
Computational Narrative Intelligence: A Human-Centered Goal for Artificial Intelligence

Mark O. Riedl

Georgia Institute of Technology
Atlanta, GA 30332, USA
riedl@cc.gatech.edu

Abstract

Narrative intelligence is the ability to craft, tell, understand, and respond affectively to stories. We argue that instilling artificial intelligences with computational narrative intelligence affords a number of applications beneficial to humans. We lay out some of the machine learning challenges necessary to solve to achieve computational narrative intelligence. Finally, we argue that computational narrative is a practical step towards machine enculturation, the teaching of sociocultural values to machines.

Author Keywords

Artificial Intelligence, Machine Learning, Narrative Intelligence, Machine Enculturation

ACM Classification Keywords

I.2.0 [Artificial Intelligence]: General

Introduction

Storytelling is an important part of how we, as humans, communicate, entertain, and teach each other. We tell stories dozens of times a day: around the dinner table to share experiences; through fables to teach values; through journalism to communicate important events, and in entertainment movies, novels, and computer games for fun. Stories also motivate people to learn, which is why they form the backbone of training scenarios and case studies at school

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.

CHI'16 Workshop on Human-Centered Machine Learning, May 8, 2016, San Jose, California, USA.

or work.

Despite the importance of storytelling as part of the human experience, computers still cannot reliably create and tell novel stories, nor understand stories told by humans. When computers do tell stories, via an eBook or computer game, they simply regurgitate something written by a human. They do not partake in the culture we are immersed in, as manifested through journalistic news articles, the movies we watch, or the books we read.

Why does it matter that computers cannot create, tell, or understand stories? Artificial intelligence has become more prevalent in our everyday lives. Soon, it will not be unusual for us to interact with more advanced forms of *Siri* or *Cortana* on a daily basis. However, when we use those systems today, we find it to be an *alien* sort of intelligence. The AI makes decisions that sometimes can be hard for us to make sense of. Their failures are often due to the fact that they cannot make sense of what we are trying to accomplish or why.

Narrative intelligence is the ability to craft, tell, understand, and respond affectively to stories. Research in *computational* narrative intelligence seeks to instill narrative intelligence into computers. In doing so, the goal of developing computational narrative intelligence is to make computers better communicators, educators, entertainers, and more capable of relating to us by genuinely understanding our needs. Computational narrative intelligence is as much about human-computer interaction as it is about solving hard artificial intelligence problems.

In this position paper, we enumerate a number of human-centered applications of computational narrative intelligence that may be of benefit to humans interacting with artificial intelligences in the future. We also discuss some of the ma-

chine learning challenges that will need to be overcome through research to achieve computational narrative intelligence. Finally, we describe how computational intelligence can provide a way forward to creating artificial intelligences that are more human-like, better at understanding their human users, and more easily comprehended by human users.

Computational Narrative Intelligence

Winston [37] argues that narrative intelligence is one of the abilities that sets humans apart from other animals and non-human-like artificial intelligences. Research in computational narrative intelligence has sought to create computational intelligences that can answer questions about stories, generate fictional stories and news articles, respond affectively to stories, and represent the knowledge contained in natural language narratives.

Given that humans communicate regularly and naturally through narratives, one of the long-standing challenges of computational narrative intelligence has been to answer questions about stories [32, 25, 36]. Question-answering is a way of verifying that a computer is able to understand what a human is saying. However, question-answering about narrative content is considered to be more challenging than fact-based question-answering due to the causal and temporal relationships between events, which can be complex and are often left implicit. One prerequisite for narrative question-answering is a better understanding of how to represent the knowledge contained in natural language narratives [4, 10, 5].

The flip-side of understanding stories is the creation of novel, fictional story content such as fairy tales and computer game plots [12, 31, 28, 38]. The obvious application of fictional story generation is entertainment. On-demand

narrative generation can maintain a continuous flow of novel content for users to engage with while customizing the content to individual preferences and demands. One may imagine serial novels, serial scripts for TV shows and movies, or serial quests and plotlines in computer games. However, note that even entertainment can convey morals and other pedagogical aspects.

