Robotics

Applied artificial intelligence (EDA132)
Lecture 16
2012-05-10
Elin A. Topp

Course book (chapter 25), images & movies from various sources, and original material

Outline

- Robots and Robotics
  - Types of robots
  - Robotics
    - Kinematics and dynamics
    - Systems (hard- and software, components)
    - Challenges (and results)
- AI in robotics
  - Mapping & Localisation
  - (Path) Planning
  - Deliberation & High level decision making and planning
- Human-Robot Interaction
  - Technical aspects
  - Cognitive / social aspects

Outline

- Robots and Robotics
  - Types of robots
  - Robotics
    - Kinematics and dynamics
    - Systems (hard- and software, components)
    - Challenges (and results)
- AI in robotics
  - Mapping & Localisation
  - (Path) Planning
  - Deliberation & High level decision making and planning
- Human-Robot Interaction
  - Technical aspects
  - Cognitive / social aspects

What is a "Robot"?

✓

✓

✓

✓

✓

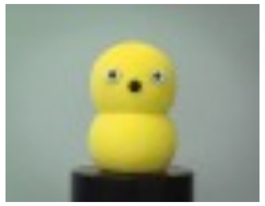

?

?

✓

• • •

Types of robots

Industrial robots vs. service robots vs. personal robots / robot toys

Static manipulators vs. mobile platforms (vs. mobile manipulators)

Mechanistic vs. humanoid / bio-inspired / creature-like

For all in common:

A robot is a physical agent in the physical world
(with all the consequences that might have... ;-)

Robot actuators - joints and wheels

6 DOF (6 "joint") arm:

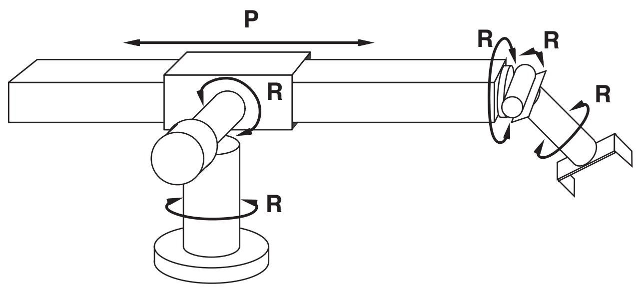

2 (3 effective) DOF synchro drive (car):

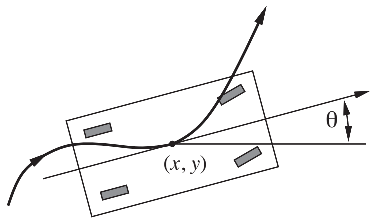

2 (3 effective) DOF differential drive (Pioneer p3dx):

3 DOF holonomic drive ("shopping cart", DLR's Justin):

Robotics - making robots do their job

Kinematics (relatively simple):

Where do I get with a certain configuration of parts / wheel movement?

Inverse kinematics (less simple, but more interesting):

How do I have to control joints and wheels to reach a certain point?

How fast, how heavy, when to brake?

Dynamics:

Make the robot move (and move stuff) without falling apart, or crashing into things

- How much payload is possible?
- How fast can I move without tipping over?
- What is my braking distance?
- How do I move smoothly? (ask the automated control people ;-)

Weight: ca 1300 kg

Payload: ca 150 kg

Video removed.
This was original research material and
can not be published due to privacy agreements
with the study participants.
Contact me (Elin) if you want to learn more
about the study.

Dynamics vs. Kinematics

Dynamics also gets you into two problems: direct and inverse dynamics.

Direct dynamics:

Given masses, external forces, position, velocities and acceleration in the joints / wheels, what forces / moments are put to the depending joints and the tool centre point (TCP)? "Rather" simply solvable, at least more or less straight forward.

Inverse dynamics (again, more interesting than direct dynamics):

While solving the *inverse kinematics* problem is nasty, but still "only" a bunch of linear equations, solving the *inverse dynamics* problem leaves you with a bunch of more or less complex differential equations.

Supporting parts: Sensors

In a predictable world, we do not need perception, but good planning and programming

As the world is somewhat unpredictable, some perception is useful, i.e., robots / robot installations need sensors.

Passive / active sensors.

Range / colour / intensity / force / direction ...

Optical / sound / radar / smell / touch ...

Most common for mobile robots: position (encoders / GPS), range (ultrasound or laser range finder), image (colour/intensity), sound

Most common for manipulators: position (encoders), force / torque, images, (range - infrared, laser RF)

Sensors on a mobile robot

Microphones (sound)

Ultrasound (24 emitters / receivers) (range)

Camera (image - colour / intensity)

Laser range finder (SICK LMS 200) (range)

Infrared (range / interruption)

Bumpers (touch)

Wheel encoders (position / pose)

System integration

Make all those components work together

Architectures, “operating systems”, controllers, programming tools ...

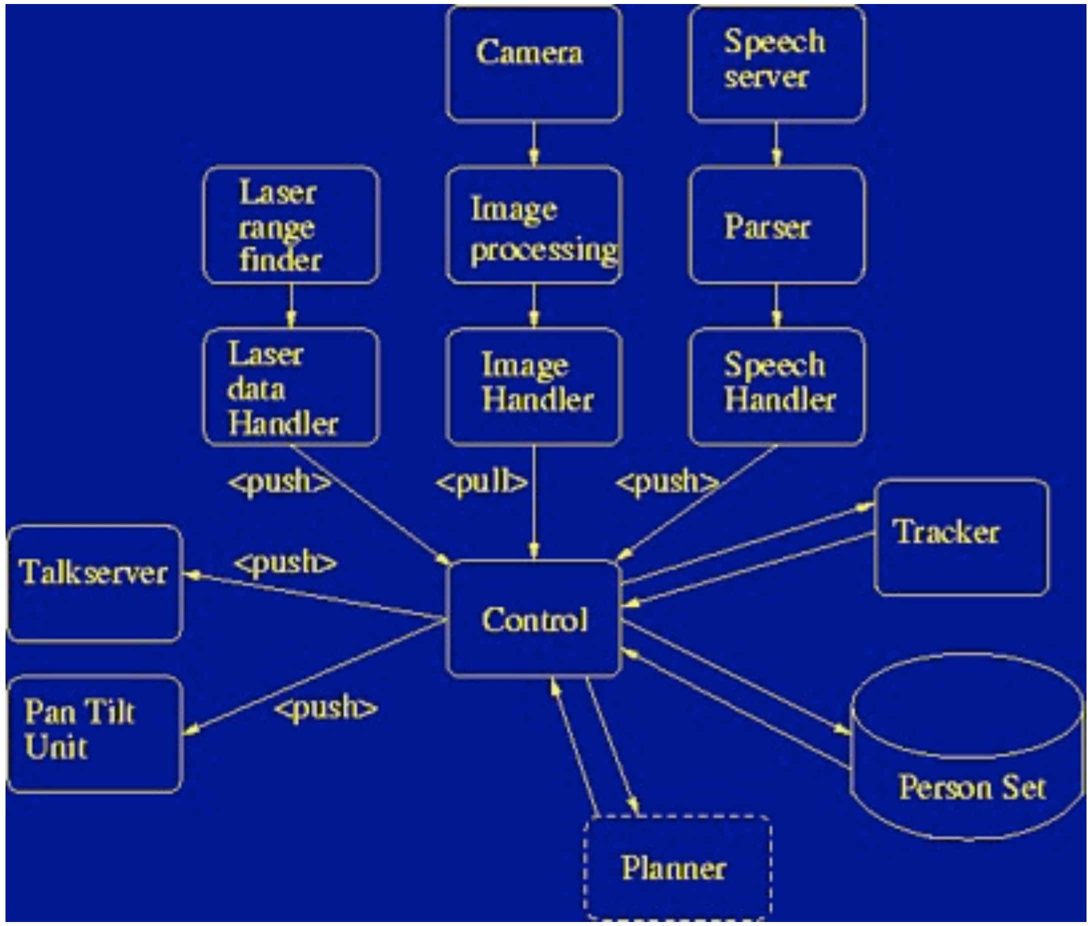

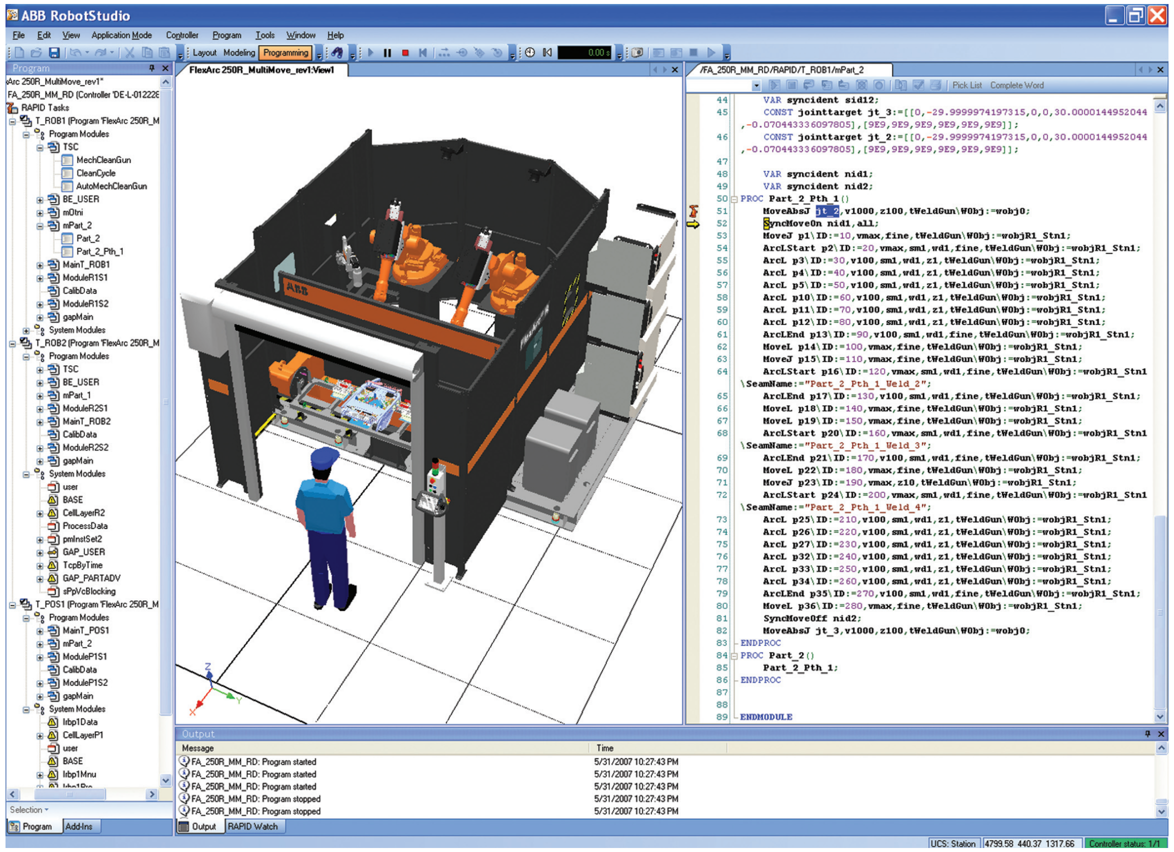

Behaviour based system architectures

Behaviour based system architectures

from sense-react (Brooks: "Planning is just a way of avoiding figuring out what to do next", 1987)

