Artificial Cognitive Systems

Module 8: Knowledge and Representation

Lecture 3: Acquiring and sharing knowledge

David Vernon
Carnegie Mellon University Africa

www.vernon.eu

Acquiring and Sharing Knowledge

- How cognitivist agents — robots in particular — share knowledge
  - Directly with humans
  - Directly with other robots
- How emergent agents — again, cognitive robots — acquire knowledge and know-how
  - By observing other agents, i.e. by learning from demonstration

Acquiring and Sharing Knowledge

Direct Knowledge Transfer — the Cognitivist Perspective

- Strength and weakness of the cognitivist paradigm:

  all cognitive agents have access to the same understanding of the world around them

Acquiring and Sharing Knowledge

Direct Knowledge Transfer — the Cognitivist Perspective

- **Strength**: the knowledge possessed by one agent is inherently compatible with the understanding mechanisms of another agent (perception, reasoning, and communication
- This strength has been widely exploited by all cognitivist systems and most hybrid systems to allow humans to embed representations of their knowledge in cognitive systems, very often as symbol-based rules.

Acquiring and Sharing Knowledge

Direct Knowledge Transfer — the Cognitivist Perspective

- **Weakness**: a programmer's knowledge is limited to what he or she knows and considers important enough to encode and embed in the cognitive agent

Acquiring and Sharing Knowledge

Direct Knowledge Transfer — the Cognitivist Perspective

- Recent advance: autonomously mine other sources of knowledge
- Thus, as one robot learns how to solve a given task, it can make that knowledge available to other robots

Acquiring and Sharing Knowledge

Direct Knowledge Transfer — the Cognitivist Perspective

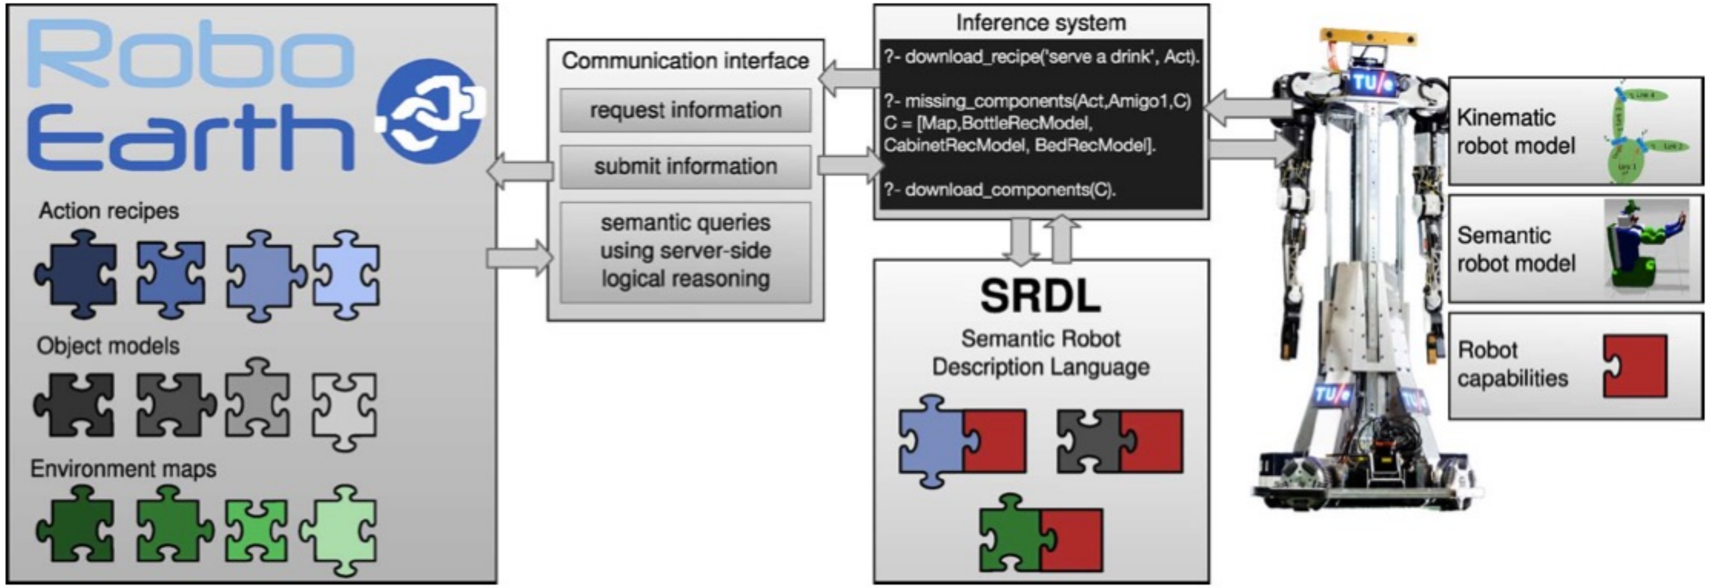

Tenorth et al. 2013

Acquiring and Sharing Knowledge

Direct Knowledge Transfer — the Cognitivist Perspective

- RoboEarth, a fast-growing world-wide open-source framework that allows robots to generate, share, and reuse knowledge and data

