# Research Statement
# PATRICK BEESON

## Overview

Intelligent machines will play an important role in solving difficult real-world problems in the years to come. Unlike other branches of artificial intelligence, mobile robotics allows investigation into how robust intelligence is tied directly to real-world experience. Having robots that are aware of their surroundings, and of how their actions affect those surroundings, is an important step towards providing machines that can assist or replace humans in complex or dangerous tasks.

I am particularly interested in how mobile robots can develop discrete concepts and categories from the continuous, high bandwidth sensations they receive when interacting with the external world. Moreover, for specific tasks, the robots should develop concepts that prove useful for communication with human users, without needing a human to provide vast amounts of domain-specific knowledge. When possible, I appeal to general theories of human cognitive abilities to motivate learning tasks and knowledge frameworks for mobile robots.

My current robotics research focuses on navigation, as I believe that spatial reasoning is a key foundational domain for developmental learning. Robot navigation research lends itself to investigating, among other things, interesting issues in perception, action, symbol grounding, and human-robot interaction (HRI). Specifically, my research investigates how low-level perceptions and actions can be abstracted by a robot into a hierarchy of representations for describing navigable space at different scales and with different ontologies. The spatial concepts abstracted by the robot are grounded in its experience with the physical environment, but they also need to have a natural correspondence to human methods for spatial navigation.

This research topic has proven fruitful in that I have developed techniques for improving the state of the art across a variety of mobile robot research areas. I have demonstrated algorithms that improve low-level, metrical robot localization [1]. I have generated algorithms for reliably detecting and describing topological places [2, 3]. I have formalized how to build large-scale metrically consistent maps from a known topological layout [4]. I have investigated machine learning techniques to discriminate between similar appearing locations [5]. Most recently, I have started evaluating the navigation performance of disabled subjects using an intelligent wheelchair platform [6].

## Hierarchical Representations of Space

The majority of my research focuses on mobile robot map-building and navigation. Most research in robot mapping focuses on methods for converting perceived distances of nearby obstacles into a unified metrical map, which describes the free and occupied regions of the environment in a single frame of reference. In contrast, my research focuses on how a mobile robot can build a hierarchy of related, yet ontologically distinct, representations of space. This hierarchy, called the Hybrid Spatial Semantic Hierarchy (HSSH), allows the robot to reason about space at different scales and abstractions. The hierarchy provides computational benefits for map-building and planning over commonly-used monolithic representations of space, robustness to incomplete knowledge, and, because it is inspired by human cognitive abilities, a natural human-robot interface.

One major contribution of my work has been to define the computational relation between *small-scale space* and *large-scale space*. Small-scale space is the local space surrounding the agent. Human representations of small-scale space are highly metrical: e.g., consider the detailed model you have of your office, even if the lights are off. Large-scale space lies beyond the sensory horizon of an agent. Human models of large-scale space are inherently less metrical and more symbolic: e.g., consider reasoning about the connectivity of highways in a city.

The hierarchy of spatial representations is built upon the robot's immediate environmental perceptions and direct access to motor outputs. From there, metrical representations of small-scale space are built. Currently, my research utilizes well-known Bayesian metrical mapping approaches to describe the transient small-scale space that scrolls past an agent as it moves through the world.

Assuming well-formed metrical models of small-scale space (largely dependent on precise localization with respect to local obstacles [1]), the next level of the hierarchy describes the symbolic description of small-scale space. Specifically, my research reliably abstracts the *local topology* of small-scale space from the local metrical model. At this level, the world is described as a set of local path fragments that enumerate the ways out of the current location—forward, turn around, turn left, turn right. One interesting point is that this abstraction gives us an elegant way to detect places: if the symbolic description of small-scale space has one path through the space, the robot is not at a place, otherwise, it must be at a place [2].

Using the symbolic description that describes places, the robot can hypothesize the global connectivity of places and paths [3]. Finally, using a topological structure of the layout of places, along with error-prone metrical information about the relative alignment between neighboring places, my colleagues and I have detailed a method for building a single, global metrical map of a large-scale space. This method is similar to greedy methods used by researchers attempting to extend the Bayesian metrical mapping techniques to large environments that contain loops; however, by using a small set of places instead of a large number of saved robot poses, our method is computationally superior to those other implementations [4].

**Bootstrap Learning**

I am also interested in a robotics problem that my colleagues and I call bootstrap learning, which is closely related to the emerging field of developmental robotics. The basic idea is to compose multiple machine learning methods, using weak but general unsupervised or delayed-reinforcement learning methods to create the prerequisites for applying stronger but more specific learning methods such as abductive inference or supervised learning. The result is a bootstrapping of novel experience into useful categories, which are then generalized and/or subdivided as needed into other useful categories.

Specifically, I have focused on bootstrap learning for place recognition, which is a kind of *self-supervised learning* of place descriptions [5]. This work is motivated by the idea that humans initially have general rules for remembering locations (e.g., the third door on the left) that become more specific and perceptually local over time (e.g., the door with the gold doorknob), thus simplifying recognition.

