

Integrating Environmental Awareness Into NPCs: Contextual Conversational Interaction in Games

Grega Radež, Ciril Bohak*

Faculty of Computer and Information Science, University of Ljubljana, Večna pot 113, 1000 Ljubljana, Slovenia

Abstract

We present a novel system that enhances non-playable characters (NPCs) utilizing a large language model (LLM) in games with contextual awareness of their surroundings, offering dynamic, environment-sensitive interactions. Traditionally, NPC rely on pre-scripted dialogue and lack awareness of their environment, limiting their responsiveness to player actions. Our system addresses this by capturing panoramic images of the NPC's surroundings and applying semantic segmentation to identify objects and their spatial positions. We generate a structured JSON representation of the NPC's environment by combining object locations with segmentation information. This data is provided as context to a LLM, enabling NPCs to incorporate spatial knowledge into their conversations with players. The result is more immersive gameplay, where NPC can reference nearby objects, landmarks, and environmental features during interactions, enhancing believability and engagement. This paper discusses the technical implementation of our system, demonstrating how integrating visual perception into NPCs can transform in-game dialogues and interactions.

Keywords

context-aware NPCs, semantic segmentation, immersive gameplay, environmental perception, dynamic interaction systems

1. Introduction

Non-playable characters (NPC's) in video games play a central role in player engagement and narrative immersion. However, despite advancements in artificial intelligence (AI) and game design, NPC interactions often remain limited to pre-scripted dialogue trees and predictable behavior patterns. These static models of interaction, while effective in certain contexts, fail to fully leverage the dynamic environments that modern game worlds offer. This lack of environmental awareness can disrupt immersion, as players encounter NPCs that seem oblivious to their surroundings [1]. While NPC AI has seen significant improvements, particularly in terms of adaptive behavior [2, 3], NPCs typically respond only to player actions rather than the context of the game world itself. In an era where game environments are increasingly realistic and complex, there is a need for NPCs that can meaningfully integrate information from their surroundings into their interactions. Players expect more from NPCs: the ability to not only respond to players but also acknowledge and react to the world around them.

HCI SI 2024: Human-Computer Interaction Slovenia 2024, November 8th, 2024, Ljubljana, Slovenia

*Corresponding author.

✉ gr1561@student.uni-lj.si (G. Radež); ciril.bohak@fri.uni-lj.si (C. Bohak)

🌐 <https://lgm.fri.uni-lj.si/ciril> (C. Bohak)

🆔 0009-0006-5558-1860 (G. Radež); 0000-0002-9015-2897 (C. Bohak)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

In this paper, we present a novel system for context-aware NPCs with the following contributions:

- **Neighborhood recognition using panoramic images with semantic segmentation**, enabling the system to capture detailed environmental data, process it using computer vision techniques, and translate it into a format understandable by LLM.
- **Creation of context-aware NPCs**, capable of referring to nearby objects, landmarks, and spatial relationships in their conversations with players, introducing dynamic, environment-sensitive interactions that move beyond static dialogue trees or simple LLM chatbot interaction.
- **Bridging the gap between game world realism and NPC behavior**, allowing NPCs to “see” and interpret their surroundings, providing a foundation for more natural, context-sensitive dialogue systems.

In the remainder of the paper, we first present related work in Section 2, followed by the presentation of our method in Section 3. Next, we present the results in Section 4 and discussion in Section 5. In Section 6, we present the conclusions and the possible future work.

2. Related Work

The challenge of enhancing NPCs with dynamic environmental awareness has long been a focal point in the fields of the game AI and interactive narrative systems. While many approaches have explored improving NPC behavior through adaptive AI and procedural dialogue, few have incorporated real-time environmental perception. This section reviews the major areas of related research, including traditional NPC dialogue systems, adaptive AI, environmental perception in game design, and relevant computer vision techniques.

2.1. NPC Dialogue Systems

Traditional NPC dialogue systems rely heavily on pre-defined branching dialogue trees, which allow limited variability based on player choices. Early works in narrative generation, such as Black and Wilensky’s story understanding model [4], focused on dynamic interaction via computational stories. However, these systems were purely narrative-driven and devoid of environmental perception, as characters were unaware of their physical surroundings.

One of the most influential projects in dynamic NPC interaction was *Façade*, developed by Mateas and Stern [5], which allowed NPCs to respond to player actions within an emergent narrative. Although the NPCs in *Façade* exhibited more flexible dialogue, their awareness was limited to the narrative’s internal logic, and they were disconnected from the physical environment around them. Feng *et al.* [6] took steps towards autonomous behavior learning of NPC, but these systems still did not consider current environmental cues.

Feng [7] investigates how behavior models of NPCs may be learned by mimicking behavior patterns of other players, how their behavior models may be adapted through interaction and feedback in a dynamic environment and how emotion of NPCs may be modeled and integrated with the behavior system and to create variations of NPCs. He still has not investigated the element of environmental context.

Basart [8] proposes an alternative approach to interactive narratives that emphasizes thematic coherence and responsiveness over traditional structural models like the Aristotelian framework. His work focuses on dynamically adapting narratives to player input, which parallels our system’s aim of enhancing NPCs through real-time environmental awareness, enabling more immersive and responsive interactions.

2.2. Adaptive NPCs and Game AI

Laird’s work on anticipatory NPCs [3] demonstrated how NPCs could predict player actions and adjust their responses, marking a significant shift in NPC behavior design. However, while these NPCs adapted to player behaviors, they were not designed to perceive or understand their surroundings. The adaptation was based on game logic and internal state modeling, limiting the NPC’s potential to reference and react to their physical environment in a meaningful way.

Yannakakis and Togelius [9] examined how procedural content generation (PCG) could shape game environments and NPC behavior based on player interaction. This approach helped create more flexible and varied interactions between NPCs and players, but the system still did not incorporate current environmental perception into the NPC’s decision-making processes. Similarly, work by Togelius *et al.* [10] explored search-based PCGs for generating environments based on player preferences but lacked focus on NPCs perception of these environments.

