

The Story So Far on Narrative Planning

Rogelio E. Cardona-Rivera¹, Arnav Jhala², Julie Porteous³, R. Michael Young¹

¹Division of Games, University of Utah, Salt Lake City, UT, USA

²Department of Computer Science, North Carolina State University, Raleigh, NC, USA

³School of Computing Technologies, RMIT University, Melbourne, Australia

{r.cardona.rivera | rmichael.young}@utah.edu, ahjhala@ncsu.edu, julie.porteous@rmit.edu.au

Abstract

Narrative planning is the use of automated planning to construct, communicate, and understand stories, a form of information to which human cognition and enaction is pre-disposed. We review the narrative planning problem in a manner suitable as an introduction to the area, survey different plan-based methodologies and affordances for reasoning about narrative, and discuss open challenges relevant to the broader AI community.

The Role of Narrative

The field of artificial intelligence has grappled with modeling story reasoning since its beginning (McCarthy 1990), due in part because this ability is thought to underlie or inevitably result from human cognition (Winston 2011).

Research on *narrative intelligence* progresses along three efforts that encompass a multitude of intertwined natural language, commonsense, and multi-agent reasoning tasks: narrative construction, communication, and understanding (Mueller 2013). Common across these efforts is the use of *automated planning* as a formal, rigorous, and common vocabulary for framing advances in the field. This is because AI planning naturally reasons over concepts – e.g., agents, objects, states, events – central to plot structure and its communication (Young 1999).

While not all narrative intelligence research *focuses* on planning, it is predominantly plan-based or plan-like. Applications include efforts to model human cognition (Cardona-Rivera et al. 2016), achieve human-level performance on language processing tasks (Martin et al. 2018), demonstrate independent creativity (Summerville et al. 2017), structure human-computer interaction (Porteous, Cavazza, and Charles 2010), explain AI rationales (Riedl 2016), and play adventure games (Hausknecht et al. 2020). Strikingly, AI research that does *not* focus on *narrative planning* is rediscovering the utility of data structures and algorithms that form its basis. For example, neuro-symbolic systems that use narrative representations outperform non-trivial baselines in commonsense reasoning (Bosselut et al. 2019; Cohen 2020).

We thus feel the time is ripe to take stock of the state-of-the-art in narrative planning. This research community

has converged upon representational and reasoning commitments necessary to account for story telling and understanding in people (Cardona-Rivera and Young 2019). These commitments are important to survey for advancing AI toward performing as robustly and flexibly as people do. At the same time, this area is relevant where we already see use of narratives, when having greater predictive control over their effects would benefit society; e.g., in narratives for personalized learning (Wang et al. 2017), rehabilitation therapy and healthcare communication (Ozer et al. 2020), and intelligence analysis (Lukin and Eum 2023). And while there is no standard for specifying narrative planning problems (Hayton et al. 2020; Shirvani and Ware 2020) there is a vibrant community of practice, whose advances and unsolved problems bear relevance to the broader AI community.

One work in recent history has surveyed narrative planning (Young et al. 2013), but its contours are imprecise in technical detail. In contrast, we survey the field in formal depth that is sufficient to precisely describe common themes. We structure our survey as a tour guided by the question:

How might a computer system tell a story?

The answer is deceptively trite: *it depends*. In unpacking this answer, we illustrate why narrative planning is so vast. Along the way, we use an extended case example to identify where our community has converged and diverged. As we discuss later, the points of convergence are centered on elements deemed necessary for a computer-generated story to be perceived *as a story* by an audience. In turn, the points of divergence concern what *layers of narrative detail* are needed, what *narrative effects* are desirable, and what *planning approaches* are used for modeling.

As a consequence of our exploration, we will cut across narrative planning systems that address the tasks of narrative understanding, construction, and communication. Finally, we briefly sketch promising open problems that might stimulate the broader community’s pursuit of new directions.

Pre-rigorous Notions in an Example

Before introducing this area in technical depth, we first present more-intuitive descriptions of key narrative concepts. We introduce a running example that makes our motivating question more specific: *how might a computer system tell a story like the one depicted in Figure 1?*



Figure 1: Panels from Marvel’s *Star Wars* #39–#44, which depict Princess Leia, Han Solo, Chewbacca, and C-3PO realizing they are in the belly of a giant space worm. The crew originally piloted the *Millennium Falcon* into the beast, thinking it was an asteroid cave where they could hide from the relentlessly-pursuing *Empire*. In this paper, we explore how an automated planning system would be able to tell such a story.

The depicted excerpt is from *Star Wars* #39–#44 (Goodwin, Williamson, and Garzon 1977), a Marvel Comics adaptation of the film *Star Wars Episode V: The Empire Strikes Back* (Kershner 1980). Despite its brevity, it is rich enough to discuss *narrative* in two senses that narrative planning models must contend with: as a *communicative act* and as a *designed artifact*.

Narrative as a Communicative Act

Figure 1 reflects an implicit communicative choice: the choice of *narration*—the narrative’s surface form realization in a medium (Hühn and Sommer 2013). While we use a comic, we could have used a film, book, or video game.

Audiences familiar with the comic’s source film may note differences in content beyond narration; namely, in the *plot*—the narrative’s virtual world (Rimmon-Kenan 2002). The comic differs from the film in its *events* (causally and purposively related changes in states-of-affairs). To wit, there are differences in the dialogue of *characters* (anthropomorphized intention-driven agents), and event *locations* (spatial context).

Plot differences reflect why the choice of narration is more than a matter of aesthetic taste: the medium constrains what plot information *can* be narrated (Elliott 2004). Above, the constraints are from the *comic form*; i.e., syntax and semantics of visual language (Martens, Cardona-Rivera, and Cohn 2020). This explains why the last panel contains dialogue to clarify the ship’s maneuver; dialogue absent from the film because the maneuver itself can be shown. The information that *is* narrated is the *discourse*—a time-ordered, intentionally-selected subset of the plot (Genette 1980).

Algorithmically telling a story akin to Figure 1 demands formalizing plot, discourse, and narration. This is often not practically feasible, which invites the choice of which layers to model and which others to control for. The choice ultimately depends on *authorial intent* (Pratt 1977)—the communicative goals (e.g., entertain, teach, inspire) to accomplish via the telling of the story. Above, the authors entertain the audience by carefully preserving ambiguity to build suspense until the last panel reveals a surprise: the crew’s “safe haven” is in fact a deadly threat.

Narrative as a Designed Artifact

Telling a story is a *design task* (Simon 1996). Methods to accomplish design tasks are not strictly “right” or “wrong,” but rather “better” or “worse” along *qualitative dimensions* that the designer cares about. In storytelling, quality is assessed relative to authorial intent. Thus, the quality of a story goes beyond just *what* is told. An author may also care about *how* it is told, *why* it is told, to *whom*, *when*, and *where*—all to varying, use-inspired degrees and extents.

