

Towards a Critical Technical Practice of Narrative Intelligence

Chris Martens

martens@csc.ncsu.edu
Computer Science Department
North Carolina State University

Gillian Smith

gmsmith@wpi.edu
Interactive Media & Game Development
Worcester Polytechnic Institute

Abstract

The practice of building intelligent narrative technologies is implicitly grounded in a sociopolitical context whose values are reflected in the artifacts we build and our narration of their use. This position paper attempts to make these implicit values explicit by naming them and conceptualizing alternatives. We argue for adopting a *Critical Technical Practice* of narrative intelligence in which explicit critical reflection on the values and metaphors a system's assumptions embed is interwoven into the design process. We demonstrate a broad application of this practice by identifying four ideals underlying narrative intelligence systems that impose and reflect systems of power: the Hero's Journey, Amalgamated Human Authorship, Social Believability, and the Holodeck. We then identify values that are marginalized by these ideals, such as narrative pluralism, amplifying underrepresented voices, representation of non-normative identities and relationships, and opportunities for reflection and interpretation. Finally, we propose imagined alternatives to existing ideals that bring these marginalized values to the center.

Introduction

When building tools and technologies to support the human practice of storytelling, our training leads us to focus on technical innovation: what data structure will we use to hold knowledge about the world? Will that knowledge be hand-authored or extracted from a dataset? What algorithms for model construction, search, and selection will determine the components of generated stories?

However, maintaining this focus comes at the expense of asking other questions. For instance, what are stories *for*, and what is the *social* value of creating technology that can replicate or model the human processes around understanding and creating them? Both the technology and mass-media storytelling industries exhibit dramatic inequalities across gender, race, and other axes of privilege in terms of whose voices are centered and whose stories are told. Without taking seriously the questions of how the technology we create can shift or amplify existing power imbalances, the narrative intelligence community is in danger of perpetuating and

deepening these inequalities. We may operate under the belief that the systems we construct are sociopolitically neutral, but in doing so, we willfully ignore our own biases. As Elizebeth Sampat puts it in her book *Empathy Engines*, “Instead of striving for neutrality, the first step to truly affecting game design is to realize that nothing is neutral: what we see as a lack of bias is our own blindness to the circumstances we live within” (Sampat 2017).

The good news is: we don't have to stop writing code and building systems in order to confront these issues. In fact, the practices of technology-building and critical reflection can mutually support one another. This position paper argues for a *critical technical practice* of Narrative Intelligence in which technical innovation is tightly coupled with an examination of implicit values throughout the design process.

Critical Technical Practice (CTP) is a proposal by Phil Agre, originally in the context of Artificial Intelligence more broadly, to pair the creation of software artifacts with critical reflection on the underlying (often implicit) metaphors and values used by its creators (Agre 1997a). These metaphors, Agre argues, not only shape the design process and characteristics of the resulting artifact, but also inform *human narration* of the artifact's operation, including scholarly discourse about its significance, capabilities, and potential roles in society. While Agre is in favor of creating working systems to demonstrate alternative metaphors and values, he argues that CTP is a *bidirectional interaction* between the act of technical implementation and that of critical reflection. The code and the critique should inform one another. Agre laments the narrowness of an AI field “constrained by the ensemble of technical schemata that operated in [its] discourse and practice.” This limitation exists because, Agre claims, the field lacks “critical tools with which to defamiliarize these ideas—to see their contingency or imagine alternatives to them.”

Motivated by the continued relevance of Agre's 1997 observations, we present a position paper arguing that Narrative Intelligence as a field needs a Critical Technical Practice. We make an effort to “defamiliarize” familiar concepts in the field of Narrative Intelligence and imagine alternatives to them, listing a broad (but shallow) set of examples specific to this field.

According to Boehner et al. (2005), CTP “consists of the following moves:”

1. Identifying the core metaphors of the field;
2. Noticing what, when working within those metaphors, remains marginalized;
3. Inverting the dominant metaphors to bring the margin to the center;
4. Embodying the alternative as a new technology.

We demonstrate the first three steps of the above outlined process by taking a broad view of Narrative Intelligence as a field, considering problems as diverse as narrative cognition, interactive story systems, and expressive simulation. However, our aim is *not* to take the fourth step by presenting an embodiment of our imagined alternatives as a new technology. Instead, we outline several potential directions for narrative intelligence research to more closely examine marginalized values, aiming to create space for critical reflection and ideation without the pressure to develop validated solutions to each problem as it arises. This approach is consistent with Agre, who “believe[s] that building things is an important way of learning about the world,” but also points out that “even if the value of critical reflection is proven only in its contribution to improved technical systems, many valuable criticisms will go unpublished if all research papers are required to present new working systems as their final result” (Agre 1997b).

