Science Considered Helpful

R. Michael Young

Entertainment Arts and Engineering Program and
School of Computing
University of Utah
Salt Lake City, UT, USA, 84103 young@eae.utah.edu
http://liquidnarrative.cs.utah.edu

Abstract. As the interactive narrative community continues to mature, discussions are beginning in which we debate the relative merits of differing methodologies, discuss priorities around classes of problems and look at epistemological questions that arise from what we perceive as limitations of our work. Horswill's *Science Considered Harmful* initiated a conversation around the role of science in the advancement of knowledge in our field, putting forward the idea that a scientific mindset restricts our ability to progress. In this paper, I respond, arguing that science, and more generally scientific rigor and the kind of results that it produces, are well served by a discourse that makes productive distinctions between such things as science and not science. In particular, I argue that such a thing as a science of narrative exists, that scientific work is an important way to advance our knowledge of computational models of narrative and that scholarly practice around interactive narrative research does not need to be viewed as only scientific or as only artistic/aesthetic.

Keywords: Interactive Narrative · Sciences of the Artificial · Theoretical Foundations · Methodology, Research and Scholarship.

1 Introduction

In 2013, Ian Horswill published a position paper [26] in which he raised concerns around the role of scientific research and, more broadly, a scientific mindset that he thought was having negative effects in the community of scholars that were using a computational lens to study interactive narrative. His purpose, as he stated, was two-fold. First, he wished to initiate a dialog around the nature of the work done by this community, and second, he wanted to advance the idea that a scientific view in this research community was *harmful* to progress. He titled his paper *Science Considered Harmful*.

I consider many of the points that Horswill makes valid. In fact, they are critical for those of us working in this community to consider as we decide upon courses of exploration for our research programs. However, I also consider many of the points he makes along the way incorrect. Further, in justifying his views, he attributes motives and context to researchers in the community that are inaccurate. At points, he advances a false dichotomy in which the community

is characterized, like a scientific version of Peer Gynt, as scientifically All or Nothing. And he characterizes scientific approaches to research with negative properties that may be characteristics of poor work but not attributable to science or to the processes of scientific inquiry in our field.

Since the publication of Horswill's paper, there has been unfortunately little dialog around the role of a scientific mindset within our community. There has been no published responses engaging the discussion he initiated and called to continue. But this conversation is of great importance to us as a field, both because the alternative methods of scholarship that Horswill promotes should not be marginalized or diminished by advocates of science *and* because a clear understanding of the distinct values of the different methodologies will help us advance all of them.

I find myself, then, in the surprising position of defending science and scientific inquiry in interactive narrative research. In this paper, I try to answer to some of the questions Horswill raised about the nature of a science of interactive narrative, basing my responses on elements of the work in our research community. But I also describe a perspective on the nature of scholarship in the space of computational narrative and the use of the term science when referring to the study of artificial, rather than natural phenomena.

In *Science Considered Harmful*, Horswill asks four questions about scientific work and then describes a contrasting way of doing scholarly work around computational narrative. His phrasing and the discussion he provides around his questions imply, at least in my own mind, that a scientific view is inappropriate when approaching computational narrative. In the rest of this paper, I try to respond to his questions with concrete answers, discuss existing scholarship that I think is relevant to support my answers, provide examples of how the principles that prompt my answers have been applied and then add to Horswill's call for further engagement on these issues.

There are two conventions that I adopted in the formatting of this paper that warrant short discussion. First, I leave out explicit citation to quotes from *Science Considered Harmful*. Throughout the paper, offset text appearing in italics indicates that the content is a quote from this source. Second, I cite a number of examples of the types of scientific work I mention in my discussion. These works are meant as exemplars; space limitations prevent me from citing the often many, many other relevant papers in any given area. I apologize for these obvious omissions.

2 What Would a Science of Narrative Even Mean?

Horswill asks this question at the start of his paper, as the header of his second section. The question's phrasing presupposes that a science of narrative is a concept that lies far from a direct definition. In my view, there is a ready definition of a science of narrative, drawn from ideas at the core of scholarly work within AI, cognitive science, design and other disciplines. This definition builds directly upon Herbert Simon's model of the *sciences of the artificial* [46]. As

Simon says, the task of an empirical science is to discover and verify invariants in the phenomena under study. A science of the artificial is the application of this pursuit to artifacts rather than to natural phenomena. An artifact can be thought of as "a meeting point, an 'interface' in today's terms between an 'inner' environment, the substance and organization of the artifact itself, and an 'outer' environment, the surroundings in which it operates. If the inner environment is appropriate to the outer environment, or vice versa, the artifact will serve its intended purpose." [46, page 6].

