Logic foundations of manipulation as game mechanics

Javier M. Torres Ramón
Brainific SA
Madrid, Spain
javier.m.torres@brainific.com

Abstract— **Several subfields in symbolic AI can be useful in game agents: dynamic epistemic logic, epistemic planning, belief revision and natural language processing. Furthermore, there is a sizeable number of papers exploring their interactions, but very few practical implementations. We present WemblAI, a toolbox that aims at collecting algorithms from these fields using a common framework. WemblAI is still largely a work in progress.**

Keywords— dynamic epistemic logic, epistemic planning, belief revision, natural language processing, multi agent systems

I. INTRODUCTION

Interactions with actors in games have previously included planning (Republique in [25]), reactive planning and utility systems (The Sims in [19], Versu in [16]), reasoning using first order logic (MKULTRA in [23] and [22]) or exclusion modal logic (Versu in [16]), and natural language generation and parsing (Façade in [27], Talk of the Town in [33]) within the game mechanics. More recent advances in symbolic AI and close fields may provide interesting mechanics and increased believability, especially when combined. The present paper shows how several of these techniques can be blended together in a single software package called WemblAI and describes new possibilities arising from their interaction: dynamic doxastic logic, belief revision, epistemic planning and natural language processing. The aim of this project is to provide a package that can cater for the needs of complex characters (e.g. RPGs or interactive fiction); the development of a sample game to demonstrate its capabilities is under study. Although some mechanisms are already in place, WemblAI is still a work in progress.

II. BACKGROUND

A. *Dynamic epistemic logic*

Dynamic epistemic logic (DEL), as described by van Ditmarsch, Hoek and Kooi in [11], combines two important aspects of agency: knowledge (and belief, as "fallible knowledge") and action consequences. In an example formulation, the language is defined as follows. Given two sets of agents A and atoms P, we will define two languages, one defining formulas and one defining actions:

$$\phi ::= p \mid \neg\phi \mid (\phi \wedge \phi) \mid Ka\phi \mid CB\phi \mid [\alpha]\phi$$

$$\alpha ::= (M, s) \mid (\alpha \cup \alpha)$$

The language thus allows us to specify several types of operators:

- Classic logic operators;
- Knowledge and common knowledge operators, which restrict the epistemic models for an agent or for a set of agents with the target proposition, and which can be replaced by belief operators to allow for incorrect information; and
- Dynamic operators, which specify an action and define the successor epistemic state (or possible states, for with nondeterministic actions) of an agent after that action has been executed, provided that the current state did in fact comply with its preconditions.

In the action language, M is defined as a tuple of: an action point (which basically identifies an action), an equivalence relation (which defines which actions outcomes are indistinguishable for each agent) and a set of preconditions in the logical language. The effects of distinguishable actions on distinguishable states as successor states will also be distinguishable for an agent.

B. *Epistemic planning*

The automated planning community started adopting dynamic epistemic logic as the base formalism for epistemic planning, a novel approach to planning under partial observability and nondeterminism, adding another aspect of agency like goals and plans. The Dagstuhl seminar on epistemic planning in [8] describes several problems to which it can be applied, some of which can be very appealing to games: cooperative problems, security games, adversary game playing and multi agent deception. Eger and Martens have already provided some examples of its uses in games in [13], [14] and [15].

C. *Belief revision*

Belief revision (described in [20] and [21]) is a useful way to formalize nonmonotonic reasoning. It organizes beliefs in a hierarchical set or base (depending on whether a belief in one layer implies beliefs in other layers) and methods to add new beliefs so the most plausible layer is always consistent, possible resulting in beliefs moving up and down the hierarchy. It appeared in the previously mentioned fields, as epistemic planning started considering plausibilities in action outcomes and Baltag and Smets defined Conditional Doxastic Logic (CDL) in [6] and [7], a generalization of well-known layered beliefs models based on their version of Dynamic Doxastic Logic. Belief revision has been extensively developed, but not incorporated into games, possibly because reasoning is perceived as complex and time consuming. Multi agent systems like Jason have adapted such algorithms as described in [1] and [2], but in this field the emphasis is put in cooperation and predictability. Fig. 1 shows two examples of different expansion methods, where accepting a new proposition results in different possible truth sets depending on the method (lower layer, p={T} and q={T} on the right vs. p={T,F} and q={T} on the left).

