Autonomous Robotics

Particle Filters

Josiah Hanna
University of Wisconsin — Madison

Some figures taken from Probabilistic Robotics and Drew Bagnell's course notes.

Learning Outcomes

After today's lecture, you will:

- Understand particle approximations of belief distributions.
- Be able to compute a robot's state estimate using a set of weighted particles.
- Understand the weighting and re-sampling schemes used by particle-based methods.

Particle Filter Applications

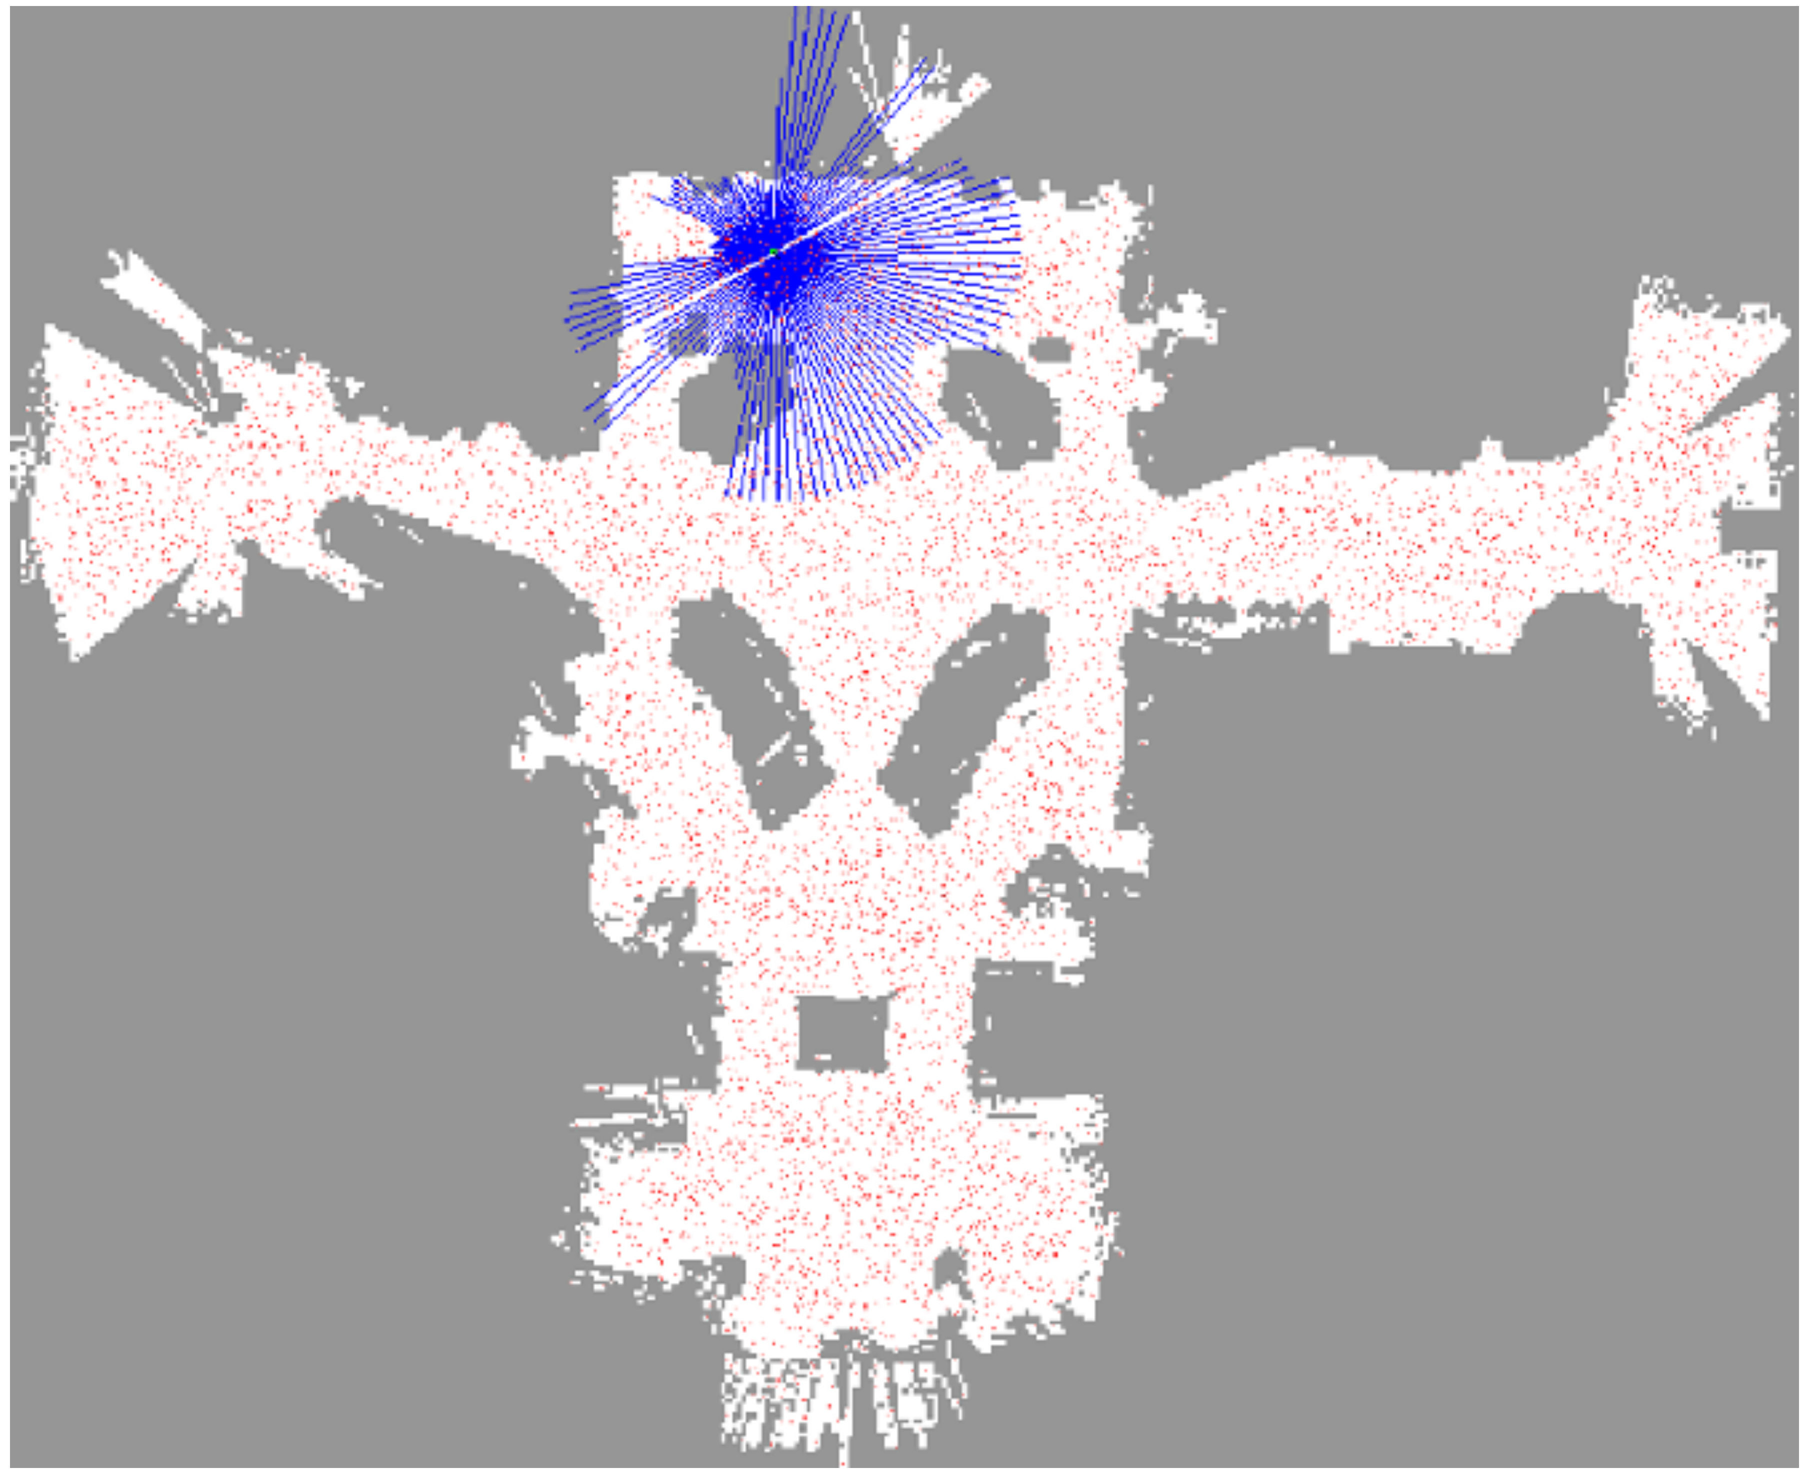

After Incorporating 65 Ultrasound Scans

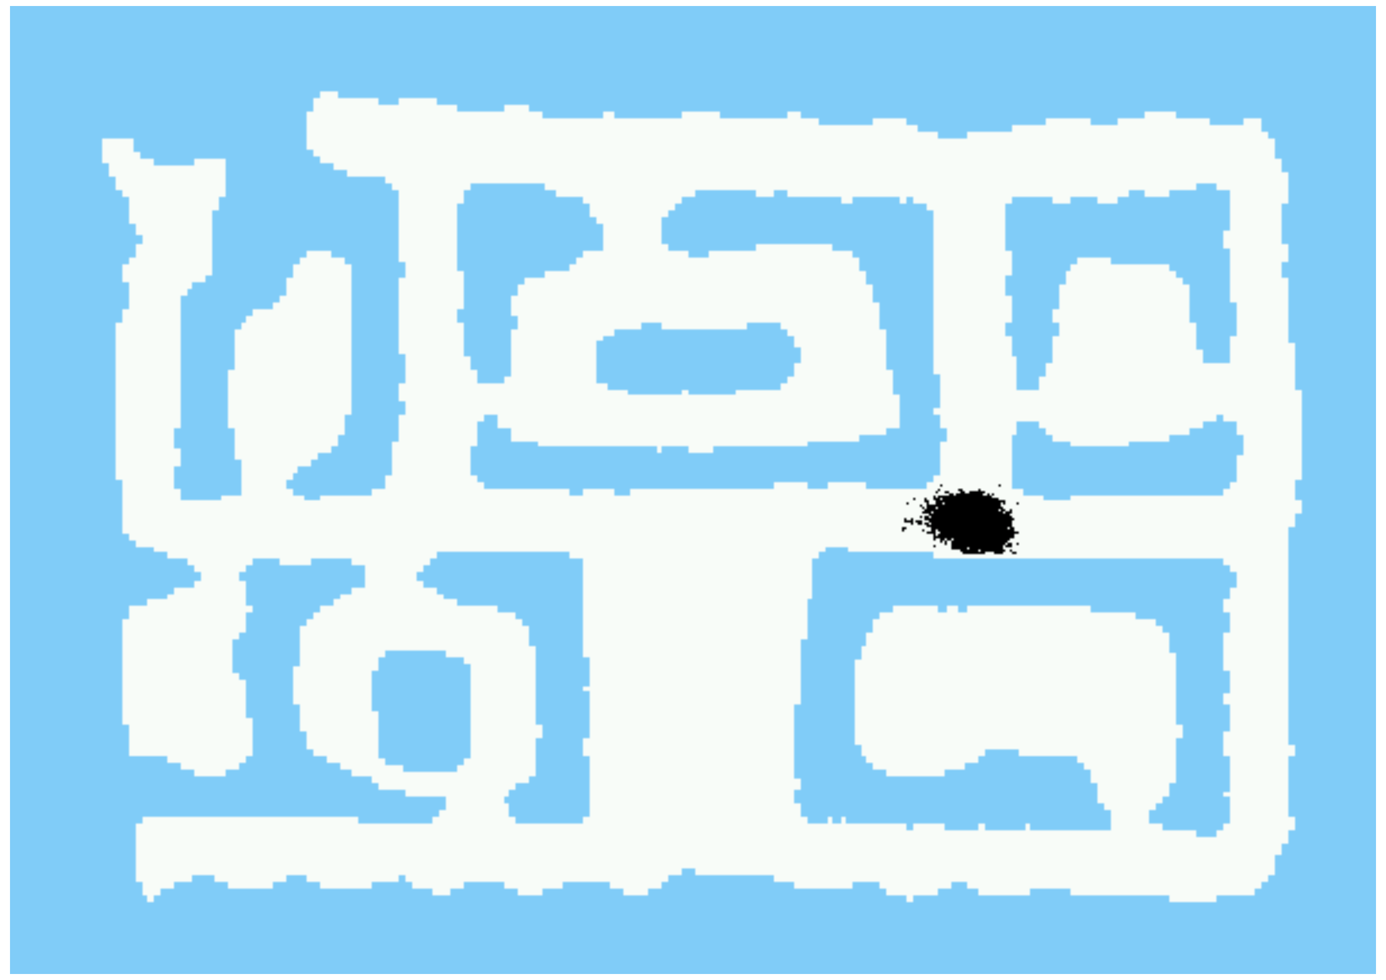

Particle Filter Applications

