

Open Design Challenges for Interactive Emergent Narrative

James Owen Ryan, Michael Mateas, and Noah Wardrip-Fruin

Expressive Intelligence Studio
Center for Games and Playable Media
University of California, Santa Cruz
{jor, michaelm, nwf}@soe.ucsc.edu

Abstract. We introduce a research framework for the design of interactive experiences in the domain of *emergent narrative*, an application area of computational narrative in which stories emerge bottom-up from the behavior of autonomous characters in a simulated storyworld. Prior work in this area has largely concerned the development and tuning of the simulations themselves from which interesting stories may reliably emerge, but this approach will not necessarily improve system performance at its most crucial level—the actual interactive experience. Looking to completed experiences, namely simulation games like *Dwarf Fortress* and *The Sims*, we identify a series of shortcomings that yield four design challenges at the level of interaction: *modular content*, *compositional representational strategies*, *story recognition*, and *story support*. In this paper, we motivate and discuss each of these design challenges and, for each, summarize prior work and propose new approaches that future work might take.

Keywords: emergent narrative · content authoring · story recognition

1 Introduction

Stories in *emergent narrative* arise bottom-up, usually from the interactions of autonomous characters in a simulated storyworld [77]. This approach to computational narrative is as old as the field itself. *Novel Writer* and *Tale-Spin*, the earliest known story generators, were emergent-narrative systems. In the former, a murder story forms from the simulated actions of a group of partygoers [40]; in *Tale-Spin*, characters in a storyworld autonomously carry out plans that may allow them to achieve their goals [56]. In recent decades, there have been a number of subsequent emergent-narrative systems, including other AI-research systems like *Virtual Storyteller* [78], *FearNot!* [13], and *Prom Week* [55], as well as commercial games like *The Sims* [54] and *Dwarf Fortress* [3]. Other character-based systems, such as the *Petz* series of digital games [28], the interactive-story system of Cavazza et al. [21], and Chris Crawford’s *Storytron* [24], might also be called emergent-narrative systems.

Many of these systems, however, are not fundamentally *interactive*. *Novel Writer* and *Tale-Spin* are only interactive in that, in each, a user may specify the initial state of the storyworld. *Virtual Storyteller* is not interactive at all (it employs an embodied agent that tells emergent stories), though later extensions to it incorporated interactivity [73, 33].

The *emergent narrative* research program, cultivated by Ruth Aylett, Sandy Louchart, and others, has typically concerned only interactive systems, most often *FearNot!*-like research systems. In this paper, we likewise limit our concerns to interactive systems, but we use the term ‘emergent narrative’ more broadly to include systems like the digital games listed above; this usage is not unprecedented [49, 82, 18, 22]. We note that nondigital storytelling systems, such as tabletop role-playing games, should also be called emergent-narrative systems [45], but in this paper we do not consider these (as our concerns herein pertain only to digital systems). So, to be clear, by ‘emergent narrative’, we mean the application area characterized by digital, fundamentally interactive systems whose narratives emerge bottom-up, typically from the richness of underlying simulations that feature autonomous characters.

A basic appeal of emergent narrative is of course the very *emergence* that it yields, producing stories that even a system’s designer might not have anticipated. Indeed, games with rich underlying simulations have often been lauded for their emergent qualities [38, 4]—this is what got *Dwarf Fortress*, which we discuss in Section 4, into the Museum of Modern Art [87]. The fundamental advantage to this bottom-up approach, however, is that it defeats the sticky issue in top-down interactive narrative of accommodating player actions [64, 9]. If the affordances given to a player are a subset of (or are coextensive with) the actions that non-player characters (NPCs) may themselves take in an underlying simulation, then an emergent-narrative system will be inherently reactive to player inputs. As such, the stories that emerge in these systems naturally incorporate player actions. Emergent narrative was articulated by Aylett as a solution to this problem of conciliating interactivity and narrativity [9], which has been called the *narrative paradox* [47, 13] and the *interaction dilemma* [60].