Related Work

Agre’s initial writing defining CTP was in the context of artificial intelligence, and he lamented being unable to get across to his colleagues the importance of examining metaphors and values (Agre 1997b). The Human-Computer Interaction (HCI) community, perhaps due to its history as a coalition of a wide variety of disciplines, including those widely considered “nontechnical,” has had more success in taking his points: in 2004, the ACM SIGCHI conference hosted a workshop on Reflection in Design (Dourish et al. 2004). The following year, Boehner et al. (2005) explicitly apply Agre’s framework to HCI, presenting case studies from Phoebe Sengers’ lab at Cornell.

We are hopeful that AI for entertainment’s similarly interdisciplinary makeup bodes well for CTP. There is certainly precedent for openly confronting methodological differences between the disciplines that comprise us, such as the dialogue between “Science Considered Harmful” (Horswill 2013) and “Science Considered Helpful” (Young 2018) in the context of building technologies for interactive narrative. Smith (2017) carries out a similar exploration of values and what they marginalize in procedural content generation, and Compton (2017) (the same year!) lays out pitfalls and opportunities when representing humans in procedural systems. Dickinson, Wardrip-Fruin, and Mateas (2017) dig deeper into the prospect of using social simulation for social justice, exploring what it would take to bring issues of marginalized identity and social activism to the front and center of an interactive narrative experience.

Finally, efforts to address implicit assumptions in AI as a field at large are resurfacing. In “Decolonial AI”, Mohamed, Png, and Isaac (2020) lay out a critique of an implicit colonizer mindset underlying much of AI research and practice,




Figure 1: The Hero’s Journey, as visually depicted on Wikipedia (fetched August 3, 2020).

and they describe framework for decolonizing it, which incorporates CTP as an essential step.

How Narrative Intelligence Systems Impose and Reflect Systems of Power

In this section, we identify several common metaphors and values prevalent in the past and present of Narrative Intelligence as a community and body of work. For each metaphor or value, we identify ideas and values that are marginalized as a result.

The Hero’s Journey and Other Story Patterns

Some of the earliest commonly-cited work in narrative intelligence is the Russian formalists’ attempts to identify common structural patterns in Russian folktales, such as Propp’s Morphology of the Folktale (Propp 2010). Computational work has found this approach useful, because story generation could be formalized as context-free grammar expansion or other means of “instantiating templates” (Thompson 2001; Bringsjord and Ferrucci 1999; Gervás 2013). This approach implies the definition of a story as literally *formulaic*, i.e. a grammar for Russian folktales reduces the essence of Russian folktale to a template in which interchangeable subjects and alternatives can be chosen to fill in the blanks.

For those of us with computer science training, the concept of interchangeable parts within a constant structural framework appeals naturally to our ideal of *abstraction*. By generalizing from examples, we can extrapolate *patterns* in the same sense as a sewing pattern, something that can be transmitted through language, then traced and replicated to materialize in different contexts, each instance still recognizable as having a common origin.

In addition to the general use of patterns appealing to particular computational metaphors, there is also much to be gleaned from the *choice* of patterns employed (and the source material that gives rise to them). For example, a common “plot template” is the *Hero’s Journey*, as popularized

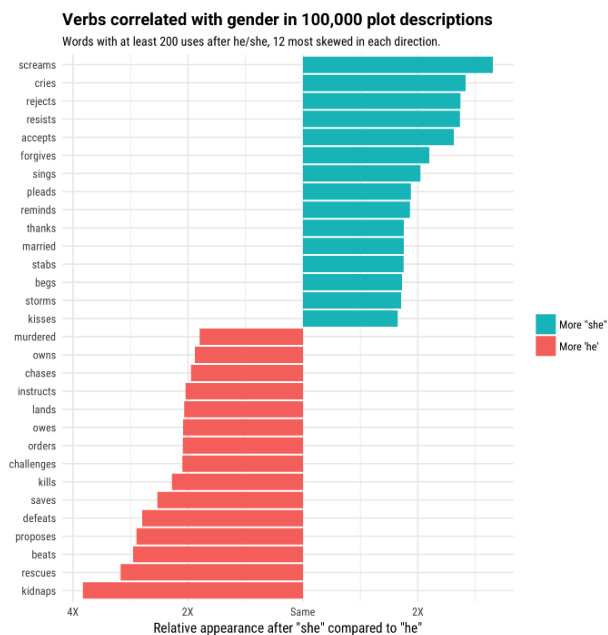


Figure 2: David Robinson’s analysis of gender pronoun-verb correlation in Wikipedia plot descriptions (Robinson 2017).

by Campbell (2008) (see Figure 1). The Hero’s Journey is a generalization of tales and lore in which a hero departs from the familiar on an adventure, faces crisis-inducing trials and tribulations, and concludes with a transformation and return to the familiar.

