
A Call for Flow Modeling in Interactive Storytelling

Vadim Bulitko

BULITKO@UALBERTA.CA

Department of Computing Science, University of Alberta, Edmonton, Alberta, T6G 2E8, Canada

David Thue

DAVIDTHUE@RU.IS

School of Computer Science, Reykjavik University, Menntavegur 1, 101 Reykjavik, Iceland

Abstract

The field of interactive storytelling aims to create a narrative experience that is tailored to the player. A variety of Artificial Intelligence (AI) methods have been used to dynamically manage the narrative to suit the player's preferences. Modern approaches tend to represent the domain of narrative discourse in a machine-readable form and then run automated planners to create a narrative that is consistent with the player's choices as well as the author's goals. The resulting planning task is frequently under-constrained and allows for many solutions. Since not every plan makes for an engaging story, the challenge lies with selecting one that will appeal to the particular player. In this paper, we conjecture that an engaging story is one that keeps the player in the psychological state of flow. Thus, an experience manager should select the narrative that is predicted to maximize the player's state of flow. We propose to use a recent computational model of flow based on matching cognitive abilities of the audience with the cognitive demands of the narrative. The model will then be combined with a recent AI interactive narrative manager. This position paper is meant to solicit comments from researchers in the field to help shape the project.

1. Introduction

When a small group of men and women sat at the fire in ancient times, the stories that they told would likely have been interactive, with members of the audience interrupting the speaker and influencing the structure of the narrative. While this interaction continued with small-scale theatre productions as well as the bedside stories we tell our children, the mass media has switched to non-interactive narrative forms such as books and motion pictures.

It is believed that feeling agency in daily life is beneficial to one's well-being (Larson, 1989) and that some players enjoy games primarily because games gives them such a feeling. Interactivity in narrative can give the audience a sense of agency and is likely to improve the quality of entertainment. Video games have been bringing interactivity back into mass market storytelling – modern productions such as BioWare's *Mass Effect* or *Dragon Age* series (BioWare Corp., 2012; BioWare Corp., 2014) feature an impressive cast of actors, a branching storyline and a number of side quests that allow the player to interact with and affect the narrative world around them.

Following the success of massive open-world games such as *Fallout 3* (Bethesda Softworks, 2008) and *Skyrim* (Bethesda Softworks, 2011), modern video games are expected to give the player a degree of narrative agency and allow them to make “their own story”. However, not every possible

sequence of events makes for an engaging narrative (VanOrd, 2014) bringing game designers back to the age-old question of what makes a good interactive story. Complicating the problem further is the fact that narrative appeal is not universal, with different members of the audience preferring different types of narrative and gameplay.

2. Related Work

In the last several decades, the problem of creating individualized narrative has been tackled with Artificial Intelligence methods (Riedl & Bulitko, 2013). A common approach is to encode the domain of narrative discourse in a formal, computer-readable format and then use automated planning methods to derive possible stories (Young et al., 2004). Once such story plans are computed, the problem is reduced to selecting the best one. Early systems such as ASD (Riedl et al., 2008) preferred the stories closest to a manually pre-authored exemplar story, regardless of player preferences. Later research explicitly modelled the player by observing his/her actions throughout the game. For instance, a system called PAST (Ramirez & Bulitko, 2014) used a player model based on Robin Laws’ player types (Laws, 2001). It engaged an automated planner whenever the player deviated from the current story plan. New narratives consistent with the player’s previous choices as well as the author’s goals would be automatically generated, and the narrative most matching the player’s type would then be presented to the player. The most recent system in this line of work, PACE, uses the player type model to infer the player’s desires over a certain set of goals (Hernandez, Bulitko, & St. Hilaire, 2014). The desires are then used with an appraisal model of emotions (Bulitko et al., 2008; Marsella & Gratch, 2009) to estimate the player’s emotions for different possible narratives. The narrative which is estimated to keep the player on a pre-authored emotion arc is then selected. This approach can be viewed as a narrative extension of the AI zombie modulator within the commercial video game *Left 4 Dead* (Valve Corporation, 2008; Booth, 2009).