Consider that in Figure 1, Han Solo’s actions do not simply emerge from his rational deliberation as an agent in a task environment. While his behavior may accomplish the authorial goal of hiding from the Empire, it is irrational with respect to leading the *Millennium Falcon*’s crew to *overall* safety. Nonetheless, his actions are *authorially* planned for precisely because of their discourse-level narrative quality: they advance the plot in a deliberately *suspenseful* way.

Ultimately, a narrative is successful to the degree it accomplishes an author’s intent. While similar to the process of anticipating a plan’s execution dynamics (as in robotic task planning; e.g., Cashmore et al. 2019), it is further complicated by the nature of storytelling: the domain of discourse is different when plans are narrated, as plan effects are manifest in the *mental states* of the audience. Thus, a narrative plan’s success is not directly observable and must be empirically assessed. For our example, the comic would succeed as a narrative artifact if it predictably elicits entertaining suspense in its audiences.

The Narrative Planning Challenge

The central challenge for developing narrative planning systems is two-fold. The first challenge is computationally modeling both senses of narrative – as communicative act and as designed artifact – in terms of structural features of plans and their construction (search) processes. This could mean modeling narratives in such a way that classical planners can produce them; e.g., via compilation (Haslum 2012; Christensen, Nelson, and Cardona-Rivera 2020). More commonly, it means developing novel planners with expanded knowledge representations and reasoning mechanisms. We focus on this latter approach.

Importantly, the choice of modeling technique presumes a researcher has identified a desired *narrative quality*—i.e., the authorial intent that should be achieved by the narrative planner during its operation. This narrative quality defines the second challenge: the empirical assessment of how well a developed narrative planner predictably elicits (in an audience) the phenomenon it purports to model. In other words, the researcher must assess their system’s success by answering the question: “to what degree does my narrative planner accomplish the authorial intent I care about?” The answer usually manifests as an experimental inquiry to determine how well a resulting narrative plan, its construction process, or both evokes the desired communicative goal in an audience. For example, given a plan-based representation of Figure 1, the SUSPENSER narrative planner (Cheong and Young 2014) predicts what discursive information would elicit suspense in an audience, as evidenced by several human-subjects studies.

Introduction to Narrative Planning

Narrative planning frames a planning agent as a *story director*. A story director must craft a *narrative plan* that represents a *narrative*: the product of a narration of a sequence of events that constitutes a trajectory through states-of-affairs. Like other planning agents, a story director crafts a plan by searching over a graph constructed from a declarative description of a task environment. In our case, the description is a *narrative problem*, isomorphic to a *classical problem*.

Revisiting Classical Problems We briefly review the classical problem representation that most narrative planners rely on: STRIPS (Fikes and Nilsson 1971), where a problem is a tuple $P = \langle \mathcal{L}, I, A, G, f_{\text{cost}} \rangle$. \mathcal{L} is a set of atomic well-formed formulae (wff) or their negations, composed from a formal language of predicate \mathbb{P} , constant \mathbb{C} , and variable \mathbb{V} symbols (and no function symbols).¹ $I \subseteq \mathcal{L}$ is an initial state that obeys the closed-world assumption, $G \subseteq \mathcal{L}$ is a set of goal conditions, and A is a set of actions.

A state is a set of literals. An action is a state-transition represented by a triple $a = \langle \text{PRE}(a), \text{ADD}(a), \text{DEL}(a) \rangle$; respectively the precondition, add, and delete lists, all subsets of \mathcal{L} . An action a is applicable in a state s if $\text{PRE}(a) \subseteq s$. Applying it results in a state $s' = (s \setminus \text{DEL}(a)) \cup \text{ADD}(a)$ and incurs a cost per the function $f_{\text{cost}}: A \rightarrow \mathbb{R}^{0+}$; unless otherwise indicated, we assume $f_{\text{cost}} = 1$.

The solution to P is a classical plan $\pi = [a_1, \dots, a_m]$. This ordered sequence of actions $a_i \in A$ transforms the problem’s initial state I to a state s_m that satisfies the goal; i.e., $G \subseteq s_m$. Its cost is computed as $c(\pi) = \sum_{a_i \in \pi} f_{\text{cost}}(a_i)$. The execution trace $(\pi, I) = [I, a_1, s_1, \dots, a_m, s_m]$ of π from I is an alternating sequence of states and actions, starting with I , such that s_i results from applying a_i to s_{i-1} .

Narrative Planning Problems and Solutions

A *narrative (planning) problem* is broadly defined using the same ingredients: \mathcal{L} , I , A , G , and f_{cost} . However, unlike its

¹Words in true-type font are predicates. When prepended with ‘?’, they are variables. When TitleCase, they are constants.

classical counterpart, a narrative problem is a formal representation of authorial intent (including the task environment with which to achieve it). It thus encodes the communicative goals that the researcher would like to accomplish via the narrative planner. These goals are the ones the story director *must* aim for during the construction of a narrative plan.

Critically, authorial intent encompasses more than just the *achievement* of conditions specified within G . The planner also tacitly encodes authorial intent via the *maintenance* of conditions *during plan construction* demanded by the researcher’s desired quality. The narrative plan construction process is therefore key: it must guarantee that the solution π exhibits the authorially-intended qualities when it (or some morphism of it) is narrated to an audience.

For example, the authorial intent of the PROVANT planner (Porteous and Lindsay 2019) is the generation of narrative that exhibits a Hollywood-style struggle between a protagonist and antagonist. This is guaranteed by more than just the statement of conditions to achieve in G : it is *also* guaranteed by its plan construction process. During operation, PROVANT rules out portions of the search space that do not conform to the canonical Hollywood form. To expand, the comic in Figure 1 would not be an output reachable by PROVANT, as it does not exhibit the desired authorial intent.

A consequence of what narrative problems mean is that (by default) classical planners are insufficient for storytelling.² To illustrate, we reconsider Han’s apparent irrationality: his actions are not what one might expect from (for example) a cooperative multi-agent belief-desire-intention (BDI) planning architecture (Rao and Georgeff 1995). In fact, his actions might never be considered for inclusion in π because they potentially thwart the intent of all other agents. What we need instead is the ability to reason about how elements added to π achieve authorial intent, sometimes to the apparent detriment of the characters within π .

The preceding considerations explain why classical notions of plan quality (e.g., cost, length) are insufficient: the anticipated audience’s reception must be a critical part of a story director’s plan quality assessment. Han’s actions are deliberately chosen because of their expected value to the *tellability* of the resulting narrative (Bruner 1991).

Points of Divergence in Narrative Planning

As mentioned, researchers typically focus on dimensions of narrative plan quality they care about. For instance, a great deal of recent work has modeled agents who obtain mistaken beliefs and act per them such that they fail. These failed actions, like the disparities of belief used to prompt them, can be used to create irony as in the SABRE planner (Shirvani, Ware, and Farrell 2017), support the illusion of theory of mind as in the IMPRACTICAL planner (Teutenberg and Porteous 2013), and build tension as in the HEADSPACE planner (Sanghrajka, Young, and Thorne 2022). Of these, HEADSPACE *would* be able to generate narratives like Figure 1: the crew’s mistaken beliefs are what lead them inside the giant space worm to begin with.

