

Why Are We Like This?: Exploring Writing Mechanics for an AI-Augmented Storytelling Game

Max Kreminski, Melanie Dickinson, Michael Mateas, Noah Wardrip-Fruin

University of California, Santa Cruz

{mkremins,mldickin,mmateas,nwardrip}@ucsc.edu

Abstract

Why Are We Like This? (*WAWLT*) is a playful, co-creative, AI-augmented, improvisational storytelling game in which one or more players explore and influence an ongoing simulation which they then glean for narrative material. It uses the recently developed simulation technology of *story sifting* (the recognition of microstories in a chronicle of simulation events), via the Felt library, to afford a new kind of playful, social, and creative writing experience. In this paper, we discuss our primary design goals: (1) using computation and interaction design to support casual player creativity, and (2) foregrounding character subjectivity as a driver for socially realistic interpersonal conflict. We further discuss how those design goals informed the system development. In particular, they led to the system features of subjective character reflection on past actions through character-centric sifting patterns, player-facing story sifting tools for querying storyworld state and history, and a set of *writing mechanics* to interface with the simulation and support playful creative writing. Examples of those writing mechanics include (1) explicit statement of system-understandable author goals, which are used to improve next action recommendations, and (2) free text editing of a malleable, textual transcript seeded by parameterized descriptions of player-selected simulation actions. We found in testing that, even in an incomplete state of development, and even among those who don't consider themselves fiction writers, *WAWLT* successfully supports player creativity. We also found that *WAWLT* affords particularly engaging play and a unique co-creative experience with two players, as opposed to just one.

Introduction

We present *Why Are We Like This?* (*WAWLT*), an AI-supported digital story construction game. In *WAWLT*, one or more players work to construct a story in a pastiche of the cozy mystery genre, 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.

One aim of *WAWLT* is to explore how computation can be used to provide support for new forms of playful creative writing practice, drawing particular inspiration from player storytelling practices that have emerged around simulation-driven games like *The Sims* and *Dwarf Fortress*. Eladhari has argued that player-created *retellings* of play experiences in interactive narrative games constitute narrative artifacts worthy of study in their own right, while Kreminski et al. have argued that simulation-driven games may be popular among players who write stories based on their play experiences because these games contain features that help players overcome

common barriers to creativity. With *WAWLT*, we aim to support similar forms of storytelling as a first-class activity, rather than an emergent “side effect” of a design not primarily intended to provide creativity support for player storytelling. The primary goal of a *WAWLT* play session is to produce a narrative retelling of the session’s events, with an underlying event structure produced through collaboration between the players and the simulation, and descriptive or elaborative prose recounting these events produced primarily by the players.

A second aim of *WAWLT* is to construct an AI-based play experience that centers a particular set of aesthetic goals involving the search for truth through narrative sensemaking. In pursuit of these goals, we draw some inspiration from features found in works of the *cozy mystery* genre, including sympathetic character motivations (even in the case of “bad guys”) and plots involving the disruption and eventual reconciliation of a tight-knit community. Unlike in some mystery stories that focus predominantly on the gradual discovery of the “hard facts” surrounding a crime, we sought to place the focus on character motivation. To achieve this goal, we set out to create a system that naturally led to the emergence of stories in which character subjectivity – the differences between different characters’ interpretations of and reactions to the same concrete situations and events – is a central thematic concern. During the course of play, both the players and the characters are attempting to fit narrative explanations or framings around the events that have transpired, and choosing actions to perform next based on a rapidly shifting narrative understanding of “the story so far”. Play therefore revolves centrally around the process of “trying on” different narrative framings, especially regarding character motivations, and developing the story based on these interpretations.

Together, these two design goals have led us to build our play experience around the recently introduced concept of *story sifting*, first proposed by Ryan et al. and further developed by Kreminski et al. Story sifting is the extraction of narratively potent sequences of events from a much larger chronicle of all the events that have taken place within a simulated storyworld. *WAWLT* applies story sifting technology in two key areas where it has not been applied before. First, sifting is used diegetically within the simulated storyworld to implement character subjectivity: characters run sifting patterns over the history of all events that have transpired, evaluate these events based on the stories they identify, and act based on these evaluations. Character conflict is therefore driven by divergent interpretations of the same concrete events. Second, sifting is also given over to the player(s), who can use the provided sifting tools to identify interesting directions in which to take their story.