In my model of the science of narrative as a science of the artificial, narratives are the artifacts that stand in relation between inner and outer environments. Inner environments are the story worlds, their characters and events. The narrative as interface is composed of narrative discourse – text, film, chant, game – that abstracts, filters and presents the narrative to the outer environment. And that outer environment, then is composed of the cognitive, affective, social and broader context of readers' narrative interaction and experience. As a science, then, computational narrative works to identify invariants that hold between some or all of the many elements of each of these three components.

The connection that I'm sketching here between Simon's sciences of the artificial and a science of narrative is relatively specific, and may be a view that is not widely held by others, even others who would characterize themselves as doing science. My intent, however, is not to be prescriptive or definitive, but to give one example of what a science of narrative would even mean.

One thing that a science of narrative does *not* mean is what Horswill asserts that it does, that the idea of science is designed to keep certain kinds of work from being done. From his perspective, Horswill feels that a communal orientation towards scientific thinking is acting like a knife to cut away parts of a scholarly community that do not adhere to an overly restrictive set of norms.

To frame what a science of narrative means in examples of existing research programs, let me first frame a broader notion of scholarship around computational narrative. Boyer [5] provides a compelling model of academic scholarship in the large, in which scholarly work is roughly broken into four distinct categories:

- *The scholarship of discovery*, involving original research expanding our knowledge of a given field.
- *The scholarship of integration* involving the processes of knowledge synthesis across disciplines, sub-fields, or time
- *The scholarship of engagement* involving the application of a scholar's expertise to real-world problems, with results that can be shared with and/or evaluated by peers
- *The scholarship of teaching and learning* involving the search for innovative teaching and learning processes in a way that affords public sharing, adoption and evaluation by others.

Boyer's model of scholarship is an integrative one; no category is exclusive of the others. As an example of a long-running research program that can be

characterized using this model, consider the work done by me and my colleagues in the Liquid Narrative Group (e.g., [42, 27, 11, 56, 36, 54]). In the work done in our lab, a science of narrative lies mostly in three of Boyer's categories. It falls mainly within Boyer's scholarship of discovery (the conventional model of scientific inquiry you might see in computer science research) as well as within the scholarship of integration (as this work in computational narrative integrates work from computer science, psychology, narrative theory, game design, cinematography, linguistics, and other disciplines). My work involves, to a lesser degree, a scholarship of application, because a smaller portion of the work there involves the construction of artifacts, either tools or narratives, that are built with sharing as the primary purpose. A model that would seem to align more with Horswill's perspective might be distributed in a kind of inverse of their efforts, with scholarship falling more directly in the third category, engagement, and also in the second, with lesser contact with the first, the scholarship of discovery.

I point out Boyer's model to make the point that perspectives on the nature of research (or, as Boyer makes the shift, *scholarship*) need not be dogmatic all-or-nothing categorizations, and that even within a single field, there are modes of scholarship that differ. Scholarship need not be colonial, exclusive or Darwinian. Academic work doesn't need to be forced into a single category to have academic or scholarly value and work in one category doesn't by its nature drive out work in others.

3 What Do We Mean by Narrative?

Horswill points out that narrative is a broad phenomenon, and I agree. He says that the definition of narrative is difficult to come to universal agreement upon, and I also agree. Narrative theorists, critical studies scholars and authors that reflect on the structure of narrative rarely specify the concepts they use with as much precision as do computer scientists, even when they agree on terms and concepts. Because precise definitions are lacking, many computational researchers that work in the area of narrative identify specific, somewhat stand-alone elements of narrative identified by previous scholars, and rely on characterizations of those elements from narrative theory that, while not universally agreed upon, are well-understood and well-situated relative to contrasting formulations.

Horswill implies that advocates of a scientific approach view narrative as a clearly defined natural kind, and because it's not, he finds their work problematic. I'm not familiar with arguments from computational narrative research that claim that there is a single model of narrative that forms a natural kind. Yes, psychologists and other proponents of Fisher's *narrative paradigm* [19] view narrative as a fundamental mode of understanding the world around us [6]. Psychologists have demonstrated empirically that much of the cognitive and neural activity involved in understanding and producing narrative is increasingly well-understood across genres and media [31]. And there are a number of computational narrative researchers that have stated that these results help motivate the

relevance of their work. But proposing a model of a narrative phenomenon is a far thing from asserting that the model is the phenomenon's comprehensive and final characterization.

