COL864: Special Topics in AI

Semester II, 2021-22

Agent Representation - I

Rohan Paul

Today’s lecture

- Last Class
    - Course Introduction
- This Class
    - Agent Representation - I

Agent Environment Interaction

- Embodied AI agent
  - Takes observations from the environment.
  - Goal or objective. Synthesizes goal-directed actions.
- Autonomous
  - You can tell the agent what to do without having to say how to do it.

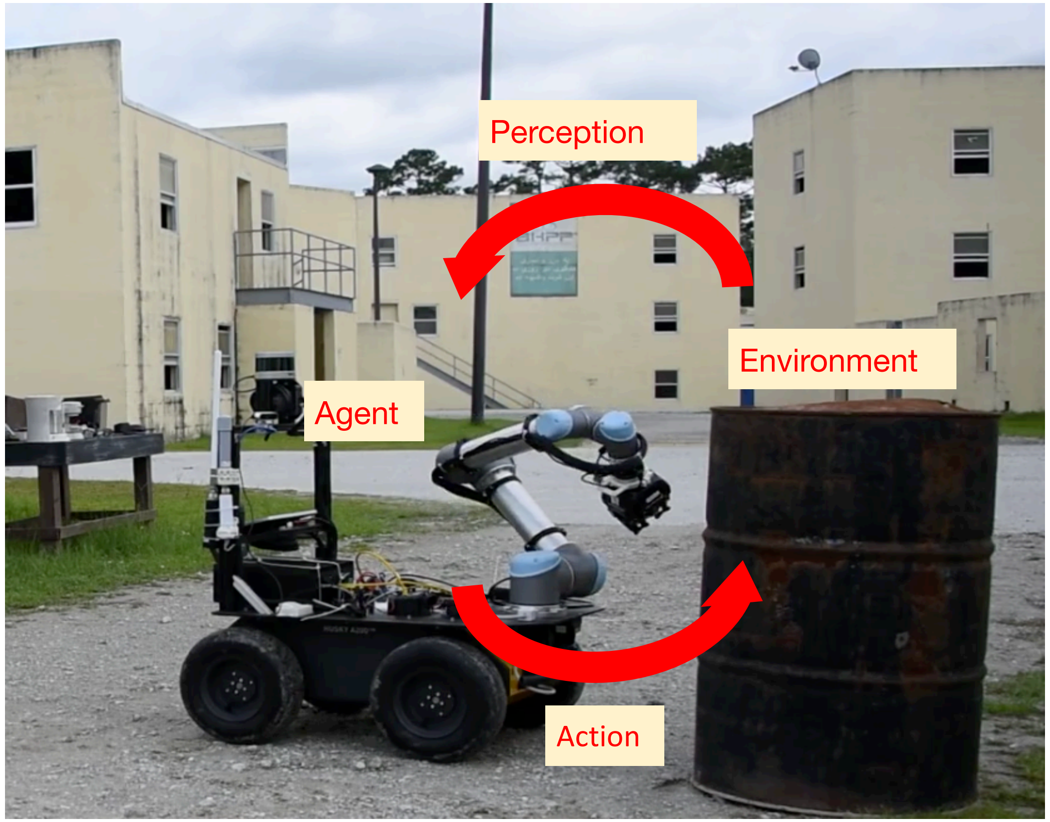

From Observations to Physical Actions

- State Estimation
  - What is the state of the agent and the environment.
- Planning
  - (High-level)Sequence of actions to reach the goal.
- Low-level Control
  - Performing each action reliably.

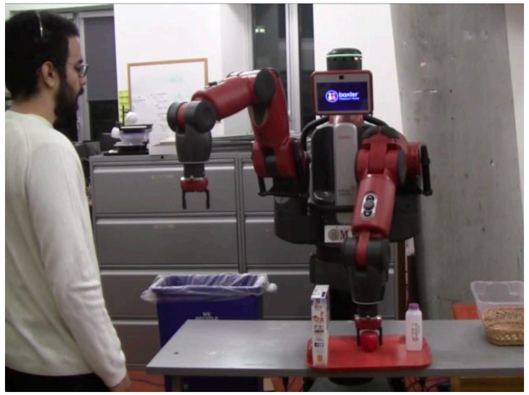

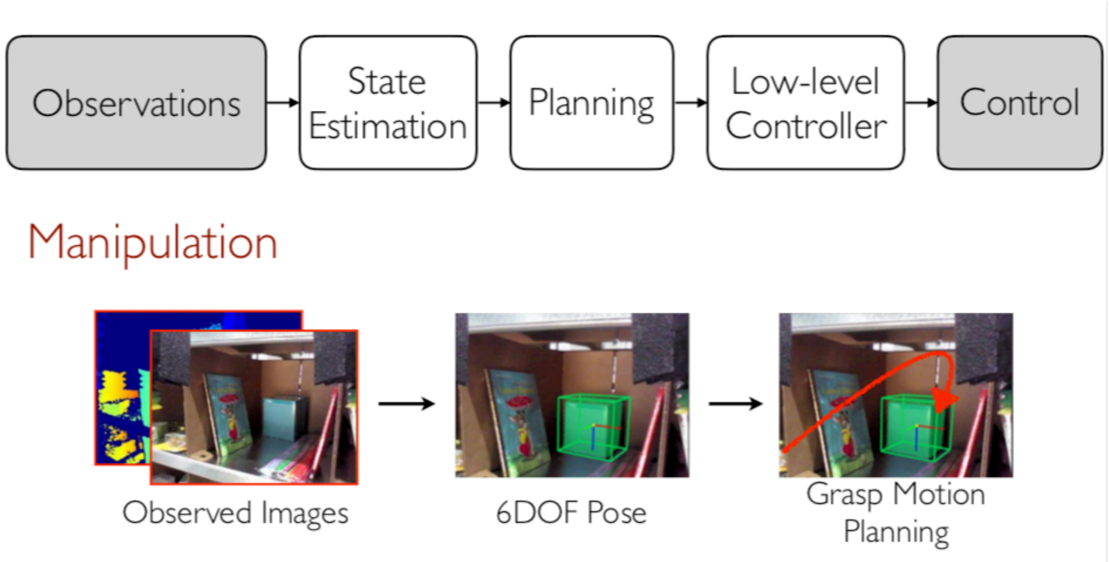

From Observations to Physical Actions

- For some environments and tasks a clear separation may not be modeled.