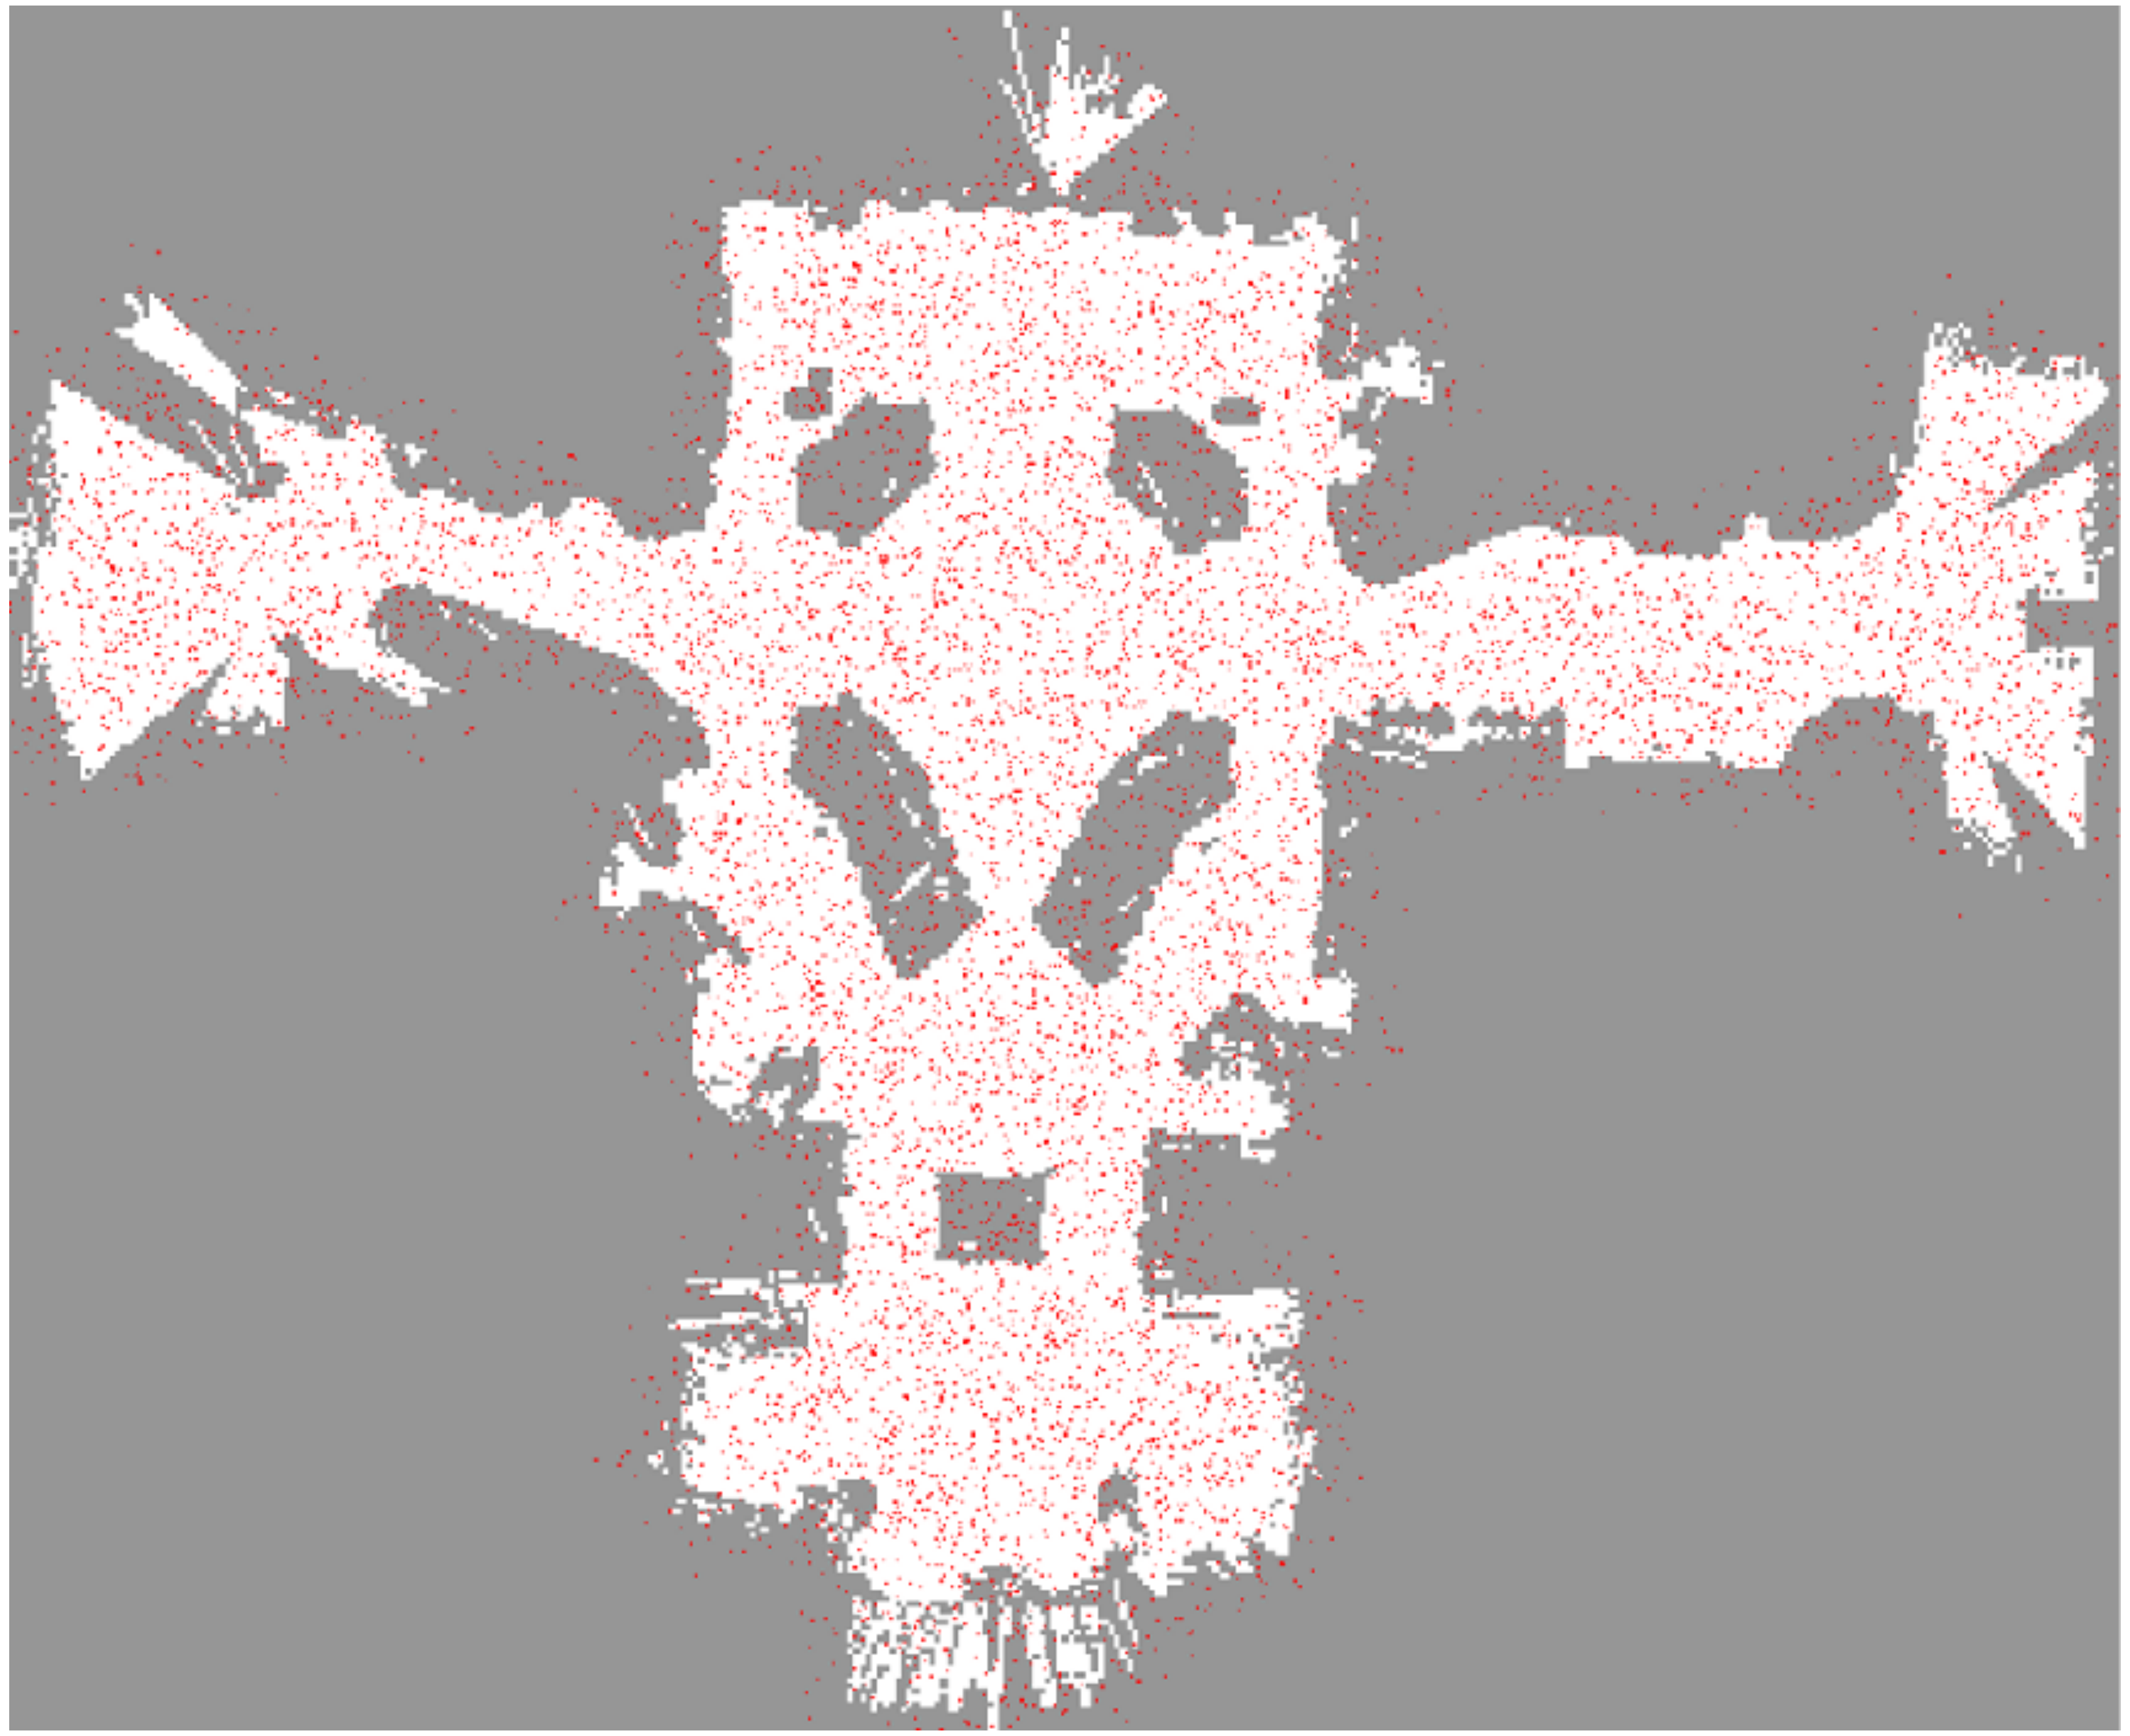

Particle Filter Applications

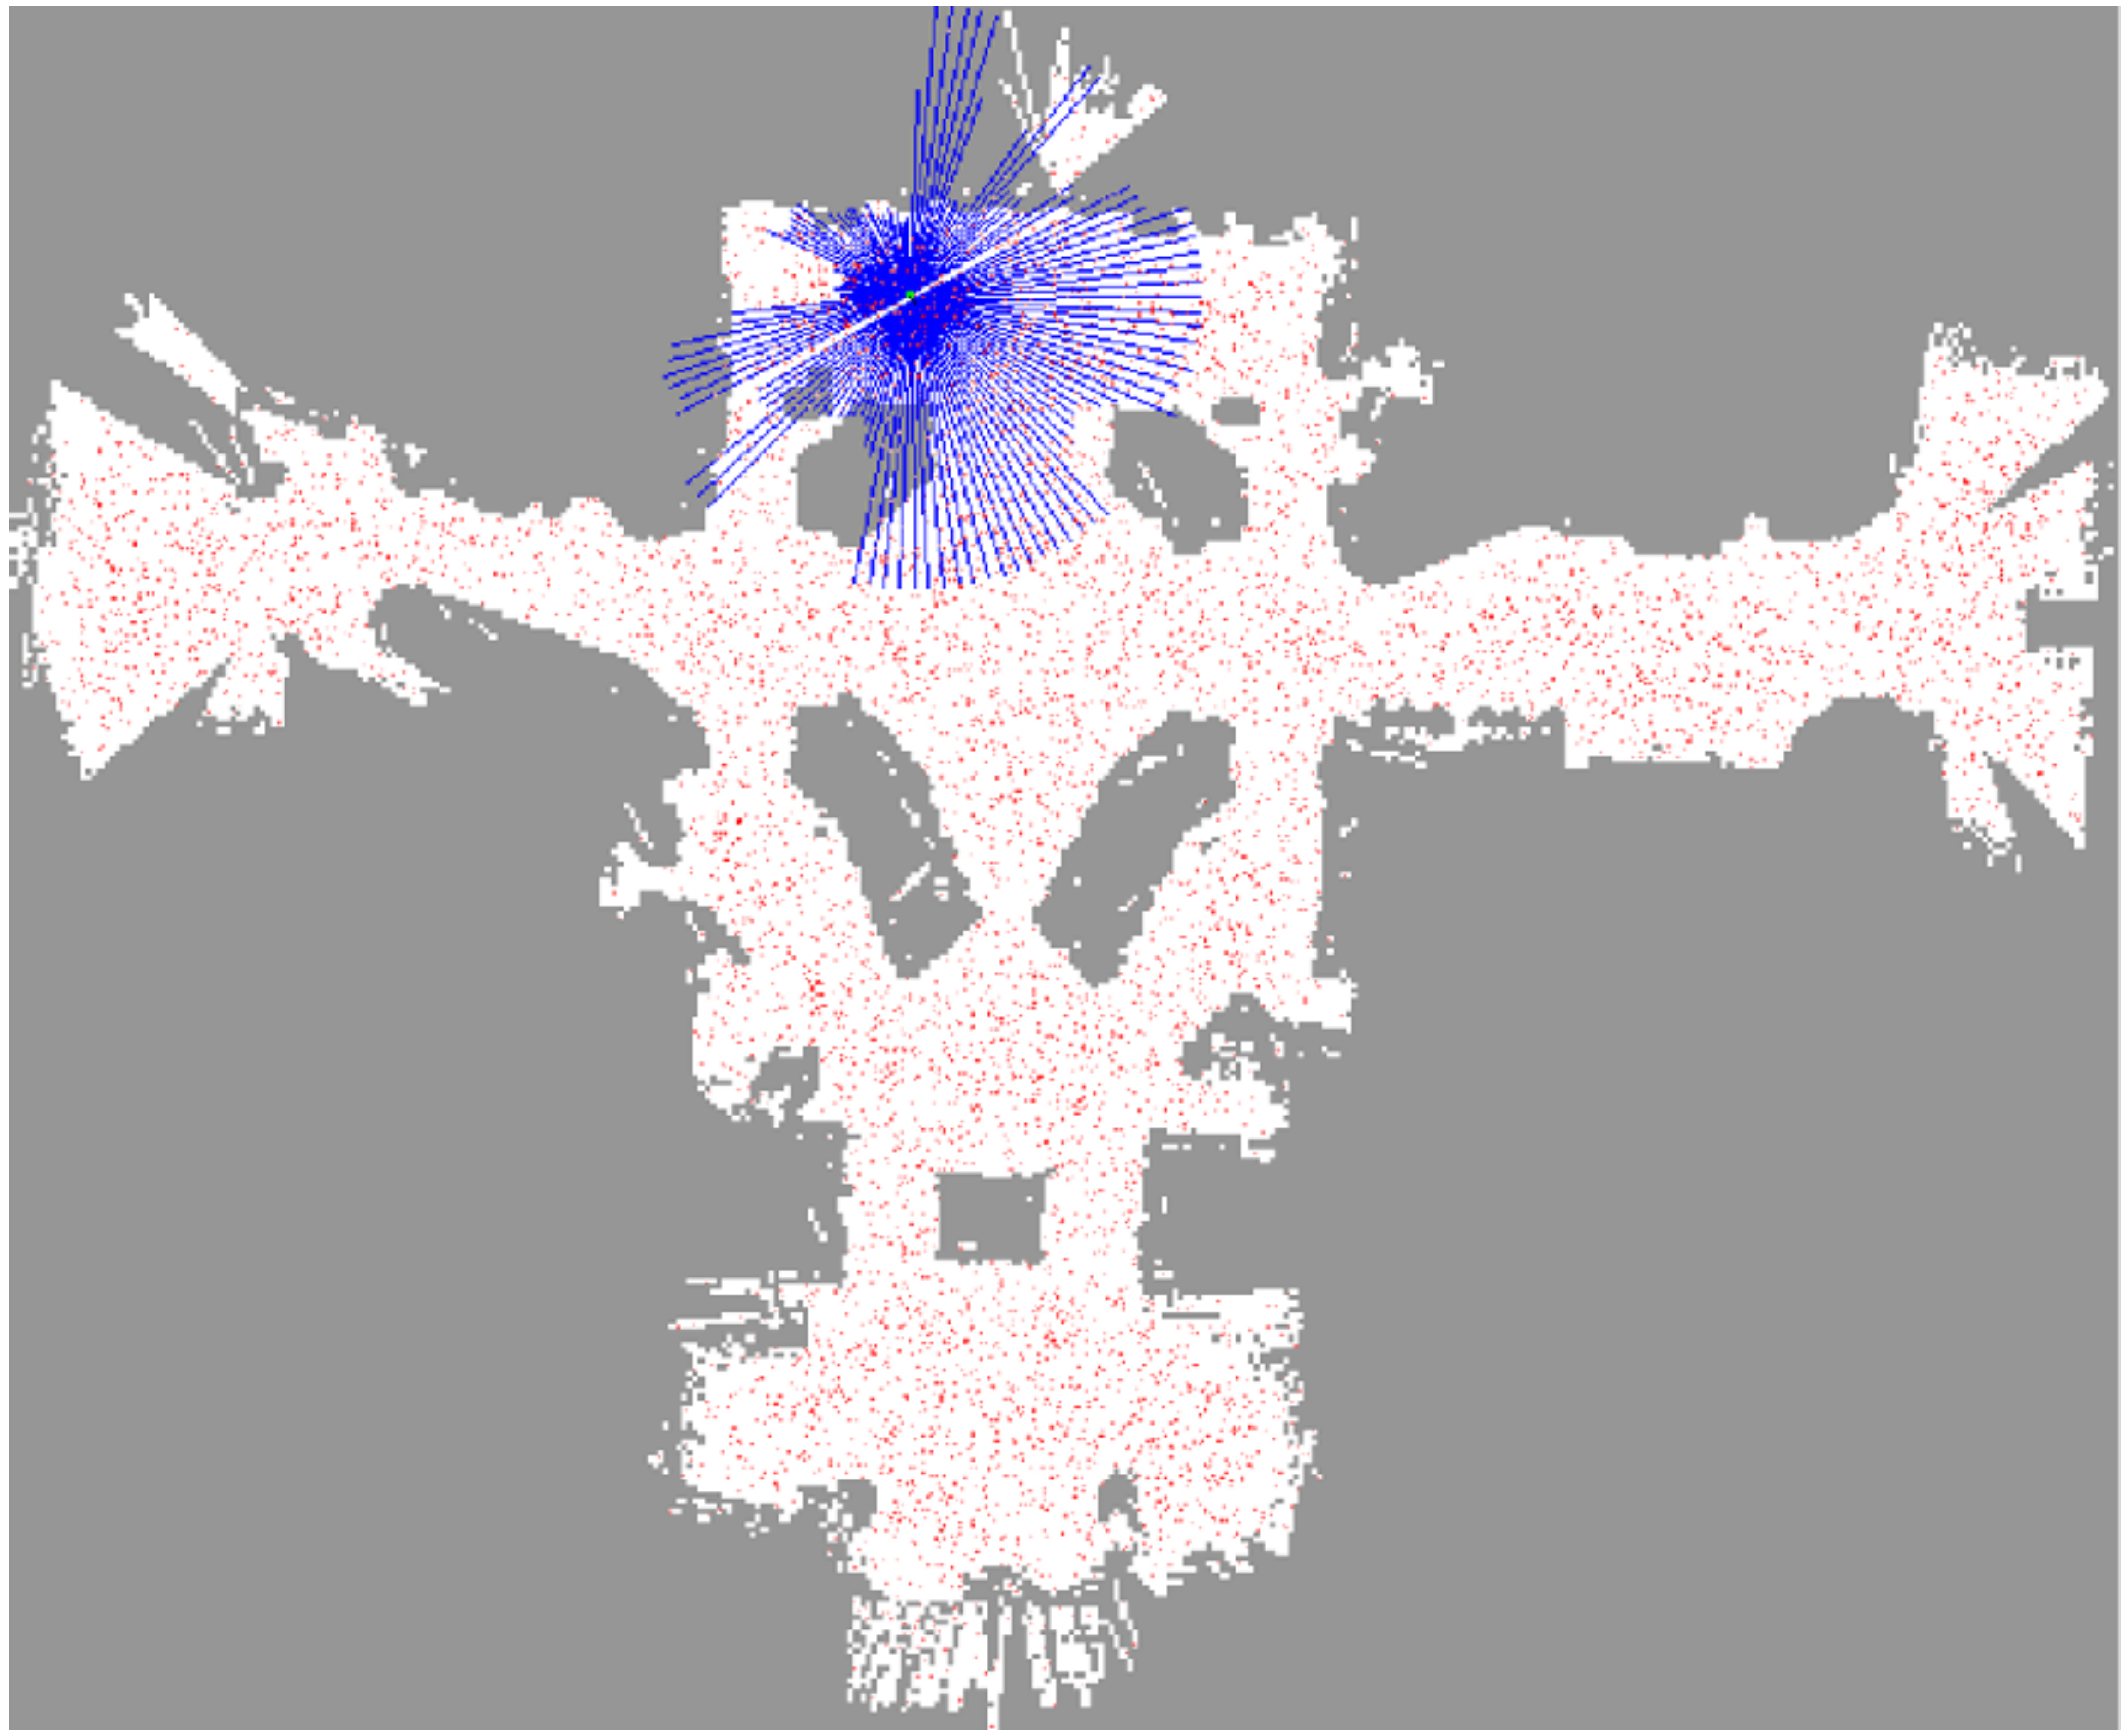

Particle Filter Applications

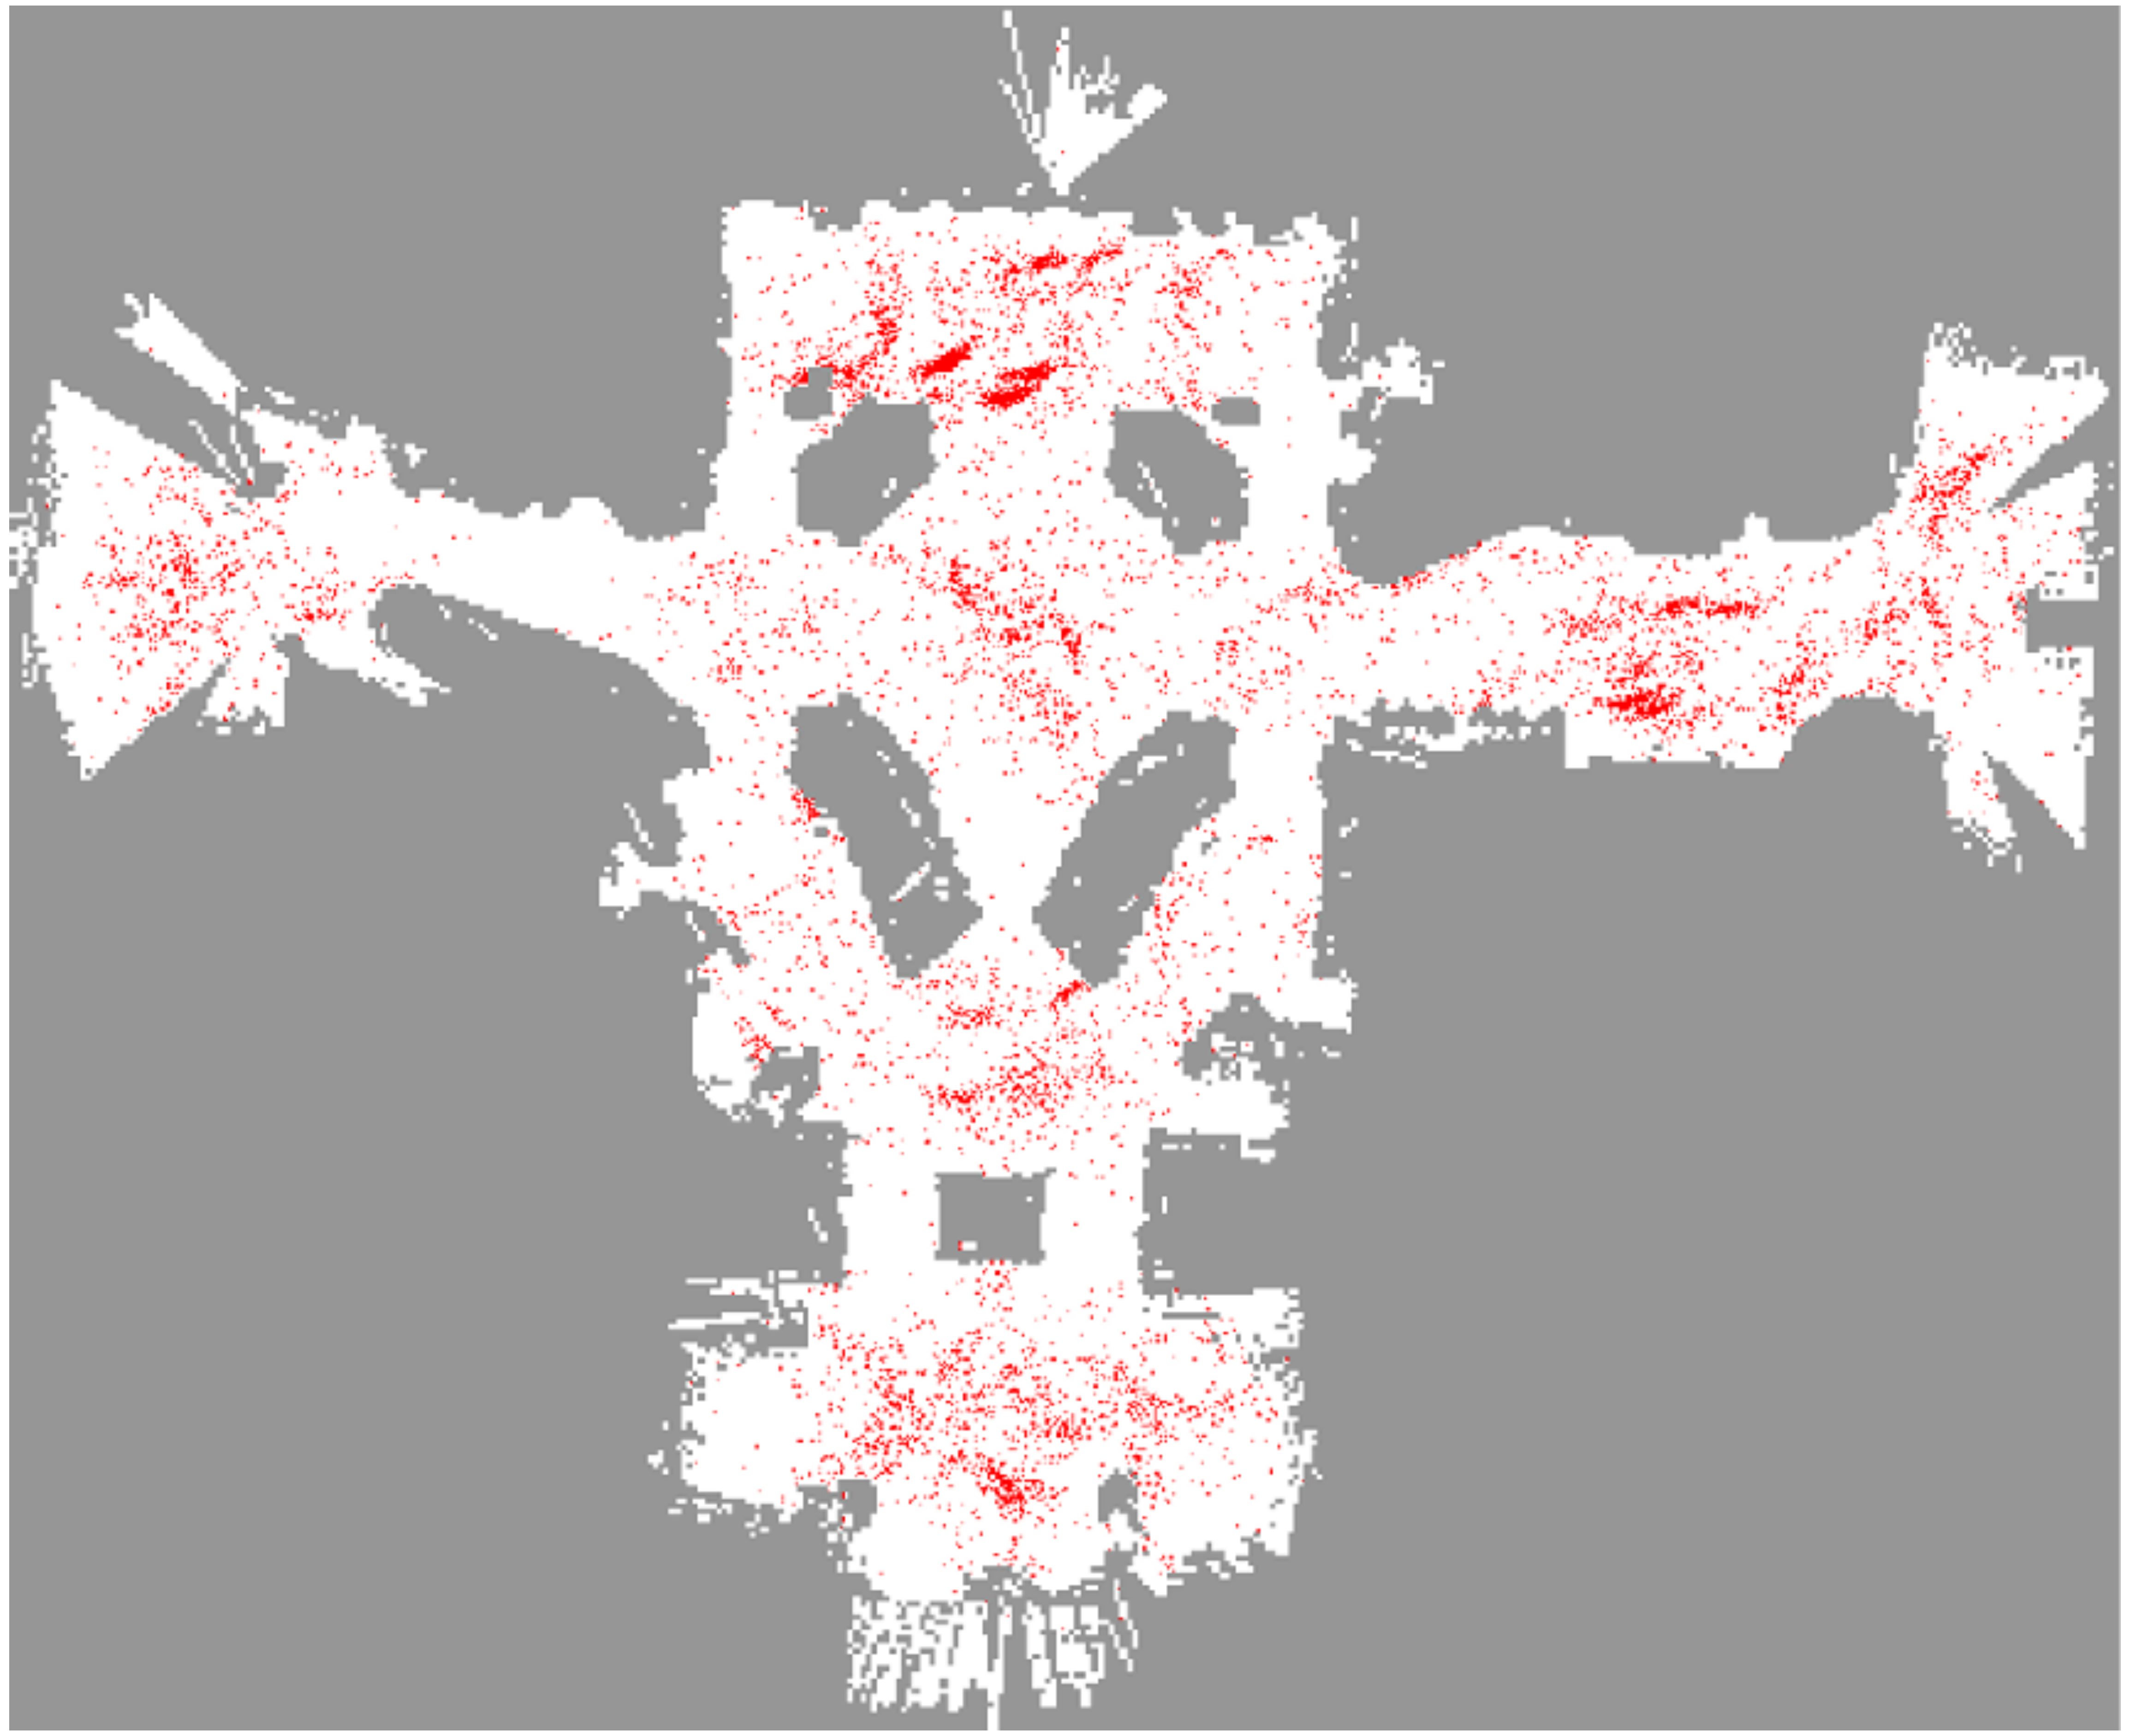

Particle Filter Applications

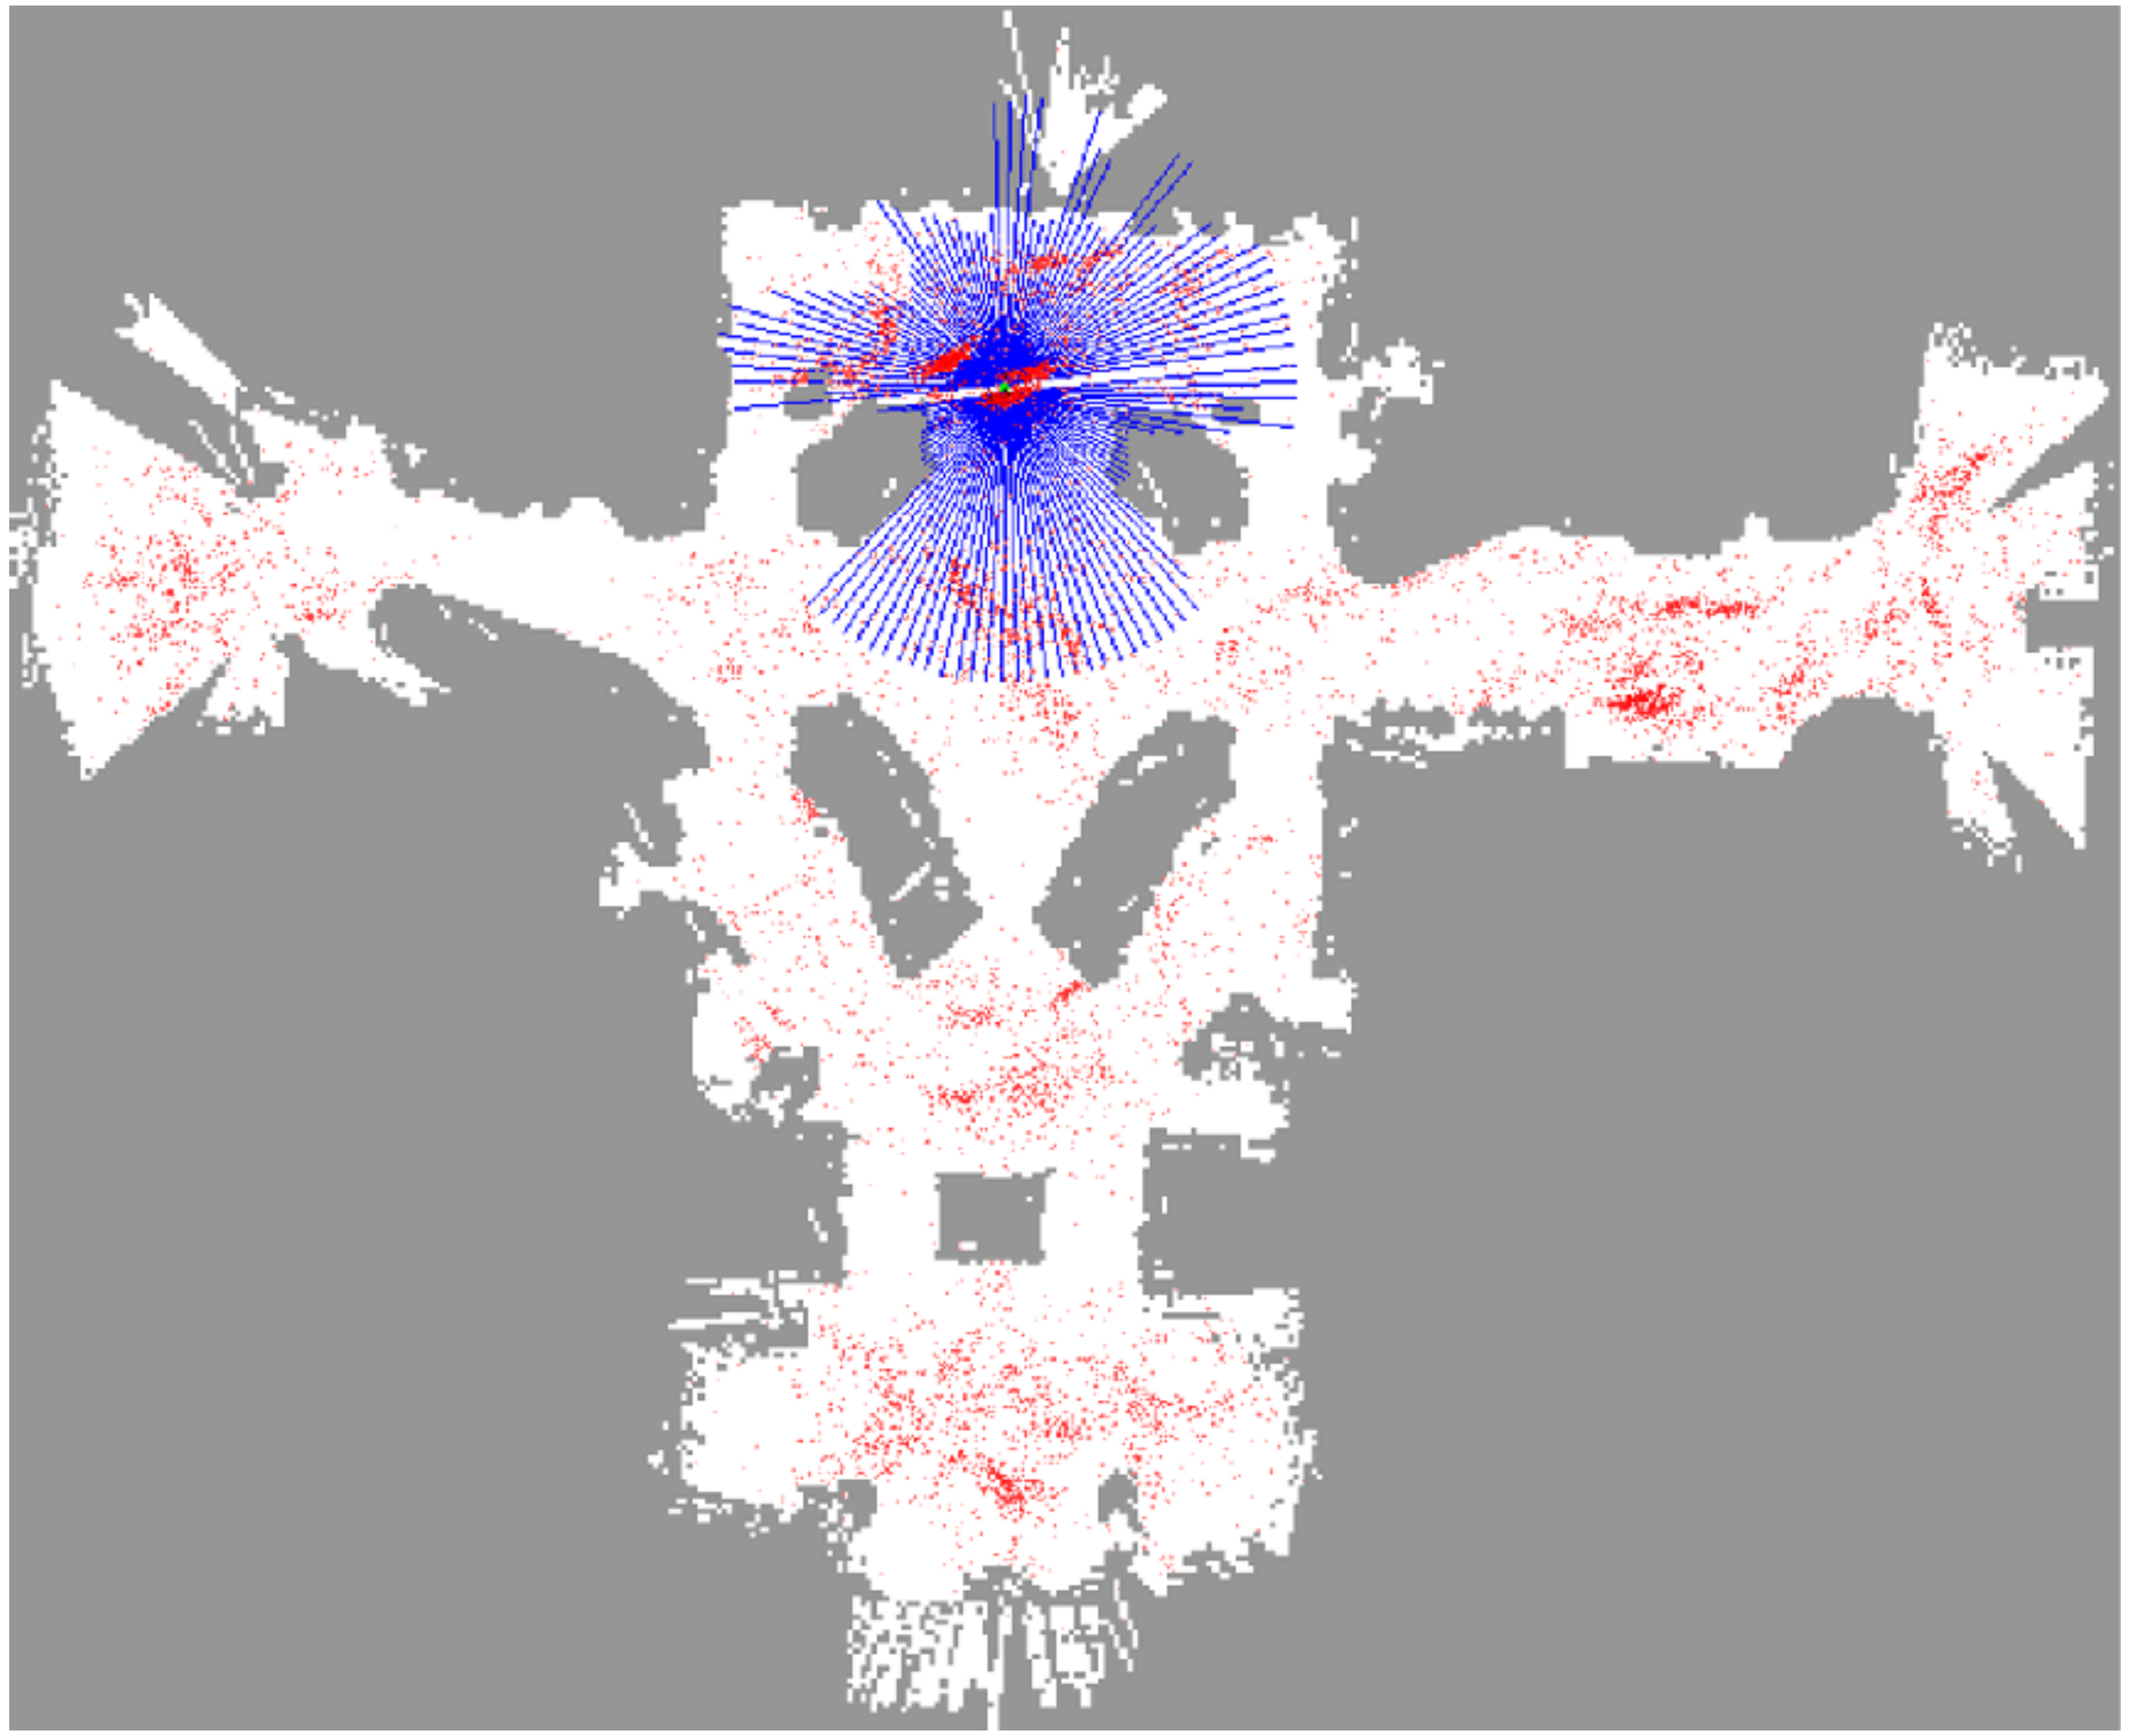

Particle Filter Applications

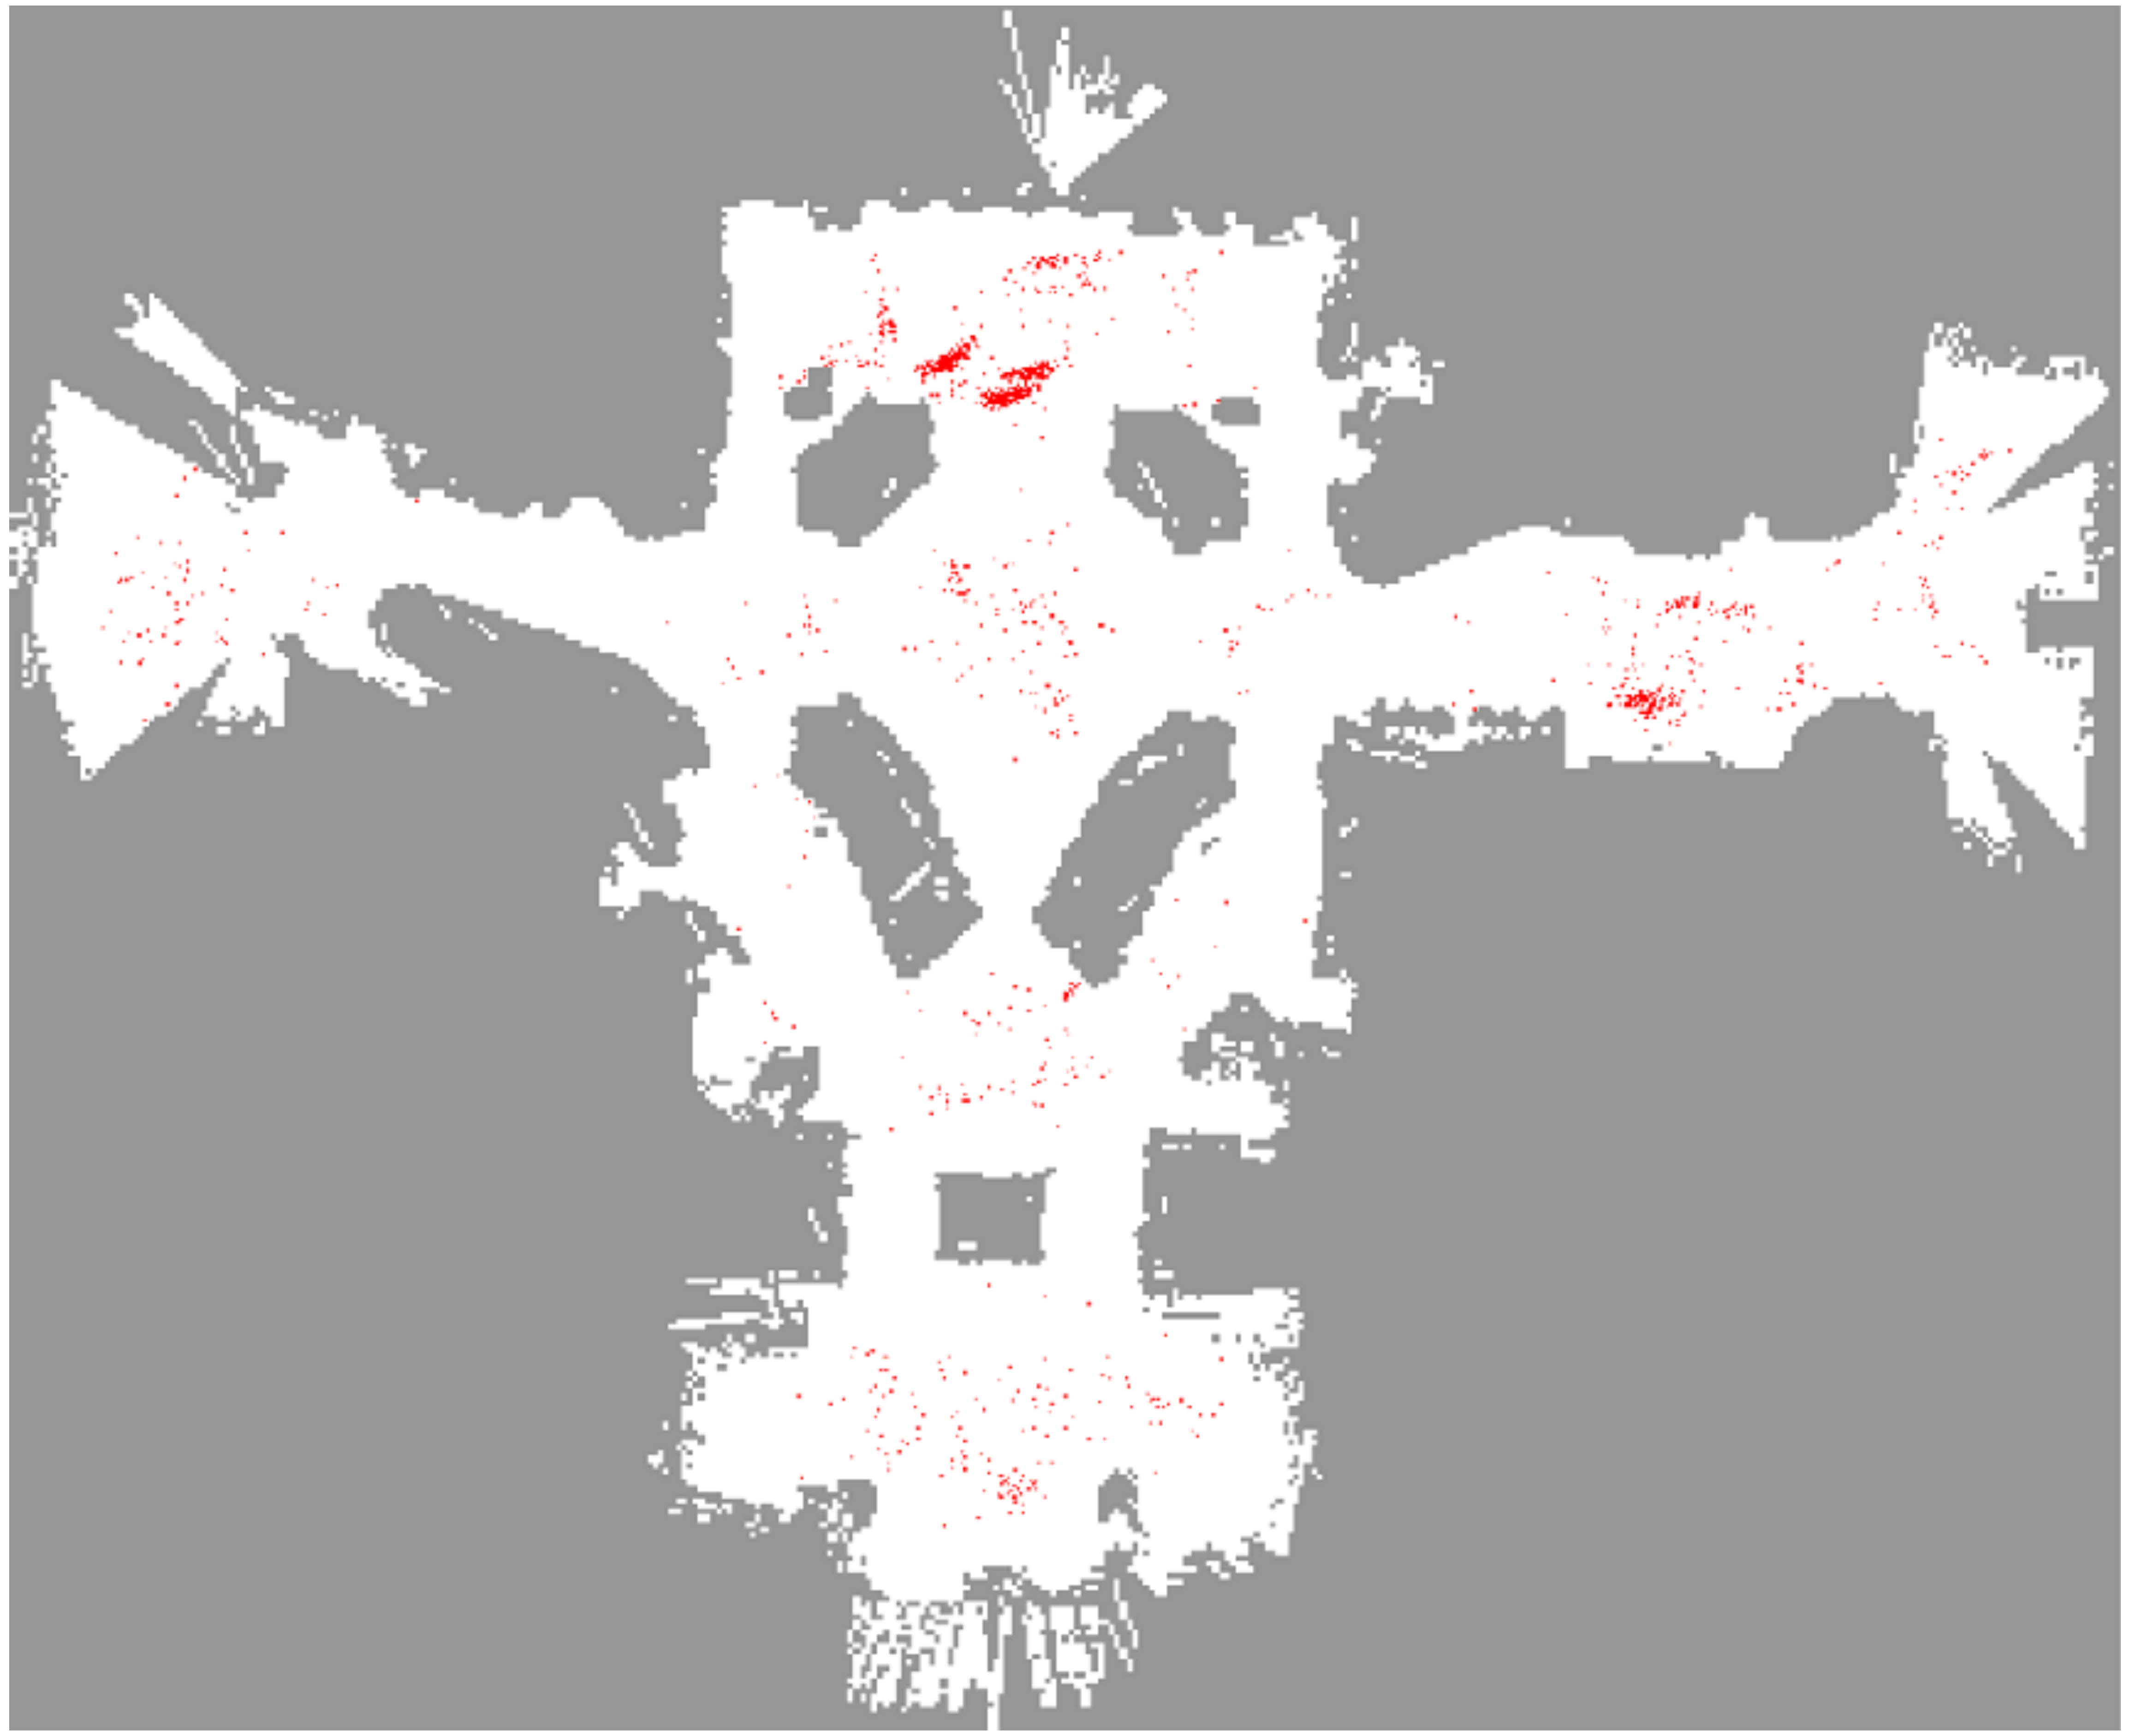

