Integrating multiple representations of spatial knowledge for mapping, navigation, and communication

Patrick Beeson
Matt MacMahon
Joseph Modayil
Aniket Murarka
Benjamin Kuipers
Department of Computer Sciences

Brian Stankiewicz
Department of Psychology

The University of Texas at Austin

Goal

- Intelligent Wheelchair
  - Provides:
    - Safe execution of commands
    - Perception
    - Communication
  - Benefits:
    - Mobility impaired
    - Visually impaired
    - Cognitively impaired

Wheelchair Research Issues

- Wheelchair Hardware
  - Sensors, power consumption, etc.
- Interface Hardware
  - Varies by disability, personal preference, etc.
- Low-level Control
  - Velocities to motor voltages, safe/comfortable acceleration
- Knowledge Representation
  - Perception, navigation, spatial concepts, mixed autonomy
- User community studies
  - Usefulness, trust, cost

Interface Goals

- “Dock at my desk.”
- “Enter restroom stall.”

- “Go to the end of the hallway.”
- “Take the next left.”
- “Go right at the ‘T’ intersection.”
- “Go to the Psychology building.”
- “Stop at the water fountain.”
- “Take the scenic route.”

Representation Independence

- We want the spatial reasoning system to be independent of:
  - Specific interface with user
  - Specific robot platform/sensors

Talk Overview

1. Knowledge Representation
2. Pilot Experiments

Current focus

- Knowledge representation should facilitate:
  - Modeling of environment
  - Safe navigation
  - Communication
  - Mixed autonomy

- High-precision control (small, precise spaces)
  - Bathroom stalls, office navigation/desk docking, etc.
- Low-precision control (large-scale spaces)
  - Obstacle avoidance in hallways, turning corners, etc.

Progress

- This talk:
  - Spatial reasoning framework
    - The Hybrid Spatial Semantic Hierarchy (HSSH)
  - Experimental results
    - Wheelchair navigation with simulated low-vision users

- Related work from our lab:
  - Natural language route instructions
  - 3D safety
  - Object / Place learning

State of the art in mobile robotics

- Mobile robot research is largely focused on SLAM (simultaneous localization and mapping).
- Most SLAM implementations create a monolithic representation of space
  - Metrical map
  - Single frame of reference
  - e.g. occupancy grids, landmark maps

Issues:

- Closing large loops
  - Heuristic
  - Long compute times
- Interaction
  - Exploring a new environment
  - Blind users
  - Planning

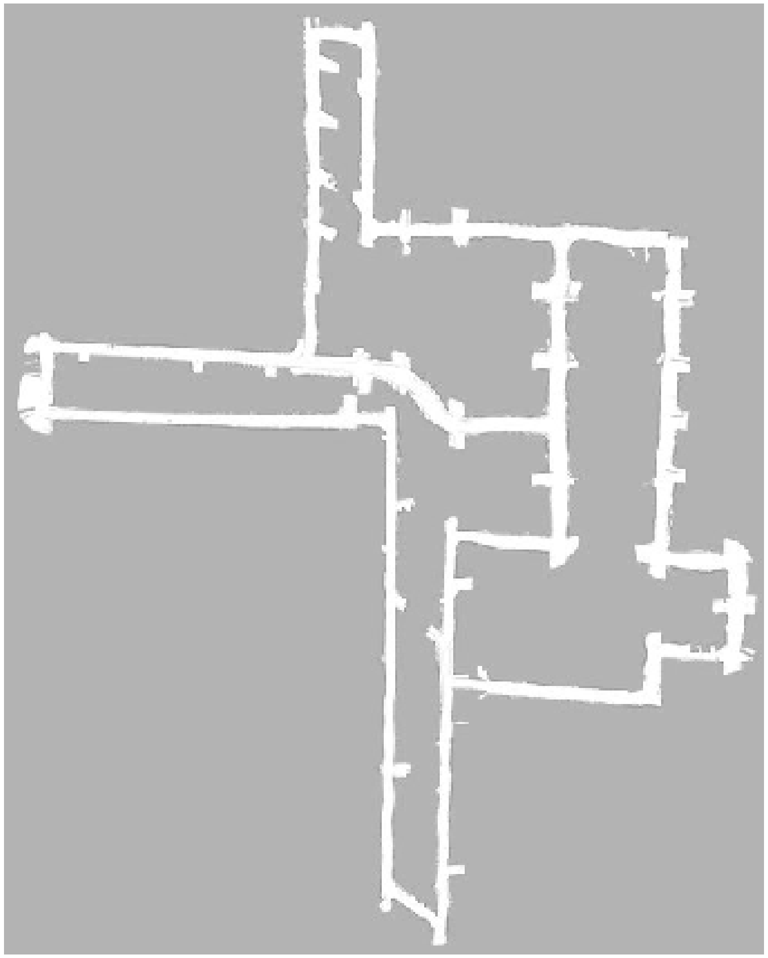

Hybrid Spatial Semantic Hierarchy

- Factor spatial reasoning about the environment into reasoning at four levels
  - local metrical – models obstacle locations in local surround
  - local topology – models symbolic structure of local surround