Nevertheless, Horswill argues that thinking about narrative as a natural kind when it is not could be problematic because trying to maintain a science around a concept with multiple disjoint definitions would mean that some methodologies employed in that effort would fail. He says that this would force us either to miss out on some methodologies by centering on one definition or compelling us to use multiple methodologies.

First, the use of multiple methodologies within a single field should not be seen a problem. In fact, a core portion of any scientific progress is the development of new methodologies to seek new answers to old problems, or to seek initial answers to problems that haven't been tackled before. Competing results, drawn from multiple methodologies addressing a common phenomenon, are *valued* for the ability to foster constructive competition between theories needed for a healthy science. Our field is young, but is just starting to see multiple methodologies used to answer core problems (e.g., data-driven [28, 24] versus declarative approaches [57] to knowledge authoring for story domains). For a scientist, this is a positive thing.

But more centrally, Horswill is wrong in his underlying assumption here that we need a completely proscribed definition of a phenomenon to begin to construct knowledge around it. As philosophers of science are quick to point out, we construct our models incrementally, and we know that those models are bound to be lacking because our understanding is almost always partial. Every area of scientific inquiry has started out in ignorance about its core phenomena and proceeded to build knowledge incrementally. We don't have to travel far to see how this approach plays out, as researchers and philosophers in the 1980s debated the nature of intelligence, artificial and natural, in long discussions (see Nilsson's summary of arguments by Dreyfus, Minsky, Penrose, Searle and others in his book *The Quest for Artificial Intelligence* [37, chapter VI]). Today, these arguments are mostly relegated to historical reading, while AI makes rapid advances in both academic and commercial contexts.

Fortunately, scholars interested in bringing scientific approaches to the exploration of narrative do not lack well-articulated (though often formally imprecise) descriptions of narrative's form and function. One methodology that has been adopted is for a researcher to select a single characterization of a narrative phenomenon, provide a computational refinement of the characterization and then use that refinement either generatively or analytically. The resulting output – a novel narrative artifact or the characterization of an existing corpus – is then evaluated and used to measure refinement.

As an example, in work that provides computational models of suspense [10, 38], researchers based their work on an existing characterization of narrative comprehension. The characterization models readers acting as problem-solvers [20] and views suspense as tension arising from a reader's anticipation of the success or failure of a protagonist's plans [15]. Computational work, then, operational-

ized these existing models, and experimental work with the resulting systems provided further support for the psychological approaches while also posing new questions about the cognitive processes involved.

4 What Do We Mean by Science?

In thinking about Horswill's question, I realize that its a question to which we should give much more consideration. As an initial framing, let me draw from the on-going discussion in the computer science (CS) community about the scientific nature of CS research. In that discussion, Peter J. Denning gives a characterization of scientific fields typically having the following structure [17]:

- they are organized to understand, exploit, and cope with a pervasive phenomenon
- they encompass natural and artificial processes of the phenomenon.
- they are codified structured body of knowledge.
- they show a commitment to experimental methods for discovery and validation.
- they value reproducibility of results.
- they advance falsifiability of hypotheses and models.
- they have an ability to make reliable predictions, some of which are surprising.

How does the computational narrative work done in our community align with these characteristics? First, I don't think that anyone would argue that narrative is not a pervasive phenomenon. And it should be clear that the work done by the computational narrative community is, to a great extent, a communal effort to understand, exploit and cope with narrative from a computational perspective.

There are clearly many people working to address natural and artificial aspects of narrative, often addressing both at the same time. Psychologists work to understand narrative understanding, from neurological bases [48] to cognitive work around things like inferencing and affective responses to narrative [33, 47], memory [41], goal tracking [32], and many others. AI researchers work to characterize artificial aspects, like the structure and form of narrative seen in cinematic lighting [18], character dialog [29] and the structure and form of plot lines [51, 21, 8]. Even some narrative theorists, who provide an almost analytic-only view on narrative, also develop a formalized and scientific views of narrative [25, 43].