These approaches have progressively distanced authors from writing the static structure of a traditional book. Instead of writing the entire narrative, interactive story designers can create a world of characters, equip them with possible actions, specify a few authorial goals (e.g., the grandmother gets eaten in “Little Red Riding Hood” (Perrault, 1697)) and let the player-controlled character loose in the world. The difficulty with these approaches lies with the assumptions that underlie their operation. For instance, ASD assumes that stories closer to the original exemplar story are most fitting for any player. How do we know if the exemplar story is fitting for a wide range of players? PAST assumes that matching the Laws-style player type at all times makes for a good narrative. Is it really so? Are these player types informative enough to tailor the narrative to the player? To which extent are they applicable across various narrative genres? PACE requires the designer to pre-author a static trajectory through the emotional space that all players will be kept on. Is there a single emotional trajectory that fits all players? If so, how can it be identified?

3. The Proposed Approach

The approaches discussed in the previous section attempted to answer the fundamental question “What makes a good interactive story?” by making a number of assumptions. While the resulting implementations have frequently been positively evaluated in practice, we feel unsatisfied by the

answers and unsure of how widely applicable these assumptions are. Thus, in the rest of the paper we describe an alternative approach based on a single psychological concept: flow. We start by giving the intuition of our proposal and then follow with algorithmic details.

3.1 An Intuitive Overview

The psychological state of flow has been linked to optimal performance in humans (Csikszentmihalyi, 1990). People in the state of flow appear not only to perform better but also to feel engaged, motivated and happy. To achieve that state, several conditions are thought to be important, including a balance of the person’s skills and the problem’s complexity, well defined goals and rules, and timely and clear feedback. In this paper, we will focus on the first condition: a good match between the person’s cognitive skills (e.g., short term memory, vocabulary, social awareness, empathy) and the cognitive complexity/challenge of following a particular narrative (i.e., the cognitive skills required of the audience).

Our primary conjecture and the answer to the question “What makes a good interactive story?” is that *good interactive narratives are the ones that maximize the player’s degree of flow¹ while the story is underway*. Given that interactive stories are often presented in a video-game-like setting, there is a connection between our conjecture and the use of flow in video game design. In fact, the concept of flow originated from psychological studies of game playing (Csikszentmihalyi, 1975) and connections between flow and games have been discussed extensively (Csikszentmihalyi, 1990; Green & Brock, 2000; Sweetser & Wyeth, 2005; Chen, 2007; Cowley et al., 2008; Baron, 2012; Koster, 2013). That being said, the innovation of our approach is twofold.

First, we propose to keep the player in the state of flow by shaping the *narrative* using an estimate of the player’s flow as an objective function. This stands in contrast to the common case of dynamically adjusting *gameplay* difficulty (e.g., by modulating zombie influx in *Left 4 Dead* (Booth, 2009)) which does not substantially alter the story being told. Consequently, while both commercial video games (Ritual Entertainment, 2006; Pagulayan et al., 2012) and academic research in dynamic difficulty adjustment (Hunicke & Chapman, 2004; Zook & Riedl, 2014; Chen, 2007) have focused on gameplay skills, we focus on the player’s cognitive skills that are specifically related to comprehending narrative (e.g., remembering minute details of a crime scene, or suspending one’s disbelief in a forest with magic fairies). This focus is supported by work that found that reading can commonly induce flow (Csikszentmihalyi, 1990), where the skills involved include narrative comprehension and visualization, empathizing with its characters, and anticipating twists in its plot (Sweetser & Wyeth, 2005; Nell, 1988). This is also supported by research on flow in games that focused on the cognitive processing involved in playing a game (Cowley et al., 2008).

Second, we propose that an AI-based experience manager should perform flow-maximizing adjustments to the narrative automatically *on-line*, as the narrative is being experienced by the player. Specifically, whenever an AI-based experience manager decides among several possible narrative segments to run next, it should estimate the degree of flow that each segment will induce in the player and then select the segment with the maximum estimated flow. We propose to employ an

1. In this paper the degree of flow refers to the frequency and/or duration and/or the depth of flow states experienced by the audience of the narrative.

explicit computational model of flow to estimate the degree of flow of a specific player given a candidate narrative segment.

This *on-line closed loop approach* is in contrast to the common practice of manually tuning a game’s difficulty curve during the development process so that ideally an average player’s gameplay skills would approximately match the game’s complexity/challenge throughout the game (known as pacing (Schreiber, 2009)). For instance, many first-person shooters and role-playing games gradually ramp up the difficulty of the enemies either by introducing more difficult enemy types as the player progresses through the story (e.g., *Fallout: New Vegas* (Bethesda Softworks, 2010)) or by increasing the difficulty of the existing enemy types (e.g., *The Elder Scrolls IV: Oblivion* (Bethesda Softworks, 2006)). Alas, creating a difficulty ramp to match every player’s skill ramp is generally impossible because different people have substantially different skill ramps (Koster, 2013). Similarly, narrative difficulty ramps are common in traditional novels where the author attempts to tune the pacing of the story to avoid overwhelming the reader or making them bored. Just like with video games, different people may have different narrative skill ramps, and this limits the appeal of a static, pre-authored narrative.

