

Learning to Listen: Critically Considering the Role of AI in Human Storytelling and Character Creation

Anna Kasunic

Carnegie Mellon University
Human-Computer Interaction Institute
Pittsburgh, PA USA
akasunic@andrew.cmu.edu

Geoff Kaufman

Carnegie Mellon University
Human-Computer Interaction Institute
Pittsburgh, PA USA
gfk@cs.cmu.edu

Abstract

In this opinion piece, we argue that there is a need for alternative design directions to complement existing AI efforts in narrative and character generation and algorithm development. To make our argument, we a) outline the predominant roles and goals of AI research in storytelling; b) present existing discourse on the benefits and harms of narratives; and c) highlight the pain points in character creation revealed by semi-structured interviews we conducted with 14 individuals deeply involved in some form of character creation. We conclude by proffering several specific design avenues that we believe can seed fruitful research collaborations. In our vision, AI collaborates with humans during creative processes and narrative generation, helps amplify voices and perspectives that are currently marginalized or misrepresented, and engenders experiences of narrative that support spectatorship and listening roles.

1 Introduction

Somebody gets into trouble. Gets out of it again. People *love* that story!
(Vonnegut, 1970)

Once upon a time, people decided it was not enough to share stories; as humans are as much a curious animal as we are a storytelling one, we sought to study them. Theories of narratology, or the study of narratives, defines and breaks down narratives according to their distinct states of action, events, and elements (Prince, 1974; Bal, 2009) and (for the most part) narratologists agree that to constitute a narrative, a text must tell a story, exist in a world, be situated in time, include intelligent agents, and have some form of causal chain of events. Narratives also usually seeks to convey something meaningful to an audience (Ryan et al., 2007). (Note: In this paper,

we are somewhat relaxed with the terms “story” and “narrative,” and will often use the two interchangeably). With the advancement of artificial intelligence (AI), research related to narratives has taken on a whole new shape and meaning; not only are there new narrative forms ripe for the studying, including more examples of stories that are interactive and branching (Ryan and Rebreyend, 2013), but there is also a whole field of research devoted to improving AI storytelling capabilities.

Although research that teaches AI to tell better stories is challenging and intriguing on many levels, research and discourse in domains like psychology, literary fiction, and even social media suggests that a) humans benefit from telling stories, so AI might serve us better if it nurtures our storytelling predilections rather than act solely as storyteller, and that b) after all these years, we as humans still have a lot to learn and improve when it comes to telling narratives. Our stories brim with issues of representation, bias, and authenticity that can dilute or negate the beneficial powers of consuming and interacting with narratives, and our social structures promote certain stories while silencing others. If we’re not careful, AI storytellers will only inherit and exacerbate these problematic patterns. In what follows, we present an opinion piece that urges researchers at the intersections of AI, natural language processing (NLP), human-computer interaction (HCI), and storytelling to envision different futures for research related to AI and storytelling. In this paper, to be clear, we do not strive to present precise quantitative or qualitative conclusions or recommendations, nor are we exhaustive in our presentation of extant storytelling research, discourse, and innovations. Rather, we seek to spark dialogues, questions, and cross-discipline exchanges around the role of AI in storytelling. To this end, we discuss existing conversations and

research from human-computer interaction, AI, and other domains concerned with storytelling; we provide anecdotal highlights from interviews we conducted with character creators; and we illuminate alternative, promising directions for AI storytelling research. We begin by outlining some of the predominant roles that AI research in storytelling has held in order to provide context for our discussion. As a framing note, readers should be aware that we will use the term AI to encompass different forms of computational approaches that may not fall within everyone’s defined scope of AI. Whatever their specific disciplines or perspectives, we humbly request that our readers relax their definitions of AI for the duration of this paper to include computational approaches, more broadly.

2 Predominant Roles of AI in Narratives

We can categorize AI work in storytelling in three ways: 1) teaching AI to generate and understand stories; 2) helping human storytellers as a co-creator; and 3) modeling story elements. Some of the earliest work in storytelling and AI focused on improving AI’s understanding of stories using scripts, or “boring little stories” in the words of the authors (Schank et al., 1975; Schank and Abelson, 1975). As technology advanced, more research attention shifted to using AI to generate stories. For example, the hypertext fiction model emerged, in which links to branching narratives allowed for branching stories and some level of direct agency over the story (Bolter and Joyce, 1987). Researchers have continued making advances in story generation, developing AI story, world, and character generators that are planning- or event-sequence-based (Fairclough and Cunningham, 2004; Lebowitz, 1987; Porteous and Cavazza, 2009; Riedl and Young, 2010; Young et al., 2004; Barber and Kudenko, 2007; Min et al., 2008). Generators may also be character-centric, in which players’ interactions with intelligent agents move the stories forward (Magerko, 2006; Swartjes and Theune, 2006; Cavazza et al., 2002). This work most commonly aims to create entertainment value, with an emphasis on learning and catering to player preferences (e.g., (Thue et al., 2007)) and is often framed in terms of applications to interactive narratives, even if the potential applications of the work could extend to narratives, in general.

However, the goals of AI research in narratives are not exclusively focused on AI story generation; some work also strives to teach machines how to be more “human” through stories (Huang et al., 2016; Riedl and Harrison, 2016). Although much less common, some prior work has also positioned AI as co-creator, encouraging and guiding humans in creating their own stories (Ryokai and Cassell, 1999; Bers and Cassell, 1998; Van Broeckhoven et al., 2015). We have also seen work in the interactive drama space in which human-allowed actions are more open-ended, positioning humans to act more like actors on a stage than characters that must conform to a limited story world (Mateas and Stern, 2003). In addition, some research in modeling stories has focused not on story generation, but on understanding (and subsequently improving) human experiences of narratives. For example, work in identifying “emotional arcs” looks at mood shifts and audience engagement in experiencing narratives (Chu and Roy, 2017; Reagan, 2017), and other work has endeavored to identify turning points in stories (Ouyang and McKeown, 2015), model the shapes of stories (Mani, 2012; Elson, 2012), and understand the relationships of characters to narrative arcs (Bamman et al., 2014b,a). These works have identified what people enjoy about stories, have learned from our existing stories, and subsequently can augment story generation efforts.

