Advances in Cognitive Systems 6 (2018) 1–18 Submitted 6/2018; published 8/2018

Subjective Experience of Plan-based Intention Revision

Adam Amos-Binks AAAMOSBI@NCSU.EDU
Laboratory of Analytic Sciences, NC State University, Raleigh, NC, USA, 27695

Randall Spain RDSPAIN@NCSU.EDU
Center for Educational Informatics, NC State University, Raleigh, NC, USA, 27695

R. Michael Young YOUNG@EAE.UTAH.EDU
School of Computing and the Entertainment Arts and Engineering Program, University of Utah, Salt Lake City, UT, USA, 84112

Abstract

Creating a story branch in response to unexpected user actions in an interactive narrative is thought to be a desirable feature of an experience manager. Adapting to the user in this manner aspires to improve a their experience. However, few robust mechanisms exist to create story branches and there is a lack of studies that measure a story branch's effect on experience. Our work addresses these limitations by using plan-based intention revision as a branching mechanism and evaluates it using established measures of narrative experience. Our results show that when measuring narrative transportation, we observe a significant main effect for group $F(1, 652) = 10.74$, $p < .01$ where individuals in the treatment condition ($M = 50.70$; $SD = 12.81$) who read narrative texts that included intention revisions reported higher levels of narrative transportation compared to individuals in the control group ($M = 47.43$; $SD = 12.92$). When measuring perceived interest, we observed a significant main effect of group $F(1, 652) = 6.43$, $p < .05$, where individuals in the treatment condition ($M = 28.51$; $SD = 8.97$) who read narrative texts that included intention revisions reported higher levels of perceived interest compared to individuals in the control group ($M = 26.67$; $SD = 9.22$). These results establish the efficacy of using intention revision as a branching mechanism for intentional planning and invites its extension to interactive narrative.

1. Introduction

Interactive narrative leverages our cognitive resources for increased engagement in both learning and entertainment contexts. Plan-based representations of narrative have progressed from non-player character agents that always achieve their intention (Riedl & Young, 2010), to agents that abandon their intention (Ware et al., 2014), and now agents that revise their intentions (Amos-Binks & Young, 2018). This range of agent behavior enables an experience manager agent to create a story branches that accommodate unexpected user actions when a story would otherwise end.

However, a long held assumption is that creating a story branch to accommodate a user's actions improves their experience and is therefore a desirable feature for experience management. Evaluations of plan-based narrative have focused on comprehension (e.g., Cardona-Rivera et al., 2016) and specific narrative devices (e.g., Ware et al., 2014) while developing and evaluating mechanisms for creating story branches has received comparatively less attention (Ramirez & Bulitko, 2014). Despite the continued research into interactive narrative and the commercial success of branching story games, it is still unclear how branching stories affect user experience. The appeal of interactive narrative relies on our proclivity to relate to stories, therefore it is necessary to investigate how the mechanisms used to create story branches affect user experience.

To test this assumption, we evaluate intention revision as a mechanism for creating story branches by measuring subject responses to the narrative transportation scale (Green & Brock, 2000) and perceived interest questionnaire (Schraw, 1997). When intention revision is used to create a story branch, it leverages our cognitive processes to think divergently about an agent's future intentions and we hypothesize this process increases narrative transportation and perceived interest. Our results show that when our plan-based formalization of intention revision is used to create story branches for simulated interactive narratives, subjects who read the resulting text-based narratives reported higher narrative transportation and perceived interest than their linear, non-branching counterparts. Additionally, the results also suggest a more specific type of intention revision (revenge) could increase transportation scale and perceived interest scores even more. Overall, these results support the assumption that branching stories can improve experience, and provide a foundation for future work investigating the effects of intention revision in an interactive narrative.

2. Previous Work

We review three areas of previous work that our approach builds upon. The first, story generation, provides a brief overview of story generation systems and representation. We focus primarily on more formal representations of story, as they lend themselves more easily to other scientific disciplines, namely cognitive psychology. Second, we review concepts of intention revision. While it is new concept, there have been attempts to leverage it in various forms. A final area is narrative experience. Here, we review the relationship between interactive narratives and measures of subjective experience.

2.1 Story Generation

As early as Schank and Abelson (1977), it was theorized that classical planning could represent story plots based on the theoretical overlaps of plot events and the action-oriented, causally-linked, and temporally-ordered properties of plans. Since these first insights, story generation systems have extended representations to capture a range of narrative features (Meehan, 1977; Porteous et al., 2010; Pérez y Pérez & Sharples, 2001).

Our focus is on intentional partial ordered causally linked (called IPOCL) planning (Riedl & Young, 2010), where intentional (goal-oriented) agents take causally linked actions towards their goals. Together, these individual agent goals reach the goal conditions specified in a planning problem. IPOCL story planning operationalized intention by generating solution plans that only contain actions on a causally connected action sequence to an agent's goal. An agent's goal is aggregated into an *intention frame*, along with a motivating plan step and a causally connected

Table 1. Steps in a belief, desire, intention agent's decision making (adapted from Rao & Georgeff (1998)).

#	Agent action
1	make an observation from the environment, ω
2	revise ***beliefs*** based on the observation ω
3	if current ***beliefs*** cause an ***intention*** to be reconsidered then
4	let ***desires*** be the results of options when considering ***beliefs*** and ***intentions***
5	let ***intentions*** be the results of deliberating ***beliefs, desires*** and ***intentions***
6	if current ***beliefs*** introduce complications to an **intention** achieving ***plan*** then
7	let ***plan*** be the results of planning with ***beliefs*** and ***intentions***

action sequence called a subplan. Both the goal and subplan are key in identifying reconsidered intentions and discussed further in Section 3.

Building on IPOCL, Ware and Young 2014 developed the Glaive planner that defined the conflict partial ordered causally linked representation (called CPOCL). CPOCL models conflict between two agents by allowing them to pursue conflicting plans. One agent drops their goal once a causal link threat is introduced to their plan. This conflict model was shown to produce high inter-rater agreement supporting CPOCL's roots in narratological concepts of conflict. In Section 3, we use the same causal link threat formalization as an indication for an agent to reconsider their intention.

