

Getting to the Point: Toward Resolving Ambiguity in Intelligent Narrative Technologies

David Thue
School of Computer Science
Reykjavik University
Menntavegur 1, Reykjavik 101, Iceland
davidthue@ru.is

Elin Carstensdottir
College of Computer and Information Science
Northeastern University
360 Huntington Ave., Boston, MA 02115
elin@ccs.neu.edu

Abstract

Despite more than three decades of work, researchers of intelligent narrative technologies have yet to settle on a lexicon that consistently describes common concepts. In some cases, certain terms are used in passing as though a common definition is understood, even though one or more different definitions can be used to understand the same term. In this paper, we present an exploratory study of how the terms “plot point” and “plot graph” have been used in the literature. We show that there exist at least two competing definitions for these terms, and we demonstrate how two recently-developed, representational frameworks can be used to clarify and distinguish between such competing definitions. On the basis of the demonstrated differences, we offer new terms to help distinguish the different uses of “plot point” and “plot graph”.

1 Introduction

Despite more than three decades of work, researchers of intelligent narrative technologies (INT) have yet to settle on a lexicon that consistently describes common concepts [KHFS10, KHFS11]. In some cases, certain terms are used in passing *as though* a common definition is understood, even though one or more different definitions could be used to understand each term, and where each definition could be considered “common” to some subset of the field. As a result, important differences between existing works remain implicit: without knowing precisely what construct(s) or concept(s) a given term is meant to convey across two related works, it is difficult to form an accurate impression of either [SBP11, SRPM12]. This hinders collaboration in our field, and creates unnecessary obstacles for new researchers who wish to join it.

One example of such a term is *plot point*; across at least twenty-five years of INT-related publications that use this term, there are at least two competing definitions for what it means and how its related concept might arise as part of a narrative experience [KWB93, RS06a]. In this paper, we demonstrate how two recently-developed representational frameworks [CSEN18, TB18] can be used to clarify and distinguish between such competing definitions, offering a concrete step toward increased unity in the field. In doing so, we hope to inspire others to pursue similar clarifications for other ambiguous terms.

The terms that we have chosen to address in this paper are *plot point* and *plot graph*. This selection was motivated by the fact that the existing INT-related literature includes two kinds of works: those that use one or both of these terms without defining them (as though a common definition is widely known), and those that clearly define these terms, but in ways that differ from one definition to the next (we have found at least

Copyright © by the paper’s authors. Copying permitted for private and academic purposes.

In: A. Jhala, M. Si, H-Y. Wu (eds.): Proceedings of the Joint Workshops on Intelligent Narrative Technologies (INT) and Intelligent Cinematography and Editing (WICED) at AIIDE’18, Edmonton, Canada, 13-NOV-2018, to appear at <http://ceur-ws.org>

two). Examples of the former kind include work by Si et al. [SMP09], Lee et al. [LML11], Swartjes [Swa10], Barber [Bar08], Rowe and Lester [RL13], Kline [Kli09], and possibly more that we have not yet found. Examples of the latter kind can be divided into at least two categories, where papers within each category broadly share the same definition, but the definitions differ from one category to the next. We will present these two categories and their definitions in Section 3.

The remainder of this paper is organized as follows. We begin by presenting the methodology that we used to gather a set of definitions for *plot point* and *plot graph* from a set of prior work. We then present the two sets of definitions that we have found. To clarify the differences between these definitions, we apply two recent, analytical frameworks to their underlying representations. We conclude the paper with some discussion and our plans for future work.

2 Methodology

Our goal was to explore the usage of *plot point* and *plot graph* in the literature and determine what definitions are being used for these terms. Subsequently, we aimed to examine and compare these definitions to see whether they differed from each other.

