Online Goal Recognition By Mirroring In Continuous Domains

Mor Vered veredm@cs.biu.ac.il
Gal A. Kaminka galk@cs.biu.ac.il
Sivan Biham

BIRC
BIU Robotics Consortium

Real World Applications

- Inferring unobserved goals, based on observed actions

- Recognising intended gestures/sketches

  (Sezgin & Davis, 2005)

- Anticipating user commands (Blaylock & Allen, 2004)
- Recognising navigational goals (Zhu, 1991)

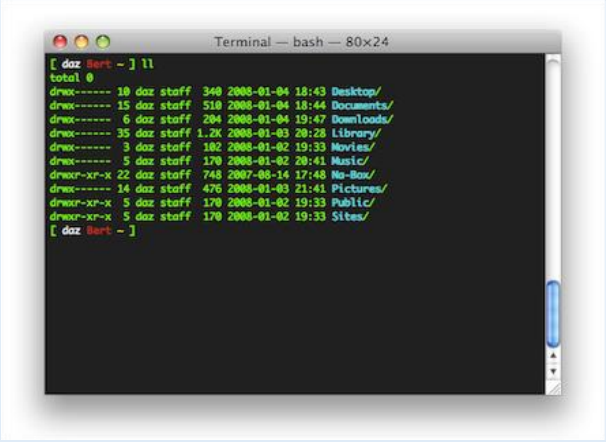

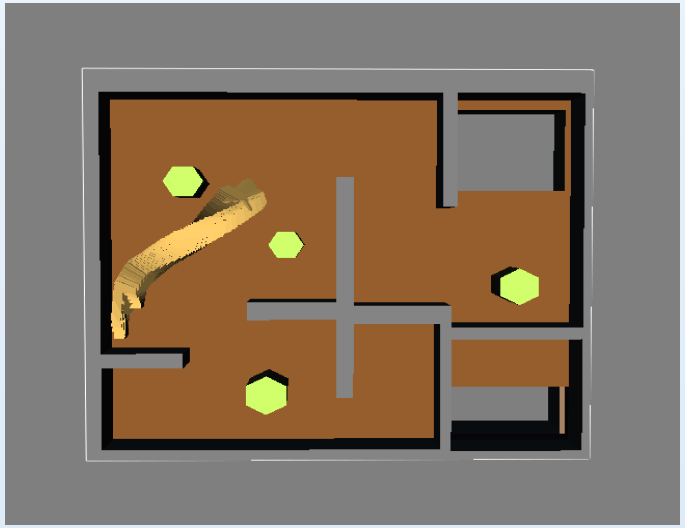

Most Past Approaches

- Dedicated plan library
  - Represents all known plans to achieve known goals
- Redundant : Separate plans for execution and recognition
- Not efficient for continuous domains
  - Where number of possible plans is potentially infinite
- Problem handling new goals
  - Must also receive all possible plans to achieve each new goal

Most Past Approaches

- Dedicated plan library
  - Represents all known plans to achieve known goals
- Redundant : Separate plans for execution and recognition
- Not efficient for continuous domains
  - Where number of possible plans is potentially infinite
- Problem handling new goals
  - Must also receive all possible plans to achieve each new goal

Plan Recognition By Planning

[Ramírez & Geffner, 2010]

- Use planner to generate plans instead of plan library
- Assumes all observations are given at once
- Discrete domains only (STRIPS)
- Fails in continuous environments

Goal Mirroring – Space Efficient Goal Recognition for Continuous Environments

- Uses an existing planner in the recognition process
- No need for library of existing plans
- Easily add new goals
- Whatever can be planned can also be recognized
- Especially efficient for complete agents

Challenges

- Continuous environments
  - Infinite plan possibility
  - Noise in observations and actions
- How to incorporate observation history as input to planner
  - [Ramírez& Geffner (2010)] changed planner domain theory
- Different planners, different representation methods
- No general recognition performance measures
  - Independent of domain, planner and problem
- Space efficiency

Algorithm

Algorithm 1 ONLINE GOAL MIRRORING $(R, planner)$

1: **for all** $g \in G$ **do**
2: $\quad \bar{m}_g \leftarrow planner(W, g, O(\emptyset))$
3: **for** $t = 0$ to T **do**
4: $\quad \Delta \leftarrow cost(O^t)$
5: $\quad$ **for all** $g \in G$ **do**
6: $\quad\quad m'_g \leftarrow planner(W, g, O^t(t))$
7: $\quad\quad score(g) \leftarrow cost(\bar{m}_g)/(\Delta + cost(m'_g))$
8: $\quad P(G|O(t)) \leftarrow \eta \cdot score(g)$