Particle Filter Applications

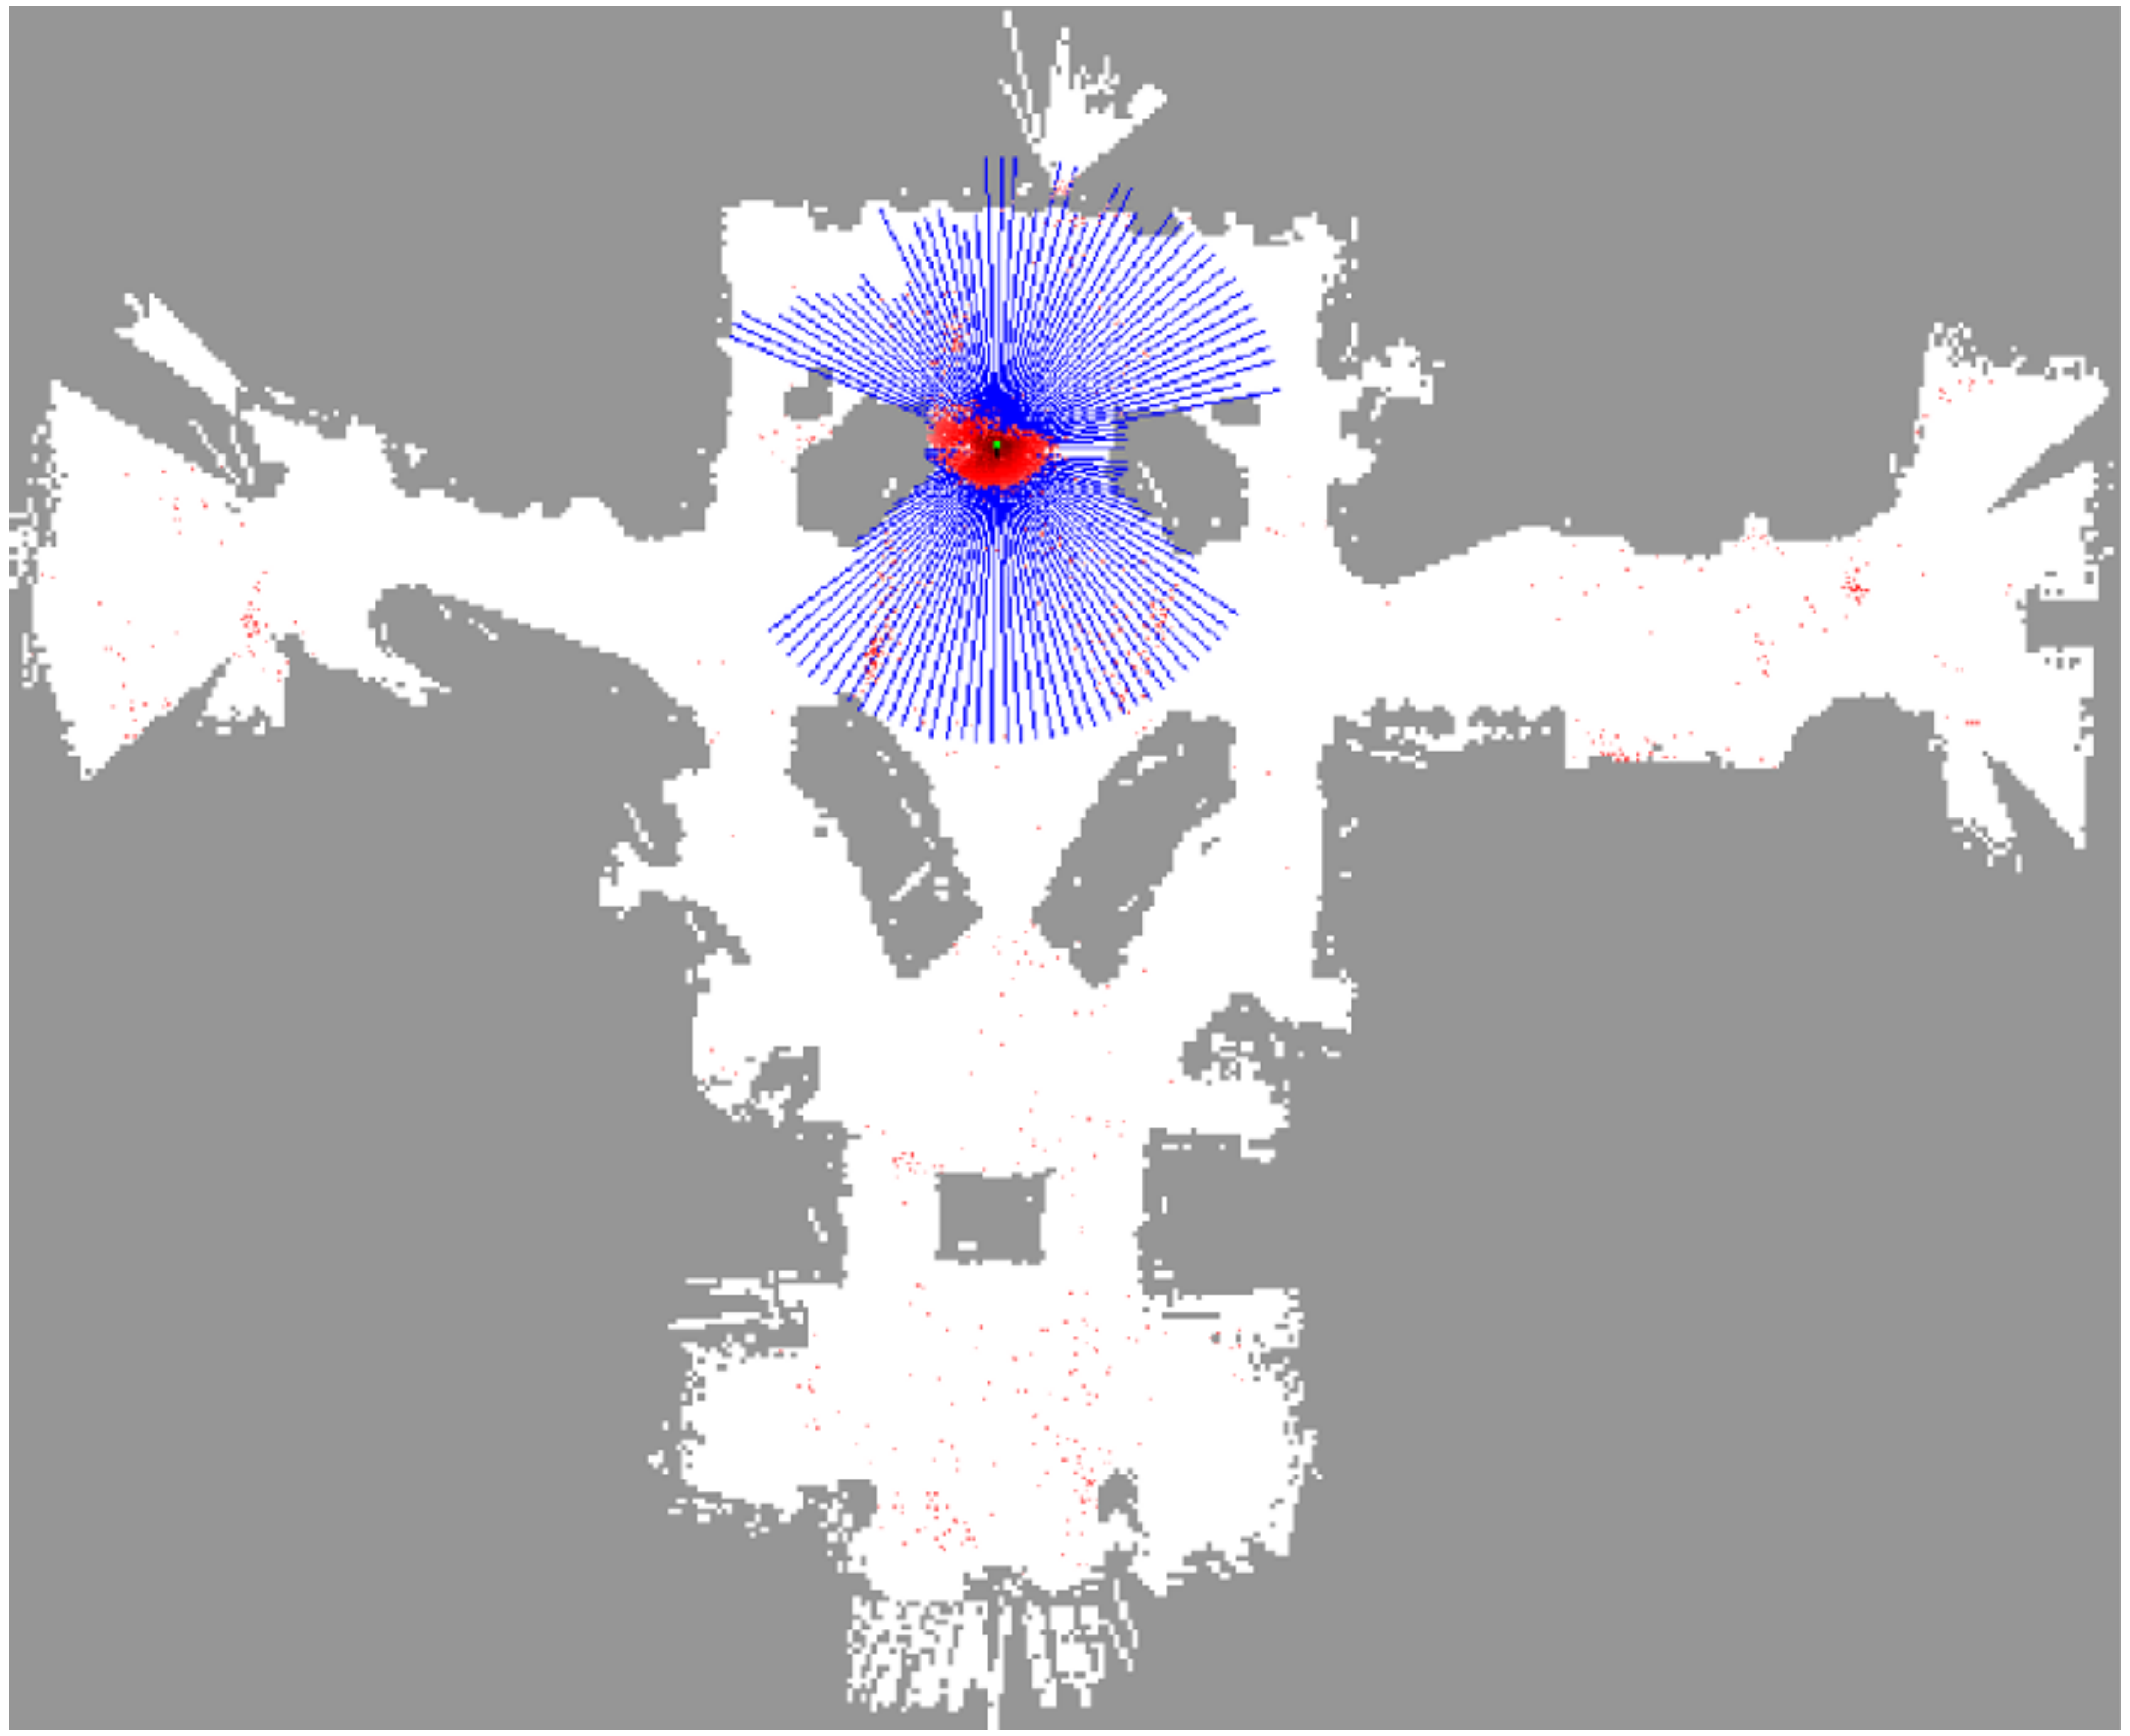

Particle Filter Applications

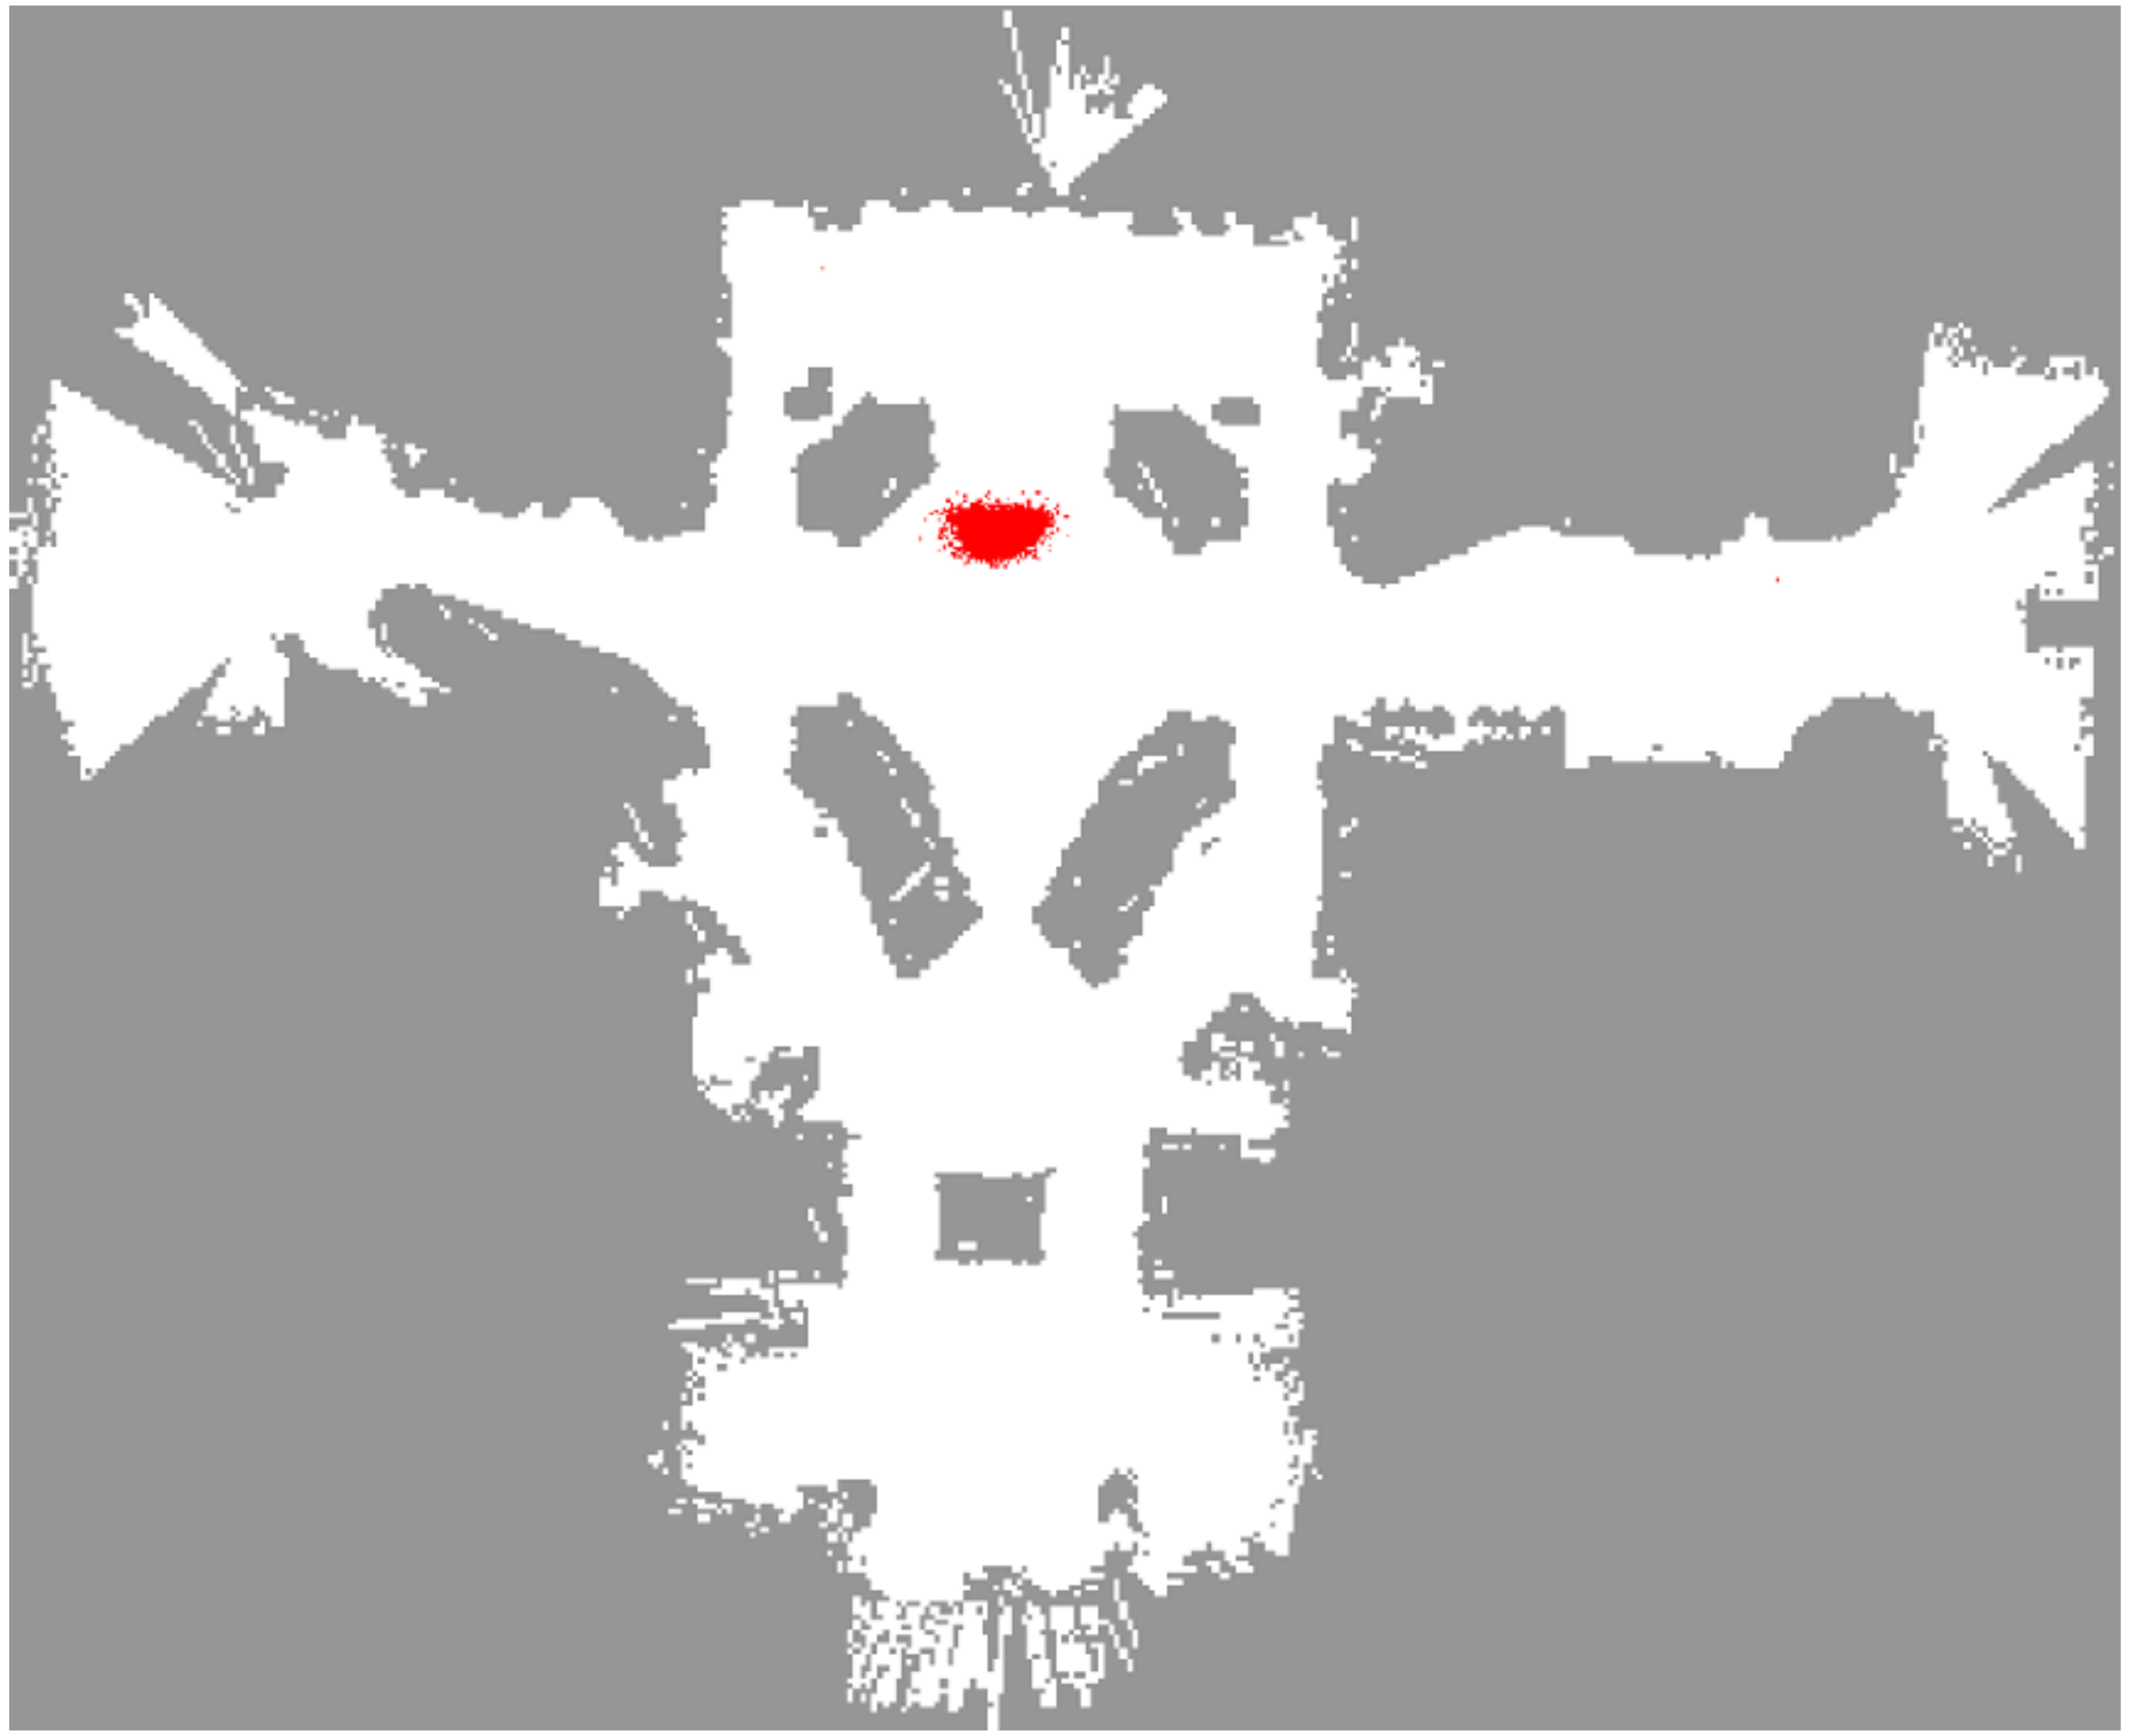

Particle-based Beliefs

- Difficult to exactly represent the robot's belief in continuous or large state spaces.
  - Bayes filter is usually intractable.
  - Kalman filter and EKF restrict to Gaussian beliefs.
- New idea: the robot represents its belief with a set of N particles.
  - Each particle represents a possible robot state.
  - Particles are explicitly or implicitly weighted based on the likelihood the robot is in a particular state.
  - Known as a **non-parametric** belief representation.

