Exploring Regression-Based Narrative Planning

Stephen G. Ware and Orion Fisher
Narrative Intelligence Lab, University of Kentucky
Lexington, Kentucky, 40508
sgware@cs.uky.edu, orion.fisher@uky.edu

Abstract

A valid and believable narrative plan must often meet at least two requirements: the author's goal must be satisfied by the end, and every action taken must make sense based on the goals and beliefs of the characters who take them. Many narrative planners are based on progression, or forward search through the space of possible states. When reasoning about goals and beliefs, progression can be wasteful, because either the planner needs to satisfy the author's goal first and then explain actions, backtracking when an explanation cannot be found, or explain actions as they are taken, which may waste effort explaining actions that are not relevant to the author's goal. We propose that regression, or backward search from goals, can address this problem. Regression ensures that every action sequence is intentional and only reasons about the agent beliefs needed for a plan to make sense.

Introduction

Narrative planning algorithms search for a sequence of actions that tell a story and that make sense for each character involved in the actions. Many search strategies have been adapted from classical planning research, including partial-order causal-link planning (Young 1999; Riedl and Young 2010; Ware and Young 2011), constraint satisfaction (Thue et al. 2016), and answer set programming (Dabral and Martens 2020; Siler and Ware 2020), to name just a few, but as in the classical planning community, many narrative planners are based on forward heuristic search though the space of states (Charles et al. 2003; Teutenberg and Porteous 2013; Ware and Young 2014; Thorne and Young 2017).

Forward search (or progression) starts at the initial state of the problem and checks which actions are possible in that state. Those actions are applied to generate the possible next states. Then any actions which are possible in those states are applied, and so on, until a valid story is discovered. Plans are constructed from start to end in order.

Narrative planning can be challenging because it places complex constraints on what action sequences are considered valid stories, and these constraints may be defined in terms of the whole sequence, or even in terms of the space of possible sequences. Consider intentionality. Narrative planners often require that every action taken by an agent contribute to a sequence of actions to achieve that agent's goal. Because goals are achieved at the end of the sequence, it is difficult to know at the beginning whether the actions an agent is taking will contribute or not.

In this paper, we propose a regression-based narrative planning algorithm that starts at the author's and agents' goals and works backwards to the initial state. Regression planning was described as early as 1975 (Waldinger 1975), but is rarely used in classical planners. We propose it is a good fit for narrative planners for two reasons:

1. Intentions are goal-directed, so searching backwards from goals ensures the planner does not spend effort considering actions that don't contribute to goals.
2. When we allow for a theory of mind (what x believes y believes, etc.), belief propositions can be infinitely nested. Regression can limit the planner to reasoning only about the beliefs that are relevant to the plan.

We begin with a description of our particular narrative planning formalism. We then present our regression algorithm and explain why it is promising. We conclude with a fully worked example to demonstrate the process.

Narrative Planning

Narrative planners have modeled many kinds of story phenomena (see Young et al. (2013) for a survey). In this paper, we build on a version of narrative planning described by Shirvani, Farrell, and Ware (2018) with these features:

- There is a system-level *author goal* that must be achieved by the end of the story.
- Agents have (possibly wrong) *beliefs* about the world and other agents. Beliefs can be arbitrarily nested, meaning there is no depth limit on the theory of mind.
- Agents have *intentions*, or personal goals. For an agent to take an action, the agent must believe the action can contribute to achieving their goal (whether or not it will).

In this section, we formally define our model of narrative planning, modifying Shirvani, Farrell, and Ware's definitions slightly to include an explicit representation of the author as an agent and to redefine intentionality without using causal links. We introduce our own version of the *Treasure Island* problem as a running example in Figure 1, which is a simplified plot of Robert Louis Stevenson's 1883 novel.

In the story, protagonist Jim Hawkins (H) finds a map that gives the location of treasure (T) buried by Captain Flint. Antagonist Long John Silver (S) is Flint's former first mate, but does not know where the treasure is buried. Hawkins lets it be known that he has the map, prompting Silver to recruit a pirate crew and sail to Treasure Island with Hawkins. There, Hawkins digs up the treasure. Both Hawkins and Silver hope to take the treasure for themselves, and Hawkins eventually succeeds.

Formally, a narrative planning problem is a tuple $\langle C, F, G, s_0, A \rangle$. C is a set of agents, F a set of state fluents, G a goal function, s_0 the initial state, and A a set of actions that change the state. Each of these is defined in the sections below.

Agents, Fluents, and Goals

C is a set of objects that represent the *agents*, (i.e. characters) in the story. All domains include the special author agent c_A that represents the author of the story. For *Treasure Island*, $C = \{c_A, H, S\}$.