In short, for the (many) story-lovers among us, it’s an exciting time to be working in AI, NLP and HCI. However, as we will discuss in the next section, there can be potential dangers in placing AI in a storytelling role, in aligning AI storytelling research too closely with interactive narratives, and in catering too much to the preferences of the humans engaging in stories and games. We now turn to a discussion of both the benefits and potential harms of storytelling as they relate to extant work on AI in narratives.

3 The Pleasures and Perils of Narratives

3.1 The Complex Pleasures of Narrative

Scholars in philosophy, psychology, anthropology, and related disciplines have characterized storytelling as fundamental to how we as humans grow, learn, develop, and process and experience the world, (Jung, 1964; Dautenhahn and Nehaniv, 1998; Sutton-Smith, 1986, 2012; Paley, 2009; Cooper, 1993). As such, our desire to en-

gage in stories and storytelling is not a “frivolous impulse, but a fundamental adaptive response” (Rose, 2012). As an experimental study from the 1940s has shown, we go so far as to ascribe narrative to situations where no narrative form exists (Heider and Simmel, 1944). Although AI research in emotional arcs recommends “happy” endings and seeks to maximize positive moods (Chu and Roy, 2017; Reagan, 2017), other research suggests that the relationship between story enjoyment and emotion is more complex. Research on benign masochism tells us that the human brain can derive pleasure from negative reactions and feelings such as sadness, fear, and disgust (Rozin et al., 2013). For example, sad films can be highly enjoyable, especially for certain groups such as female and younger viewers (Oliver, 1993, 2003; Mares et al., 2008). Similarly, research has found that mixed emotional experiences, such as experiencing both happiness and sadness rather than just one or the other, can be beneficial to one’s physical health (Hershfield et al., 2013). Thus, AI work that focuses on maximizing human enjoyment may overemphasize “sunny” experiences of narratives, and by focusing on pleasure rather than growth, may favor stories with narratives that fail to challenge and aid in human development.

Entertainment through narratives can be a valuable end goal in itself, but it can also have other, attendant advantages. According to transportation theory, we can achieve immersion in narrative worlds through identification with characters and perceptions of plausibility or the “suspension of disbelief,” in which we view narrative worlds and character actions as authentic (Green et al., 2004, 2003; Tesser et al., 2005). Not only does this transportation lead to enjoyment, but it can also enable perspective taking and belief change, (Kaufman and Libby, 2012; Berns et al., 2013). It can positively transform us, e.g. leading us to personality growth and maturation (Djikic et al., 2009b), with potentially higher transformative effects on attitudes for those who are resistant to change, or have diminished emotionality (Dal Cin et al., 2004; Djikic et al., 2009a). Reading fiction has been shown to improve the ability to attribute mental states to oneself and others (known as Theory of Mind), an important cognitive foundation for complex social relationships (Kidd and Castano, 2013), and reading narratives can lead us to feel psychologically connected to groups of characters, increasing feel-

ings of belongingness and subsequently leading to greater feelings of satisfaction and more positive mood (Gabriel and Young, 2011).

3.2 On “Listening” Versus “Agentic” Narrative Forms

The jury is still out, however, on which forms of media provoke higher levels of transportation, transformation, and enjoyment. Here, we find it useful to separate “listening” forms of narrative from “agentic” forms of narrative, and we will use these terms throughout the remainder of the paper. We characterize “listening” narratives as positioning the consumer of the narrative in a more passive role, listening or watching the story rather than making direct decisions that define or shape the characters, the plot, or the narrative world. “Listening” narratives would include traditional, typically non-branching narratives such as films, novels, and short stories. We define “agentic” narratives as stories in which the narrative consumer has some level of direct agency over characters, plot, or other story elements. In our definition, “agentic” narratives are akin to interactive narratives, and this aligns with other definitions of interactive narratives. For example, agency in the context of interactive narratives can be said to occur when the world “responds expressively and coherently to our engagement with it” (Murray, 2004). Accordingly, we prefer to stress the idea of agency over interaction because we do not think that “listening” narratives preclude interaction. Indeed, other researchers and designers ostensibly share our view of agency in narratives as nebulously interactive. For example, Persuasive Games has produced experiences that question the notion of how much traditional game definitions of agency truly allows players to interact with and engage in a narrative (Bogost, 2005, 2006), and other researchers have chosen to use the term “agency play” rather than agency to suggest that interactive narratives require more expansive notions of interactivity (Harrell and Zhu, 2009). As we consider the future of narrative expression and consumption, we can consider the possibility of listening narratives that allow narrative consumers to interact without directly making decisions about narrative or character arcs and shapes.

Indeed, studies have found that traditional “listening” forms such as books and movies may be equally or even more engaging and transforma-

tional than “agentic” narratives (Oh et al., 2014; Green et al., 2008; Jenkins, 2014). However, other studies have found that enjoyment and identification can be higher in agentic narratives (Hefner et al., 2007; Elson et al., 2014; Hand and Varan, 2008). It may also be that certain types of interaction may have varying costs and benefits. For example, a study allowing for character customization actually decreased narrative engagement and enjoyment, showing that the qualities of a narrative, not character agency, might be more important (Green et al., 2004).

We make no attempt to argue for “listening” over “agentic” forms of engagement with narratives, but we do believe that listening forms warrant more attention. For example, advancements in AI mechanisms for branching narratives needn’t apply to exclusively agentic forms of narrative; we can also think in terms of new listening experiences that are simultaneously branching and non-agentic (listening). A large body of work on branching narratives could be translated into new forms of listening rather than agentic digital media; for example, researchers have shown how their work in story generation techniques like plot graphs can be applied to branching (not necessarily agentic) narratives (Li et al., 2013; Guzdial et al., 2015). However, to our knowledge, there does not exist a wide range of practical applications of advancements in story generation and branching narrative work to new, listening forms of narrative.

