SLAM with Sparse Sensing

Kris Beevers and Wes Huang

Department of Computer Science
Rensselaer Polytechnic Institute
{beevek,whuang}@cs.rpi.edu

May 17, 2006

Motivation

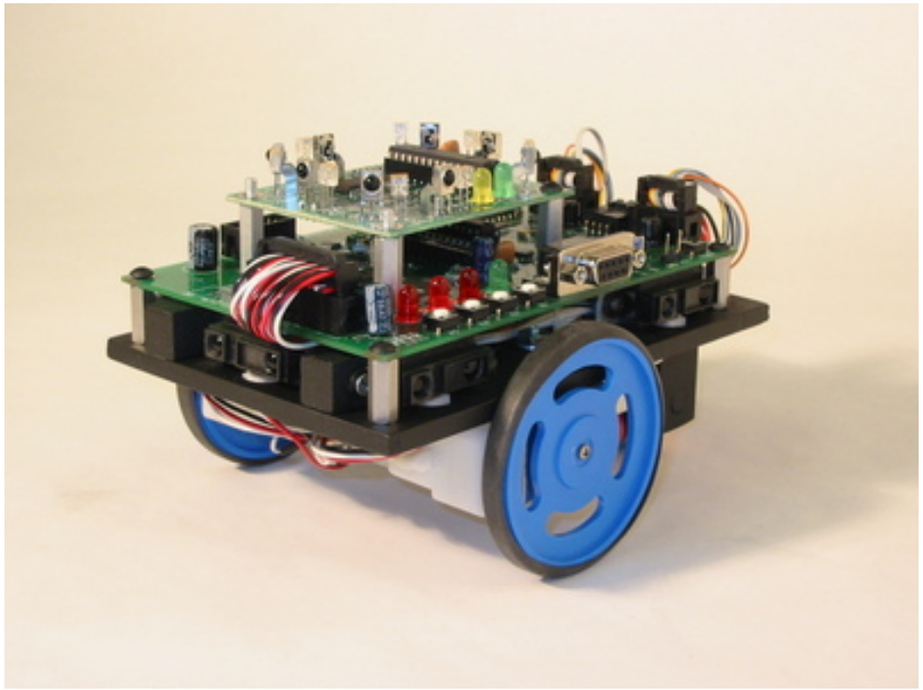

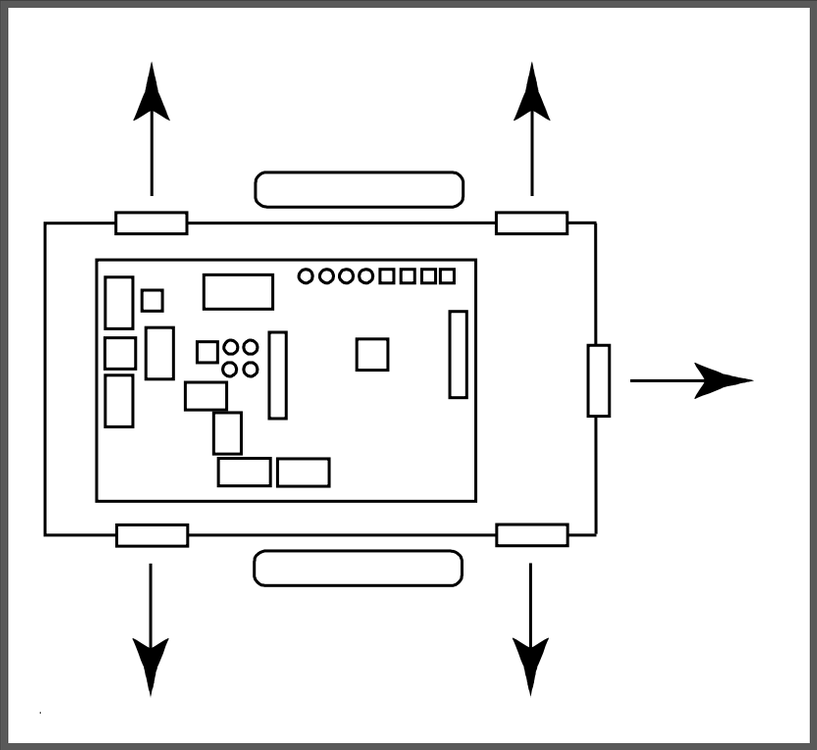

Goal: mapping/navigation with **limited sensing and computation**

The general problem

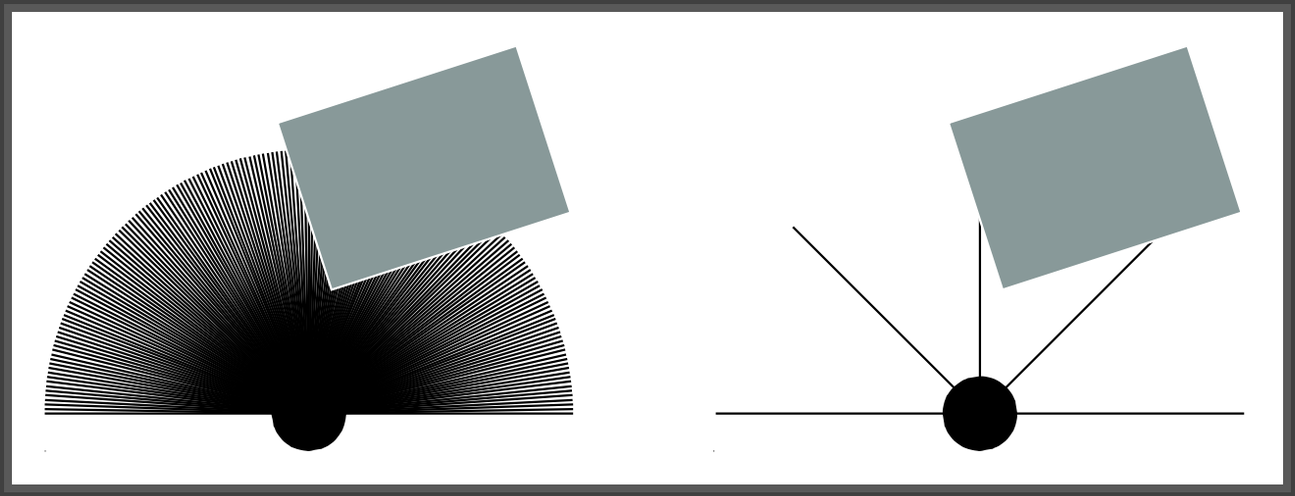

- Cost: laser rangefinder: US $5000 vs. IR array: US $50
- Low-power, small, lightweight
- Consumer robots, disposable robots, . . .