2.2 Intention Revision

The use of agent intentions in narrative plan representations is grounded in the Belief Desire Intention theory of mind. Beliefs are facts an agent believes to be true, desires are states of the world an agent wants to be true, and intentions are those desires an agent is committed to make true through action. A concept of intention was first theorized by Bratman (1987) and later formalized for logical agents by Cohen and Levesque (1990). This work lead to the development of belief, desire, intention agents and their decision making abilities (Katsuno & Mendelzon, 1991).

Within the belief, desire, intention research community there has been substantial research on belief revision and update (e.g., Schank, 1986), while there has been only cursory investigations on the cascading effects of belief changes to other mental states, namely intention. This is likely due to the still evolving and contested definition of exactly what an intention is, as well as intention's interdependence on other belief, desire, intention components. As part of an investigation into intention revision logic, Van der Hoek (2007) formalized intention revision in linear time logic based on Table 1.

More specifically, intention revision is concerned with the *reconsider* function (line 4) and its relationship to obtaining new observations (line 2). The *reconsider* function is viewed as a costly process, while new observations (a propositional or temporal formula) are obtained with relative ease, making reconsidering at every possible observation unfeasible. The authors do not specify exactly when agents may reconsider, but that observation and enablement of previously unachievable goals alone are not sufficient for a rational agent to reconsider their intentions. On the other hand, when observations are made that make a current intention unachievable, the agent would be well served to *reconsider* and execute lines 5-8 to develop a new plan for an achievable goal. This was operationalized in a plan-based model of intention revision by Amos-Binks and Young (2018) where causal link threats cause an agent to reconsider an intractable goal and initiate a revision.

Fendt and Young (2017) used the IPOCL representation to develop an algorithm that initiates an intention revision on the agent's part when they receive belief updates. They extended this initial work to a more specific example using suspense as the catalyst for intention revision. Another, more applied, area for intention revision is accommodation (Amos-Binks, 2017). Accommodation is a mediation strategy used in interactive narrative when a human player takes an action that breaks the current plan and forces non-player characters to revise their intentions.

2.3 Narrative Experience

The concept of a *drama manager* that would intelligently manipulate a virtual world and is driven by experience quality originated with Bates (1992). Research in graphics, camera control, non-player character behavior, and player modeling has led to the emergence of more general experience manager agents to manage these mechanisms during interactive narrative gameplay.

Using a plan-based representation for an interactive narrative's plot can lead to users taking unexpected actions (exceptional actions, Robertson & Young (2014)) that break the current story plan. An experience manager agent will accommodate the exceptional action by creating an new story branch in accordance with a measure of experience quality. The assumption that justifies performing accommodations is that it increases a player's sense of agency in gameplay. Research by Ramirez and Bulitko (2014) tested this assumption by executing accommodations that created story branches in accordance with player-types within a plan-based interactive narrative of Little Red Riding Hood. This player-style accommodation strategy was evaluated using an instrument for effectance and enjoyment of interactive narratives (Vermeulen et al., 2010). The experiments found no interaction was found between gaming skill and experience management, but overall accommodating exceptional actions increased fun and agency. This study, while important in its investigation of a long held assumption, still leaves a question about the reasons why accommodations increase agency.

There has been widespread use of the presence questionnaire and immersive tendency questionnaire (Witmer & Singer, 1998) in virtual environments. While our ultimate goal is to develop a robust mechanism for creating a story branches in an interactive environment, our goals in this paper are to investigate the more basic effects of creating story branches using plan-based intention revision in narrative texts. This allows us to isolate the interactive effects in future work but limits the applicability of the presence questionnaire and immersive tendency questionnaire.

To evaluate the effects of plan-based intention revision in narrative texts, we turn to established instruments of experience of narrative text. The transportation scale (Green & Brock, 2000) is motivated by Gerrig's theory of narrative transportation (1993). Transportation is the extent a subject becomes involved with the narrative world and its protagonists. The transportation scale is composed of eleven items (plus one additional item per character in the story) from which three subscales (cognitive, affective, imagery) are identified. Perhaps most importantly, the transportation scale was sensitive to manipulations of text. The authors altered a violent text ("Murder at the Mall") to one without violence and retained much of the structure and theme, and found the transportation scale was sensitive to the change.

Educational psychology has investigated the general concept of student interest as it relates to learning (Hidi & Renninger, 2006). Interest guides children to learn and is a psychological state that emerges from a person interacting with an activity. Situational interest is an attractive psychological state that is short-lived, common across individuals, and within a context. Schraw (1997) proposed a model that distinguishes six sources of situational interest to develop the six subscales that comprise the sources of interest questionnaire. Additionally, the authors developed the perceived interest questionnaire to identify an overall feeling of situational interest in a text. The perceived interest questionnaire focuses exclusively on the reader's assessment of their own feelings of interest while the sources of interest questionnaire assesses the text's content and structure.

3. Generating Plans with Intention Revision

Intentional planning systems generate action sequences that reach the goal conditions of a planning problem. Within these plans, intentional agents either achieve or drop their goals but are limited as they do not adopt new goals when they fail. To address this limitation, we extend intentional planning to generate action sequences with a formal definition of intention revision. We use a running example, *Jail*, to first provide concrete examples of the basic definitions of intentional plans. Second, we continue with the same example to give definitions for plan-based intention revision. Finally, we outline an algorithm for generating plans containing intention revision.

3.1 Intentional Planning

We use definitions for intentional planning based on Riedl and Young's work on IPOCL (2010). Intentional planning differs from classical planning with one key additional constraint on the solutions, all steps in a solution plan must be causally linked to achieving at least one agent's goal (happenings are *fate*'s intention). This causally linked set of actions are referred to as an agent's subplan to achieve their goal. An agent, their subplan and goal are all captured in an aggregate structure called an intention frame that is included in the definition of an intentional plan.

Intentional planning uses agents and their goals to solve intentional planning problems. Our example in Figure 1 has two agents, Smith and the warden, who pursue goals:

Definition 1 (Agent) An agent is a symbol that uniquely identifies a goal-oriented agent.

