Multiple Spatial Ontologies in Humans and Robots

Benjamin Kuipers
Computer Science & Engineering
University of Michigan

Human Cognitive Maps

- Humans use multiple representations
  - Topological maps of large-scale space
    - Metrical errors and distortions are common.
    - Topological errors are rare.
  - Some spatial knowledge is metrical.
    - Multiple frames of reference
  - Individual variation is everywhere:
    - with developmental stage
    - with experience in an environment
    - with individual cognitive style
  - Lynch, 1960; Piaget & Inhelder, 1967, . . .

Inspiration for Computational Models

- A computational model must be *sufficient* to produce the behavior it hopes to explain.
  - Therefore, it must have multiple representations.
  - It must also be capable of learning a cognitive map from observations, and using it to navigate.

- Knowledge of space must be grounded in perception and action.
  - A computational model of mind, including perception and action, is by definition a *robot*.

Scales and Ontologies of Space

- Distinguish *scales of behavioral space*.
  - **Small-scale space**
    - Within the agent's sensory horizon
  - **Large-scale space**
    - Beyond the agent's sensory horizon
- Distinguish *ontologies for spatial maps*.
  - **Metrical mapping**:
    - Within a single frame of reference, define location, heading, pose, distance, and shape.
  - **Topological mapping**:
    - Places, paths, and regions are related by connectivity, order, and containment.

Spatial Semantic Hierarchies

- **The Basic SSH**:
  - Even without knowledge of sensors, hill-climbing control laws and distinctive states can define places, leading to topological maps.
- The Hybrid SSH:
  - Often, we do know what the sensors are sensing.
  - Use well-understood local mapping methods, and build the place abstraction and topological maps on top of that.

Distinctive States

- A *distinctive state* (location plus orientation) is the isolated fixed-point of a hill-climbing control law.

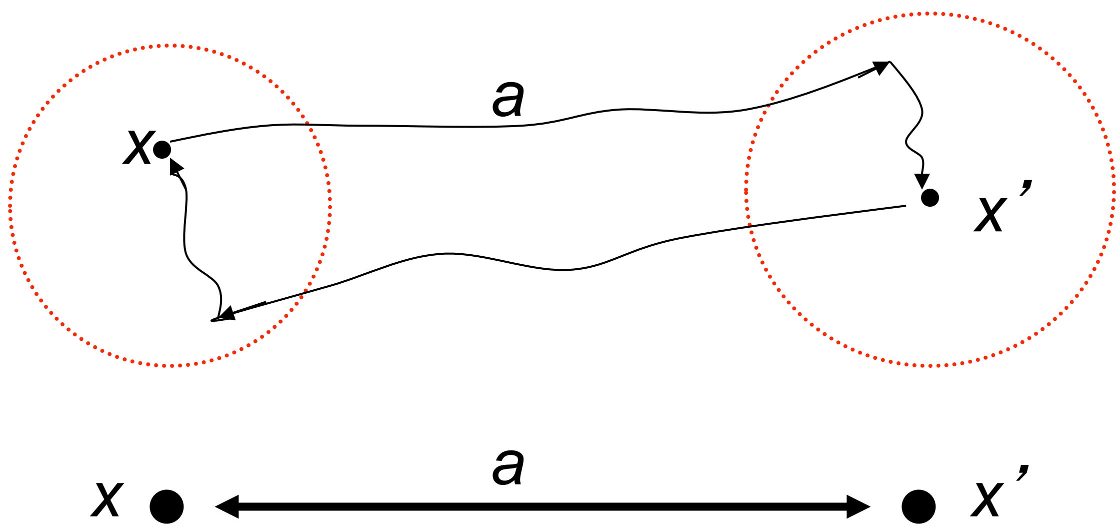

  – *High-level concepts* (places) *can be abstracted from the behavior of low-level control laws, which operate at the pixel level.*

Distinctive States

- Between distinctive states, actions are *functionally deterministic*
  - if all final-state uncertainty is contained within every initial-state basin of attraction.
- Supports abstraction from continuous to discrete state space.
  - Hill-climbing eliminates cumulative position error.

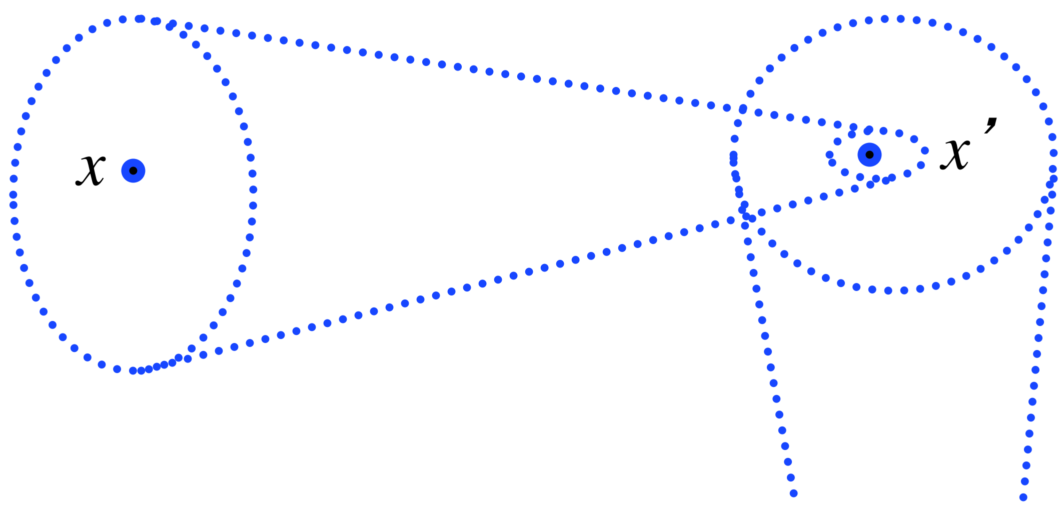

Topological Abstraction

- A control law defines an attractor
  - that represents its basin of attraction

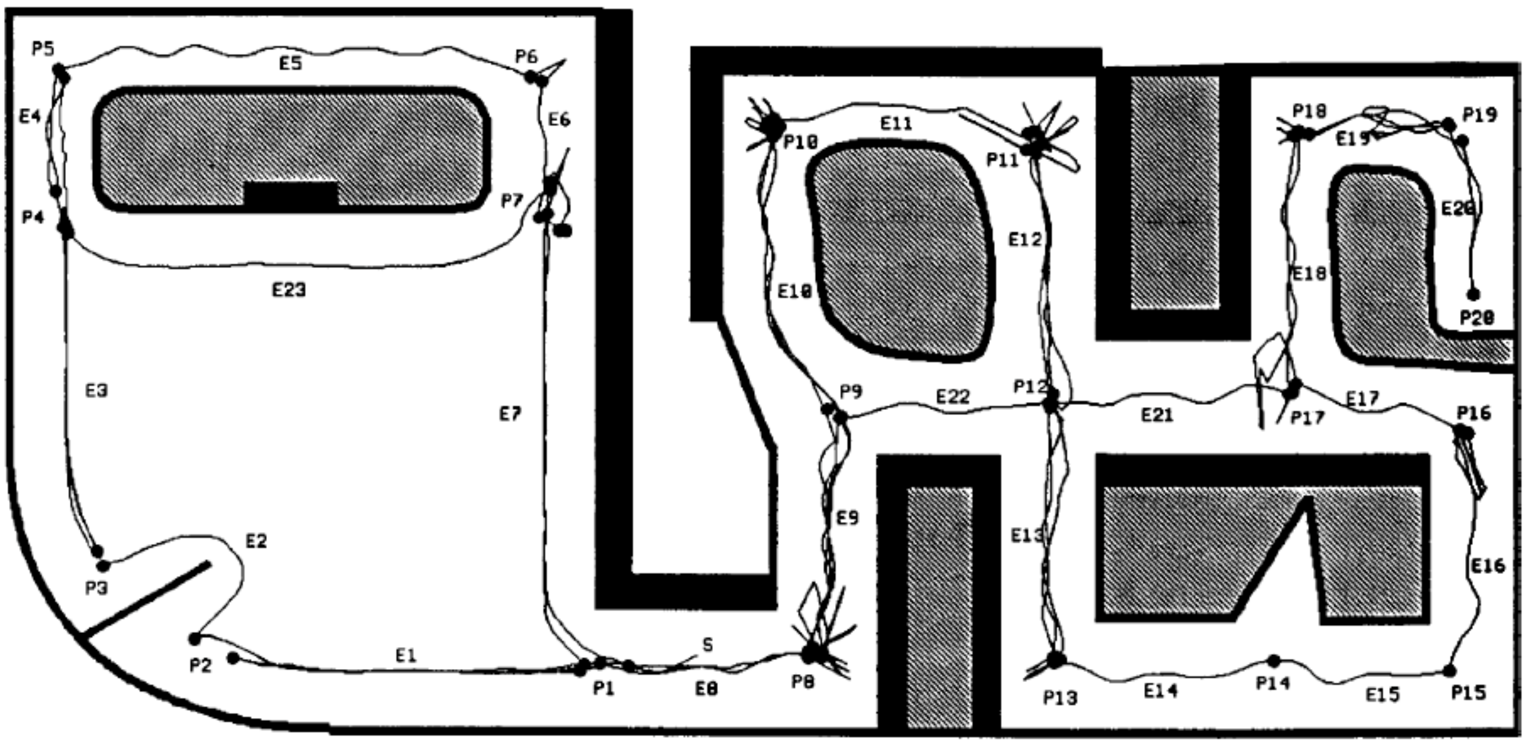

[Kuipers & Byun, JRAS, 1991]

Topological Abstraction

- A small, finite graph concisely represents the structure of behaviors in a continuous space.

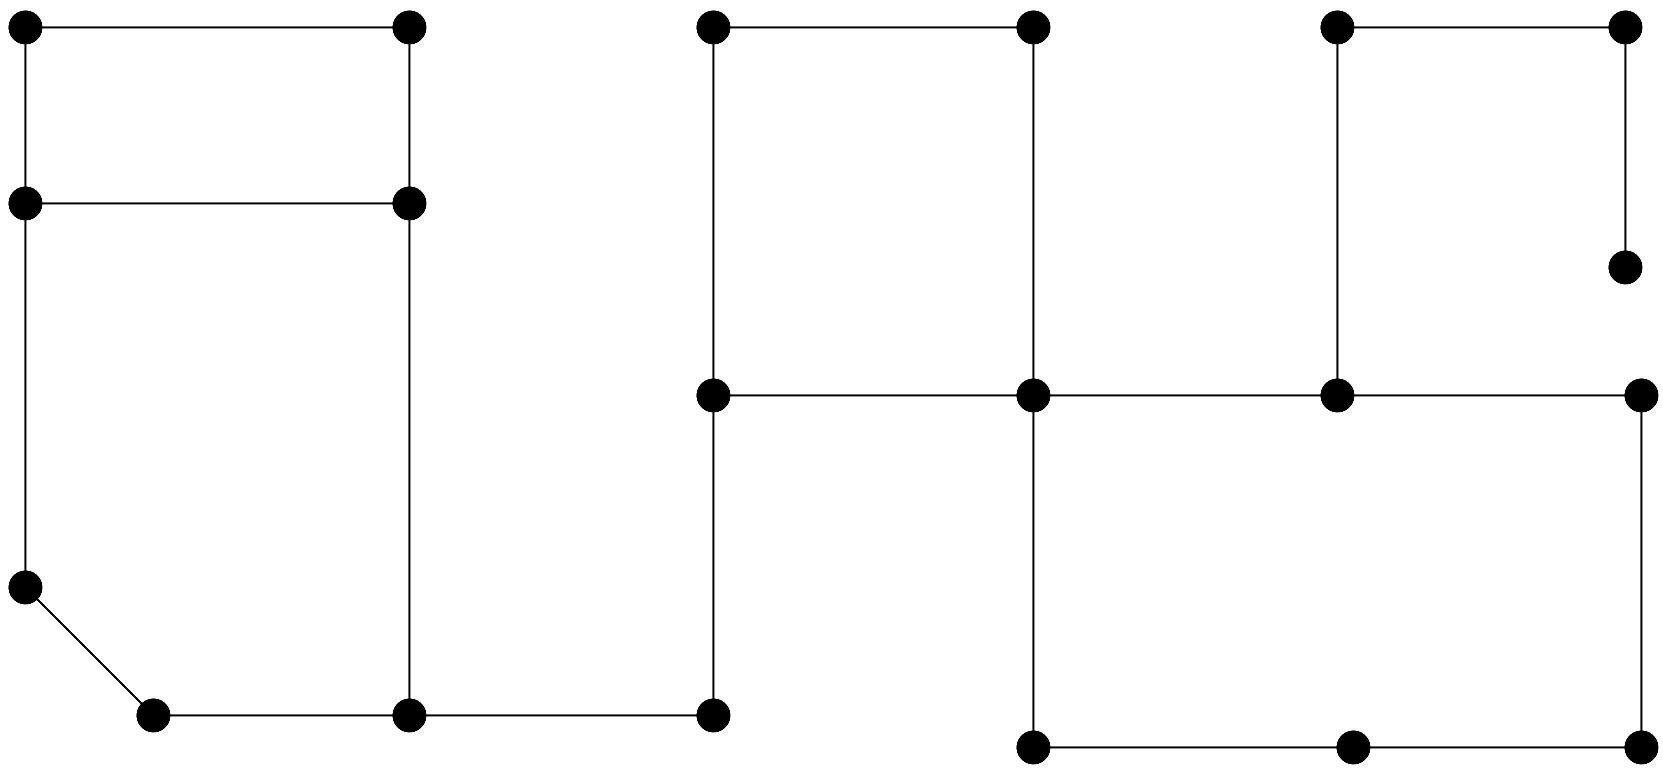

Spatial Semantic Hierarchies

- The Basic SSH:
  - Even without knowledge of sensors, hill-climbing control laws and distinctive states can define places, leading to topological maps.
- **The Hybrid SSH**:
  - Often, we *do* know what the sensors are sensing.
  - Use well-understood local mapping methods, and build the place abstraction and topological maps on top of that.