We chose to review a set of papers published in the field of intelligent narrative technologies. We selected an initial, convenient sample from papers published in three interactive narrative research venues: INT, ICIDS, and AIIDE. We first reviewed the papers in the initial sample for mentions or use of “plot point” and “plot graph”. For each mention that we found, we collected information regarding its definition and/or usage. If the definition or usage cited prior work in the literature, we extended our initial sample by adding the paper that was cited; we processed each added paper in the same way as those that were initially in our sample. We excluded papers from further analysis if they included either “plot point” and “plot graph” *without* defining either term, giving context or explanation of what they referred to, or referencing other work to explain their meaning.

By the second round of analysis, our sample only included papers that offered explicit definitions for *plot point* or *plot graph*, either by defining one or both terms directly, or by citing and adopting an earlier definition. For each definition, we analyzed its use as a concept or structure in the context of the work, and, if applicable, how it was implemented within a system or a framework. Once we had collected this data, we analyzed the gathered definitions and approaches, toward identifying the specifics of how plot points and plot graphs are represented, initiated, executed, and terminated.

3 Competing Definitions

By applying the methodology that we described in Section 2, we have identified two sets of definitions for *plot point* and *plot graph*. They include works by Kelso et al. [KWB93], Weyhrauch [Wey97], Nelson & Mateas [NM05, NRIM06, NMRI06], and Sharma et al. [SMmR07, SOMR10] in one set, and works by Riedl & Stern [RS06a, RS06b], Riedl et al. [RSDA08], and Li et al. [LLUAR12, LLUJR13] in another set. We discuss each set in turn.

3.1 Player-focused Plot Points

The earliest definition of *plot graph* in the context of INT research seems due to Kelso et al.’s work on the Oz Project [KWB93]. They defined a plot graph as a directed, acyclic graph where the nodes represent “*events and situations that are the major moments of the story*” and the edges represent ordering constraints over the nodes; a node’s event cannot begin until all nodes that precede it in the graph have been finished. Following this work, Weyhrauch later used a plot graph to represent ordering constraints over “User Moves” in an interactive narrative [Wey97]. Each User Move consisted of one or more granular interactive sequences (e.g., a specific conversation with a Non-Player Character (NPC)), all of which could be recognized automatically by Weyhrauch’s experience manager, Moe. To the best of our knowledge, the nodes of a Kelso et al. plot graph were first described as *plot points* by Mateas [Mat97], in reference to the User Moves in Weyhrauch’s work: “*The Oz drama manager [Weyhrauch 97] controls a story at the level of plot points. Plot points are “important moments” in a story.*” (Weyhrauch and Mateas were also members of the Oz Project).

To summarize, Kelso et al., Weyhrauch, and Mateas viewed a plot point as an abstract player action (a “User Move”), which the player could enact by performing any one of a set of predetermined, granular interactive sequences. At any point in time, the plot graph defined the set of abstract actions that could (potentially) be enacted by the player, based on the precedence rules encoded in the graph. Later research that adopted these definitions includes work by Nelson et al. [NM05, NMRI06, NRIM06] as well as Sharma et al. [SMmR07,

SOMR10]. Given that the plot points in this definition are meant to represent (abstract) player actions, we will refer to them as *player-focused plot points* henceforth.

3.2 Character-focused Plot Points

A second definition for *plot point* can be found in Riedl & Stern’s work on the Automated Story Director (ASD) [RS06a, RS06b], which uses an Artificial Intelligence (AI) planner to automatically generate a tree of contingency plans. Each (partially-ordered) plan represents a potential, alternative version of an author-created story, and ASD adopts a logical alternative plan whenever a player action invalidates part of the story’s current plan. Riedl & Stern refer to each plan step as a *plot point*: “*Plan steps, which in this application represent plot points [...]*” [RS06b]. Each of these plot points represents one of two things: a directive to one or more *non-player characters* (NPCs) in the story’s world [RS06a, RS06b], or an objective (pertaining to the story world’s state) that a *player-controlled* character will ideally achieve through the granular actions they perform [RSDA08]. Each NPC directive is expressed as an objective that the receiving NPC should adopt and pursue through its own granular actions in the story world. Regardless of whether or not any plot point is happening, the player can freely observe the story world and perform granular actions therein.

