Localization 3: Histogram filters and particle filters

Particle filtering

Posterior distributions: Particle filtering

Key idea: Represent the posterior distribution as a collection of **samples** called **particles**, each with a state x and a weight ω.

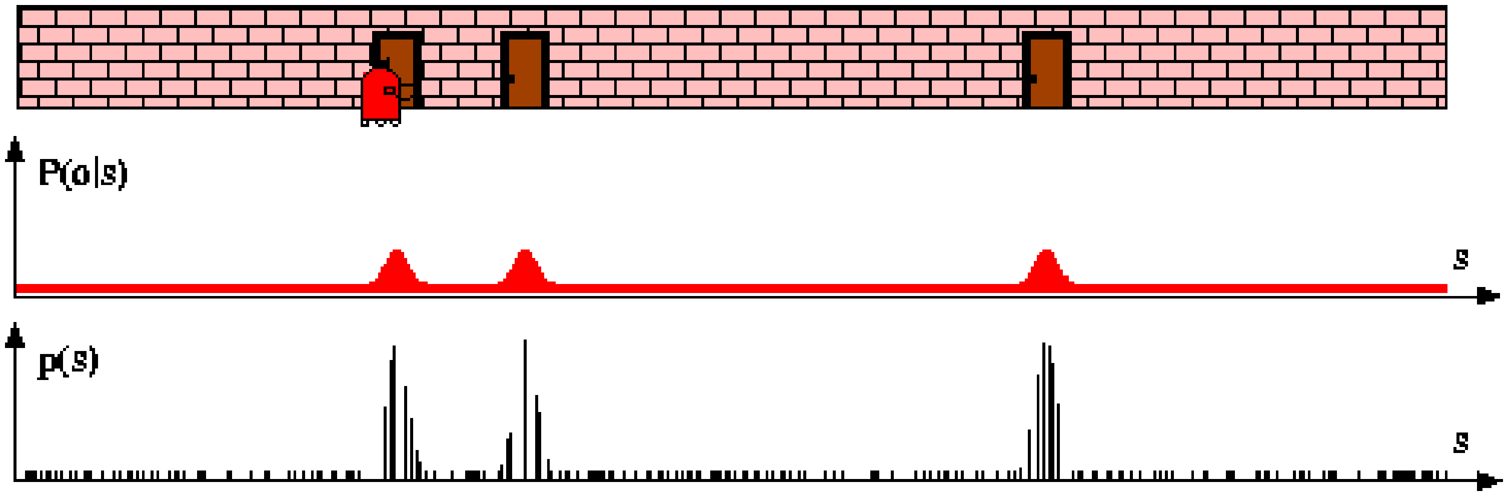

Particle filter: Initialization

Start with a set χ of M particles based on the robot's initial knowledge.

Each particle is a single state.

Particle filter: Actions

When an action u_k is executed, update χ via **forward projection**.

For each particle $\left(x_{k-1}^{(m)}, \omega_{k-1}^{(m)}\right)$ in χ_{k-1}, choose a state at random according to the transition model:

$$P(x_k | x_{k-1}^{(m)}, u_k).$$

Particle filter: Observations

When an observation y_k is received, assign **weights** $\omega_k^{(m)}$ to the particles.

One simple weighting approach: The weight of a particle at $x_k^{(m)}$ is proportional to the observation model:

$$P(y_k \mid x_k^{(m)})$$

Key idea: Particles more consistent with the observation get higher weights.

Particle filter: Resampling

Occasionally, we **resample** the particles to concentrate them in areas more likely to be the correct state.

Repeat M times: Randomly choose (with replacement) a particle, using probabilities proportional to the weight of each particle.

