due to its ability to capture nonlinear relationships between learner characteristics and agent personality preferences. To train a preference-based MLP, we first conducted a controlled user study to collect learner–agent interaction data. The model underpins the design of the personality matching mechanism for automatically selecting suitable educational agent personalities. All participants were undergraduate students recruited from the local campus. A total of fifteen students (S1–S15) participated in the study, including seven males and eight females, with ages ranging from 18 to 20 years (M = 18.7, SD = 0.7). The study was approved by the Institutional Review Board (IRB) of the Hong Kong University of Science and Technology (HKUST). Informed consent was obtained from all participants before the study. The experimental setup was based on a workstation (i9-13980HX CPU, 32 GB RAM) with the Meta Quest Pro headset. At the beginning of the session, participants provided demographic information and completed the BFI-2-XS questionnaire [15] to assess their personality profiles, which were used as training labels to measure the learning outcome across different personality conditions of agents. Participants then donned a VR headset and experienced three instructional contexts in a randomized order, with questionnaires administered between sessions.

Previous study demonstrated that users were able to reliably distinguish among eight personality configurations along the dimensions of Openness, Extraversion, and Agreeableness, and that conditions A, B, and C were rated significantly higher than the other configurations. Specifically, three personality conditions were designed: A (high O, low E, high A), B (low O, high E, high A), and C (high O, high E, high A). Accordingly, these three conditions were selected as candidate agents for further investigation of user–agent personality matching. Within each instructional context, participants experienced all three agent personality conditions (A, B, and C). The agent's verbal expressions and nonverbal behaviors were systematically manipulated to embody the corresponding personality condition, which was demonstrated in Figure 2. Across the three conditions, agent personalities were instantiated through systematic variations in verbal expressiveness (e.g., tone and rhetorical style) and nonverbal behaviors (e.g., gesture amplitude and motion dynamics), with Condition A being relatively restrained, Condition B emphasizing energetic expressiveness, and Condition C integrating both high expressiveness and supportive social cues.

After each context, participants rated their perceived agent personality and learning experience using the LOES-S scale [4]. The resulting learner-agent interaction data were used to support the design and instantiation of the preference-based personality matching mechanism described in the following section.

2.2 Personality Matching Mechanism

To adapt avatar personalities to individual learners, the system employs a preference-based personality matching mechanism. The design is motivated by the fundamental challenge that an 'optimal' avatar personality cannot be directly observed or measured. In addition, learner outcomes and subjective experiences can only be obtained through interactions with a finite set of predefined avatar personality configurations. This motivates a learning-based formulation that reframes personality matching as a utility maximization problem rather than direct personality estimation. Accordingly, the system models learner-agent personality matching as a utility maximization problem, rather than regressing an abstract optimal personality parameter.

Learner personality is represented as a three-dimensional vector

$$\mathbf{p}_u = (O_u, E_u, A_u), \tag{1}$$

where O, E, and A denote Openness, Extraversion, and Agreeableness, respectively, normalized to the range $[0, 1]$. Similarly, each avatar personality is represented as

$$\mathbf{p}_a = (O_a, E_a, A_a). \tag{2}$$

For a given learner-avatar pairing, the system estimates a scalar learning utility that reflects the expected quality of the learning experience under that avatar's personality. Internally, this estimation is implemented through a parametric utility function

$$\hat{y} = g_\theta(\mathbf{p}_u, \mathbf{p}_a), \quad \hat{y} \in [0, 1]. \tag{3}$$

that jointly considers learner and avatar personality traits. $g_\theta(\cdot)$ is implemented as an MLP to predict learning utility from combined learner–avatar personality features. The input to the function is a six-dimensional vector,

$$\mathbf{x} = [O_u, E_u, A_u, O_a, E_a, A_a], \tag{4}$$

and the output is a predicted utility value $\hat{y} \in [0, 1]$.

