Proceedings, The Thirteenth AAAI Conference on
Artificial Intelligence and Interactive Digital Entertainment (AIIDE-17)

A Possible Worlds Model of Belief for State-Space Narrative Planning

Alireza Shirvani, Stephen G. Ware, Rachelyn Farrell
Narrative Intelligence Lab, Department of Computer Science,
University of New Orleans, New Orleans, LA 70148, USA
{ashirvan, sgware, rfarrell}@uno.edu

Abstract

What characters believe, how they act based on those beliefs, and how their beliefs are updated is an essential element of many stories. State-space narrative planning algorithms treat their search spaces like a set of temporally possible worlds. We present an extension that models character beliefs as epistemically possible worlds and describe how such a space is generated. We also present the results of an experiment which demonstrates that the model meets the expectations of a human audience.

Introduction

Narrative planning (Young et al. 2013) has proven a valuable tool for generating and adapting stories in interactive virtual environments. They have the potential to manage story spaces too large for human authors to anticipate entirely at design time, and can adapt the narrative to each individual user. Several recent narrative planners (Teutenberg and Porteous 2013; Ware and Young 2014) have used state-space search so as to incorporate recent advances in heuristic search from the classical planning community.

Riedl and Bulitko (2013) call narrative planners *strong story* systems because they focus on the author's requirements then find explanations for character actions. In contrast, *strong autonomy* systems focus on rich agent simulation then coordinate agents to ensure the author's constraints. Strong story systems may be preferable in some situations, but to match the richness of strong autonomy systems they must include a model of agent beliefs. How agents act based on their (possibly wrong) beliefs and how their beliefs change is an essential part of many stories.

We present a model of agent beliefs suitable for strong-story state-space narrative planners. We treat the search space of the problem as a map of temporally possible worlds, and to this we add epistemically possible worlds to represent what agents believe. We describe how this space is generated and present an evaluation to demonstrate that it matches the expectations of a human audience.

Related Work

Research on agent beliefs is far too vast to survey completely here, so we divide it into groups in an attempt to motivate the unique challenge faced by strong-story narrative planners.

Classical planning assumes a single agent in a fully observable environment. Ample prior work exists on modeling beliefs in real world, multi-agent, partially observable environments (e.g. Petrick and Bacchus 2002). However, narrative planning seeks only the *illusion* of reality, and is free to exploit the omniscience and omnipotence of the author in the virtual world while generating the story (Ware and Young 2010). Real world planners tend to assume that agents are all cooperative (Grosz and Kraus 1996) or all competitive (De Rosis et al. 2003), that ignorance or wrong beliefs are bad (Bolander and Andersen 2011), or that the planner's job is to find a solution that works despite wrong beliefs (Hoffmann and Brafman 2006).

Likewise, ample previous work exists for modeling agent beliefs in strong autonomy systems (Bates, Loyall, and Reilly 1992, Pynadath and Marsella 2005, Ryan et al. 2015, and many more). Though not directly applicable to strong story systems, the theories they are based on influenced our model, especially doxastic model logic (Linsky 1968), Kripke structures (1963), truth maintenance systems (Doyle 1979), and beliefs in situation calculus (Demolombe and del Pilar Pozos Parra 2010).

Various narrative planners have modeled some aspect of belief, either of the audience or characters, to achieve specific outcomes like suspense (Cheong and Young 2015), mediation (Robertson and Young 2015) or deception (Christian and Young 2004). These models are typically in service of some other phenomena and not intended as general solutions to the belief problem.

Teutenberg and Porteous (2015) implement a general model of character belief in their IM-PRACTical planner. They limit recursive beliefs to 1 level; the planner can reason about what is true (0 levels) and what x believes (1 level), but to reason about what x believes y believes or deeper (2+ levels) the planner relies on a single state called the *popular belief* shared by all agents. Brinke, Linssen, and Theune's Virtual Storyteller (2014) is also limited to 1 level. While this is sufficient for most stories, we seek a model which supports arbitrary recursion, where the size of the search space grows proportionately to the amount of recursion used.

Beliefs across Possible Worlds

As an example we use a version of *Treasure Island* (Stevenson 1919) in Figure 1. Jim Hawkins acquires a treasure map but must enlist the help of the pirate Long John Silver to to get it. Silver believes the treasure does not exist, but Hawkins spreads a rumor causing him to believe it is on the island. Silver takes Hawkins to the island in the hopes of claiming the treasure for himself, but Hawkins digs it up and makes off with it.