Some analogies may be useful here. In describing his affinity for both traditional (listening) and interactive (agentic) narratives, storyteller and SIMS creator Will Wright likens traditional narratives to a roller coaster, and games and interactive narratives to a dirt bike (Rose, 2012). Neither experience is necessarily “better” than the other; whether the driver or the passenger, each has their unique benefits and affordances. Moreover, just as the invention of the roller coaster enabled new forms of “riding”, listening forms of narrative needn’t necessarily be “traditional.” Just as humans enjoy, learn, and develop through social interaction, we also have much to gain by spectatorship such as the popular pastime of “people-watching.” Just as there is value in active thinking, there is also value in meditating (watching our thoughts pass by without directly engaging). We argue that by focusing so much scholarly attention

on agentic forms of narrative, we may be missing out on ways to use technology to engender new ways of listening. Technological advancements in branching narratives, for example, could be realized by means of listening narratives rather than agentic narratives; we can consider the potential benefits of narratives that allow a multitude of experiences and paths, without conceding choice or agency to the listener/spectator.

Lastly, although existing efforts in AI narrative generation would suggest that humans benefit primarily from the experience of receiving narratives, the *telling* of stories can be highly beneficial. They can help us develop resilience, (East et al., 2010), provide therapeutic benefits (Block and Leseho, 2005; Carlick and Biley, 2004; Chelf et al., 2000; Pennebaker, 1997), and activate imaginative processes that are key to human growth and development (Harris, 2000). Thus, as with the narrowness of focus on interactive (agentic) rather than listening forms of narrative, prioritizing AI’s role as storyteller misses valuable opportunities for empowering humans as storytellers.

3.3 The Potential Harms of Existing Narratives

In addition, as we consider AI-enhanced storytelling experiences, we need to be mindful that our starting points—the story frames that we use to train AI— may unintentionally marginalize, misrepresent, and altogether exclude many groups. Just as recent work in natural language processing has criticized and sought ways to rectify the amplification of negative biases (e.g. gender biases) in NLP techniques (Dwork et al., 2012; Zhao et al., 2017; Bolukbasi et al., 2016; Voigt et al., 2017), AI storytelling has the potential to amplify and exacerbate issues of bias and diversity, which in turn excludes certain individuals from experiencing the potential benefits of story engagement. For example, AI story generators that learn from existing narrative corpora may learn that straight white male characters are best suited to be protagonists or figures of power, and that genderqueer and people of color should occupy sidekick roles. These stereotypes may persist regardless of the identification of the human author; for example, a study of online fan-fiction found that gendered stereotypes were highly common, and perpetuated by both male and female-identifying authors (Fast et al., 2016). Thus, machine learning models may

also learn from patterns of speech and role characterizations that stereotype certain groups, decreasing character authenticity. This is especially worrisome because research demonstrates that narrative persuasion is less effective if people cannot identify with the characters (So and Nabi, 2013; Ritterfeld and Jin, 2006; Slater et al., 2006; Gillig and Murphy, 2016).

Many popular films fail the classic Bechdel test, which simply specifies that the movie must 1) feature at least two women, 2) that these women must talk to each other, and 3) that their conversation must concern something other than a man (Selisker, 2015), and fare even worse on new NLP techniques that assess power differentials between men and women in movies (Sap et al., 2017). Issues of representation in film highlighted by the 2014 and 2015 Oscars, in which all awardees were white, sparked a social media firestorm under the hashtag #OscarsSoWhite (Syed, 2016; Borum Chattoo, 2018), and brought to further light a history of under-representation in film, with only 6.4% of all awardees since 1929 (1,688) being non-white (Berman, 2016). Minority groups are often under-, mis-, or negatively represented in film and other forms of narrative (Okoye, 2016; Smith, 2009; Hooks et al., 2006). In writing communities, gendered violence under the dominance of a “straight male cisgender patriarchy” and exclusion of black and brown writers from major literary publications has spawned a wave of debate and protest about exclusion, marginalization, and the silencing of voices (Tsay et al., 2015; Groom, 2015). Such issues suggest that AI may be more useful to us as an aid that can help identify biases and stereotypes, and amplify muffled voices, rather than a generator that replicates and extends our existing, problematic narratives.

In the next section, we cull anecdotal excerpts from a series of interviews we conducted with individuals deeply involved with some form of character creation to reveal existing pain points in human’s attempts to avoid and address issues of biasing, stereotyping, under-representation, and misrepresentation. Our interviewees’ discussions suggest concrete, specific ways in which AI can aid humans in improving some of the more problematic elements of our existing storytelling efforts.

4 An Exploration of Challenges in Character Creation

The one thing about being a dude and writing from a female perspective is that the baseline is, you suck.

As author Junot Díaz’s quote above (Rosenberg, 2012) points out, creating characters can be an intractable challenge. Unless we are in the rare case of writing a story that is only about the self, with no secondary characters, creating the “Other”— a character who is different from oneself along one or multiple dimensions (Shaw and Ward, 2005)— is inevitable. As humans, we define ourselves along multiple axes of identity, including gender, sexuality, race/ethnicity, class, nationality, health, disability, education, and passions/interests, to name a few; some axes may be especially salient for some individuals, and inconsequential for others. We posit that anxiety and uncertainty about how to authentically and sensitively create characters who are Other can hinder both the creative process and the narrative experience, as inauthentic characters can also impede character identification and subsequent narrative transportation and enjoyment. Understanding the ways in which human character creators approach and grapple with creating characters that are Other can provide insights into where the most crucial needs lie, and how we might design AI systems to assist with rather than model human stories. To this end, we conducted interviews to explore and better understand the space of character creation and its attendant pain points. Below, we present anecdotal highlights of interviews as they pertain to insights into needs for assistive AI; a full presentation of our methodology and qualitative analysis processes can be found elsewhere (more information available upon request).

We conducted qualitative, semi-structured interviews with 14 individuals with deep involvement in character creation, including novelists, short story writers, poets, journalists, television and game writers, actors, directors, and role-playing gamers, game masters, and designers (including both tabletop and live-action role-playing games). These character creators, recruited with the help of professors in relevant departments at our local university, ranged in age from 19 to 62 (average of 45), and held education levels from “some college” to PhD. All spoke English as a primary language, and primarily were born and

raised in the U.S. Eight identified as male, five as female, and one as non-binary; 11/14 identified as white, one as black, one as Native American, and one as Asian. For several of our participants, aspects of the narrative creation process constituted their full-time occupation or activity—e.g. writer, videographer, professor of drama or literature, and game designer—whereas others pursued narrative and/or character creation as a passion, hobby or pastime while also holding another occupation, such as secretary, civil servant, human resources coordinator, or student. With IRB approval, we audio-recorded the interviews (each lasting roughly 40-70 minutes), transcribed, and qualitatively analyzed using open-coding techniques to identify patterns across our participants.

