Wrap-up of Sensing/Perception + Representational Issues

October 8, 2002

Class Meeting 14

Polly: Horswill vision-based robot

Toto: Mataric mapping robot

Today’s Objectives

- Part I: Wrap-Up of Sensing/Perception:
  - Case Study: Multiple Sensors/Sensor Fusion for Outdoor Elevation Map Building

- Part II: Representational Issues for Behavioral Systems:
  - To understand working definitions for knowledge and knowledge use
  - To explore qualities of knowledge representation
  - To understand what types of knowledge may be representable for use within robotic systems
  - To determine the appropriate role of world and self-knowledge within behavior-based robotic systems
  - To study several representational strategies developed for use within behavior-based systems

Case Study: Distributed Heterogeneous Sensing for Outdoor Multi-Robot Localization, Mapping, and Path Planning

- Application Objective: Moving Multi-Robot Teams Outdoors
- Applications require robots to work over period of time in same outdoor area (i.e., not just "pass through")
- Applications require efficiency of navigation, "optimal" path planning

Example applications:

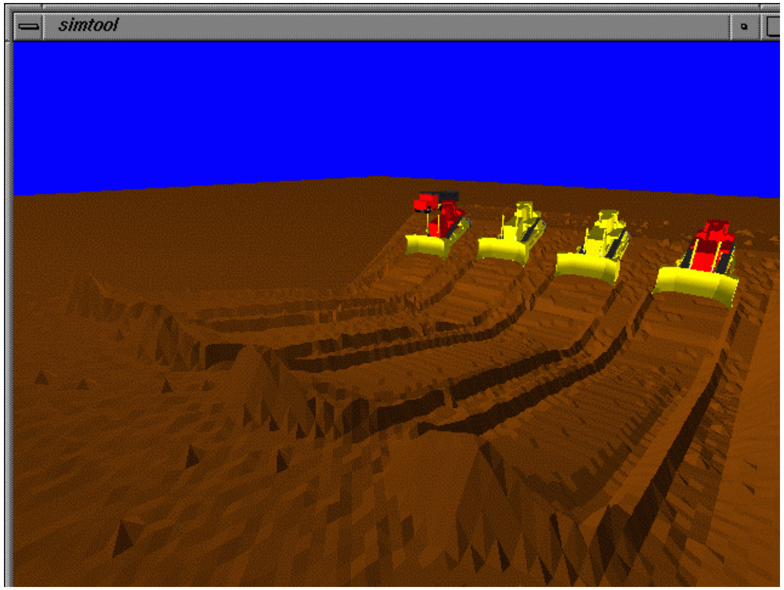

Surface coal mining

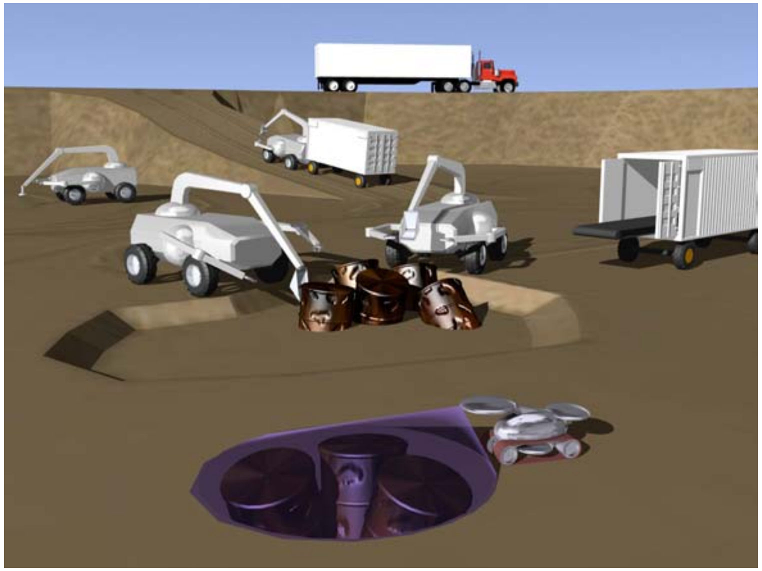

Hazardous waste cleanup

Case Study Objective: Teams of Robots Operating in Outdoor Environments w/o Significant Setup Time

- Application of robot teams to site security, surveillance and reconnaissance, etc.

ATRV-mini robots at ORNL

Case Study: Mission Objective -- Plan Multi-Robot Paths to Optimize Performance

- Example motivation application:
  - Perimeter patrols for security
  - Desire to minimize infrastructure setup
  - Need for robots to move along highest visibility paths, dependent upon the number of available robots
    - Not easy to derive manually
    - May change if team composition or mission changes
- However, can’t plan optimal paths without decent maps
- Obtaining “good enough” maps not straight-forward:
  - Digital Elevation Maps (DEMs) usually of too low resolution
  - Human surveys/sensor scans time consuming/tedious
- Approach:
  - Teams of robots working together to build “good enough” maps, which are then used for multi-robot path planning

Case Study: Overall Highest-Level Schematic

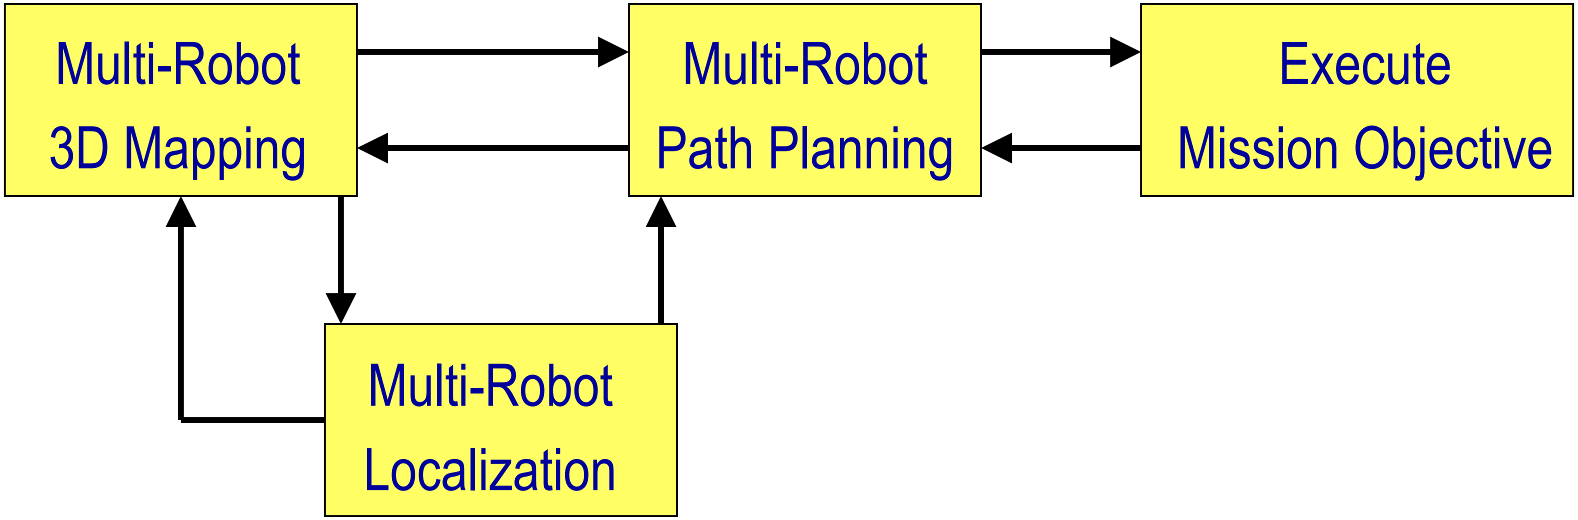

Strategy: Use Distributed Sensing/Positioning to Improve Localization, 3D Mapping, Path Planning

- Basic assumption: DGPS sometimes available, but frequently obstructed/degraded due to trees, buildings, multi-pathing, etc.

- Allow robots to take advantage of relative positioning to improve localization

- Allow robots to coordinate relative pose information to cooperatively build terrain map

- Multi-robot coordination/revision of paths to satisfy multiple objectives

- Information from a variety of sensory sources needed

- Sensory data must be fused to help interpretation of information

Case Study: Robot Team and Experimental Setup

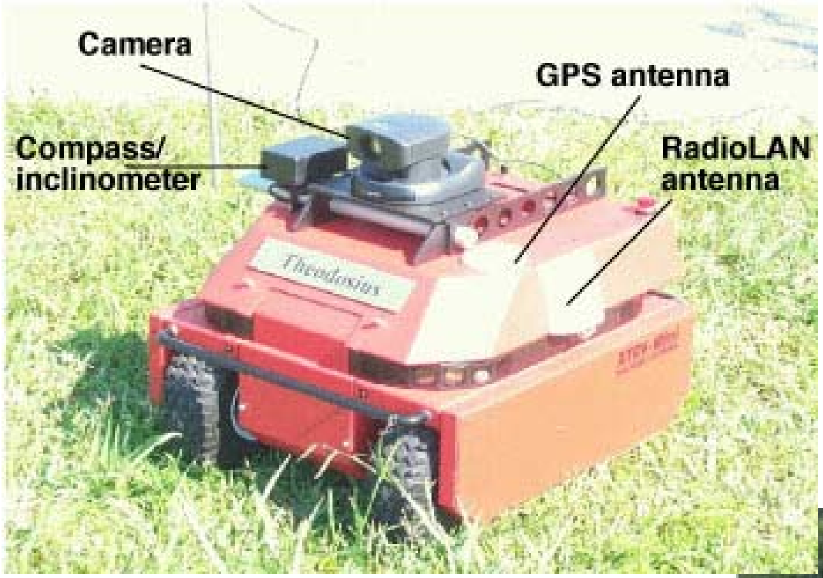

- Robot Team: 4 ATRV-mini robots (Manuf: RWI/iRobot)
  - Named (after Roman Emperors): Augustus, Constantine, Theodosius, Vespasian