Classical and Narrative Planning

A classical planning problem (Russell and Norvig 2010) is defined as $\langle s_0, g, A\rangle$, where s_0 is the initial state, g is a goal expression, and A is a set of actions. For every action $a \in A$, $\textsc{pre}(a)$ denotes the preconditions of a, propositions which must be true immediately before the action is taken, and $\textsc{eff}(a)$ denotes the effects of a, propositions which becomes true after the action. A valid plan is a sequence of ground actions which achieves g. We assume Helmert's multi-valued planning task representation (2006), though our model is also compatible with traditional predicate logic.

The state space for this problem is a graph whose nodes are states and whose edges are actions. An edge $n_1 \xrightarrow{a} n_2$ exists from node n_1 to node n_2 via action a if the preconditions of a are met in n_1 and applying the effects of a to n_1 results in n_2. Considered as a narrative space, this graph represents what narratologists refer to as the set of possible future worlds (Ryan 1991; Bruner 1986), so we refer to these edges as *temporal edges*, drawn in black in Figure 1 (the states on the left column are the actual world states). A state-space planner begins with a graph composed of only the node for s_0 and, during search, expands the space by adding temporal edges until reaching a state where g holds.

We define a narrative planning problem as $\langle s_0, g, A, C\rangle$, where s_0, g, and A are defined as above, and C is a set of characters, special constants which represent agents that can have intentions and beliefs. For every action $a \in A$, we define $\textsc{con}(a)$ to be the set of characters $\in C$ who must consent to take that action. We define $\textsc{obs}(a)$ to be the set of characters $\in C$ who observe the action when it occurs[1].

We define the state space of this problem as above, but with the addition of *epistemic edges*. An epistemic edge $n_1 \xrightarrow{c} n_2$ exists from node n_1 to node n_2 (possibly the same node) for every character $c \in C$ and means that when the world is in state n_1 character c believes that world to be in state n_2. Epistemic edges are drawn in red in Figure 1.

The modal predicate $b(c, p)$ represents whether character c believes proposition p. To evaluate $b(c, p)$ at some node n, follow the epistemic edge from n for character c and then evaluate p. Consider n_2 in Figure 1, after Hawkins spreads the rumor. TB means the treasure is buried. The expression $b(H, b(S, TB))$ means "Hawkins believes that Silver believes the treasure is buried." To evaluate $b(H, b(S, TB))$ at n_2, we follow the edge $n_2 \xrightarrow{H} n_2$ to find what Hawkins believes, then $n_2 \xrightarrow{S} n_3$ to find what Hawkins believes Silver believes, and we see that TB is true; Hawkins believes that Silver believes the treasure is buried.

[1]Methods for choosing $\textsc{obs}(a)$ have been given by Christian and Young (2004), Brinke, Linssen, and Theune (2014), Teutenberg and Porteous (2015), and others. Our model makes no particular commitment to how $\textsc{obs}(a)$ is chosen.

Adding beliefs introduces two changes to the classical formalism. First, a state is no longer uniquely identified by its fluent values. n_0 and n_2 have the same fluent values, but different beliefs, specifically in n_0 we see $b(S, TN)$ but in n_2 we see $b(S, TB)$. Second, it is sometimes possible for an action to occur when a character wrongly believes its preconditions are not met. We call this kind of temporal edge a *surprise edge*. In n_6 Silver wrongly believes the treasure does not exist, so he is surprised to observe Hawkins digs it up $\left(n_6 \xrightarrow{dig} n_9\right)$, an action Silver believed impossible.

Expanding the Narrative Space

We now describe how the narrative space is expanded. The following procedures allow any node to be expanded at any time, which facilitates search algorithms that expand only those parts of the space that are needed to generate a solution (a focus for our future work).

We use the notation $\alpha(a, n)$ to mean the node reached by following the temporal edge for action a from node n. $\alpha(a, n)$ means "the state after action a is taken in n." In Figure 1, $\alpha(rumor, n_0) = n_2$. We use the notation $\beta(c, n)$ to mean the node reached by following an epistemic edge for character c from node n. $\beta(c, n)$ means "when in state n, c believes the state is $\beta(c, n)$." In Figure 1, $\beta(S, n_0) = n_1$.

