Plan and Goal Recognition in the Real World

Felipe Meneguzzi†

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

Scotland, September, 2017

Table of Contents

Introduction

- Recognizing plans and goals of others is a critical ability for intelligent interaction:
  - important for humans/agents working in the same environment
  - increasingly important as we build more intelligent systems
- Overall area of Plan, Activity and Intent Recognition
  - Activity recognition: recognizing meaningful activities from low-level sensor data
  - Plan/Intent/Goal recognition: recognizing intentional higher-level sequences of activities

Introduction

- **Goal Recognition** is the task of recognizing agents' goal that explains a sequence of observations of its actions;
  - Related to plan recognition, i.e. recognizing a *top-level* action
  - A specific form of the problem of abduction
- Approaches to goal and plan recognition divided into roughly two types:
  - Plan-library based (*classical* plan recognition)
  - Domain-theory based (plan recognition as planning, or PRAP)

Flavors of Recognition Formalism

Plan Library

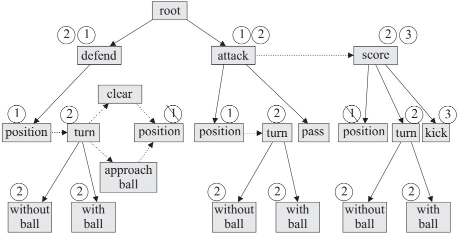

Domain Theory (PRAP)

```
(define (domain grid)
(:requirements :strips :typing)
(:types place shape key)
(:predicates (conn ?x ?y - place)
             (key-shape ?k - key ?s - shape)
             (lock-shape ?x - place ?s - shape)
             (at ?r - key ?x - place )
             (at-robot ?x - place)
             (locked ?x - place)
             (carrying ?k - key)
             (open ?x - place)
)

(:action unlock
:parameters (?curpos ?lockpos - place ?key - key ?shape
:precondition (and (conn ?curpos ?lockpos) (key-shape ?k
                   (lock-shape ?lockpos ?shape) (at-robo
                   (locked ?lockpos) (carrying ?key))
:effect (and
(open ?lockpos) (not (locked ?lockpos)))
)

(:action move
:parameters (?curpos ?nextpos - place)
:precondition (and (at-robot ?curpos) (conn ?curpos ?nex
:effect (and (at-robot ?nextpos) (not (at-robot ?curpos)
)

(:action pickup
:parameters (?curpos - place ?key - key)
:precondition (and (at-robot ?curpos) (at ?key ?curpos))
:effect (and (carrying ?key)
   (not (at ?key ?curpos)))
)
```

An example of Activity Recognition

An example of Activity Recognition

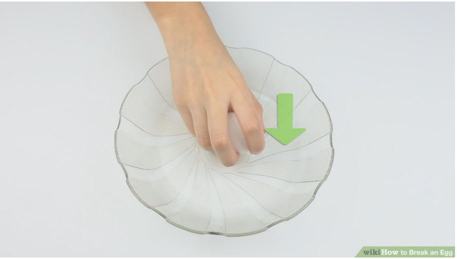

An example of Activity Recognition

An example of Activity Recognition

`breaking egg`

An Example of Goal/Plan Recognition

from Miquel Ramirez's thesis

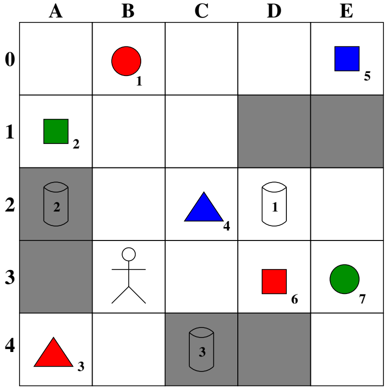

Wooden pieces $p_1, p_2, \ldots p_n$

Pieces have shapes and colors

Bins $b_1, b_2, \ldots, b_n$

The possible **goals** the trainer expected to pursue:

1. Store all triangles in b_1
2. Store all spheres in b_2
3. Store all cubes in b_3
4. Store red objects in b_2
5. Store green objects in b_3
6. Store blue objects in b_1

An Example of Goal/Plan Recognition

from Miquel Ramirez's thesis

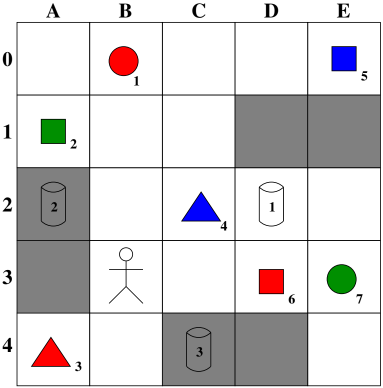

Wooden pieces $p_1, p_2, \ldots p_n$
Pieces have shapes and colors
Bins $b_1, b_2, \ldots, b_n$

One possible *plan* for the trainer to achieve task #1
(store all triangles in b_1):

1. Walk from B3 into A4
2. Pick p_3 up
3. Walk from A4 into B3
4. Walk from B3 into C2
5. Pick p_4 up
6. Throw p_3 into b_1
7. Throw p_4 into b_1

An Example of Goal/Plan Recognition

from Miquel Ramirez's thesis

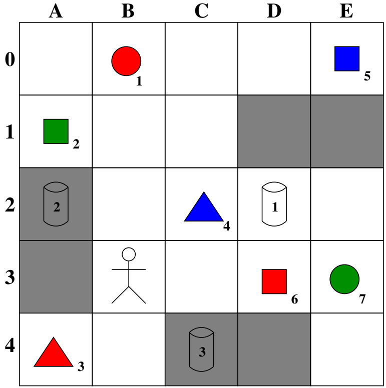

Wooden pieces $p_1, p_2, \dots p_n$
Pieces have shapes and colors
Bins $b_1, b_2, \dots, b_n$

If sensors miss 70% of *walk* actions and half *pick* and *drop* actions, we may only see:

1. Pick p_3 up
2. Walk from A4 into B3

An Example of Goal/Plan Recognition

from Miquel Ramirez's thesis

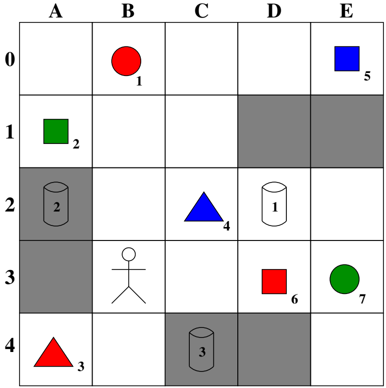

Wooden pieces $p_1, p_2, \ldots p_n$

Pieces have shapes and colors

Bins $b_1, b_2, \ldots, b_n$

If sensors miss 70% of *walk* actions and half *pick* and *drop* actions, we may only see:

1. Pick p_3 up
2. Walk from A4 into B3

Here, we could deduce either task #1 or #4 (store all red objects in b_2), as other tasks are less *likely*.

Table of Contents

Motivation

- In this work, we use a **planning domain definition** to represent agent behavior and environment properties;
- Previous approaches involve multiple calls to a modified planner.
- Our main contribution is twofold:
  - We **obviate the need to execute a planner multiple times** for recognizing goals; and
  - We develop novel goal recognition heuristics that **use planning landmarks**.
- We show that our approaches are **more accurate** and **orders of magnitude faster** than Ramírez and Geffner's approach.