  - Global world models of objects, environments, actions, all linked to semantic information

  - Different representations are used

    - images
    - point clouds
    - 3-D models for object
    - maps and coordinates for environments
    - human-readable action recipes for actions and task descriptions

Acquiring and Sharing Knowledge

Direct Knowledge Transfer — the Cognitivist Perspective

- RoboEarth is effectively a world wide web for robots:
  - exchange knowledge among each other
  - benefit from the experience of other robots
  - customizing that knowledge to suit their own particular circumstances
- It is this ability to customize that particularly distinguishes RoboEarth

Acquiring and Sharing Knowledge

Direct Knowledge Transfer — the Cognitivist Perspective

- It is not just a repository of object, environment, and task data:
  - it includes also the semantic knowledge that encodes the meaning of the content in terms of the relationships between the various entities
  - in a way that allows the robot to decide if an object model will be useful in some given task

Acquiring and Sharing Knowledge

Direct Knowledge Transfer — the Cognitivist Perspective

- Cognitive robots might also use normal WWW resources (for humans)

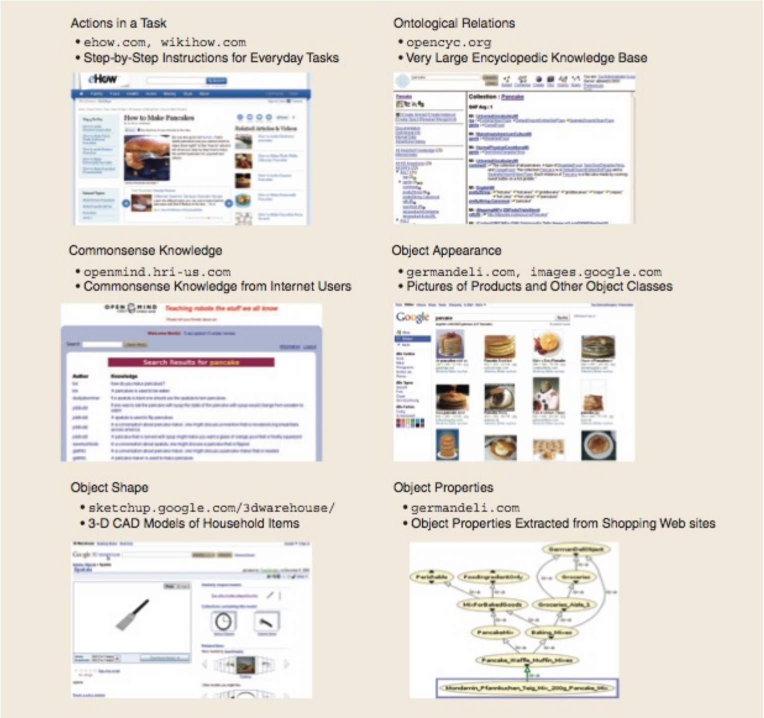

Tenorth et al. 2011

Acquiring and Sharing Knowledge

Direct Knowledge Transfer — the Cognitivist Perspective

Benefits of "Cloud Robotics"

- Provides a shared knowledge database
  - Organizes and unifies information about the world in a format usable by robots
- Offloads heavy computing tasks to the cloud
  - Cheaper, lighter, easier-to-maintain hardware (akin to desktop PC vs. a thin-client "netbook")
  - Longer battery life
  - Less need for software pushes/updates
  - CPU hardware upgrades are invisible & hassle-free
- Skill / Behavior Database
  - reusable library of "skills" or behaviors that map to perceived task requirements / complex situations.
  - Data-mining the history of all cloud-enabled robots