Definition 2 (Agent Goal) Is a logical sentence that identifies a desired world-state of an agent.

Goals are represented by the *intends (agent, goal)* predicate. In our example, Smith takes action in pursuit of exoneration (*intends (Smith, exonerated(Smith))*) while the warden takes an action to help Smith, *intends (warden, hasTrial(Smith))*, see the domain and problem for full details (Fig. 2) .

The agent who executes an action is called the consenting agent. In Figure 1, Smith executes the *MakeFriends*, *Embezzle*, *RequestTrial*, and *Testify* actions. This is reflected in our Action definition:

Definition 3 (Action) Action A consists of preconditions that must be satisfied before execution, $\textsc{Pre}(A)$, effects that result, $\textsc{Eff}(A)$, and a consenting agent, $\textsc{agent}(A)$, who performs the action.

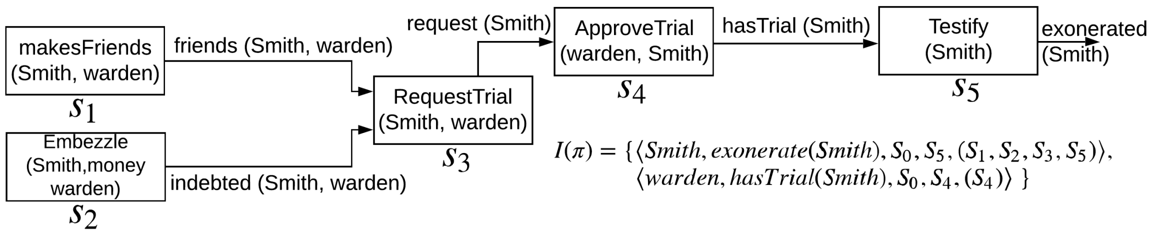

Figure 1. Depiction of the example *Jail* intentional plan (π) including its intention frames, $\mathcal{I}(\pi)$.

Preconditions are literals in a state space whose conjunction must evaluate to true *before* an action's execution. An action's effects are literals whose conjunction evaluates to true *after* A is executed.

An action's name, parameter list, preconditions, effects, and consenting agent describe an *action schema*. An action schema (see Λ in Figure 2) can be instantiated into steps by grounding the free variables and result in the plan steps $s_1 - s_5$ in Figure 3. An agent's goal-oriented actions are within the context of an intentional plan:

Definition 4 (Intentional plan) An intentional plan π is $\langle S, B, O, L, I\rangle$ where the set of steps is S, B the set of binding constraints on the variables of S, O the partial ordering of steps in S, L the set of causal links joining steps in S, and finally I, the intention frame set that define agent subplans.

Figure 1 is an intentional plan. We use typical formal definitions for POCL planning (e.g., Penberthy & Weld, 1992) for bindings, orderings, and causal links. Bindings indicate that two variables must unify (or not be allowed to). Ordering constraints indicate the execution order of plan steps. Causal links record that an effect p of an action a satisfies action b's precondition u and are indicated by the edges labeled with predicates connecting the plan steps in Figure 1.

Definition 5 (Causal links) A causal link, $s \xrightarrow{p} u$, is a tuple $\langle s, p, u\rangle$ where s, u are actions and p is a literal. A causal link records that p is both an effect of s and satisfies the precondition in u.

Causal links are the edges that connect the steps to one another in Fig. 1. The final element of an intentional plan are intention frames. Intention frames structure intentional plan elements into goal-oriented behavior of individual agents.

Definition 6 (Intention Frame) An intention frame is a tuple $\mathcal{I} = \langle a, g, m, \sigma, T\rangle$ where a is an agent, g is a's goal, motivating step $m \in S$ with the effect $\neg g$, the satisfying step $\sigma \in S$ with g as an effect. A subplan for a to achieve g is a set of steps $T \subseteq S$ that a consents to, each step shares at least one causal link to another step in T, and achieves g. Steps in T occur after m and before σ.

Figure 1 shows the two intention frames for *Smith* and *warden*. Finally, intentional plans are generated to solve planning problems, the plan in Figure 1 solves the problem defined in Figure 2.

Definition 7 (Planning problem) A planning problem Φ is a five-tuple $\langle \mathcal{I}, \mathcal{G}, \mathcal{A}, \mathcal{O}, \Lambda\rangle$ where $\mathcal{I}$ and $\mathcal{G}$ are conjunctions of true literals in the initial and goal state respectively, $\mathcal{A}$ the set of symbols referring to agents, $\mathcal{O}$ the set of symbols referring to objects, and Λ a set of action schemata.

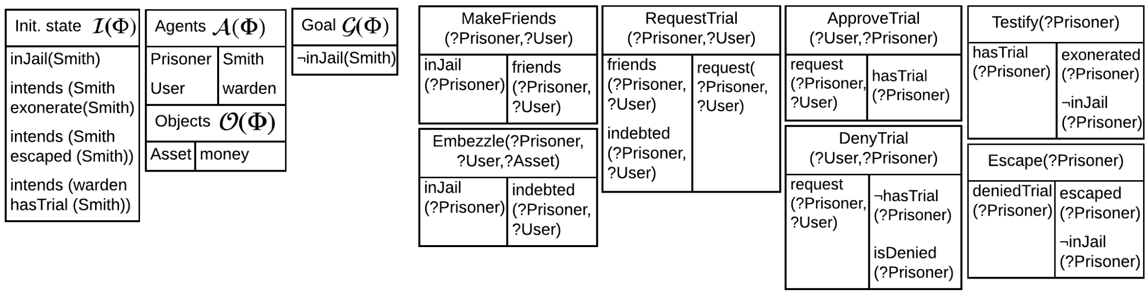

Figure 2. Representation of the *Jail* planning problem (Φ) used by an intentional planner to generate the original *Jail* intentional plan, π.

3.2 Intention Revision

During execution of a plan-based IN, a step is labeled as executed if we have updated its effects in the execution state, where the execution state is a set of consistent, non-modal, ground literals. We use executed steps to determine active intentions.

Definition 8 (Active Intention) An active intention is part of the current plan, $i \in I(\pi)$ where at least one step of the subplan is executed and the satisfying step, $\sigma(i)$ is not executed. The active intentions of a plan are indicated by $I^a(\pi)$.

