A Plan Optimality Monitoring Approach to Detect Commitment Abandonment

Ramon Fraga Pereira†, Nir Oren‡, and Felipe Meneguzzi†

†Pontifical Catholic University of Rio Grande do Sul, Brazil
ramon.pereira@acad.pucrs.br
felipe.meneguzzi@pucrs.br

‡University of Aberdeen, United Kingdom
n.oren@abdn.ac.uk

COIN@AAMAS, 2017

May, 2017

Introduction

- Determining whether an agent is actually executing steps towards a goal (or has abandoned it), is important when:
  - multiple agents are trying to achieve joint goals, or
  - agents are committed for achieving goals for each other.
- **Commitment abandonment**: situation in which an agent switches from executing the actions of one plan that achieves the consequent it is committed to, to executing actions from another plan;
- We develop a **domain-independent** approach based on **planning techniques** to:
  - detect sub-optimal steps; and
  - infer whether an agent will honour a commitment

Background: Commitments

- A **commitment** `C(DEBTOR, CREDITOR, antecedent, consequent)` formalizes that the agent `DEBTOR` commits to agent `CREDITOR` to bring about the consequent if the antecedent holds;
- The **antecedent** and **consequent** conditions of a commitment are conjunctions or disjunctions of events and possibly other commitments;
- In this paper, we aim to **monitor** the `DEBTOR`'s behaviour (i.e., sequence of actions) to **detect** if this agent is **individually committed to carrying out a plan to achieve** the consequent for the `CREDITOR`.

Background: Planning, Heuristics, and Landmarks

Definition (Planning)

A planning instance is represented by a triple $\Pi = \langle \Xi, \mathcal{I}, G \rangle$, in which:

- *$\Xi = \langle \Sigma, \mathcal{A} \rangle$ is the* **domain definition**, *and consists of a finite set of* **facts** *Σ and a finite set of* **actions** *$\mathcal{A}$ (action costs = 1);*
- *$\mathcal{I}$ and G represent the* **planning problem**, *in which $\mathcal{I} \subseteq \Sigma$ is the* **initial state**, *and $G \subseteq \Sigma$ is the* **goal state**.

- **Heuristics** are used to estimate the cost to achieve a particular goal. In this work, we use **domain-independent heuristics**;

Definition (Landmarks)

Given a planning instance $\Pi = \langle \Xi, \mathcal{I}, G \rangle$, a **fact** *(or* **action***) L is a landmark in Π iff L must be* **satisfied** *(or* **executed***) at some point along all valid plans that achieve G from $\mathcal{I}$.*

Background: Fact Partitioning

Pattison and Long ("**Domain Independent Goal Recognition**". In STAIR, 2010) classify facts into a set of mutually exclusive fact partitions. We use such partitions to infer whether certain observations are consistent with a particular goal state, as follows:

- **Strictly Activating** is a type of fact that can never be added by any action unless defined in the initial state;
- **Unstable Activating** is a type of fact that that once deleted, cannot be re-achieved;
- **Strictly Terminal** is a type of fact that once added, cannot be deleted.

Background: Commitment Abandonment Problem

Definition (**Commitment Abandonment Problem**)

- Domain definition (Properties and Actions) $\Xi = \langle \Sigma, \mathcal{A} \rangle$;
- Commitment C, in which C(DEBTOR, CREDITOR, At, Ct), DEBTOR is the debtor, CREDITOR is the creditor, At is the antecedent condition, and Ct is the consequent;
- Initial state $\mathcal{I}$, *s.t.*, At $\subseteq \mathcal{I}$ (when begins the monitoring process);
- An observation sequence $O = \langle o_1, o_2, ..., o_n \rangle$, representing a full observable plan execution; and
- Threshold θ, representing the percentage of sub-optimal actions that the DEBTOR agent can deviate to achieve the consequent state Ct.

- The solution for a commitment abandonment problem is whether an observation sequence O has deviated more than θ from the optimal plan to achieve the consequent Ct of commitment C.

Monitoring Plan Optimality

- We use plan optimality monitoring techniques from the literature to detect sub-optimal steps (Pereira *et al.* "**Monitoring Plan Optimality using Landmarks and Domain-Independent Heuristics**". In PAIR@AAAI, 2017.);
- This approach combines **planning techniques**, *i.e.*, landmarks and domain-independent heuristics.
  - It uses **landmarks** to obtain information about **what cannot be avoided** to achieve a goal G; and
  - It uses **heuristics** to analyse possible **plan execution deviation**.

Analyzing Plan Execution Deviation

- If an observation o_i results a state s_i, we consider a **deviation from a plan** to occur if $h(s_{i-1}) < h(s_i)$.

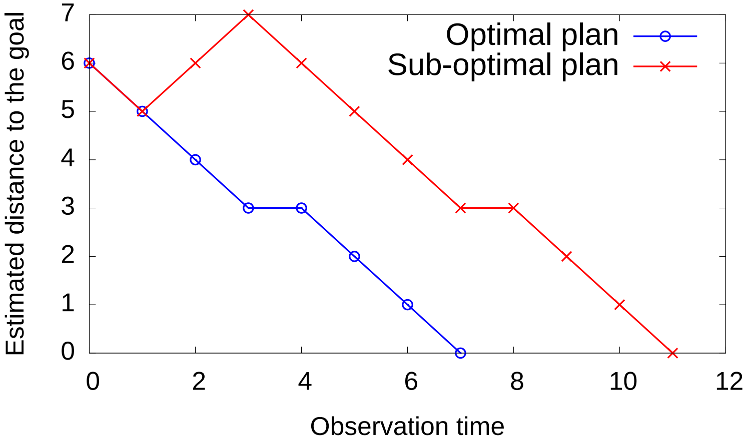

Predicting Non-regressive Actions via Landmarks

- To predict which actions could be executed in the next observation, we **estimate the distance to the closest landmarks** (using h_{max}) from the current state to the extracted landmarks $\mathcal{L}$, and select the following actions:
  - For every fact landmark $l \in \mathcal{L}$ with $h_{max}(l) = 0$, we select actions $a \in \mathcal{A}$ such that $l \in pre(a)$; and
  - For every fact landmark $l \in \mathcal{L}$ with $h_{max}(l) = 1$, we select actions $a \in \mathcal{A}$ such that $pre(a) \in$ current state and $l \in eff(a)^{+}$;
- Predicted actions **may reduce the distance** to the **monitored goal** and **next landmarks**.