²The storytelling limitations of classical planners are not simply addressed by non-classical (e.g., conformant) approaches.

The choice of which phenomenon to model is the major point of divergence in the narrative planning community. Different researchers seek to model different narrative-theoretic dimensions, each with unique rationales to justify why a given dimension matters for narrative plan quality. There is no consensus on which dimensions matter *most*. We further contend: because storytelling is a design task, there is no universal set of “most important” narrative-theoretic dimensions in characterizing stories. Transitively, there is a rich (potentially infinite) set of storytelling forms that researchers may reason about via planning systems.

Reasoning about diverse narrative phenomena requires adapting the narrative problem representation, planning process, or – more commonly – both. In the cited examples, the problem representation is expanded in different ways. In HEADSPACE, the wff in \mathcal{L} , I , and G are expanded to admit statements about positive and negative character beliefs, and actions in A must specify belief-based PRE(a), ADD(a), DEL(a) lists. This is like IMPRACTICAL, which further distinguishes conscious(?char) and at(?char, ?location) in \mathcal{L} to reason about characters who witness others doing things at given locations, and transitively, characters who can predict when *others* do so, to afford generating stories in which characters deceive one another.

Points of Convergence in Narrative Planning

While the diverse motivations for modeling narratives as plans has precluded a standard specification of narrative problems, the narrative planning community has tacitly converged on several “fundamental particles.” These include characters, locations, and actions as distinguished plot concepts. These concepts are distinguished because they are privileged in human cognition as the basis for *mental models*—our mental simulation of possible worlds—and *event models*—our mental simulation of sequences of events. Both mental models and event models underlie our ability to make sense of stories (Cardona-Rivera and Young 2019). We describe other key points of convergence below.

Several narrative planners distinguish *outcome* G_o from *trajectory* G_t goals within G . The former are akin to classical planning ones: desired outcomes for the story’s end. The latter are similar to state-space *landmarks* (Hoffmann, Porteous, and Sebastia 2004), plan-space *islands* (Hayes-Roth and Hayes-Roth 1979), and search *control knowledge* (Kvarnström and Doherty 2000). Whereas these offer *guidance* to non-narrative planners, narrative problems admit trajectory goals to afford more-direct expression of authorial intent over solution narrative plans (Riedl 2009).

Several narrative planners further partition A into two narratological classes. *Happenings* A_h are actions that can occur without reason, e.g., an accident. *Non-happenings*, on the other hand, must be intended (Bratman 1987)—they are carried out by characters in service of goals they have adopted, and are termed *intentional* actions A_I . A happening is isomorphic to a STRIPS-style action, whereas an intentional action is a quadruple $a = \langle \text{PRE}(a), \text{ADD}(a), \text{DEL}(a), \text{CHA}(a) \rangle$, where PRE(a), ADD(a), and DEL(a) are as before and CHA(a) is a list of terms in \mathbb{C} that denote plot characters.

Another point of broad convergence is the role of a narrative plan’s *causal coherence*, that each action that takes place has its preconditions satisfied. This causal backbone is critical for audiences to derive *temporal* sequences—people cannot easily understand stories without a spatio-temporal frame (Radvansky and Zacks 2014). A narrative plan’s causal coherence contributes the tacitly-valued quality of *comprehensibility*, that the resulting narrative can be understood by audiences. Several studies have demonstrated that causally-coherent narrative plans rendered as textual or filmic media can themselves be used to predict average human answers to comprehension questions about the plan’s constituent actions (Christian and Young 2004; Cardona-Rivera et al. 2016). For a given action $a_i \in \pi$, we can predict people’s responses to the questions *Why / How did a_i happen?*, *What enabled a_i to happen?*, and *What was the consequence of a_i ?*

The last point of broad convergence concerns another tacitly-valued dimension of narrative plan quality: *believability*, that the resulting narrative does not thwart the audience’s *willing suspension of disbelief* or sense of being in a fictional world (Holland 2003). Believability is what motivated partitioning A into A_h and A_I in the first place. This partition is needed because *ceteris paribus* a story director will search to satisfy G , without regard to whether all $a_i \in \pi$ are believable for characters to execute. For example, suppose our narrative problem describes part of our comic, with Leia and Han, aboard the Millennium Falcon. Further suppose the author specifies $G_o = \{(\text{not } (\text{conscious}(\text{Han})))\}$. A plausible narrative plan π narrated via templated text is:

- (1) Leia picks up a blaster. (2) Leia stuns Han.

This plan is causally coherent; i.e., the blaster must be picked up in order to stun Han. However, is this plan believable as a story? Perhaps, but nothing in π would structurally justify that. Why Leia stuns Han is not clear. Arguably, we need more context to understand how Leia’s actions are believable. In other words, the plan lacks *motivational coherence*: the plan contains actions that do not appear to be motivated by anything. The IPOCL planner was developed to address this concern (Riedl and Young 2010), under the rationale that *intentions* are a distinguishing feature of anthropomorphic activity (Bates 1994; Dennett 1989) and that intentions provide motive to act (Bratman 1987). IPOCL expanded the wff in \mathcal{L} , I , and G to admit statements about character intentions. These afford the expression of *character* goals as modal sentences of the form *intends*(c , g_c), where c represents any term from \mathbb{C} that denotes a plot character and $g_c \in \mathcal{L}$ represents a condition the character c intends to accomplish. These intentions constrain the search space of the story director: an action $a \in A_I$ can only be considered for expanding the search space if it is possible to make a part of a character c ’s sub-plan $\pi_{g_c} \subseteq \pi$ that accomplishes g_c . Today, narrative planners largely take intentional actions as a given,³ as we do for the remainder of this paper.

³State-of-the-art planners such as HEADSPACE, SABRE, and IMPRACTICAL all reason about character intentions as sub-parts of their primary modeling purpose.

The Different Layers of Narrative

Fully modeling a narrative as we have defined would require formalizing all layers we have discussed thus far: plot, discourse, and narration. This task is often not practically feasible, leading researchers to pragmatically choose *which* layers to model. Each layer re-casts the classical problem representation to mean different narrative concepts. Below, we review different framings and extract common themes.