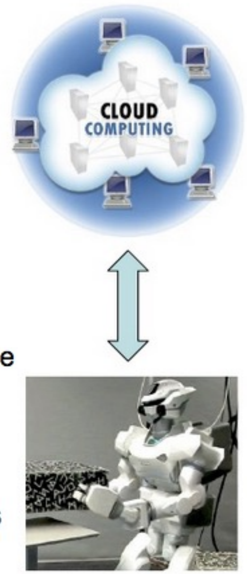

James Kuffner (IEEE Conference on Humanoid Robotics in 2010)
http://www.scribd.com/doc/47486324/Cloud-Enabled-Robots

Acquiring and Sharing Knowledge

Cloud Robotics and Automation

What if robots and automation systems were not limited by onboard computation, memory, or software? Rather than viewing robots and automated machines as isolated systems with limited computation and memory, "Cloud Robotics and Automation" considers a new paradigm where robots and automation systems exchange data and perform computation via networks. Extending earlier work that links robots to the Internet, Cloud Robotics and Automation builds on emerging research in cloud computing, machine learning, big data, open-source software, and major industry initiatives in the "Internet of Things", "Smarter Planet", "Industrial Internet", and "Industry 4.0."

http://goldberg.berkeley.edu/cloud-robotics/

Acquiring and Sharing Knowledge

Direct Knowledge Transfer — the Cognitivist Perspective

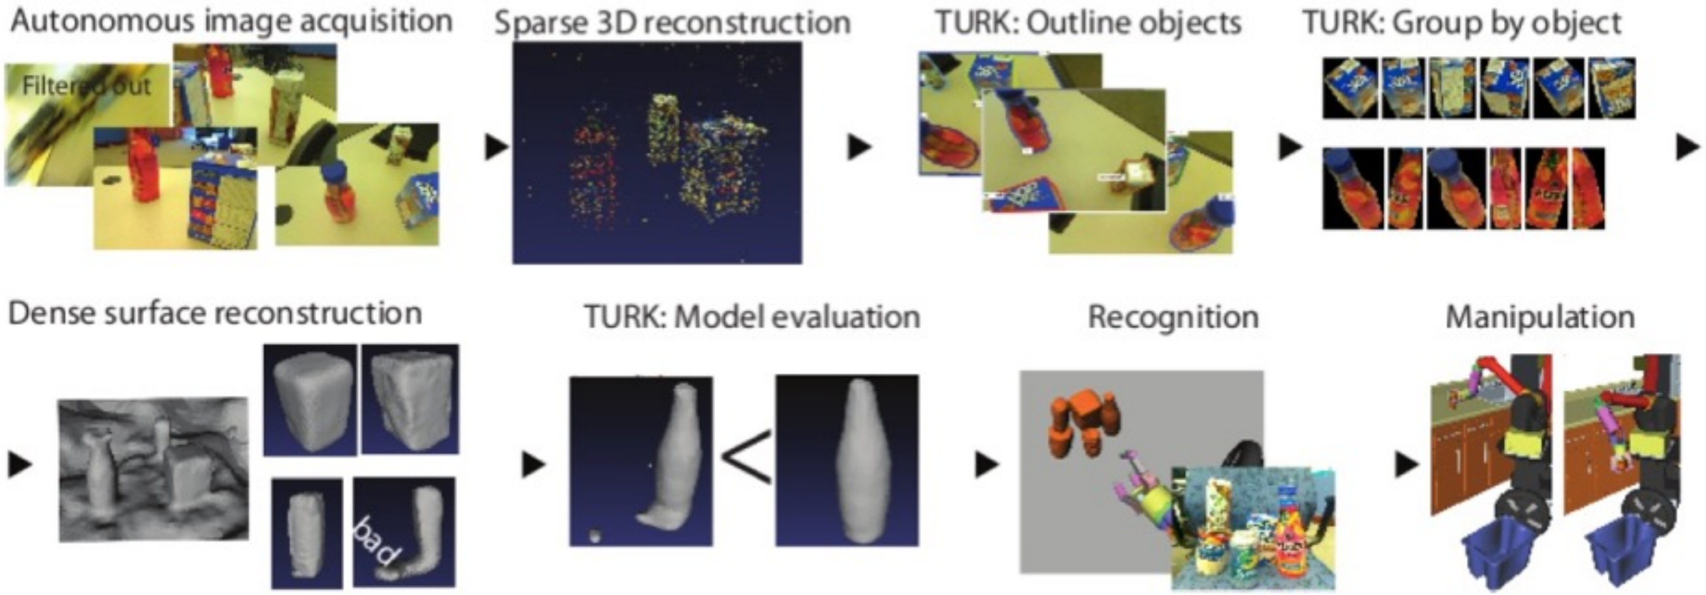

A cloud robot system that incorporates Amazon's Mechanical Turk to "crowdsource" object identification to facilitate robot grasping. Goldberg and Kehoe 2013