In later work, Li et al. [LLUAR12, LLUJR13] presented a method for automatically learning a partially-ordered graph of granular characters actions from simple, crowd-sourced stories. They referred to this graph as a plot graph, and its ordering represented a precedence relationship between each pair of connected actions.

While none of Riedl, Stern, et al. nor Li et al. used both *plot point* and *plot graph* to refer to their work in a single paper, we describe them together here because they share the property that the nodes of their graphs (either ASD plan steps or granular character actions) were meant to describe what *arbitrary* characters should do next in the story; they might be NPCs, or they might be controlled by the player. This intention is notably different from how plot points were used in the Oz Project and the work that followed, where they represented only *player* actions; non-player actions were represented in another way (e.g., via “Moe Moves” [Wey97]).

In summary, the work following Riedl & Stern’s use of *plot points* viewed them as a way to specify idealized character goals in a story, either as an objective that a character should adopt [RS06a], or as a direct action that a character should perform [LLUJR13]. The plans generated by Riedl & Stern’s Automated Story Director are partially-ordered graphs of the elements that they call plot points, which leads us to view them as plot graphs, even though the authors never named them as such. Li et al.’s plot graphs describe the order in which the actions described by the graphs nodes can be performed. Given that the plot points in this definition are meant to represent arbitrary character actions (either by NPCs or players), we will refer to them as *character-focused plot points* for the rest of this paper.

3.3 Summary of Definitions

Here we summarize the two sets of competing definitions that we have found.

- A *player-focused plot point* is a pattern of player behaviour that can be recognized during a specific part of a story.
- A *player-focused plot graph* is a partial ordering over player-focused plot points that determines the part of the story during which each plot point’s pattern can be recognized.
- A *character-focused plot point* is a goal that should be realized by a (player or non-player) character.
- A *character-focused plot graph* is a partial ordering over character-focused plot points that determines when each plot point’s NPC goal will be assigned or its player goal can be recognized.

4 Analysis via the Joint Perspective

To provide unified view of the differences between the player-focused and character-focused representations that we discussed in Section 3, we will analyze both representations using Thue and Bulitko’s *joint perspective* [TB18]. The joint perspective views an interactive experience (e.g., a narrative game) and an associated experience manager [RSDA08] *jointly* – that is, as a single, more complicated game. In doing so, it ensures that all aspects of how the manager and game relate to one another can be directly discussed in a language that is free from the manager-specific terms that would otherwise be necessary. This freedom makes it simpler to compare concepts across different representations for experience managers and the games that they target. This makes the joint

perspective useful for our current task, since an experience manager is an integral part of all but two of the papers in our final sample (the cited work by Li et al. does not explicitly involve any manager [LLUAR12, LLUJR13]).

To compare competing sets of definitions for plot points and plot graphs using the joint perspective, one must first translate the definitions and their related concepts into the joint perspective’s representation. In this representation, the player is an agent in a Factored-state Markov Decision Process (FMDP) [CS11]. Each state of the FMDP represents: (i) the current state of the game that will be managed, (ii) a set of *tuneable parameters* that the manager can edit, and (iii) any internal state or features that the manager requires to operate. Each FMDP action represents an action that the player can perform. The transition function of the FMDP represents both the mechanics of the target game and the policy of the manager. The mechanics determine each new state of the game based on its previous state and the tuneable parameters. The manager’s policy determines new values for the tuneable parameters, based on its observations of each current state and the actions that the player performs. Since the transition function is the only element of an FMDP that is capable of representing computation, it must also be the case that any computations pertaining to NPC behaviour must also be represented therein [TB18]. Each NPC action manifests as one or more changes to the FMDP’s state over time.