This pattern embeds a number of culturally relative assumptions, particularly in Campbell’s more detailed breakdown of each segment of the journey, including *gender as binary*, *social functionalism*, and *male as primary/default*. The eponymous Hero is always male (male as default), and women appear only in elements of his journey in specific roles, such as the Goddess who aids his quest, and the Temptress who tries to derail him. The central male character takes on all the active roles in the story, solving problems with decisiveness, aggression, and violence; women play much more passive roles, such as fleeting assistants or irrelevant distractions. This prescription of role to gender is a form of *social functionalism*, a philosophy that argues that the moral function of society relies on people following the roles that match their innate characteristics. This use of gender roles both embeds the cultural values of the authors who created these stories, and reinforces those values as cultural standard by embedding them in trans-generational myths.

The Hero’s Journey does not exist in a vacuum: it comprises the mythologies that precede it and informs the stories that follow it. The same cultural assumptions can be discovered in modern (literary and film) media, as revealed by David Robertson’s analysis of Mark Riedl’s Wikipedia Plot Summaries dataset (Robinson 2017): see Figure 2 to see which types of action are more commonly associated with the pronouns “he” and “she” in these stories.

A second set of values reflected in the Hero’s Journey and

its fellow narrative templates pertains to their embedding in Western culture, taking Western assumptions about narrative for granted as universal. For instance, the assumption that *all narratives must have conflict*, which drives, for instance, the CPOCL project (Ware and Young 2011), is not universal (Hunter 2016). For instance, the Japanese *Kishōtenketsu* story form consists of exposition, development, a “turn” or twist, and reconciliation, with no apparent need for conflict.

Along similar lines, Cherny (2018) writes about the story template system *Plotto* (Cook 2011). While investigating its potential for story generation, Cherny identified overt racism and misogyny, such as the event template “Female protagonist, of an inferior race, rescues male protagonist, of a superior race.” Cherny identifies three top-level categories for Plotto event nodes—enterprise, love and courtship, and married life—visualizes them by their connections as successors or predecessors to their nodes, and counts the in-/out-degree for each node. The node of highest degree is a “Love and Courtship” template: “B-0, rescued from an accident by A-0, whom she does not [sic] know, falls in love with him.” In other words, the “gratitude romance” trope, which undermines expectations of autonomy and consent in real human relationships. Although Cherny herself uses the “datedness” of Plotto as a reason not to publish the dataset, other projects in our community have made use of its structure, facing a need to reconcile the problematic content after the fact (Eger et al. 2015; Eger and Mathewson 2018).

In summary, narrative templates, especially those based on culturally dominant pretexts, marginalize the following values:

- Narrative Pluralism: the co-existence of differing human cultures and the non-universality of any given culture’s approach to storytelling
- Social Role Autonomy: the idea that one’s function within a story, as in life, is not pre-determined by attributes like gender and race that one is born with.

Amalgamating Human Authorship

A number of computational narrative researchers explicitly cite Roger Schank’s concept of *social scripts* (Schank and Abelson 1975) as inspiration for their concept of narrative intelligence. Script theory posits that human behavior is governed by the specific situations they find themselves in, such as being in a restaurant, and that these situations each have associated mental “scripts” for appropriate action within them, which people learn through socialization. For instance, when sitting with a friend in a restaurant, we understand such actions as ordering food, eating, chatting with our friend, and paying the bill, to be in-scope (in a particular order), and we would not, without special provocation, go and take food from someone else’s table to satisfy our hunger (as might otherwise be expected in purely goal-driven models of intelligent behavior).

Once the challenges of extensive domain authoring, maintenance, and debugging became apparent for creating narrative templates as above, some people got the idea to try to assemble this kind of “commonsense knowledge” about

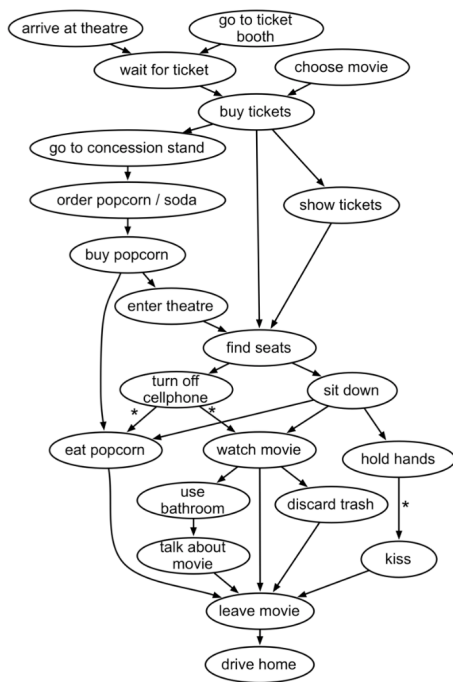


Figure 3: A learned plot graph for the “movie date” scenario from Li et al. (2012a)’s narrative crowdsourcing project.