Players in *WAWLT* spend much of their time investigating the history of the simulated world using story sifting tools. However, unlike in many mystery games, their primary aim in acquiring information is not to solve a particular puzzle, but to locate sites of narrative potential that they can then develop in their stories. The central pleasure of play is not that of discovery but that of co-authorship, especially in moments of unexpected convergence between plot threads established earlier in the play session. And even when players do find themselves sifting through the history in order to discover the answer to a particular question, these are rarely if ever questions about *what* concretely happened. Instead, the players’ deeper investigations

tend to focus on questions of *why* this particular character chose to commit this particular action at this point in time.

Related Work

The current iteration of *WAWLT* was conceived first and foremost as an analogue to tabletop story construction games, which attempt to provide scaffolding for a collaborative storytelling process between a group of human players. Robbins's *Microscope* in particular offers a wide variety of creativity support features, particularly the *palette* – a way for players to collectively negotiate what they do and do not want to see in the story – and a recursive story structure that enables players to “dive deeper into” any part of the story that they would like to further flesh out. Alder's *The Quiet Year* involves the collective production through play of a physical artifact, namely a map of the world that the players have created, which players can take with them after play as a reminder of the play session. And Reed's *Archives of the Sky* provides mechanisms for structuring character conflict around *values* held both by individual player characters and the larger society in which they exist. All of these features have directly inspired design elements in the current version of *WAWLT*.

WAWLT is a *mixed-initiative co-creative* system (Liapis et al.), and can be viewed as a *casual creator* (Compton and Mateas) for cozy mystery stories set in a particular context. Other casual creators for storytelling, such as Samuel et al.'s *Writing Buddy*, and other mixed-initiative co-creative systems intended to be used in a storytelling context, such as Stefnisson and Thue's *Mimisbrunnur*, have provided valuable design inspiration for *WAWLT*, but have not fully embraced the use of a fine-grained simulated storyworld in the way that we aim to here. The same is true of language-model-driven co-creative writing processes, such as the Botnik *Predictive Writer* app, Roemmele and Gordon's *Creative Help* system, and the writing practices described by Manjavacas et al. and Sloan.

WAWLT is built around *story sifting* in both its implementation and its design, making central use of the story sifting and simulation engine Felt. 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. Story sifting, originally known as “story recognition”, was first proposed as an open design challenge for interactive emergent narrative by Ryan et al., and further refined by Ryan in his dissertation, *Curating Simulated Storyworlds*.

Several existing play experiences make use of story sifting technology in some way, but none of them attempt to center story sifting as a player-facing game mechanic as we aspire to in *WAWLT*. Garbe's *Dwarf Grandpa* runs sifting retroactively on the history of a *Dwarf Fortress* world to pull out and highlight the stories of certain kinds of vampires. Samuel et al.'s *Bad News*,

an interactive theater experience with both human and computational components, involves a process of live sifting in which a human “wizard” (notionally one of the performers, rather than a member of the audience) interacts with a Python console to pull out interesting information about the simulated storyworld in which the story is set and feed this information in real time to the human actor portraying the simulated characters. Kreminski et al.’s *Cozy Mystery Construction Kit* also makes use of story sifting, albeit without centering it to the extent that we attempt to here.

Design

In *WAWLT*, one or more players work together to construct a story in a pastiche of the cozy mystery genre, supported by an AI system that provides players with inspiration and keeps the story moving forward even when the players are unsure what should happen next. *WAWLT* stories are set in a snowed-in research conference at a remote venue populated by a mix of researchers, permanent staff, and tourists. Characters and their initial relationships are procedurally generated at the start of a play session. Each character holds two randomly selected *values* from a pool of eight or nine, representing the things that this character views as particularly important, and these values – in conjunction with character personality traits – restrict the story sifting patterns that this character can use to make sense of the world.

Because different characters have access to different sifting patterns, they will end up telling themselves different stories about the events that have transpired in the world so far, even without any direct modeling of character knowledge phenomena – and, therefore, will end up acting in conflicting ways based on their conflicting evaluations of the same evidence. The closed environment of the conference venue acts as a pressure cooker, exacerbating initially minor tensions between characters until a variety of plausible motivations exist for characters to commit severe crimes.

