Knowledge Goal Recognition for Interactive Narratives

Cory Siler[1], Stephen G. Ware[1]

[1]Narrative Intelligence Lab, Department of Computer Science, University of Kentucky, Lexington, KY, USA 40506

Abstract

Player goals in games are often framed in terms of achieving something in the game world, but this framing can fail to capture goals centered on the player's own mental model, such as seeking the answers to questions about the game world. We use a least-commitment model of interactive narrative to characterize these *knowledge goals* and the problem of *knowledge goal recognition*. As a first attempt to solve the knowledge goal recognition problem, we adapt a classical goal recognition paradigm, but in our empirical evaluation the approach suffers from a high rate of incorrectly rejecting a synthetic player's true goals; we discuss how handling of player goals could be made more robust in practice.

Keywords

goal recognition, interactive narrative, narrative planning, question answering

1. Introduction

Goal recognition is the task of inferring the intentions behind an agent's actions. When the agent in question is a human game-player, it can serve as a form of player modeling [1] that helps the system predict what the player will do next. Proposed applications have included tailoring procedurally-generated quests to a player's preferences in an adventure game [2], assessing the player's understanding in an educational application [3], or detecting when the player's actions threaten to derail the authorial intent in a story-focused experience [4].

Goal recognition has been studied extensively in a games context [2], but the work so far has largely centered around goals about the state of the game world: *achievement goals* to make a particular fact about the world state true or *maintenance goals* to prevent a fact from being undone [5]. Although many goals in games fit into this framework—e.g., obtaining an item, getting to a location, or defeating an adversary are achievement goals; keeping a character alive is a maintenance goal—players may also pursue goals that cannot be expressed solely in terms of world state.

For instance, Ram [6] defines *knowledge goals* as the intentions of an agent to extend or organize its own mental structures. Knowledge goals encompass players' desire to explore the game world, uncover mysteries about past occurrences, and explain unexpected findings. They are central to genres such as mystery games [7] as well as tutoring and training systems [8]. A goal recognition approach that does not account for exploratory behavior is liable to fail even when dealing with an accomplishment-focused player, since a knowledge goal can be instrumental to a world-state goal: Obtaining the item involves figuring out where to find it, and staying alive involves figuring out which characters have harmful intent.

In interactive narrative games, the line between knowledge goals and non-knowledge goals is further blurred: As gameplay progresses, the player extends their model of the story by discovering new information through their interactions while at the same time their choices constrain the range of possible stories that could emerge. In an interactive narrative architecture using an *experience manager*—an artificially intelligent agent that controls the non-player elements of the game to adapt to the player's decisions—the experience manager may have its own goals for the story, reflecting the game designers' intent for the player's experience. An experience manager that recognizes the player's goals can find the overlap with its own goals and guide the story down a path that satisfies both [9, 10].

This paper's contributions are as follows. First, we provide a framework for characterizing knowledge goals and knowledge goal recognition in an interactive narrative environment. Because the player may have limited awareness of how the game will respond to their decisions, classical frameworks that make strong assumptions about an agent's model of environment dynamics are inadequate. Instead, we draw on a formal model of discourse from semantics and pragmatics that has the asking and answering of questions as its basic operations [11]. Analogous to how a question prompts the respondent to extend the body of information mutually known to both parties in the dialogue, a player acting in an interactive narrative game prompts the game to extend the mutually-known body of information about the story. Our model can capture traditional goal types, but also knowledge goals reflecting implied questions for the story to answer; we define these goals with reference to a cognitive model of literal, spoken questions [12].

Second, we present a preliminary study of algorithms

AIIDE Workshop on Experimental Artificial Intelligence in Games, October 08, 2023, University of Utah, Utah, USA

cory.siler@uky.edu (C. Siler); sgware@cs.uky.edu (S. G. Ware)

for identifying a player's knowledge and achievement goals from the player's actions. We adapt a planning-based goal recognition paradigm [13] into our framework to define these algorithms, and empirically evaluate them on synthetic player agents. Our experiments reveal important shortcomings in the algorithms; robust goal recognition for diverse goal types remains an open problem, so we conclude by discussing a research agenda for addressing it.

2. Related Work

Our model of interactive narrative draws from others that treat the story world as only "existing" as far as the player is aware of it, rather than simulating the entire world. By modeling the player's knowledge of the story so far, an experience manager can delay decisions about properties and events outside the player's perception and use those decisions as tools to adjust the course of the story in response to the player's actions. This idea has been the basis of approaches to preventing player derailment of experience manager goals [14], saving computational resources by deferring [15] or shortcutting [16] how off-screen events are decided, and increasing the depth [17] and diversity [18] of generated stories.

Horswill [19] introduces the term *story state* to refer to the evolving set of design decisions about a story over the course of its creation, whether the creation involves the changing decisions of an external author or improvisation of story-world background within an interactive narrative during the unfolding of the narrative events themselves. The state-transition model we use in our framework operates on a form of story state.

