

Hloockey: An application of story sifting to Blaseball-like emergent narrative

Isaac Karth
Liminal Entertainment, Inc.
Los Angeles, CA, USA
isaac.karth@gmail.com

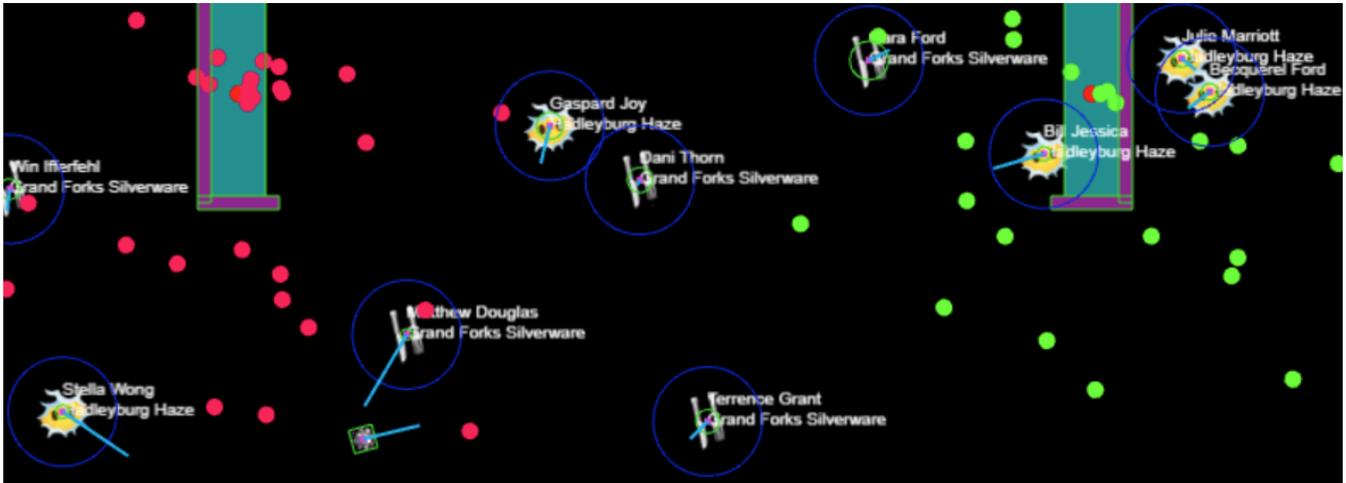


Figure 1: A game of Hloockey in progress. Players are represented with icons, with a blue circle indicating their hlocky stick reach and a cyan vector indicating current motion. Hadleyberg Haze vs. Grand Forks Silverware. The green and red dots indicate attempted shots.

ABSTRACT

Hloockey is a Blaseball-inspired sports simulation that uses story sifting to generate narratives out of the transcript of events generated by the simulation. Hloockey was created for National Novel Generation Month (NaNoGenMo) as an emergent simulation of continuous real-time events, tracking the collision logics to generate a raw transcript. The Hloockey game rules deliberately emphasize the characteristics of the underlying physics engine, including its flaws, in order to maximize the emergent characteristics of the simulation. The system uses an existing domain-specific story-sifting language to conduct partial-match story sifting, extracting narrative arcs from the simulated events. It synthesizes the sifted story patterns into a broadcast commentary of the sports games, narrating the tournament as it plays out to its conclusion.

CCS CONCEPTS

• **Applied computing** → *Media arts*; • **Computing methodologies** → **Natural language generation**; **Information extraction**; Discourse, dialogue and pragmatics.

KEYWORDS

story sifting, text generation, blaseball, emergent narrative, sports commentary, natural language generation

ACM Reference Format:

Isaac Karth. 2024. Hloockey: An application of story sifting to Blaseball-like emergent narrative. In *Foundation of Digital Games '24: 19th International Conference on the Foundations of Digital Games, Games and Demo Track, May 21–24, 2024, Worcester, MA USA*. ACM, New York, NY, USA, 3 pages.

1 INTRODUCTION

Hloockey is a simulated sport that is almost entirely unlike Hockey, live-narrated narrated by a commentator. 30 teams compete in a series of the physics-simulated sports games, vying to win the Hloockey tournament as violence escalates.

Hloockey was originally created in 2023, for the 10th anniversary National Novel Generation Month (NaNoGenMo 2023). NaNoGenMo is an annual, month-long event, where participants write code that writes a novel. Previous novels have included adventure-game simulations,¹ a transcription of all of the colors in Georges Seurat's *Un dimanche après-midi à l'Île de la Grande Jatte*,² fanfiction-inspired melodrama,³ Nick Montfort's realignment of Shakespeare,⁴ and a great many reworkings of Moby Dick.

Hloockey combines two related traditions: Blaseball-inspired sports simulation and story sifting. to discover emergent narrative patterns in simulation histories.