The field of interactive digital storytelling is relatively new. Initial steps grew out of the critical mass catalyzed roughly in 1999 at the AAAI Fall Symposium on Narrative Intelligence [34]. The field has continued to codify its knowledge through traditional publication venues that are intentionally broad in scope (e.g., the International Conference on Interactive Digital Storytelling, the Workshop on Computational Models of Narrative), games-focused (e.g., the conference of the Digital Interactive Games Research Association, the Foundations of Digital

Games conference), AI-centric (e.g., IEEE conference on Computational Intelligence in Games, The AAAI conference on AI and Interactive Digital Entertainment, the workshop on Intelligent Narrative Technologies), or others. Within, take an example of planning-based methods to characterize generative models of plot [39], where much of the work from my group has been directed. Just there, there is the creation of planning based methods [57], the identification of limitations to produce intentional planning [42],the identification of efficient algorithms to produce the same plans [23], the extension to characterize conflict in stories [54], the extension to express character personality through choice of actions [3] and most recently extensions to address character belief and mistaken actions [52], with work also here by Porteous [51], Shrivani and his collaborators [44], and interactions between character beliefs and intention management [55].

The field, as Horswill points out, has a strong commitment to experimental work, particularly in evaluation. Work done by psychologists is particularly experimental in both discovery and validation. As one example, Graesser and Franklin's QUEST model of question-answering in the context of stories [22] is a computational cognitive model of narrative – essentially, a model of narrative understanding that uses a computational model (in a broad sense) of narrative knowledge and the processes that make use of it. The QUEST model makes predictions about how people answer questions about stories after reading them, and makes specific predictions about that behavior that have been born out in years of experimental testing. Because those predictions are both specific and shown to be reliable, researchers on the computational side of narrative generation have been able to use them as proxies for effective human comprehension in a wide range of work (e.g., [42, 27, 1]).

Much of the work on the generation of narrative also relies on experimental methods, typically involving human subjects and characterizations of the efficacy of new algorithms to build narrative artifacts or experiences (e.g., [35, 58, 30]). Other work uses experimental evaluation, but relies on analytic or statistical means for validation (e.g., [53, 9]). The falsifiability of experimental designs is a critical element in the progression of our understanding in the science of narrative. Unexpected or negative results prompt theory revision, and give rise to new experiments to gauge how proposed new models match the interactions of the Simonian tripartite view of narrative. As one example, Christian and Young's work [11] sought to map plan-based narrative structures to cognitive models used to comprehend cinematic presentations of stories. Their work showed support only for certain types of story structures, and this limitation then prompted follow-on work by Cardona-Rivera and his collaborators [7], who sought to better understand Christian and Young's results by performing comparable experiments over a range of varying mappings and cognitive models.

While not an issue pointed to by Horswill, the lack of emphasis on reproducibility appears to me as a scientific weakness of the work done in computational narrative. Because our area is so new, researchers have a rich field of open questions to choose from and no real canon of successful methods to build upon

or to challenge. Consequently, the draw to explore a novel method of one's own devising is so great and the value of exploring undiscovered country so high that we're not stopping to do the hard work around reproducability that already lies within our reach (see the work of Tearse,*et al* [50], in the production of *Minstrel Remixed* as a notable counter-example).

As for our science leading towards surprising predictions, surprise is certainly in the eye of the beholder. Is it a surprise that story structures, plan data structures and mental models storylines align to such an extent that the data structures can predict cognitive processing during story comprehension [42, 11, 27, 36, 7]? It was not to me, although to some it might be. As another example, narrative psychologist Radvansky has developed a cognitive model of how narrative consumers build mental models of a story world and its unfolding events, called the *event horizon* model [40]. This model makes accurate predictions about, among other things, the way that readers manage their story memories when the current location in a story change during reading. Radvansky predicted that these changes would also be prompted when the readers themselves moved from one location to another. Surprisingly, at least to me, these predictions were verified [41].

5 On the Relation of Cases to Methodology in Computational Narrative

> *If you can't gather data across all possible designs and users, then the alternative is to look as deeply as possible at designs you do examine; to glean as much as possible from a specific encounter with a specific piece. So art training and theory are traditionally focused on examining specific pieces, specific cases. And as such, it is in some ways closer to the practice found in law and business than in engineering.*

Horswill makes an important point about the relationship between artistic and/or critical training and case-oriented methods used in law and business. I disagree with the premise of the paragraph, that if you can't gather data across all possible designs and users, then *the* alternative (emphasis mine) is to examine a small number of instances. But let's take this argument to a not-so-extreme extension. If physicists are not able to gather data across all possible instances of a phenomenon, then Horswill would have them resort only to studying a small number of cases. When the space of a phenomenon includes all of time and space, physical scientists can't come close to this bar. Studying cases is often a great way to conduct exploratory work, and physical scientists working in an area where knowledge is just emerging may do just that. But it's with an eye towards theory formation, which is a key part of the characterization of experimental computer science and its iteration spelled out by Cohen [13] and others. As Cohen and Howe [14] say when describing the role of evaluation in AI research, " For the individual, evaluation can tell us how and why our methods and programs work and, so, tell us how our research should proceed. For the community, evaluation

expedites the understanding of available methods and, so, their integration into further research." This is as true for the study of interactive digital storytelling as it is for the field of AI.