Baikadi et al. [7] present a machine learning model for recognizing player goals where the player may not be aware of all game-supplied goals at the start, and may take exploratory actions to discover new gameplay goals for themself. Interactive narratives are modeled as a graph of *narrative discovery events* where story-driving information is revealed; with the testbed of the *Crystal Island* [20] educational mystery game, Baikadi et al. allude to the idea we build on here of handling knowledge-seeking and objective-achieving in a unified framework. However, their approach uses training data to build domain-specific goal recognition models, while ours uses domain-independent planning to try to recognize goals in the absence of training data.

Goal recognition has been explored in a planning-centric narrative context before, albeit focused on world-state models of planning. Farrell and Ware [21] take a narrative generation framework that models story characters' beliefs and intentions [22], and build upon it to identify the intentions and beliefs of an existing agent from its actions. Cardona-Rivera and Young [9] present algorithms to recognize a player's intentions for the narrative trajectory, and also propose that interactive narrative *players* predict an *experience manager's* intentions for the narrative trajectory, and that plan recognition algorithms can serve as a proxy for how players make these predictions.

Meneguzzi and Pereira [13] give a survey of planning-oriented approaches to goal recognition in general. They taxonomize the approaches by the type of environment (stochasticity/determinism and complete/incomplete information), extent of the goal recognizer's information (complete awareness of the target agent's actions vs. missing or noisy information), the target agent's behavior (whether it plans optimally and whether it tries to thwart goal recognition), and form of the solution (whether candidate goals are assigned a probability or qualitative ordering for how likely they are to be the true goal, or else a binary accept/reject decision). Recently active goal elicitation has been proposed by Amos-Binks and Cardona-Rivera [23] where the recognizer affects the agent's environment.

Building on the analogy of interactive narrative as a dialogue between player and game [24], player and author [10], or player and co-participants [25], our framework borrows from a model of explicit dialogue from semantics and pragmatics by Roberts [11] with foundations from Stalnaker [26]. We adapt the concept of the *common ground*, the set of propositions that dialogue participants mutually accept as true, and the progression of a dialogue as a sequence of *moves* from among two types, *assertions* that add to the common ground (analogous to our observation sets) and *questions* that define which subsequent moves are relevant (analogous to our player actions). However, a key difference between our assumptions and those of the dialogue models is that moves in an interactive narrative can constrain or determine which facts are true rather than simply revealing static facts that were already true.

Another perspective on explicit question-asking comes from Rothe et al. [27], who ask what humans are likely to ask in the context of the game *Battleship*. They define a "good" question in terms of maximizing information gain. By studying players empirically, they conclude that people tend to choose the most informative questions when presented with a list of question options, but not when generating their own questions from scratch. Their environment is restricted compared to our interactive narrative focus: They assume that the *Battleship* players are asking questions solely to optimize their chances of winning the game, and informativeness of questions translates directly to increased ability to win, whereas we consider knowledge goals that sometimes reflect questions asked simply for their own sake.

3. A Model of Interactive Narrative and Goals

In Section 3.1, we propose a representation for how information is revealed over the progression of interactive narrative. In Section 3.2, we define a class of knowledge goals with respect to this framework.

3.1. Interactive Narrative Domains

At a high level, we model an interactive narrative domain as a state-transition model similar to a Markov decision process, but nonstochastic: there is a known set of possible outcomes for taking a given action in a given situation, but no assignment of probabilities to outcomes.

The domain is a tuple $\langle P, S, A, L, O \rangle$. P is a universe of propositions. S is a universe of proposition sets $s \subseteq P$. We call s the *common ground* in reference to the discourse model by Stalnaker [26], as it represents the information about a story mutually known between the player and game at a given time during a playthrough. Besides not being self-contradictory, we place no general restrictions on the contents of a common ground, although we propose a more restricted implementation later in this section. A common ground functions like a state, but unlike a world state which only tracks facts that are true in the present moment of the story, a common ground in describes the story as a whole and will only grow over time; if a propositional representation of world state needs to be tracked within the model, the propositions should be defined to contain time indices or other ordering constraints that distinguish the past of the story from the present in which the player is currently interacting.

A is a set of player actions. $L(s)$ is a function mapping a common ground $s \in S$ to the set of player actions that are legal from that common ground. $O(s, a)$ is a function that maps a common ground $s \in S$ and action $a \in L(s)$ to a set of possible resultant *observation sets*, each of which takes the form of a proposition set $o \subseteq P$ where $(s \cup o) \in S$.

We use this formalism to model the evolution of a player's knowledge over the course of a playthrough of an interactive narrative game. At any given time, the current common ground s encompasses all of the facts revealed to the player about the story so far. When the player chooses an action a, they know the result will be some observation set in $O(s, a)$, but they cannot necessarily predict which one. This can model game architectures where the actual results of actions are predetermined but depend on information unknown to the player (and therefore unmodeled in the common-ground representation), but also architectures that use least- or late-commitment experience management or dynamic procedural content generation where that information is altogether undecided by the system until it is needed. After a is taken, the game determines and reveals the chosen observation set $o \in O(s, a)$. This encompasses the direct results of the player action as well as anything else that happens in the story world before the player is able to act next (e.g., NPC actions). The new common ground then becomes $s' = (s \cup o)$.