Acquiring and Sharing Knowledge

Direct Knowledge Transfer — the Cognitivist Perspective

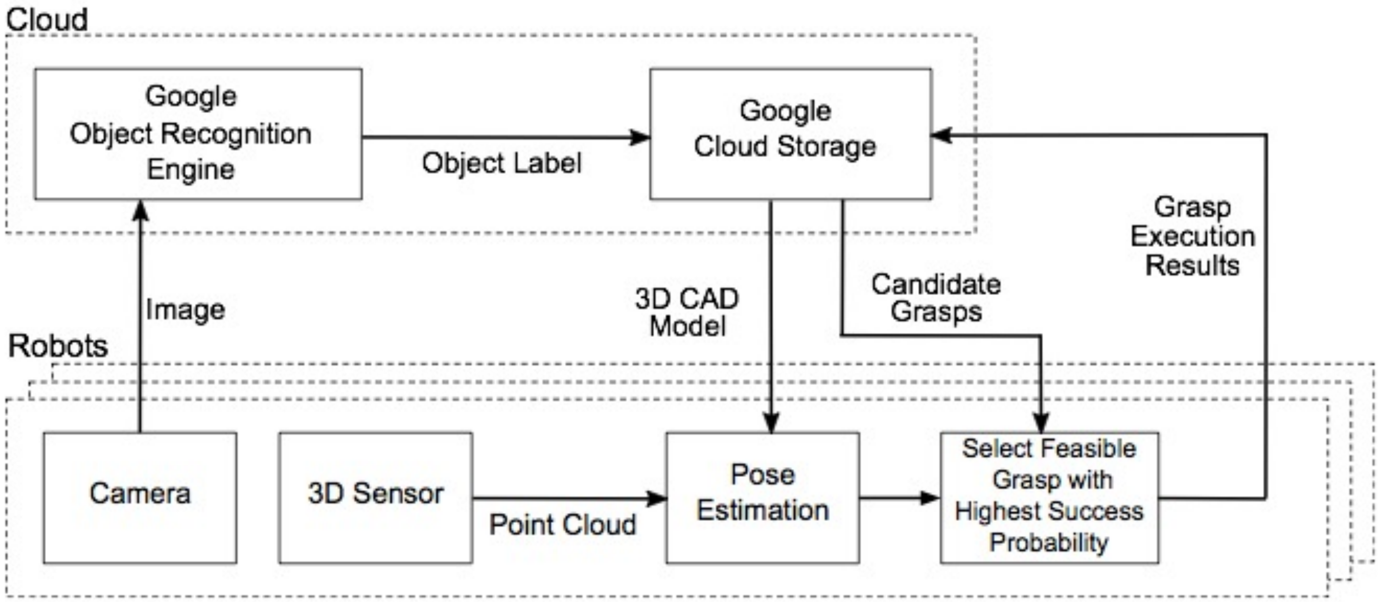

Goldberg and Kehoe 2013

Acquiring and Sharing Knowledge

Learning from Demonstration — the Emergent Perspective

- We do not have direct access to the knowledge of a cognitive system
- Such knowledge is specific to that particular agent
- It is acquired through development and learning

Acquiring and Sharing Knowledge

Learning from Demonstration — the Emergent Perspective

– This approach to imparting knowledge to a cognitive robot

  especially know-how in the form of procedural and episodic knowledge

  is referred as

  - Learning from Demonstration
  - Programming by Demonstration

Acquiring and Sharing Knowledge

Learning from Demonstration — the Emergent Perspective

- Learning from demonstration is relevant to
  - Cognitivist
  - Emergent traditions
- There are four variants of the general approach
- Just one of these is applicable to emergent cognitive robotics

Acquiring and Sharing Knowledge

Learning from Demonstration

- A cognitive robot needs to know what action is appropriate given the current state of its world

- The goal of the exercise is to learn a mapping from world state to action: **an action policy**

Acquiring and Sharing Knowledge

Learning from Demonstration

- To learn this action policy, a teacher provides a series of examples of the world state and the associated actions
- Often, the world state isn't directly accessible to the robot, so it has to depend on its observations
- The learned action policy then tells the robot what action to take based on its observations

Acquiring and Sharing Knowledge

Learning from Demonstration

- During learning, knowledge is transferred to the robot
- That transfer may be explicit or implicit
  - Explicit if the robot has direct access to the teacher's knowledge
    - the motion and world state information that characterize the actions being taught
  - Implicit if the robot has to infer it

