Algorithmic Robotics and Motion Planning

Introduction

Dan Halperin
School of Computer Science
Tel Aviv University

Fall 2019-2020

[Dolce & Gabbana 2018 handbag collection]()

Today’s lesson

- basic terminology
- fundamental problems
- robotics vs. automation
- review of the major course topics
- course mechanics

As time permits:

- the Roomba in the café, combinatorics and algorithms

Robots, take I

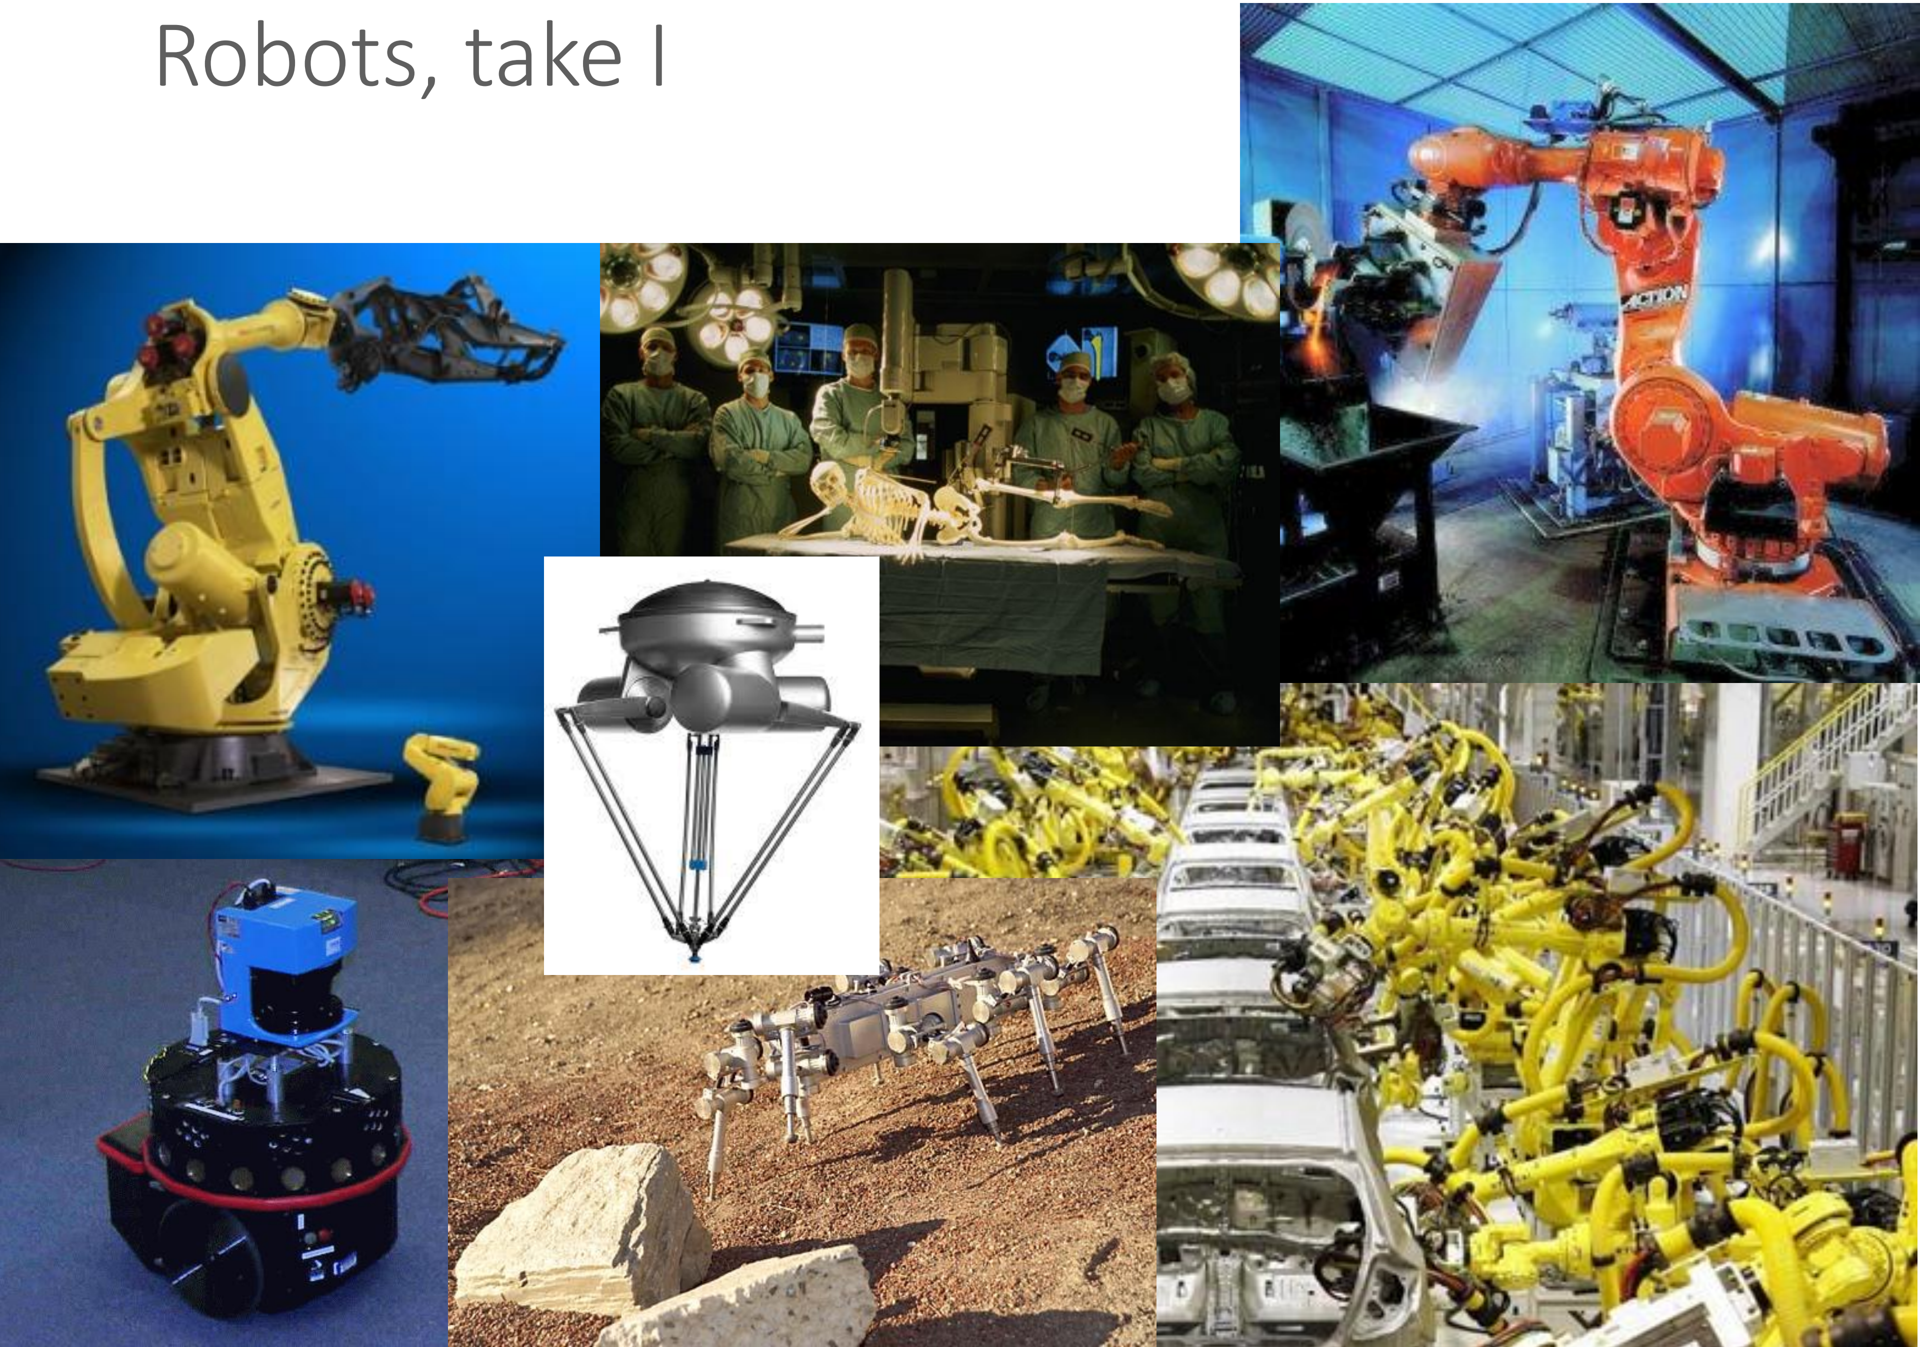

An extremely brief history of robotics

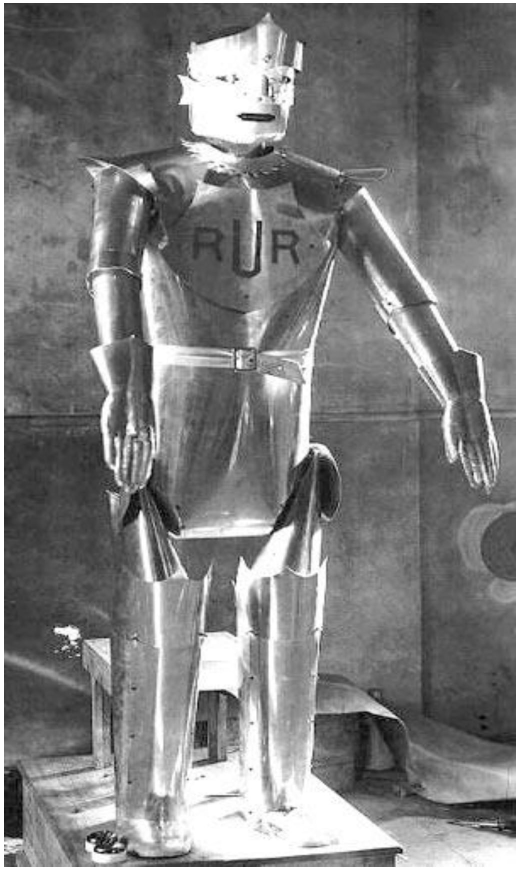

The RUR robot which appeared in an adaption of Czech author Karel Capek's *Rossum's Universal Robots.*
Circa 1930's.