Computing Achieved Landmarks

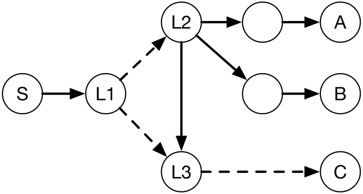

- Our heuristics require identifying which fact landmarks have been achieved during the observed plan execution for every candidate goal $G \in \mathcal{G}$;
- For every candidate goal $G \in \mathcal{G}$:
  - Extract *ordered* landmarks for G;
  - Use achieved landmarks of G in preconditions and effects of every observed action $o \in O$;
  - Under partial observability, we deal with missing actions by inferring that predecessors of observed landmarks must have been achieved;

Landmark-Based Goal Completion Heuristic

- Goal Completion h_{gc} aggregates the percentage of completion of each sub-goal into an overall percentage of completion for all facts of a candidate goal;

$$h_{gc}(G, \mathcal{AL}_G, \mathcal{L}_G) = \left(\frac{\sum_{g \in G} \frac{|\mathcal{AL}_g \in \mathcal{AL}_G|}{|\mathcal{L}_g \in \mathcal{L}_G|}}{|G|} \right) \tag{1}$$

where:

- $\mathcal{AL}_G$ achieved landmarks for goals in G
- $\mathcal{L}_G$ all landmarks for goals in G

Landmark-Based Uniqueness Heuristic (1 of 2)

- Our second heuristic computes **landmark uniqueness**: inverse frequency of a landmark within landmarks for candidate goals:

$$L_{Uniq}(L, \mathcal{L}_{\mathcal{G}}) = \left(\frac{1}{\sum\limits_{\mathcal{L} \in \mathcal{L}_{\mathcal{G}}} |\{L | L \in \mathcal{L}\}|} \right) \quad (2)$$

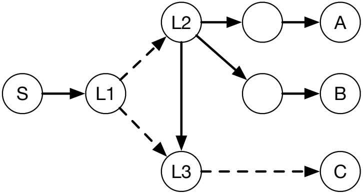

$$L_{Uniq}(L2) = 1/2$$
$$L_{Uniq}(L1) = 1/3$$
$$L_{Uniq}(L3) = 1$$

Landmark-Based Uniqueness Heuristic (2 of 2)

- Our second heuristic, called h_{uniq}, estimates the goal completion of a candidate goal G by calculating the ratio between the sum of the uniqueness value of the achieved landmarks of G and the sum of the uniqueness value of all landmarks of G;

$$h_{uniq}(G, \mathcal{AL}_G, \mathcal{L}_G, \Upsilon_{uv}) = \left(\frac{\sum_{\mathcal{A}_L \in \mathcal{AL}_G} \Upsilon_{uv}(\mathcal{A}_L)}{\sum_{L \in \mathcal{L}_G} \Upsilon_{uv}(L)} \right) \tag{3}$$

where:

- Υ_{uv} is a table of uniqueness values
- $\mathcal{AL}_G$ achieved landmarks for goals in G
- $\mathcal{L}_G$ all landmarks for goals in G

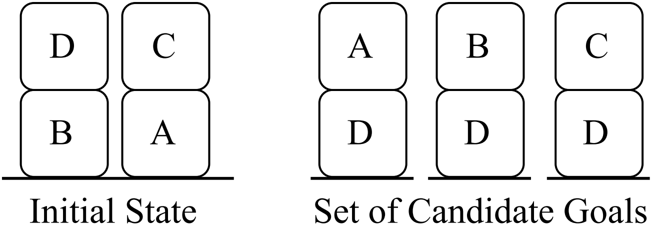

- Observations:
  - `(unstack D B)`; and
  - `(unstack C A)`.
- The real goal is: `(and (ontable D) (on C D) (clear C))`

Example (2 of 4)

Achieved Landmarks in Observations:

- (and (ontable D) (clear A) (on A D)), 5 out of 8:
  - [(clear A)], [(clear A) (ontable A) (handempty)], [(on C A) (clear C) (handempty)], [(holding D)], [(clear D) (on D B) (handempty)]
- (and (ontable D) (clear B) (on B D)), 4 out of 7:
  - [(clear B)], [(ontable B) (handempty)], [(on D B) (clear D) (handempty)], [(holding D)]
- (and (ontable D) (clear C) (on C D)), 5 out of 7:
  - [(clear C)], [(clear C) (on C A) (handempty)], [(clear D) (holding C)] [(clear D) (on D B) (handempty)], [(holding D)]