Acquiring and Sharing Knowledge

Learning from Demonstration

There are two steps to this transfer:

1. Recording of the teacher's actions

2. Determining the correspondence of this data with the robot's embodiment

Acquiring and Sharing Knowledge

Learning from Demonstration

- Either step can be direct or indirect

Direct

The data can be used by the robot without modification

Indirect

The data has to be mapped to the robot

i.e. the information required by the robot to reproduce the action has to be inferred

Acquiring and Sharing Knowledge

Learning from Demonstration

- Since there are two stages
  - Recording teacher actions
  - Determining correspondence with robot embodiment
- Each with two options
  - Directly used
  - Indirectly mapped
- There are four different types of learning from demonstration

Acquiring and Sharing Knowledge

Learning from Demonstration

- There are four different types of learning from demonstration

    1. Teleoperation
    2. Shadowing
    3. Sensors on Teacher
    4. External Observation

Acquiring and Sharing Knowledge

Learning from Demonstration

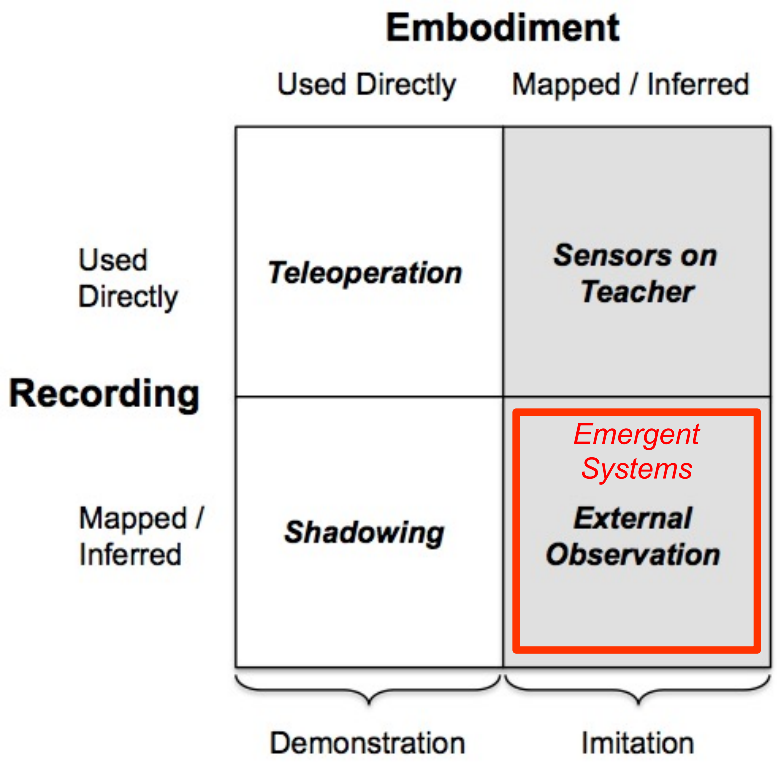

In the top row, the required teaching data is recorded and can be used directly by the robot;

In the bottom row, the required data is available indirectly and so has to be inferred by the robot.

In the left column the teaching data matches the robot's body and can be used directly.

In the right column the teacher's body doesn't match the robot's body and therefore the teaching data has to be mapped to the robot's body.

[Argall et al. 2009]

Acquiring and Sharing Knowledge

Example: human directly controls mobile robot and robot records the control data

Example: human uses gestures to teach a **mobile** robot and human gesture data is recorded for analysis by robot