F is a finite set of state *fluents*, each with an associated finite domain D_f. Each fluent $f \in F$ is like a variable that can be assigned exactly one value from D_f at any moment in time. The proposition $f = v$ means that fluent f has value $v \in D_f$. In Figure 1, the fluent T represents the treasure's location, which can be buried on the island (B), unknown (N), dug up on the island (I), or in the possession of Hawkins (H) or Silver (S). We use the shorthand TB to mean "the treasure is buried on the island." The constant N, for unknown, is simply a value and has no special semantics here.

We define a simple logical language which allows three kinds of propositions p, expressed by this grammar:

$$p := f = v \mid b(c, p) \mid p \wedge p$$

The first kind, $f = v$, is defined above. The modal proposition $b(c, p)$ means that some non-author agent $c \in C$ believes proposition p to be true (where p can be any proposition, including another belief). We also allow conjunctions, $p \wedge p$. We assume this equivalency:

$$b(c, p \wedge q) \leftrightarrow b(c, p) \wedge b(c, q)$$

These three kinds of propositions are sufficient to describe our model, though our implementation (currently under development) includes additional features like negation, disjunction, first order quantifiers, and conditional effects.

In this simplified model, it is often convenient to use a concept of membership in a proposition. Some proposition p, as defined above, is a member of a proposition q if p is itself a conjunction of any number of conjuncts from q. This is denoted by $q \models p$.

G is a function $G(C) \rightarrow p$ that defines the goal proposition of every agent $c \in C$. $G(c_A)$ is the *author's goal*, a proposition which must be true at the end of the story. For *Treasure Island*, $G(c_A) = TH$, meaning Hawkins has the treasure. Hawkins and Silver both want the treasure; $G(H) = TH$ and $G(S) = TS$.

For simplicity, we define every agent to have exactly one goal for the whole story, though in our implementation agents can have multiple goals which can be adopted or dropped during the story.

States and Actions

A state must be able to determine the truth value of any proposition. It must define a value for every fluent, plus every agent's beliefs about the values of every fluent, plus their beliefs about others' beliefs, and so on infinitely.

Two functions are needed to define a state. For some state s and some fluent f, let $V(s, f) \rightarrow D_f$ be the value of that fluent in that state. For some agent c and some state s, let $\beta(c, s)$ be the state that represents agent c's beliefs in s. In other words, when the world state is s, agent c believes the world state is actually $\beta(c, s)$.

The proposition $f = v$ holds in state s when $V(s, f) = v$. The proposition $b(c, p)$ holds in state s when p holds in $\beta(c, s)$.

The author agent c_A does not have wrong beliefs about the state of the world, so we define $\beta(c_A, s) = s$ for all states.

Note that β is a function, which implies that every agent commits to a specific (but possibly wrong) belief about every fluent. This requirement simplifies problems significantly, but means we cannot represent uncertainty (where an agent could hold one of several sets of beliefs). We have found this a useful tradeoff in practice, though others have found it valuable to model uncertainty (Mohr, Eger, and Martens 2018).

As an analog to the use of notation for membership of propositions, it is helpful to succinctly indicate that a proposition p holds true in a state s (equivalently, p is satisfied by s). This is indicated by $s \vdash p$.

s_0 is the *initial state* of the narrative planning problem. It describes the initial values of all fluents and all initial agent beliefs.

In *Treasure Island*, the treasure is initially buried on the island, TB, and Hawkins believes this. Using Shirvani, Farrell, and Ware (2018)'s extension to the closed world assumption, we do not need to explicitly state $b(H, TB)$; this is assumed because TB is true and Hawkins has no explicitly stated belief that contradicts it. Silver does not know the treasure's location, so $b(S, TN)$ must be explicitly stated. Hawkins believes Silver does not know where the treasure is, $b(H, b(S, TN))$, but this also is assumed by the closed world assumption and does not need to be stated. It is equivalent to say that $b(S, TN)$ holds in s_0 and to say that TN holds in $\beta(S, s_0)$.

The set A is the set of all *actions* that could be taken in a narrative planning problem. Every action $a \in A$ has a precondition, $\text{PRE}(a)$, a proposition that must hold in the state immediately before a occurs, and an effect, $\text{EFF}(a)$, a proposition which becomes true in the state immediately after a occurs.