Example (3 of 4) - h_{gc}

Landmark-Based Goal Completion Heuristic

- `(and (ontable D) (clear A) (on A D))`:
  - Goal Completion: 0.7222
- `(and (ontable D) (clear B) (on B D))`:
  - Goal Completion: 0.6666
- `(and (ontable D) (clear C) (on C D))`:
  - **Goal Completion: 0.7777 (highest estimated value)**

Landmark-Based Uniqueness Heuristic

- `(and (ontable D) (clear A) (on A D))`, $\text{Total}_{Uniq} = 5.5$:
  - `[(clear A)] = 1`, `[(clear A) (ontable A) (handempty)] = 1`, `[(on C A) (clear C) (handempty)] = 0.5`, `[(holding D)] = 0.3333`, `[(clear D) (on D B) (handempty)] = 0.3333`
  - $h_{uniq} = 3.1666 \,/\, 5.5 = 0.5757$
- `(and (ontable D) (clear B) (on B D))`, $\text{Total}_{Uniq} = 5$:
  - `[(clear B)] = 1`, `[(ontable B) (handempty)] = 1`, `[(on D B) (clear D) (handempty)] = 0.3333`, `[(holding D)] = 0.3333`
  - $h_{uniq} = 2.6666 \,/\, 5 = 0.5333$
- `(and (ontable D) (clear C) (on C D))`, $\text{Total}_{Uniq} = 4.5$:
  - `[(clear C)] = 1`, `[(clear C) (on C A) (handempty)] = 0.5`, `[(clear D) (holding C)] = 1`, `[(holding D)] = 0.3333` `[(clear D) (on D B) (handempty)] = 0.3333`
  - $h_{uniq} = 3.1666 \,/\, 4.5 = 0.71$

Recognized `(and (ontable D) (clear C) (on C D))` **with:** $h_{uniq} = 0.71$

Experiments and Evaluation

- We evaluate our heuristics over datasets with 15 planning domains (6 of these domains from original Ramírez and Geffner paper):
  - Blocks-World, Campus, Depots, Driver-Log, Dock-Worker-Robots, Easy-IPC-Grid, Ferry, Intrusion-Detection, Kitchen, Logistics, Miconic, Rovers, Satellite, Sokoban, and Zeno-Travel;
- These datasets contain hundreds of goal recognition problems, varying the observability (10%, 30%, 50%, 70%, and 100%);
- We compared our heuristics against the original approach of Ramírez and Geffner (Plan Recognition as Planning. IJCAI, 2009), which is their fastest and most accurate approach;

Experiments and Evaluation - ROC Space (1 of 2)

- Results of our heuristics use threshold $\theta = 20\%$;
- We compare Ramírez and Geffner's approach over ROC space, which shows the trade-off between TPR and FPR;
- We aggregate multiple domains and plot these goal recognition results in ROC space.

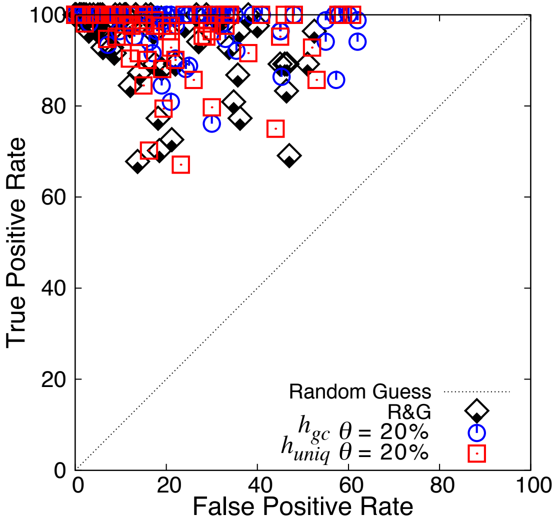

True Positive Rate
False Positive Rate
0
20
40
60
80
100
Random Guess
R&G
h_{gc} θ = 20%
h_{uniq} θ = 20%