We asked our participants to describe their processes of creating or embodying their characters, what informs the development of their characters, on what axes they identify with or diverge from their characters, and what conflicts or hesitations they have in creating or embodying certain kinds of characters. These responses validated existing research on the benefits of storytelling as a source of joy and growth. As one participant put it, “you get shaped by these stories that touch you, and by the sources that touch you. And I think you develop greater empathy. I think you become a better human being though that” (p2). The results of our interviews offer key insights about how AI systems can support humans in character creation and storytelling efforts, which we can organize into three themes: 1) the distinct ways in which different participants struggled and dealt with inauthenticity concerns; 2) the pain points participants discussed about giving voice to characters from under-represented groups; and 3) the impact of collaboration on character creation.

4.1 Concerns about Inauthentic Characters

Our participants generally fell into one of two camps when it came to relationships between their characters and their self-identity. Either a) they specifically chose characters they viewed as similar to themselves, operating under the adage, “write what you know” or b) they grounded themselves in notions of universality, seeking to find elements of themselves in characters that were seemingly highly divergent from themselves. Yet both groups expressed feelings of discomfort with

and anxiety about representing different viewpoints, suggesting opportunities for AI to assist with authentic character representation.

Participants that fell into group A explained that certain character decisions were outside their comfort zones, e.g. role-playing a character of an opposite gender (p9, p12). Another participant stated, “I’m very careful about running mental illnesses that I don’t and have never have. Because I’m sensitive enough to portrayals of my own, that I kind of don’t want to screw that up.” (p13). Other participants echoed this sentiment of certain stories or portrayals being “not my story to tell” (p4).

Participants in group B felt it was very important to include diverse characters in their stories and games in order to be more inclusive, but many of these participants expressed concern that despite their best efforts, they might be misrepresenting characters that identified in ways differently from themselves. They wanted to remain inside the lines of what they felt, as one participant worded it, “cultural appreciation” rather than “cultural appropriation” (p9). One role-player was initially hesitant to move outside her own identity. She had slowly branched into different ethnicities, genders, and sexualities, but had lingering apprehension regarding whether her portrayals were ethical and authentic, saying “Hopefully I’m not being horribly insulting to anyone of that ethnicity or sexuality while playing them. I hope I’m not. I think I’m not. I think I’m doing it relatively sensitively (p11). Thus, AI could be helpful in flagging characters that might be cause for concern by perpetuating certain stereotypes or offenses.

4.2 Issues of Representation

We have given examples in this paper of narrative exclusion along lines of gender and race/ethnicity; concerns related to these topic arose often in our interviews, and suggests opportunities for AI to offer additional support. For example, one of our participants, a drama director, said he made a purposeful decision to cast racially diverse actors in his plays (p3).

Yet even among those for whom improving the representation of certain under-represented groups is a priority, there can be conflict over *how* we should represent such groups. For example, one participant (p1) spoke of the controversies in TV writing communities around what it means to

write authentic characters of color. He spoke of a panel he participated in about TV representations of people of color in which many of the panelists were sharply divided on the questions: Is it okay for a character to be universal in identity, such that someone of a different race could conceivably play that character? Or ought characters be steeped in the specifics of their social identities and contexts? Participants also brought up issues of exclusion along axes that are often overlooked. For example, one participant discussed issues with neurotypical privilege, explaining that collaborative storytelling games are often exclusionary because they require players to pool from a relatively common pool of narrative tropes, meanings and interpretations that are not easy accessible to those who are neurodivergent (p8). These concerns about and disagreements on how to authentically and sensitively represent different groups indicate that AI that could serve to assist humans in grappling with and reflecting on these issues.

4.3 Impact of Collaboration on Character Creation

Where writers of novels or short stories may be more likely to develop narratives relatively autonomously and in isolation, other media lend themselves to highly collaborative environments, such as role-playing games (where the game master and the role-playing actors interact to shape the narrative), writing for the stage or screen (where writers interface with actors that embody their characters), and video game writing (where it is common for large teams to collaborate). Participants in narrative media with more collaborative development processes spoke enthusiastically of how actors and other characters had reshaped their understandings of their own characters. For example, a playwright discussed how interactions with actors often reshaped not just a character, but a whole play (p4). Similarly, a screenwriter-participant (p7) discussed completely revising a major scene after an actress revealed she couldn't "in her wildest dreams" imagine taking the action assigned to her. The interviewee stated that he often gains invaluable insights from actors, and explained that the relationship between actors and their characters are symbiotic; if an actor can't feel they can be true to a character, then everything will fall apart. Role players and game masters spoke of how in-

teractions with other characters shaped their understandings of their own characters, and affected the decisions they made in the game. Where writers in less collaborative contexts do research and seek guidance from those they feel may have more expertise or insight, in these more collaborative contexts, characters are not just created; they are constantly negotiated and re-negotiated. Through these processes of negotiation, our participants explained that they felt their characters took on more authentic, lifelike forms. However, not all narrative media are structured to automatically support such forms of collaboration and feedback, and not all narrators and character creators have access to social circles that can enable such collaboration. Intelligent agents that can play similar roles to human collaborators (e.g., other role players and actors) could provide critical, transformative feedback to creators and narrativists that work in relative isolation.

In sum, our interviews indicate that AI could be helpful as a storytelling assistant or co-creator by offering practical assistance (e.g. flagging misrepresentations), providing support for reflection on representation, and by taking on character embodiment roles that are usually assumed by humans in collaborative creation contexts.