- Once, for each possible Goal *g*
  - Calculate *directPlan*$_g$ from *start* to *g* by running planner
- For each incremental observations
  - Calculate *newPlan*$_g$, using planner, from current state to *g*
  - *Current Cost= cost(newPlan*$_g$ *+ observations seen so far)*
    - *Cost function – domain dependant*
  - *ratio = cost(directPlan*$_g$ */ newPlan*$_g$)

    (consistent with studies on human rational intentionality bias)

Navigational Goal Recognition

Task : identify goal location of an object observed moving in a 3D continuous world

Using 4 ***off the shelf*** planners **RRT***, **TRRT**, **RRTConnect**, **KPIECE1**, Cubicles env. and robot (OMPL)[Sucan, Moll, & Kavraki (2012)]

- Selected 11 points arbitrarily
- Generated observed paths from each point to all others
- 110 recognition problems

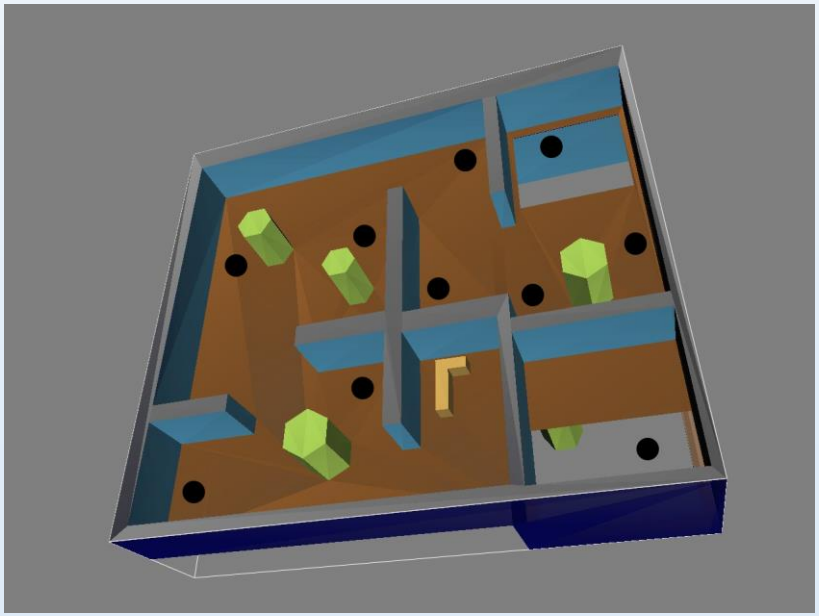

Measuring Recognition Results

Convergence Ratio

- Time the recognizer converged to the correct hypothesis
- Measured by number of current rankings from the end