3.2 Algorithmic Details

We propose to extend the narrative management framework of PACE (Hernandez, Bulitko, & St. Hilaire, 2014) with a computational model of flow that is based on the balance between the player’s skills and the problem’s complexity (Bulitko, 2014).

As with PACE, our proposed AI experience manager takes the narrative space expressed as the set S of narrative states, the set A of narrative actions that the player may perform and the world dynamics p which links the narrative state and the narrative actions. It also takes a set S^\dagger of terminal narrative states and a complexity function \bar{c} . The complexity function maps any narrative state to a vector of m numbers: $\bar{c} : S \rightarrow [0, 1]^m$, where each number indicates the degree of cognitive skill that is required from the audience to engage with that narrative state. For instance, a story with many related characters may have a narrative state with the complexity of $(0.8, 0.1)$, where the values indicate that a high skill (0.8) in mapping people’s names and relations is required from the player, but that their ability to solve logical puzzles would not be taxed (0.1). The same m dimensions are also used to represent the player/audience’s cognitive skills $\bar{\sigma}$ as modelled by the AI experience manager. The model is initialized to some prior in line 2 of Algorithm 1. We discuss ways to define the complexity function in Section 3.3.

As long as the player has not reached a terminal state (line 3) the AI manager presents the current narrative state s_t to the player (e.g., the player controlling Red encounters a wolf in the forest) and collects the player’s action a_t (e.g., the player chooses to shoot the wolf). The player’s cognitive skill model is then updated (e.g., friend/foe identification skill is raised) in line 6. We discuss mappings from the player’s action to their skills in Section 3.3. In line 7, the AI manager computes candidate narrative continuations in the same way as ASD, PAST and PACE: by invoking an automated planner with the current world dynamics given by p . Each of the narrative candidates n_j produced by the planner is consistent with the narrative formed so far and satisfies the authorial goals. In our running example, there may be two narrative alternatives computed by the planner: n_1 brings in a brother of the murdered wolf while n_2 employs a magic fairy to resurrect the wolf. Both

Algorithm 1: Flow-maximizing Narrative Management

inputs: narrative space (S, A, p) , narrative start state s_1 , narrative final states $S^\dagger \subset S$,
 complexity function \bar{c}

```

1  $t \leftarrow 1$ 
2 initialize player's skill model  $\bar{\sigma}_1$ 
3 while  $s_t \notin S^\dagger$  do
4     present narrative state  $s_t$  to the player
5     collect the player's narrative action  $a_t$ 
6     update the player's skills  $\bar{\sigma}_{t+1}$  from  $a_t$ 
7     compute narrative candidates  $\{n_j\}$  from  $s_t, a_t, p$ 
8     for each  $n_j$  do
9         estimate the resulting flow  $f_j$  from  $\bar{\sigma}_{t+1}, \bar{c}(n_j)$ 
10    select the highest flow:  $j^* \leftarrow \arg \max_j f_j$ 
11    select the next desired narrative state:  $s_{t+1} \leftarrow n_{j^*}|_1$ 
12    update the world dynamics  $p$  so that  $s_t \xrightarrow{a_t} s_{t+1}$ 
13     $t \leftarrow t + 1$ 