Rather than predicting an idealized personality target, this formulation captures and learns how different users respond to different avatar personalities based on observed interactions. The learning utility integrates three complementary aspects of the learning experience: learning effectiveness, perceived quality, and engagement. All scores are linearly normalized to the range $[0, 1]$, and combined as

$$y = w_l * y_{\text{learning}} + w_q * y_{\text{quality}} + w_e * y_{\text{engagement}}. \tag{5}$$

In this study, the three hyperparameters are set to be equal ($w_l = w_q = w_e = \frac{1}{3}$) to avoid biasing the utility estimation towards any single subjective dimension.

The utility model is trained using all observed learner–avatar interaction records, where each instance corresponds to a triplet $(\mathbf{p}_u, \mathbf{p}_a, y)$ consisting of the learner personality, the avatar personality, and the aggregated learning utility. The model employs the loss function as minimizing the mean squared error (MSE) between the predicted utility and the observed utility:

$$\mathcal{L} = \|g_\theta(\mathbf{p}_u, \mathbf{p}_a) - y\|^2. \tag{6}$$

Training is performed using the AdamW optimizer with random mini-batches, and early stopping is applied to prevent overfitting. Regularization techniques, including dropout and weight decay, are employed to improve generalization and prevent overfitting. This formulation avoids explicitly regressing toward an unobservable "optimal" avatar personality, allowing the model to directly learn personality matching relationships from preference-based interaction signals.

To select the most suitable personality for each new learner's personality $\mathbf{p}_u$, the system estimates the utility of each candidate avatar individually. It computes the predicted utility for each candidate avatar $\mathbf{p}_{a_i}$ drawn from a predefined personality set (i.e., Conditions A, B, and C) and selects the avatar that maximizes the estimated utility:

$$a^* = \arg\max_{a_i \in \mathcal{A}} g_\theta(\mathbf{p}_u, \mathbf{p}_{a_i}). \tag{7}$$

Detailing the experiment, the utility function $g_\theta(\cdot)$ is instantiated as a preference-based MLP with a learning rate of 0.001 trained on empirical data collected from a controlled user study. The dataset was split into training and test sets by uniformly random sampling at the ratio of 8:2. The MLP was trained for up to 400 epochs, and the model achieved a low validation error (training MSE = 0.0111), indicating that the MLP was able to reliably learn user preference patterns for personality matching. In evaluating the results on the test set, model performance was assessed using MSE between the predicted utility values and the ground-truth preference scores. The test MSE was 0.0106, providing preliminary evidence that the model captured learner–agent personality interaction patterns within this dataset.

By grounding personality matching in observed learner preferences and restricting selection to predefined personality configurations, the mechanism provides a transparent and stable basis for adaptive personality selection under limited data conditions.

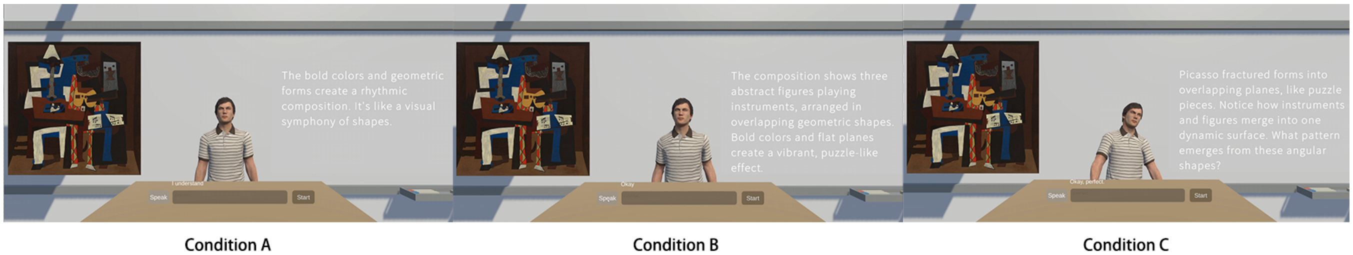

Figure 2: The different agent personalities are presented in the art context. From left to right, the conditions correspond to A (high O, low E, high A), B (low O, high E, high A), and C (high O, high E, high A).