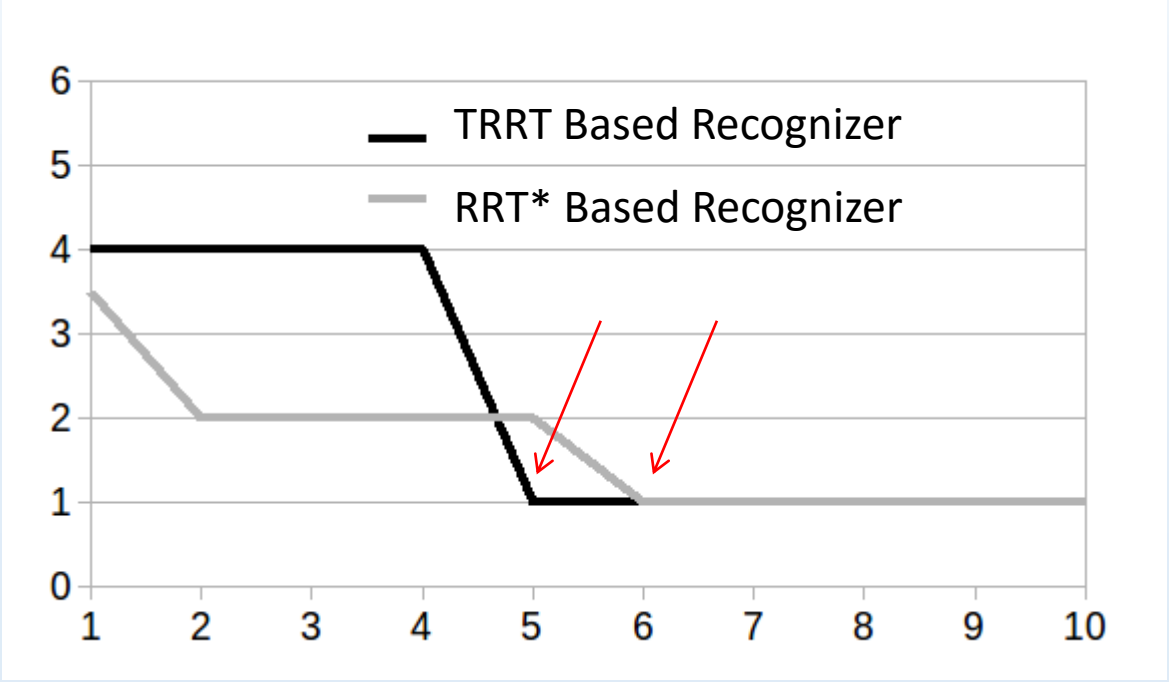

Measuring Recognition Results

AUC – Area Under Curve

- Greater area means the recognizer ranked the correct hypothesis lower
- False positive measure
- Indication as to uncertainty