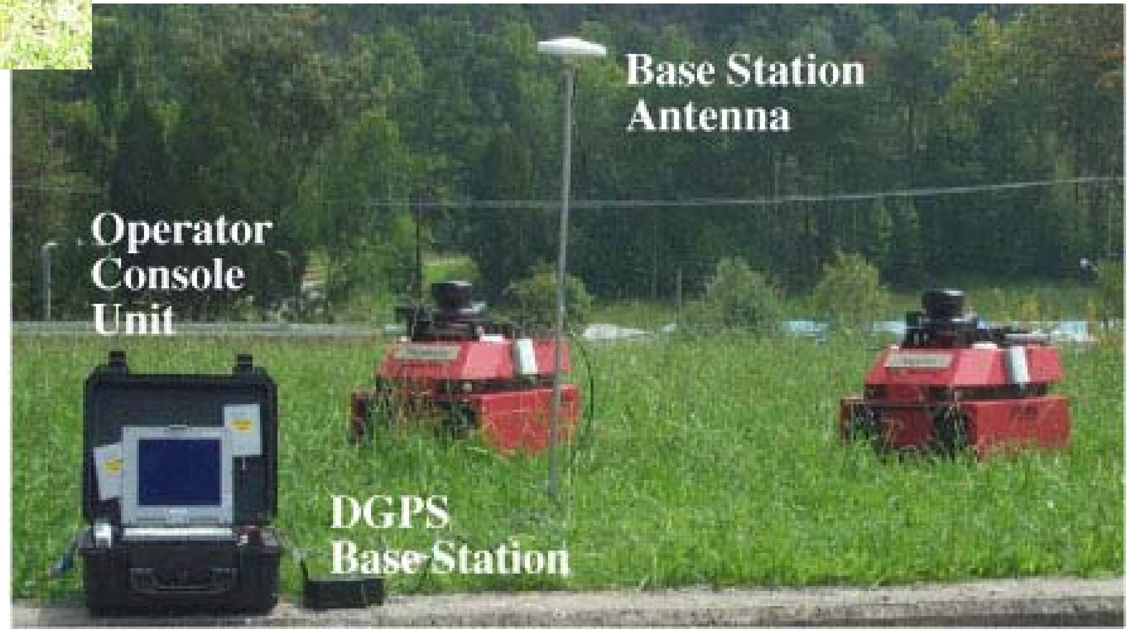

- Sensors:
  - 2 robots: PTZ camera
  - 2 robots: SICK laser
  - Compass/inclinometer
  - DGPS
  - Sonar

Case Study: Multi-Robot Path Planning

- Calculating optimal paths for all robots simultaneously is computationally expensive

- For now, path planning issue:
  - Given assigned starting and goal positions, find optimal path to goal

- Approach:
  - Plan optimal independent paths for each robot
    - Cost is function of obstacles, distance, terrain slope, path smoothness:
  - Search for inter-robot collisions along paths
  - Define optimal velocity profiles to enable robots to follow paths while eliminating collisions
    - Cost is function of collisions, N-dimensional distance, robot idle time, prioritized penalty for giving way

Details will come in a later class

Case Study: Multi-Robot Localization

- Approach: Extended Kalman Filter based on multi-robot relative localization
- Similar to Roumeliotis and Bekey, 2000, except in ours:
  - Kinematic model of robots is nonlinear
  - No absolute positioning system assumed consistently available to give relative pose information
  - Robots traverse on uneven and unstructured outdoor terrains
- When DGPS unavailable, use laser- or vision-based determination of relative positioning

Details to come in a later class

Case Study: When DGPS Degraded, Use Relative Positioning

Vision-based relative positioning:

Laser-based relative positioning:

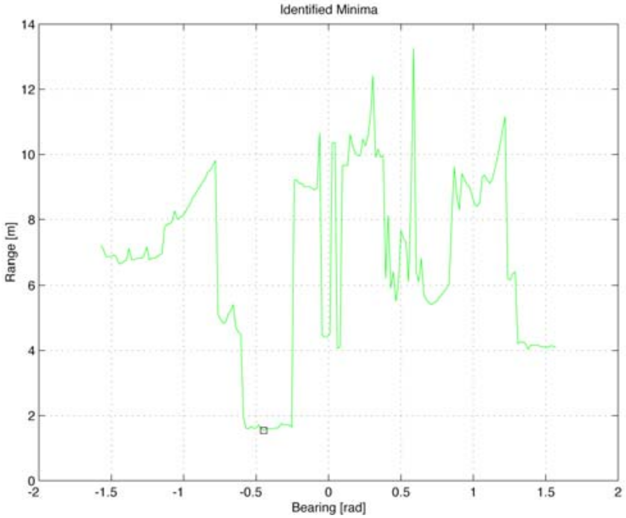

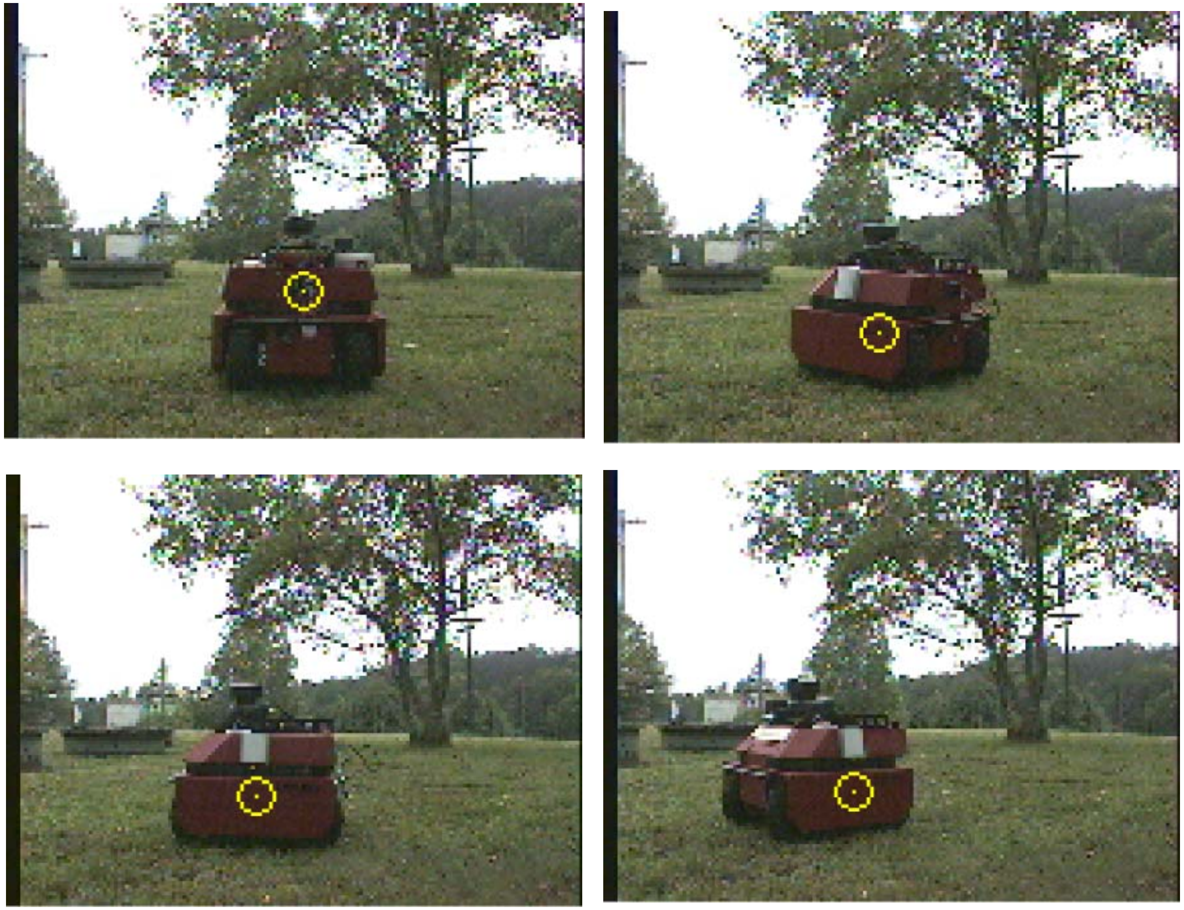

Case Study: Results of Cooperative Localization

EKF estimated robot paths:

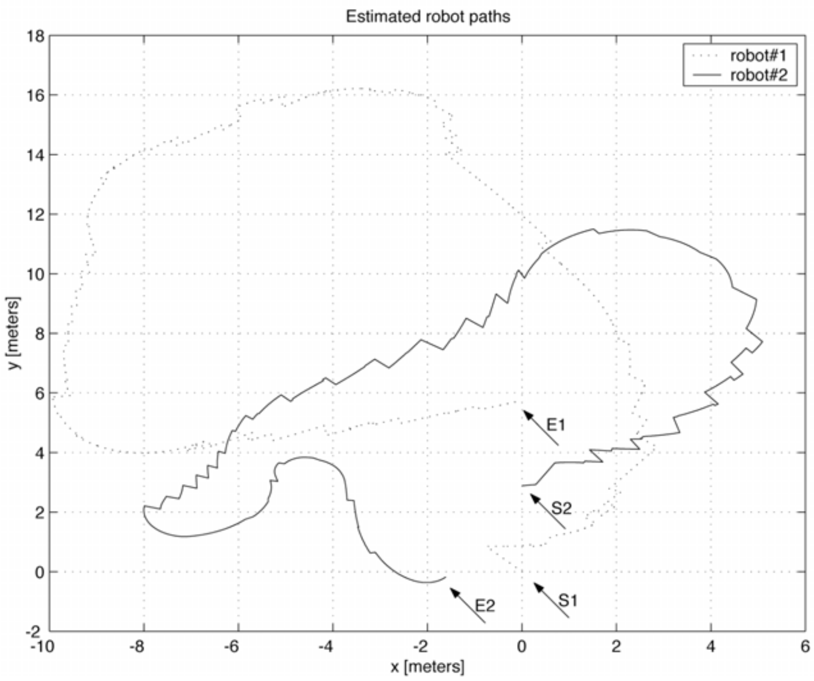

External corrections from relative positioning:

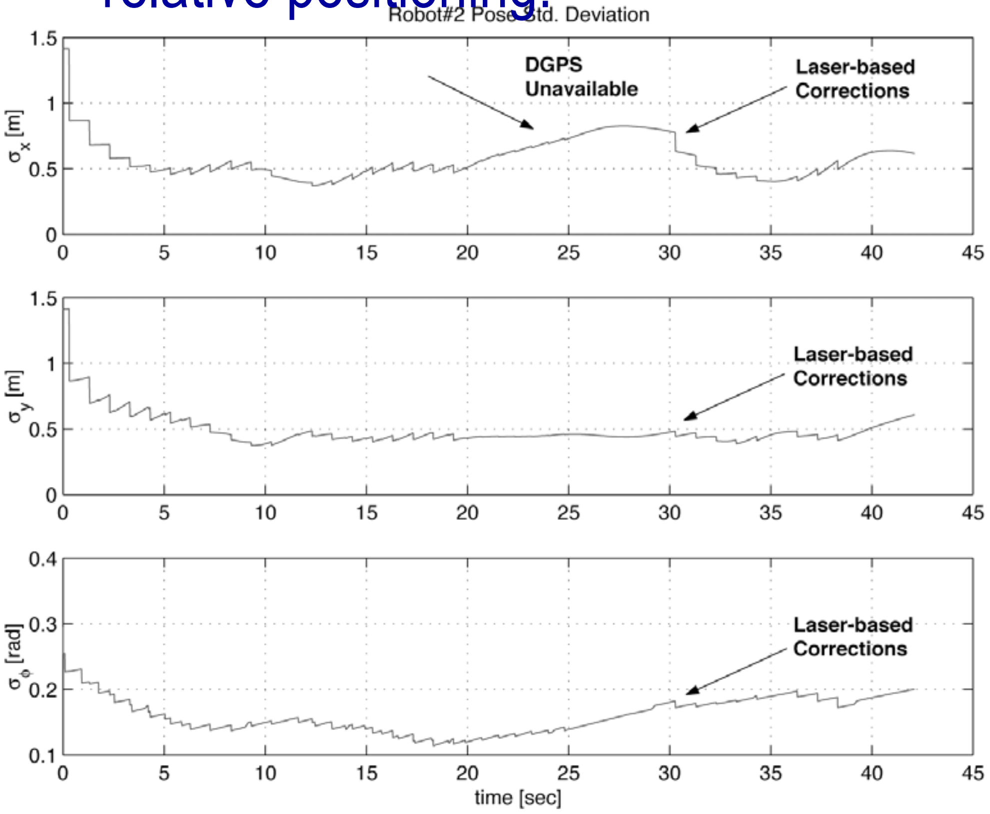

Case Study: Multi-Robot 3D Terrain Mapping

- Approach:
  - Depth-from-camera-motion (adaptation of Matthies *et al.*, 1989) to obtain depth ranges to features in environment
  - Relative pose of robots associated with depth information
  - Elevation gradient of terrain determined by fusing DGPS altitude info with vertical displacements from robot pitch inclinometer
  - Depth and elevation info registered with covariances (which provide confidence of measurements)
  - Depth map updated with high-confidence information

Case Study: Overall Terrain Mapping Scheme

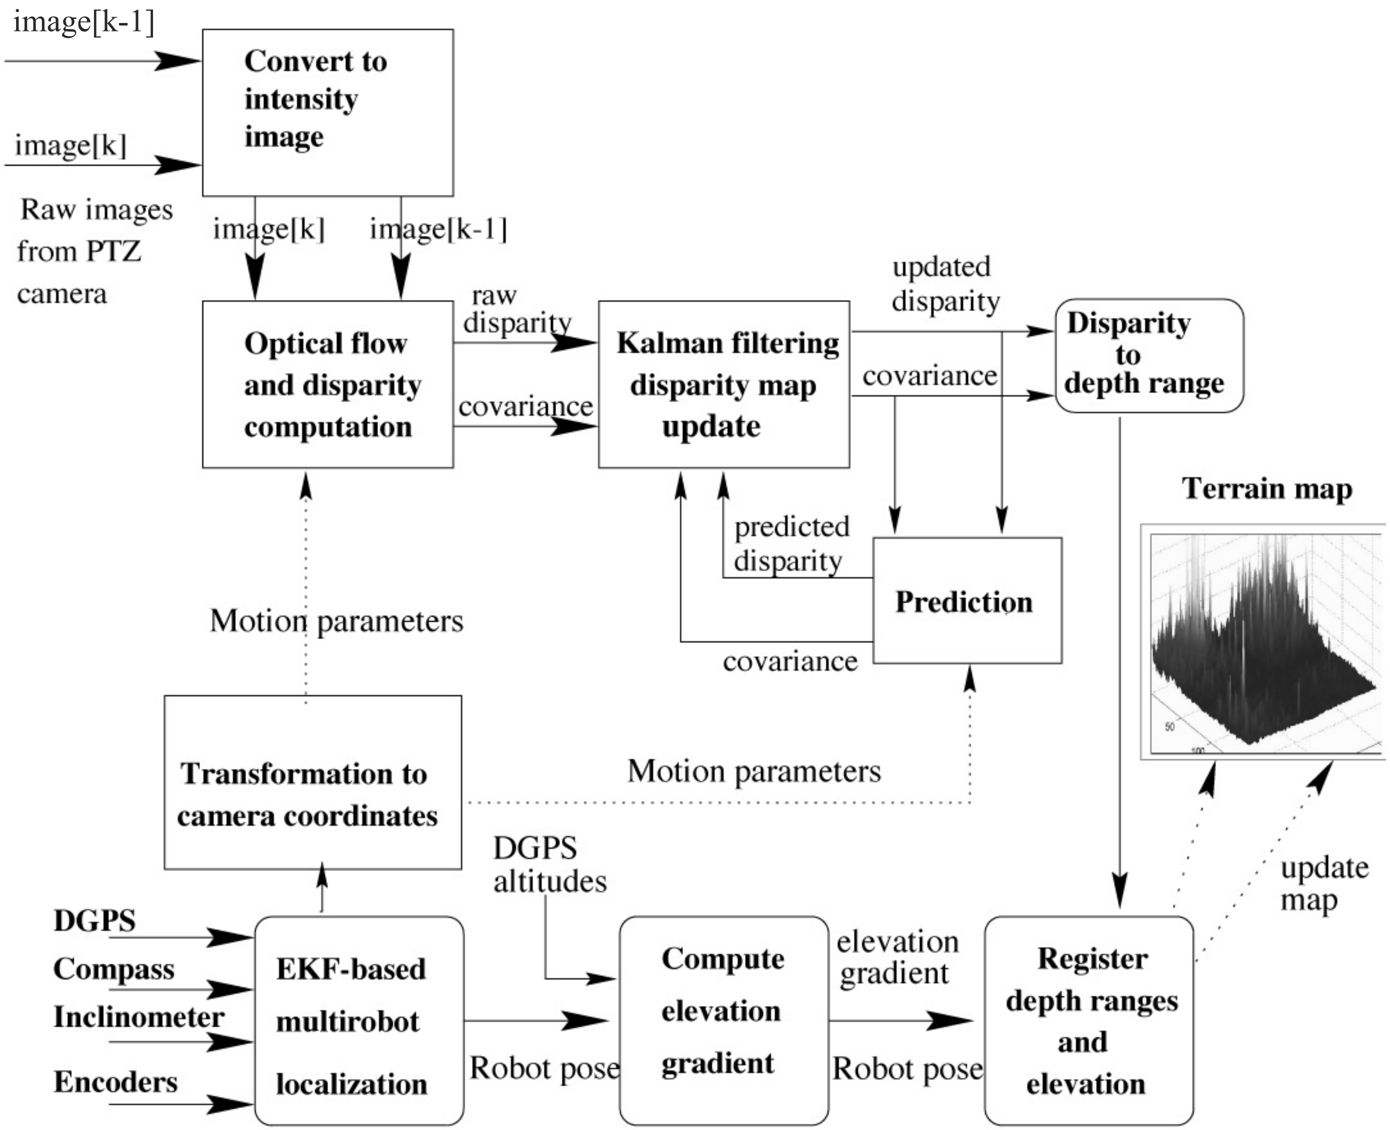

image[k-1]

image[k]

Raw images from PTZ camera

Convert to intensity image

image[k]

image[k-1]