At the start of each chapter, players select a set of initial author goals for what they would like to accomplish narratively in this chapter. The list of author goals is:

1. Involve *character* in plot
2. Cast suspicion on *character*
3. Dispel suspicion on *character*
4. Escalate tension between *character* and *character*
5. Defuse tension between *character* and *character*
6. Escalate tension between *value* and *value*
7. Defuse tension between *value* and *value*
8. Introduce false lead
9. Dismiss false lead

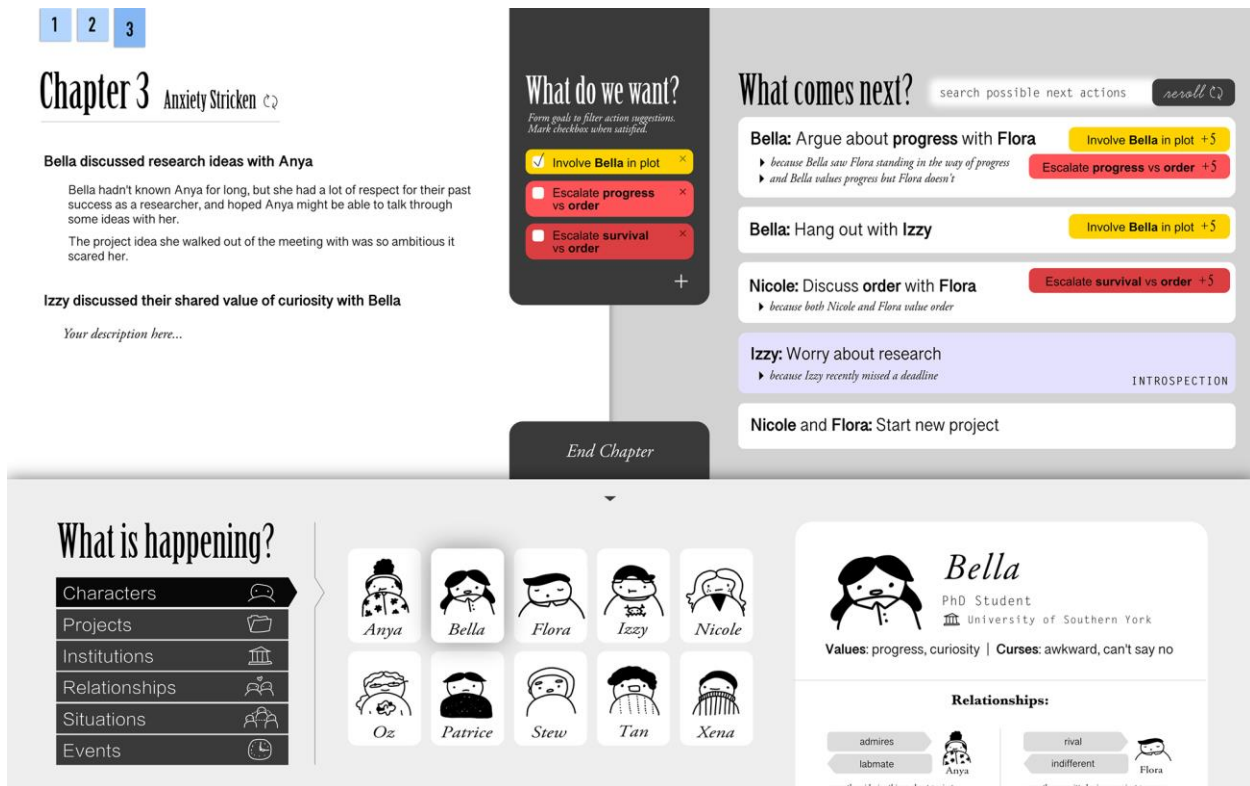
Italics in goal names indicates slots into which the players may drop any matching storyworld entity – currently, either a specific character or a value held by one or more characters within the storyworld. As play progresses, players periodically revise their author goals to reflect the things they would next like to accomplish within the story, allowing the system to provide them with better suggestions for the current context.



(Fig. 1: The author goal selection interface.)

Next, players select actions for the characters to perform, drawing primarily from a set of up to five suggested actions surfaced by the system, which is continually reassessing (based on a lightweight social simulation and some other simulation rules) what the characters might want to do next. Players can also choose to reroll the suggested actions to get a new set of suggestions. The simulation takes the players' author goals into account when making suggestions, attempting to surface actions that would help to advance these goals.

