Proceedings of the Twenty-First AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE 2025)

Setting the Stage: Using Knowledge Graphs for Environmental Context for Social Characters

Arunpreet Sandhu
University of California, Davis
1 Shields Ave,
Davis, CA 95616
asisandhu@ucdavis.edu

Abstract

In this paper, I present a preliminary approach for encoding location data for CiF-like social characters in a to alleviate some of the authorial burden associated with authoring these types of interactive characters. To ease some of the authorial burden, I discuss how decoupling character actions and storing them within the environment can lead to a lowering of authorial complexity as well as treating the environment as a first-class entity can lead to more nuanced character performances. To better understand what is being decoupled, I describe how the term environment may be too vague when regarding interactive characters and their performances, and instead extending concepts from theater and dramaturgy, such as stage and props. From there I describe their functionality and possible encoding for props and stages.

Introduction

The authoring of interactive characters is an immense authorial burden (Chen, Nelson, and Mateas 2009; Garbe 2020) for writers and designers because of the possible interaction space of the player. Developing characters that are both believable and responsive to player interaction is challenging because it requires the effective capture of nuanced performances while maintaining consistency within the game environments to ensure believable characters (Mateas 1999). To address the authorial burden, character systems and narrative systems have been created to reduce the difficulty of designing believable characters for interactive experiences, systems such as A Behavior Language (ABL) (Mateas and Stern 2002), Comme il Faut (CiF) (Mccoy et al. 2011), and Praxis (Evans and Short 2014).

Most computational character systems prioritized character representation and behavior, addressing various aspects as primary salient representations (e.g., planning character actions, modeling social dynamics, simulating cognitive processes). An additional representational layer of the environment and its objects could be incorporated. In Facade, the ABL-powered characters Trip and Grace (Mateas and Stern 2005) are in an apartment, and the CiF-powered high school students in Prom Week (Mccoy et al. 2013) are primarily at their high school before the setting changes in the final act to the prom night. These character systems still needed to nest their characters in space; thus, if characters must have performances, then they need a stage for those performances.

The importance of encoding more information into the environment is due to enhancing character believability and addressing authorial burden; therefore, location data and information should have the same degree of discrete systems-making as other computational narrative and character systems. By reframing the environment and objects as a stages, sub-stages, and props and treating each as a container of additional social or meta-functionality through a queryable graph structure, I believe this would give more authorial leverage for writing complex narrative systems. The rest of this paper will focus on my preliminary breakdown of the environment, tentative implementations by extending CiF to be environment-aware to prototype a more robust future system

Previous Work

In 2018, I worked on a design-enabled WaveFunctionCollapse (Sandhu, Chen, and McCoy 2019), which would add in rules that would execute depending on the state of the generative artifact. If a locked door was generated, then a key should be put somewhere within the level. I began looking towards other schools of thought when it came to space–such as architecture. After familiarizing myself with architectural design philosophies, including Christopher Alexander's spatial concepts CITE, I authored a poster paper focused on design patterns in levels and environments, as well as potential methods for encoding them into procedural algorithms, such as WaveFunctionCollapse (Sandhu and McCoy 2019). The focus of my third project in procedural level generation focused on generating 2D horror maps that are generated based on genre conventions (Sandhu, Mitchell, and Mccoy 2021). TileTerror was made possible with the work of my co-author, Kyle Mitchell.

As for my character and procedural narrative research, I was involved in the development of a system, named Sunset Valley (Liao et al. 2023), which focused on the flow of gossip in a social network. For Sunset Valley, I utilized novelist methods of introducing characters through direct interaction or by report (Card 1988). Furthering my analysis of novelist techniques, I published a paper centered on narrative generation via MICE, a method utilized by novelists for story creation, where MICE was an acronym for generating four different plots—Milieus, inquiry, Character, and Event. My most recent published work examined the relationship between procedural narrative and procedural content generation (Sandhu and McCoy 2023).

Prospective Idea

Characters do not act in front of a white room[1] devoid of context; their actions are grounded in the digital world that they are embedded in. Yet, often the digital environment is generally considered set dressing, with a few interactable objects scattered throughout the environment. While there are ways to ensure characters seem involved in their settings, such as pathfinding algorithms to more scripted interactions such as cutscenes, a deeper representation of the environment may still be undiscovered. Which brings me to the core of my prospective idea, extending the dramaturgical metaphor of CiF to the environment by considering objects and the environment as props and stages. In order to enhance character believability, the environments in which they exist must be first class entities, much like social relationships and social practices. I do not believe this is too far-flung of an idea as game designers have already been considering the importance of environmental consideration, such as environmental storytelling (Carson 2004) to using scenography within their practice (Kemenczy 2014), to constructing "smart zones" (de Sevin, Chopinaud, and Mars 2019) for ambient NPCs.

