A Stochastic Belief Management Architecture for Agent Control

Gavin Rens, Tommie Meyer, Deshen Moodley

Centre for AI Research, CSIR Meraka, South Africa
University of Cape Town, Dept. Comp. Sci., South Africa
{grens,tmeyer,deshen}@cs.uct.ac.za

Architectures for Generality and Autonomy, 2017

Outline

Introduction

Elements and Components

- POMDP theory.
- Logic languages.
- Belief change theory.
- Dealing with ignorance (via entropy optimization).
- Learning and planning assumed present (no detail).
- Observation stream processing.

Conceptual Model of Architecture

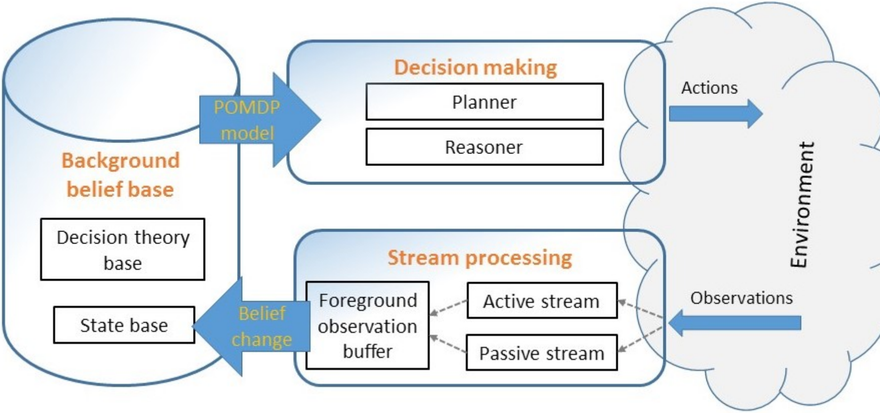

Architecture Components in Detail

Observation Streams

- Passive stream
  - periodic extraction
  - unsolicited
  - unclear observation cause
- Active stream
  - asynchronous
  - associated with intention
  - observation in context of action

The Background Belief Base

- Decision Theory Base
  - environment models (transition, observation & event functions)
  - *Stochastic Decision Logic* (SDL) [Rens et al., 2015]
  - possibly incomplete specifications
- State Base
  - probabilistic constraints on possible states
  - P-logic [Zhuang et al., 2017]
  - possibly incomplete specification

The Foreground Observation Buffer

- Observations are recorded in the order which they were received.
- $Z = (z_1, z_2, \ldots, z_n)$
- Combining observations from passive and active streams.
- z_i is ϕ or (a, ϕ) or $p(\phi) \bowtie t$,
  - where $\phi \in L_P$

The Foreground Observation Buffer

- The SB is modified by z_1 via a belief change operation.
- z_1 removed from Z.
- Then z_2 modifies the SB, and so on.

Decision Making

- Reasoner for query answering.
- Asking whether some query posed to the agent logically follows from the BG BB.
- No implementations. (?)
- Not optimized/analyzed with respect to efficiency.

Decision Making

- Planner for action/policy recommendation.
- Online POMDP algorithm employing finite horizon forward search.
- Such algos require a single current (root) belief state and a fully specified POMDP model.
- Entropy optimization (if necessary).

General Belief Change

Dimensions of Belief Change

- Belief revision ($*$) $\neq$ Belief update (Δ).
- *Revision*: for change w.r.t. epistemic (static) info.
- *Update*: for change w.r.t. ontic (physical/dynamical) info.
- Active vs Passive.
- Adequacy of transition model specification.

Appropriate Belief Representation

- Focus on belief management of the state base (SB).
- P-revision (of p-logic) is fine for *revision*.
- Something else required for *update*.
- Update operators typically applied to pd's, not sets of constraints.
- Several candidate methods to extract representative *credal sets* from SB.

Operators and Operator Selection

- Every observation is assumed *marked* or *raw*.
- Marked: accompanied by degree of belief info.
- Raw: no extra info.
- Epistemic observations are always marked (from both streams).
- Ontic observations are always raw (from both streams).