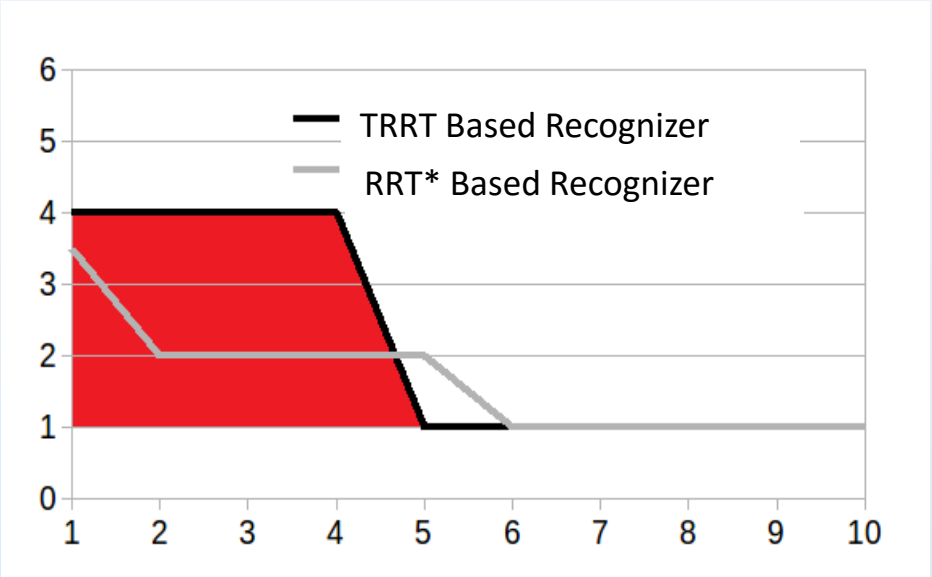

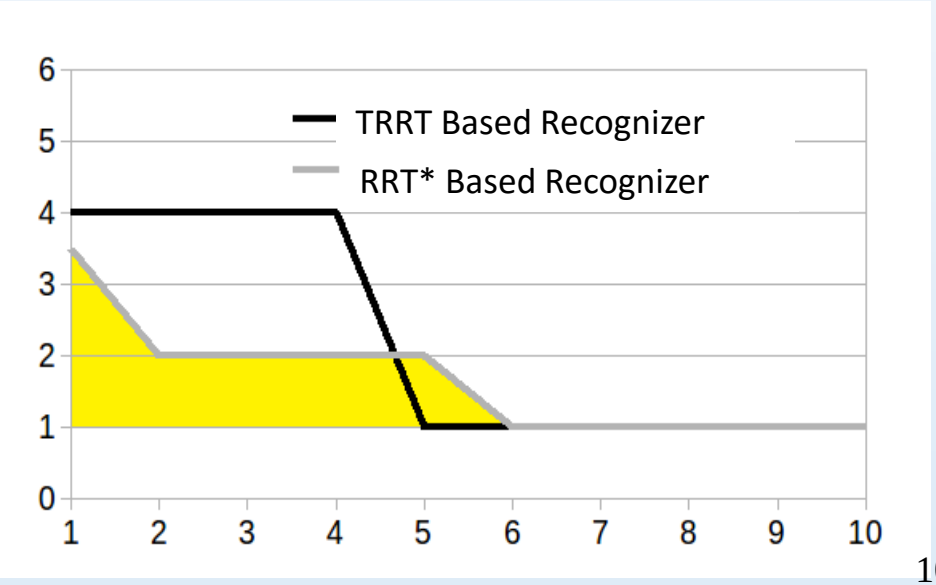

Measuring Recognition Results

Ranked First

- Amount of times ranked first – not consecutively
- Measure of reliability

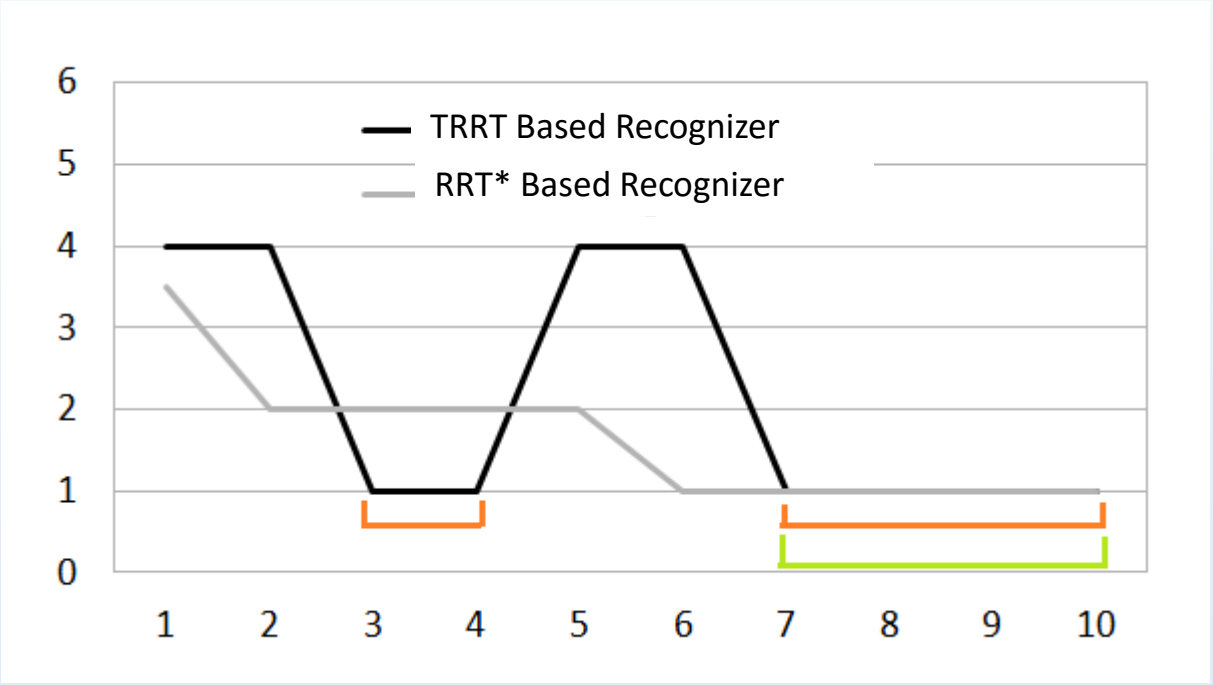

Recognition Results

Planner Comparison

- TRRT and RRT* better
- TRRT, RRT* produce paths closer to optimal

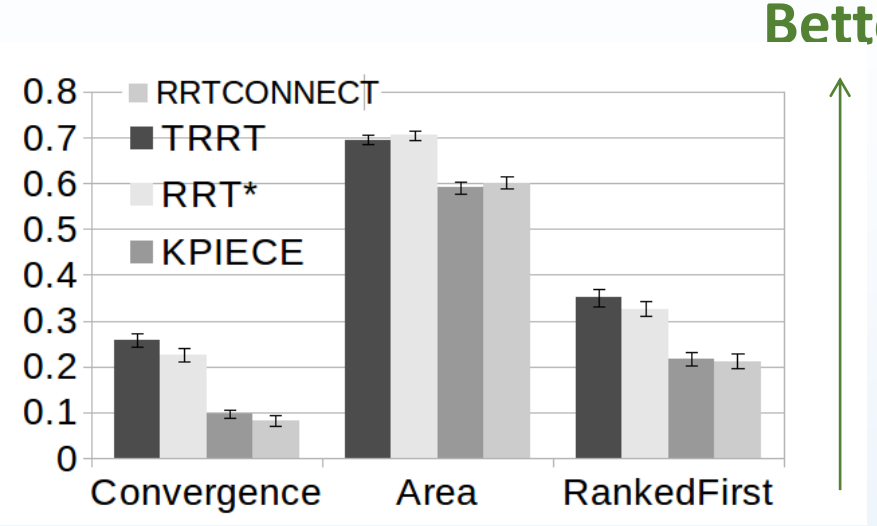

