AI: Representation and Problem Solving

Particle Filtering

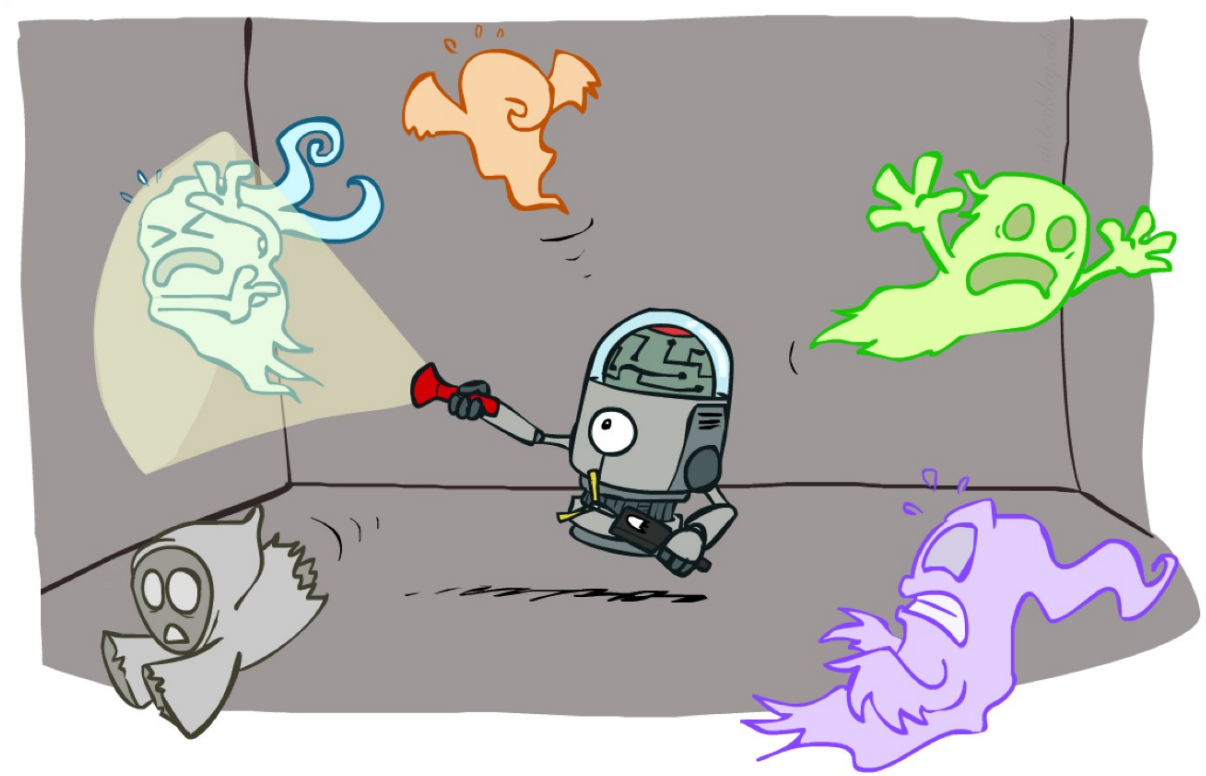

Instructor: Pat Virtue

Slide credits: CMU AI and http://ai.berkeley.edu

Poll 1

When sampling this Bays net using likelihood weighted sampling, what distribution do we have when we multiply fraction of counts times of the weight of the samples?

Write your distribution in terms of *X* and *e* starting with "P("

