On LTL_f Synthesis and Goal Formation/Revision

Yves Lespérance

York University, Toronto, Canada

AAAI Spring Symposium 2023, San Francisco, March 27-29, 2023

Goal Reasoning in Autonomous Agents

- In autonomous agents there is a growing interest in goal reasoning:

  "intelligent systems may benefit from deliberating about, and changing their active goals when warranted." [D.Aha 18]

- Change goals when circumstances change

- Goal-driven architectures [Molineaux et al 10]

Mental Attitudes & Rational Balance

- Philosophy of Mind has examined relationship between different mental attitudes: belief, desire, intention, & action
- Need “rational balance”
- In AI, have BDI logics/theories, e.g. [Cohen & Levesque 90], [Rao & Georgeff 91], [KARO framework]
- Bratman’s key idea: intentions support resource bounded reasoning by acting as a filter on “practical reasoning” [Bratman 87]
  - normally, we only adopt a new intention if it is consistent with existing intentions and beliefs
  - but an intention can be given up for another intention if utility of doing so is high, i.e. have a filter override

Challenges for Goal Deliberation/Planning in Autonomous Agents

- Environment may be very dynamic, so plans may need to change
- Agent may acquire new desires or get new requests over time
- Agent may have incomplete knowledge, making planning difficult

Desiderata for Rational Agent Architecture

- Agent progressively commits to desires & refines/revises intentions over time
- Agent's plans may be abstract; no need to consider fully detailed plans until it is feasible
- Need to ensure that agent's intentions remain consistent among themselves and with its beliefs
- Can reason about possible interactions between plans/goals and add temporal constraints to prevent interference
- Environment is typically nondeterministic, so agent must do strategic reasoning and consider likely contingencies

Simple Example Application: Courier Agent

- Have courier driver agent that must decide on what orders to pickup and drop off, what route to take, when to recharge/refuel, etc.
- Traffic and road conditions can change
- New orders may be received and orders may be cancelled; they may have different deadlines and priorities/service guarantees
- Agent needs to revise goals and make plans as it operates
- Different levels of abstraction (which interact):
  - Which orders to serve
  - Route planning
  - Low-level navigation

LTL_f Synthesis and Goal Formation/Revision

- LTL_f provides a rich language for expressing temporally extended goals
- LDL_f supports more procedural goal specifications
- Build on recent work on LTL_f/LDL_f synthesis & FOND planning to support intention consistency checking/maintenance
- In particular, build on techniques for synthesizing maximally permissive strategies [Zhu & De Giacomo 22] for temporally extended goals, that support postponing commitments
- Possibly also use for intention refinement

Synthesis of Maximally Permissive Strategies for LTL_f Goals [Zhu & De Giacomo 22]

- Represent using nondeterministic strategies
- For safety goals, can do anything while remaining in winning region
- In general, need to avoid infinite deferral: MPS includes any strategy that remains in winning region initially & eventually switches to a non-deferring strategy
- Computing MPS is not harder than plain LTL_f synthesis

Khan & Lesperance's Formal Model of Goal Dynamics

- Goals/desires/intentions are sets of infinite paths
- PGoals
  - represent desires
  - organized in priority list, possibly infinite (max priority being 0)
  - may be organized in a subgoal hierarchy
- CGoals
  - represent intentions, i.e. desires that the agent is committed to pursue
  - must be consistent with one another and with agent's knowledge
  - are computed from PGoals: the largest set of highest priority PGoals that are consistent with knowledge and among themselves
- Actions *adopt* and *drop* update PGoals, their priority, and their parent (implicitly updating CGoals):
  - *adopt*(ψ, phi, m), where the PGoal ψ is adopted at level m relative to a parent PGoal ϕ (which must exist and be at higher priority than m)
  - *drop*(ϕ) drops the PGoal ϕ and all its subgoals from the priority list
- Other models [van Riemsdijk et al 09, Winikoff 11]

Agent Abstraction [Banihashemi et al 17, 23]

- Defines notions of sound/complete abstraction between an abstract action theory/domain model and a concrete one (based on Situation Calculus & Golog)
- Assumes one has defined a mapping between them
  - high-level fluent is mapped to a low-level state formula
  - high-level atomic action is mapped to a Golog programs over low-level actions and fluents
- Sound/complete abstraction between HL and LL theories is defined in terms of bisimulation between their models relative to mapping
- Sound abstraction provides guarantees that HL strategy to achieve a goal can always be refined to a LL strategy to achieve it at LL
- Can also be used to generate high-level explanations of system execution
- Handles nondeterministic domains

Related Work: Intention Progression Problem

- Intention Progression Problem [Castle-Green et al 20] in autonomous agents area focusses on how an agent operating in a dynamic environment should select
  - which subplan to adopt
  - which concurrent plan to advance
- Assumes that relevant plans have been precomputed into a goal/plan tree
- Agent must analyse how its intentions may interact & how environment may interfere
- Competitions have been held
- Monte-Carlo tree search techniques have shown promise [Yao et al 21]

Related Work: Partial-Order Planning [Barrett & Weld 94]

- Focusses on finding plans to achieve a set of atomic goals without committing to a full temporal ordering on the actions
- A plan is a set of actions together with a set of temporal constraints on their execution ordering
- Causal links keep track of dependence of a given action on a prior action to achieve a precondition or context condition
- Add actions and introduce temporal constraints to handle dependencies
- Causal links can also be used in execution monitoring and replanning
- Caveat: usually less efficient than total-order planning

Conclusion

- Goal formation/revision is an interesting and challenging area where advances in LTL_f synthesis (and FOND planning) could be applied
- Need to ensure consistency between intentions (& beliefs)
- Need to support abstract plans & avoid early commitment to detailed plans
- Need to ensure robustness in dynamic and incomplely known worlds

Issues to Investigate

- Modularity and compositionality in LTL_f synthesis
- Handling imperfect information in LTL_f synthesis
- Intention refinement & goal formation
- Handling conditional commitments, e.g. agent intends to ship item when it receives payment
- Using models at different levels of abstraction for different tasks