Papers in this area have typically introduced new systems [13, 16, 6] or more generally concerned the engineering of systems [46, 11, 85, 81, 14], though others have discussed emergent narrative through the lens of narrative theory [48, 47], creativity [42, 62, 10], and improvisation [76, 77]. One class of contributions to the discipline has concerned the particular *design practices* that are employed in the domain [50, 49, 75]. Most prominent among these practices is the iterative design cycle of crafting a simulation from which certain kinds of stories may emerge, testing for whether those stories do emerge, hypothesizing why they do not, revising the simulation according to these hypotheses, and repeating these steps until the desired stories reliably emerge. To facilitate this procedure, Suttie et al. introduced the notion of *intelligent narrative feedback*, where a system gives real-time feedback about how well the simulation is satisfying authorial goals [75]; earlier, an authoring tool with a similar purpose was proposed [43].

We contend, however, that this iterative loop fails to capture system performance at its most crucial level—the actual interactive experience. It is only when a player is brought into the loop that an emergent-narrative system materializes as a work of interactive storytelling. Iteratively tuning a simulation will certainly produce a better simulation, but a finely tuned simulation will not necessarily yield a well-crafted interactive experience. We subscribe to the stance articulated elsewhere that AI systems that are intended to support new types of interactive experiences cannot be truly appraised except through actual implemented experiences that get built atop them [52, 74, 41, 84].

Let us look then to implemented interactive experiences in this domain, which have largely been games like *The Sims*, *Dwarf Fortress*, and *Prom Week*. We assert that these experiences, though wildly successful in many respects, exhibit a series of major shortcomings that, in turn, raise significant design challenges for emergent-narrative systems. First, systems from which narratives emerge tend to yield these by virtue of a complex underlying simulation, but often the richness of this simulation (and the very stories that emerge from it) is not made apparent to the player. This problem has prevailed since the earliest emergent-narrative systems and, fittingly, has been called the *Tale-Spin effect* [83]. We believe it will continue to persist until our first two design challenges are met: we must invent new authoring strategies for yielding *modular content* and we must develop new *compositional representational strategies*. Next, systems from which narratives emerge are typically unable to discern those narratives from the uninteresting event sequences that more commonly appear. When this happens, emergent stories may not get showcased by the system, and in turn they may go unnoticed by the player. This problem represents a nascent task area and our third challenge, *story recognition*, for which little work has yet been done. Lastly, even if a system is able to recognize some emergent story, how should it showcase it? We call this last challenge *story support*, and conceive of systems that could showcase (and support) emergent stories as they are unfolding.

In this paper, we introduce a research framework for the design of interactive experiences in the domain of emergent narrative. Specifically, we motivate and discuss each of the above design challenges and, for each, summarize relevant prior work and propose new approaches that future work might take. Because our concerns are at a level of system design (the player experience) above what has typically been discussed (simulation crafting) in prior work directly situated in this area, we more often refer to promising work in other domains. Likewise, because we are most interested in widely disseminated interactive experiences, we use systems like *Dwarf Fortress* and *The Sims*, rather than *FearNot!* and *Virtual Storyteller*, as examples. It is our hope that this paper will encourage new projects that may advance the medium of emergent narrative by taking on these difficult challenges.

2 Challenge 1: Modular Content

By *content*, we mean here the material used to express the underlying system state to the player, such as the animations and Simlish icons in *The Sims* or the textual dialogue in *Prom Week*. To the degree that such content is instantial (*i.e.*, non-procedural or ‘canned’), authorial burden grows monotonically with a system’s state space. This is intuitive: the more states a system can get into, the more states that must be expressible by the system, and thus the more content that must be authored for it. So while, for instance, games with huge state spaces are often lauded for this very enormity [38, 4, 87], that same property yields a serious authoring challenge: when a system cannot express many or most of the states that its simulation yields, the latter will appear to the player much simpler than it in fact is—this is the *Tale-Spin* effect, mentioned above [83]. Because its systems often employ rich simulations that exhibit huge state spaces, this problem is endemic in emergent narrative. Note that there are other forms of authored material, such as the Edith scripts associated with objects in *The Sims* and volition rules in *Prom Week*, that already exhibit modularity and thus avoid this blow-up of authorial burden. But a player’s experience of this authored procedural material (modularized simulation elements) is gated by the *expressive content*, which unfortunately tends to be instantial.

