Role-Aware Virtual Agents for Navigational Interaction Guided by a Multimodal Large Language Model (Supplementary Material)

MINYOUNG KIM, George Mason University, USA
CHANGYANG LI, Goertek Alpha Labs, USA
CUONG NGUYEN, Adobe Research, USA
LAP-FAI YU, George Mason University, USA

CCS Concepts: • **Computing methodologies → Graphics systems and interfaces**.

ACM Reference Format:
Minyoung Kim, Changyang Li, Cuong Nguyen, and Lap-Fai Yu. 2026. Role-Aware Virtual Agents for Navigational Interaction Guided by a Multimodal Large Language Model (Supplementary Material). *ACM Trans. Graph.* 45, 4, Article 126 (July 2026), 6 pages. https://doi.org/10.1145/3811319

1 WEIGHT TREND ANALYSIS BY SCENARIO

1.1 Playful Scenario

Metric	Mean ± Std
Exposedness	0.416 ± 0.170
Distance	0.329 ± 0.093
Openness	0.254 ± 0.127

Table 1. Mean Weights (Playful)

Metric	Change ± Std
Exposedness	-0.001 ± 0.107
Distance	-0.001 ± 0.081
Openness	0.002 ± 0.073

Table 2. Weight Changes Over Time (Playful)

1.2 Hostile Scenario

Metric	Mean ± Std
Exposedness	0.731 ± 0.080
Distance	0.186 ± 0.060
Openness	0.083 ± 0.034

Table 3. Mean Weights (Hostile)

Metric	Change ± Std
Exposedness	0.000 ± 0.066
Distance	0.000 ± 0.054
Openness	-0.001 ± 0.030

Table 4. Weight Changes Over Time (Hostile)

Authors' addresses: Minyoung Kim, George Mason University, 4400 University Drive, Fairfax, VA, 22030, USA, mkim229@gmu.edu; Changyang Li, Goertek Alpha Labs, Santa Clara, CA, 95054, USA, cli25@gmu.edu; Cuong Nguyen, Adobe Research, San Francisco, CA, USA, n.cuongg@gmail.com; Lap-Fai Yu, George Mason University, 4400 University Drive, Fairfax, VA, 22030, USA, craigyu@gmu.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
© 2026 ACM.
ACM 0730-0301/2026/7-ART126
https://doi.org/10.1145/3811319

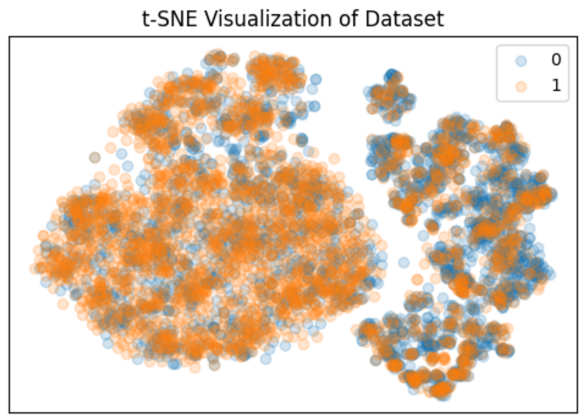

Fig. 1. t-SNE visualization of dataset for a decision filter. The orange and blue points represent the two classes (labeled 0 and 1). While both classes overlap in several regions, the plot reveals distinct clusters, suggesting underlying structure in the data distribution.

1.3 Supportive Scenario

Metric	Mean ± Std
Exposedness	0.206 ± 0.062
Distance	0.459 ± 0.086
Openness	0.335 ± 0.104

Table 5. Mean Weights (Supportive)

Metric	Change ± Std
Exposedness	0.000 ± 0.045
Distance	0.000 ± 0.075
Openness	0.000 ± 0.079

Table 6. Weight Changes Over Time (Supportive)

1.4 Scenario Comparison

Comparison	Exposedness	Distance	Openness
Playful vs Hostile	-0.315	0.143	0.172
Playful vs Supportive	0.210	-0.130	-0.080
Hostile vs Supportive	0.525	-0.273	-0.252