Action preconditions and effects should not be contradictions. For example, an action may not have the precondi-

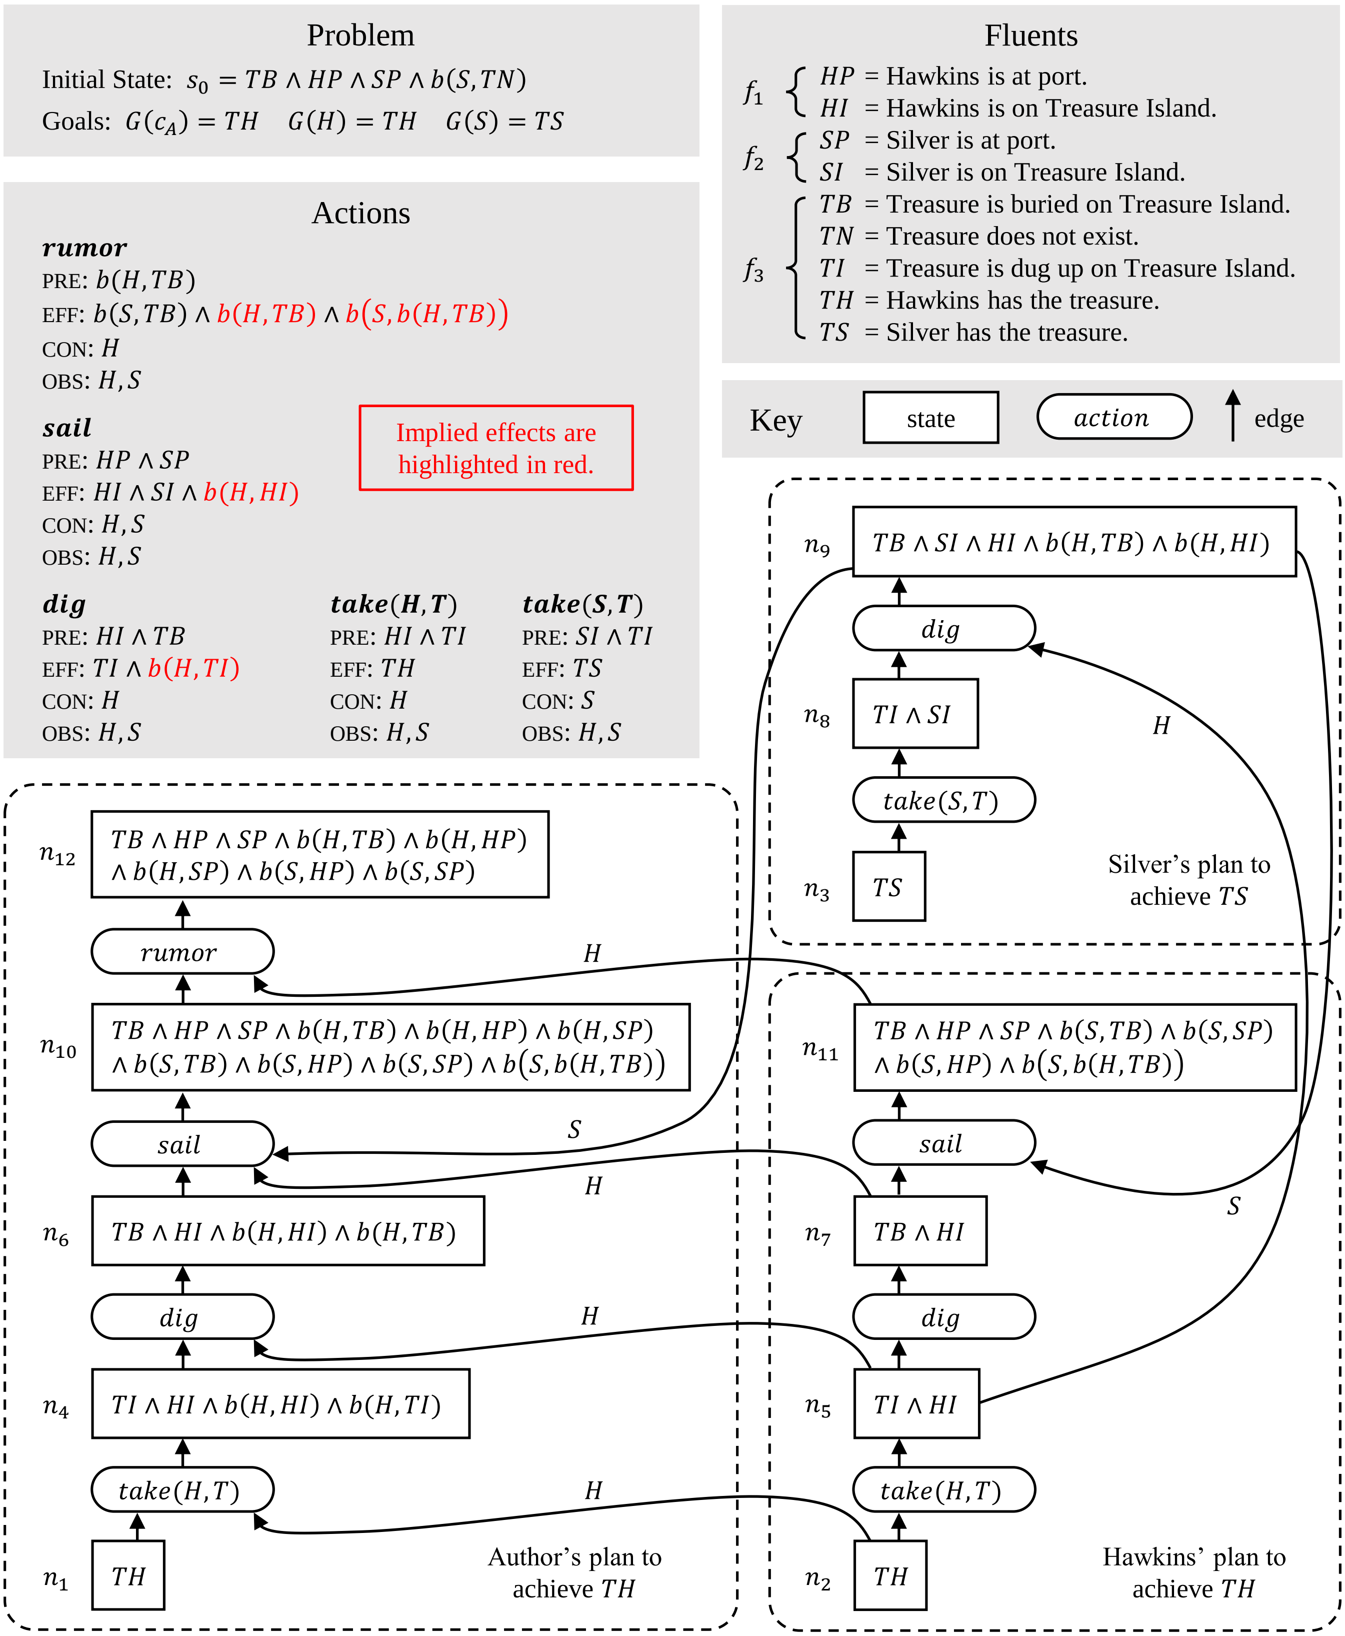

Figure 1: An example problem and example regression search space.

tion $TB \wedge TN$, indicating that the treasure is both buried and nonexistent. Since a fluent may only have one value at a time, this is considered contradictory. This rule also applies to beliefs. For example, an action cannot have the precondition $b(S, TB) \wedge b(S, TN)$, which would indicate that Silver holds two beliefs about the treasure's location. To formally indicate that a proposition p is not a contradiction, we write $p \not\models \bot$.

Actions also define $\text{CON}(a)$, a set of 0 to many *consenting agents*, who must have a reason to take the action. Not every agent involved in an action is necessarily a consenting agent. Consider the $rumor$ action. Silver's beliefs are modified, so he is involved, but he is a passive participant. Only Hawkins needs a reason to take this action, so $\text{CON}(rumor) = \{H\}$. Actions that happen by accident (i.e. actions agents cannot anticipate) should have only the special author agent c_A as the consenting character, which means only the author needs a reason for it to occur.

Finally, every action a defines $\text{OBS}(a)$, a set of 0 to many *observing agents*, which are non-author agents who see the action occur and update their beliefs accordingly. Because $\beta(c_A, s) = s$ by definition, the author effectively observes every action.

