

The Sabre Narrative Planner: Multi-Agent Coordination with Intentions and Beliefs

Extended Abstract

Stephen G. Ware
Narrative Intelligence Lab
University of Kentucky
Lexington, Kentucky
sgware@cs.uky.edu

Cory Siler
Narrative Intelligence Lab
University of Kentucky
Lexington, Kentucky
jcsi225@uky.edu

ABSTRACT

Narrative planning algorithms coordinate the virtual agents of interactive environments. Sabre is a single, centralized, omniscient decision maker that solves a multi-agent problem. It has a system-level author goal it must achieve, but every action taken by an agent must make sense according to that agent's individual intentions and limited, possibly wrong beliefs. We describe our motivation for solving such problems and the difficulties of comparing our planner to existing systems.

KEYWORDS

Narrative; Intention; Belief; Planning; Social Simulation

ACM Reference Format:

Stephen G. Ware and Cory Siler. 2021. The Sabre Narrative Planner: Multi-Agent Coordination with Intentions and Beliefs: Extended Abstract. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021)*, Online, May 3–7, 2021, IFAAMAS, 3 pages.

1 MOTIVATION

Planning was originally formulated as a problem where one decision maker coordinates all resources and anticipates a sequence of actions to achieve a goal. It assumes a fully observable, discrete, deterministic world. These assumptions rarely hold in the real world, and multi-agent systems have been developed to model many decision makers, partial observability, continuous state features, nondeterminism, etc.

Virtual environments that feature interactive stories, like video games, training simulations, intelligent tutoring systems, and virtual reality therapy, present an interesting case that can benefit from a combination of single and multi-agent approaches. These systems typically invite the player to take on the role of one character while the system controls a host of virtual non-player characters. These virtual characters need to act like realistic agents with their own goals, limited observations, and possibly wrong beliefs. However, as long as the illusion of realism is maintained, the system is free to use a centralized storytelling agent that is omniscient and omnipotent in the virtual world to coordinate these virtual characters to achieve the system-level author goals for the aesthetic or pedagogic structure of the story.

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online. © 2021 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Narrative planning is an interesting and challenging problem because it can be performed by a single agent but needs to generate a solution befitting a multi-agent system, what Durfee calls centralized planning for distributed plans [5]. A traditional planner can ensure the system-level goal is met but does not reason about agent realism. A traditional multi-agent system can ensure realistic behavior by giving agents limitations like partial observability they don't actually suffer in a virtual world, but the system must be coordinated toward the author's goals.

Our previous work [18] motivated this model of centralized agent-based planning. Using five narrative test domains, we enumerated the total number of legal action sequences and counted how many (1) featured only intentional character behavior, (2) achieved the author goal, and (3) did both—i.e. achieved the author goals via intentional character behavior. Results varied by domain, but it was always the case that either set #1 or #2 or both were significantly larger than #3, indicating the value of centralized coordination of realistic-seeming agents.

2 SABRE'S FEATURES

Sabre is a forward-chaining state-space narrative planning framework that models both the intentions of the author (i.e. the system designer's constraints) as well as the intentions and beliefs of each virtual character. This abstract briefly describes how the problems it solves differ from previous planners and why this makes it difficult to compare our framework to others.

Kybartas and Bidarra [8] survey approaches to narrative generation. While several researchers have used planning to control interactive narrative systems [12], Sabre explicitly incorporates computational models of character and narrative structure into the algorithm. Young et al. [25] survey other similar narrative planners.

These systems face a tradeoff: as the richness of the model increases the scope of problems that can be solved decreases. Sabre is designed for coordinating narratives in virtual worlds that typically feature 2 to 5 characters. There is a system-level author goal that must be met. Characters have their own goals and beliefs, and they observe change. Character goals can synergize or conflict with others, with the author, and with the player. A typical plan is between 4 and 12 actions. We briefly survey other systems that handle similar problems to highlight how Sabre is similar or different.

Intention

Riedl and Young [13] introduced intentional planning, which defines both author and character goals. In classical planning, actions

define preconditions and effects. Intentional planning annotates each action with a list of zero to many characters who are intentionally taking the action. A valid plan must accomplish the author goal but can only be composed of actions that contribute to the goals of the characters who take them, e.g. a character wants medicine, so he goes to the market and buys it. Sabre uses Ware et al.’s [22, 23] extension to this model which allows failed plans and conflict. A character action no longer needs to be part of a successfully executed plan to achieve their goal; as long as such a plan exists, the action is reasonable, even if that plan is never actually executed. A character who wants medicine can go to the market even if their plan to buy it is interrupted. Rather than traditional propositional goals, Sabre uses utility functions for the author and each character that allow for complex preferences over states.