```

of them satisfy the authorial goal of Red's grandmother being eaten and Red subsequently deceived. For each of the computed narrative candidates, line 9 estimates the degree of the player's flow if they were to experience that continuation. We describe a way compute this estimate in Section 3.3. Once the flow is estimated for each narrative candidate, the index j^* of the flow-maximizing candidate is determined in line 10, the next narrative state is set to the first state of the narrative n_{j^*} in line 11, and the dynamics of the world p are updated so that the player's action a_t indeed leads to that state in line 12 (Thue & Bulitko, 2012). For example, to select between narratives n_1 and n_2 , the AI manager will first compute the complexity of each. Suppose that the cognitive complexity of the wolf's brother narrative n_1 is $\bar{c}(n_1) = (0.7, 0.7, 0.1)$, where the three dimensions are friend/foe identification skill, fighting ability, and the ability to suspend disbelief. Meanwhile, suppose that the resurrecting fairy narrative n_2 has a complexity of $\bar{c}(n_2) = (0.1, 0.1, 0.7)$ (since the player might have to suspend their disbelief in the existence of fairies). Next, the AI manager will examine the model of the player's skills that it has constructed thus far (say, $\bar{\sigma}_{t+1} = (0.8, 0.9, 0.5)$), and then use it to estimate the player's flow for each candidate narrative. The flow induced by the narrative n_1 will be $f_1 \approx 1/(0.4583 + \xi)$, whereas $f_2 \approx 1/(1.0817 + \xi)$. Thus narrative n_1 is estimated to give the player a higher degree of flow and so will be selected by the AI manager.

3.3 Defining Flow, Complexity, and Skill

Selecting narrative to maximize the player's estimated degree of flow critically depends on the definition of flow and, more specifically, on the definitions of the skill and complexity functions related to narrative comprehension.

Several models of flow have been suggested (Weber et al., 2009; Bulitko & Brown, 2012; Moneta, 2012; Klasen et al., 2012; Bulitko, 2014). As a first step, we propose to use a simple flow model based solely on the balance of the player’s skills $\bar{\sigma}_{t+1}$ and the complexity of the narrative candidate $\bar{c}(n_j)$. The model was previously evaluated in a synthetic domain (Bulitko & Brown, 2012; Bulitko, 2014) and, in our context, becomes:

$$f_j = \frac{1}{\|\bar{\sigma}_{t+1} - \bar{c}(n_j)\| + \xi}$$

where $\|\cdot\|$ is the 2-norm distance: $\|\bar{x} - \bar{y}\| = \sqrt{\sum_{i=1}^m (x_i - y_i)^2}$ and ξ is a small positive constant to keep f_j finite when the player’s skills exactly match the narrative complexity (i.e., $\bar{\sigma}_{t+1} = \bar{c}(n_j)$). Note that n_j is a sequence of narrative states computed by the automated planner. In the formula above we assume that $\bar{c}(n_j)$ returns the complexity of the first narrative state of n_j and ignores the remainder of the sequence. More generally, the degree of flow can be computed along a multi-state narrative trajectory with a possible discounting of the flow estimated for more distant future states.

A basic approach to modeling the player’s narrative skills is to manually annotate each action available to the player with a vector of deltas to the player’s skill vector, similarly to the approach taken in our previous work on modeling player preferences (Thue et al., 2007; Thue et al., 2011; Ramirez & Bulitko, 2014; Hernandez, Bulitko, & St. Hilaire, 2014). To validate such annotations, one could run a user study in which the narrative experience is occasionally interrupted and the player’s narrative comprehension skills are measured with questionnaires or tests.

There are several ways to define the cognitive complexity of a narrative segment. A basic approach is to manually annotate all narrative events with a complexity vector. This is similar to manually annotating narrative encounters with player type suitability in PaSSAGE (Thue et al., 2007; Thue et al., 2011) and PAST (Ramirez & Bulitko, 2014). A more advanced approach would be to present possible narrative events to a variety of players whose narrative-comprehension skills had been measured ahead of time. Then, for each such player, one could measure his/her comprehension of the specific event that was presented to them. The cognitive complexity of the narrative event could then be data-mined from the collected measurements. For instance, adopting the unimodal assumption of Bulitko (2014), we can form a corpus of narrative-comprehension skills for all test players who sufficiently comprehended a narrative event and then take per-dimension minimum. To illustrate, suppose we had three test players whose narrative skills were pre-measured as $\bar{\sigma}_1 = (0.1, 0.2, 0.3)$, $\bar{\sigma}_2 = (0.4, 0.5, 0.6)$, $\bar{\sigma}_3 = (0.9, 0.2, 0.7)$ where the three dimensions are friend/foe identification skill, fighting ability, and the ability to suspend disbelief. Suppose the first player did not demonstrate a sufficient comprehension of a narrative event n whereas the other two players did. Then the complexity of n is the per-dimension minimum of $\bar{\sigma}_2$ and $\bar{\sigma}_3$: $\bar{c}(n) = (0.4, 0.2, 0.6)$.