Local Metrical Mapping Works

- In small-scale space, modern laser-based SLAM methods work extremely well.
  - Great progress with visual SLAM, too.

	Metrical Mapping	Topological Mapping
Small-scale space	Local SLAM	
Large-scale space		

Global Metrical Mapping Is Hard

- Within a single global frame of reference over large-scale space, errors accumulate.
  - Sufficiently large loops can always be a problem.

	Metrical Mapping	Topological Mapping
Small-scale space	Local SLAM	
Large-scale space	Cumulative errors Scalability	

Problem: Closing Large Loops

Raw Odometry

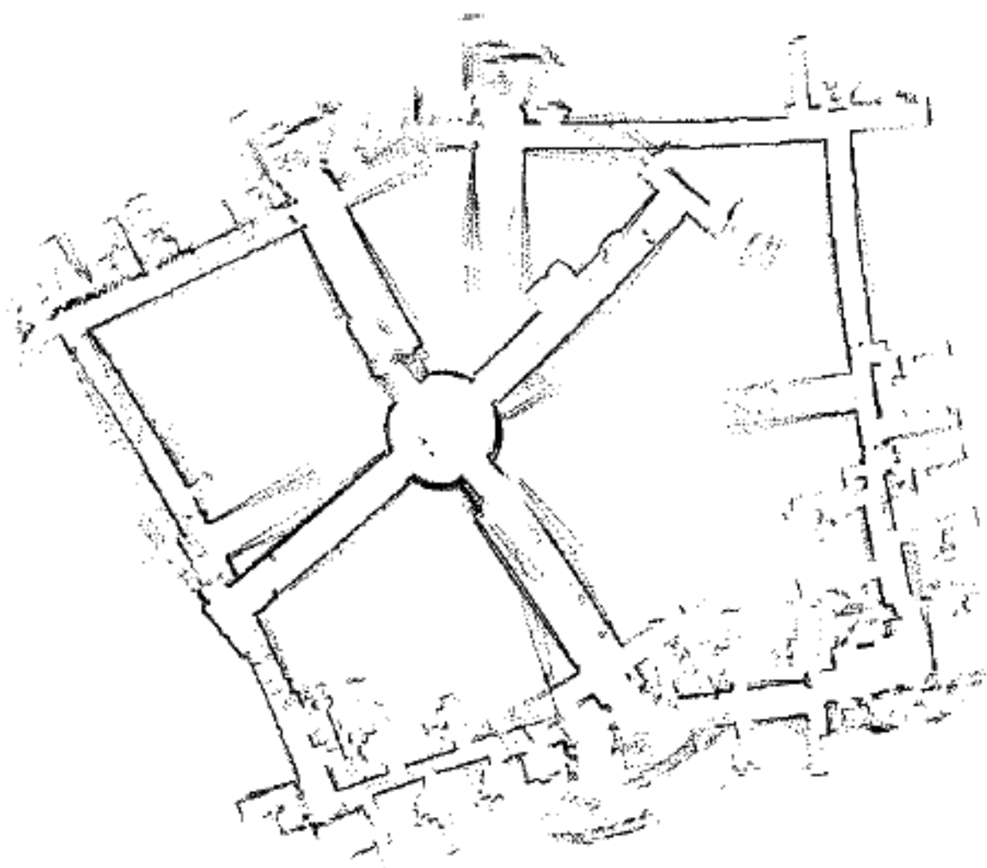

SLAM Corrected Odometry

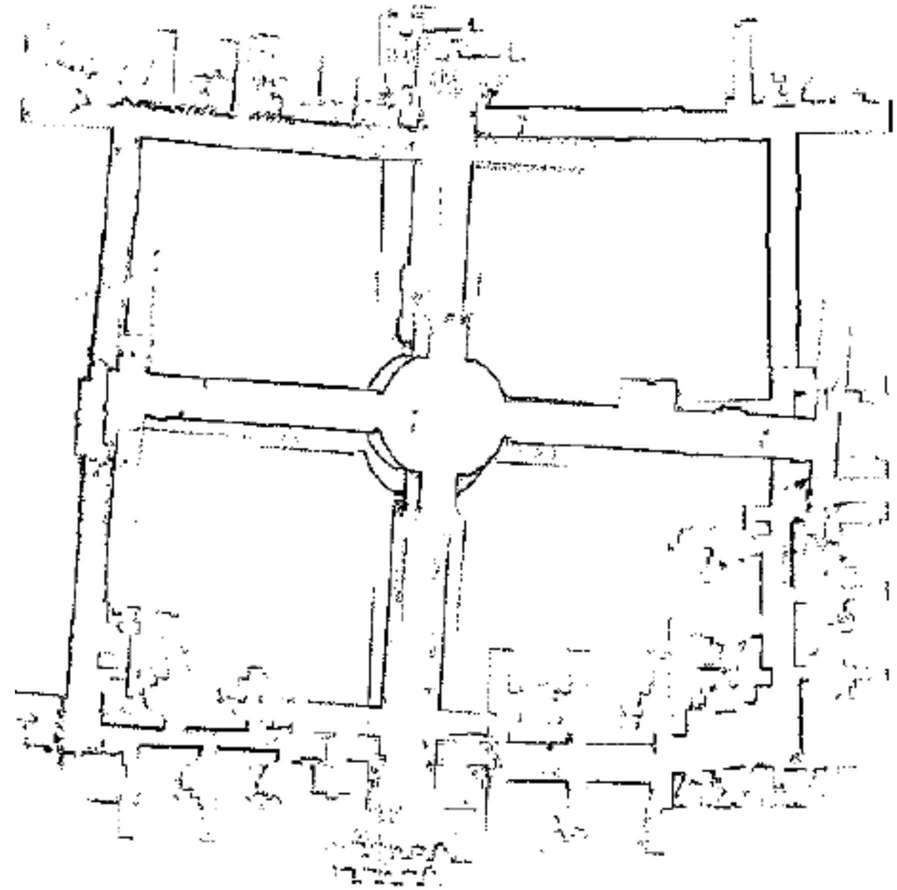

Local matching can find false, but locally optimal, loop closures

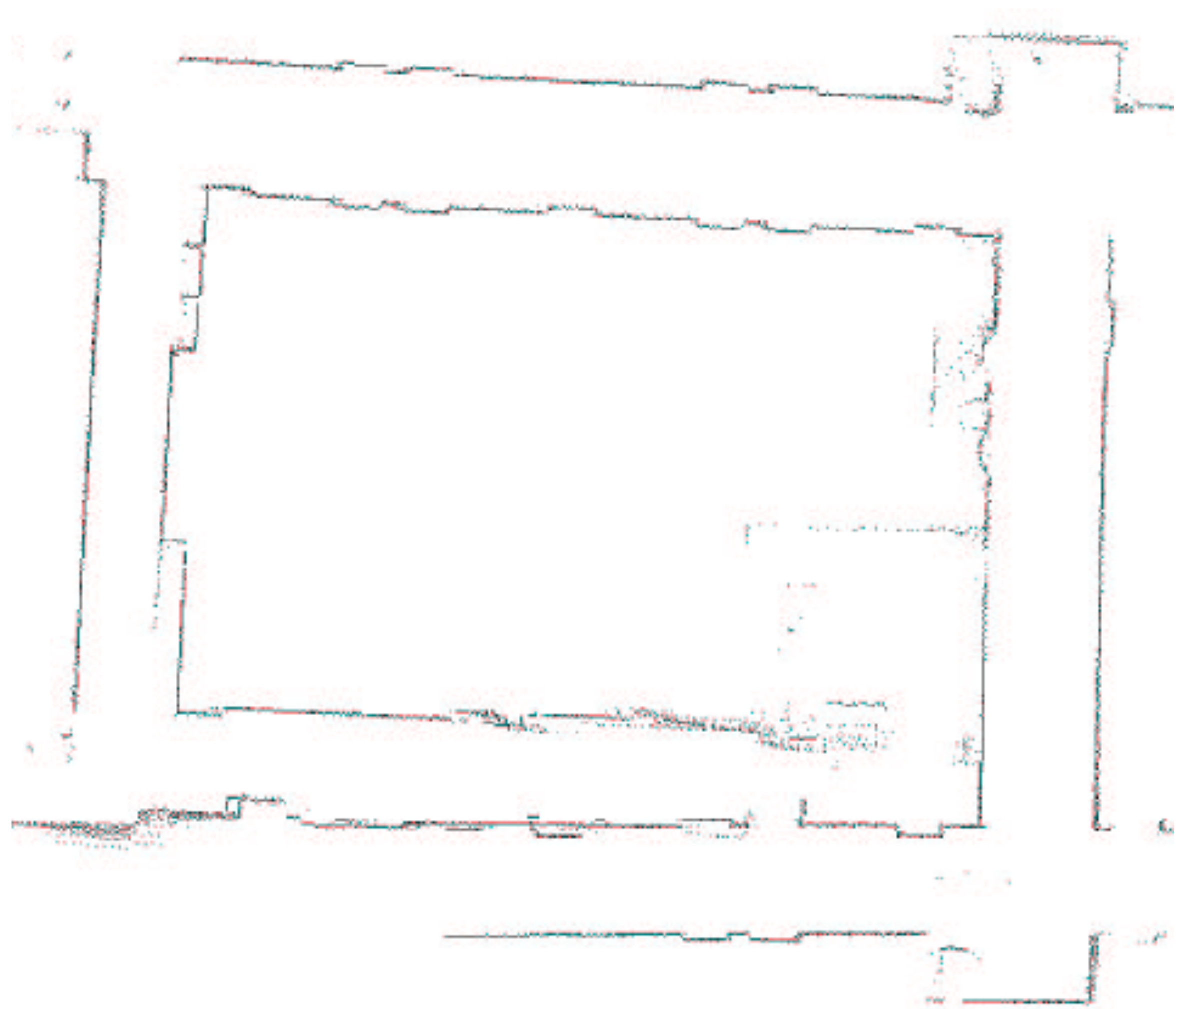

The Local Perceptual Map (LPM)

- **Local SLAM** in a bounded, fixed-sized map
  - The LPM scrolls keeping the agent near center.
  - Incremental update has $O(1)$ complexity.
  - The local map includes no “large loops”.

- The LPM is useful for:
  - Planning safe and comfortable local motion
  - Avoiding collisions with static and dynamic obstacles
  - Analyzing qualitative local decision structure in a place neighborhood.