Research by Csepregi [11] demonstrates the potential of context-aware LLM-based NPCs in role-playing games to enhance player engagement through dynamic, context-aware dialogues, improving conversation quality and immersion. Similarly, Vidrih and Mayahi [12] highlight how generative AI is used in marketing by companies like Google and Netflix to create personalized and immersive consumer experiences. Both studies align with our work, where LLM-driven NPCs equipped with real-time environmental awareness enhance game immersion through dynamic, personalized interactions.

2.3. Environmental Awareness in Game AI

Concurrent environmental perception has been a significant area of development in autonomous robotics and AI systems. Thrun *et al.* [13] demonstrated how AI systems could use sensors to navigate complex environments in autonomous vehicles, showcasing the potential of real-time awareness. However, while these techniques have found success in robotics, their application to NPCs in gaming is limited, particularly in integrating such perception into conversational and interactive behavior.

Long *et al.* [14] pioneered the use of fully convolutional network (FCN) for semantic segmentation, allowing AI systems to understand and classify every pixel in a visual scene. This kind of fine-grained visual understanding could be directly applicable to NPCs in gaming, enabling them to “see” objects and spatial relationships in their environment. Kirilov *et al.* [15] extended these ideas with panoptic segmentation, which combines both semantic and instance segmentation to provide a richer understanding of scenes, making it possible for NPCs to identify not only objects but also their spatial arrangement.

Recent advancements in visual navigation have focused on enabling agents to map novel environments and leverage spatial reasoning to efficiently reach goals. Notably, the work by

Marza *et al.* [16] introduces reinforcement learning (RL) agents equipped with auxiliary tasks that enhance spatial perception, improving their ability to understand spatial relationships and object characteristics. By incorporating these tasks, the system significantly boosts navigation performance, even rivaling agents with ground-truth maps. This focus on spatial reasoning and perception aligns with our goal of enhancing NPCs with real-time environmental awareness, allowing them to interpret and interact with their surroundings more effectively.

2.4. Context-Aware NPC Interaction

The idea of context-aware computing was introduced by Schilit *et al.* [17], who explored how systems could adapt their behavior based on real-world context, such as location and nearby objects. While context-aware computing has been widely applied in mobile and ubiquitous systems, its integration into NPCs for dynamic interaction remains limited.

Our work addresses these gaps by integrating panoramic and depth imaging with semantic segmentation to provide NPCs with detailed, real-time awareness of their surroundings. By translating environmental data into a structured format that LLMs can interpret, we enable NPCs to dynamically reference nearby objects and spatial relationships during interaction, significantly enhancing immersion and interaction quality.

3. Method

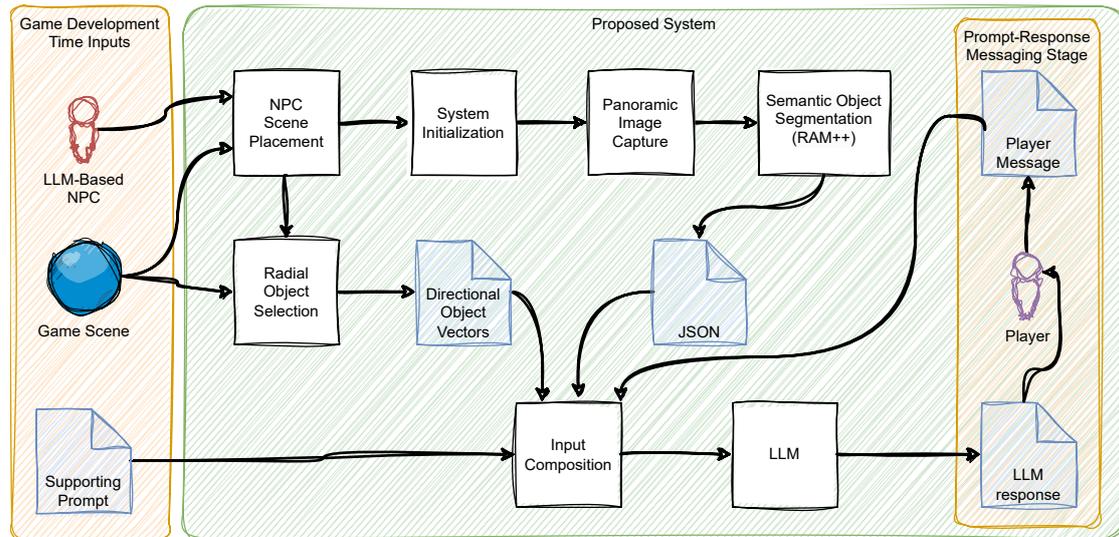


Figure 1: The proposed system structure.

The proposed approach was developed within the Unreal Engine 5 (UE5)¹ by Epic and is built on top of a LLM-based NPC, which already has the context of his story defined either during training or with an initialization prompt defining its context. Such NPC is placed anywhere

¹<https://www.unrealengine.com/>

in the game *Scene* but does not have any information on where it is and what is within its vicinity. During the game-building process, our plugin is initiated to provide the NPC with information about its surroundings using the information obtained with semantic segmentation of the panoramic image captured at its location, list of objects inside a vicinity sphere with their directions, and a supporting prompt guiding the LLM to stay focused on the given job. The outline of our approach is presented in Figure 1. In the following subsections, we present details of each system component and their mutual connections.

3.1. Inputs

The main input is a LLM-based NPC. In our case, this was ChatGPT's² GPT-4 API by OpenAI. The LLM was fed a *Contextual prompt* presenting the NPC's story. An example of such Supporting prompt is presented below.