When the players select an action for a character to perform, its effects are realized in the simulation, and a short textual description of the action is generated from a parameterized template string attached to the action. This description is appended to the end of a textual transcript recounting the story so far. Players edit this textual transcript to further describe or elaborate on the events of this play session in prose form. By the conclusion of the play experience, through continual revision of the transcript, the players will have produced a full written story detailing the events of their play session. We see author goal selection; player selection of the next story action from an ongoing simulation, which then feeds back into influencing the simulation; and freeform prose editing within a more rigid plot-structural "skeleton" as examples of *writing mechanics*, which are uniquely afforded by *WAWLT*'s tight integration of underlying system and experience design.



(Fig. 2: The main WAWLT interface, with the running transcript of “the story so far” on the left and action suggestions on the right. Bold text in the transcript is system-generated, non-bold text is authored by the players.)

Emplotment

In *WAWLT*, by arming our simulated characters with story sifting, we aim to realize a computational model of Ricouerian *emplotment*: the process of narrative sensemaking by which agents narrativize the events of their own lives (Atkins). *WAWLT* characters “tell themselves stories about” the events that have transpired in the storyworld so far, and the stories they tell themselves influence the actions they then go on to take. For instance, one *WAWLT* character might consider multiple instances of another character being rude or short-tempered and conclude that this character is a bully, then resolve to treat them with less kindness afterwards. At the same time, another character might consider these same instances of rudeness alongside instances that suggest the rude character is currently under a lot of stress, then conclude that the rude character is going through a tough time and should be treated with *more* kindness rather than less. Players can then leverage these conflicting explanations of the same events as a potential driver of conflict between the characters in question.

After some experimentation, we found that this application of story sifting to implement character subjectivity naturally lends itself to the emergence of stories about misunderstandings and misreadings between characters. The sifting patterns that a character can use to make sense of the world are dependent on that character's personality. Because each character has access to different sifting patterns, characters often misinterpret the actions of other characters

in systematically biased ways. An action that one character might have intended as a gesture of comfort might be interpreted by its recipient as an expression of pity, and by an outside onlooker as an act of sarcastic mockery. Characters then take actions based on their biased interpretations of the world, sometimes leading to the emergence of conflict even between characters whose motivations might in truth be entirely compatible.

Thematically, this meshes well with our view of the mystery genre. *WAWLT* characters are searching for truth in a world where things are rarely as they first appear – not because of any explicit attempt to conceal the truth on the part of any particular character, but because every character has a biased and limited perspective on the world, and the question of “what really happened” is in many ways too subjective to receive a single wholly correct answer.

Moreover, because there are so many ways that any given action can be interpreted by other characters, the system can easily identify and surface a diversity of ways to dial the tension between two specific characters up and down at the player's will. In testing *WAWLT* with prospective users, we found that the system's ready provision of plausible ways to increase inter-character conflict can lead to a gleeful kind of *schadenfreude* in players, similar to the delight that some authors report that they experience when they find ways to “make their characters suffer.”

Implementation

WAWLT is implemented as a frontend-only web application, built using HTML, CSS, and JavaScript. It uses the Felt JavaScript library for simulation and story sifting. Internally, Felt uses the DataScript library – a JavaScript implementation of Datalog – to store and execute queries against the simulation state. For more details on the AI architecture of *WAWLT*, see (Kreminski et al. 2020).

All simulation actions are defined as Felt actions, which consist of several distinct parts:

- A Felt query over the current simulation state, defining the action's bindings (“roles”) and preconditions.
- A tagline string into which the values of bound variables can be substituted to produce a short human-readable description of this action, which serves as a prompt for the user to edit and expand.
- A function that takes a set of variable bindings and returns an event object, which is added to the simulation state database as a record of this action.
 - This event object contains a list of zero or more effects, which describe other ways in which the simulation state should be updated when this event is performed.
- An optional base weight corresponding to this action's narrative significance (or rule salience), such that – on an author-defined basis – some actions can be marked as