Plot Planning

The bulk of the work in narrative planning has focused on modeling plot, in which the narrative plan is meant to represent the plot structure: a sequence of actions taken by characters in the story that evolve the virtual world from its initial configuration to an author-desired one. Here, the narrative problem is similar to a multi-agent planning (MAP) one (Brafman and Domshlak 2013; Belle et al. 2023). This is because characters exhibit intention dynamics as they are orchestrated by the story director toward achieving G . But unlike MAP, it is plausible for characters to *intentionally* conflict, under the rationale that this phenomenon features prominently in global Western narratives (Herman, Manfred, and Ryan 2010). Whereas plans with resource or coordination conflicts would be ruled out in MAP, they are made possible by design in the GLAIVE planner (Ware and Young 2014): its search space allows the story director to find plans where characters adopt intentions that are mutually exclusive, relaxing the IPOCL requirement that a character c 's subplan π_{g_c} necessarily achieve g_c . In other words, characters can take actions toward goals, but fail to accomplish them.

Plot-level plans ought to be sound with respect to the story world domain and problem in which they take place. But because these task environments are virtual, they can themselves be modified to suit particular storytelling goals. For instance, the INITIAL STATE REVISION (ISR) algorithm (Riedl and Young 2005) partitions I into true, false, and undetermined sentence sets whose combinations determine the plot's set of *alternative possible worlds* (Ryan 1991). While this is similar to I being an open-world state, ISR shifts undetermined sentences into true or false as convenient to accomplish storytelling goals, which only makes sense per the synthetic nature of the plot's virtual world.

Because quality is assessed with respect to an audience, plot plans must be narrated in some way in order to empirically evaluate whether the plot model achieves its intended effect. Today, most narrative planners tacitly follow Reiter and Dale's (2000) Natural Language Generation (NLG) pipeline, which starts with (plot) *Content Determination*, is followed by (discourse) *Content Structuring*, and ends with (narration) *Linguistic Realization*. Plot planners typically have a perfunctory NLG pipeline, to control for any spurious discourse and narration effects on comprehension due to natural language. But because AI planning is itself an effective model of NLG, a separate line of research has focused on narrative discourse planning.

(Narrative) Discourse Planning

As in conventional discourse planning, narrative plans are meant to reflect the informational and intentional structure

of a discourse (Grosz and Sidner 1986): a sequence of *communicative* actions that are intended to evolve the mental state of the audience from its (initial) state prior to experiencing the discourse to an author-desired one.

Thus, when representing discourse, the narrative problem is similar to the more general discourse planning one (Garoufi 2014). This is because the story director relies on the framing of utterances as *speech actions* and treats communication as a goal-oriented process in the space of audience beliefs (Cohen and Perrault 1979). But unlike its more general formulation (cf. Gatt and Krahmer, 2018), narrative discourse planning does not assume that the beliefs of the audience monotonically increase over time. Nor does it assume that the goal must be a belief that the audience should obtain at the *end* of the discourse. In fact, the audience's *belief dynamics* – i.e., the trajectory of belief expansion, contraction, and revision operations (cf. Alchourrón, Gärdenfors, and Makinson, 1985) – is a key determinant of *narrative coherence*; i.e., that the artifact is received *as* a narrative (Herman 2013). This quality is particularly important in human-AI applications: information parsed by people as a story is better comprehended and better retained by them relative to non-story information (Fisher and Radvansky 2018).

In the pipeline approach, the input plot plan π_{plot} represents the knowledge base and informs the material communicative goals that the narrative discourse planner will strive for. Thus, elements of π_{plot} remain fixed and are reified: they become part of the language \mathcal{L} for the narrative *discourse* problem. Narrative discourse planning operates in belief space, and thus \mathcal{L} also admits statements about belief—these are used to specify the expected audience's mental state before narration I , the author-desired mental state outcome for the plot's telling G , and the communication actions A that will effect changes in audience belief.

Thus, for narrative discourse planning, the actions in A tacitly specify a model of *belief-based* narrative-theoretic phenomena: they manipulate the presentation of plot details to elicit particular belief dynamics in audiences, which in turn result in particular narrative effects. Examples include the use of *staging* to direct audience attention as in MISER (Matthews et al. 2017) and the elicitation of audience *inferencing* via INFER (Niehaus and Young 2014).

Narration Planning

Narrative discourse planners are often tightly-coupled to the narration, such that a solution narrative discourse plan is itself the realization in a given medium. As a result, narration planning is relatively under-explored; the FIREBOLT cinematic realizer (Thorne et al. 2019) is a notable exception.

Other work in this area has sought to break-away from the NLG pipeline, under the rationale that the storytelling's *form* (i.e., discourse and narration) cannot be separate from its (plot) *content* (Elliott 2004). Work includes adapting the direction of the pipeline via specification of plot-level landmarks based on required discourse “snapshots” as in PLOTSHOT (Cardona-Rivera and Li 2016), as well as doing away with the pipeline altogether in order to co-evolve plot and discourse (coupled to narration) simultaneously as in BIPOCL (Winer and Young 2016).

Three Key Narrative Planning Paradigms

Techniques to solve narrative problems are tightly-coupled to a problem representation—the one needed to formalize the particular class of narrative phenomena of interest to the modeler. A fundamental assumption they share is that storytelling is computationally well-modeled as a search process.

At the same time, different planning paradigms afford *different* ways of modeling this search process and transitively, how to *think* about modeling narrative phenomena. Different paradigms offer different spaces that shape the story director’s range of generatable stories, or *expressive range* (Summerville 2018). We cover three broadly-used paradigms.

Modeling Narrative via Plan-space Planning

Plan-space search operates over a graph in which nodes represent partial-plans and arcs represent plan refinement operations (Kambhampati, Knoblock, and Yang 1995). For example, in POCL planning (Weld 1994), refinements are introduced to guarantee that (1) no preconditions remain unsatisfied (recorded via *causal links*), and (2) no action in the plan could be ordered such that it threatens to undo (establish the opposite condition of) a causal link between two other actions. These respectively are *Open Condition* (OC) and *Threatened Link* (TL) flaws.

Plan-space narrative planning affords modeling story phenomena in terms of *narrative-theoretic* plan construction flaws and fixes. That is, a modeler must add (to OC and TL) new classes of flaws and fixes that in some way capture a dimension of plan quality with narrative import. For example, IPOCL introduced three flaws to model the intention dynamics that make character actions more believable when narrated. When IPOCL adds an action $a \in A_I$ to fix some other flaw, it must guarantee that a is at some point added to a sub-plan $\pi_{g_c} \subseteq \pi$ that accomplishes $g_c, \forall c \in \text{CHA}(a)$; while a remains un-added, the partial-plan (under refinement) has an *Unknown Intent* (UI) flaw. If any sub-plan π_{g_c} is not preceded by an action that establishes the effect $\text{intends}(c, g_c)$ for the sub-plan’s $c \in \mathbb{C}$, the partial-plan has an *Open Motivation* (OM) flaw, fixable by adding such an action. And if any two sub-plans $\pi_{g_c}^1, \pi_{g_c}^2 \subseteq \pi$ assert opposite sub-goals (i.e., $g_c^1 = \neg g_c^2$) the partial-plan has a *Threatened Intent* (TI) flaw, fixable by ordering $\pi_{g_c}^1$ after $\pi_{g_c}^2$ or *vice-versa*.