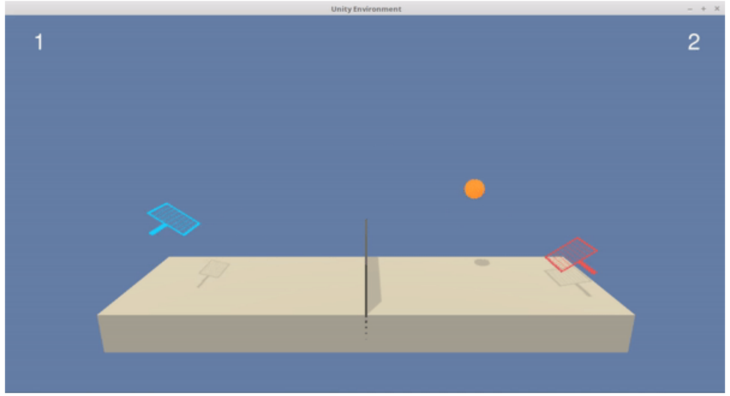

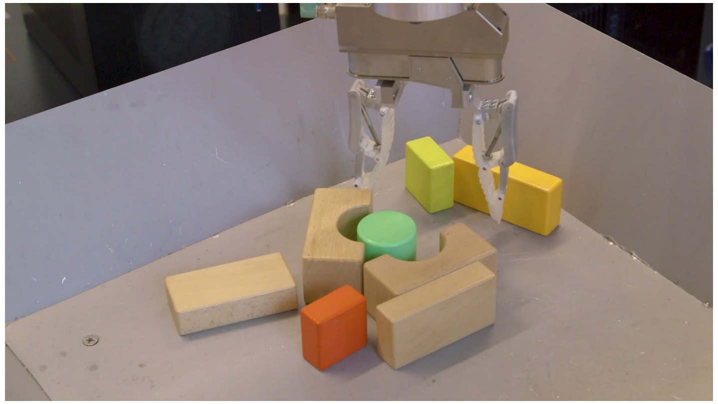

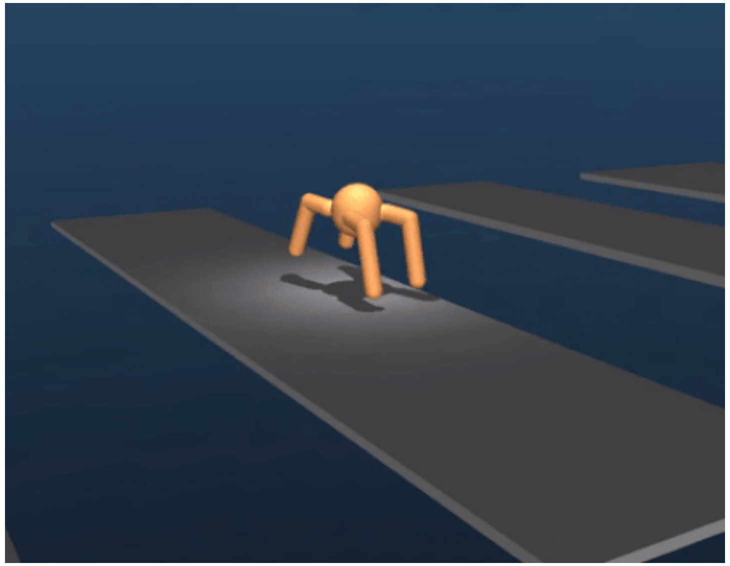

Another example: AI Habitat https://aihabitat.org/

Uncertainty

- Imagine an unmanned vehicle (land, air, manipulation etc.) in operation.
- How does the vehicle make decisions about what to do next?
- Things we might be uncertain about:
  - If the vehicle runs its propulsion system (or motors), what will happen? Is the vehicle working or not?
    - The world is stochastic.
  - Where is the vehicle?
    - Noisy observations.
  - What is around the vehicle?
    - The world is partially unobservable.

Any decision-making model must tackle the uncertainty and complexity as above.

Physical Interactions: Generating Movement

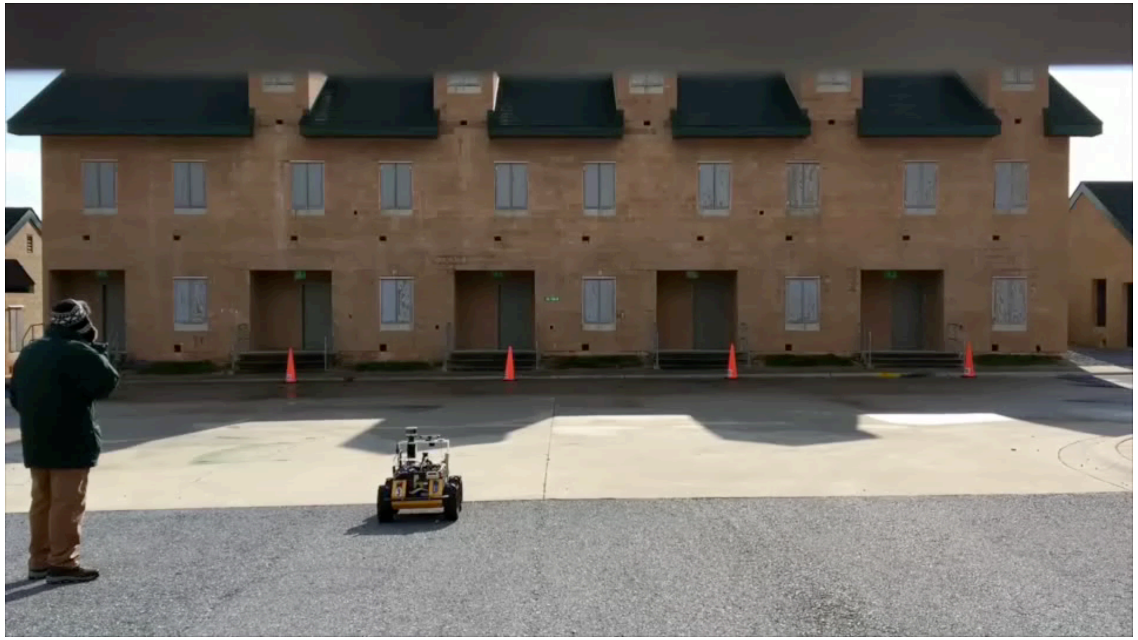

Coordinate Frames

- Coordinate frame
  - Global
  - On the agent.
- Pose, a particular point on its body, (e.g., x, y, z, heading etc.) with respect to a coordinate frame.
- State changes occur due to actions the agent takes.
- Need a way to determine the pose w.r.t. any coordinate frame.
- Relationship maintained in coordinate transforms.

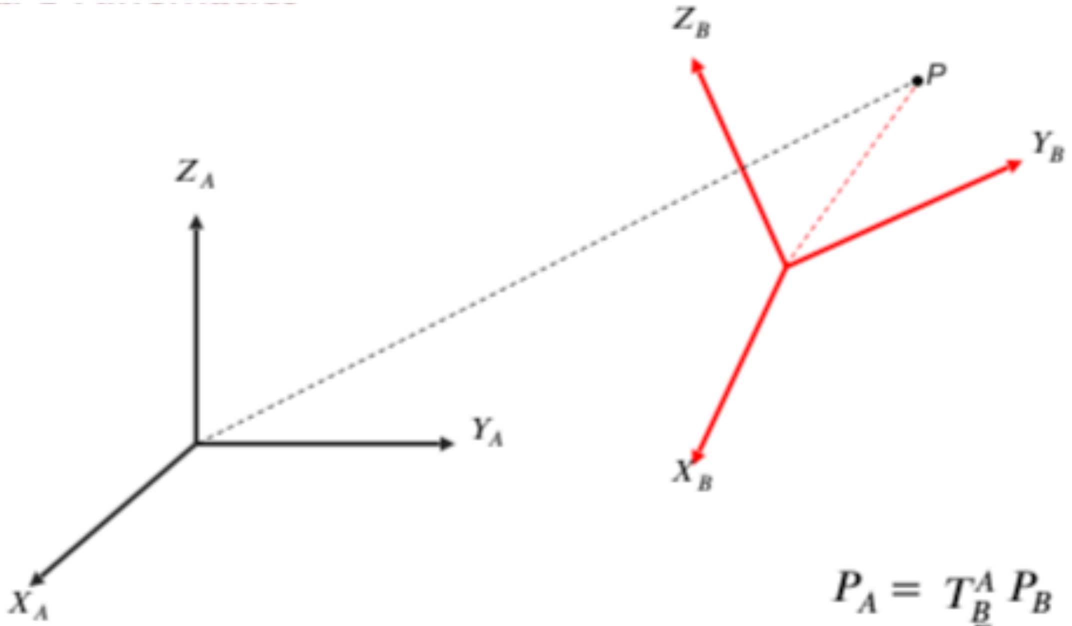

$$P_A = T_B^A P_B$$

$$T_B^A = \begin{bmatrix} r_{11} & r_{21} & r_{31} & \Delta x \\ r_{12} & r_{22} & r_{32} & \Delta y \\ r_{13} & r_{23} & r_{33} & \Delta z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

How does movement occur?

- Simple agents
  - Point-like or disk-like agents
- Linear and angular velocities change the state from current to the next time step.
- Position and heading direction is changed.