NASA's Curiosity, 2011 →

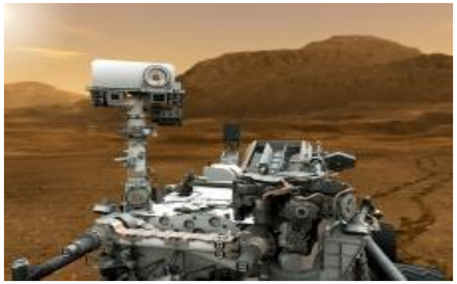

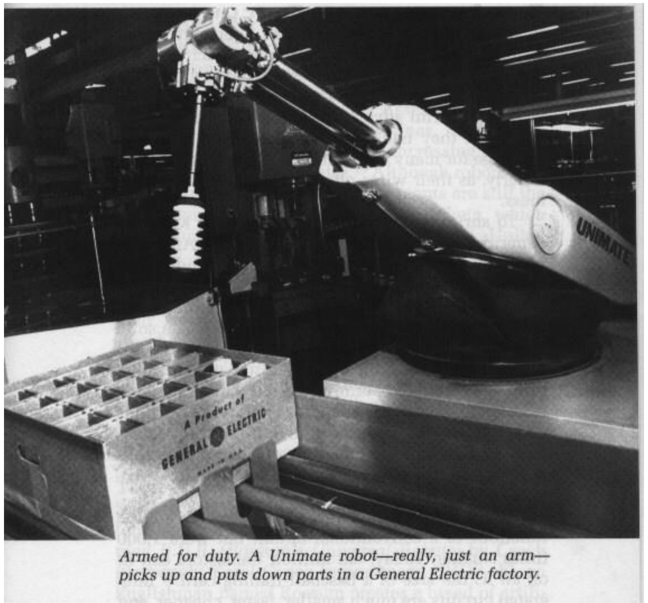

Armed for duty. A Unimate robot—really, just an arm—picks up and puts down parts in a General Electric factory.

UNIMATE becomes the first industrial robot in use. It was used at the General Motors factory in New Jersey. 1961.

Honda's ASIMO, 2002

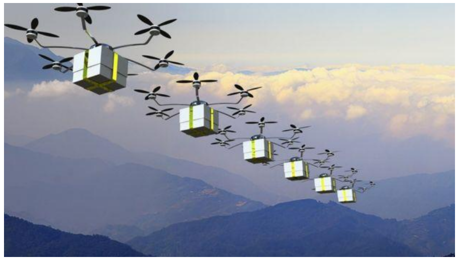

Robotics and robots

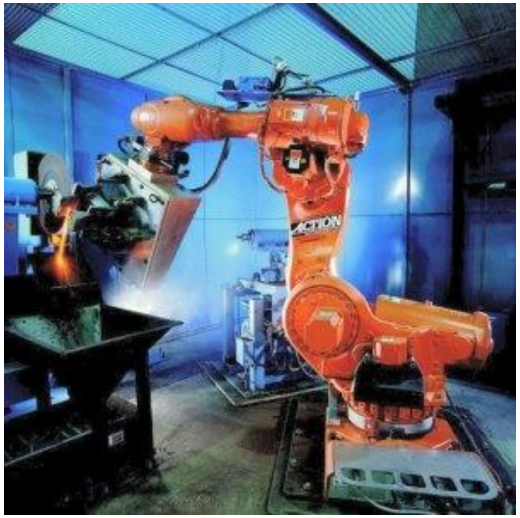

[https://robots.ieee.org/learn/]

What is a robot?

A robot is an autonomous machine capable of sensing its environment, carrying out computations to make decisions, and performing actions in the real world.

!?

Robotics and robots

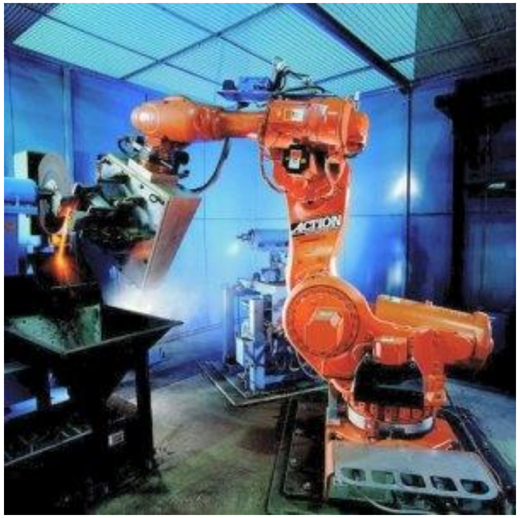

!?

Here it will be interesting if

- it is autonomous (at least in part), and
- it has non-trivial motion and/or manipulation capabilities

Motion planning:
the basic problem

Let B be a system (the robot) with k degrees of freedom moving in a known environment cluttered with obstacles. Given free start and goal placements for B decide whether there is a collision free motion for B from start to goal and if so plan such a motion.

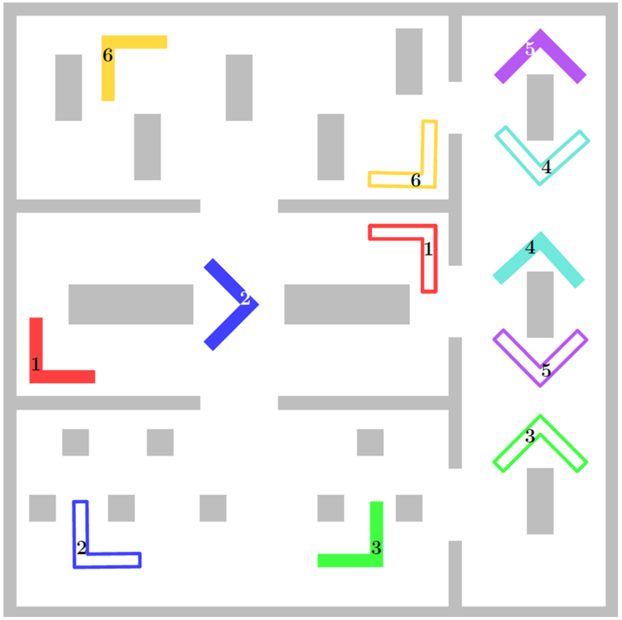

Example I: The Roomba in the café

A disc moving among discs

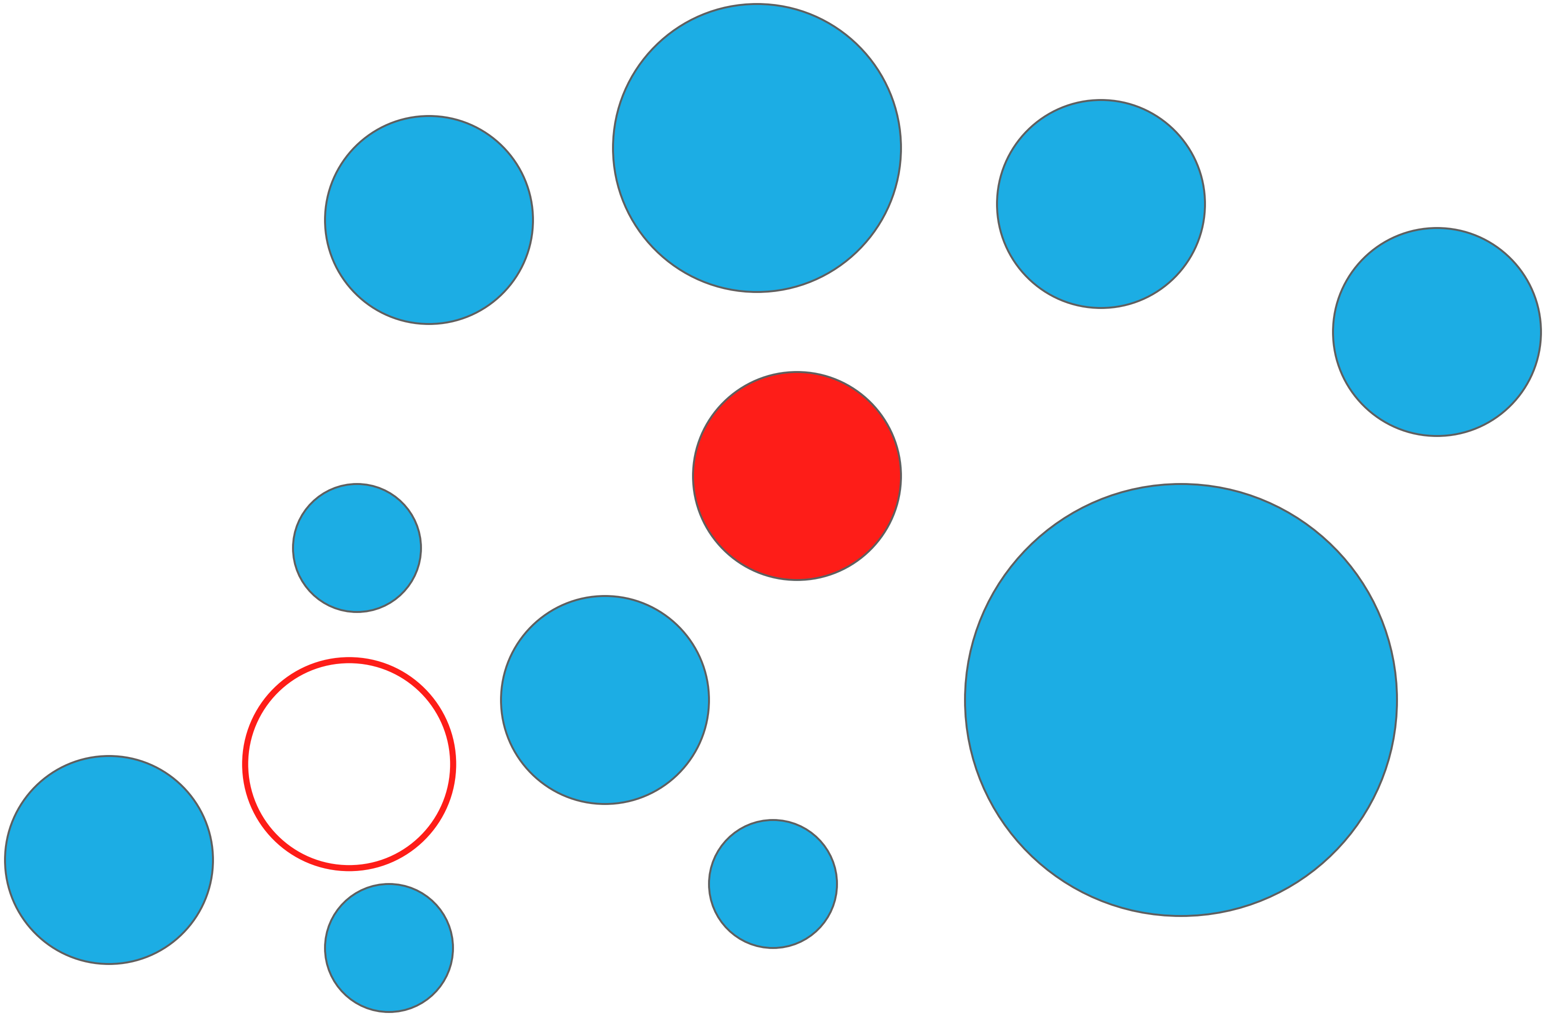

Example II: Oskar's cube

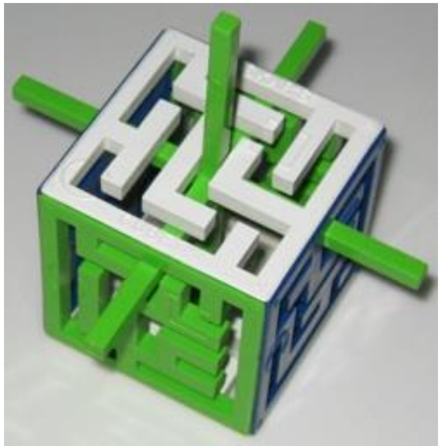

- MP with 3 translational dofs
- Hint: Scientific American, Sep 1988 issue
- Jay's Oskar's cubes