Theory of Mind

Virtual Storyteller [2], HeadSpace [21], and Christensen et al.’s extension of Glaive [3] all use a 1 layer theory of mind, meaning they reason about what is true, and what each character believes is true, but not what x believes y believes, and so on. IMPRACTical [19] uses a 1 layer model, and past that defers to a shared global set of popular beliefs. Zunshine [26] argues that multiple levels are essential for fiction, and our previous work [16] showed a multi-level theory of mind is needed to model deception, cooperation, anticipation, and surprise. Sabre places no arbitrary limit on the depth of theory of mind.

Si and Marsella’s Thespian [17] models multi-level theory of mind, and Ryan et al.’s Talk of the Town [14] models characters that observe, misremember, and lie. These and others like them are multi-agent systems; agents have true partial observability and leverage little or no centralized planning to coordinate the story. Conversely, some centralized planning algorithms model agent beliefs for real world situations [11] that make assumptions unhelpful for narrative problems, e.g. that agents always cooperate [6], always compete [4], that ignorance is always bad [1], etc.

Intention + Theory of Mind

Sabre builds on our previous work [15] which demonstrated that reasoning about intention and a multi-layer theory of mind together produce more realistic agent behavior. At least one other system, Ostari [9], models agent intentions and a multi-layer theory of mind for centralized planning. However, Ostari models true uncertainty, e.g. agents can have complex disjunctive beliefs, e.g. the merchant is at the market or in her house. The high cost of modeling all doxastically accessible possible worlds limits the scope of problems Ostari can solve. Sabre allows arbitrarily nested and wrong beliefs, but agents must always commit to specific beliefs, e.g. a character can believe the merchant is in the market when she is actually in her house, but cannot entertain several possibilities about the merchant’s location. We have found this an acceptable tradeoff between model complexity and problem scope for our needs.

Observations and Triggers

Sabre supports full ADL syntax [10] for defining preconditions and effects, including conjunction, disjunction, conditional effects, and first order quantifiers. Actions specify conditions under which they

are observed, e.g. when someone buys medicine in the market, everyone in the market sees it happen. When a character observes an action, they update their beliefs about the current state; otherwise their beliefs do not change. Sabre also supports a modal *believes* predicate (which can be infinitely nested) so that preconditions can require certain beliefs and effects can modify beliefs explicitly. Finally, Sabre supports triggers, which also have preconditions and effects but which must occur when they can. Triggers are often used for observations, e.g. when characters x and y are both in the market but x believes y is somewhere else, x now believes y is in the market. Triggers are similar to axioms in PDDL planners [7], but operate directly on state fluents instead of special derived fluents. They are also similar to sensory actions in systems that mix planning, sensing, and execution [24], in the sense the Sabre can plan for a character to discover something unexpected, e.g. a character goes to the market to buy medicine and discovers that a guard is also at the market.

3 FAIR COMPARISONS

Sabre has a unique set of features: it is a centralized planner, reasons about intentions and beliefs, imposes no limit on theory of mind, and allows wrong beliefs but without uncertainty. We find these features important for our interactive narratives, but they make it difficult to benchmark against other algorithms.

Intentional planners [13, 19, 22, 23] do not reason about agent beliefs, so either Sabre will be burdened reasoning about belief when it is not required, or the problems will be unsolvable by the intentional planners—hardly a fair comparison. A similar argument applies to planners that limit theory of mind to 1 layer [3, 20, 21]. In any case, only Glaive [23] and Christensen et al.’s extension to it [3] has available implementations or test domains. The same applies for multi-agent story generation systems with limited centralized planning [14, 17]. Either Sabre is burdened with the author’s goal, or those systems would need to run many times until they happen to achieve the author’s goal, and they were not designed for that.

Ostari [9] is centralized, reasons about both intentions and beliefs, has been used to generate stories, and has a publicly available implementation. However, it reasons about true uncertainty (i.e. over all doxastically possible worlds). As a test, the smallest example from one of our test domains was implemented and tested in Ostari, but it runs out of memory before it can reason about a plan long enough to be a solution. Again, this comparison seems unfair.

We have curated a suite of benchmark narrative planning domains from several authors, adapting them to include intention and belief or remove uncertainty as necessary. Sabre’s performance on this suite is outside the scope of this abstract, but we intend to release a full implementation along with these benchmarks. Since implementations of similar systems are rarely available, we intend to establish Sabre’s performance via an ablation study, where various features are disabled to demonstrate the tradeoff between model richness and problem scope.

ACKNOWLEDGMENTS

This work was supported by the U.S. National Science Foundation under Grant No. IIS-1911053. Opinions, findings, and conclusions are those of the authors and do not reflect the views of the NSF.