The following sections of this paper will be dedicated to presenting definitions of roles, props, stages, and substages, and possible implentantion strategies will be described. First roles will be explained as they are a foundational step for enabling the stage and props; while important, they are not the crux of my interests, which are stages and props, which I will describe and then give a possible encoding for each. CiF will be mentioned much through this, but I treat it as a prototyping tool for the eventual stage-centric system I wish to build. I also take inspiration from Dormans work on generating spaces for missions (Dormans and Bakkes 2011) and Ludoscope's graph system (Protsenko, Dormans, and van Rozen 2025) for encoding environment information.

Roles To start, roles can be understood as a compilation of information on how a specific "role" should be executed, similar to character archetypes in Role-Playing Games (RPGs), such as, Larian Studio's Baldur's Gate III (Larian Studios 2023). An example would be a "tank" in an RPG focuses on absorbing a significant amount of damage, while a healer heals their allies, and the Rogue or Fighter focus on attacking as frequently as possible. These roles are distinct from each other, containing its own strengths and weaknesses. But they contain their own procedures and actions, such as a healer waiting behind their teammates as they are often suspectible to enemy attacks, in the back line, healing the tank. However, these roles need not be limited to conventional RPG classes. They can also be social roles, which can be encoded in CiF.

[1]The "white room syndrome" occurs in novels when a character or characters are not grounded in their environments and their actions seem to occur in a "white room."

The architecture of CiF will be excluded from this discussion, but it should be noted that CiF enables the bundling of rules and actions into data structures that function in a manner similar to roles. These collections are generic (for reusability reasons), which allow authors to devise a "parent" role alongside a "tank" role within the same game. The divergence between the two roles necessitates this point that a role need not be only combat. Instead of tanking the damage, the parent might have to speak to their child. CiF can fulfill both of these roles and, what's more engaging, a character can incorporate both roles concurrently. As for the need of roles, there must be some interface between characters, the stage, and props. A character with the role of a tank might find weapons far more valuable than a character with a parent role, but a character with both roles might find armor far more necessary than a weapon as they want to get back home to their child. Both the armor and sword are objects, but they are also props.

Props An example of a prop within a fantasy world, like Bethesda's Oblivion Remastered (Bethesda Game Studios and Virtuos 2025) [2], would be a sword. Dependent upon the player's sword quality, a blacksmith could comment, offering repair or replacement if it is a rusty blade. If instead the player has a legendary sword, then the blacksmith might show reverence or awe. Even though the two objects are of the same class (a sword) their social responses and implications should not be the same type of response. To further clarify, there is no functional difference between a prop and an object, except for the prop's extension into the meta-space of social interactions and character/social roles.

Concerning prop encoding, this represents an expansion of smart objects (Kallmann and Thalmann 1998), which encompass properties, procedures, and interactions within the object, primarily to allocate computational complexity from characters to objects, thus distributing procedures and allowing for a more reusable design. If smart objects were intended to help the animation process, then props are to decentralize the social complexity of social characters to objects as well. Rather than encoding all characters and their interfacing mechanisms with all props, only those with the correct social role can successfully interact with a prop that also contains information about that prop. For instance, a tank knows how to best use a shield, while a swordsman knows how to wield a blade versus a blacksmith repairing it. By choosing the object's affordances based on role, the prop becomes modular–much in the same way a smart object does, but now with a focus towards character believability rather than functionality. A key point, however, is that there needs to be some interface layer between the character and prop, hence the need for roles.

Stages Consider this example of two characters, A and B, where A has just been broken up with, while B is A's friend and has been tasked to get them out of the house. Undeterred

[2]Oblivion contains another great example of mechanical ways that props lead to more believable behaviors. Stolen items must be sold to a fence, thus reinforcing a player's choices, possibly leading to a deeper sense of role immersion.

by the request, B suggests they go to the cafe that A likes, and as they arrive, A appears more emotionally distraught than before. Concerned, B asks what's wrong, to which A responds with that their ex-partner would take them on dates to the same cafe. The scene closes with B feels bad now for being so crass and not remembering something so important.