[oskarvandeventer.nl]

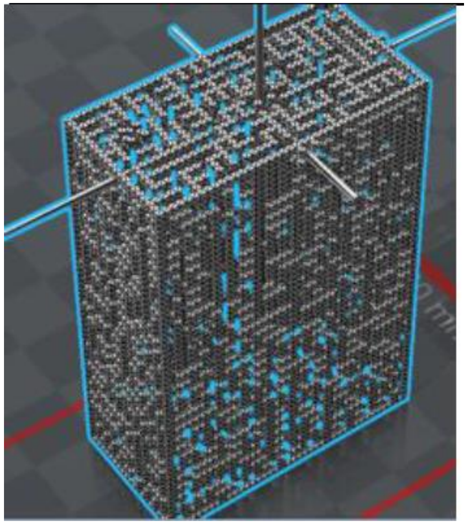

Terminology

- Workspace
- Configuration space (state space)
- Degrees of freedom (dofs)

Degrees of freedom

- a polygon robot translating in the plane
- a polygon robot translating and rotating
- a spatial robot translating and rotating
- industrial robot arms
- many robots

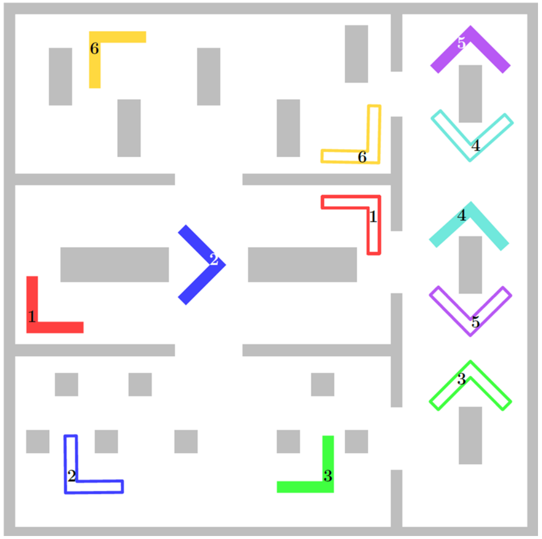

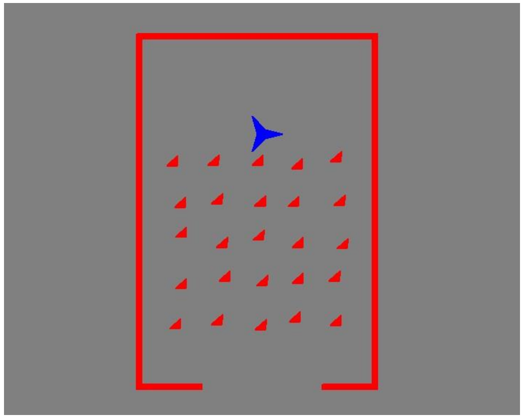

Configuration space

of a robot system with k degrees of freedom

- C-space, for short
- also known as state space
- the space of parametric representation of all possible robot configurations
- C-obstacles: the expanded obstacles
- the robot -> a point
- k-dimensional space
- point in configuration space: free, forbidden (, semi-free)
- path -> curve