> *Although the goal is in some sense to develop a sufficient understanding to be able to invert the mapping, and predict how a hypothetical design might affect a viewers experience, that understanding is nonetheless always partial and contingent, so the actual creation process is iterative, with each iteration forming a case to be examined and used in the formation of the next iteration.*
>
> *This presents a serious challenge for AI work in support of the arts because computer science is traditionally, and understandably, resistant to work on cases. One cannot generally get a Ph.D. for writing a program; one gets a Ph.D. for developing a general theory that will help others write some class of programs. But I see no alternative but to relax this constraint for AI work that have specifically aesthetic goals. We have to be open to the case because the case is all we have.*

AI work to support the arts is one view of what a science of narrative does: supporting the arts. But a science is geared towards an increased understanding of a phenomenon, rather than a particular engineering approach, which might be an application of that understanding to a problem like creation of art. Framing the scholarship under Boyer's model, Horswill's view seems to be centered here on a scholarship of application, but a science of narrative affords scholarship of discovery and integration as well.

Further, not all computational work on narrative has explicitly aesthetic goals. [1] But the idea that understanding how a hypothetical design might affect a viewer's experience, is precisely what Simon talks about in characterizing design as a science of the artificial [45].

6 Horswill's Conclusion

> *My concern is that we're at risk of letting our methods drive our inquiries rather than letting our inquiries drive our methods. I have on several occasions been told by students at conferences that they would like to make a game or artwork that demonstrates the ideas from their thesis but that that's not science.*
> *That makes me want to shoot myself.*

Those students are correct. Building an artifact is, on its own, not science (or rather, not doing science). There are a lot of things my students wish to

[1] Of course, one might argue that not all work in the arts has exclusively aesthetic goals (for example, see [4]).

do while being students, [2] and some of those, while amazing, intellectually demanding and capable of yielding insight into research questions, just are not science. That doesn't mean that these activities would be inappropriate, or that they would not lead to increased knowledge about narrative. It means that they would not be direct contributors on their own to the idea of forming hypotheses about the relevant invariants and testing them out computationally. Building an artifact is not anti-science. Building an artifact that demonstrates the range of relationships you're exploring in your research *is* science, just like building a software simulation of a physical phenomenon might be part of a physicist's exploration of a physical science.

> *If we limit ourselves to using only analytics and psychology-style human subjects experiments, we are at best crippling ourselves and at worst deceiving ourselves with a false sense of rigor. In the worst case, we could end up creating (unironically) a kind of interactive narrative version of Komar and Melamid's Peoples Choice series (Wypijewski, 1997) in which they commissioned a market research company to study the characteristics most- and least-wanted in paintings by different demographic groups, then painted the results. Their United States: Most Wanted Painting (1994) is a realist landscape incorporating trees, water, mountains, George Washington, a happy family, and a pair of deer.*

I object to two points Horswill makes here. First, its unclear to me where the rigor applied in the methods themselves falls short. Second, Horwill's suggestion that use of these methods will lead us to overly specific methods that are too narrowly focused seems to fly in the face of the history of analytic and experimental science. Psychologists make heavy use of controlled human-subjects experiments, and they've validated a significant set of principals that apply broadly (rather than learning how to model the 300 freshmen psychology students that participated in any given experiment).

> *I recognize there is Realpolitik to science. We must all make ourselves intelligible and respectable to our departments, Deans, hiring committees, and funders; and that may require us to use methods and metrics we don't always believe in. But that doesn't mean we have to limit ourselves to only using those methods.*

Scholars should never use methods and metrics that they don't believe in. The integrity of a researcher (including their sincere confidence in their work) is critical to the believablity of results on which progress in a field is based. Belief in the outcomes of research, dependent on the methods and metrics employed, is at the core of any researcher's work. There is a hint here in what Horswill is saying that the savvy and successful academic will engage in *Realpolitik*, pitching the science story *in order to* appeal to a tenure process or a funding review. If I do science, and I work in a system where scientific results are rewarded, is it

[2] Aside from Fortnite.

correct to automatically characterize me as *conforming to get ahead* and then damn me because of it? That would effectively damn scientific scholarship across a good portion of the academy. Horswill paints with too broad a brush here. By implying that academics, especially junior ones, can be forgiven for presenting as scientists in order to play the tenure game, the sympathetic argument admits without examination the notion that all those folks are well-meaning, worthy of our sympathy, but disingenuous.