This example captures why I think the environment needs to be treated as a first-class entity, but also shows the power of the stage. Without the environment as a first-class entity, encoding information regarding A and their ex-partner's cafe dates would be complex, involving inefficient searches through character action histories to find out which cafe A and their ex-partner would go to. Instead, by encoding the history and the actions within the cafe, it would be relatively easy to pull out the information of "A is at the cafe, their partner is now their ex, A should feel melancholy now."

Despite its potential for oversimplification, this represents a possible outcome within the cafe, assuming the cafe contained its own set of information. Furthermore, should A begin to weep, what would be the response of the patrons? Maybe there would be a series of character look ups within the system and cross-reference each role's reaction, however, the environment could contain more than just history. Stages could also contain their own procedures and practices that can be agnostically encoded from either character or narrative. Now the cafe-goer doesn't need to know about what to do, instead the cafe contains this information.

The cafe in the above example is at its essence, what I consider a stage to be, a container of location-specific procedures, actions, and interactive objects, or a meta-location encompassing data pertaining to a tangible location (e.g., a single screen in a 2D game, or the dimensions within a 3D game), in addition to all props located therein, along with any additional stages contained within that location (substages). Thus if props give additional social role functionality to objects, then the stage gives additional social actions to the environment.

An example of procedures with included props would be, "have a drink," which is different depending on the location. At the cafe, the drink might be a latte while at a bar it is a beer. This process is similar to Smart Zones (Howard 2019), which are areas or zones that try and invoke the ambient life of a location, such as tavern goers, through encapsulation role actions and procedures within a zone collection. The difference between Smart Zones and the proposed solution within this paper is that I wish to focus more on the social character rather than creating ambient characters that only exist to create living scenes. However, a stage is not solely comprised of these attributes or collections, nor is it a singular stage; a stage could be many stages combined.

Substages Stages[3] can contain stages within them. Subordinate stages–or substages–are exactly like stages; however, they are housed within the framing of another stage. Like acamera focusing on a bar conversation between two people in a tavern. Though the tavern is there, the two characters and their performance are the focus. While it might seem unnecessary to care about the tavern, note that the tavern still flavors the interaction, the two characters could be drinking a strong drink versus if they were speaking at the bartop of a cafe. Though the interaction may remain the same, the setting can change the conversation's mood.

Procedural Generation and Graph Encodings To encode stages into a data structure, a straightforward way would be to use ideas from procgen research, like using graphs to represent spatial data The method for generating dungeons and worlds within Unexplored focus on creating Directed Acyclic Graphs (DAG), which is the basis for all the dungeons within Unexplored. Additionally, Cooper's Sturgeon-GRAPH (Cooper 2023) is another system, which uses a WaveFunctionCollapse-like method to generate location/level graphs. With graphs, my prospective system would gain additional advantages such as encoding each digital prop as a node with edges that define which characters in the scene can use said prop; additionally graphs can be queried to find information that might not be readily available to a designer (assuming a mixed-initiative tool of my system was created). These preliminary graphs could contain information such as interactables on the stage, the characters that are present, and so forth.

As for the graphs themselves, they would begin with a root node, either the world or stage, and in a more tree-like structure descend from abstract concepts like "cafe" or "tavern" to property nodes such as the owner of the cafe or tavern to then terminating with prop nodes such as the range of chairs, the decor, static and dynamic props, and the different exit and entrance points. At the terminating node, characters can also be added to the location node to show which characters are currently there. Additionally, there would be CiF-like sets of production rules and procedures that would apply to certain node types, like cafes and taverns. These location rules would describe the actions or procedures a character could take. Finally, both the graph and the production collection would be shared with other AI systems, much like how smart objects are shared with other systems. By sharing graphs, different systems can work together to create more nuanced character performances and more consistent narratives. A current weakness of this approach is that there is currently no strong ontology for what should be contained within a stage graph, and it is still in its preliminary stages.

Conclusion

With the additional salience layer of the environment and its objects, the characters can have even more nuanced performances through knowing about their stage or the props in hand. To implement both the props, stages, and sub-stages, I borrow from previous works of smart objects, location graph encodings, and CiF to prototype a system that cares for the stage in a first-class way. I also describe why character roles are still important as they act as an interface layer between the character and the stage as well as the props themselves. With this system, I hope to explore different avenues of design spaces, such as computational environmental storytelling or computational scenography.

[3]While I use the word stages and substages, sets and stages may also work, however, more interrogation between the two sets of terms is required.

References

Bethesda Game Studios; and Virtuos. 2025. The Elder Scrolls IV: Oblivion Remastered. Remastered version of the 2006 video game.