Initialization The space begins with n_0, which reflects the initial values of all fluents in the actual world (i.e. the omniscient author's beliefs). We use the classical *closed world assumption*, which says anything not explicitly stated true is assumed false. We also assume that for any proposition p:

$$\forall c \in C\, (b(c, p) \rightarrow \forall d \in C\, (b(c, \beta(d, p))))$$

This means that, unless explicitly stated otherwise by the author, characters believe other characters believe the same things they do. The author states that Silver believes the treasure does not exist, so Silver assumes Hawkins believes the same. Note that multiple nodes may be required to represent the initial state (in Figure 1, these are n_0 and n_1).

Expanding Nodes New nodes in the space are generated via Algorithm 1, which describes where epistemic edges should point when adding a new temporal edge. Consider adding the temporal edge $n_4 \xrightarrow{dig} n_7$. Assume only nodes $n_0 ... n_6$ have been generated so far, that the search is currently visiting n_4, and that it wants to add a temporal edge from n_4 for the action dig—that is, we want to compute $\alpha(dig, n_4)$. The edge does not yet exist, so n_7 is created to be the head of this new edge (Algorithm 1, line 3). Initially each epistemic edge for n_7 points to the node the edge pointed to in the parent node; that is, $\beta(H, n_7) = n_4$ and $\beta(S, n_7) = n_5$ (line 4). Then, for every character who observes dig (both H and S), their epistemic edges are updated to point to the state that would result from taking the dig

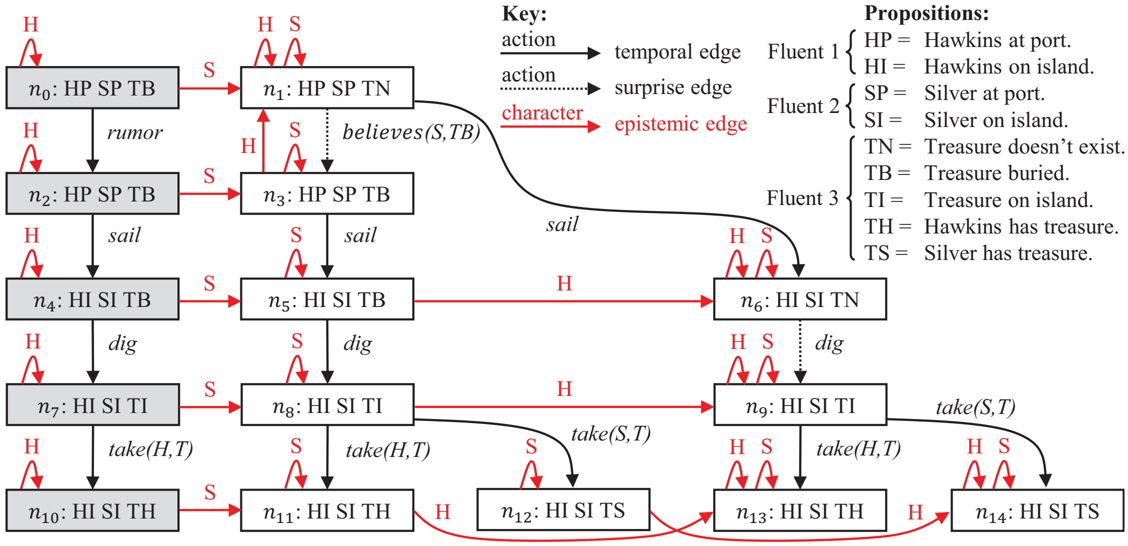

Initial State:	
HP SP TB b(S,TN)	
Author's Goal:	TH
Hawkins' Goal:	TH
Silver's Goal:	TS

Actions:	***rumor***	***believes(S,TB)***	***sail***	***dig***	***take(H,T)***	***take(S,T)***
PRE:	TB	*false*	HP SP	HI TB	HI TI	SI TI
EFF:	b(S,TB)	TB	HI SI	TI	TH	TS
CON:	H	*none*	S	H	H	S
OBS:	H	S	H S	H S	H S	H S

Figure 1: A narrative search space graph for the plot of *Treasure Island* (Stevenson 1919).

1: **procedure** EXPAND(a, n) ▷ To compute $\alpha(a, n)$.
2: **if** $\alpha(a, n)$ does not exist **then**
3: Let n_* be a new node.
4: $\forall c \in C$ let $\beta(c, n_*) = \beta(c, n)$.
5: $\forall o \in$ OBS(a) let $\beta(o, n_*) = \alpha(a, \beta(o, n))$.
6: Let $\alpha(a, n) = n_*$.
7: Impose EFF(a) on n_*.
8: **end if**
9: **end procedure**

Algorithm 1: Adding a temporal edge to the search space.

action in the state those characters believe the world is currently in (line 5). This requires computing $\alpha(dig, n_5)$ and $\alpha(dig, n_6)$ recursively, which causes nodes n_8 and n_9 to be created. Finally, the effects of dig are imposed on n_7, so the fluent TB is changed to TI (line 6).