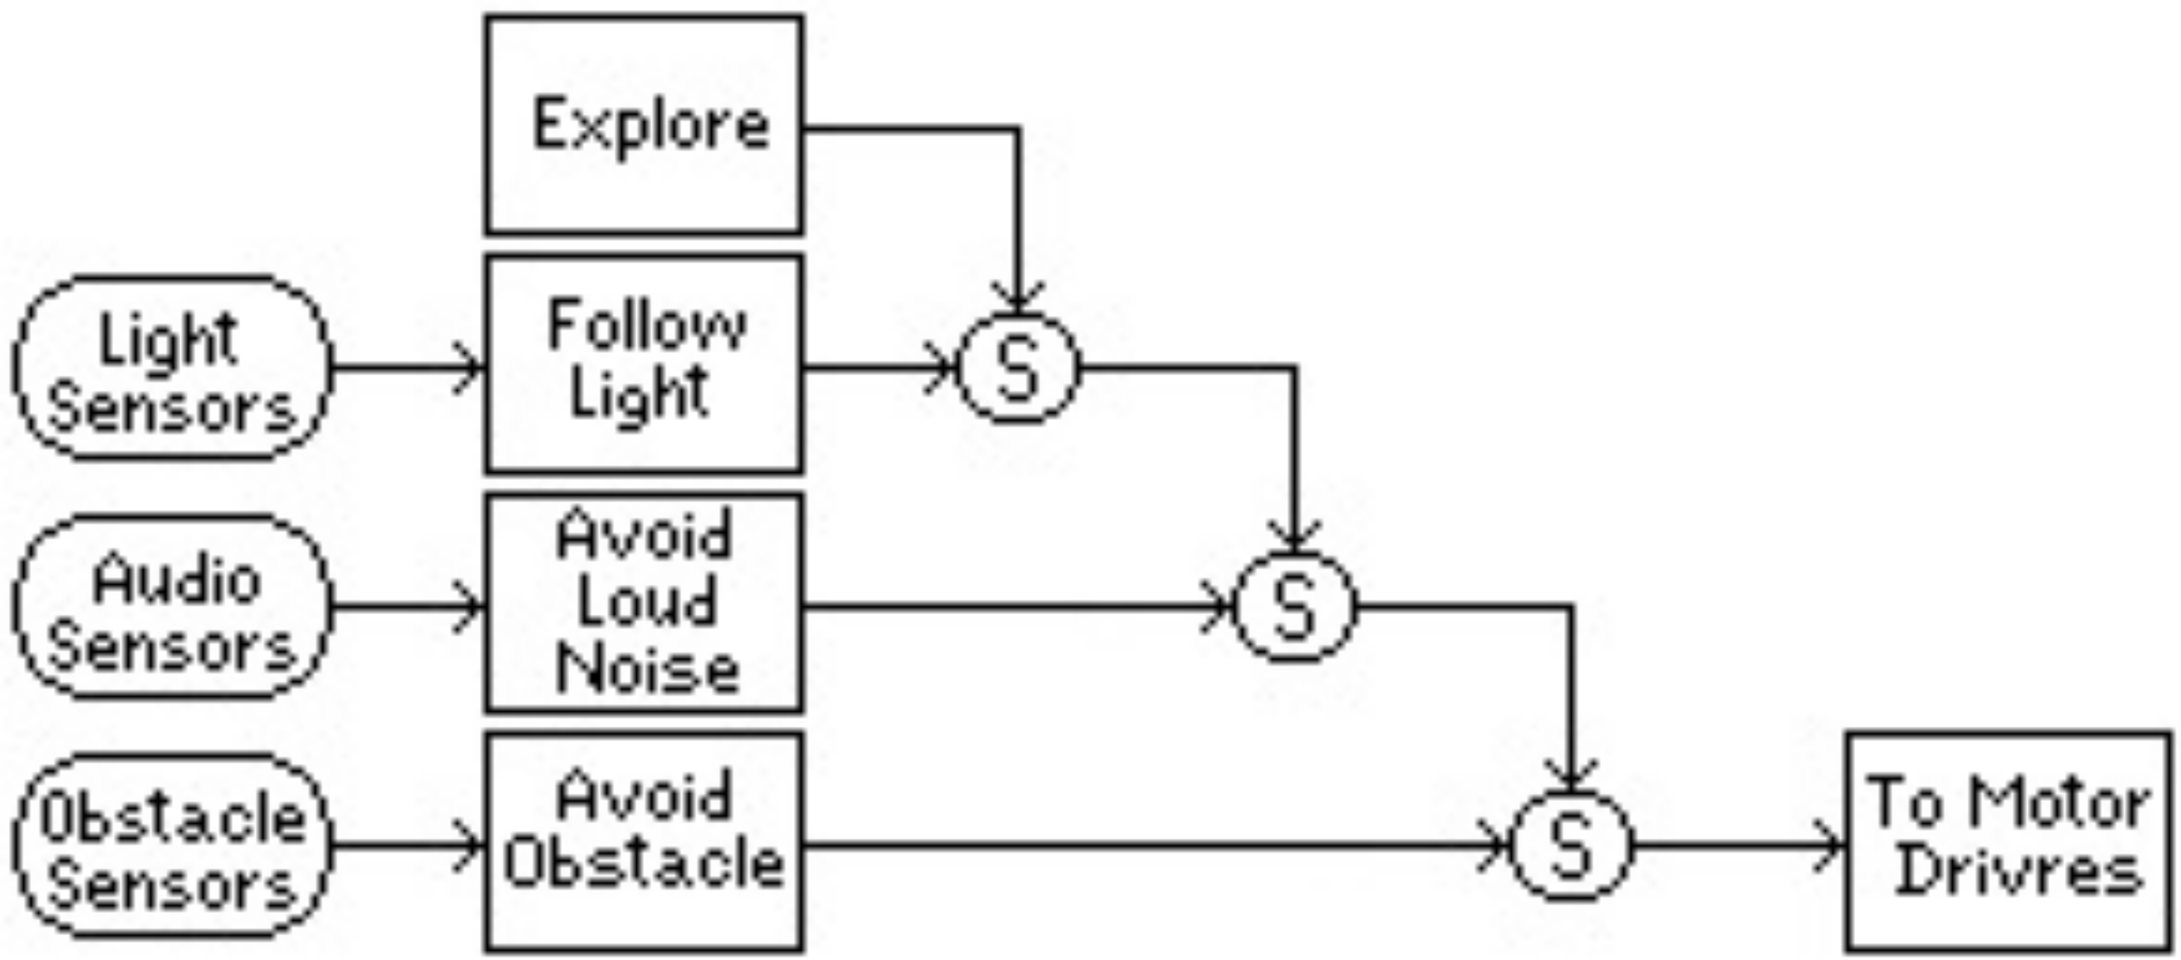

Behaviour based system architectures

from sense-react (Brooks: "Planning is just a way of avoiding figuring out what to do next", 1987)