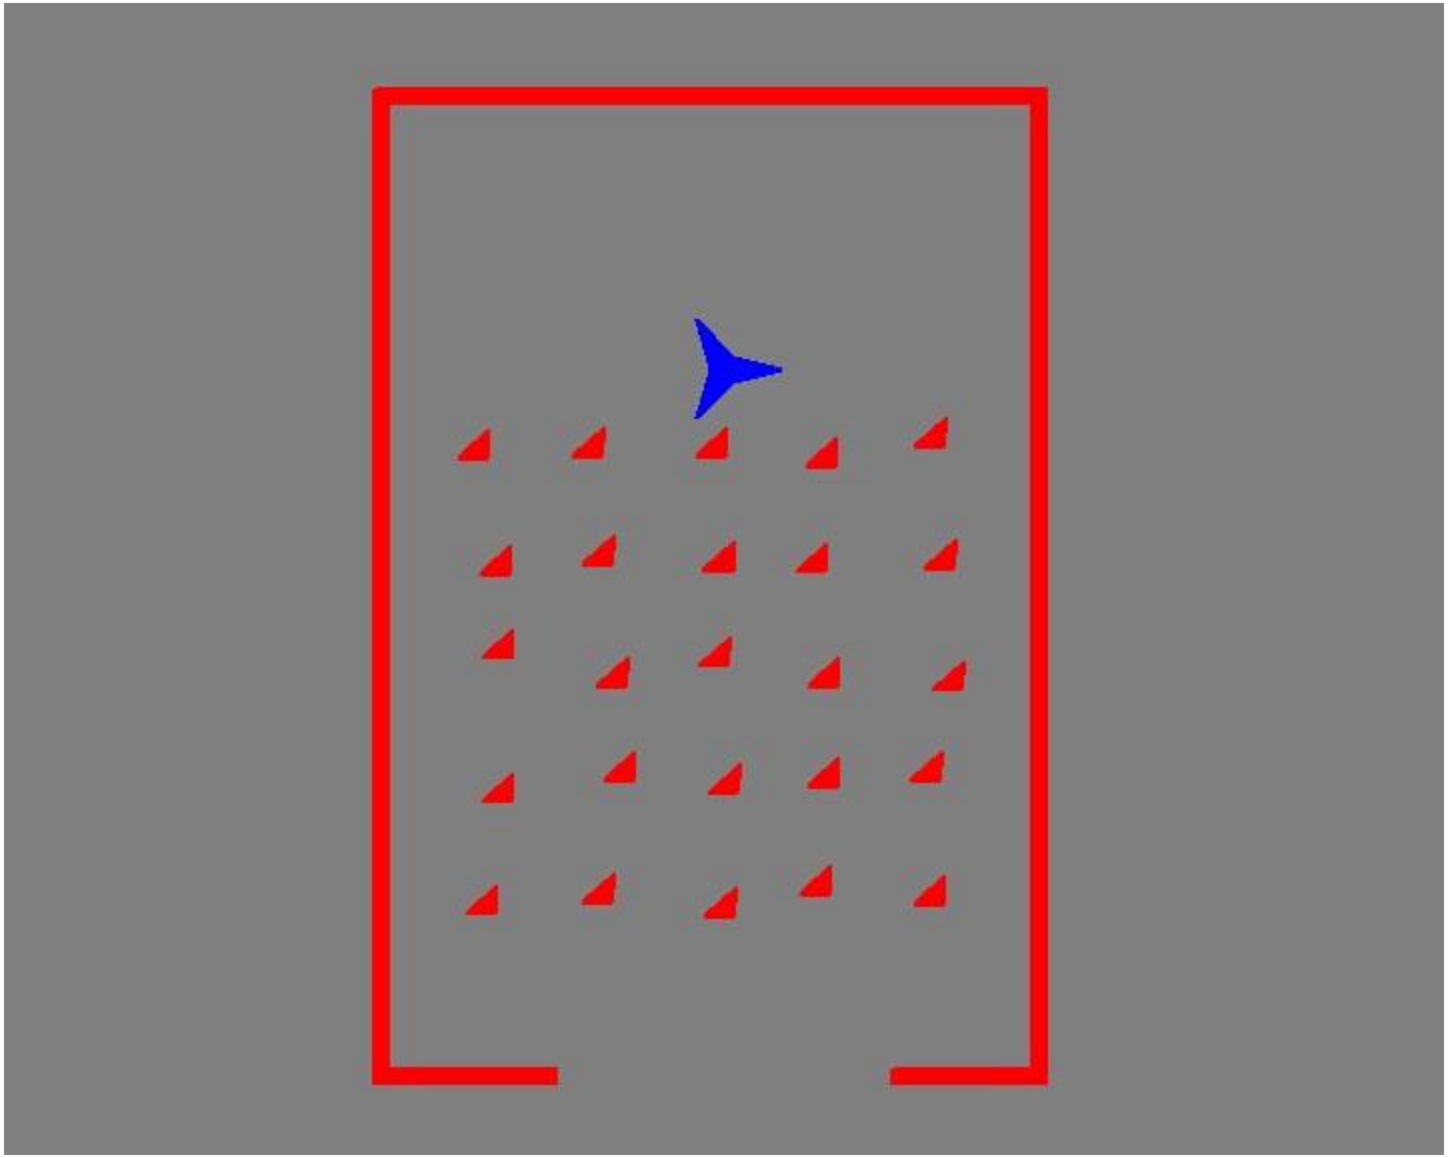

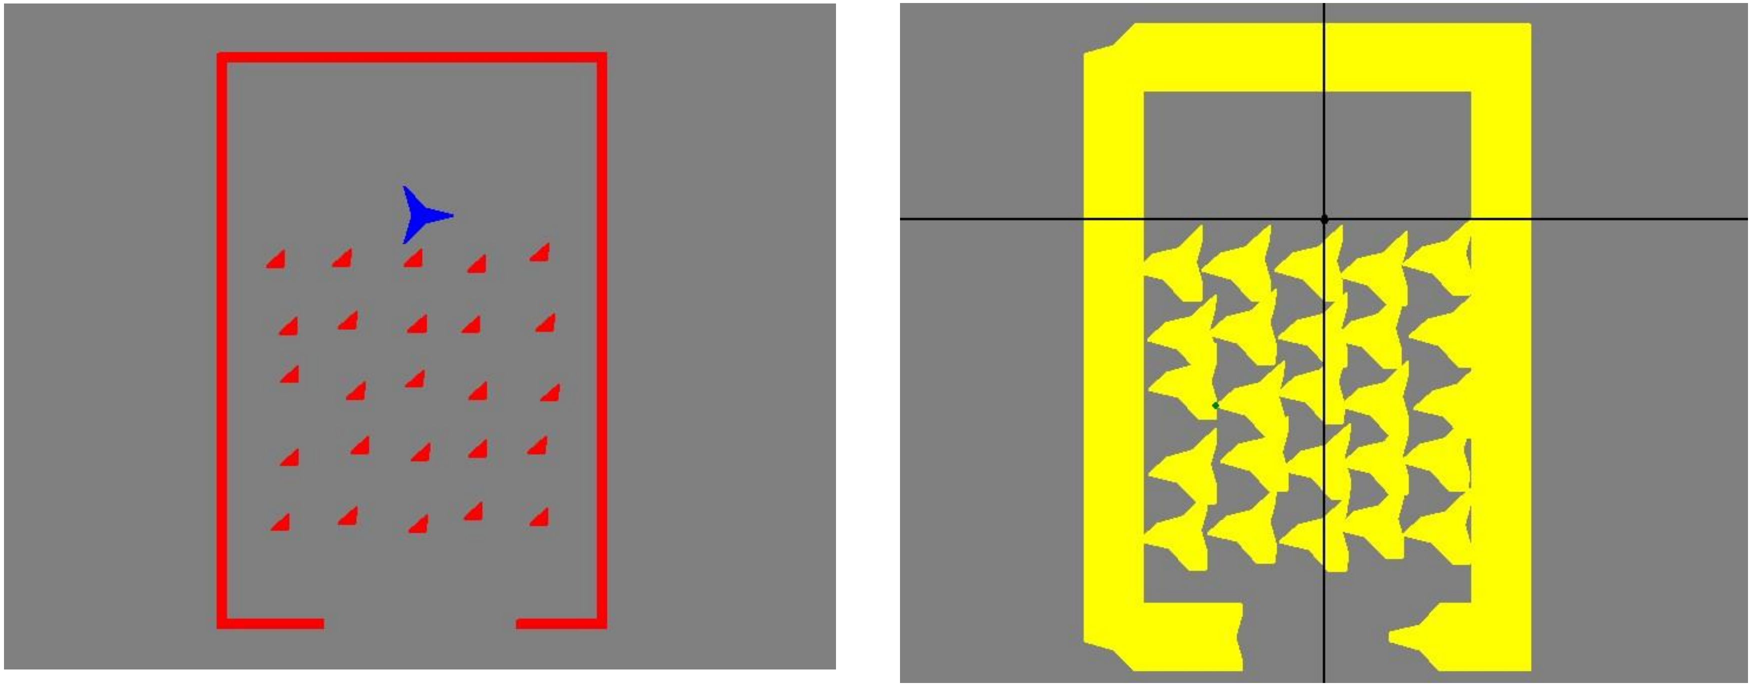

MOVE

C-obstacles

Q - a polygonal object that moves by translation

P - a set of polygonal obstacles

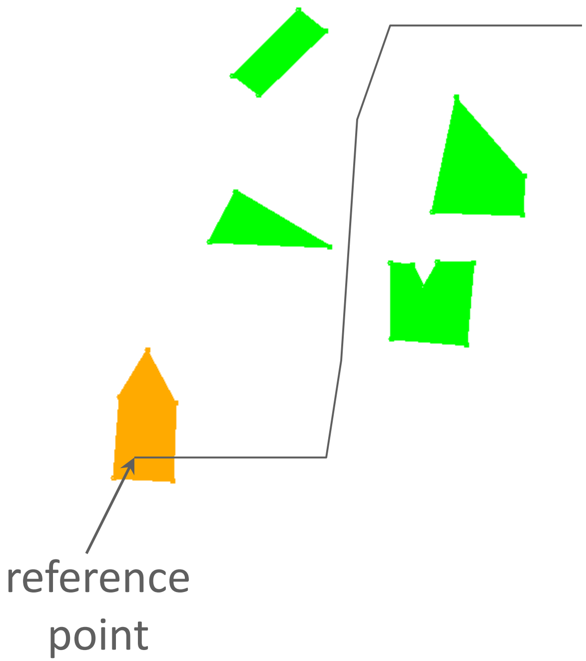

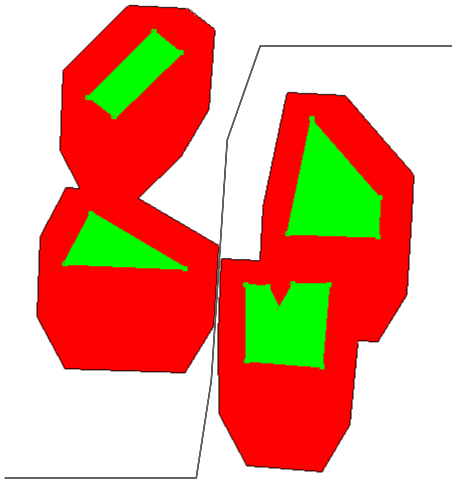

Minkowski sums

and translational C-obstacles

a central tool in geometric computing applicable to motion planning and other domains

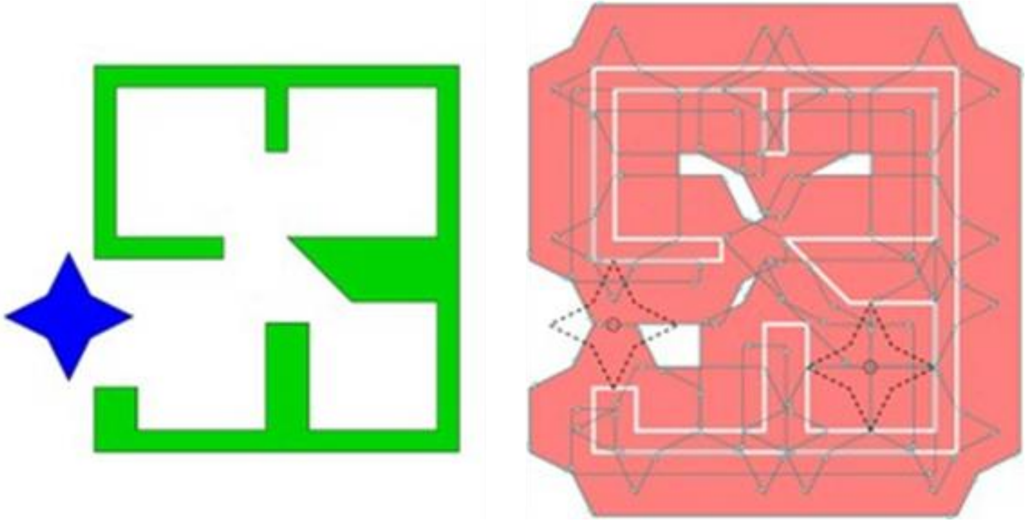

More complex systems

new designs, multi-robot systems, and other moving artifacts have many more dofs

Example III: the α puzzle

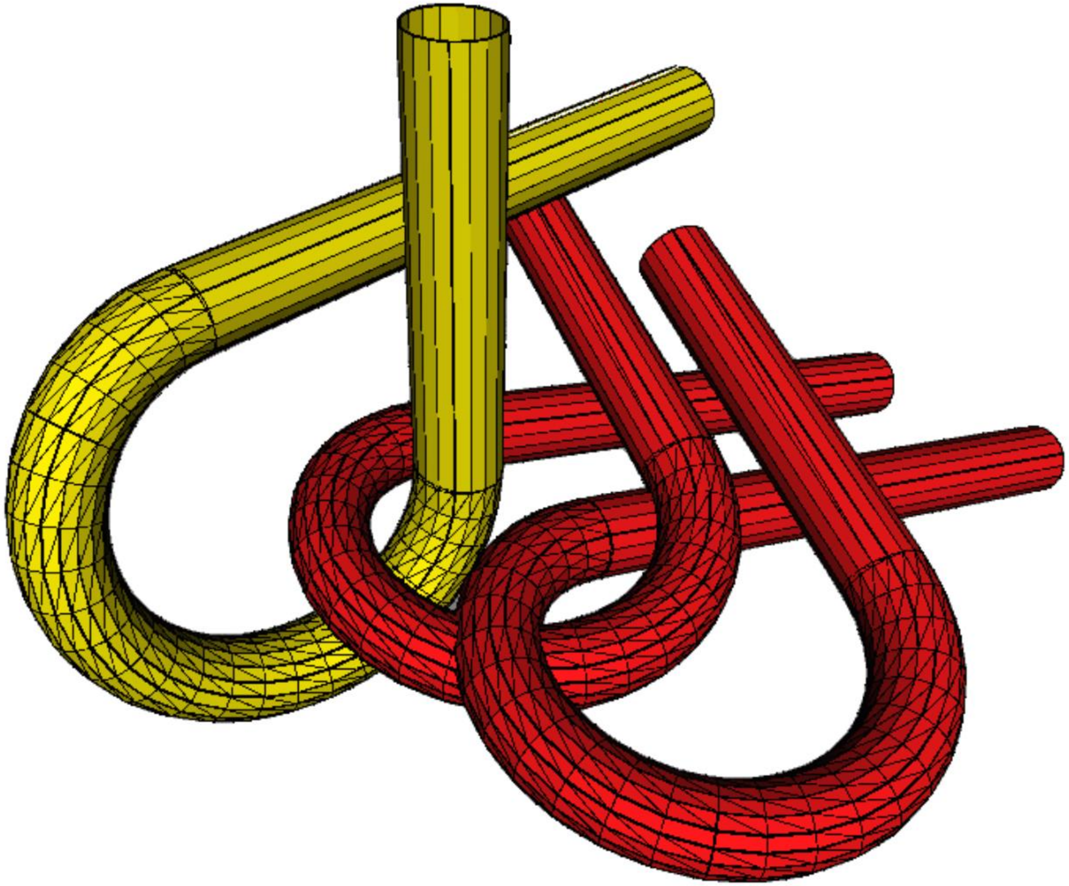

Types of solutions

- exact
- probabilistic
- hybrid
- heuristic

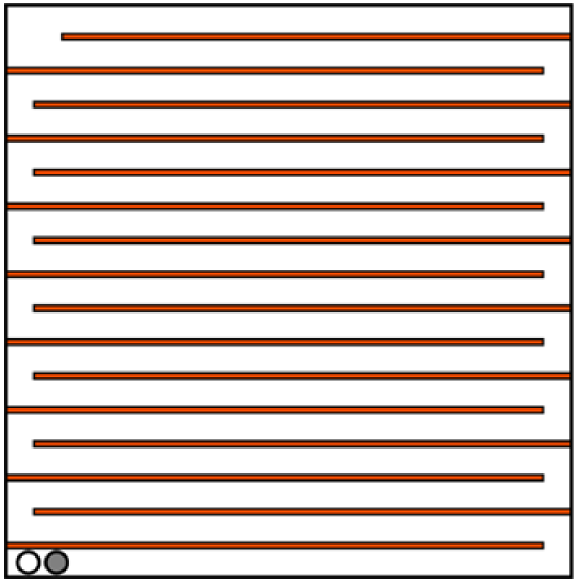

- major components in practical solutions: nearest-neighbor search, collision detection

