

Accurately Determining Intermediate and Terminal Plan States Using Bayesian Goal Recognition



David Pattison and Derek Long

University of Strathclyde, Glasgow
G1 1XH, UK

david.pattison@cis.strath.ac.uk

GAPRec Workshop
ICAPS 2011, Freiburg

12th June, 2011

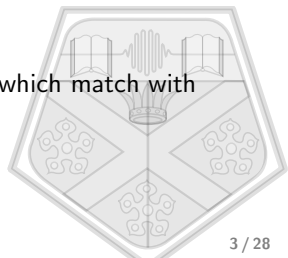
Overview

- 1 Recognition without Libraries
- 2 Results
- 3 Conclusions and Future Possibilities



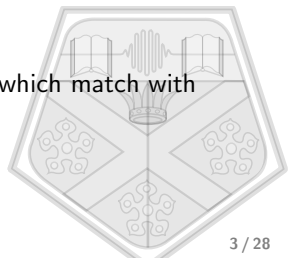
The *de facto* (and defined) standard

- Traditional GR/PR makes use of libraries
 - Collection of known goals/plans
 - Hand coded or generated
 - Plans through state space
 - Specialised to one subject
 - Represented as HTNs
- Recognition
 - Probabilistic/Bayesian
 - Weights hand coded or automated
 - Observe actions and map to X plans from library which match with varying probabilities



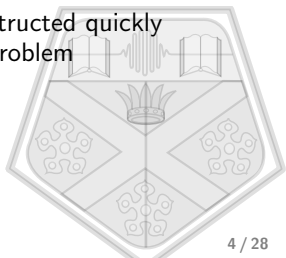
The *de facto* (and defined) standard

- Traditional GR/PR makes use of libraries
 - Collection of known goals/plans
 - Hand coded or generated
 - Plans through state space
 - Specialised to one subject
 - Represented as HTNs
- Recognition
 - Probabilistic/Bayesian
 - Weights hand coded or automated
 - Observe actions and map to X plans from library which match with varying probabilities
- But what if there is nothing to map to?

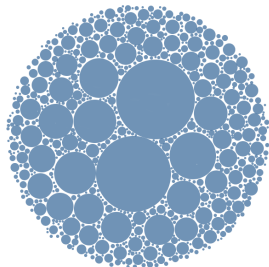


Recognition without Libraries

- Goal Recognition as Planning
 - “Planning” in the sense of not doing any planning
- Planning and Recognition mirror one-another
 - Goal Recognition also uses Propositions, Actions, States and Goals
 - So why not try to link the two?
 - Recognition systems have no common language, but Planning has PDDL
 - By working with PDDL, any problem can be constructed quickly
 - Use recent Planning advances in solving the GR problem
 - *heuristic convergence*
- No plan/goal library
 - Try to automatically detect lost information



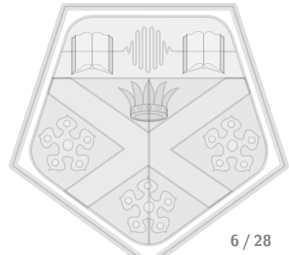
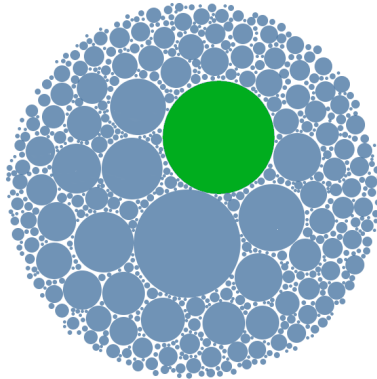
Problem Formulation



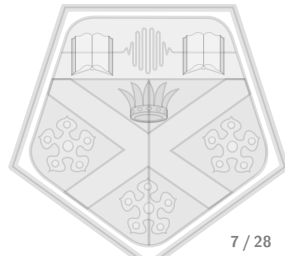
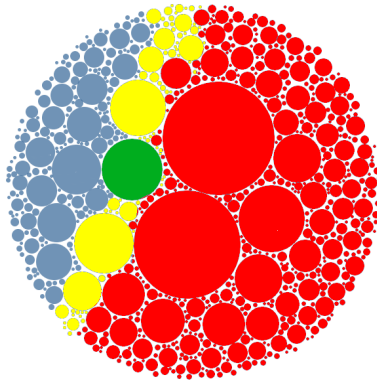
- No libraries
- Any domain
- No pre-compilation
- Any (valid) fact conjunctions can be goal
- Use Planning representation for *goal space*
 - Cannot hope to enumerate the true goal space
 - Goal Space \mathcal{H} = domain's reachable facts
 - Assume independence between facts
 - No explicit conjunctions (yet)
 - Standard mutex detection
- Also analogous to Particle Filtering and Fault Diagnosis