I do strongly agree with Horswill that my methods do not need to be your methods, and that valid knowledge can be gained by a diverse set of approaches. Horswill's tack, however, is to call for us to use methods and metrics that are not scientific. While adopting this approach might fit for some scholars' work, for communities of researchers in computational narrative that *are* scientific, those methods stand apart. They don't address many of Denning's properties that scientific fields rely upon. As such, they have limited power to advance relevant types of knowledge. Specifically, they are limited in their means of validation for general claims, they are challenged to support reproducibility, they may not advance hypotheses nor contribute to a process where wrong hypotheses can be falsified, and they don't particularly serve to make reliable predictions about narrative phenomena.

7 Let Me Sum Up

> *In the meantime, I look forward to the day I stop hearing the phrase "I'd like to do X, but that's not science."*

To pursue science is not to deny art. When your day job is to do science, however, you have to pick your battles.

Further, to privilege either science or art over the other in the broader community of computational media researchers would be unjustifiable, short-sighted and would have all of the negative effects that Horswill is concerned about. But to advance artistic methods to be on par with or to serve as replacements for scientific ones *within the scientific community working on narrative* would significantly constrain the power of the claims we can make and the knowledge that we would gain. One might conclude from Horswill's paper that scientific approaches in our field ignore the art of narrative. This is hardly the case. Scientific work is all about digging into artistic elements of narrative. The artful scientific work done by Swanson and his collaborators on learning photo composition preferences from gameplay [49] is a great example. out of many others (e.g., [2, 12, 16]).

Narrative art and narrative science both exist. They are both individuals in their own right, capable of standing apart one from another. But they are also capable of standing together without straining their individual natures. Rather than call for scientific scholars to dial back those properties of their work that bring some of its strongest internal values, I suggest that we form new discussions that stand apart from our artistic or scientific contributions, but that build

upon them. As a potential model, consider SIGGRAPH's inclusion of artistic contributions in the conference's Art Papers track and the Art Gallery format that leverage novel computational advances. Other approaches, like the calls for the ICIDS and FDG conferences, make clear that submissions characterizing critical, computation or aesthetic work are appropriate. By creating these and more venues for dialog, we engage scientists and artists in the creation of a shared body of creative knowledge.

I am in agreement with Horswill on many, many things about our common professional endeavor: the beauty of LaTeX, the primacy of emacs, the use of tabs over spaces. More seriously, I agree with him completely about the critical role in society that we all play as scholars and the special character of work around computational media, play and narrative.

Simon's framing of a science of the artificial is particularly useful for those of us working in computational methods for interactive narrative. As we increase our understanding of the invariants in narrative elements – what relations hold between the design of story worlds, their telling and the resulting experiences of our readers, viewers and players – we can leverage that knowledge to build algorithms that themselves create novel types of playable narrative experiences. Scholarship that contributes to the science of narrative need not come only from the scholarship of exploration, but should be welcomed from work that ranges across all four of Boyer's categories (and, potentially, others).

Scholar of computational narrative, do not shy away from science. In many cases, it has been shown to be helpful.

8 Acknowledgments

This material is based upon work supported in whole or in part with funding from the Laboratory for Analytic Sciences (LAS). Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the LAS and/or any agency or entity of the United States Government.