Robots, take II

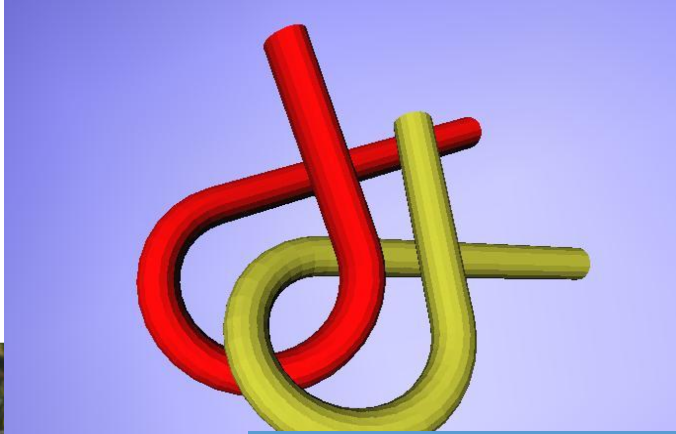

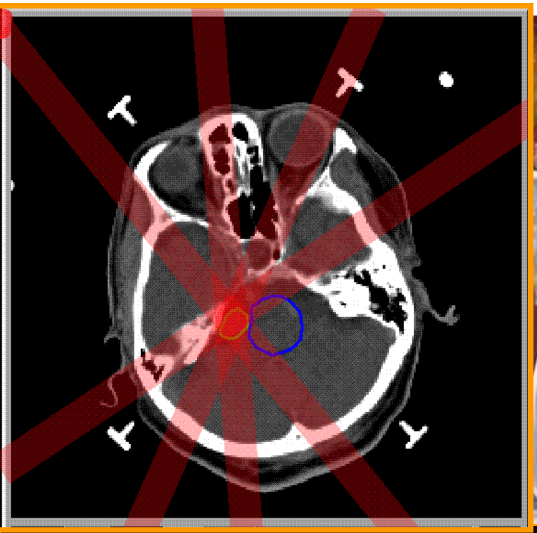

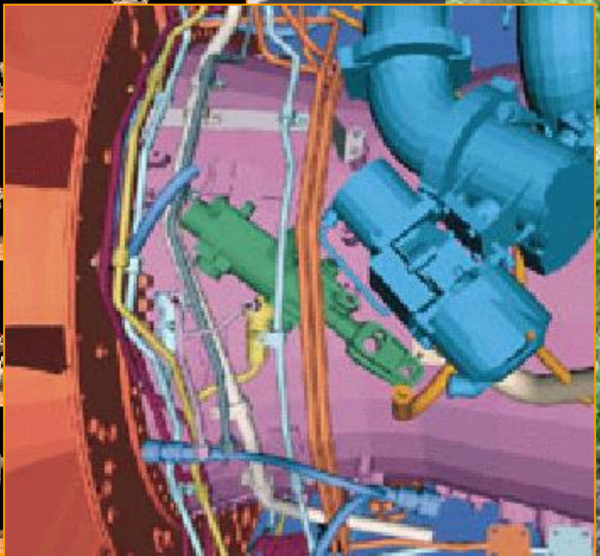

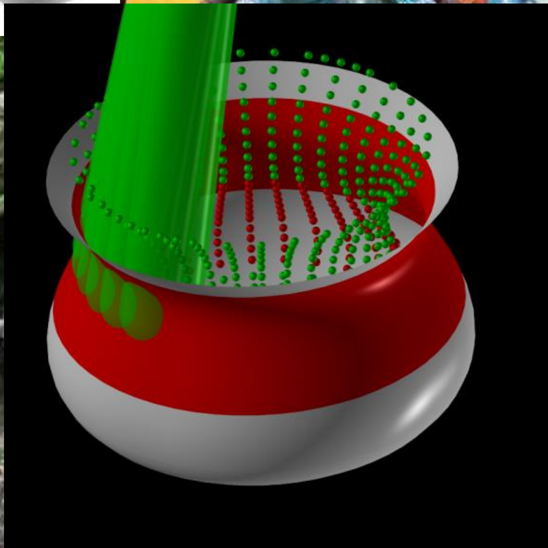

Beyond the basic MP problem

- moving obstacles
- multiple robots
- movable objects
- uncertainty
- nonholonomic constraints
- dynamic constraints
- . . .

Multiple robots

[home.ustc.edu.cn/~hxiangli]

[cbsnew]

[flow free]

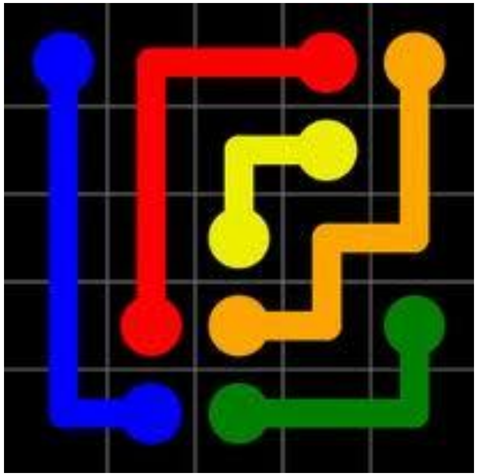

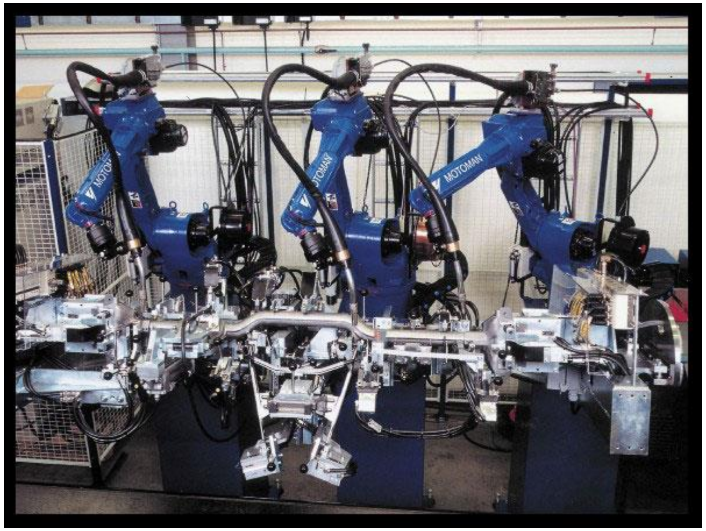

[IccRobotics.com]

[autonomy.cs.sfu.ca]

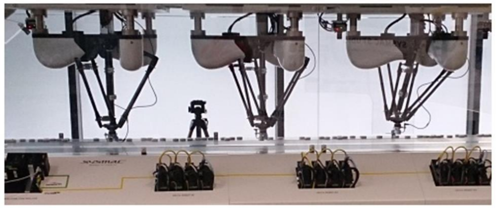

Problem

Given two Roomba’s, each has to move from given start to goal positions, no obstacles. What are the joint shortest paths (minimum total length)?

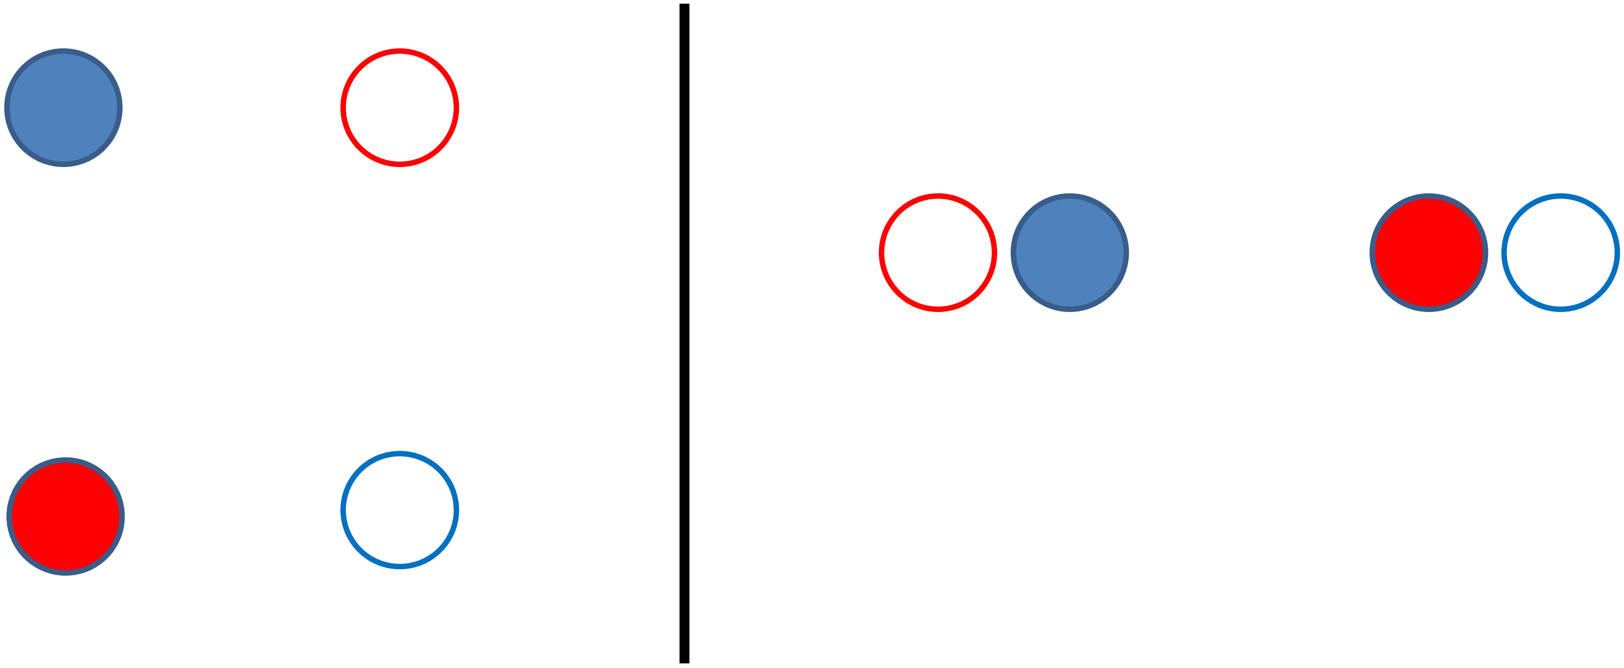

Path quality

- length
- clearance
- combined measures
- minimum energy
- Minimum time
- …
- hard even in simple settings