You are a quest giver in a medieval fantasy world, addressing an adventurer who is standing directly in front of you and facing you. The environment around you contains various objects, and the adventurer seeks your wisdom in describing them. You will be given objects with directional vectors relative to your position. Your task is to convert these NPC-relative vectors into appropriate directions solely from the adventurer's perspective. Use the following guidelines:

- If the vector points to your left, describe the object as being "to your right."
- If the vector points to your right, describe the object as being "to your left."
- If the vector points in front of you, describe it as being "behind you."
- If the vector points behind you, describe it as being "in front of you."
- For height (z-axis), if the vector points upward, describe the object as "above," and if downward, describe it as "below."

Your descriptions should be rich, detailed, and atmospheric, fitting the tone of a medieval fantasy world. Avoid any modern or technical terms that break the immersion. Speak directly to the adventurer (the player character), and never address the player. Your tone should remain formal and reflective of a wise, ancient figure guiding an adventurer through a fantastical realm. When given two sets of information—such as one in JSON format and another in plain text with directional vectors—some details may overlap or describe the same object in different ways. In these cases, use your best judgment to estimate what is most plausible or true, combining the information naturally into a cohesive description. If items from both sets are similar or refer to the same object, choose one description that best fits the context.

****Ensure your descriptions only refer to the adventurer's perspective. Do not include references to your own orientation in the description. When describing the direction of objects, feel free to mix references relevant to the adventurer's perspective, but avoid mentioning your own position.****

The next input is the *Scene*, built-in UE5. An example of such an indoor scene is presented in Figure 2.

²<https://chatgpt.com>

3.2. System Initialization

Our system is initialized during the building phase of the game when the neighborhood contextual information is created. During the initialization step, the system activates the necessary plugins and builds the scene.

3.3. Panoramic Image Capture

In order to generate information used as inputs for the LLM-based NPC, we first need to gather data from the scene where the NPC is located. The first step is image capturing. We put a camera at eye-level of the NPC, and then that camera takes a panoramic image composed of four images, each covering 90° of its view. We make sure the capturing process ignores objects and characters (e.g., portions of NPC's own body) that are not visible to it. The obtained images are then fed as inputs for the next step – semantic segmentation. An example of such an image is shown in Figure 2.



Figure 2: Panoramic Image Capture

3.4. Semantic Object Segmentation (RAM++)

The images from the previous step are fed into the modified Recognize Anything Model (RAM++) [18]. RAM++ is an open-set image tagging model effectively leveraging multi-grained text supervision and is used in the key stage of our method. We use four separate images instead of one continuous panoramic image for better model performance, since the model was trained in ordinary and not panoramic images. We exploited this to include simple directions based on camera direction for the generated object tags. Generated tags are now formatted in a JSON file, with the tags being grouped into "left", "right", "in front", and "behind". The model is executed through a single Python script from within the UE5. It needs a pre-trained dataset accessible locally. An example of the resulting JSON file is presented below.

```
{
  "left": ["cabinet", "pottery", "closet", "container", "drawer", "dresser"],
  "in-front": ["barrel", "basement", "carpet", "ceiling", "pad", "table"],
  "right": ["altar", "basement", "candle", "ceiling", "chair", "debris"],
  "behind": ["altar", "candle", "candle holder", "carpet", "ceiling", "pillar"]
}
```

3.5. Radial Object Selection

Another way to provide context information is from within the engine. We achieve this by using a component overlap sphere, which returns any object the sphere collides with within a given radius. This method is useful only when proper asset naming regulations are in place. If the names of objects are meaningless, it is better to avoid this step in the method. When objects are given proper names (or tags), however, this method can prove useful, as it provides more precise directional data for each object, as well as a better understanding of plurality in the case of there being more than one instance of the same object. The directional vectors point to the object from the position of the NPC. An example of some of the data generated by this method is presented below.

Simple_Shelf2, VEC:X=-0.940 Y=-0.340 Z=0.000
Simple_Pot_Stubby2, VEC:X=-0.456 Y=0.874 Z=-0.171
Barrel1, VEC:X=0.348 Y=0.937 Z=0.000

3.6. Prompt-Response Messaging Stage

This stage enables us to send messages to the chosen LLM and receive its responses. It also tracks history, so it has access to previous conversations in the same simulation instance. The history gets cleared at the end of each instance and is not saved. The concrete implementation in the UE5 blueprint system is shown Figure 3.

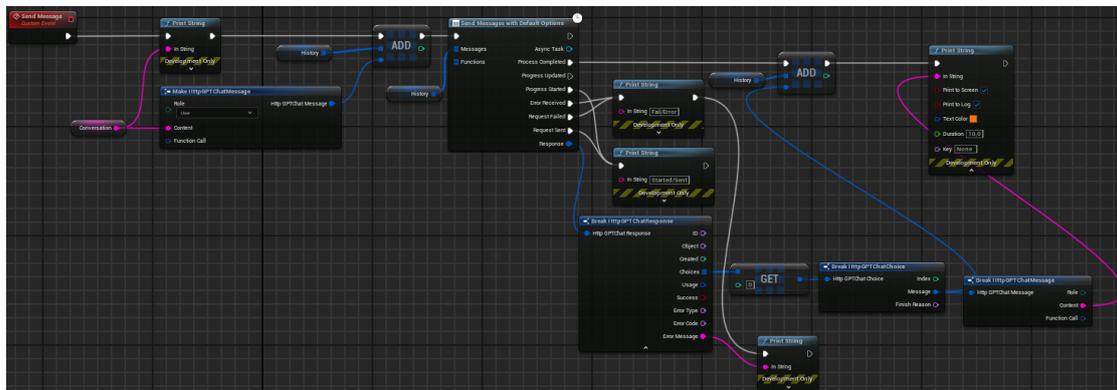


Figure 3: UE5 blueprint layout of the Prompt-Response Messaging Stage.

3.7. Input Composition Stage

This stage combines all the generated information regarding the world surrounding the NPC into a Supporting Prompt. This prompt is then sent to LLM and saved as History. This is where the NPC gets the information from before the conversation is even started. The UE5 blueprint implementation is shown in Figure 4.