This modeling strategy is attractive in that it directly shapes the underlying search space in a way that facilitates providing theoretical guarantees about the space of solutions. Flaws and their fixes respectively identify partial-plans that would not be solutions to a given narrative-theoretic problem class, and the algorithmic means to refine partial-plans such that they *do* become solutions (or fail in the attempt). The full modeling strategy then is to empirically demonstrate that narrative plans with the structural quality to-be-preserved (via flaw detection and refinement) do in fact elicit (in audiences) a particular narrative-theoretic effect of interest to the modeler. The drawback of plan-space narrative planners lies in their performance, but plan-space heuristics which seek to offset that penalty, such as those codified by VHPOP (Younes and Simmons 2003), give insights into potential directions.

Modeling Narrative via Hierarchical Planning

Hierarchical formalisms are varied, but broadly share the property of being more expressive and complex than classical planning: STRIPS-style *primitive* actions are complemented with isomorphic more-abstract *compound* actions that require *decomposition*, or associated sub-plan (Bercher, Alford, and Höller 2019). Decomposition happens via *methods* representing sub-goals that require sub-plans, to be further decomposed until all compound actions are reduced to primitives. The idea is that a decomposition method $d = \langle a_C, \pi_C \rangle$ maps a compound action a_C to a sub-plan π_C that depends on (i.e., has preconditions that relate to) $\text{PRE}(a_C)$ and contributes to (i.e., has effects that relate to) $\text{EFF}(a_C) = \text{ADD}(a_C) \cup \text{DEL}(a_C)$ (Bercher et al. 2016).

Hierarchical narrative planning affords modeling story phenomena in terms of narrative-theoretic recursive specifications, akin to grammar rewriting rules. This is well-suited for representing a wide variety of story phenomena that depend on *abstraction*, across the narrative layers. In plot, for example, character intentions can be straightforwardly codified (Cavazza, Charles, and Mead 2002): a compound action a_C may assert $\text{intends}(c, g_c) \in \text{ADD}(a_C)$ with its decomposition being the sub-plan π_{g_c} that achieves it. Another example is DARSHAK (Jhala and Young 2010), which uses hierarchies to bridge cinematic discourse and narration: a compound action’s effects $\text{EFF}(a_C)$ represent discourse-layer information about the plot being filmed, and a method represents a *cinematic idiom* that identifies a narration-layer sub-plan of film shots thought to achieve $\text{EFF}(a_C)$.

Modeling via hierarchical formalisms is attractive due its potential for *authorial leverage* (Chen, Nelson, and Mateas 2009): that is, hierarchical story directors afford users significant power to define narrative plan quality aligned to their authorial intent, across all narrative layers. Like in other hierarchical planners, authors may introduce elements of “advice” to a hierarchical story director: decomposition methods afford ways to encode typical (not necessarily optimal) action sequences or *scripts* (Schank and Abelson 1975). For example, this is used in the GDPOP planner (Winer and Cardona-Rivera 2018) to suggest idiomatic film edits during the construction of a cinematic narration plan. A hierarchical representation also provides a mechanism to exercise authorial leverage via specification of non-classical temporally-extended goals. For example, partially-ordered compound tasks have been used to scaffold the structure of novel variants of the TV drama *Friends* (Cavazza, Charles, and Mead 2002). Related approaches to such non-classical goal specification are discussed below in the context of heuristic search planning.

Modeling Narrative via Heuristic Search

Heuristic search planning, the dominant current approach, plans via forward search through a state space, wherein states are evaluated on the basis of general, domain-independent heuristics (Bonet and Geffner 2001). In this paradigm, the challenge is to discover how best to encode narrative-theoretic phenomena in order to be able to leverage the efficient performance of such planners.

One strategy has been to encode narrative phenomena as constraints within the narrative domain or as control knowledge to scaffold the structure of narrative. In the NETWORKING system (Porteous, Charles, and Cavazza 2013), the use of constraints scaffolds construction of story genre-consistent sub-plans that enable characters to realize their intentions. It has also featured as control knowledge to guide story development, as with MADAME BOVARY, the *affective* storytelling system (Pizzi and Cavazza 2007). Such approaches have been shown to work well in practice, but they fail to capture more general narrative-theoretic phenomena.

Hence, other work has looked to develop *heuristics* that directly reason about narrative-theoretic properties. This is the case for GLAIVE (Ware and Young 2014), a forward search planner which generates motivationally coherent narrative plans, as discussed earlier. GLAIVE introduces a heuristic that incorporates character intentions via a *goal graph* constructed via possible execution traces of plans that contain intentional actions. This goal graph is then used alongside an FF-style plan graph (Hoffmann and Nebel 2001) to calculate the heuristic estimate of states in search.

Compilation Techniques for Narrative Planning While we focus on modeling narrative via extensions to planning constructs, compilations do merit attention. A compilation is a systematic remodelling of the narrative problem that affords solving it via a classical planner. The potential of this approach was demonstrated by Haslum, who was motivated to leverage state-of-the-art heuristic search planners available at the time. Clearly, there are other benefits as well: a common Planning Domain Description Language (McDermott et al. 1998), and access to a greater number of general solvers and heuristics. Modeling narrative via compilation is thus a specific variant of modeling via heuristic search.

Haslum (2012) introduced two compilations that targeted motivational coherence, first addressed by the IPOCL planner as discussed earlier. The META PLANNING compilation is not domain-independent, and is not guaranteed to produce a semantically equivalent representation. That is, a classical plan that solves the compiled problem is not guaranteed to be motivationally coherent in the IPOCL sense. Haslum (p. 392) notes: “although most of the time they will be.”

In contrast, the JUSTIFICATION TRACKING compilation *is* guaranteed to be semantically equivalent. However, this equivalence comes at a price—the resulting classical problem is *exponentially* larger, in terms of three factors in the original narrative problem: (1) the number of possible character intentions codified via $\text{intends}(c, g_c)$; (2) the number of intentional actions $a \in A_I$ relevant to the respective intentions; and (3) the number of characters $\text{CHA}(a)$ that feature in the respective intentional actions a . When tested on the original benchmark problem by Riedl and Young, this compilation allowed forward-chaining A* with LM-Cut heuristic (Helmert and Domshlak 2009) to outperform IPOCL’s run time by three orders of magnitude. However, these speedups are not always guaranteed (Teutenberg and Porteous 2013). In fact, relative to FF applied to this compilation, IMPRACTICAL is two orders of magnitude faster, and GLAIVE is three orders of magnitude faster.

To our knowledge, the only other work that has pursued narrative compilation is by Christensen, Nelson, and Cardona-Rivera (2020), who proposed that character beliefs (like those in HEADSPACE) can be compiled into a *simpler* narrative problem that GLAIVE can solve. But because GLAIVE is narrative-theoretic, the full potential of narrative compilation to *classical* planning remains unknown. We note however that the purpose of compilation is less on modeling narrative phenomena *per se*, and more for the benefit of more tractably and flexibly solving narrative problems.