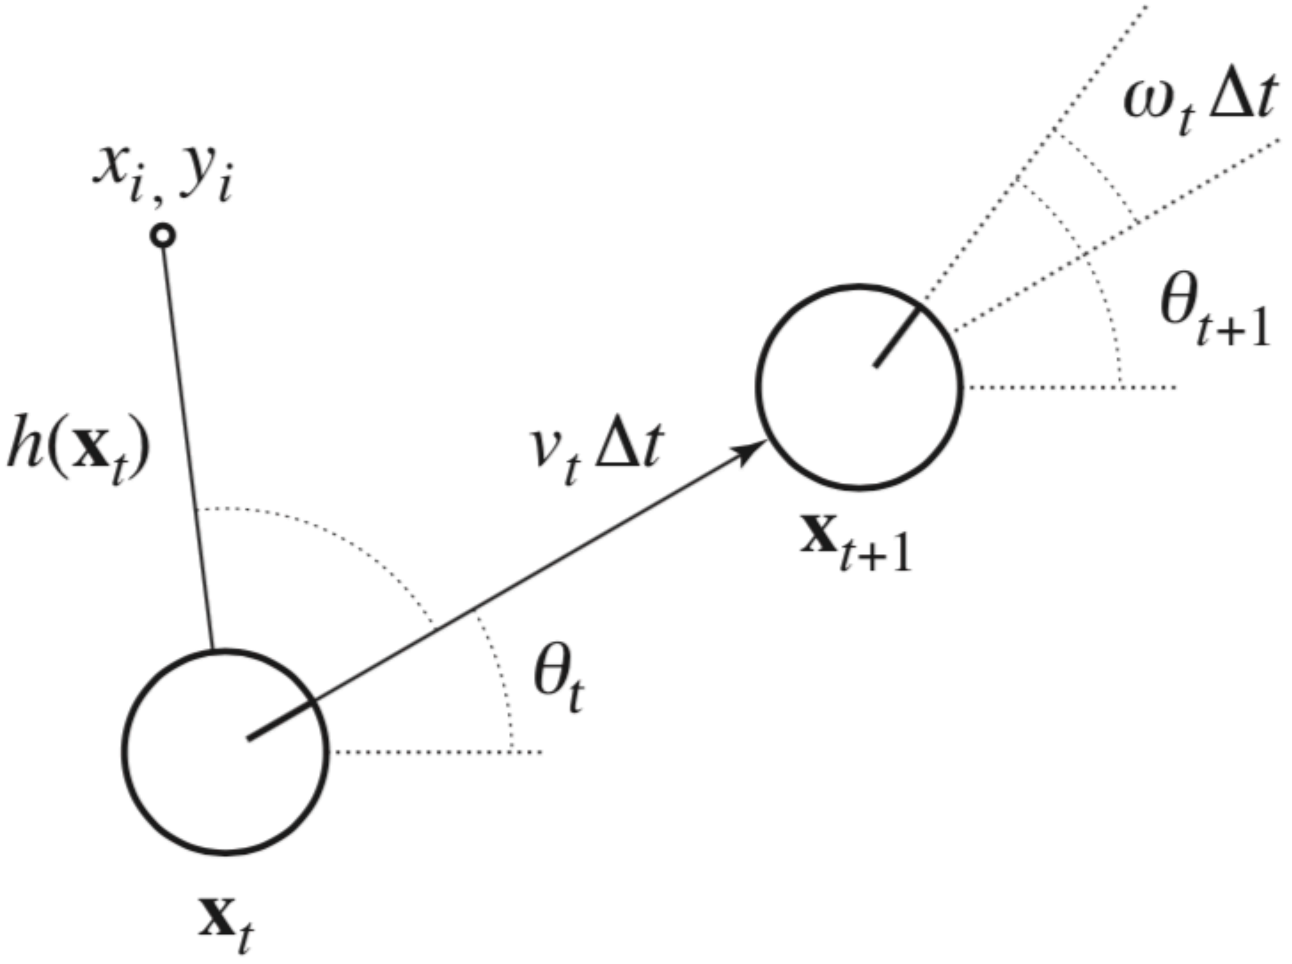

More complex physical agents

- An agent capable of manipulation
  - Simply a number of links connected via joints.

How is movement generated ?

- Joints typically have motors that exert torque.
- Also called actuation Typically, abstracted away for algorithmic purposes.

Internal mechanics of joints

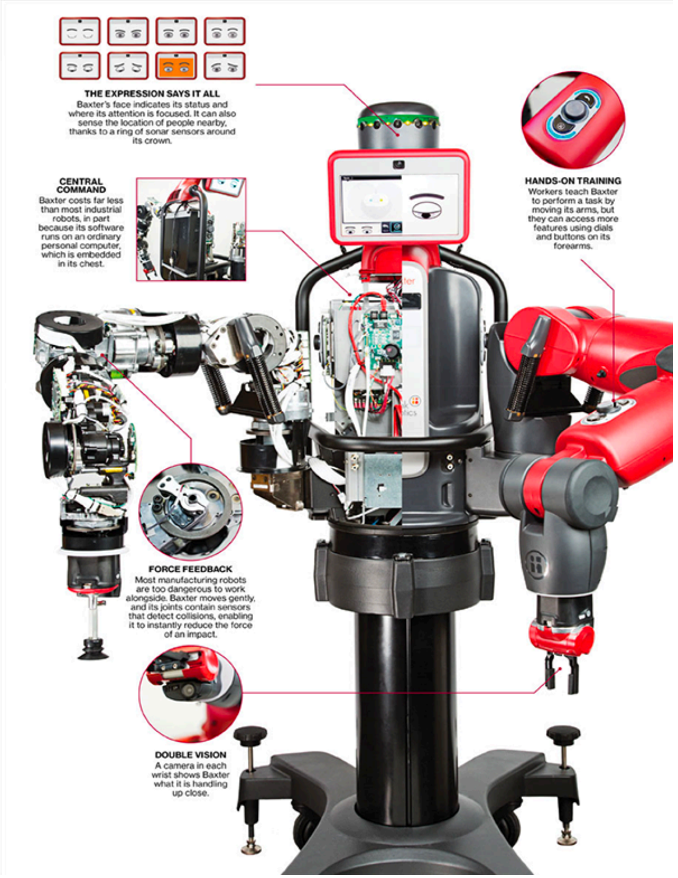

DC motor

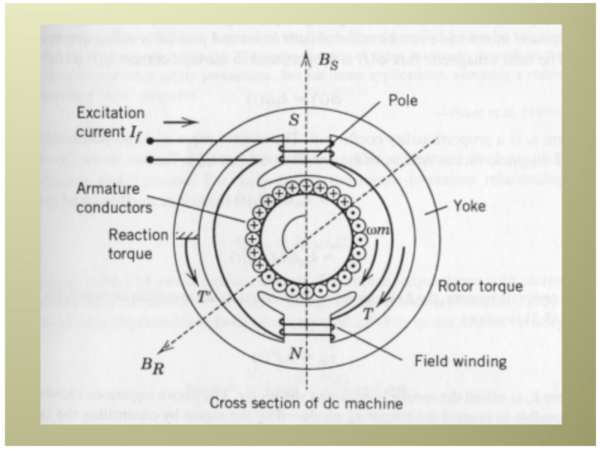

Stepper motor

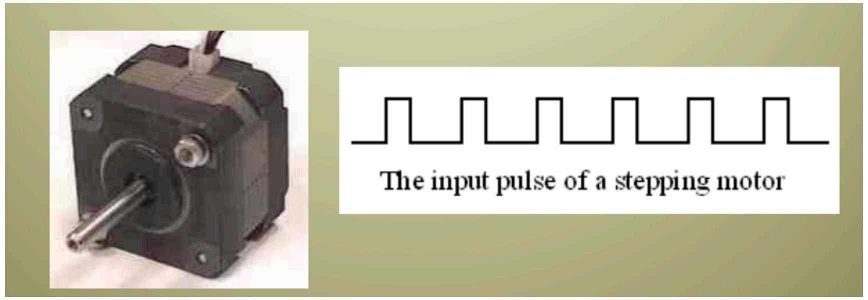

Composing Transforms

- An agent capable of manipulation
    - Simply a number of links connected via joints.
- Each link is associated with a. coordinate frame.
- The coordinate frames can be composed.