This issue is best alleviated by more content, but for games with particularly massive state spaces, it is not feasible to author by hand as much content as would be needed. Consider the case of *L.A. Noire*, a AAA game that does *not* have a massive state space, but for which a team of authors produced a script exceeding 2,000 pages [32]. This accumulation of dialogue is commensurate to that of ten full-length motion pictures—all this for an interactive experience that is very linear. The game, inasmuch as its narrative is concerned, cannot get into that many states—how else could it have a *script*? If this is the authorial burden commanded by *L.A. Noire*, imagine what is commanded by a system that can get into an essentially infinite number of states. The fact of the matter is that, in emergent narrative, standard authoring practice provides an author no hope for producing a base of content sufficient to give even marginal coverage to a massive state space.

2.1 Potential Approaches

We need to invent new, smarter authoring practices that are tailored to accommodate the massive state spaces and nonlinear properties that are characteristic of emergent narrative. We propose that such a scheme should exhibit two properties: first, with relatively little effort, a human should be able to produce many units of content that each express specific aspects of underlying state; second, these units of content should be largely independent of one another, but in a way that still affords sequencing and recombination.

Let us ground this out in an example. *Prom Week* [55] is a simulation game that features human-authored dialogue exchanges between NPCs. Each of these dialogue exchanges is composed of five to ten lines of dialogue and has on it a set

of precondition rules that specify what must be true about the storyworld for the exchange to be enacted [70]. For instance, a precondition rule might enforce that a dialogue exchange only be enacted when the two characters who would carry it out are dating. But only some subset of the lines of dialogue that make up this exchange are likely to actually depend on the characters dating, which means that all the exchange’s other lines (that do not have this dependency) could have very well been used in cases where the characters were not dating. This represents misspent authoring effort, in that this content gets deployed less often than it could be. In [71], we worked to break down *Prom Week* dialogue exchanges into their constituent lines of dialogue (and then annotated the lines for their individual dependencies) in order to facilitate them getting deployed more often in a way that would increase coverage of the game’s state space. An obstacle arose, however, in that many lines strictly depended on others (*e.g.*, due to anaphora) and thus could not be used independently of them.

From this exploration, we impart two general lessons that correspond to the content desiderata we outlined above: good content in emergent narrative should express one or few specific aspects of system state and should be as context-independent as possible, while still supporting content sequencing and recombination. This way, small units that each express individual things may be flexibly recombined into larger content units that express all those things simultaneously. Moreover, by being sequencable, content units may be used in a coordinated way that affords *emergent linear experiences* (*e.g.*, by having units reference earlier ones). Content should also be explicitly marked up for what it expresses about underlying state. When this is done, a system may be able to reason about its current state at any time in order to select content (or recombine content into a larger unit) that expresses that state. Currently, we are developing an authoring tool whose design is specifically informed by these two lessons [72].

3 Challenge 2: Compositional Representational Strategies

Related to the severe authorial burden endemic to emergent narrative is another challenge that pertains to the actual *deployment* of content. Rather than the makeup of individual content units, we are concerned here with how such units may be deployed simultaneously in a coordinated expression of underlying state. In systems underpinned by simulations whose primitives (and the procedures that operate on them to evolve state) are easily expressible, system state can be expressed *compositionally* by deploying content units that collectively express all active simulation primitives and procedures. This is classically the case in systems that simulate Newtonian physics. Let us consider *Angry Birds* [68], a game with a minor physics simulation whose primitives are the positions of game objects and whose procedures are the forces of classical mechanics. While *Angry Birds* can get into a lot of states, it has no trouble expressing any of them because these primitives and procedures happen to be easily expressible using *graphical logics*, which afford a particularly feasible *representational strategy*. By ‘representational strategy’, we mean a strategy for expressing internal processes

at the surface (usually through the deployment of expressive content). Graphical logics are a brand of *operational logics* (abstract procedural tropes) by which a system may model things like movement and collision quite easily [53]. While graphical logics *can* be used rhetorically to make different kinds of meaning [79, 80]—this has been done compellingly in games like *Passage* [67] and *Dys4ia* [7]—they really work best for expressing state in simulations of physical spaces.