    } Small-scale models

  - global topology – models global symbolic structure of entire environment
  - global metrical – models global layout of obstacle locations

    } Large-scale models

    - Largely unnecessary, but often useful if it exists
- Each level has its own ontology / language
  - Inspired by human cognitive behaviors
- More robust and efficient than a single, monolithic representation, but also more useful to provide human-robot interaction.
  - Better than a single, large occupancy grid representation

HSSH Diagram

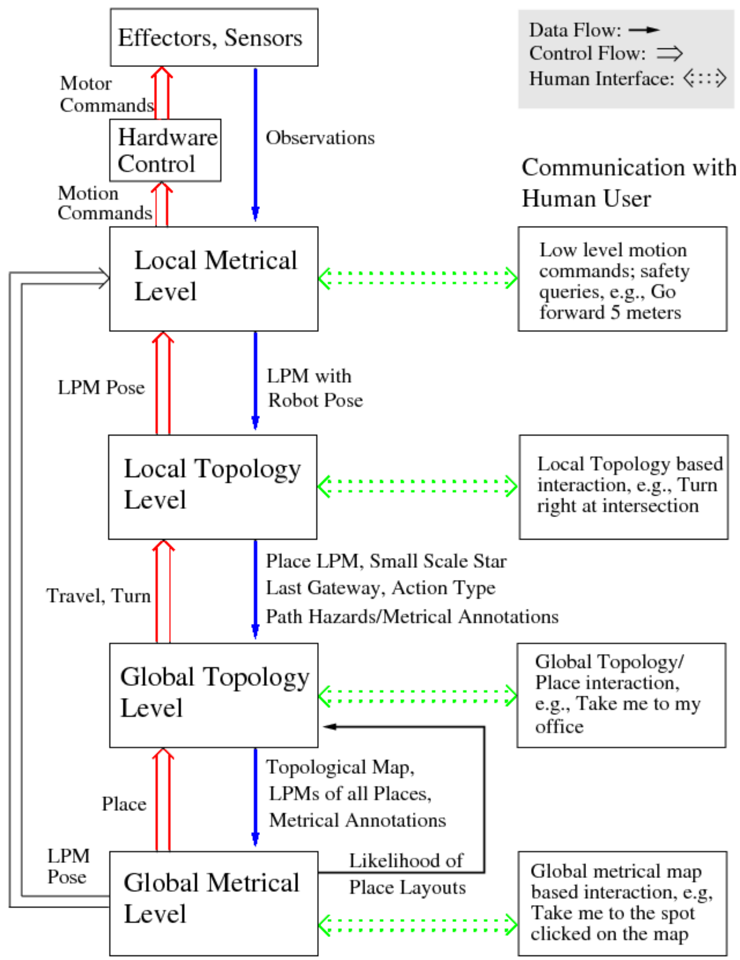

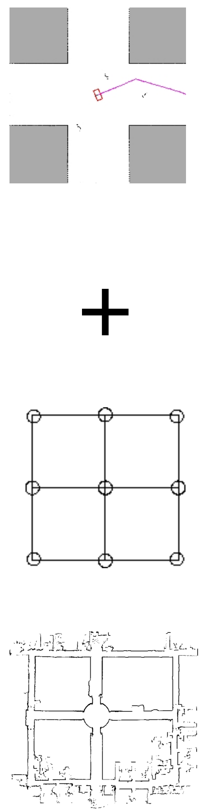

Local Metrical Level

- Environment is modeled as a bounded metrical map of small-scale space within the agent's perceptual surround.
  - Scrolls with the agent's motion
  - Not tied to a global frame of reference.
  - Useful for "situational awareness" of the immediate surround.