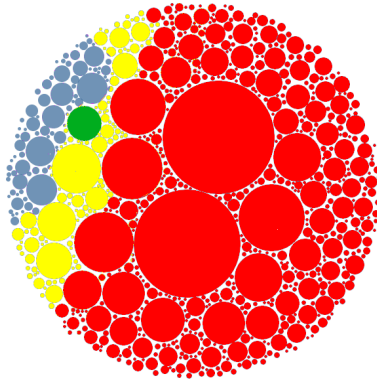
Plan movement through state-space



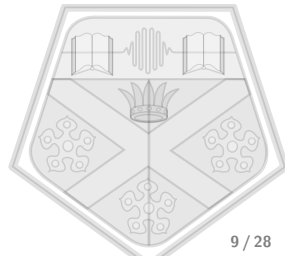
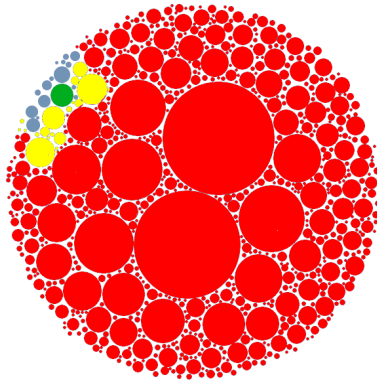
Plan movement through state-space



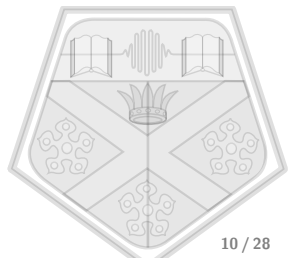
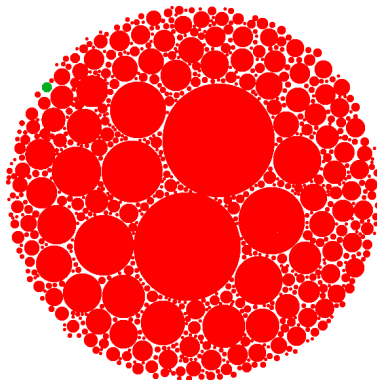
Plan movement through state-space



Plan movement through state-space



Plan movement through state-space



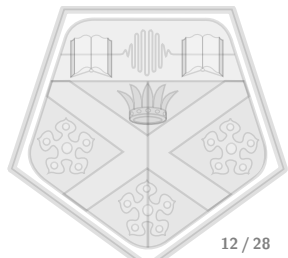
Assumptions and Relaxations

- Plan is totally-ordered
 - Can be taken from anywhere- created or parsed in from known results
 - We use IPC3/IPC5 results
- Fully observable
 - No hidden actions
- No assumption about “intelligence” of plan
- No knowledge of plan steps remaining
- **Anything** can be a goal, and a goal can be made up of anything
 - Conjunctions are common in Planning, but uncommon in Recognition



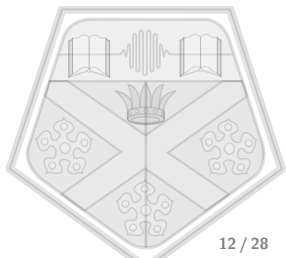
Step 1 – Putting the Vitamins back in

- Cue strange orange juice analogy...



Step 1 – Putting the Vitamins back in

- Cue strange orange juice analogy...
- PDDL domain inputs are flat and dull
- But once instantiated, structure is rich, albeit hard to find



Step 1 – Putting the Vitamins back in

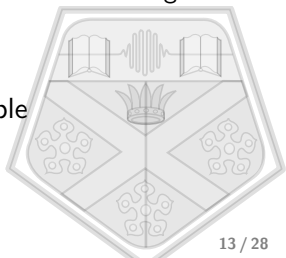


- Cue strange orange juice analogy...
- PDDL domain inputs are flat and dull
- But once instantiated, structure is rich, albeit hard to find
- Domain Transition Graphs, Causal Graphs, Static Facts, Relaxed Plans, Heuristic Estimates, Sampling



Domain Analysis

- Predicate Partitioning
 - Grounding process produces all goals
 - So try and categorise them to find those which are very likely and those which are less likely
- Causal Graph Leaf-Nodes
 - Exist only to be altered, so adjust probabilities of facts containing them appropriately
- Produce **initial probability distribution** over \mathcal{H}
- But of course a manual distribution is still possible