Here, a robot uses unsupervised clustering techniques to eliminate *perceptual variability* (the same location looks different on separate occasions) by potentially increasing *perceptual aliasing* (multiple locations look similar). This clustering, along with physical exploration, facilitates topological map-building, which eventually assigns a unique identifier to each location. This unique identifier allows the robot to perform supervised learning of locations from raw sensory images. This reclaims important discriminatory features that may have been washed out as noise earlier. Eventually the robot learns to immediately recognize a location from a raw sensory image, ignoring the noise and looking for the discriminating features.

**Intelligent Wheelchair**

Most recently, I have become interested in evaluating the claim that the Hybrid Spatial Semantic Hierarchy provides a straightforward human-robot interface for a mixed autonomy robotic platform. Specifically, I am involved in evaluating the performance of an intelligent wheelchair that uses the HSSH to represent space and infer commands from the human user.

The user has the ability to control the wheelchair at a variety of levels. At the motor level, the user drives the wheelchair like a normal chair (using a joystick), and no autonomy is given to the wheelchair. At the control level, the wheelchair uses the local metrical description of space to revise the user's joystick commands if they are unsafe. At the command level, the wheelchair treats the world like a graph of decision points, moving between places, conveying the local topological structure, and accepting the path to exit on as its command. At the goal level, the user gives a high-level goal (like "Go to my office.") and the robot acts autonomously to navigate to the goal.

Preliminary experiments show that low-vision subjects perform better at navigating the wheelchair when delegating certain decisions to the intelligent agent [6].

## Future Research

I am enthusiastic about continuing research in the realm of mobile robot navigation and developmental learning. Low-level knowledge about perception, action, and space is a crucial component of any agent situated in the real world. Additionally, the real world contains people, so a robot's knowledge should be guided in order to facilitate easy communication with humans.

The bootstrap learning research is an important component in this vision as it provides a robustness that is often too complex to simply program in. Whenever sensors fail, or a new sensor is added, we would like the robot to adapt. Similarly, I believe that bootstrap learning can improve transfer learning, where the domain changes slightly, by allowing new concepts and behaviors to be built-up from reliable low-level commonsense knowledge, grounded in the robot's own perceptual abilities.

There is still interesting research on the topological map-building problem. Though I have provided a few algorithms for abstracting the continuous to the metrical to the symbolic, the problem of detecting and describing useful places, is still a rich research area (e.g., consider using a water tower next to a highway as a landmark). Additionally, topological map-building is often viewed as a logic-based research area, as places and paths are symbolic entities. However, I believe that incorporating noisy metrical knowledge along with uncertain symbolic descriptions, will allow us to pursue topological inference using a Bayesian framework, further enhancing the robustness of topological navigation.

Finally, the intelligent wheelchair platform allows for some very interesting work. From a robotics point of view, vision processing should allow us to model small-scale space of outdoor walkways. Identifying the bounds of sidewalks and crosswalks on campus, allows us to use the same metrical/topological navigation algorithms that have been demonstrated using range sensors in indoor environments. From a human-robot interaction perspective, the intelligent wheelchair platform allows us to continue research into how humans conceptualize space as well as how the robot and disabled users can communicate about space. There are many interesting issues on mixed autonomy that need to be examined, along with investigations into the benefits of such an intelligent assistant on a variety of disabled populations.

I believe we will soon begin to see more autonomy in our everyday transportation vehicles and more autonomous vehicles combined with human exploration (both outer space and deep ocean exploration). In these domains, we will need robots that can reason at a variety of abstractions in order to efficiently model the world, to handle incomplete knowledge, to communicate complex information in a variety of ways, and to understand human commands in the current context. I hope to contribute to this effort in the years to come.

## References


[1] Patrick Beeson, Aniket Murarka, and Benjamin Kuipers. Adapting proposal distributions for accurate, efficient mobile robot localization. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2006.

[2] Patrick Beeson, Nicholas K. Jong, and Benjamin Kuipers. Towards autonomous topological place detection using the Extended Voronoi Graph. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2005.

[3] Benjamin Kuipers, Joseph Modayil, Patrick Beeson, Matt MacMahon, and Francesco Savelli. Local metrical and global topological maps in the Hybrid Spatial Semantic Hierarchy. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2004.

[4] Joseph Modayil, Patrick Beeson, and Benjamin Kuipers. Using the topological skeleton for scalable global metrical map-building. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2004.

[5] Benjamin Kuipers and Patrick Beeson. Bootstrap learning for place recognition. In *National Conference on Artificial Intelligence (AAAI)*, 2002.

[6] Patrick Beeson, Matt MacMahon, Joseph Modayil, Aniket Murarka, Benjamin Kuipers, and Brian Stankiewicz. Integrating multiple representations of spatial knowledge for mapping, navigation, and communication. In *AAAI Spring Symposium Series, Interaction Challenges for Intelligent Assistants*, Stanford, CA, 2007.