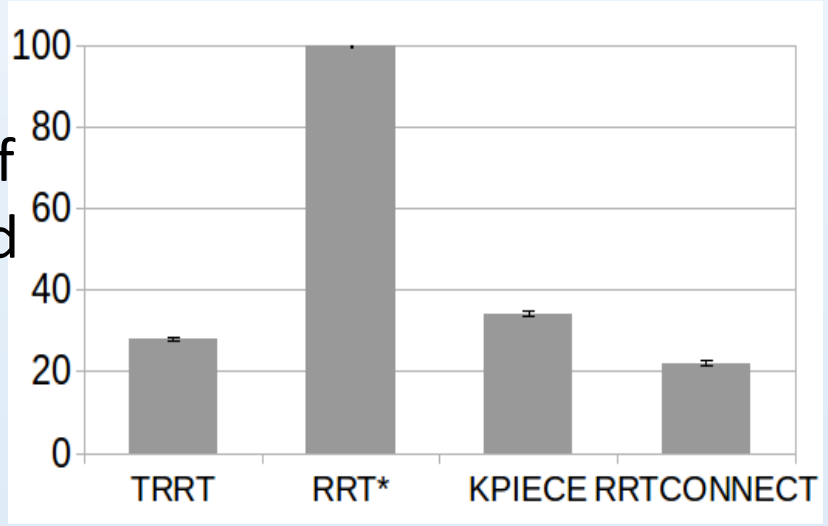

Mean Running Time of the Planners

- Planner limited to 1 sec.
- Mean over 110 problems
- RRT* uses all time allotted

Recognition Results

Sensitivity to Recognition Difficulty

- Added 9 goal points, 380 recognition problems
- Added in close proximity to existing points - clusters
- TRRT more robust in Convergence

