Keystone Rescue

Jacky Baltes
University of Manitoba
Winnipeg, Canada

Email: jacky@cs.umanitoba.ca
WWW: http://www.cs.umanitoba.ca/~jacky

Outline

- Introduction to the Keystone Rescue
  - Hardware
  - Motivation
- Egomotion estimation
- Wall/line detection
- Map building
- Conclusion
- Comments about the competition

Introduction

- Entered rescue competition in 2002, '03.
- Based on cheap robotic platforms, originally used for soccer
  - Hummer: 25 MHz 68332MHz, 1MB RAM, CMOS camera, toy car chassis. 1 FPS
  - Rescue2: 300MHz Xscale, 64MB RAM, CMOS Webcam, differential drive, home-built based on CMU Corky. 5 FPS
  - PT Cruiser: 300MHz Pentium laptop with Webcam. Teleoperated.

Robots

Hummer

Rescue2

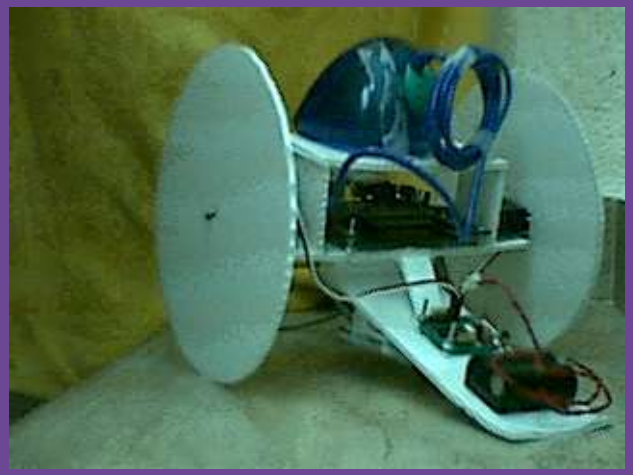

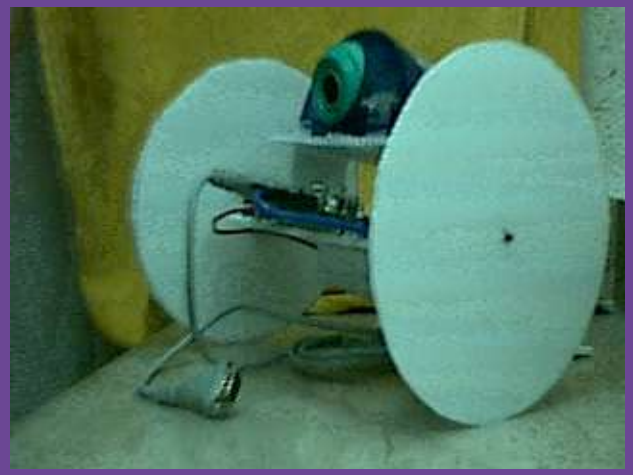

Introduction

- Not a swarm architecture
  - No confidence measure for when an area has been searched completely
  - Expect to perform sophisticated exploration and mapping
- Emphasis on computer vision research
  - Camera is the only sensor
  - No odometry

Egomotion Estimation

- Based on optical flow analysis
- Motion of objects in the scene and estimate of current motion are combined to form an estimate of the motion of the robot
- Replaces odometry