Physical agent

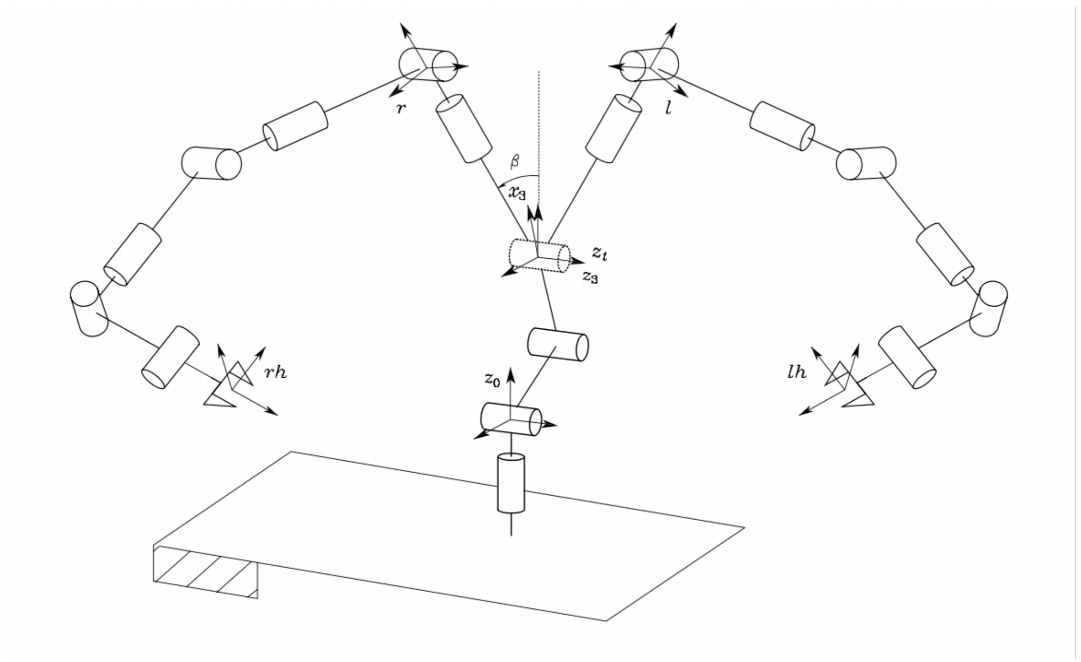

Kinematic chains

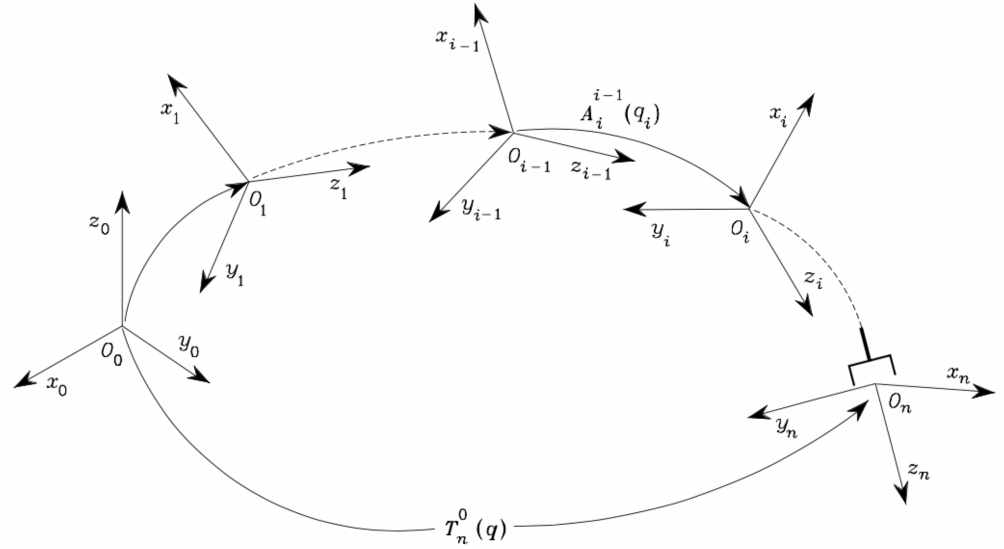

Planning Motions

- Given
  - Initial state
  - Goal state
- How to generate actions?