via hybrid-deliberative (e.g., Arkin's "AuRA", 1990) and event-based systems

Behaviour based system architectures

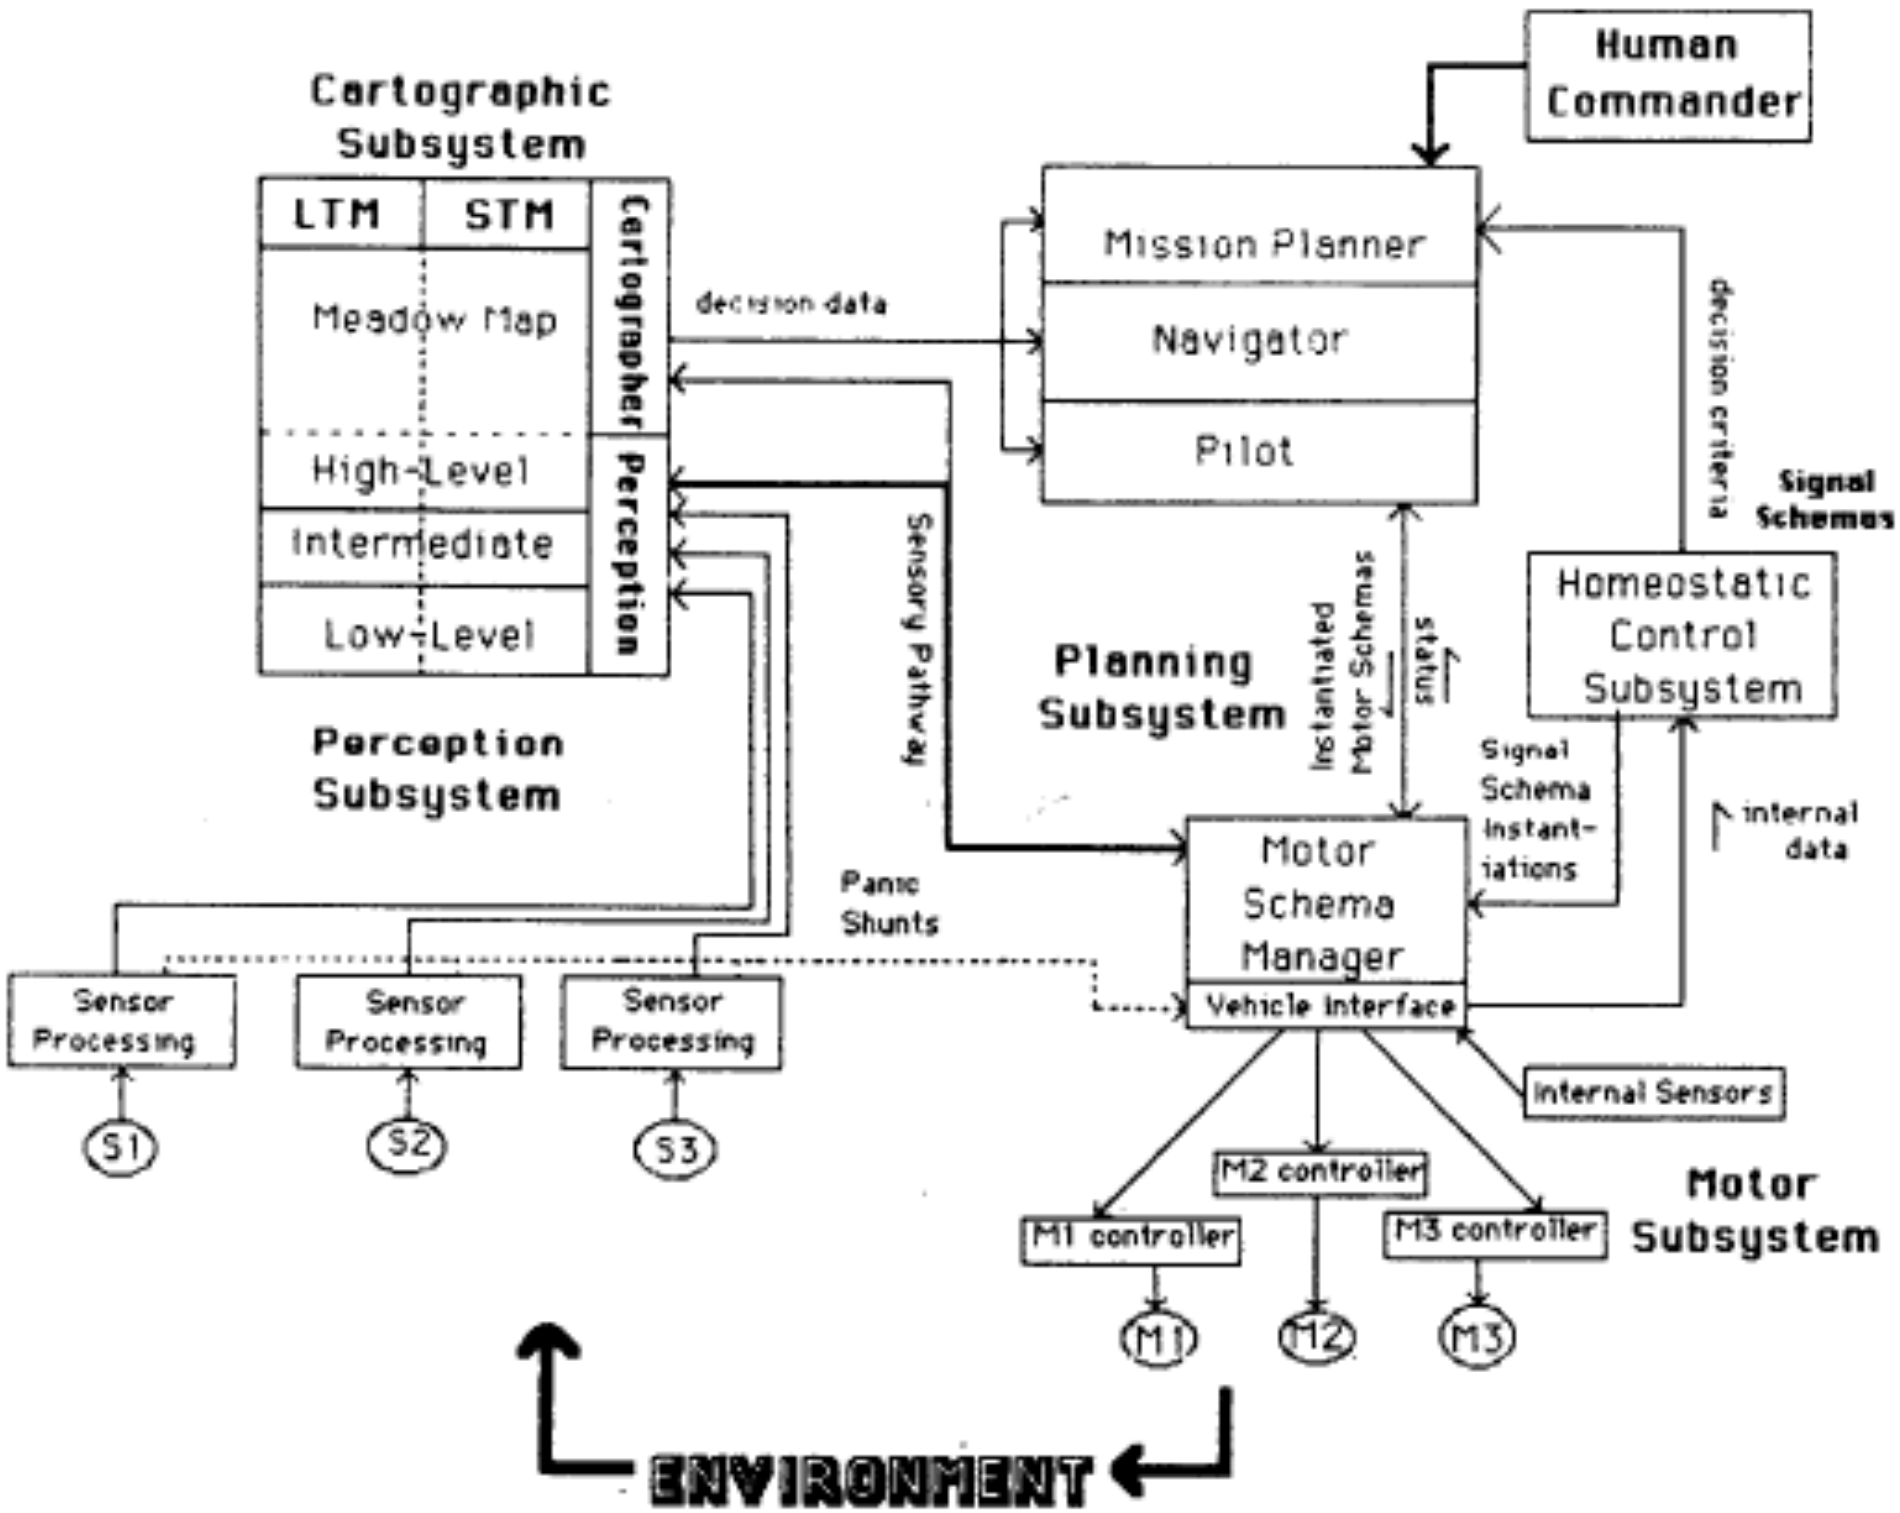

Behaviour based system architectures

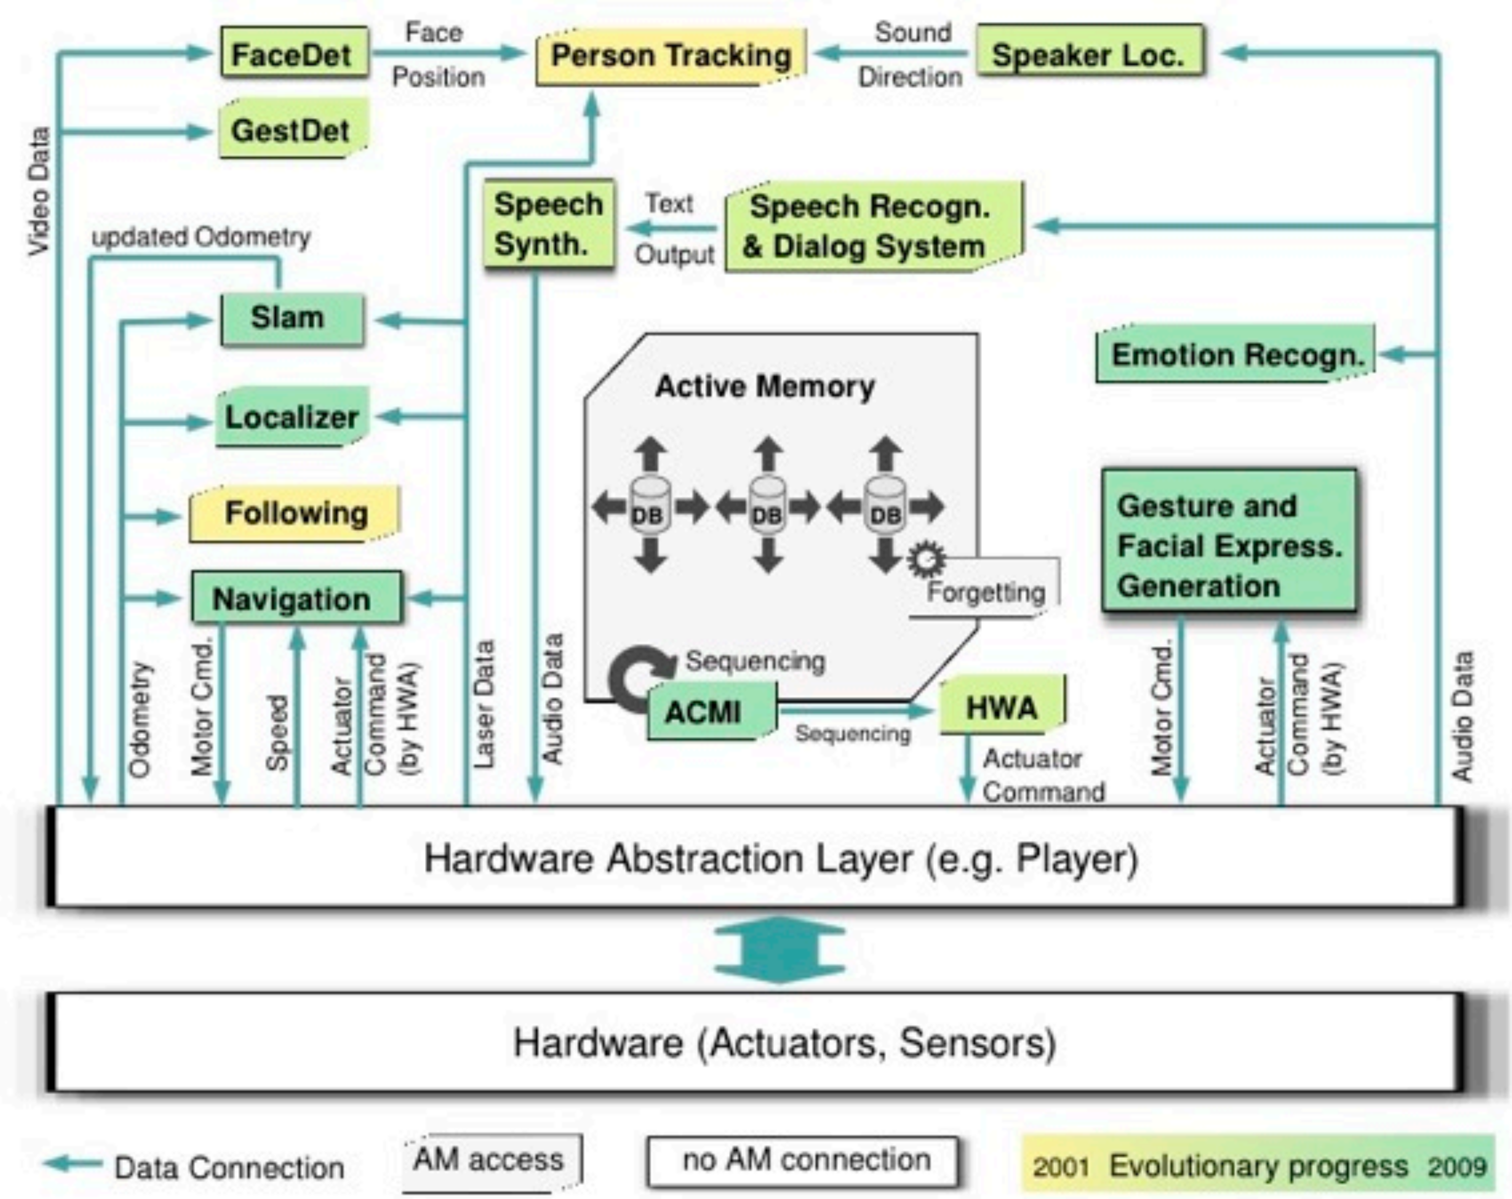

to "cognitive" architectures (memory & event based, e.g., T.P. Spexard, 2009)

Do the right thing at the right time...

Make industrial robots more flexible, interactive, easy to program (get some of the "behaviour"- and cognition idea into them)

Make mobile service robots more precise, go from research code to applications!

How far have we come?

Videos removed. Can be found at:

http://www.youtube.com/watch?v=fBtZ6EA2fpI

http://www.youtube.com/watch?v=R6pPwP3s7s4

http://www.youtube.com/watch?v=SOESSCXGhFo

Do the right thing at the right time...

Make industrial robots more flexible, interactive, easy to program (get some of the "behaviour"- and cognition idea into them)

Make mobile service robots more precise, go from research code to applications!

How far have we come?

DARPA Urban Challenge 2007

(almost there ;-)

Videos removed. Can be found at:

http://www.youtube.com/watch?v=fBtZ6EA2fpI

http://www.youtube.com/watch?v=R6pPwP3s7s4

http://www.youtube.com/watch?v=SOESSCXGhFo

Do the right thing at the right time...

Make industrial robots more flexible, interactive, easy to program (get some of the "behaviour"- and cognition idea into them)

Make mobile service robots more precise, go from research code to applications!

How far have we come?

Videos removed. Can be found at:

http://www.youtube.com/watch?v=fBtZ6EA2fpI

ABB robots and their precision... 2009

http://www.youtube.com/watch?v=R6pPwP3s7s4

http://www.youtube.com/watch?v=SOESSCXGhFo

Do the right thing at the right time...

Make industrial robots more flexible, interactive, easy to program (get some of the "behaviour"- and cognition idea into them)

Make mobile service robots more precise, go from research code to applications!

How far have we come?

Videos removed. Can be found at:

http://www.youtube.com/watch?v=fBtZ6EA2fpI

http://www.youtube.com/watch?v=R6pPwP3s7s4

http://www.youtube.com/watch?v=SOESSCXGhFo

DLR's Justin catches balls... 2011

Outline

- Robots and Robotics
  - Types of robots
  - Robotics
    - Kinematics and dynamics
    - Systems (hard- and software, components)
    - Challenges (and results)
- AI in robotics
  - Mapping & Localisation
  - (Path) Planning
  - Deliberation & High level decision making and planning
- Human-Robot Interaction
  - Technical aspects
  - Cognitive / social aspects

Mapping

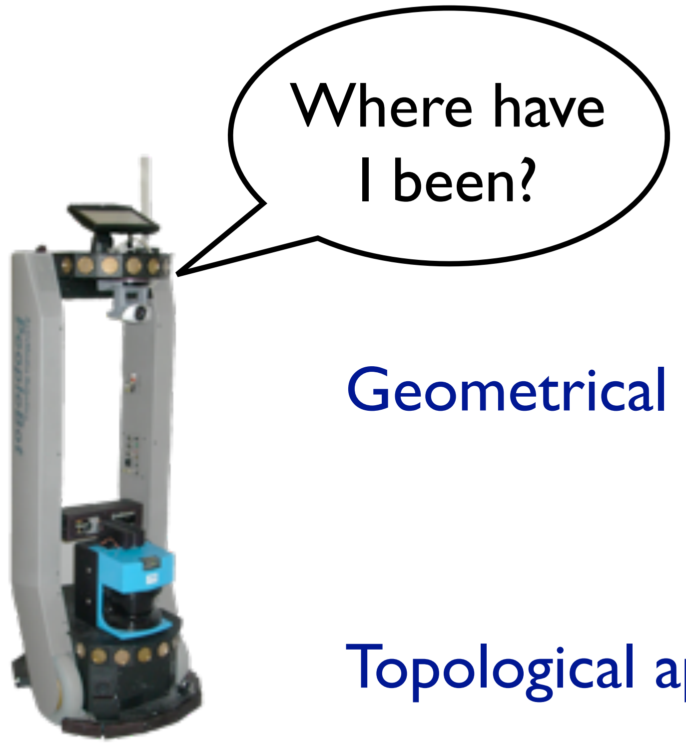

Geometrical approaches

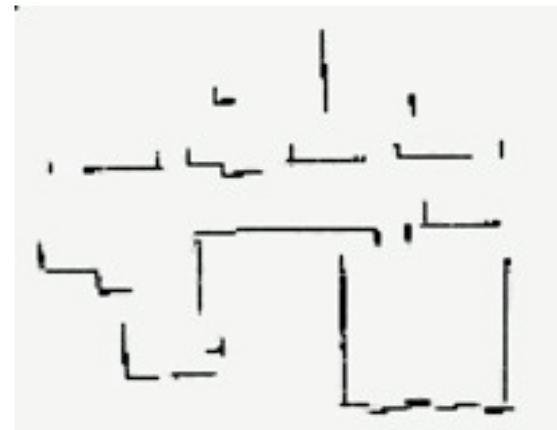

Topological approaches

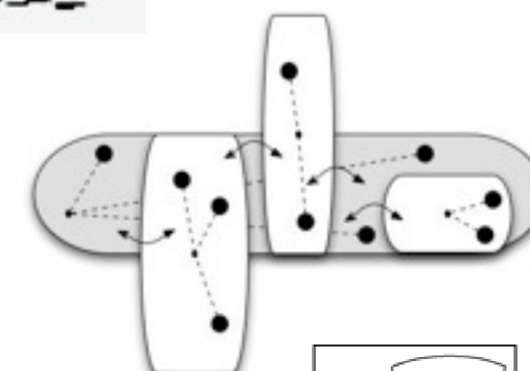

Occupancy grid approaches (e.g., Sebastian Thrun)

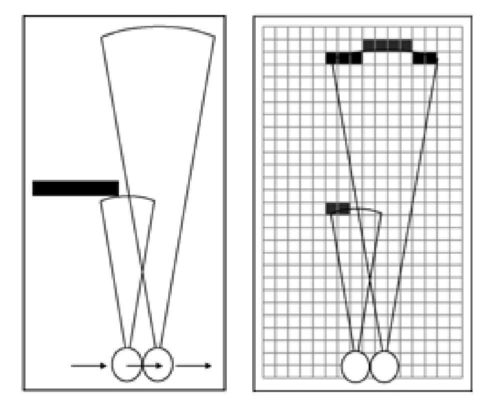

(Hybrid approaches)

Localisation

E.g., Monte Carlo Localisation (S. Thrun)

Video removed. Check Sebastian Thruns homepage,
if you are interested in MC-localisation:
http://robots.stanford.edu

Data filters for state estimation

0. Represent state, identify system function

1. Estimate / predict state from model applying the function

2. Take a measurement

3. Update state according to model and observation (measurement)

Used for position tracking, detection of significant changes in a data stream, localisation ...