Computational narrative intelligences can also create plausible sounding—but fictional—stories that might happen in the real world [33, 20]. While not meant to be entertaining, the generation of plausible real world stories provides a strong, objective measure of general computational intelligence. Plausible real-world story generation can be used to generate virtually unlimited scenarios for skill mastery in training simulations [39]. Computational narrative intelligences could engage in forensic investigations by hypothesizing about sequences of events that have not been directly observed. Virtual agents, such as virtual health coaches, can appear more life-like and create rapport with humans by sharing fictional vignettes and gossip [2].

Computational narrative intelligence also brings computers one step closer to understanding the human experience and predicting how humans will respond to narrative content. Automated journalists generate narrative texts about real world events and data such as sports and financial reports (e.g., [1]). Automated journalists may benefit from narrative intelligence when determining how best to convey a narrative to different audiences. Going beyond journalism, it is important to note that humans can have very visceral emotional responses to stories. Understanding how the human will interpret and respond to narrative situations has important implications if we wish for computers to avoid accidentally making people upset or anxious. Computers may one day intentionally attempt to induce pleasure, or create a

sense of suspense [26] in both entertainment and journalistic contexts.

Finally, narrative can be used to explain the behavior of artificial intelligences. Any process or procedure can be told as a narrative, so it follows that an AI can describe the means by which it came to a conclusion or the reasons why it performed an action by couching its explanation in narrative terms. As part of an explanatory process, narratives can convey counterfactuals—what would have happened if circumstances had been different. We hypothesize that narrative explanation will be more easily understood by non-expert human operators of artificial intelligence since the human mind is tuned for narrative understanding.

Machine Learning Challenges

Automated story understanding and automated story generation have a long history of pursuit in the field of artificial intelligence. Until recently, most approaches used hand-authored formal models of the story world domain the generator or understander would operate in [32, 23, 18, 25, 13, 31]. This made *open-domain* narrative intelligence—the sort employed by humans—intractable due to knowledge engineering bottlenecks. More recent approaches use machine learning to attempt to automatically acquire and reuse domain models from narrative corpora on the Internet [22, 4, 33] and from crowdsourcing [20].

There are at least four primary challenges related to learning domain models from narrative corpora and using them to create stories or explanations. First, human-written narratives are written to be consumed by other humans. We use *theory of mind* to infer what others are likely to already know and adjust our storytelling accordingly. Thus, human-written narratives collected into a corpus often leave out elements that are assumed to be commonly shared knowl-

edge among other humans but possibly not known by computers. For example, a news corpus may have a story about bank robbery, but that story only has the points that make it unique from other bank robberies and “newsworthy.” In some sense, all stories interesting enough to tell, or to have been told, are outliers from each other, making patterns hard to detect. A machine learning system would never learn about the aspects of the domain model that are most common to all bank robberies.

Stories are often told to highlight an unexpected obstacle or event in an otherwise typical situation. By virtue of telling a story of this sort, one may infer the counterfactual as the norm [14]. Crowdsourcing allows for greater control of the narrative content and can be used to acquire a corpus of *typical* stories about situations at the desired level of granularity [20]. Many children’s books and television shows teach expectations for common situations such as going to a doctor’s office or what to expect on the first day of school.

Second, narrative intelligence is closely associated with commonsense reasoning. It is necessary for both narrative understanding and for narrative generation. Humans learn commonsense knowledge and reasoning through a lifetime of experiences in the real world. Learning commonsense knowledge as been an ongoing challenge in AI and machine learning.

Commonsense knowledge in the form of declarative facts and procedures will be essential in comprehending narratives. Research into automated commonsense knowledge acquisition includes [19, 21, 7, 24]. Images and video also implicitly capture commonsense knowledge (e.g., things fall downward, people kick balls but not bricks, etc.) [35] and techniques that jointly learn from stories with accompanying video or illustrations may provide key insights.

Third, natural language stories written by humans for humans make abundant use of metaphors and metonymy [17]. Decoding the meaning of metaphors and metonymy requires high-level semantic comprehension of the narratives collected into a machine learning corpus.

