

Dynamic Epistemic Logic in Game Design

Javier M. Torres

Brainific SL
javier.m.torres@brainific.com

Abstract

Dynamic Epistemic Logics are a set of multimodal logics that deal with knowledge and change. We argue that the theory for this formalism is mature enough to start a practical implementation, while at the same time having a sizable amount of theoretical expansions. We also claim that these two factors make it an attractive formalism for new features in games featuring an agent's internal model, like planning for gathering information, acting without revealing one's goals or feeding possibly false information to a set of agents to influence their beliefs, augmenting already existing mechanisms.

Introduction

NPCs in computer games usually react to stimuli in a fairly scripted way, for example using static behavior trees. Planning actions for the future or using memories in a complex way are the exception and not the rule, not to speak about manipulating other agents. This paper explores the use of the Dynamic Epistemic Logic (DEL) (van Ditmarsch et al. 2006) to describe an agent's knowledge, allowing complex inferences about actions and others' knowledge that should give way to emergent narrative situations, using a model checking approach. It has been used to explain inductive puzzles like the "muddy children problem" and to incorporate sensing actions in planning algorithms (Bolander and Andersen 2011). DEL includes both epistemic operators ("A knows that p") and dynamic operators ("action phi brings about p or q") in addition to classical logic operators. DEL allows expressing and solving "inductive puzzles", where the sole expression of ignorance or certainty from an agent does provide information to other agents listening to them.

Although other authors offer different interpretations (like the algebraic approach in Baltag et al. 2007) for their own versions of a dynamic epistemic logic, we will focus on the formalism mentioned above. Some further refinements from the DEL literature are also presented here.

DEL in (Some) Detail

The DEL described by van Ditmarsch, Hoek and Kooi in van Ditmarsch et al. 2006, is a multimodal logic based on a possible worlds interpretation that consists of the product of an epistemic logic, and an action model. The epistemic logic allows us to reason about what an agent knows (including knowledge of itself or other agents) or what all agents know (common knowledge). The action model allows us to produce new epistemic models according to the actions compatible with each possible state, when the action preconditions are met; and the results for each action, possibly nondeterministic and unobservable. Informally, the language combines atomic propositions with the following constructs:

Classical logic operators, like negation or conjunction;

Knowledge and common knowledge operators, which express the notion that an agent or group of agents knows a logic sentence to be true in every world they consider possible, although the actual situation cannot be fully observed; and

Dynamic operators, which specify an action and define (possibly with nondeterministic actions) the successor epistemic (knowledge) states of an agent after that action has been executed, provided that the current state did in fact comply with its preconditions.

DEL follows the line of the logic of Public Announcements, which defines modal operators expressing updates in knowledge. However, it includes more fine grained actions that can both cause factual and epistemic changes, and have nondeterministic effects that may cause the number of epistemic states to actually increase after an action.

Recent developments in model checking through Binary Decision Diagrams like the SMCDEL project (van Benthem et al. 2015) may allow more efficient reasoning in DEL. So far, the few existing tools (like the DEMO model checker in van Eijck 2007) directly create the epistemic models and check the formulas through them. Although model checking and satisfiability may have very tough complexity in terms of space and time (Aucher and

Schwarzentruber 2012), games may be designed with a low enough number of elements to avoid an explosion in the computing power to solve these problems.

DEL is, however, still in active development, with its own theoretical problems with awkward solutions. Its authors discuss some of the more theoretically complex details in the book, and some authors (e.g Herzig 2016) have highlighted the shortcomings of the current status of DEL formalisms.

Related work

Theoretical approaches

The idea of managing internal models of other agents' internal models (the psychological concept of a "theory of mind") is not new. Actually, any adversarial modelling can be considered as producing some kind of "theory" of the adversary's "model". For example, some authors have used POMDPs to model the other agent's behavior (Baker et al. 2011), whereas others (Bosse et al. 2007) propose a two-level Belief-Desire-Intention (BDI) model of the beliefs and intentions of other agents. A more complex model can be found in the PsychSim project (Marsella, Pynadath, and Read 2004), with several psychological constructs modelled into the system based mainly on utility mechanisms (Pynadath and Marsella 2007).