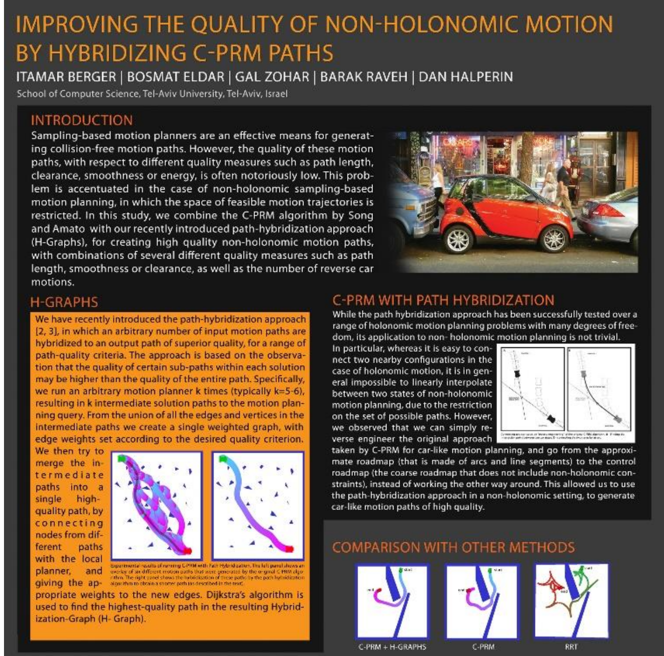

IMPROVING THE QUALITY OF NON-HOLONOMIC MOTION BY HYBRIDIZING C-PRM PATHS

ITAMAR BERGER | BOSMAT ELDAR | GAL ZOHAR | BARAK RAVEH | DAN HALPERIN

School of Computer Science, Tel-Aviv University, Tel-Aviv, Israel

INTRODUCTION

Sampling-based motion planners are an effective means for generating collision-free motion paths. However, the quality of these motion paths, with respect to different quality measures such as path length, clearance, smoothness or energy, is often notoriously low. This problem is accentuated in the case of non-holonomic sampling-based motion planning, in which the space of feasible motion trajectories is restricted. In this study, we combine the C-PRM algorithm by Song and Amato with our recently introduced path-hybridization approach (H-Graphs), for creating high quality non-holonomic motion paths, with combinations of several different quality measures such as path length, smoothness or clearance, as well as the number of reverse car motions.

H-GRAPHS

We have recently introduced the path-hybridization approach [2, 3], in which an arbitrary number of input motion paths are hybridized to an output path of superior quality, for a range of path-quality criteria. The approach is based on the observation that the quality of certain sub-paths within each solution may be higher than the quality of the entire path. Specifically, we run an arbitrary motion planner k times (typically k=5-6), resulting in k intermediate solution paths to the motion planning query. From the union of all the edges and vertices in the intermediate paths we create a single weighted graph, with edge weights set according to the desired quality criterion. We then try to merge the intermediate paths into a single high-quality path, by connecting nodes from different paths with the local planner, and giving the appropriate weights to the new edges. Dijkstra's algorithm is used to find the highest-quality path in the resulting Hybridization-Graph (H- Graph).

C-PRM WITH PATH HYBRIDIZATION

While the path hybridization approach has been successfully tested over a range of holonomic motion planning problems with many degrees of freedom, its application to non- holonomic motion planning is not trivial. In particular, whereas it is easy to connect two nearby configurations in the case of holonomic motion, it is in general impossible to linearly interpolate between two states of non-holonomic motion planning, due to the restriction on the set of possible paths. However, we observed that we can simply reverse engineer the original approach taken by C-PRM for car-like motion planning, and go from the approximate roadmap (that is made of arcs and line segments) to the control roadmap (the coarse roadmap that does not include non-holonomic constraints), instead of working the other way around. This allowed us to use the path-hybridization approach in a non-holonomic setting, to generate car-like motion paths of high quality.

COMPARISON WITH OTHER METHODS

C-PRM + H-GRAPHS | C-PRM | RRT

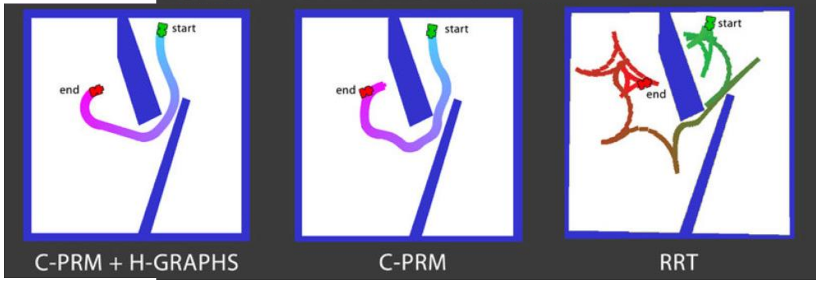

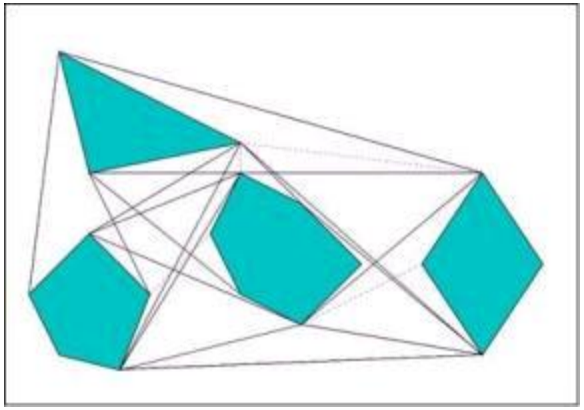

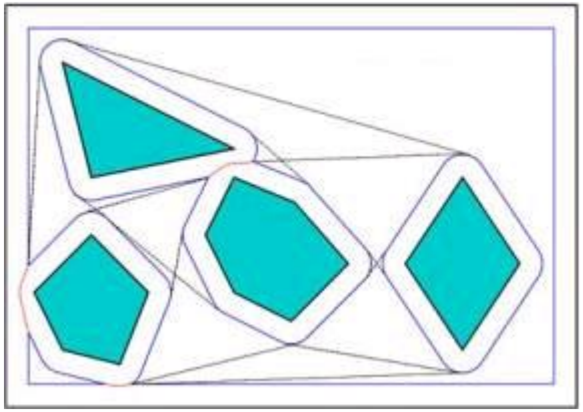

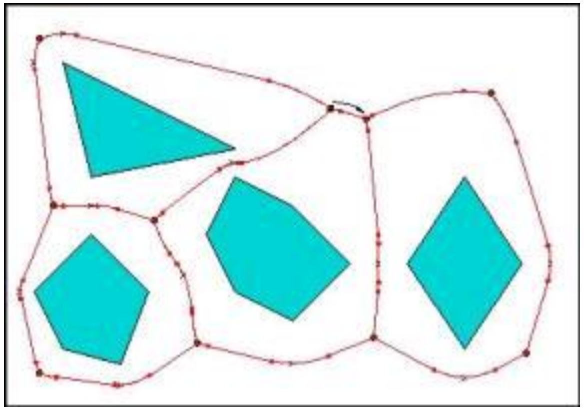

Problem

A (point) robot is moving in the plane in the vicinity of a (point) source of nuclear radiation. The cots per unit distance is inversely proportional to the clearance from the source of radiation

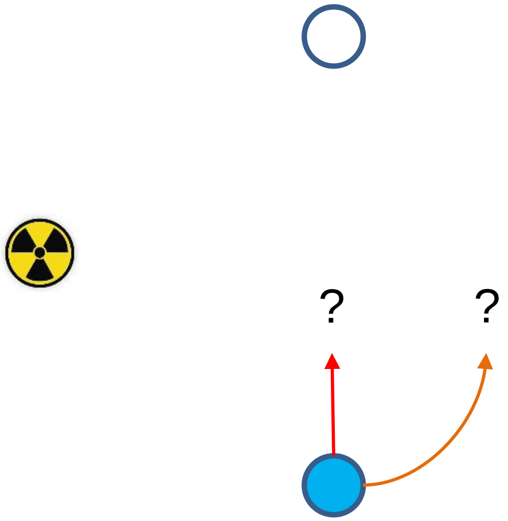

typical in robotics:
multi-objective optimization

Kinematics

- link
- joint
- base
- tcp
- kinematic chain
- direct kinematics
- inverse kinematics

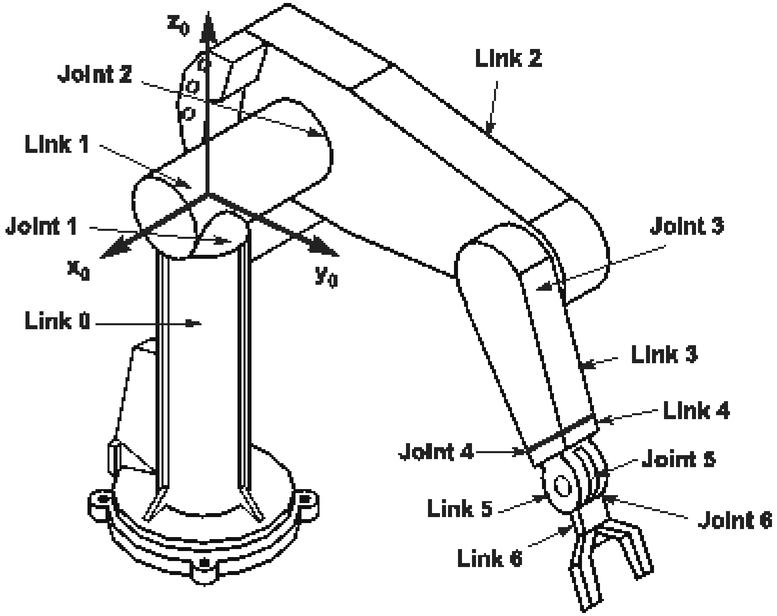

Inverse kinematics

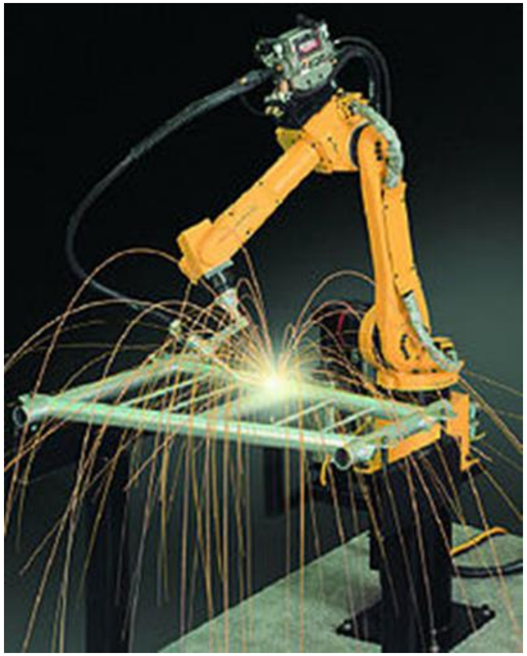

[Fanuc Iberia]