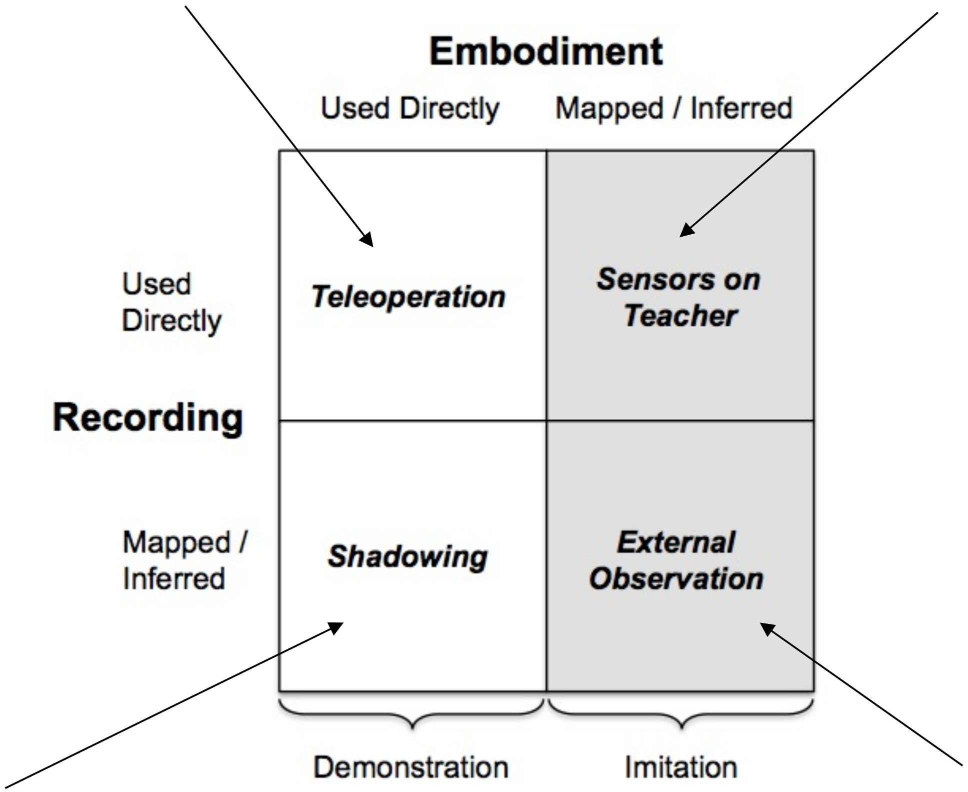

Example: human uses gestures to teach a dexterous **humanoid** robot and robot observes human, interpreting movement

Example: human uses gestures to teach a **mobile** robot and robot observes human, interpreting movement

Acquiring and Sharing Knowledge

Learning from Demonstration

- Teleoperation
  - A teacher directly operates the robot learner
  - The robot's sensors records the required actions
  - The recording is direct
  - The teacher and learner embodiments correspond exactly (they are in fact the same body, i.e. the robot), no further mapping is required

Kristoffer Öfjäll and Michael Felsberg. Biologically inspired online learning of visual autonomous driving. In Proceedings of the British Machine Vision Conference. BMVA Press, 2014.

Acquiring and Sharing Knowledge

Learning from Demonstration

- Shadowing
  - The robot doesn’t have direct access to the teachers action data
  - Instead, it attempts to mimic teachers actions (e.g. by tracking the teacher’s action) and it records its own movements through its own sensors
  - Thus, the recording is indirect
  - However, the data recorded corresponds directly to the robot’s own embodiment

Acquiring and Sharing Knowledge

Learning from Demonstration

- Sensors on teacher
  - The teacher's movements are recorded directly on the teacher and the learning robot has direct access to this data
  - Thus, the recording is explicit and direct
  - However, the embodiment of teacher and learner robot are different so this data has to be used indirectly and mapped to the particular embodiment of the learner robot

Acquiring and Sharing Knowledge

Learning from Demonstration

- External observation

  - Both the recording and the embodiment correspondence have to be inferred
  - The robot observes the teacher with sensors that are external to the teacher (in contrast to the sensor on teacher type)
  - These could be on the robot itself or arranged around the robot in a sensorized environment
  - Furthermore, the teacher and robot embodiments differ so that the observed data, i.e. the teacher's movements, has to be mapped to the robot's frame of reference

Acquiring and Sharing Knowledge

Learning from Demonstration

- The term demonstration is used to refer to the first two types:

  Teleoperation

  Shadowing

- Imitation refers to either the third or fourth type

  Sensors on teacher

  External Observation

Acquiring and Sharing Knowledge

Learning from Demonstration

- For emergent cognitive systems, only the fourth type — external observation imitation — is applicable
- Why?
- Direct knowledge transfer (in either steps of the knowledge transfer process) is not possible

Acquiring and Sharing Knowledge

Learning from Demonstration

- There are two options for which the demonstration — the pairs of world states or observations of world states and associated actions — is presented to the learning robot

  - Batch learning
    - The action policy is learned after all the data has been presented
  - Interactive learning
    - The action policy is up dated as the training data is progressively presented to the learning robot
    - The teacher can correct the learned action policy or to suggest ways of improving it, i.e. the teacher can also coach the robot

Recommended Reading

Vernon, D. Artificial Cognitive Systems – A Primer, MIT Press, 2014; Chapter 8.