


# Accelerating Reinforcement Learning on a Robot by Using Subgoals in a Hierarchical Framework

Bart van Vliet

Wouter Caarls

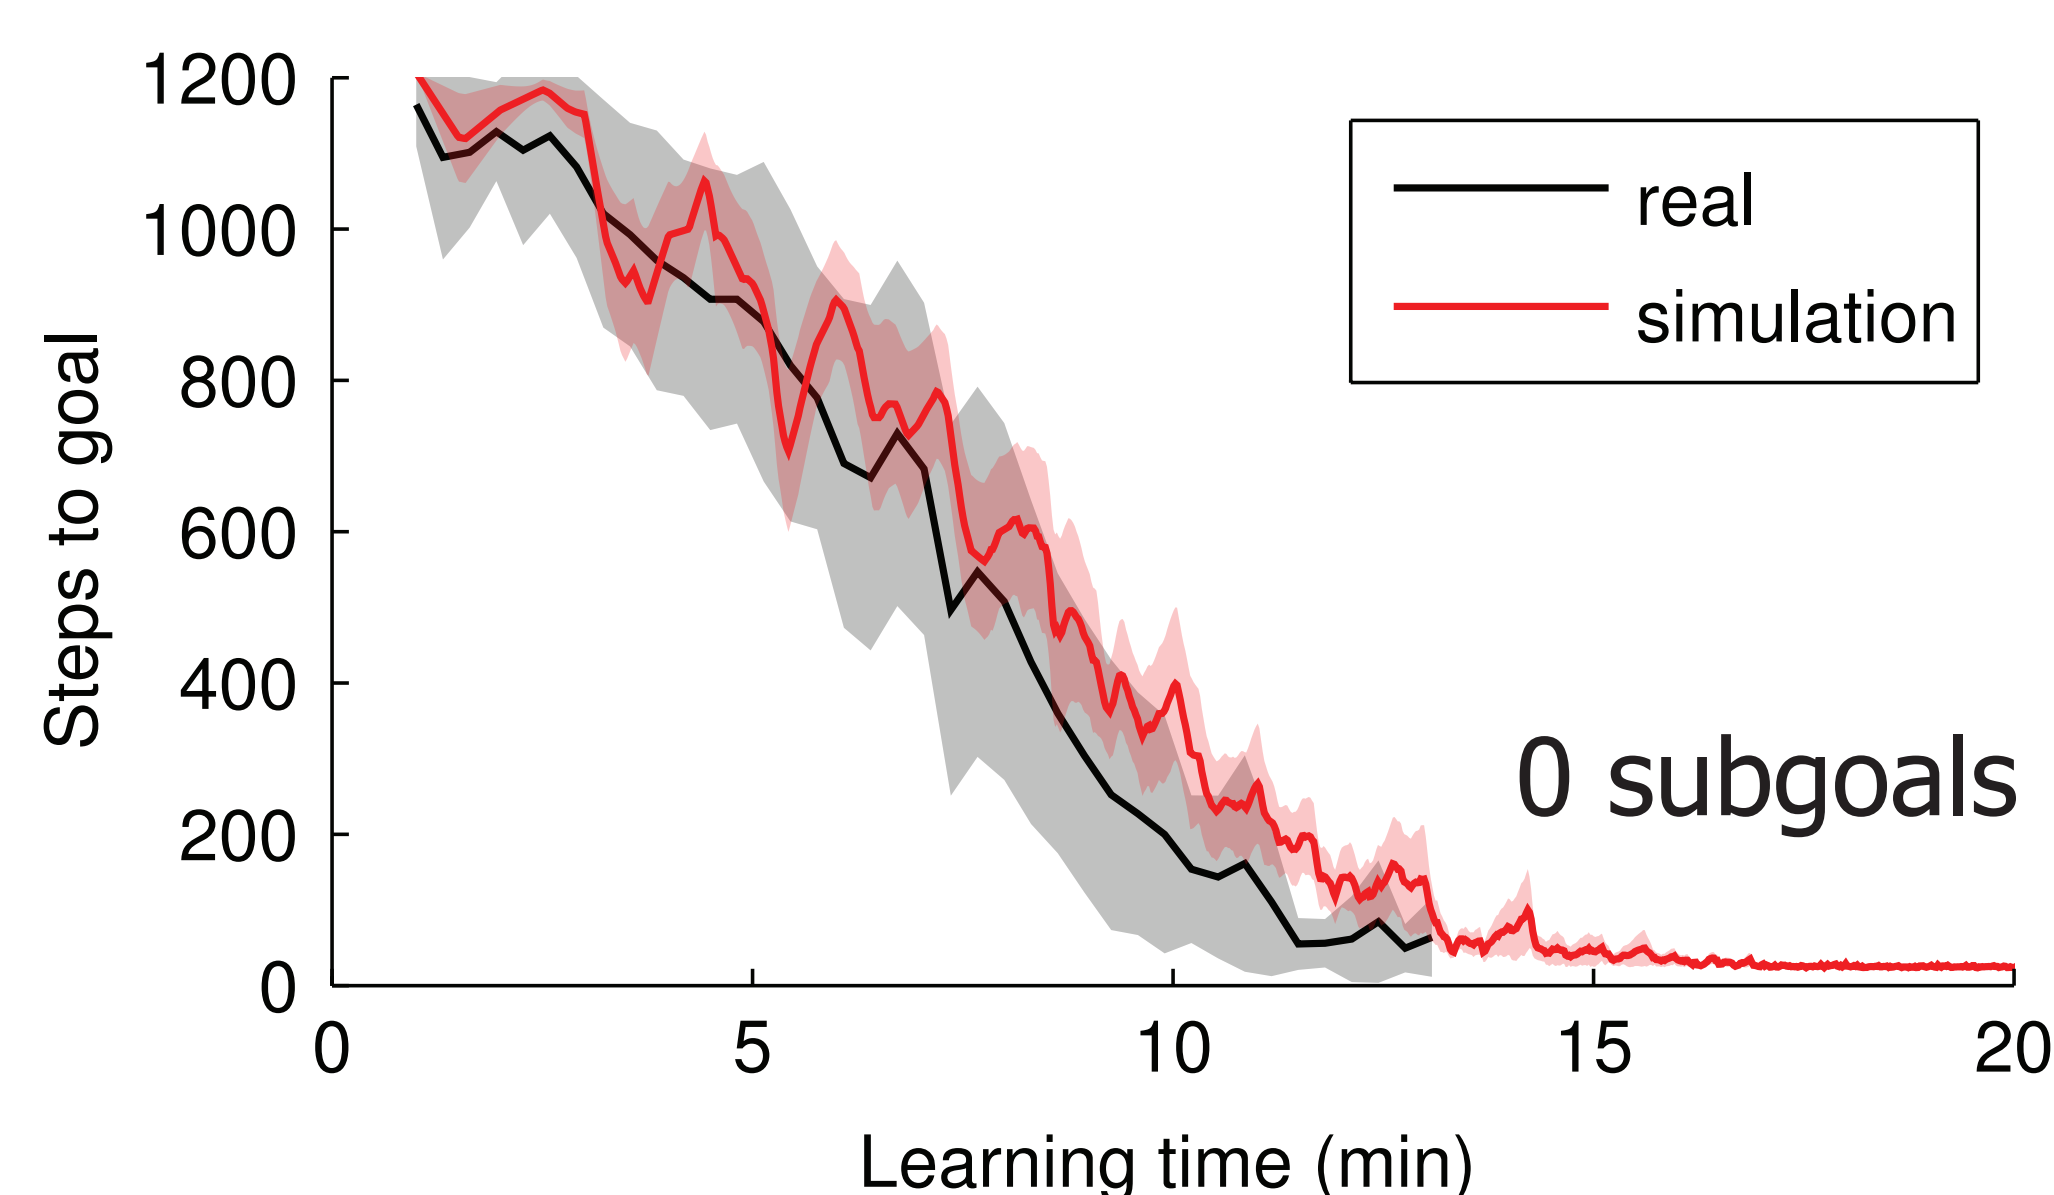