4.1 Translating Player-focused and Character-focused Plot Points

Recall that a player-focused plot point is a set of granular, interactive sequences, each of which can be recognized as enacting an abstract player action that is associated with the plot point (Section 3). From the joint perspective, the set of interactive sequences corresponds to a set of alternating state/action sequences through the FMDP. Since the FMDP’s transition function is the only part of the joint perspective representation that can perform computation [TB18], the process of recognizing a plot point’s interactive sequences must be representable by some part of this transition function. Indeed, Weyhrauch defines each such “recognizer” as an executable script [Wey97]. This script begins executing when all prior plot points in the plot graph have been executed (initiating the plot point), and succeeds when a particular pattern of player interaction is observed (terminating the plot point). Since each recognizer considers prior sequences of game states and player actions, there must also be some extra-game state factors in the FMDP that represent a bounded history of the player’s experience.

Similarly to player-focused plot points, a character-focused plot point occurs generally as a set of alternating state/action sequences through an FMDP, since the player remains free to perform granular actions whenever any plot point is happening (in Riedl et al.’s work). Although Li et al.’s work does not specifically address player interaction, their plot points can still be represented from the joint perspective using the same construct; any single player action is still a (short) state/action sequence in the FMDP, and any single NPC action would occur across one or more states in such a sequence. Each character-focused plot point could be initiated when its prior plot points in the plot graph have terminated, but the initiation process varies depending on whether the plot point specifies an NPC directive or a desired player goal.

If a plot point is an NPC directive, then it is initiated by sending the directive to the target NPCs. Since this process involves some computation, it must be represented as part of the FMDP’s transition function. Furthermore, since each NPC behaviour must include some computation, it must be represented in the joint perspective as some part of the FMDP’s transition function. To connect the assignment of a directive to the behaviour of an NPC, one can create a set of state factors for each NPC such that (i) the assignment process can modify all of the factors, and (ii) those factors can be tested as part of determining the NPC’s current behaviour.

If the plot point is a desired player goal, then it is initiated by starting a “daemon” [RS06a, RS06b] to detect the change in world state that the goal represents. Similarly to a player-focused “recognizer”, such a daemon can be represented as part of the FMDP’s transition function, but with the simplification that no history of the player’s experience is needed, since a daemon only checks the current state of the game.

4.2 Translating Player-focused and Character-focused Plot Graphs

A player-focused plot graph is a directed, acyclic graph of player-focused plot points that restrict which points can occur in terms of which other points have already occurred; a plot point has “occurred” as soon as one of its sequences has been recognized. From the joint perspective, the occurrence of each plot point can be represented by adding a Boolean state factor to the FMDP’s state; the moment that a plot point has been recognized, its corresponding factor becomes true; otherwise it remains false. To represent and enforce the plot graph’s ordering constraints, the FMDP’s transition function could be made to (i) test the factors that represent plot point occurrences whenever it computes a new state for the target game and (ii) respect the plot graph’s rules by only computing “occurred” for the state factors of plot points whose ancestors in the plot graph have occurred

and whose interaction pattern has been recognized. This would guarantee that each plot point could only be recognized (and thus marked as having occurred) in a way that respects the ordering of the given plot graph.

A character-focused plot graph describes a partial order over character-focused plot points, in terms of both (i) the execution of an NPC directive or (ii) the recognition of a desired player goal. From the joint perspective, it can be represented somewhat similarly to a player-focused plot graph. A Boolean state factor can be used to represent the satisfaction of either an NPC directive or a desired player goal. The set of such factors (across all of a plot graph’s points) can then be used by the FMDP’s transition function to either (i) set an NPC’s directive (as an in-game state factor) or (ii) “unlock” the part of the transition function that recognizes the satisfaction of a player-desired goal (similarly to how the occurrence of player-focused plot points can be tested).

4.3 Comparison

Our application of Thue and Bulitko’s joint perspective has revealed a further difference between player-focused and character-focused plot points, beyond their respective foci on strictly player actions or general character actions. Specifically, with Riedl et al.’s definition of character-focused plot points, the process of recognizing player actions is simpler than the “recognizers” used by player-focused plot points; Riedl et al.’s (character-focused) daemons only monitor the current world state, while Weyhrauch’s (player-focused) recognizers could consider a history of prior states and actions. The complication involved in the latter approach can be seen in the work that followed by Nelson et al. [NM05, NMRI06, NRIM06], which opted to implement recognizers that only tested the current state of the game (rather than a history of states and actions); Thue and Bulitko present and consider this point in further detail [TB18].