Denavit-Hartenberg 1955, Pieper-Roth 1969

Inverse kinematics, a simple example

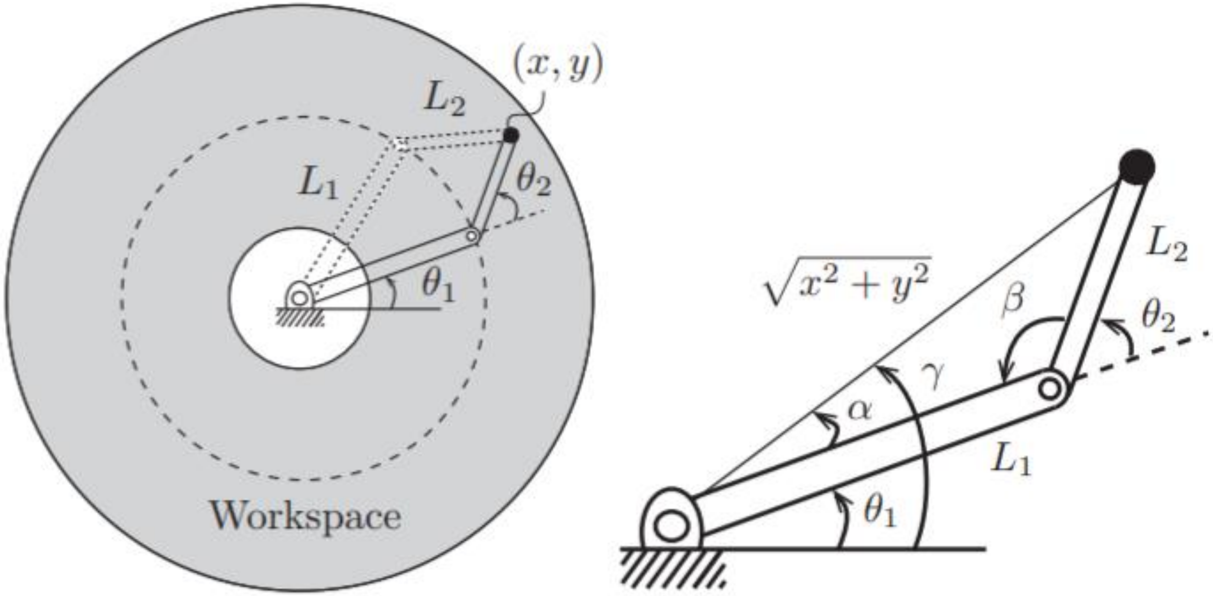

(a) A workspace, and lefty and righty configurations.

(b) Geometric solution.

Figure 6.1: Inverse kinematics of a 2R planar open chain.

[Modern Robotics, Lynch-Park, Cambridge UP]

Large kinematic structures

SWIMMING SNAKE ROBOT

Algorithmic robotics and automation

typically structured predictable environment

slightly less structured environment

looking toward unpredictable environments; lifelong planning

Q: is the cloth always below the line through the two fingers?

Cluterred environments

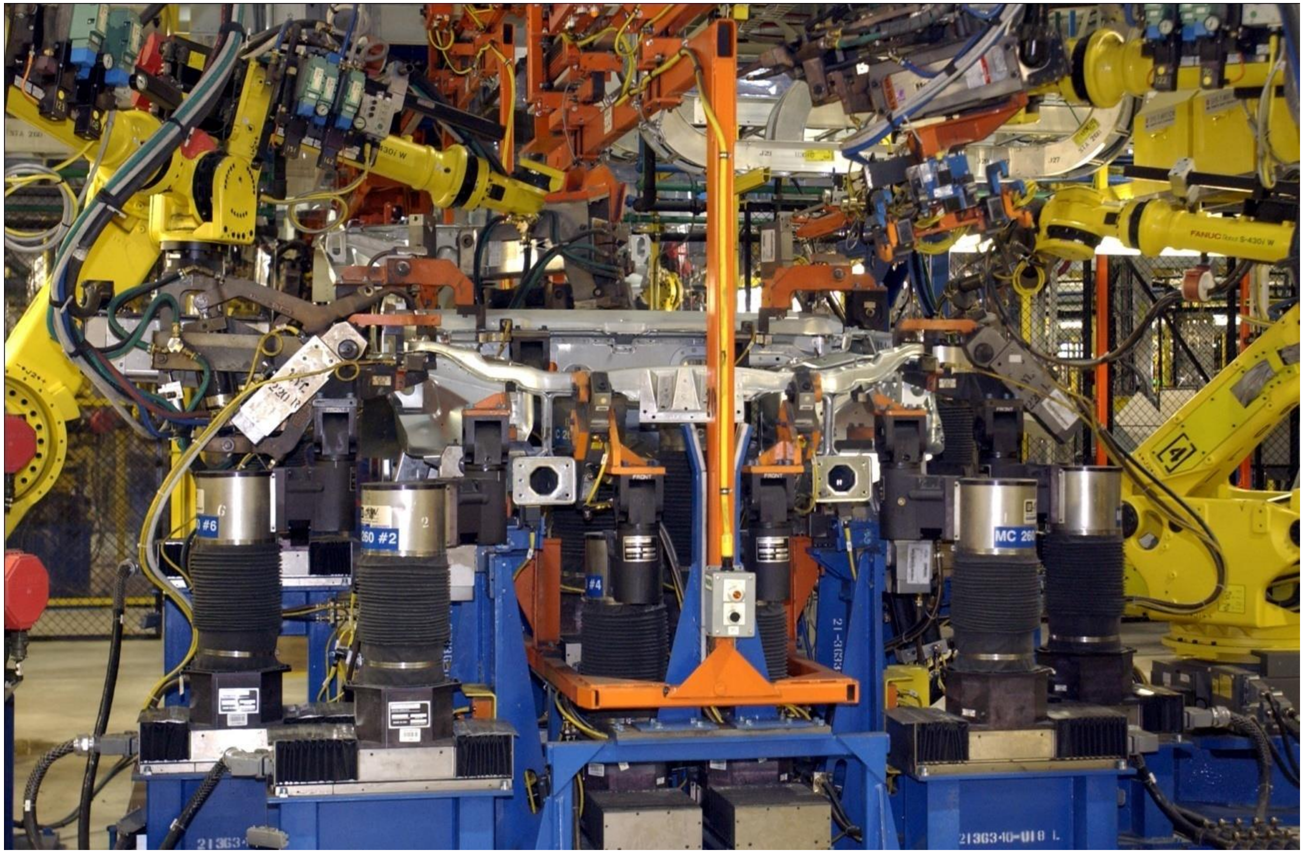

Algorithmic robotics and automation

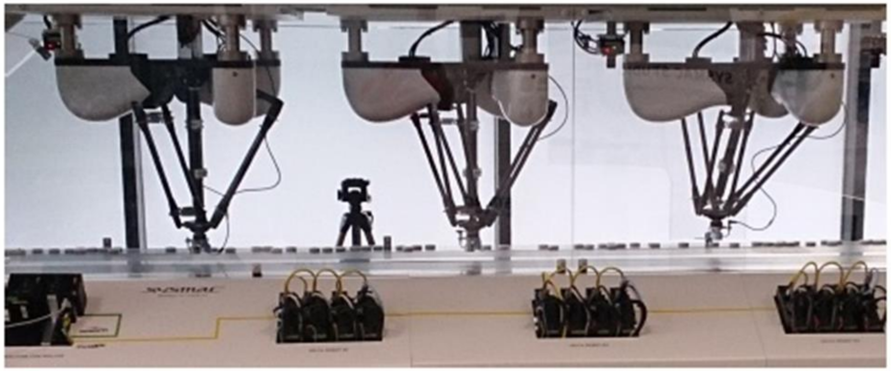

Packaging: collision detection in tight settings

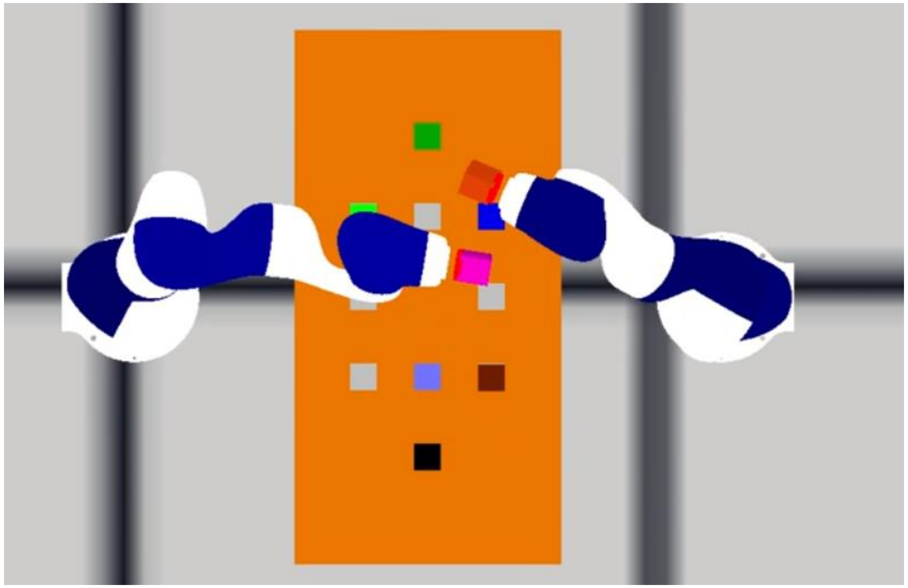

Dual arm object rearrangement

Algorithmic robotics, sensorless manipulation

Example:

the parallel jaw gripper [Goldberg]

