

Computer as Audience: A Strategy for Artificial Intelligence Support of Human Creativity

Mark O. Riedl and Brian O’Neill
GVU Center
School of Interactive Computing
Georgia Institute of Technology
{riedl, boneill}@cc.gatech.edu

ABSTRACT

Tools exist for creating practically every type of artistic, creative, or communicative digital artifact, including pictures, music, video, and computer animation. While these tools make it technically feasible to produce creative content, they do not guarantee that the creator will produce work that is valued by a community as possessing quality or meaning. Can an intelligent system augment non-expert creative ability? We present a metaphor for designing intelligent systems that support human creativity: *computer-as-audience*. We motivate the computer-as-audience metaphor with discussion of support of human creative story authoring.

INTRODUCTION

Access to a nearly ubiquitous medium for information exchange – the Internet – and greater access to tools for media content production have led to a cultural phenomenon of user-generated content sharing. Tools exist for creating practically every type of artistic, creative, or communicative digital artifact, including pictures, music, video, and computer animation. While these tools make it technically feasible to produce creative content, they do not guarantee that the creator will produce work that is valued by a community as possessing quality or meaning. Therefore, the development of tools that support authors in creating purposeful content plays an important role in enabling and improving media content production. Tools that support the creation of content, in contrast to those that focus on providing the technical ability to create a media artifact, are especially valuable when the creation process is prohibitively costly, difficult, or time-consuming.

We posit that artificial intelligence systems that are aware of the human user’s creative intentions and that are knowledgeable about the artistic domain (storytelling, pictures, music, video, animation, etc.) can effectively work with non-expert human authors to increase the value of their creative artifacts. The question, then, is: Can an intelligent system augment non-expert creative ability? Lubart [7] enumerates four ways in which computer interfaces can support creativity:

- **Computer as nanny:** The computer provides organizational and classification services and performs routine operations on behalf of the user.

- **Computer as pen-pal:** The computer facilitates brainstorming with functionality that captures and transmits to collaborators the user’s thoughts.
- **Computer as coach:** The computer is knowledgeable about the process and can offer suggestions and stimulate creativity.
- **Computer as colleague:** The computer forms half of a human-computer team by contributing to the solution.

When considering *intelligent* systems that support creativity in humans, the focus is on the last two metaphors: computer-as-coach and computer-as-colleague. The computer-as-coach metaphor is used extensively in intelligent tutoring systems [18]. However, it is not always the case that a computational system that supports creativity is also trying to teach a user how to be creative or how to produce some artistic artifact through a particular technique. Our intuition is that intelligent creativity support systems should augment the user’s ability to create valuable artifacts. How can this be done? One way is to create an intelligent system that acts as a colleague, teaming with the user to create an artifact together. A computer-as-colleague system implements an expert system that can contribute to the creation of the artifact equally with the user, or complement the user’s ability by tackling rote tasks (see [2] for an overview of mixed-initiative systems). The computer-as-colleague metaphor makes a lot of sense for completing difficult tasks with well-defined goals.

Is the computer-as-colleague metaphor appropriate for intelligent creativity support systems? We posit that the answer to this question depends on the goals of the human creator with regard to the intended use of the artifact to be created. If the human creator is interested in creating an artifact as a means to an end – that is, the human creator is creating an artifact to solve a problem – then the human creator is most likely willing to accept the assistance from a computational system that is operating at the level of a peer. However, we focus on the case where the human wants to be the sole creator – to have made, or have appeared to have made, all the creative decisions.

We suggest a fifth category of ways in which computer interfaces can support creativity: **computer as audience**. A computer-as-audience system simulates the recipient of a creative artifact: a community, an individual viewer, reader, watcher, etc. More importantly, a computer-as-audience sys-

tem could conceivably incorporate this feedback into the creative process itself in stages or as a continual watch-over-the-shoulder approach.

INTELLIGENT STORY AUTHORIZING SUPPORT

Due to the complexity, subtlety, and nuance that can be expressed through narrative, story authoring – especially fictional story authoring – can be challenging to non-experts and non-professionals. However, the existence of storytelling rings, fan-fiction, and machinima suggests that there is a high degree of interest among non-experts in creating and sharing story content. In this section, we describe our current plans to research and develop intelligent systems to support creative story authoring. Our approach implements the computer-as-audience metaphor that we introduced previously.