5 Analysis via the Interaction Model and Interaction Maps

To provide a view of the differences between player-focused and character-focused plot points and plot graphs, we will examine them through the lens of user experience and interaction design. Specifically, we will consider how they differ in terms of what kinds of experiences they afford the player, as a result of both their definitions and the constraints they place on system design. To do so, we will use the *Interaction Model for Interactive Narratives* [CKSEN17] and *Interaction Maps* [CSEN18].

The Interaction Model for Interactive Narratives is a framework for describing and studying the design-related aspects of how players interact with interactive narrative games. The model has four factors: Structure, Narrative Mechanics, Interaction, and User Experience, each of which has its own sub-constructs and unique considerations. The model is abstract and meant for qualitative analysis and reasoning about interaction design. To augment the descriptive capabilities of the Interaction Model, Carstensdottir et al. [CSEN18] developed Interaction Maps, a descriptive, computational representation of player interaction, to describe interaction that occurs from one moment to the next.

An interaction map is a graph-based representation of the interaction that a player can have within an interactive narrative game. Interaction maps have been successfully implemented and used for automated structural analysis of interactive narratives [PCS⁺18].

An individual interaction map corresponds to one interactive narrative artifact, and is meant to “represent and capture the full set of possibilities of player interactions within an interactive narrative at any given time.” [CSEN18] Each map contains multiple *interaction units*, each unit representing a single interaction opportunity that is available to the player. An interaction map can be considered to be a graph of interaction units, and each interaction unit is a sub-graph that contains the specific elements of the interaction opportunity that the unit represents. Each interaction opportunity corresponds to a single cycle of perception, action, and feedback (i.e., the core interaction cycle).

Each interaction unit can contain four types of nodes: *interaction point*, *option*, *feedback*, and *event*. *Interaction points* represent the interaction opportunity itself and its context. This includes information given to the player directly, such as dialog or interface information, and its presentation. *Options* correspond to the actions or choices that are available to the player as a part of an interaction opportunity. Thus, the number of options corresponds directly to the number of interaction possibilities, and includes all available actions – even those not explicitly presented to the player. *Feedback* represents information directly caused by or related to an action initiated by the player. As such, feedback is linked to interaction points or options in the interaction map, depending on when it is conveyed. *Events* also provide information about the story, but represent information that is not given as a result of (or in response to) player interaction. Events, in other words, represent information or content that is delivered regardless of what the player does.

5.1 Translating Player-focused and Character-focused Plot Points

According to the definition of a player-focused plot point, a single User Move can consist of one or more granular interactive sequences (i.e., sequences of player actions). An interaction unit represents an individual player action. As such, a User Move can be defined as sequences of interaction units that can form subgraphs within the larger interaction map. While the interaction map, as defined, is not strictly a directed acyclic graph, it can be structured as such. Thus, a player-focused plot point can be represented as a subgraph of one or more interaction units. Similarly, a character-focused plot point is defined as a character goal, which we can also represent as a sequence of interaction units (we discuss this further in Section 5.2). Such a sequence can also be represented as a subgraph of one or more interaction units.

5.2 Translating Player-focused and Character-focused Plot Graphs

Using the Interaction Map representation, it is relatively straightforward to represent a player-focused plot graph as an interaction map, as both representations are centered on the actions of the player within the story. Specifically, a player-focused plot graph can be represented as an interaction map that contains one subgraph for each plot point. In this sense, the plot graph is a set of connected sub-graphs that can be combined to represent all possible plans available within the system.

For a character-focused plot graph, the player is able to take action at any time (see Section 3.2). Thus, there must be a continuous chain of interaction units, where each interaction point contains information about the world state, and where the set of options within each unit depends on the current state. We assume that “Waiting” is always included as an option, since the player will not be continuously taking action despite having the opportunity to do so.