human behavior through crowdsourcing and statistical data analysis. For instance, Chambers and Jurafsky (2008) induce script-like schemata called *narrative event chains* from raw newswire text, following the grammatical role (subject or object) of a chosen “protagonist” and the actions (identified through verbs) that involve them. Orkin and Roy (2009) describe their project, *The Restaurant Game*, which collects data from players in a restaurant scenario who are given a fixed set of actions to apply, and extracts a Markov chain model. Li et al. (2012b) use Amazon’s “Mechanical Turk” system to collect natural language stories from users to satisfy a scenario prompt, such as going on a date to a movie, written in a particular format that makes it easy to parse into events (subject, verb, object) and learn script-like “plot graphs” from statistically aggregated responses, which they go on to use for story generation (Li et al. 2013).

This last example, the Scheherazade project, is especially interesting in the central metaphor the authors use in the title of the first paper: “Learning to Tell Stories about Social Situations” Li et al. (2012a). They state (emphasis added): “By leveraging the crowd and its collective understanding of social constructs, we can learn a potentially unlimited range of scripts regarding **how humans generally believe real-world situations unfold**. We seek to apply this script-like knowledge to the generation of believable stories that involve common social situations or the direct engagement of virtual characters in social behaviors.” The form of social intelligence they aim to replicate computationally is exactly that found in the statistical average of “the crowd,” along



Figure 4: A screenshot of Rimworld’s interface depicting characters changing their opinions of one another.

with whatever ingrained biases and cultural assumptions the humans making up this crowd may carry—which is evident in the plot graph that emerges for a prompt in which “John and Sally” go on a date (see Figure 3). Although the authors suggest that one could use such a model to subvert common expectations, generating a more narratively-interesting form of surprise, in practice, their uses of the crowd-learned model for story generation and interaction straightforwardly replicate the model’s statistical distribution (Li et al. 2013; Li and Riedl 2015). These projects lead us to the conclusion that even when cultural biases are not hard-coded into a template like the Hero’s Journey, human exposure to those templates, both for fictional and real-world situations, biases them towards these predictable patterns. Furthermore, by collecting data from anonymous human participants to generate and evaluate story segments, these systems obscure the human role in authorship and the human creative labor that underlies these models of story.

In summary, the amalgamation of human authorship into statistical models marginalizes the following values:

- Underrepresented Storytellers: amplifying the voices and creative decisions of people whose stories are seldom told.
- Authorial Voice: the ability to attribute meaning to the identity and labor of human story authors.

Social Believability

One common approach to interactive storytelling is *social simulation*: the use of AI to simulate story characters that embody a systemic model of social interaction. Examples include commercial games such as The Sims, Rimworld, and Animal Crossing, in addition to academic work such as the *Comme il Faut* engine (McCoy et al. 2010) and corresponding game *Prom Week* (McCoy et al. 2011), Ryan and Mateas (2019)’s *Talk of the Town* town population simulator and corresponding mixed-reality game *Bad News* (Samuel et al. 2016), and Azad and Martens (2019)’s Lyra system for simulating character opinions. Social simulation systems like these typically represent characters in a story as instances of a common data structure, along with consistent and systemic rules for interaction. Under such rule systems, the possibility space for social interaction depends on the state held in the data structures for the interacting characters, and each potential interaction can update that state.

The concept of *simulated societies* necessarily embeds sociological theories chosen by the programmer. When these theories are based on intuition rather than research, we might call them *folk sociological theories*: for example, Rimworld (Sylvester and Ludeon Studios 2016) simulates romantic attraction with highly specialized code based on character gender (Lo 2016) (see Figure 4). Even when following generally accepted sociological theories of affinity, however, we see that systems in which gender and relationships are explicitly represented tend to use those variables to govern social interaction in culturally-relative, *ad hoc* ways. Whether to permit same-gender or polyamorous relationships requires an *explicit choice* on behalf of the designer, and the decision is realized as part of the inherent logic of the world—e.g. whether the player is given the option to “romance” a character or not—rather than a contingent boundary made of laws and social customs, the navigation of which is often central for queer and polyamorous people in real life. In her AIIDE 2015 keynote, Mitu Khandaker (Khandaker 2015) discusses her own struggle with simulating romance in her game *Redshirt*, concluding that simulating human behavior in this vein is “always going to be a bit reductive.” Perhaps, she suggests, the goal of *social believability* is inherently flawed.

In summary, the ideal of social believability marginalizes the following values:

- Non-Normative Representation: representing non-socially-normative people and their experiences with chafing against norms (e.g. gender/sexual minorities, people of color, people with disabilities)
- Procedural Inquiry: purposefully limiting social simulations to legible systems of expression as a means of interrogating what you can and cannot say within that system.