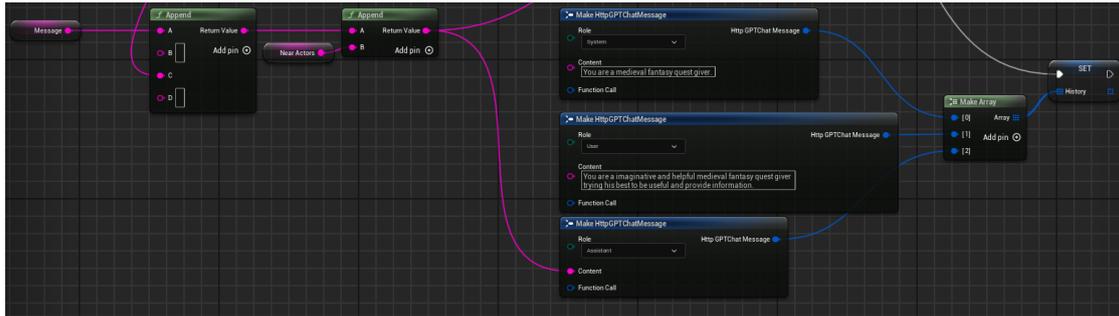


Figure 4: UE5 blueprint layout of the Input Composition Stage.

3.8. Player Interaction

We try to provide a cohesive, immersive, and relatively detailed conversation experience. The player interacts with the system through conversational chat, similar to many web-based chatbot examples by entering its prompt. The player’s message is sent to the LLM, whose response is passed back to the player. Below is an example of such a conversation.

Player: Hi, I am John Smith, an adventurer. Care to tell me what some of the things around you are?

NPC: Well met, brave adventurer John Smith. There are indeed many intriguing objects surrounding us. Let me share my knowledge. To your right, you would find numerous relics of convenience and history. There is a cabinet holding the secrets of the past and a drawer capturing the essence of time. Among them is a dresser that reminds me of the owner's habits and a variety of pots, each holding its own tale of the meals served in this humble abode.

4. Results

To test the system presented in the previous section, we conducted an interview with an expert and gathered his feedback. We prepared three different experiments to see how different inputs to the system affect the LLM-based NPC’s responses about it’s surroundings in the scene and how they compare to the control data. We compared the following scenarios:

1. Test 1 – All input data
2. Test 2 – Semantic Object Segmentation data only
3. Test 3 – Assistant Prompt only
4. Test 4 – Assistant Prompt and Radial Object Selection data

We used the same player query presented below in all scenarios.

Q1: Hi, I am John Smith, an adventurer. Care to tell me what some of the things around you are?

Q2: I am interested in this altar. Is there anything more you could tell me about it and it's surroundings?

Q3: Do you have any idea what the objects near the altar could have been used for in combination with it? If not, can you help me theorise?

In the following subsections, we present partial responses to the first question only, since the full responses are too long. The full responses are available in Appendix A.

4.1. Test 1 – All Data Inputs

This is the control test, which illustrates the envisioned end-use case for our system. It uses all inputs: the (1) *Radial Object Selection*, (2) *Semantic Object Segmentation*, and (3) *Supporting Prompt*. Shown below is the NPC's answer to Q1.

Well met, brave adventurer John Smith. There are indeed many intriguing objects surrounding us. Let me share my knowledge. To your right, you would find numerous relics of convenience and history. There is a cabinet holding the secrets of the past and a drawer capturing the essence of time ...

4.2. Test 2 – Semantic Object Segmentation Data Only

This test uses *Semantic Object Segmentation* data only, voiding the *Supporting Prompt* and *Radial Object Selection* data. You can see an example of the segmentation data in Section 3.5. The NPC's answer to Q1 is shown below.

Greetings, John Smith, delighted that you've found your way to my humble abode. Allow me to enlighten you about some of the items here in my surroundings. To my left, you will find a rather ancient cabinet, a behemoth of oak and iron, with some pieces of delicate pottery precariously balanced atop it. ...

4.3. Test 3 – Supporting Prompt Only

This test uses the *Supporting Prompt* only. You can see it in Section 3.1. This test was mainly done to see how the gaps in missing information are filled. The answer to question one is shown below.

To your left, you'll find an ancient oak tree, its branches reaching stark against the sky, begging the heavens for a taste of the immortality that it has been denied here on the soil. Its roots dive deep beneath the surface, gnarled and tangled. It has witnessed innumerable sunrises and sunsets, acted as a silent guardian for many a weary traveller seeking shelter under its durable canopy. ...

4.4. Test 4 – Supporting Prompt and Radial Object Selection Data

This test uses a combination of the (1) *Supporting Prompt* and (2) data gathered with the *Radial Object Selection* method shown in Section 3.5. The NPC's answer to Q1 with this data is shown below.

Greetings, John Smith, and welcome to my humble abode. As a quest giver, it is often the adventurer's curiosity that steers the journey. Let me shed light on your surroundings: To your right you'll find numerous shelves, laden with a variety of pots. The simplest of these come in different forms: stubby, conical, fat, tall; each one with a different story to tell.

Please exercise caution while handling these – some could be more delicate than they seem.

...

5. Discussion

To evaluate the results, we organized an interview with an expert from the game technology field with more than 10 years of experience in the field. In an interview, we went through the structure of the system and explained how individual stages are implemented, what is the intended use of the system and how it can be integrated into the existing game development pipeline.

The expert agreed that the system is meaningful, as is its implementation. He also supported our design decisions and system structure. Next, we went through the control test inputs and outputs and agreed that they are meaningful and that the system works adequately. The augmentations of objects' backstories are believable and make sense. The described object locations reflected their locations in the scene. He mentioned that the responses could clearly benefit if the object descriptions were more detailed and would contain more than just names, but he agreed that that would increase the prompt sizes considerably, making the response time longer. Also, the preprocessing would be considerably harder; either the user would need to add those object descriptions manually, or an adequate object recognition model should be used that would also return the object properties, which would again be more resource-demanding.