Differently from the player-focused plot graph, a character-focused plot graph allows NPCs to be given goals to realize. The realization of NPC goals can be modeled as events in an interaction map, where the system can have an NPC take action within an interaction point while the player is waiting. To understand how this can be represented in the interaction map, consider the following example. If we assume that the player is engaging with a continuous chain of interaction units as previously described, and if the player has been Waiting within one of those units, an experience manager might want to progress the narrative by having an NPC accomplish a goal. The system could insert an event into that chain of interaction units, creating one possible variation of a interaction chain (i.e., a sub-graph of the larger plot graph).

5.3 Comparison

When the two plot graph definitions are compared using both the Interaction Map and the Interaction model, the differences between them become more concrete – particularly in the types of user experiences they afford.

The most noticeable difference is the flexibility that the character-focused plot graph affords the player in terms of the number of possible variations of experience. This is largely due to the player continuously being able to take action within the story. When considering what kind of experience this affords, it might seem like an ideal approach; the player is free to perform any action and an experience manager is in place to ensure that the experience remains cohesive and engaging. However, allowing for such freedom might reduce the player’s capacity to reason about how their actions advance the narrative, as the manner in which the narrative is progressed might not be transparent or consistent. Player confusion might occur, for example, when the same actions or sequence of actions have different consequences in similar contexts within the world. The player would have insufficient information to form a functioning mental model of how the underlying progression works, and would not be able to reason about how their behaviour would affect the narrative. In addition, the system might assign goals to either the player or an NPC and then ensure that this goal be achieved by sending an NPC to do so, without the player being aware of either the goal or its importance.

Alternatively, player-focused plot graphs can be more constrained in terms of actions and rely more on precedence constraints (whether they are continuous or not will depend on the implementation, but all implementations that we are aware of were not continuous). This affords more authorial control than the character-focused plot graph, as the former relies on the player to engage with the plot points by performing actions, while the latter allows for more flexibility by deploying NPCs as needed to further the current plot.

Table 1: A comparison of two ways of defining plot points and plot graphs, based on the papers cited.

	Player-focused	Character-focused
Plot Point	Represents something that the <i>player</i> does that’s narratively important.	Represents something that (NPC or player) <i>characters</i> do that’s narratively important.
Plot Graph	A directed graph specifying precedence constraints over plot point recognition.	A directed graph that orders directed NPC behaviour and desired player behaviour in time.
Initiation	When all prior plot points in the graph have terminated, recognition of the given plot point becomes possible.	When all prior plot points in the graph have terminated, the given plot point either directs one or more NPCs or becomes recognizable.
Execution	The player can perform actions and see new states (exploring a subgraph of the game content).	The player can perform actions and see new states (exploring a subgraph of the game content), and/or NPCs pursue given directives.
Termination	Detected automatically from the granular sequence of prior world states and player actions. It can fail to terminate if the required pattern does not occur.	Detected automatically as the satisfaction of a goal given to NPCs (or expected of the player) at initiation time. It can abort if the player invalidates the effects of prior plot points.
Papers	[KWB93, Wey97, NM05, NRM06, NMRI06, SMmR07, SOMR10]	[RS06a, RS06b, RSDA08, LLUAR12, LLUJR13]

6 Discussion and Future Work

In this paper, we explored whether there was a difference in how the terms *plot point* and *plot graph* are used in the literature, and we did find such a difference. Specifically, while the original definitions by the Oz Project and the work that followed adopted a more player-focused view, the later definitions by Riedl et al. adopted a more character-focused view. Furthermore, our analysis via Thue and Bulitko’s “joint perspective” revealed that, as defined, the occurrence of player-focused plot points could potentially be more difficult to recognize than the satisfaction of player goals within character-focused plot points. Our analysis via the Interaction Model and Interaction Map representations showed a difference between the two different views in terms of their flexibility, which could have implications on authoring and the user experience, depending on implementation.

Table 1 shows a summary of our findings regarding player-focused and character-focused plot points and plot graphs, along with the papers that we identified that use one of the two sets of definitions.