Belief propositions can be explicitly stated in preconditions and effects. Consider the $rumor$ action. Its precondition is that Hawkins believe the treasure is buried on the island, $b(H, TB)$, and its effect is that Silver now believes the treasure is buried on the island, $b(S, TB)$. See Shirvani, Ware, and Farrell (2017) for full details on how effects are imposed on states.

Actions can have implied effects which are not explicitly authored but which still result from the action. Some of these implied effects are marked in red in Figure 1. They can happen in two ways.

The first implied effects are from *surprise actions*. It is possible for agents to observe actions they do not believe are possible. For example, if Silver does not know the treasure's location (i.e. he believes $\text{PRE}(dig)$ is false), he would be surprised to see Hawkins dig it up. When a surprise action happens, agents first update their beliefs to correct wrong beliefs and then observe the effects. We accomplish this by copying any preconditions that remain unchanged into the effects of an action. Formally,

$$\forall a, p : (\text{PRE}(a) \models p) \wedge (p \wedge \text{EFF}(a) \not\models \bot) \rightarrow \text{EFF}(a) \models p$$

Consider the $rumor$ action. Its precondition is $b(H, TB)$, and Hawkins' belief about the treasure is not changed by the action's effect, so this action implicitly also has the effect $b(H, TB)$. This is important, because when Silver hears the rumor, he not only believes the treasure is buried on the island, he also believes Hawkins believes this.