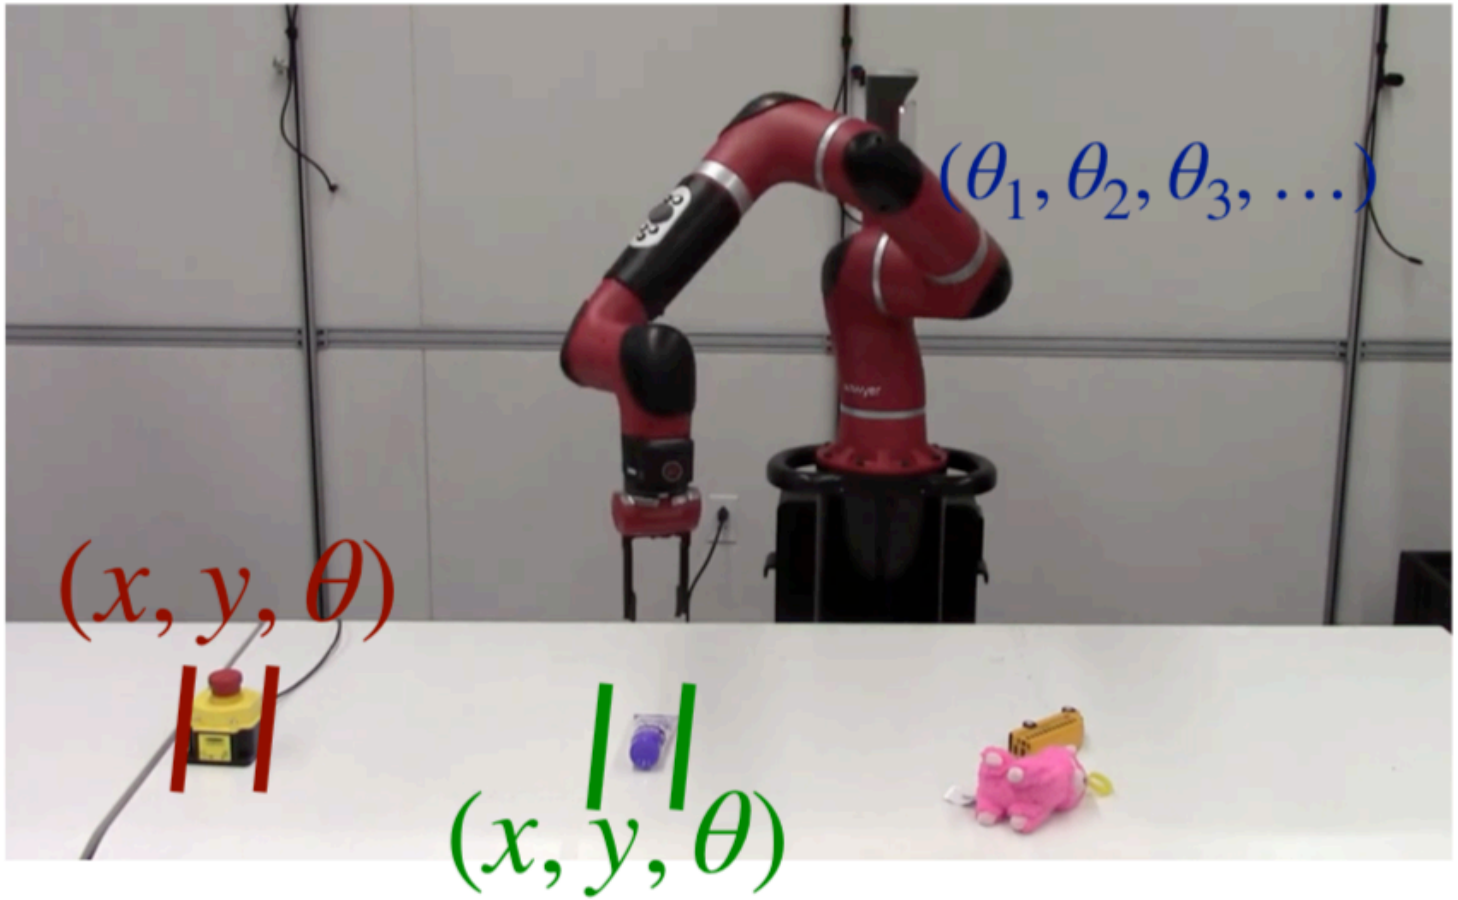

Task Space and Configuration Space

- Task Space
  - Workspace in which the agent operates.
  - May be populated with obstacles.
- Configuration space (C-space)
  - Parameter space of the robot.
  - Space spanned by its allowable degrees of freedom.
- C-space obstacles
  - Non allowable configurations in the C-space.
  - Includes all configurations where the robot collides with the obstacle.

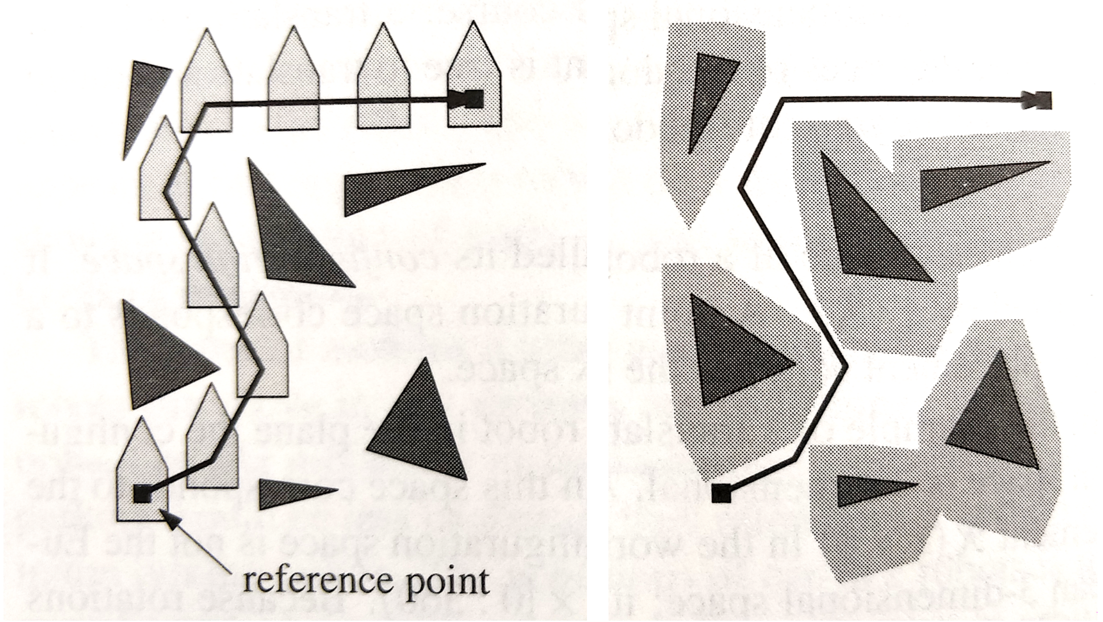

Task space (left) and configuration space (right) for a translating planar agent (cannot rotate).

Task Space and Configuration Space

Task space

- Easier to describe the task.
- Problem: Not all points in the task space may be reachable
- There are physical constraints of the agent’s embodiment.

Configuration space (C-space)

- Each point satisfies the intrinsic physical constraints.
- Each point is attainable as it corresponds to a valid configuration of the agent.

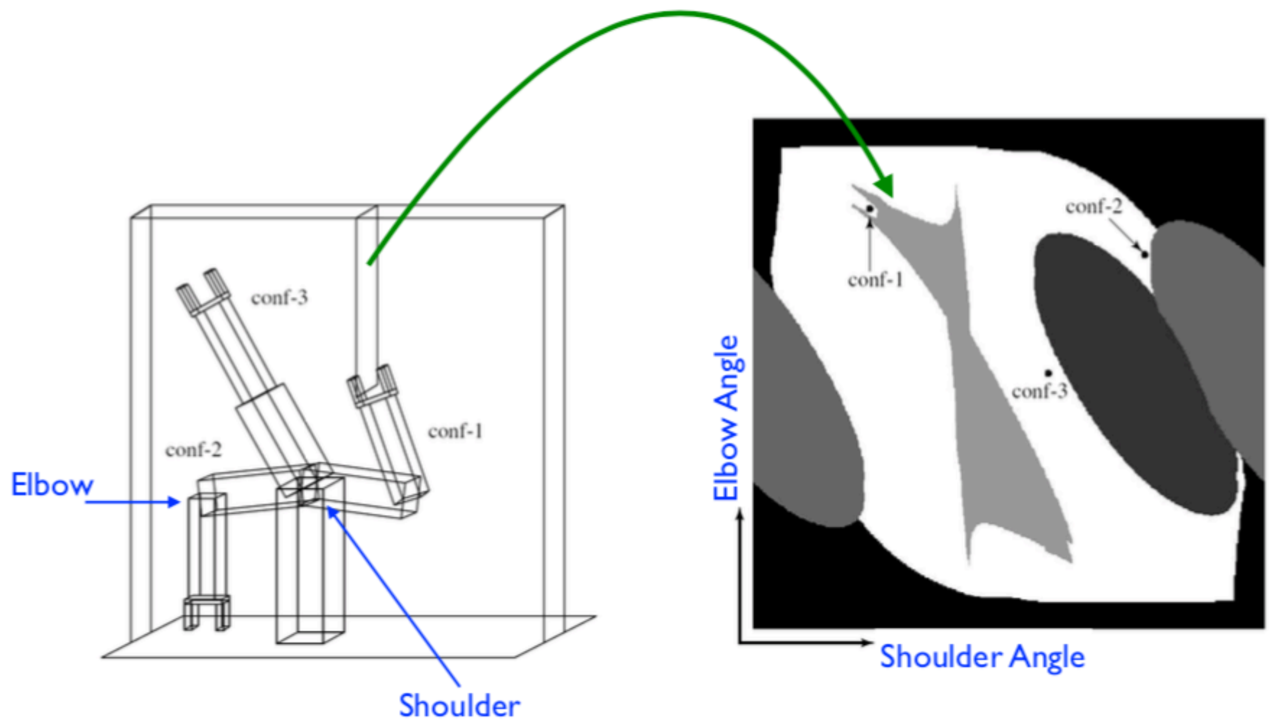

Figure: AIMA Ch 25.4

Holonomicity

- Robot specified in terms of parameters (degree of freedom).
- Holonomic
  - if the number of local degrees of freedom of movement equals the number of global degrees of freedom.
- Non-holonomic otherwise.