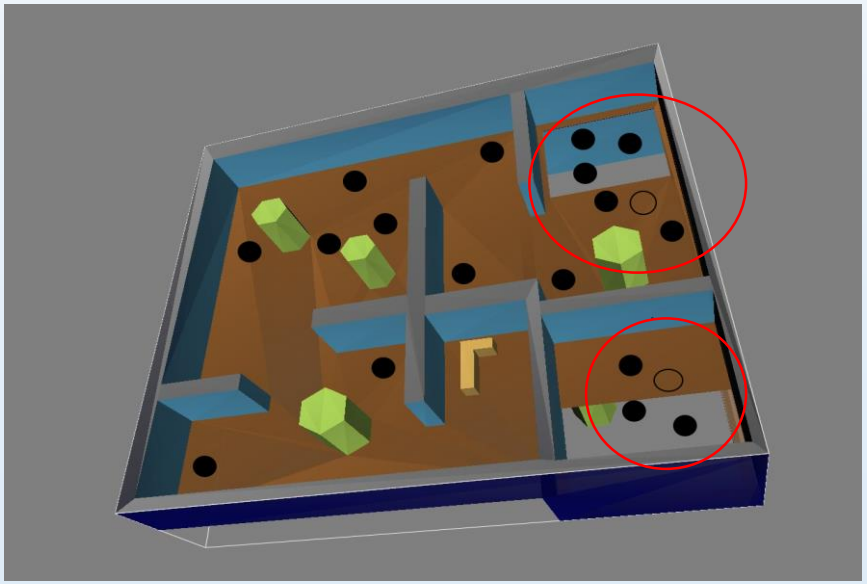

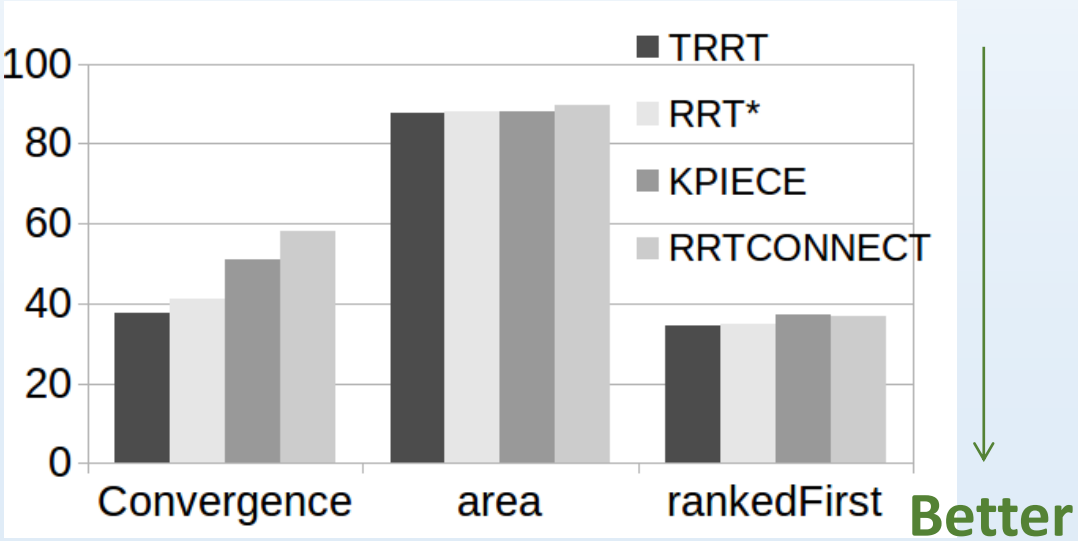

Deterioration Percent

Recognition Results

Goal Mirroring vs Hidden Markov Model

- Discretized the environment
  - Robot-sized cells, each one represented by a state
- HMM training data : 20 paths generated by optimal– RRT*
- Standard MATLAB HMM package
- Mirroring on-par with HMM

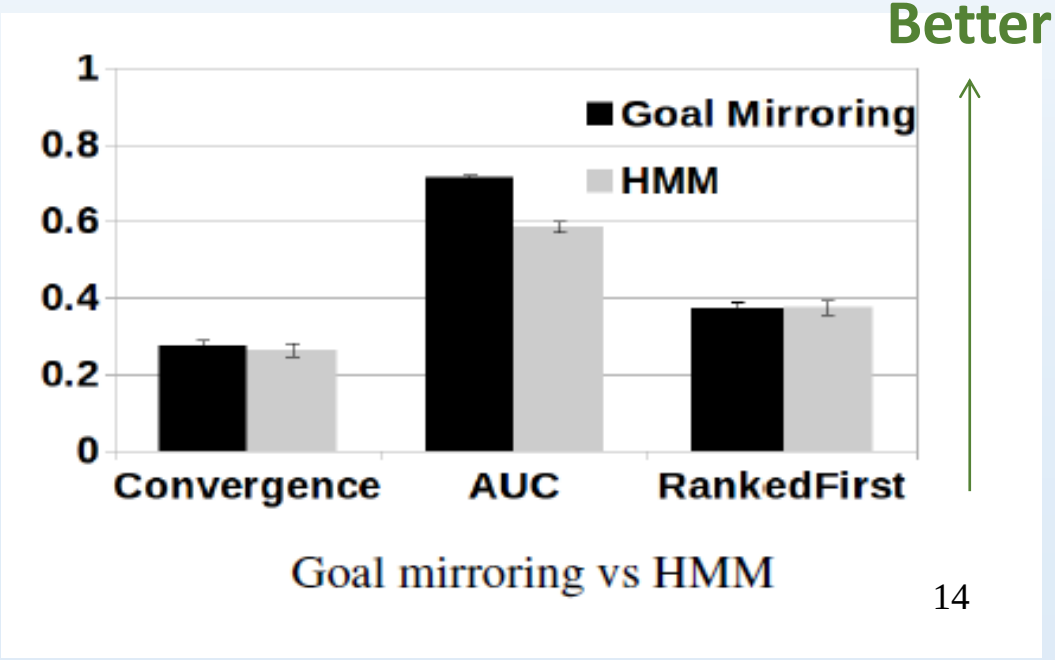

Goal mirroring vs HMM

Recognition Results

Comparison of Different Heuristics

- Different ranking heuristic : ratio vs. difference
- Will not work in continuous env.