The second kind of implied effects are from observations. When a character observes an action, they believe its effects have occurred. Consider $sail$. It has the effect that Hawkins is on the island, HI, and Hawkins observes this action, so it implicitly has the effect $b(H, HI)$. Formally:

$$\forall c, a, p : c \in \text{OBS}(a) \wedge (\text{EFF}(a) \models p) \rightarrow (\text{EFF}(a) \models b(c, p))$$

Valid Narrative Plans

We use the function α to denote the state after a sequence of actions. In state s, let $\alpha([a_1, a_2, ..., a_n], s)$ denote the state of the world after taking those n actions in order from state s. α is only defined if the preconditions of those actions are satisfied immediately before they occur; that is $\text{PRE}(a_1)$ holds in s, and $\text{PRE}(a_2)$ holds in $\alpha([a_1], s)$, etc.

A sequence of actions is a valid story when it achieves the author's goal and when every action can be explained by the beliefs and intentions of the agents who take them.

In a state s, an action a_1 is *explained for* agent c iff there exists a sequence of actions $[a_1, a_2, ..., a_n]$ such that:

1. $\alpha([a_1, a_2, ..., a_n], \beta(c, s))$ is defined.
2. $\alpha([a_1, a_2, ..., a_n], \beta(c, s)) \vdash G(c)$.
3. All actions in $[a_2, a_3, ..., a_n]$ are explained.
4. Unless $c = c_A$, no action has c_A as a consenting agent.
5. No strict subsequence of those actions also meets these same 5 criteria.

In other words, it makes sense for agent c to take action a_1 if and only if, according to c's beliefs about what the current state is, c can imagine a reasonable sequence of actions starting with a_1 that achieves c's goal (items 1 to 3). Item 4 means that accidental actions can only be explained for the author; agents cannot plan for them to happen. Item 5 expresses the idea that the plan the agent imagines should not contain unnecessary or redundant actions.

Some of the actions in the plan may be actions the agent must consent to, which we term actions the agent *takes*. Some of the actions are those that the agent *anticipates* will be taken by other agents. This definition of anticipation is drawn from Shirvani, Ware, and Farrell (2017). Anticipation must be justified in the same way that we ensure actions to be taken are justified: it must be founded on those actions being explained for the agents who take them, according to the anticipating agent's beliefs. In *Treasure Island*, Hawkins anticipates that Silver will choose to sail to the island if he believes his goal can be achieved by acquiring the treasure. This anticipated action is necessary to explain how Hawkins' choice to spread the rumor of the treasure—to take the rumor action—is justified.

Note that the explanatory action sequence only needs to exist; it does not actually have to occur in the story. Silver is willing to sail to the island because he hopes to take the treasure, even if he never actually succeeds in executing this plan. This is Ware and Young's (2014) model of conflict. It is important to note that explaining an action is, itself, a planning problem. The high cost of explaining actions is one of the motivations to use regression planning, which we discuss in the following sections.

In a state s, an action a_1 is *explained* (in general) iff it is explained for every agent $c \in \text{CON}(a_1)$. In other words, an action makes sense when it makes sense for every agent who takes it.

Finally, we can define that a sequence of actions $[a_1, a_2, ..., a_n]$ is a valid solution to the narrative planning problem iff:

- $\alpha([a_1, a_2, ..., a_n], s_0)$ is defined.
- $\alpha([a_1, a_2, ..., a_n], s_0) \vdash G(c_A)$.
- All actions are explained.

Progression

Progression, or forward search, begins at the initial state s_0 and generates possible futures until a state is discovered where the author's goal $G(c_A)$ holds. A classical planner is finished once this node is discovered because any path to the goal is a valid solution.

Progression is difficult for intention-based narrative planners, like ours, because solutions must meet two requirements: the author's goal is achieved *and* every action is explained. Not every path to the goal is a solution. Planners like Glaive (Ware and Young 2014) first search for sequences that achieve the author's goal and then try to explain the actions in the sequence. Significant work is wasted when an action cannot be explained. Glaive's heuristic tries to account for the number of yet-unexplained actions in its calcuations, but this is only effective in some cases.

Recent work on the density of narrative planning solutions (Siler and Ware 2020) suggests it may be valuable to do progression the other way—the planner tries to explain an action immediately after taking it, and when it cannot be explained, that branch of the search can be pruned. This guarantees that any path to the author's goal is a solution, but this approach risks wasting significant work by explaining actions that are not relevant to achieving the author's goal. IMPRACTical (Teutenberg and Porteous 2013) uses an explain-first approach, but actions are explained using heuristics, so it cannot guarantee every action in the final solution will be explained.

Regression

Regression, or backward search, starts at the goal $G(c)$ and generates plans from end to start until one is found that can be executed in the initial state s_0.

