

Why Are We Like This?: The AI Architecture of a Co-Creative Storytelling Game

Max Kreminski

mkremins@ucsc.edu

University of California, Santa Cruz
Santa Cruz, California

Michael Mateas

mmateas@ucsc.edu

University of California, Santa Cruz
Santa Cruz, California

Melanie Dickinson

mldickin@ucsc.edu

University of California, Santa Cruz
Santa Cruz, California

Noah Wardrip-Fruin

nwardrip@ucsc.edu

University of California, Santa Cruz
Santa Cruz, California

ABSTRACT

We present *Why Are We Like This?* (WAWLT), a mixed-initiative, co-creative storytelling game in which two players develop a story transcript by selecting and editing actions to perform and narrativize in an ongoing simulation. In this paper, we lay out the major technical features of WAWLT's AI architecture—including story sifting via Datalog queries, social simulation, action suggestions, and player-specified but system-understandable author goals—and discuss how these features work together to produce a play experience that facilitates player creativity.

CCS CONCEPTS

• **Software and its engineering** → **Interactive games**; • **Applied computing** → **Computer games**.

KEYWORDS

story sifting, social simulation, emergent narrative, storytelling games, casual creators, player creativity, mixed-initiative co-creativity

ACM Reference Format:

Max Kreminski, Melanie Dickinson, Michael Mateas, and Noah Wardrip-Fruin. 2020. *Why Are We Like This?: The AI Architecture of a Co-Creative Storytelling Game*. In *International Conference on the Foundations of Digital Games (FDG '20)*, September 15–18, 2020, Bugibba, Malta. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3402942.3402953>

1 INTRODUCTION

Why Are We Like This? (WAWLT) is an AI-supported digital storytelling game. In WAWLT, one or more players (ideally two) work to construct a story, supported by an AI system that serves to provide players with inspiration and keep the story moving forward, even when players are unsure what should happen next.

WAWLT's current setting involves a small community of researchers who are temporarily stranded during a symposium at a remote venue. Stories revolve around the conflicts that emerge

between these characters due to misunderstandings and shared history. Characters in WAWLT use *story sifting patterns* to make narrative sense of the world. Because different characters have access to different sifting patterns, they tell themselves different stories about the events that have transpired—and these conflicting understandings lead them to act in conflicting ways. By choosing which of several possible narrative frames each character will adopt for the same events, and how characters will act on the world based on their understanding of past events, players guide the evolution of these emergent conflicts.

WAWLT represents an example of AI-based game design [6] inspired by the study of existing player storytelling practices [5, 12] in simulation-driven games. In designing WAWLT, we set out to provide creativity support features that would help players overcome four major barriers to creativity documented in [13]: fear of the blank canvas; fear of judgment; writer's block; and perfectionism. Further design inspiration was drawn from tabletop storytelling games [2, 19, 20], and from the AI-augmented improvisational theater experience *Bad News* [26]. WAWLT aims to support player storytelling practices by providing players with intelligent plot direction suggestions, drawn from a live social simulation and guided by player utterances in a machine-understandable *intent language*.

This paper briefly describes the WAWLT AI architecture and the roles of its key subsystems in supporting mixed-initiative storytelling. For more on the design of WAWLT, see [10].

2 RELATED WORK

WAWLT is a mixed-initiative co-creative [14] casual creator [4] for storytelling. Many of its storytelling mechanics were inspired by similar mechanics from tabletop storytelling games [2, 19, 20]. Other casual creators for storytelling, such as *Writing Buddy* [25], and other mixed-initiative co-creative storytelling systems, such as *Mimisbrunnur* [28], have provided valuable design inspiration for WAWLT, but have not fully embraced the use of a fine-grained simulated storyworld as we aim to here.

The same is true of co-creative writing processes driven by language models [3, 15, 21, 27]. Moreover, language model-based systems are particularly flawed from a creativity support perspective due to their lack of an explicit world model. Because of this lack, they frequently go off track or make suggestions that clearly contradict previous statements, forcing users to spend time and energy repeatedly reminding them of established facts via prompting

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
FDG '20, September 15–18, 2020, Bugibba, Malta
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8807-8/20/09.
<https://doi.org/10.1145/3402942.3402953>