As a more specific language for representing the common ground, we consider the knowledge representation from QUEST [12], which has seen prior use for modeling audience reasoning about narratives [28] and can encode the causal relationships [29] and character intentions [30] commonly emphasized in plan-based narrative generation.

A QUEST knowledge structure (QKS) is a directed graph where nodes are annotated with semantic information and where nodes and arcs each have one of several predefined *types*; see Graesser et al. [31] for an extended account of types and their constraints. We focus on a few types: *event nodes* which correspond to character actions or happenings in the world, *state nodes* which correspond to something being true in the world, *consequence arcs* which express a causal relationship between two nodes, *goal nodes* which define in-story character goals (distinct from our model of player goals; we omit these from our examples for brevity and clarity), and *outcome arcs* showing motivation of event nodes by goal nodes, and *reason arcs* linking goal nodes together as character plans.

To relate this to the abstract model from above, we can define the propositions in P to indicate the existence of QKS nodes and edges, so a common ground in $s \in S$ corresponds to a QKS. When the player takes an action a in s, each observation set in $O(s, a)$ will include at minimum a new event node expressing that a took place and consequence arcs to that event node from prior events or facts that made the player action possible.

We describe an example of how we represent a common-ground change in a hypothetical adventure game. To start, the player is informed that their character is at their cottage and that a bandit has just broken into the cottage and left with some stolen money. These facts make up the common ground s. We illustrate an initial QKS representation in Figure 1, including a state node reflecting the player's location, and a network of state and event nodes reflecting the burglary backstory. (Depending on the architecture, the game may have predetermined where the coin and the bandit went *after* the burglary, but this information is not yet part of the story as far as the player is aware so it is not modeled in the common ground.)

The player is presented with a choice of actions to go to one of the two other locations in the game world—the market and the camp. The player chooses to go to the market (action a). From their perspective, there may be one of multiple outcomes (the full range of possibilities makes up $O(s, a)$): They may encounter the bandit there,

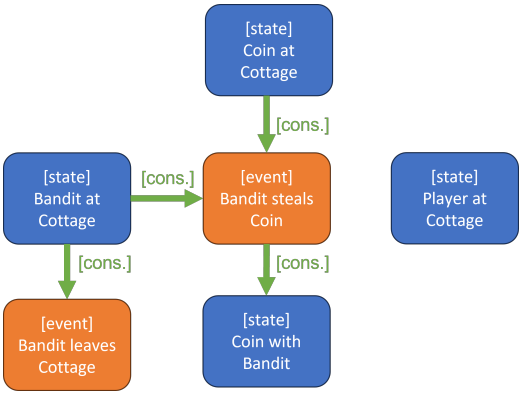

Figure 1: Example of a common ground in the QKS representation.

and may or may not witness the bandit spending the stolen money there, or else they may not find the bandit and therefore conclude that the bandit went to the camp instead. These map to candidates for how to update the common ground, as illustrated in Figure 2.

The game mechanics resolve what the actual outcome should be: For instance, the player is informed that they see the bandit buying a potion with the money. We update the QKS to include the corresponding observation set to represent the new situation, as with the top-right QKS in Figure 2. The observation set adds information both forward in time, such as the event node for the player's action of traveling, and backward in time, such as the consequence arc reflecting the past occurrence of the bandit's arrival at the market.

3.2. Goals

Abstractly, we define a player goal as a formula over propositions in P. We say that a goal G is satisfied in common ground s if $s \models G$. In the QKS model, this translates to a goal specifying what nodes and arcs can be added to make the QKS satisfactory. This representation does not lose generality over a world-state model since it can specify world-state goals using state nodes for the desired facts, but we focus this section on how it can be used to define a certain class of knowledge goals.

An advantage of the QKS representation is that it can make direct use of QUEST's *question-answering* procedures to determine whether the present common ground answers certain kinds of questions about the story. For instance, a question of the form "How did *[event/state]* happen?" can be answered by following consequence arcs backward from the node; a question of the form "What are the consequences of *[event/state]*?" can be answered by following consequence arcs forward from the node; or a question of the form "Why did the character want (via *[goal node]*) to do *[event linked by outcome arc to the goal node]*?" can be specified by following reason arcs forward from the goal node. The QUEST cognitive model predicts that among the neighboring nodes returned this way, humans will rate nearer neighbors as better answers than more distant ones.

In the situation where we ask a QUEST-style question about a common ground that does not yet contain an answer, we can model the question as a goal to reach a common ground where the answer exists. For instance, consider the example from Section 3.1. The player could have the question: "What are the consequences of the bandit having the coin?" The question starts out unanswered in the initial QKS of Figure 1; but by going to the Market, the player may witness the bandit using the coin as in the top-right QKS of Figure 2, in which case a consequence arc from the bandit-has-coin state node now exists and the question is answered.