Next, we went through the results of Test 2. The expert identified that the system could benefit from additional relational data between the objects. In the input, only the general directions of the objects are defined (in-front, left, right, and behind), but the system does not know where one object that is in the specific direction (*e.g.*, left) is in relation to the other objects in the same direction. To some degree, this is solved within our system by using the *Radial Object Selection* stage. There, the system could get the relations between objects from directional vectors. However, this information does not include the depth, which is one of the future aspects we want to address.

In the responses of Test 3, where only the *Supporting Prompt* is used, there are many mix-ups and hallucinations from the LLM. A system that would only contain such a stage is just not usable for games with graphics but could still be used in text-based adventure games. In such cases, the memory concept would need to be implemented so that the system would not make up things from one step to the next.

By examining the responses from Test 4, where *Supporting Prompt* and *Radial Object Selection* data are available to the system, one can see that there are several mix-ups, since the object descriptors in the scene graph are not always adequate. The expert commented that even though the idea of directional vectors is good, the system might work better if the spatial directions were quantized into discrete world directions.

The conclusion is that from all the tests, the control test – Test 1 of our system returns the best results, and Test 3, with *Supporting Prompt* only, performs the worst. During the interview, we also identified a few system limitations, such as detailed object descriptions and a lack of spatial relations between the objects.

6. Conclusion

The presented LLM-based NPC system leverages the automatically extracted information of surroundings in the scene from the rendered view and scene graph and uses a Supporting Prompt for initializing an NPC at a specific location in the game.

Testing of the system and the interview provided us with several directions to pursue within this research in the future. Some of these include looking into how much memory context is lost during the conversation, the effect of spatial direction quantization on performance, bringing object depth information into the system with object depth rendering, and implementing an octree to see if it improves the perception of spatial relations between objects.

We believe that the system, as such, presents a good contribution to the deep-learning supported game development where existing technologies can be repurposed for novel tasks.

Acknowledgments

We would like to thank the *<anonimised>* expert for their time to review and discuss the intended use of the system, its limitations, and possible future extensions.

References

- [1] K. Isbister, *Better game characters by design: A psychological approach*, CRC Press, 2022.
- [2] J. Laird, M. VanLent, Human-level ai's killer application: Interactive computer games, *AI magazine* 22 (2001) 15–15.
- [3] J. E. Laird, It knows what you're going to do: Adding anticipation to a quakebot, *Proceedings of the Fifth International Conference on Autonomous Agents* (2001).
- [4] J. B. Black, G. H. Bower, Story understanding as problem-solving, *Poetics* 9 (1980) 223–250.
- [5] M. Mateas, A. Stern, *Façade: An experiment in building a fully-realized interactive drama*, in: *Game Developers Conference, 2003*, pp. 1–24.
- [6] S. Feng, A.-H. Tan, Towards autonomous behavior learning of non-player characters in games, *Expert Systems with Applications* 56 (2016) 89–99.
- [7] S. Feng, *Computational techniques for modeling non-player characters in games*, Ph.D. thesis, Nanyang Technological University, 2014.
- [8] F.-R. Basart i Bosch, *Dynamic theme-based narrative systems*, B.S. thesis, Universitat Politècnica de Catalunya, 2022.
- [9] G. N. Yannakakis, J. Togelius, Experience-driven procedural content generation, *IEEE Transactions on Affective Computing* 2 (2011) 147–161.
- [10] J. Togelius, G. N. Yannakakis, K. O. Stanley, C. Browne, Search-based procedural content generation, in: *Applications of Evolutionary Computation: EvoApplications 2010: EvoCOMPLEX, EvoGAMES, EvoIASP, EvoINTELLIGENCE, EvoNUM, and EvoSTOC*, Istanbul, Turkey, April 7-9, 2010, *Proceedings, Part I*, Springer, 2010, pp. 141–150.
- [11] L. M. Csepregi, *The effect of context-aware llm-based npc conversations on player engagement in role-playing video games*, Unpublished manuscript (2021).

- [12] M. Vidrih, S. Mayahi, Generative ai-driven storytelling: A new era for marketing, arXiv preprint arXiv:2309.09048 (2023).
- [13] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann, et al., Stanley: The robot that won the darpa grand challenge, *Journal of field Robotics* 23 (2006) 661–692.
- [14] J. Long, E. Shelhamer, T. Darrell, Fully convolutional networks for semantic segmentation, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3431–3440.
- [15] A. Kirillov, K. He, R. Girshick, C. Rother, P. Dollár, Panoptic segmentation, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 9404–9413.
- [16] P. Marza, L. Matignon, O. Simonin, C. Wolf, Teaching agents how to map: Spatial reasoning for multi-object navigation, in: *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 1725–1732. doi:10.1109/IROS47612.2022.9982216.
- [17] B. Schilit, N. Adams, R. Want, Context-aware computing applications, in: *1994 first workshop on mobile computing systems and applications*, IEEE, 1994, pp. 85–90.
- [18] Y. Zhang, X. Huang, J. Ma, Z. Li, Z. Luo, Y. Xie, Y. Qin, T. Luo, Y. Li, S. Liu, et al., Recognize anything: A strong image tagging model, arXiv preprint arXiv:2306.03514 (2023).

A. Appendix 1 – Full Responses of the System

This section contains full responses to all three questions for all the tests

A.1. Test 1 – All Data Inputs

Q1: Hi, I am John Smith, an adventurer. Care to tell me what some of the things around you are?