Example

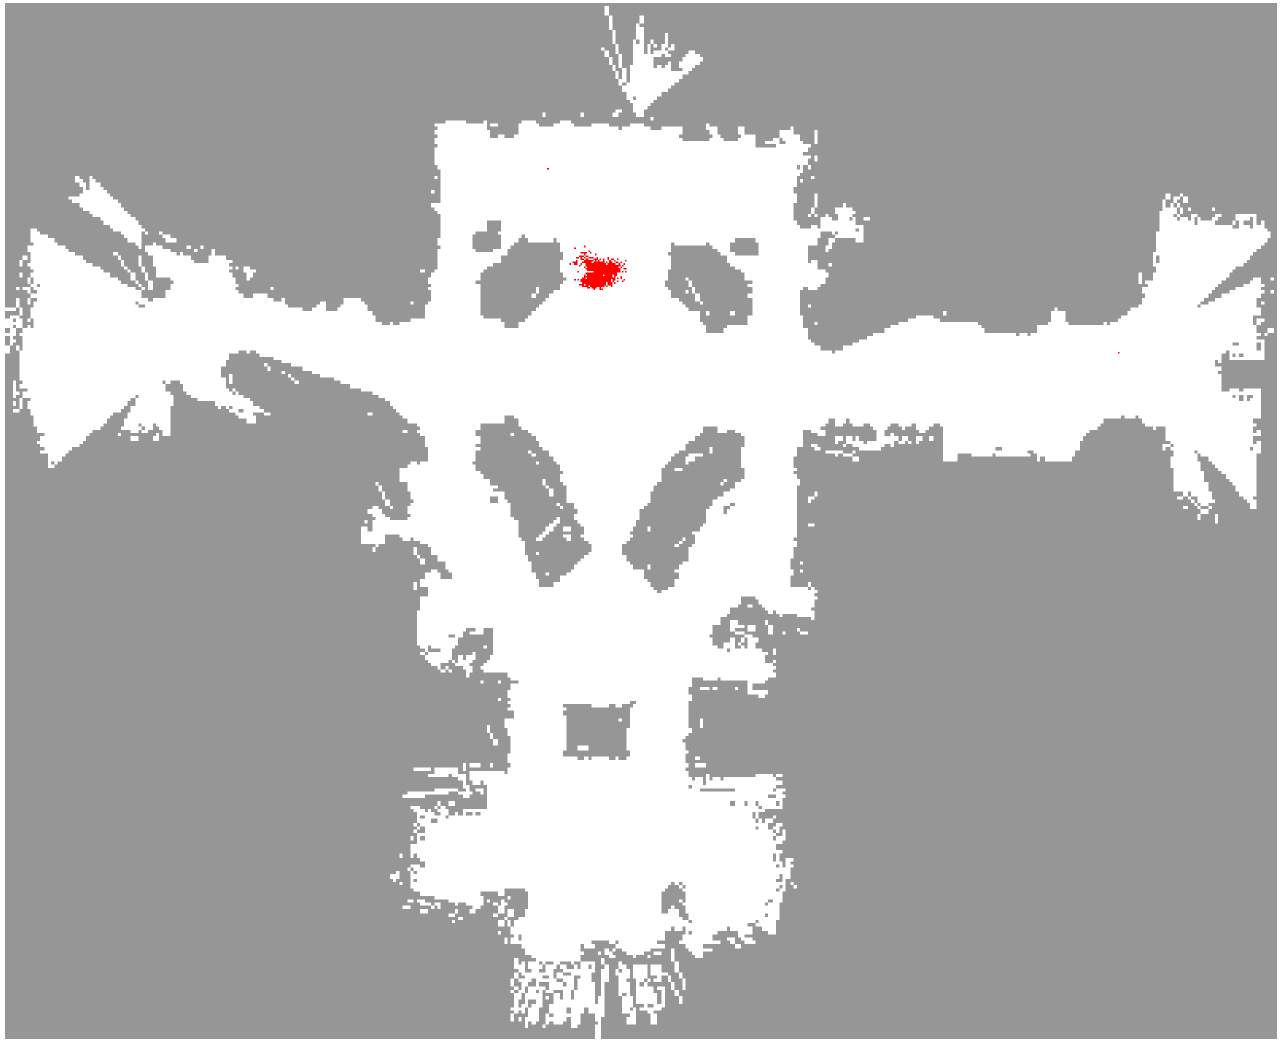

[Wolfram Burgard, TUN]