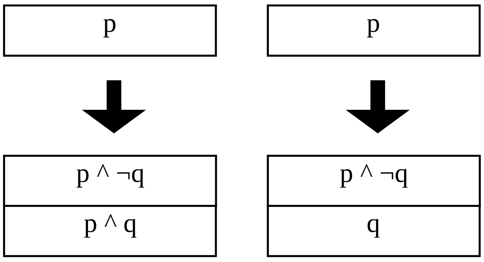

Fig. 1. Belief revision: moderate expansion (left) vs conservative expansion (right)

D. Symbolic Natural Language Processing (NLP)

The use of symbolic tools for NLP analysis and synthesis like the Grammatical Framework described in [31]. can take these algorithms one step beyond: at-the-core formal statements can be interchanged using a controlled language giving the illusion of natural language, as has been done with description logics in [26]. The use of parsers with ambiguous representations allows a player to influence a character by feeding them with information that can be interpreted in different ways. The ambitious effort of the Grammatical Framework project already provides a huge library of terms and grammatical constructions in several languages, with ties to the WordNet project (in [30]).

E. Integration in the WemblAI package

The author started the *WemblAI*[1] and *heroes-cant*[2] projects to collect and implement software that can be embedded in games and perform dynamic doxastic logic reasoning and epistemic planning under belief revision for more complex NPCs allowing more rewarding interactions. The amount of algorithms described in literature is impressive, but few of them have software implementations. For example, Girlando, Negri et al. developed a sequent-calculus-based system to check sentences in CDL in [18]. Rott describes in [32] nearly 30 methods for revision of belief sets. A dynamic doxastic logic tableau was provided by Aucher and Schwarzentruber in [3] while establishing its complexity, and at least Andersen, Bolander et al. in [3] and Huang, Fang et al. in [24] provide epistemic planning algorithms, a simple breadth-first search ranked by action plausibility and an optimized planner with a custom epistemic representation. A notable mention goes to the DEMO model checker software by van Eijck in [16].

III. ARCHITECTURE

In WemblAI, an agent's mind consists of the following elements:

- A historic list of timestamped observations. Observations can be direct tests of atomic ontic (ie facts) propositions ("Saw p", the agent perceives that p is true, or "Watched B a", the agent perceives that B is performing action a) or records of sentences from other agents ("Told B (Believes C q)", agent B tells the agent that agent C believes that q is true, "Asked C p", agent C asks the agent whether p is true, or "Requested D z", where agent D requests the agent to take action z).
- A store of conjoined sentences as current beliefs, ordered by plausibility and keeping the source that originated the belief (another agent, past observations or even statistical inference). These sentences are written in dynamic doxastic logic, hence incorporating doxastic ("believes") and dynamic ("after action") modalities. States are only expanded when needed during reasoning and planning. A CNF is built from selected sentences, and it is preferred to DNF to defer state replication in stages like planning as much as possible. There are some data structures dealing with doxastic sentences already available as in [24], but this prevents handling sentences like "B(a) ¬q v [α]p" that target at both doxastic and dynamic modalities.
- A cache of dynamic operators and beliefs about other agents' minds. This cache is only valid as long as the appropriate support is held, and it may be necessary to rebuild it every time a new datum is received by the agent. When reasoning, the beliefs about the other agent's beliefs are expanded into full blown instances of this same architecture. Likewise, dynamic operators in the agent's mind are used to expand "next states" in planning. Note that keeping operators in a declarative representation allows the agent to dynamically update the consequences ad plausibilities of a certain action, be it through observation and inference, or information exchange with other agents. Keeping track of the validity of the elements in this cache is right now an open issue.

In Fig. 2, a sample agent mind is shown, containing, from top to bottom, the list of timestamped observations (higher timestamp is most recent, at the bottom), hierarchical beliefs with sources (higher plausibility at the bottom) and a cached model of C's mind (unknown observations, p is believed with unknown plausibility related to other beliefs).