The BDI logic mentioned above (Rao and Georgeff 1995) has traditionally been proposed as a framework for multiagent systems. It does include the notion of belief, and two goal-related notions, those of desires (things that the agent would like to see true) and intentions (some subset of desires). However, reasoning within the implementations of logics in this family tends to be fairly limited; for example, the dMARS architecture (d'Inverno et al. 2004) just selects plans as intentions according to some beliefs (actually facts) from a larger set of plans (the desires). It does not address the issue of selecting an action due to its epistemic actions (so that other plans could be possibly selected) or of using actions like "public announcements" that may influence the actions of other agents.

Closely related, the topic of belief revision deals with changes in the epistemic state of agents. Belief revision is an active research area: even when new knowledge comes, an agent needs to decide which inconsistent beliefs to dump to accommodate to the new situation. This means that an agent may need to retract some inferences and create some others, beyond a mere update of a knowledge base. A short revision of the subject can be found in van Ditmarsch 2006.

Games with knowledge and reasoning mechanics

First of all, we must emphasize the DEL in itself does not render the existing following approaches obsolete. It rather allows a certain kind of reasoning (about actions and mental states) that provides new gameplay possibilities to all these examples, enriching the experiences described here, but also creates new opportunities.

Versu (Evans and Short 2014, Evans and Short 2015) implements an open-ended narrative simulation based on two main concepts: 1) a deontic exclusion logic, which expresses different concepts like goals and norms through the use of a deontic "should" operator, and in which "rules" of the game are implemented with a temporal "next" operator; and 2) utility based agents that select the action that yields the higher utility among all possible actions. The deontic operator feels very similar to actions in a dynamic logic, but in the case of DEL it is possible to reason about actions and their consequences, whereas in the former case they are simply executed, changing the model. Utility based agents also lack the ability to plan the consequences of their actions in a more distant future.

In the case of *Façade* (Mateas and Stern 2005, Mateas and Stern 2007), a natural language interpreter based on a rule system feeds a drama system that plays "beats", a set of a number of "joint dialogue behaviours" specified in the ABL ("A Behaviour Language") language for the NPCs present. These behaviours can affect the game state while executing. However, again these behaviours are reactive with very little planning.

Ian Horswill has illustrated a creative use of classical logic in *MKULTRA* (Horswill 2015). In this case, the player injects some beliefs that cause the NPC to carry out some reasoning and reach a conclusion, which should advance the player's goals. This allows a more oblique approach to solving puzzles. However, still NPCs do not plan their actions or reflect on others' knowledge. We believe that *Prolog* is a double-edged tool, in the sense that solves the problem of logic reasoning for the designer, but does not easily allow adding new fundamental constructs or logic tools (e.g. a modal logic). The development of new tools for handling logic like the Z3 SMT solver (De Moura and Bjørner 2008) should be explored for new logic-related developments.

A use of a Belief-Desire-Intention logic is demonstrated in dialogue-based instruction for sales (Muller 2012). It is again a good example of a logic formalism where plans can be formulated, and a knowledge base can be checked and updated, but where little plan creation is actually done, and the agent does not reflect on what knowledge does it have or lack and how it affects its plans.

On the other hand, James Ryan's *Juke Joint* (Ryan 2016) is remarkable in using some kind of Hebb-like rule between nodes in a semantic network like structure. Even

though it still does not reason about the agents' own knowledge, it models how certain items become more salient in an agent's mind in response to external stimuli or association with other mental items.

The game on which Juke Joint is based, "Talk of the Town" (Ryan et al 2015), is also remarkable due to its extensive treatment of belief dynamics in characters. These dynamics are varied (propagation, fabrication, misremembrance...), but the contents of the beliefs will not be reflected upon or used in reasoning further than eliciting a certain appraisal. Interestingly there are many pieces in this game (evidence, salience or propagation, to name a few) that can be added to a basic DEL logic to model richer situations.