Detecting Sub-Optimal Steps

- To detect sub-optimal steps (actions) in observation sequence O for a monitored goal G, we combine the techniques we developed and filter with the following condition:
  - **An observed action** $o \in O$ **is considered sub-optimal if**: $o \notin$ set of predicted actions AND $(h(s_{i-1}) < h(s_i))$.

Commitment Abandonment Detection Approach

- We monitor the sequence of actions of a `DEBTOR` to infer whether it **will abandon a commitment**
  - Observed sequence should achieve the consequent from a state in which the antecedent holds
- We use a **threshold** θ, representing the percentage of sub-optimal actions that the `DEBTOR` agent can deviate to achieve the consequent it is committed to, *i.e.*, a percentage of actions that `CREDITOR` agent agrees to deviate from the optimal.

Determining Commitment Abandonment using Plan Optimality Monitoring and Fact Partitioning

Our approach determines that a DEBTOR agent has abandoned a commitment it is committed to if any one of three conditions is true:

1. *Strictly Activating* facts that we extracted are not in the initial state;
2. we observe the evidence of any *Unstable Activating* and/or *Strictly Terminal* facts during the execution of actions in the observations; or
3. the number of observed sub-optimal steps is greater than θ defined by the CREDITOR.

Experiments and Evaluation (1 of 2)

- We evaluate our approach over 8 planning domains, most of which are inspired by real-world scenarios;
  - **Precision:** percentage of the abandoned commitments inferred that were actually abandoned (quality);
  - **Recall:** percentage of actually abandoned commitments inferred by the approach (quantity);
  - **F1-score:** harmonic mean between *Precision* and *Recall*.
- We use 6 domain-independent heuristics:
  - h_{adjsum}, $h_{adjsum2}$, $h_{adjsum2M}$, h_{combo}, h_{ff}, and h_{sum};
- We manually generated the dataset from medium and large planning problems, generating plans that either abandoned (ultimately went to a different goal) or did not abandon their corresponding goals/consequent, varying the number of sub-optimal actions.

Experiments and Evaluation (2 of 2)

Domain	$\|O\|$	T	Precision θ (0% / 5% / 10%)	Recall θ (0% / 5% / 10%)	F1-score θ (0% / 5% / 10%)
DRIVER-LOG (30) $h_{adjsum2M}$	20.0	0.83	1.0 / 1.0 / 1.0	1.0 / 1.0 / 1.0	1.0 / 1.0 / 1.0
DEPOTS (30) $h_{adjsum2}$	18.6	1.79	1.0 / 1.0 / 1.0	1.0 / 1.0 / 0.8	1.0 / 1.0 / 0.88
EASY-IPC-GRID (30) h_{ff}	17.3	0.95	1.0 / 1.0 / 1.0	1.0 / 1.0 / 1.0	1.0 / 1.0 / 1.0
FERRY (30) $h_{adjsum2}$	13.5	0.38	1.0 / 1.0 / 1.0	1.0 / 0.8 / 0.8	1.0 / 0.88 / 0.88
LOGISTICS (30) $h_{adjsum2}$	21.0	0.56	1.0 / 1.0 / 1.0	1.0 / 1.0 / 1.0	1.0 / 1.0 / 1.0
SATELLITE (30) $h_{adjsum2M}$	23.5	5.4	0.8 / 1.0 / 1.0	0.8 / 0.6 / 0.6	0.8 / 0.75 / 0.75
SOKOBAN (30) h_{combo}	22.8	5.2	0.83 / 1.0 / 1.0	1.0 / 0.6 / 0.6	0.91 / 0.75 / 0.75
ZENO-TRAVEL (30) $h_{adjsum2}$	10.0	1.1	1.0 / 1.0 / 1.0	0.8 / 0.8 / 0.8	0.88 / 0.88 / 0.88

- $|O|$ is the average number of observed actions in a plan execution;
- T is the average monitoring time (in seconds); and
- θ is threshold value varying at 0%, 5%, and 10%.

Related Work

- Geib and Goldman. "**Recognizing Plan/Goal Abandonment**". In IJCAI, 2003;
- Kafali *et al.* "**GOSU: Computing GOal SUpport with Commitments in Multiagent Systems**". In ECAI, 2014; and
- Kafali and Yolum. "**PISAGOR: A Proactive Software Agent for Monitoring Interactions**". In Knowledge and Information Systems, 2016.

Conclusions

- **Contribution:**
  - We formalized the commitment abandonment problem using planning;
  - Our approach is domain-independent and require minimal domain knowledge; and
  - We show that our approach has high accuracy (very good results) in almost all domains (apart from SATELLITE).
- **Limitations:**
  - We only deal with full observability;
  - Our approach assumes a centralized monitor;
- **Future Work:**
  - Detect commitment abandonment using multiple monitors; and
  - Deal with partial observability (*i.e,* missing observations).

Thank you!
Questions?