A few research projects have attempted to use metaphor in the automated generation of stories [15, 34]. Hobbs [16] lays out three general approaches to understanding metaphors: transferring properties from one entity to another, mapping aspects of one thing to another by inference, or mapping aspects of one thing to another by analogy. Analogical mapping has received the most attention in computational narrative intelligence [8, 30, 38].

Fourth, creating stories requires a model of *creativity* as process that transcend straightforward pattern learning. The space of all possible, tellable, and interesting stories is vast. This is one explanation for why the generation of stories via sampling from recurrent neural networks trained on narrative corpora has not fared well to date because of the complexity of human-written narratives and the need for very large training sets. Further, stories make use of long-term causal connections between events that have not been easy to model; long term dependencies mean that stories, and the process of creating stories, are non-Markovian. However, some recent progress has been made in using long short-term memory neural nets that can extract script-like representations from text [27].

Descriptions of human creativity emphasize the blending of two or more mental models to create new concepts. The appeal of *conceptual blending* [9] is the invention of concepts that might never have existed in a data set or even the real world. Conceptual blending shares similarities to unsupervised transfer learning, a critical area of research in machine learning. One example in the domain of creativity

is the blending of two neural nets trained on different aspects of art [11]. However, an equivalent approach has not yet been found for story generation.

Solving these challenges will be necessary in order to achieve a complete, open-domain, computational narrative intelligence that is trained from narrative corpora. In some cases, the challenges are those associated with semantic-level natural language processes. However note that television, movies, dramatic plays, comic books, and illustrated children's books can also be sources of valuable data and require integrated natural language processing and machine vision.

Machine Enculturation

In addition to the applications described earlier—and assuming the above challenges can be met—computational narrative intelligence may present a way forward toward *machine enculturation* [29]. Machine enculturation is the act of instilling social norms, customs, values, and etiquette into computers so that they can (a) more readily relate to us and (b) avoid harming us (physically or psychologically) or creating social disruptions. In a perfect world, humanity would come with a user manual that we could simply scan into a computer. However, for any sufficiently complex domain, such as the real world, manually encoding a comprehensive set of values or rewards in order to recreate sociocultural behavior is intractable.

If sociocultural values are not easily instilled in artificial intelligences, perhaps they can be learned. Instead of a user manual we have the collected works of fiction by different cultures and societies. This collected works give us examples with which to teach an artificial intelligence the “rules” of our societies and cultures. These stories includes the fables or allegorical tales passed down from generation to

generation, such as the tale of George Washington confessing to chopping down a cherry tree. Fictional stories meant to entertain can be viewed as examples of protagonists existing within and enacting the values of the culture to which they belong, from the mundane—eating at a restaurant—to the extreme—saving the world.

Stories are an effective means of conveying complex tacit and experiential knowledge that implicitly encodes social and cultural values [3]. Humans do not need to be trained to communicate via storytelling, nor be trained to decode the knowledge contained within narratives. Computers will likely require require human-level narrative comprehension to mine social and cultural values from fictional and non-fictional narrative texts because those values are rarely made explicit.

The actions of characters in stories can be viewed as demonstrations of socioculturally appropriate behavior under hypothetical situations. Unlike demonstrations, which occur in the environment that the artificial intelligence will operate in, narratives may be more general. This presents some new challenges. Stories written in natural language can contain events and actions that are not executable by an artificial intelligence. Stories are written by humans for humans and thus make use of commonly shared knowledge, leaving many things unstated. Stories frequently skip over events that do not directly impact the telling of the story, and sometimes also employ flashbacks, flashforwards, and anachrony which may confuse an artificial learner. However, learning from narratives can make certain things easier. Stories can make explicit the normally unobservable mental operations and thought processes of characters. Written stories make dialogue more explicit in terms of whom is speaking, although some ambiguity remains [6] and comprehension of language is still be an open challenge.