¹<https://github.com/dariusk/NaNoGenMo/issues/39>

²<https://github.com/NaNoGenMo/2022/issues/50>

³<https://github.com/dariusk/NaNoGenMo-2015/issues/11>

⁴<https://github.com/NaNoGenMo/2016/issues/132>

2 SIMULATION

Blaseball (2020-2023)⁵ was a COVID-era baseball-like browser game developed and run by The Game Band [11]. In it, simulated teams play a horror-tinged version of baseball, with a week-long season the culminates in a weekend playoff tournament. Rather than each human observer having their own team, there was one global league that the human players all observed (and made bets on with in-game currency). Due to the ongoing developer orchestration *Blaseball* has been characterized as a perennial game [8], where the ongoing interaction between fans and developers affects the unfolding narrative. For our purposes, that interaction is related to the most important element of *Blaseball*: the way that fans found the narratives by observing (and indirectly influencing) the emergent results of the simulation.

Whereas *Blaseball* is presented to observers as a turn-based experience, *Hockey* is deliberately surfaced as a real-time simulation, leaning on the emergent results to generate narratively interesting configurations of events.

The core *Hockey* simulation is deliberately reductive, modeling players as akin to slightly-steered billiard balls on a low-friction surface. Two teams of players face off on the rink, attempting to knock a puck (another physics object) into their team’s goal (a static zone blocked off by walls on three sides).

Players have attributes—such as fear, distractibility, and grit—which affect their performance on the ice rink. The very-low-friction environment means that their simulated behavior tends to resemble that of a perfectly spherical hockey player. In addition to attempting to get the puck in the right goal, the players also come into conflict with each other; conflict that escalates as tempers flare and grudges build over the course of the tournament.

The flaws of the physics simulation are deliberately included in the game design. Clipping outside the bounds of the rink is possible at high speeds, and rather than fixing this it was incorporated as a feature, with the penalty box replaced by having absent players wander in the void beyond until recalled to the rink, echoing some of *Blaseball*’s existential horror elements.

3 SIFTING

One drawback of emergent narratives is that too much happens: many emergent simulations derive much of their depth from the sheer number of events that occur, making it difficult for an outside observer to discern any coherent narrative in the chaos. Story sifting [15] is one way to corral the content into manageable patterns. In contrast to top-down narrative generation, story sifting takes a bottom-up approach [7, 16].

There have been many different approaches to implementing story sifting, including using a domain-specific language [4], Dynamic Time Warping [9], and ChatGPT [12]. It has been applied to identifying narratives during live performances in *Bad News* [17] and for identifying potential stories for a drama manager [1].

In this case, we use story sifting to generate sports commentary for the simulated game. Sports commentary generation in commercial games are massive endeavors that can involve tens of thousands of recorded voice lines [18]. Despite its common use in commercial sports videogames, as a research topic commentary generation has

⁵<https://blaseball.com/>

```
(pattern scoreGoalOtherTeam
  (event ?e2 where
    event: sportsGoalScored
    target: ?puck
    actor: ?actor
    actorTeam: ?team2
    scoringTeam: ?team1
    (not= ?team1 ?team2)
    (not= ?team2 puck)))
```

Figure 2: An example of a Winnow sifting pattern. `scoreGoalOtherTeam` captures the variables when a `sportsGoalScored` collision event happens where the player making the shot is not on the scoring team, indicating that they shot the puck in the wrong goal.

been comparatively under-studied. However, there are existing research projects for generating commentary for game as far apart as *League of Legends* [14] and real-world baseball games [3].

Commentary for non-sports games such as *Minecraft* [2] and *Getting Over It with Bennet Foddy* [10] have been learned by training on video footage of human commentators. This suggests an intriguing possibility. If we invert this process, the audience for game commentary is no longer confined to humans: being able to generate a transcript for a game is potentially useful as a further source of training data for future AI interpretations of that game.

The emergent nature of the sports simulation means that a translation is required. In order to make the physics simulation tractable, it must be read [6]. The simulation is quantified by translating the continuous events of the game into a quantized timeline of events. The collision logics [19] of the simulation are at the core of the operational logics that package its rhetorical strategy and process [13]. This gives us a transcript of the basic events which we can use as the raw material to generate the announcer’s commentary.

A raw transcript doesn’t make for a compelling narrative. In order to translate it into something more comprehensible, *Hockey* uses Winnow [5] to implement story sifting. Winnow is a domain-specific language for incremental story sifting, allowing for the detection of partial story patterns while the simulation is running (Fig. 2).

The primary window on the players’ behavior is captured in the running log of collision events. The system runs a sifting pattern on the event log and feeds the results to a commentator, who turns them into utterances, describing the play-by-play of the game.

ACKNOWLEDGMENTS

Special thanks to Max Kreminski for team names and advising on story sifting implementation feedback.