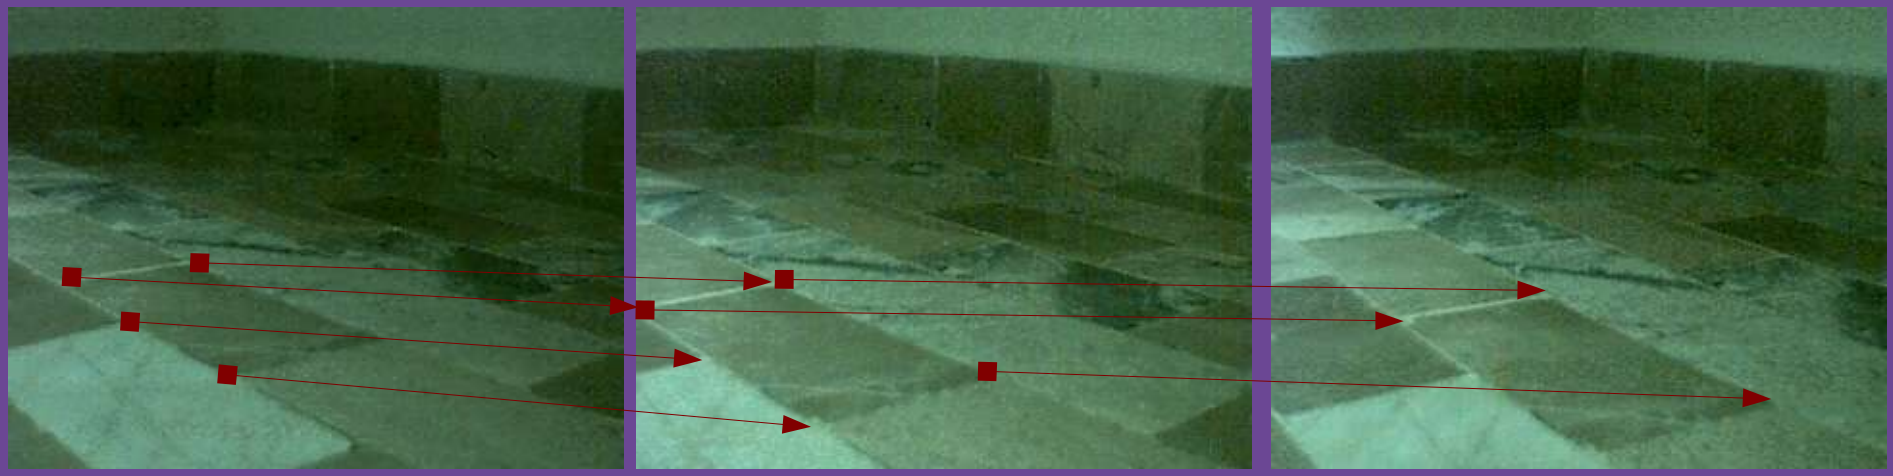

Egomotion Estimation Using Line Segments

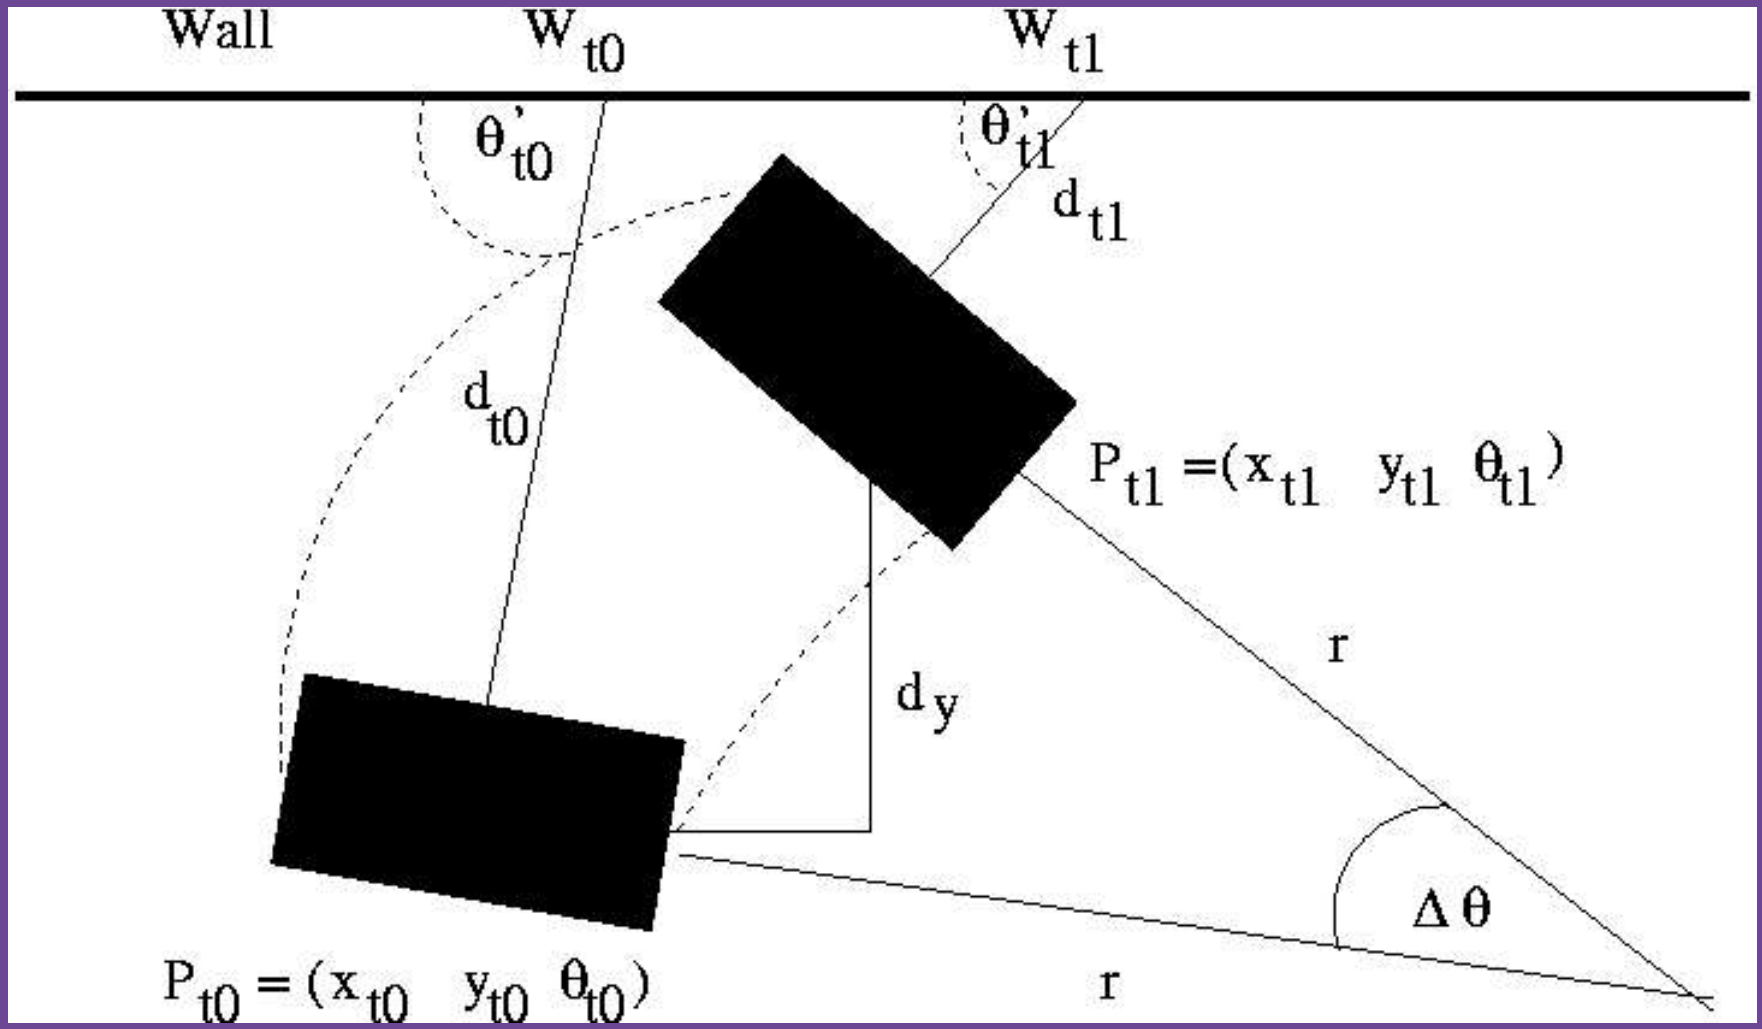

Egomotion Estimation From Lines

- Angle alone okay to determine orientation
- Angle and distance from the wall sufficient to compute differential drive kinematics

