



Exploring the Union Between Procedural Narrative and Procedural Content Generation

Arunpreet Sandhu^(✉) and Joshua McCoy

UC Davis, Davis, CA 95616, USA
{asisadhu, jamccoy}@ucdavis.edu

Abstract. Level Generation and Narrative Generation have often been separated from each other, even though they are both forms of Procedural Content Generation (PCG). The union of these between level generation and narrative generation, however, is under-explored given their history, and by exploring the union between these two topics will lead to new forms of PCG-enabled gameplay and stories. To address this under-exploration, we list both narrative and level generators and a description of a possible combination of these systems through shared knowledge representation structures. We also lay out various evaluation dimensions for these combined systems, which include level generation evaluation, expressive range, procedural narrative evaluation, and believable character evaluation. Finally, we describe the possible new frontiers that this approach enables, such as generating areas based on character personality traits.

Keywords: Procedural Content Generation · Emergent Narrative · Environmental Storytelling

1 Introduction

Procedural Content Generation (PCG) is an umbrella term containing both level generation and narrative generation techniques; yet the union between these two has been rarely explored. In this paper, we explore the union between these two PCG areas, and how the union enables indexical storytelling—a sibling to environmental storytelling.

PCG often—in academics—describes processes that help designers save time and have more resources for development. Cook [1], however, proposes another way to describe PCG, stating: “A better justification for PCG is that it enables new kinds of game design, by changing the player’s relationship with the game’s systems, or allowing the developers to work at a scale that would be otherwise impractical.” Cook furthers this view by demonstrating how PCG gives players a different experience due to the PCG-oriented gameplay design [2]. Games such as *Caves of Qud* [3], *Dwarf Fortress* [4], or *Ultima Ratio Regum* [5] are exemplars of how PCG allows for a scale of gameplay that would not be achievable otherwise.

This type of PCG viewpoint described by Cook also captures emergent narrative as described by Ryan [6], where world’s simulation creates a narrative, but only in the retelling of the events does it become a story; Ryan calls his approach a curatorial

approach to simulated story worlds. Ryan continues by saying that this curation can be done procedurally as well, which has led to systems focused on sifting through the simulated worlds, looking for emergent narrative, such as *Winnow* [7] and *Felt* [8]. These systems often focus on characters and character actions, but there is an additional layer of believability that pertains to space, as these characters must exist within a virtual environment, and how characters act within that space reveals more information about them that can help convince the player that these characters are believable. Take for example a man dressed in armor standing in front of a medieval gate with others dressed similarly, all at attention.

While there was no mention of a “city guard,” it is not too difficult to make the leap that the character described is a city guard, protecting the gate with their fellow guards¹. While much of the example has the character doing an action, it’s only within the context of the world around them does the belief of, “this character is a guard,” get further grounded. Yet, this isn’t the only function that space has as it can help tell the player what type of experience the virtual world holds for them; an eerie abandoned town has a much different character than a vibrant thriving city. The story world can help to improve the audience’s believability of character action by the way the world is presented to the player. This presentation of the story world and how the player experiences, is to us, environmental storytelling; where the story world conveys information about the story through the world. This, we believe, is at the core of environmental storytelling, which Carson states in his piece on environmental storytelling [9]: “One of the trade secrets behind the design of entertaining themed environments is that the story element is infused into the physical space a guest walks or rides through.”

While Carson at the time was describing theme parks, suffusing places with the story element for visitors to traverse, his description of environmental storytelling has been seen within games. Game worlds often attempt to enrich the story through its design, like in *Bioshock Infinite*’s city design of Columbia—the location that the majority of the game’s story takes place in. Totten [10]—in his book, *Architect’s Approach To Level Design*—describes portions of Columbia as: “buildings and islands form large-scale rhythms that draw players from one climactic battle to another. However, players who explore individual structures will find richly developed reward spaces and embedded narrative spaces.” Levels—and the world for that matter—offer a chance to convey a narrative through their construction, props, and visuals, much like Carson’s environmental storytelling; but adhering to Carson’s approach might be a pitfall for a game developer as it does not capture the entire experience; it fails to capture how players interact with space. Instead, Fernández-Vara [11] argues that games should adopt an indexical style to environmental story telling.