but

- Short range
- **Low spatial resolution (“sparse”)**
- Impossible to extract features from a single scan

Previous work

Topological mapping with limited sensing: Acar et al. (2001); Tovar et al. (2003); Huang and Beevers (2005)

Bearing-only SLAM: Deans and Hebert (2000); Bailey (2003); Solá et al. (2005)

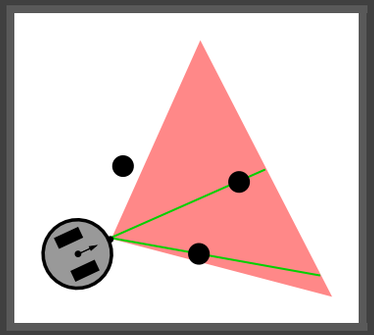

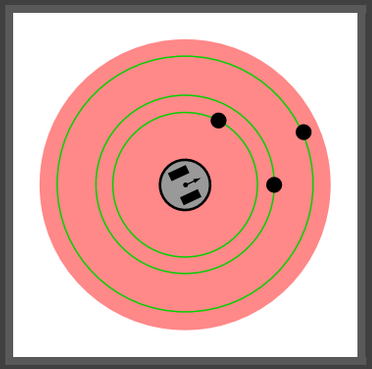

Range-only SLAM (with RF beacons): Kantor and Singh (2002); Kurth (2004); Djugash et al. (2005)

SONAR-based SLAM: Wijk and Christensen (2000); Zunino and Christensen (2001); Tardós et al. (2002); Leonard et al. (2002)

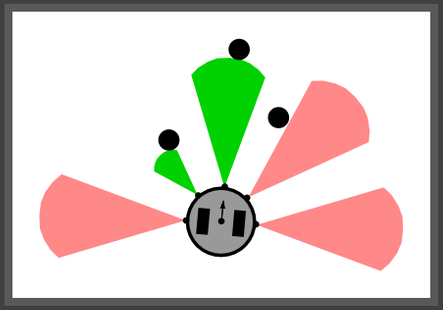

SLAM with multiscans

- Basic idea: extract features using scans from consecutive poses
- Treat groups of m poses as a single "**multiscan**"
- Enough data for feature extraction
- Tradeoff: pose uncertainty contributes to measurement uncertainty
- Full formulation in the paper

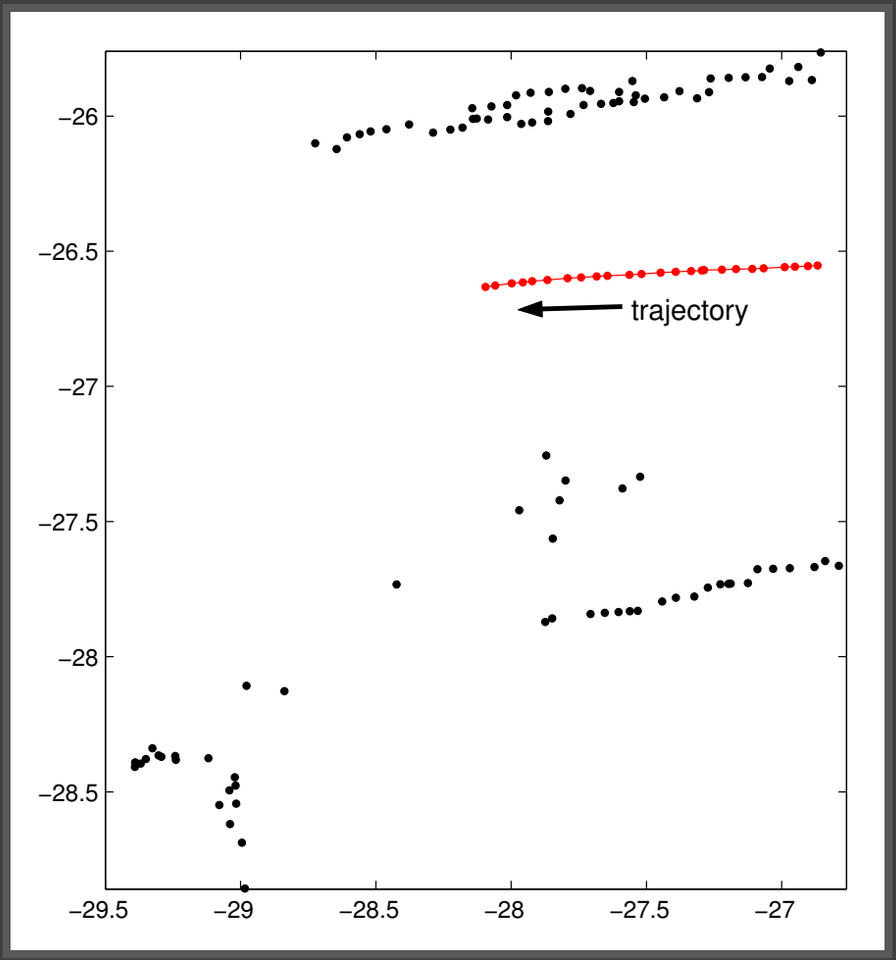

SLAM with multiscans

Modified Rao-Blackwellized particle filter (RBPF):

```
for k = 1 ... done do
    for t = k ... k + m do
        for all particles p^i do
            x_r^i(t) ~ p(x_r(t) | x_r(t − 1), u(t − 1))
        end for
    end for
    for all particles p^i do
        Extract features from z = [z(k + 1)  ...  z(k + m)]^T
        Find correspondences with x_f^i(k − m)
        Update map and initialize new features
        ω^i = p(z | x_r^i(k + 1), ..., x_r^i(k + m), x_f^i(k − m))
    end for
    If necessary, resample particles according to the ω^i s
end for
```