Exploration and Navigation

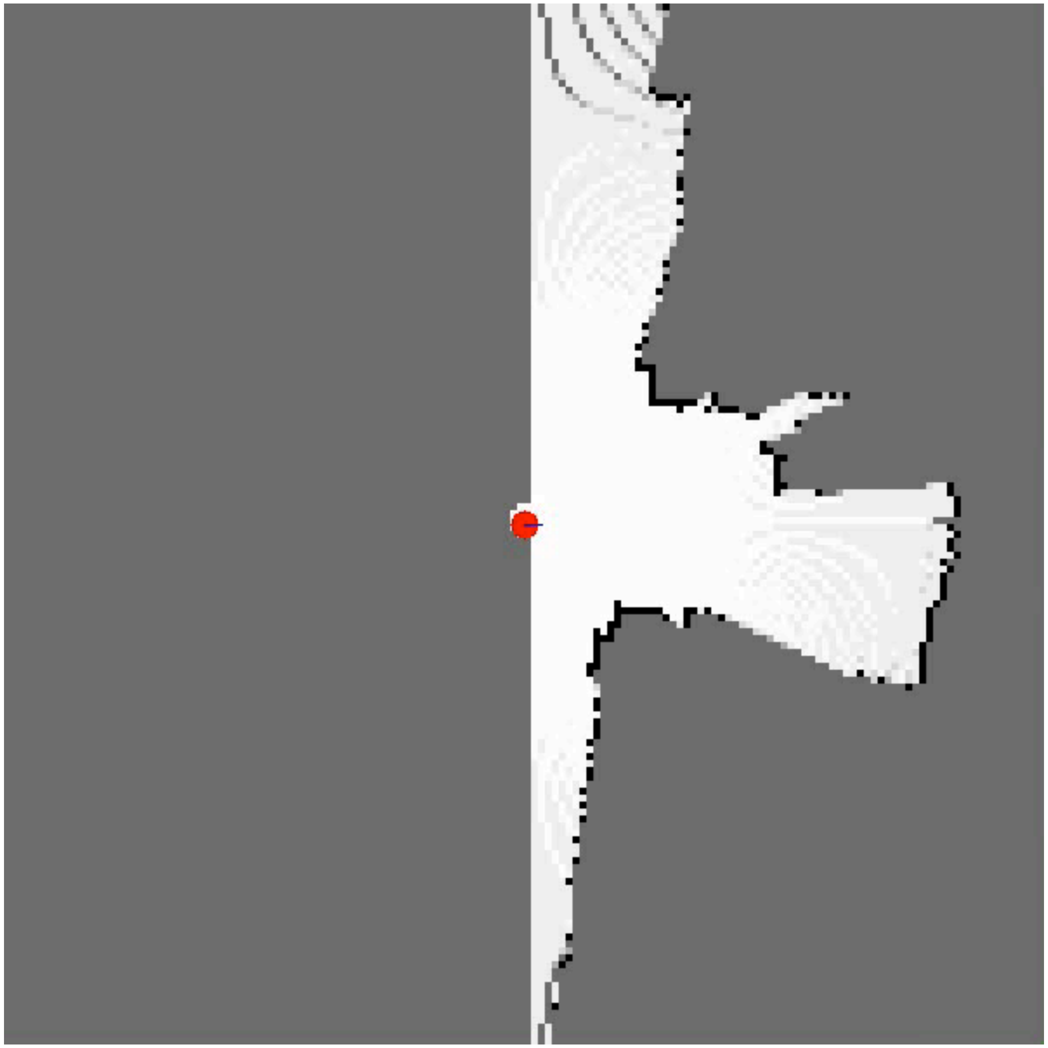

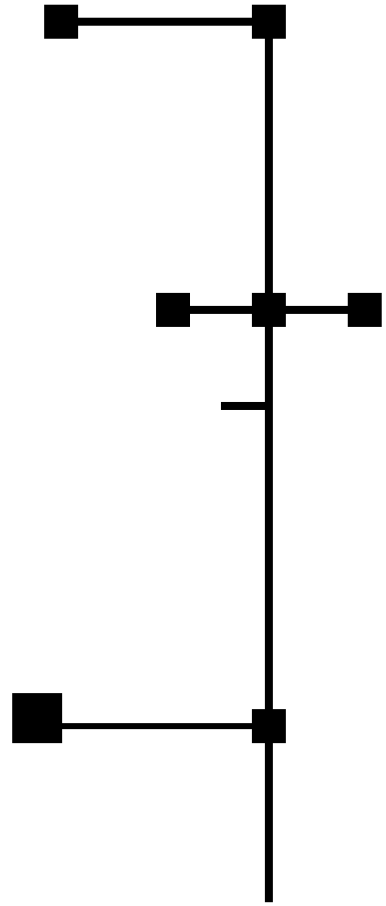

Identify the Local Topology

- Identify the local decision structure of each place neighborhood.
  - Travel experience as graph exploration

	Metrical Mapping	Topological Mapping
Small-scale space	Local SLAM →	Local decision structure
Large-scale space		

Local Decision Structure

- Identify *gateways* and *path fragments*
  - 2 gateways & 1 path fragment ⇒ on a path
  - Otherwise ⇒ at a place neighborhood

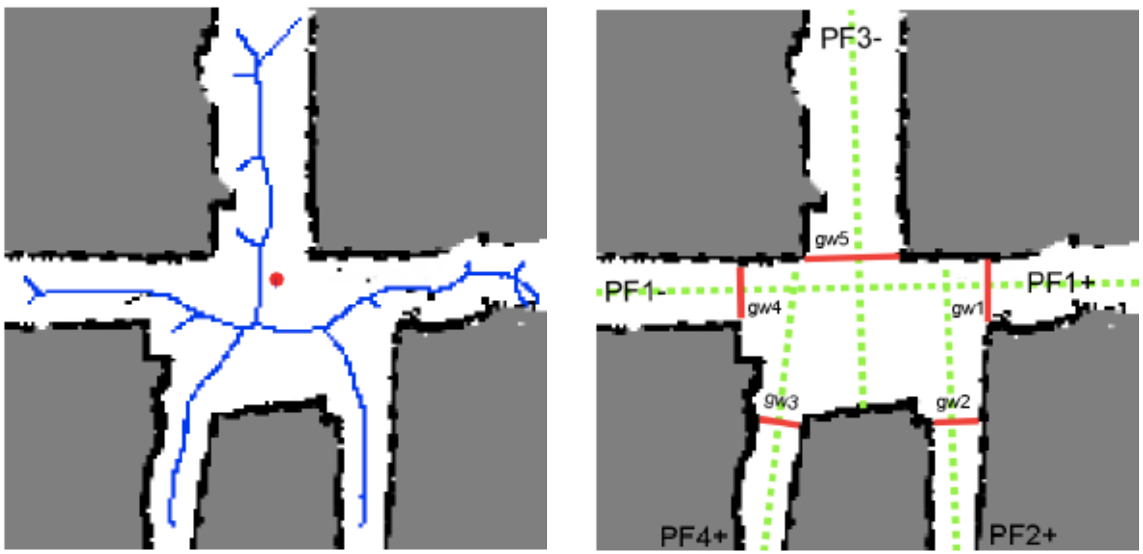

in small-scale space

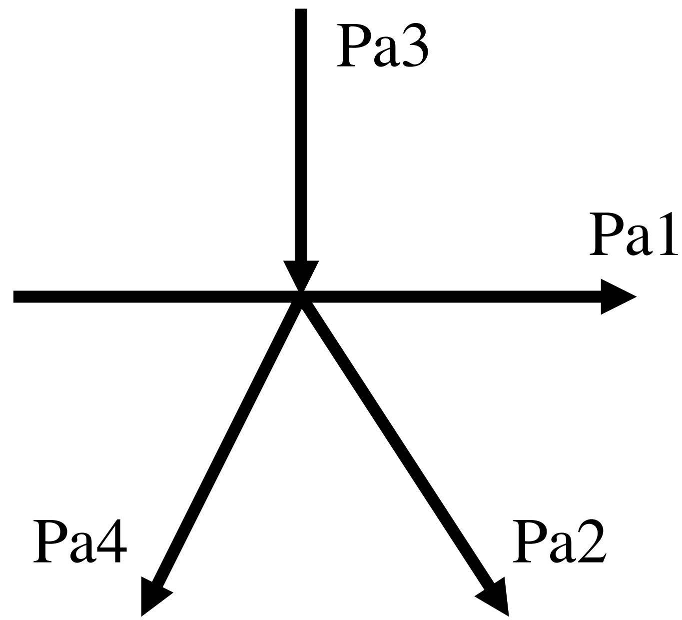

in large-scale space

Gateways

- A gateway is a transition between a *travel action* and a *place neighborhood*
  - i.e., between a trajectory-following control law and a local perceptual map.
  - Transitions can be *inbound* or *outbound*.
  - Gateways are detected from local properties of the environment and the conditions on the control law.