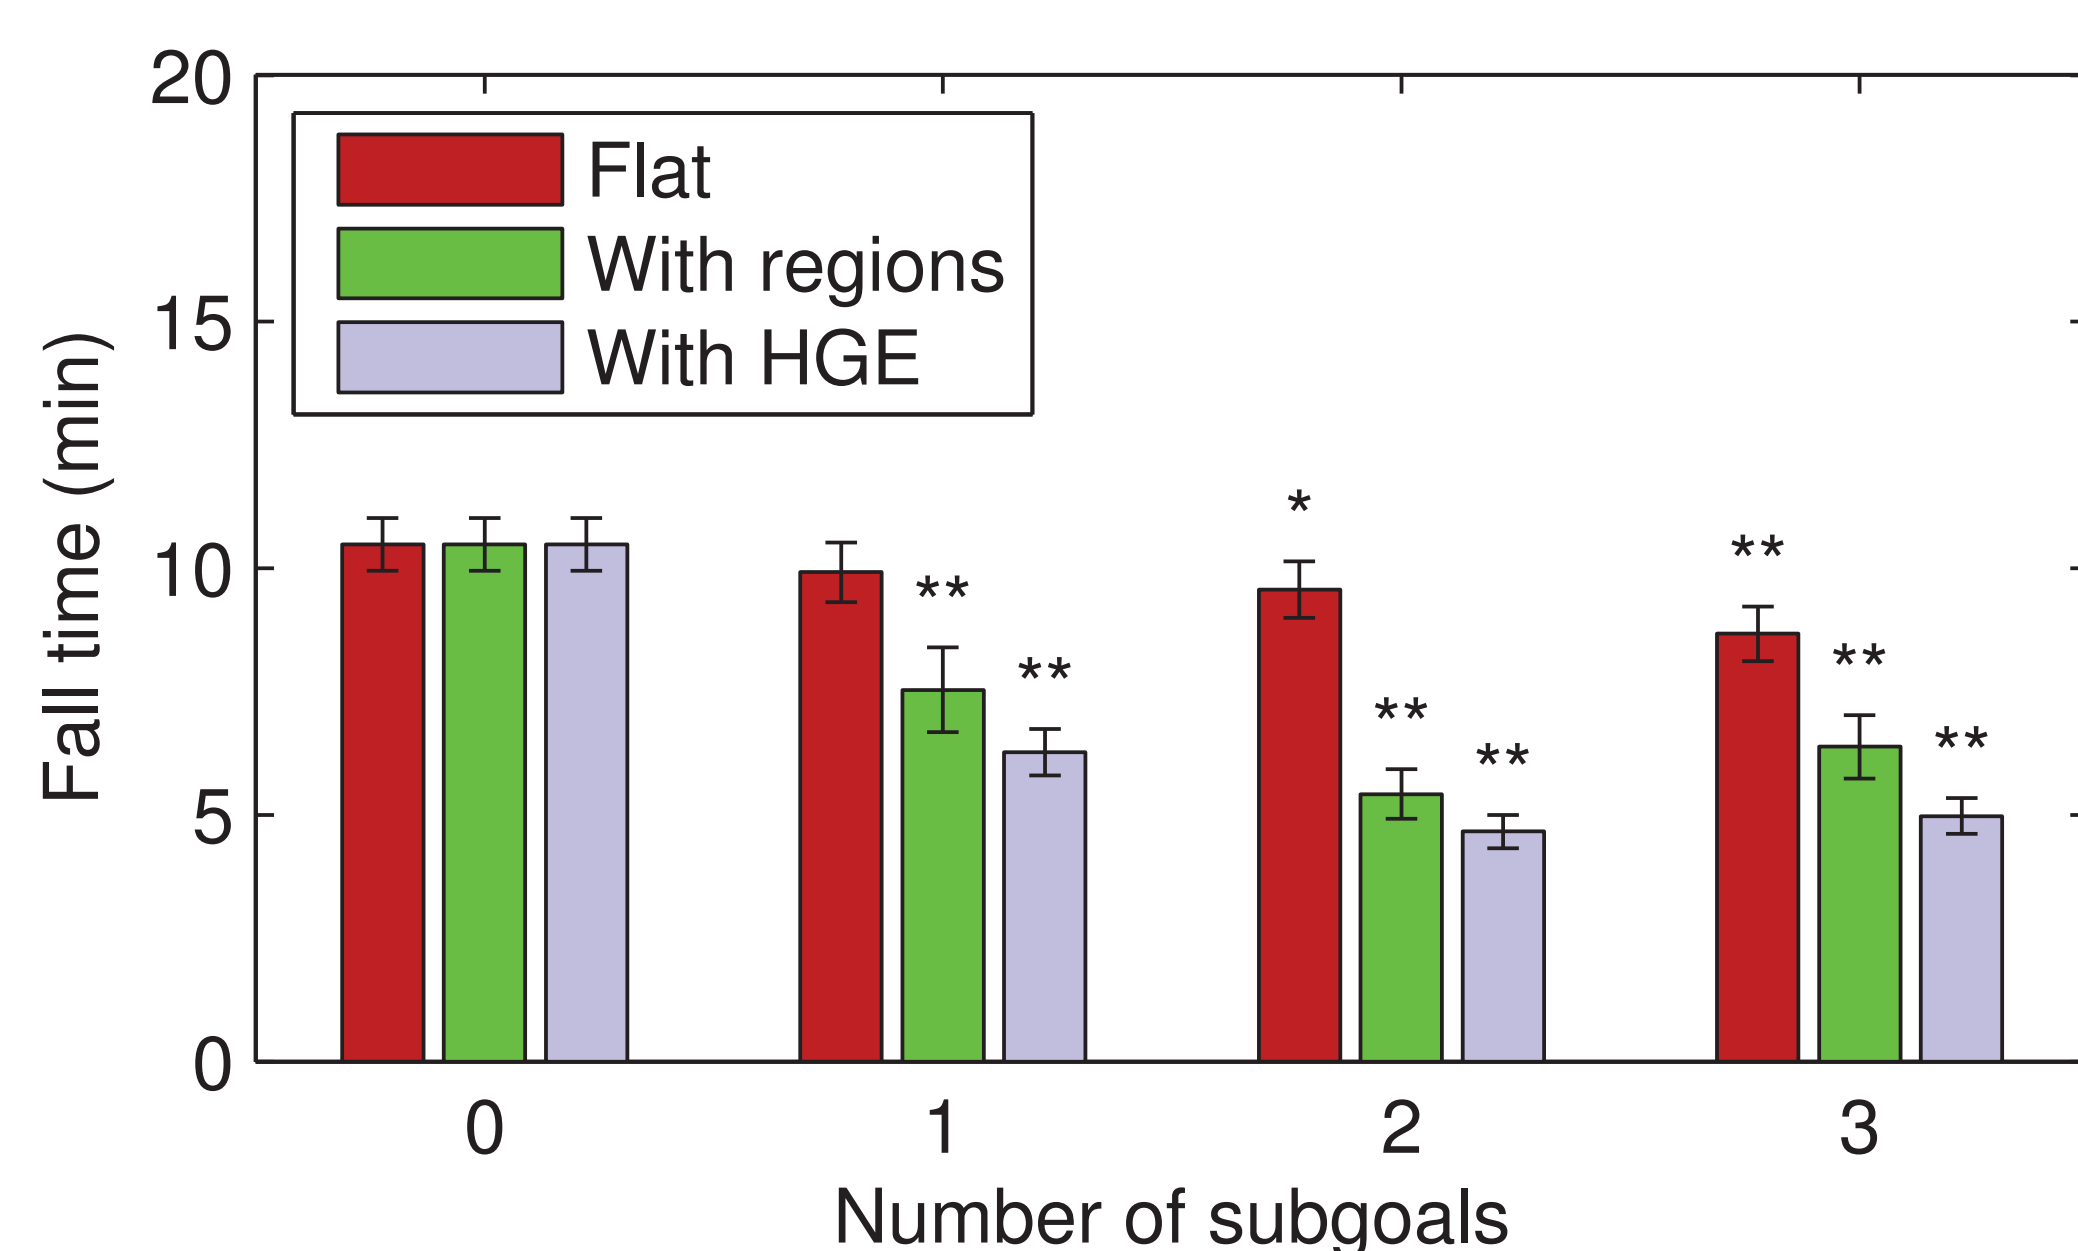
Erik Schuitema