REFERENCES

- [1] Ben Clothier and David E. Millard. 2023. Awash: Prospective Story Sifting Intervention for Emergent Narrative. In *Interactive Storytelling*, Lissa Holloway-Attaway and John T. Murray (Eds.). Springer Nature Switzerland, Cham, 187–207.
- [2] Matthew Guzdial, Shukan Shah, and Mark Riedl. 2018. Towards Automated Let’s Play Commentary. arXiv:1809.09424 [cs.AI]
- [3] Byeong Jo Kim and Yong Suk Choi. 2020. Automatic baseball commentary generation using deep learning. In *Proceedings of the 35th Annual ACM Symposium on*

- Applied Computing* (Brno, Czech Republic) (SAC '20). Association for Computing Machinery, New York, NY, USA, 1056–1065. <https://doi.org/10.1145/3341105.3374063>
- [4] Max Kreminski, Melanie Dickinson, and Michael Mateas. 2021. Winnow: a domain-specific language for incremental story sifting. In *Proceedings of the AAAI conference on artificial intelligence and interactive digital entertainment*, Vol. 17. 156–163. Issue 1.
- [5] Max Kreminski, Melanie Dickinson, and Michael Mateas. 2021. Winnow: A Domain-Specific Language for Incremental Story Sifting. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* 17, 1 (Oct. 2021), 156–163. <https://doi.org/10.1609/aiide.v17i1.18903>
- [6] Max Kreminski, Isaac Karth, and Noah Wardrip-Fruin. 2019. Generators that read. In *Proceedings of the 14th International Conference on the Foundations of Digital Games*. ACM, 1–7.
- [7] Max Kreminski, Noah Wardrip-Fruin, and Michael Mateas. 2023. Authoring for story sifters. In *The Authoring Problem: Challenges in Supporting Authoring for Interactive Digital Narratives*. Springer, 207–220.
- [8] Bjarke Alexander Larsen and Elin Carstensdottir. 2021. Wrestling with Destiny: Storytelling in Perennial Games. In *Interactive Storytelling*, Alex Mitchell and Mirjam Vosmeer (Eds.). Springer International Publishing, Cham, 236–254.
- [9] Wilkins Leong, Julie Porteous, and John Thangarajah. 2022. Automated Sifting of Stories from Simulated Storyworlds. *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence AI and Arts*, 7.
- [10] Chengxi Li, Sagar Gandhi, and Brent Harrison. 2019. End-to-end let's play commentary generation using multi-modal video representations. In *Proceedings of the 14th International Conference on the Foundations of Digital Games* (San Luis Obispo, California, USA) (FDG '19). Association for Computing Machinery, New York, NY, USA, Article 76, 7 pages. <https://doi.org/10.1145/3337722.3341870>
- [11] Todd Martens. 2020. It's a beautiful summer for 'Blaseball.' Yes, 'Blaseball,' where America's pastime meets 'D&D'. <https://www.latimes.com/entertainment-arts/story/2020-08-24/blaseball-online-browser-game-dungeons-dragons>
- [12] Gonzalo Méndez and Pablo Gervás. 2023. Using ChatGPT for story sifting in narrative generation. In *Proceedings of The 14th International Conference on Computational Creativity*.
- [13] Joseph C. Osborn, Noah Wardrip-Fruin, and Michael Mateas. 2017. Refining operational logics. In *Proceedings of the 12th International Conference on the Foundations of Digital Games* (Hyannis, Massachusetts) (FDG '17). Association for Computing Machinery, New York, NY, USA, Article 27, 10 pages. <https://doi.org/10.1145/3102071.3102107>
- [14] Noah Renella and Markus Eger. 2023. Towards automated video game commentary using generative AI. In *Proceedings of the AIIDE workshop on Experimental AI in Games*.
- [15] James Ryan. 2018. *Curating Simulated Storyworlds*. Ph. D. Dissertation. UC Santa Cruz.
- [16] James Owen Ryan, Michael Mateas, and Noah Wardrip-Fruin. 2015. Open design challenges for interactive emergent narrative. In *International Conference on Interactive Digital Storytelling*. Springer, Springer, Copenhagen, Denmark, 14–26.
- [17] Ben Samuel, Adam Summerville, James Ryan, and Liz England. 2021. A Quantified Analysis of Bad News for Story Sifting Interfaces. In *Interactive Storytelling*, Alex Mitchell and Mirjam Vosmeer (Eds.). Springer International Publishing, Cham, 142–156.
- [18] Dan Treadway. 2014. How are commentators programmed into sports video games? <https://web.archive.org/web/20140829000304/http://www.si.com/extramustard/2014/08/21/how-are-commentators-programmed-sports-video-games>
- [19] Noah Wardrip-Fruin. 2018. Beyond Shooting and Eating: Passage, Dys4ia, and the Meanings of Collision. *Critical Inquiry* 45, 1 (2018), 137–167. <https://doi.org/10.1086/699587> arXiv:<https://doi.org/10.1086/699587>