The Holodeck

While not unique to interactive narrative, the obsession with “immersion” that sees virtual reality as its logical conclusion pervades discourse in our community as well. In particular, the Star Trek fiction of the *Holodeck* has long been a metaphor for the goal of intelligent interactive narrative experiences (Murray 2017), in which ideals like *you feel like you’re really there* and *you have total control over the story* are extolled.

Critique of this ideal is not new or unique to our field, either. Salen, Tekinbaş, and Zimmerman (2004) define the **immersive fallacy** in digital games as “the idea that the pleasure of a media experience lies in its ability to sensually transport the participant into an illusory, simulated reality.” The implication is that game designers often take this idea for granted without considering alternatives. In interactive narrative, even when we aren’t considering immersion in the form of sensory simulation such as VR, we are often assuming a related ideal in terms of suspension of disbelief when it comes to the concept of *social believability* in our story characters, or more general *narrative believability*—the idea that a generated story is most successful when its reader or participant loses track of their grounding in reality and is fully convinced of the veracity of the fictional world.

In other words, the immersive fallacy is the endless pursuit of narrowing the “aesthetic distance” between the player and simulation.

The immersive fallacy marginalizes the following three values, the first two of which come from Mitu Khandaker’s AIIDE 2015 keynote (Khandaker 2015):

- Double Consciousness: a media effect in which subjects hold both the simulated reality and their bodily reality simultaneously in mind, maintaining aesthetic distance.
- Reflection: through aesthetic distance, we create opportunities for the player to infer, interpret, and otherwise actively engage in meaning-making with an interactive experience, rather than passively absorbing it.
- Resistance to Player Whim: asking the player to learn the language and rules of a system that works differently from the reality they are familiar with, rather than attempting to shape the world according to player expectations.

Bringing Marginalized Values to the Center

What are stories for, and why are they so important to being human? They help us imagine alternatives to the way things were in the past or might be in the future; they help us make sense of our own lives and difficult situations we have faced; they help us see alternative versions of ourselves, and by placing ourselves in the shoes of a story character, we learn new ways of emotionally processing and problem solving.

Intelligent narrative technologies have the potential and opportunity to do all of these things in new and powerful ways: authoring tools that democratize technology, enabling more authors to tell stories about the world through their unique perspectives; supporting the human imagination through procedural generation of unexpected configurations; narrative generation systems that ask players to make choices and reflect on their consequences. To achieve this potential, we posit that we need to reexamine what human forms of *narrative intelligence* we really want to emulate with our technology. For example, a college literature professor might assert that *narrative intelligence* means the ability to interpret ambiguous stories, to situate the meaning of the story with respect to concepts of authorship (who wrote the story?), representation (who is shown in the story doing what kinds of things?), and historical context (what was going on in the world when and where the author wrote this story?). We might value a reader’s ability to draw parallels with present-day events, with other narrative works, and with their own personal experiences.

In this section, we return to the lists of marginalized values at the end of each subsection above and identify alternative ideals and research goals that would center them.

Expanding the Canon

In support of the values undermined by using common story patterns as a basis for narrative generation, we propose to redefine what we consider the “canon” of narrative reference work, and to rework our definitions of narrative on the basis of these findings. We need to surface the implicit canon (e.g. the Hero’s Journey, Plotto, and Aesop’s Fables are just