Optical flow and disparity computation

raw disparity

covariance

Kalman filtering disparity map update

updated disparity

covariance

Disparity to depth range

predicted disparity

covariance

Prediction

Terrain map

Motion parameters

Transformation to camera coordinates

Motion parameters

DGPS altitudes

DGPS

Compass

Inclinometer

Encoders

EKF-based multirobot localization

Robot pose

Compute elevation gradient

elevation gradient

Robot pose

Register depth ranges and elevation

update map

Case Study: Preliminary Results of Mapping Approach

Actual scene

Depth map

Depth covariance

Augustus:

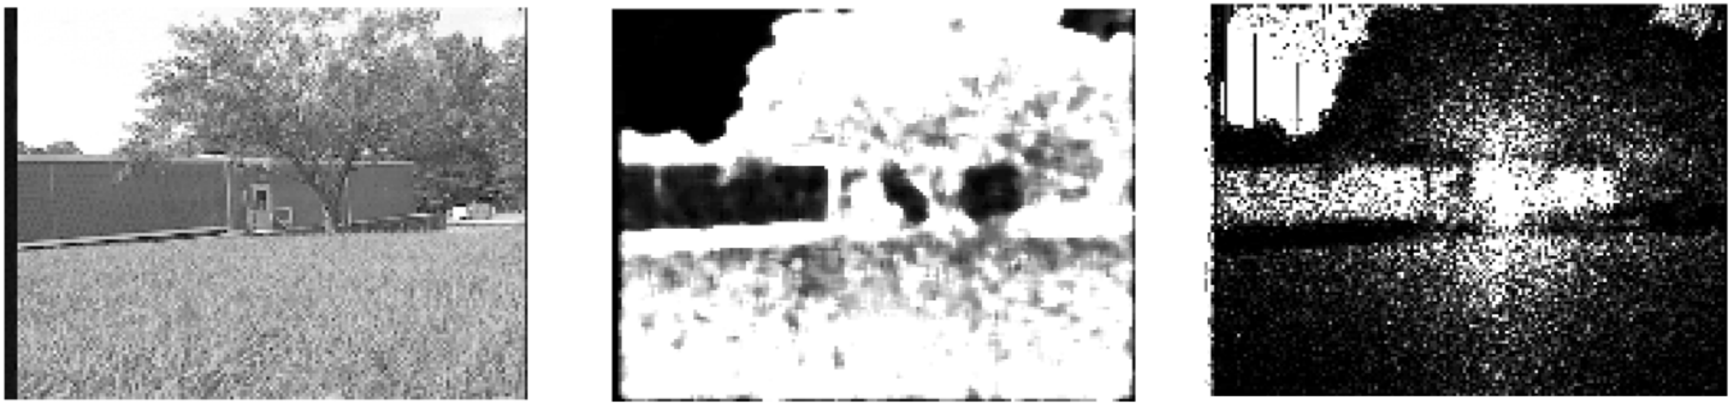

Theodosius:

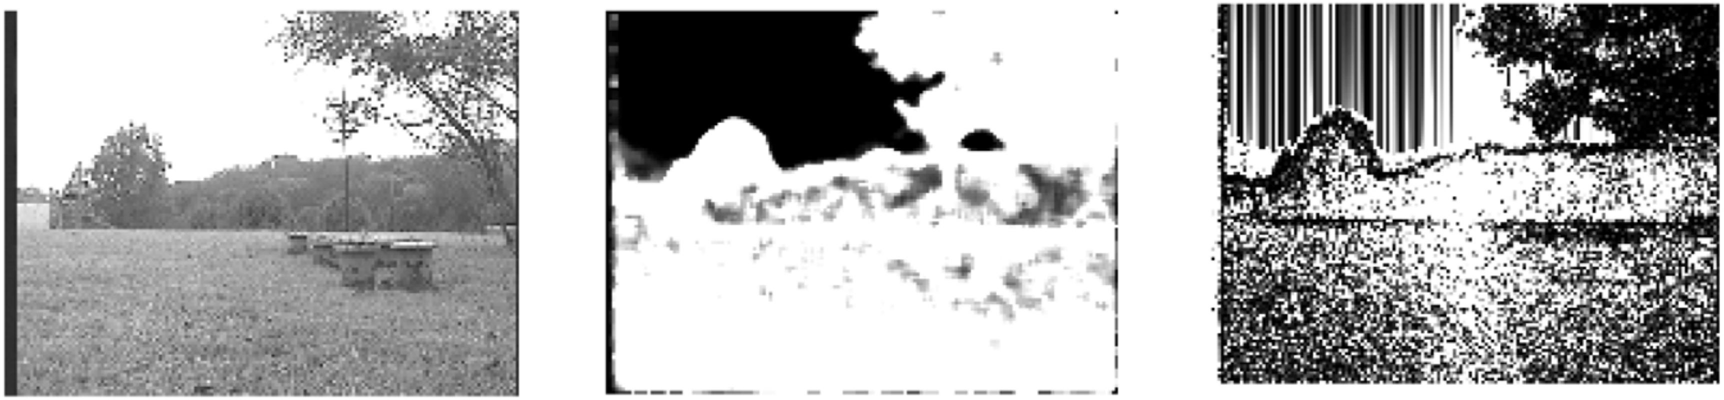

Case Study: Preliminary Results of Mapping Approach

Partially updated terrain map from two robot explorations:

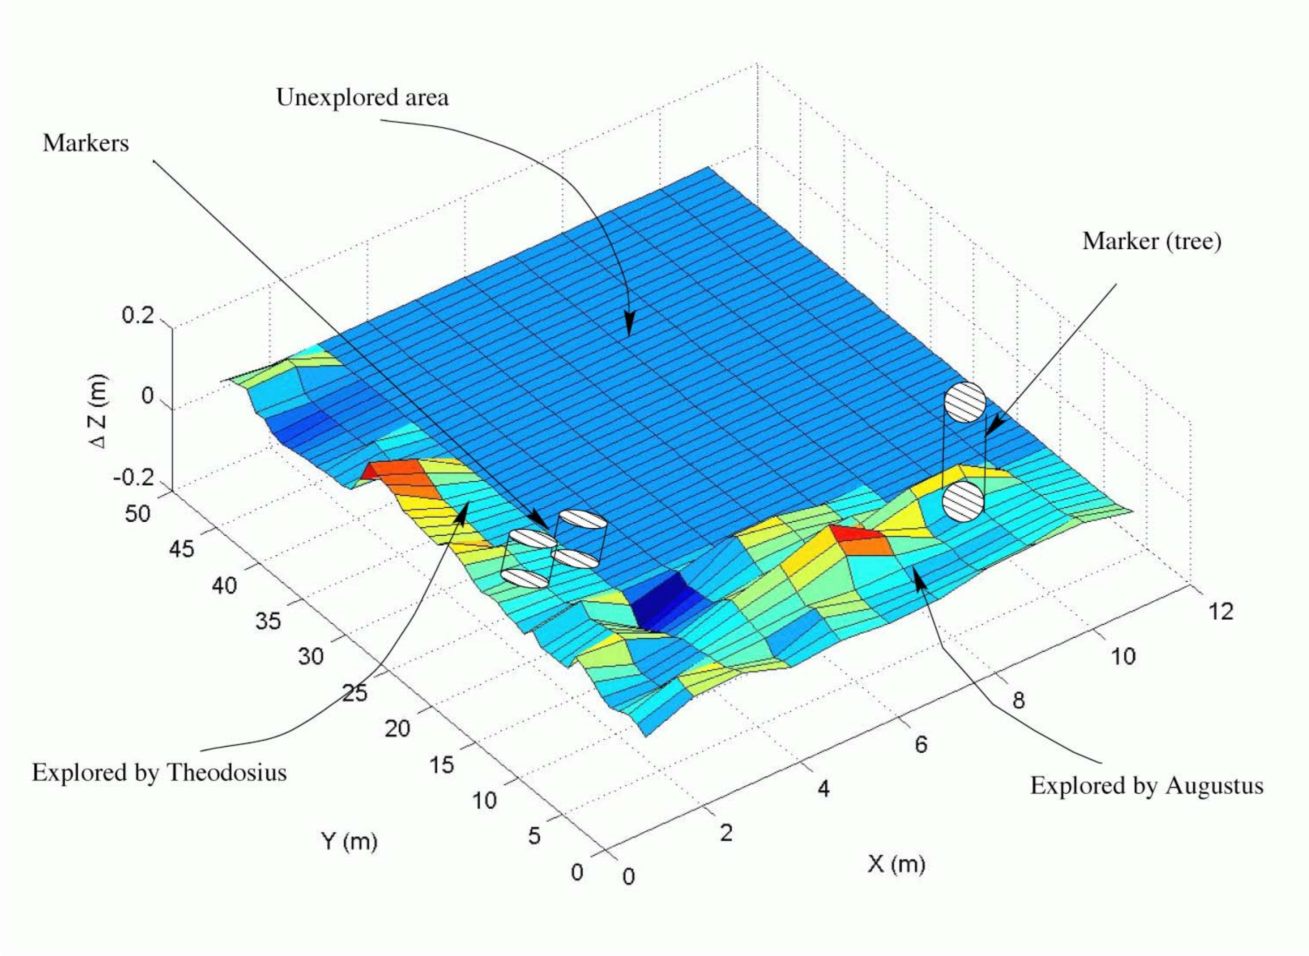

Case Study: Summary

- Case Study Objective:
  - Development of localization, mapping, and path planning tools enabling multi-robot teams to operate in outdoor environments quickly, without need for extensive human setup time.
- Multiple sensors used:
  - DGPS
  - CCD Cameras
  - Compass
  - Inclinometer
  - Encoders
  - Laser range scanner
- Sensory information had to be merged in multiple ways to obtain desired map knowledge

Part II: Representational Issues for Behavioral Systems

- Objectives:
  - To understand working definitions for knowledge and knowledge use
  - To explore qualities of knowledge representation
  - To understand what types of knowledge may be representable for use within robotic systems
  - To determine the appropriate role of world and self-knowledge within behavior-based robotic systems
  - To study several representational strategies developed for use within behavior-based systems

What is Knowledge?