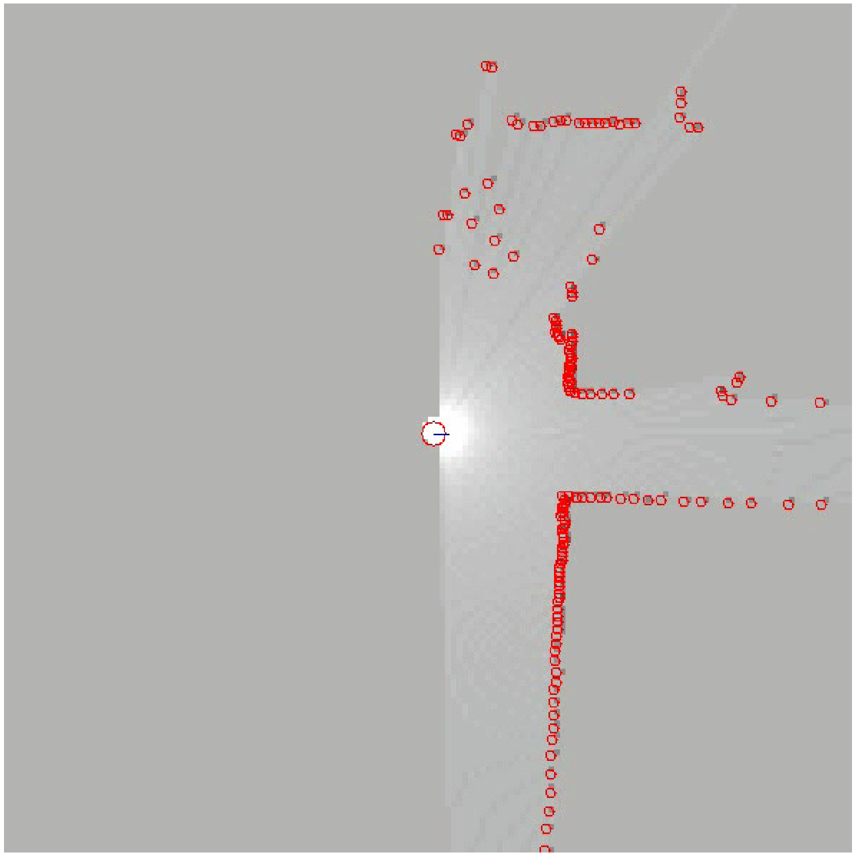

Local Metrical Control

- Driver uses the joystick.
  - Robot checks commands against the local map for safety.
- Driver may specify a target or direction of motion within the local map.
  - Robot plans hazard-avoiding motion toward that target.

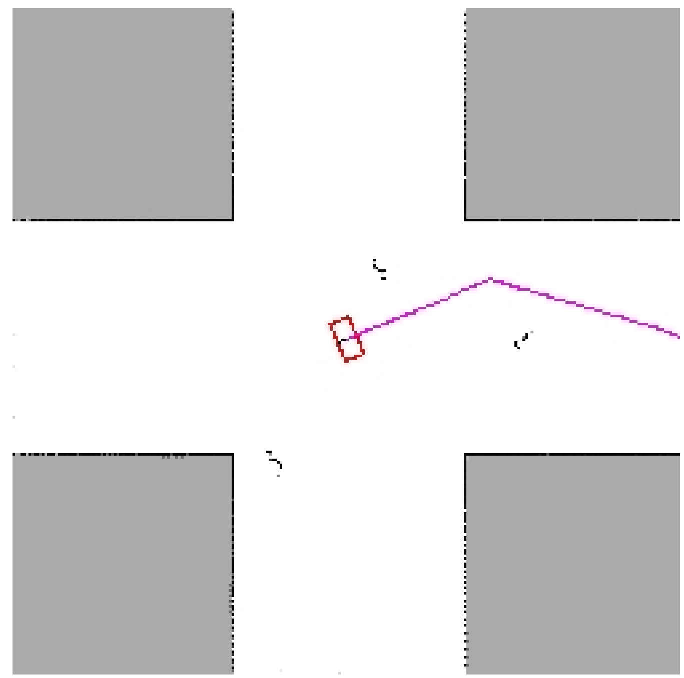

Local geometry → local topology

- Compute “gateways”

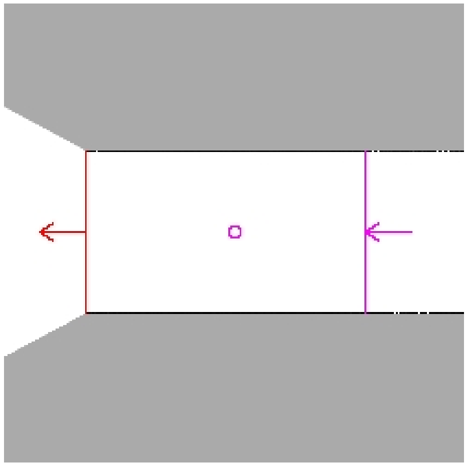

- Gateways help define “places”

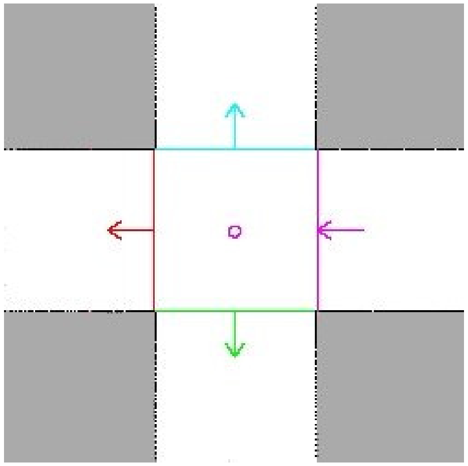

Local Topology Level

- Environment is modeled as a set of discrete decision points, linked by actions
  - Turn selects among options at a decision point
  - Travel moves to the next decision point.