Detect and Describe a Place

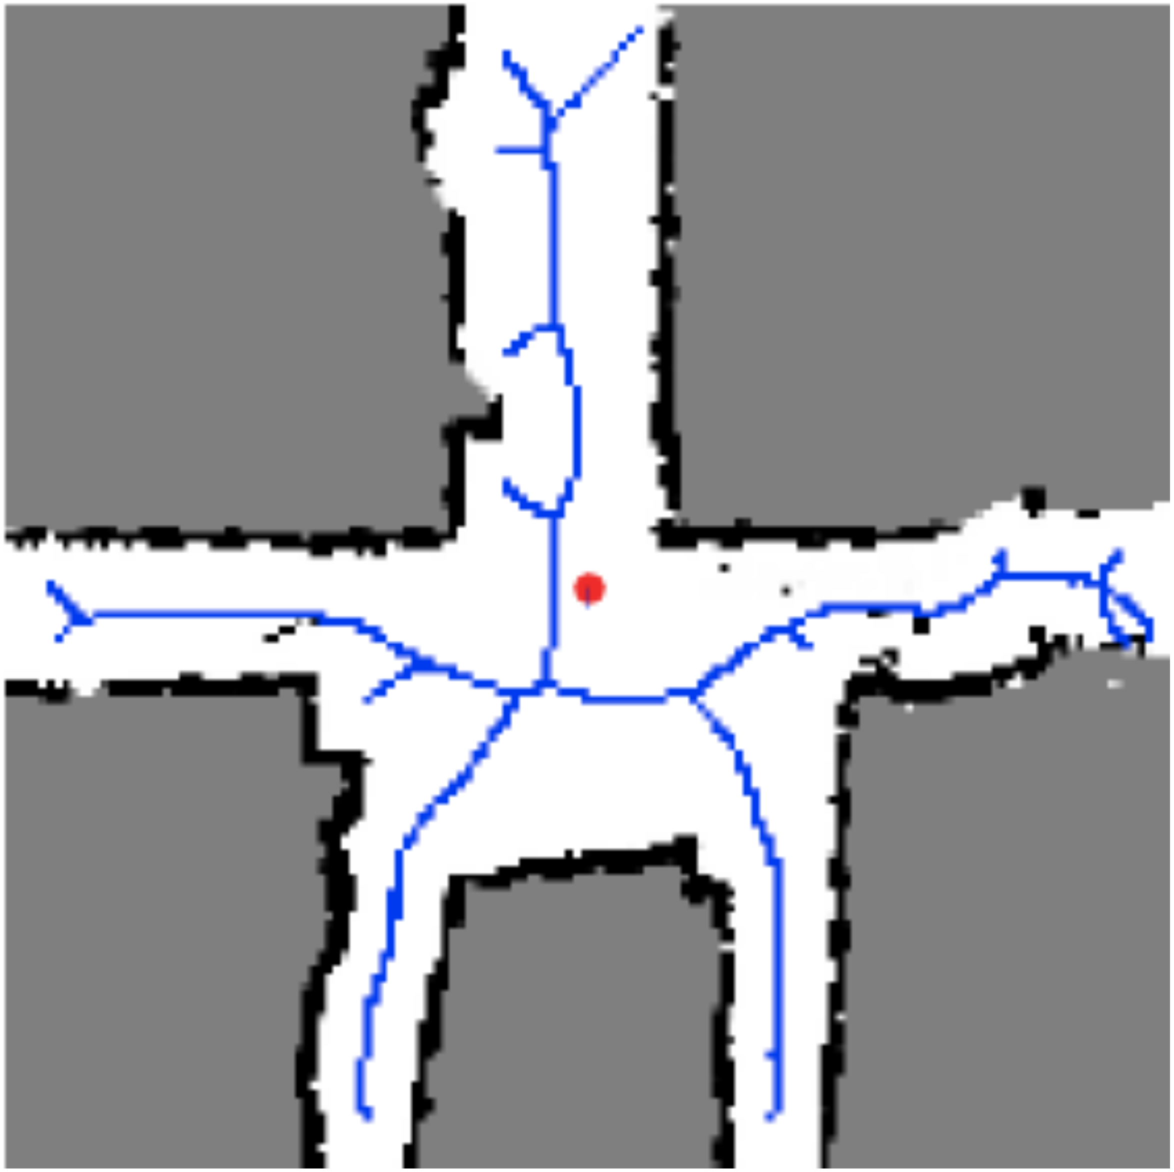

Identify Constrictions

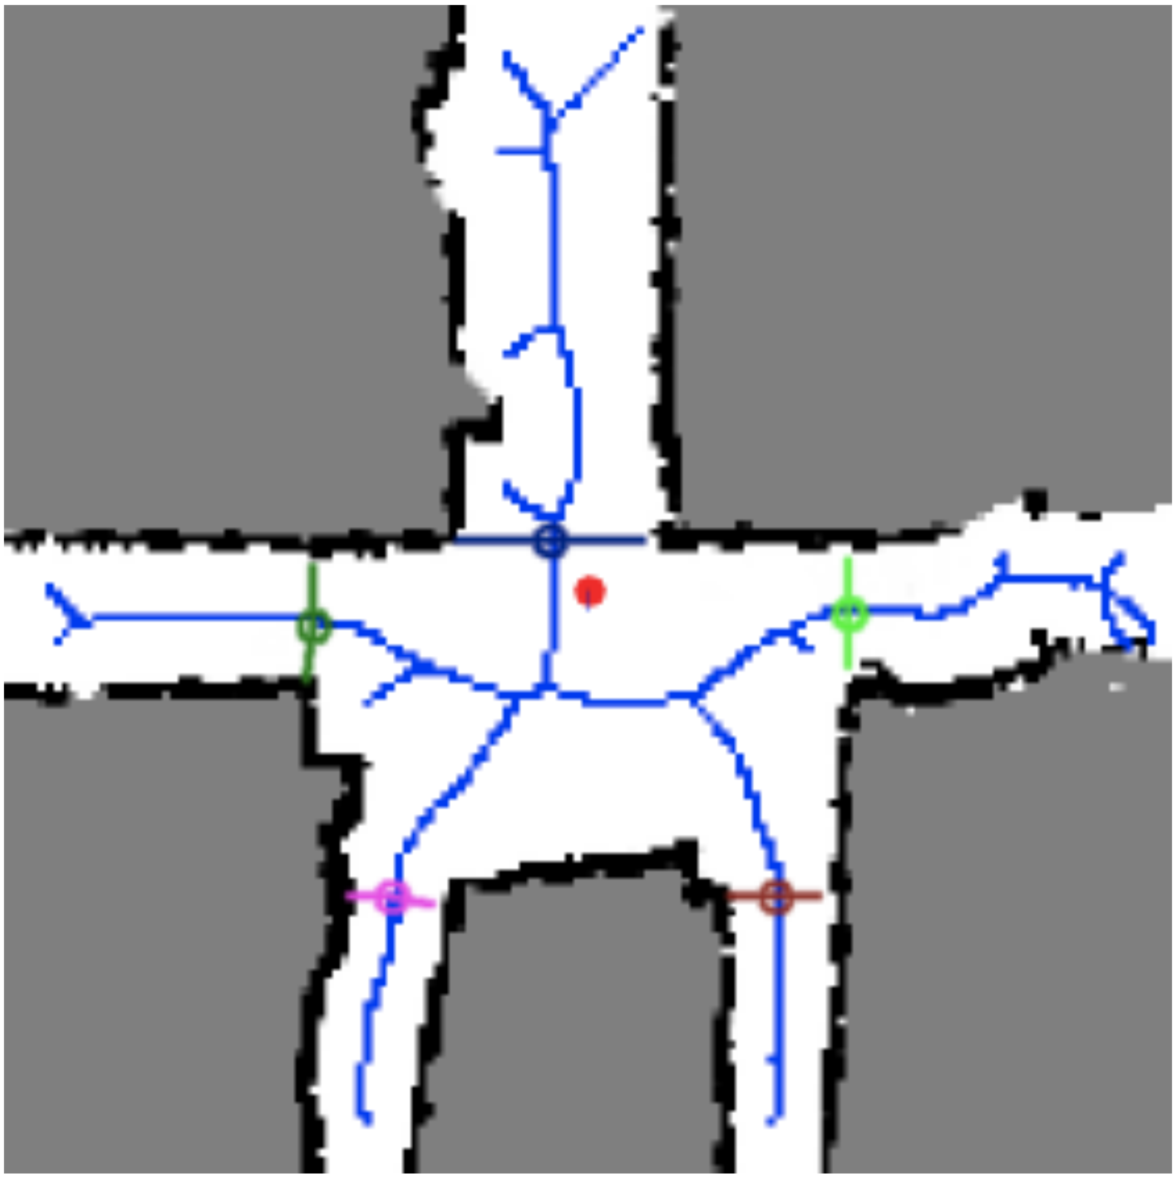

Define Gateways

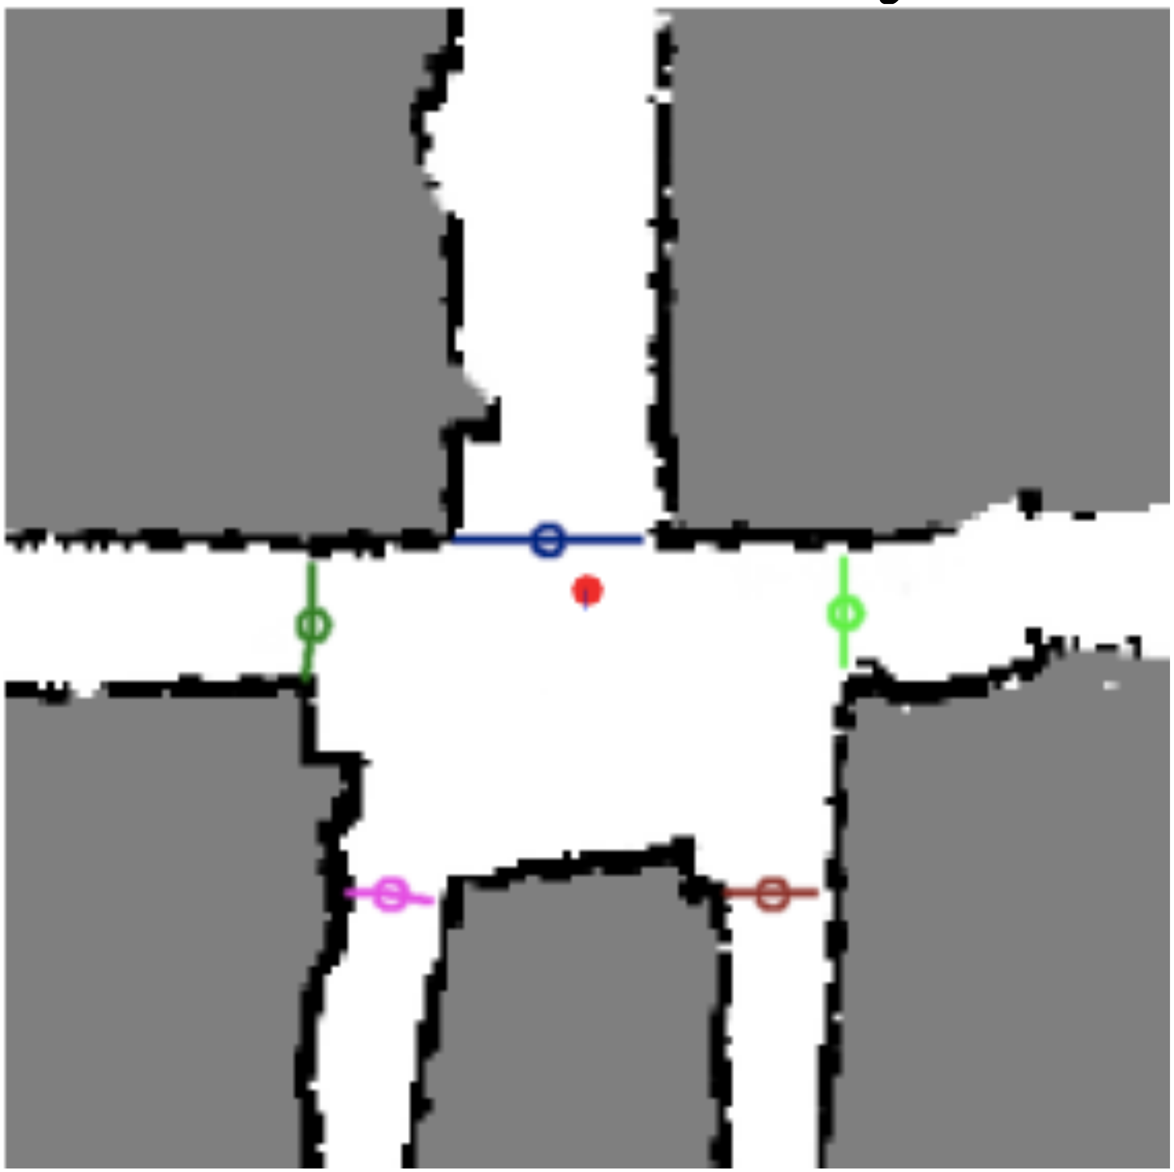

Define Local Path Fragments

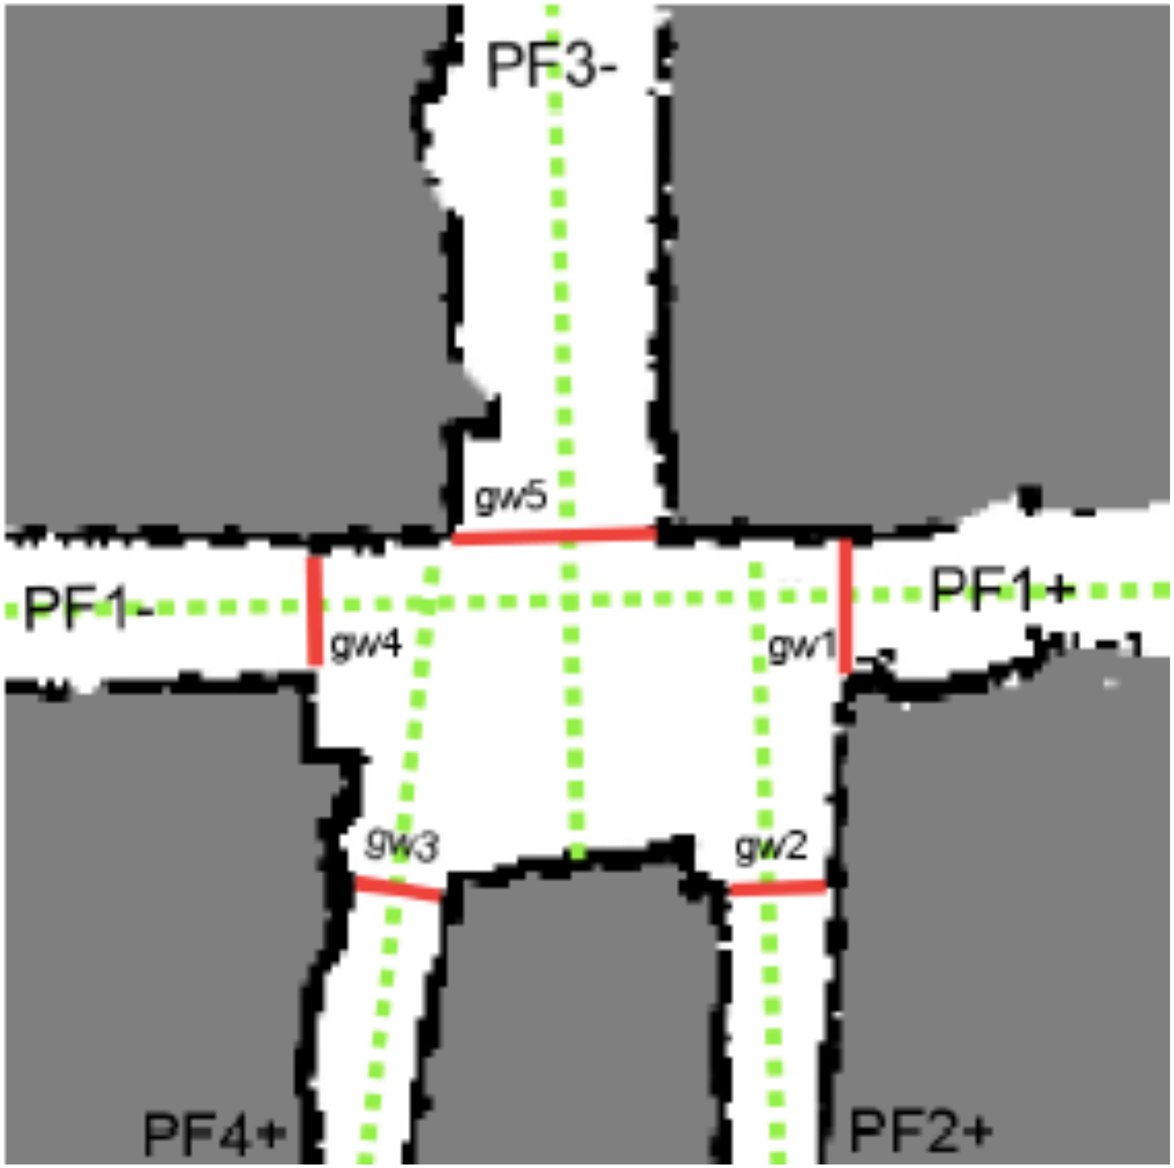

Local Topology Description

- The *small-scale star* is a circular order of path fragments, gateways, and control laws.