Consider Hawkins' goal of acquiring the treasure, represented by TH. Only the $take(H, T)$ action has this as an effect. We can regress Hawkins' goal TH over $take(H, T)$ by removing the action's effects from the proposition and adding the action's preconditions. The result is the proposition $TI \wedge HI$. In other words, if we can find a state where the treasure is dug up and Hawkins is on the island, Hawkins would have a way to achieve his goal—the plan $take(H, T)$.

The search space for the regression is the space of *valid* and *supported* agent propositions, represented by $\langle c, p \rangle$ for $c \in C$. The proposition p represents a goal that, if satisfied, indicates that the agent's goal $G(c)$ may be accomplished by continuing to follow some (potentially empty) plan. The criteria of being valid and supported define two key aspects of the search process, which come together to ensure that the plan which follows from each such proposition is *explained*. These nodes in the search space are connected by actions over which the regression is performed. A node $\langle c, q \rangle$, which was generated by the regression of $\langle c, p \rangle$ over action a is *valid* iff:

1. $q \not\models \bot$.
2. any state satisfying q satisfies $\textsc{pre}(a)$, so a can be taken.
3. $\textsc{eff}(a)$ partially satisfies p, formally: $\exists l : p \models l \wedge \textsc{eff}(a) \models l$.
4. p holds after applying $\textsc{eff}(a)$ to q. $\forall r\ \textsc{eff}(a) \models r$ and $(r \wedge p \not\models \bot)$.

Between nodes of the same agent an action edge represents a step the agent plans to take, or anticipates will be taken. Between nodes of distinct agents the action edge represents evidence for the anticipation of that action. Consider node n_{10} in Figure 1. This node is a valid regression for a node also owned by the author, n_6. It also contains the necessary beliefs to be *supported* by nodes n_7 and n_9. A node $\langle c, q \rangle$ generated by expanding a node with action a is supported if a regression can be found for at least one node for every agent in the consenting set except for c. That is, given γ is the regression function defined in Algorithm 1:

$$\forall c_{other} \in (\textsc{con}(a) - \{c\}) \quad [\exists \langle c_{other}, p_{other} \rangle : q \models b(c_{other}, \gamma(a, p_{other}))]$$

When a node is supported, this indicates that all beliefs necessary to ensure explainability of the actions leading out of that node are present, in particular those which provide reason to anticipate the actions consenting agents will take.

Algorithm

The regression of a single proposition over an action is given by the procedure $\gamma(a, p)$ in Algorithm 1. This function returns the simplest proposition required for the action to be acceptable for any plan continuing from that point, or it signals failure.

Algorithm 1 $\gamma(a, p)$

```
Let a be an action, p is a proposition.
if (∃l : EFF(a) ⊨ l ∧ p ⊨ l) ∧ (∀r, EFF(a) ⊨ r ∧ (r ∧ p ⊭ ⊥)) then
    Let q be PRE(a).
    ∀l, p ⊨ l : Let q be q ∧ l iff EFF(a) ⊭ l
    if q ⊭ ⊥ then
        return q
    else
        return failure
    end if
else
    return failure
end if
```

The process of expanding the regression graph is given by the procedure SEARCH(C, G, A, s_0) in Algorithm 2. If this function returns, it provides a plan that satisfies the author's goal, starting from the initial state.

Search starts with the set of nodes $\{\langle c, G(c) \rangle : c \in C\}$ (line 3). The SEARCH algorithm is an iterative expansion of the search space which proceeds by choosing a node to expand (line 5) and an action to expand it with (line 9), then

Algorithm 2 SEARCH(C, G, A, s_0)

1: C is the set of agents, G is a function of agents to agent goals, A is the set of actions, and s_0 is the initial state.
2: Let X be $\emptyset$
3: $\forall c \in C$: Let X be $X \cup \langle c, G(c) \rangle$
4: **loop**
5:     Choose a node $\langle c, p \rangle \in X$.
6:     **if** $(c = c_A) \wedge (s_0 \vdash p)$ **then**
7:         **return** the path from $\langle c, p \rangle$ to $\langle c_A, G(c_A) \rangle$
8:     **else**
9:         Choose an action $a \in A$.
10:         Let p_{new} be $\gamma(a, p)$.
11:         **for** $c_{other} \in \text{CON}(a) : c_{other} \neq c$ **do**
12:             Choose a node $\langle c_{other}, p_{other} \rangle \in X$ such that $\gamma(a, p_{other})$ does not fail.
13:             Let p_{new} be $p_{new} \wedge b(c_{other}, \gamma(a, p_{other}))$
14:         **end for**
15:         **if** $\langle c, p_{new} \rangle$ not redundant for $\langle c_A, G(c_A) \rangle$ **then**
16:             Let X be $X \cup \{\langle c, p_{new} \rangle\}$
17:         **end if**
18:     **end if**
19: **end loop**

choosing nodes from the plans of consenting agents to establish support for the action (line 12). All such chooses are non-deterministic.