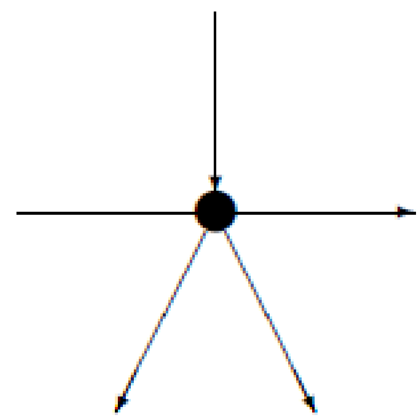

Local Topology Control

- Driver specifies turn actions at decision points.
  - Turning actually corresponds to selecting a gateway location and performing control at the local metrical level.
  - Travel moves from a gateway to the next place.

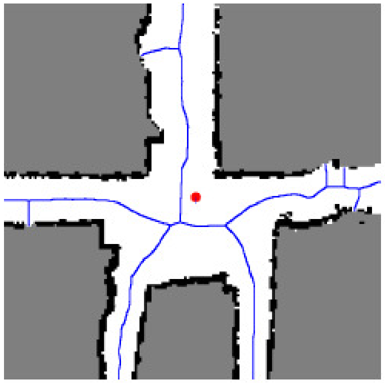

(a)

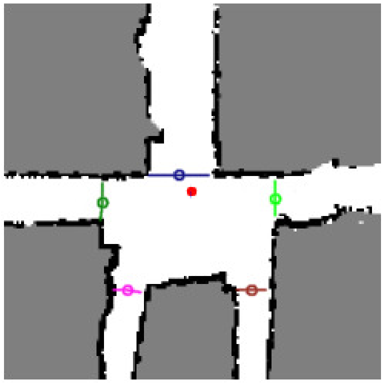

(b)

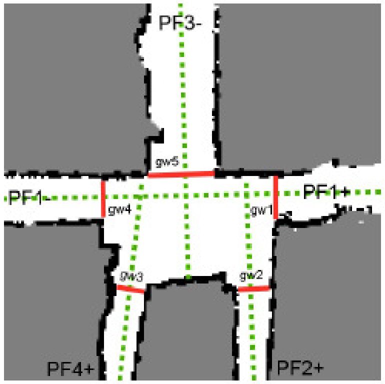

(c)

Small-scale star description from PF1+

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

(d)

Local topology → global topology

- Detect loop closures based on matching local topology and local metrical models.

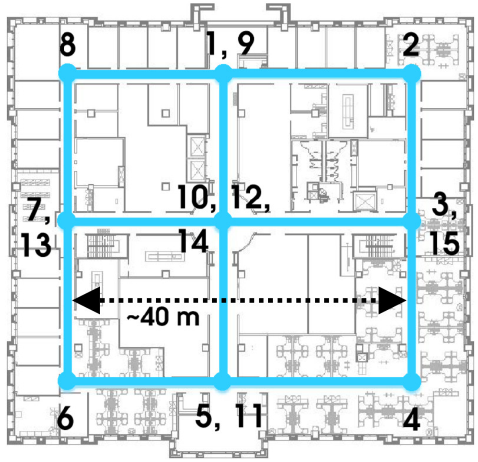

- Build tree of possible topological maps and use simplest model as current best guess.

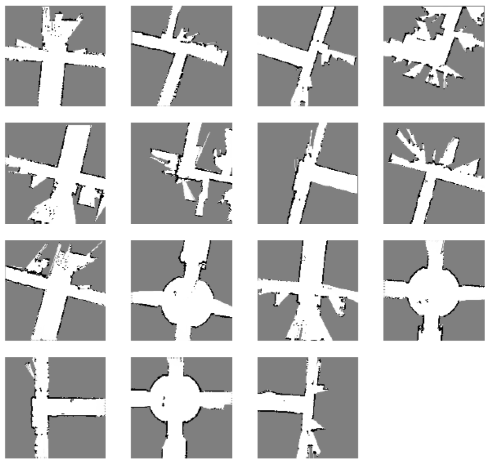

Global Topology Level

- Environment is modeled as a network of places, on extended paths, contained in regions
  - Efficient route planning in large environments
    - graph search