PF1+	(gw1,out) & (gw4,in)	Midline
PF2+	(gw2,out)	Midline
PF3+	(gw5,in)	DeadEnd
PF4+	(gw3,out)	Midline
PF1-	(gw4,out) & (gw1,in)	Midline
PF4-	(gw3,in)	DeadEnd
PF3-	(gw5,out)	Midline
PF2-	(gw2,in)	DeadEnd

Local Topology Description

- The *large-scale star* describes the place with distinctive states and directed paths.

ds1	Pa1, +	
ds2	Pa2, +	
ds3	Pa3, +	Endpoint
ds4	Pa4, +	
ds5	Pa1, -	
ds6	Pa4, -	Endpoint
ds7	Pa3, -	
ds8	Pa2, -	Endpoint

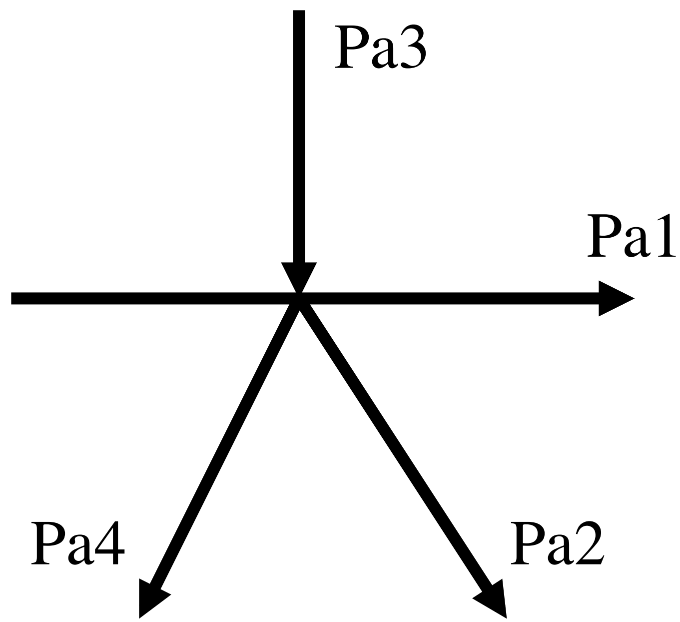

Decision Structure Abstraction

- A Turn action follows a trajectory through the local place neighborhood.

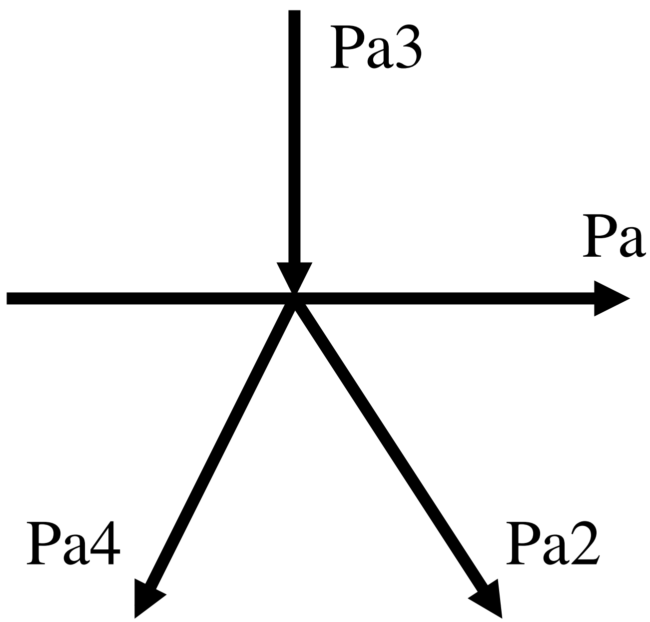

in large-scale space

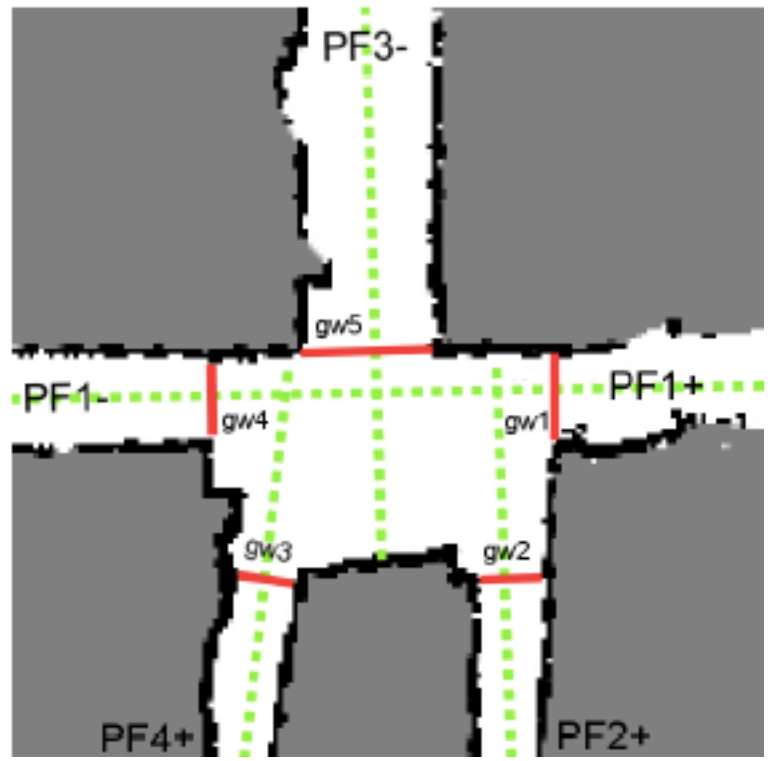

in small-scale space

Does a place abstraction always exist?

- Not in truly pathological environments
  - open ocean (but what about the Puluwat navigators?)

  or with pathological sensors
  - video snow

- **Conjecture: Yes**, with sufficiently rich sensors in a sufficiently rich environment.
  - office environments
  - campus/urban indoor/outdoor environments

Build the Global Topological Map

- Decide when and how loops are closed
  - When does the next place match a previous place?
- Build a tree of all possible topologies

	Metrical Mapping	Topological Mapping
Small-scale space	Local SLAM	Local decision structure
Large-scale space		Global topological map

Build the Global Topological Map

- Define a tree of *all possible* topological maps consistent with exploration experience.
  - They are the leaves of this tree.
- For each new action+observation
  - If the map predicts the observation, *OK*.
  - If it contradicts the observation, *prune it*.
  - Otherwise, *branch* on maps with new edges:
    - All possible loop-closing hypotheses
    - One hypothesis of a brand-new place
  - Identify the current best map.

Building the Tree of Maps

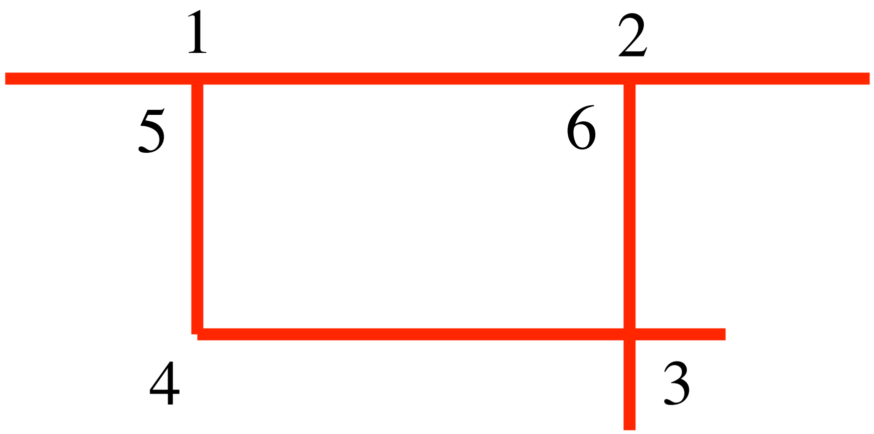

Tree of Maps (1)

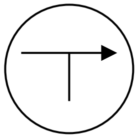

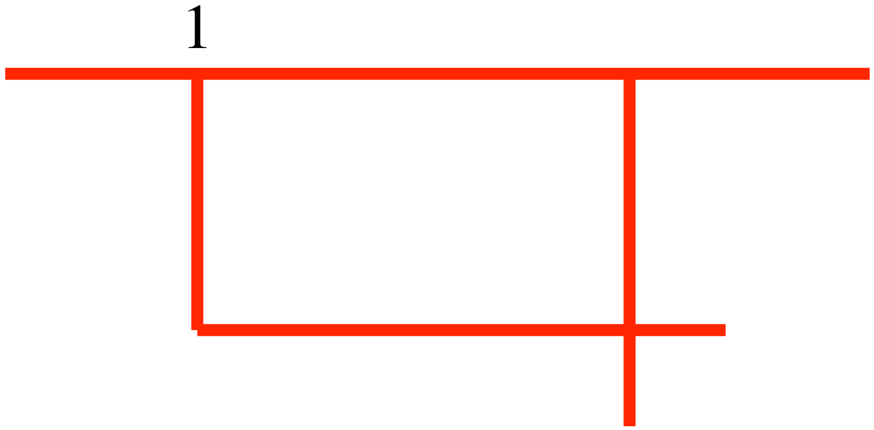

Tree of Maps (2)

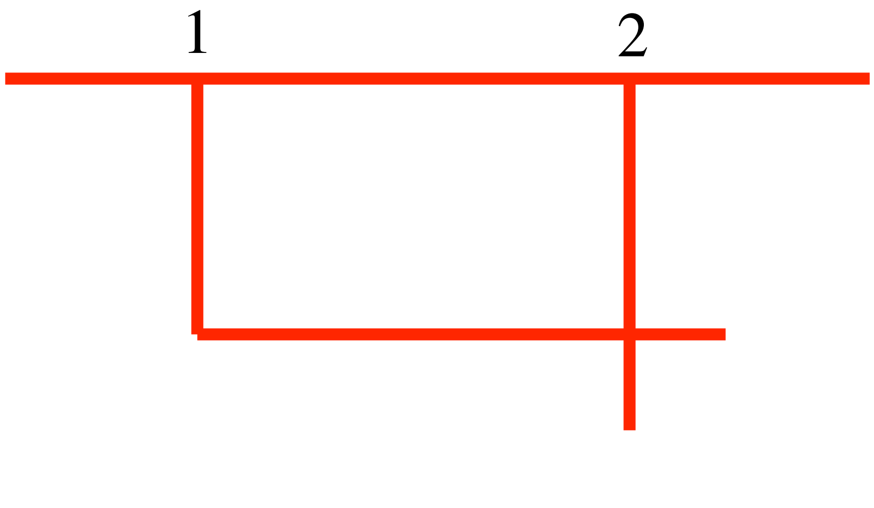

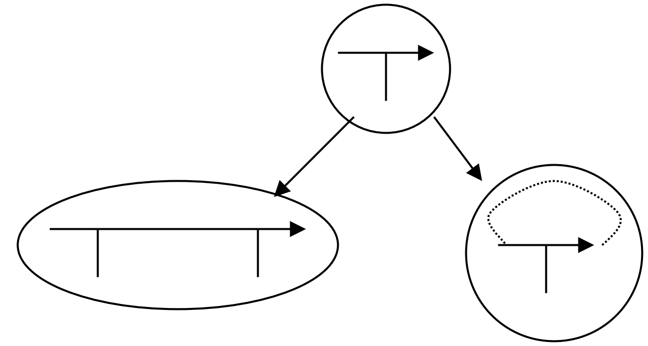

Tree of Maps (3)

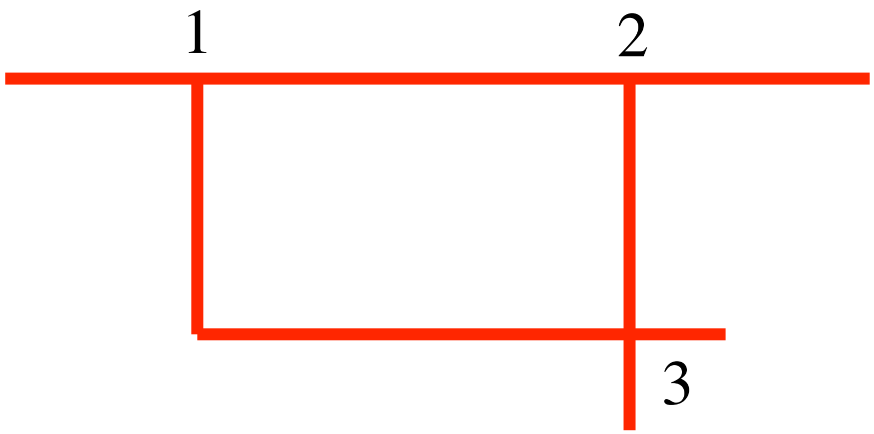

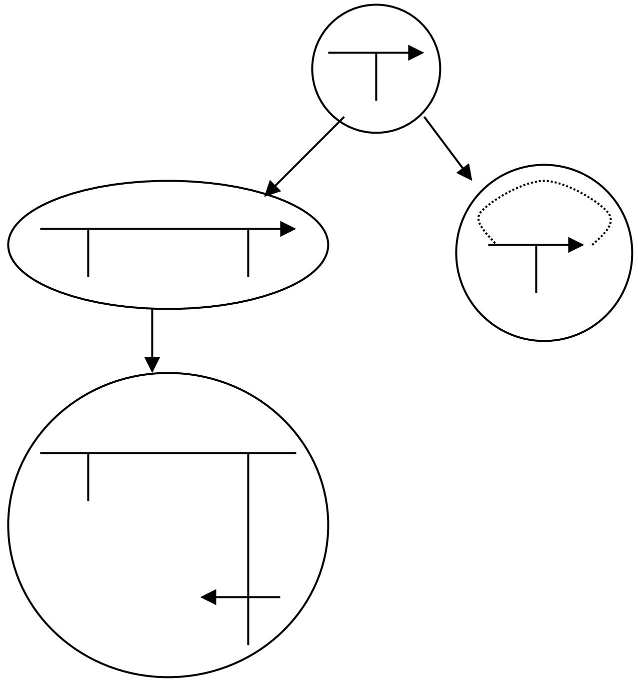

Tree of Maps (4)