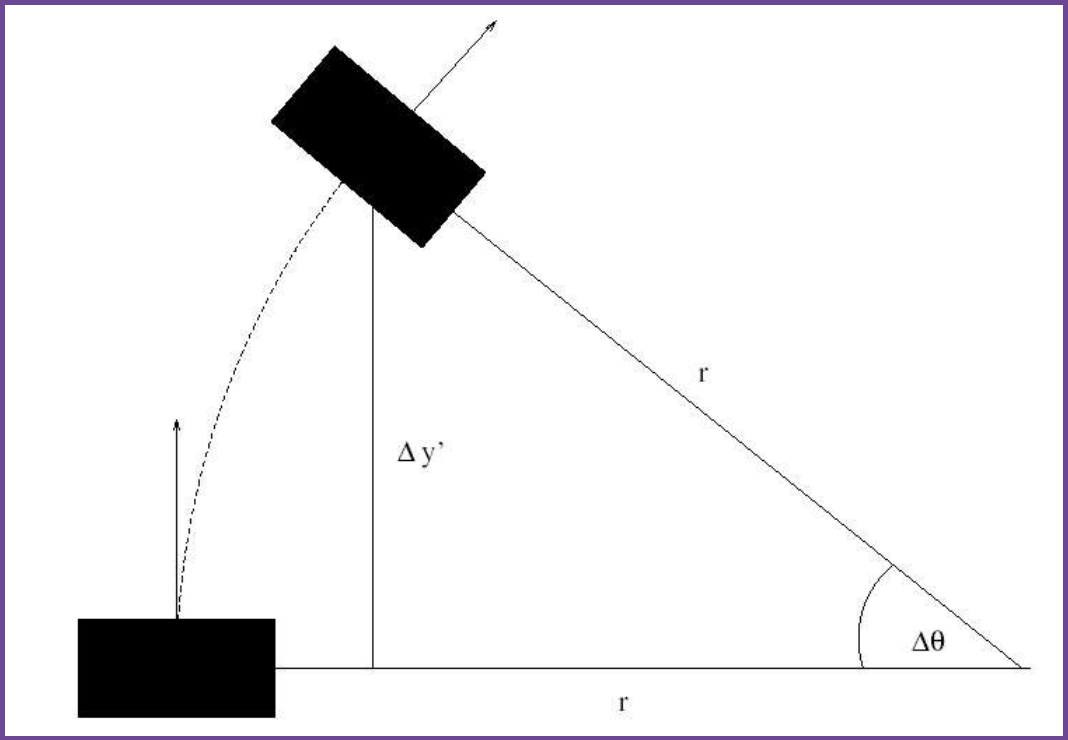

$$r = \frac{\Delta y'}{\sin(\theta_2' - \theta_1')}$$

$$v = \frac{r * (\theta_2' - \theta_1')}{\Delta t}$$

Detecting Walls in Robotic Soccer

- Easier in the robotic soccer domain
  - Color of the wall is known
  - Only object of this color in the domain
- Still not easy, because of occlusions
- "Drill" upwards, find two consecutive segments that coincide. Conservative.

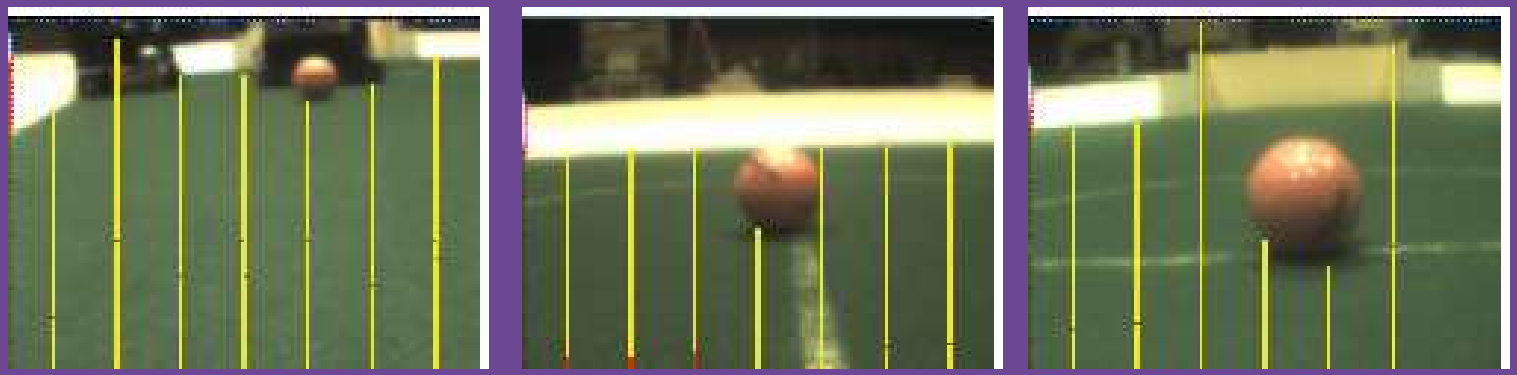

Detecting Walls in Robotic Rescue

- Lines are more stable than single feature points
- Edge detection, followed by line detection (Hough transform).
- Multiple lines
- Lines close to the robot
- Conservative parameters (Sufficient conditions)