“higher priority” or “more strongly motivated” than others, and will be floated closer to the top of the action suggestions list as a result.

```
Felt.registerAction('worryAboutOthersProjectDrama', {
  tagline: '?n1: Notice that ?n2 is struggling with project "?projname"',
  baseWeight: 'low',
  where: [
    // there's an active project to which c2 is a contributor
    '?proj "state" "active"',
    '?proj "projectContributor" ?c2',
    // c1 is a character who likes c2
    '(likes ?c1 ?c2)',
    // three increaseProjectDrama events involving c2 and proj
    '?e1 "tag" "increaseProjectDrama"', '?e1 "actor" ?c2', '?e1 "project" ?proj',
    '?e2 "tag" "increaseProjectDrama"', '?e2 "actor" ?c2', '?e2 "project" ?proj',
    '?e3 "tag" "increaseProjectDrama"', '?e3 "actor" ?c2', '?e3 "project" ?proj',
    '< ?e1 ?e2 ?e3)',
    // extra info for display purposes
    '?proj "projectName" ?projname',
    '?c1 "name" ?n1',
    '?c2 "name" ?n2'
  ],
  event: (vars) => ({
    actor: vars.c1,
    target: vars.c2,
    project: vars.proj,
    effects: [
      {type: 'addImpression', source: vars.c1, target: vars.c2, value: 2}
    ],
    text: `${vars.n1} became concerned about ${vars.n2}'s struggles with project "${vars.projname}"`,
    tags: ['projects']
  })
});
```

(Fig. 3: An example action definition. In this introspection action, a character reflects on three past instances of a second character struggling with a particular project, and forms a weak positive impression of this character born of sympathy to their struggles.)

Author goals are used to evaluate and rank potential actions by their current relevance to the story conditions the players are trying to achieve, so that more relevant actions can be surfaced more prominently as suggestions. (In this context, a “potential action” is a data structure containing both a pointer to an action definition and a set of currently valid variable bindings for this action in the context of the current simulation state.) Each author goal’s heuristic function takes a possible action as an argument and returns a numerical score indicating this potential action’s suitability for this author goal.

Testing

We tested *WAWLT* with five different players: three as solo players, and two as a pair engaging with a single instance of *WAWLT* simultaneously on a single computer. Initially, each player was given a brief introduction to the project and the different parts of the user interface. Players were then instructed to think aloud during their interaction with the game for 5-15 minutes at the player's discretion.

Broadly speaking, we found that the current version of *WAWLT* – though incomplete – nevertheless already supports player creativity in some of the intended ways, and is capable of producing an enjoyable play experience. Players had little difficulty making use of the game's primary mechanics once they were introduced. All players, even those who initially struggled to make the pieces of their story fit together into a larger storyline, eventually found themselves excited or curious to discover what would happen next in the story. All players also expressed overall enjoyment of the play process. Four of five players (including both of the paired players and two of three solo players) reported some sense of ownership over the story they produced through play. Moreover, the paired players in particular expressed a great degree of enjoyment of the play experience; desire to continue working on the story (to such an extent that they were vocally disappointed that they could not continue at the conclusion of the play session); and feeling that what happened in the storyworld was somehow “real”.

Nevertheless, there were also some significant points of confusion among players. Four of five players expressed some desire to learn more about the backgrounds of each individual character when they first became relevant to the story, suggesting that the addition of character information cards should be a top priority for future design. Four players (including both of the paired players) reported a sense of directionlessness at least once during the play process, suggesting that the system's action recommendations were not always sufficient to provide players with a sense of narrative structure. In the paired-players group, both players initially assumed that author goals were intended primarily to be used by the system to filter and prioritize action suggestions, without realizing that they were also intended to be used as a way to encourage multiple simultaneous human players to negotiate intended story directions. Debriefing after the testing also indicated that three players (including both of the paired players) at some point forgot that the author goals existed, although the paired players “rediscovered” the author goals when a minor creative conflict briefly emerged between them.

With regard to the framing of the project, the success of the paired players in particular suggests to us that we should be further playing up and leaning into the multiplayer angle. Testing has clearly shown that the creativity support features we provide in *WAWLT* are, like their counterparts in tabletop story construction games, perhaps useful for individual players, but especially transformative when helping to scaffold and structure co-creativity in a multiplayer context, where negotiation between players regarding the content and direction of the story they intend to tell becomes a central part of play.

Conclusions

We present *Why Are We Like This?*, an AI-augmented playful co-creative storytelling game in which players work together with an AI system to write a story in a pastiche of the cozy mystery genre. In the process of creating *WAWLT*, we discovered several *writing mechanics* – including author goal selection, goal-dependent character action suggestions, and freeform prose editing within a more rigid plot-structural “skeleton” – that we use to provide support for player creativity. We have found through testing that the current design of *WAWLT*, though incomplete, effectively supports player creativity, especially by helping to structure and mediate a multiplayer storytelling process.