E.g., particle filters (Monte Carlo), Kalman filters

Particle filter

1. Represent possible positions by samples (uniform distribution) $\boldsymbol{x} = (x, y, \theta)$

2. Estimate movement / update samples according to assumed robot movement + noise

3. Take a measurement z

4. Assign weights to samples according to posterior probabilities (Bayes!) $P(x_i \mid z)$

5. Resample (pick "good" samples, use those as new "seeds", redistribute in position space and add some noise), continue at 2.

Kalman filter

Represent posterior with a Gaussian.

Assume linear dynamical system
(F, G, H system matrices, u measurement, v, w, gaussian noise)

$$x(k+1) = F(k)\, x(k) + G(k)\, u(k) + v(k)$$ (state estimate)

$$y(k+1) = H(k)\, x(k) + w(k)$$ (output)

1. Predict based on last estimate:

$$x'(k+1 \mid k) = F(k)\, x'(k \mid k) + G(k)\, u(k) + v(k)$$

$$y'(k+1 \mid k) = H(k)\, x'(k+1 \mid k) + w(k)$$

2. Calculate correction based on prediction and current measurement:

$$\Delta x = f(y(k+1), x'(k+1 \mid k))$$

3. Update prediction:

$$x'(k+1 \mid k+1) = x'(k+1 \mid k) + \Delta x$$

Mapping & Localisation: Chicken & Egg?

Simultaneous localisation and mapping (SLAM)

While building the map, stay localised!

Use filters to “sort” landmarks:

- Known? Update your pose estimation!
- Unknown? Extend the map!

SLAM example

FastSLAM (D. Haehnel)

Video removed. Check Sebastian Thruns homepage, if you are interested in SLAM:
http://robots.stanford.edu

Path / trajectory planning

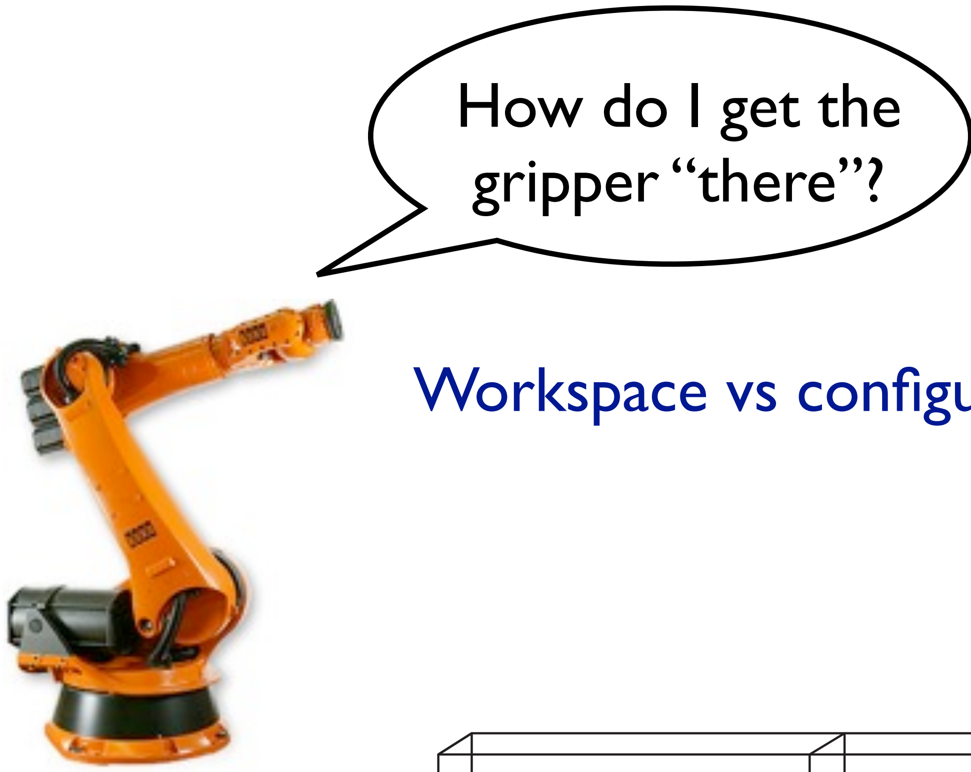

Workspace vs configuration space

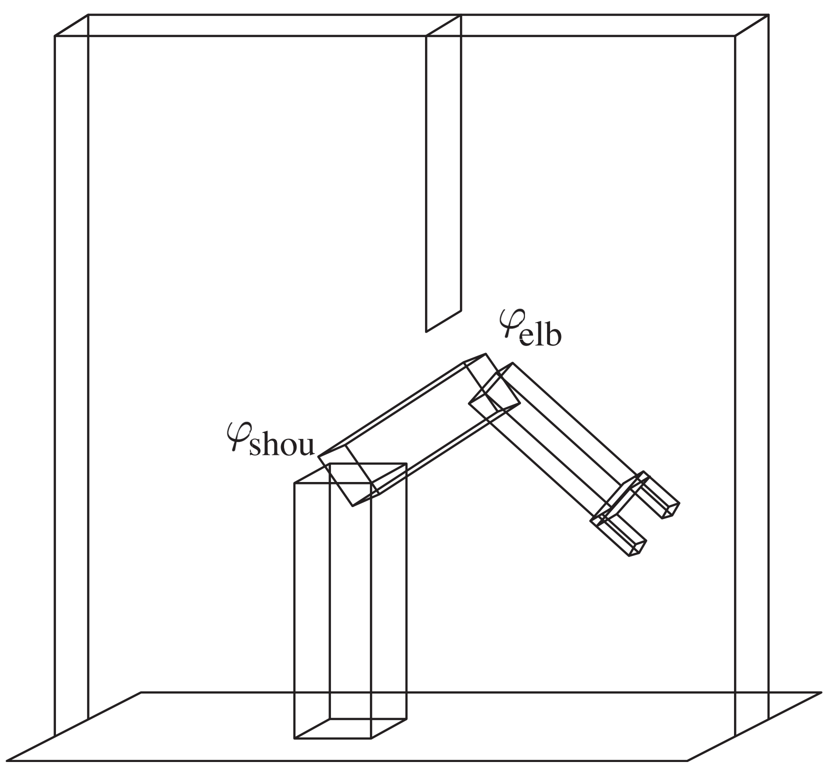

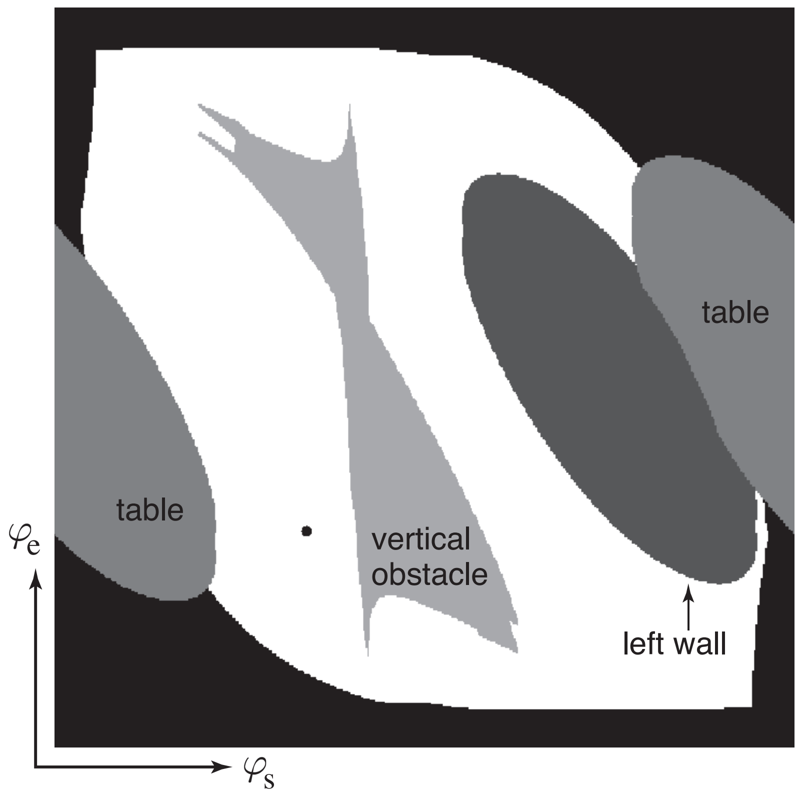

Cell decomposition

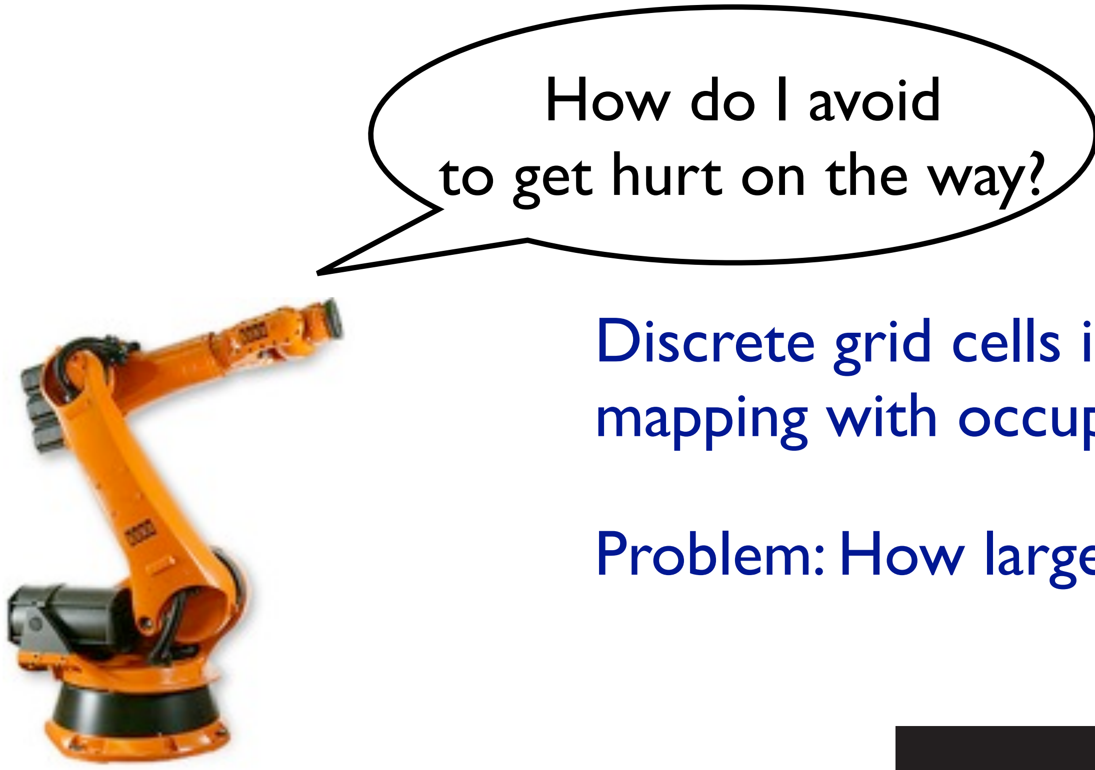

Discrete grid cells indicating "free" vs "occupied" (compare mapping with occupancy grids!)

Problem: How large should cells be?

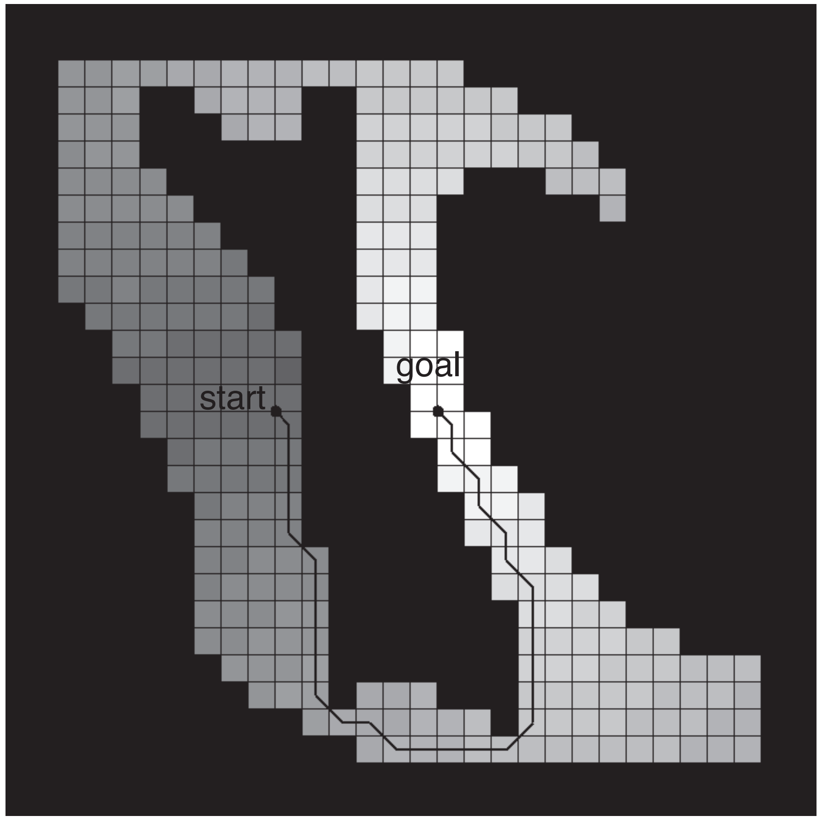

Potential fields

How do I avoid
to get hurt on the way?