Imposing Effects Two special cases may arise when imposing the effects of an action: imposing a surprise action and imposing belief effects.

When imposing a surprise action, such as when adding edge $n_6 \xrightarrow{dig} n_9$, we first impose the preconditions of the action (except the constant $false$, which cannot be imposed) and then impose the effects. When a character observes an action they believed was impossible, they first update their beliefs to allow the action to be possible and then update their beliefs based on the action's effects.

Imposing belief effects is done through dummy actions. To impose the belief $b(c, p)$ for character c about proposition p on some node n, we create a dummy action d such that PRE$(d) = false$, EFF$(d) = p$, CON$(d) = \emptyset$, and OBS$(d) = c$. This action is then imposed on node $\beta(c, n)$ via Algorithm 1, and $\beta(c, n) = \alpha(d, \beta(c, n))$. Consider adding the edge $n_0 \xrightarrow{rumor} n_2$, assuming only nodes n_0 and n_1 have been generated so far. Adding the edge creates n_2. To impose the effect $b(S, TB)$ on n_2, we create the dummy action $believes(S, TB)$ and impose it on $\beta(S, n_2)$, which at this time is n_1. This causes node n_3 to be created, and the epistemic edge $\beta(S, n_2)$ is updated to point to n_3. In other words, after Hawkins spreads the rumor about the treasure, Silver will change his belief from TN to TB, and Hawkins will know that. In Figure 1, the dummy action $believes(S, TB)$ is not part of the authored domain; it is the dummy action generated to impose $b(S, TB)$, and its details are only shown in Figure 1 for clarity. Note that these dummy belief update actions are always surprise actions because their precondition is $false$.

Valid Narrative Plans

For a classical planner, any path through the search space from s_0 to a state where g holds is a valid plan. Narrative planning imposes additional constraints on the solution—every character must believe that every action they take will contribute to one of their goals. We ensure this using the causal model of intentionality first defined by Riedl and Young (2010) then adapted for state-space planning by Ware and Young (2014). We reproduce only those definitions needed to explain how belief affects the model.

We say two actions a_1 and a_2 are *causally linked* via proposition p if $p \in$ EFF(a_1), $p \in$ PRE(a_2), and no action occurring between a_1 and a_2 has an effect which negates p[2]. Traditionally, a_2 must be a temporal descendant of (i.e. after) a_1, but we allow causal links to traverse epistemic edges also. When the path from a_1 to a_2 crosses an epistemic edge for character c, every precondition q of subsequent actions should be treated as $b(c, q)$. This means the *rumor* action from $n_0 \to n_2$ is causally linked to *dig* from $n_5 \to n_8$ because *rumor* has effect $b(S, TB)$ and all paths from n_2 to n_5 cross an epistemic edge for S, so *dig* is treated as having precondition $b(S, TB)$. In other words, Hawkins spreading the rumor that the treasure was buried enabled Silver to believe *dig* would eventually be possible.

A *causal chain* is an alternating sequence of actions and propositions $\langle a_1, p_1, ..., a_n, p_n \rangle$ such that for all i from 1 to $n-1$, action a_i is causally linked to a_{i+1} via p_i. An *intentional chain* for character c to achieve character goal g_c is a causal chain $\langle a_1, p_1, ..., a_n, g_c \rangle$ such that c intends g_c before a_1 and until a_n, no action is a surprise action, no proposition is repeated, and a no proposition appears in the chain with a proposition that negates it. This is based on Ware and Young's definition, but with one important relaxation—we do not require that every action in an intentional chain have character c as a consenting character.

This modified definition allows us to directly adopt Ware and Young's concept of an explained action. An action a is *explained* iff, for every consenting character $c \in \text{CON}(a)$, a is on an intentional chain for c and every other action in that chain is also explained. A *valid narrative plan* is one that achieves the author's goal and for which all actions are explained.