Experiments and Evaluation - Recognition Time

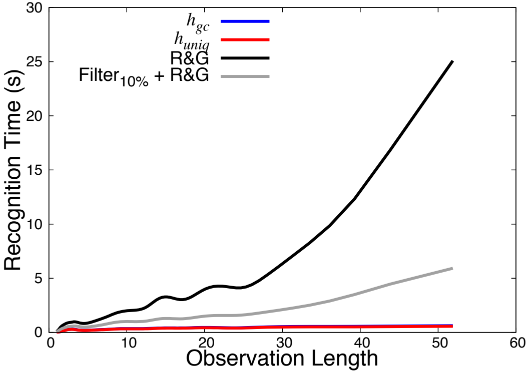

Experiments and Evaluation - Recognition Time with Noise

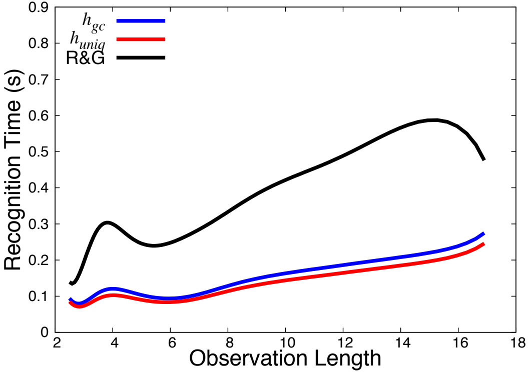

Contributions and Limitations

- **Contribution so far:**
  - Use planning landmarks for goal recognition;
  - Obviate the need to run a planner during goal recognition, resulting in much faster and highly accurate recognition; and
  - Robust dataset to evaluate goal recognition algorithms
- **Limitations:**
  - Sensitive to the presence of landmarks; and
  - Low accuracy with very few observations, *i.e.*, 10% of observability;

Table of Contents

Motivation for Efficient Online Goal Recognition

Most goal recognition approaches using domain models have three key limitations:

1. assumption of a discrete state-space in a PDDL-like formalism
   - not viable for use with path planning scenarios
2. assume all access to all observations at once
   - approaches do not consider the time to recognition
3. need to call a planner multiple times per goal to rank hypotheses
   - PRAP is computationally expensive, impractical for long plans

Online vs. Offline Plan Recognition

- Offline plan recognition:
  - All observations received at once;
  - Observations may be incomplete or noisy;
  - One-shot recognition;
- Online plan recognition:
  - Observations received incrementally;
  - Observations may be incomplete or noisy;
  - Objective is to recognize goal as soon as possible, without the full observation sequence

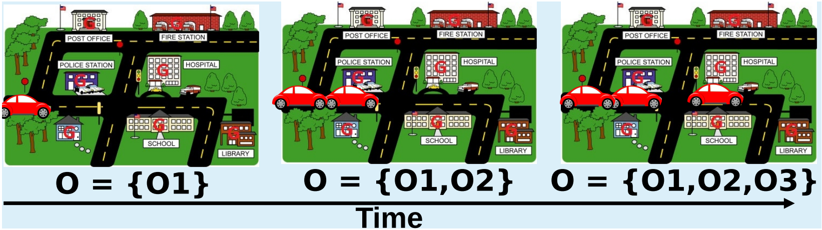

Efficient Online Goal Recognition

Our approach:

- is efficient for online goal recognition;
- works in both discrete and continuous domains;
- minimizes planner calls;
- reasons about landmarks to minimize the number of goal hypotheses;
- returns reliable goal ranking as soon as possible