Given the exploratory nature of this work, there are some limitations to consider. The sample of papers chosen for analysis in this paper was based on a conveniently available sample, and thus is not necessarily representative of the field as a whole. We did not include the previous work that might be relevant to this exploration that a more comprehensive and systematic search would have uncovered.

Although this work is exploratory (and thus not yet comprehensive), it nevertheless illustrates that multiple, substantively different definitions exist for two terms that are used commonly *without* explicit definitions in the field of Intelligent Narrative Technologies. In addition to making it difficult for new researchers to approach our field, having divergent definitions for common terms put up barriers to many productive conversations that *could* be taking place. This highlights the importance of developing a shared vocabulary in the field, or at least the importance of defining key terms explicitly. We intend to expand the scope of our analyses going forward, and we encourage others in the field to use our approach to resolve other ambiguous terms.

Acknowledgements

We would like to thank the reviewers of this work for their thoughtful and insightful comments.

References

- [Bar08] Heather Barber. *Generator of Adaptive Dilemma-based Interactive Narratives*. PhD thesis, Department of Computer Science, The University of York, 2008.
- [CKSEN17] Elin Carstensdottir, Erica Kleinman, and Magy Seif El-Nasr. Towards an interaction model for interactive narratives. In Nuno Nunes, Ian Oakley, and Valentina Nisi, editors, *Interactive Storytelling*, pages 274–277, Madeira, Portugal, 2017. Springer International Publishing.

- [CS11] Doran Chakraborty and Peter Stone. Structure learning in ergodic factored MDPs without knowledge of the transition function’s in-degree. In Lise Getoor and Tobias Scheffer, editors, *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 737–744, New York, NY, June 2011. ACM.
- [CSEN18] Elin Carstensdottir and Magy Seif El-Nasr. Interaction Maps for Interactive Narratives. Technical Report NU-CCIS-TR-2018-001, College of Computer and Information Science, Northeastern University, Boston, MA, August 2018.
- [KHFS10] Hartmut Koenitz, Mads Haahr, Gabriele Ferri, and Tonguc Ibrahim Sezen, editors. *The ICIDS 2010 Workshop Towards a Shared Vocabulary for Interactive Digital Storytelling*. Springer Berlin / Heidelberg, 2010.
- [KHFS11] Hartmut Koenitz, Mads Haahr, Gabriele Ferri, and Tonguc Ibrahim Sezen, editors. *The ICIDS 2011 Workshop Towards a Unified Theory for Interactive Digital Storytelling: Classifying Artifacts*. Springer Berlin / Heidelberg, 2011.
- [Kli09] Daniel Kline. Bringing interactive storytelling to industry: Designing a reactive narrative encounter system. In *Proceedings of the 5th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE’09)*, pages 113–118, Palo Alto, CA, 2009. AAAI Press.
- [KWB93] Margaret Thomas Kelso, Peter Weyhrauch, and Joseph Bates. Dramatic presence. *Presence: The Journal of Teleoperators and Virtual Environments*, 2(1):1–16, 1993.
- [LLUAR12] Boyang Li, Stephen Lee-Urban, Darren Scott Appling, and Mark O Riedl. Automatically learning to tell stories about social situations from the crowd. In *the LREC 2012 Workshop on Computational Models of Narrative*, 2012.
- [LLUJR13] Boyang Li, Stephen Lee-Urban, George Johnston, and Mark O. Riedl. Story generation with crowd-sourced plot graphs. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, pages 598–604, Bellevue, WA, USA, 2013. AAAI Press.
- [LML11] Seung Y Lee, Bradford W Mott, and James C Lester. Director agent intervention strategies for interactive narrative environments. In *International Conference on Interactive Digital Storytelling*, pages 140–151. Springer, 2011.
- [Mat97] Michael Mateas. An oz-centric review of interactive drama and believable agents. Technical Report CMU-CS-97-156, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA., 1997.
- [NM05] Mark J. Nelson and Michael Mateas. Search-based drama management in the interactive fiction anchorhead. In *Proceedings of the First Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE)*, pages 99–104. AAAI Press, Marina del Rey, California, 2005.
- [NMRI06] Mark J. Nelson, Michael Mateas, David L. Roberts, and Charles L. Isbell. Declarative optimization-based drama management in interactive fiction. *IEEE Computer Graphics and Applications*, 26(3):33–41, 2006.
- [NRIM06] Mark J. Nelson, David L. Roberts, Charles L. Isbell, Jr., and Michael Mateas. Reinforcement learning for declarative optimization-based drama management. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems, AAMAS ’06*, pages 775–782, New York, NY, USA, 2006. ACM.
- [PCS⁺18] Nathan Partlan, Elin Carstensdottir, Sam Snodgrass, Erica Kleinman, Gillian Smith, Casper Hartevelde, and Magy Seif El-Nasr. Exploratory Automated Analysis of Structural Features of Interactive Narrative. In *Proceedings of the 14th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, Edmonton, AB, Canada, November 2018. AAAI Press. To appear.
- [RL13] Jonathan Rowe and James Lester. A modular reinforcement learning framework for interactive narrative planning. In *Proceedings of the 6th Workshop in Intelligent Narrative Technologies*, pages 57–63, Boston, MA, USA, 2013. AAAI Press.