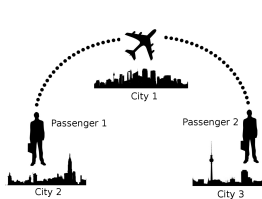
Step 2 – Plan Observation

- Action is fed into recogniser
- Get *heuristic estimate* to all $f \in \mathcal{H}$
 - Further actions needed to achieve f
 - If decreasing, fact is possibly goal
 - If increasing, fact is probably not goal
- Use heuristic results to increase/decrease probability if f being a goal w.r.t. mutually-exclusive facts
- Over time, some facts will become highly likely to be goals
 - ... or at least be in final state
- Heuristic estimates used to update goal probabilities using Bayes'



Heuristic Bayesian Updates

- After each observation, a subset of the search-space will be closer
- The amount of work performed by an action w.r.t G is



$$W(G|O) = \begin{cases} \frac{1}{|\bar{G}_{mutex}^{nearer}|} & \text{if } h_t(G) < h_{t-1}(G), \\ \frac{1}{|\bar{G}_{mutex}^{nearer}|} & \text{if } h_t(G) = h_{t-1}(G) = 0, \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

- Give a *bonus* to facts which remain true

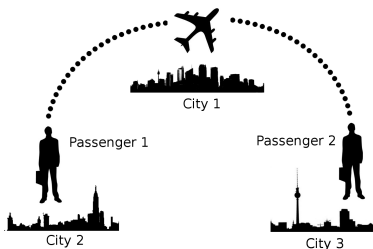
Example of $W(G)$ with and without bonus

Table: Without bonus

	at p2 c1	at p2 c2	at p2 c3	in plane p2
1	0.33	0.33	0	0.33
2	0.33	0.33	0	0.33
3	0.5	0.5	0	0
4	1	0	0	0
5	0	0	0	0
6	0	0	0	0
7	0	0.33	0.33	0.33
8	0	0	0	0

Table: With bonus

	at p2 c1	at p2 c2	at p2 c3	in plane p2
1	0.25	0.25	0.25	0.25
2	0.33	0.33	0	0.33
3	0.33	0.33	0	0.33
4	1	0	0	0
5	1	0	0	0
6	1	0	0	0
7	0.25	0.25	0.25	0.25
8	1	0	0	0



- Goal: Passenger 1 and Passenger 2 at City 1
- $W(G)$ associated with Passenger 2

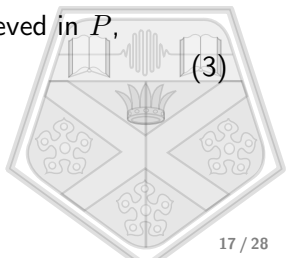
Is O relevant if G is goal

- Feed into *conditional probability*

$$P(O|G) = \lambda * W(G|O) * S(G) + (1 - \lambda) * \frac{1}{1 + |mutex(g)|} \quad (2)$$

- Stability $S(G)$ indicates how often a fact flicks from true to false

$$S_t(G) = \begin{cases} 1 & \text{if } G \text{ unachieved in } P, \\ \frac{|Obs| - G_t^{true}}{\sum G_i^{true}} & \text{otherwise} \end{cases} \quad (3)$$



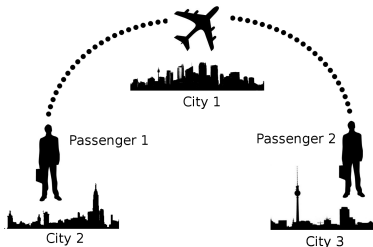
Example of $P(G | A)$ with and without bonus

Table: Without bonus

	at p2 c1	at p2 c2	at p2 c3	in plane p2
init	0.25	0.25	0.25	0.25
1	0.25	0.25	0.25	0.25
2	0.32	0.32	0.05	0.32
3	0.33	0.33	0.01	0.33
4	0.89	0.05	0.00	0.05
5	0.89	0.05	0.00	0.05
6	0.89	0.05	0.00	0.05
7	0.63	0.18	0.00	0.18
8	0.63	0.18	0.00	0.18

Table: With bonus

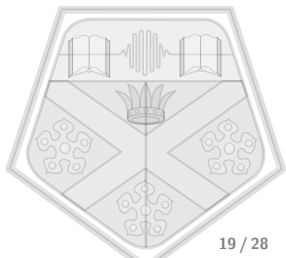
	at p2 c1	at p2 c2	at p2 c3	in plane p2
init	0.25	0.25	0.25	0.25
1	0.25	0.25	0.25	0.25
2	0.32	0.32	0.05	0.32
3	0.33	0.33	0.01	0.33
4	0.89	0.05	0.00	0.05
5	0.99	0.00	0.00	0.00
6	1.00	0.00	0.00	0.00
7	1.00	0.00	0.00	0.00
8	1.00	0.00	0.00	0.00