2.3 Multimodal Educational Agent Design

Following the personality matching mechanism described above, the system first determines the avatar personality among conditions A, B, and C that best aligns with an individual learner. Once a target avatar personality is selected, this personality profile is treated as a fixed control signal that conditions the agent's multimodal expression throughout the instructional interaction.

The selected personality does not alter the pedagogical structure or instructional flow. Instead, it modulates how the educational agent communicates, including its verbal style, body movement, and facial expression. To achieve this behavior, the system utilizes a multimodal agent architecture comprising coordinated behavioral, affective, and language generation components, which collectively translate abstract personality traits into coherent and human-like embodied expressions in VR.

The proposed system is implemented as a modular multimodal educational agent composed of three tightly coupled components: a behavioral modulation layer, an affective modulation layer, and a language generation layer. Together, these components enable the agent to produce temporally aligned verbal and nonverbal behaviors that consistently reflect a given personality profile during instructional interactions in VR.

The behavioral modulation layer governs full-body motion and gesture generation. Personality traits defined in the OCEAN model are mapped to Laban Movement Analysis (LMA) effort parameters, which control spatial extent, motion dynamics, and rhythm. These parameters are continuously updated and applied to the Unity Animator, allowing personality traits to directly influence gesture amplitude, pacing, and expressiveness during instruction.

The affective modulation layer controls facial expression and emotional continuity. Based on both the agent's personality traits and the emotional intent of the dialogue, the system selects from a set of prototypical facial expressions and applies temporal smoothing and cross-fading to ensure natural transitions. This layer maintains emotional coherence across speech segments and nonverbal actions, enabling stable yet expressive affective behavior over time.

The language generation layer serves as the cognitive core of the agent. A large language model generates instructional utterances conditioned on personality-specific prompting rules, which define linguistic tone, verbosity, politeness, and emotional intensity. Each generated response is accompanied by structured metadata specifying the dominant emotion and an associated animation cue. These outputs are parsed by the Unity runtime and used to synchronize speech, facial expression, and gesture execution. To maintain interaction continuity, language generation is processed asynchronously and buffered prior to animation playback, enabling nonverbal behavior synthesis to proceed without blocking the interaction flow.

A predefined teaching script orchestrates the interaction flow. The script specifies when the agent delivers explanations, poses questions, provides feedback, or transitions between instructional phases. Importantly, the script controls interaction timing and structure, while personality-related variations are introduced exclusively through the multimodal modulation layers. This separation ensures that differences in learner experience arise from personality expression rather than from changes in instructional structure. Through the coordination of these components, the system realizes a coherent multimodal educational agent in which language, gesture, and facial expression are jointly shaped by personality traits and executed in real time within a VR environment.

3 Limitations and Future Work

This work demonstrates the feasibility of personality-aware recommendation for embodied VR educational agents. Several considerations warrant discussion. First, the study employed subjective learning utility measures (LOES-S), capturing perceived engagement and instructional quality rather than objective knowledge gains. Second, the dataset was modest in scale, and the MLP component is best understood as an exploratory modeling effort illustrating potential computational matching rather than a fully validated predictive solution. Future research may extend this framework through larger samples, objective assessments, and more explainable modeling strategies.

4 Conclusion

In this paper, we proposed a personality-aware educational agent system aimed at improving learning performance by matching learners with educational agents exhibiting suitable personality profiles in immersive learning environments. To support personality matching, we collected empirical data through user studies that captured self-assessments of learning, quality, and engagement with educational agents expressing different personality traits. Furthermore, we integrated the recommendation mechanism with a multimodal educational agent capable of expressing personality through coordinated verbal and nonverbal behaviors. Together, the recommendation-based personality matching and multimodal personality expression form a unified framework for personalized and personality-aware educational agents. Future work may extend this framework by incorporating additional personality dimensions, broader learning contexts, and longitudinal evaluations to explore the impact of personality-aware adaptation on learning outcomes.