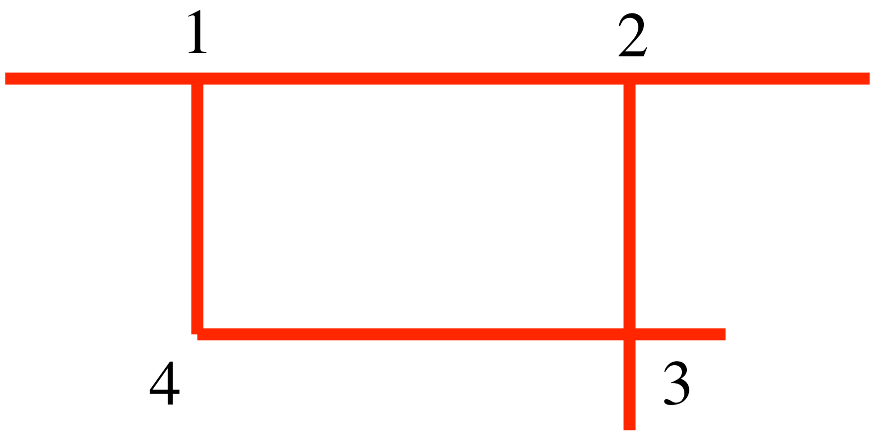

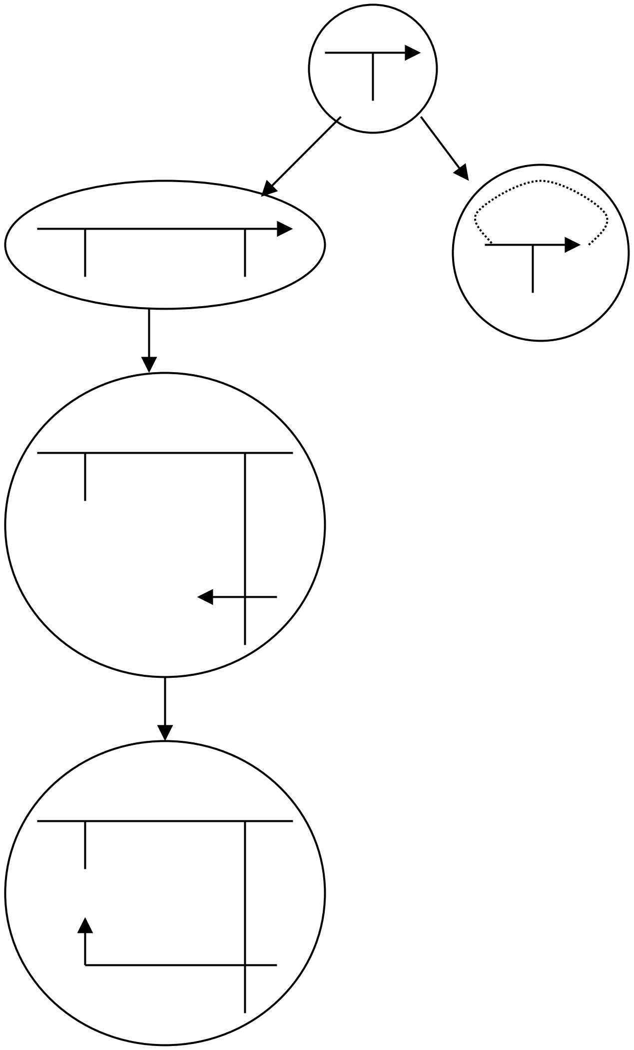

Tree of Maps (5)

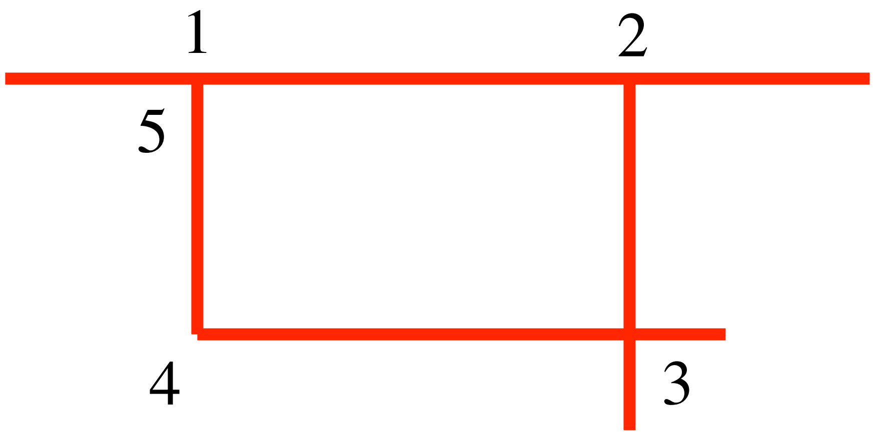

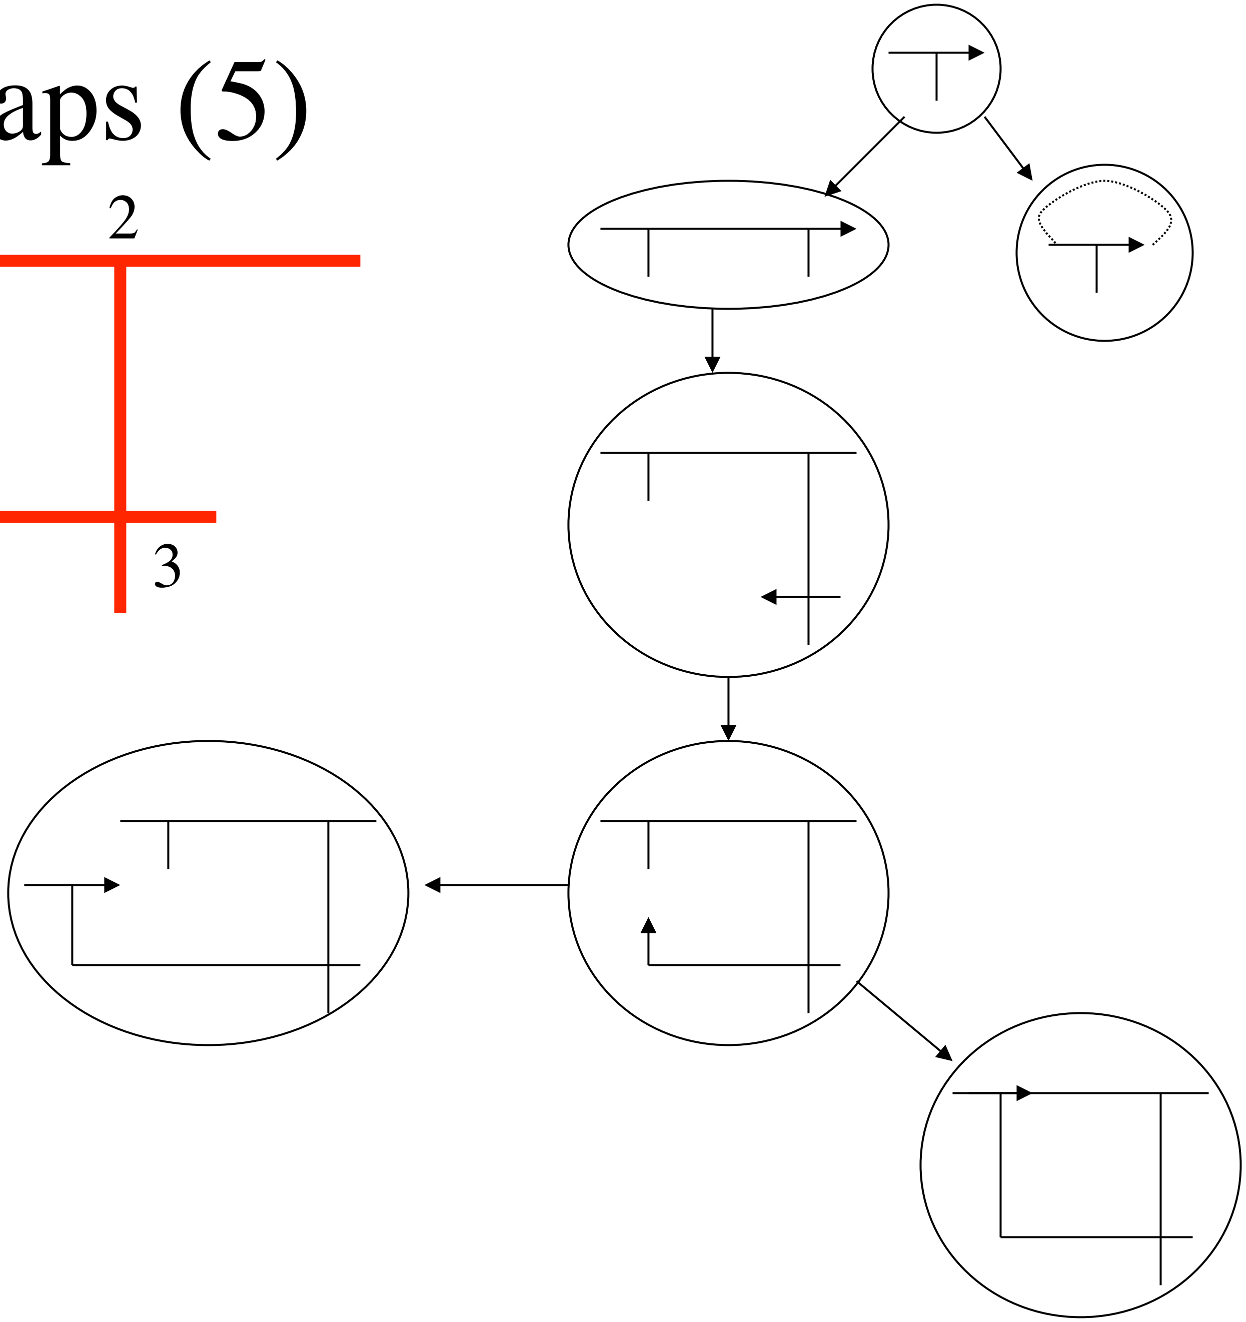

Tree of Maps (6)