5 Discussion

As we move forward towards new forms of media, narrative, and interaction, we urge scholars to take a step back and question the *whys* of AI in storytelling. Based on existing research and current trends in AI and other domains invested in narratives, and informed by the qualitative interviews our team conducted with character creators, we recommend a reorientation towards how we conceptualize the role of AI in storytelling. Yes, we can keep moving towards a future in which AI becomes more and more adept at human forms of storytelling. But is that the preferred future? As we consider the joys and benefits humans experience by engaging in storytelling, the shortcomings of our current narrative forms and processes that can exclude groups and dilute the transformative power of narratives, and both the struggles and affordances of different processes of character creation, we see several potential branches that future AI narrative systems can grow. Here, we return to the idea of "listening": we envision AI that better listens to human storytellers and assists us as co-

creators, and AI-assisted narrative forms that enable “listening” rather than “agentic” engagement. We give examples of specific starting ideas we have for 1) designing AI to support human storytellers, and 2) investigating “listening” rather than “agentic” forms of narrative that we hope will inspire the growth of new branches of AI narrative research. We acknowledge that the current state of computational powers renders some of our suggestions only feasible through at least partial Wizard of Oz approaches; we consider these ideas as starting points to guide future research and scientific advancements. Although we think current paths of AI research merit continued work and investigation, we believe that these alternate paths of inquiry and design are at least equally promising and important.

5.1 AI and Crowd-Powered Feedback Mechanisms for Human Storytellers

As we saw from our literature review and our interviews, humans *enjoy* storytelling; they grow, learn, and heal from it. However, as we saw in our interviews, creating authentic characters can be a challenging and emotionally fraught task, and as we saw from discussions of (mis)representation and stereotyping in film and literature, the stories that humans currently create are not the ideal models for AI to emulate. Thus, instead of expending all our effort on teaching AI to tell stories, we can divert some of our energies to using AI to help humans tell the stories they may struggle to tell. AI has already been used to model emotional arcs in narrative, and to identify bias in a number of domains. AI could be leveraged to better identify potential problems that could dilute authenticity and stymie narrative transportation, such as exclusion and stereotyping. There are a number of approaches that already use crowd-powered “mini-corpora” to teach AI how to generate narrative, (Guzdial et al., 2015; Li et al., 2012, 2013, 2014; Purdy and Riedl, 2016) but this existing work does not seek to improve experiences or engage crowdworkers in meaningful forms of storytelling. Taking cues from work in improving crowd workers experiences by inducing curiosity (Law et al., 2016), we might further consider crowd-powered feedback mechanisms could allow both story creators and crowd workers to engage in and benefit from stories in different ways.

For example, taking a character-centric ap-

proach, AI systems could prove useful in identifying when characters begin to fall into traps of stereotypes or implausibility. We could consider training models using a combination of existing corpora and crowdsourcing; as a secondary benefit, we could design crowdsourcing studies such that they engage crowd workers in meaningful storytelling that contributes to larger, concretized goals so that crowd workers are also benefiting by consuming and ultimately contributing to revisions of narratives. Studies could invite participation from those who most identify with marginalized and underrepresented populations, and are thereby able to speak to concepts of authenticity around specific identities (including axes of identity that our interview participants highlighted, such as mental illness and neurodivergence). We could then apply these models to new narratives and use them to generate feedback for narrative creators (e.g., flagging certain depictions that are deemed to be inauthentic or insensitive); AI systems provide prompts and exercises to encourage reflective or creative practices to psychologically and creatively grapple with these challenges.

Considering the dynamics of collaborative character creation processes in domains like role playing games and acting, there could be opportunities to re-purpose advances we’ve made in intelligent agents. During the narrative and character creation processes, creators could engage with intelligent agents that take on certain roles in the story, and give feedback on elements that feel inauthentic or incohesive. For creators that work in relative isolation, AI could simulate the more collaborative creative atmospheres native to role-playing and acting environments, in which characters can quite literally talk back and generate thoughts of their own, thereby re-shifting and re-shaping aspects of the characters and of narratives as a whole. Under such a scenario, humans would still be the primary storytellers, just as playwrights are still the people writing the script even if they decide to make revisions and edits based on feedback they receive from actors. Although AI cannot yet simulate human intelligence to the degree that such ideas would require, we can think of ways we could use crowdsourcing and/or partial Wizard of Oz approaches to achieve similar ends and provide guidance for future goals of AI.

5.2 Innovating and Exploring “Listening” Narrative Forms

Technological progress has spawned innovation in the field of interactive narratives and narrative games, particularly in the realm of video games. However, we argue that the scholarly energy around interactive narratives might be occluding potential for technological innovation in “listening” forms of narrative in which consumers are watchers or spectators rather than active agents in the narrative. As discussed previously, research has shown that both interactive and “traditional” narrative media have positive impacts, and under certain circumstances, “traditional” narratives may be even more effective for producing particular outcomes, such as narrative transportation. However, there is room to explore what “non-traditional” listening narratives could look like and produce.

As a starting point to this path of inquiry, we could leverage existing NLP research in style transfer, which uses neural networks to learn stylistic elements of a corpus, and apply the style schema onto new texts (Kabbara and Cheung, 2016; Shen et al., 2017; Carlson et al., 2017; Fu et al., 2017; Han et al., 2017). Narratives that seek to persuade, shift opinions, or otherwise transform readers are not always successful. For example a study exposing youth to stories of LGBTQ people found that where LGBTQ youth felt more hopeful, hetero and cis youth felt more negative attitudes after the narrative exposure (Gillig and Murphy, 2016). Here, we could begin to think about how we could transform stories that could better achieve their narrative end goals (e.g., changing attitudes) in ways that better speak to different groups of readers. Automated style transfer while maintaining diegetic plausibility and coherency could be one way to achieve this, and is worth further exploring. Again, we acknowledge that given the current state of computational sophistication, such an idea would require at least partial “Wizard of Oz” approaches.