Open Challenges

While many open challenges are common to narrative and non-narrative planning, others are driven by the differentiation required for plans to serve as narrative artifacts. We discuss three of the most significant: expanding expressive range, using desired cognitive effects to drive generation, and performing ethically conscientious narrative planning.

Increased Expressive Range There is a need to increase the expressive range of narrative planning systems. Their expressive range – the breadth of structural features of output artifacts – must differ from those of conventional planning because the structural features of narrative impose a set of distinct properties not readily accountable for by conventional planning knowledge representations. For example, plot structures contain actions that fail, explicit conflict between the actions and goals of their agents, and controlled increases in potential but unrealized threats to a plan’s execution. Narrative discourse elements may work to intentionally obfuscate aspects of their domain of discourse from a reader or viewer in the audience. They may intentionally drive false beliefs or they may carefully curate and promote uncertainty, prompt repeated belief revision around specific concepts, play off of the relationships between the time of story events and the time of their telling, and focus as much on the trajectory of cognitive and emotional states during the experience of a reader as the set of beliefs they hold at the end of a narrative experience.

Cognitive-driven Narrative Generation There is a need to strengthen our knowledge of the connection between narrative generation and narrative comprehension. In the field of cognitive psychology, narrative comprehension has long been studied by building and refining cognitive models (e.g., Graesser and Franklin 1990). These models generally posit that human readers progressively construct a mental concept graph during narrative comprehension that represents events in a story and their causal, temporal and spatial relationships, character intentions, and other elements with parallels to the contents of plans and their search spaces. So similar are the parallels between representations in the two fields that narrative planning researchers have adapted the experimental methods used to validate these cognitive models for use by psychologists in evaluating the motivational coherence of generated plans (Riedl and Young 2010). However, little work has been done to directly take in to account comprehension during narrative generation. Further, no work to date has accounted for the expected cognitive burden involved in comprehending a planner-generated story.

Ethically Conscientious Narrative Planning Narratives are powerful communicative tools (Suzuki et al. 2018): they educate, inspire, entertain, and effect behavioral change in long-lasting ways. As a result, narrative planners that can carefully scaffold the construction of mental and event models has great promise, but also great potential for harm (Coman and Aha 2017). Disinformation campaigns, radicalization narratives, and the codification of stereotype knowledge could potentially be amplified through the computational support that narrative planners offer.

We believe there is potential to explore plan-based user models that simulate how groups of users might interpret the output of narrative planners, and thereby drive the construction of a narrative plan toward solutions that are safer than sibling outcomes in the space of alternatives. This would be possible by using existing work that already predicts normative answers to questions about stories (Cardona-Rivera et al. 2016). We foresee the use of predicted comprehension to prune plans in the search space that do not satisfy a “comprehension constraint;” i.e., some (to be discovered) metric that codifies how likely it is that the intended comprehension of the generated story actually manifests in the audience.

We also note that narrative planning might help us structurally understand (in terms of action, state, and heuristic representations) what makes harmful stories so effective, which may equip us to systematically guard against them.

Summative Review and Conclusion

We have reviewed narrative planning, covering the rationales and methodologies employed-to-date for modeling a narrative’s plot, discourse, and narration. Our survey discussed knowledge representation and reasoning affordances for narrative in technical depth, with the hope that we can concretely connect to others within the automated planning community and beyond.

Several modern human-centered AI systems use narrative to structure interaction, resulting in skill-development being more motivating, learning more fun, and intelligent systems more understandable. Despite a reliance on narrative, systems that are *not* plan-based or plan-like lack computationally-precise understanding of how to structure stories toward an eventual desired comprehension by people. This results in ad-hoc accounts, loss of generality, and re-inventing the wheel; problems that narrative planners are designed to guard against (Young 1999). We thus expect that narrative planning will continue to play a significant role in the creation of theories and systems that aim to characterize narrative intelligence writ large, and thereby help address a foundational problem within artificial intelligence.

Acknowledgements

This material is based upon work supported by the U.S. National Science Foundation (Grant #2046294) and the U.S. Air Force Office of Scientific Research (Award FA9550-20-1-0355). We thank Stephen G. Ware and our anonymous reviewers who offered tremendously helpful feedback during peer review. Figure 1 is included under fair use exemption of U.S. Copyright Law and is restricted from further use.

References

- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic*, 510–530.
- Bates, J. 1994. The role of emotion in believable agents. *Communications of the ACM*, 37(7): 122–125.
- Belle, V.; Bolander, T.; Herzig, A.; and Nebel, B. 2023. Epistemic planning: Perspectives on the special issue. *Artificial Intelligence*, 316: 103842.
- Bercher, P.; Alford, R.; and Höller, D. 2019. A Survey on Hierarchical Planning - One Abstract Idea, Many Concrete Realizations. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*.
- Bercher, P.; Höller, D.; Behnke, G.; and Biundo, S. 2016. More than a Name? On Implications of Preconditions and Effects of Compound HTN Planning Tasks. In *Proceedings of the 22nd European Conference on Artificial Intelligence*, 225–233.
- Bonet, B.; and Geffner, H. 2001. Planning as heuristic search. *Artificial Intelligence*, 129(1-2): 5–33.
- Bosselut, A.; Rashkin, H.; Sap, M.; Malaviya, C.; Celikyilmaz, A.; and Choi, Y. 2019. COMET: Commonsense Transformers for Automatic Knowledge Graph Construction. arXiv:1906.05317.
- Brafman, R. I.; and Domshlak, C. 2013. On the Complexity of Planning for Agent Teams and its Implications for Single Agent Planning. *Artificial Intelligence*, 198: 52–71.
- Bratman, M. 1987. *Intention, Plans, and Practical Reason*. Cambridge, MA, USA: Harvard University Press.
- Bruner, J. 1991. The Narrative Construction of Reality. *Critical Inquiry*, 18(1): 1–21.
- Cardona-Rivera, R. E.; and Li, B. 2016. PLOTSHOT: Generating Discourse-constrained Stories around Photos. In *Proceedings of the 12th AAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 2–8.
- Cardona-Rivera, R. E.; Price, T. W.; Winer, D. R.; and Young, R. M. 2016. Question Answering in the Context of Stories Generated by Computers. *Advances in Cognitive Systems*, 4: 227–246.
- Cardona-Rivera, R. E.; and Young, R. M. 2019. Desiderata for a Computational Model of Human Online Narrative Sensemaking. In *Proceedings of the AAI Spring Symposium on Story-enabled Intelligence*.
- Cashmore, M.; Coles, A.; Cserna, B.; Karpas, E.; Magazzini, D.; and Ruml, W. 2019. Replanning for situated robots. In *Proceedings of the 29th International Conference on Automated Planning and Scheduling*, 665–673.
- Cavazza, M.; Charles, F.; and Mead, S. J. 2002. Character-Based Interactive Storytelling. *IEEE Intelligent Systems*, 17(4): 17–24.
- Chen, S.; Nelson, M.; and Mateas, M. 2009. Evaluating the Authorial Leverage of Drama Management. In *Proceedings of the 5th AAI Conference on Artificial Intelligence and Interactive Digital Entertainment*.