SLAM with multiscans

Modified Rao-Blackwellized particle filter (RBPF):

```
for k = 1 ... done do
    for t = k ... k + m do
        for all particles p^i do
            x_r^i(t) ~ p(x_r(t) | x_r(t − 1), u(t − 1))
        end for
    end for
    for all particles p^i do
        Extract features from z = [z(k + 1)  ...  z(k + m)]^T
        Find correspondences with x_f^i(k − m)
        Update map and initialize new features
        ω^i = p(z | x_r^i(k + 1), ..., x_r^i(k + m), x_f^i(k − m))
    end for
    If necessary, resample particles according to the ω^i's
end for
```

SLAM with multiscans

Modified Rao-Blackwellized particle filter (RBPF):

```
for k = 1 ... done do
    for t = k ... k + m do
        for all particles p^i do
            x_r^i(t) ~ p(x_r(t) | x_r(t − 1), u(t − 1))
        end for
    end for
    for all particles p^i do
        Extract features from z = [z(k + 1)  ...  z(k + m)]^T
        Find correspondences with x_f^i(k − m)
        Update map and initialize new features
        ω^i = p(z | x_r^i(k + 1), ..., x_r^i(k + m), x_f^i(k − m))
    end for
    If necessary, resample particles according to the ω^i's
end for
```

Feature extraction with expected trajectory

- Extract features once using *expected trajectory* over multiscan:

$$\mathbf{E}\left[\mathbf{x}_r(k)\right] = \overline{\mathbf{x}}_r(k) = \left[\overline{x}_r(k-m+1) \quad \ldots \quad \overline{x}_r(k)\right]^T$$

- Transform features for each particle
- Saves a lot of computation
- Tradeoff:
  - large m: more data for extraction
  - small m: better approximation of per-particle extraction
- More details in the paper

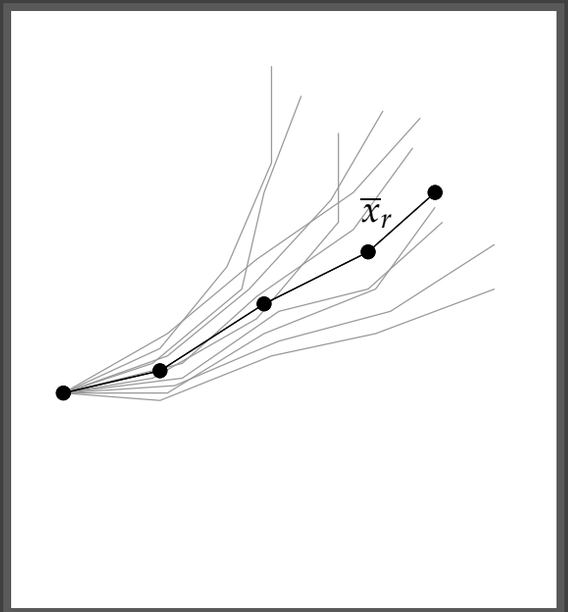

- Alternative: group particles into “strata,” do extraction per-strata

Feature extraction with expected trajectory

- Extract features once using *expected trajectory* over multiscan:

$$\mathbf{E}\left[\mathbf{x}_r(k)\right] = \overline{\mathbf{x}}_r(k) = \left[\overline{x}_r(k-m+1) \quad \ldots \quad \overline{x}_r(k)\right]^T$$

- Transform features for each particle
- Saves a lot of computation
- Tradeoff:
  - large m: more data for extraction
  - small m: better approximation of per-particle extraction
- More details in the paper

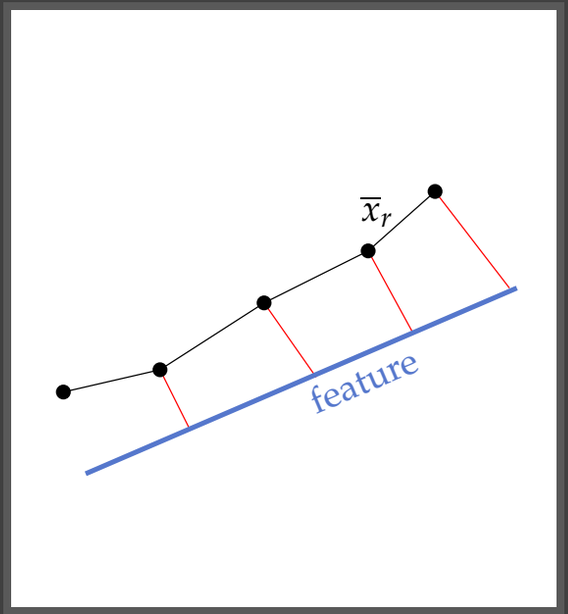

- Alternative: group particles into “strata,” do extraction per-strata

Feature extraction with expected trajectory

- Extract features once using *expected trajectory* over multiscan:

$$\mathbf{E}\left[\mathbf{x}_r(k)\right] = \overline{\mathbf{x}}_r(k) = \left[\overline{x}_r(k-m+1) \quad \ldots \quad \overline{x}_r(k)\right]^T$$

- Transform features for each particle
- Saves a lot of computation
- Tradeoff:
  - large m: more data for extraction
  - small m: better approximation of per-particle extraction
- More details in the paper