Another aspect of reasoning that could be merged with DEL is defeasible reasoning, as described by Joseph Blass and Ian Horswill (Blass and Horswill 2015). An alternative to defeasible logic might take the form of a probabilistic approach to modal logic, as described below, but defeasible logic is definitely simpler to handle.

Comme il Faut (McCoy et al 2014) is an authoring system to describe complex social situations. The agents' internal state is not so important as the social relations and their dynamics in its goal of producing a narrative, which is basically implemented with a forward rule system. We must highlight the use of a "memory" or "history" system whereby the outcome of a social exchange may depend on whether certain facts or qualities may have happened in the past, an uncommon characteristic in game AI in narratives.

Making a mental model of the other agent's mind and reproducing their reasoning is crucial both to detect a lie and to lie successfully, a kind of agent behavior that is not well explored in games. Although it is already implemented in a game mentioned above (Ryan et al 2015), a more detailed description of the goals and mechanisms of lying as a way of manipulating another agent was described by Henrique Reis (Reis 2012), although we think the mechanisms used were not appropriate for the task. A formal approach to this human characteristic has been described in the context of DEL (van Ditmarsch et al 2012), also linking it to game theory.

Other relevant projects include the "Logical Agents" project using Temporal Action Logic to guide agents (University of Linköping 2014), although it seems a dead project; and the "Ace Attorney" series of games (Capcom 2001) and the related "Aviary Attorney" game (Sketchy Logic Games 2014), which use investigation and interrogation as their main mechanics.

An alternative to runtime inferences, if the process proves too costly in computation, may be to use it as an authoring tool to create scripts in which certain logic properties are preserved. For example, plans can be generated in the design process that include sensing actions to obtain information the agent may lack, and just use those plans

without changes during the game. In this area, we can highlight several systems. For example, Chris Martens has used linear logic to create traces from formal descriptions of a game design (Martens 2015). Adam Smith and Mark Nelson (Smith et al 2010) use event calculus, a formalism for describing factual change alternative to dynamic logic, again as a way to model gameplay. It is interesting to notice that, in the planning arena, both dynamic logic and event calculus have been considered to improve the planning process. Adam Smith later applies (Smith and Mateas 2011) the capabilities of answer set programming to procedural content generation of game elements. In contrast with the formalisms used in these examples, a DEL based tool would be more concerned with the contents of the mind of an NPC, and how it plans its actions and infers their consequences.

Application to games

DEL is specially well suited to inductive puzzles, where the lack of knowledge of an agent is in itself a useful hint. However, this is just a plain application of proof or model methods to verify the truth of a certain DEL formula. Of course, these puzzles depend on the skill of the user to reason "manually" about these formulae.

A more interesting scenario is feeding the appropriate information to a DEL-reasoning agent so it reaches some kind of conclusion itself, as a sort of "proof guiding" process. We can call this game mechanic "manipulation", and it would include intimidating, lying, cajoling, exaggerating or otherwise subverting information. We will describe a couple of scenarios, framed in the game design patterns described in Treanor et al 2015:

The player detective asks questions to different agents. The agents declare what they have done in the past in relation to a mystery, in such ways that some actions and answers only make sense if they have some knowledge that only some kind of role (e.g. the "thief") may have had ("But how could Mr. Cartwright resist the strangling? – Ha! How did you know he was strangled? I never mentioned it!"). It would be an example of "AI as Villain", since the gameplay is asymmetric; the AI may reveal more than a human would do to make the puzzle easier.

The player strategist must make empire-wide "moves", taking actions that allow them to advance their goals while at the same time having plausible ambiguous explanations that prevent swift reactions from the other empires ("We admit that we camp near your border, but the miners in Alfheim are prey to wild animals, and it is only natural that we establish a garrison of armored knights there."). This mechanic could be interesting for 4X or Diplomacy-style games where direct aggressive action is not always the best option, basically presenting "AI as Adversary".