Indexical storytelling allows players to discover a narrative by ordering what they find in the game world, like exploring the streets of Columbia, or collecting pieces of information for a murder mystery. Here the player can construct their own narrative, one which isn’t reliant on visiting a place first. Instead, they may pursue the goal in whichever manner they deem fit. We believe that the combination of different PCG systems and areas, such as narrative and level generators, can make indexical storytelling possible.

¹ We assume that there is no deeper social game being played in this example, such as a group of rogues pretending to be guards to fool others, or any other kind of social engineering.

In this paper, we contribute an exploration of the combination of narrative and level generators. First, we describe both PCG algorithms that focus on generating levels, then we describe narrative generation techniques. We then describe the benefits of combining both narrative generation and level generation algorithms together, which come from a system that we are currently developing but have yet to test. We also give possible ways to evaluate this approach.

2 Related Work

Kybartas [12] conducted a survey on narrative generation techniques that touches on story world generation, but the description given differs from the world generation we discuss in this paper when discussing level-oriented PCG. Instead, Kybartas describes space as the world in which the narrative happens, whereas level and world generators prioritize creating game worlds care more for game mechanics than story constraints. Games such as *Spelunky* [13] or *Dead Cells* [14] exemplify this viewpoint by generating levels solely for gameplay. Yet, an example of space being used for story telling would be more akin to how drama managers often use space. *Façade* [15], a short interactive drama, has its characters (Trip and Grace) interact objects within their apartment, using these objects as conversation with the player.

As for level generators, their focus is often on generating intractable game worlds that are conducive to gameplay. Often level generators are search-based approaches, grammars, rules, and others, which have all been well-documented in *Procedural Content Generation in Games* [16]. Other techniques that have become well known are constraint solvers like WaveFunctionCollapse (WFC) [17], which allows for constraint satisfaction and statistical learning to generate levels. In this approach, we use WFC as it shares a knowledge representation abstraction that is cohesive to how narrative generators represent story beats and WFC has been used to generate non-gameplay constraints such as genre [18] or aesthetics [19]. We still list out other algorithms as they are still valid candidates for combining with narrative generators.

As for narrative generators, these include planners—such as Sabre [20] and Glaive [21], which causal planning to ensure narrative coherency—storylet systems such as Lume [22], and story sifters like Winnow [7] and Felt [8]. There are many other approaches, as described in Kybartas’s survey, but for our use case we have found these three are a solid foundation for our integrated framework. While there has been much work on narrative planners, their usual graph-based knowledge representation seemed the most appealing for our approach. Storylet systems, on the other hand, treat narrative beats as composable units to generate coherent narratives. They often require story history to use as a world fact to then input as preconditions to each narrative beat; they will then choose an appropriate narrative beat that fills the story requirements. These systems require a critical mass of generative material before the systems generate believable narratives [22]. Story sifters, such as Winnow and Felt, look for interesting narrative locations while the world is being created. The output of these systems could then be connected to other approaches/systems for narrative generation.

Interesting, the closest system that treats indexical storytelling generation comes from Dormans’s work, which looks at space and tries to generate it with a cohesive “feel”

such as generating a temple within the game [23]. Other related work that attempts to tackle indexing the world comes from the Liquid Narrative Lab has research on indexing the world for better understanding of narrative constraints [24] that is similar to the indexical storytelling that is discussed in this paper. The closest related work we have been able to find is Bidarra’s work on generating narrative worlds that use a story as it’s starting position [25]. We differentiate ourselves by using the world generation as *input* for the story—as in not constraining our generative process with an already generated story. We instead approach this more with a curatorial view of emergent narrative.

3 Emergent Story World Generation

“Another example of “cause and effect” is the use of what I call “Following Saknussem.” Derived from the story *Journey to the Center of the Earth* by Jules Verne. In Verne’s story the main characters follow a trail of symbols scratched into subterranean walls by their adventuring predecessor, a sixteenth century Icelandic scientist, Arne Saknussem. In this way, the game player is pulled through the story by following “bread crumbs” left behind by a fictitious preceding game character. Whether you create notes scattered throughout your environments, or have the game player follow the destructive path of some dangerous creature, “cause and effect” elements will only heighten the drama of the story you are trying to tell!” – Carson [9]

The above quote is effectively what we wish to achieve with this approach; a world in which the characters can modify the world—and thus the story world as well². Often, however, the issue with this embedding of “bread crumbs” is that a designer placed them in the world—there was not a character who performed the act procedurally. This designed story often means the audience will only have one way to discover story, which Fernández-Vara would argue as environmental storytelling and not indexical storytelling. Fernández-Vara states that game spaces should not have this singular way through them, but allow the player to piece together information, indexing the world as they wish. By indexing the space, the player creates a narrative that is unique to their play trace, giving the player a more memorable experience that can then be recounted to others. Fernández-Vara’s version of storytelling is our goal, to create indexed environmental stories for players to create their own stories by interacting with the environment.