Our relaxed definition of an intentional chain allows one agent to anticipate the actions of another. Consider the path $n_3 \to n_5 \to n_8 \to n_{12}$, which has this intentional chain:

$$S \text{ achieves } TS \text{ via } \langle sail, HI, dig, TI, take(T, S), TS \rangle$$

Silver intends to sail to the island and let Hawkins dig up the treasure, then Silver will take it for himself. Here we see Silver anticipate what Hawkins will do and incorporate that into his plan, even though Silver is not a consenting character for *dig*. To summarize, you can include another agent's action in your plan if you can explain why that agent would take that action. In this case, Silver expects Hawkins will *dig* because he knows Hawkins also wants the treasure.

Surprise actions cannot be in intentional chains because characters cannot anticipate them. The path $n_1 \to n_6 \to n_9 \to n_{14}$ *does not* produce the above intentional chain because *dig* is a surprise action which Silver did not expect to be possible. Only after *rumor* changes his beliefs can he imagine a plan to get the treasure.

This concept of anticipated actions is used to compute the actual plan (the story that is narrated) by considering the author as an omniscient character who believes n_0. When the author can anticipate an action, the author believes it will be possible (and since the author's beliefs are accurate, it will be), and it must have an explanation for all character who consent, therefore it should appear believable to the audience (Riedl and Young 2010). In Figure 1 the actual plan is $n_0 \to n_2 \to n_4 \to n_7 \to n_{10}$, but other nodes like n_{12} represent possible worlds that had to be expanded to explain, for example, why Silver consented to *sail* (because he planned to take the treasure, though this never actually happens).

Evaluation

We claim that our model has two important advantages. First, it accurately simulates how people process beliefs and second, it does so more accurately than models lacking nested beliefs. We demonstrate these features using an empirical human subject evaluation.

We compare our model of belief with two others, one without the notion of belief and one with a single layer. We say the former has 0 layers of belief, where all agents are considered omniscient. Classical planners call this the *full observability* assumption. The latter generalizes agents' belief of others' by a shared state among all agents. For IMPRACTical (Teutenberg and Porteous 2015), this shared state is called the popular belief, which is updated based on public actions. We call this the 1+1 layer of belief model. Our model with 2 or more layers of belief is denoted 2+.

We also sketch a proof that our 2+ model generates a superset of the stories produced by these other models and the resulting stories make sense to human readers.

Empirical Evaluation

To evaluate the 2+ model, we used it to produce three stories representing different narrative elements including deception, cooperation, anticipation, and surprise. The stories were presented to human subjects who answered questions about agents' beliefs. These answers were used to measure the accuracy of each model.

Stories We generated each story as a plan by hand using the 2+ model. Each plan was translated into text by converting its actions into sentences using a simple natural language template for the actions. Figure 2 presents summaries of the stories. These three stories were specifically selected to depict the benefits of epistemic edges.

In *Homecoming*, Mike forms a plan to meet Jenny at the school because he believes that Jenny believes that a party is happening at the school and he expects her to go there. Using the 0 model, everyone is aware of the actual location of the party, so invitations cannot be explained. Using 1+1, Mike cannot reason about what he believes that Jenny believes, since inviting someone to a party is not a public action and does not update the popular belief.

In *The Forty Thieves*, the thieves use their mark to indicate the location of the treasure. Alibaba anticipates their attack and takes advantage of their belief by marking the guardhouse. Using the 0 model, the thieves directly attack Alibaba's house. Using 1+1, marking Alibaba's door is not a public action and thus Alibaba cannot anticipate thieves' attacking the marked house.

Finally, in *The Most Wanted*, Sheriff William infers that to arrest Jack, he should trick him into robbing the wagon.

[2] In a multi-valued planning task, the concept of negation is complicated. See Helmert's (2006) discussion.

Expecting this, he hides in the wagon to surprise and arrest Jack. Using the 0 model, Jack always knows where the money is and Sheriff William's hiding can't be explained. This story can also be generated using 1+1, similar to Teutenberg and Porteous's Othello example (2015).

Experiment Design An interactive webpage was presented to participants via Amazon's Mechanical Turk crowd-sourcing platform. One story was presented to each participant. The page consisted of an introduction, the story and belief questions, and finally five comprehension questions after the story was removed from the screen. Figure 3 shows an image of the page asking readers what Sheriff William believes that Jack believes after he has told Jack about the money.

The three stories are divided into sequences of simple steps (temporal edges), and each sequence is presented to the participant one step at a time. Figure 2 summarizes the steps of each sequence in a numbered list. Before the first (the initial state) and after each sequence, participants are asked multiple-choice questions about the agents' beliefs for 1-2 layers. Figure 2 shows the questions of each story. All questions had the same answer options, and options were arranged randomly.