- [RS06a] M Riedl and A Stern. Failing believably: Toward drama management with autonomous actors in interactive narratives. In *3rd International Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE 2006)*, volume 4326/2006, pages 195–206, Darmstad, Germany, December 4-6 2006. Springer.
- [RS06b] Mark O. Riedl and Andrew Stern. Believable agents and intelligent story adaptation for interactive storytelling. In *3rd International Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE 2006)*, pages 1–12, Darmstad, Germany, December 4-6 2006. Springer.
- [RSDA08] Mark O. Riedl, Andrew Stern, Don Dini, and Jason Alderman. Dynamic experience management in virtual worlds for entertainment, education, and training. In H. Tianfield, editor, *International Transactions on Systems Science and Applications*, volume 4, pages 23–42, Glasgow, 2008. SWIN Press.
- [SBP11] Nicolas Szilas, Thomas Boggini, and Paolo Petta, editors. *The ICIDS 2011 Workshop on Sharing Interactive Digital Storytelling Technologies*, 2011.
- [SMmR07] Manu Sharma, Manish Mehta, Santiago Ontañón, and Ashwin Ram. Player modeling evaluation for interactive fiction. Technical report, AIIDE 2007 Workshop on Optimizing Player Satisfaction, Palo Alto, California: AAAI Press, June 2007.
- [SMP09] Mei Si, Stacy C. Marsella, and David V. Pynadath. Directorial control in a decision-theoretic framework for interactive narrative. In Ido A. Iurgel, Nelson Zagalo, and Paolo Petta, editors, *Interactive Storytelling*, pages 221–233, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- [SOMR10] Manu Sharma, Santiago Ontañón, Manish Mehta, and Ashwin Ram. Drama management and player modeling for interactive fiction games. *Computational Intelligence*, 26(2):183–211, 2010.
- [SRPM12] Nicolas Szilas, Stefan Rank, Paolo Petta, and Wolfgang Mueller, editors. *The ICIDS 2012 Workshop on Sharing Interactive Digital Storytelling Technologies*, 2012.
- [Swa10] Ivo Martinus Theodorus Swartjes. *Whose story is it anyway? How Improv Informs Agency and Authorship of Emergent Narrative*. PhD thesis, Faculty of Electrical Engineering, Mathematics & Computer Science, University of Twente, Enschede, The Netherlands, 2010.
- [TB18] David Thue and Vadim Bulitko. Toward a Unified Understanding of Experience Management. In *Proceedings of the 14th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE’18)*, Edmonton, AB, Canada, 2018. AAAI Press. To appear.
- [Wey97] P. Weyhrauch. *Guiding Interactive Drama*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 1997.