- Knowledge (like "intelligence"): notoriously difficult to define

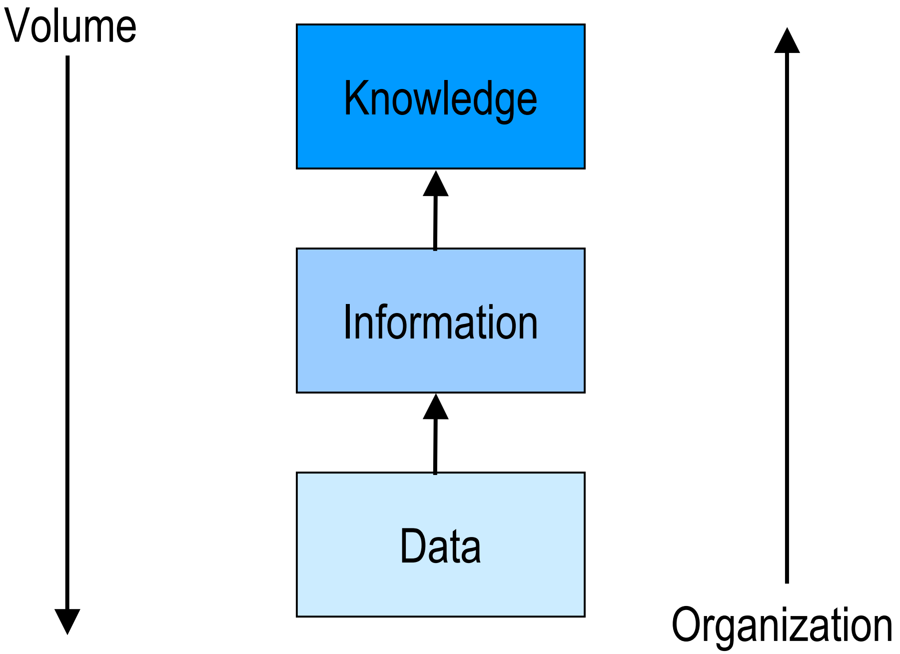

Definitions of Knowledge

- Knowledge (Turban 1992): Understanding, awareness, or familiarity acquired through education or experience. The ability to use information.

- Knowledge representations (Steels 1995): Physical structures which have correlations with aspects of the environment and thus have predictive power for the system.
  - Environmental correlation:
    - Temporal durability/persistance (e.g., short term, long term)
    - Nature of correlational mapping (e.g., metric, relational)
  - Predictive power:
    - If no need to predict, then can rely entirely on what is sensed (i.e., reactive)

Key issue: "Sensing" vs. "Representing"

Tradeoffs for Knowledge Use

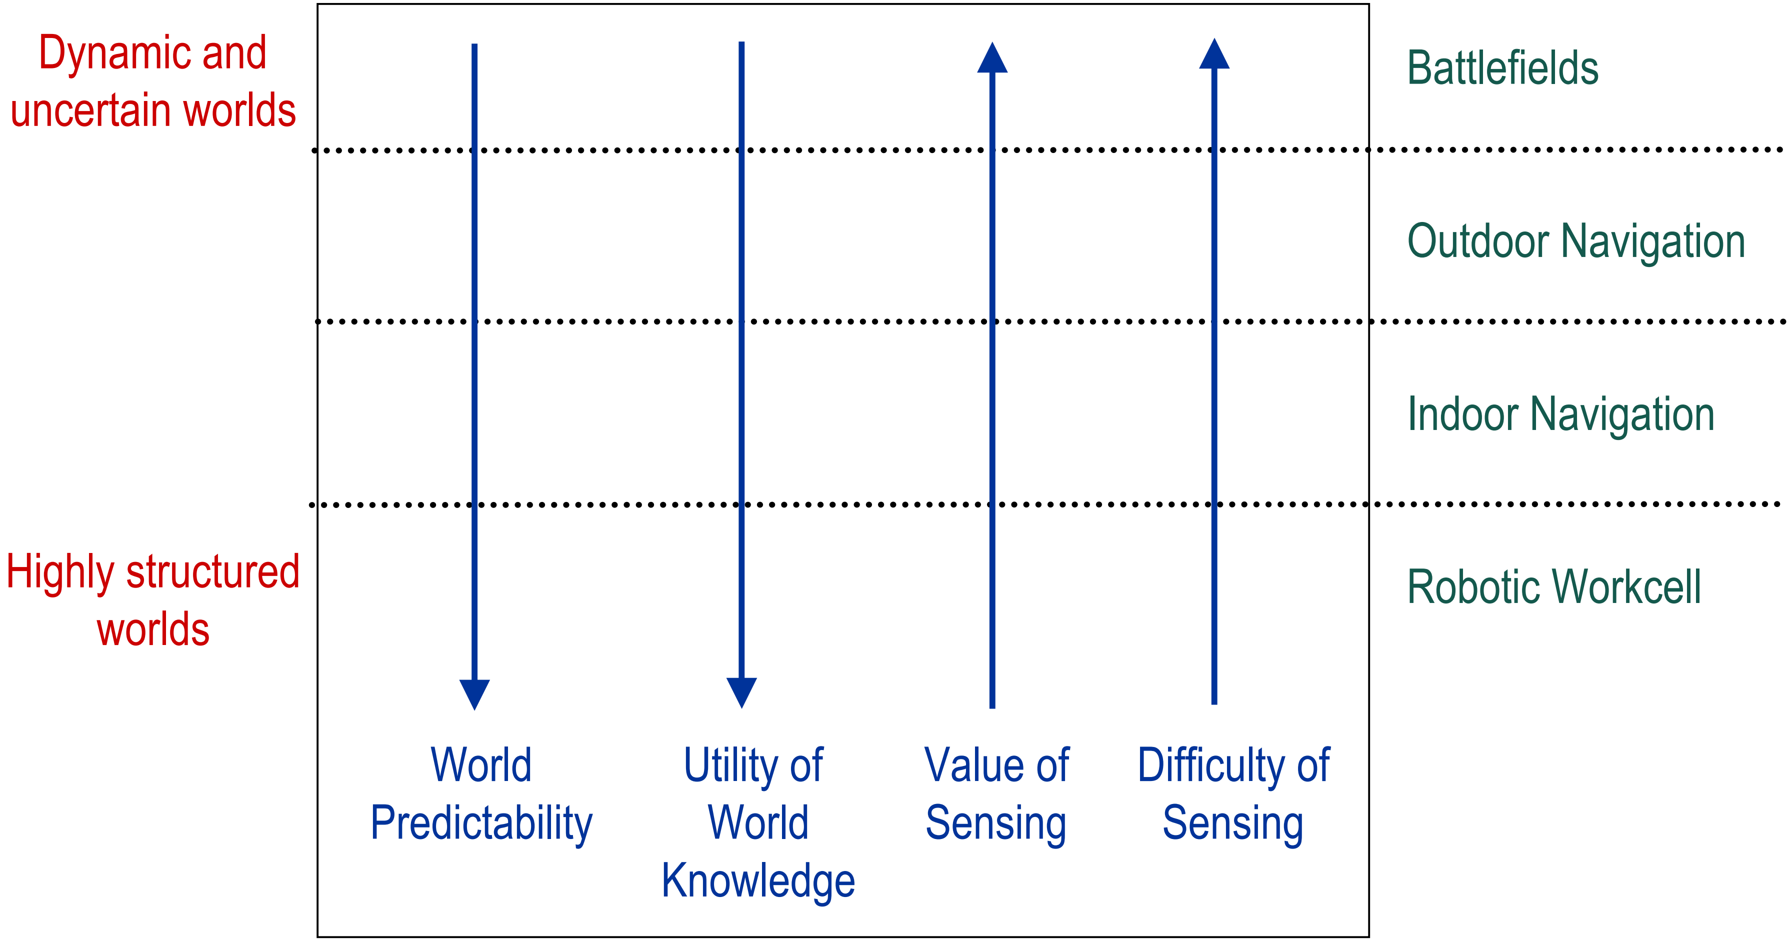

Dynamic and uncertain worlds

Highly structured worlds

Battlefields

Outdoor Navigation

Indoor Navigation

Robotic Workcell

World Predictability

Utility of World Knowledge

Value of Sensing

Difficulty of Sensing

Considerations

- When world changes rapidly, stored knowledge potentially becomes obsolete quickly
- However, continuous sensing is not free (computationally); prefer to minimize sensing process
- Issue: maintaining accurate correlation between robot's position in world and its representational point of view
  - For spatial location, this is called *localization*
  - *"Where am I?"*
  - Purely reactive systems do not address this issue

Taxonomy of Knowledge Representations

- Explicit: symbolic, discrete, manipulable knowledge representations typical of traditional AI

- Implicit: knowledge that is non-explicit, but reconstructable and can be made explicit through procedural usage.