Repelling potential field around obstacle keeps the TCP in free space.

Path is found by minimising path length and potential simultaneously

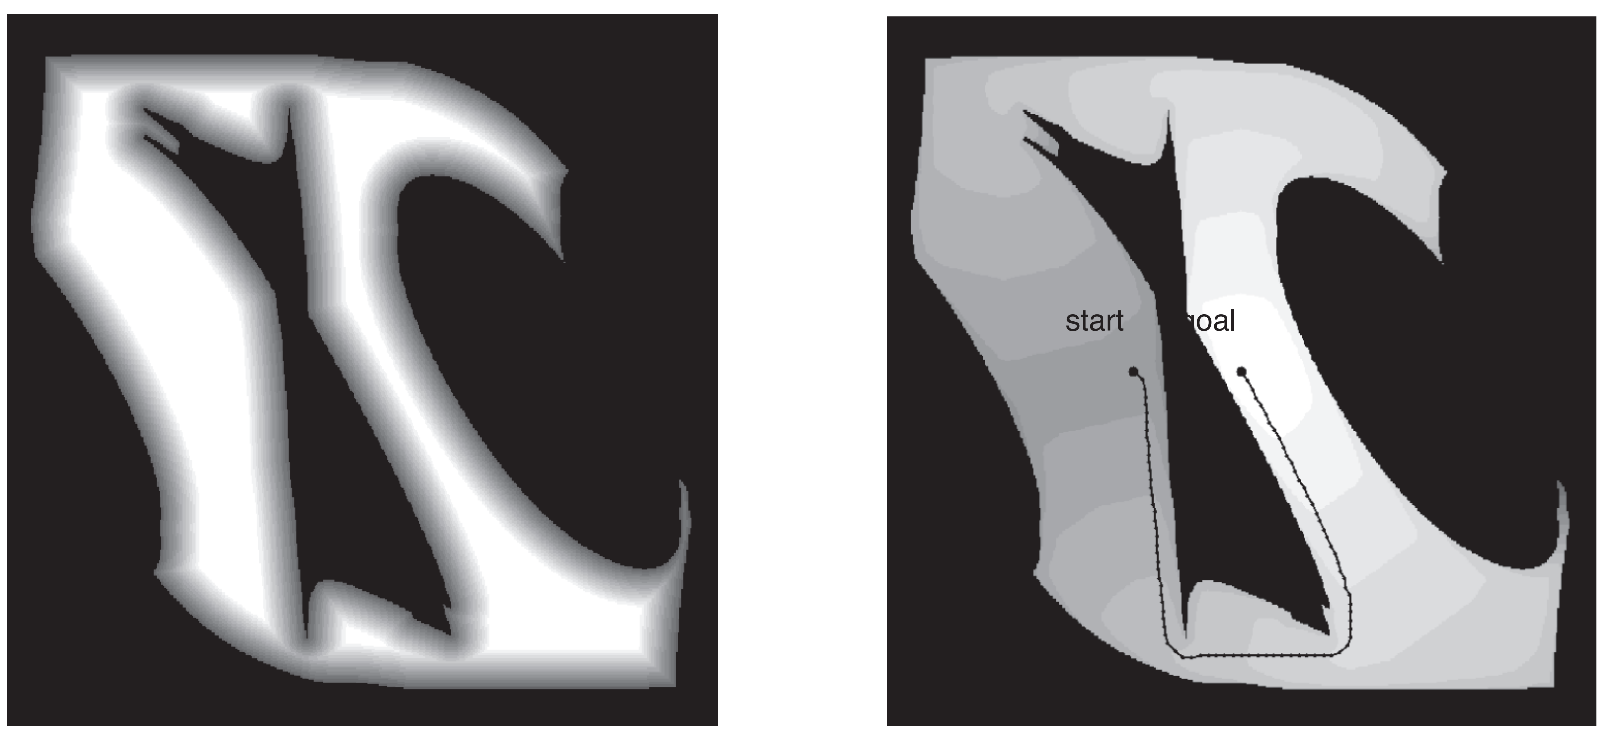

Voronoi graph

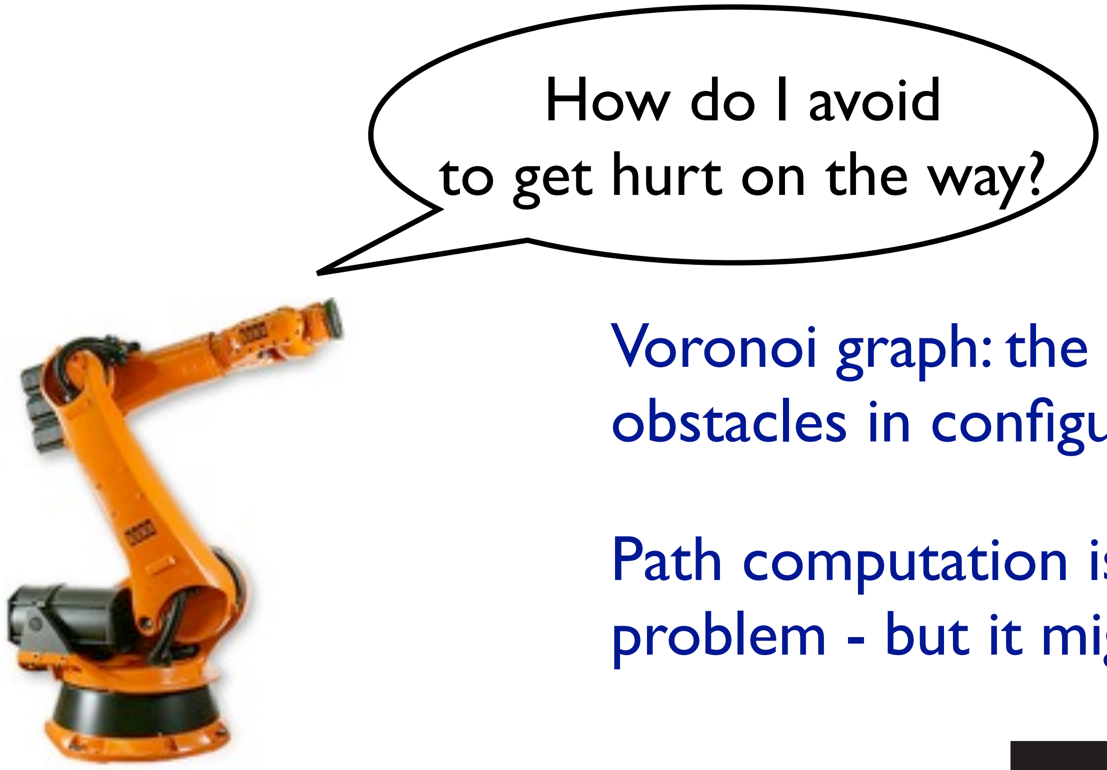

Voronoi graph: the set of points equidistant to two or more obstacles in configuration space

Path computation is boiled down to a simple graph-search problem - but it might not give the shortest path in space...

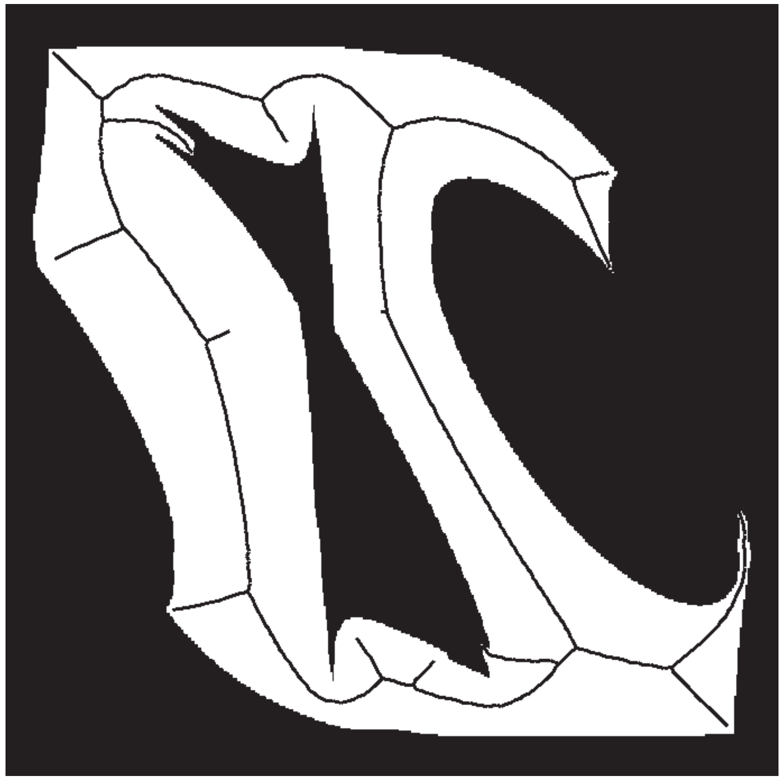

Planning movement under uncertainty?

Not knowing anything about the surroundings and simply following instructions might "hurt"

Apply "careful" exploration strategies and consider "emergency braking" (obstacle avoidance)

"Decide" on the fly, based on gathered information!

Deliberation in a navigation system

Deliberation in a navigation system

A robotic system might have several goals to pursue, e.g.,

Deliberation in a navigation system

A robotic system might have several goals to pursue, e.g.,

- Explore the environment (i.e., visit as many areas as possible and gather data) and build a map

Deliberation in a navigation system

A robotic system might have several goals to pursue, e.g.,

- Explore the environment (i.e., visit as many areas as possible and gather data) and build a map
- Use a certain strategy (e.g., follow the wall to the right)

Deliberation in a navigation system

A robotic system might have several goals to pursue, e.g.,

- Explore the environment (i.e., visit as many areas as possible and gather data) and build a map
- Use a certain strategy (e.g., follow the wall to the right)
- Do not bump into things or people on the way

Deliberation in a navigation system

A robotic system might have several goals to pursue, e.g.,

- Explore the environment (i.e., visit as many areas as possible and gather data) and build a map
- Use a certain strategy (e.g., follow the wall to the right)
- Do not bump into things or people on the way
- Go "home" for recharging in time

Deliberation in a navigation system

A robotic system might have several goals to pursue, e.g.,

- Explore the environment (i.e., visit as many areas as possible and gather data) and build a map
- Use a certain strategy (e.g., follow the wall to the right)
- Do not bump into things or people on the way
- Go "home" for recharging in time

Behaviours (e.g., as used by Arkin) can take care of each of the goals separately

Deliberation in a navigation system

A robotic system might have several goals to pursue, e.g.,

- Explore the environment (i.e., visit as many areas as possible and gather data) and build a map
- Use a certain strategy (e.g., follow the wall to the right)
- Do not bump into things or people on the way
- Go “home” for recharging in time

Behaviours (e.g., as used by Arkin) can take care of each of the goals separately

Particular perception results can be fed into a control unit for decision making

Deliberation in a navigation system

A robotic system might have several goals to pursue, e.g.,

- Explore the environment (i.e., visit as many areas as possible and gather data) and build a map
- Use a certain strategy (e.g., follow the wall to the right)
- Do not bump into things or people on the way
- Go "home" for recharging in time

Behaviours (e.g., as used by Arkin) can take care of each of the goals separately

Particular perception results can be fed into a control unit for decision making

This decision making unit (deliberation process) can assign weights (priorities) to the behaviours depending to the sensor data.

Deliberation in a navigation system

A robotic system might have several goals to pursue, e.g.,

- Explore the environment (i.e., visit as many areas as possible and gather data) and build a map
- Use a certain strategy (e.g., follow the wall to the right)
- Do not bump into things or people on the way
- Go "home" for recharging in time

Behaviours (e.g., as used by Arkin) can take care of each of the goals separately

Particular perception results can be fed into a control unit for decision making

This decision making unit (deliberation process) can assign weights (priorities) to the behaviours depending to the sensor data.

E.g., when battery level sensor reports a certain level, only the "going home" behaviour and immediate obstacle avoidance are allowed to produce control output, exploring and wall following are ignored.

More complex decisions / plans

More complex decisions / plans

If the system does not only involve one robot with several "competencies", but several robots with partly overlapping, partly complementary abilities, the decisions are to be taken to another dimension:

More complex decisions / plans

If the system does not only involve one robot with several “competencies”, but several robots with partly overlapping, partly complementary abilities, the decisions are to be taken to another dimension:

- Given a task, what do I need to know to fulfill it?

More complex decisions / plans

If the system does not only involve one robot with several “competencies”, but several robots with partly overlapping, partly complementary abilities, the decisions are to be taken to another dimension:

- Given a task, what do I need to know to fulfill it?
- Do I know these things?

More complex decisions / plans

If the system does not only involve one robot with several "competencies", but several robots with partly overlapping, partly complementary abilities, the decisions are to be taken to another dimension:

- Given a task, what do I need to know to fulfill it?
- Do I know these things?
- Given I know what to do, do I have the means (robot) to do it?