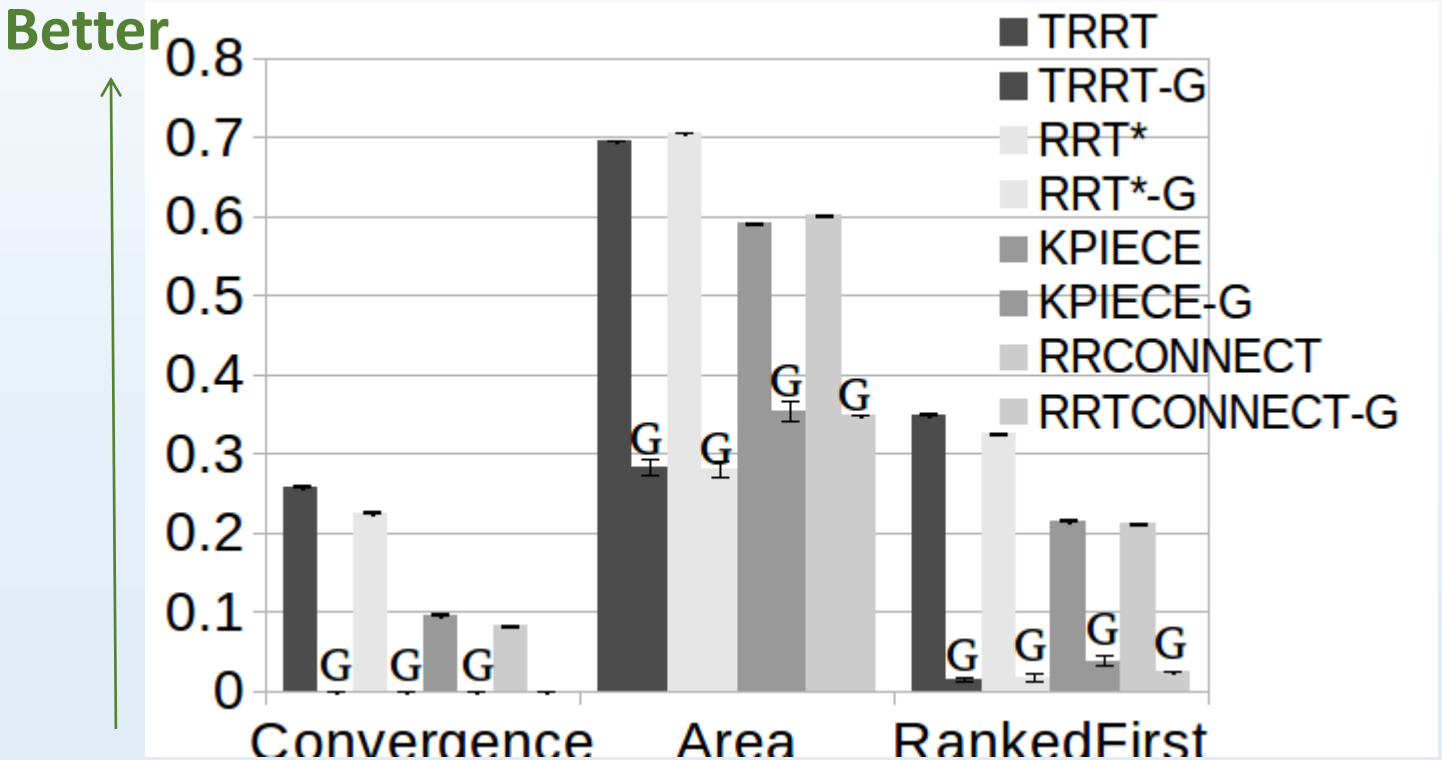

Conclusions : Online Goal Mirroring

- Continuous domains
- Uses planner to generate recognition hypotheses
- Shown that two factors impact recognition success
  - Optimality of planner used
  - Ranking heuristic
- Goal Mirroring preferred when less data is available and when possibilities are infinite.
- Further results in paper

Bar-Ilan University
אוניברסיטת בר-אילן

BIRC
BIU Robotics Consortium

Mor Vered **veredm@cs.biu.ac.il**