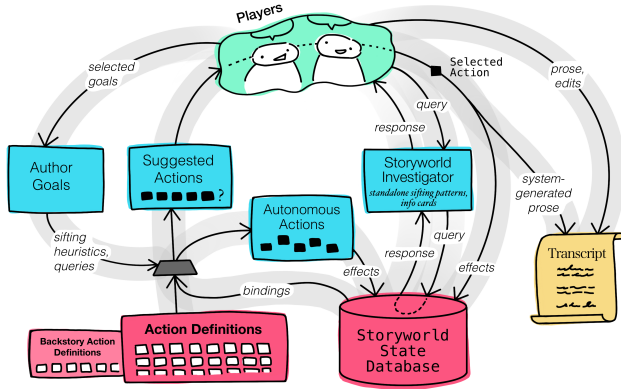


Figure 1: An overall system diagram of WAWLT, showing the important modules and data flows. Pink subsystems (action definitions and the storyworld state database) consist of inert data; blue subsystems (author goals, suggested actions, autonomous actions, and the storyworld investigator) act on this data; and the transcript emerges from player actions over the course of play.

techniques. This distracts from the useful creativity support features they provide (suggestions about “what happens next”) and exacerbates the problem of maintaining consistency, which even experienced authors may already find difficult on their own.

WAWLT is built around *story sifting* [22, 23] in both its implementation and its design, making central use of the story sifting and simulation engine Felt [11]. Story sifting approaches to emergent narrative attempt to address the challenges of narrative generation through curation: by allowing a simulated storyworld to run, generating a massive chronicle of mostly-uninteresting simulated events, and then searching within this chronicle for patterns of narratively compelling events, it is possible to provide players with compelling stories or microstories without baking knowledge of how to tell a compelling story directly into the simulation engine itself. Several existing play experiences [8, 9, 26] make use of story sifting technology to some extent, but we believe WAWLT is the first to center story sifting as a player-facing game mechanic.

3 EXAMPLE PLAY SESSION

At the start of a play session, two players sit down and generate a new WAWLT scenario. The game generates a storyworld state database containing an initial cast of characters and institutions and some basic relationships between them. It then performs *backstory simulation* to quickly generate a history for this cast.

Control is handed off to the players, who are prompted to choose some *author goals* and a subset of all the characters at the symposium to participate in the first *scene*. The players don’t yet know anything about the history of the storyworld, so they pick a couple of characters with interesting-sounding names and traits to participate in the first scene. They also select a single author goal: “cast suspicion on Vincent”, one of the two characters they selected.

The system prompts the players with five different *action suggestions*, all of which could motivate Vincent to pursue revenge against another character. The players discuss which of these actions to perform, then select an action in which Mikayla, the other character in the scene, makes a disparaging comment about Vincent’s research. A terse system-generated description of this action is added to the *transcript*, with an editable text box below it in which the players can write a more detailed description. The action’s *effects* are also run, updating the storyworld state database.

The players want to learn more about Mikayla and Vincent, so they use the *characters* tab of the *storyworld investigator* to look up Mikayla’s and Vincent’s *information cards*. They discover that Mikayla and Vincent are both doctoral students in the same lab who had previously worked on a major project together, but eventually both left the project over personal differences. This information helps the players develop a clearer picture of the relationship between Mikayla and Vincent up until this point, and allows them to write dialogue in their description of the insult action that makes reference to these past events.

After selecting several more system-suggested actions for Mikayla and Vincent to perform within this scene, the players decide that they have accomplished their current author goal of casting suspicion on Vincent. They end the scene, and are prompted to pick characters and author goals for a new scene. In the meantime, characters not participating in the first scene have performed a number of *autonomous actions*, guided by the players’ author goals. The players use the storyworld investigator to view these background events, and use this to guide their selection of participants and author goals for the second scene.

The players disagree briefly about whether the next scene should focus on establishing a conflict between two *characters* (Alex and Rashida) or between two *values* (comfort and safety). Eventually, they choose to compromise, selecting three author goals: one goal that explicitly focuses action suggestions on the establishment of a values conflict, and two goals that instruct the system to involve both Alex and Rashida in the plot as much as possible.

As the play session continues, the players’ focus shifts from investigating the history of the storyworld and opening up new possibilities toward pulling threads together to bring the story to a satisfying conclusion. Correspondingly, in early scenes, the players choose author goals focused on escalating tension and introducing new conflicts, but in later scenes, they select goals that steer action suggestions toward reconciliation between characters instead.

