

Overview, agents, environments, typical components

CSC752 Autonomous Robotic Systems

Ubbo Visser

Department of Computer Science
University of Miami

August 25, 2025



UNIVERSITY OF MIAMI
ROBOCANES



Outline


- 1 Autonomous robots
- 2 Agents
- 3 Environments
- 4 Agent types
- 5 Typical components
- 6 Example soccer robot
- 7 Example service robot



From automation...



...to autonomy



Waymo Autonomous Vehicle



Spot Robot



Adult Humanoids, RoMeLa, UCLA & NimbRo, U Bonn



Da Vinci Surgical Robot



Solo, 3D Robotics

Course goals

Course goals

- To learn the *theoretical*, *algorithmic*, and *implementation* aspects of important techniques for robot autonomy. Specifically, you will
 - 1 Gain a fundamental knowledge of the autonomy stack
 - 2 Be able to apply such knowledge in applications / research by using ROS
 - 3 Devise novel methods and algorithms for robot autonomy

Autonomous robots

Robot

A robot is a autonomous system which exists in the physical world, can sense its environment and can act on it to achieve some goals.

Autonomous robot

- An autonomous robot acts on its own decisions.
- It is not **directly** controlled by humans.
- Take an appropriate action for any given situation.

Robots

- **Situatedness**

Agents are strongly affected by the environment and deal with its immediate demands (not its abstract models) directly.

- **Embodiment**


Agents have bodies, are strongly constrained by those bodies, and experience the world through those bodies, which have a dynamic with the environment.

Agents

Agent definitions

- Russell und Norvig
 - "An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors."
- Winston
 - "The study of the computations that make it possible to perceive, reason, and act."

Agents and environments



Agents interact with environments through sensors and actuators.

- Perception, perception sequences
- Agent function (abstract)
- Agent program (concrete)

Good behavior: rationality

- Rational agent
 - A rational agent is one that does the right thing...
 - First approximation, we will say that the right action is the one that will cause the agent to be most successful.
 - Problem: How and when do we decide whether or not the agent was successful?
- Performance measures
 - Subjective
 - Agent evaluates himself.
 - Objective
 - Evaluation done by observer: he defines standards for being successful in the environment.
 - Example: soccer agent.


Good behavior

- Omniscience and rationality
 - An omniscient agent knows the effects of its actions and can act accordingly.
 - But: who knows it all? → theoretical.
 - Rationality: expected success based on things that can be perceived.
- Rationality based on
 - The performance measure that defines the criterion of success.
 - The agent's prior knowledge of the environment.
 - The actions that the agent can perform.
 - The agent's percept sequence to date.

Ideal rational agent

- An ideal rational agent...
 - For each possible percept sequence, a rational agent should select an action that is expected to maximize its performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has.
- Autonomy
 - Inherent knowledge.
 - A system is autonomous, if its behavior is determined by its own experience.


The see-think-act cycle knowledge



Source: R. Siegwart et al. 2015, ETH Zürich, Lecture series on Autonomous Mobile Robots


Motion control and kinematics

- Wheel types and its constraints
 - Rolling constraint
 - no-sliding constraint (lateral)
- Motion control




$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix} = f(\dot{\phi}_1 \cdots \dot{\phi}_n, \theta, geometry)$$

$$\begin{bmatrix} \dot{\phi}_1 \\ \vdots \\ \dot{\phi}_n \end{bmatrix} = f(\dot{x}, \dot{y}, \dot{\theta}) \quad ?$$




The see-think-act cycle knowledge






Sensing

■ Laser scanner


- time of flight



■ Camera



Information extraction



- Filtering / Edge Detection

- Keypoint Features
 - features that are reasonably invariant to rotation, scaling, viewpoint, illumination
 - FAST, SURF, SIFT, BRISK, ...





Image from [Rosten et al., PAMI 2010]

- Keypoint matching
 - BRISK example




The see-think-act cycle knowledge




Localization


- **SEE:** The robot queries its sensors
→ finds itself next to a pillar
- **ACT:** Robot moves one meter forward
 - motion estimated by wheel encoders
 - accumulation of uncertainty
- **SEE:** The robot queries its sensors
again → finds itself next to a pillar
- **Belief update (information fusion)**



The see-think-act cycle knowledge




Cognition, where am I going, and how?




Cognition, where am I going, and how?


- Global path planning
 - Graph search



- Local path planning
 - Local collision avoidance



The see-think-act cycle knowledge



Environments

Agent Type	Performance Measure	Environment	Actuators	Sensors
Taxi driver	Safe, fast, legal, comfortable trip, maximize profits	Roads, other traffic, pedestrians, customers	Steering, accelerator, brake, signal, horn, display	Cameras, sonar, speedometer, GPS, odometer, accelerometer, engine sensors, keyboard

- PEAS for an automated taxi
 - P: performance measure
 - E: environment
 - A: actuators/effectors
 - S: Sensors

Environments

- Fully observable vs. partially observable
- Deterministic vs. stochastic
- Episodic vs. Sequential
- Static vs. Dynamic
- Discrete vs. Continuous
- Single-agent vs. Multi-agents