- Goal: Passenger 1 and Passenger 2 at City 1
- $P(G | A)$ associated with Passenger 2

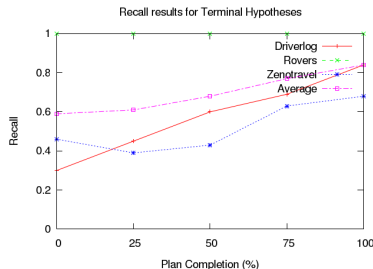
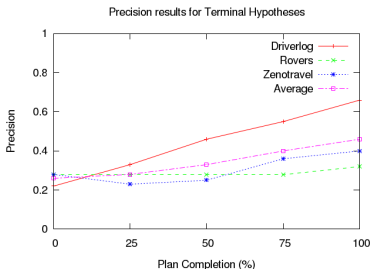
Step 3 – Hypotheses

- Now have a new probability distribution over \mathcal{H}
- Pull out highest probability facts to form *terminal goal hypothesis*



Step 3 – Hypotheses

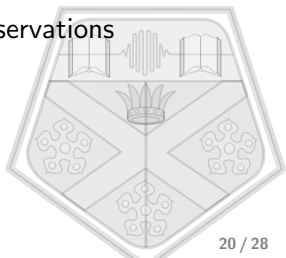
- Now have a new probability distribution over \mathcal{H}
- Pull out highest probability facts to form *terminal goal hypothesis*



Domain	P = 0%	P = 25%	P = 50%	P = 75%	P = 100%
Driverlog	0.22/0.3	0.33/0.45	0.46/0.6	0.55/0.69	0.66/0.84
Rovers	0.28/1	0.28/1	0.28/1	0.28/1	0.32/1
Zenotravel	0.28/0.46	0.23/0.39	0.25/0.43	0.36/0.63	0.4/0.68
Average	0.26/0.59	0.28/0.61	0.33/0.68	0.4/0.77	0.46/0.84

A Step Further

- But we would also like to have hypotheses for non-goal *intermediate states*
- So *estimate* the number of steps remaining based on what the final goal is expected to be
- Can then generate a hypothesis for n further observations

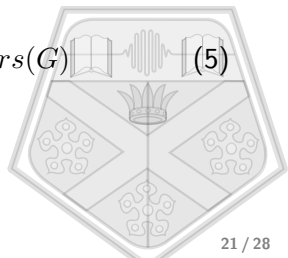


Estimating Intermediate Goals

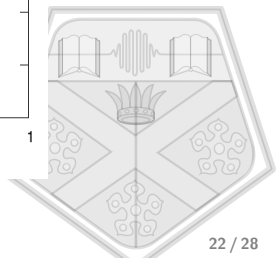
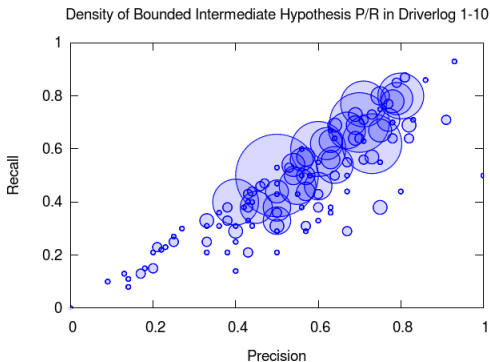
- Estimate whether G will be true in n steps
- Clearly linked to whether action which achieves it will be observed

$$P^n(A) = \begin{cases} 0 & \text{if } h(A_{pre}) > n, \\ \max P(f) \quad \forall f \in A_{add} & \text{otherwise} \end{cases} \quad (4)$$

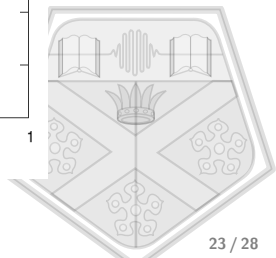
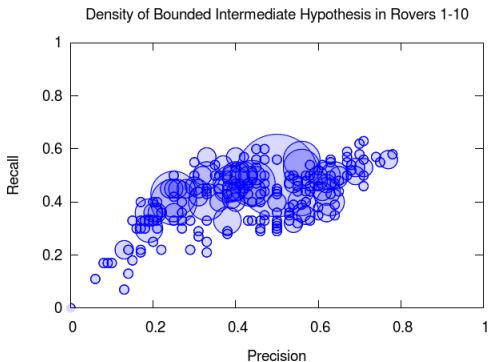
$$P^n(G) = \max P^n(A) \quad \forall A \in \text{achievers}(G) \quad (5)$$