Landmarks in Continuous Domains

We need a notion of landmark in continuous domains

- Redefine landmarks as areas surrounding goals
  - Goals – Black dots
  - Surrounding Rectangles – continuous landmark areas
- To reach a goal the observed motion must intersect (go through) the corresponding landmark area.
- In this work, landmark areas roughly correspond to rooms partitioned as rectangular Voronoi diagrams
  - Other notions of numeric landmarks may apply (e.g. Scala et al. IJCAI 2017)

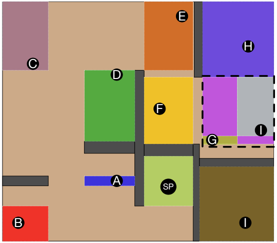

Online Recognition with Landmarks

- Generate the ordered set of achieved landmarks
- Maintain the group of goals eliminated due to landmarks
- For every observation:
  - Check if it "achieved" a landmark
  - If observations backtrack, re-instate goals
- Rank goals using the landmark completion heuristic h_{gc}

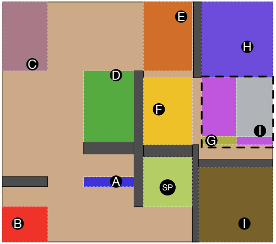

Goal Mirroring with Landmarks

Combines landmark reasoning with goal mirroring

- Compute landmarks and optimal plans for all goals
- For every observation:
  - Compute plan prefix, and for every goal
    - Either prune goals that have **passed** the last landmark; or
    - Compute plan suffix (from last observation) using planner
    - Compute **cost ratio** between prefix+suffix and optimal plan

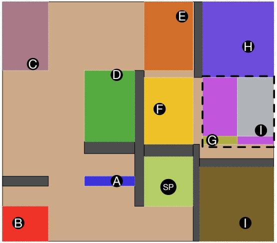

- Rank unpruned goals based on a **normalized cost ratio**
  - Ranks $P(g_k \mid O)$ using a normalizing factor $\eta 1/\sum_{g_k \in G} rank(g_k)$
  - Approximates $P(g \mid O) = \eta \sum_{g_k \in G} P(O \mid g_k) P(g_k)$ for all goals, assuming $P(g_k) = 0$ for pruned goals

Continuous Evaluation

- Cubicles environment and robot (OMPL)
- 11 points spread evenly over the environment
- 220 problems

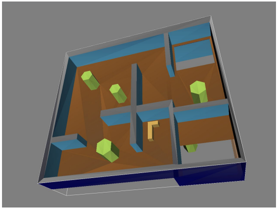

Discrete Evaluation

- Dataset expanded from Ramirez and Geffner's original work
- Domains extracted from the IPC competition
- Hundreds of goal recognition problems

Domains

- Blocks-World
- Campus
- Depots
- Driver-Log
- Dock-Worker-Robots
- Easy-IPC-Grid
- Ferry
- Intrusion-Detection
- Kitchen
- Logistics
- Miconic
- Rovers
- Satellite
- Sokoban; and
- Zeno-Travel