Task Environment	Observable	Deterministic	Episodic	Static	Discrete	Agents
Crossword puzzle	Fully	Deterministic	Sequential	Static	Discrete	Single
Chess with a clock	Fully	Strategic	Sequential	Semi	Discrete	Multi
Poker	Partially	Strategic	Sequential	Static	Discrete	Multi
Backgammon	Fully	Stochastic	Sequential	Static	Discrete	Multi
Taxi driving	Partially	Stochastic	Sequential	Dynamic	Continuous	Multi
Medical diagnosis	Partially	Stochastic	Sequential	Dynamic	Continuous	Single
Image-analysis	Fully	Deterministic	Episodic	Semi	Continuous	Single
Part-picking robot	Partially	Stochastic	Episodic	Dynamic	Continuous	Single
Refinery controller	Partially	Stochastic	Sequential	Dynamic	Continuous	Single
Interactive English tutor	Partially	Stochastic	Sequential	Dynamic	Discrete	Multi

Example: standard problems: Chess vs. RoboCup

	Chess	RoboCup
Environment	static	dynamic
State change	with each move	always (real time)
Information access	given	incomplete
Sensors	symbolic	not symbolic
Control	central	distributed

- A RoboCup environment is a partially observable, stochastic, dynamic, continuous, multi-agent environment.
- Real-time.

Robots and uncertainty

- Uncertainty is a key property of existence in the physical world.
- Physical sensors provide limited, noisy, and inaccurate information.
- Physical effectors produce limited, noisy, and inaccurate action.
- The uncertainty of physical sensors and effectors is not well characterized, so robots have no available a priori models.


Robots and uncertainty

- A robot can not accurately know the answers to the following:
 - Where am I?
 - Where are my body parts, are they working, what are they doing?
 - What did I just do?
 - What will happen if I do X?
 - Who/what are you, where are you, what are you doing, etc.?
 - ...

Agent types

Agents types

Spectrum of robot control:



From "Behavior-Based Robotics" by R. Arkin, MIT Press, 1998


Types of agent programs

We outline four basic kinds of agent programs that embody the principles underlying almost all intelligent systems:

- Simple reflex agents
 - condition-action rules
- Model-based reflex agents
 - internal states
- Goal-based agents
 - explicit goals, more flexible
- Utility-based agents
 - explicit utility functions, degree of happiness

Types of agent programs


Simple reflex agents



- Actions based only on the current percept
- No history

Types of agent programs


Model-based reflex agents



- History for partial access of environment
- Internal states
- Update needs two kinds of knowledge:
 - How does the world function without agent
 - What kind of effects does agent have on environment
→ Model of the world

Types of agent programs


Goal-based agents



- Model-based and goal-oriented agent
- Goal helps select actions
- Combination of goal and feasible actions
- Selection sometimes easy, most of the time difficult
→ search, planning

Types of agent programs


Utility-based agents



- Goal-orientated sometimes not enough → e.g. various paths to Rome
- Priority with utility value
- Utility function as mapping between state and a real number
- Advantages with goal conflicts and uncertainty

Types of agent programs

Learning agents



- Learning element for improvement
- Performance element for selection of external actions
- Critique: performance of agent?
- Problem generator for exploration


Typical components

Typical components

- Previous agent types from "S. J. Russell and P. Norvig. Artificial Intelligence: A Modern Approach."
- Focus so far on decision-making.
- Usually there are other parts in the architecture of an autonomous robot.


Typical components

The model-based reflex agent:




Typical components


Split into a modeling and behavior:




Typical components




Typical components



Typical components



Typical components



Example soccer robot


Example soccer robot



- What do these robots do?
- Same categories:
 - Perception
 - Modeling
 - Behavior
 - Motions
 - Control

Modeling

- Self-localization
- Estimate orientation of the robot (standing/lying)
- Ball tracking
- Opponent tracking



Behavior

- Decide what to do based on
 - current world model,
 - team communication,
 - role,
 - current plan,
 - internal state,
 - ...
- Select actions (e.g. "walk forward", "left kick")

Motion & control

- Motion:
 - Walking, kick, stand-up, ...
 - Set an angle for each joint.
 - Calculate trajectories, inverse kinematics, balancing, ...
 - Execute static angle sequences.
- Control
 - Move joints to the target positions.


Example service robot

Example service robot, HSR TMC



- What does this robot do?
- Same categories:
 - Perception
 - Modeling
 - Behavior
 - Motions
 - Control

Perceptions



From image processing:

- household objects
- categorization (kitchen items, food items...)
- maps
- environment segmentation

Other:

- current joint angles
- battery state
- IMU
- ...

Motion & control

- Motion:
 - Moving using wheels, pick-and-place, interact with humans...
 - Set an angle for each joint.
 - Calculate trajectories, inverse kinematics, balancing, ...
 - Execute static angle sequences.
- Control
 - Move joints to the target positions.

Acknowledgement

Acknowledgement

Some of the slides for this course have been prepared by Andreas Seekircher. I took some material from the ASL Lab at ETH Zürich <https://asl.ethz.ch/>.