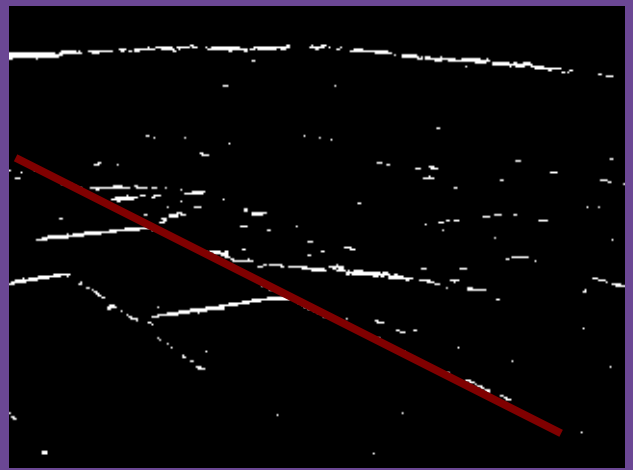

SLAM

- Optical flow used to generate a map of the environment
- Chicken and egg problem with egomotion estimation
- No maps are better than wrong maps
- Create sets of local 2D maps (100cm by 100cm)
- Seperated by links and long traversals
- Use maps to determine if a position has been visited previously (In the future: add images)

Local Maps

- Occupancy grids for 100cm by 100cm, 10cm resolution

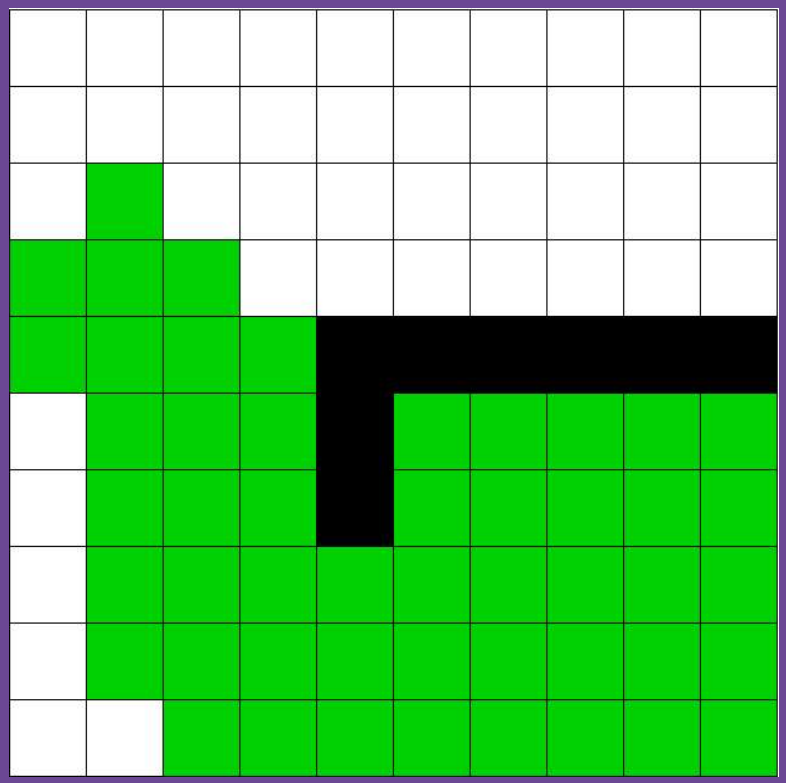

Detecting Victims

- Skin color predicate
  - Based on R,G,B, R-G, R-B, G-B thresholds
  - Normalized RGB thresholds (e.g., r/(r+g+b+1))
  - Ignore highlights
  - Compactness of the region
- A lot of work in the face recognition community

Conclusion

- Performance in the competition was poor
- Hardware failures => More robust hardware
- However, will be poor in the foreseeable future
  - Optical flow in real-time is hard
  - Easier if higher FPS. Focus on important regions
    - In front
    - At the horizon

Comments about the Competition

- Rules committee for the rescue competition needs to be revised. No activity for a long time
- Points for additional sensors should be allocated only
  - for when a victim was first identified using a given sensor
  - Multiple sensors should allow you to find more victims
  - Similar reasoning to why we took away the multiple robot multipliers
- Rules should be flexible enough to allow for fully autonomous solutions to at least enter