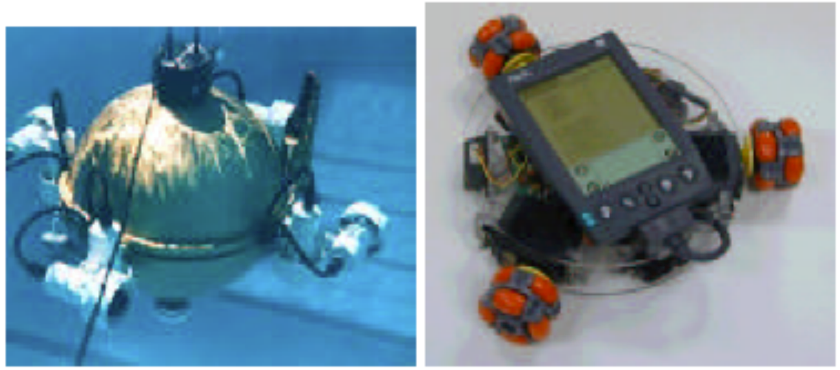

Holonomic (ODIN, University of Hawaii and PPRK (CMU).

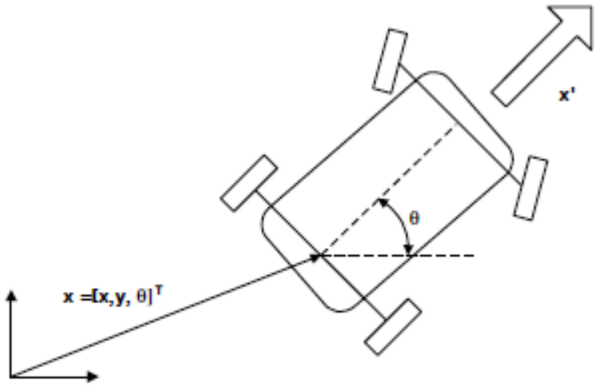

Non-holonomic (car cannot turn in place).

From Configuration Space to Task Space

- Forward kinematics
  - From the parameters specifying the agent (theta parameters) determine the position of the end-effector.
  - Compose the transformations one by one.

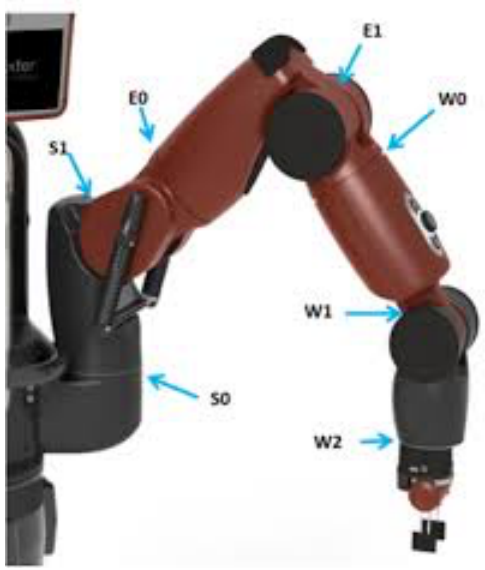

Forward Kinematics

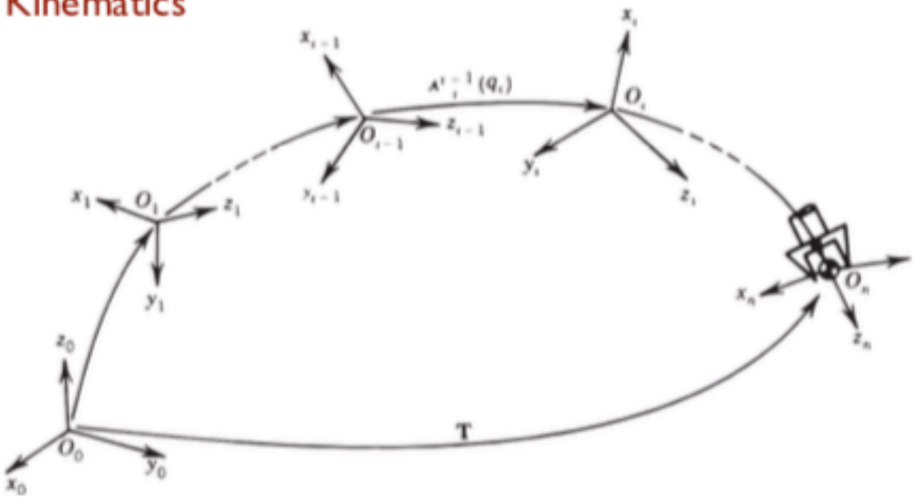

$$T = T_1^0(\theta_1)\, T_2^1(\theta_2)\ldots \quad T_{n-1}^{n-2}(\theta_n) \quad T_n^{n-1}(\theta_n)$$

Maps configuration space to work space

$$= \begin{bmatrix} r_{11} & r_{21} & r_{31} & \Delta x \\ r_{12} & r_{22} & r_{32} & \Delta y \\ r_{13} & r_{23} & r_{33} & \Delta z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$x = f(\theta) = f(\theta_1, \theta_2, \ldots, \theta_{n-1}, \theta_n)$$

From Task Space to Configuration Space

- Inverse kinematics
  - Determining the setting of the parameters to yield the required end effector position.
  - Given the final position of the end-effector (end point of the arm), determine the theta parameters for the arm.

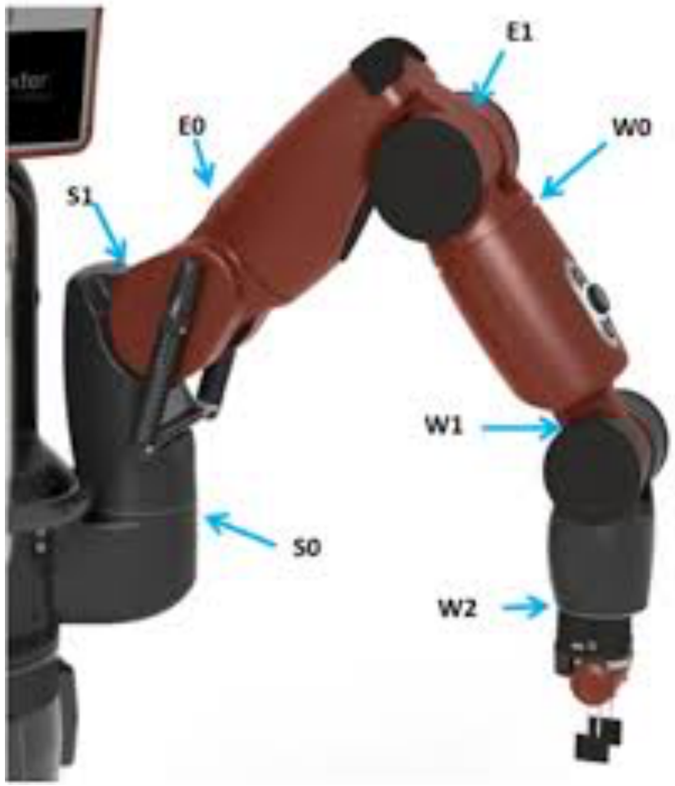

From Task Space to Configuration Space

- Inverse kinematics
  - Solve the inverse problem to obtain theta
    The function f is often non-linear.

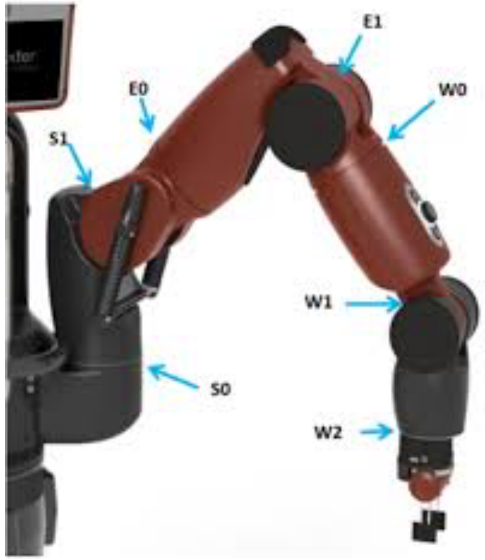

Forward Kinematics

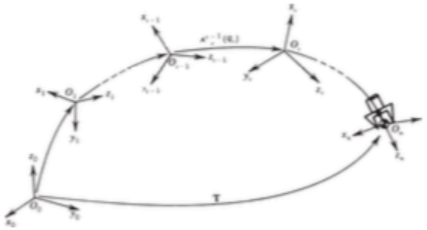

$$x = f(\theta)$$

Maps configuration space to work space

Inverse Kinematics

Solve for θ_d in:

$$x_d - f(\theta_d) = 0$$

Typically done using Newton-Raphson method.