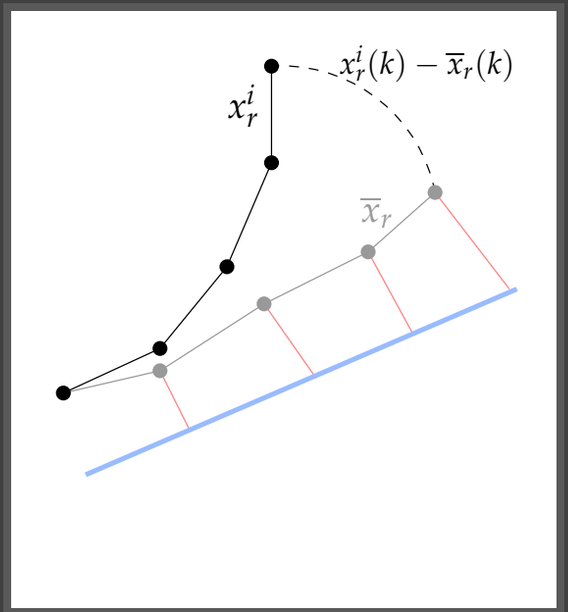

- Alternative: group particles into “strata,” do extraction per-strata

Feature extraction with expected trajectory

- Extract features once using *expected trajectory* over multiscan:

$$\mathbf{E}\left[\mathbf{x}_r(k)\right] = \overline{\mathbf{x}}_r(k) = \left[\overline{x}_r(k-m+1) \quad \ldots \quad \overline{x}_r(k)\right]^T$$

- Transform features for each particle
- Saves a lot of computation
- Tradeoff:
  - large m: more data for extraction
  - small m: better approximation of per-particle extraction
- More details in the paper

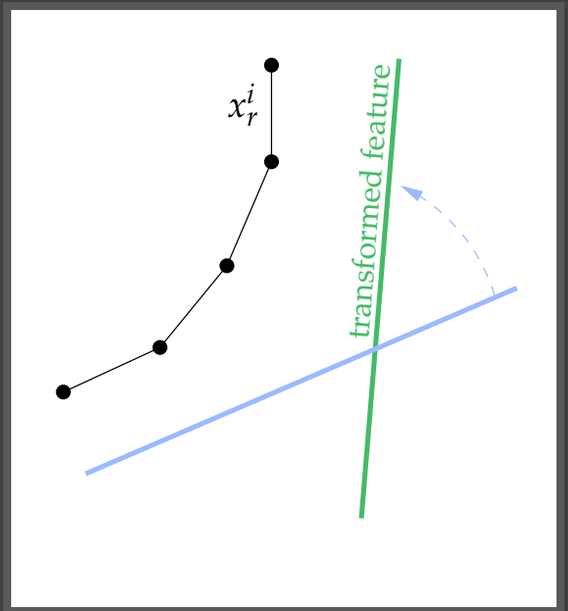

- Alternative: group particles into “strata,” do extraction per-strata

Feature extraction with expected trajectory

- Extract features once using *expected trajectory* over multiscan:

$$\mathbf{E}\left[\mathbf{x}_r(k)\right] = \overline{\mathbf{x}}_r(k) = \left[\overline{x}_r(k-m+1) \quad \ldots \quad \overline{x}_r(k)\right]^T$$

- Transform features for each particle
- Saves a lot of computation
- Tradeoff:
  - large m: more data for extraction
  - small m: better approximation of per-particle extraction
- More details in the paper

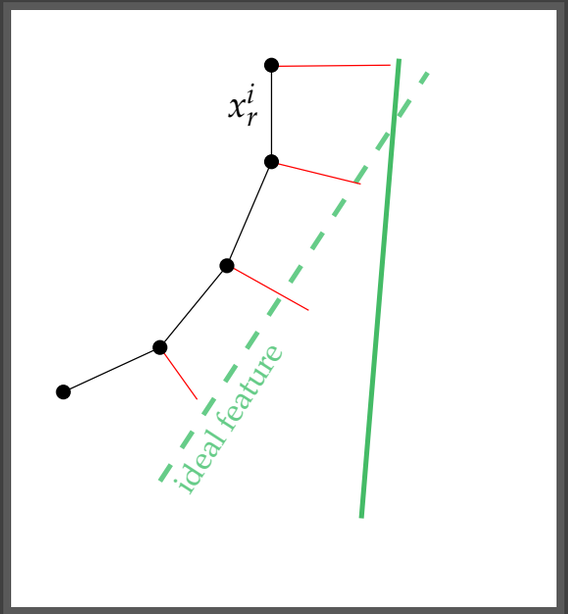

- Alternative: group particles into “strata,” do extraction per-strata

Results: implementation details

- Implementation:
  - RBPF similar to FastSLAM 1.0 (Thrun et al., 2004)
  - Adaptive resampling ($\hat{N}_{\text{eff}}$) (Grisetti et al., 2005)
- Line and line segment features
  - Data segmentation with adaptive agglomerative clustering
  - See the paper
- Datasets:
  - Data from Radish (Howard and Roy, 2003, `radish.sf.net`)
  - Laser datasets (), "sparsified" ()
  - $\{0^\circ, 45^\circ, 90^\circ, 135^\circ, 180^\circ\}$