More complex decisions / plans

If the system does not only involve one robot with several "competencies", but several robots with partly overlapping, partly complementary abilities, the decisions are to be taken to another dimension:

- Given a task, what do I need to know to fulfill it?
- Do I know these things?
- Given I know what to do, do I have the means (robot) to do it?
- If yes, which one?

More complex decisions / plans

If the system does not only involve one robot with several "competencies", but several robots with partly overlapping, partly complementary abilities, the decisions are to be taken to another dimension:

- Given a task, what do I need to know to fulfill it?
- Do I know these things?
- Given I know what to do, do I have the means (robot) to do it?
- If yes, which one?
- Given different steps and parts of a task, can things be done in parallel?

More complex decisions / plans

If the system does not only involve one robot with several “competencies”, but several robots with partly overlapping, partly complementary abilities, the decisions are to be taken to another dimension:

- Given a task, what do I need to know to fulfill it?
- Do I know these things?
- Given I know what to do, do I have the means (robot) to do it?
- If yes, which one?
- Given different steps and parts of a task, can things be done in parallel?
- By which robot?

More complex decisions / plans

If the system does not only involve one robot with several "competencies", but several robots with partly overlapping, partly complementary abilities, the decisions are to be taken to another dimension:

- Given a task, what do I need to know to fulfill it?
- Do I know these things?
- Given I know what to do, do I have the means (robot) to do it?
- If yes, which one?
- Given different steps and parts of a task, can things be done in parallel?
- By which robot?
- What if something goes wrong with one part of the plan? Does this affect the whole task execution, or only one of the robots?

More pancakes...

TU Munich (Prof. Michael Beetz)

Video removed.
If you are interested in how
Rosie and James really manage to make pancakes,
have a look at TU Munichs YouTube channel:

http://www.youtube.com/iasTUMUNICH

Outline

- Robots and Robotics
  - Types of robots
  - Robotics
    - Kinematics and dynamics
    - Systems (hard- and software, components)
    - Challenges (and results)
- AI in robotics
  - Mapping & Localisation
  - (Path) Planning
  - Deliberation & High level decision making and planning
- Human-Robot Interaction
  - Technical aspects
  - Cognitive / social aspects

HRI - going beyond pressing buttons

HRI - going beyond pressing buttons

Human-Robot Interaction is quite new as a research field of its own

HRI - going beyond pressing buttons

Human-Robot Interaction is quite new as a research field of its own

Like AI and Robotics themselves it is quite multidisciplinary

HRI - going beyond pressing buttons

Human-Robot Interaction is quite new as a research field of its own

Like AI and Robotics themselves it is quite multidisciplinary

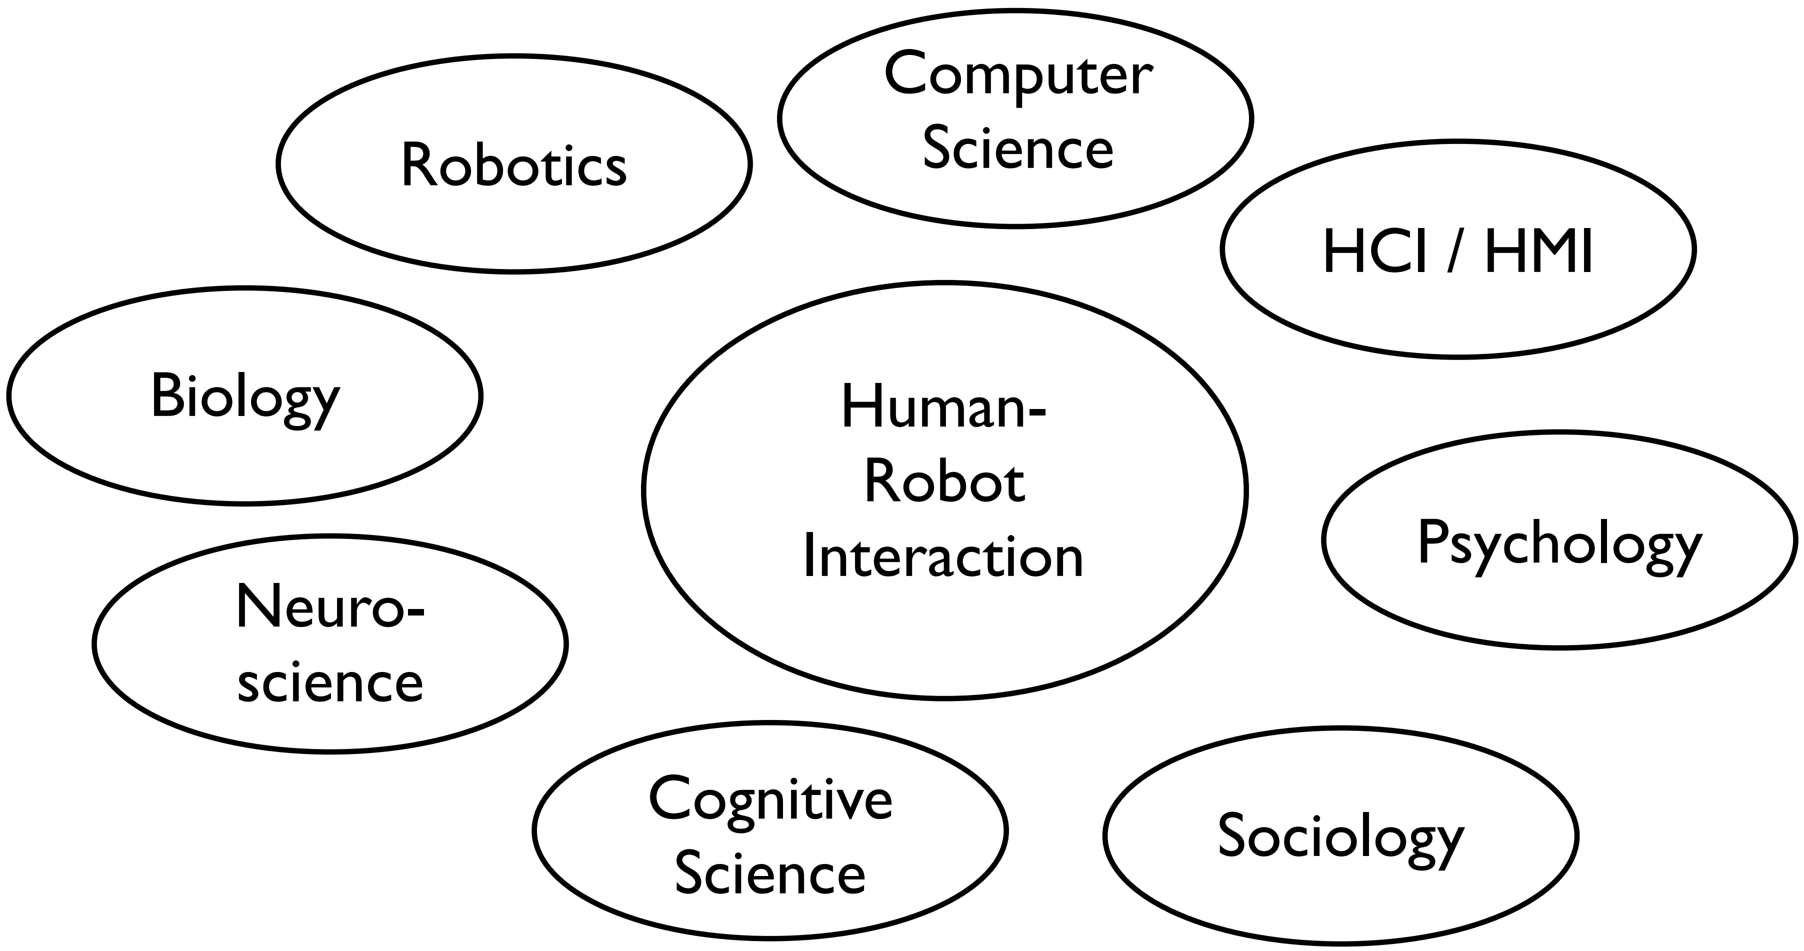

Human augmented mapping - an example for work in HRI

- Integrate robotic and human environment representations

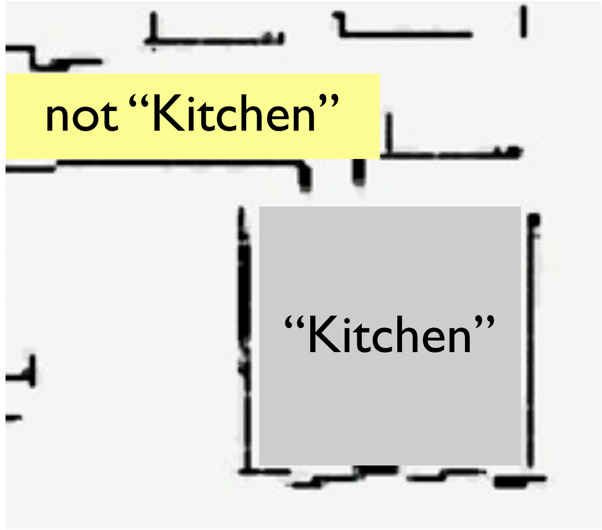

- Home tour / guided tour as initial scenario

Human augmented mapping - architectural overview

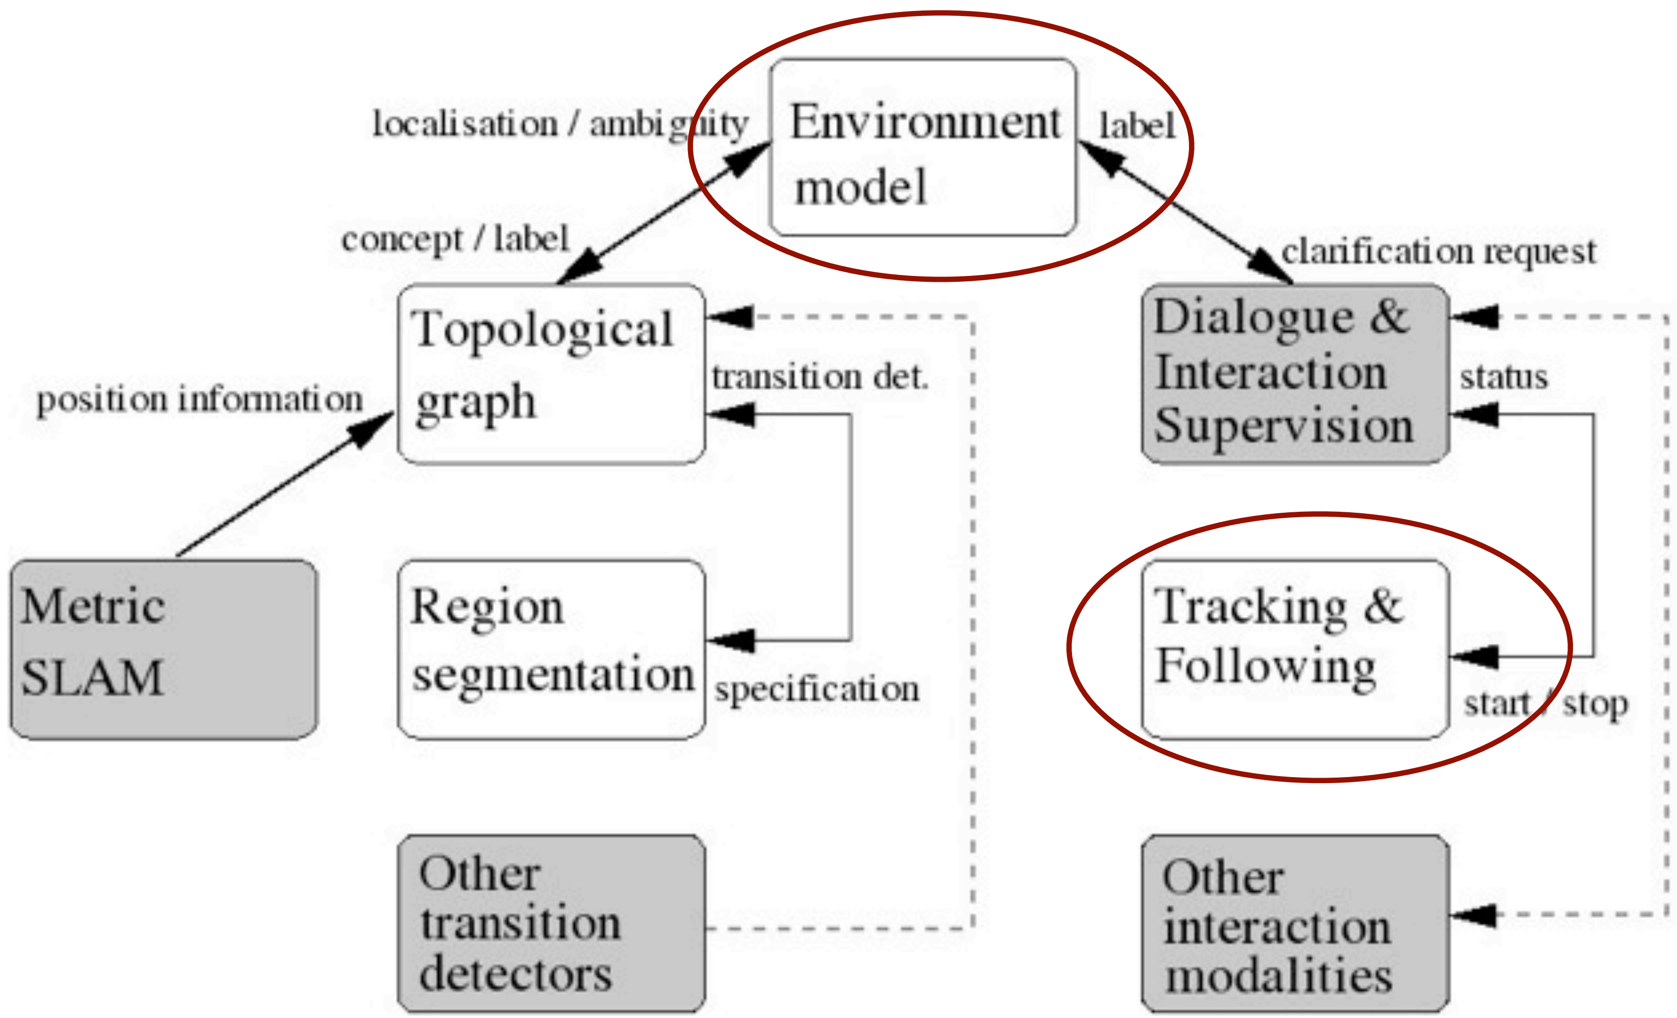

HRI techniques - tracking for following

Videos (tracker animation and corresponding video) removed.
Please contact me (Elin) if you are interested.