In Figure 1, Smith's intention of $exonerated(Smith)$ is active from $s_1 - s_4$, until he executes the satisfying step, *Testify* (s_5). Active intentions are useful for identifying reconsidered intentions and support our definition of intention revision. During execution of a plan-based IN, the player agent (the warden) can take actions that introduce causal link threats that prevent a non-player agent from achieving their goal.

Definition 9 (Causal link threat) A causal link threat occurs when a causal link is established $s \xrightarrow{p} u$, and some other step w has effect $\neg p$ and could be executed after s but before u. Executing w in this interval means the precondition q of u is no longer satisfied by s and u will not execute.

In Figure 3, the player agent executes the *DenyTrial* step (s'_4) instead of the planned *ApproveTrial* (s_4), introducing a causal link threat $\neg hasTrial(Smith)$ to the *Testify* action that is part of Smith's *exonerate(Smith)* intention. We refer to an action that introduces a causal link threat at execution time as an exceptional action.

Definition 10 (Exceptional Action) An exceptional action s'_t executed at time t by the user agent, $agent(s'_t) = user$, where one of its effects, $e \in \textsc{Eff}(s'_t)$, introduces a causal link threat to a precondition of a future step $\textsc{Pre}(s_u)$ in the current plan π where $t \leq u$.

This exceptional action causes Smith to reconsider his $exonerate(Smith)$ intention.

Definition 11 (Reconsidered Intention) A reconsidered intention, $\langle \mathcal{I}, \epsilon \rangle$, where $\mathcal{I}_R$ is an active intention and ϵ a literal that introduces a causal link threat to two linked steps in the subplan, $T(\mathcal{I})$.

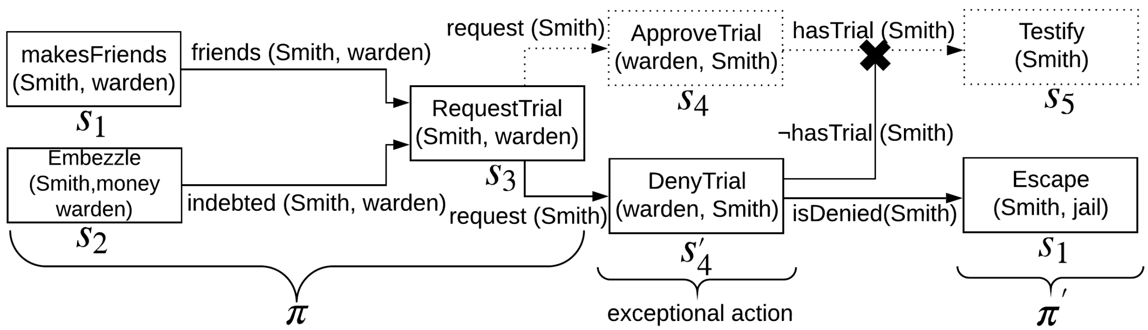

Figure 3. While executing π, a story branch (π') is created in response to an exceptional action S'_4.

Once an intention is reconsidered, an agent will deliberate whether the goal is worth pursuing. Cohen and Levesque (1990) prescribe that an agent should only drop a goal after achieving it or when the agent believes the goal is unachievable. We are interested in the latter:

Definition 12 (Unachievable Goal) A goal is unachievable, g_u, if using a agent's belief state as the initial state, no subplan to achieve $g(\mathcal{I}_R)$ exists.

In our plan-based representation, agents maintain a belief state of their environment represented as sets of consistent, non-modal, ground literals. Agents update their belief state by observing the effects of actions. After the *DenyTrial* action, Smith believes exoneration is unachievable as no subplan exists to achieve exoneration. This belief leads Smith to drop this goal and adopt an intention with an achievable goal *escaped(Smith)* and subplan (s'_1, *Escape* in π') that also solves the problem:

Definition 13 (Intention Revision) An intention revision is $\langle \mathcal{I}_R, \mathcal{I}' \rangle$ where $\mathcal{I}_R$ is an active intention $g(\mathcal{I}_R)$ is unachievable and $\mathcal{I}'$ is an intention frame where $g(\mathcal{I}') \neq g(\mathcal{I}_R)$, and $c(\mathcal{I}') = c(\mathcal{I}_R)$.

3.3 Intention Revision Generation

Our goal is to generate new intentional plans in response to unexpected user actions (exceptional actions) that initiate a *reconsidered intention*. When the new plan is combined with the current plan, it will form our definition of an *intention revision* (Def. 13). We simulate exceptional actions during plan execution and restrict new intentional plans to ones that use intention revision to branch from the previous plan. Once we execute the final step of a plan, we aggregate all the intentional plans (branches) we have generated into a plan-set. With this plan-set, we generate simple text of non-interactive stories that use intention revision as a branching mechanism. These texts are used to evaluate how intention revision affects subject experience in Section 4.

Recall our original intentional plan and planning problem in Figures 1 and 3, respectively, that we use as inputs π and Φ to the method in Table 2. In an interactive narrative, Smith would be a (non-player) character agent and the warden, the user's agent. After some initialization in lines 1 to 2, we enter the main loop and execute the character agent actions s_1-s_3 of π in lines 5 to 7.

Table 2. Pseudocode for introducing an intention revision to the user agent while executing an intentional plan, creating a set of story branches (new plans) as the output.