- Cheong, Y.-G.; and Young, R. M. 2014. Suspenser: A Story Generation System for Suspense. *IEEE Transactions on Computational Intelligence and AI in Games*, 7(1): 39–52.
- Christensen, M.; Nelson, J. M.; and Cardona-Rivera, R. E. 2020. Using Domain Compilation to add Belief to Narrative Planners. In *Proceedings of the 16th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*.
- Christian, D. B.; and Young, R. M. 2004. Comparing Cognitive and Computational Models of Narrative Structure. In *Proceedings of the 19th National Conference on Artificial Intelligence*, 385–390.
- Cohen, P. R. 2020. Back to the Future for Dialogue Research: A Position Paper. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*.
- Cohen, P. R.; and Perrault, C. R. 1979. Elements of a Plan-Based Theory of Speech Acts. *Cognitive Science*, 3(3): 177–212.
- Coman, A.; and Aha, D. W. 2017. Cognitive Support for Rebel Agents: Social Awareness and Counternarrative Intelligence. In *Proceedings of the 5th Conference on Advances in Cognitive Systems*.
- Dennett, D. C. 1989. *The Intentional Stance*. MIT Press.
- Elliott, K. 2004. Literary film adaptation and the form/content dilemma. *Narrative Across Media: The Languages of Storytelling*, 220–243.
- Fikes, R. E.; and Nilsson, N. J. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3-4).
- Fisher, J. S.; and Radvansky, G. A. 2018. Patterns of Forgetting. *Journal of Memory and Language*, 102: 130–141.
- Garoufi, K. 2014. Planning-Based Models of Natural Language Generation. *Language and Linguistics Compass*, 8(1): 1–10.
- Gatt, A.; and Kraemer, E. 2018. Survey of the State of the Art in Natural Language Generation: Core Tasks, Applications and Evaluation. *Journal of Artificial Intelligence Research*, 61: 65–170.
- Genette, G. 1980. *Narrative Discourse: An Essay in Method*. Cornell University Press.
- Goodwin, A.; Williamson, A.; and Garzon, C. 1977. *Star Wars #39–#44* [Comic]. Marvel.
- Graesser, A. C.; and Franklin, S. P. 1990. QUEST: A cognitive model of question answering. *Discourse Processes*, 13(3): 279–303.
- Grosz, B. J.; and Sidner, C. L. 1986. Attention, Intentions, and the Structure of Discourse. *Computational Linguistics*, 12(3): 175–204.
- Haslum, P. 2012. Narrative Planning: Compilations to Classical Planning. *Journal of Artificial Intelligence Research*, 44: 383–395.
- Hausknecht, M.; Ammanabrolu, P.; Côté, M.-A.; and Yuan, X. 2020. Interactive Fiction Games: A colossal adventure. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, 7903–7910.
- Hayes-Roth, B.; and Hayes-Roth, F. 1979. A Cognitive Model of Planning. *Cognitive Science*, 3(4): 275–310.
- Hayton, T.; Porteous, J.; Ferreira, J. F.; and Lindsay, A. 2020. Narrative Planning Model Acquisition from Text Summaries and Descriptions. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*.
- Helmert, M.; and Domshlak, C. 2009. Landmarks, critical paths and abstractions: what’s the difference anyway? In *Proceedings of the 19th International Conference on Automated Planning and Scheduling*, 162–169.
- Herman, D. 2013. *Storytelling and the Sciences of Mind*. MIT Press.
- Herman, D.; Manfred, J.; and Ryan, M.-L. 2010. *Routledge Encyclopedia of Narrative Theory*. Routledge.
- Hoffmann, J.; and Nebel, B. 2001. The FF Planning System: Fast Plan Generation Through Heuristic Search. *Journal of Artificial Intelligence Research*, 14: 253–302.
- Hoffmann, J.; Porteous, J.; and Sebastia, L. 2004. Ordered landmarks in planning. *Journal of Artificial Intelligence Research*, 22: 215–278.
- Holland, N. N. 2003. The willing suspension of disbelief: A neuro-psychoanalytic view. *PsyArt*.
- Hühn, P.; and Sommer, R. 2013. Narration in Poetry and Drama. In Hühn, P.; Pier, J.; Schmid, W.; and Schönert, J., eds., *the living handbook of narratology*. Hamburg, Germany: Hamburg University.
- Jhala, A.; and Young, R. M. 2010. Cinematic Visual Discourse: Representation, Generation, and Evaluation. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(2): 69–81.
- Kambhampati, S.; Knoblock, C. A.; and Yang, Q. 1995. Planning as refinement search: a unified framework for evaluating design tradeoffs in partial-order planning. *Artificial Intelligence*, 76(1): 167–238.
- Kershner, I. 1980. *Star Wars V: The Empire Strikes Back* [Film]. Lucasfilm Ltd.
- Kvarnström, J.; and Doherty, P. 2000. TALplanner: A temporal logic based forward chaining planner. *Annals of Mathematics and Artificial Intelligence*, 30: 119–169.
- Lukin, S. M.; and Eum, S. 2023. SEE&TELL: Controllable Narrative Generation from Images. In *Proceedings of the Workshop on Creative AI Across Modalities at the 37th AAAI Conference on Artificial Intelligence*.
- Martens, C.; Cardona-Rivera, R. E.; and Cohn, N. 2020. The Visual Narrative Engine: A Computational Model of the Visual Narrative Parallel Architecture. In *Proceedings of the Annual Conference on Advances in Cognitive Systems*.
- Martin, L. J.; Ammanabrolu, P.; Wang, X.; Hancock, W.; Singh, S.; Harrison, B.; and Riedl, M. O. 2018. Event representations for automated story generation with deep neural nets. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*.
- Matthews, J.; Charles, F.; Porteous, J.; and Mendes, A. 2017. MISER: Mise-En-Scene Region Support for Staging Narrative Actions in Interactive Storytelling. In *Proceedings of*