- Tacit: knowledge embedded within the system that existing processes cannot reconstruct

- *Symbolic* systems: use explicit knowledge
- *Sub-symbolic* systems: use implicit or tacit knowledge

Symbol Grounding Problem

- Symbol grounding problem: refers to the difficulty in connecting the meaning (semantics) of an arbitrary symbol to a real world entity or event.
  - Degeneracy is often recursive or circular (symbols used to describe symbols)
- For humans (and behavior-based robots), meaning is derived from interactions with objects in the world → not intrinsic to the objects themselves

Types of Knowledge

- Spatial world knowledge: an understanding of the navigable space and structure surrounding the robot
- Object knowledge: categories or instances of particular types of things within the world
- Perceptual knowledge: information regarding how to sense the environment under various circumstances
- Behavioral knowledge: an understanding of how to react in different situations
- Ego knowledge: limits on the abilities of the robot's actions within the world (e.g., speed, fuel, etc.) and on what the robot itself can perceive (e.g., sensor models)
- Intentional knowledge: information regarding the agent's goals and intended actions within the environment – a plan of action.

Another categorization: Based on Durability

- Persistent knowledge:
  - A priori information about robot's environment that can be considered relatively static for mission's (or task's) duration
  - Allows for pre-conceived ideas of robot's relationship with world
  - E.g., object knowledge, models of free space, ego model of robot itself
  - Knowledge base: long-term memory (LTM)

- Transitory knowledge:
  - Acquired dynamically as robot moves through world
  - Knowledge base: short-term memory (STM)
  - Typically forgotten (fades) as robot moves away from locale where informaiton was gathered

Time Horizon of Knowledge

Transitory Knowledge

Persistent Knowledge

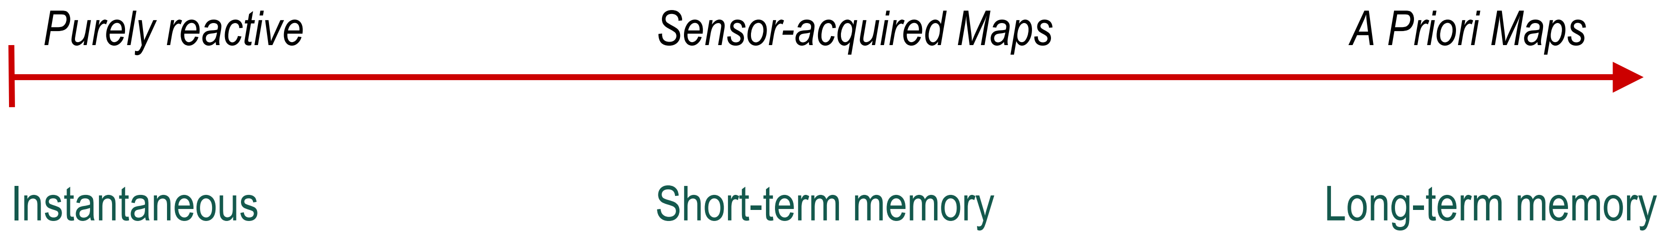

Purely reactive

Sensor-acquired Maps

A Priori Maps

Instantaneous

Short-term memory

Long-term memory

Time Horizon

Representational Knowledge for Behavior-Based Systems

- Short-term behavioral memory

- Long-term memory maps:
  - Sensor-derived maps
  - A priori map-derived representations

Short-Term Behavioral Memory

- Advantages of behavioral memory:
  - Reduces need for frequent sensor sampling in reasonably stable environments
  - Provides recent information to guide robot that is outside of its sensory range

- Characteristics:
  - Used in support of a single behavior (usually obstacle avoidance)
  - Representation directly feeds behavior rather than tying it to a sensor
  - Transitory: representations are constructed, used while the robot is in the environment, and then discarded

Behavioral Memory

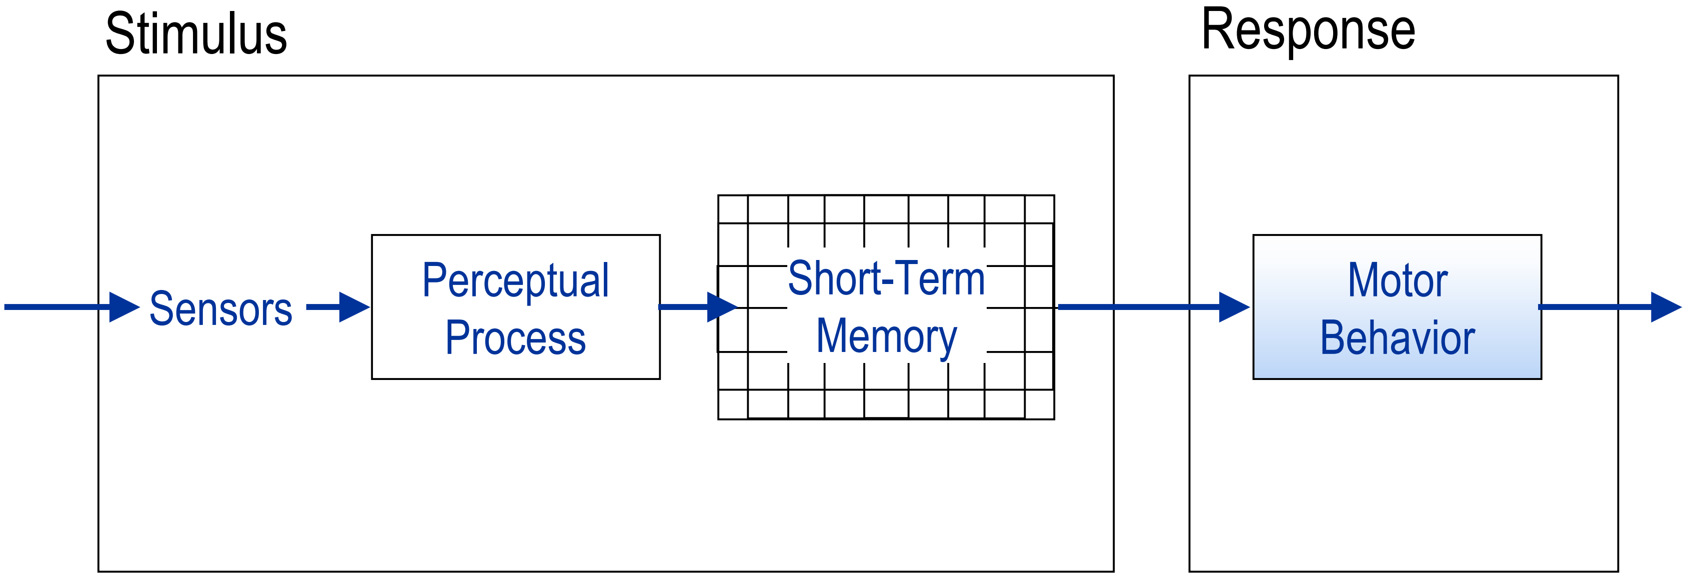

Stimulus

Response

Sensors

Perceptual Process

Short-Term Memory

Motor Behavior

Grid Representation

- Grid representation is common for behavioral memory
- Grids vary in the following ways:
  - Resolution: amount of area each grid unit covers
  - Shape: most frequently square, but could also be others, such as radial sectors
  - Uniformity: all grid cells same size, or size may vary.
    - Most common variable-sized grid methodology: quadtrees (recursive decompositions of free space)

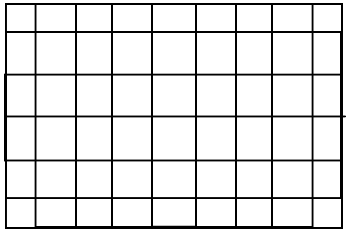

Regular grid

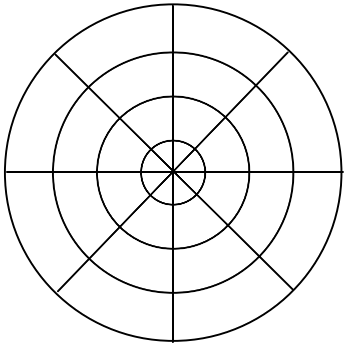

Sector grid

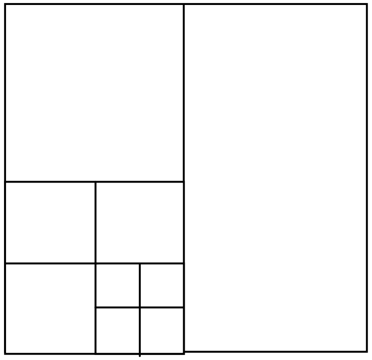

Quadtree

Long-Term Memory Maps

- Persistent information useful for advising behavioral control regime
- Origin of map:
  - From sensors onboard robot
  - From information gathered independently of robot (e.g., remote sensors)
- Typical encodings:
  - Metric: absolute measurements and coordinate systems used
  - Qualitative: salient features and their relationships (spatial or temporal) representated

Issues with Long-Term Memory Maps

- Disadvantages:
  - Data may be untimely (i.e., world changed)
  - Localization needs to be conducted (nontrivial)

- Advantage:
  - Can provide guidance beyond horizon of immediate sensing

Sensor-Derived Maps

- Provide information directly gleaned from robot's experiences in world
- Often advantageous to use qualitative representations instead of metric representations due to:
  - Inherent inaccuracies in robot motion and sensor readings
- Hallmark of qualitative navigational techniques:
  - Distinctive places: Regions of the world that have characteristics that distinguish them from their surroundings
    - E.g., symmetry, abrupt discontinuities in sensor readings, unusual constellations of sensor readings, point of maximum or minimum sensor reading
  - Once identified, can be used later for lower-level control
  - Can be easily integrated to behavior
    - E.g., "move forward until abrupt discontinuity occurs on right, then switch to a move-through-door behavior"