Pieter Jonker



## LEO

- Autonomous 2d bipedal robot
- 7 degrees of freedom
- Robust



## Results

- More subgoals increase learning speed
- Reduced end performance
  - Partially cancelled out by HGE
- Simulation and real tests agree for 0 subgoals
- Real tests slower for 2 subgoals
  - Small goal area
  - Sensor noise and backlash

## Humanoid robots

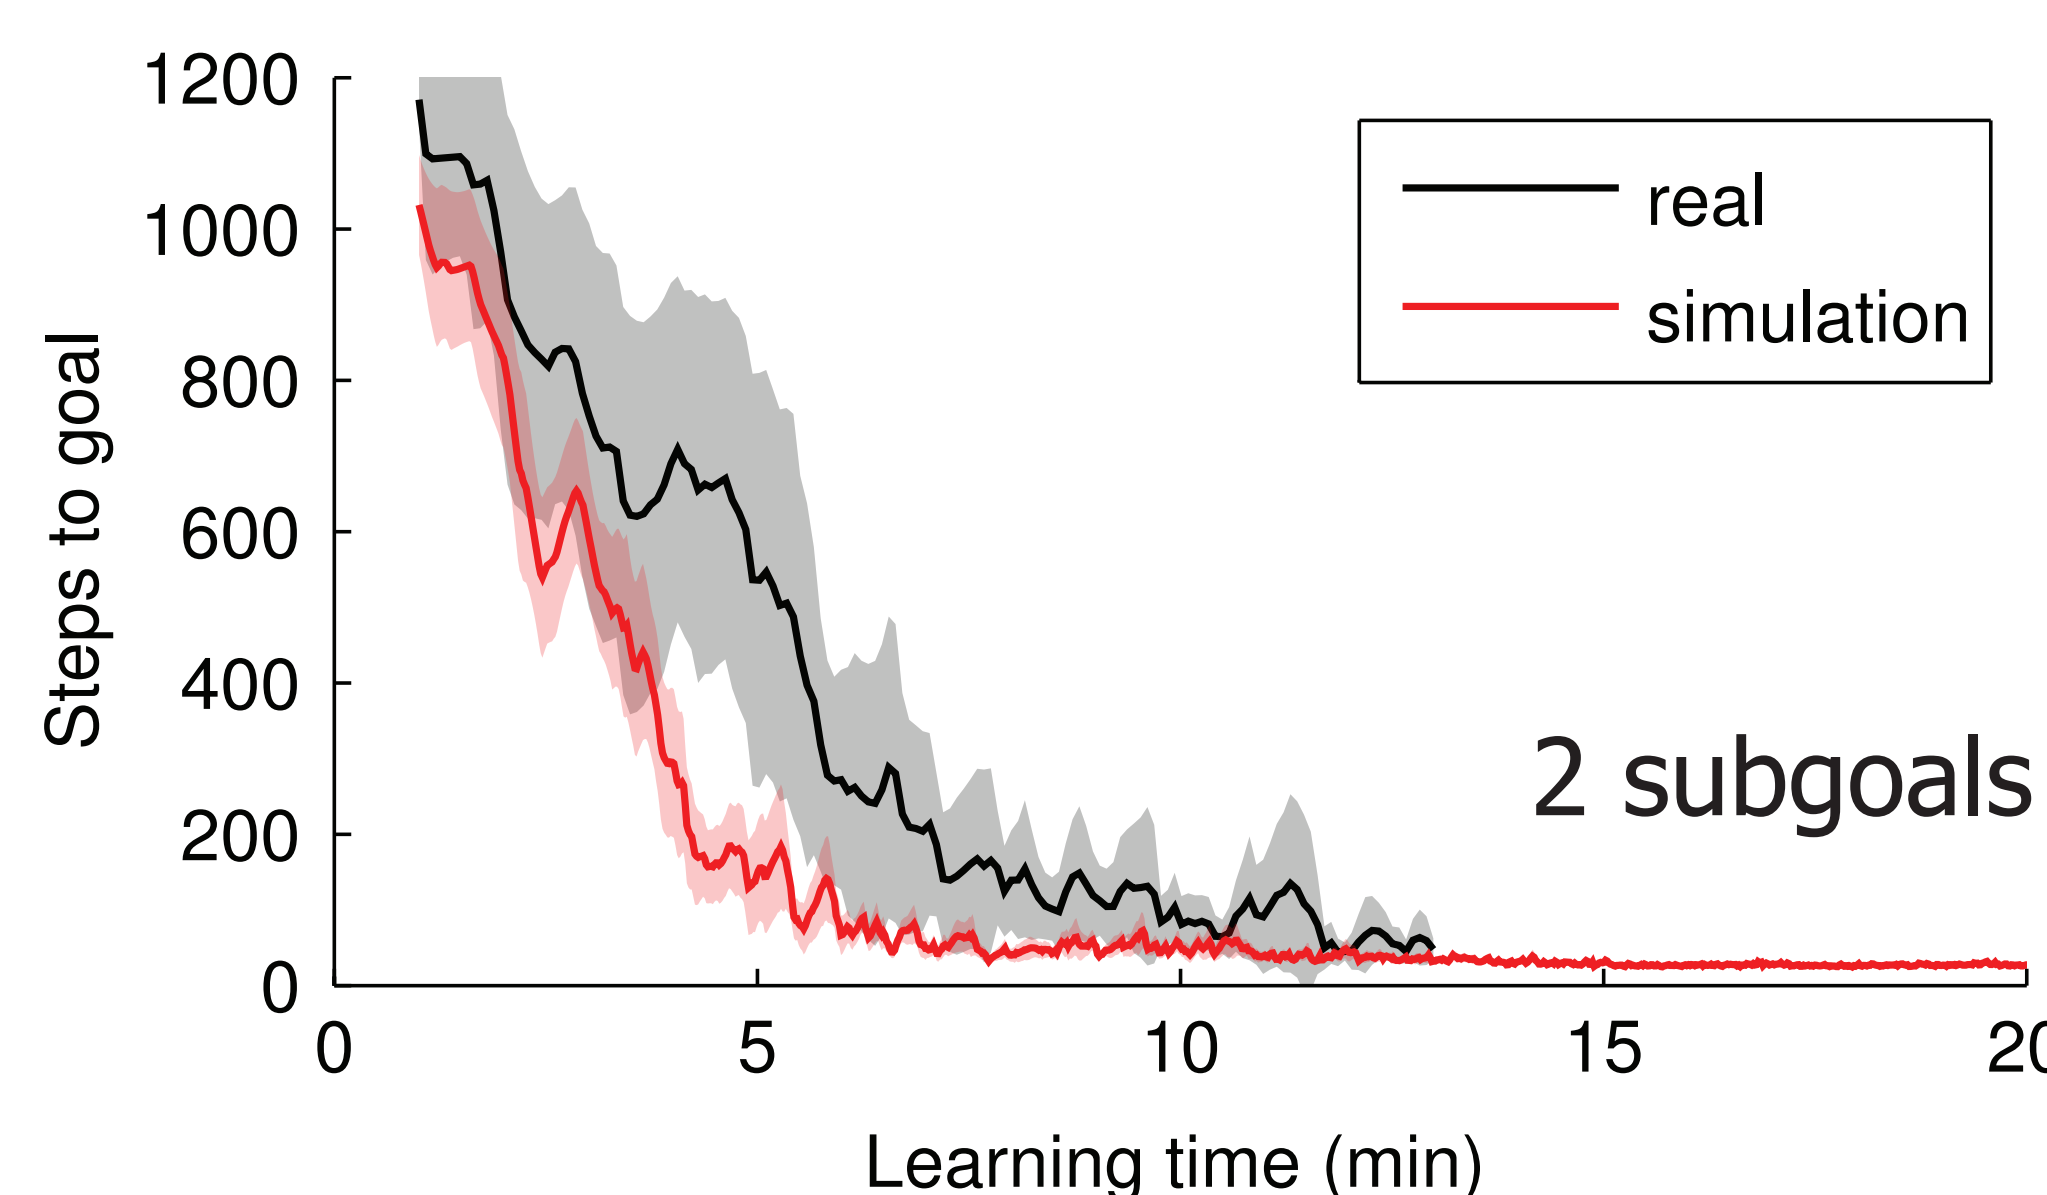
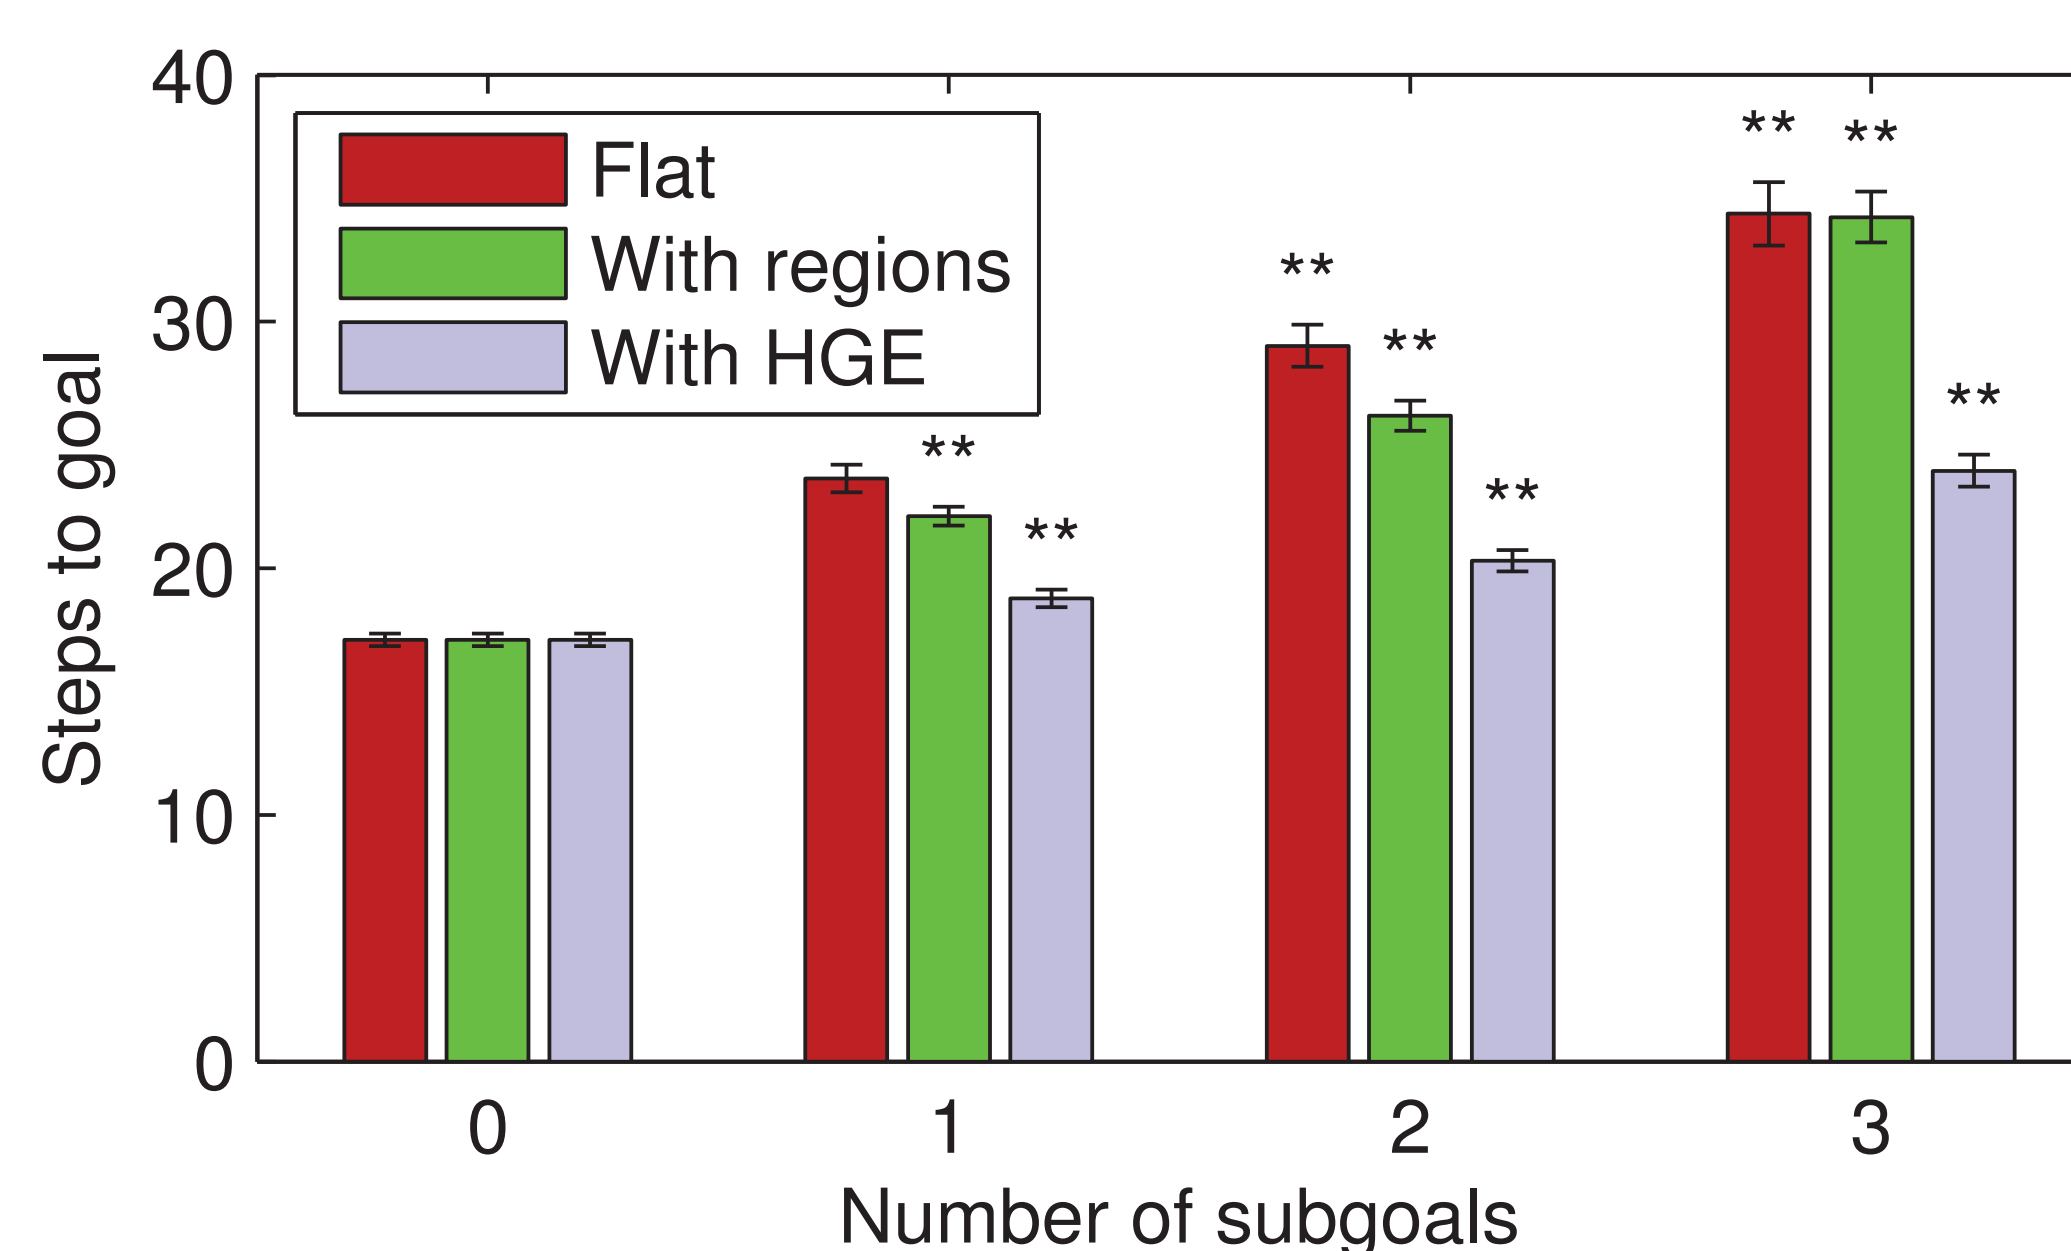
- Hard to model, hard to control
- Use reinforcement learning
- Naive learning is slow
- Add prior knowledge in the form of subgoals

## Hierarchical reinforcement learning

- MAXQ framework
  - Extended to MAXQ-Q( $\lambda$ ) with tile coding
- Subtasks, each reaching a subgoal
  - Execute only in a certain state space region
- Root task selects subtasks




## Hierarchical greedy execution

- Root task can interrupt subtasks
- Beneficial when subgoals do not lie on optimal path
- "Cutting corners"



## Towards optimal end performance

- All goals updating
  - Subtasks can learn while not in control
- Last subtask eventually learns entire problem
- Optimal end performance possible



## Experiment

- Stair step-up
- Goal reward, time penalty
- Subgoals near path to goal