In terms of achieving indexical storytelling, we suggest that there needs a common knowledge representation that both the narrative system and world/level generator share. As in, with a graph-based approach, such as a narrative planner and WaveFunctionCollapse (which can be considered a constraint graph solver); these two systems have the same type of knowledge representation of graphs³. The issue then is connecting the two systems through this shared representation.

² While Carson speaks of characters and their actions, we consider this apart of the narrative as divorcing character from narrative and narrative from character is a tenuous task.

³ We do assume that the knowledge each graph contains is different; the narrative planner’s graph would contain narrative beats and links between them while WFC would treat each tile as a graph or possibly a higher level structure, such as a dungeon graph where each node is a room and each link is a path.

Currently, our approach focuses on the knowledge representation that both narrative generation and level generation share. We do so by using a graph-based representation between level generation (through WFC) and a narrative planner. For our narrative planner, we use space generation as a pre-condition for narrative nodes. As for a location beat to occur, the location must exist for that beat to occur, and it must occur in that location. As in if the narrative system's beat generation requires a temple (akin to one as described in Dormans's *Theory of The Place* [23]) then there must be a temple that is generated within the world. As for this type of generation, currently we are generating all locations as we test the connection between narrative planner and level generator.

We are experimenting with the inverse, as in if the level/world generator generates a larger structure (like a temple) that the narrative planner may not know about—as in the level generator understands this higher order pattern⁴ (like a port, or a highway, or some landmark) that the narrative planner does not contain—then then this information should be passed into the narrative generator, allowing the planner to take this new landmark and possibly generate stories for that location.

Another solution would be to give both the narrative and level generation an intermediate representation (IR), much like how compilers create an IR structure from source code. By having an intermediate symbol set, any system can target this approach, assuming the other system has implemented a similar symbol set representation. This approach is similar to how EM-Glue [26] connects system together, but we are describing is more akin to a Language Server Protocol approach⁵.

If the IR approach was taken, then other systems could be introduced to the generative process such as including character AI systems to flavor the generative process by using character traits to help decorate NPC interiors⁶. Examples of character AI systems are Fatima Modular [27] or PsychSim [28], which could drive both the character's decision making, but also inform the generation systems through the symbol set that a character would like to move from their home to another, which could then be generated for the characters. This generation step could even be constrained by social rules through a system like *Comme Il Faut* [29]. By including this step, our system explores questions such as how does one map personality to place, or social implications into levels. While we think this approach would be more extendable, we are placing this IR approach as future work as we do not know *what* the IR structure would look like until we build one possible system and evaluate it.

As for evaluations, we are still determining the best approach, as our approach is still being iterated on. Since this approach combines both narrative and level generation techniques, currently we are considering using both level generation evaluations, such as those proposed by Smith [30]. We also believe this system should be evaluated according to the believable character dimension, mostly due to our simulation/curatorial approach. Currently we are deciding between Jhala's description of believable character metrics [31] and those described by Aljammaz [32].

⁴ As described in the previous footnote, WFC's graph could contain a "temple" node. As for generating this temple, our approach could do something as TileTerror's generation process.

⁵ <https://microsoft.github.io/language-server-protocol/>

⁶ An example of this type of character to location generation would be akin to how a room for a disorganized person would look versus an organized person. This.

Currently we are leaning towards Jhala's as there is more one-to-one correspondence between Totten's level design book. Still though, there is another dimension of evaluation akin to an expressive range analysis of narrative worlds, much like the work of Kybartas [33]. Yet, unlike Kybartas's analysis where conflict of character states is the evaluation metric, the world's total narrative conflict potential could be used as a metric. As for which evaluation would be the best, we are still deciding as we continue to work on our approach, however, we are almost certain that a suite of metrics will be needed.