Examples of Distinctive Places

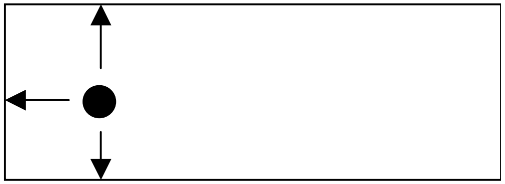

End-of-hall (3-way symmetry)

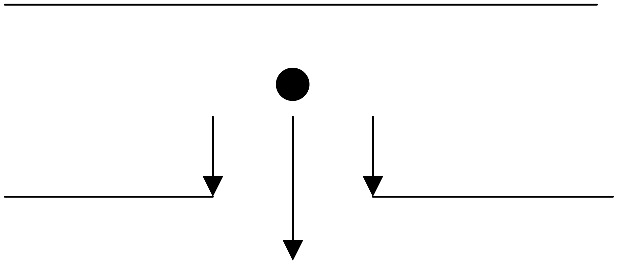

Doorway (abrupt depth discontinuity)

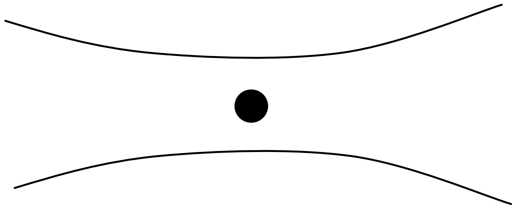

Hallway constriction (depth minimum)

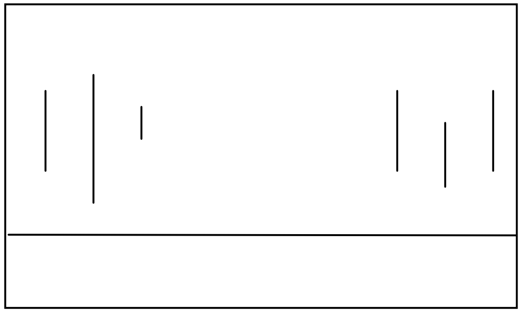

Visual constellations
(unique feature patterns)

Example of Qualitative Maps

- Landmarks:
  - Derived from sonar, using features that are stable and consistent over time
  - E.g., right walls, left walls, corridors
- Add spatial relationships connecting various landmarks via graph construction
- Subsumption-style approach (Mataric 1992):

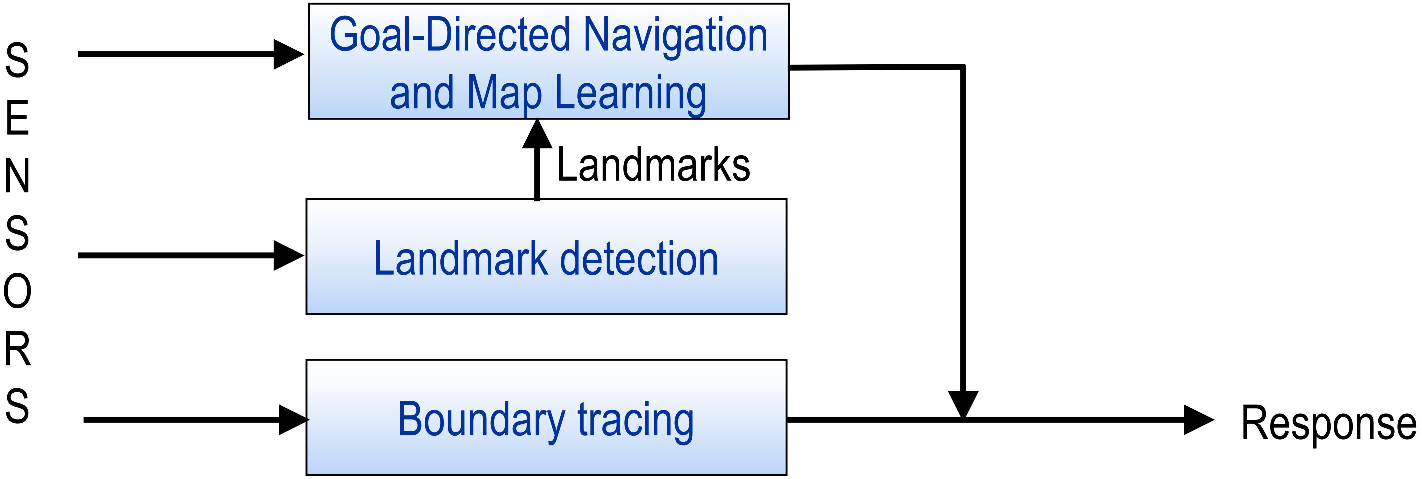

Another Example of Sensor-Derived Map

- Metric map: absolute distances given
- Created by sensor fusion of multiple sensor scans (e.g., laser)
- *(More discussion on this approach in mapping discussions later this term)*

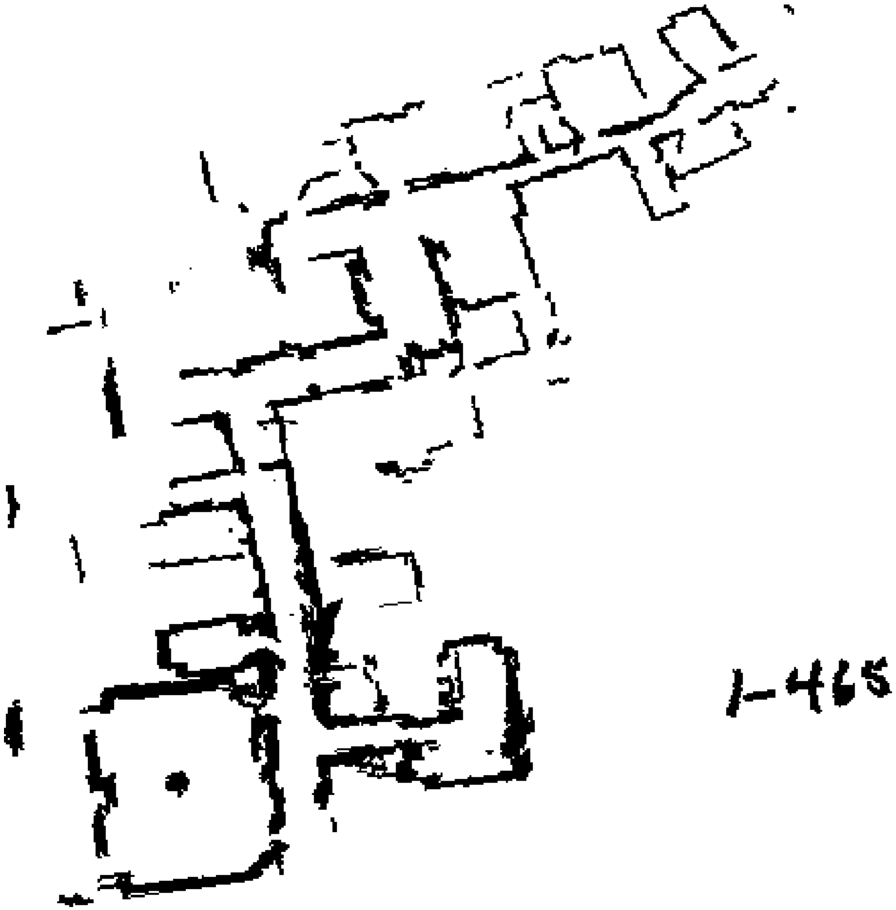

A Priori Map-Derived Representations

- Constructed from data obtained independently from the robotic agent itself.

- Reasons for using this type of map:
  - May be easier to compile data directly without forcing robot to travel through entire world ahead of time
  - May be available from standard sources such as Defense Mapping Agency or U.S. Geographical Survey, etc.
  - Precompiled sources of information may be used (e.g., blueprints, floorplans, roadmaps, etc) that only need to be encoded for robot's use

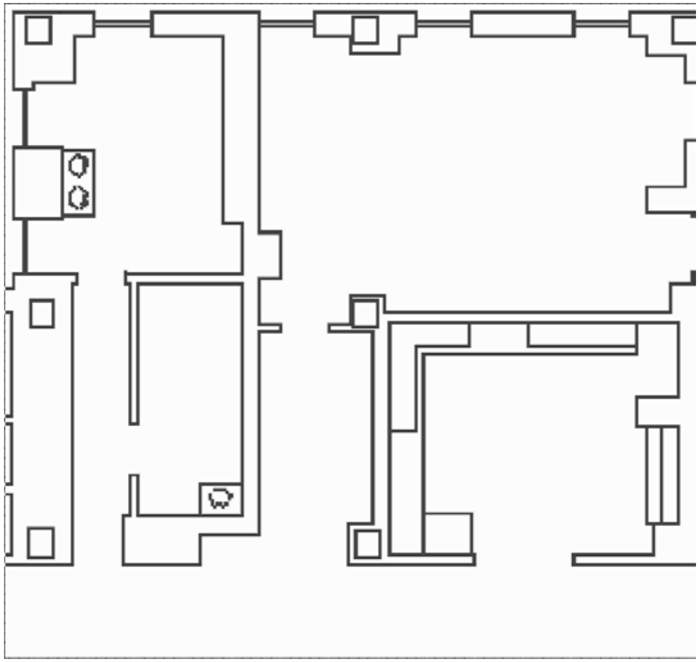

Example *a priori* map: building floor plan

Example of A Priori Maps: Internalized Plans

- Map of environment containing known obstacles, terrain info, goal location provided in a grid-based format from a digital terrain map

- Cost associated with each grid cell based on mission criteria, e.g.:
  - Traversability
  - Visibility
  - Ease of finding landmarks
  - Impact on fuel consumption, etc.