The following mechanisms are being implemented:

- The doxastic dynamic logic tableau from [5]. Unfortunately, this tableau needs to be extended to consider ontic, and not only epistemic, action postconditions. Z3 (in [29]) and other options are being considered as a final step in non-modal sentences to accelerate reasoning, or alternatively as a source of algorithms, like using CDCL (Conflict Driven Clause Learning) as opposed to the former DPLL (Davis-Puttnam-Logemann-Loveland) algorithm.
- The epistemic planning algorithm from Andersen and Bolander in [3], which in essence performs a breadth-first search on the state tree. [24] incorporates the PrAO planning algorithm from [34] and will be inspected as a possible improvement on this algorithm.
- Part of the belief revision algorithms from [32], mainly expansions, updates and contractions, without

---

[1] https://bitbucket.org/brainific/wemblai/src/master/

[2] https://bitbucket.org/brainific/heroes_cant/src/master/

considering right now more complex alternatives that take into account the current contents of the belief set.

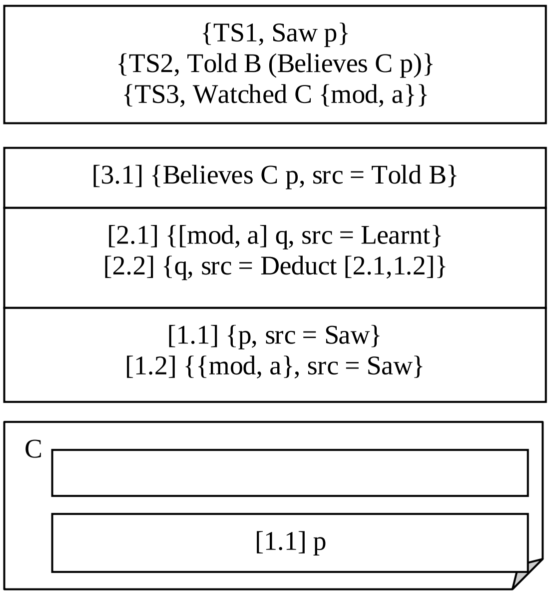

Fig. 2. Sample contents of an agent's mind

- A reversible parser implemented using the Grammatical Framework that generates controlled language sentences from logic sentences. This parser can now represent dynamic ("if Aisha tells someone that BoYang has a dagger then Aisha will be a thief") and doxastic ("BoYang believes that Aisha doubts that BoYang has the shield ") sentences in a controlled language.

if Aisha tells someone that BoYang has a dagger then Aisha will be a thief

DynamicMod1 (ConcrAction Aisha (TellAct Someone (PresentSentence State (Has ADet Weapon BoYang Dagger)))) (IsA Aisha Thief)

BoYang believes that Aisha doubts that BoYang has the shield

PresentSentence State (DoxasticMod BoYang 0 Believe (PresentSentence State (DoxasticMod Aisha 2 Doubt (PresentSentence State (Has TheDet Armor BoYang Shield)))))

Fig. 3. Sample parsing of two (correct but nonsensical) sentences including doxastic and dynamic operators in *heroes_cant* formal representation

IV. DISCUSSION

The epistemic planner and the tableau together should be enough to face games with epistemic aspects like Cluedo or Hanabi, taking advantage of the ability to reason over own and others' epistemic states. In Cluedo, a player can keep track of what combinations other players are asking and infer possible initial epistemic states for them (what cards they know) from these actions and their goal of finding out the correct combination. There are some areas that still need further research, like event specification; for example, different versions of DEL allow different ways to combine actions, e.g. only serial composition vs. branching, or even concurrent actions like in [11].

When implementing a planning algorithm, some issues also come to light. Planning algorithms often use add/delete lists as in STRIPS, i.e. updating an atomic value when the result of an action is applied. However, a postcondition as an arbitrary dynamic doxastic logic formula is harder to handle. It may be handled like a radical update, placing the formula at the base of the belief set in a similar way to an add/delete effect, but since according to [3] postconditions maybe be also plausibility ordered, it is not evident how the different outcomes may be accommodated. This plausibility ordering may also guide replanning stages if the most plausible outcome did not turn out in the end. This would be equivalent to contingent planning (or a behavior tree solving the equivalent planning problem) but with the added benefit of keeping more related information about the alternate outcomes.

Belief revision allows for interesting game mechanics when mixed with the previous two aspects. For example, and like MKULTRA ([23]), a player may want to change the belief set of an agent so that they reach a decision beneficial to the player. However, this may not be done directly, but instead presenting information or evidence that first must earn the trust of the other agent. As belief layers are relative, it may also be easier to devalue current sources of information, so that other conflicting information is used instead. This method does not even rely on presenting information on the matter at hand.