Machine enculturation may give us a way forward toward achieving artificial intelligences that understand humans better, and make themselves more comprehensible—less alien—to humans. Further, agents and robotics that act in accordance with social values will naturally avoid situations where humans will be harmed or inconvenienced whenever possible. Harrison and Riedl [29] describe a technique, *learning from stories* (LfS), for emulating human behavior expressed in simple, crowdsourced narratives. It is proof of concept that machine enculturation may be feasible via machine learning over a corpus of stories.

Conclusions

Narrative intelligence is central to many of the things we as humans do, from communication to entertainment to learning. Narrative is also an effective means of storing and disseminating culture. In this position paper we argue that future artificial intelligences should be instilled with computational narrative intelligence so that they can act like humans, or understand human wants, needs, and desires. Artificial intelligences instilled with computational narrative intelligence may be more effective at communicating with humans and explaining their behavior. Finally, computational narrative intelligence may be a practical step towards machine enculturation.

References

- [1] Nicholas Allen, John Templon, Patrick Summerhays McNally, Larry Birnbaum, and Kristian Hammond. 2010. *StatsMonkey: A Data-Driven Sports Narrative Writer*. In *Computational Models of Narrative: Papers from the AAAI Fall Symposium (Technical Report FS-10-04)*, Mark Finlayson (Ed.). AAAI Press.
- [2] Timothy Bickmore, Daniel Schulman, and Langxuan Yin. 2009. Engagement vs. Deceit: Virtual Humans with Human Autobiographies. In *Proceedings of the 2009 International Conference on Intelligent Virtual Agents*.
- [3] J. Bruner. 1991. The Narrative Construction of Reality. *Critical Inquiry* 18, 1 (1991), 1–21.
- [4] Nathanael Chambers and Dan Jurafsky. 2008. Un-supervised learning of narrative event chains. In *Proceedings of 46th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*.
- [5] David K. Elson. 2012. *Modeling Narrative Discourse*. Ph.D. Dissertation. Columbia University.
- [6] David K. Elson and Kathleen McKeown. 2010. Automatic Attribution of Quoted Speech in Literary Narrative. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*.
- [7] Oren Etzioni, Anthony Fader, Janara Christensen, Stephen Soderland, and Mausam. 2011. Open Information Extraction: the Second Generation. In *Proceedings of the 2011 International Joint Conference on Artificial Intelligence*.
- [8] Brian Falkenhainer, Ken Forbus, and Dedre Gentner. 1989. The Structure-Mapping Engine: Algorithms and Examples. *Artificial Intelligence* 41 (1989), 1–63.
- [9] Gilles Fauconnier and Mark Turner. 1998. Conceptual Integration Networks. *Cognitive Science* (1998), 133–187.
- [10] Mark Finlayson. 2011. *Learning Narrative Structure from Annotated Folktales*. Ph.D. Dissertation. Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology.
- [11] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. 2015. A Neural Algorithm of Artistic Style. *CoRR* abs/1508.06576 (2015). <http://arxiv.org/abs/1508.06576>