Works Cited

Alder, Avery. *The Quiet Year*. 2013, buriedwithoutceremony.com/the-quiet-year. Accessed 4 Jan. 2020.

Atkins, Kim. “Paul Ricoeur (1913-2005).” *The Internet Encyclopedia of Philosophy*. www.iep.utm.edu/ricoeur. Accessed 5 Jan. 2020.

Botnik. “Predictive Writer.” botnik.org/apps/writer. Accessed 5 Jan. 2020.

Compton, Kate and Michael Mateas. “Casual Creators.” *Proceedings of the Sixth International Conference on Computational Creativity*, 2015, pp.228-235.

Eladhari, Mirjam Palosaari. “Re-Tellings: The Fourth Layer of Narrative as an Instrument for Critique.” *International Conference on Interactive Digital Storytelling*. Springer, Cham, 2018.

Garbe, Jacob. “Simulation of History and Recursive Narrative Scaffolding.” 11 Feb. 2018, project.jacobgarbe.com/simulation-of-history-and-recursive-narrative-scaffolding. Accessed 4 Jan. 2020.

Kreminski, Max, et al. “Cozy Mystery Construction Kit: Prototyping Toward an AI-Assisted Collaborative Storytelling Mystery Game.” *Proceedings of the Fourteenth International Conference on the Foundations of Digital Games*, 2019.

Kreminski, Max, et al. “Evaluating AI-Based Games Through Retellings.” *Proceedings of the Fifteenth Conference on Artificial Intelligence and Interactive Digital Entertainment Conference*, 2019, pp.45-51.

Kreminski, Max, et al. “Felt: A Simple Story Sifter.” *International Conference on Interactive Digital Storytelling*, 2019, pp.267-281.

- Kreminski, Max, et al. "Generative Games as Storytelling Partners." *Proceedings of the Fourteenth International Conference on the Foundations of Digital Games*, 2019.
- Kreminski, Max, et al. "Why Are We Like This?: The AI Architecture of a Co-Creative Storytelling Game." *Proceedings of the Fifteenth International Conference on the Foundations of Digital Games*, 2020.
- Liapis, Antonios, et al. "Can Computers Foster Human Users' Creativity? Theory and Praxis of Mixed-Initiative Co-Creativity." *Digital Culture & Education* 8.2 (2016): 136-153.
- Manjavacas, Enrique, et al. "Synthetic Literature: Writing Science Fiction in a Co-Creative Process." *Proceedings of the Workshop on Computational Creativity in Natural Language Generation (CC-NLG 2017)*, 2017, pp.29-37.
- Reed, Aaron A. *Archives of the Sky*. 2018, archivesofthesky.textories.com. Accessed 4 Jan. 2020.
- Robbins, Ben. *Microscope: A Fractal Role-Playing Game of Epic Histories*. lamemage.com/microscope. Accessed 4 Jan. 2020.
- Roemmele, Melissa and Andrew S. Gordon. "Creative Help: A Story Writing Assistant." *International Conference on Interactive Digital Storytelling*, 2015, pp.81-92.
- Ryan, James. *Curating Simulated Storyworlds*. 2018. University of California, Santa Cruz, PhD dissertation.
- Ryan, James, et al. "Open Design Challenges for Interactive Emergent Narrative." *International Conference on Interactive Digital Storytelling*, 2015, pp.14-26.
- Samuel, Ben, et al. "Bad News: An Experiment in Computationally Assisted Performance." *International Conference on Interactive Digital Storytelling*, 2016, pp.108-120.
- Samuel, Ben, et al. "The Design of Writing Buddy: A Mixed-Initiative Approach Towards Computational Story Collaboration." *International Conference on Interactive Digital Storytelling*, 2016, pp.388-396.
- Sloan, Robin. "Writing with the Machine." May 2016, www.robinsloan.com/notes/writing-with-the-machine. Accessed 5 Jan. 2020.
- Stefnisson, Ingibergur Sindri and David Thue. "Mimisbrunnur: AI-Assisted Authoring for Interactive Storytelling." *Proceedings of the Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2018, pp.236-242.