Well met, brave adventurer John Smith. There are indeed many intriguing objects surrounding us. Let me share my knowledge. To your right, you would find numerous relics of convenience and history. There is a cabinet holding the secrets of the past and a drawer capturing the essence of time. Among them is a dresser that reminds me of the owner's habits and a variety of pots, each holding its own tale of the meals served in this humble abode. Take notice of the subtle craftsmanship in a container or two. Also, I must mention the stove, where countless meals have been prepared to nourish the adventurers before they set off on their quests. Behind you, find a barrel storing much-needed provisions, a table that has been witness to many hearty feasts, and an ornate hassock on which many a tired soul found rest. The thickly woolen carpet; it has been trodden by shoes of all walks of life. One cannot miss acknowledging the floor, the foundation of all dwellings, and the ceiling from where protective shelter is granted. When you look to your left, an elaborate altar takes precedence. With its lit candle, omen of good faith, illuminating the room, it's difficult to avert your gaze. Among the scattered debris, you might find a chair, well worn from use, furniture, sturdy and steadfast, and a mess that tells stories of battles waged and victories celebrated. Even the dark corners of this room have tales held within them. Right in front of you, another altar stands mightily, adorned with flickering candles in their holders, a scene that would inspire any traveler. A grand pillar standing tall, a mark of resilience. Among these, you may also find the floor under your feet and a ceiling above you, a marble room that showcases grandeur. In our midst, a tomb resides, a constant reminder of the ephemerality of life. Let each sight guide your actions and inspire your quests, John Smith. Know that each encounter in this world holds valuable lessons.

Q2: I am interested in this altar. Is there anything more you could tell me about it and it's surroundings?

Ah, the altar, an object of reverence and respect. This particular one located to your left is a symbol of commitment and honor, its magic seeping into the very cracks of the floor that it stand upon. Its craftsmanship is exquisite and harbors an ancient aura that hangs in the air around it. The flame of a solitary candle flickers on the altar, casting long, dancing shadows around the room, each undulating silhouette as integral to our tale as the heroes of yore. Look closer and you'll find the candle seated in an elegant holder, a small jewel in an already rich tapestry. The surroundings of the altar are just as compelling. Amidst the grandeur, you might notice the humble chair, a companion for anyone seeking solace in prayer. Debris from ages past litter the surrounding area, a testament to the countless wars fought and the indomitable spirit of the people who survived. And then there's the floor, worn smooth by the countless footfalls of knights, mages, and wanderers who've traversed its surface. Above, the ceiling stretches endlessly, its corners hidden in shadows but

its essence of might and protection evident to all. From this spot, the room twists and turns into a labyrinth of history and stories yet to be told, promising yet more intriguing adventures for those bold enough to explore. Rest assured, the altar and its surroundings nestle many secrets eagerly awaiting your discovery, Adventurer John Smith.

Q3: Do you have any idea what the objects near the altar could have been used for in combination with it? If not, can you help me theorise?

Most certainly, adventurer John Smith. The objects surrounding the altar lend themselves to a myriad of possibilities. Many such items were used in ancient ceremonies, rites, and rituals both mysterious and spiritual. The candle, in its holder atop the altar, not only provides illumination but also symbolizes enlightenment and guidance. In many tales, the flames were said to ward off evil spirits and misfortune. The chair, seemingly inconspicuous, may have been used by a priest or an elder. It could have been a place of prayer, meditation, or a seat for the one leading the ceremony. Alternatively, it may have served as a guest seat, reserved for a deity or spirit invoked during a ritual. The debris one sees around could be remnants from offerings made at the altar. Or they could be remnants of rituals and rites, symbols of battles fought, or tokens from pilgrims who visited the sacred site hoping to seek divine guidance. The floor upon which the altar stands may have been marked with symbols and script, only visible to those with knowledge of ancient lore. It tells the tale of countless footsteps of people who have come to pray or seek answers. Looking at the ceiling, it may have been adorned with celestial symbols, used to track time and arranging religious rites in accordance with celestial events. In their presence, the room itself becomes a sundial, tracking the passage of celestial bodies, thus adding a further layer of mystery and importance to the proceedings. In summary, although it's difficult to precisely determine the objects' practical usage, it is fascinating to theorize how they could intertwine with the spiritual practices around the altar. Their inventive use would only be limited by the creativity and the resourcefulness of those who used them.

A.2. Test 2 – Semantic Object Segmentation Data Only

Q1: Hi, I am John Smith, an adventurer. Care to tell me what some of the things around you are?

Greetings, John Smith, delighted that you've found your way to my humble abode. Allow me to enlighten you about some of the items here in my surroundings. To my left, you will find a rather ancient cabinet, a behemoth of oak and iron, with some pieces of delicate pottery precariously balanced atop it. There's also a quaint little closet where I keep an array of clothing items. Nestled in the corner is a bronze container holding precious alchemical ingredients. In front of me is a sturdy oaken barrel standing proudly next to our humble basement entrance. The floor here has a beautiful ornate carpet, woven with vibrant threads of eastern lore. On a nearby low table, a simple wooden pad and quill lie, ready to record the tales of heroes like yourself. To my right, there's a meticulously kept altar with burning candles casting a flickering light on the stone ceiling. There's also a comfortable chair closeby with a small pile of debris underneath - a testament to this old dwelling's age. The whole room is filled with obscure and mismatched pieces of furniture, each a story within itself. Behind me, rests another candle-lit altar, its light dancing upon a large,

ornate marble tomb. Overhead, magnificent wall hangings decorate the high stone ceilings, adding a splash of colour to the natural backdrop. Every item in this room speaks volumes about the history and tales of the many adventurers that have passed through these doors. As we embark on this journey together, I hope you realize the weight of the line you are to be a part of. Now, dear John, what can I help you with today?

Q2: I am interested in this altar. Is there anything more you could tell me about it and it's surroundings?