Story authoring appears to be extremely amenable to computer-as-audience because story reading employs numerous cognitive processes [3]. Indeed, we know quite a bit about the types of inferencing behaviors that readers of narratives engage in (c.f., [4]). Because of the high degree of cognitive engagement of readers, a computer model of an audience – a reader model – can draw on understood psychological and artificial intelligence principles. Natural language understanding goes back to at least the early 1960’s and story understanding work goes back to at least the early 1970’s. However, there is a fundamental difference between story understanding and computer-as-audience for human story authoring. Story understanding systems algorithmically construct knowledge structures representing an input story and then demonstrate comprehension by answering questions about the story. Computer-as-audience implies that understanding informs a process of constructive critical feedback.

An intelligent story authoring support system must also be able to comprehend human-authored narrative content. However, instead of merely demonstrating comprehension, such a system must use this knowledge to provide constructive feedback during the creative process.

Toward Modeling Audience

Ultimately, an intelligent creativity support system built on the computer-as-audience metaphor should express the same reactions, sentiments, and emotions as a human audience. In our current work, we focus on objective metrics such as comprehensibility. Arguably, any story will fail to achieve desired effects on an audience – in this case readers – if it is not comprehensible. Currently, we base our work off of the QUEST model of question-answering for stories [5]. However, instead of modeling question-answering behavior, we use the models and processes to guide an artificial intelligence in the comprehension of human-authored stories.

QUEST [5] is a psychological model of question answering that simulates the question-answering performance of humans when responding to open-class questions about narrative content. Specifically, QUEST models encode the answers to why, how, when, enablement, and consequence-type questions. This type of model can be used to illus-

Once there was a Czar who had three lovely daughters. One day the three daughters went walking in the woods. They were enjoying themselves so much that they forgot the time and stayed too long. A dragon kidnapped the three daughters. As they were being dragged off, they cried for help. Three heroes heard the cries and set off to rescue the daughters. The heroes came and fought the dragon and rescued the maidens. Then the heroes returned the daughters to their palace. When the Czar heard of the rescue, he rewarded the heroes.

Figure 1. The Czar’s Daughters story.



Figure 2. A sample QUEST cognitive model.

trate how people build cognitive representations of stories, and the manner in which these cognitive representations capture certain relationships between narrative events and the perceived goals of characters. QUEST represents stories as directed graphs of plot elements, thereby capturing reader knowledge about story events, story states, and character goals. Directed links signify the causal relationships between story events and the intentionality relationships between events and character goals. To predict human question answering behavior, QUEST traces the graph using traversal routines for each type of question. Figure 1 shows a short story in which a czar’s three daughters are kidnapped by a dragon and eventually rescued by heroes, and Figure 2 shows a portion of the cognitive model that results from reading the story.

QUEST operates on complete representations of narratives, and predicts question-answering behavior of humans. To incorporate the cognitive model of story into an intelligent story authoring support system, we developed a system called ReQUEST [13]. ReQUEST serves as a model of a reader, keeping a graphical data structure of the story-so-far as it would be interpreted by a reader. ReQUEST uses the QUEST model of question-answering to identify gaps in reader understanding and to ask questions to fill those gaps. As the author creates the story, the ReQUEST algorithm attempts to resolve unclear aspects of the story. ReQUEST does this in much the same way as a human reader, querying its memory of story events for possible explanations. The algorithm itself does not understand what it reads, but tries to make sense of the story based on causal and character-goal structural relationships. Because the story is incomplete at most

times during authoring, the ReQUEST reader model may be incapable of making these connections, and therefore poses questions to the user in order to draw out these connections.

There are two types of questions in the ReQUEST algorithm:

- **“Why” Questions** seek the cause of their associated event or state. Why questions are answered in terms of character goals or events that motivate characters to form goals. Why questions typically take the form “Why did [character] want to perform [some action]?”
- **“Consequence” Questions** seek the resulting action from some event or state. That is, what does a particular event cause to happen? Consequence questions typically take the form “What was the consequence of [some event]?”