Illustration of Particle-Based Beliefs

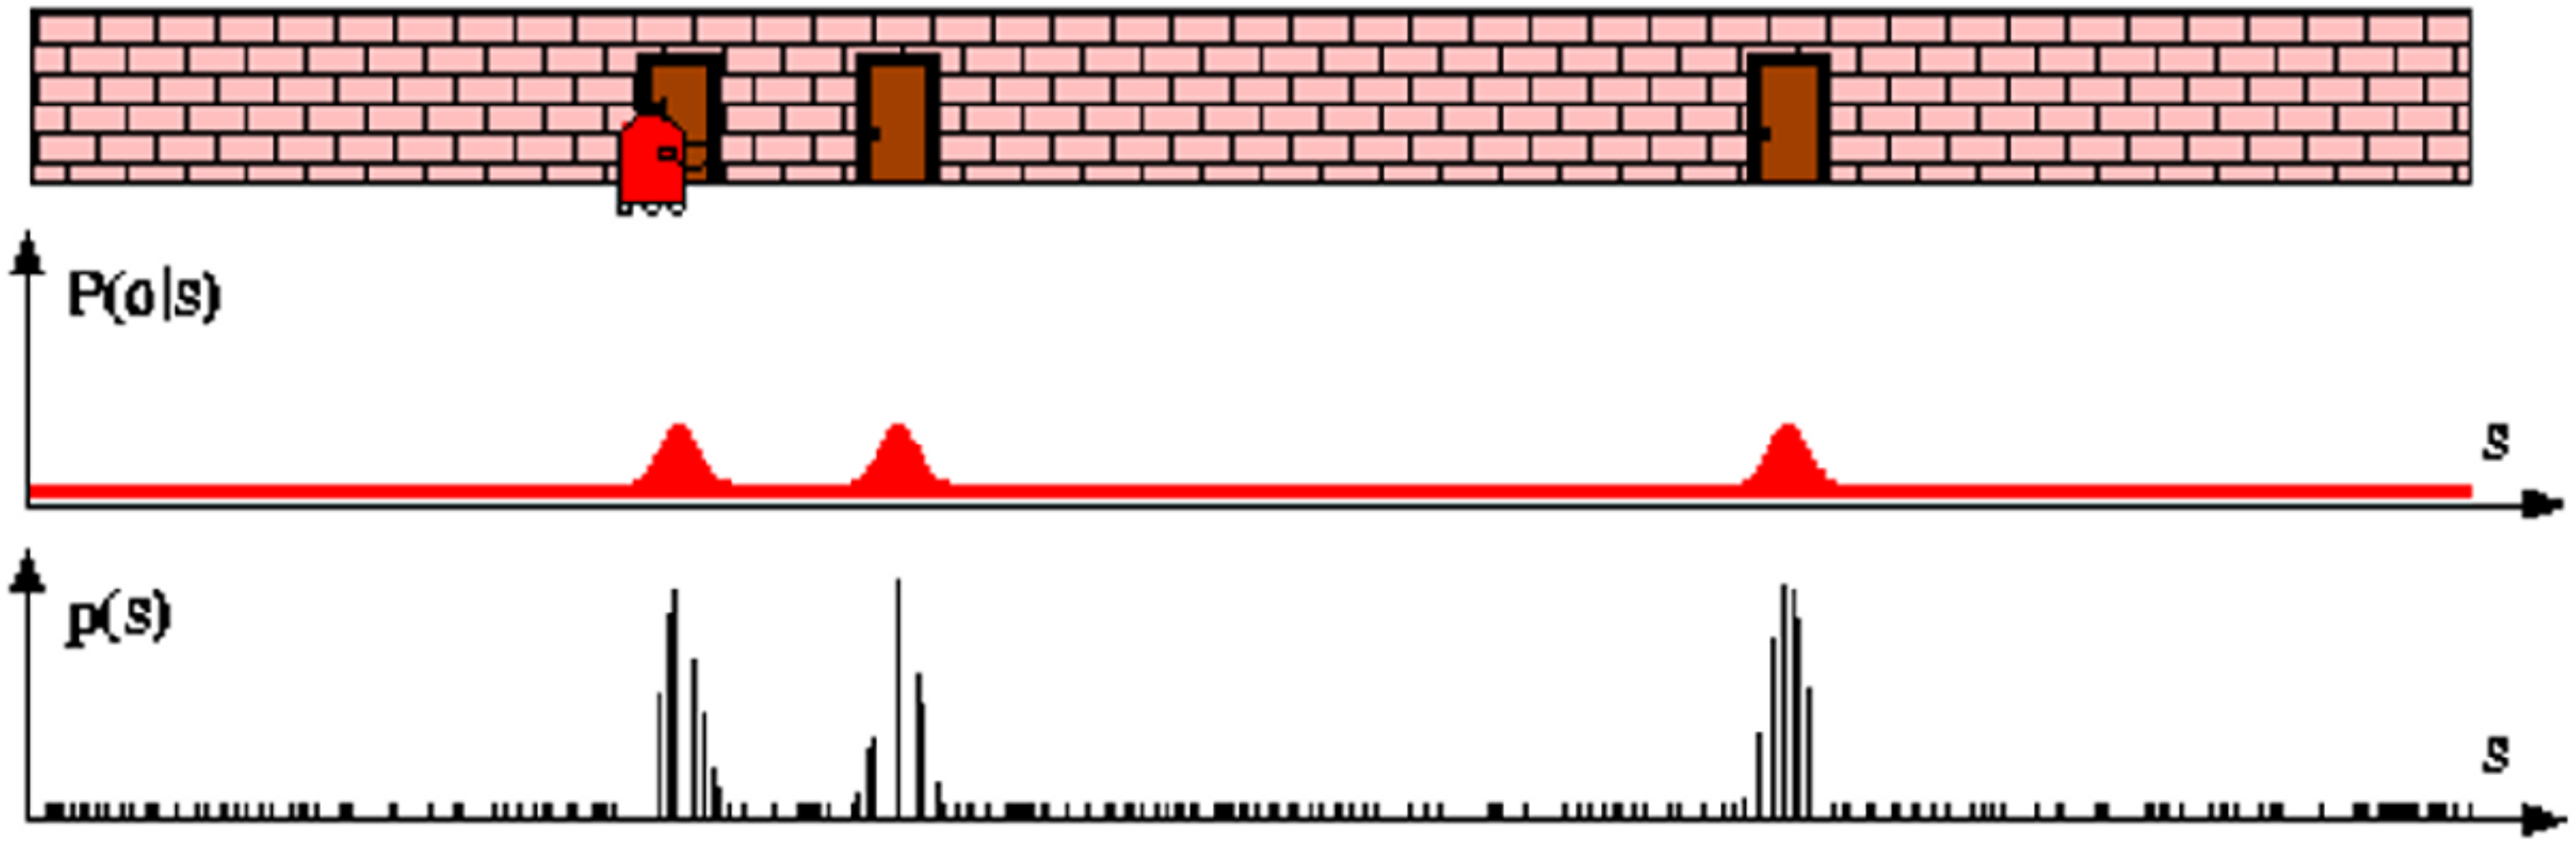

Importance Sampling

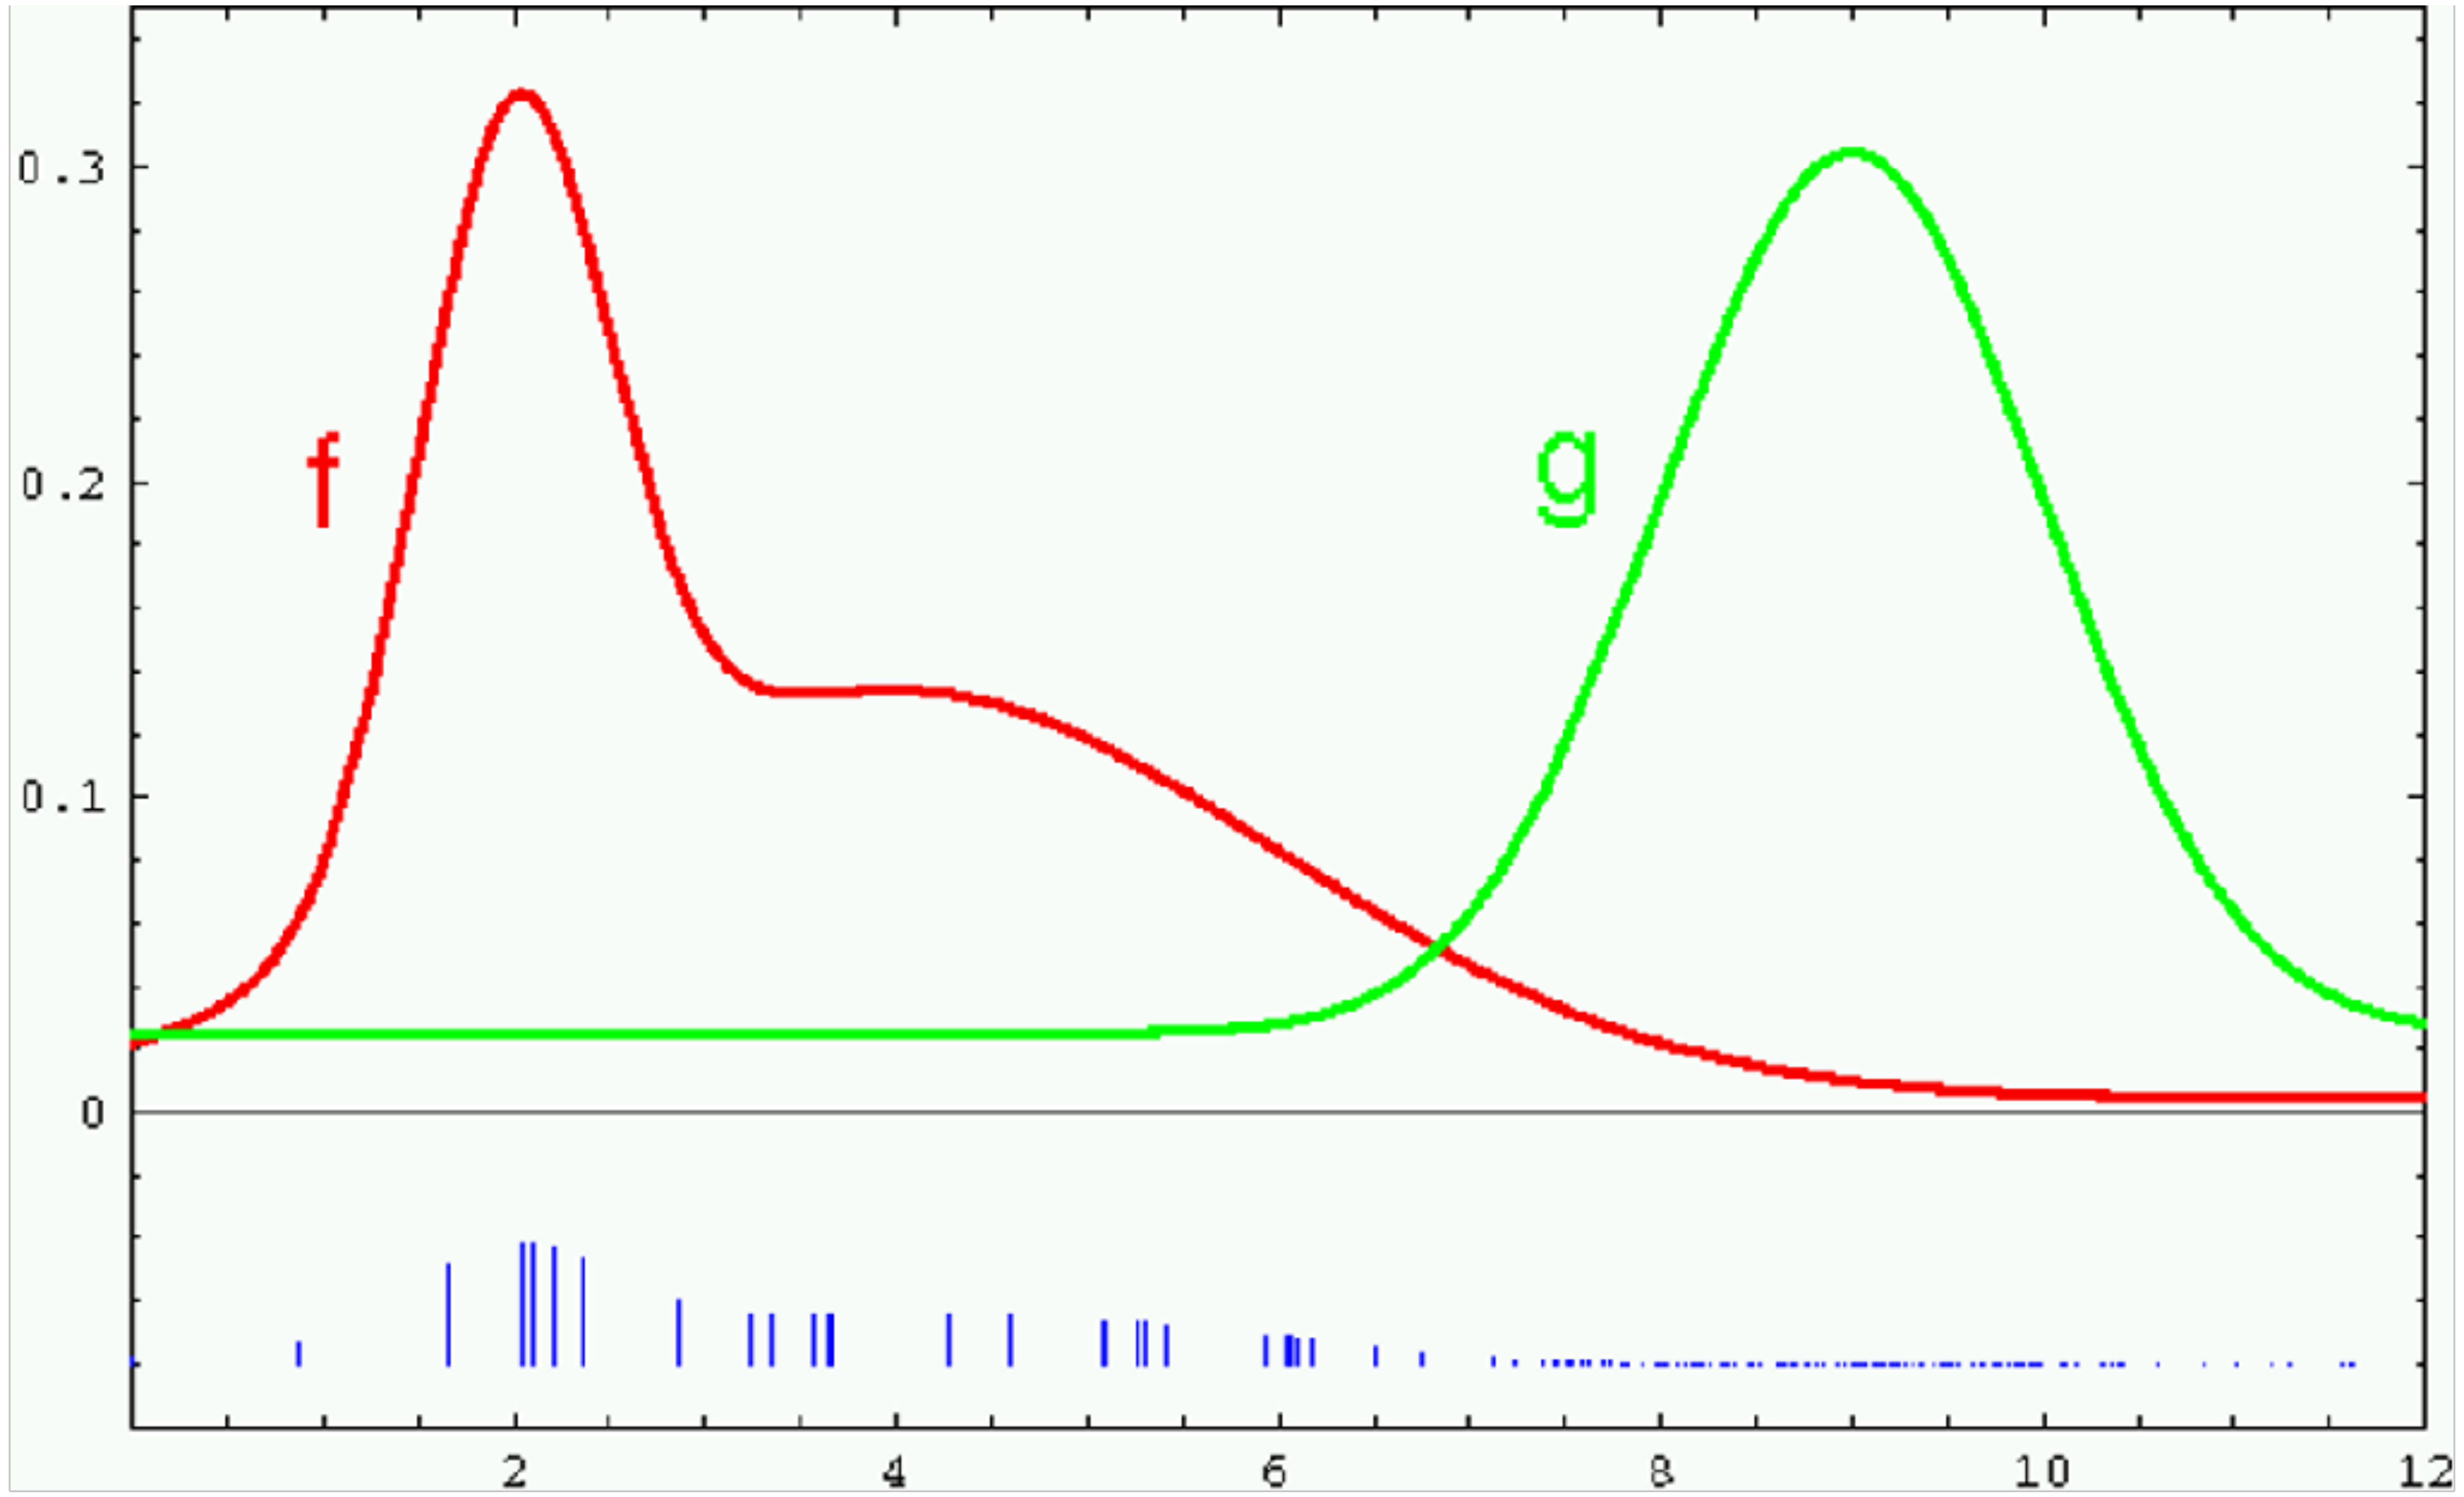

Weight samples: $w = f / g$

Computing the Mean State

- In state estimation, you often want summary statistics:
  - What state is the robot expected to be in?
  - How spread out is the robot's belief.
- Expected value of robot's state is a weighted average:

$$\mu_t = \sum_{i=1}^{N} w_i x_t^i \qquad \texttt{bel}(x_t) = \sum_{i=1}^{N} w_i \cdot \mathbf{1}\{x_t^i = x_t\}$$

Normalized Importance Sampling

- Belief is represented by a set of particles, $\{(x_i^t, w_i)\}$.
- Robot takes action u_t and then observes z_t.
- Update particles:
  - $x_{t+1}^i \sim p(\cdot \mid x_t^i, u_t)$
  - $w_i \leftarrow w_i * p(z_t \mid x_{t+1}^i)$
  - Normalize weights so that $\sum_{i=1}^{N} w_i = 1$.

Intuition: Each particle represents a path through the state space and weights represent the plausibility of the path.

Problem: Most paths become unlikely very quickly.

Effective sample size:

$$\frac{1}{\sum_{i=1}^{N} (w_i)^2}$$

Particle Filters

- Belief is represented by a set of particles, $\{(x_i^t, w_i)\}$.
- Robot takes action u_t and then observes z_t. Set $w_i \leftarrow \frac{1}{N}$.
- Update particles:
  - $x_{t+1}^i \sim p(\cdot \mid x_t^i, u_t)$
  - $w_i \leftarrow w_i * p(z_t \mid x_{t+1}^i)$
  - Normalize weights so that $\sum_{i=1}^{N} w_i = 1$.
  - Sample N new particles (with replacement) to form a new particle set.

Particle Filters

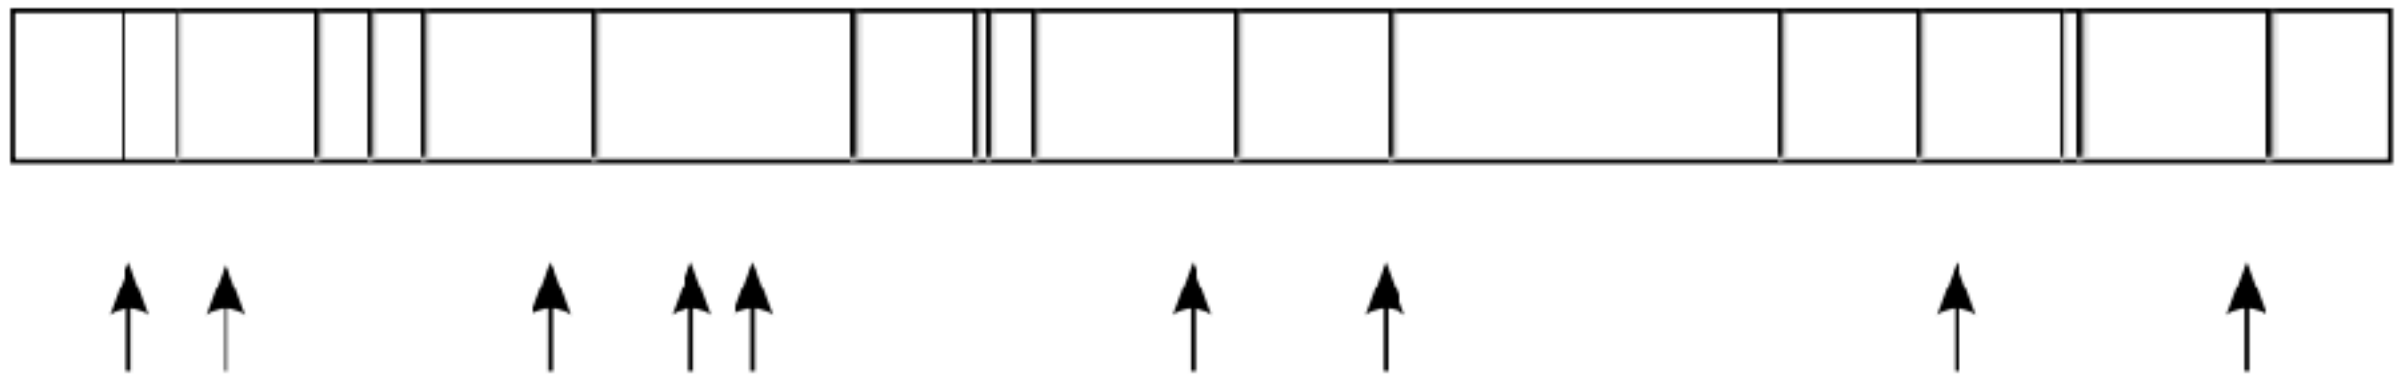

Figure 1: Resampling

- Survival of the fittest particles.

- Randomly select particles according to their weights.

- Reset weights to $\frac{1}{N}$ after each iteration. Why?