Acknowledgments

This work was supported by grants from the National Natural Science Foundation of China (No. 62502375).

References

[1] Sabarish Babu, Evan Suma, Tiffany Barnes, and Larry F Hodges. 2007. Can immersive virtual humans teach social conversational protocols?. In *2007 IEEE Virtual Reality Conference*. IEEE, 215–218.

[2] Adrian Furnham and Tomas Chamorro-Premuzic. 2005. Individual differences in students' preferences for lecturers' personalities. *Journal of Individual Differences* 26, 4 (2005), 176–184.

[3] Yan Ru Guo and Dion Hoe-Lian Goh. 2015. Affect in embodied pedagogical agents: Meta-analytic review. *Journal of Educational Computing Research* 53, 1 (2015), 124–149.

[4] Robin H Kay and Liesel Knaack. 2009. Assessing learning, quality and engagement in learning objects: the Learning Object Evaluation Scale for Students (LOES-S). *Educational Technology Research and Development* 57, 2 (2009), 147–168.

[5] Mark R Leary and Ashley Batts Allen. 2011. Personality and persona: Personality processes in self-presentation. *Journal of personality* 79, 6 (2011), 1191–1218.

[6] Hai Li, Wanli Xing, Chenglu Li, Wangda Zhu, Bailing Lyu, Fan Zhang, and Zifeng Liu. 2025. Who should be my tutor? Analyzing the interactive effects of automated text personality styles between middle school students and a mathematics chatbot. In *Proceedings of the 15th international learning analytics and knowledge conference*. 910–917.

[7] Tze Wei Liew and Su-Mae Tan. 2016. Virtual agents with personality: Adaptation of learner-agent personality in a virtual learning environment. In *2016 Eleventh International Conference on Digital Information Management (ICDIM)*. IEEE, 157–162.

[8] Harry G Murray, J Philippe Rushton, and Sampo V Paunonen. 1990. Teacher personality traits and student instructional ratings in six types of university courses. *Journal of educational psychology* 82, 2 (1990), 250.

[9] Clifford Nass and Youngme Moon. 2000. Machines and mindlessness: Social responses to computers. *Journal of social issues* 56, 1 (2000), 81–103.

[10] Kristine L Nowak and Frank Biocca. 2003. The effect of the agency and anthropomorphism on users' sense of telepresence, copresence, and social presence in virtual environments. *Presence: Teleoperators & Virtual Environments* 12, 5 (2003), 481–494.

[11] Jocelyn Parong and Richard E Mayer. 2018. Learning science in immersive virtual reality. *Journal of educational psychology* 110, 6 (2018), 785.

[12] Heinrich Peters, Moran Cerf, and Sandra C Matz. 2024. Large language models can infer personality from free-form user interactions. *arXiv preprint arXiv:2405.13052* (2024).

[13] Jaziar Radianti, Tim A Majchrzak, Jennifer Fromm, and Isabell Wohlgenannt. 2020. A systematic review of immersive virtual reality applications for higher education: Design elements, lessons learned, and research agenda. *Computers & education* 147 (2020), 103778.

[14] Pia Rosander, Martin Bäckström, and Georg Stenberg. 2011. Personality traits and general intelligence as predictors of academic performance: A structural equation modelling approach. *Learning and individual differences* 21, 5 (2011), 590–596.

[15] Christopher J Soto and Oliver P John. 2017. The next Big Five Inventory (BFI-2): Developing and assessing a hierarchical model with 15 facets to enhance bandwidth, fidelity, and predictive power. *Journal of personality and social psychology* 113, 1 (2017), 117.