Example

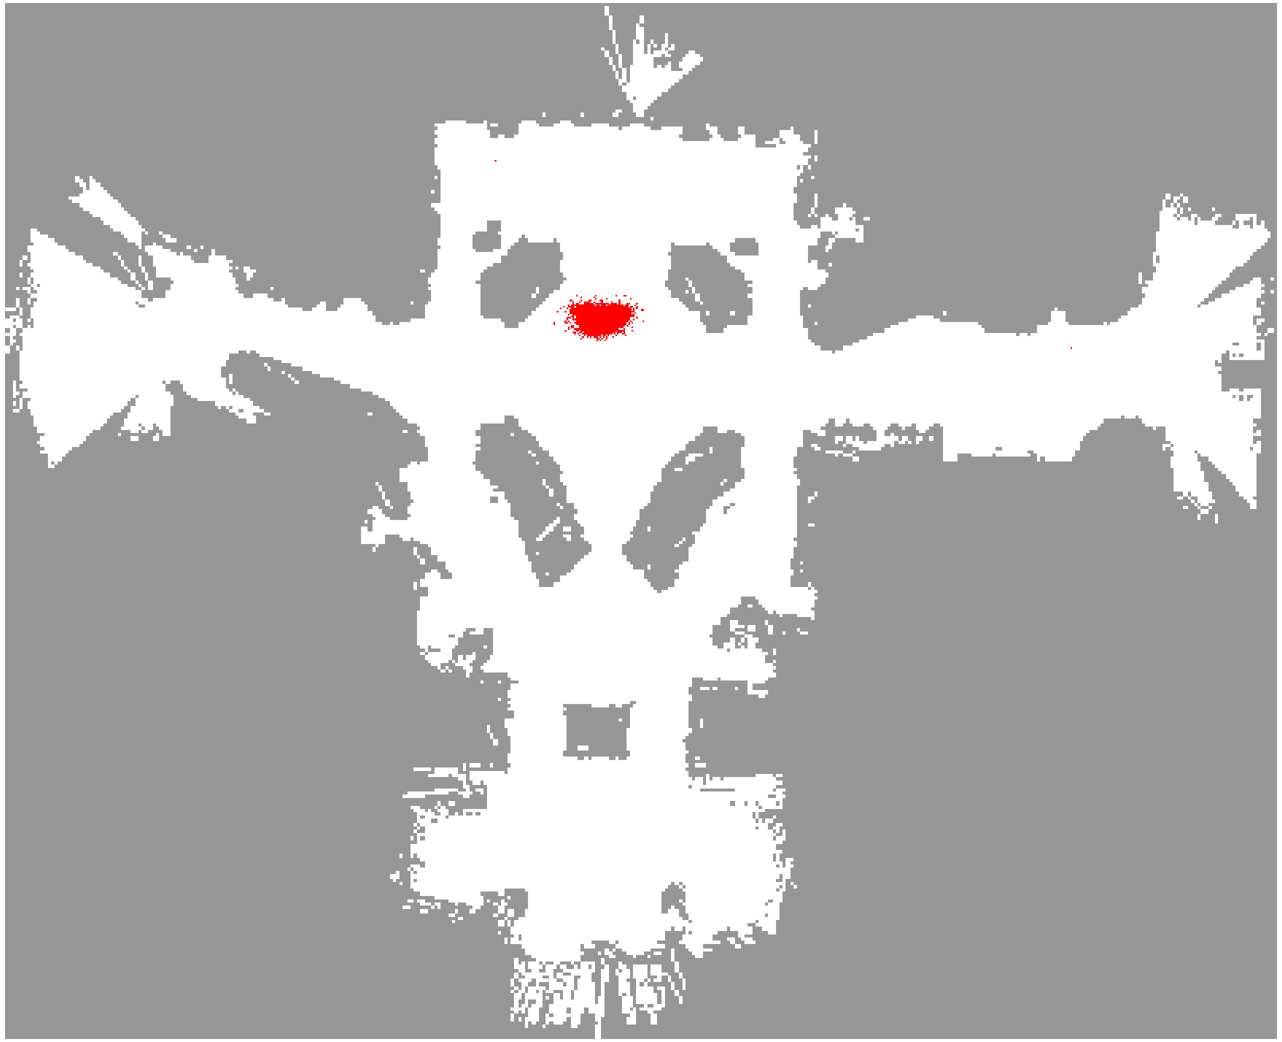

[Wolfram Burgard, TUN]