	METHOD: executePlan INPUT: An intentional plan π that solves Φ, a planning problem. OUTPUT: A set of plans Π that create a branching story using intention revision.
1	let ***currentState*** be equal to the initial state of Φ
2	let ***hasBranched*** be False
3	while ***hasBranched*** is False:
4	let ***step*** be the next unexecuted step from π
5	if the AGENT(***step***) is not the user agent then
6	add EFF(***step***) to ***currentState***
7	label ***step*** as executed
8	if the AGENT(***step***) is the user agent then
9	let ***step*** be applicable action in Φ such that an EFF(***step***) introduces a causal link threat
10	add ***step*** to π
11	add EFF(***step***) to ***currentState***
12	label ***step*** as executed
13	let Φ' be a new planning problem from Φ with its initial state set to the ***currentState***
14	if a solution plan, π', to Φ' exists such that π' and π form an intention revision
15	let Π be the results of executePlan(π', Φ')
16	let ***hasBranched*** be True
17	if there is no such π'
18	Dead end
19	return the union of π and Π

When we reach the first user action, s_4 in π, we enter the *If* block at line 8. Instead of executing s_4 (ApproveTrial), in line 9 we identify an applicable exceptional action s'_4 (DenyTrial) that introduces a causal link threat, $\neg hasTrial(Smith)$ to s_5 (Testify). This causal link threat causes our non-player agent Smith, to reconsider his intention of being exonerated.

After adding s'_4 to the plan and executing the action (lines 10-12), at line 13 we create a new planning problem (Π') with the $currentState$ as the initial state and other elements equal to the original planning problem (Φ). In line 14, we search for a solution plan π' to Π', such that it satisfies our definition of an intention revision. For our example in Figure 3, we find a plan with a single action (*Escape*) that comprises a single intention frame. Finally, we continue execution with π' as the current plan using a recursive call to $executePlan$ and mark the current plan as having branched 15-16). The recursive call enables further branching using the same exceptional user action-intention revision mechanism and conveniently aggregates all story branches together as a plan-set (line 19) when the final step is executed.

3.4 Limitations of Intention Revision

There are two limits to plan-based intention revision whose technical challenges we do not address explicitly in this paper. The first is determining the existence of a plan that contains an agent's

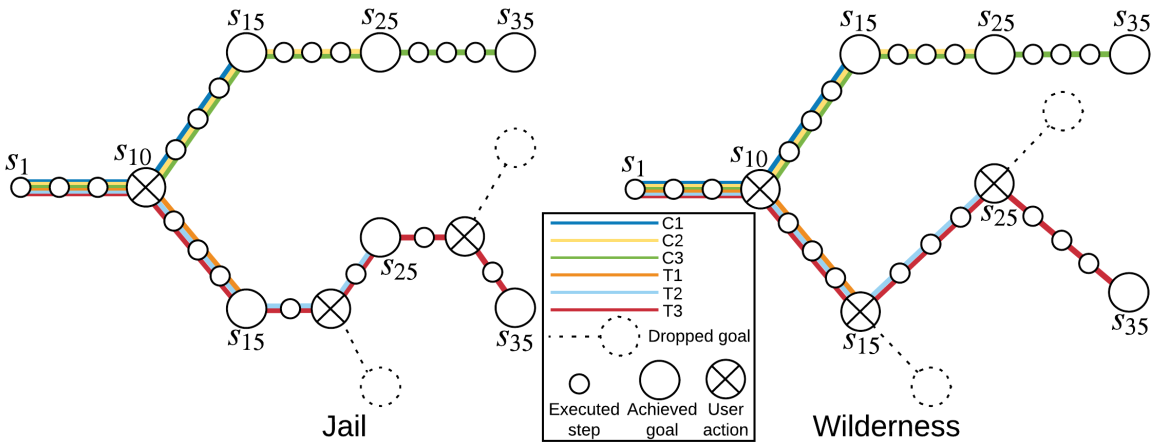

Figure 4. A figure with the two intention revision story structures we used in our experiment.

reconsidered intention. The lack of existence of such a plan supports dropping the reconsidered intention. A second challenge is finding a plan that contains a new intention for an agent who has dropped an intention. This challenge is made more complex as it requires a plan that continues the active intentions from the original plan. A typical approach to address both challenges would be to create a plan-library using diverse planning.

4. Experimental Evaluation

Story branches are an important feature of interactive narrative systems as they enable an Experience Manager to create a story branch in response to user actions that would otherwise end the story. In this section, we describe an experiment who's goal is to establish the efficacy of using intention revision as a branching mechanism. Our experiment investigates how plan-based intention revision affects narrative transportation and perceived interest narrative texts. These results enable us to isolate the effects in an interactive narrative in future work.

4.1 Experimental Design

Story Structures. We designed a longer version of our *Jail* example, as a planning problem in PDDL, Φ_{Jail}, and generated an initial solution plan, π^{C3}_{Jail}, as input to the method in Table 2. This initial solution is for our control group where non-player character agents achieve three intentions without issue and does not contain intention revisions. This type of plan represents the story from an interactive narrative where the user would have executed three actions consistent with π^{C3}_{Jail}. Our algorithm in Table 2 introduces three user actions that require Smith to revise his intentions and generates plans in response, that when aggregated together form π^{T3}_{Jail}. This type of plan represents the story from an interactive narrative where a user chooses actions that require an experience manager to generate a new story branch using intention revision. Both π^{C3}_{Jail} and π^{T3}_{Jail} were designed to

Table 3. Intentional plan steps transformed into story statements read by the pi_{Jail}^{T1} group.

#	Statement
S1	Smith is in jail and wants to be exonerated.
S2	Smith finds an escape plan.
S3	Smith discards the escape plan.
S4	Smith complies with demands from the warden.
S5	Smith files taxes for the warden.
S6	Smith embezzles money for the warden.
S7	Smith makes friends with some inmates.
S8	Smith learns his friend has some exoneration evidence.
S9	Smith requests help with his exoneration from the warden.
S10 (χ)	The warden denies the request to be exonerated.
S11	Smith retrieves the escape plan.
S12	Smith takes a spoon from the cafeteria.
S13	Smith hangs artwork in his cell.
S14	Smith digs a tunnel with the spoon behind the artwork.
S15	Smith escapes through the tunnel.

allow variants that contain a different number of achieved and revised intentions. The variants are π_{Jail}^{C1}, π_{Jail}^{C2} and π_{Jail}^{T1}, π_{Jail}^{T2} and are represented in Figure 4.

Finally, we designed a second planning problem to account for alternative intention revision structure. In π_{Jail} treatment groups, non-player character agents revise their intentions and subsequently achieve the new revised intention. Conversely, in our second intention revision structure, *"Wilderness"* (π_{Wild}), the character agent must revise their new intention three times before they finally achieve their final intention. The difference in structure of both stories is captured in Fig. 4.