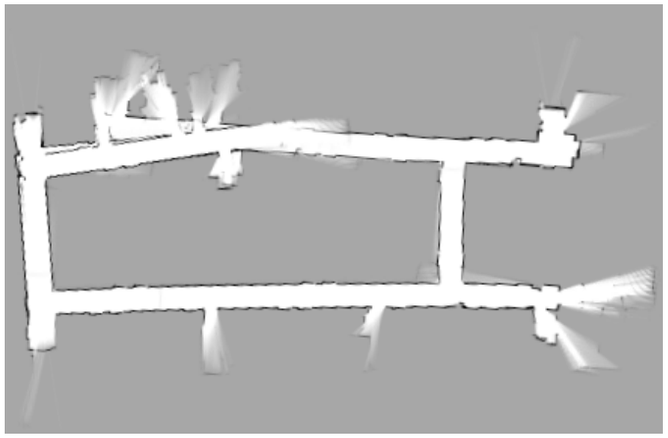

Raw data

USC SAL

$\sim$ 39 m $\times$ 20 m

100 particles

Line features

Max range 5 m

$m = 50$

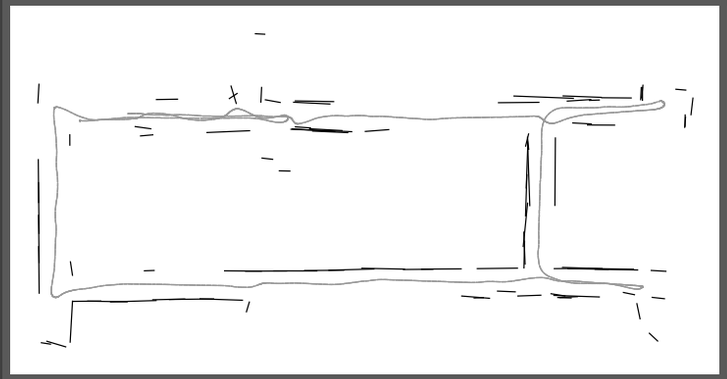

Landmark map (✳)

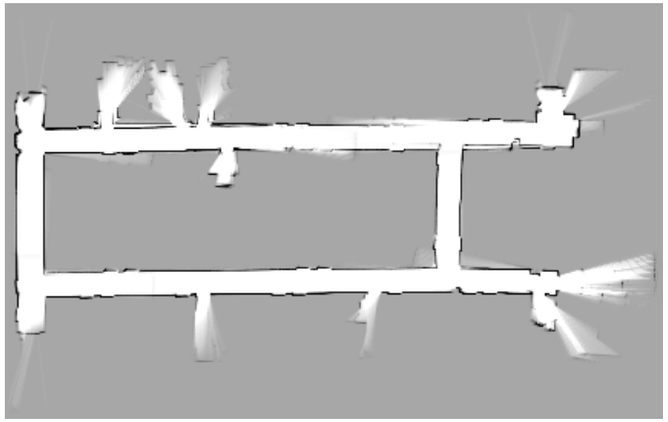

Occupancy map (✳)

Dataset from Radish courtesy of Andrew Howard.

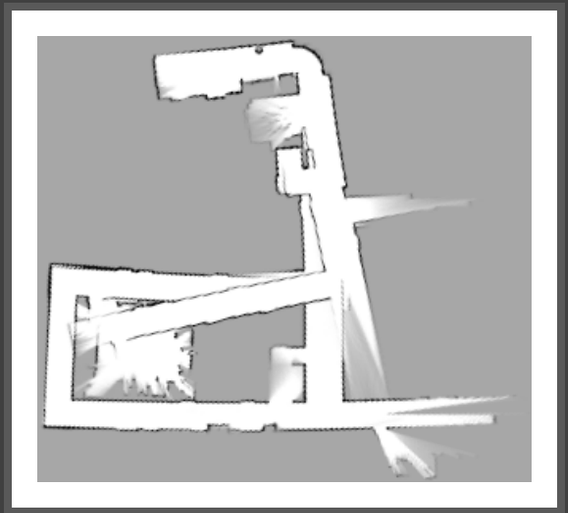

Raw data

CMU NSH

$\sim$ 25 m $\times$ 25 m

600 particles

Segment features

Max range 3 m

$m = 40$

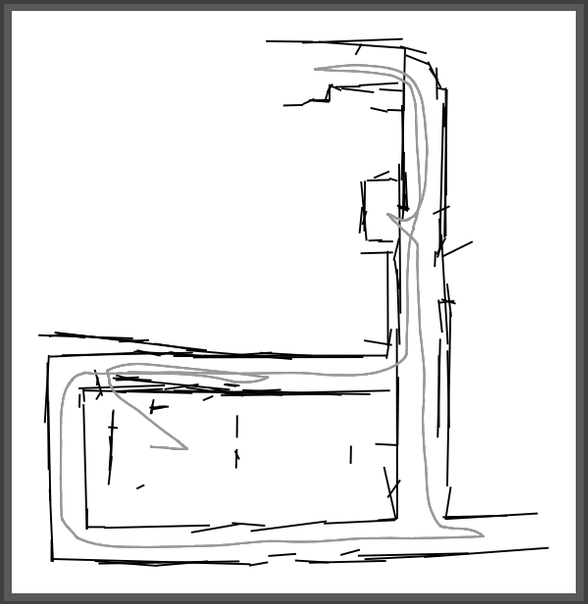

Landmark map ( )

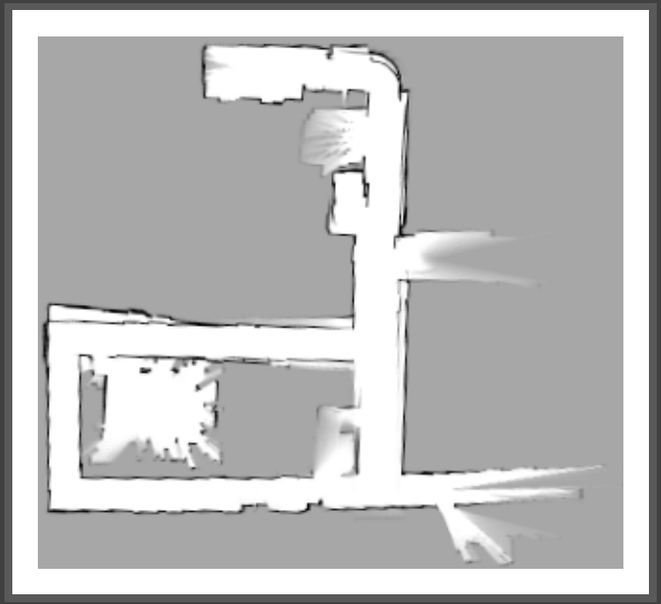

Occupancy map ( )

Dataset from Radish courtesy of Nick Roy.