Example

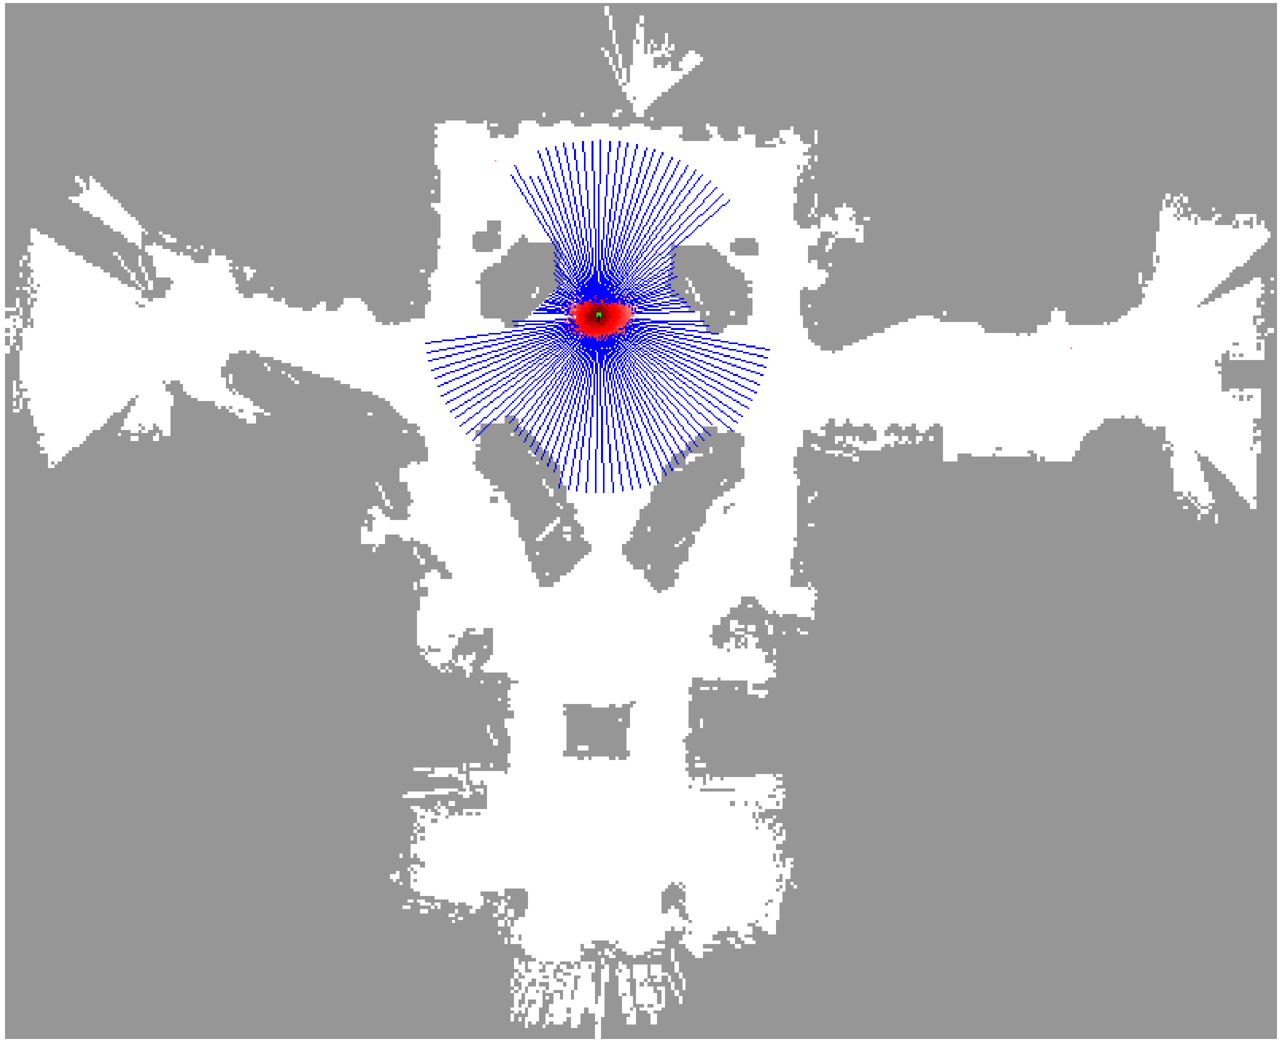

[Wolfram Burgard, TUN]