Preliminary studies of the ReQUEST algorithm [13, 10] suggest that this is a promising approach, but also indicate that reader modeling alone is not enough. In particular, ReQUEST has no notion of appropriate timing with regard to when it asks questions. This has implications for the human author’s ability to maintain creative flow [1].

Toward Incorporating Audience into the Creative Process

What is missing from our preliminary work on creativity support using the computer-as-audience approach is a model for how to integrate a simulated audience into the creative authoring process itself in a way that is constructive instead of destructive. This is a challenging task; it is not clear that there is any real-world analogue to study to determine appropriate system behavior.

We are turning to models to creativity for insight into how to solve this problem. A particularly relevant piece of research models writing as design. Sharples [15] proposed that the writing process occurs in cycles of *engagement* and *reflection*. The actual production of text, as the realization of new concepts, takes place within the engagement phase. The reflection phase includes review, organization, and revision of the content.

We hypothesize that intelligent support of story authoring must also flow in cycles of engagement and reflection. Furthermore, we hypothesize that the computational system’s cycles of engagement and reflection must be coordinated with the human author’s cycles of engagement and reflection. That is, the intelligent system must be able to detect what mode of creation the human author is in and respond accordingly.

Currently, ReQUEST seems more aligned with Sharples’ notion of reflection. A preliminary study [10] provides strong evidence to support this belief. The next step in our research is to determine how to recognize when the human author is in a reflective mode. This may be as simple as noticing whether the human is authoring forward – e.g., is moving the story forward through concept creation – or is authoring backward – e.g., has returned to existing concepts or is filling in details between previously generated concepts. Future research will address this issue.

Should the Audience Model also be an Expert?

In the previous section, we reviewed the engagement-reflection model [15] of human creative writing. It is our belief that a computer-as-audience support system can contribute to reflection cycle of human creative story authoring. But what about contributing to the engagement cycle of story authoring? This is an important question because it is not at all clear that a computer-as-audience should be responsible for providing input into the creation and instantiation of new concepts. Doing so may diminish the human author’s perception that he or she is the creator.

An expert system that can contribute to the engagement cycle of story authoring is a *story generator*. Most intelligent story generation systems are designed to supplant the human author instead of work with the human author. Notably, the MEXICA system [11] uses Sharples’ engagement-reflection model to generate narrative content. MINSTREL [16] models story generation as creative retrieval and transformation of concepts in case-based reasoning. Other systems model story generation as problem-solving, particularly planning (c.f., [8, 6, 14, 12]).

A story generator can be considered an expert system when integrated into a creativity support system for story authoring. However, note that, as with all mixed-initiative problem-solving systems, the expert system must be augmented with models of cooperation with human users [2]. That is, a mixed-initiative system must be able to determine *when* and *how* to engage with the user. Current work on story generation systems assume no human-in-the-loop and thus attempt to produce complete narratives. However, we are not currently addressing the question of how to integrate story generation capabilities into our intelligent story authoring tool. The exception to this is in cases where we require the intelligent tool to perform inference about the story-so-far, e.g., to hypothesize about directions that the story might go.

Should the Audience Model also be a Coach?

There are numerous formulaic processes for writing stories, especially screenplays. For example, the *Save The Cat*TM method, which professes that a screenplay is told in 15 “beats” – themes – and goes so far as to suggest which pages the beats should fall on. Further, software applications such as *Dramatica*TM, *Power Structure*TM, and *Truby’s Blockbuster*TM, enforce organizational methods and plot structures based on other common screenwriting techniques. It is possible to teach these processes as scaffolds for the writing process. *Coached problem solving* [17] is an approach to intelligent tutoring (c.f., [18]) in which the system collaborates with the learner to solve a problem with a known solution. Coached problem solving could be used to teach process, as long as a process – such as those used in the software tools above – exists. Some argue that it is also possible to teach creativity [9].