Raw data

Stanford

$\sim$ 64 m $\times$ 56 m

1000 particles

Segment features

Max range 5 m

$m = 18$

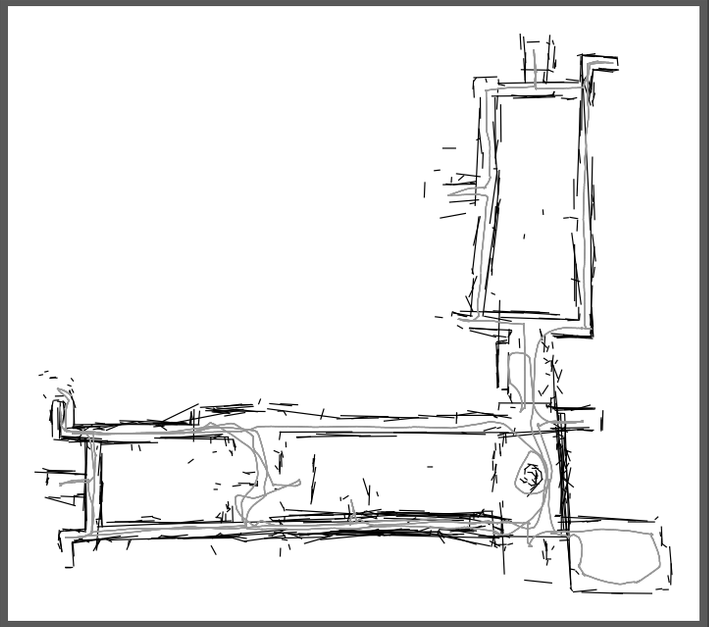

Landmark map ( )

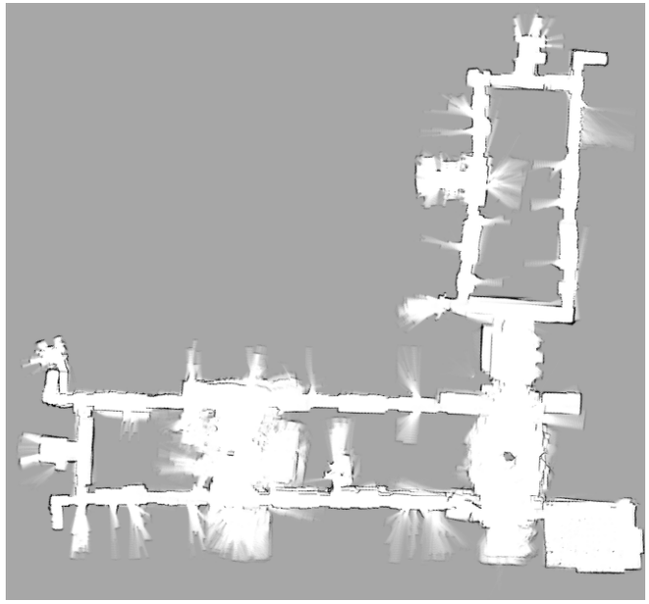

Occupancy map ( )

Dataset from Radish courtesy of Brian Gerkey.

Comparison

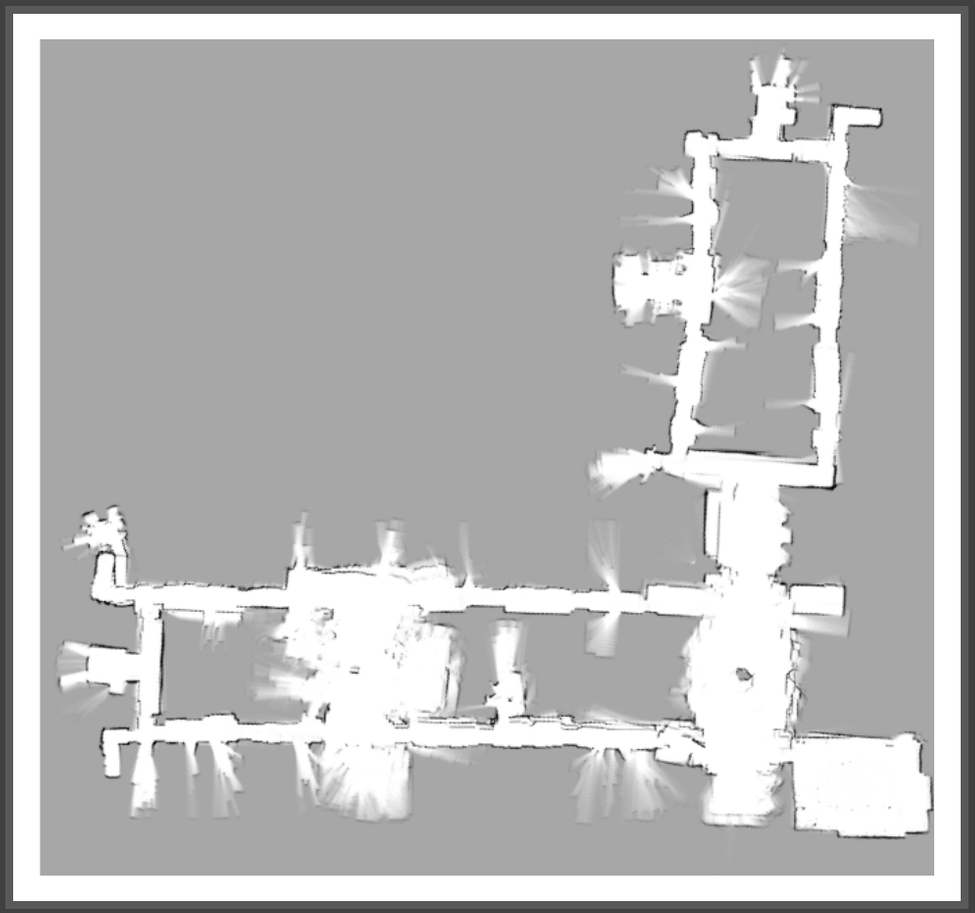

Multiscan SLAM ( )