Example

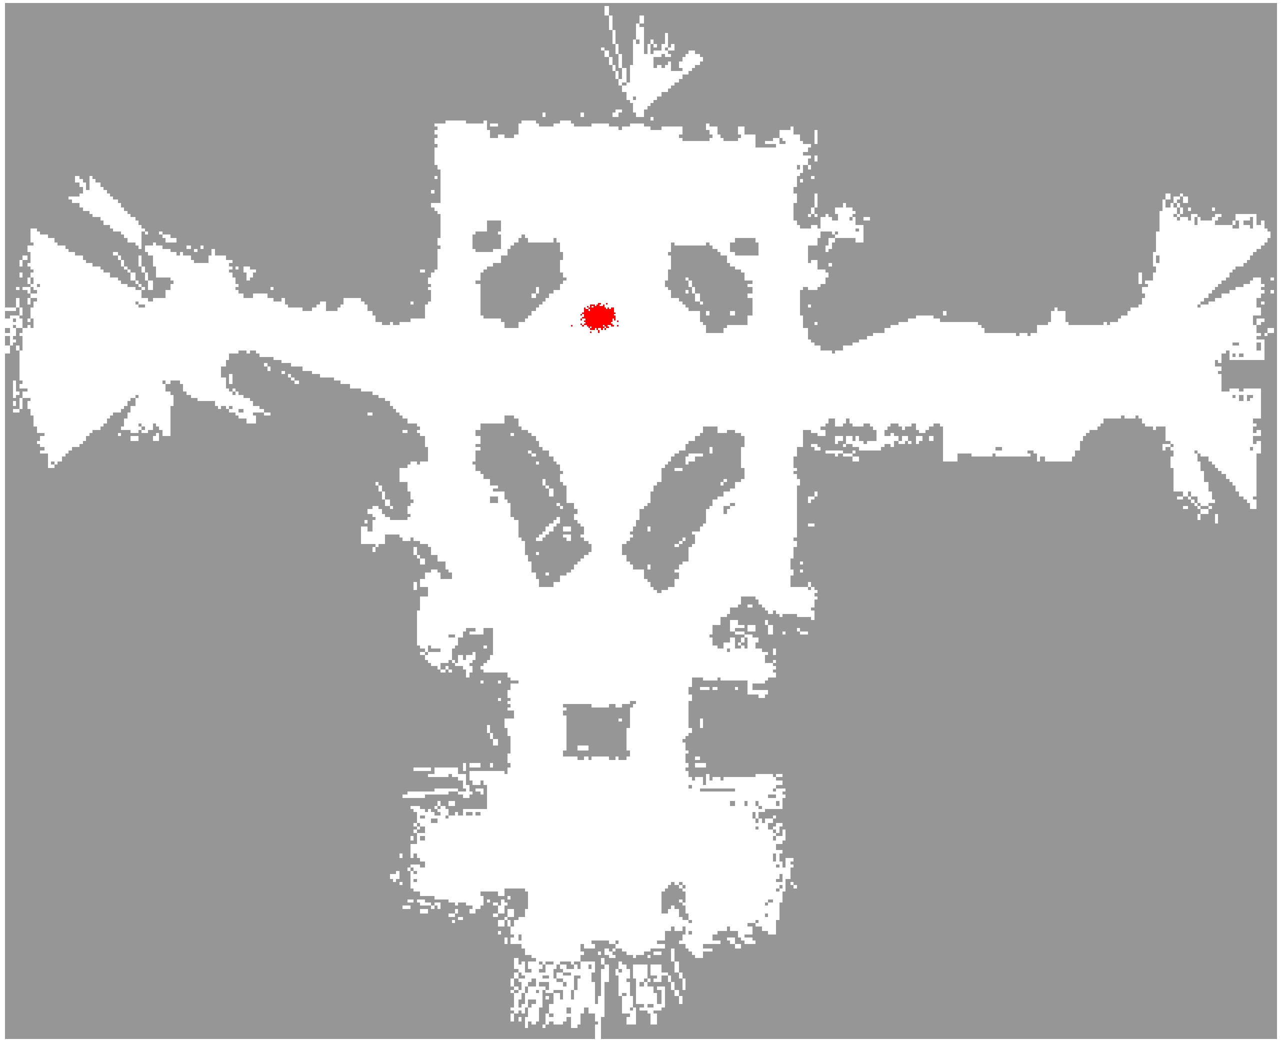

[Wolfram Burgard, TUN]