The player agent listens to an account from an NPC agent. The player then finds out conflicting evidence demonstrating that the NPC agent is lying; however, the player needs to keep the NPC ignorant about this fact to avoid the NPC from changing their strategy, so a reliable counterplan can be devised (“The regent said the assassins wanted them dead, but all the weapons seemed to use blank shots. Play along with them for the moment, but keep your eyes peeled!”). This gameplay would again be an example of “AI as Adversary” or “AI as Villain”.

However, basic DEL logic has been enriched in several ways by different authors. Some mechanics combining DEL and other existing techniques and formalisms are:

Planning algorithms. As mentioned before, planning algorithms can be applied in dynamic epistemic scenarios, for example to plan actions that reduce uncertainty (Boland and Andersen 2011). An “AI as Adversary” in an RTS game may plan scouting (sensing) actions as a way of reducing the uncertainty of possible plans, instead of resorting to replanning or planning with alternatives.

Reasoning about goals. Manipulation involves biasing someone’s reasoning taking into account their goals. Goals have been formalized in such logics as BDI (Rao and Georgeff 1995) or stit (“see-to-it-that”) logics (Horty and Belnap 1995). The combination of DEL with goal logics may enable mechanics of an “AI as a Guided Adversary” pattern where an enemy is overcome by feeding such information that, when combined with their internal state, will cause the enemy to act in a way as to advance the player’s goals. It would be similar in its goal as MKULTRA (Horswill 2015), but using subtler mechanisms: instead of feeding mental facts, actions could be suggested that, when included in an adversary’s plan, causes beneficial side effects. In the subtlest cases, information could be presented almost casually as background scenery. For example, a guard can be persuaded to leave a locked room to ask for food for a starving person, only if the perception that the person asking for food is incapable of entering the room (e.g. is not a wizard with an Open Locks spell) is established in their mind, by using poor clothes, a lowly accent, or similar elements. Comparing to MKULTRA, the “mindbending” aspect is replaced with a more difficult “plausibility” requirement. On the other hand, an “AI as Adversary” that lies may convey information about goals indirectly, so that the fact that an agent issued a specific lie would only make sense if the agent had a specific goal.

Probabilistic DEL, already explored by Kooi in (Kooi 2003). It expands the logic of DEL with information about the probabilities of certain facts as calculated by the observing agent, so by feeding certain information we can make an agent consider certain actions more probable than others, and therefore probably influence a planning process that depends on these estimations. It is somewhat different from injecting new information in that only the perceived probability of the current actions

is modified; on the more complex implementations, these probabilities are not modified directly but by pointing out to certain evidence and hiding other. By making certain actions seem more or less probable, we can “defeat” an “AI as Adversary” making them take less than optimal actions. But we can also use the “AI is Guided” approach if we want to help an agent succeed in spite of their preconceptions, as in the case of a stubborn child.

Temporal logic. At least, the use of a PAST operator allows an agent to reason about past events (or even past knowledge, as in “B knew that p”). Different authors (van Ditmarsch et al. 2011, Renne et al. 2016, Hoshi, and Yap 2009) describe possible ways in which temporal operators can be added to a DEL logic. The “AI as Co-creator” pattern could fit this blend of logics, by creating agent traces (similar to the detective games we have mentioned before) that only make sense if at some point the agent has or lacks certain knowledge. A detective player observing a broken window can infer that this fact is consistent with the sentence “the thief did not know the front door was unguarded”, which may rule out certain suspects; the authoring tool would highlight a proof of this sentence to the story writer.

Interrogative logic. Hintikka’s interrogative logic can be combined with DEL (Hamami 2015, van Benthem and Minică 2012) to describe scenarios in which questions are planned to obtain the truth of a proposition. Settings with “AI is Guided” patterns may benefit from agents that ask their own questions to their player “gods”, and may even evolve into “AI as Spectacle” if the player is interested in causing outrageous inferences from their agents with weird or conflicting answers.