4 ARCHITECTURE

4.1 Storyworld State Database

The state of the WAWLT storyworld is stored in a DataScript [18] database managed by the Felt [11] story sifting and simulation engine. Storyworld entities represented in the database include characters, projects, institutions, events, relationships, and impressions. The latter two of these are especially interesting, and are described at greater length here:

A **relationship** entity stores one character’s view of another, including both *impressions* formed of that character’s actions and *role relationship* information (for instance, in the case of an advisor/student relationship or a marriage). It also contains a numeric

charge value representing the source character’s overall attitude toward the target, with positive values reflecting a positive attitude and vice versa. There are two relationship entities for every pair of characters, one pointing in each direction.

An **impression** represents a source character’s evaluation of a target character based on a particular introspection event. Each impression stores a pointer back to its *cause*: the event that led to its creation. One character’s overall charge toward another is given by the sum of the numerical scores associated with the source character’s impressions of the target. A single relationship may be defined by up to three positive and three negative impressions; stronger impressions displace weaker ones over time, and some actions (“gossip actions”) allow impressions to be communicated from one character to another. Many simulation-driven narrative systems [7, 17, 24] model character knowledge phenomena at a fine-grained level, tracking per-character awareness of many individual facts about the world. In *WAWLT*, we instead assume that all characters know about every event, but that most characters only care about events that are somehow directly relevant to them. Character *knowledge* is thus replaced in our system by subjective impressions, which are fewer in number and more reliably narratively consequential than granular facts about the world.

At the beginning of each play session, we initialize a cast of 10 characters with random traits and probabilistically set role relationships (“works for”, “advises”, etc) between characters where appropriate. We then run *backstory simulation* to flesh out several years of character life history and shared intellectual and social history among the cast. This bootstrapping of social context mitigates *fear of the blank canvas* [4, 13] in players by providing an interesting starting scenario, with many sites of narrative potential for players to develop through play. Backstory simulation draws from a pool of longer-timescale backstory actions that aren’t available during gameplay, such as “join institution”, “write paper together”, “publish a controversial book”, “take industry position”, and “take leave to parent”. After initial setup, subsequent simulation actions draw from a larger pool of shorter-timescale actions that make sense to happen in the course of a few days, on location at the symposium.

4.2 Author Goals

Author goals provide players with a *player intent language* [16] that they can use to explicitly communicate their current creative intent to the system. This intent language is one of the main novel features of the *WAWLT* architecture.

Examples of supported author goals include “Involve *character* in plot as *role*”; “Cast suspicion on *character*”; “Defuse tension between *character* and *character*”; and “Introduce false lead”. Italics in goal names indicate parameter slots. For instance, if players want to focus on actions involving a particular character, they can add this character to a *character* slot in any goal that provides one. Alternatively, players can also leave this slot empty, in which case the system will treat it as a wildcard that stands for “any character.”

Custom author goals that allow players to specify finer-grained constraints on action suggestions are also available. When specifying a custom author goal, players can use a lightweight textual query language to (for instance) prioritize actions that contain specific substrings in their taglines, feature specific characters, or

belong to specific categories (such as “introspection actions” or “actions that involve projects”). These criteria can also be combined in arbitrary ways.

Author goals are used to rank all currently possible actions in order to provide players with action suggestions. Each author goal is associated with a heuristic function that takes a possible action as an argument and returns a numerical score representing the relevance of this possible action to this author goal. To provide players with action suggestions, *WAWLT* first generates a list of all currently possible actions, then evaluates these possible actions against the set of currently active author goals. Every active author goal contributes a score to each possible action, and the list of possible actions is sorted by the total combined score from all active author goals, so that the most goal-relevant actions appear closest to the top of the suggestions list. This provides players with support in navigating the space of currently possible actions, which may contain hundreds of possible actions at any given point.

Author goals are also intended to help players negotiate their intent with one another by making this intent explicit. Because players have to agree on what author goals to select, and because they’re prompted to adjust their author goals at the start of every new scene, the system encourages players to regularly discuss their intent with one another, and players may have to explicitly argue for the things they want to have happen in the story. In this way, author goals can function similarly to the “palette” mechanism in tabletop storytelling games like *Microscope* [20], which requires players to explicitly discuss what they do and don’t want to include in the story they are creating together.

4.3 Action Suggestions

Character actions in *WAWLT* are defined as Felt [11] actions. They have preconditions, which take the form of Felt sifting patterns, and effects, which update the database when the action is performed.