Mechanical Turk provides notoriously noisy data, and this study was especially difficult given that it asks people to think about thinking. Therefore, we only considered a response valid if the participant correctly answered all comprehension questions. For *Homecoming*, *The Forty Thieves*, and *The Most Wanted*, we collected 21, 23, and 70 sets of valid responses respectively.

Results The results were first analyzed to investigate participant agreement. Using Krippendorff's α (Krippendorff 2012), inter-rater reliability for *Homecoming* was $\alpha = 0.79$, and for *The Forty Thieves* $\alpha = 0.70$, representing moderately high agreement. For *The Most Wanted*, $\alpha = 0.20$, and while any positive value represents agreement, this story showed low agreement.

Next, for each question, we determine which answer is considered to be correct by human subjects. The binomial exact test (Howell 2012) with Bonferroni correction (Holm 1979) was used to investigate if participants significantly agreed on an answer for each question. We say participants agreed when $p \leq 0.05$. Consequently, 4 questions (about 6%) were removed from further analysis because participants did not agree on a right answer (all with p-values of 1.00). Note that, using this method, it is possible for two answers to be significant if subjects were divided about equally between them. This occurred for 1 question in *The Most Wanted*, so we considered either answer correct.

The accuracy of each model is calculated as the number of questions that model answers the same as human subjects divided by the total of number of questions (which was 62). The contingency tables that allow us to compare our model to 0 and 1+1 are given in Table 1. The 2+ model achieves the highest accuracy for all stories, with 95% accuracy overall, whereas the 1+1 and 0 models achieved 54% and 40% respectively. Also note that the 1+1 model was not strictly better than the 0 model.

Homecoming (24 questions)

1. Ed finds out about a party at school and invites his friend Mike.
2. Mike invites Jenny to the party because he has a crush on her.
3. Ed discovers the party is actually at a fraternity house.
4. Ed calls Mike to inform him the party is at the fraternity house.
5. Mike does not have Jenny's phone number, so he goes to school, waits for Jenny to arrive, and then informs her the party is at the fraternity house.

Where does Mike believe the party is?
Where does Jenny believe the party is?
Where does Mike believe Jenny believes the party is?
Where does Ed believe Mike believes the party is?

The Forty Thieves (24 questions)

1. Alibaba overhears forty thieves tell their boss, Jafar, of a treasure hidden in a cave.
2. Alibaba goes to the cave, gets the treasure, and takes it to his house.
3. Jafar goes to the cave to find the treasure missing. The thieves search the town and find it in Alababa's house. Alibaba notices the thieves placing a mark on his house.
4. The thieves report to Jafar that the treasure is in Alibaba's house, which is marked.
5. Alibaba removes the mark from his house and marks the town`s guardhouse. Jafar and the thieves arrive in town and see the mark.
6. Jafar and the thieves break into the guardhouse and are arrested.

Where does Alibaba believe the treasure is?
Where does Jafar believe the treasure is?
What does Alibaba believe the thieves believe?
What do the thieves believe that Jafar believes?

The Most Wanted (18 questions)

1. Sheriff William wants to trick a local gunman named Jack. The Sheriff loads some money from the town bank into a wagon.
2. Sheriff William goes to the saloon to inform Jack that the money is in the wagon.
3. Sheriff William returns to the bank and takes the money out of the wagon.
4. Sheriff William climbs into the wagon himself.
5. The wagon sets off to the desert. Jack follows it and holds it up at gunpoint. He opens the wagon to find Sheriff William waiting inside to arrest him.

What does Jack believe is in the wagon?
What does Jack believe William believes?
What does William believe Jack believes?

Figure 2: Summaries of the stories used in the evaluation.

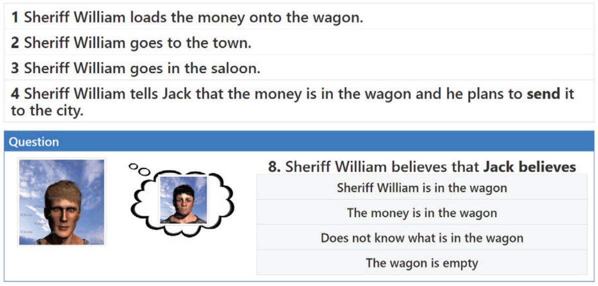

Figure 3: An image of a belief question after four steps

		0		1+1	
		Correct	Incorrect	Correct	Incorrect
2+	Correct	6	18	15	9
	Incorrect	0	0	0	0