Formally, let s be the QKS representation of the current common ground. We define a *QUEST question goal* as a tuple $\langle n_q, t_r, d_r, t_a \rangle$, where n_q is an existing node in s, d_r is an arc direction among *incoming* or *outgoing*, t_r is an arc type, and t_a is a node type. We say that a state $s' \supseteq s$ satisfies $\langle n_q, t_r, d_r, t_a \rangle$ if the QKS for s' contains an arc of type t_r going direction d_r from n_q such that the node on the other end of the arc has type t_a. n_q constitutes the subject of the question (e.g., an event/state node in the *how* case) and the others define the form of an answer (e.g., incoming consequence arc from a state/event node).

4. Goal Recognition

Suppose we have a log of actions the player has taken during an interactive narrative. The log may end before the player has completed any identifiable goals, but an observer—e.g., a game designer analyzing the playthrough in hindsight or an experience manager trying to adapt to the player for later interactions—may nonetheless need to reason about the player's intentions. How do we identify possible player goals, including knowledge goals, motivating the actions?

Our first attempt to solve this problem is a *goal recognition as planning* [32] approach: We model the player with an artificial agent which we call an *agent model*, hypothesize that the player has a specific goal, simulate the agent model's reasoning about how to pursue that goal, and determine whether the agent could have chosen the same course of action that the player did in the logs; if so, we conclude the player had that goal.

However, agent models for existing goal-recognition-as-planning approaches prescribe behavior in a deterministic or stochastic environment, whereas our framework treats the game as a nondeterministic, nonstochastic envi-

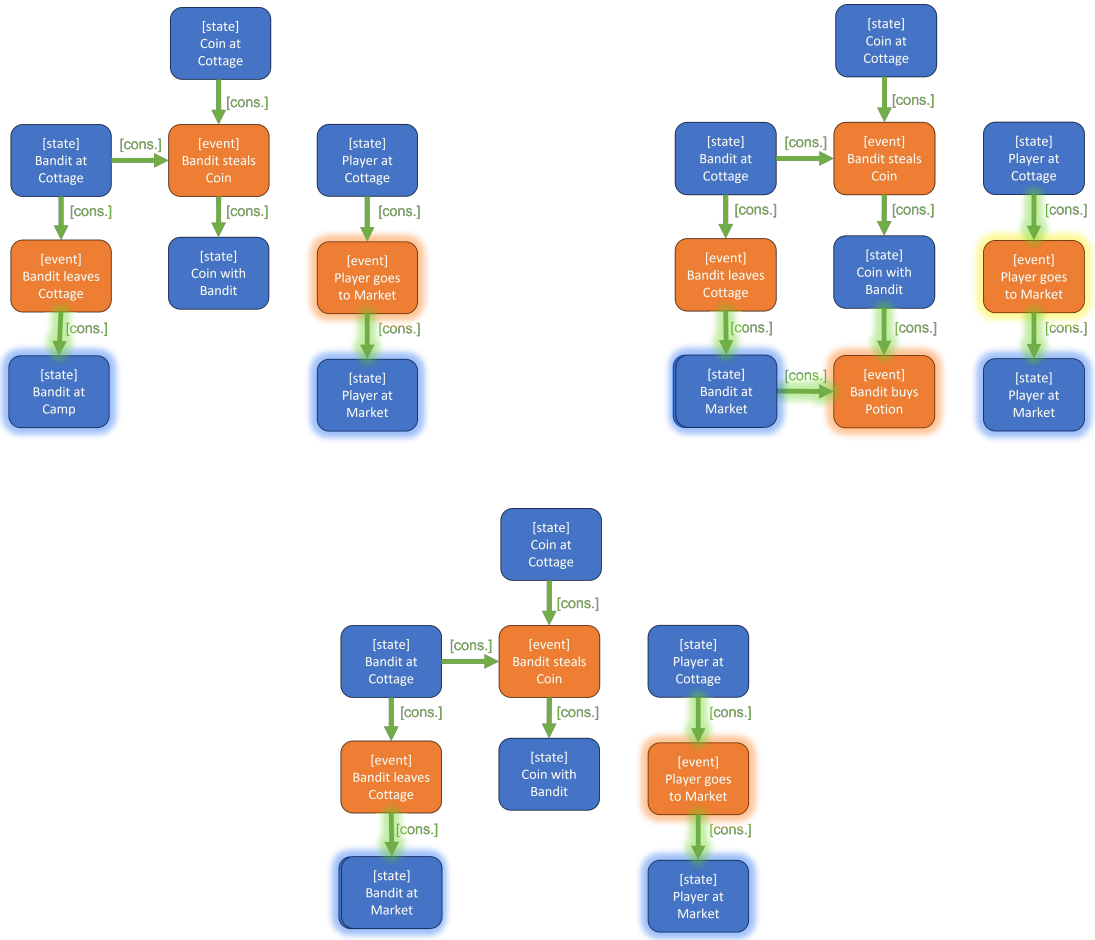

Figure 2: Example of possible observation sets (highlighted) and resulting QKS structures from one player action after the QKS in Figure 1.

ronment: players know the range of possible observation sets that could result from their action but have no reliable way of anticipating which specific observation set will be chosen. An agent model now needs to account for how a player might handle this unpredictability.