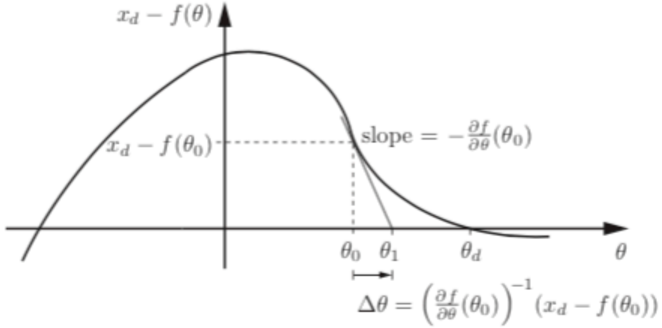

Find configuration(s) that map to a given work space point

Example: Moving the Agent

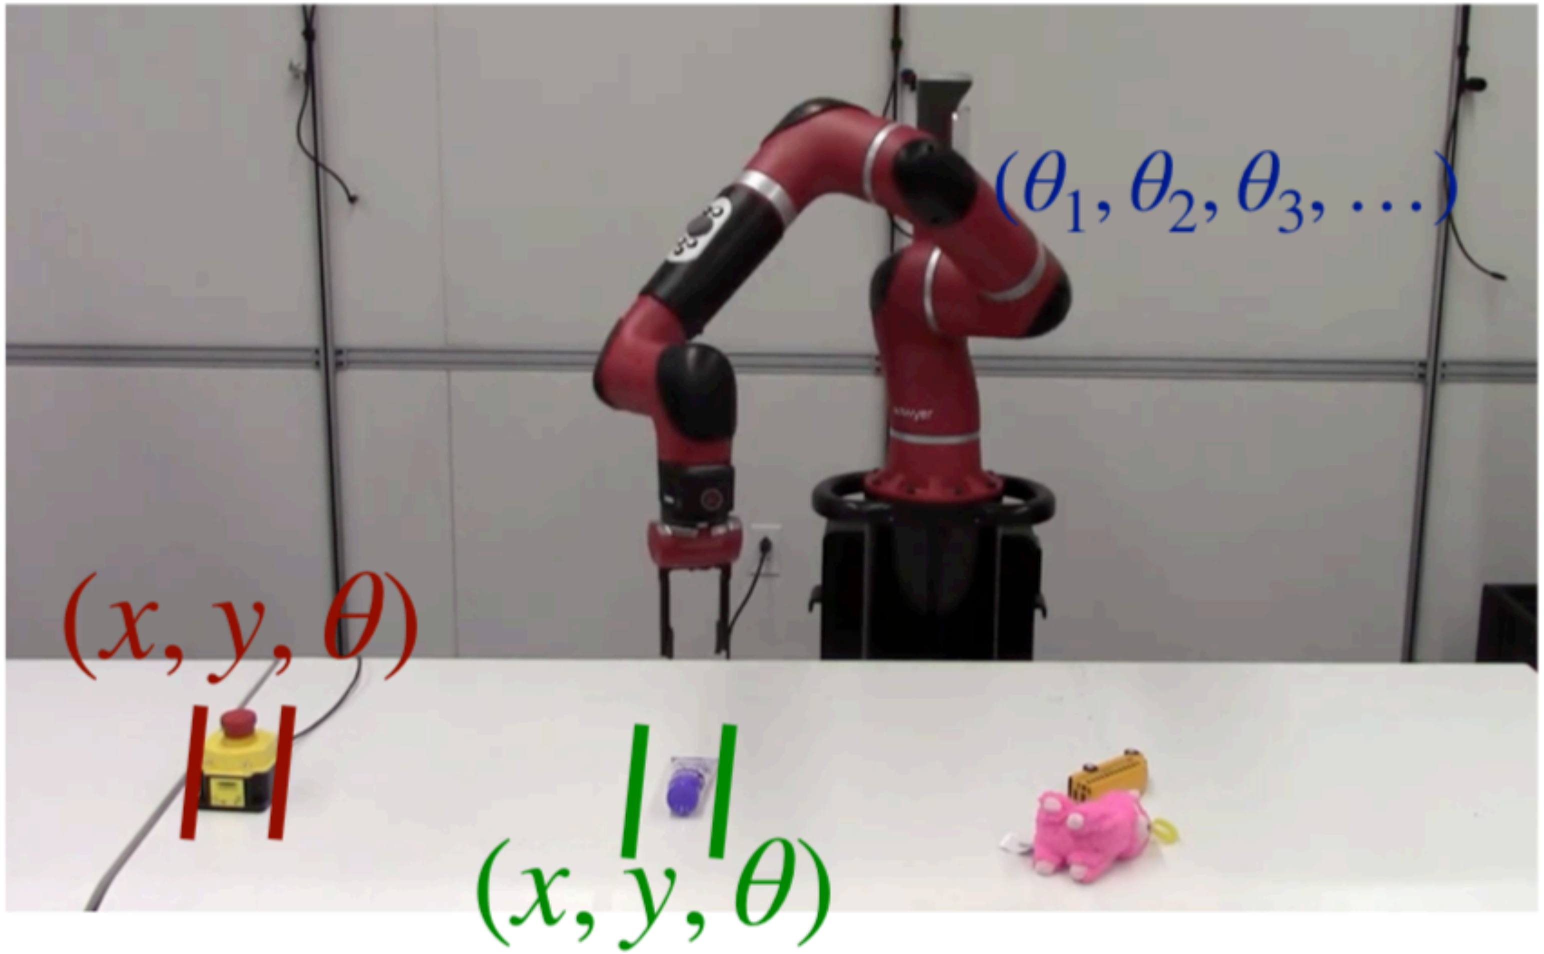

Initial Configuration

1. Task space to Configuration space

Initial configuration

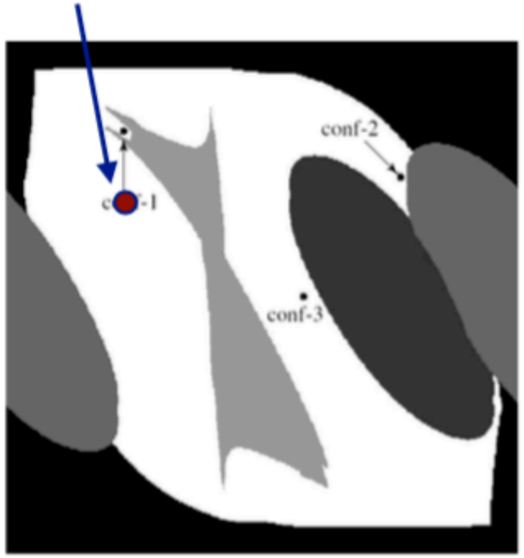

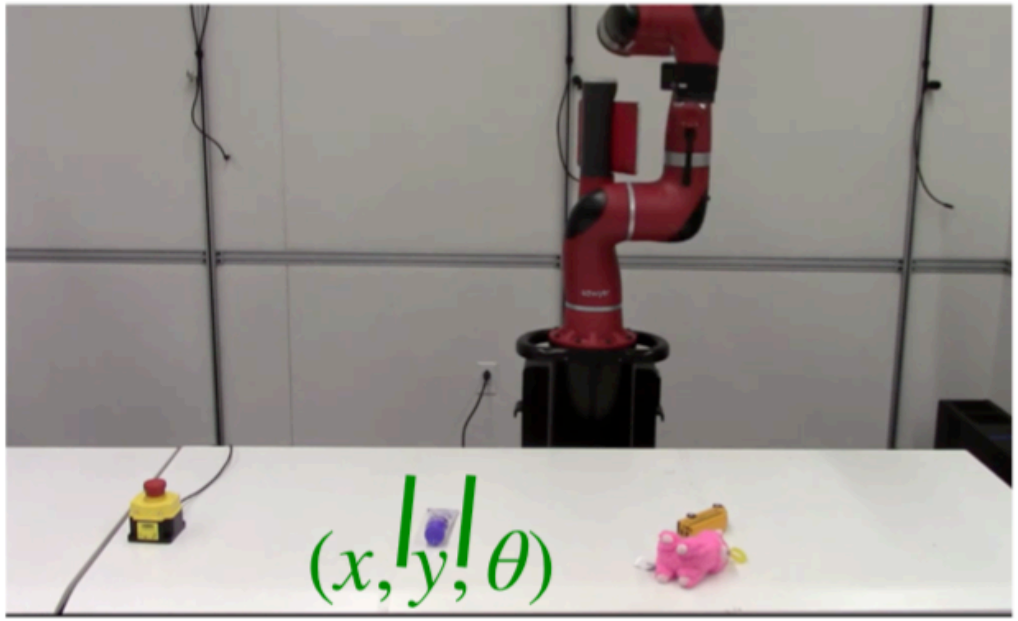

Goal Configuration

1. Task space to Configuration space

Initial configuration