Example

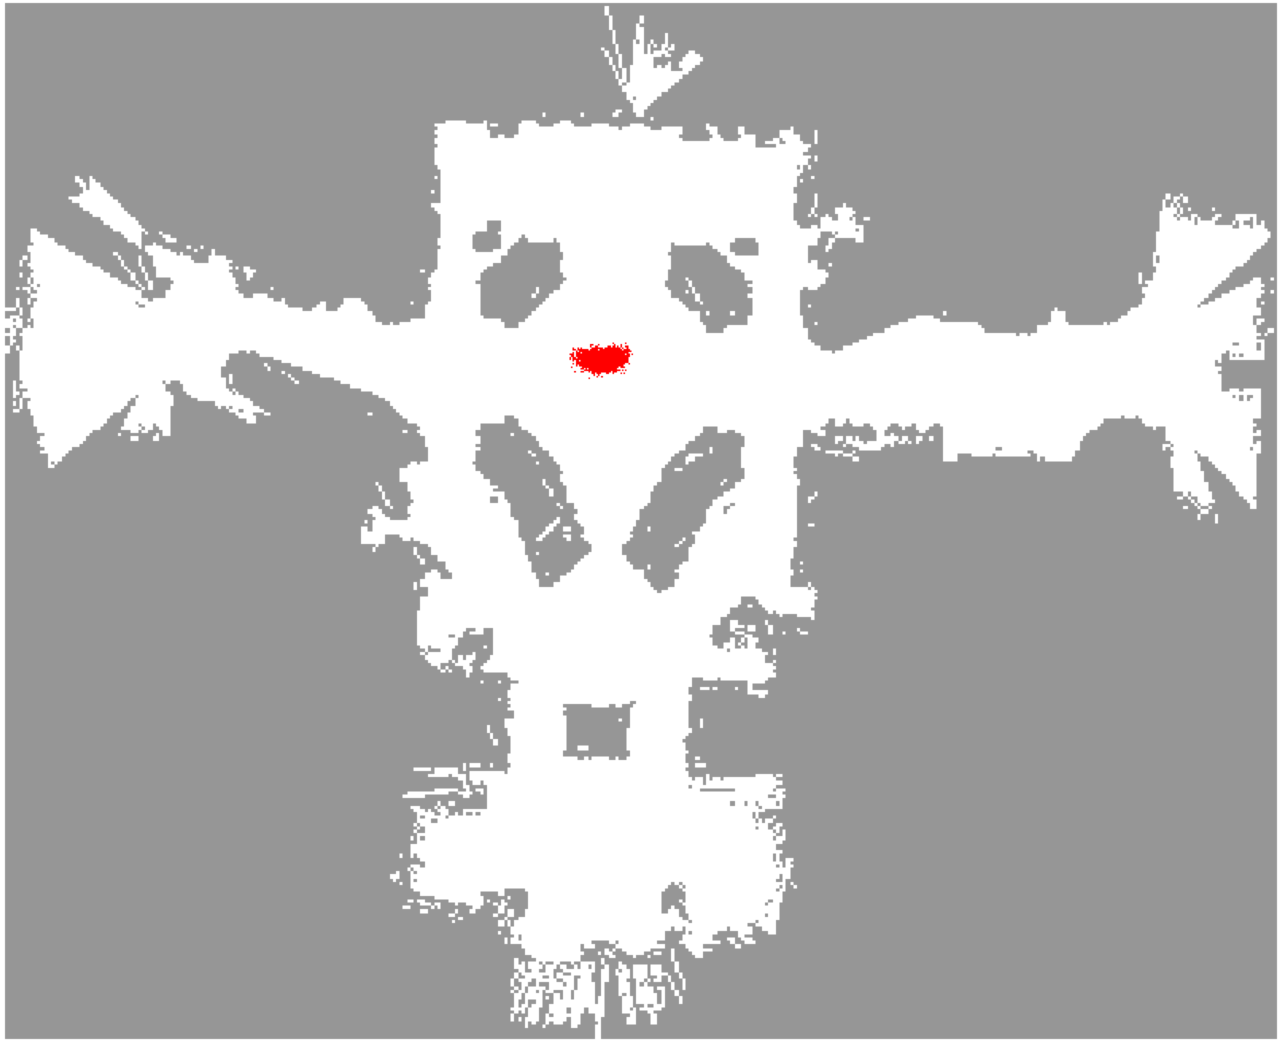

[Wolfram Burgard, TUN]