In our framework, we define a goal recognition problem instance as a tuple $\langle D, T, C \rangle$. D is an interactive narrative domain $\langle P, S, A, L, O, M \rangle$ as defined in Section 3.1. T is a trajectory consisting of a sequence $s_1, a_1, s_2, a_2, \cdots s_n, a_n$, where $s_i \in S$ and $a_i \in A$ for $i = 1$ to n. This represents (chronologically) the common grounds that the player has experienced so far and the actions the player took in response. C is a set of candidate goals. M is an agent model as elaborated later in this section.

Assume T comes from a game log, presenting a snapshot of an in-progress playthrough where the player is acting toward some goal in C but has not yet achieved it. The solution to a goal recognition problem is the set $C' \subseteq C$ of candidate goals such that an agent modeled by M acting in domain D pursuing any goal $g \in C'$ could produce trajectory T.

Algorithm 1 sketches the goal-recognition-as-planning process. For each player action so far in T, it checks the consistency of that action with each goal; assume the subroutine VERIFY is a search process that returns whether the action could be selected by the agent model. (We check each action individually instead of the whole sequence of actions at once because an agent may plan with the expectation of a certain observation set but receive a different observation set in actuality and have to revise its plan.)

We spend the rest of this section discussing specific agent models that the goal recognizer could assume.

Algorithm 1 Goal recognition for the common ground

Input: Domain D, common-ground/player-action trajectory T, candidate goals C, agent model M

Output: A set of goals $C' \subseteq C$ that an agent modeled by M could have been pursuing if it took the action sequence in T

$C' \leftarrow C$
for all $s_i \in T$ **do**
  for all $g \in C'$ **do**
    if $\neg$VERIFY(D, s_i, a_i, g, M) **then**
      $C' \leftarrow C' \setminus \{g\}$
return C'

First we propose goal recognition using an *optimistic planning* agent model where the agent plans for the best case, hoping that its action will result in a specific observation set that gets it closer to the goal. Given a current common ground s_i, for an optimistic-planning agent with goal g, the agent can take an action a_i iff there exists some hypothetical plan $a_i, s_{i+1}, a_{i+1}, \cdots a_j, s_j$ where s_j satisfies g; we also require the plan to be nonredundant in that no strict subsequence of $a_i, a_{i+1}, \cdots a_j$ also satisfies g. This definition is based on Sabre's character model [22, 33].

We also propose an *adversarial planning* agent model that plans for the worst case, trying to act according to a policy that can eventually satisfy the goal even when its actions result in the worst-case observation sets. For some goal g, define a *safe* common ground s as (base case) one that satisfies g or (recursively) for which there exists an action $a \in L(s)$ such that all outcomes in $O(s, a)$ result in safe common grounds. Given a current common ground s_i, for an adversarial-planning agent with goal g, the agent can take an action a_i if all possible resulting common grounds are safe. However, because this definition alone could easily result in situations where the agent would have *no* valid action choices defined (because the agent will eventually have to take an action where at least one possible outcome could prevent the goal), we generalize this definition—we model an agent who believes the observation sets are chosen uniformly at random, and the agent follows an expectimax-style [34] policy that it thinks will maximize the worst-case *probability* of satisfying the goal. Define the score EXPECTIMAX(s) of a common ground s for goal g as 1 if g satisfies s, or 0 if g can never be satisfied from s (e.g., because s is a leaf in a finite tree of possible trajectories); otherwise, define EXPECTIMAX(s) as the average score of common grounds reached from choosing a best action, $\max_{a \in L(s)} \frac{\sum_{o \in O(s,a)} \text{EXPECTIMAX}(s \cup o)}{|O(s,a)|}$. An adversarial-planning agent can take an action a_i if a_i maximizes the average in this manner.

5. Experiments

There are many risks to the robustness of a goal recognition model when applied to real human players: the player acting toward a goal outside of the candidates considered, changing goals, behaving in a non-goal-directed way, missing or misunderstanding information the model assumes is available to them. This preliminary study considers synthetic players that do not yet incorporate these risks, but we acknowledge the importance of human factors for our future work.

There is a wide spectrum of ways even idealized artificial agents can handle the nondeterministic environments of our framework, as shown by the contrast between the highly risk-taking optimistic-planning agent model and the highly risk-averse adversarial-planning agent model described in Section 4. A mismatch between the agent model assumed by the goal recognizer and the decision-making criteria of the actual player can result in wrong conclusions about the player goals—false positives where a candidate goal is wrongly attributed to the player, and false negatives where the player's actual goal is wrongly rejected as a possibility.

Our experiment seeks to quantify the error-proneness of goal recognition that assumes one agent model when the player acts according to another agent model. By using an optimistic planner as the "real" player and trying to identify that player's goals using the opposite extreme of an adversarial-planning goal recognizer, and the reverse, we aim to establish upper bounds on goal recognition error before human factors are applied.