Speech acts. DEL deals with the basic declarative speech act, and some extensions, as seen in the previous bullet, deal also with questions. Indirect speech acts like excuses could be modelled accurately in the exchange of DEL sentences. The definitions from existing agent communication languages (like the FIPA ACL) might serve as an inspiration for these models, as well as the models found in the field of computational semantics. We believe that this NLG/NLP capabilities would make any of the other described mechanics feel more human. Both MKULTRA and “Talk of the Town” use NLG in this way. Agents (both virtual and human) may also plan to use different strategies, like convincing, confusing or threatening, depending on their adversary, but using NLP/NLG techniques instead of selecting a “mode” in an interaction window. A design pattern of “AI as Role-model” (although definitely much less related to DEL) can be used in serious games to teach pragmatics in a foreign language (even in a fantasy, invented language!); DEL may make the content more interesting than talking about plain facts (“Did they divorce last month? I really had no idea... I know I’m a close friend to their lawyer and we often dine together, but I assure you I have been too busy lately to talk about other people’s lives.”).

DEL can be used as the mind mechanics for an NPC player in runtime, but given the state of the current imple-

mentations the best option at first may be to create specific content for games using DEL as an authoring tool (as discussed for the mix with temporal logics). Many of the previous mechanics become examples of the “AI as Co-creator” pattern where the product is a plan (with actions that take into account uncertainty or work to reduce it), or a scripted dialogue or behavior that would hint at the possession of some knowledge or lack thereof.

Conclusion

The DEL formalism is richer than logics used in previous works to reason about actions and knowledge, and enables different reasoning mechanisms like deduction or planning that have already been addressed in theory. We have shown some mechanics based on this reasoning about knowledge that may provide new twists to existing game genres. We are certain that the current computing power and advances in automated reasoning will make it possible to use techniques previously deemed intractable to deliver these new experiences.