Actions include both *external actions*, in which characters act on the world outside of themselves, and *introspection actions*, in which characters reflect on past events through the lens of a particular narrative frame. External actions might include “discuss a shared value with Rashida”, “insult Vincent”, or “sabotage Alex’s experiment”; introspection actions might include “speculate that Lea dislikes me” or “conclude that Mikayla is a mean person.” Introspection actions often produce *impressions* that influence the introspecting character’s attitude toward another character. By separating out character reasoning into a category of first-class actions that players can observe and perform directly, and that are added to the transcript for players to elaborate on within their stories, we hope to expose character reasoning to the players explicitly—and, thereby, to mitigate the *Tale-Spin effect* [29], in which a system is mistaken by players as less intelligent than it actually is due to insufficient exposure of the internal processes.

The action suggestion interface displays the five most highly ranked possible actions for characters who are involved in the current scene. Possible actions in the action suggestion interface are displayed alongside some details about why this action is currently possible, including a display of which active author goals contribute to the surfacing of this action suggestion; a tree of previous actions that causally contributed to this action; and a short description of

any other preconditions for this action, such as the presence of certain character traits on the action's protagonist.

The player may use a search bar to filter action suggestions without changing author goals. This search bar uses the same light-weight textual query language used in custom author goals.

4.4 Autonomous Actions

Characters who aren't involved in the current scene can perform actions autonomously in the background. Autonomous actions are influenced by author goals, but chosen via weighted random selection from several of the higher-scoring possible actions for offscreen characters, rather than selected by the players. Autonomous actions are intended to help mitigate writer's block by ensuring that the storyworld will always continue to develop in significant ways, even if the players can't think of anything interesting to do in a particular scene. We are inspired here by the design property of *incrementality* [13], which was found to support player storytelling in existing simulation-driven games.

4.5 Storyworld Investigator

The *storyworld investigator* provides players with fine-grained tools for investigating the history and current state of the storyworld while writing their stories. The investigator is divided into several tabs, each of which displays a complete list of all instances of a certain type of storyworld entity (characters, relationships, projects, institutions, and events) and lets the player view *information cards* containing more detailed information about these entities. Information cards are linked together with hyperlinks to enable rapid exploration of the web of storyworld entities. For instance, while viewing the information card for a particular character, links under the "Relationships" section allow rapid navigation to information cards containing more detailed information about this character's relationships with other characters, including the impressions that shape each character's perception of the other.

4.6 Transcript

The transcript holds a running record of "the story so far": an *artifact of play* [1] containing short system-generated summaries of all player-accepted action suggestions since the start of the play session, interleaved with more detailed player-generated prose descriptions of these actions. By allowing players to annotate events with their own descriptions, we aim to provide support for *extrapolative narrativization* [12]: a player storytelling behavior in which players seize on and elaborate minor details in the stories they tell about their play experiences, regardless of whether these details are explicitly modeled in the simulation. Editable free-text descriptions of events give players a place to decide for themselves what aspects of an event are most important, and to establish and reference recurring story elements not modeled in the computational system.

5 CONCLUSION

In this paper, we presented the AI architecture of *Why Are We Like This?* (WAWLT), an AI-supported storytelling game intended to provide explicit support for the kinds of player storytelling practices seen in many simulation-driven games. WAWLT makes extensive use of story sifting, both to implement character subjectivity

(through introspection actions) and to provide players with tools for investigating the history and current state of the storyworld (through the storyworld investigator). WAWLT also supports player storytelling practices by providing players with intelligent plot direction suggestions, drawn from an ongoing social simulation and guided by player utterances in a machine-understandable intent language, realized here in the form of *author goals*.