$$\mu_t = \sum_{i=1}^{N} w_i x_t^i$$

Particle Filter Illustration

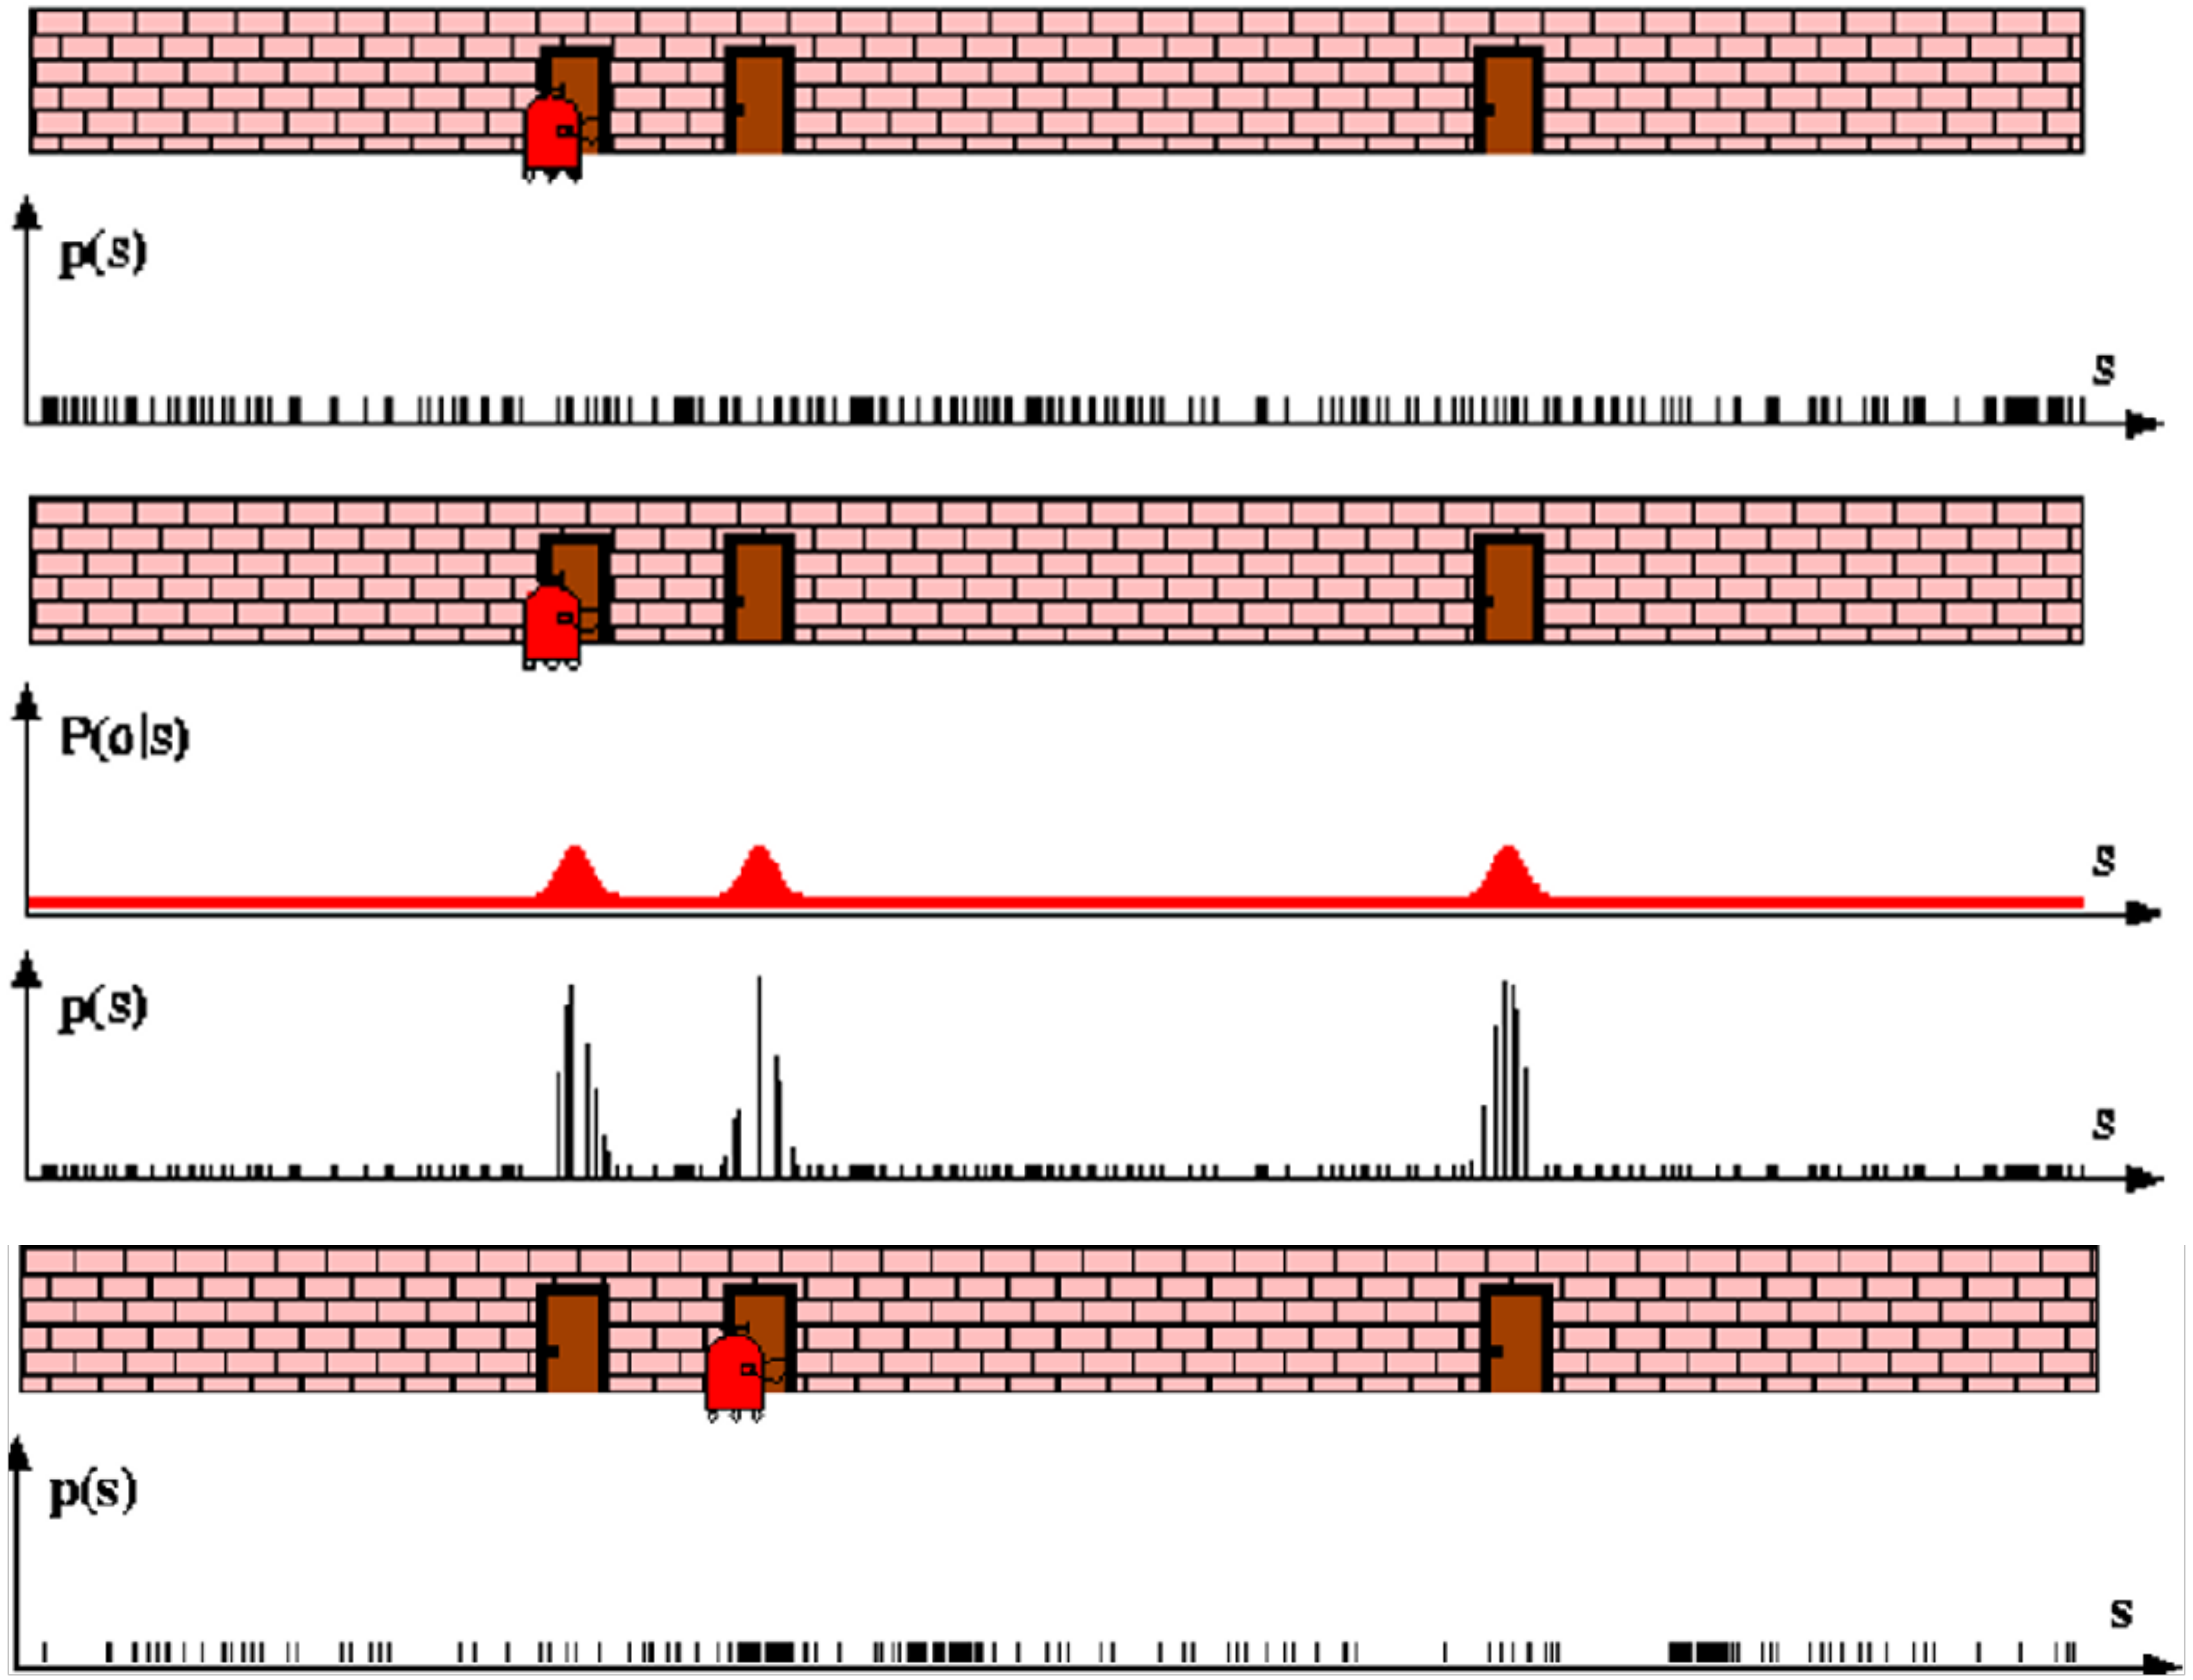

Particle Filter Illustration

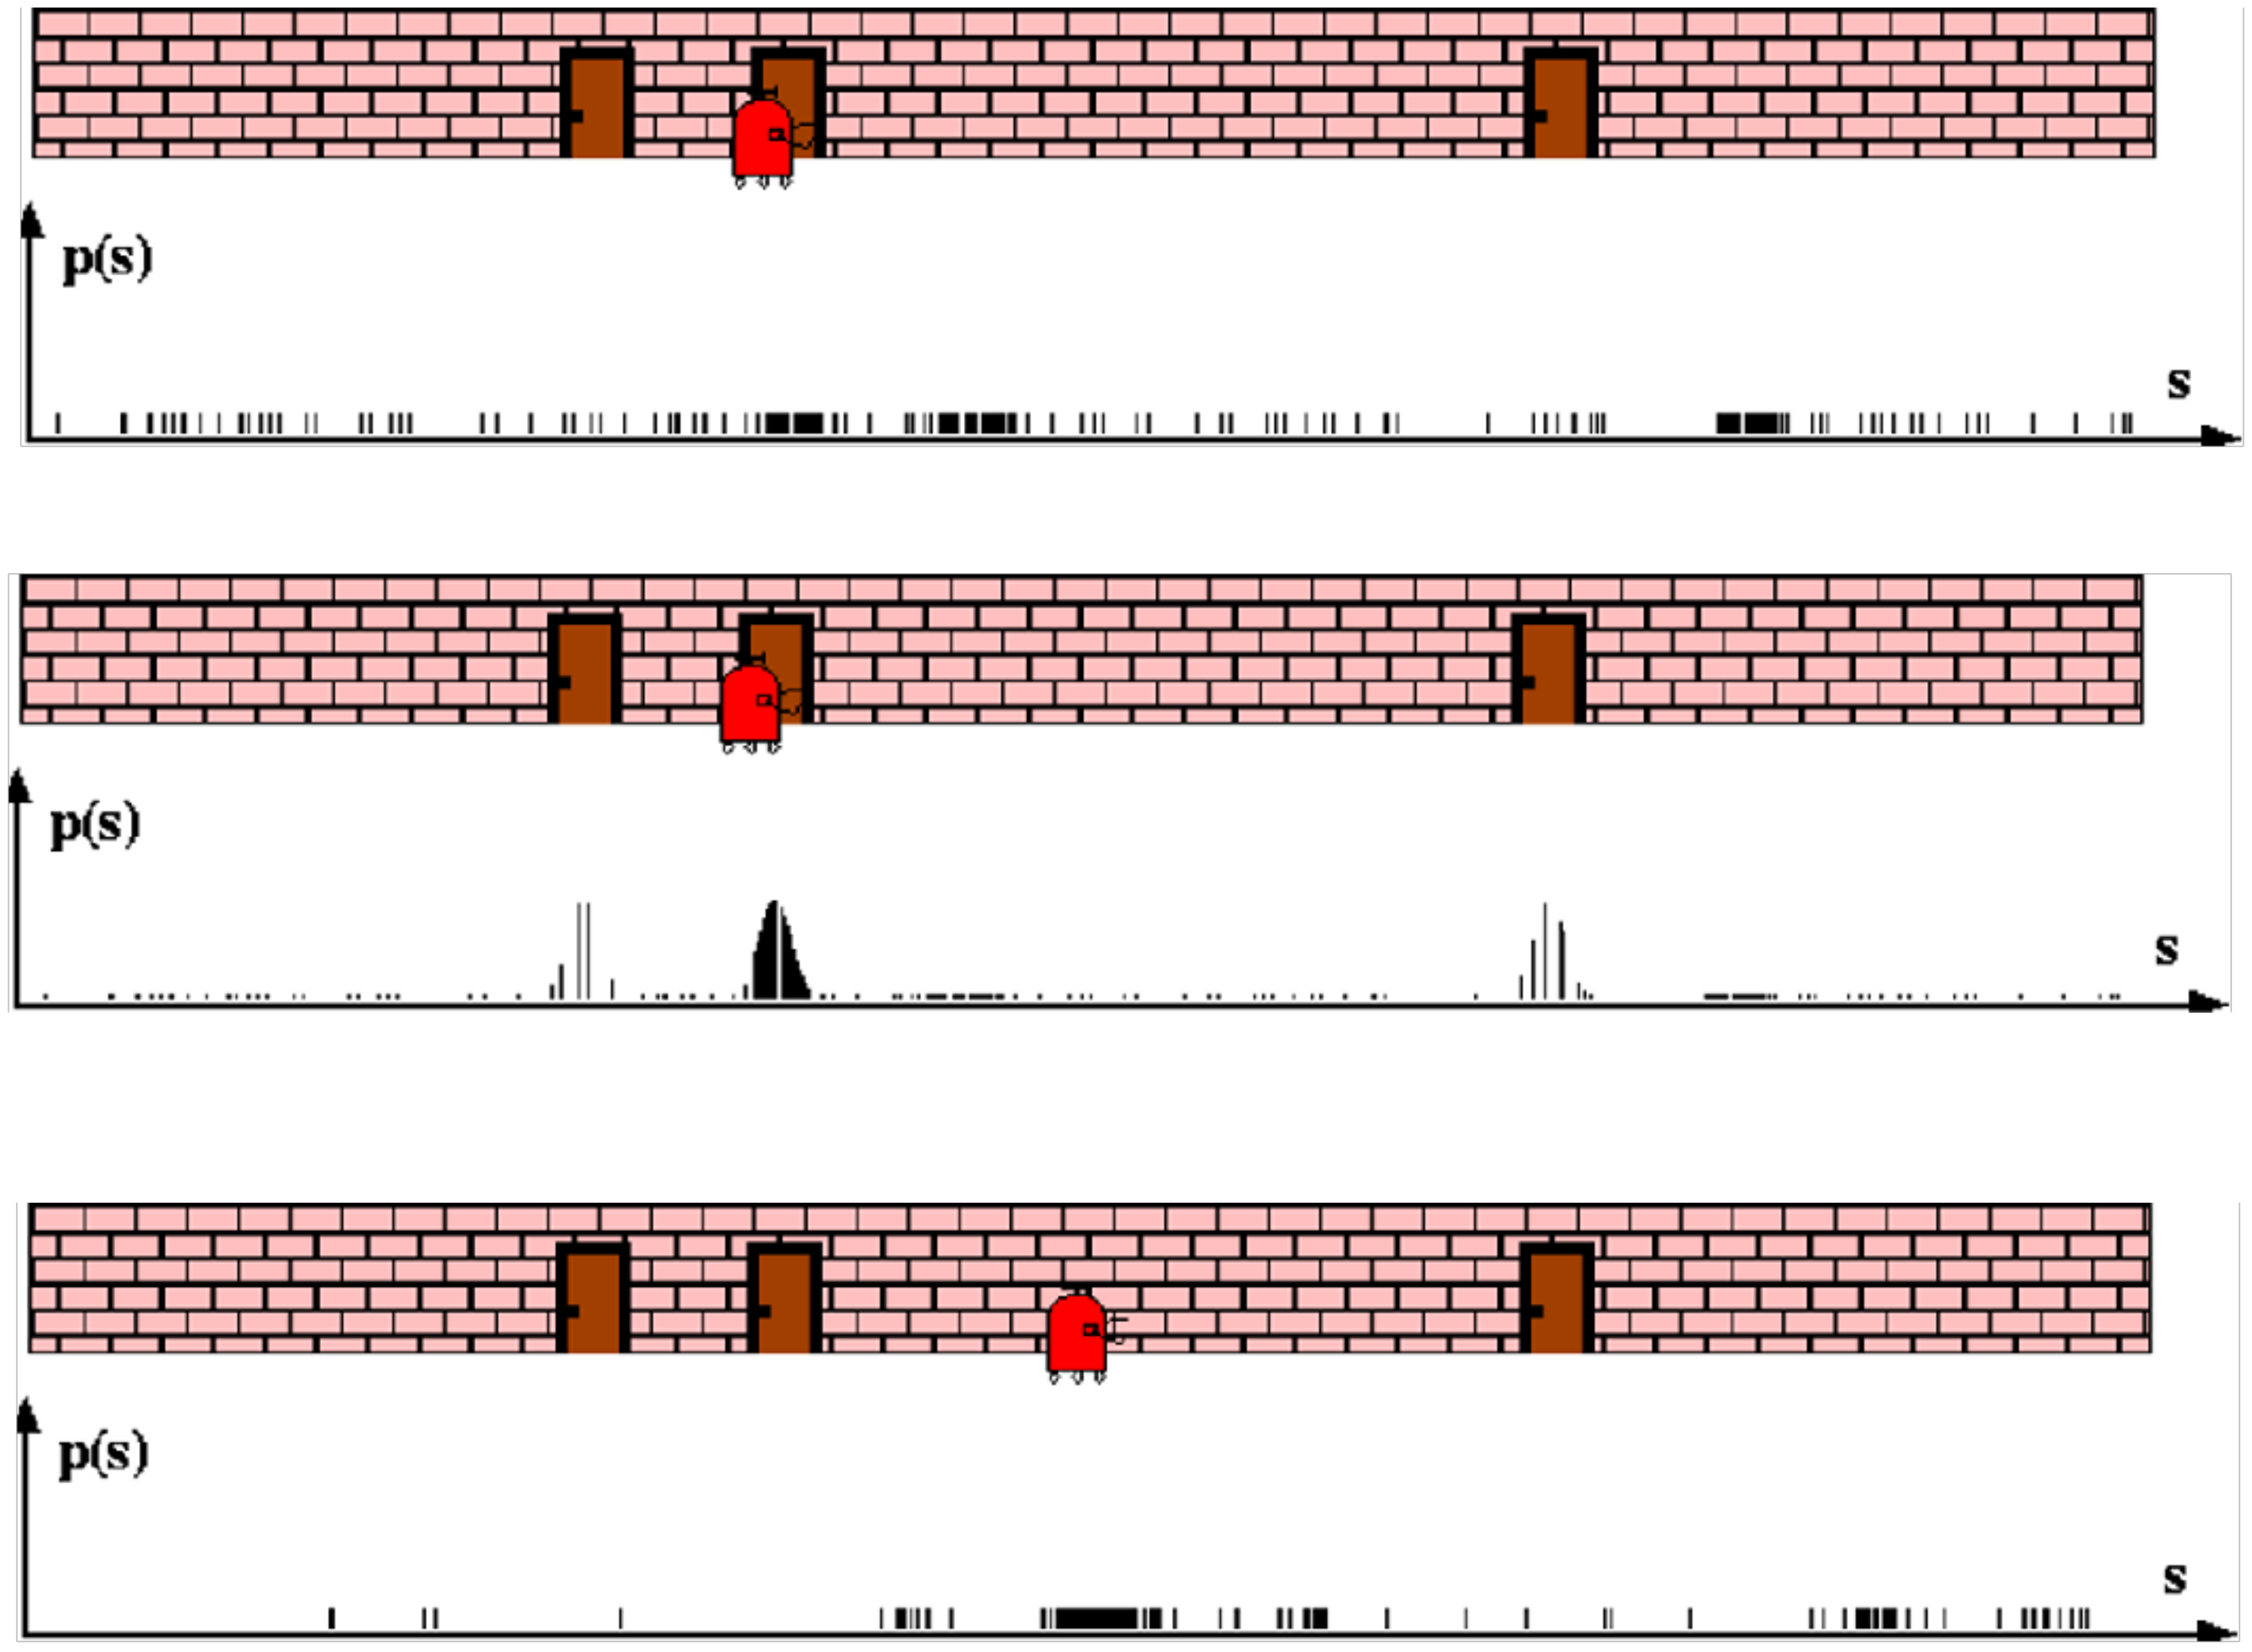

Low Variance Resampler

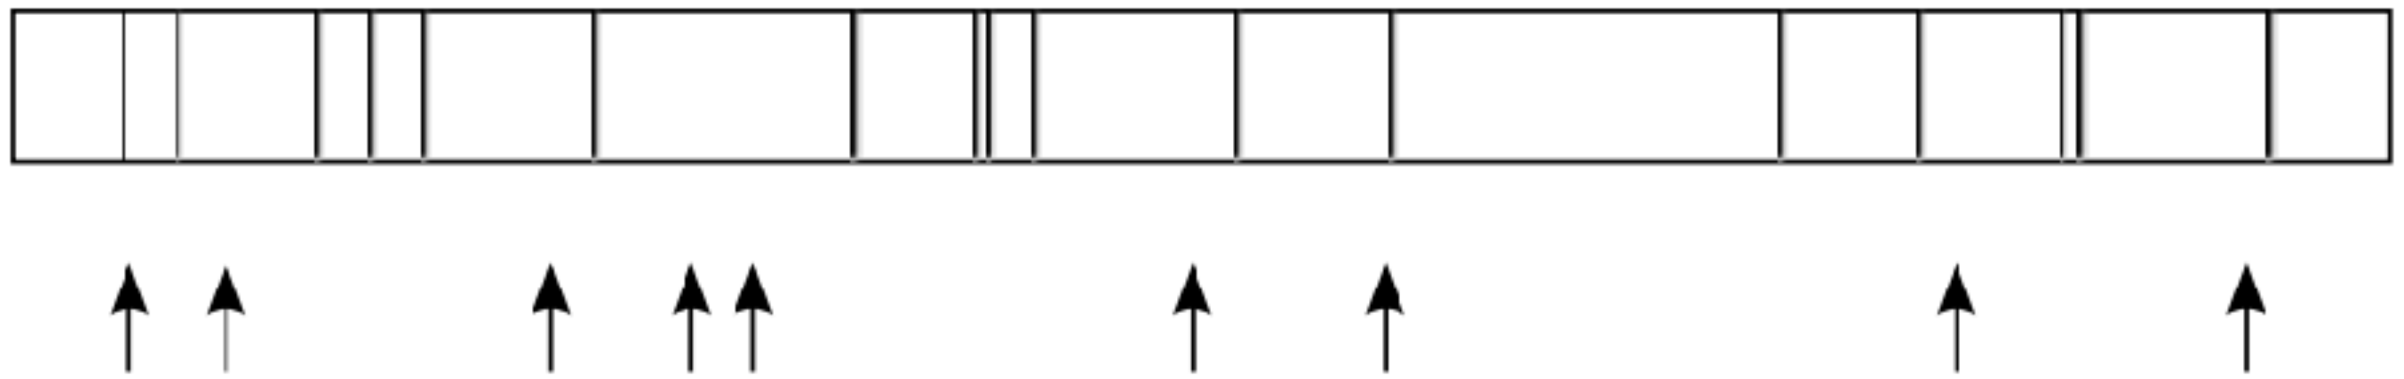

Figure 1: Resampling

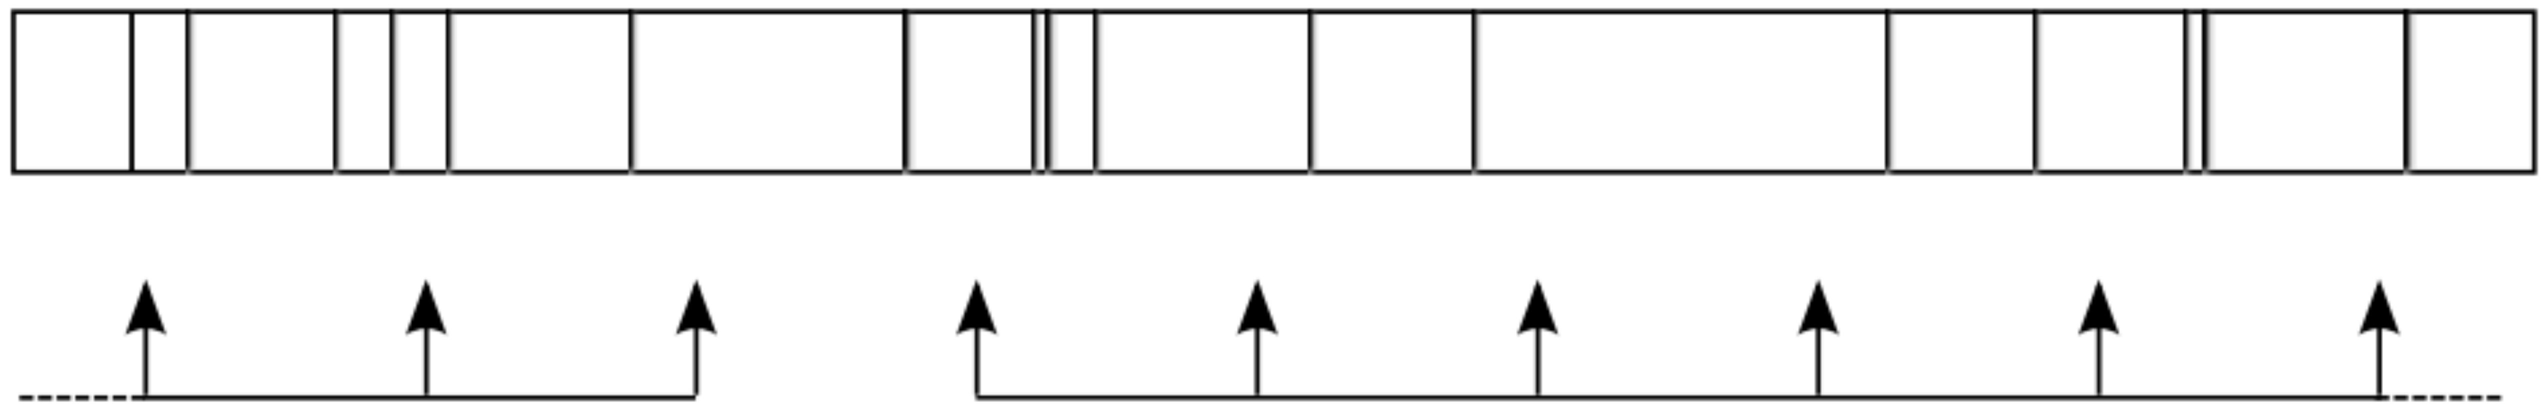

Advantages / Disadvantages

- Particle filters:
  - Can be used for continuous state spaces.
  - Can approximate any belief distribution (compare to Kalman/EKF).
  - Scale with computation
- But…