We are not planning to incorporate any coaching behavior into our authoring system because we believe that teaching a specific process may interfere with the ability of the human to be creative. However, we have discovered in our prelim-

inary study of ReQUEST [10] that users learn to model and predict the circumstance under which ReQUEST initiates questions and change their authoring style to preempt the audience model. That is, users implicitly learn to model the behavior of the ReQUEST algorithm. Because ReQUEST operates as a model of an audience, we can conclude that authors also implicitly learned to model an audience.

CONCLUSIONS

We have taken a first step toward a novel type of computational creativity support system, based on a new metaphor for human-computer interaction: the computer-as-audience. The computer-as-audience approach is especially useful in the domain of story authoring because of the relatively well-understood nature of cognitive processes used in reading and comprehension. However, open questions remain. What cognitive processes of reading and comprehension are most valuable for generating feedback? How can we capitalize on models of creativity, engagement, and reflection to avoid interrupting creative flow? Does computer-as-audience subsume computer-as-colleague and computer-as-coach, or are they orthogonal?

We believe that intelligent authoring support tools – especially for story authoring – will become more important in the relatively near future. The Web 2.0 philosophy continues to shape users’ online experiences. Particularly, users are transitioning from traditional consumers of content created by experts to producers and sharers of creative artifacts. We anticipate increased story creation activity, which will create a demand for intelligent tools that support non-expert creative activities.

REFERENCES

1. M. Csikszentmihalyi. *Creativity: Flow and the Psychology of Discovery and Invention*. HarperCollins, New York, 1996.
2. G. Ferguson and J. Allen. Mixed-initiative systems for collaborative problem solving. *AI Magazine*, 28(2):23–32, 2007.
3. R. Gerrig. *Experiencing Narrative Worlds: On the Psychological Activities of Reading*. Yale University Press, New Haven, 1993.
4. A. Graesser, M. Singer, and T. Trabasso. Constructing inferences during narrative text comprehension. *Psychological Review*, 101(3), 1994.
5. A. C. Graesser, K. L. Lang, and R. M. Roberts. Question answering in the context of stories. *Journal of Experimental Psychology: General*, 120(3):254–277, 1991.
6. M. Lebowitz. Story-telling as planning and learning. *Poetics*, 14, 1985.
7. T. Lubart. How can computers be partners in the creative process: Classification and commentary on the special issue. *International Journal of Human-Computer Studies*, 63:365–369, 2005.
8. J. Meehan. *The Metanovel: Writing Stories by Computer*. PhD thesis, Yale University, 1976.
9. R. Nickerson. Enhancing creativity. In R. Sternberg, editor, *Handbook of Creativity*. Cambridge University Press, 1999.
10. B. O’Neill and M. O. Riedl. Engagement vs. reflection: A study of user interfaces for an intelligent story authoring assistant. Unpublished manuscript, available at: <http://www.cc.gatech.edu/riedl/pubs/request-ui-study.pdf>, 2008.
11. R. Pérez y Pérez and M. Sharples. Mexica: A computer model of a cognitive account of creative writing. *Journal of Experimental and Theoretical Artificial Intelligence*, 13, 2001.
12. M. O. Riedl. Vignette-Based Story Planning: Creativity Through Exploration and Retrieval. In *Proceedings of the 5th International Joint Workshop on Computational Creativity*, 2008.
13. M. O. Riedl, J. P. Rowe, and D. K. Elson. Toward Intelligent Support of Authoring Machinima Media Content: Story and Visualization. In *Proceedings of the 2nd International Conference on Intelligent Technologies for Interactive Entertainment (INTETAIN)*, 2008.
14. M. O. Riedl and R. M. Young. An Intent-Driven Planner for Multi-Agent Story Generation. In *Proceedings of the 3rd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2004.
15. M. Sharples. *How We Write: Writing as Creative Design*. Routledge, 1999.
16. S. Turner. *The Creative Process: A Computer Model of Storytelling*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1994.
17. K. VanLehn. Conceptual and meta learning during coached problem solving. In *Proceedings of the 3rd International Conference on Intelligent Tutoring Systems*, 1996.
18. E. Wenger. *Artificial Intelligence and Tutoring Systems: Computational and Cognitive Approaches to the Communication of Knowledge*. Morgan Kaufmann, San Francisco, 1987.