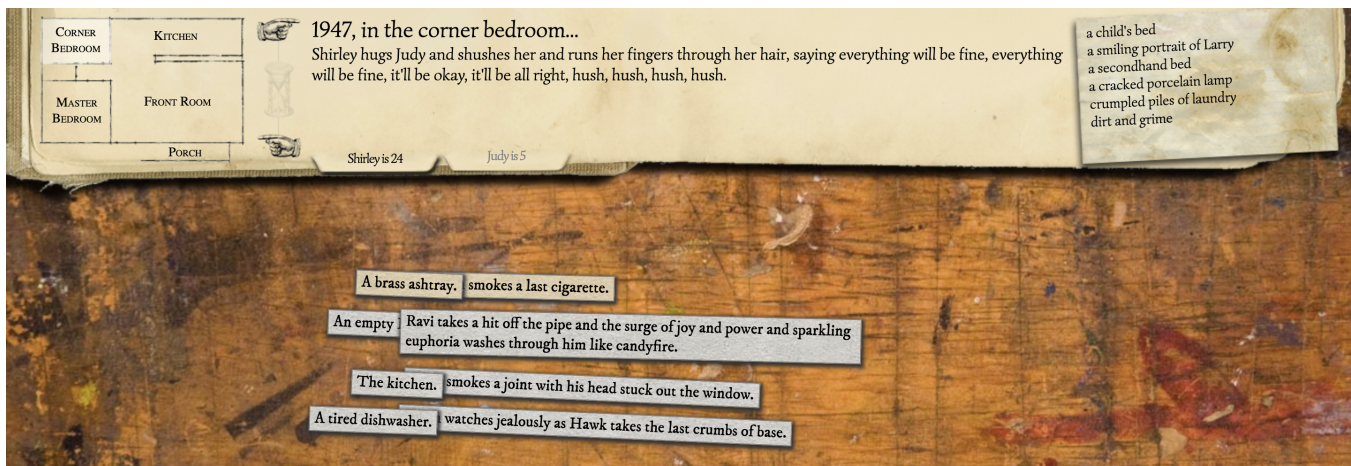


Figure 5: A screenshot of *18 Cadence* (Reed 2013). Source: <http://18cadence.textories.com/index.html?id=200>

a few), and we need to recognize their contingency by seeking counterexamples to their structural assumptions. What would a practice of narrative intelligence that centers the value of *narrative pluralism* look like?

One potential approach is to work actively to make narrative intelligence a more international field, forming more collaborations with researchers from Africa, Asia, South America, and members of indigenous nations across the world. Another is to encourage inquiry about our own tastes and influences as narrative researchers, consider the cultural origins of the stories we love, and surface the contingency of the patterns we take for granted. A third strategy is to change our approach to pedagogy when we are teaching students from narrative cultures that differ from ours: rather than stating the Western canon as a given that they must learn, we must give students the pedagogical “permission” to invent new typologies and definitions based on the stories they grew up with.

Computational Media Literacy

What would it look like to take a crowdsourced dataset, or a model like GPT-3 that is compiled from a huge number of different English text resources, and model their tropes and biases *for the sake of* analyzing and subverting them? Would it actually work to take Scheherazade’s plot graph model and simply invert the optimization function, choosing the *least* expected event over the most expected, and argue that the result inverts common social expectations, or would we need a more systematized model for tropes and expectation? Supposing we could identify bias in datasets or text input, could we build a mixed-initiative system with an interface similar to Tom Forth’s “gender bias calculator” for recommendation letters (Forth 2018), that provides feedback to authors on their use of narrative tropes that reinforce harmful real-world stereotypes?

More generally, we propose a model of narrative intelligence that takes into account the fact that our understanding of any given story depends on the thousands of stories we have encountered before it. Through our lives, we develop

not just linguistic literacy, but also *media literacy*, expectations about the conventions and tropes present in a given genre or artistic movement. Authors depend on this literacy in order to subvert it: if they know their audience expects one thing, they can strategically pull the rug out from under them to say something new about that trope. What would it look like to design a computational narrative generation system that could reason about and implement similar strategies?

Enabling Reflective Co-Creation

In interactive narrative experiences, neither the player nor the game designer has sole creative ownership over the story. However, most systems attempt to lean strongly one way or the other, either asking the player to (merely) bear witness to a well-defined story constructed by the designer, or providing a fully sandbox-mode simulation in which the player is free to tell any story they can with the simulation mechanics.

We envision a model for granting a player a sense of autonomy without turning the dial all the way to “omnipotence.” This kind of experience requires a push-and-pull between player and system in terms of asserting authorial voice; the story emerges from a negotiation over their contested overlap. Models for this kind of collaborative negotiation can be found in tabletop storygames, particularly those without a central story-manager or game-master, such as *Fiasco* (Morningstar 2009) and *Microscope* (Robins 2011). Examples in digital games exist as well, including much of Aaron Reed’s work, such as the game *18 Cadence* (Reed 2013), which offers a “drafting” interface for assembling a story from fragments taken across space and time in the history of a residence at the titular address (see Figure 5). This kind of interactive storytelling requires re-imagining the mode of interaction, replacing the assumption that the player will perform the role of a character within the story with a request of them to play the role of co-author.

Even within the more conventional player-as-story-character mode of interaction, however, we have opportunities to enable more active participation in shaping the story. One potentially fruitful avenue is to prompt the player to-

wards autobiography, as recently explored for the purpose of mental health recovery for post-traumatic stress disorder (Kerr, Deane, and Crowe 2019). This paper explores how the *life story model of identity and intentional change theory*, two concepts in personal psychology related to how people conceptualize themselves and the challenges they encounter, can be incorporated into a play experience that asks them to project their personal struggles onto fictional characters.

Resisting Legibility in Social Simulation

The concept of *legibility*, as defined by Scott (1998), refers to the power in being able to reduce complex phenomena to a discrete set of interpretable symbols. A desire for legibility is often what drives reductionist choices in social simulation: the assignment of binary (or otherwise finitely-enumerated) gender to characters (so that we can write code that branches based on it), the labeling of sexual orientation, the definition of personality as a vector of 5 real numbers. Incidentally, as Scott (1998) argues, it is also how political states assert control over human beings through surveillance.