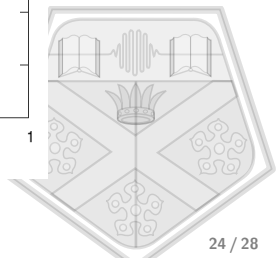
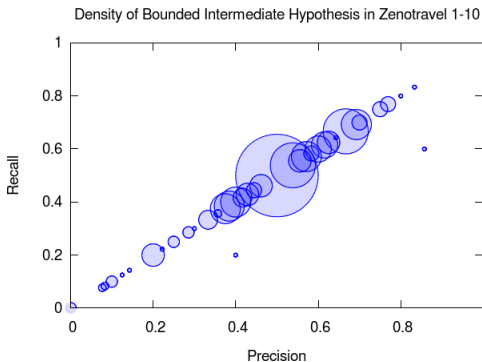
Intermediate Results- Driverlog



Intermediate Results- Rovers

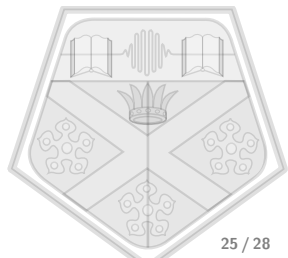


Intermediate Results- Zenotravel



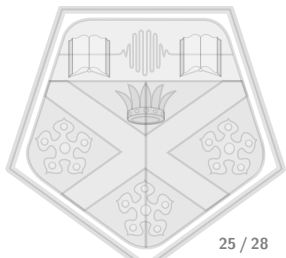
Conclusions

- Presented a new formulation of Goal Recognition as a Planning task, which does not rely on libraries



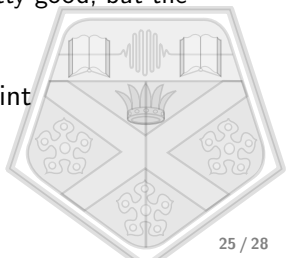
Conclusions

- Presented a new formulation of Goal Recognition as a Planning task, which does not rely on libraries
- How well are Plan Libraries replaced?



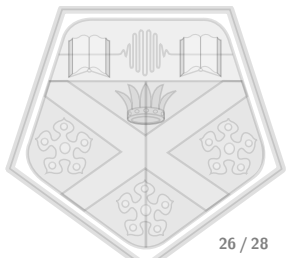
Conclusions

- Presented a new formulation of Goal Recognition as a Planning task, which does not rely on libraries
- How well are Plan Libraries replaced?
 - ① Structure- largely done
 - ② Prediction- Good results for both intermediate and terminal results
 - ③ Abstraction- None really. Could be learned from domains, or probable conjunctions generated at runtime
 - ④ Termination- Intermediate state estimates are pretty good, but the estimation itself is too short
 - Probably heavily linked to heuristic choice
- **Backwards compatibility** not broken at any point
 - Known goal conjunctions can still be added
 - Known plans still applicable
 - Probability weightings still applicable



Extensions

- The move into PR seems natural



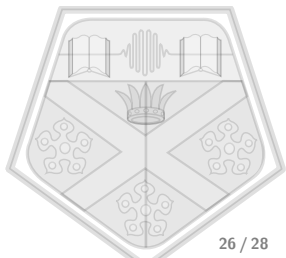
Extensions

- The move into PR seems natural
- Bringing Planning and PR closer together



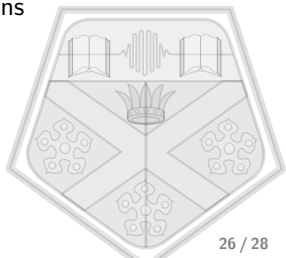
Extensions

- The move into PR seems natural
- Bringing Planning and PR closer together
- Convergence



Extensions

- The move into PR seems natural
- Bringing Planning and PR closer together
- Convergence
 - Instead of storing plans in a library, generate them at runtime
 - Use of landmarks, inference, deduction in next action-prediction
 - “Heuristic learning” from previous plan observations
 - Macro-Actions \Rightarrow high-level concepts?
 - Domain-learning/extension
 - Conjunction learning- genetic techniques



Thank you for your attention

- Questions/comments?



Coffee Break



- Resume at 11.00