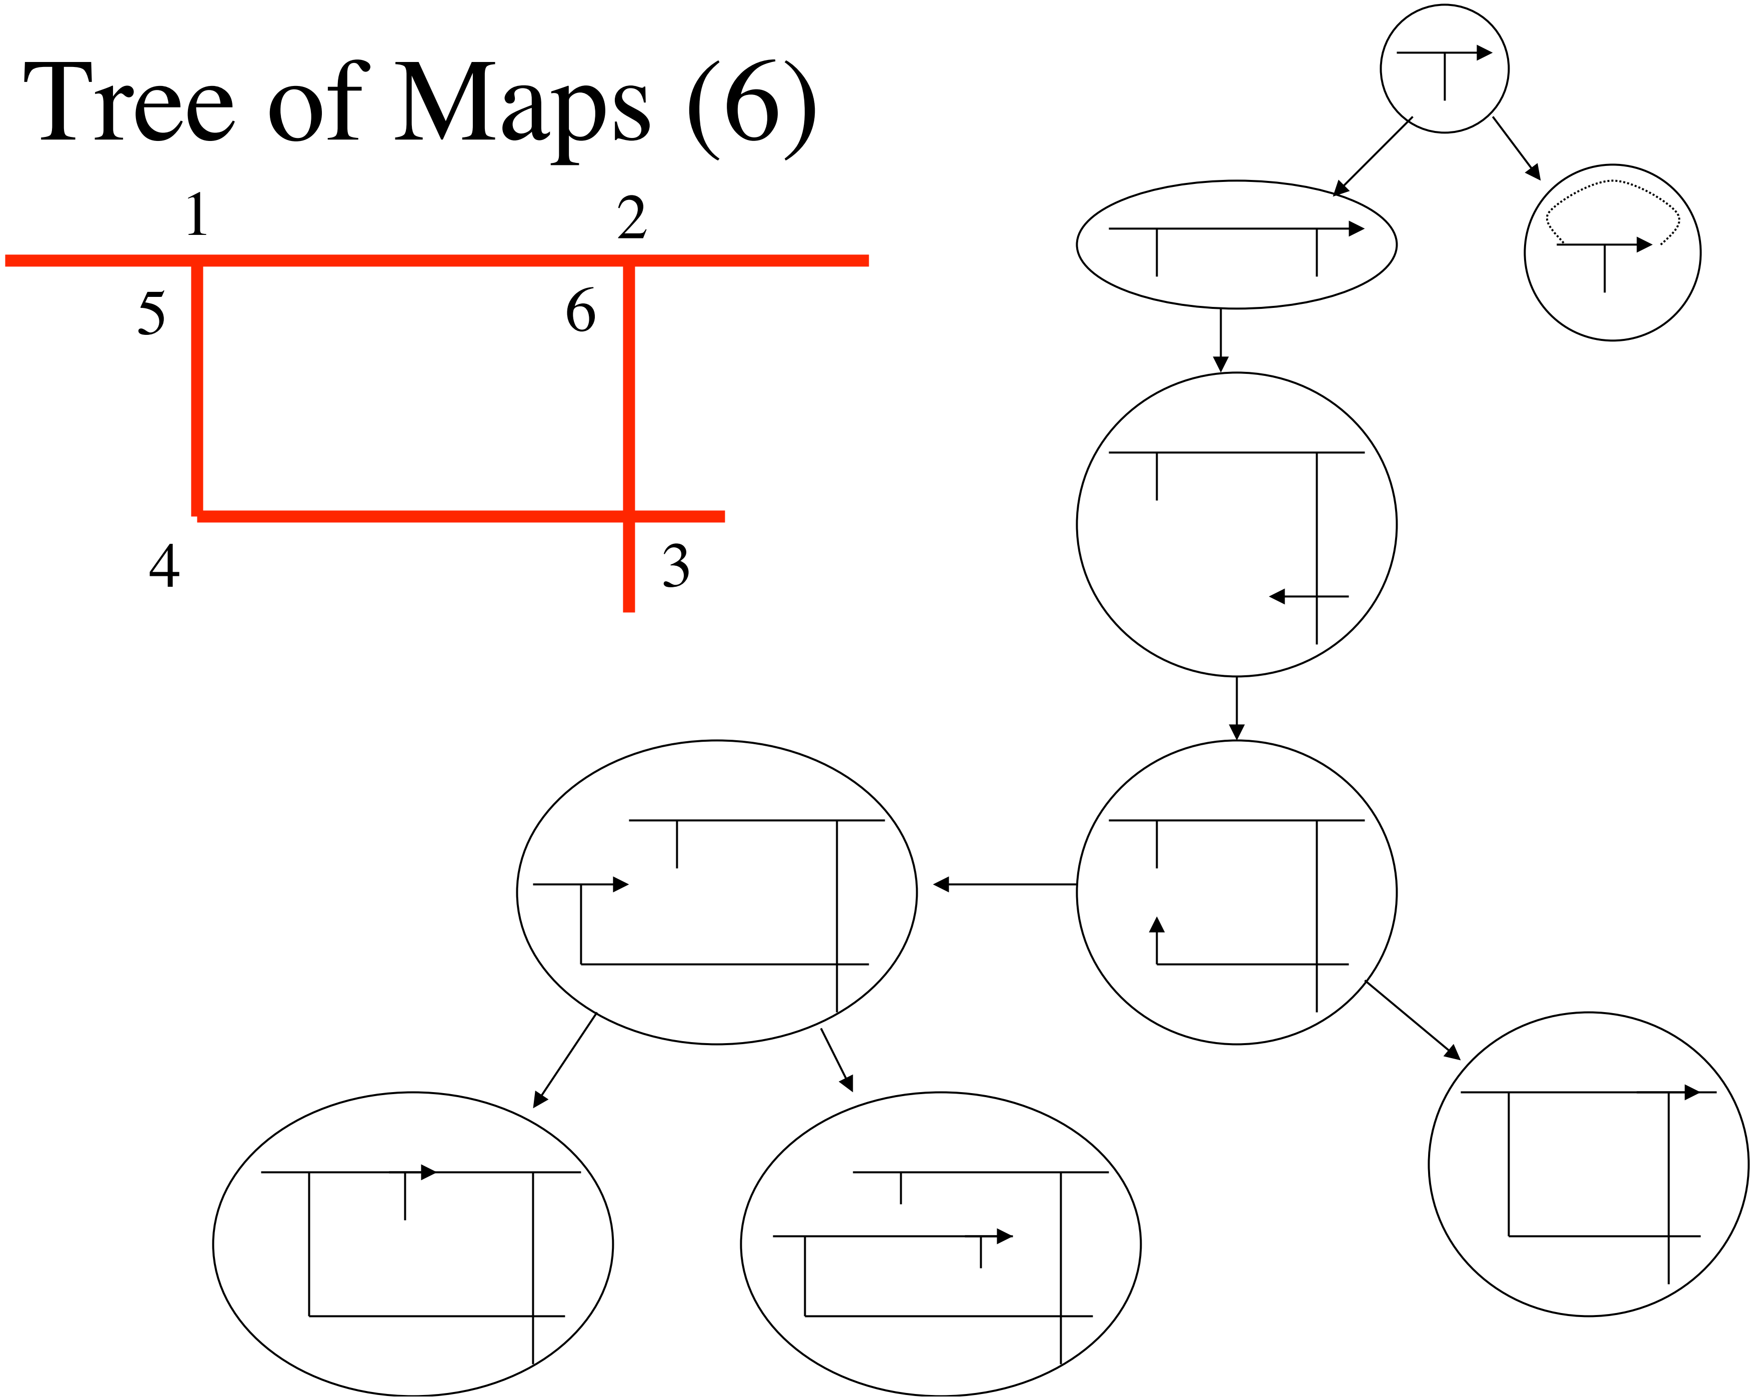

Rank the Consistent Maps

- The tree is **guaranteed** to contain the true map
  - All consistent maps are created.
  - Only inconsistent ones are deleted.
- Each map is a distinct loop-closing hypothesis.
  - Rank the consistent maps by simplicity (# places)
  - and/or likelihood, $p(odometry \mid layout)$.
- Use the current best map for planning.
  - Remember the tree.
  - The current best map could be refuted.

Plausible maps may be wrong

- Especially in Boston!

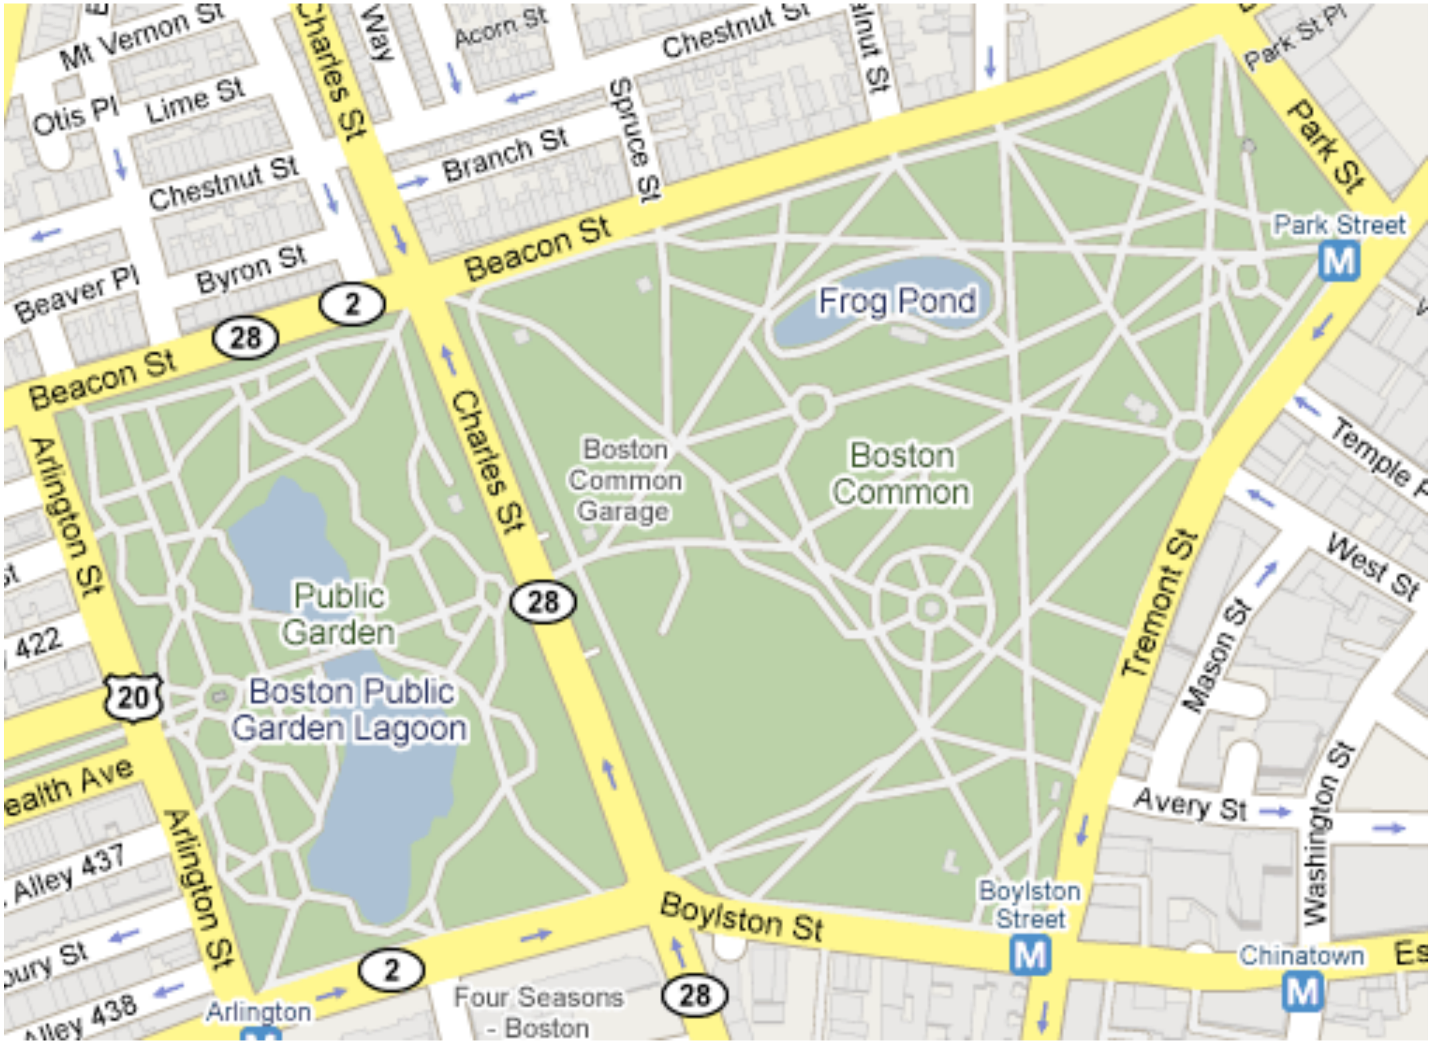

Global Metrical Map

- Use the topological map as a skeleton.
  - Lay out places in a single global frame of reference.
  - Fill in the details from local places and segments.

	Metrical Mapping	Topological Mapping
Small-scale space	Local SLAM	Local decision structure
Large-scale space	Global metrical map	Global topological map

Estimating Place Layout

- Local displacements propagate to global place layout.
  - Loop-closings are especially helpful.
- Relaxation converges quickly to a maximum likelihood layout.

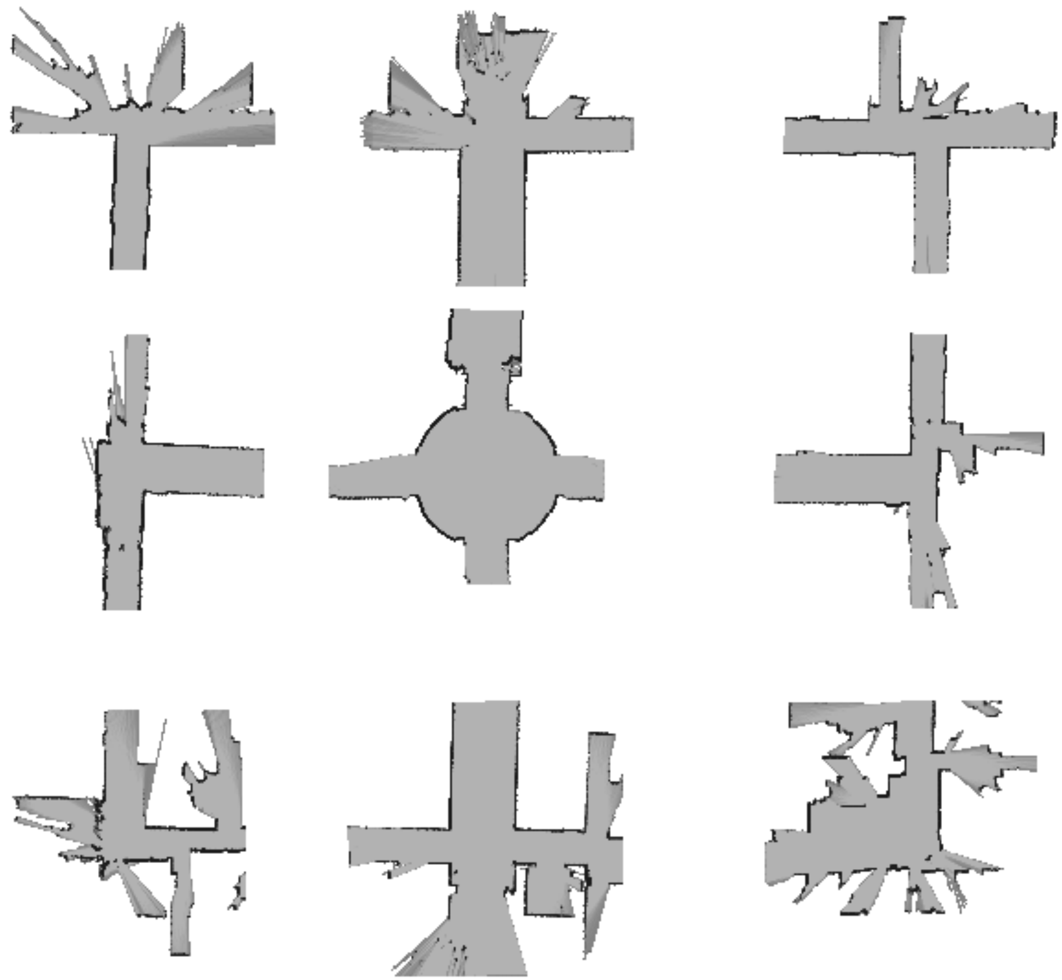

Estimating Robot Poses

- Given a max likelihood place layout
- and the trajectory of robot poses
- define a fixed anchor pose each time the trajectory passes through a place neighborhood
- correct the odometry in each segment.