So what representational strategy would afford compositional expression of state in a simulation that models, say, character internal state, as would be typical in emergent narrative? This is an open question, and it appears that until new (or currently immature) representational strategies are developed specifically for expressing things like this, experiences underpinned by simulations of such will not be able to express state compositionally. This means that, lacking such a strategy, designers of these experiences must author content units that essentially are each meant for big chunks of system state. But the better a content unit expresses a particular compound state, the more dependent it becomes on that state, and so the less generalizable it becomes to other states and the less often it can be deployed, as we discussed in the last section.

3.1 Potential Approaches

This is perhaps the hardest challenge we present in this paper, but we propose one potential compositional representational strategy that employs *generative dialogue* and *procedural animation*. A number of projects have explored using techniques from natural language generation to facilitate generative dialogue in interactive narrative systems [20, 69, 63, 51, 35]. A handful of other projects have explored expressive procedural character animation [57, 36]. If each character in an underlying simulation could express her internal state at the surface through bespoke dialogue or animation, the system could express *its* underlying state compositionally by enacting all such procedures for all (on-screen) characters. Another compelling approach would be to utilize a modality that conventionally has *not* been used to express character internal states: sound. Indeed, Ian Horswill is currently exploring this in the experimental game *MKULTRA* [34].

4 Challenge 3: Story Recognition

Systems whose narratives emerge from simulations currently have no way of discerning the very stories they support. To illustrate this, let us consider perhaps the most lauded work of emergent narrative, *Dwarf Fortress* [3]. A roguelike set in procedurally generated fantasy worlds, *Dwarf Fortress* is best known for its rich underlying simulation. This game is often talked about in terms of the stories that emerge from this simulation (both by fans [1] and academics [39]) and its creator, Tarn Adams, has called *Dwarf Fortress* a “story generator” [86]. Undoubtedly, its simulation generates narratives—but does it know that it does?

There is typically no explicit narratological modeling whereby narratives in simulationist systems get composed; this would constitute a top-down approach

to story generation, but simulations work bottom-up. As in other emergent-narrative systems with sufficiently complex underlying simulations, stories happen in *Dwarf Fortress* only incidentally; they are remarkable event sequences among a huge boiling stew of *things happening*. While humans who play experiences like *Dwarf Fortress* are capable of recognizing which event streams are storylike, the system itself is not. We call this challenge *story recognition*: how does one make a system that can discern stories embedded in the morass of data that its simulation produces?¹

4.1 Potential Approaches

Story recognition is related to *story understanding*, in which a system processes an event sequence and attempts to *understand* the story it represents by attributing explanations to the events that account for causal and temporal relations among them. Though related, story understanding is a different task than story recognition. Systems that do the former operate over storylike event sequences as a given, while systems doing the latter must *excavate* such sequences from larger accumulations of data. In this sense, story recognition may be thought of as a preliminary task to story understanding—a story’s event sequence must first be recognized and compiled before it can be understood, though of course a system that does story recognition will probably have to operationalize a rudimentary notion of event temporality and causality, *i.e.*, *storiness*.² To further distinguish the tasks, work in story understanding has typically focused on artifacts like newswire stories [25], while in story recognition the central artifact is a simulation’s event stream. Still, it is likely that techniques or insights from work on story understanding could inform the task of story recognition.

We are aware of preliminary work on story recognition outside the domain of interactive narrative. A handful of recent projects have explored recognizing storylike event sequences from structured representations of sports games (*e.g.*, baseball box scores) [5, 15, 44]. Pablo Gervás has been actively exploring the generation of stories from event-sequence data in chess [29]. In [31], event logs constituting the lives of characters in a social simulation are curated by the system to produce stories that interweave character biographies. (We build on lessons from this preliminary work in our proposal at the end of this section.)