Tracking for following - issues

Confusion user - bystander:

Robot might follow a bystander

No error reported

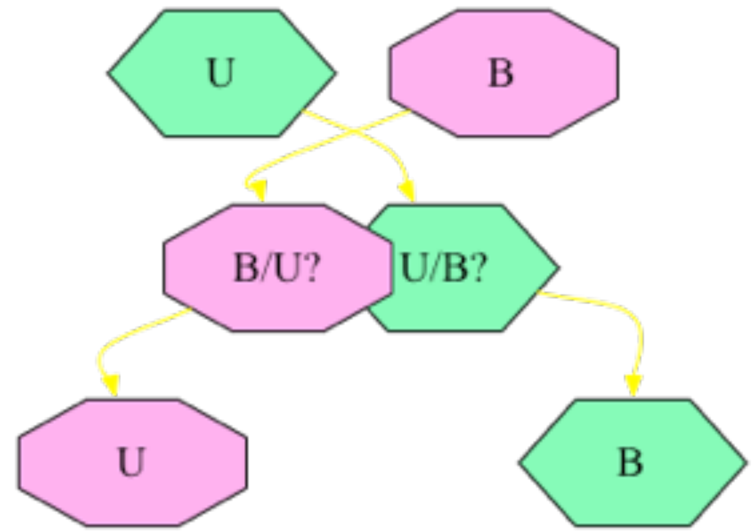

Loss of the user

No person to follow

Error handling is possible - depending on the strategy of user choice

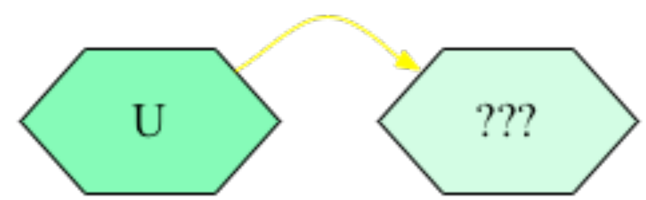

Tracking for following - approach

- Detect persons by filtering laser range data for respective patterns (legs, body sized shapes)
- Assign flags (walking, static, user ...) to targets
- Sample based Joined Probabilistic Data Association Filters (Schulz *et al.* 2001) for tracking (particle filters!)
  - Designed to keep track of multiple targets
  - Approach capable of handling the critical situations
  - Accept static targets for tracking

HRI and cognition - environment model

- Finding an environment representation that fits
  1. a human
  2. a hierarchical robotic mapping system
- Evaluating model and methods both empirically and with user studies

HRI and cognition - representing an indoor environment

- Object: Small item that can be manipulated (cup, plate, remote control)
- Location: The area from where a large, not as a whole manipulated object is reachable/visible (sofa, fridge, pigeon-holes). Also "the place where the robot is supposed to do something or look for objects"
- Region: A container for one or several locations. Offers enough space to navigate (rooms, corridors, delimited areas in hallways)
- Partially hierarchical representation of space

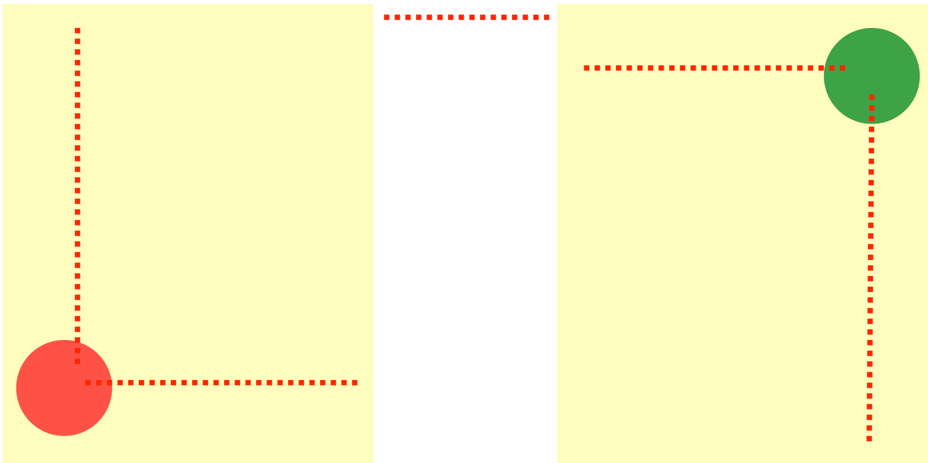

HRI and cognition - representing an indoor environment

- Object: Small item that can be manipulated (cup, plate, remote control)
- Location: The area from where a large, not as a whole manipulated object is reachable/visible (sofa, fridge, pigeon-holes). Also "the place where the robot is supposed to do something or look for objects"
- Region: A container for one or several locations. Offers enough space to navigate (rooms, corridors, delimited areas in hallways)
- Partially hierarchical representation of space

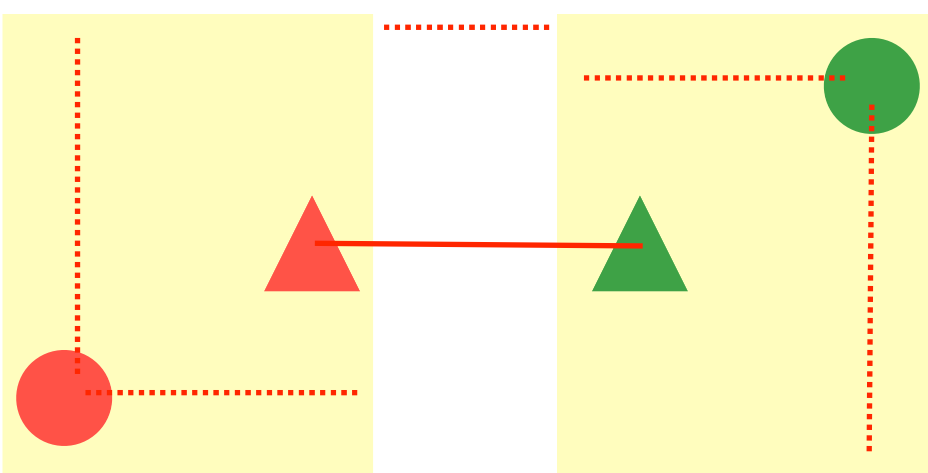

HRI and cognition - representing an indoor environment

- Object: Small item that can be manipulated (cup, plate, remote control)
- Location: The area from where a large, not as a whole manipulated object is reachable/visible (sofa, fridge, pigeon-holes). Also "the place where the robot is supposed to do something or look for objects"
- Region: A container for one or several locations. Offers enough space to navigate (rooms, corridors, delimited areas in hallways)