- [12] Pablo Gervás. 2009. Computational Approaches to Storytelling and Creativity. *AI Magazine* 30, 3 (2009), 49–62.
- [13] Pablo Gervás, Belen Díaz-Agudo, Federico Peinado, and Raquel Hervás. 2005. Story Plot Generation based on CBR. *Journal of Knowledge-Based Systems* 18, 4–5 (2005), 235–242.
- [14] Jonathan Gordon and Lenhart Schubert. 2011. Discovering Commonsense Entailment Rules Implicit in Sentences. In *Proceedings of the TextInfer 2011 Workshop on Textual Entailment*.
- [15] Raquel Hervás, Francisco Pereira, Pablo Gervás, and Amílcar Cardoso. 2007. Cross-Domain Analogy in Automated Text Generation. In *Proceedings of the 3rd Workshop on Computational Creativity*.
- [16] Jerry Hobbs. 1992. Metaphor and Abduction. In *Communication from an Artificial Intelligence Perspective: Theoretical and Applied Issues*, A. Ortony, J. Slack, and O. Stock (Eds.). Springer-Verlag.
- [17] George Lakoff. 1987. *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. University of Chicago Press.
- [18] Michael Lebowitz. 1987. Planning stories. In *Proceedings of the 9th Annual Conference of the Cognitive Science Society*. 234–242.
- [19] Doug Lenat. 1995. Cyc: A Large-Scale Investment in Knowledge Infrastructure. *Commun. ACM* 38 (1995).
- [20] Boyang Li, Stephen Lee-Urban, George Johnston, and Mark O. Riedl. 2013. Story Generation with Crowdsourced Plot Graphs. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*. Bellevue, Washington.
- [21] Hugo Liu and Push Singh. 2004. ConceptNet—a practical commonsense reasoning tool-kit. *BT Technology Journal* 22, 4 (2004), 211–226.
- [22] Neil McIntyre and Mirella Lapata. 2010. Plot Induction and Evolutionary Search for Story Generation. In *The 48th Annual Meeting of the Association for Computational Linguistics*. 1562–1572.
- [23] James R. Meehan. 1977. TALE-SPIN: An Interactive Program that Writes Stories. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence*. 91–98.
- [24] T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohamed, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, and J. Welling. 2015. Never-Ending Learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*.
- [25] Erik Mueller. 2004. Understanding script-based stories using commonsense reasoning. *Cognitive Systems Research* 5, 4 (2004), 307–340.
- [26] Brian C. O’Neill and Mark O. Riedl. 2014. Dramatis: A Computational Model of Suspense. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*.
- [27] Karl Pichotta and Raymond Mooney. 2015. Learning Statistical Scripts with LSTM Recurrent Neural Networks. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*.
- [28] Julie Porteous, Marc Cavazza, and Fred Charles. 2010. Applying Planning to Interactive Storytelling: Narrative Control using State Constraints. *ACM Transactions on Intelligent Systems and Technology* (2010).
- [29] Mark Riedl and Brent Harrison. 2016. Using Stories to Teach Human Values to Artificial Agents. In *Proceedings of the 2nd International Workshop on AI, Ethics and Society*.

- [30] Mark O. Riedl. 2010. Case-Based Story Planning: Creativity Through Exploration, Retrieval, and Analogical Transformation. *Minds and Machines* 20, 4 (2010).
- [31] Mark O. Riedl and R. Michael Young. 2010. Narrative Planning: Balancing Plot and Character. *Journal of Artificial Intelligence Research* 39 (2010), 217–268.
- [32] R. Schank and R. Abelson. 1977. *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures*. Lawrence Erlbaum Associates.
- [33] Reid Swanson and Andrew Gordon. 2012. Say Anything: Using Textual Case-Based Reasoning to Enable Open-Domain Interactive Storytelling. *ACM Transactions on Interactive Intelligent Systems* 2, 3 (2012), 16:1–16:35.
- [34] Tony Veale. 2014. Coming Good and Breaking Bad: Generating Transformative Character Arcs For Use in Compelling Stories. In *Proceedings of the 2014 International Conference on Computational Creativity*.
- [35] Ramakrishna Vedantam, Xiao Lin, Tanmay Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. Learning Common Sense Through Visual Abstraction. In *Proceedings of the 2015 International Conference on Computer Vision*.
- [36] Jason Weston, Sumit Chopra, and Antoine Bordes. 2015. Memory Networks. In *Proceedings of the 2015 International Conference on Learning Representations*.
- [37] Patrick H. Winston. 2011. The Strong Story Hypothesis and the Directed Perception Hypothesis. In *Advances in Cognitive Systems: Papers from the 2011 AAAI Fall Symposium (Technical Report FSS-11-01)*, Pat Langley (Ed.). AAAI Press.
- [38] Jichen Zhu and Santiago Ontañón. 2014. Shall I Compare Thee to Another Story: An Empirical Study of Analogy-Based Story Generation. *IEEE Transactions on Computational Intelligence and AI in Games* (2014).
- [39] Alexander Zook, Mark O. Riedl, Heather K. Holden, Robert A. Sottilare, and Keith W. Brawner. 2012. Automated Scenario Generation: Toward Tailored and Optimized Military Training in Virtual Environments. In *Proceedings of the 7th International Conference on the Foundations of Digital Games*. Raleigh, North Carolina, 164–171.