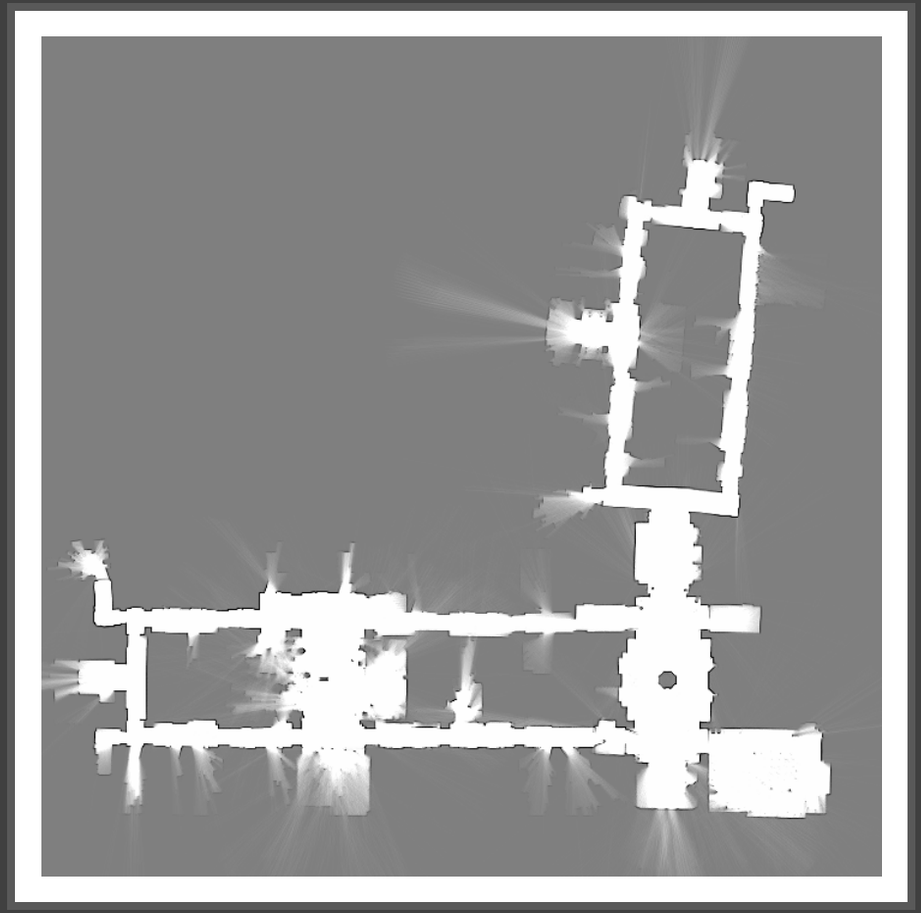

Scan-matching SLAM ( )

Scan-matching result from Radish courtesy of Brian Gerkey.

Discussion

- SLAM in large indoor environments is possible with
- Complexity: same as normal RBPF

- Augmenting sensing with odometry: lots of uncertainty
  - Unpeaked measurement distributions
  - Simple remedy: more particles
  - Other approaches: exploiting prior knowledge, better PF samples, etc.
- How to pick m?
  - Tuning parameter
  - Adaptively based on accumulated uncertainty

Summary

- Contributions:
  - **Algorithm for RBPF SLAM with sparse sensing** ( )
    - Group data from consecutive poses into “multiscans”
    - Do feature extraction on multiscans
  - **Approximations for efficiency**
    - Feature extraction using expected trajectory
  - **Simple feature extraction for multiple-pose data**
    - Segmentation based on agglomerative clustering
- Results: successful mapping in several large, real indoor environments