Thus, the goal of this integrated system is to bridge the gap between narrative and world generation and create an appropriate evaluation for the system. We also aim to create some more generalized language for this system so other AI systems can be included in the generation process. Currently, we are still understanding and exploring this space through our approach. By combining these two areas of content generation, we hope to create more believable worlds on the scale that PCG can provide, while maintaining narrative coherency given the player's actions.

4 Conclusion

In this paper, we describe the differences between level generation and narrative generation and the gap between these two research areas. Next, we describe environmental storytelling, then indexical storytelling, and how our system aims to approach a procedural indexing system. We then describe a framework for using WFC and a narrative planner to create an integrated system. We then describe a series of evaluations that could be used to determine if the system's output is credible. This connection between narrative and level generators, we believe, will expose new potential spaces and scale of play and story that can only be possible through PCG.

References

1. Cook, M.: Generative Forensics: Procedural Generation and Information Games. *ArXiv*. (2020). Accessed 01 Aug 2023. <https://www.semanticscholar.org/paper/bbdc570e38f9db14538f967ba62dcd5f91ea5b08>
2. Smith Nicholls, F., Cook, M.: 'That darned sandstorm': a study of procedural generation through archaeological storytelling. *Proceedings. 18th International Conference Found. Digit. Games*, pp. 1–8 (2023). <https://doi.org/10.1145/3582437.3587207>
3. Freehold Games. "Caves of Qud." (2019)
4. Bay 12 Games. "Dwarf Fortress." (2006)
5. Johnson, M.: *Ultima Ratio Regum*. Independent (2023)
6. Ryan, J.: *Curating Simulated Storyworlds*. UC Santa Cruz (2018). Accessed 23 Aug 2023. <https://escholarship.org/uc/item/1340j5h2>
7. Kreminski, M., Dickinson, M., Mateas, M.: Winnow: a domain-specific language for incremental story sifting. *Proceedings AAAI Conference Artif. Intell. Interact. Digit. Entertain.* **17**(1), 1 (2021). <https://doi.org/10.1609/aiide.v17i1.18903>
8. Kreminski, M., Dickinson, M., Wardrip-Fruin, N.: Felt: a simple story sifter. In: *Interactive Storytelling: 12th International Conference on Interactive Digital Storytelling, ICIDS 2019, Little Cottonwood Canyon, UT, USA, November 19–22, 2019, Proceedings*, Springer-Verlag, Berlin, Heidelberg, pp. 267–281 (2019). https://doi.org/10.1007/978-3-030-33894-7_27