Performance Results

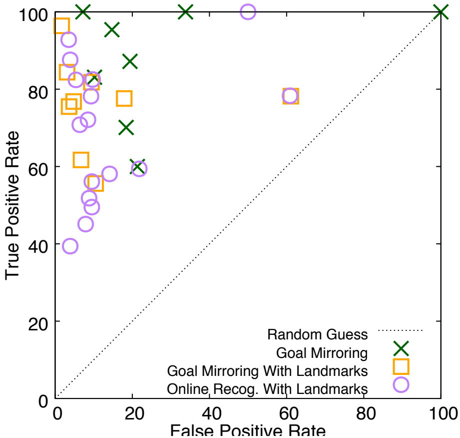

Efficiency Results

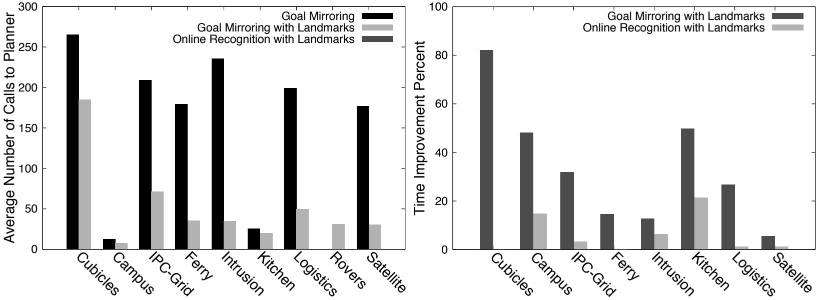

Contributions and Limitations

- **Contribution so far:**
  - Extended de idea of landmarks for continuous domains; and
  - Developed online algorithms able to recognize plans in discrete and continuous domains;
  - **Very** efficient in both discrete and continuous domains.
- **Limitations:**
  - Naive notion of spatial landmarks;
  - Much better performance on discrete domains.

Table of Contents

Plan Recognition using Video Data

- **Plan recognition**
  - Task of recognizing the plan (i.e., the sequence of actions) the observed agent is following in order to achieve his intention (Sadri, 2012)
- **Activity recognition**
  - The task of recognizing the independent set of actions that generates an interpretation to the movement that is being performed (Poppe, 2010)
  - Such task is particularly challenging in the real physical world
- Much research effort focuses on activity and plan recognition as separate challenges;
- We develop a hybrid approach that comprises both activity and plan recognition;
- The approach infers, from a set of candidate plans, which plan a human subject is pursuing based exclusively on fixed-camera video.

Poppe, R. A survey on vision-based human action recognition. Image and Vision Computing 28(6), pp. 976–990, 2010.
Sadri, Fariba. Intention Recognition in Agents for Ambient Intelligence: Logic-Based Approaches.

Ambient Intelligence and Smart Environments, pp. 197-236, 2012.

- **Conceptually divided in two main parts**
  - CNN-based activity recognition (CNN)
  - CNN-backed symbolic plan recognition (SBR)

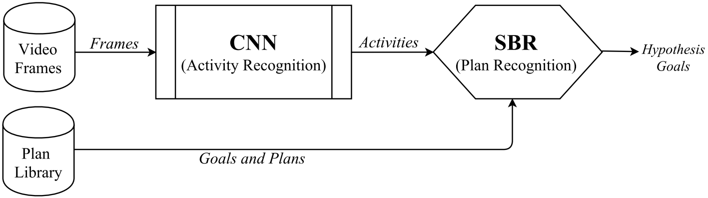

Experiments: Dataset

- ICPR 2012 Kitchen Scene Context based Gesture Recognition dataset (KSCGR)
- **5 recipes for cooking eggs in Japan**
  - Ham and Eggs, Omelet, Scrambled-Egg, Boiled-Egg and Kinshi-Tamago
  - Each recipe is performed by 7 subjects (5 in training set, 2 in testing set)
- **9 cooking activities composes the dataset**
  - Breaking, mixing, baking, turning, cutting, boiling, seasoning, peeling, and none

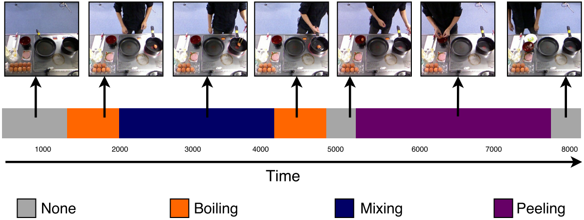

Summary of the Results

Conducted experiments on two levels:

- Activity Recognition
  - Accuracy lower than 50% (in 9-label classification) for infrequent activities
  - Very good accuracy to identify "no-action"
- Overall Plan Recognition
  - Low accuracy for overall plan recognition using plan-libraries