We could also think of how “listening” to AI-powered on-demand storytelling could soothe, heal, and transport in times when we cannot access human-generated stories, or in time-sensitive situations when the narrative specifications we are seeking are not readily met by existing, available stories. In the midst of a bad break-up, an episode of depression, an anxiety attack, a death of a loved

one, a school or work-related failure, or any number of upsetting or traumatic experiences (potentially including the stresses of authentically representing “Other” characters in narratives), it may be difficult to reach out to others, and mustering the energy to actively engage in an agentic narrative might be too daunting. Instead, a listening form of narrative could be more helpful. An AI-powered listening narrative system could encourage certain emotional, psychological, or behavioral responses, such as allowing individuals to shift to more realistic and optimistic perspectives, or motivating individuals to reach out to friends, family or health support staff. It could be tailored to the specific situation or instance in which extra support is needed, could learn from one’s engagement with other narratives to cater to personal preferences of narrative style, content, and characters, and could take various media forms, such as text, audio (including more musical or sound-oriented narratives), video, augmented reality, or virtual reality. For example, multi-modal sensing could allow for branching even in the absence of explicit listener choice, such as using the listener’s nonverbal or physiological responses to make decisions or to alter the course or trajectory of the story. Given the current limitations of AI, early iterations could sample from existing corpora that have been studied to produce specific psychological or behavioral reactions.

We see these research and design suggestions as mere starting points to inspire more interesting ideas and conversations. We look forward to further discussing the opinions and ideas we’ve laid out in this paper, and collaborating with others who share our passions for narrative and exploring the limits and potentials of technology. In other words:

To be continued. . .

6 Acknowledgements

A special thanks goes to the National Science Foundation’s Graduate Research Fellowship Program for their support of this work, and to Diyi Yang, Mark Riedl, Anjalie Field, Judeth Oden Choi, and all our reviewers, all of whom provided feedback, ideas, and references to literature that we earnestly but imperfectly tried to incorporate into this final version.

References

- Mieke Bal. 2009. *Narratology: Introduction to the theory of narrative*. University of Toronto Press.
- David Bamman, Brendan O'Connor, and Noah A Smith. 2014a. Learning latent personas of film characters. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, page 352.
- David Bamman, Ted Underwood, and Noah A Smith. 2014b. A bayesian mixed effects model of literary character. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 370–379.
- Heather Barber and Daniel Kudenko. 2007. A user model for the generation of dilemma-based interactive narratives. In *Workshop on Optimizing Player Satisfaction at AIIDE*, volume 7.
- Eliza Berman. 2016. [See the entire history of the oscars diversity problem in one chart](#).
- Gregory S Berns, Kristina Blaine, Michael J Prietula, and Brandon E Pye. 2013. Short-and long-term effects of a novel on connectivity in the brain. *Brain connectivity*, 3(6):590–600.
- Marina Umaschi Bers and Justine Cassell. 1998. Interactive storytelling systems for children: using technology to explore language and identity. *Journal of Interactive Learning Research*, 9(2):183.
- Laurie Block and Johanna Leseho. 2005. listen and i tell you something: Storytelling and social action in the healing of the oppressed. *British Journal of Guidance & Counselling*, 33(2):175–184.
- Ian Bogost. 2005. [Airport insecurity](#).
- Ian Bogost. 2006. [Disaffected!](#)
- Jay David Bolter and Michael Joyce. 1987. [Hypertext and creative writing](#). In *Proceedings of the ACM Conference on Hypertext, HYPERTEXT '87*, pages 41–50, New York, NY, USA. ACM.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in Neural Information Processing Systems*, pages 4349–4357.
- Caty Borum Chattoo. 2018. Oscars so white: Gender, racial, and ethnic diversity and social issues in us documentary films (2008–2017). *Mass Communication and Society*, pages 1–27.
- A Carlick and Francis C Biley. 2004. Thoughts on the therapeutic use of narrative in the promotion of coping in cancer care. *European Journal of Cancer Care*, 13(4):308–317.
- Keith Carlson, Allen Riddell, and Daniel N. Rockmore. 2017. [Zero-shot style transfer in text using recurrent neural networks](#). *CoRR*, abs/1711.04731.
- Marc Cavazza, Fred Charles, and Steven J Mead. 2002. Character-based interactive storytelling. *IEEE Intelligent systems*, 17(4):17–24.
- Jane Harper Chelf, Amy MB Deshler, Shauna Hillman, and Ramon Durazo-Arvizu. 2000. Storytelling: A strategy for living and coping with cancer. *Cancer Nursing*, 23(1):1–5.
- Eric Chu and Deb Roy. 2017. Audio-visual sentiment analysis for learning emotional arcs in movies. *arXiv preprint arXiv:1712.02896*.
- Patsy Cooper. 1993. *When Stories Come to School: Telling, Writing, and Performing Stories in the Early Childhood Classroom*. ERIC.
- Sonya Dal Cin, Mark P Zanna, and Geoffrey T Fong. 2004. Narrative persuasion and overcoming resistance. *Resistance and persuasion*, pages 175–191.
- Kerstin Dautenhahn and Chrystopher Nehaniv. 1998. Artificial life and natural stories. In *International Symposium on Artificial Life and Robotics*, pages 435–439. Citeseer.
- Maja Djikic, Keith Oatley, Sara Zoeterman, and Jordan B Peterson. 2009a. Defenseless against art? impact of reading fiction on emotion in avoidantly attached individuals. *Journal of Research in Personality*, 43(1):14–17.
- Maja Djikic, Keith Oatley, Sara Zoeterman, and Jordan B Peterson. 2009b. On being moved by art: How reading fiction transforms the self. *Creativity Research Journal*, 21(1):24–29.
- Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226. ACM.
- Leah East, Debra Jackson, Louise O'Brien, and Kathleen Peters. 2010. [Storytelling: an approach that can help to develop resilience](#). *Nurse Researcher (through 2013)*, 17(3):17–25.
- David K Elson. 2012. *Modeling narrative discourse*. Columbia University.
- Malte Elson, Johannes Breuer, James D Ivory, and Thorsten Quandt. 2014. More than stories with buttons: Narrative, mechanics, and context as determinants of player experience in digital games. *Journal of Communication*, 64(3):521–542.
- Chris Fairclough and Pádraig Cunningham. 2004. Ai structuralist storytelling in computer games. Technical report, Trinity College Dublin, Department of Computer Science.