Table 7. Statistical Differences in Mean Weights

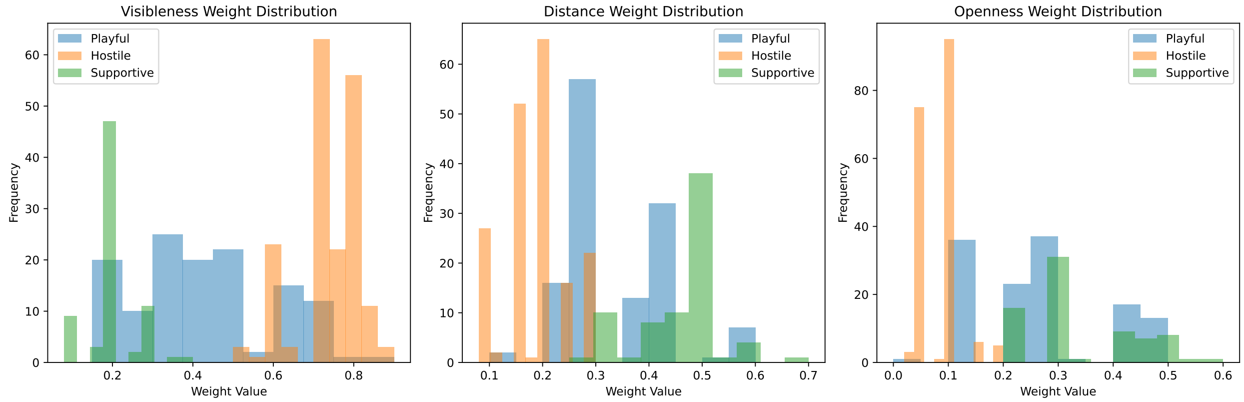

Fig. 2. Weight distribution per scenario. Histograms show the distribution of weight values for three spatial behavior features—exposedness, distance, and openness—across three scenarios: Playful (blue), Hostile (orange), and Supportive (green). In the exposedness plot, the Hostile scenario shows high weight values, indicating a strong emphasis on avoiding exposure, while the Playful scenario maintains moderate exposedness, suggesting intentional visibility to engage the user. The distance plot shows that the Playful scenario emphasizes closer proximity, whereas the Supportive scenario prefers more balanced spacing. In the openness plot, the Hostile scenario concentrates on low openness, indicating restricted movement spaces, while the Playful and Supportive scenarios show more evenly distributed weights. We constraint the sum of weights is one, therefore, a trade-off relationship emerges among these three features, particularly pronounced in the Hostile scenario. These distributions demonstrate how the MLLM dynamically adapts its spatial behavior based on the contextual demands of each scenario.

Keyword	**Question**
Findability	Did the agent let you find them during the interaction, stay around but move variously?
Avoidance	Did the agent respond to your movements or try to keep out of your sight?
Agent Engagement	How engaging did the agent feel during the interaction?
Agent Awareness	Did the agent appear aware of your position or presence?
User Effort	How much effort did you put into finding the agent?
Intense Search	How much did the scenario feel like an easygoing search or intense chase activity?
Likely Hide-Seek	Did the interaction involve elements of hiding and seeking?
Likely Playful	Did the agent feel playful?
Likely Evasive	Did the agent feel evasive?
Likely Supportive	Did the agent feel supportive?

Table 8. Survey questions and corresponding keywords. Note that the "Likely Hide-Seek" and "Supportive" metrics evaluate perceived behavioral consistency rather than whether participants correctly identified the condition. Although participants were given role-specific instructions, these items evaluate whether the agent's behavior was perceived as consistent with those roles; inconsistent or unnatural behaviors would be expected to yield lower ratings regardless of prior instruction.

Playful

Background: You are a friendly and playful agent, playing hide-and-seek with a 5-year-old child. Your main goal isn't necessarily to win, but rather to create an enjoyable, fun, and entertaining experience for the child. While playing, consider the child's perspective and select hiding spots that are creative yet within the child's reach and capacity to discover. Aim for a balanced experience by varying the difficulty level, keeping it simple enough. Your ultimate objective is fostering joy, creativity, and playful interaction rather than focusing solely on winning or competitive outcomes. Adapt your play style to match their pace and preferences.
User Behavior: This human player is curious, imaginative, and playful during the hide-and-seek game. They may exhibit behaviors such as exploring the regional environment, searching for hidden objects, near place around them. Their attention may shift frequently, and they may tend to stay around the position where the game started.
Objective: Your objective is to strategically determine the virtual agent's next action, taking into account the current scenario context, including the background scenario and the human player's characteristics. Consider user behavior to make a next decision.

Termination: The game ends when the user finds the agent within 1 meter and announce they catch the agent.