Contributions and Future Work

- We developed a hybrid architecture for activity and plan recognition
- Pipeline includes:
  - A CNN for activity recognition that feeds directly into:
  - a modified (SBR) approach that uses the CNN to index activities in the plan library
- Approach limited by the plan library in the plan recognizer
- Next steps:
  - Employ other deep learning architectures such as Long-Short Term Memory networks (LSTM) and 3D CNNs
  - Use a more flexible approach for plan recognition, such as PRAP
  - Explore object recognition to provide additional clues of the activity that is being performed

Demo video: `https://youtu.be/BoiLjU1vg3E`

Table of Contents

Papers reporting these results

FRAGA PEREIRA, Ramon; MENEGUZZI, Felipe. **Landmark-based Plan Recognition.** ECAI, 2016.

PEREIRA, Ramon F.; OREN, Nir; and MENEGUZZI, Felipe. **Landmark-Based Heuristics for Goal Recognition**. AAAI, 2017.

PEREIRA, Ramon F.; OREN, Nir; and MENEGUZZI, Felipe. **Monitoring Plan Optimality using Landmarks and Domain-Independent Heuristics**. PAIR Workshop@AAAI, 2017.

GRANADA, Roger L.; PEREIRA, Ramon F.; MONTEIRO, Juarez; BARROS, Rodrigo; RUIZ, Duncan; and MENEGUZZI, Felipe. **Hybrid Activity and Plan Recognition for Video Streams**. PAIR Workshop@AAAI, 2017.

PEREIRA, Ramon F.; OREN, Nir; and MENEGUZZI, Felipe. **Detecting Commitment Abandonment by Monitoring Plan Execution**. AAMAS, 2017.

MONTEIRO, Juarez; GRANADA, Roger; BARROS, Rodrigo and MENEGUZZI, Felipe. **Deep Neural Networks for Kitchen Activity Recognition**. IJCNN, 2017.

VERED, Mor; PEREIRA, Ramon F.; MAGNAGUAGNO, Maurício C.; KAMINKA, Gal; and MENEGUZZI, Felipe. **Online Goal Recognition Combining Landmarks and Planning**. GRW@IJCAI, 2017.

Summary

- We progressively relaxed many assumptions about plan recognition:
  - Domain knowledge
  - Quality of observations
  - Exclusively discrete domains
  - Precise domain knowledge
- We illustrated applications of these techniques:
  - Real world video-data
  - Multiagent systems

Future Directions

- Plan Recognition with Domain Theories
  - Use different landmark extraction algorithms;
  - Extend landmark-based heuristics to temporal and non-uniform-cost domains
  - Experiment with more advanced notions of numeric landmarks (e.g. Scala et al.)
- Applications of Plan Recognition
  - Use object recognition techniques (deep learning) to generate fact observations in video
  - Couple the above with plan recognition in domain theories
  - Do plan recognition in latent space

Thanks and Acknowledgement

People involved in this research

- Ramon Fraga Pereira (PhD Student)
- Mor Vered (PhD Student, Bar Ilan University)
- Maurício Magnaguagno (PhD Student)
- Juarez Monteiro (MSc Student)
- Roger Granada (Postdoc)
- Gal Kaminka (Bar Ilan University)
- Nir Oren (University of Aberdeen)
- Rodrigo Barros (PUCRS colleague)
- Duncan Ruiz (PUCRS colleague)

Institutions

- The Scottish Informatics and Computer Science Alliance (SICSA) Distinguished Visiting Fellowship (DVE)
- Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq) – PQ Fellowship

Master Class Plug

If this talk was interesting and you want to know more, please come to:

Plan Recognition Master Class

University of Aberdeen – 16th October 2017

We will cover:

- Detailed algorithms
- Worked out examples
- Plan recognition with incomplete **domains**
- Much more

Thank you!
Questions?