We generated goal recognition problem instances as follows: To derive the domain D, we started with depth-limited, tree-structured *story graphs* [35] from a narrative planning [36] problem, generated using the Sabre narrative planner [33]; these graphs consist of nodes representing world states and edges representing player or non-player actions, annotated with information such as whether the player observed a given non-player action. We restructured the story graphs to alternate between branching on a choice of player actions and branching on a choice of non-player macro-actions containing any number of non-player actions. At each node, we used previously-proposed mappings [29, 30] to derive a QKS equivalent of the story so far. We also used an approach similar to Robertson and Young [37] and Fisher et al. [38] to allow uncertainty about which of multiple story-graph nodes the player was in, due to possible unobserved past events; we derived the final QKS common ground representing the player knowledge by taking the maximal subgraph that is shared by the original stories.

To obtain a trajectory T of player actions so far, we sampled and truncated goal-satisfying playthroughs given a goal g and agent model M. We manually defined the set of candidate goals C for the domain.

As a source for our domain, we used the narrative planning problem from the *Grandma* adventure game used by Ware et al. [39]. We retained the same characters, actions, locations, and items, but modified the initial state and NPC goals to create the initial setup as follows: Known to the player in the initial QKS, three actions have already happened in the backstory: The bandit character has stolen a sword and a coin from the player character's house and left the house. The merchant character is at the market and the guard character is at the bandit's camp, both locations reachable from a crossroads reachable from the cottage. Unknown to the player and thus not represented in the initial QKS, the bandit intends to use the sword to kill the guard and/or rob the merchant, and/or use the coin to buy from the merchant.

We defined four possible goals that the synthetic player would try to satisfy and that the goal recognizer would try to distinguish between as the candidate goal set C: achievement goals to get the stolen coin and to get the stolen sword, and knowledge goals in the form of QUEST question goals for "Why did the bandit steal the coin?" and "Why did the bandit steal the sword?" These goals overlap in some of the player actions that can be used in the course of satisfying them, e.g., following the bandit can support any of the goals, but diverge in others, e.g., killing the bandit enables taking back the stolen items but eliminates opportunities to watch the bandit's plans unfold and learn their intentions.

We generated the narrative planning problem's story graph to a fixed depth of 6 steps, based on available computation time, and converted it to the graph of QKS common grounds as described above. We then collected all trajectories for each agent model and each goal and analyzed them in the following process:

Suppose the real player is an optimistic-planning agent and the goal recognizer assumes an adversarial-planning agent model, or the reverse. Let T be a trajectory; let C be all the goals, let C'_{true} be the set of goals for which the player generated T and let C'_{output} be the set of goals identified by the goal recognizer for T. We counted goals in $C'_{true} \cap C'_{output}$ as true positives for which the goal recognizer would correctly identify a player goal for T; $C'_{true} \setminus C'_{output}$ was false negatives for which the recognizer failed to identify a goal that was consistent with the true player model; $C'_{output} \setminus C'_{true}$ was false positives for which the recognizer identified a goal that was actually inconsistent with the true player model; and $C \setminus (C'_{true} \cup C'_{output})$ was true negatives where the recognizer correctly did not identify a goal that would have been inconsistent with the true player model.

We show confusion matrices for the results across all goals and trajectory lengths: Figure 3 shows the performance of a goal recognizer assuming an adversarial-planning agent model if the actual player acts like an optimistic-planning agent model, and Figure 4 shows the reverse. We show standard measures of performance for a test to distinguish positive from negative cases: sensitivity (how often the recognizer concluded the player had the goal, given that the goal was consistent with the true player model), and specificity (how often the recognizer concluded the player did not have the goal, given that the goal was inconsistent with the true player model).

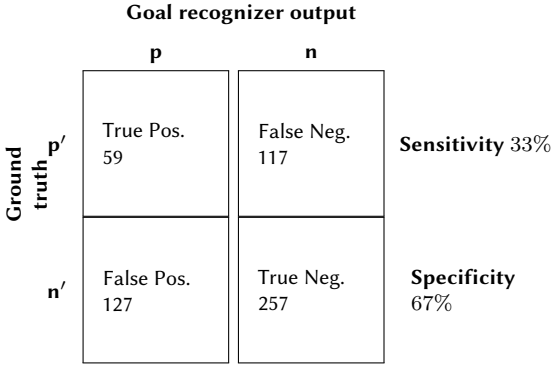

Figure 3: Adversarial-planning goal recognizer on optimistic-planning player over all goals and trajectories

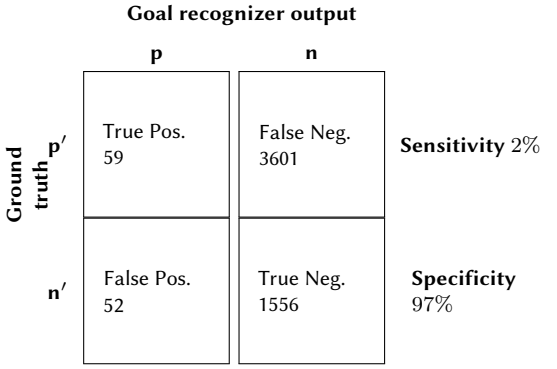

Figure 4: Optimistic-planning goal recognizer on adversarial-planning player over all goals and trajectories