Blending in NLP, the previous devaluation may be performed using ad hominem methods, where a specific rhetoric skill is needed to cause the other agent's reasoning engine to disregard some source. A low value in such skill may result in the player's information to be devalued instead. The player must include a well-formed sentence that represents a negative valuation of a source known to the agent. The devalued belief can be an ontic formula (some real state), a doxastic formula (what other agent thinks) or even a dynamic formula (the outcomes of an action), which clearly affects the result of a planning algorithm. To change the cognitive status of an actor and possibly cause a change in action the player must first probe how they update their bases: do they trust external sources more, or rather their own reasoning? And do they trust all their sources equally? If a new information comes in, is all contradictory information thrown out, or some effort is spent reconciling past information as well?

Symbolic systems in restricted languages can be fairly efficient as the Grammatical Framework has shown. It also provides all alternative parsings of a sentence given a certain grammar; validation using categories; and a two-tier architecture that allows easy replacement of linearizations for multilingual development (including formal languages). Interactions between agents can be linearized in a controlled English for an illusion of natural language, but the grammar can also be used to fill in sentences in a word-by-word approach, or making use of templates, both in a suitable middle ground between scripts and full-fledged NLP and easier to use for novel players. The use of an explicit "deep structure" also makes it more manageable than a trainable NLP deep learning model.

One important drawback to keep in mind is the complexity of the algorithms involved. For example, the satisfiability problem of the tableau used is NEXPTIME-hard. In [3] the plan existence problem of multi agent epistemic planning (i.e. whether a plan exists for a multi-agent planning task) is proven to be undecidable, and the author also notes that the plan existence problem in conditional planning with nondeterministic actions for partially observable domains is 2-EXP-complete. However, three circumstances may be helpful: 1) the increasing computing power and memory available, 2) the need to keep situations reasonably small and manageable for humans to solve them, and 3) the option to resort to other heuristics if a solution cannot be found in a reasonable time (a.k.a. “smoke and mirrors”).

Finally, the development of a sample game showcasing these aspects is being studied. Regarding the base mechanics for this game, the mix of doxastic logic, planning, belief revision and natural language lends itself well to games with high uncertainty, many actors and complex interactions with information exchange. A modified version of the Blades in the Dark RPG or the Honor among Thieves boardgame seem appropriate alternatives.

V. CONCLUSION

In this paper, we have presented an overview of subfields within symbolic AI that can provide meaningful capabilities to agents in games: dynamic epistemic logic, epistemic planning, belief revision and NLP. After reviewing the existing literature, we find many algorithms that share themes and subjects but very few implementations. We have started the WemblAI project to collect and provide a common framework to the algorithms in these areas, so their combination may be the basis for believable agent behavior with depth.

ACKNOWLEDGMENT

The author would like to thank Zans Mihejevs and Claudia Doppioslash for their ongoing discussion. Also, he would like to thank Santiago Ontañón for his review and support.

REFERENCES

[1] Alechina, N., Bordini, R. H., Hübner, J. F., Jago, M., & Logan, B. (2006, May). Belief revision for AgentSpeak agents. In Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems (pp. 1288-1290). ACM.

[2] Alechina, N., Bordini, R. H., Hübner, J. F., Jago, M., & Logan, B. (2006, May). Automating belief revision for agentspeak. In International Workshop on Declarative Agent Languages and Technologies (pp. 61-77). Springer, Berlin, Heidelberg.

[3] Andersen, M. B., & Bolander, T. (2015). Towards theory-of-mind agents using automated planning and dynamic epistemic logic.

[4] Andersen, M. B., Bolander, T., & Jensen, M. H. (2015). Don't plan for the unexpected: Planning based on plausibility models. Logique et Analyse, 58(230), 145-176.

[5] Aucher, G., & Schwarzentruber, F. (2012). On the Complexity of Dynamic Epistemic Logic (Extended Version) (Doctoral dissertation, INRIA).

[6] Baltag, A., & Smets, S. (2008). The logic of conditional doxastic actions. Texts in Logic and Games, Special Issue on New Perspectives on Games and Interaction, 4, 9-31.

[7] Baltag, A., & Smets, S. (2008). A qualitative theory of dynamic interactive belief revision. Logic and the foundations of game and decision theory (LOFT 7), 3, 9-58.