Mitu Khandaker describes wrestling with how to “represent love” in the game *Redshirt*, and unlike the author of *Rimworld*, opts to make some of its workings based on a hidden variable with no connection to the explicit character traits (Khandaker 2015). This choice suggests an alternative in *resisting legibility*, or mindfully choosing when *not* to simulate a complex phenomenon in detail and instead treat it as an uninterpretable black box. Khandaker argues for acknowledging that “modeling social dynamics of humans is obviously always going to be a bit reductive” and aiming instead for “conscientious reductionism,” “aspiring towards **procedural believability** rather than **social believability**” (Khandaker 2015).

“Simulation can’t simulate everything, of course,” she argues. “The things we leave out when we’re designing a game are often more interesting and more expressive than the things that we put in. Play operates on that gap. It entices us to fill it in. It’s the same principle [of cartoon abstraction as identified by] Scott McCloud in *Understanding Comics* (McCloud 1993)—the less fidelity we put into something, the more room there is for us to fill the gap with ourselves.” In this spirit, we argue for conscientious consideration of when *not* to represent something as complex as gender and sexual orientation in social simulations. This doesn’t mean that these simulations can’t be interpreted as representing those concepts—quite the opposite. For instance, if we don’t explicitly assign gender to our characters, but we do assign them physical features (e.g. body shape, facial features, hair style, clothing, makeup), we can learn something interesting about how players make gender-related inferences about these characters.

Conclusion

In this position paper, we have argued for narrative intelligence as a field to engage in Critical Technical Practice, and against the requirement of proof in the form of a working system as a process for resisting the “fallacy of alternatives.” We outlined the steps for doing so and demonstrated the first

three. We took a shallow but broad pass at identifying dominant metaphors and values in the current-day technical practice of narrative intelligence and discussed several alternative values and metaphors, which we believe have the power to expand our imagination of what this field can achieve.

While this paper attempts to overview the breadth of ways in which narrative intelligence can benefit from CTP, to really do justice to the critical thought required, an entire paper could be written about each identified value (and corresponding alternatives) described above. Likewise, any given technical project could also go through a much more in-depth application of the steps of Critical Technical Practice. We therefore conclude by asking researchers in the INT community to learn the vocabulary of critique, build relationships with their humanist colleagues, and ultimately to engage differently with their technical work, creating a culture in which considering the values reflected by each technical project is not an afterthought, but a built-in and continual part of the design process.

Acknowledgements

The authors gratefully acknowledge the reviewers for engaging thoughtfully with this work and helping us identify interesting lines of related work.

References

- Agre, P. 1997a. *Computation and human experience*. Cambridge University Press.
- Agre, P. 1997b. Lessons learned in trying to reform AI. *Social science, technical systems, and cooperative work: Beyond the Great Divide (1997)* 131.
- Azad, S., and Martens, C. 2019. Lyra: Simulating believable opinionated virtual characters. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 15, 108–115.
- Boehner, K.; David, S.; Kaye, J.; and Sengers, P. 2005. Critical technical practice as a methodology for values in design. In *CHI 2005 Workshop on quality, values, and choices*, 2–7.
- Bringsjord, S., and Ferrucci, D. 1999. *Artificial intelligence and literary creativity: Inside the mind of BRUTUS, a storytelling machine*. Psychology Press.
- Campbell, J. 2008. *The hero with a thousand faces*, volume 17. New World Library.
- Chambers, N., and Jurafsky, D. 2008. Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, 789–797.
- Cherny, L. 2018. “PLOTTO”: Generating truly offensive stories since 1928. Medium.com. Fetched July 10, 2020.
- Compton, K. 2017. Little procedural people: playing politics with generators. In *Proceedings of the 12th International Conference on the Foundations of Digital Games*, 1–2.
- Cook, W. 2011. *PLOTTO: the master book of all plots*. Tin House Books.
- Dickinson, M. L.; Wardrip-Fruin, N.; and Mateas, M. 2017. Social simulation for social justice. In *Thirteenth Artificial*

Intelligence and Interactive Digital Entertainment Conference.

Dourish, P.; Finlay, J.; Sengers, P.; and Wright, P. 2004. Reflective HCI: Towards a critical technical practice. In *CHI'04 extended abstracts on Human factors in computing systems*, 1727–1728.

Eger, M., and Mathewson, K. W. 2018. dairector: Automatic story beat generation through knowledge synthesis. *arXiv preprint arXiv:1811.03423*.

Eger, M.; Potts, C. M.; Barot, C.; and Young, R. M. 2015. Plotter: operationalizing the master book of all plots. In *Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference*.

Forth, T. 2018. Gender bias calculator. Personal website. fetched July 10, 2020.