Each expansion produces nodes which describe the conditions under which the plan—the chain of actions leading back to the node $\langle c, G(c) \rangle$ for that same agent—will succeed. These nodes also explain participation of all consenting agents for each action to be taken. The search concludes when a node is found which is both owned by the author and satisfied by the initial state (line 7).

Recall that the sequence used to explain an an action should not contain unnecessary or redundant actions (e.g. sailing back and forth to the island before digging up the treasure). For now, we define a node $\langle c, p \rangle$ to be redundant when it has an ancestor node $\langle c, q \rangle$ such that $p \models q$. In other words, a plan is redundant when it ends with a sequence of actions that would also achieve the goal and would apply in all of the same states (and possibly more).

As an example, consider regressing node n_{12} over $rumor$. This represents the obviously redundant story:

$$\{rumor, rumor, sail, dig, take(H, T)\}$$

Hawkins spreading the rumor that he has the map twice is possible, but unnecessary, because the proposition produced by this regression would be exactly the same as the proposition for n_{12}.

Note that a node $\langle c, p \rangle$ is *not* redundant when it has an ancestor node $\langle c, q \rangle$ such that $q \models p$. The proposition for node n_{12} is a strict subset of the proposition for n_{10}, but spreading the rumor is not necessarily redundant, because the plan represented by node n_{12} may apply in some states where n_{10} does not apply, e.g. any state where $b(S, TN)$ holds—Silver believes the treasure does not exist.

This definition of redundant plans is not as robust as ones used in some progression planners like Glaive (Ware and Young 2014). Improving this check is an area for future work.

Worked Example

Looking at Figure 1 in more detail, we can see how the algorithm takes shape. Initially, we begin our search at the goals for each agent: Silver, Hawkins, and the author. Any of these would be effective choices for our first expansion, but we choose to expand the author's goal, n_1: Hawkins has the treasure.

We compute the regression of TH over $take(H, T)$: $\gamma(take(H, T), TH) = TI \wedge HI$. If the treasure is on the island, and so is Hawkins, we can use $take(H, T)$ to accomplish the author's goal. The resulting node is *valid*, but we must also ensure that the node is supported by finding a regression over $take(H, T)$ from a node owned by Hawkins, the consenting agent of $take(H, T)$. n_2 serves our purpose, and the regression is also $TI \wedge HI$. From the perspective of the author, this is our expectation of what the agent needs to think is true of the world in order to take the action, as opposed to what the true state of the world is. Therefore, this proposition is added as a belief: $b(H, TI \wedge HI) = b(H, TI) \wedge b(H, HI)$. This is conjoined with $TI \wedge HI$ to get the final result. Regardless of whether he is correct, Hawkins believes that n_4 will put him in the position to take the treasure. Since he is correct, the author can accomplish that goal as well.

The next regression in the author's sequence will be the regression of the proposition for n_6 over the action dig, but we can only expand a node if we can find a regression for it *and* for a node from every consenting agent as well as the current one. In this case, we must first expand n_2 (Hawkins' goal to have the treasure) to get n_5 (Hawkins' belief that he can eventually get the treasure if he is on the island and it is too) and now we have everything necessary to produce n_6 in the same way that we did for n_4. When preforming this regression over dig, we must be sure to remove the implied effect $b(H, TI)$, as we preform this regression from n_4, to avoid the contradiction of Hawkins believing the treasure is buried and excavated at the same time.

The process continues as we consider the dig actions for the author and Hawkins, and perform those expansions. Then prior to being able to consider the $sail$ action, which requires Silver's consent, we must expand upon Silver's plan until his search space has a proposition which can be regressed over the $sail$ action. We find that we can perform a regression of his goal over $take(S, T)$, and then regress over the action dig. Hawkins is the only agent who must consent to dig, so Silver must expect that Hawkins will have reason to dig. This is an instance of anticipation. Anticipating the dig action provides an explanation for why Silver should consent to a $sail$ action, if it left the world in a state fitting n_9.

The most complicated proposition for this example is the result of the regression of $TB \wedge HI \wedge b(H, HI) \wedge b(H, TB)$ over $sail$. $sail$ requires consent from both Hawkins and Silver, so we must retrieve their regression results as well, and add their beliefs. The final proposition is given by: $\gamma(sail, TB \wedge HI \wedge b(H, HI) \wedge b(H, TB)) \wedge b(S, \gamma(sail, TB \wedge SI \wedge HI \wedge b(H, TB) \wedge b(H, HI))) \wedge b(H, \gamma(sail, TB \wedge HI)))$. Included in this, as an example of nested belief, is Silver's belief that Hawkins believes the treasure is buried—and therefore Hawkins will seek to dig up the treasure and give Silver the chance to take it. n_{11} is determined in much the same way, but only needs consideration of Hawkins' and Silver's goals, not the author's. n_{12} is expanded in the same way as the others.