- Gradient field computed over entire map from start point to goal point with minimum cost direction represented within each cell to get to goal

- Gradient field represents internalized plan, since it contains preferred direction of motion to accomplish the mission's goals

Example of Internalized Plan (Payton 1991)

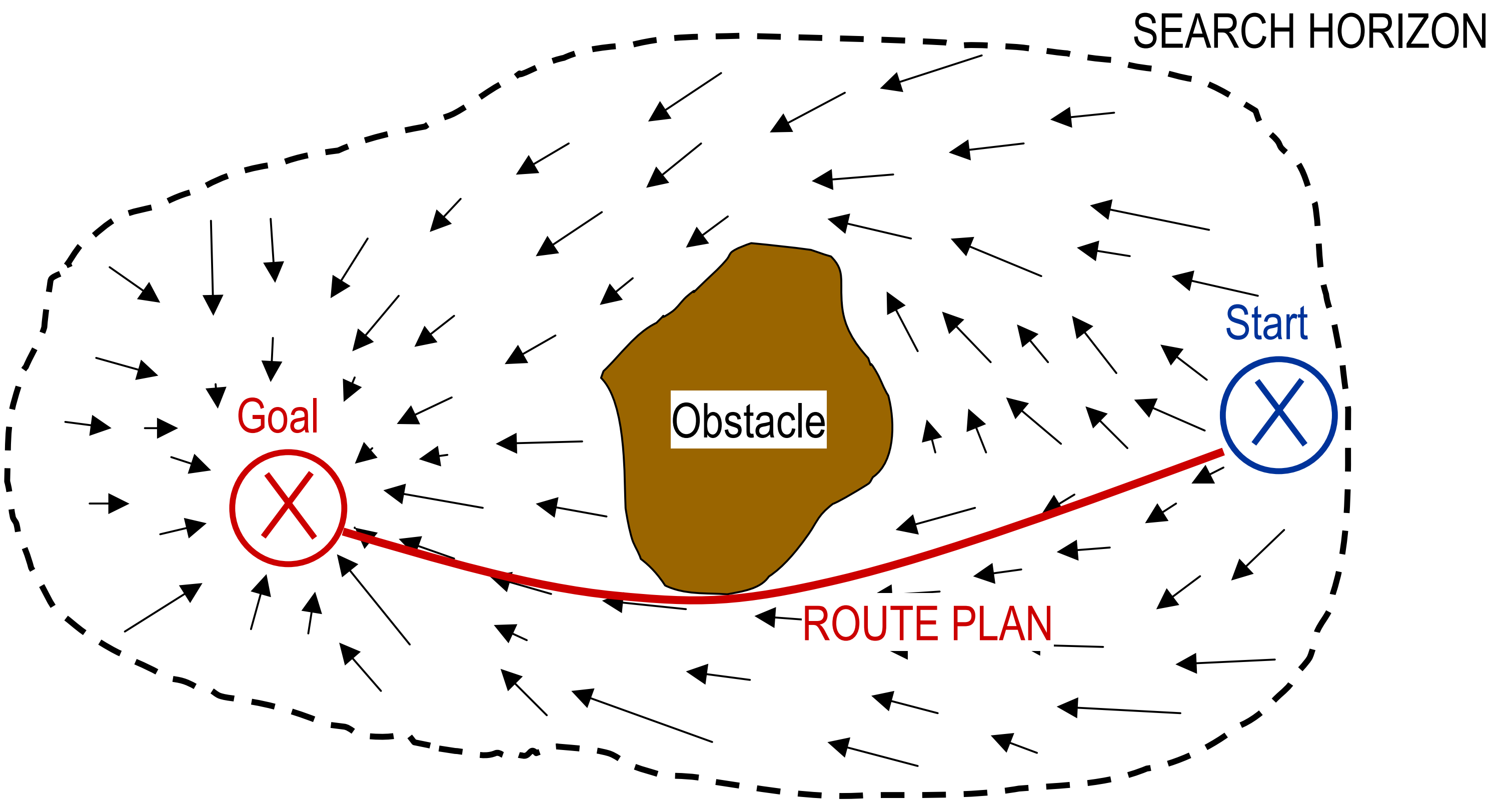

Behavior Control Using Internalized Plans

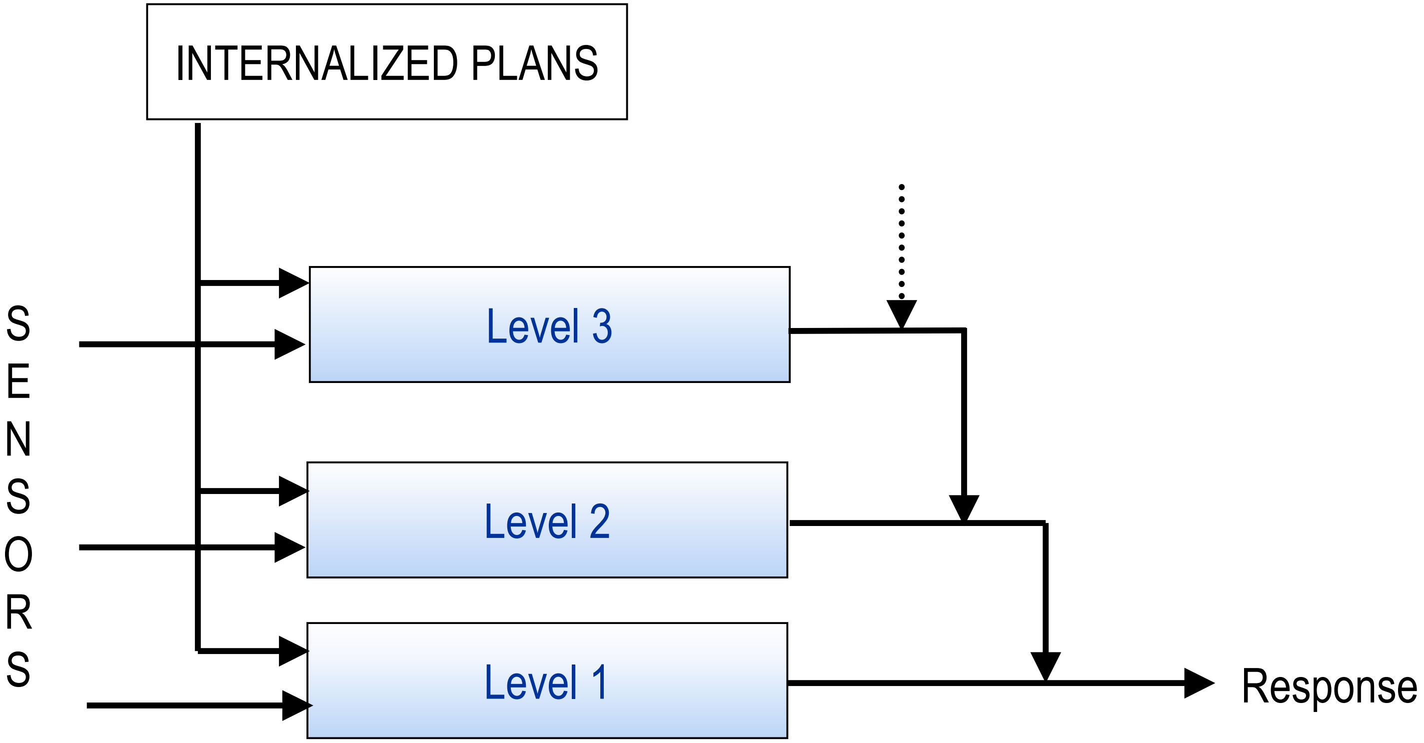

Summary of Representational Issues for Behavioral Systems

- The more predictable the world is, the more useful knowledge representations are
- Two important characteristics of knowledge include its predictive power and the need for the information stored to correlate with the environment in some meaningful way
- Knowledge can be characterized in three primary forms:
  - Explicit
  - Implicit
  - Tacit
- Knowledge can be further characterized according to its temporal durability:
  - Transitory
  - Persistent

Summary of Representational Issues (con't.)

- Using representational knowledge has several potential drawbacks within behavior-based systems:
  - Stored information may be inaccurate or untimely
  - Robot must localize itself within the representational framework for the knowledge to be of value
- Representational knowledge's primary advantage lies in its ability to inject information beyond robot's immediate sensory range into the robotic control system
- Examples of explicit representational knowledge:
  - Short-term behavioral memory
  - Sensor-derived maps
  - A priori map-derive representations

Summary of Representational Issues (con’t.)

- Short-term behavioral memory: extends behavioral control beyond the robot’s immediate sensing range, and reduces demand for frequent sensory sampling
- Grid-based representations often used for short-term behavioral memory
- Long-term maps are either metric or qualitative
- Notion of distinctive places is central to use of sensor-derived maps.
- A priori map-derived representations offer robot information regarding places where it has never been before.
- Internalized plans inject a priori grid-based map knowledge directly into a behavior-based control system.

Preview of Next Class (Tuesday, Oct. 15th)

- Hybrid deliberative/reactive systems