Operators and Operator Selection

	ontic		epistemic
	transition info?		
	✓	✗	
active	Δ^{atv}_{trn}	Δ^{atv}_{dst}	$*$
passive	Δ^{psv}_{trn}	Δ^{psv}_{dst}	$*$

Operators and Operator Selection

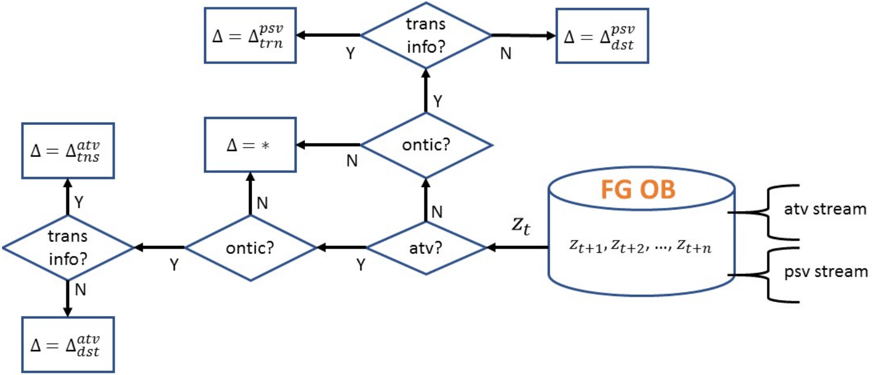

Conclusion and Future Work

Conceptual Model of Architecture

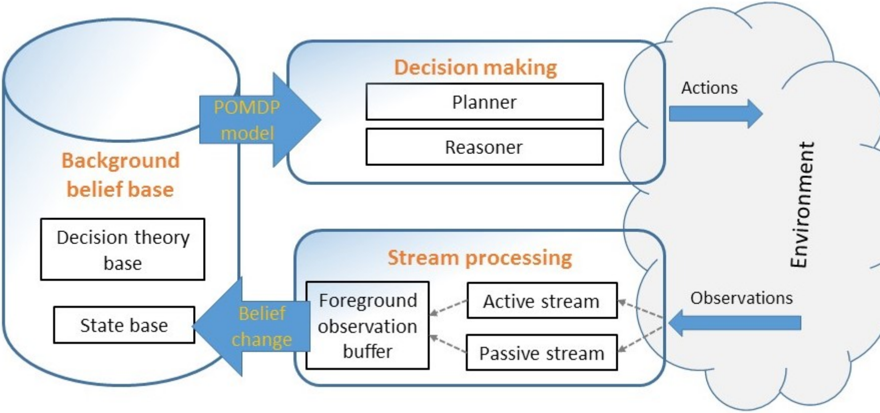

Summary

1. We proposed a coherent framework for general agent knowledge management and decision-making under uncertainty and ignorance.

2. We presented a means to deal with two kinds of steams of observations: *active* and *passive*, where observations from different streams are dealt with differently.

3. We presented a technique for belief maintenance which takes into account whether beliefs should be revised (due to erroneous beliefs) or updated (due to changes in the environment).

Remarks

- Preliminary proposal for an architecture.
- Only a basic framework.
- Generate a discussion and ideas for its improvement.

Remarks

- Many ways to sophisticate the architecture.
- Environment model learning.
- New states learning.
- Case-based reasoning.
- Etc., etc.

Future Work

- Newest findings in probabilistic/possibilistic belief update.
- Issues with knowledge/memory embeddedness.
- W.r.t. real-time planning,
  - Hybrid POMDP-BDI agent architecture [Rens & Moodley, 2017].
  - Partially Observable Monte-Carlo Planning (POMCP) [Silver & Veness, 2010].
- Current work to be incorporated:
  - More robust, generally applicable belief revision,
  - More robust, generally applicable belief update.

G. Rens, T. Meyer, D. Moodley

Centre for AI Research, CSIR Meraka, South Africa
University of Cape Town, Dept. Comp. Sci., South Africa

A Stochastic Belief Management Architecture for Agent Control

Architectures for Generality and Autonomy, 2017