9. Environmental Storytelling: Creating Immersive 3D Worlds Using Lessons. <https://www.gamedeveloper.com/design/environmental-storytelling-creating-immersive-3d-worlds-using-lessons-learned-from-the-theme-park-industry>. Accessed 23 Aug 2023
10. Totten, C.W.: An Architectural Approach to Level Design (2014). <https://api.semanticscholar.org/CorpusID:62473540>
11. Clara, F.-V.: Game spaces speak volumes: indexical storytelling. In: DiGRA - proceedings of the 2011 DiGRA international conference: Think design play, DiGRA/Utrecht School of the Arts (2011). <http://www.digra.org/wp-content/uploads/digital-library/Game-Spaces-Speak-Volumes.pdf>
12. Kybartas, B., Bidarra, R.: A survey on story generation techniques for authoring computational narratives. *IEEE Trans. Comput. Intell. AI Games* **9**(3), 239–253 (2017). <https://doi.org/10.1109/TCIAIG.2016.2546063>
13. Mossmouth Games. “Spelunky.” (2008)
14. Motion Twin. “Dead Cells.” (2017)
15. Mateas, M., Stern, A.: *Facade* (2005)
16. Shaker, N., Togelius, J., Nelson, M.J.: *Procedural Content Generation in Games: A Textbook and an Overview of Current Research*. Springer (2016). <https://doi.org/10.1007/978-3-319-42716-4>
17. Karth, I., Smith, A.M.: WaveFunctionCollapse is constraint solving in the wild. *Proceedings 12th International Conference Found. Digit. Games*, pp. 1–10 (2017). <https://doi.org/10.1145/3102071.3110566>
18. Sandhu, A., Mitchell, K., McCoy, J.: TileTerror: a system for procedurally generating 2D horror Maps. In: *AIIDE Workshops* (2021)
19. Charity, M., Togelius, J.: Aesthetic bot: interactively evolving game maps on twitter. *Proc. AAAI Conf. Artif. Intell. Interact. Digit. Entertain.* **18**(1), 1 (2022). <https://doi.org/10.1609/aiide.v18i1.21943>
20. Ware, S.G., Siler, C.: Sabre: a narrative planner supporting intention and deep theory of mind. *Proc. AAAI Conf. Artif. Intell. Interact. Digit. Entertain.*, **17**(1), 1 (2021). <https://doi.org/10.1609/aiide.v17i1.18896>
21. Ware, S.G., Young, R.M.: Glaive: a state-space narrative planner supporting intentionality and conflict. In: *Proceedings of the 10th AAAI International Conference on Artificial Intelligence and Interactive Digital Entertainment*, pp. 80–86 (2014)
22. Mason, S., Stagg, C., Wardrip-Fruin, N.: Lume: a system for procedural story generation. In: *Proceedings of the 14th International Conference on the Foundations of Digital Games*, ACM, San Luis Obispo California USA, pp. 1–9 (2019). <https://doi.org/10.1145/3337722.3337759>
23. J. D. B. October 26 and 2021, *The Theory of The Place: A level design philosophy for Unexplored 2*. Game Developer (2021). <https://www.gamedeveloper.com/blogs/the-theory-of-the-place-a-level-design-philosophy-for-unexplored-2>. Accessed 23 Aug 2023
24. Fisher, M.: Narrative planning in large domains through state abstraction and option discovery. *Proc. AAAI Conf. Artif. Intell. Interact. Digit. Entertain.*, **18**(1), 1 (2022). <https://doi.org/10.1609/aiide.v18i1.21979>
25. Balint, J.T., Bidarra, R.: Procedural generation of narrative worlds. *IEEE Trans. Games* **15**(2), 262–272 (2023). <https://doi.org/10.1109/TG.2022.3216582>
26. Mori, G., Thue, D., Schiffel, S.: EM-Glue: a platform for decoupling experience managers and environments. *Proc. AAAI Conf. Artif. Intell. Interact. Digit. Entertain.*, **18**(1), 1 (2022). <https://doi.org/10.1609/aiide.v18i1.21972>
27. Dias, J., Mascarenhas, S., Paiva, A.: FATiMA modular: towards an agent architecture with a generic appraisal framework. Bosse, T., Broekens, J., Dias, J., Van Der Zwaan, J. (eds.), *In Lecture Notes in Computer Science*, **8750**. Springer International Publishing, Cham, pp. 44–56 (2014). https://doi.org/10.1007/978-3-319-12973-0_3

28. Marsella, S.C., Pynadath, D.V., Read, S.J.: PsychSim: Agent-Based Modeling of Social Interactions and Influence p. 6
29. McCoy, J., Treanor, M., Samuel, B., Mateas, M., Wardrip-Fruin, N.: Comme il Faut: A System for Authoring Playable Social Models, p. 6
30. Canossa, A., Smith, G.: Towards a Procedural Evaluation Technique: Metrics for Level Design
31. Gomes, P., Paiva, A., Martinho, C., Jhala, A.: Metrics for character believability in interactive narrative. In: Interactive Storytelling, Koenitz, H., Sezen, T.I., Ferri, G., Haahr, M., Sezen, D., Çatak, G. (eds.), Springer International Publishing, Cham, pp. 223–228 (2013). https://doi.org/10.1007/978-3-319-02756-2_27
32. Aljammaz, R., Wardrip-Fruin, N., Mateas, M.: Towards an understanding of character believability. In: Proceedings of the 18th International Conference on the Foundations of Digital Games, ACM, Lisbon Portugal, pp. 1–9 (2023). <https://doi.org/10.1145/3582437.3582466>
33. Kybartas, B., Verbrugge, C., Lessard, J.: Expressive Range Analysis of a Possible Worlds Driven Emergent Narrative System. In: Rouse, R., Koenitz, H., Haahr, M. (eds.) Interactive storytelling, pp. 473–477. Springer International Publishing, Cham (2018)