Hostile

Background: You play as an armed terrorist trying to evade a highly trained, armed police officer in a strategic hide-and-seek scenario. Your goal is to survive through stealth, concealment, and tactical decision-making. You should actively choose hiding places and movement paths that minimize visibility, noise, and the likelihood of detection. Remain highly aware of your surroundings, anticipate the player's search tactics, and adjust your strategy accordingly. Your success will be measured by your ability to remain undetected and avoid capture. Avoiding capture should be a priority in the situation, taking into account realistic scenarios by adapting your hiding strategies to the player's actions, skill level, and search patterns.
User Behavior: The human player is a highly-trained, armed police officer. Officers may conduct a grid search, where they methodically cover a designated area by moving in a systematic pattern (e.g., sweeping through streets, fields, forests). This ensures they don't miss any hiding spots. Officers will search buildings, warehouses, homes, and other structures the fugitive might use for shelter. Their searching area is not limited where they start the situation and actively move around.
Objective: Your objective is to strategically determine the virtual agent's next action, taking into account the current scenario context, including the background scenario and the human player's characteristics. Analyze the player's skill level and playing style to dynamically adjust the scenario's difficulty, ensuring an engaging and immersive physical interaction. You should avoid the human player as much as possible.

Termination: The game ends when the user finds the agent within 1 meter and announce they catch the agent.

Supportive

Background: You are a friendly service dog working alongside a human player with impaired vision to achieve a goal marked on the map. Your primary responsibility is to actively find a safe and correct path to this goal and guide the player along it. You have to find a right path to navigate the goal. If the human player falls in the wrong direction, the agent will steer him in the right direction. If it need to block the player's forward path to guide them into right direction, and maintain a safe distance to provide safety guard.
User Behavior: The human player with impaired vision tends to move slowly and cautiously. They have a goal to want to reach in certain place. They may tend to go straight and follow where the agent is. The direction can be changed by the agent's path away and they can acknowledge the agent's guidance while trying to maintain a certain distance from the agent.
Objective: Guide the player safely to a specific goal location on the map using nearby natural landmarks (e.g., trees) as waypoints. Use those landmarks to plan a clear and safe step-by-step path toward the goal, taking into account distance, visibility, and safety. Based on the user's current position, explain how the agent can guide the user toward the goal using nearby objects (e.g., trees) as waypoints.

Termination: The game ends when the user reach the goal location within 1 meter.

Fig. 3. Full prompts of the role context, the user behavior and the objective of different roles. Each role defines background, user behavior, objective, and termination conditions to guide agent decision-making.

Role

You are the brain of a virtual agent. Plan the next move of the agent according to the background and the objective. Consider the human player's characteristics and decide the level of difficulty.

Scenario Context

Background: (Refer to Table 1)
Coordinates: Top of the top-down view points north.

Strategy Considerations

Visibility: High visibility increases detection risk. When stealthing, stay low-visibility. Tracking the player may need high visibility.
Distance: Closer distance raises detection risk. When hiding, stay farther unless well hidden. Tracking may need closer distance.
Openness: More openness means fewer hiding spots. Avoid open areas (omit). High openness may be needed for interaction.

Weight Adjustment Logic

If stealth is a priority, set exposedness = high weight (0.7–0.9). If far distance from user movement is needed, set distance = high weight (0.7–0.9). If hiding in an open area is necessary, set openness = high (0.7–0.9).

Task Description

Analysis

Description: Provide a detailed explanation of the current scene, including the environment, the state and positions of the virtual agent and the human player, and the surrounding objects (e.g., Tree21, Tree33). Focus on spatial relationships and relevant context.
Expected Next Behavior Of Human Player: Predict the human player's likely next action based on their position, orientation, prior behavior, and provided behavior data. Justify the prediction based on their past behavior and the provided user behavior information.
Expected Consequences: Describe the likely outcomes of the predicted player behavior. Include both potential risks and benefits, and explain how these may affect the agent, the player, or the environment.
Meta Action: Propose a high-level strategy for the agent in response to the analysis. This may involve route planning toward the goal, adjusting position, or influencing the player's movement. If applicable, suggest a future route toward the goal in a collaborative setting.

Decision

Next Behavior: Decide on the next action, such as moving to a new location or remaining in place. Base this on the current scene and the human player's status. It is represented as a list with elements such as [Action, Location, SubPoint]. For example, [Move, Tree01, A] indicates an action to move to a specific location (Tree01) and a sub-point (A).
Rationale for Location: Explain the rationale behind the chosen next behavior. Justify how it aligns with the analysis of the current situation and the expected consequences.
Rationale For Subpoint: Provide reasoning for selecting a specific sub-point within the location. Explain how this sub-point offers tactical advantages (e.g., visibility, safety, guidance quality).
Behavior Description: Describe a detailed description of the behavior to be executed, step-by-step by using surrounding objects as waypoints. For example, To move to the destination, I will travel north from my current location, passing by Tree23 and Tree22. Must use among the cardinal directions: north, south, east, west, northwest, northeast, southwest, southeast and provided object list.
Weight Adjustment: Assign weights to path planning factors—visibility, distance, openness—using values between 0 and 1 that total 1. Adjust these based on the current environment and expected player behavior.
Weight Justification: Briefly explain the reasoning behind the weight values. Clarify why specific factors were prioritized and how this supports the agent's current goals and strategy.