Exploiting prior knowledge

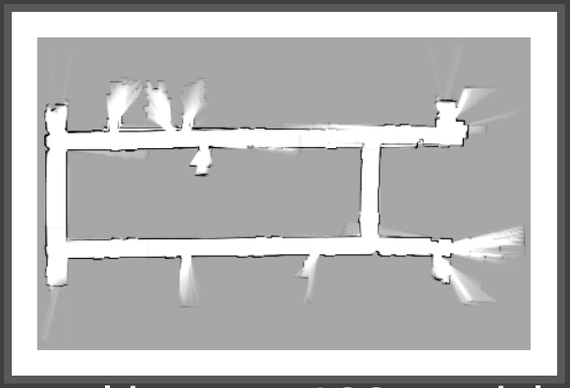

Plain multiscans: 100 particles

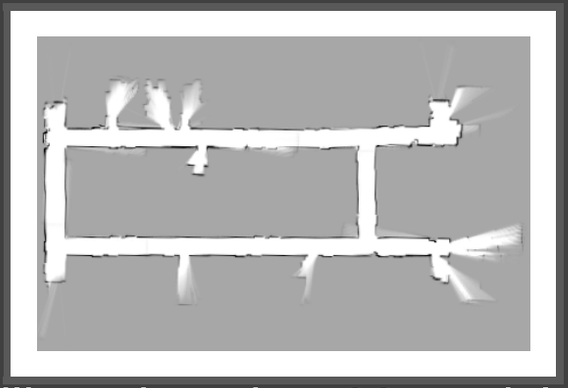

Rectilinearity prior: 20 particles

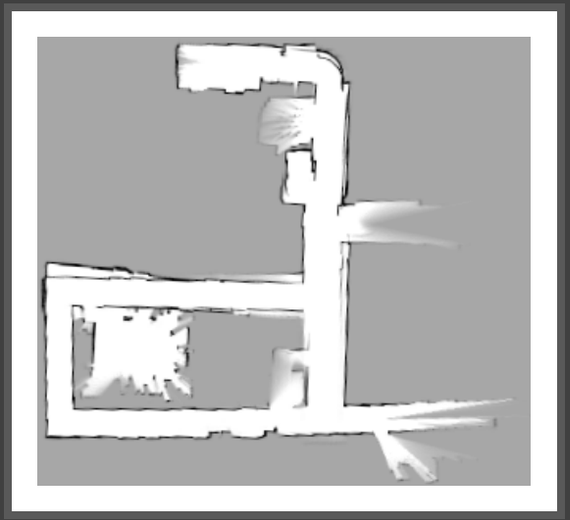

Plain multiscans: 600 particles

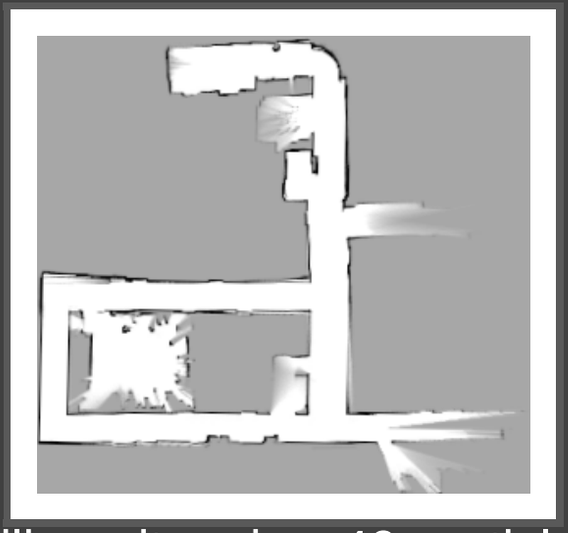

Rectilinearity prior: 40 particles

K. Beevers and W. Huang. Inferring and enforcing relative constraints in SLAM, *WAFR 2006*, to appear.

Thank you!

Leonard et al. (2002)

- "Mapping partially observable features from multiple uncertain vantage points"
- Basic idea:
  - *At every timestep*, extract features from data of last m poses
  - Add m-step pose history to state, filter with EKF
- Differences of our approach:
  - Only acquire enough data to observe feature parameters infrequently: do extraction/filtering every m timesteps
  - Tradeoff: less frequent filtering **but** much less computation
  - RBPF: enables approximations like extraction with expected trajectory
- Can combine the two approaches to play with the tradeoff
  - Every p timesteps, extract features using last m scans

References

E.U. Acar, H. Choset, and P.N. Atkar. Complete sensor-based coverage with extended-range detectors: A hierarchical decomposition in terms of critical points and Voronoi diagrams. In *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots & Systems*, pages 1305–1311, October 2001.

T. Bailey. Constrained initialisation for bearing-only SLAM. In *Proceedings of the 2003 IEEE International Conference on Robotics & Automation*, pages 1966–1971, September 2003.

M. Deans and M. Hebert. Experimental comparison of techniques for localization and mapping using a bearings only sensor. In *Proceedings of the Seventh International Symposium on Experimental Robotics (ISER)*, pages 395–404. Springer-Verlag, December 2000.

J. Djugash, S. Singh, and P. Corke. Further results with localization and mapping using range from radio. In *Proceedings of the Fifth International Conference on Field and Service Robotics*, Australia, July 2005.

Giorgio Grisetti, Cyrill Stachniss, and Wolfram Burgard. Improving grid-based SLAM with Rao-Blackwellized particle filters by adaptive proposals and

selective resampling. In *Proc. of the IEEE Intl. Conf. on Robotics and Automation*, pages 2443–2448, 2005.

Wesley H. Huang and Kristopher R. Beevers. Topological mapping with sensing-limited robots. In M. Erdmann et al., editors, *Algorithmic Foundations of Robotics VI*, pages 235–250. Springer, 2005.

G.A. Kantor and S. Singh. Preliminary results in range-only localization and mapping. In *Proceedings of the 2002 IEEE International Conference on Robotics & Automation*, volume 2, pages 1818–1823, May 2002.

D. Kurth. Range-only robot localization and SLAM with radio. Master's thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, May 2004.

J.J. Leonard, R. Rikoski, P. Newman, and M. Bosse. Mapping partially observable features from multiple uncertain vantage points. *Intl. Journal of Robotics Research*, 21(10):943–975, October 2002.

J. Solá, A. Monin, M. Devy, and T. Lemaire. Undelayed initialization in bearing only SLAM. In *Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots & Systems*, pages 2499–2504, August 2005.

J.D. Tardós, J. Neira, P.M. Newman, and J.J. Leonard. Robust mapping and localization in indoor environments using SONAR data. *Intl. Journal of Robotics Research*, 21(4):311–330, April 2002.

S. Thrun, M. Montemerlo, D. Koller, B. Wegbreit, J. Nieto, and E. Nebot. FastSLAM: An efficient solution to the simultaneous localization and mapping problem with unknown data association. *Journal of Machine Learning Research*, 2004. to appear.

B. Tovar, S. LaValle, and R. Murrieta. Optimal navigation and object finding without geometric maps or localization. In *Proceedings of the 2003 IEEE International Conference on Robotics & Automation*, Taipei, Taiwan, September 2003.

O. Wijk and H.I. Christensen. Triangulation based fusion of sonar data with application in robot pose tracking. *IEEE Transactions on Robotics and Automation*, 16(6):740–752, 2000.

G. Zunino and H.I. Christensen. Navigation in realistic environments. In M. Devy, editor, *9th Intl. Symp. on Intelligent Robotic Systems*, Toulouse, France, 2001.