Desired configuration

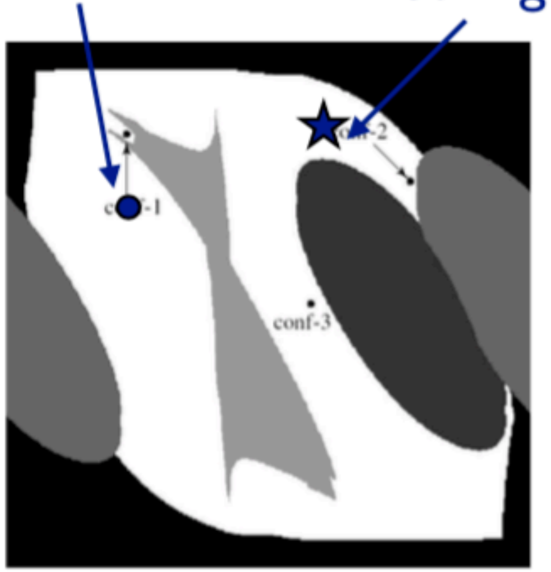

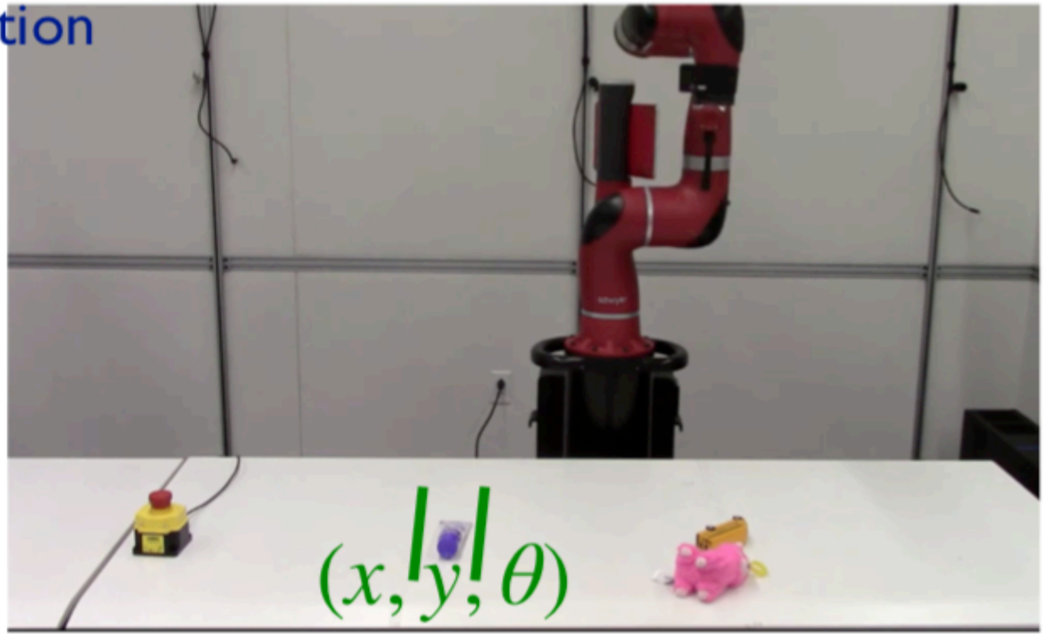

Finding a feasible path

1. Task space to Configuration space

2. Configuration space trajectory

Initial configuration

Desired configuration

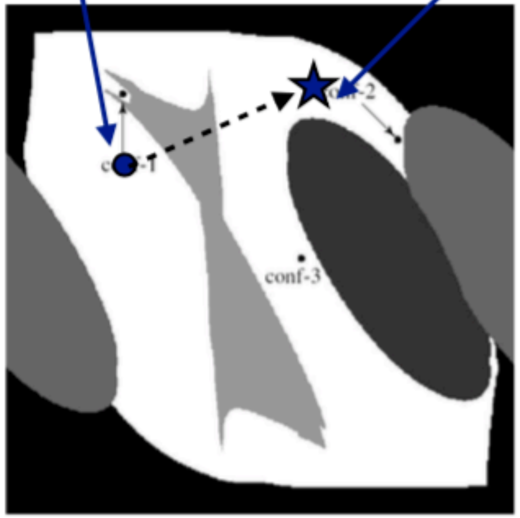

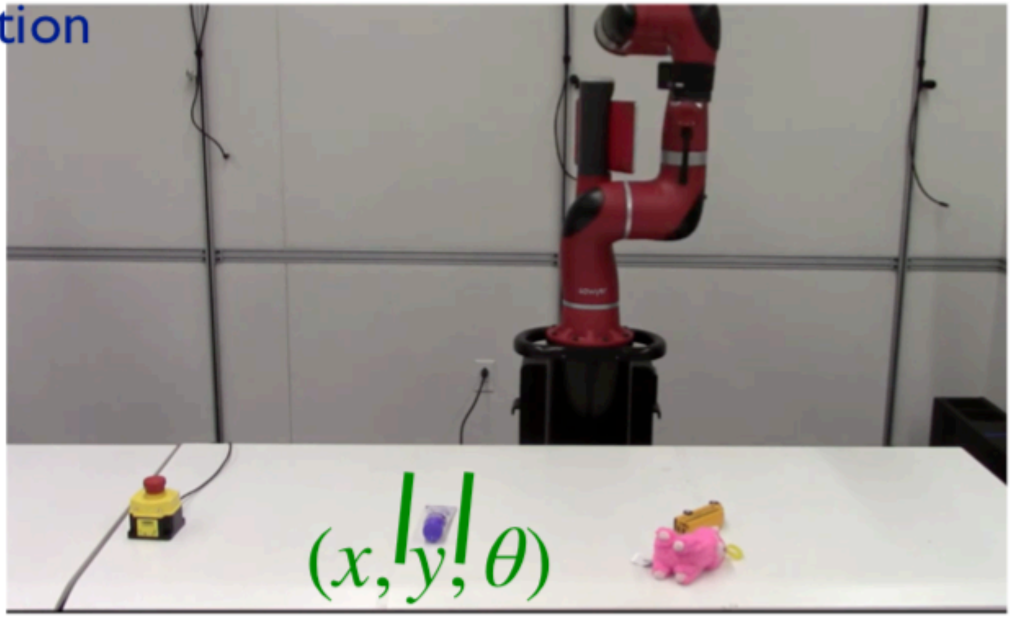

Finding a feasible path

1. Task space to Configuration space

2. Configuration space trajectory