4. Testbeds for Empirical Evaluation

Once we have implemented our approach, we will first evaluate it in the context of an interactive, AI-managed narrative such as an interactive version of the “Little Red Riding Hood” story (Thue et al., 2007; Riedl et al., 2008; Ramirez & Bulitko, 2014). The presentation can be via a text-only format (Ramirez & Bulitko, 2014), a full 3D game world (Thue et al., 2011) or a series of

still images (Hernandez, Bulitko, & St. Hilaire, 2014) (see Figure 1). The cognitive model of the player's skills will be updated from the player's input in the game (e.g., dialogue choices or other actions). In authoring the narrative space and the cognitive skill/complexity annotations on player actions and narrative segments, we will use the same process that we followed when creating our previous testbeds for PaSSAGE (Thue et al., 2007; Thue et al., 2011), PAST (Ramirez & Bulitko, 2014) and PACE (Hernandez, Bulitko, & St. Hilaire, 2014).

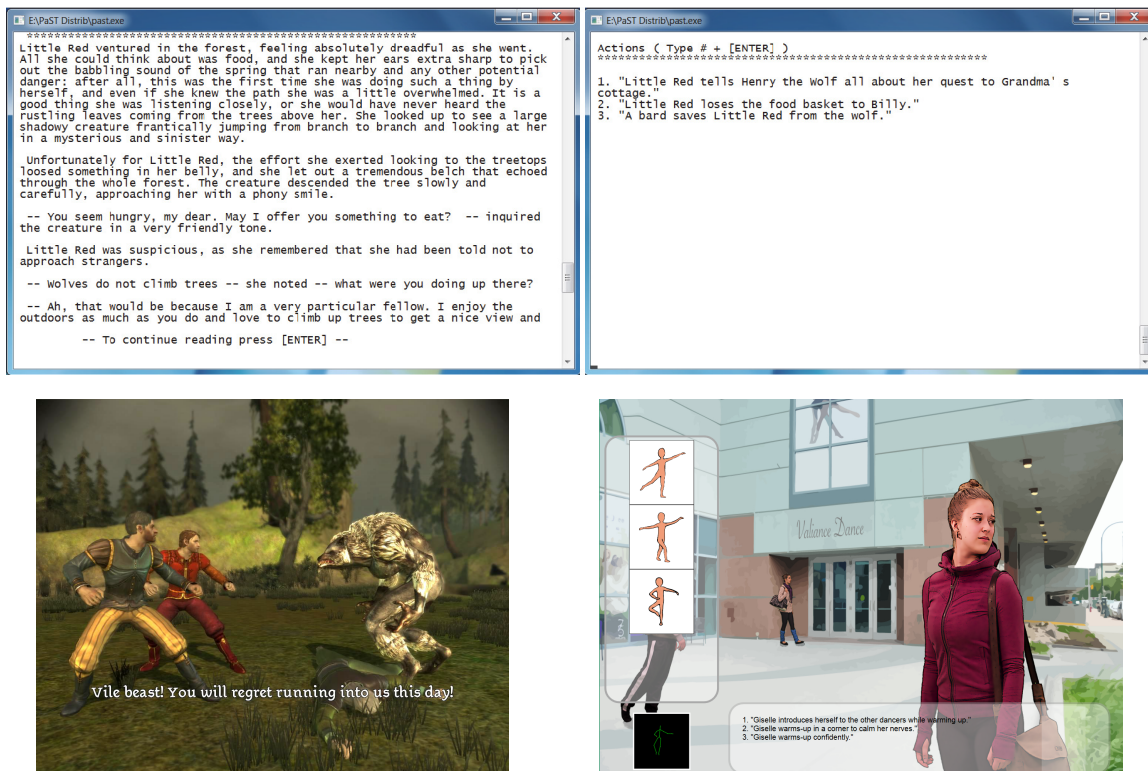


Figure 1. **Top row:** a text-based presentation of narrative in PAST and the player's choices (reproduced from (Ramirez & Bulitko, 2014)). **Bottom row:** a presentation of narrative in a 3D video game (left) or as still images (right) (reproduced from (Riedl & Bulitko, 2013; Hernandez, Bulitko, & St. Hilaire, 2014)).

The players in the experimental condition will experience an AI-managed story with our proposed flow estimate as the objective function. Their post-experience responses (e.g., enjoyment of the story) will be compared to those in the control condition (e.g., with random narrative candidate selection). This is a common approach for evaluating experience managers that we have used over the last eight years (Thue et al., 2007; Ramirez & Bulitko, 2014). We will attempt to complement questionnaire-based data about the overall experience with specific measurements of the degree of flow that is experienced by the participants directly using either questionnaires (Moneta, 2012) or fMRI (Klasen et al., 2012).