Both of the player-recognizer combinations had worse-than-random sensitivity and better-than-random specificity; the recognizers skewed toward correctly rejecting candidate goals when the player did not have those goals, but failing to detect the true goals. The difference was especially strongly pronounced in a goal recognizer that assumed an optimistic-planning agent model when the actual player used the adversarial-planning agent model; that is, when considering a goal that the player actually had, the optimistic-planning goal recognizer was highly likely to erroneously reject that goal. These instances came from the fact that our optimistic-planning agent model attempts to be as efficient as possible by avoiding

actions that could be redundant to the goal, while the adversarial-planning model accepts longer paths in favor of safety; the simulated adversarial-planning player often took actions that were unexplainable to the optimistic-planning recognizer because there was a more direct route available. This experiment suggests that strict assumptions about agent efficiency—which are common in existing goal recognition approaches—are too brittle in practice, and future goal recognition approaches should be designed to handle cautious or meandering players.

6. Conclusions

This paper highlighted an underexplored class of goals important to interactive narratives—player goals to fill the gaps in their knowledge about the story so far. We extended goal recognition to these goals by defining a planning framework over the space of player mental models rather than over the space of world states, drawing on representations of discourse and question-answering from linguistics and cognitive science.

Accurate algorithms for knowledge goal recognition are still an open problem. An approach based on simulating a hypothetical player and comparing its decisions to the real player's can easily fail to detect goals of a player whose playstyle does not match the algorithm's assumptions. However, the desiderata for a goal recognition algorithm depend on how that algorithm will be used. For instance, high-specificity but low-sensitivity goal recognition could be acceptable for an experience manager whose objective is to find a small handful of the player's interests and use them to offer the player mutually-beneficial opportunities. Conversely, low-specificity but high-sensitivity goal recognition can still be useful for a highly-improvisational experience manager deciding when to fix "plot holes" in its stories that may be exposed by player knowledge goals [40].

Reasoning about player goals will ultimately require considering the context that goals come from. In the case of knowledge goals, aside from offering models, the literature we reference goes on to emphasize that reasoning effectively about questions requires understanding *why* they were asked: Roberts [11] and Ram [41] frame basic questions as part of strategies to answer higher-level questions, and Graesser et al. [42] and Ram [6] discuss the functions of questions to support the asker's goals and explain anomalous findings. In the long term, we aim to take theories of when knowledge goals are likely to occur, and integrate them with mechanisms for confirming those knowledge goals from a player's actions and for using this information to shape the story in concert with the player.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. IIS-2145153. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

[1] G. N. Yannakakis, P. Spronck, D. Loiacono, E. André, Player modeling, Artificial and Computational Intelligence in Games (2013).

[2] B. Harrison, S. G. Ware, M. W. Fendt, D. L. Roberts, A survey and analysis of techniques for player behavior prediction in massively multiplayer online role-playing games, IEEE Transactions on Emerging Topics in Computing 3 (2014) 260–274.

[3] E. Y. Ha, J. P. Rowe, B. W. Mott, J. C. Lester, Recognizing player goals in open-ended digital games with Markov logic networks, in: G. Sukthankar, R. Goldman, C. Geib, D. Pynadath, H. Bui (Eds.), Plan, Activity and Intent Recognition: Theory and Practice, 2014, pp. 289–311.

[4] J. Harris, R. M. Young, Proactive mediation in plan-based narrative environments, IEEE Transactions on Computational Intelligence and AI in games 1 (2009) 233–244.

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

[6] A. Ram, Knowledge goals: A theory of interestingness, in: Cognitive Science Society Annual Conference, 1990, pp. 206–214.

[7] A. Baikadi, J. Rowe, B. Mott, J. Lester, Generalizability of goal recognition models in narrative-centered learning environments, in: International Conference on User Modeling, Adaptation, and Personalization, 2014, pp. 278–289.

[8] A. Gómez, L. Márquez, H. Zapa, M. Florez, Gda-based tutor module of an intelligent tutoring system for the personalization of pedagogic strategies, in: International Conference on Intelligent and Interactive Systems and Applications, Springer, 2020, pp. 742–750.

[9] R. Cardona-Rivera, R. Young, Symbolic plan recognition in interactive narrative environments, in: AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, volume 11, 2015, pp. 16–22.