Our experiments focus on the causal and intentional structure of intention revision, therefore we desire to minimize discourse effects (such as re-ordering steps for dramatic effect). After linearizing the partial orderings of a plan's steps, we apply a simple template,$\langle agent1, action, agent2\rangle$, that generates a single story sentence for each plan step. Templates are perhaps the most basic form of natural language generation as they omit intermediate representations (e.g., sentence plans) used by more robust systems (Reiter & Dale, 1997). Instead, template methods map non-linguistic input (e.g., plan steps) directly to text structure (e.g., story statements). An example of the output from this approach is in Table 3.

Our experiment was designed to answer the four following research questions:

R1. How does plan-based intention revision affect narrative transportation and perceived interest?

R2. How does plan length affect narrative transportation and perceived interest?

R3. How does plan length interact with the number of intention revisions to affect narrative transportation and perceived interest?

R4. How does the story structure affect narrative transportation and perceived interest?

The answers to these questions will provide insights into how appropriate it is to use intention revision as a branching mechanism for interactive narratives.

Subjects. Data were collected from 664 participants who were recruited on the CrowdFlower crowd sourcing platform and paid $0.40 USD to read a narrative text and complete the narrative transportation scale and the perceived interest questionnaire. Results of a power analysis indicated that the sample size was adequate for finding our desired effects at 0.05 significance level.

Procedures. Upon logging into the online experiment, participants were randomly assigned into either the control or treatment groups, and then assigned to read either the Jail or Wilderness story. The control group read a variation of the story where the character agents achieved every intention they committed to. Conversely, the treatment group read stories where the character agent's original intention was made impossible and had to revise their intentions. In addition, participants were randomly assigned to read a story length of either 15, 25, or 35 statements.

Participants began the experiment by completing a brief demographic survey and then read their assigned story, which was presented as a single page of text. After reading their respective stories, participants completed the fourteen item transportation scale and the twelve item perceived interest questionnaire.

We hypothesize that when subjects read a story statement that introduces the possibility of an intention revision, it will compel them to think divergently about a story character's future intentions. This activity requires recalling story facts and reasoning over them, increasing the degree to which they become involved and interested with the plot and characters. Both the transportation scale and perceived interest questionnaire are designed to measure such effects and are widely used with high degree of success (Section 2.3).

4.2 Experimental Results

Prior to computing inferential statistical analyses, data were inspected and coded to identify missing values, to identify outliers and to ensure that the data were normally distributed. No extreme values were identified. Research questions were tested using a 2 (group: treatment vs. control) x 2 (story structure: Jail vs. Wilderness) x 3 (story length: 15, 25, 35 statements) between subjects ANOVA. Unless otherwise noted, all analyses were computed using a critical alpha value of $p = .05$. For all data, Levenes tests confirmed homogeneity of variance. Means and standard deviations for the different conditions are presented in Table 4.

R1. How does plan-based intention revision affect narrative transportation and perceived interest?

Narrative Transportation. Our first research question asked whether plan-based intention affected perceptions of narrative transportation. To address this question we compared narrative transportation scores between participants in the treatment and control conditions. Results of our omnibus 2x2x3 ANOVA showed a main effect for group $F(1, 652) = 10.74$, $p < .01$ where individuals in the treatment condition ($M = 50.70$; $SD = 12.81$) who read narrative texts that included intention revisions reported higher levels of narrative transportation compared to individuals in the control group ($M = 47.43$; $SD = 12.92$).

Table 4. The table below contains the mean Transportation Scale (TS) and Perceived Interest Questionnaire (PIQ) scores, with variance in parentheses, for the control (C) and treatment (T) experimental groups of each Jail and Wilderness story variant.

		Jail				Wilderness			
		15 statements	25 statements	35 statements	Total	15 statements	25 statements	35 statements	Total
TS	C	46.33 (10.54)	48.84 (14.06)	52.62 (13.25)	49.27 (12.92)	44.53 (12.09)	46.26 (13.75)	46.48 (12.58)	45.70 (12.72)
	T	48.62 (11.65)	53.46 (12.50)	54.66 (11.91)	52.21 (12.24)	45.72 (12.32)	49.04 (12.16)	52.86 (14.33)	49.17 (13.23)
PIQ	C	27.49 (8.25)	26.56 (10.42)	28.35 (10.08)	27.45 (9.62)	24.50 (8.19)	26.96 (9.67)	26.61 (8.60)	25.93 (8.80)
	T	28.91 (8.77)	31.63 (7.74)	30.11 (8.33)	30.20 (8.32)	24.72 (8.72)	26.70 (9.43)	29.02 (9.42)	26.79 (9.30)

Perceived Interest. Regarding perceived interest, a significant main effect of group $F(1, 652) = 6.43$, $p < .05$ was observed that showed individuals in the treatment condition ($M = 28.51$; $SD = 8.97$) who read narrative texts that included intention revisions reported higher levels of perceived interest compared to individuals in the control group ($M = 26.67$; $SD = 9.22$) who read narrative texts that did not include intention revisions.

R2. How does plan length affect narrative transportation and perceived interest?

Narrative Transportation. Our second research question asked whether story length impacted perceptions of narrative transportation. To address this question we compared mean narrative transportation scores between participants who read stories that were either 15, 25, or 35 statements long. Results showed a significant main effect of story length F(2, 652) = 10.14, p < .001). Post hoc analyses showed that perceptions of narrative transportation increased significantly as story length increased such that participants who read the 35 statement narrative text ($M = 51.64$, $SD = 13.33$) and 25 statement text ($M = 49.54$; $SD = 13.29$) reported higher levels of narrative transportation compared to the participants who read the 15 statement text ($M = 46.27$, $SD = 11.73$).

Perceived Interest. Regarding perceived interest, a significant main effect of plan length was observed $F(2, 652) = 3.34$, $p < .05$. Post hoc analyses showed that participants who read the 35 statement narrative text ($M = 28.52$, $SD = 9.14$) reported higher levels of perceived interest in the story compared to the participants who read the 15 statement text ($M = 26.34$, $SD = 8.64$).