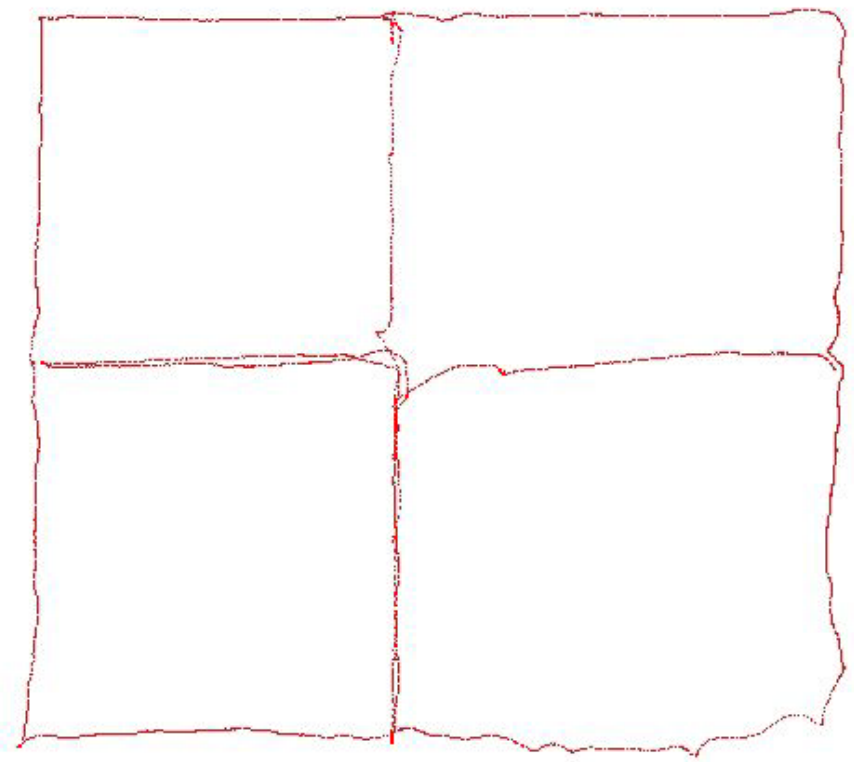

Global SLAM with new poses

- Use the corrected odometry to do SLAM in the global frame of reference.

- Or just treat the odometry as correct, and build the map.

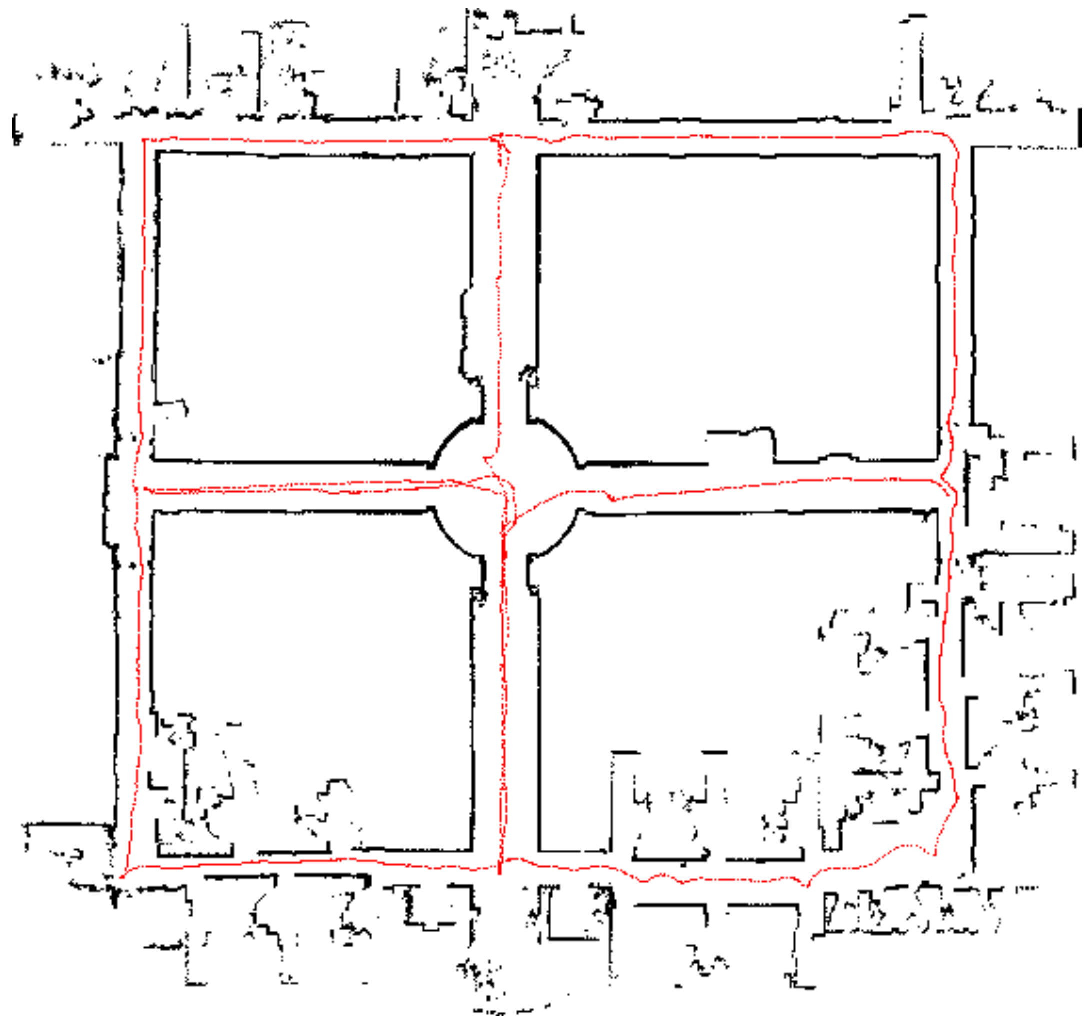

The Global Metrical Map

- The result is an accurate map in the global frame of reference.
- Cumulative error is eliminated by the topological map.
- More experience can be added locally to reduce any remaining errors.

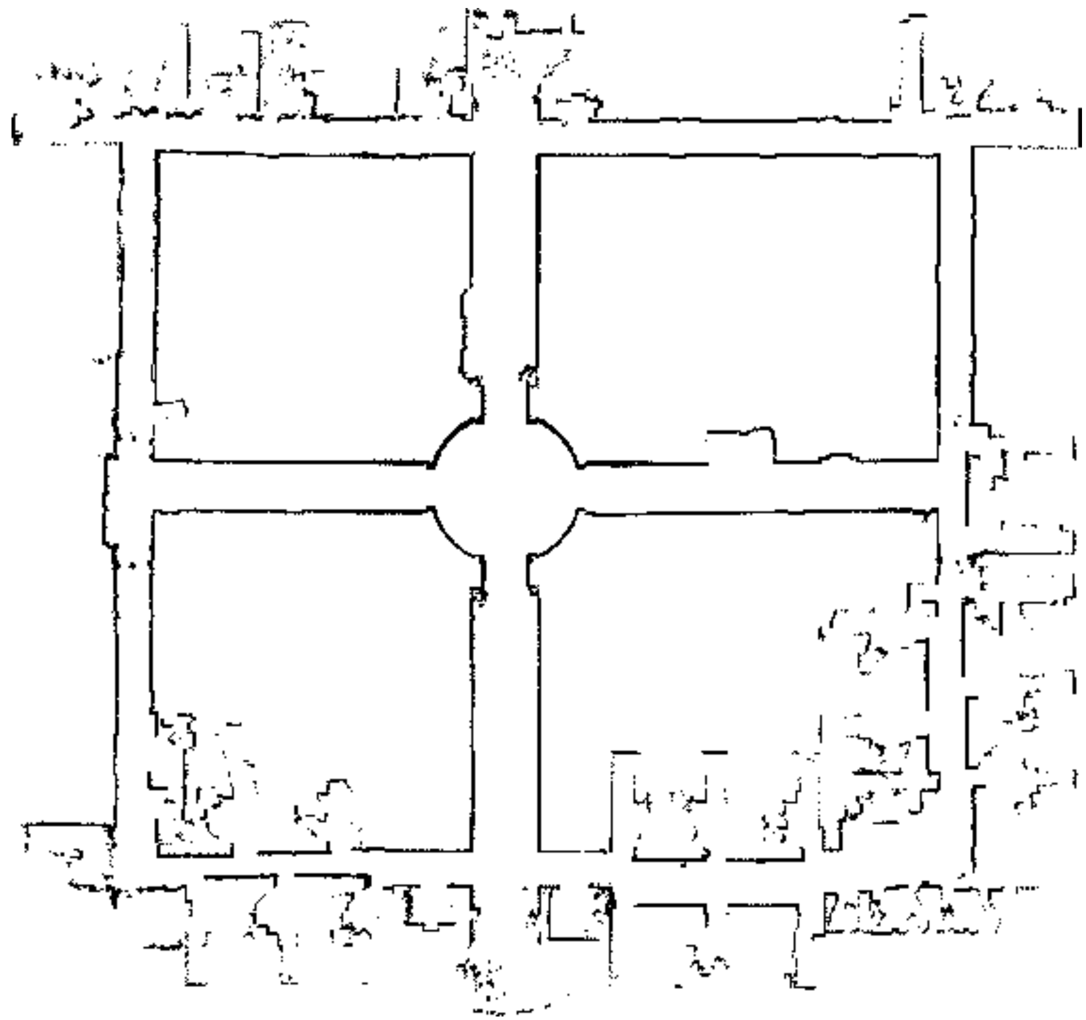

What have we got?

- Four representations for navigable space
  - Agent can learn them, or be told

	Metrical Mapping	Topological Mapping
Small-scale space	Local map for safe motion	Well separated decision points
Large-scale space	Heuristics to guide planning	Scalable map for route planning

Three-Tier Behavior Architecture

- *Deliberative planning*:
  - **Global topological map** defines the search space for route planning.
  - **Global metrical map** provides search heuristics.
- *Task sequencing*:
  - **Local decision structure** determines transitions between travel and turn actions
- *Continuous control*:
  - **Local perceptual map** provides world model for safe local motion planning.

Human-Robot Interaction

- Different kinds of human instructions map to different spatial knowledge representations

	Metrical Mapping	Topological Mapping
Small-scale space	*"Go there"* *"To my desk"*	*"Turn right"* *"Second left"*
Large-scale space	Select map point	*"To the kitchen"* *"Doctor's office"*

Controlling the Robot Wheelchair

Play from QuickTime Player

Lessons Learned

- Multiple representations are unavoidable
  - "Semantic Hierarchy" of representations
  - Large-scale and Small-scale space
  - Metrical and Topological representations
  - and others

- Multiple representations are useful
  - Reasoning can be more flexible and robust
  - Allow different kinds of sensors to contribute
  - Allow different kinds of human communication

- The human cognitive map provides guidance

References

- Park & Kuipers. **Feedback motion planning via non-holonomic RRT* for mobile robots**. IROS, 2015.
- Park, Johnson & Kuipers. **Robot navigation with model predictive equilibrium point control**. IROS, 2012.
- Beeson, Modayil & Kuipers, **Factoring the mapping problem: Mobile robot map-building in the Hybrid Spatial Semantic Hierarchy**. IJRR, 2010.
- Kuipers, **An intellectual history of the Spatial Semantic Hierarchy**. In Jefferies & Yeap (edited volume), Springer, 2008
- Remolina & Kuipers, **Towards a general theory of topological maps**. AIJ, 2004.
- Kuipers, **The Spatial Semantic Hierarchy**. AIJ, 2000.
- http://eecs.umich.edu/~kuipers/research/ssh/papers.html