Gervás, P. 2013. Propp's morphology of the folk tale as a grammar for generation. In *2013 Workshop on Computational Models of Narrative*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Horswill, I. D. 2013. Science considered harmful. In *Ninth Artificial Intelligence and Interactive Digital Entertainment Conference*.

Hunter, M. 2016. From conflict to concord: Lessons from the mouse. *Etudes*.

Kerr, D. J. R.; Deane, F. P.; and Crowe, T. P. 2019. Narrative identity reconstruction as adaptive growth during mental health recovery: A narrative coaching boardgame approach. *Frontiers in Psychology* 10:994.

Khandaker, M. 2015. The Eliza affect. Keynote talk, AIIDE 2015. Recorded and uploaded to YouTube (https://www.youtube.com/watch?v=_CNEzBf7gKM). Fetched August 3, 2020.

Li, B., and Riedl, M. 2015. Scheherazade: Crowd-powered interactive narrative generation. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.

Li, B.; Lee-Urban, S.; Appling, D. S.; and Riedl, M. O. 2012a. Automatically learning to tell stories about social situations from the crowd. In *the LREC 2012 Workshop on Computational Models of Narrative*.

Li, B.; Lee-Urban, S.; Appling, D. S.; and Riedl, M. O. 2012b. Crowdsourcing narrative intelligence. *Advances in Cognitive Systems* 2(1).

Li, B.; Lee-Urban, S.; Johnston, G.; and Riedl, M. 2013. Story generation with crowdsourced plot graphs. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*.

Lo, C. 2016. How rimworld's code defines strict gender roles. Rock Paper Shotgun.

McCloud, S. 1993. Understanding comics: The invisible art. *Northampton, Mass.*

McCoy, J.; Treanor, M.; Samuel, B.; Tarse, B.; Mateas, M.; and Wardrip-Fruin, N. 2010. Authoring game-based interactive narrative using social games and *comme il faut*. In *Proceedings of the 4th International Conference & Festival of the Electronic Literature Organization: Archive & Innovate*, volume 50. Citeseer.

McCoy, J.; Treanor, M.; Samuel, B.; Mateas, M.; and Wardrip-Fruin, N. 2011. Prom week: social physics as gameplay. In *Proceedings of the 6th International Conference on Foundations of Digital Games*, 319–321.

Mohamed, S.; Png, M.-T.; and Isaac, W. 2020. Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology* 1–26.

Morningstar, J. 2009. Fiasco. Bully Pulpit Games.

Murray, J. H. 2017. *Hamlet on the holodeck: The future of narrative in cyberspace*. MIT press.

Orkin, J., and Roy, D. 2009. Automatic learning and generation of social behavior from collective human gameplay. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 385–392.

Propp, V. 2010. *Morphology of the Folktale*, volume 9. University of Texas Press.

Reed, A. 2013. 18 Cadence: a story kit. Self-published.

Robins, B. 2011. Microscope. Lame Mage Productions.

Robinson, D. 2017. Gender and verbs across 100,000 stories: a tidy analysis. Personal Blog: <http://varianceexplained.org/r/tidytext-gender-plots/>. Fetched August 3, 2020.

Ryan, J., and Mateas, M. 2019. Simulating character knowledge phenomena in talk of the town. *Game AI Pro* 3:433–448.

Salen, K.; Tekinbaş, K. S.; and Zimmerman, E. 2004. *Rules of play: Game design fundamentals*. MIT press.

Sampat, E. 2017. *Empathy engines: Design games that are personal, political, and profound*. Elizabeth Sampat.

Samuel, B.; Ryan, J.; Summerville, A. J.; Mateas, M.; and Wardrip-Fruin, N. 2016. Bad news: An experiment in computationally assisted performance. In *International Conference on Interactive Digital Storytelling*, 108–120. Springer.

Schank, R. C., and Abelson, R. P. 1975. Scripts, plans, and knowledge. In *IJCAI*, volume 75, 151–157.

Scott, J. C. 1998. *Seeing like a state: How certain schemes to improve the human condition have failed*. Yale University Press.

Smith, G. 2017. What do we value in procedural content generation? In *Proceedings of the 12th International Conference on the Foundations of Digital Games*, 1–2.

Sylvester, T., and Ludeon Studios. 2016. RimWorld. Digital download on Steam.

Thompson, S. 2001. *Motif-Index of Folk-Literature: A Classification of Narrative Elements in Folk Tales, Ballads, Myths, Fables, Mediaeval Romances, Exempla, Fabliaux, Jest-Books, and Local Legends*, volume 1. Indiana University Press.

Ware, S. G., and Young, R. M. 2011. Cpocl: A narrative planner supporting conflict. In *Seventh artificial intelligence and interactive digital entertainment conference*.

Young, R. M. 2018. Science considered helpful. In *International Conference on Interactive Digital Storytelling*, 21–35. Springer.