- the 16th International Conference on Autonomous Agents and Multiagent Systems, 782–790.
- McCarthy, J. 1990. An example for natural language understanding and the AI problems it raises. *Formalizing Common Sense: Papers by John McCarthy*, 355.
- McDermott, D.; Ghallab, M.; Howe, A.; Knoblock, C.; Ram, A.; Veloso, M.; Weld, D.; and Wilkins, D. E. 1998. PDDL—The Planning Domain Definition Language. Technical report, Yale Center for Computational Vision and Control, New Haven, CT, USA.
- Mueller, E. T. 2013. Computational Models of Narrative. *Sprache und Datenverarbeitung: International Journal of Language Processing, Special Issues on Formal and Computational Models of Narrative*, 37(1–2): 11–39.
- Niehaus, J.; and Young, R. M. 2014. Cognitive Models of Discourse Comprehension for Narrative Generation. *Literary and Linguistic Computing*, 29(4): 561–582.
- Ozer, E. M.; Rowe, J.; Tebb, K. P.; Berna, M.; Penilla, C.; Giovanelli, A.; Jasik, C.; and Lester, J. C. 2020. Fostering engagement in health behavior change: Iterative development of an interactive narrative environment to enhance adolescent preventive health services. *Journal of Adolescent Health*, 67(2): S34–S44.
- Pizzi, D.; and Cavazza, M. 2007. Affective Storytelling Based on Characters’ Feelings. In *Proceedings of the AAAI Fall Symposium on Intelligent Narrative Technologies*.
- Porteous, J.; Cavazza, M.; and Charles, F. 2010. Applying planning to interactive storytelling: Narrative control using state constraints. *ACM Transactions on Intelligent Systems and Technology*, 1(2): 1–21.
- Porteous, J.; Charles, F.; and Cavazza, M. 2013. Network-ING: Using Character Relationships for Interactive Narrative. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*.
- Porteous, J.; and Lindsay, A. 2019. Protagonist vs Antagonist PROVANT: Narrative Generation as Counter Planning. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems*, 1069–1077.
- Pratt, M. L. 1977. *Toward a Speech Act Theory of Literary Discourse*. Indiana University Press.
- Radvansky, G. A.; and Zacks, J. M. 2014. *Event Cognition*. Oxford University Press.
- Rao, A. S.; and Georgeff, M. P. 1995. BDI Agents: From Theory to Practice. In *Proceedings of the 1st International Conference on Multiagent Systems*, 312–319.
- Reiter, E.; and Dale, R. 2000. *Building Natural Language Generation Systems*. Cambridge University Press.
- Riedl, M. O. 2009. Incorporating Authorial Intent into Generative Narrative Systems. In *Proceedings of the AAAI Spring Symposium on Intelligent Narrative Technologies*, 91–94.
- Riedl, M. O. 2016. Computational Narrative Intelligence: A Human-centered Goal for Artificial Intelligence. In *Proceedings of the Workshop on Human-Centered Machine Learning at SIGCHI Conference on Human Factors in Computing Systems*.
- Riedl, M. O.; and Young, R. M. 2005. Open-World Planning for Story Generation. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*.
- Riedl, M. O.; and Young, R. M. 2010. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research*, 39: 217–268.
- Rimmon-Kenan, S. 2002. *Narrative Fiction: Contemporary Poetics*. Routledge.
- Ryan, M.-L. 1991. Possible Worlds and Accessibility Relations: A Semantic Typology of Fiction. *Poetics Today*, 12(3): 553–576.
- Sanghrajka, R.; Young, R. M.; and Thorne, B. 2022. HeadSpace: Incorporating Action Failure and Character Beliefs into Narrative Planning. In *Proceedings of the 18th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 171–178.
- Schank, R. C.; and Abelson, R. P. 1975. Scripts, plans, and knowledge. In *Proceedings of the 4th International Joint Conference on Artificial Intelligence*.
- Shirvani, A.; and Ware, S. G. 2020. Camelot: A Modular Customizable Sandbox for Visualizing Interactive Narratives. In *Proceedings of the 12th Workshop on Intelligent Narrative Technologies at the 16th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*.
- Shirvani, A.; Ware, S. G.; and Farrell, R. 2017. A Possible Worlds Model of Belief for State-Space Narrative Planning. In *Proceedings of the 13th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*.
- Simon, H. A. 1996. *The Sciences of the Artificial*. MIT Press, 3rd edition.
- Summerville, A. 2018. Expanding Expressive Range: Evaluation Methodologies for Procedural Content Generation. In *Proceedings of the 14th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 116–122.
- Summerville, A.; Martens, C.; Harmon, S.; Mateas, M.; Osborn, J.; Wardrip-Fruin, N.; and Jhala, A. 2017. From Mechanics to Meaning. *IEEE Transactions on Games*, 11(1): 69–78.
- Suzuki, W. A.; Feliú-Mójer, M. I.; Hasson, U.; Yehuda, R.; and Zarate, J. M. 2018. Dialogues: The science and power of storytelling. *Journal of Neuroscience*, 38(44): 9468–9470.
- Teutenberg, J.; and Porteous, J. 2013. Efficient Intent-Based Narrative Generation Using Multiple Planning Agents. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, 603–610.
- Thorne, B. R.; Winer, D. R.; Barot, C.; and Young, R. M. 2019. Firebolt: A system for automated low-level cinematic narrative realization. In *Proceedings of the 12th International Conference on Interactive Digital Storytelling*, 333–342.
- Wang, P.; Rowe, J. P.; Min, W.; Mott, B. W.; and Lester, J. C. 2017. Interactive Narrative Personalization with Deep Reinforcement Learning. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 3852–3858.

- Ware, S. G.; and Young, R. M. 2014. Glaive: A State-Space Narrative Planner Supporting Intentionality and Conflict. In *Proceedings of the 10th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 80–86.
- Weld, D. S. 1994. An Introduction to Least Commitment Planning. *AI Magazine*, 15(4): 27.
- Winer, D.; and Young, R. M. 2016. Discourse-driven narrative generation with bipartite planning. In *Proceedings of the 9th International Natural Language Generation Conference*, 11–20.
- Winer, D. R.; and Cardona-Rivera, R. E. 2018. A Depth-Balanced Approach to Decompositional Planning for Problems where Hierarchical Depth is Requested. In *Proceedings of the 1st Hierarchical Planning Workshop at the 28th International Conference on Automated Planning and Scheduling*, 1–8.
- Winston, P. H. 2011. The Strong Story Hypothesis and the Directed Perception Hypothesis. In *AAAI Fall Symposium on Advances in Cognitive Systems*.
- Younes, H. L. S.; and Simmons, R. G. 2003. VHPOP: Versatile Heuristic Partial Order Planner. *Journal of Artificial Intelligence Research*, 20: 405–430.
- Young, R. M. 1999. Notes on the use of plan structures in the creation of interactive plot. In *Proceedings of the AAAI Fall Symposium on Narrative Intelligence*, 164–167.
- Young, R. M.; Ware, S.; Cassell, B.; and Robertson, J. 2013. Plans and Planning in Narrative Generation: A Review of Plan-Based Approaches to the Generation of Story, Discourse, and Interactivity in Narratives. *Sprache und Datenverarbeitung: International Journal of Language Processing, Special Issue on Formal and Computational Models of Narrative*, 37(1–2): 67–77.