One approach would be to assume that the trace of player actions in a system playthrough will typically represent a storylike sequence, since it may be expected that players will tend to interact with the storyworld in interesting or meaningful ways. In another project, however, we found that players of *Prom Week* cannot be relied upon to produce storylike gameplay, which works against this intuition [8]. However, by generalizing story recognition to mean making inferences about gameplay event-sequence data, prior work on both *goal recognition* and *gameplay summarization* becomes useful for resolving this problem.

¹ We have adapted this term from that of an analogous task in computer vision, *object recognition*, in which discrete objects are identified in image data [37].

² What the most useful notion of storiness is will likely vary by system and application.

In goal (and plan) recognition, a system attempts to infer a player’s goals (and plans) from her playing behavior [23, 30]. This is often done as part of a *player adaptation* scheme, wherein a system dynamically adapts gameplay to improve player satisfaction [88]. In gameplay summarization, a system curates the event sequence constituting a gameplay session in order to summarize it [61, 26, 8]. These tasks differ from that of story recognition in that they are concerned with gameplay itself, not emergent stories that may be embedded in a gameplay session (or in parts of a simulation that are removed from player input). Still, we propose that they are similar enough to potentially inform our task here.

Finally, we propose a method whereby a human specifies interesting event sequences that the system will then attempt to match against sequences generated by its simulation. Indeed, a variant of this method has actually been employed in *The Sims 2*, as we discuss in the next section [17]. One concern here is that an author may not wish to specify precise event sequences, but rather higher-level patterns with major events that must remain constant and allowances for extraneous events that may be interspersed among those. For this, we believe *Playspecs*, a new formalism that supports regular expressions for play-trace data, could prove quite useful [58]. Relatedly, prior work in story understanding provides knowledge frameworks for expressing at higher levels of abstraction desired properties of storylike sequences. We are currently exploring this approach.

5 Challenge 4: Story Support

The last challenge that we will present is an extension of the previous one. Once a system has recognized a storylike event sequence that has emerged from its simulation, what should it do with it?

One thing the system could do is *retell* the story, either after gameplay or during it, perhaps by having NPCs conversationally relate stories about in-game events. This task would be quite related to prior work on gameplay summarization [61, 26, 8], but in this case the impetus is on summarizing a particularly interesting subset of gameplay events, not an entire gameplay sequence.

Perhaps more interestingly, a system could recognize partial event sequences *as they are happening* in order to influence the simulation or gameplay in a way that would support the emergence of interesting stories (*i.e.*, interesting completions of those partial event sequences). We will call this design challenge that of developing methods for *story support*. This problem may appear to be very similar to *narrative planning* [65], *drama management* [66], or *game mastering* [60, 12, 27, 19]—these are certainly related notions (the latter might even be called an approach to story support), but we mean to encourage techniques that do not significantly compromise the integrity of the simulation, but rather gently nudge it toward the emergence of certain desirable event sequences. In other words, story support is not about enforcing some degree of narrativity on the simulation, but rather subtly manipulating it to facilitate and foreground emergent stories. Indeed, simulationist systems are appreciated for the emergent quality of their narratives [38, 4, 87], not necessarily the well-formedness of these

stories according to narratological or dramatic concerns, and so we encourage the development of techniques that may promote emergence above other concerns.

5.1 Potential Approaches

One notable approach to story support was featured in *The Sims 2* [17]. In that game, characters have *fears* and *aspirations* that are expressed by tree data structures that specify event sequences by which certain fears or aspirations may be realized. For instance, a teen character may aspire to have her first kiss, and so the event sequence may include events like meeting, becoming romantically interested in, and eventually kissing another character. As gameplay proceeds, the system parses its event stream to check for partial completions of the sequences specified in characters' trees, and if a match is found, the system may nudge the simulation toward the event that would lead to that sequence's completion. Once a given sequence culminates, the system triggers content that will showcase the story it represents. We note that this particular aspect of the game was received favorably by critics [59, 2].