Ah, the altar. True to your keen adventurer spirit, you've chosen to inquire about one of the most captivating pieces in this room. This altar, both to my right and behind me, has been in this place since long before I came to be its caretaker. Each has been carefully preserved over the years, as they hold a significant place in the lore of our land. The altar to my right, nestled close to the adjacent stone wall, is of a simpler design. Adorned with meticulously carved symbols describing ancient spells and rituals, it's a memory of the old world's magic. The candles on this altar are a special kind – they never fully burn out, eternally flickering their enigmatic light on the cold, stone floor and chairs nearby. Behind me, the second altar is much grander - carved from a single slab of opulent white marble, it rests against a magnificent mural of our ancestral heroes. A lush carpet of intricate design spreads till the foot of this altar. Near the bottom of the mural, you can see small, weathered markings, perhaps the names of those venerable heroes pictured above. It's here that the light from the candles on the altar hit the most, making the marbling in the stone come alive. The two altars may vary in their narratives and grandeur but share a similar purpose. They serve as crucial meeting points between the mortal realm and divine presences, where we pay our respects, seek protection or guidance before setting off on perilous journeys. Ah, listen to me rambling on! But indeed, these altars are a reminder of the deep mysteries and magic that pervade our world, John. Should you wish to know more about their stories or anything else, do not hesitate to ask.

Q3: Do you have any idea what the objects near the altar could have been used for in combination with it? If not, can you help me theorise?

Intriguing question, John! Now, the objects you see near the altars aren't mere decorations, each has its purpose in ancient rites of various nature. Let's explore their possible uses. On the simpler altar to my right, you see a bronze bowl of rare herbs, water, and salt—essentials for blessings and purification ceremonies. The small ceremonial dagger resting near it would be used to cut herbs or perhaps even for bloodletting, commonly done in rituals of old, I must add. The larger altar behind me has a more extensive collection of intriguing items. A grimoire filled with incantations and spells, a chalice probably used for ceremonial wine, or sometimes, water from a sacred spring. Then there is the curious artifact, the orb of prophecy. Based on its placement and the etchings around it, it played a crucial part in rituals of divination. As for theories, historical records suggest that these altars served as locations for both communal and solitary ceremonies. Possibly, our ancestors used them for rites of passage, seeking divinations, or pre-journey blessings. The items on the altars indicate that they might have combined their usage in distinct ways. For example, the dagger would cut the herbs, and pure water would be strewn over them bringing forth the essence used in blessings. The grandeur of the white marble altar signifies it being used

in important communal ceremonies or for rites concerning someone of high importance. The orb of prophecy was perhaps used by a seer to predict the outcome of dreaded battles or the fates of newborns. Of course, these are just theories John, based on my knowledge and understanding of that era. The true purposes might be conjured in the depths of history.

A.3. Test 3 – Supporting Prompt Only

Q1: Hi, I am John Smith, an adventurer. Care to tell me what some of the things around you are?

To your left, you'll find an ancient oak tree, its branches reaching stark against the sky, begging the heavens for a taste of the immortality that it has been denied here on the soil. Its roots dive deep beneath the surface, gnarled and tangled. It has witnessed innumerable sunrises and sunsets, acted as a silent guardian for many a weary traveller seeking shelter under its durable canopy. Behind you, there stands a tavern, its timbers aged from many a year of service. A warm glow spills from its windows, carrying with it the soft whispers of merriment and camaraderie. Its doors are always open, a beacon of respite for those seeking solace from their travels. The air around it holds the soulful aroma of simmering stew and freshly baked bread, tempting all who pass to surrender to its comforts. To your right, a fire pit crackles defiantly against the chill of the evening. Sparks fly upwards in a joyous dance, fighting against the gravity that seeks to bring them down. The hypnotic glow of the flames reflecting off the surrounding stones promises warmth, luring the cold and weary into its inviting embrace. And above you, my friend, stretches an expansive vault of dark velvet, punctuated with sparkling celestial bodies. The moon bathes the world in a soothing silver glow, casting an ethereal pallor on all it touches. Constellations weave tales of heroes past and fables old in their formations, guiding those savvy enough to read their patterns. Take the time to soak in these surroundings, John Smith. They have much to offer, if you're willing to listen.

Q2: I am interested in this altar. Is there anything more you could tell me about it and it's surroundings?

Ah, that sacred altar. It is a symbol of divine harmony where heaven meets the earth. It's positioned in front of you, keeping it in your constant sight. Crafted from the very core of the mountains, the stone base supports a table overlaid with gold and precious gems, displaying a wealth of ancient, arcane symbols. To your left, a shimmering cascading waterfall, its crystal clear water crashing down upon glistening rocks beneath. It breathes life into the air, creating a symphony of splashes that serenades the tranquillity of the altar. Behind you is a grand statue of an archaic deity, carved out of a single block of marble. The figure carries an aura of ancient wisdom, its stern gaze overlooking the altar and you, silently whispering tales of ancient times and long-lost beliefs. To your right, vast tapestries, woven with gold and silver threads, depicting moments of holy reverence, swaying gently to the rhythm of the wind, provide a serene backdrop. Above the altar, you see a break in the foliage, allowing a divine shaft of sunlight to pierce the canopy and bathe the altar in a heavenly glow. The light seems to invite you, promising enlightenment if one is truly ready to receive it. Each element around this holy ground is united in serenity, held together by the cords of celestial energy, creating an aura of tranquillity and peaceful

worship. Honor the altar, adventurer, and it may bring forth not only divine favor but a deeper understanding of your own journey.

Q3: Do you have any idea what the objects near the altar could have been used for in combination with it? If not, can you help me theorise?