REFERENCES

- [1] Lea Albaugh, April Grow, Chenxi Liu, James McCann, Gillian Smith, and Jennifer Mankoff. 2016. Threadstading: playful interaction for textile fabrication devices. In *2016 CHI Extended Abstracts*. 285–288.
- [2] Avery Alder. 2013. The Quiet Year. buriedwithoutceremony.com/the-quiet-year.
- [3] Botnik. 2017. Predictive Writer. botnik.org/apps/writer.
- [4] Kate Compton and Michael Mateas. 2015. Casual creators. In *Proc. International Conference on Computational Creativity*. 228–235.
- [5] Mirjam Palosaari Eladhari. 2018. Re-tellings: the fourth layer of narrative as an instrument for critique. In *Proc. ICIDS*. Springer, 65–78.
- [6] Mirjam P Eladhari, Anne Sullivan, Gillian Smith, and Josh McCoy. 2011. *AI-based game design: enabling new playable experiences*. Technical Report. UC Santa Cruz Baskin School of Engineering, Santa Cruz, CA.
- [7] Richard Evans and Emily Short. 2013. Versu—a simulationist storytelling system. *IEEE TCAIG* 6, 2 (2013), 113–130.
- [8] Jacob Garbe. 2018. Simulation of history and recursive narrative scaffolding. project.jacobgarbe.com/simulation-of-history-and-recursive-narrative-scaffolding.
- [9] Max Kreminski, Devi Acharya, Nick Junius, Elisabeth Oliver, Kate Compton, Melanie Dickinson, Cyril Focht, Stacey Mason, Stella Mazeika, and Noah Wardrip-Fruin. 2019. Cozy Mystery Construction Kit: prototyping toward an AI-assisted collaborative storytelling mystery game. In *Proc. FDG*.
- [10] Max Kreminski, Melanie Dickinson, and Noah Wardrip-Fruin. 2020. Why Are We Like This?: Exploring Writing Mechanics for an AI-Augmented Storytelling Game. In *Proc. Electronic Literature Organization*.
- [11] Max Kreminski, Melanie Dickinson, and Noah Wardrip-Fruin. 2019. Felt: a simple story sifter. In *Proc. ICIDS*. Springer, 267–281.
- [12] Max Kreminski, Ben Samuel, Edward Melcer, and Noah Wardrip-Fruin. 2019. Evaluating AI-based games through retellings. In *Proc. AIIDE*, Vol. 15. 45–51.
- [13] Max Kreminski and Noah Wardrip-Fruin. 2019. Generative games as storytelling partners. In *Proc. FDG*.
- [14] Antonios Liapis, Georgios N Yannakakis, Constantine Alexopoulos, and Phil Lopes. 2016. Can computers foster human users' creativity? Theory and praxis of mixed-initiative co-creativity. *Digital Culture & Education* 8, 2 (2016), 136–153.
- [15] Enrique Manjavacas, Folger Karsdorp, Ben Burtenshaw, and Mike Kestemont. 2017. Synthetic literature: writing science fiction in a co-creative process. In *Proc. Computational Creativity in Natural Language Generation (CC-NLG)*. 29–37.
- [16] Chris Martens and Matthew A Hammer. 2017. Languages of play: towards semantic foundations for game interfaces. In *Proc. FDG*. 32–41.
- [17] James R Meehan. 1977. TALE-SPIN, an interactive program that writes stories. In *Proc. IJCAI*. 91–98.
- [18] Nikita Prokopov. 2014. DataScript. github.com/tonsky/datascript.
- [19] Aaron Reed. 2018. Archives of the Sky. archivesofthesky.textories.com.
- [20] Ben Robbins. 2011. Microscope: A Fractal Role-Playing Game of Epic Histories. lamemage.com/microscope.
- [21] Melissa Roemmele and Andrew S Gordon. 2015. Creative Help: a story writing assistant. In *Proc. ICIDS*. Springer, 81–92.
- [22] James Ryan. 2018. *Curating Simulated Storyworlds*. Ph.D. Dissertation. UC Santa Cruz.
- [23] James Owen Ryan, Michael Mateas, and Noah Wardrip-Fruin. 2015. Open design challenges for interactive emergent narrative. In *Proc. ICIDS*. Springer, 14–26.
- [24] James Owen Ryan, Adam Summerville, Michael Mateas, and Noah Wardrip-Fruin. 2015. Toward characters who observe, tell, misremember, and lie. In *Proc. AIIDE*.
- [25] Ben Samuel, Michael Mateas, and Noah Wardrip-Fruin. 2016. The design of Writing Buddy: a mixed-initiative approach towards computational story collaboration. In *Proc. ICIDS*. Springer, 388–396.
- [26] Ben Samuel, James Ryan, Adam J Summerville, Michael Mateas, and Noah Wardrip-Fruin. 2016. Bad News: an experiment in computationally assisted performance. In *Proc. ICIDS*. Springer, 108–120.
- [27] Robin Sloan. 2016. Writing with the machine. robinsloan.com/notes/writing-with-the-machine.
- [28] Ingibergur Sindri Stefnisson and David Thue. 2018. MimiSbrunnur: AI-assisted authoring for interactive storytelling. In *Proc. AIIDE*. 236–242.
- [29] Noah Wardrip-Fruin. 2009. *Expressive Processing: Digital Fictions, Computer Games, and Software Studies*. MIT Press.