[8] Baral, C., Bolander, T., van Ditmarsch, H., & McIlrath, S. (2017). Epistemic planning (Dagstuhl seminar 17231). In Dagstuhl Reports (Vol. 7, No. 6). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

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

[10] Bolander, T., & Gierasimczuk, N. (2015, October). Learning actions models: Qualitative approach. In International Workshop on Logic, Rationality and Interaction (pp. 40-52). Springer, Berlin, Heidelberg.

[11] van Ditmarsch, H. P., Van Der Hoek, W., & Kooi, B. P. (2003). Concurrent dynamic epistemic logic. In Knowledge contributors (pp. 105-143). Springer, Dordrecht.

[12] van Ditmarsch, H., van Der Hoek, W., & Kooi, B. (2007). Dynamic epistemic logic (Vol. 337). Springer Science & Business Media.

[13] Eger, M. (2017, September). Intentional Agents for Epistemic Games. In Thirteenth Artificial Intelligence and Interactive Digital Entertainment Conference.

[14] Eger, M. (2018). Intentional Agents for Doxastic Games.

[15] Eger, M., & Martens, C. (2017, September). Practical specification of belief manipulation in games. In Thirteenth Artificial Intelligence and Interactive Digital Entertainment Conference.

[16] van Eijck, J. (2007, May). a demo of epistemic modelling. In Interactive Logic. Selected Papers from the 7th Augustus de Morgan Workshop, London(Vol. 1, pp. 303-362).

[17] Evans, R., & Short, E. (2014). Versu—a simulationist storytelling system. IEEE Transactions on Computational Intelligence and AI in Games, 6(2), 113-130.

[18] Girlando, M., Negri, S., Olivetti, N., & Risch, V. (2018). The logic of conditional beliefs: neighbourhood semantics and sequent calculus.

[19] Graham, D. (2013, September). An Introduction to Utility Theory. In Game AI Pro 1 (pp. 47-60). CRC Press.

[20] Gärdenfors, P. (Ed.). (2003). Belief revision (Vol. 29). Cambridge University Press.

[21] Gärdenfors, P., Rott, H., Gabbay, D. M., Hogger, C. J., & Robinson, J. A. (1995). Belief revision. Computational Complexity, 63, 6.

[22] Horswill, I. D. (2018, September). Postmortem: MKULTRA, An Experimental AI-Based Game. In Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference.

[23] Horswill, I. D. (2015, September). MKULTRA. In Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference.

[24] Huang, X., Fang, B., Wan, H., & Liu, Y. (2018). A general multi-agent epistemic planner based on higher-order belief change. arXiv preprint arXiv:1806.11298.

[25] Lafeta, P. (2014, July). Stealth AI on Tablets in République . Game/AI Conference 2014.

[26] Marginean, A. (2017). Question answering over biomedical linked data with grammatical framework. Semantic Web, 8(4), 565-580.

[27] Mateas, M., & Stern, A. (2003, March). Façade: An experiment in building a fully-realized interactive drama. In Game developers conference (Vol. 2, pp. 4-8).

[28] Miller, G. A. (1995). WordNet: a lexical database for English. Communications of the ACM, 38(11), 39-41.

[29] De Moura, L., & Bjørner, N. (2008, March). Z3: An efficient SMT solver. In International conference on Tools and Algorithms for the Construction and Analysis of Systems (pp. 337-340). Springer, Berlin, Heidelberg.

[30] Princeton University "About WordNet." https://wordnet.princeton.edu/ Princeton University. 2010.

[31] Ranta, A. (2011). Grammatical framework: Programming with multilingual grammars (Vol. 173). Stanford: CSLI Publications, Center for the Study of Language and Information.

[32] Rott, H. (2009). Shifting priorities: Simple representations for twenty-seven iterated theory change operators. In Towards mathematical philosophy (pp. 269-296). Springer, Dordrecht.

[33] Ryan, J., Mateas, M., & Wardrip-Fruin, N. (2016, September). Characters who speak their minds: Dialogue generation in Talk of the Town. In Twelfth Artificial Intelligence and Interactive Digital Entertainment Conference.

[34] To, S. T., Son, T. C., & Pontelli, E. (2011, March). Contingent planning as and/or forward search with disjunctive representation. In Twenty-First International Conference on Automated Planning and Scheduling