  - Only approximate belief.
  - Limited for high-dimensional state-spaces.

Loss of Diversity

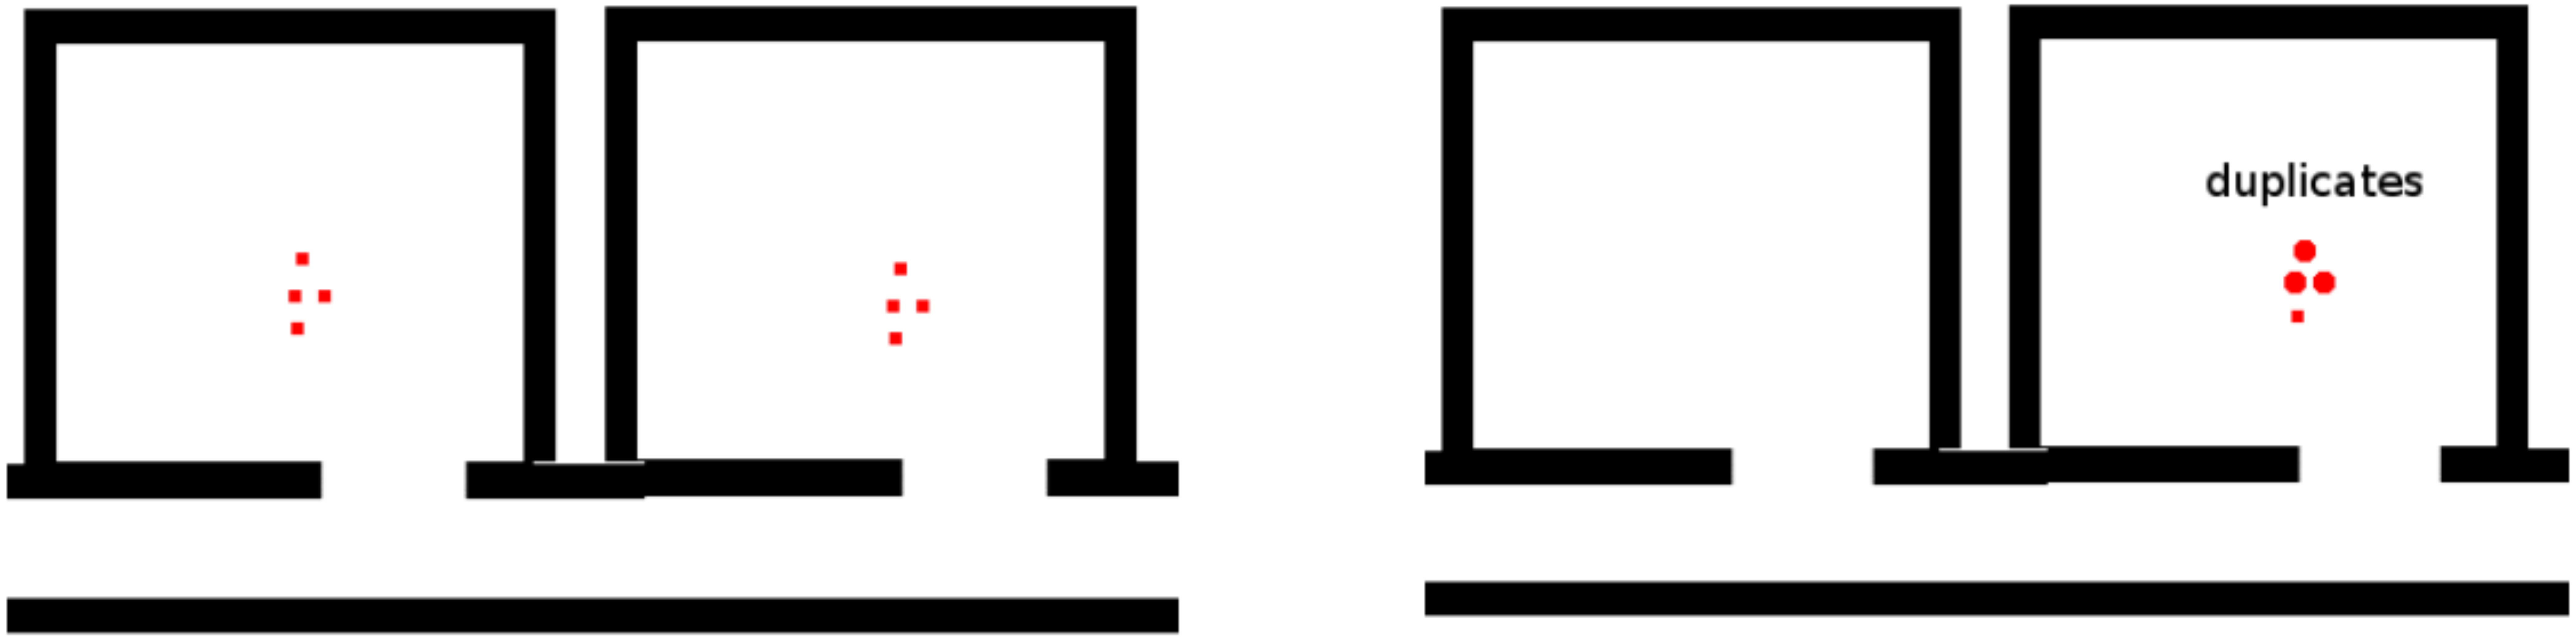

Good observation models are bad?

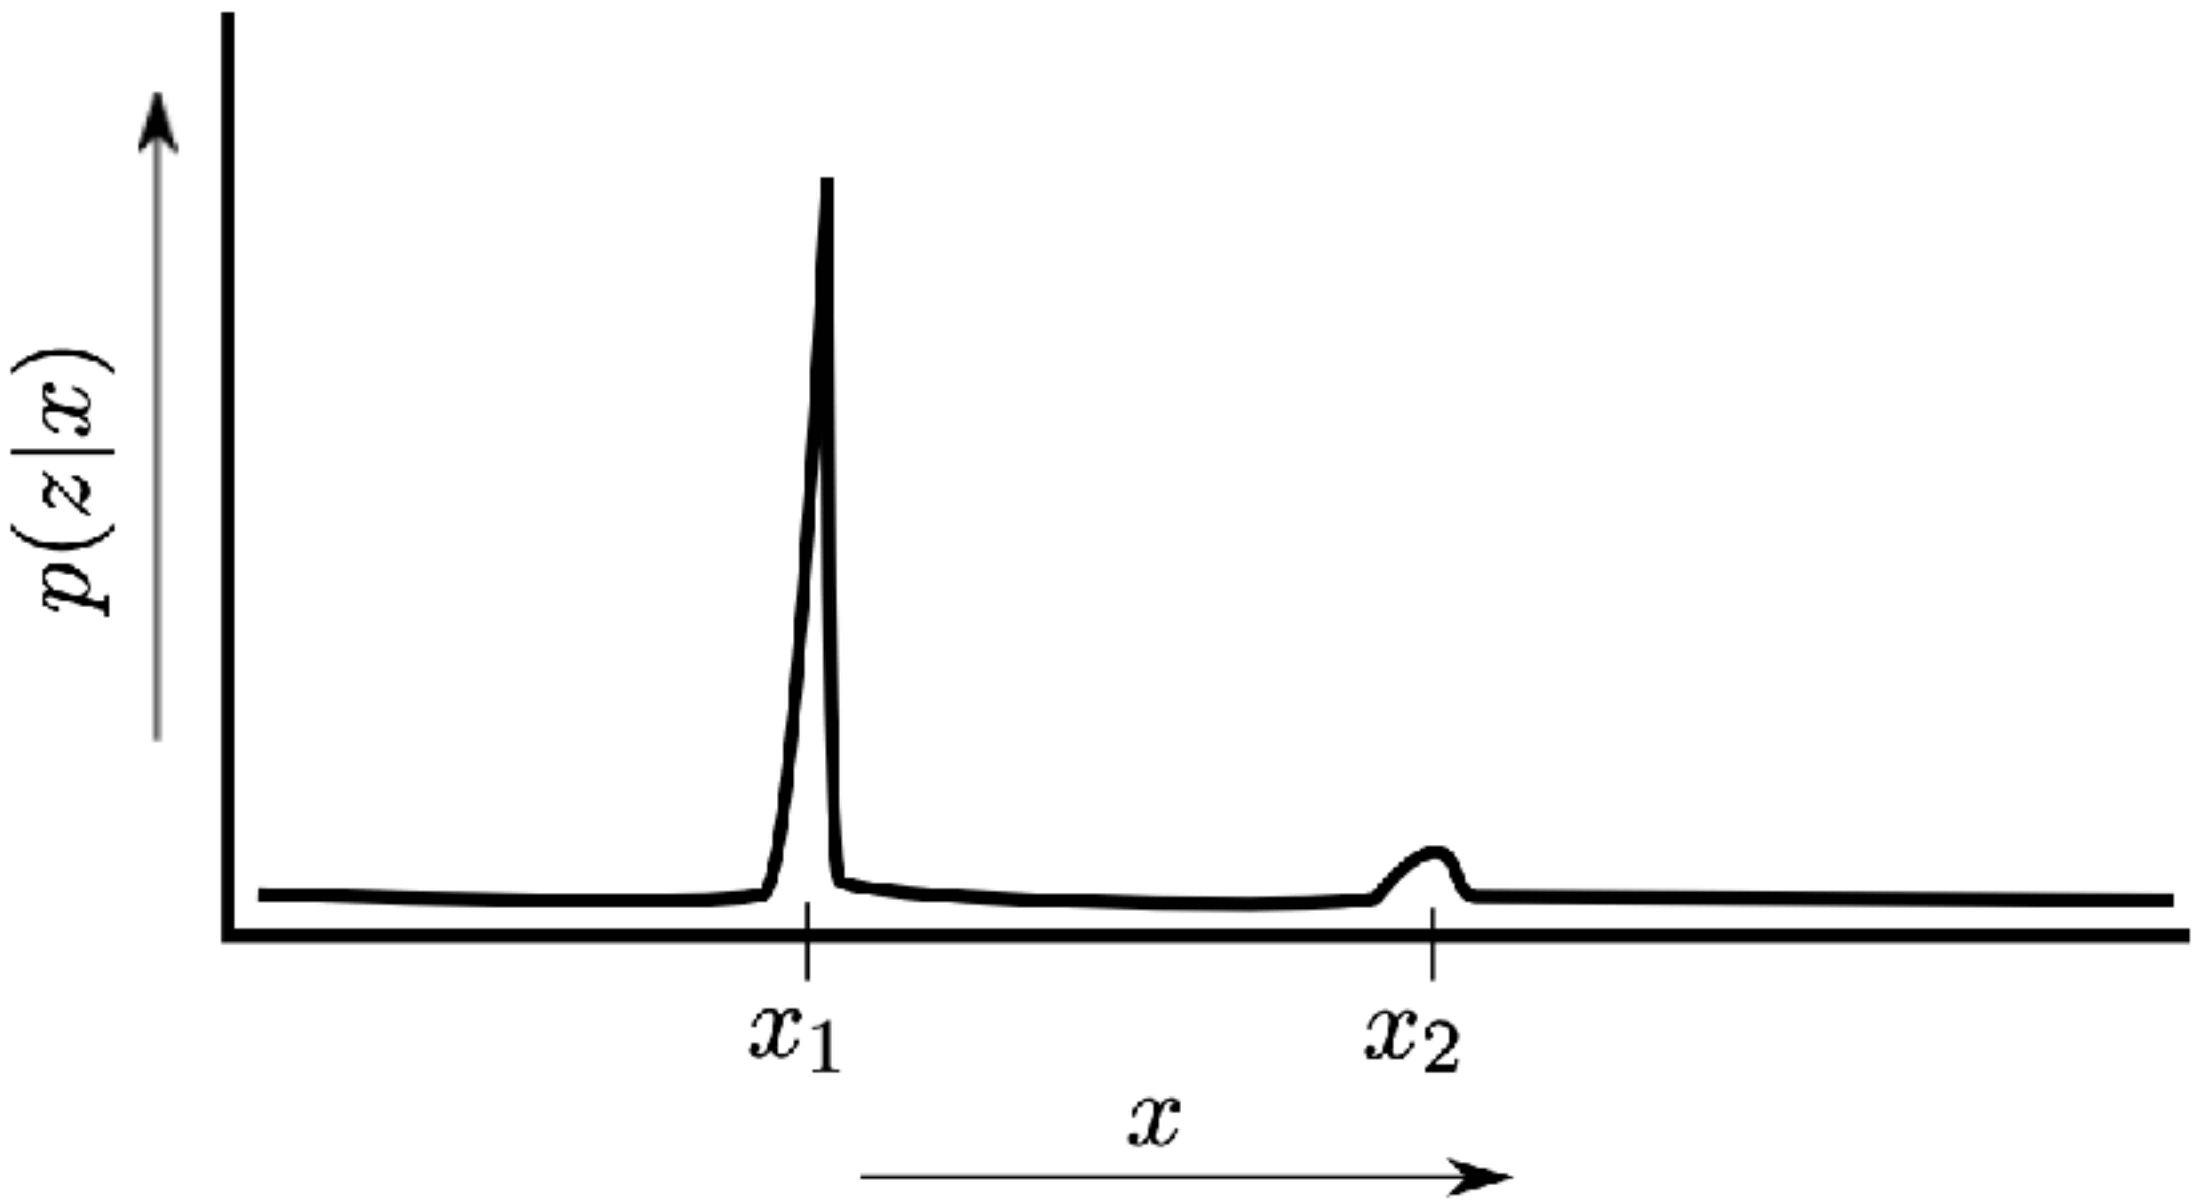

Summary

- Saw examples of particle-based belief representations.
- Discussed differences between NIS and particle filters.
- Discussed pitfalls and remedies for particle filters.

Action Items

- Work on programming assignment #2.
- Read on SLAM for next week; send a reading response by 12 pm on Monday.