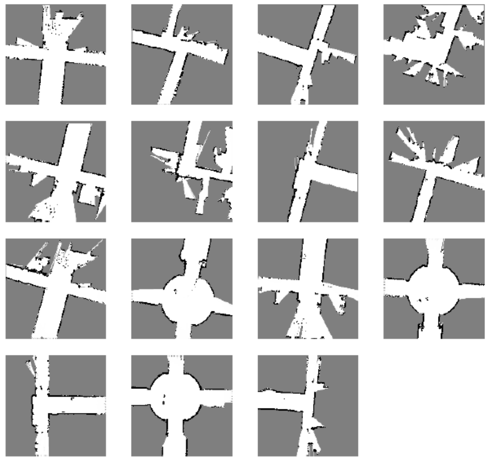

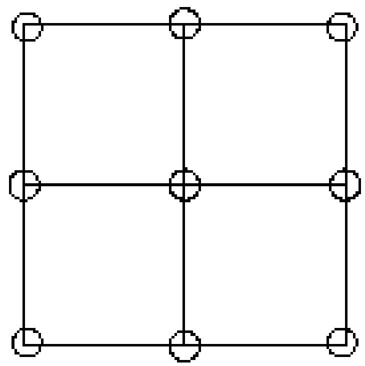

Global Topology Control

- Driver specifies a destination place in a topological map, by name or in a schematic diagram (like a subway map).

1. Robot plans a route to that goal
2. Route is translated into a sequence of local topology travel/turn commands
3. Route is executed by hazard-avoiding control laws in the local metrical model

Global topology → global metrical

- Use local metrical information between topological places to find global metrical layout of places.

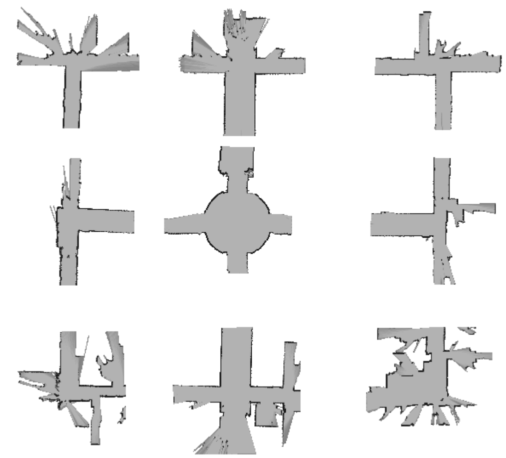

- Build global metrical map on top of the topological skeleton.
  - More computationally efficient than other methods

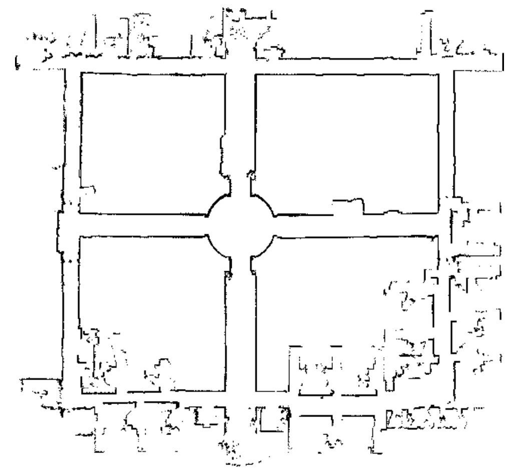

Global Metrical Level

- Environment has a geometric model in a single global frame of reference.
  - Useful for route optimization when available, but not necessary for large-scale navigation.

Control

- Driver clicks on a global metrical map
  - Robot plans a route to that destination in the topological map, then completes its route in the local metrical model.

- Driver specifies a saved destination that may not correspond to a "place", but has a location in the global map (e.g., "Go to the charger.").

Talk Overview

1. Knowledge Representation
2. Pilot Experiments

Background

- Wheelchair software is written for and tested on actual robotics platforms.
- To safely simulate disabled users, we port the code to a virtual environment.
  - Also useful for safely evaluating new ideas.

VR Setup

- Wheelchair software runs on "virtual wheelchair" in a virtual 3D maze environment.
  - Human avatars act as obstacles.
  - Virtual "laser scanner" at shin height
  - Users eye level at about chest height
- We test two perceptual conditions
  - Normal vision
  - Degraded vision