At every step the algorithm compares expanded author nodes against the initial state, though we have left out mention of this until now. When n_{12} is compared with the initial state, we see that we have satisfied the needs of the problem—keeping in mind that, unless explicitly stated otherwise in the initial state, we assume that each agent has an accurate belief of the world.

We propose that regression planning has three major advantages:

- By searching backward from goals, we ensure action sequences are intentional. There is still a risk that search effort will be wasted exploring sequences which can never be possible, but regression addresses the two criteria problem described in the previous section. Heuristic search can prioritize sequences that can reach the initial state, and once such a sequence is found, it is guaranteed to be a solution, with no additional constraint checking required afterwards.
- With no limit imposed on the model's theory of mind, it can be difficult to know which beliefs are relevant to an agent's plan. Shirvani, Ware, and Farrell's (2017) model, on which we build, spends much effort generating all changes to beliefs that result from actions, many of which are not relevant. Regression reasons only about the beliefs which are needed to make a plan work.
- Narrative planners are often used in interactive systems where the narrative is replanned frequently. A regression plan expresses only the requirements needed to ensure it will work, so plans found this way can be easily reused in many states. Consider node n_5 in Figure 1. Hawkins has a plan to get the treasure in any state where the proposition $TI \wedge HI$ holds, which might be multiple states during the lifetime of an interactive story.

Conclusions and Future Work

The algorithm we detail here presents a method to manage intention and belief in narrative planning problems in a single search process, with no requirement to check that actions are explained after reaching the author goal. By the nature of the search space, nodes are only added to the search if the action being used for the regression is fully explained.

Our implementation of the algorithm is in development, and will be tested a suite of benchmark narrative planning problems to determine the experimental performance of the method. We also intend to develop and test heuristics to guide the regression effectively. Heuristics like the one used by Glaive are complicated because they attempt to account for the number of yet-unexplained steps in a plan. Since every node produced by our regression planner is represents a valid plan, a heuristic only needs to estimate the distance between the initial state and a node's proposition.

Acknowledgments

This research was supported by the National Science Foundation under Grant No. IIS-1911053.

References

Charles, F.; Lozano, M.; Mead, S.; Bisquerra, A. F.; and Cavazza, M. 2003. Planning formalisms and authoring in interactive storytelling. In *Proceedings of the conference on Technologies for Interactive Digital Storytelling and Entertainment*.

Dabral, C., and Martens, C. 2020. Generating explorable narrative spaces with answer set programming. In *Proceedings of the 16th AAAI international conference on Artificial Intelligence and Interactive Digital Entertainment*. (forthcoming).

Mohr, H.; Eger, M.; and Martens, C. 2018. Eliminating the impossible: a procedurally generated murder mystery. In *Proceedings of the 5th Experimental AI in Games workshop at the 14th AAAI international conference on Artificial Intelligence and Interactive Digital Entertainment*.

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

Shirvani, A.; Farrell, R.; and Ware, S. G. 2018. Combining intentionality and belief: revisiting believable character plans. In *Proceedings of the 14th AAAI international conference on Artificial Intelligence and Interactive Digital Entertainment*, 222–228.

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 international conference on Artificial Intelligence and Interactive Digital Entertainment*, 101–107.

Siler, C., and Ware, S. G. 2020. A good story is one in a million: solution density in narrative generation problems. In *Proceedings of the 16th AAAI international conference on Artificial Intelligence and Interactive Digital Entertainment*. (forthcoming).

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

Thorne, B. R., and Young, R. M. 2017. Generating stories that include failed actions by modeling false character beliefs. In *Proceedings of the 10th workshop on Intelligent Narrative Technologies at the 13th AAAI international conference on Artificial Intelligence and Interactive Digital Entertainment*.

Thue, D.; Schiffel, S.; Árnason, R. A.; Stefnisson, I. S.; and Steinarsson, B. 2016. Delayed roles with authorable continuity in plan-based interactive storytelling. In *Proceedings of the 9th International Conference on Interactive Digital Storytelling*, 258–269.

Waldinger, R. 1975. Achieving several goals simultaneously. Technical report, Stanford University.

Ware, S. G., and Young, R. M. 2011. CPOCL: a narrative planner supporting conflict. In *Proceedings of the 7th AAAI international conference on Artificial Intelligence and Interactive Digital Entertainment*, 97–102.

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).

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.

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.