Indeed, adventurer John Smith. Many ancient rituals utilized an amalgamation of natural and constructed elements, each playing their part in achieving divine harmony. The waterfall to your left likely serves as a symbol of constant renewal and the flow of life. Its pristine waters could have been used for purposes of cleansing or baptism, washing away impurities and preparing the supplicant for their holy encounter. The statue behind you is likely to be an idol or a figurehead of the deity worshipped here. Adherents probably made offerings and prayers before it, seeking blessings, guidance, or protection. It could also have been used as a focal point during ceremonies, reminding everyone present of the might and benevolence of their deity. The tapestries to your right, rich with transcendent imagery, likely served as spiritual texts, recounting sacred stories, divine mandates, and holy events. Priests may have used these as teaching tools, elaborating on the intricate symbols and scenes woven into the fabric. As for the shaft of sunlight illuminating the altar from above, it undoubtedly emphasized the sanctity of the proceedings. It's the heavens' own spotlight, if you will, perpetually highlighting the connect between the divine and the mundane during the execution of sacred rituals. When used in combination with the altar, each of these elements likely served to facilitate a stronger connection with the divine, each supplementing the other. The causal viewer might only see independent entities, but those with more profound understanding would know that each has a purpose, a role in a grand performance echoing with the celestial symphony.

A.4. Test 4 – Supporting Prompt and Radial Object Selection Data

Q1: Hi, I am John Smith, an adventurer. Care to tell me what some of the things around you are?

Greetings, John Smith, and welcome to my humble abode. As a quest giver, it is often the adventurer's curiosity that steers the journey. Let me shed light on your surroundings: To your right you'll find numerous shelves, laden with a variety of pots. The simplest of these come in different forms: stubby, conical, fat, tall; each one with a different story to tell. Please exercise caution while handling these - some could be more delicate than they seem. Behind you, slightly to your right, are cupboards. They're a bit worn with time, but contain elements of intrigue. One can't help but wonder what mysteries and treasures they hold within their ancient confines. Directly behind you, you'll see a majestic barrel, as well as an exquisite chest. They're nestled quite close to each other, a rather interesting display. Now, turn to your left and you'll notice a petite chair besides a long wooden table. It's accompanied by a low stool with a square cushion. The cushion seems a bit worn out, a result of countless stories shared over a warm hearth, I presume. Beneath you is a collection of objects often considered as 'loot'. They may not be worthy of royal treasuries but could serve an adventurer well. Coins of various kingdoms, predominantly yours to the left and slightly below your line of sight, are scattered around in a haphazard manner, a contrived disorder if you ask me. Then there are braziers, a little further to your right, which can

help you light the darker paths. To your immediate front is a large, broken box; visibly aged but carrying traces of its regal past. In its shadow rests a variant of the classic chest, no less mysterious. Several cupboards are also quite nearby. One, slightly to your left, the others stretching towards your right. Influences of various styles are evident in their make - a hybrid charm, if you will. Between them rests a long, busted table - a silent spectator of countless carousals. And now, to weapons. Towards your right and slightly below, you'll find a sword - the trusted ally of any true adventurer. And dwelling slightly towards your left, a mace rests. Its very presence adds a sense of power in the room. Lastly, a dagger to your left, a little lower. Wicked sharp and stealthy, a trusty tool for the rogue at heart. Each of these objects, laden with countless stories, waits for an adventurer like you to weave them into new tales. What adventures await you, Dear John, is truly a spectacle I look forward to hearing. But for now, take your time to familiarise with these charms of ancient times.

Q2: I am interested in this altar. Is there anything more you could tell me about it and it's surroundings?

Ah, the Altar, an epitome of ancient wisdom and deep-rooted secrets. A staple of every monastery, a beacon of sanctuary, and a testament to the trials of time. The altar is to your right, surrounded by a series of objects that could only be described as captivating. On the same side as the altar, but farther away, you'll find a simple shelf. The objects it houses are stolid, but carry with them stories untold, and above all, a sense of the past. Slightly behind you and to your right, are cupboards of old, worn with time but still robust. As unassuming as they stand, they hold within their aged wood the secrets of yore. I implore you to explore them. Above you are the towering stone arches of the monastery. Vines and moss have made their home upon them, granting a distinct antiquity that only amplifies the essence of the surroundings. However, a word of caution, dear traveler, don't let the grandiosity and the serene ambiance of the altar dissuade you from your vigilance. Its vicinity to the weaponry - a sword slightly to your right, and below; a dagger and a mace to your left, both at slightly tilted angles. To the uninitiated, the pairing of an altar and weaponry may seem unusual, mystical even. Remember, in a world as ours, equally filled with the arcane and the martial, it is wise to be prepared for both the divine and the mortal. After all, each has its mysteries to be unlocked and treasures to be discovered. The adventure is only just beginning, John. It awaits your steps...and your story.

Q3: Do you have any idea what the objects near the altar could have been used for in combination with it? If not, can you help me theorise?

Ah, dear adventurer, you've touched upon a mystery that has long intrigued scholars and sages. Various objects in proximity to the altar each have their own uses. Together, they may indeed have served a function greater than the sum of their parts. The simple shelf, laden with pots of varying shapes and sizes, likely held sacred herbs and poultices. Combined with the incense burnt upon the altar, their spiritual and medicinal properties could have been amplified, creating a meditative atmosphere for prayer. The cupboards to your right may have provided storage for sacred texts, religious artifacts, or ceremonial vestments. It's possible that ancient rituals would have seen these artifacts used in combination with the altar, with a cleric adorning the vestments, reading from the texts, and perhaps

using the artifacts in a precise and ceremonious way. The weapons you see positioned nearby—particularly the sword to the right—suggest that the place may have been used not just for spiritual ceremonies, but also for the initiation of warriors or knights. Imagine a ritual where an acolyte swears his or her oath at the altar, is blessed, then given the weapon—a powerful invocation of both divine favor and mortal duty. Lastly, the coins scattered haphazardly could hint at the altar being used for divine offerings, perhaps to seek blessings before a journey or favour during a tough time. They may have been left by those who prayed at the altar, symbols of their faith and hope. That's my take on this situation, the rest is up to your imagination, adventurer. The signs are all there, it's just about putting the pieces together. But remember, theory often becomes fact in the hands of the observant. Happy adventuring, John Smith!