REFERENCES

- [1] Thomas Bolander and Mikkel Birkegaard Andersen. 2011. Epistemic planning for single-and multi-agent systems. *Journal of Applied Non-Classical Logics* 21, 1 (2011), 9–34.
- [2] Hans ten Brinke, Jeroen Linssen, and Mariët Theune. 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.
- [3] Matthew Christensen, Jennifer M. Nelson, and Rogelio E. Cardona-Rivera. 2020. Using domain compilation to add belief to narrative planners. In *Proceedings of the 16th AAAI international conference on Artificial Intelligence and Interactive Digital Entertainment*. 38–44.
- [4] Fiorella De Rosis, Valeria Carofiglio, Giuseppe Grassano, and Cristiano Castelfranchi. 2003. Can computers deliberately deceive? A simulation tool and its application to Turing’s imitation game. *Computational Intelligence* 19, 3 (2003), 235–263.
- [5] Edmund H. Durfee. 2001. *Distributed problem solving and planning*. Lecture Notes in AI, Vol. 2086. 118–149.
- [6] Barbara Grosz and Sarit Kraus. 1996. Collaborative plans for complex group action. *Artificial Intelligence* 86, 2 (1996), 269–357.
- [7] Malte Helmert. 2006. The Fast Downward planning system. *Journal of Artificial Intelligence Research* 26 (2006), 191–246.
- [8] Quinn Kybartas and Rafael Bidarra. 2016. A survey on story generation techniques for authoring computational narratives. *IEEE Transactions on Computational Intelligence and Artificial Intelligence in Games* 9, 3 (2016), 239–253.
- [9] Henry Mohr, Markus Eger, and Chris Martens. 2018. Eliminating the impossible: a procedurally generated murder mystery. In *Proceedings of the Experimental Artificial Intelligence in Games workshop at the 14th AAAI international conference on Artificial Intelligence and Interactive Digital Entertainment*.
- [10] Edwin P. D. Pednault. 1987. Formulating multiagent, dynamic-world problems in the classical planning framework. In *Reasoning About Actions & Plans*. Elsevier, 47–82.
- [11] Martha E. Pollack. 1986. A model of plan inference that distinguishes between the beliefs of actors and observers. (1986), 207–214.
- [12] Julie Porteous, Marc Cavazza, and Fred Charles. 2010. Applying planning to interactive storytelling: Narrative control using state constraints. *ACM Transactions on Intelligent Systems and Technology* 1, 2 (2010), 1–21.
- [13] Mark O. Riedl and R. Michael Young. 2010. Narrative planning: balancing plot and character. *Journal of Artificial Intelligence Research* 39, 1 (2010), 217–268.
- [14] James Owen Ryan, Adam Summerville, Michael Mateas, and Noah Wardrip-Fruin. 2015. Toward characters who observe, tell, misremember, and lie. In *Proceedings of the Experimental Artificial Intelligence in Games workshop at the 13th AAAI international conference on Artificial Intelligence and Interactive Digital Entertainment*. 56–62.
- [15] Alireza Shirvani, Rachelyn Farrell, and Stephen G. Ware. 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.
- [16] Alireza Shirvani, Stephen G. Ware, and Rachelyn Farrell. 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.
- [17] Mei Si and Stacy C. Marsella. 2014. Encoding Theory of Mind in character design for pedagogical interactive narrative. *Advances in Human-Computer Interaction* (2014).
- [18] Cory Siler and Stephen G. Ware. 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*. 123–129.
- [19] Jonathan Teutenberg and Julie Porteous. 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.
- [20] Jonathan Teutenberg and Julie Porteous. 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.
- [21] Brandon R. Thorne and R. Michael Young. 2017. Generating stories that include failed actions by modeling false character beliefs. In *Proceedings of the Intelligent Narrative Technologies workshop at the 13th AAAI international conference on Artificial Intelligence and Interactive Digital Entertainment*. 244–251.
- [22] Stephen G. Ware and R. Michael Young. 2011. CPOCL: a narrative planner supporting conflict. In *Proceedings of the 7th AAAI international conference on Artificial Intelligence and Interactive Digital Entertainment*. 97–102.
- [23] Stephen G. Ware and R. Michael Young. 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.
- [24] Daniel S. Weld, Corin R. Anderson, and David E. Smith. 1998. Extending graphplan to handle uncertainty & sensing actions. In *The fifteenth national AAAI conference on Artificial Intelligence*. 897–904.
- [25] R. Michael Young, Stephen G. Ware, Bradley A. Cassell, and Justus Robertson. 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 (2013), 41–64.
- [26] Lisa Zunshine. 2006. *Why we read fiction: theory of mind and the novel*. Ohio State University Press.