References

1. Amos-Binks, A., Spain, R., Young, R.M.: Subjective Experience of Intention Revision. In: Advances in Cognitive Systems (2018)
2. Argamon, S., Dubnov, S., Jupp, J., et al.: Style and meaning in language, art, music and design. In: Proceedings of the 2004 AAAI Symposium (2004)
3. Bahamón, J.C., Barot, C., Young, R.M.: A goal-based model of personality for planning-based narrative generation. In: AAAI. pp. 4142–4143 (2015)
4. Bay, M.: Transformers 6: The Last Knight. Paramount Pictures, Murphy, Don and DeSanto, Tom and diBonaventura, Lorenzo and Bryce, Ian (2017)
5. Boyer, E.: scholarship reconsidered: Priorties of the professoriate. Carnegie Foundation for the Advancement of Teaching, New York, NY (1990)
6. Bruner, J.: The narrative construction of reality. Critical inquiry **18**(1), 1–21 (1991)
7. Cardona-Rivera, R.E., Price, T.W., Winer, D., Young, R.M.: Question answering in the context of stories generated by computers. Advances in Cognitive Systems **4**, 227–246 (2016)
8. Cavazza, M., Charles, F., Mead, S.: Character-based interactive storytelling. Intelligent Systems, IEEE **17**(4), 17–24 (2002)
9. Chambers, N., Jurafsky, D.: Unsupervised learning of narrative event chains. In: ACL. vol. 94305, pp. 789–797 (2008)
10. Cheong, Y., Young, R.: A computational model of narrative generation for suspense. In: Working Notes of the AAAI 2006 Workshop on Computational Aesthetics. pp. 8–14 (2006)
11. Christian, D., Young, R.M.: Comparing cognitive and computational models of narrative structure. In: Proceedings of the National Conference on Artificial Intelligence. pp. 385–390. American Association of Artificial Intelligence, AAAI, Menlo Park, CA (2004)
12. Cohen, H.: The further exploits of aaron, painter. Stanford Humanities Review **4**(2), 141–158 (1995)
13. Cohen, P.R.: Empirical methods for artificial intelligence, vol. 139. MIT press Cambridge, MA (1995)
14. Cohen, P.R., Howe, A.E.: How evaluation guides ai research: The message still counts more than the medium. AI magazine **9**(4), 35 (1988)
15. Comisky, P., Bryant, J.: Factors involved in generating suspense. Human Comm. Research **9**(1), 49–58 (1982)
16. Dannenberg, R.B., Raphael, C.: Music score alignment and computer accompaniment. Communications of the ACM **49**(8), 38–43 (2006)
17. Denning, P.J.: The science in computer science. Commun. ACM **56**(5), 35–38 (May 2013). https://doi.org/10.1145/2447976.2447988, http://doi.acm.org/10.1145/2447976.2447988
18. El-Nasr, M.S., Horswill, I.: Automating lighting design for interactive entertainment. CIE **2**(2), 15–15 (2004)
19. Fisher, W.R.: Human communication as narration: Toward a philosophy of reason, value, and action. University of South Carolina Press (1989)
20. Gerrig, R.: Experiencing narrative worlds: On the psychological activities of reading. Yale Univ Press (1993)
21. Gervás, P., Díaz-Agudo, B., Peinado, F., Hervás, R.: Story plot generation based on CBR. Knowledge-Based Systems **18**(4), 235–242 (2005)
22. Graesser, A.C., Franklin, S.P.: Quest: A cognitive model of question answering. Discourse processes **13**(3), 279–303 (1990)

23. Haslum, P.: Narrative planning: Compilations to classical planning. Journal of Artificial Intelligence Research (2012)
24. Hayton, T., Porteous, J., Ferreira, J.F., Lindsay, A., Read, J.: Storyframer: From input stories to output planning models. In: Proceedings of the Workshop on Knowledge Engineering for Planning and Scheduling (KEPS), at The 27th International Conference on Automated Planning and Scheduling (ICAPS). AAAI (2017)
25. Herman, D.: Storytelling and the Sciences of Mind. MIT press (2013)
26. Horswill, I.: Science considered harmful. In: Proceedings of AIIDE (2013)
27. Jhala, A., Young, R.M.: Cinematic visual discourse: Representation, generation, and evaluation. IEEE Transactions on Computational Intelligence and Artificial Intelligence in Games **2**, 69–81 (2010)
28. Li, B., Lee-Urban, S., Johnston, G., Riedl, M.: Story generation with crowdsourced plot graphs. In: AAAI (2013)
29. Lin, G.I., Walker, M.A.: All the world's a stage: Learning character models from film. In: AIIDE (2011)
30. Lukin, S.M., Reed, L., Walker, M.A.: Generating sentence planning variations for story telling. In: SIGDIAL Conference. pp. 188–197 (2015)
31. Magliano, J.P., Loschky, L.C., Clinton, J., Larson, A.M.: Is reading the same as viewing? an exploration of the similarities and differences between processing text- and visually based narratives. In: Miller, B., Cutting, L., McCardle, P. (eds.) Unraveling the Behavioral, Neurobiological, and Genetic Components of Reading Comprehension, chap. 8, pp. 78–90. Brookes Publishing Co., Baltimore, MD (2017)
32. Magliano, J.P., Taylor, H.A., Kim, H.J.J.: When goals collide: Monitoring the goals of multiple characters. Memory & Cognition **33**(8), 1357–1367 (2005)
33. Magliano, J., Dijkstra, K., Zwaan, R.: Generating predictive inferences while viewing a movie. Discourse Processes **22**(3), 199–224 (1996)
34. Mateas, M., Sengers, P.: Narrative intelligence, Papers from the 1999 Fall Symposium. Tech. rep., AAAI Press, Menlo Park, CA (1999), Technical Report FS-99-0
35. Mawhorter, P., Mateas, M., Wardrip-Fruin, N.: Generating relaxed, obvious, and dilemma choices with dunyazad. In: AIIDE. pp. 58–64 (2015)
36. Niehaus, J., Young, R.: A method for generating narrative discourse to prompt inferences. In: Workshop on Intelligent Narrative Technologies 3 (2010)
37. Nilsson, N.J.: The quest for artificial intelligence. Cambridge University Press (2009)
38. O'Neill, B., Riedl, M.: Dramatis: A computational model of suspense. In: AAAI. pp. 944–950. AAAI'14 (2014), http://dl.acm.org/citation.cfm?id=2893873.2894021
39. Porteous, J.: Planning technologies for interactive storytelling. In: Nakatsu, R., Rauterberg, M., Ciancarini, P. (eds.) Handbook of Digital Games and Entertainment Technologies. Springer, Singapore (2016)
40. Radvansky, G.A.: Across the event horizon. Current Directions in Psychological Science **21**(4), 269–272 (2012). https://doi.org/10.1177/0963721412451274, https://doi.org/10.1177/0963721412451274
41. Radvansky, G.A., Copeland, D.E.: Walking through doorways causes forgetting: Situation models and experienced space. Memory & cognition **34**(5), 1150–1156 (2006)
42. Riedl, M., Young, R.M.: Narrative planning: balancing plot and character. JAIR **39**(1), 217–268 (2010)
43. Ryan, M.L.: Possible worlds, artificial intelligence, and narrative theory. Indiana University Press (1991)