6 Conclusion

We have introduced a research framework for the design of interactive experiences in the domain of *emergent narrative*, an application area of computational narrative in which stories emerge bottom-up from the behavior of autonomous characters in a simulated storyworld. Prior work in this area has largely concerned the development and tuning of the simulations themselves from which interesting stories may reliably emerge, but we have argued that this will not necessarily improve system performance at its most crucial level—the level of interactive experience. Looking to completed experiences in this area, namely games like *Dwarf Fortress* and *The Sims*, we assert a series of major shortcomings that yield four design challenges. First, the richness of underlying simulations (and the very stories that emerge from them) are often not made apparent to the player. We have argued that it will continue to persist until our first two design challenges are met: we must invent new authoring strategies for yielding *modular content* and we must develop new *compositional representational strategies*. Next, systems from which narratives emerge are typically unable to discern those narratives from the uninteresting event sequences that more commonly appear. When this happens, emergent stories may not get showcased by the system, and in turn they may go unnoticed by the player. This problem represents a nascent task area and our third challenge, *story recognition*, which we have outlined. Lastly, even if a system is able to recognize some emergent story, how should it showcase it? We have called this last design challenge *story support* and forecasted systems that may be capable of supporting and showcasing emergent stories as they are unfolding. It is our hope that this paper will spur new projects that may advance the medium of emergent narrative by taking on these difficult challenges.

References

1. Dwarf Fortress stories. <http://dfstories.com/>, accessed: 2015-06-11
2. Adams, D.: The Sims 2 review. IGN (2004)
3. Adams, T., Adams, Z.: Slaves to Armok: God of Blood Chapter II: Dwarf Fortress (2006)
4. Alexander, L.: Spector: Go emergent. Gamasutra (2013)
5. Allen, N.D., et al.: Statsmonkey: A data-driven sports narrative writer. In: Proc. CMN (2010)
6. Alvarez-Napagao, S., et al.: Socially-aware emergent narrative. In: Agents for Educational Games and Simulations (2012)
7. Anthropy, A.: Dys4ia (2012)
8. Antoun, C., Antoun, M., et al.: Generating natural language retellings from Prom Week play traces. In: Proc. PCG in Games (2015)
9. Aylett, R.: Narrative in virtual environments-towards emergent narrative. In: Proc. Narrative Intelligence Symposium (1999)
10. Aylett, R., Louchart, S.: I contain multitudes: creativity and emergent narrative. In: Proc. Creativity & Cognition (2013)
11. Aylett, R., et al.: Unscripted narrative for affectively driven characters. Computer Graphics and Applications (2006)
12. Aylett, R., et al.: Managing emergent character-based narrative. In: Proc. Intertain (2008)
13. Aylett, R.S., et al.: Fearnot! In: Proc. IVA (2005)
14. Bevensee, S.H., Schoenau-Fog, H.: Conceptualizing productive interactivity in emergent narratives. In: Proc. ICIDS (2013)
15. Bouayad-Agha, N., Casamayor, G., Wanner, L.: Content selection from an ontology-based knowledge base for the generation of football summaries. In: Proc. ENLG (2011)
16. Brom, C., et al.: Emohawk. In: Proc. ICIDS (2009)
17. Brown, M.: The Power of Projection and Mass Hallucination: Practical AI in The Sims 2 and Beyond. In: Proc. AIIDE [Invited Talk] (2006)
18. Cardoso, P., Carvalhais, M.: Breaking the game: The traversal of the emergent narrative in video games. Science and Technology of the Arts (2013)
19. Carpentier, K., Lourdeaux, D.: Diegetization. In: Proc. ICIDS (2014)
20. Cavazza, M., Charles, F.: Dialogue generation in character-based interactive storytelling. In: Proc. AIIDE (2005)
21. Cavazza, M., Charles, F., Mead, S.J.: Character-based interactive storytelling. Intelligent Systems (2002)
22. Chauvin, S., et al.: An out of character approach to emergent game narratives. Proc. FDG (2014)
23. Cheng, D.C., Thawonmas, R.: Case-based plan recognition for real-time strategy games. In: Proc. Game-On (2004)
24. Crawford, C.: Storytron—interactive storytelling (2006)
25. DeJong, G.: An overview of the frump system. Strategies for NLP (1982)
26. Dominguez, M., Young, R.M., Roller, S.: Automatic identification and generation of highlight cinematics for 3d games. In: Proc. FDG (2011)
27. Figueiredo, R., et al.: Emergent stories facilitated. In: Proc. ICIDS (2008)
28. Frank, A., Stern, A., Resner, B.: Socially intelligent virtual Petz. In: Proc. Socially Intelligent Agents (1997)