R3. How does plan length interact with the number of intention revisions to affect narrative transportation and perceived interest?

Narrative Transportation. Our third research question asked if intention revisions significantly affected narrative transportation and if this relation depended on story length. To address this question we examined the interaction between group and story length. Results failed to show a significant interaction $F(2, 652) = .60$, *n.s.* Inspection of the results shows that means trended in the expected direction such that narrative transportation grew stronger for the treatment group compared to the control group as story length increased.

Perceived Interest. We conducted a similar analysis to determine whether group and story length interacted to impact perceived interest. Results of this analysis revealed there was no significant interaction between group and story length $F(2, 652) = .48$, *non-significant*.

R4. How does the story structure affect narrative transportation and perceived interest?

Narrative Transportation. To address the last research question we computed a planned comparison that examined whether story structure affected perceptions of narrative transportation. Participants in the control condition were excluded from this analysis. Results showed that participants who read the *Jail* story with intention revisions reported higher levels of narrative transportation (M = 52.21; SD = 12.24) than participants who read the *Wilderness* story with intention revisions (M = 49.17; SD = 13.23), $F(1, 332) = 4.99$, $p < .05$.

Perceived Interest. Regarding the impact of IR style on perceived interest, a planned comparison showed participants who read the *Jail* text with intention revisions reported higher levels of perceived interest (M = 30.20; SD = 8.32) than participants who read the *Wilderness* text (M = 26.79; SD = 9.30), $F(1, 332) = 12.75$, $p < .001$.

5. Discussion and Conclusions

Creating a story branch in response to user actions is an attractive feature of interactive narrative systems. Formalizing plan-based intention revision and using it as a story branching mechanism would ensure that non-player agents make believable behavior changes. Using this formalization, we investigated the long held assumption that branching stories improve user experience by evaluating the effects generated branching stories have on widely-cited instruments, the transportation scale and perceived interest questionnaire.

To answer our four research questions (R1-4, Section 4.1), we generated a set of branching stories by simulating player actions that disrupt (non-player) character agent intentions who then revise their intentions. After translating this plan-set into text, human subjects were assigned to conditions and read stories that either included or excluded intention revisions, varied in length, and varied in story structure. Subjects then completed the transportation scale and perceived interest questionnaire. We answered R1 by showing that the treatment group reported significantly higher narrative transportation and perceived interest scores than the control group across two story structures. In response to R2, our analysis showed that perceptions of narrative transportation and perceived interest increased linearly as story length increased. Our analysis for R3 showed that there was no interaction effect between group and story length for either narrative transportation or perceived interest, though the results suggested narrative transportation grew stronger for longer stories that included intention revisions. A complicating factor in showing the interaction effect was we inadvertently introduced an agent intention that could be characterized as revenge, increasing transportation in our control groups. Finally, to address R4 we found that the Jail story, in which the story structure allowed the main character to achieve multiple goals led to higher perceptions of narrative transportation and perceived interest. Overall, we conclude that using intention revision as a story branching mechanism is an effective strategy as it both continues a story when it would otherwise end and actually *increases* subjective experience.

When considering alternative explanations for the observed effect, we consider it plausible that the results come from overcoming greater obstacles. That is, an agent's revised intentions contain goals that subjects perceive as lower probability or difficult to achieve. These goals would cast the story's outcome in doubt, introducing adversity to the agent, and increasing subject engagement without revising intentions. To experimentally distill the relationship between this type of goal and

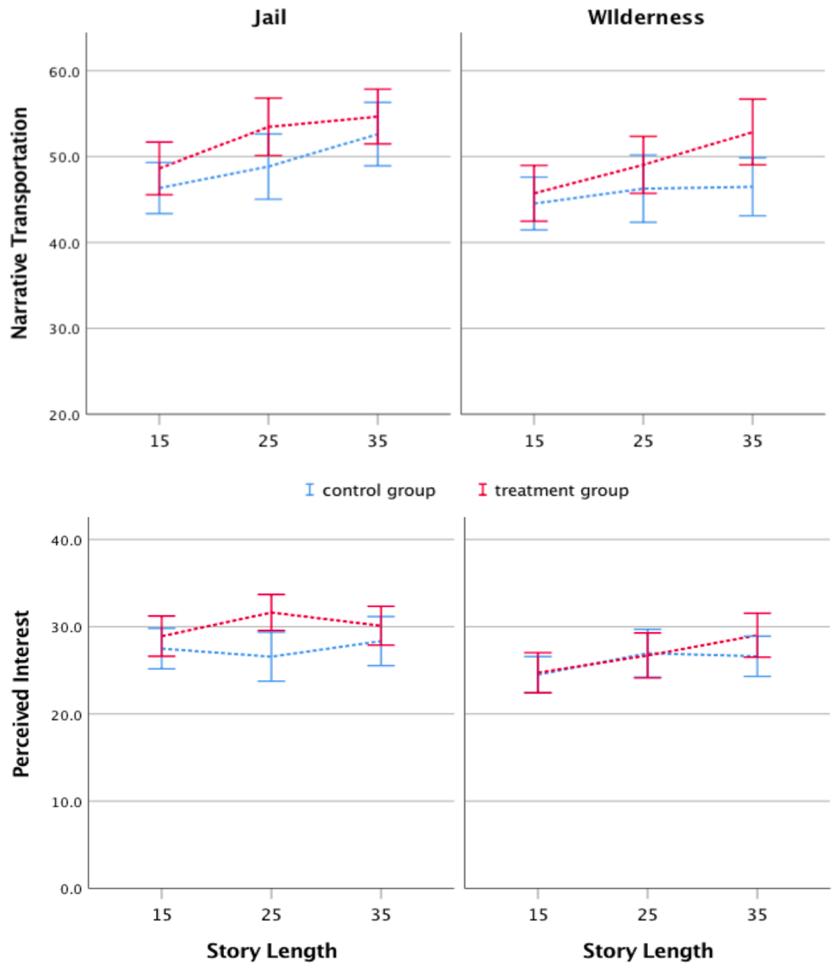

Figure 5. The mean narrative transportation and perceived interest scores of each experimental group for both story structures.

intention revision, we could create two control groups. One where it is explicitly stated that the goal is hard and adversity will ensue. The second would state the goals are easy to achieve and no adversity is expected. These control groups' results can be compared to a treatment group where the difficulty is not stated to subjects and from another treatment group where intention revision is introduced. Ideal results would show no difference between the control groups and the experimental groups where difficulty is not stated, confirming that the semantics of the goals are perceived as intended. If there is a measurable effect between adversity and intention revision we would expect to observe a difference in transportation and perceived interest scores between the control group with adversity and intention revision groups.