We will also consider evaluating this approach in intelligent training systems and on-line educational courses. For the former, we have partnered with a medical hospital and have been developing

a virtual-reality-based training system for neonatal resuscitation. Once the testbed is completed, we will evaluate whether keeping the trainee in a state of flow by dynamically modifying the training scenario can lead to a higher training effect. For the latter, we are partnering with researchers in on-line education to implement dynamic shaping of material in a massively open on-line course (MOOC) to maximize the student's degree of flow. Again, we will attempt to run user studies to evaluate the training effect of this approach.

5. Future Work

We have proposed a way to use a computational model of flow within an AI experience manager to select between automatically planned narratives. The natural next step is to actually implement this approach. To do so, several aspects of the approach need to be instantiated. First, the m dimensions describing the player's skills and the narrative complexity must be defined. We expect studies of reader engagement (Busselle & Bilandzic, 2009) to be informative for this step. Second, the player's actions must be mapped to updates in the player's skill model (line 6 in the algorithm). Third, the complexity function \bar{c} must be defined for all narrative states. We plan to work with reading psychologists and draw from research on transportation, absorption, immersion and engagement (Green & Brock, 2000; Green, 2004). Finally, more complex models of flow (Moneta, 2012) can be studied in place of the simplistic model that we presented above.

6. Conclusions

We proposed to apply the concept of flow in the context of AI-managed interactive storytelling. We conjectured that automatically shaping the player's experience toward maximizing his or her sense of flow can lead to a better narrative experience. We further proposed a specific computational model of the player's flow and a mechanism to shape the narrative towards maximizing the predicted flow. Potential applications include video games, intelligent training systems, and online education.

As this is a position paper, we welcome any feedback on the hypothesis as well as our proposed solution approach. We hope that such feedback will shape our implementation of the approach.

Acknowledgements

This research has been supported by the National Science and Engineering Research Council. Feedback from Mark Riedl, Marisa Bortolussi, Michael Young and Candace Jane Dorsey is appreciated.

References

- Baron, S. (2012). Cognitive flow: The psychology of great game design. *Gamasutra*.
- Bethesda Softworks (2006). The Elder Scrolls IV: Oblivion. http://www.elderscrolls.com/games/oblivion_overview.htm.
- Bethesda Softworks (2008). Fallout 3. <http://fallout.bethsoft.com/>.
- Bethesda Softworks (2010). Fallout: New Vegas. <http://fallout.bethsoft.com/eng/games/fnv-overview.php>.

- Bethesda Softworks (2011). The Elder Scrolls V: Skyrim. <http://www.elderscrolls.com/skyrim>.
- BioWare Corp. (2012). Mass Effect Trilogy. <http://masseffect.bioware.com/about/trilogy/>.
- BioWare Corp. (2014). Dragon Age Series. <http://www.dragonage.com/>.
- Booth, M. (2009). The AI systems of Left4Dead.
- Bulitko, V. (2014). Flow for meta control. *The seventh conference on Artificial General Intelligence*. Quebec City, QC, Canada.
- Bulitko, V., & Brown, M. (2012). Flow maximization as a guide to optimizing performance: A computational model. *Advances in Cognitive Systems*, 2, 239–256.
- Bulitko, V., Solomon, S., Gratch, J., & van Lent, M. (2008). Modeling culturally and emotionally affected behavior. *The Fourth Artificial Intelligence for Interactive Digital Entertainment Conference* (pp. 10–15). The AAAI Press.
- Busselle, R., & Bilandzic, H. (2009). Measuring narrative engagement. *Media Psychology*, 12, 321–347.
- Chen, J. (2007). Flow in games (and everything else). *Communications of the ACM*, 50, 31–34.
- Cowley, B., Charles, D., Black, M., & Hickey, R. (2008). Toward an understanding of flow in video games. *ACM Computers in Entertainment*, 6, 20:1–20:27.
- Csikszentmihalyi, M. (1975). Play and intrinsic rewards. *Journal of Humanistic Psychology*, 15, 41–63.
- Csikszentmihalyi, M. (1990). *Flow: The psychology of optimal experience*. New York, New York: Harper and Row. The first edition.
- Green, M. C. (2004). Transportation into narrative worlds: The role of prior knowledge and perceived realism. *Discourse Processes*, 38, 247–266.
- Green, M. C., & Brock, T. C. (2000). The role of transportation in the persuasiveness of public narratives. *Journal of Personality and Social Psychology*, 79, 700–721.
- Hernandez, S. P., Bulitko, V., & St. Hilaire, E. (2014). Emotion-based interactive storytelling with artificial intelligence. *Proceedings of the 10th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*.
- Hunicke, R., & Chapman, V. (2004). AI for dynamic difficult adjustment in games. *Proceedings of the Challenges in Game AI Workshop, Nineteenth National Conference on Artificial Intelligence*.
- Klasen, M., Weber, R., Kircher, T. T. J., Mathiak, K. A., & Mathiak, K. (2012). Neural contributions to flow experience during video game playing. *Social Cognitive and Affective Neuroscience*, 7, 485–495.
- Koster, R. (2013). *Theory of fun for game design*. O'Reilly Media, Inc.
- Larson, R. (1989). Is feeling “in control” related to happiness in daily life? *Psychological Reports*, 64, 75–84.
- Laws, R. (2001). Robin’s laws of good GMing. Steve Jackson Games.
- Marsella, S. C., & Gratch, J. (2009). EMA: A process model of appraisal dynamics. *Journal of Cognitive Systems Research*, 10, 70–90.