References

- Aucher, Guillaume, and François Schwarzentruher. On the Complexity of Dynamic Epistemic Logic (Extended Version). Diss. INRIA, 2012.
- Baker, Chris, Rebecca Saxe, and Joshua Tenenbaum. "Bayesian theory of mind: Modeling joint belief-desire attribution." *Proceedings of the Cognitive Science Society*. Vol. 33. No. 33. 2011.
- Baltag, Alexandru, Bob Coecke, and Mehrnoosh Sadzadeh. "Epistemic actions as resources." *Journal of Logic and Computation* 17.3 (2007): 555-585.
- van Benthem, Johan, and Ștefan Minică. "Toward a dynamic logic of questions." *Journal of Philosophical Logic* (2012): 1-37.
- van Benthem, Johan, et al. "Symbolic model checking for dynamic epistemic logic." *International Workshop on Logic, Rationality and Interaction*. Springer, Berlin, Heidelberg, 2015.
- Blass, Joseph A., and Ian D. Horswill. "Implementing Injunctive Social Norms Using Defeasible Reasoning." *AIIDE 2015 Joint Workshop on Intelligent Narrative Technologies and Social Believability in Games*. 2015.
- Bolander, Thomas, and Mikkel Birkegaard Andersen. "Epistemic planning for singleand multi-agent systems." *Journal of Applied Non-Classical Logics* (2011).
- Bosse, Tibor, Zulfiqar A. Memon, and Jan Treur. "A two-level BDI-agent model for theory of mind and its use in social manipulation." *Proceedings of the AISB 2007 Workshop on Mindful Environments*. Vol. 4. 2007.
- Capcom. "Ace Attorney". 2001.
- H.P. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2006.
- van Ditmarsch, Hans, Wiebe van der Hoek, and Ji Ruan. "Connecting dynamic epistemic and temporal epistemic logics." *Logic Journal of IGPL* 21.3 (2011): 380-403.
- Van Ditmarsch, Hans, et al. "On the Logic of Lying." *Games, actions and social software* 7010 (2012): 41-72.
- van Eijck, Jan. "a demo of epistemic modelling." *Interactive Logic* (2007): 303.
- Evans, Richard, and Emily Short. "Versu—a simulationist storytelling system." *IEEE Transactions on Computational Intelligence and AI in Games* 6.2 (2014): 113-130.
- Evans, Richard P., and Emily Short. "The AI Architecture of Versu." (2015).
- Hamami, Yacin. "The interrogative model of inquiry meets dynamic epistemic logics." *Synthese* 192.6 (2015): 1609.
- Herzig, Andreas. "Dynamic Epistemic Logics: Promises, Problems, Shortcomings, and Perspectives." *Logical Reasoning and Computation: Essays dedicated to Luis Fariñas del Cerro* (2016): 119.
- Horswill, Ian. "MKULTRA." *Proc. AIIDE* (2015).
- Horty, John F., and Nuel Belnap. "The deliberative stit: A study of action, omission, ability, and obligation." *Journal of philosophical logic* 24.6 (1995): 583-644.
- Hoshi, Tomohiro, and Audrey Yap. "Dynamic epistemic logic with branching temporal structures." *Synthese* 169.2 (2009): 259-281.
- d'Inverno, Mark, et al. "The dMARS architecture: A specification of the distributed multi-agent reasoning system." *Autonomous Agents and Multi-Agent Systems* 9.1 (2004): 5-53.
- Kooi, Barteld P. "Probabilistic dynamic epistemic logic." *Journal of Logic, Language and Information* 12.4 (2003): 381-408.
- Marsella, Stacy C., David V. Pynadath, and Stephen J. Read. "PsychSim: Agent-based modeling of social interactions and influence." *Proceedings of the international conference on cognitive modeling*. Vol. 36. 2004.
- Martens, Chris. "Ceptre: A Language for Modeling Generative Interactive Systems" *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* (2015)
- Mateas, Michael, and Andrew Stern. "Writing Façade: A case study in procedural authorship." *Second person: Role-playing and story in games and playable media* (2007): 183-208.
- Mateas, Michael, and Andrew Stern. "Structuring Content in the Façade Interactive Drama Architecture." *AIIDE*. 2005.
- McCoy, Joshua, et al. "Social story worlds with comme il faut." *IEEE Transactions on Computational intelligence and AI in Games* 6.2 (2014): 97-112.
- De Moura, Leonardo, and Nikolaj Bjørner. "Z3: An efficient SMT solver." *Tools and Algorithms for the Construction and Analysis of Systems* (2008): 337-340
- Muller, Tijmen Joppe, et al. "Glengarry Glen Ross: Using BDI for Sales Game Dialogues." *AIIDE*. 2012.
- Playable.ai 2016. "Façade". www.interactivestory.net.
- Pynadath, David V., and Stacy Marsella. "Minimal mental models." *AAAI*. 2007.
- Rao, Anand S., and Michael P. Georgeff. *Formal models and decision procedures for multi-agent systems*. Melbourne, Australia: Australian Artificial Intelligence Institute, 1995.
- Reis, Henrique Daniel Santarém. *Lie to me: Lying virtual agents*. Diss. Master's thesis, Universidade Técnica de Lisboa, 2012.
- Renne, Bryan, Joshua Sack, and Audrey Yap. "Logics of temporal-epistemic actions." *Synthese* 193.3 (2016).

Ryan, James Owen, et al. "Toward characters who observe, tell, misremember, and lie." Proc. Experimental AI in Games 2 (2015).

Ryan, James, et al. "Juke joint: characters who are moved by music." Proceedings of Experimental AI in Games (2016).

Sketchy Logic Games. "Aviary Attorney". 2014.

<http://aviaryattorney.com/>

Sadrzadeh, Mehrnoosh. Actions and resources in epistemic logic. Diss. University of Quebec at Montreal, 2006.

Smith, Adam M., Mark J. Nelson, and Michael Mateas. "Ludocore: A logical game engine for modeling videogames." Computational Intelligence and Games (CIG), 2010 IEEE Symposium on. IEEE, 2010.

Smith, Adam M., and Michael Mateas. "Answer set programming for procedural content generation: A design space approach." IEEE Transactions on Computational Intelligence and AI in Games 3.3 (2011): 187-200.

Treanor, Mike, et al. "AI-based game design patterns." (2015).

University of Linköping. The Logical Agents Project. 2014

<http://www.ida.liu.se/divisions/aiics/projects/logicalagents.en.shtml>

The Versu team. "Versu". www.versu.com.