There are two immediate areas for future work. The first is to evaluate an interactive narrative that allows users to choose actions instead of simulating exceptional actions. This enables us to measure interaction effects of intention revision and an interactive interface have on the transportation scale and perceived interest questionnaire. The second is a plan-based model of revenge. We observed an increase in transportation in both the control and treatment conditions when revenge was inadvertently introduced. We view revenge as a more specific type of intention revision and formalizing it would allow further experimentation. Lastly, plan-based intention revision enables a longer term goal of integrating intentional agent behavior into cognitive architectures that can manage goal lifecycles (e.g., Cox et al., 2017).

Acknowledgements

This material is based upon work supported in whole or in part with funding from the Laboratory for Analytic Sciences. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Laboratory for Analytic Sciences and/or any agency or entity of the United States Government.

References

Amos-Binks, A. (2017). Problem formulation for accommodation support in plan-based interactive narratives. *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*. San Francisco, CA: AAAI Press.

Amos-Binks, A., & Young, R. M. (2018). Plan-based intention revision. *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*. New Orleans, LA: AAAI Press.

Bates, J. (1992). Virtual reality, art, and entertainment. *Presence: Teleoperators and Virtual Environments*, *1*, 133–138.

Bratman, M. (1987). *Intention, plans, and practical reason*. Palo Alto, CA: Center for the Study of Language and Information.

Cardona-Rivera, R. E., Price, T. W., Winer, D. R., & Young, R. M. (2016). Question answering in the context of stories generated by computers. *Advances in Cognitive Systems*, *4*.

Cohen, P. R., & Levesque, H. J. (1990). Intention is choice with commitment. *Artificial Intelligence*, *42*, 213–261.

Cox, M. T., Dannenhauer, D., & Kondrakunta, S. (2017). Goal operations for cognitive systems. *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*. San Francisco, CA: AAAI Press.

Fendt, M. W., & Young, R. M. (2017). Leveraging intention revision in narrative planning to create suspenseful stories. *Transactions on Computational Intelligence and AI in Games*, *9*, 381–392.

Gerrig, R. J. (1993). *Experiencing narrative worlds*. New York, NY: Routledge.

Green, M. C., & Brock, T. C. (2000). The role of transportation in the persuasiveness of public narratives. *Journal of Personality and Social Psychology*, *79*, 701–721.

Hidi, S., & Renninger, K. A. (2006). The four-phase model of interest development. *Educational Psychologist*, *41*, 111–127.

Katsuno, H., & Mendelzon, A. O. (1991). On the difference between updating a knowledge base and revising it. *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*. Cambridge, MA: Morgan Kaufmann.

Meehan, J. R. (1977). TALE-SPIN: An interactive program that writes stories. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*. Cambridge, MA: Morgan Kaufmann.

Penberthy, J. S., & Weld, D. (1992). UCPOP: A sound, complete, partial order planner for ADL. *Proceedings of the International Conference on Knowledge Representation and Reasoning*. Cambridge, MA: Morgan Kaufmann.

Porteous, J., Cavazza, M., & Charles, F. (2010). Applying planning to interactive storytelling: Narrative control using state constraints. *Transactions on Intelligent Systems and Technology*, *1*, 1–21.

Pérez y Pérez, R., & Sharples, M. (2001). MEXICA: A computer model of a cognitive account of creative writing. *Journal of Experimental & Theoretical Artificial Intelligence*, *13*, 119–139.

Ramirez, A., & Bulitko, V. (2014). Automated planning and player modelling for interactive storytelling. *Transactions on Computational Intelligence and AI in Games*, *7*, 375–386.

Rao, A. S., & Georgeff, M. P. (1998). Decision procedures for BDI logics. *Journal of Logic and Computation*, *8*, 21–26.

Reiter, E., & Dale, R. (1997). Building applied natural language generation systems. *Natural Language Engineering*, *3*, 57–87.

Riedl, M. O., & Young, R. M. (2010). Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research*, *39*, 217–268.

Robertson, J., & Young, R. M. (2014). Gameplay as on-line mediation search. *Proceedings of the Tenth Artificial Intelligence and Interactive Digital Entertainment Conference EXAG Workshop*. Raleigh, NC: AAAI Press.

Schank, R. (1986). *Explanation patterns: Understanding mechanically and creatively*. Hillsdale, NJ: Lawrence Erlbaum.

Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals and understanding: An inquiry into human knowledge structures.*. Oxford, England: Lawrence Erlbaum.

Schraw, G. (1997). Situational interest in literary text. *Contemporary Educational Psychology*, *22*, 436–456.

Van Der Hoek, W., Jamroga, W., & Wooldridge, M. (2007). Towards a theory of intention revision. *Synthese*, *155*, 265–290.

Vermeulen, I. E., Roth, C., Vorderer, P., & Klimmt, C. (2010). Measuring user responses to interactive stories: Towards a standardized assessment tool. *Proceedings of the Third Joint International Conference on Interactive Digital Storytelling*. Edinburgh, Scotland: Springer.

Ware, S. G., & Young, R. M. (2014). Glaive: A state-space narrative planner supporting intentionality and conflict. *Proceedings of the Tenth Artificial Intelligence and Interactive Digital*

Entertainment Conference. Raleigh, NC: AAAI Press.

Ware, S. G., Young, R. M., Harrison, B., & Roberts, D. L. (2014). A computational model of plan-based narrative conflict at the fabula level. *Transactions on Computational Intelligence and AI in Games*, *6*, 271–288.

Witmer, B. G., & Singer, M. J. (1998). Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and Virtual Environments*, *7*, 225–240.