- Moneta, G. (2012). On the measurement and conceptualization of flow. In S. Engeser (Ed.), *Advances in flow research*, 23–50. Springer New York.
- Nell, V. (1988). The psychology of reading for pleasure: Needs and gratifications. *Reading Research Quarterly*, 23, 6–50.
- Pagulayan, R., Keeker, K., Fuller, T., Wixon, D., Romero, R., & Gunn, D. (2012). User-centered Design in Games. In J. Jacko (Ed.), *The human computer interaction handbook: Fundamentals, evolving technologies, and emerging applications*, chapter 34. Taylor & Francis. Third edition edition.
- Perrault, C. (1697). Le petit chaperon rouge. In *Histoires ou contes du temps passé, avec des moralités: Contes de ma mère l’oye*. Paris.
- Ramirez, A., & Bulitko, V. (2014). Automated planning and player modelling for interactive storytelling. *IEEE Transactions on Computational Intelligence and AI in Games*.
- Riedl, M., & Bulitko, V. (2013). Interactive narrative: An intelligent systems approach. *Artificial Intelligence magazine*, 34, 67–77.
- Riedl, M. O., Stern, A., Dini, D., & Alderman, J. (2008). Dynamic experience management in virtual worlds for entertainment, education, and training. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning* (pp. 23–42). Glasgow: SWIN Press.
- Ritual Entertainment (2006). SiN Episodes: Emergence. <http://www.sinepisodes.com/>.
- Schreiber, I. (2009). Game design concepts: Level 7: Decision-making and flow theory.
- Sweetser, P., & Wyeth, P. (2005). GameFlow: A Model for Evaluating Player Enjoyment in Games. *ACM Computers in Entertainment*, 3.
- Thue, D., & Bulitko, V. (2012). Procedural game adaptation: Framing experience management as changing an MDP. *Proceedings of the 5th Workshop in Intelligent Narrative Technologies*.
- Thue, D., Bulitko, V., Spetch, M., & Romanuik, T. (2011). A computational model of perceived agency in video games. *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE)* (pp. 91–96). Palo Alto, California, USA: AAAI Press.
- Thue, D., Bulitko, V., Spetch, M., & Wasylishen, E. (2007). Interactive storytelling: A player modelling approach. *Proceedings of the Third Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE)* (pp. 43–48). Palo Alto, California: AAAI Press.
- Valve Corporation (2008). Left 4 Dead. <http://www.l4d.com/>.
- VanOrd, K. (2014). Far Cry 4 review. *GameSpot*.
- Weber, R., Tamborini, R., Westcott-Baker, A., & Kantor, B. (2009). Theorizing flow and media enjoyment as cognitive synchronization of attentional and reward networks. *Communication Theory*, 19, 397–422.
- Young, R. M., Riedl, M. O., Branly, M., & Jhala, A. (2004). An architecture for integrating plan-based behavior generation with interactive game environments. *Journal of Game Development*, 1, 54–70.
- Zook, A. E., & Riedl, M. O. (2014). Temporal game challenge tailoring. *Computational Intelligence and AI in Games, IEEE Transactions on*, 1–11 (in press).