- Partially hierarchical representation of space

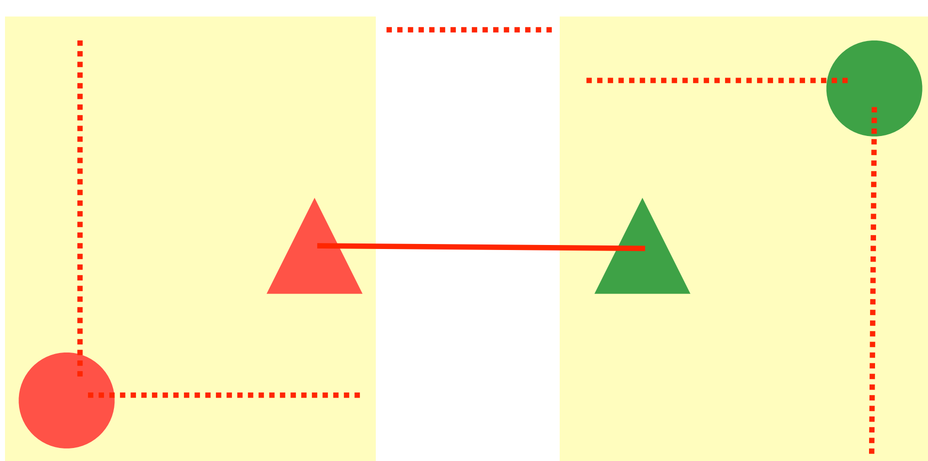

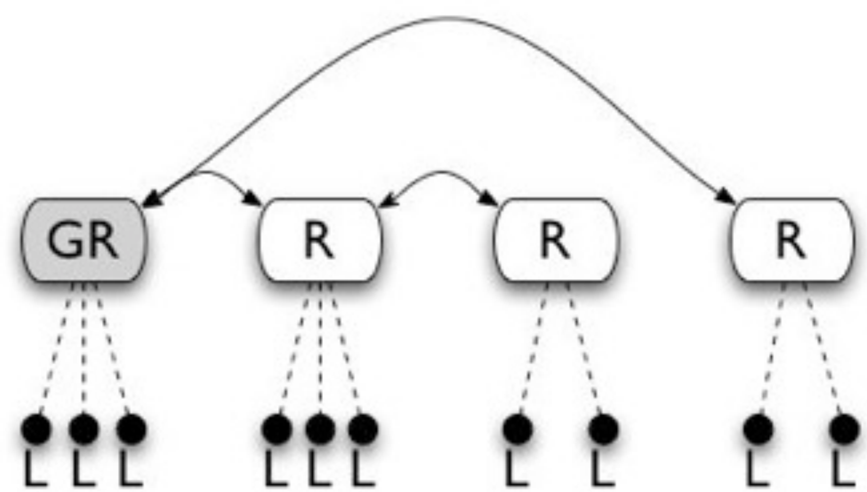

HRI and cognition - implementation of the model

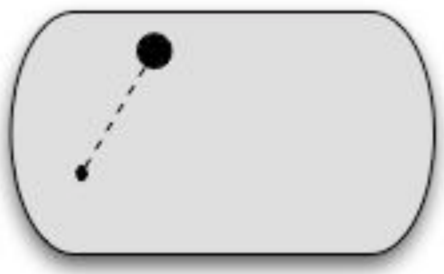

HRI and cognition - implementation of the model

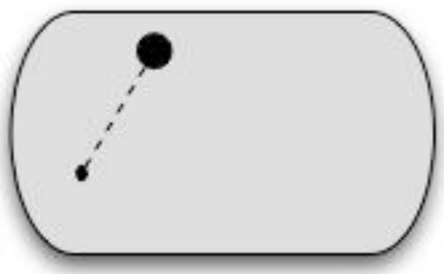

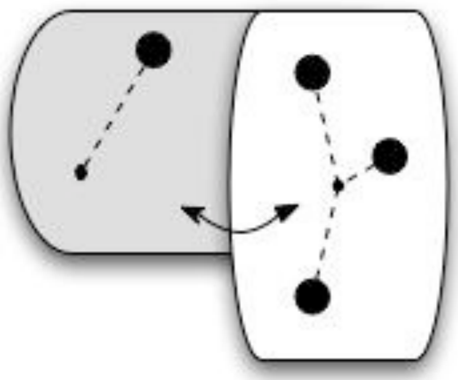

HRI and cognition - implementation of the model

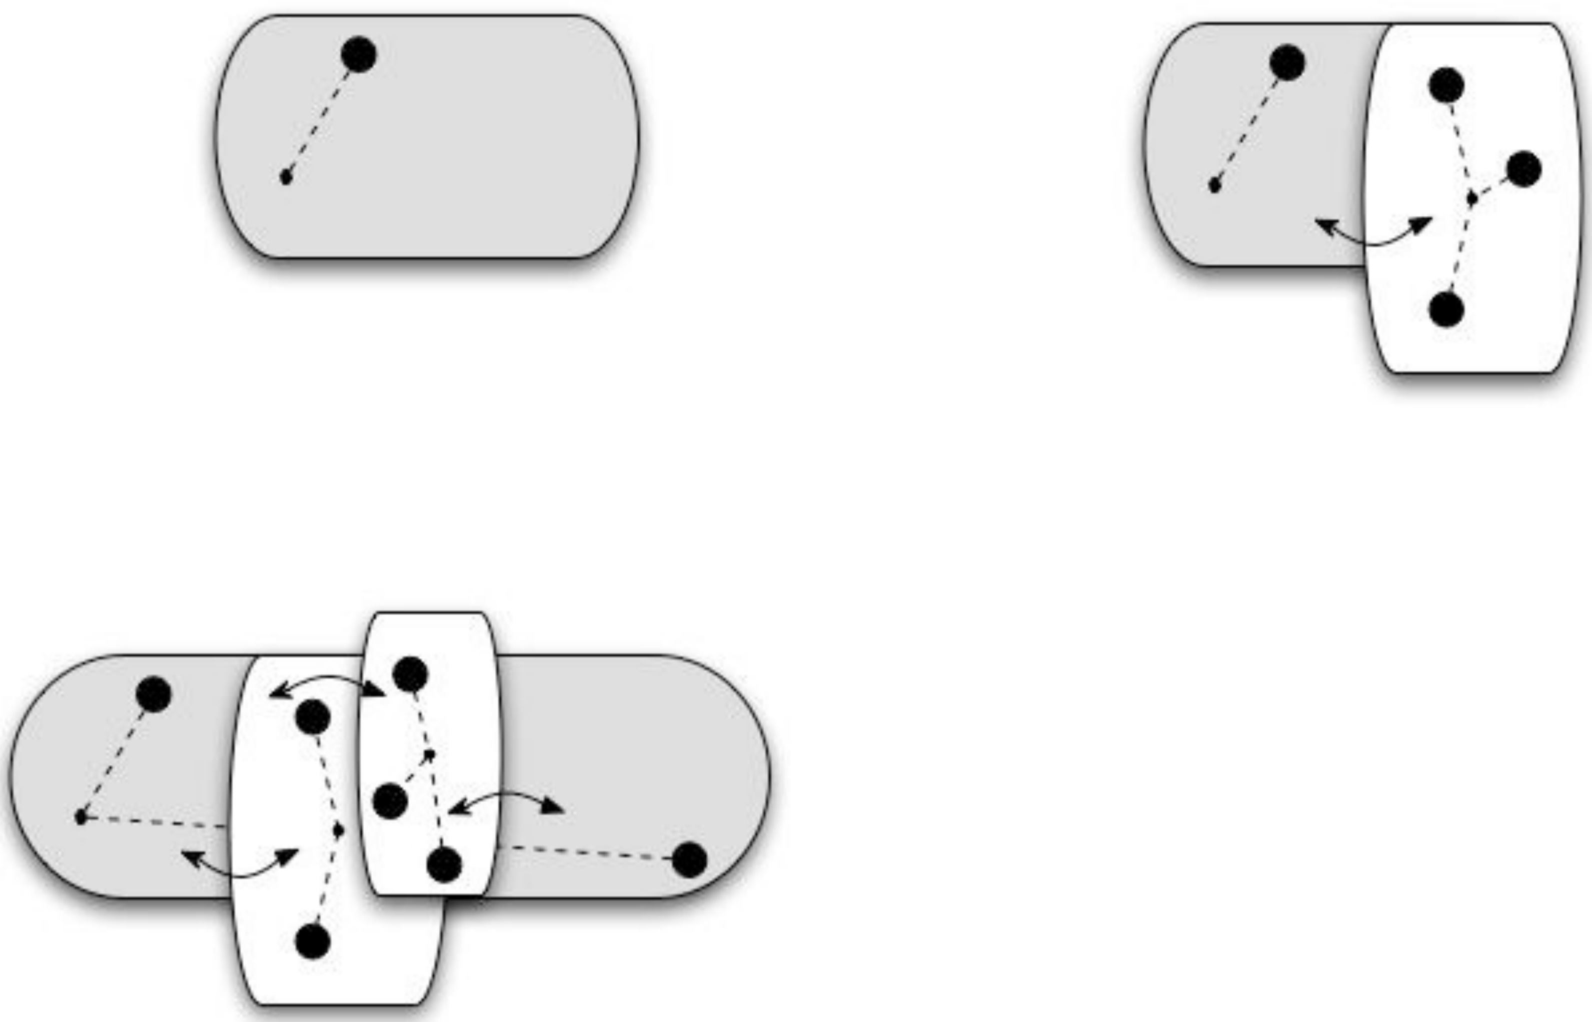

HRI and cognition - implementation of the model

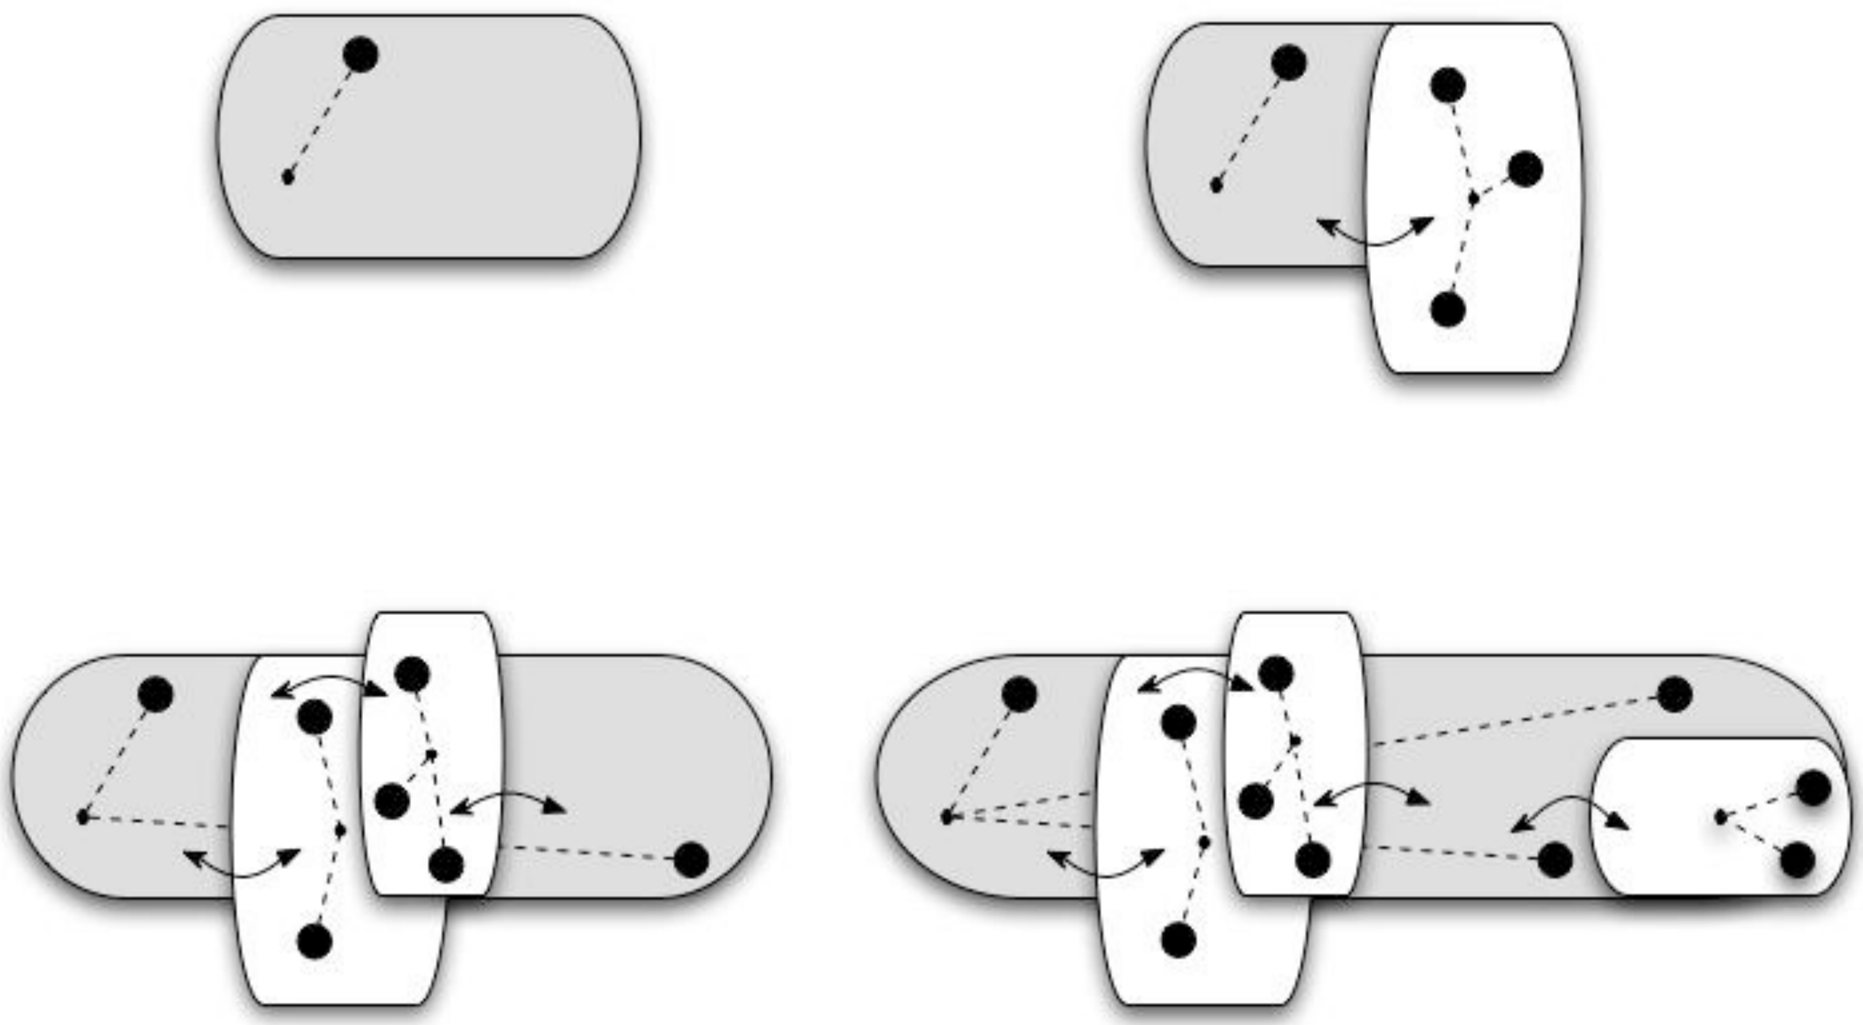

Interaction patterns - do they exist?

Can we repeatedly, with several subjects, in a clearly designed set-up, observe any structure, frequent strategies, "interaction patterns", that correspond to the spatial categories *Region*, *Location*, and *Object* when people present an indoor environment to a mobile robot?

User Study

Video removed. This was original research material and
can not be published due to privacy agreements with the study participants.
Contact me (Elin) if you want to learn more about the study.

Interaction patterns!

Patterns seem to exist - classification / interpretation of user "behaviour" is subject to current work.

First choice of approach: Bayesian networks!

Robotics and Semantic Systems @Lunds Universitet

- Master's projects (Ex-jobb) in AI, NLP, Robotics (mapping, software, cognitive modeling...), HRI
  - Internal (research oriented) or external (industry related)
  - International through project partners (depends of course on formalities as well ;-)
- Lab visit in the Robotlab in M-huset
- Contact us: Jacek, Pierre, Elin or other members of the group: Klas Nilsson, Mathias Haage, Sven Gestegård Robertz

Resources

Images on slide 4, top row left to right: ABB Flexpicker (www.abb.se/), Performance PeopleBot "Minnie", (original material / mobilrobots.com), Pioneer P3dx "Snoopy" (original material / mobilrobots.com), Honda Asimo (world.honda.com/ASIMO). Bottom row, left to right: KUKA KR 180-2 (Series 2000) (www.kuka.com), (top:) Keepon (www.beatbots.org), (bottom) Leonardo (MIT Media Lab, web.media.mit.edu/~cynthiab/research/research.html), iCub (en.wikipedia.org/wiki/ICub)

Images on slide 6 are part of the course book resources (http://aima.cs.berkeley.edu, robot arm and synchro drive illustration) or stem from MobileRobots Inc (http://mobilerobots.com, the Pioneer p3dx) and from DLR (http://www.dlr.de/rm/desktopdefault.aspx/tabid-3758/, "Rollin' Justin").

Illustrations on slide 12 are original material (simple system for user detection and speech handling), and obtained from ABB's product description for their software "ABB RobotStudio", http://www.abb.se

Architecture illustrations on slide 13-15 are taken from the following publications:

R. A. Brooks, "Planning is just a way of avoiding figuring out what to do next", Technical report, MIT Artificial Intelligence Laboratory, 1987

R. Arkin, "Integrating Behavioural, Perceptual, and World Knowledge in Reactive Navigation", in Robotics and Autonomous Systems 6, 1990

T.P. Spexard, M. Hanheide, "System Integration Supporting Evolutionary Development and Design", in proceedings of Human Centered Robotic Systems, 2009

Images (apart from robot) on slide 16 are: original material (line map illustration); original material (topological structure illustration); part of Sebastian Thruns article "Learning Occupancy Grid Maps with Forward Sensor Models", in Autonomous Robots 15, 2003 (occupancy grid illustration).

Images (apart from the robots) on slides 23 to 26 are part of the course book resources (aima.cs.berkeley.edu).

Images on slides 32 to 40 are original material and can to a large extent be found in my doctoral dissertation.