A Strategic Epistemic Logic for Bounded Memory Agents

Sophia Knight*

Introduction In [1] we presented an ATL-style epistemic logic for agents with arbitrary equivalence relations on histories, called euATL. Here we review this logic and discuss its applications to problems concerning bounded memory agents. euATL is unique and relevant for problems in resource-bounded agents since it allows us to model and reason about systems where agents each have different, *arbitrary* equivalence relations on histories, and can be aware of their own past actions. While partial information strategic logics with memoryless or perfect recall agents have already been studied, in our systems we can model a situation, e.g. where a subset of agents have bounded finite memory and another subset has unbounded memory. Since we allow arbitrary equivalence relations on histories including actions, we can also reason about situations where an agent loses all its information after entering a certain state, loses its memory after taking a certain action, remembers half of the previous states, etc. Thus, euATL is a practical logic for discussing situations where agents' memory is bounded, even in complex ways.

The logic euATL is defined on the models *epistemic concurrent game structures*, consisting of agents, states, and an equivalence relation on states for each agent, where at each state, every agent chooses an available action and the next state is chosen deterministically by the combination of these actions.

We assume that besides the agents' relations on states, the systems also have equivalence relations on histories for each agent. By letting agents have *arbitrary* equivalence relations on histories, we allow several interesting situations: we can model a perfect recall agent by letting him distinguish any pair of histories that have a pair of distinguishable states, or differ in actions he took, and we can model memoryless agents by basing their relation only on the last state in the history. Similarly we can model finite memory agents, agents who always forget a certain state, etc. By including actions in histories, we allow agents to remember, or forget, their own actions, rather than only remembering past states. We define strategies for agents as usual: a function assigning an available action for each possible history, respecting the agent's equivalence relation.

Now we can discuss euATL. The syntax is

$$\phi ::= p \mid \neg\phi \mid \phi \vee \phi \mid K_i\phi \mid C_A\phi \mid \langle\langle A\rangle\rangle \bigcirc \phi \mid \langle\langle A\rangle\rangle \Box\phi \mid \langle\langle A\rangle\rangle \phi \mathcal{U} \phi$$

Booleans, knowledge and common knowledge are interpreted as usual. The operator $\langle\langle A\rangle\rangle \bigcirc \phi$ means "the agents in group A have a strategy to make ϕ true at the next state, based on their knowledge." Thus, the strategy for each agent must succeed not only at the current history but at all other histories that the agent considers possible as well.

Applications There are many scenarios where it is practical to model agents with different memory abilities in one system. For example in a system with some friendly, bounded memory agents and other adversarial agents with unknown memory abilities, we could model the friendly agents as limited memory agents, and the adversarial agents as perfect recall agents in order to represent the worst case scenario, which would be practical for verifying security properties in a system. Also, allowing arbitrary equivalence relations on histories gives a great deal of flexibility in modelling agents with their memory bounded in interesting ways: e.g. we can model an agent whose memory fills up after he has seen n states, so he only remembers the first n states and gains no new information after this, or we could model a system where agents' actions affect their memory, such as an agent who forgets everything after performing a certain action, or even an agent who forgets the first or last state after performing some action. Thus, our systems allow us a great deal of flexibility in modelling agents with bounded abilities.

References

[1] H. van Ditmarsch and S. Knight. Partial information, knowledge and uniform strategies. Under submission, `http://www.loria.fr/~sknight/Strategy_Journal.pdf`, 2015.

*LORIA, CNRS, Université de Lorraine, France