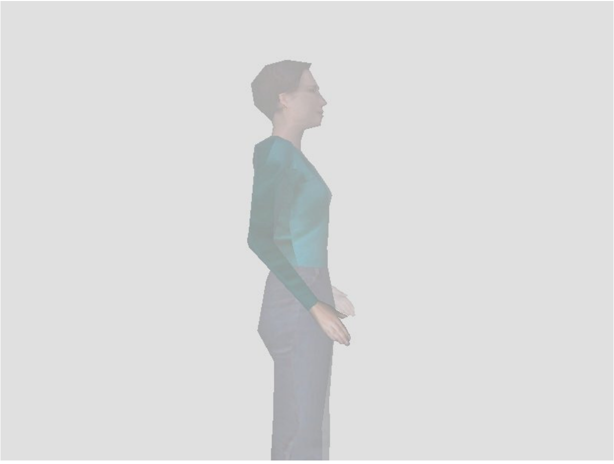

Pilot Study Interfaces

- 3 Navigation interfaces:
  - Manual (Joystick)
    - No intelligence
    - Joystick directly commands motion
  - Control (Joystick)
    - Uses local metrical model
    - Throttles velocities in hazard situations
    - Disregards unsafe actions
  - Command (GUI Interface)
    - Commands local topology level
    - “Go to next decision point”, “turn left”, etc.
  - Not tested:
    - Topological / Global metrical navigation

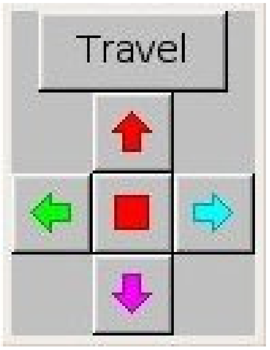

Experimental Questions

- Effect of Degraded Vision
  - Does reducing the visual information by adding fog make the task more difficult?
- Benefit of Assisted Joystick Control
  - Is performance better with local metrical control (collision avoidance)?
- Benefit of Local Topology Navigation
  - Does the navigation improve by using local topology knowledge in the wheelchair?
    - User gives discrete commands
    - Wheelchair performs navigation between decision points

Experiment Details

- 4 conditions
  - *Normal vision: Manual interface* (no safety)
  - *Degraded vision: Manual interface* (no safety)
  - *Degraded vision: Control interface* (safety)
  - *Degraded vision: Command interface* (decision graph w/ safety)

- 3 subjects
  - Each subject made 5 runs in each condition
    - 20 total runs
  - 20 runs were randomized for each subject

Experimental Details cont'd

- A run consisted of moving between 5 randomly chosen locations in the environment.
  - Natural language feedback
- Subjects knew environment beforehand
  - Avatars were randomly distributed for each run.