29. Gervás, P.: Composing narrative discourse for stories of many characters: A case study over a chess game. *Literary and Linguistic Computing* (2014)
30. Ha, E., et al.: Goal recognition with markov logic networks for player-adaptive games. In: *Proc. AIIDE* (2011)
31. Hassan, S., et al.: A computer model that generates biography-like narratives. In: *Proc. Computational Creativity* (2007)
32. Helgeson, M.: L.A. Noire: Rockstar resets the bar with its upcoming crime thriller. *Game Informer* (2010)
33. Hoek, M., Theune, M., Linssen, J.: Generating game narratives with focalization and flashbacks (2014)
34. Horswill, I.: 4.4 Sonification of character reasoning. *Artificial and Computational Intelligence in Games: Integration* (2015)
35. Horswill, I.D.: Architectural issues for compositional dialog in games. In: *Proc. GAMNLP* (2014)
36. Horswill, I.D.: Lightweight procedural animation with believable physical interactions. *Computational Intelligence and AI in Games* (2009)
37. Jain, R., Kasturi, R., Schunck, B.G.: *Machine vision*. McGraw-Hill New York (1995)
38. Juul, J.: The open and the closed: Games of emergence and games of progression. In: *Proc. CGDC* (2002)
39. Kelly, C.F.: *Dwarf Fortress: Narratives of Multiplicity and Deconstruction*. Bachelor's thesis, University of Amsterdam (2013)
40. Klein, S., et al.: *Automatic novel writing: A status report*. University of Wisconsin (1973)
41. Koenitz, H.: Five theses for interactive digital narrative. In: *Proc. ICIDS* (2014)
42. Kriegel, M., Aylett, R.: Emergent narrative as a novel framework for massively collaborative authoring. In: *Proc. IVA* (2008)
43. Kriegel, M., et al.: An authoring tool for an emergent narrative storytelling system. In: *Proc. INT* (2007)
44. Lareau, F., Dras, M., Dale, R.: Detecting interesting event sequences for sports reporting. In: *Proc. ENLG* (2011)
45. Louchart, S., Aylett, R.: Solving the narrative paradox in VEs. In: *Proc. IVA* (2003)
46. Louchart, S., Aylett, R.: Emergent narrative, requirements and high-level architecture. In: *Proc. Hellenic Conference on Artificial Intelligence* (2004)
47. Louchart, S., Aylett, R.: The emergent narrative theoretical investigation. In: *Proc. Narrative and Interactive Learning Environments* (2004)
48. Louchart, S., Aylett, R.: Narrative theory and emergent interactive narrative. *Continuing Engineering Education and Life Long Learning* (2004)
49. Louchart, S., et al.: Authoring emergent narrative-based games. *Game Development* (2008)
50. Louchart, S., et al.: Purposeful authoring for emergent narrative. In: *Proc. ICIDS* (2008)
51. Lukin, S.M., Ryan, J.O., Walker, M.A.: Automating direct speech variations in stories and games. In: *Proc. GAMNLP* (2014)
52. Mateas, M.: *Expressive ai: A hybrid art and science practice*. Leonardo (2001)
53. Mateas, M., Wardrip-Fruin, N.: Defining operational logics. *Proc. DiGRA* (2009)
54. Maxis: *The Sims*. Electronic Arts (2000)
55. McCoy, J., et al.: Prom Week. In: *Proc. FDG* (2013)
56. Meehan, J.R.: *Tale-spin*. In: *IJCAI* (1977)
57. Neff, M., Fiume, E.: Methods for exploring expressive stance. *Graphical Models* (2006)