[10] S. G. Ware, Mutual implicit question answering for shared authorship: a pilot study on player ex-

pectations, in: Intelligent Narrative Technologies Workshop, volume 10, 2017, pp. 259–265.
[11] C. Roberts, Information structure in discourse: Towards an integrated formal theory of pragmatics, Semantics and Pragmatics 5 (2012) 1–69.
[12] A. C. Graesser, S. E. Gordon, L. E. Brainerd, QUEST: A model of question answering, Computers & Mathematics with Applications 23 (1992) 733–745.
[13] F. R. Meneguzzi, R. F. Pereira, A survey on goal recognition as planning, in: International Joint Conference on Artificial Intelligence, 2021, pp. 4524–4532.
[14] J. Robertson, R. M. Young, Perceptual experience management, IEEE Transactions on Games 11 (2018) 15–24.
[15] B. Sunshine-Hill, N. I. Badler, Perceptually realistic behavior through alibi generation, in: AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, 2010.
[16] L. Flores, D. Thue, Level of detail event generation, in: International Conference on Interactive Digital Storytelling, volume 10, 2017, pp. 75–86.
[17] I. Swartjes, E. Kruizinga, M. Theune, Let's pretend I had a sword: Late commitment in emergent narrative, in: International Conference on Interactive Digital Storytelling, volume 1, 2008, pp. 264–267.
[18] D. Thue, S. Schiffel, T. Þ. Guðmundsson, G. F. Kristjánsson, K. Eiríksson, M. V. Björnsson, Open world story generation for increased expressive range, in: International Conference on Interactive Digital Storytelling, 2017, pp. 313–316.
[19] I. Horswill, Retcon: a least-commitment story-world system, in: Experimental AI in Games Workshop, volume 9, 2022.
[20] J. Rowe, B. Mott, S. McQuiggan, J. Robison, S. Lee, J. Lester, Crystal Island: A narrative-centered learning environment for eighth grade microbiology, in: International Conference on Artificial Intelligence in Education Workshop on Intelligent Educational Games, 2009, pp. 11–20.
[21] R. Farrell, S. G. Ware, Narrative planning for belief and intention recognition, in: AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, volume 16, 2020, pp. 52–58.
[22] A. Shirvani, R. Farrell, S. G. Ware, Combining intentionality and belief: revisiting believable character plans, in: AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, volume 14, 2018, pp. 222–228.
[23] A. Amos-Binks, R. E. Cardona-Rivera, Goal elicitation planning: Acting to reveal the goals of others, in: Advances in Cognitive Systems, 2020.
[24] R. E. Cardona-Rivera, R. M. Young, Games as conversation, in: AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, volume 10, 2014, pp. 2–8.
[25] B. Magerko, W. Manzoul, M. Riedl, A. Baumer, D. Fuller, K. Luther, C. Pearce, An empirical study of cognition and theatrical improvisation, in: ACM Conference on Creativity and Cognition, volume 7, 2009, pp. 117–126.
[26] R. C. Stalnaker, Assertion, in: Syntax and Semantics, volume 9, New York Academic Press, 1978, pp. 315–332.
[27] A. Rothe, B. M. Lake, T. M. Gureckis, Do people ask good questions?, Computational Brain & Behavior 1 (2018) 69–89.
[28] A. C. Graesser, K. L. Lang, R. M. Roberts, Question answering in the context of stories, Journal of Experimental Psychology: General 120 (1991) 254.
[29] D. B. Christian, R. M. Young, Comparing cognitive and computational models of narrative structure, in: National Conference of the American Association for Artificial Intelligence, volume 19, 2004, pp. 385–390.
[30] R. E. Cardona-Rivera, T. Price, D. Winer, R. M. Young, Question answering in the context of stories generated by computers, Advances in Cognitive Systems 4 (2016) 227–245.
[31] A. C. Graesser, P. J. Byrne, M. L. Behrens, Answering questions about information in databases, Questions and Information Systems (1992) 229–252.
[32] M. Ramírez, H. Geffner, Plan recognition as planning, in: International Joint Conference on Artificial Intelligence, 2009, pp. 1778–1783.
[33] S. G. Ware, C. Siler, Sabre: A narrative planner supporting intention and deep theory of mind, in: AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, volume 17, 2021, pp. 99–106.
[34] D. Michie, Game-playing and game-learning automata, in: Advances in Programming and Non-Numerical Computation, 1966, pp. 183–200.
[35] M. O. Riedl, R. M. Young, From linear story generation to branching story graphs, in: AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, 2005, pp. 111–116.
[36] M. O. Riedl, R. M. Young, Narrative planning: balancing plot and character, Journal of Artificial Intelligence Research 39 (2010) 217–268.
[37] J. Robertson, R. M. Young, A model of superposed states, in: AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, volume 12, 2016, pp. 65–71.
[38] M. Fisher, C. Siler, S. G. Ware, Intelligent de-escalation training via emotion-inspired narrative planning, in: Intelligent Narrative Technologies Workshop, 2022.
[39] S. G. Ware, E. T. Garcia, M. Fisher, A. Shirvani, R. Farrell, Multi-agent narrative experience man-

agement as story graph pruning, IEEE Transactions on Games (2022).
[40] C. Siler, Open-world narrative generation to answer players' questions, in: AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, volume 18, 2022, pp. 307–310.
[41] A. Ram, A theory of questions and question asking, Journal of the Learning Sciences 1 (1991) 273–318.
[42] A. C. Graesser, W. Baggett, K. Williams, Question-driven explanatory reasoning, Applied Cognitive Psychology 10 (1996) 17–31.