(a) Homecoming

		0		1+1	
		Correct	Incorrect	Correct	Incorrect
2+	Correct	11	9	11	9
	Incorrect	0	0	0	0

(b) The Forty Thieves

		0		1+1	
		Correct	Incorrect	Correct	Incorrect
2+	Correct	6	9	8	7
	Incorrect	2	1	0	3

(c) The Most Wanted

Table 1: Contingency tables for the 0, 1+1, and 2+ models

Analysis After Further Noise Reduction Much of the observed disagreement occurred over the initial state, especially for *The Most Wanted*. When participants do not agree on the initial configuration of the world, which was explicitly described for each story, that disagreement is likely to propagate throughout the story. As a further filter on noise, we re-ran the above analysis to remove participants who did not answer the initial state questions as intended (the initial state questions were then excluded from analysis). This reduced inter-rater reliability for *The Forty Thieves* to $\alpha = 0.65$ and *Homecoming* to $\alpha = 0.76$, but increased *The Most Wanted* to $\alpha = 0.37$. This also explains why 70 valid responses were collected in total for *The Most Wanted* above—we needed to ensure at least 20 valid responses based on this filter, and used all data in both analyses.

Using this filter, we discarded 5 questions (9%) which had no agreed-upon answer, and no question had multiple right answers. The overall accuracy for the 2+, 1+1, and 0 models were 100%, 49%, and 44% respectively. These results further support our model.

Formal Evaluation

We claim our 2+ model of belief generates a superset of stories generated by the 0 and 1+1 models. Due to space limitations, we present only the sketch of the formal proof.

Our 2+ model generates all stories the 0 layer model generates because planning without belief is a special case of planning with belief. When all characters accurately know the initial state and everyone is an observing character on all actions, we simulate the 0 model's full observability assumption. Our model generates some stories the 0 model cannot. We offer the stories in Figure 2 as examples. Consider *The Most Wanted*. The 0 model has no way to explain why Jack robs the wagon; it only makes sense when we consider that he wrongly believes the money is in the wagon.

Similar reasoning applies for 1+1. 2+ can generate stories where all epistemic paths of length 2 or more lead to the same state. When this does not hold, it can generate more.

Our final claim is that stories generated by 2+ remain believable. Previous studies (Riedl and Young 2010; Ware and Young 2015) validated that the causal model of intentionality creates believable character behavior. Our model states that when character A believes character B *could* and *would* take an action, it is reasonable for A to anticipate B's action. A study to demonstrate this is planned as future work.

Limitations and Future Work

Our results demonstrate that the 2+ model captures subjects' understanding of belief with high accuracy, higher than the other two models. However, this evaluation was designed to test stories with two layers of reasoning about belief, so the results only support our claim that the 2+ model outperforms 0 and 1+1 models on these types of stories. Indeed, 0 or 1 layer is probably sufficient in many scenarios.

We must reemphasize that Amazon's Mechanical Turk does not provide an ideal population. Specifically, of 255 responses, 195 were not usable because people failed to answer the initial or comprehension questions correctly. A high percentage of readers demonstrated a notable difficulty in processing nested beliefs. For instance, a many subjects considered *what Sheriff William believes* to be the same as *what Jack believes Sheriff William believes*.

This study only asked participants to follow epistemic edges, but one strength of our model is that paths of mixed epistemic and temporal edges are meaningful. For example, asking why Alibaba clears the mark from his house and marks the guard's house would require readers to follow temporal edges after epistemic ones. This is how agents employ nested beliefs in anticipating the plans of others, and we intend to investigate this in a future study.

Finally, our ultimate goal is to develop a fast planning algorithm which generates the same solutions as those defined by this model. We emphasize *fast* because the speed of planning algorithms is an ongoing research topic and our model has considerably increased the size of the state-space by introducing beliefs. Our future work will also focus on addressing this issue, perhaps by enhancing the 2+ model to generate fewer states or by developing a planning algorithm that intelligently selects only those nodes which need to be expanded to generate a solution.

Acknowledgements

This research was supported by NSF award IIS-1647427.

References

Bates, J.; Loyall, A. B.; and Reilly, W. S. 1992. An architecture for action, emotion, and social behavior. In *in Proceedings of the European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, 55–68.

Bolander, T., and Andersen, M. B. 2011. Epistemic planning for single and multi-agent systems. *Journal of Applied Non-Classical Logics* 21(1):9–34.