44. Shirvani, A., Ware, S.G., Farrell, R.: A possible worlds model of belief for state-space narrative planning. In: Proceedings of the 13th AAAI International Conference on Artificial Intelligence and Interactive Digital Entertainment. pp. 101–107 (2017)
45. Simon, H.A.: The science of design: Creating the artificial. Design Issues **4**(1/2), 67–82 (1988), http://www.jstor.org/stable/1511391
46. Simon, H.A.: The sciences of the artificial. MIT press (1996)
47. Singer, M., Ferreira, F.: Inferring consequences in story comprehansion. Journal of Verbal Learning and Verbal Behavior **22**, 437–448 (1983)
48. Speer, N.K., Zacks, J.M., Reynolds, J.R.: Human brain activity time-locked to narrative event boundaries. Psychological Science **18**(5), 449–455 (2007)
49. Swanson, R., Escoffery, D., Jhala, A.: Learning visual composition preferences from an annotated corpus generated through gameplay. In: CIG. pp. 363–370. IEEE (2012)
50. Tearse, B., Mateas, M., Wardrip-Fruin, N.: Minstrel remixed: a rational reconstruction. In: INT3 '10: Proceedings of the Intelligent Narrative Technologies III Workshop. p. 1–7. ACM, ACM, New York, NY, USA (2010). https://doi.org/http://doi.acm.org/10.1145/1822309.1822321
51. Teutenberg, J., Porteous, J.: Incorporating global and local knowledge in intentional narrative planning. In: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems. pp. 1539–1546. International Foundation for Autonomous Agents and Multiagent Systems (2015)
52. Thorne, B., Young, R.M.: Generating stories that include failed actions by modeling false character beliefs. Tech. rep., University of Utah (2017), eAE TR 2017-001
53. Valls-Vargas, J., Zhu, J., Ontanon, S.: Error analysis in an automated narrative information extraction pipeline. IEEE TCAIG (2016)
54. Ware, S.G., Young, R.M., Harrison, B., Roberts, D.L.: A computational model of narrative conflict at the fabula level. IEEE TCIAIG **6**(3), 271–288 (2014)
55. Young, R.M.: Sketching a generative model of intention management for characters in stories:adding intention management to a belief-driven story planning algorithm. Tech. rep., University of Utah (2017), eAE TR 2017-002
56. Young, R.M.: Story and discourse: A bipartite model of narrative generation in virtual worlds. Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems **8**(2), 177–208 (2010)
57. Young, R.M., Riedl, M.O., Branly, M., Jhala, A., Martin, R., Saretto, C.: An architecture for integrating plan-based behavior generation with interactive game environments. Journal of Game Development **1**(1), 51–70 (2004)
58. Yu, H., Riedl, M.O.: Optimizing players expected enjoyment in interactive stories. In: AIIDE. pp. 100–106 (2015)