Card, O. S. 1988. *Characters & Viewpoint (Elements of Fiction Writing)*. Writer's Digest Books.

Carson, D. 2004. Environmental Storytelling: Creating Immersive 3D Worlds Using Lessons Learned from the Theme Park industry.

Chen, S.; Nelson, M.; and Mateas, M. 2009. Evaluating the Authorial Leverage of Drama Management. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 5(1): 136–141. Number: 1.

Cooper, S. 2023. Sturgeon-GRAPH: Constrained Graph Generation from Examples. *Proceedings of the 18th International Conference on the Foundations of Digital Games*.

de Sevin, E.; Chopinaud, C.; and Mars, C. 2019. Smart Zones to Create the Ambience of Life. In *Game AI Pro 360*. CRC Press.

Dormans, J.; and Bakkes, S. C. J. 2011. Generating Missions and Spaces for Adaptable Play Experiences. *IEEE Transactions on Computational Intelligence and AI in Games*, 3: 216–228.

Evans, R.; and Short, E. 2014. Versu—A Simulationist Storytelling System. *IEEE Transactions on Computational Intelligence and AI in Games*, 6(2): 113–130.

Garbe, J. 2020. *Increasing Authorial Leverage in Generative Narrative Systems*. Ph.D. thesis, UC Santa Cruz.

Howard, P. 2019. *What is Scenography?* London: Routledge, 3 edition. ISBN 9781138504783.

Kallmann, M.; and Thalmann, D. 1998. Modeling Objects for Interaction Tasks. In *Computer Animation and Simulation*.

Kemenczy, T. 2014. Scenography of Kentucky Route Zero. In *Game Developers Conference (GDC)*. GDC Vault.

Larian Studios. 2023. Baldur's Gate 3. [Video game]. PC, PlayStation 5, Xbox Series X/S.

Liao, F.; Gelardi, G.; Mitchell, K.; Sandhu, A.; and McCoy, J. 2023. Sunset Valley: A Case Study in Computational Gossip. *2023 IEEE Conference on Games (CoG)*, 1–5. Conference Name: 2023 IEEE Conference on Games (CoG) ISBN: 9798350322774 Place: Boston, MA, USA Publisher: IEEE.

Mateas, M. 1999. An Oz-Centric Review of Interactive Drama and Believable Agents. In *Artificial Intelligence Today*.

Mateas, M.; and Stern, A. 2002. A Behavior Language for Story-Based Believable Agents. *IEEE Intelligent Systems*, 17(4): 39–47. Place: USA Publisher: IEEE Educational Activities Department.

Mateas, M.; and Stern, A. 2005. Facade.

Mccoy, J.; Treanor, M.; Samuel, B.; Reed, A. A.; Mateas, M.; and Wardrip-Fruin, N. 2013. Prom Week: Designing past the game/story dilemma. In *International Conference on Foundations of Digital Games*.

Mccoy, J.; Treanor, M.; Samuel, B.; Wardrip-Fruin, N.; and Mateas, M. 2011. Comme il Faut: A System for Authoring Playable Social Models. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*.

Protsenko, D.; Dormans, J.; and van Rozen, R. 2025. Maintenance in Procedural Level Design: Lessons from Ludoscope. *Proceedings of the 20th International Conference on the Foundations of Digital Games*.

Sandhu, A.; Chen, Z.; and McCoy, J. 2019. Enhancing wave function collapse with design-level constraints. *Proceedings of the 14th International Conference on the Foundations of Digital Games*, 1–9. Conference Name: FDG '19: The Fourteenth International Conference on the Foundations of Digital Games ISBN: 9781450372176 Place: San Luis Obispo California USA Publisher: ACM.

Sandhu, A.; and McCoy, J. 2019. A framework for integrating architectural design patterns into PCG. *Proceedings of the 14th International Conference on the Foundations of Digital Games*, 1–5. Conference Name: FDG '19: The Fourteenth International Conference on the Foundations of Digital Games ISBN: 9781450372176 Place: San Luis Obispo California USA Publisher: ACM.

Sandhu, A.; and McCoy, J. 2023. Exploring the Union Between Procedural Narrative and Procedural Content Generation. volume 14384, 255–262. Cham: Springer Nature Switzerland. ISBN 978-3-031-47657-0 978-3-031-47658-7. Book Title: Interactive Storytelling Series Title: Lecture Notes in Computer Science.

Sandhu, A.; Mitchell, K.; and Mccoy, J. 2021. TileTerror: A System for Procedurally Generating 2D Horror Maps.