- Ethan Fast, Tina Vachovsky, and Michael S Bernstein. 2016. Shirtless and dangerous: Quantifying linguistic signals of gender bias in an online fiction writing community. In *ICWSM*, pages 112–120.
- Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. 2017. Style transfer in text: Exploration and evaluation. *arXiv preprint arXiv:1711.06861*.
- Shira Gabriel and Ariana F Young. 2011. Becoming a vampire without being bitten: The narrative collective-assimilation hypothesis. *Psychological Science*, 22(8):990–994.
- Traci Gillig and Sheila Murphy. 2016. Fostering support for lgbtq youth? the effects of a gay adolescent media portrayal on young viewers. *International Journal of Communication*, 10:23.
- Melanie C Green, Timothy C Brock, and Geoff F Kaufman. 2004. Understanding media enjoyment: The role of transportation into narrative worlds. *Communication Theory*, 14(4):311–327.
- Melanie C Green, Sheryl Kass, Jana Carrey, Benjamin Herzig, Ryan Feeney, and John Sabini. 2008. Transportation across media: Repeated exposure to print and film. *Media Psychology*, 11(4):512–539.
- Melanie C Green, Jeffrey J Strange, and Timothy C Brock. 2003. *Narrative impact: Social and cognitive foundations*. Taylor & Francis.
- Kia Groom. 2015. [A call to arms: Bite the hand that feeds you](#).
- Matthew Guzdial, Brent Harrison, Boyang Li, and Mark Riedl. 2015. Crowdsourcing open interactive narrative. In *FDG*.
- Mengqiao Han, Ou Wu, and Zhendong Niu. 2017. Un-supervised automatic text style transfer using lstm. In *National CCF Conference on Natural Language Processing and Chinese Computing*, pages 281–292. Springer.
- Stacey Hand and Duane Varan. 2008. Interactive narratives: Exploring the links between empathy, interactivity and structure. In *European Conference on Interactive Television*, pages 11–19. Springer.
- D Fox Harrell and Jichen Zhu. 2009. Agency play: Dimensions of agency for interactive narrative design. In *AAAI spring symposium: Intelligent narrative technologies II*, pages 44–52.
- Paul L Harris. 2000. Understanding childrens worlds: The work of the imagination. *Oxfo Blackwell*.
- Dorothee Hefner, Christoph Klimmt, and Peter Vorderer. 2007. Identification with the player character as determinant of video game enjoyment. In *Entertainment computing–ICEC 2007*, pages 39–48. Springer.
- Fritz Heider and Marianne Simmel. 1944. An experimental study of apparent behavior. *The American journal of psychology*, 57(2):243–259.
- Hal E Hershfield, Susanne Scheibe, Tamara L Sims, and Laura L Carstensen. 2013. When feeling bad can be good: Mixed emotions benefit physical health across adulthood. *Social psychological and personality science*, 4(1):54–61.
- Bell Hooks et al. 2006. *Black looks: Race and representation*. Academic Internet Pub Inc.
- Ting-Hao Kenneth Huang, Francis Ferraro, Nasrin Mostafazadeh, Ishan Misra, Aishwarya Agrawal, Jacob Devlin, Ross Girshick, Xiaodong He, Pushmeet Kohli, Dhruv Batra, et al. 2016. Visual storytelling. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1233–1239.
- Keenan M Jenkins. 2014. *Choose your own adventure: Interactive narratives and attitude change*. Ph.D. thesis, The University of North Carolina at Chapel Hill.
- Carl Gustav Jung. 1964. *Man and his symbols*. Laurel.
- Jad Kabbara and Jackie Chi Kit Cheung. 2016. Stylistic transfer in natural language generation systems using recurrent neural networks. In *Proceedings of the Workshop on Uphill Battles in Language Processing: Scaling Early Achievements to Robust Methods*, pages 43–47.
- Geoff F. Kaufman and Lisa K. Libby. 2012. Changing beliefs and behavior through experience-taking. *Journal of Personality and Social Psychology*, 103(2):1–19.
- David Comer Kidd and Emanuele Castano. 2013. Reading literary fiction improves theory of mind. *Science*, 342(6156):377–380.
- Edith Law, Ming Yin, Joslin Goh, Kevin Chen, Michael A. Terry, and Krzysztof Z. Gajos. 2016. Curiosity killed the cat, but makes crowdwork better. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, CHI '16*, pages 4098–4110, New York, NY, USA. ACM.
- Michael Lebowitz. 1987. Planning stories. In *Proceedings of the 9th annual conference of the cognitive science society*, pages 234–242.
- Boyang Li, Stephen Lee-Urban, Darren Scott Appling, and Mark O Riedl. 2012. Crowdsourcing narrative intelligence. *Advances in Cognitive Systems*, 2(1).
- Boyang Li, Stephen Lee-Urban, and Mark Riedl. 2013. Crowdsourcing interactive fiction games. In *FDG*, pages 431–432.
- Boyang Li, Mohini Thakkar, Yijie Wang, and Mark O Riedl. 2014. Data-driven alibi story telling for social believability. *Social Believability in Games*.