58. Osborn, J., et al.: Playspecs: Regular expressions for game play traces. In: Proc. AIIDE (2015)
59. Park, A.: The Sims 2 review. GameSpot (2004)
60. Peinado, F., Gervás, P.: Transferring game mastering laws to interactive digital storytelling. In: Proc. TIDSE (2004)
61. Pizzi, D., et al.: Automatic generation of game level solutions as storyboards. Computational Intelligence and AI in Games (2010)
62. Rank, S., et al.: Creativity in configuring affective agents for interactive storytelling. In: Proc. ICCG (2012)
63. Reed, A.A., et al.: A step towards the future of role-playing games: The spyfeet mobile rpg project. In: Proc. AIIDE (2011)
64. Riedl, M.O., Young, R.M.: From linear story generation to branching story graphs. Computer Graphics and Applications (2006)
65. Riedl, M.O., Young, R.M.: Narrative planning: Balancing plot and character. Artificial Intelligence Research (2010)
66. Roberts, D.L., Isbell, C.L.: A survey and qualitative analysis of recent advances in drama management. Systems Science and Applications (2008)
67. Rohrer, J.: Passage (2007)
68. Rovio Entertainment: Angry Birds. Chillingo (2009)
69. Rowe, J.P., Ha, E.Y., Lester, J.C.: Archetype-driven character dialogue generation for interactive narrative. In: Proc. IVA (2008)
70. Ryan, J., Walker, M.A., Wardrip-Fruin, N.: Toward recombinant dialogue in interactive narrative. In: Proc. INT (2014)
71. Ryan, J.O., et al.: Combinatorial dialogue authoring. In: Proc. ICIDS (2014)
72. Ryan, J.O., et al.: Toward natural language generation by humans. In: Proc. INT (2015)
73. Silva, A., Raimundo, G., Paiva, A.: Tell me that bit again... bringing interactivity to a virtual storyteller. In: Virtual Storytelling (2003)
74. Stern, A.: Embracing the combinatorial explosion. In: Proc. ICIDS (2008)
75. Suttie, N., et al.: Theoretical considerations towards authoring emergent narrative. In: Proc. ICIDS (2013)
76. Swartjes, I., Vromen, J.: Emergent story generation: Lessons from improvisational theater. In: Proc. INT (2007)
77. Swartjes, I.M.T.: Whose story is it anyway?: how improv informs agency and authorship of emergent narrative. Ph.D. thesis (2010)
78. Theune, M., et al.: The virtual storyteller (2003)
79. Treanor, M., et al.: Proceduralist readings: How to find meaning in games with graphical logics. In: Proc. FDG (2011)
80. Treanor, M., et al.: The micro-rhetorics of Game-O-Matic. In: Proc. FDG (2012)
81. Truesdale, J., et al.: Suitability of modelling context for use within emergent narrative. In: Proc. ICIDS (2013)
82. Walsh, R.: Emergent narrative in interactive media. Narrative (2011)
83. Wardrip-Fruin, N.: Expressive Processing (2009)
84. Wardrip-Fruin, N., Mateas, M.: Envisioning the future of computational media: The final report of the media systems project (2014)
85. Weallans, A., Louchart, S., Aylett, R.: Distributed drama management: beyond double appraisal in emergent narrative. In: Proc. ICIDS (2012)
86. Weiner, J.: Where do dwarf-eating carp come from. New York Times (2011)
87. Winslow-Yost, G.: SimCity's evil twin. The New Yorker (2013)
88. Yannakakis, G.N., Hallam, J.: Real-time game adaptation for optimizing player satisfaction. Computational Intelligence and AI in Games (2009)