Brinke, H. t.; Linssen, J.; and Theune, M. 2014. Hide and Sneak: story generation with characters that perceive and assume. In *Proceedings of the International Conference on Artificial Intelligence and Interactive Digital Entertainment*, 174–180.

Bruner, J. S. 1986. *Actual minds, possible worlds*. Harvard University Press.

Cheong, Y.-G., and Young, R. M. 2015. Suspenser: a story generation system for suspense. *IEEE Transactions on Computational Intelligence and Artificial Intelligence in Games* 7(1):39–52.

Christian, D., and Young, R. M. 2004. Strategic deception in agents. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, 218–226.

De Rosis, F.; Carofiglio, V.; Grassano, G.; and Castelfranchi, C. 2003. Can computers deliberately deceive? a simulation tool and its application to Turing's imitation game. *Computational Intelligence* 19(3):235–263.

Demolombe, R., and del Pilar Pozos Parra, M. 2010. A simple and tractable extension of situation calculus to epistemic logic. *Foundations of Intelligent Systems* 191–200.

Doyle, J. 1979. A truth maintenance system. *Artificial intelligence* 12(3):231–272.

Grosz, B. J., and Kraus, S. 1996. Collaborative plans for complex group action. *Artificial Intelligence* 86(2):269–357.

Helmert, M. 2006. The fast downward planning system. *Journal of Artificial Intelligence Research* 26:191–246.

Hoffmann, J., and Brafman, R. I. 2006. Conformant planning via heuristic forward search: a new approach. *Artificial Intelligence* 170(6):507–541.

Holm, S. 1979. A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics* 65–70.

Howell, D. C. 2012. *Statistical methods for psychology*. Cengage Learning.

Kripke, S. A. 1963. Semantical considerations on modal logic. *Mathematical Logic Quarterly* 9(5-6):67–96.

Krippendorff, K. 2012. *Content analysis: An introduction to its methodology*. Sage.

Linsky, L. 1968. On interpreting doxastic logic. *The Journal of Philosophy* 500–502.

Petrick, R. P. A., and Bacchus, F. 2002. A knowledge-based approach to planning with incomplete information and sensing. In *Proceedings of the International Conference on Artificial Intelligence Planning and Scheduling*, 212–222.

Pynadath, D. V., and Marsella, S. C. 2005. Psychsim: Modeling theory of mind with decision-theoretic agents. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1181–1186.

Riedl, M. O., and Bulitko, V. 2013. Interactive narrative: an intelligent systems approach. *AI Magazine* 34(1):67–77.

Riedl, M. O., and Young, R. M. 2010. Narrative planning: balancing plot and character. *Journal of Artificial Intelligence Research* 39(1):217–268.

Robertson, J., and Young, R. M. 2015. Interactive narrative intervention alibis through domain revision. In *Proceedings of the workshop on Intelligent Narrative Technologies*.

Russell, S., and Norvig, P. 2010. *Artificial Intelligence: a modern approach*. Prentice Hall, third edition.

Ryan, J. O.; Summerville, A.; Mateas, M.; and Wardrip-Fruin, N. 2015. Toward characters who observe, tell, misremember, and lie. *Proceedings of the workshop on Experimental AI in Games*.

Ryan, M.-L. 1991. *Possible worlds, artificial intelligence, and narrative theory*. Indiana University Press.

Stevenson, R. L. 1919. *Treasure Island*.

Teutenberg, J., and Porteous, J. 2013. Efficient intent-based narrative generation using multiple planning agents. In *Proceedings of the International Conference on Autonomous Agents and Multi-agent Systems*, 603–610.

Teutenberg, J., and Porteous, J. 2015. Incorporating global and local knowledge in intentional narrative planning. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, 1539–1546.

Ware, S. G., and Young, R. M. 2010. Rethinking traditional planning assumptions to facilitate narrative generation. In *Proceedings of the AAAI Fall Symposium on Computational Models of Narrative*, 71–72.

Ware, S. G., and Young, R. M. 2014. Glaive: a state-space narrative planner supporting intentionality and conflict. In *Proceedings of the 10th AAAI International Conference on Artificial Intelligence and Interactive Digital Entertainment*, 80–86. (awarded Best Student Paper).

Ware, S. G., and Young, R. M. 2015. Intentionality and conflict in The Best Laid Plans interactive narrative virtual environment.

Young, R. M.; Ware, S. G.; Cassell, B. A.; 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, Special Issue on Formal and Computational Models of Narrative* 37(1-2):41–64.