Example

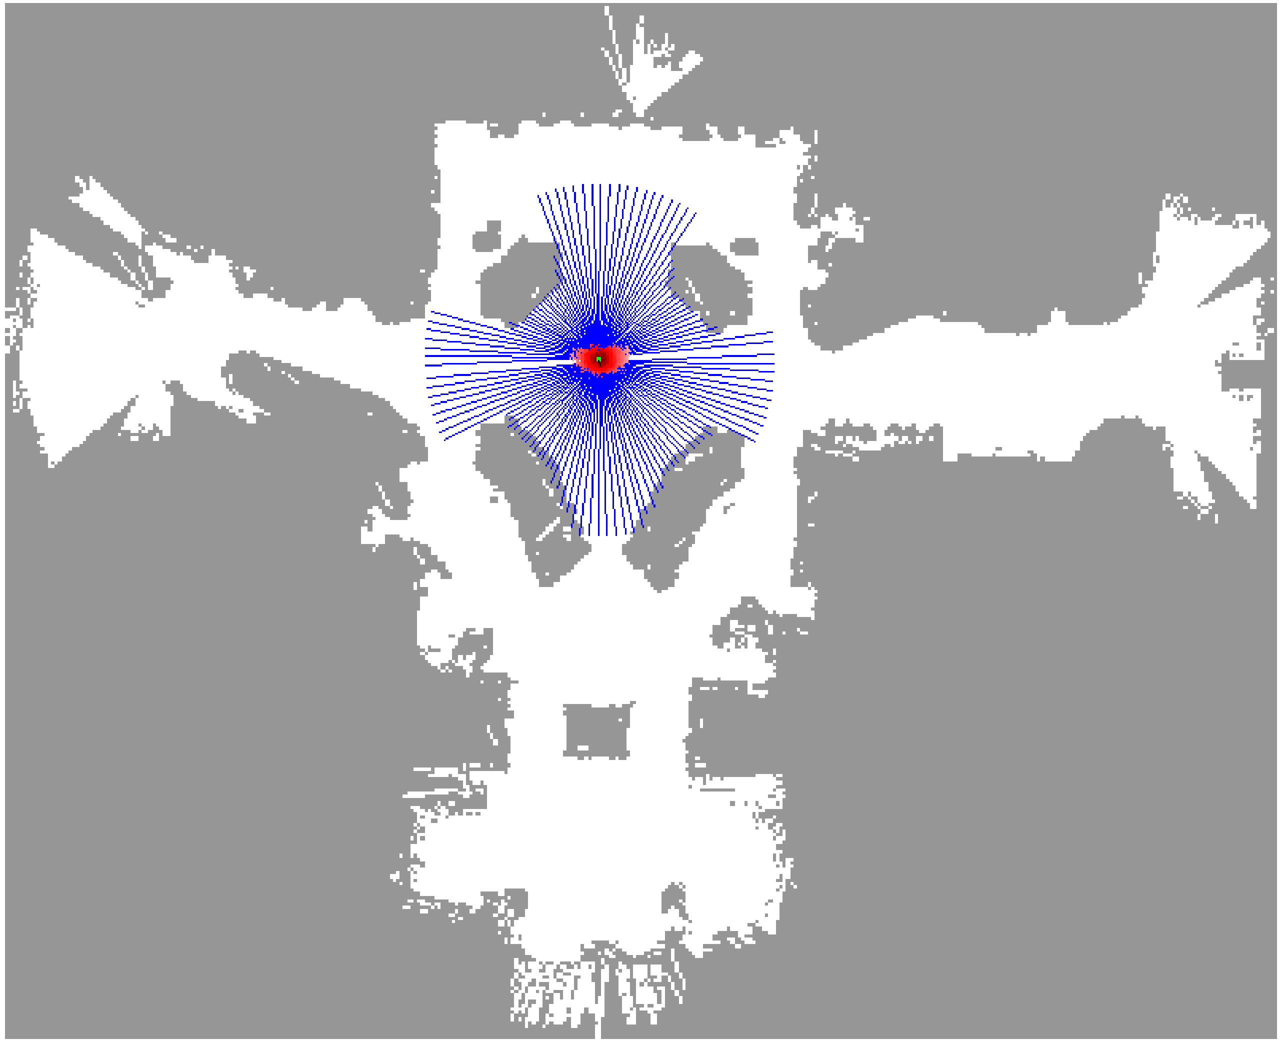

[Wolfram Burgard, TUN]