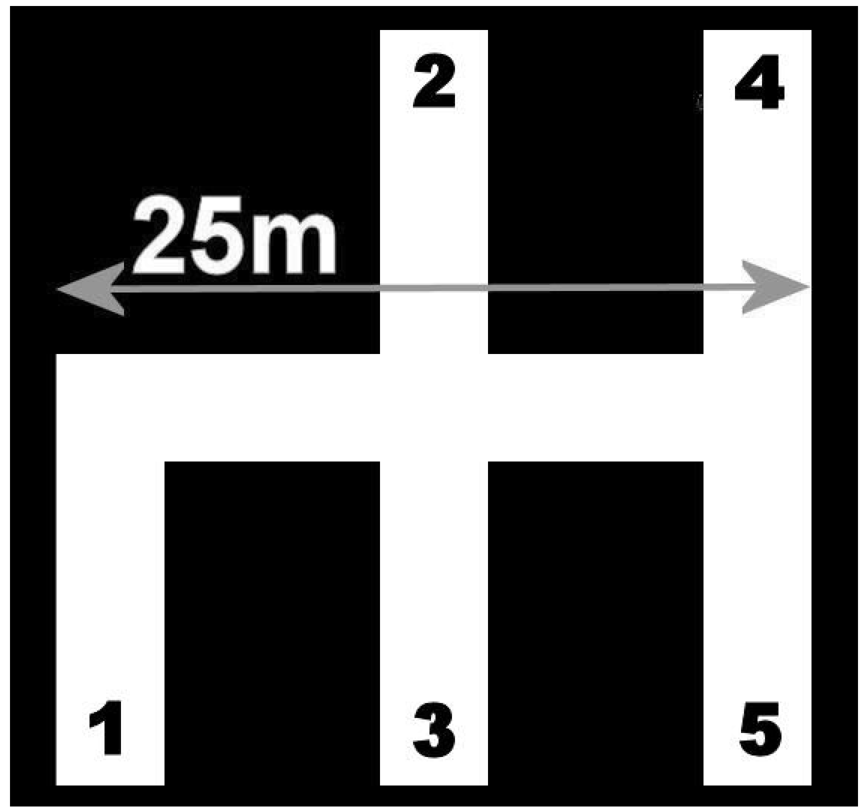

Qualitative Results

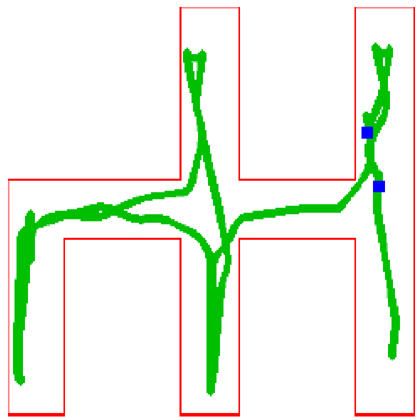

Normal:Manual

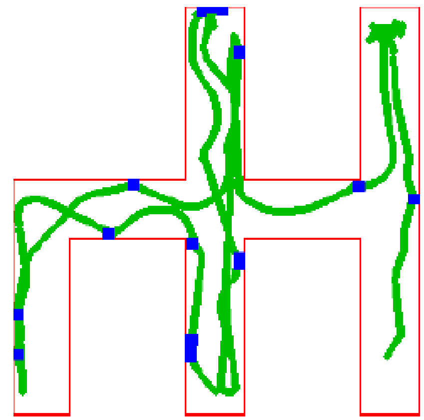

Degraded:Manual

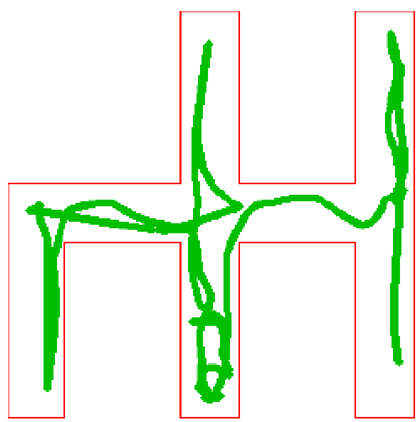

Degraded:Control

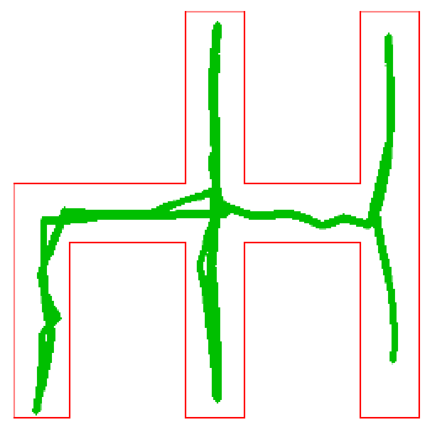

Degraded:Command

Quantitative Results

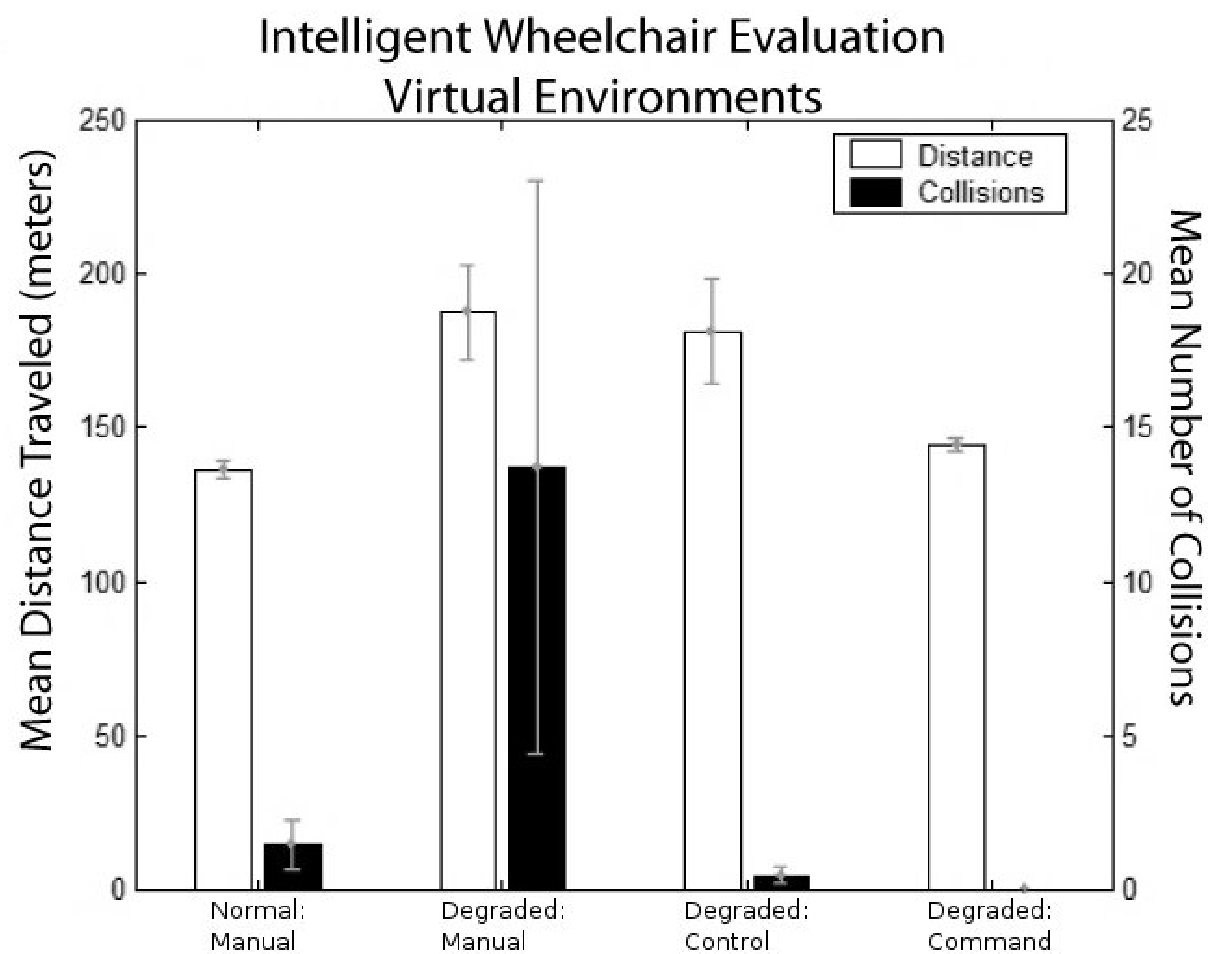

Future Work (Robot)

- Evaluate global topological navigation
  - User decides final location
  - Fully autonomous navigation by robot
  - Larger environments
- Evaluate interface devices with intelligent wheelchair platform
  - Force-feedback joystick