VIDEO

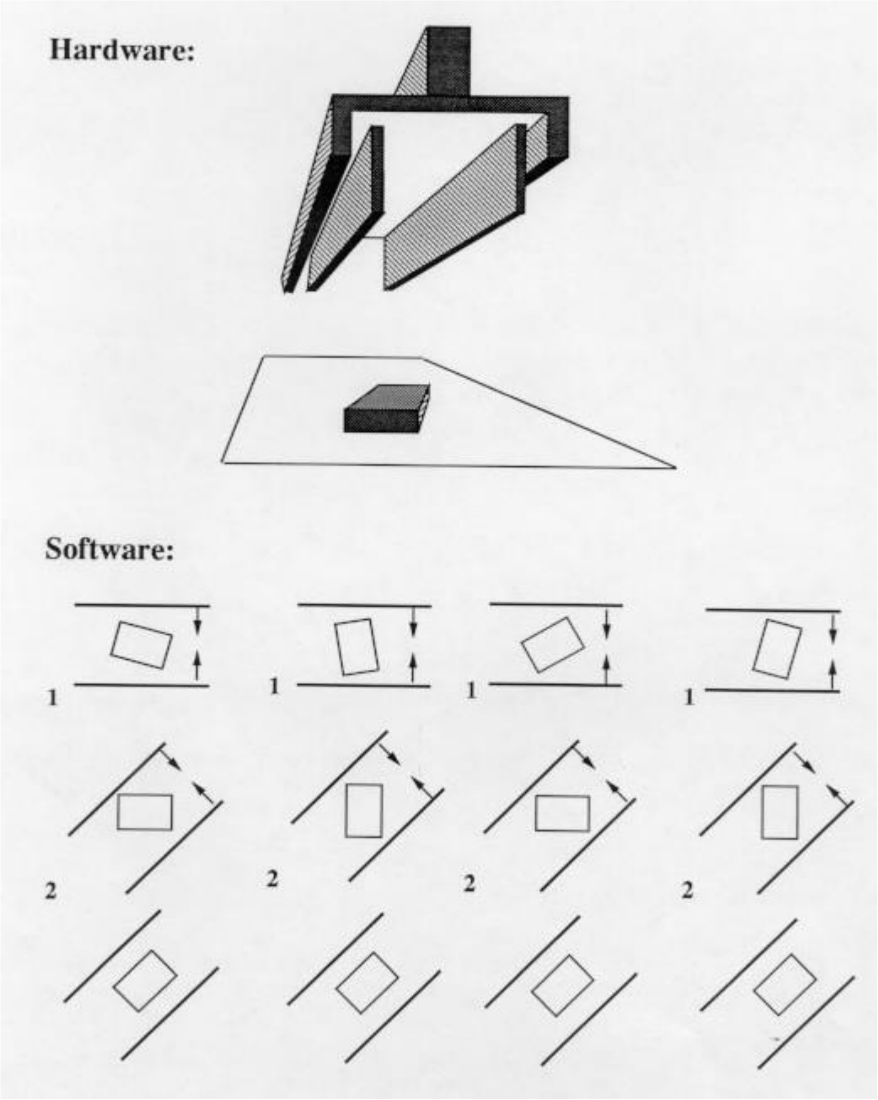

About the course

Setting your expectations

The course at a glance

The main themes

Algorithmic foundations

- Part I: Complete (exact) methods
- Part II: Sampling-based methods
- Part III: Multi-robot motion planning

Robotics at large

- Students mini-talks
- Guest lectures

Algorithmic foundations

- Part I: Complete (exact) methods
  - Arrangements, Minkowski sums, visibility graphs, Voronoi diagrams, Collins decomposition
- Part II: Sampling-based methods
  - Roadmaps, single vs. multi-query structures, probabilistic completeness, asymptotic optimality, collision detection
- Part III: Multi-robot motion planning
  - Hardness, labeled vs. unlabeled, separation assumptions, exact algorithms, SB planners

Guest lectures

- **Guy Hoffmnan**, Cornell, 23.12.19: TED
Designing Robots and Designing with Robots: New Strategies for an Automated Workplace

- **Ilana Nisky**, BGU, 6.1.20:
Haptics for the Benefit of Human Health

- **Oren Salzman**, Technion, 30.12.19:
Asymptotically-Optimal Inspection Planning with Application to Minimally-Invasive Robotic Surgery

- **Aviv Tamar**, Technion, 2.12.19:
Machine Learning in Robotics

- **Lior Zalmanson**, TAU, 25.11.19:
Trekking the Uncanny Valley --- Why Robots Should Look Like Robots?

The course at a glance

Additional topics, as time permits

- SLAM
- ROS
- Robot kinematics
- Large kinematic structures

The course at a glance

Setting your expectations, I

Algorithmic foundations

- Part I: Complete (exact) methods
- Part II: Sampling-based methods
- Part III: Multi-robot motion planning

Robotics at large

- Students mini-talks
- Guest lectures

Course mechanics

- requirements (% of the final grade):
  - assignments (40%)
  - mini talk (10%)
  - final project (50%)

- assignment types:
  - () theory
  - (p) programming, solo
  - (p2) programming, you can work and submit in pairs

- office hours: by appointment

Tailor the tasks to your interests (in part)

- 40% fixed: the assignments
- 60% adaptable: mini talk and final project

Course team

- Instructor: Dan Halperin
- Teaching assistant: Michal Kleinbort
- Grader: Yair Karin
- Software help: Michal Kleinbort, Nir Goren

Background knowledge

Setting your expectations, II

- Basic formal prerequisites: Algorithms, Data structures, Software1
- This course vs. Computational Geometry:
  - knowledge of some tools at the “API level”
  - basic reading: CG book by de Berg et al, Chapters 1&2
  - needed material will be discussed in the recitation
- Programming:
  - Python
  - some C++ might be unavoidable–we aim to provide Python bindings to C++ code, where possible
  - support will be provided in the recitation and in office hours

Main class vs recitation

- Main class, Monday 16-19, mandatory attendance
- Recitation, Monday 19-20, optional

topics of recitation: support, computational geometry tools, software tools

Mini talks

- 15 minutes
- or, 30 minutes for two students together
- topic of your choice; requires approval
- references to various up-to-date sources follow
- preferably involving more than one robot
- deadline for selecting a topic: November 25th, 2019

Final project

- compact
- topic of your choice; requires approval
- algorithm+experiments, but other options possible
- various projects will be proposed by the course team
- preferably involving more than one robot
- deadline for selecting a topic: January 5th (Sunday!), 2020

Course site

http://acg.cs.tau.ac.il/courses

Algorithmic Robotics and Motion Planning, Fall 2019-2020

includes bibliography, lesson summary, assignments and more

Conferences and journals

- Conferences

  ICRA, IROS, RSS, WAFR, …
- Journals
  - IJRR (International journal of Robotics Research),
  - IEEE TOR (Transactions on Robotics),
  - IEEE RA-L (Robotics and Automation Letters),
  - IEEE TASE (Transactions on Automation Science and Engineering),
  - Autonomous Robots,
  - …
- New conference on multi-robot systems: MRS

Bibliography I
Books

- Planning Algorithms, Steve LaValle, Cambridge University Press, 2006 (free online)
- Robot Motion Planning, Jean-Claude Latombe, Kluwer , 1991, later Springer
- Modern Robotics, Kevin Lynch and Frank Park, Cambridge University Press, 2017 (free online)
- Principles of Robot Motion: Theory, Algorithms, and Implementations, Choset et al, MIT Press, 2005

  in particular Chapter 7
- Computational Geometry: Algorithms and Applications, de Berg et al, 3rd Edition, Springer, 2008

Bibliography II
Surveys

- Sampling-Based Robot Motion Planning, Oren Salzman, Communications of the ACM, October 2019
- Sampling-Based Robot Motion Planning: A Review, Elbanhawi and Simic, IEEE Access, 2014 (free online)
- Robotics, Halperin, Kavraki, Solovey, in Handbook of Computational Geometry, 3rd Edition, 2018
- Algorithmic Motion Planning, Halperin, Salzman, Sharir, Handbook of Computational Geometry, 3rd Edition, 2018

Before the end, a little more history

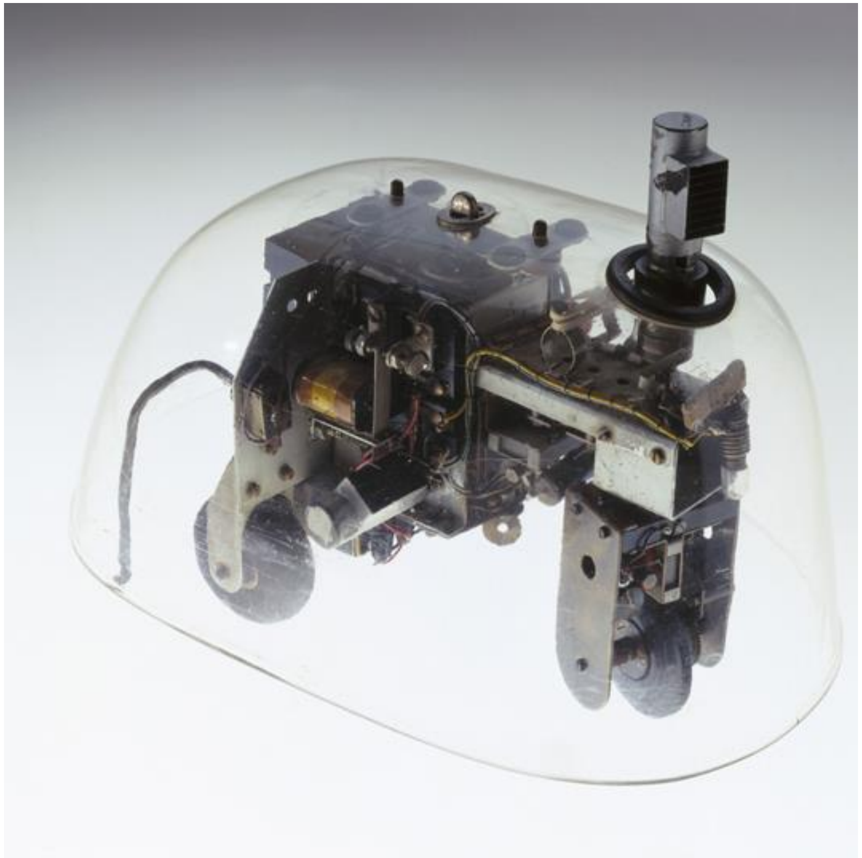

- Grey Walter's tortoises ~1948
- Turing’s visit to the Science Museum 1951

THE END