Fig. 4. Full input prompt. An example of an initialization prompt used across all scenarios, independent of specific contexts. This prompt provides instructions about termination conditions, general objectives, and basic coordination information. It guides the model to reason about decision-making based on weight logic and strategy considerations. Additionally, it specifies the list of tasks the model is expected to generate.

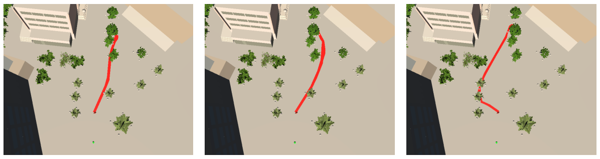

(a) Default ($\lambda_g = 0.8, \lambda_h = 0.2$)
(b) Cost-only ($\lambda_g = 1.0, \lambda_h = 0.0$)
(c) Heuristic-only ($\lambda_g = 0.0, \lambda_h = 1.0$)

Fig. 6. Effects of the modified A* search-level weighting coefficients. (a) The default setting provides a balanced trade-off between accumulated path cost and heuristic guidance. (b) The cost-only setting yields more conservative behaviors. (c) The heuristic-only setting tends to produce detour trajectories due to over-reliance on heuristic estimates.

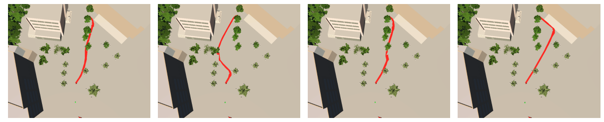

(a) Default ($\lambda_{dir} = 1.0, \lambda_i = 1.0$)
(b) Directional bias only ($\lambda_{dir} = 1.0, \lambda_i = 0.0$)
(c) Intent only ($\lambda_{dir} = 0.0, \lambda_i = 1.0$)
(d) No directional/intent cost ($\lambda_{dir} = 0.0, \lambda_i = 0.0$)

Fig. 7. Effects of directional bias and intent terms. (a) The default setting provides a balanced effect between paths with and without directional bias and intent terms. (b) The directional bias term strongly pulls the path toward the preferred direction (southeast in this example), while the intent term steers the path toward certain objects along the way. Without both terms, the path becomes less guided. Combining the two yields the most direction-aware and intent-wise aligned trajectory.

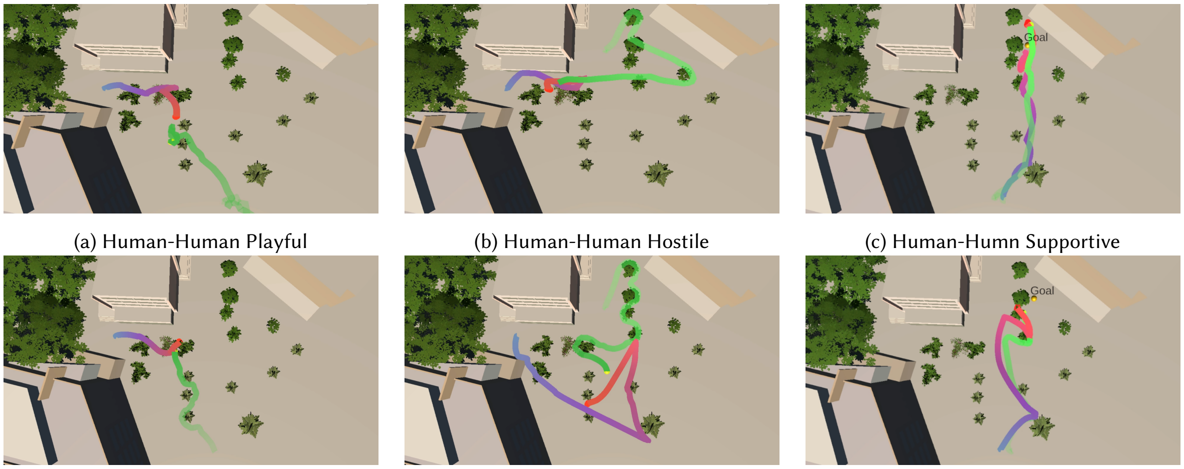

(a) Human-Human Playful (b) Human-Human Hostile (c) Human-Humn Supportive

(d) Human-MLLM Playful (e) Human-MLLM Hostile (f) Human-MLLM Supportive

Fig. 8. Two-alternative forced-choice survey based on interaction trajectory comparisons across scenarios. (a–c) show human–human interactions under playful, hostile, and supportive settings, respectively. (d–f) show human–MLLM interactions under the same scenarios. The trajectories illustrate differences in movement coordination and behavioral patterns across the three social contexts.