  - Touch screen
  - Natural language
- High-precision control
  - Create 2½ D local metrical models from vision.

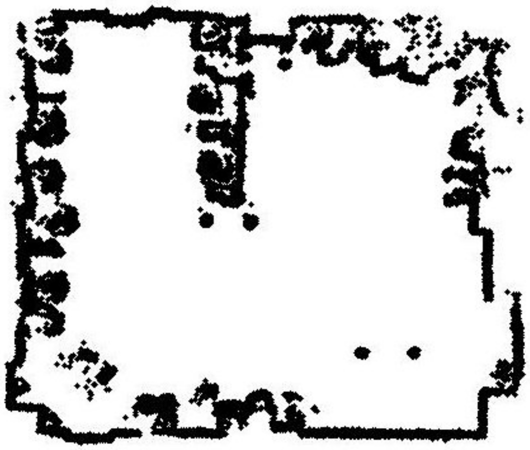

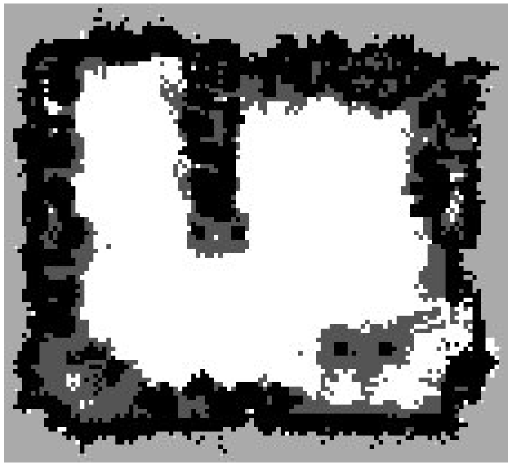

Future Work (VR)

- Continue low-vision experiments
  - Better simulation of low-vision
  - Using real wheelchair and head-mounted VR display
- Other measurements
  - Cognitive load
  - Stress
- Evaluate wheelchair for users with other disabilities
  - Fully blind
  - Quadriplegic
  - Memory loss / Alzheimer's

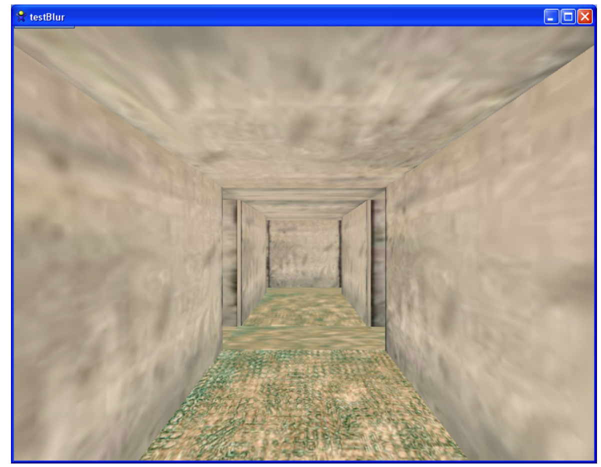

The End

http://www.cs.utexas.edu/~robot