- Brian Magerko. 2006. Player modeling in the interactive drama architecture. *Department of Computer Science and Engineering, University of Michigan*.
- Inderjeet Mani. 2012. Computational modeling of narrative. *Synthesis Lectures on Human Language Technologies*, 5(3):1–142.
- Marie-Louise Mares, Mary Beth Oliver, and Joanne Cantor. 2008. Age differences in adults’ emotional motivations for exposure to films. *Media Psychology*, 11(4):488–511.
- Michael Mateas and Andrew Stern. 2003. Façade: An experiment in building a fully-realized interactive drama. In *Game developers conference*, volume 2, pages 4–8.
- Wook-Hee Min, Eok-Soo Shim, Yeo-Jin Kim, and Yun-Gyung Cheong. 2008. [Planning-integrated story graph for interactive narratives](#). In *Proceedings of the 2Nd ACM International Workshop on Story Representation, Mechanism and Context*, SRMC ’08, pages 27–32, New York, NY, USA. ACM.
- Janet Murray. 2004. From game-story to cyberdrama. *First person: New media as story, performance, and game*, 1:2–11.
- Jeeyun Oh, Mun-Young Chung, and Sangyong Han. 2014. The more control, the better? *Journal of Media Psychology*.
- Summer Okoye. 2016. [The black commodity](#).
- Mary Beth Oliver. 1993. Exploring the paradox of the enjoyment of sad films. *Human Communication Research*, 19(3):315–342.
- M.B. Oliver. 2003. Mood management and selective exposure. In Jennings Bryan, David Roskos-Ewoldsen, and Joanne Cantor, editors, *Communication and emotion: Essays in honor of Dolf Zillmann*. Routledge.
- Jessica Ouyang and Kathleen McKeown. 2015. Modeling reportable events as turning points in narrative. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2149–2158.
- Vivian Gussin Paley. 2009. *A child’s work: The importance of fantasy play*. University of Chicago Press.
- James W Pennebaker. 1997. Writing about emotional experiences as a therapeutic process. *Psychological science*, 8(3):162–166.
- Julie Porteous and Marc Cavazza. 2009. Controlling narrative generation with planning trajectories: the role of constraints. In *Joint International Conference on Interactive Digital Storytelling*, pages 234–245. Springer.
- Gerald Prince. 1974. *A grammar of stories: An introduction*, volume 13. Walter de Gruyter.
- Christopher Purdy and Mark O Riedl. 2016. Reading between the lines: Using plot graphs to draw inferences from stories. In *International Conference on Interactive Digital Storytelling*, pages 197–208. Springer.
- Andrew J Reagan. 2017. *Towards a science of human stories: using sentiment analysis and emotional arcs to understand the building blocks of complex social systems*. Ph.D. thesis, The University of Vermont and State Agricultural College.
- Mark O Riedl and Brent Harrison. 2016. Using stories to teach human values to artificial agents. In *AAAI Workshop: AI, Ethics, and Society*.
- Mark O Riedl and Robert Michael Young. 2010. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research*, 39:217–268.
- Ute Ritterfeld and Seung-A Jin. 2006. Addressing media stigma for people experiencing mental illness using an entertainment-education strategy. *Journal of health psychology*, 11(2):247–267.
- Frank Rose. 2012. *The art of immersion: How the digital generation is remaking Hollywood, Madison Avenue, and the way we tell stories*. WW Norton & Company.
- Alyssa Rosenberg. 2012. [From Lena Dunham to Junot Díaz: How to write people who aren’t you](#).
- Paul Rozin, Lily Guillot, Katrina Fincher, Alexander Rozin, and Eli Tsukayama. 2013. Glad to be sad, and other examples of benign masochism. *Judgment and Decision Making*, 8(4):439.
- Marie-Laure Ryan and A-L Rebreyend. 2013. From narrative games to playable stories. *Nouvelle revue desthétique*, (1):37–50.
- Marie-Laure Ryan et al. 2007. Toward a definition of narrative. *The Cambridge companion to narrative*, pages 22–35.
- Kimiko Ryokai and Justine Cassell. 1999. Storymat: a play space for collaborative storytelling. In *CHI’99 extended abstracts on Human factors in computing systems*, pages 272–273. ACM.
- Maarten Sap, Marcella Cindy Prasettio, Ari Holtzman, Hannah Rashkin, and Yejin Choi. 2017. Connotation frames of power and agency in modern films. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2329–2334.
- Roger C Schank and Robert P Abelson. 1975. Scripts, plans, and knowledge. In *IJCAI*, pages 151–157.
- Roger C Schank et al. 1975. Sam—a story understander. Research report no. 43.
- Scott Selisker. 2015. The Bechdel test and the social form of character networks. *New Literary History*, 46(3):505–523.

- Nisi Shawl and Cynthia Ward. 2005. *Writing the Other*. Seattle, Washington: Aqueduct Press.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. In *Advances in Neural Information Processing Systems*, pages 6833–6844.
- Michael D Slater, Donna Rouner, and Marilee Long. 2006. Television dramas and support for controversial public policies: Effects and mechanisms. *Journal of Communication*, 56(2):235–252.
- Zadie Smith. 2009. Speaking in tongues. *The New York review of books*, 26:1–16.
- Jiyeon So and Robin Nabi. 2013. Reduction of perceived social distance as an explanation for media’s influence on personal risk perceptions: A test of the risk convergence model. *Human Communication Research*, 39(3):317–338.
- Brian Sutton-Smith. 1986. Children’s fiction making.
- Brian Sutton-Smith. 2012. *The folkstories of children*. University of Pennsylvania Press.
- Ivo Swartjes and Mariët Theune. 2006. A fabula model for emergent narrative. In *International Conference on Technologies for Interactive Digital Storytelling and Entertainment*, pages 49–60. Springer.
- Jawad Syed. 2016. Oscars so white: an institutional racism perspective. *Counterpunch*.
- Abraham Tesser, Joanne V Wood, and Diederik A Stapel. 2005. *Building, defending, and regulating the self: A psychological perspective*. Psychology Press.
- David Thue, Vadim Bulitko, Marcia Spetch, and Eric Wasylishen. 2007. Interactive storytelling: A player modelling approach. In *AIIDE*, pages 43–48.
- Tyler Tsay, Katherine Frain, and Serafima Fedorova. 2015. [Best & worst literary moments of 2015](#).
- Frederik Van Broeckhoven, Joachim Vlieghe, and Olga De Troyer. 2015. Using a controlled natural language for specifying the narratives of serious games. In *International conference on interactive digital storytelling*, pages 142–153. Springer.
- Rob Voigt, Nicholas P Camp, Vinodkumar Prabhakaran, William L Hamilton, Rebecca C Hetey, Camilla M Griffiths, David Jurgens, Dan Jurafsky, and Jennifer L Eberhardt. 2017. Language from police body camera footage shows racial disparities in officer respect. *Proceedings of the National Academy of Sciences*, page 201702413.
- Kurt Vonnegut. 1970. [The shape of stories](#).
- R Michael Young, Mark O Riedl, Mark Branly, Arnav Jhala, RJ Martin, and CJ Saretto. 2004. An architecture for integrating plan-based behavior generation with interactive game environments. *Journal of Game Development*, 1(1):51–70.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2017. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. *arXiv preprint arXiv:1707.09457*.