Particle filtering for passive global localization

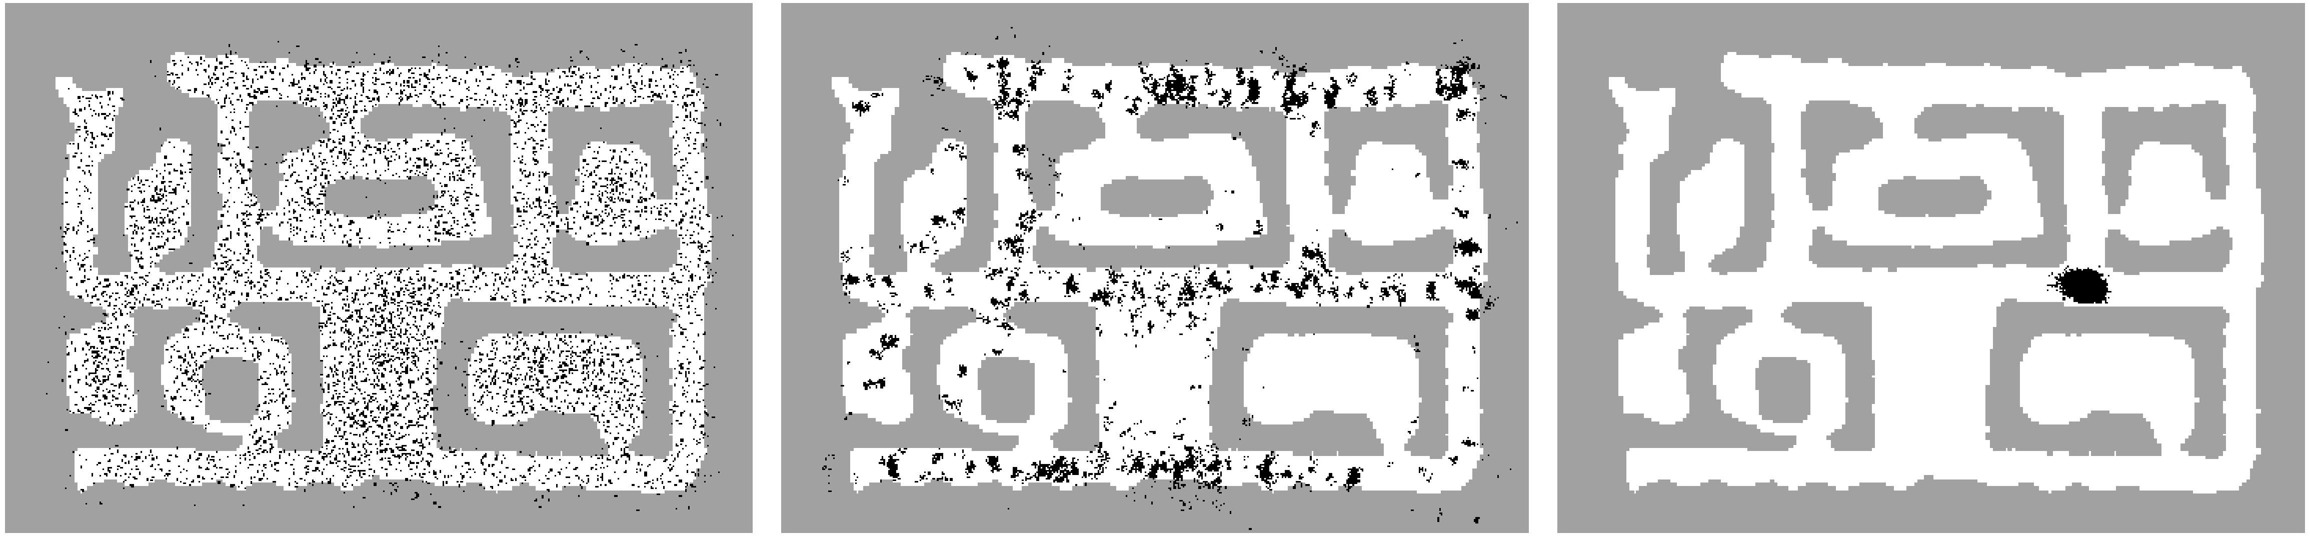

Particle filtering for active global localization