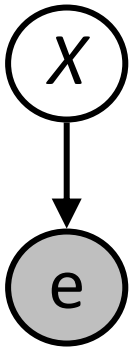

Particle Filtering

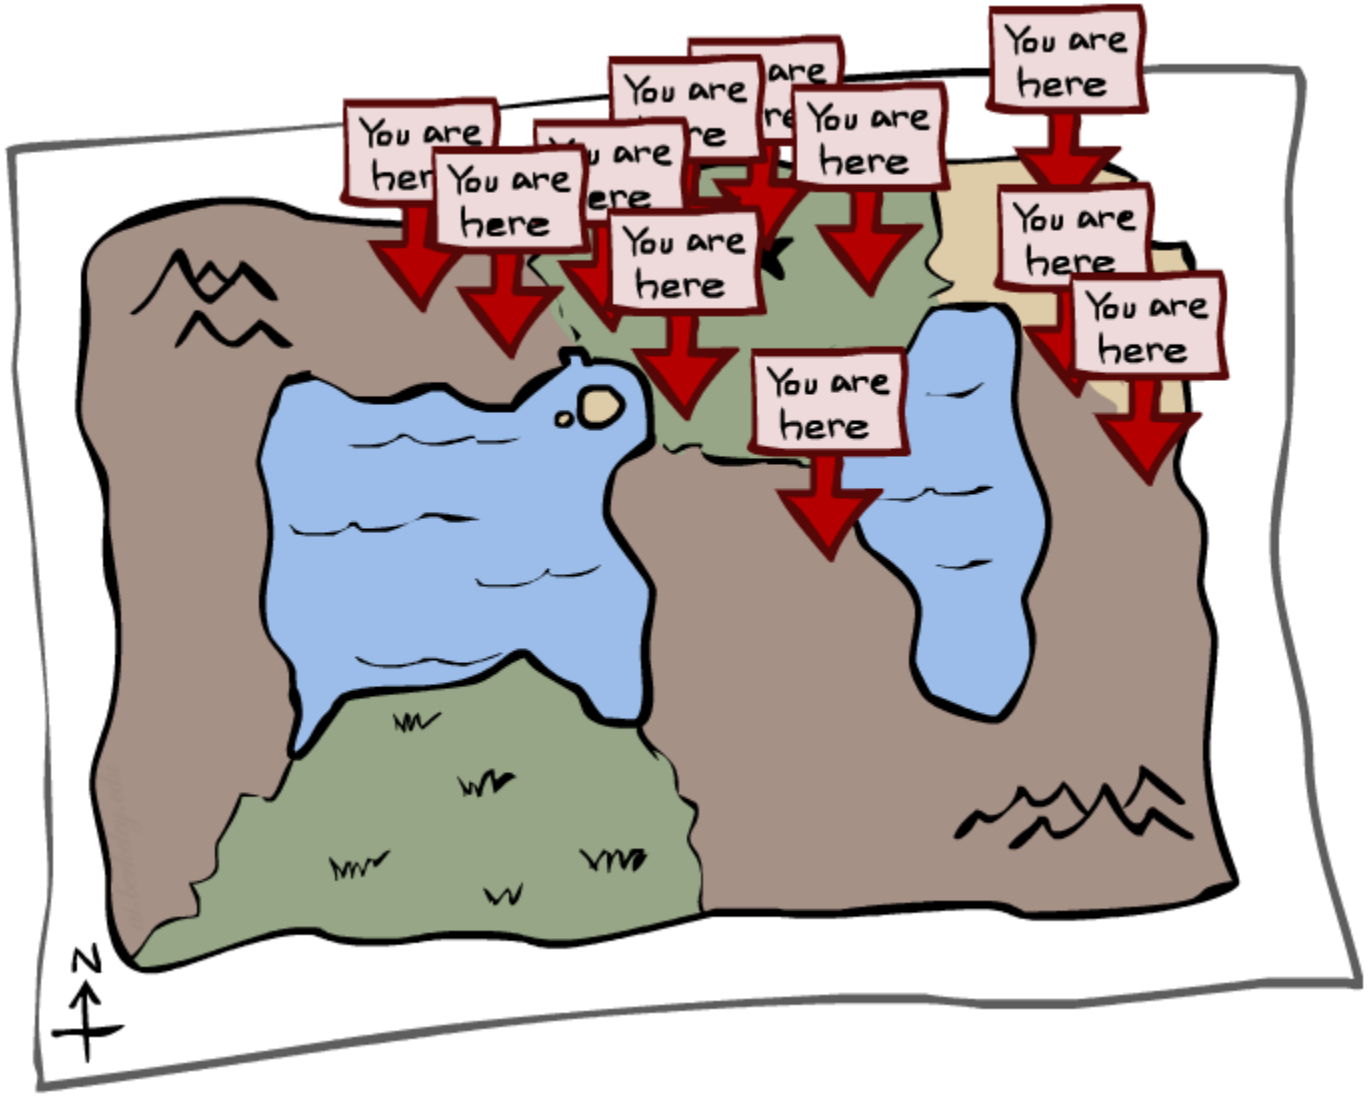

Summary: Likelihood-weighted sampling for HMMs

Particles: track samples of states rather than an explicit distribution

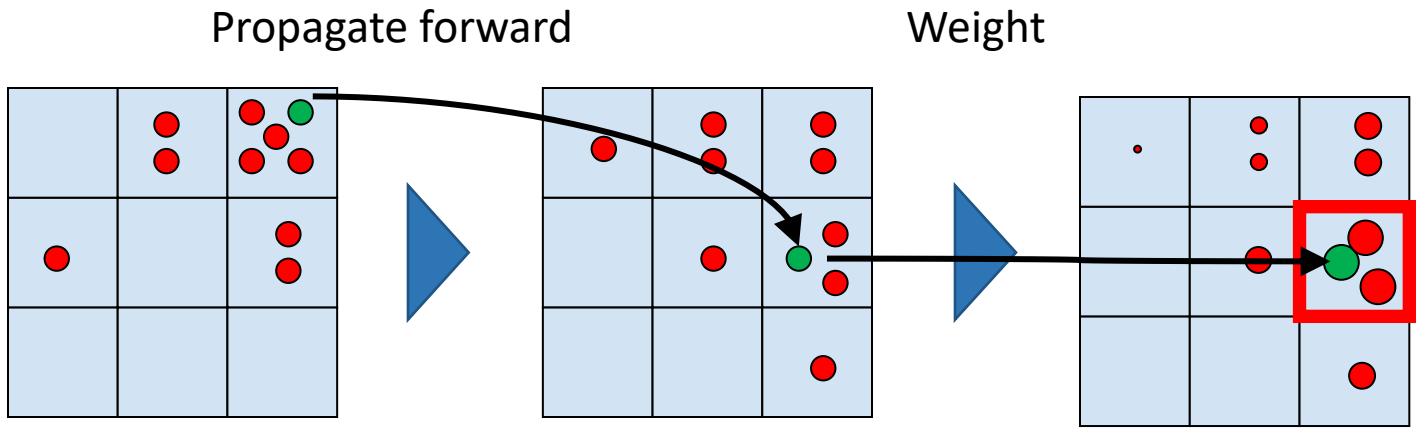

Particle Filtering

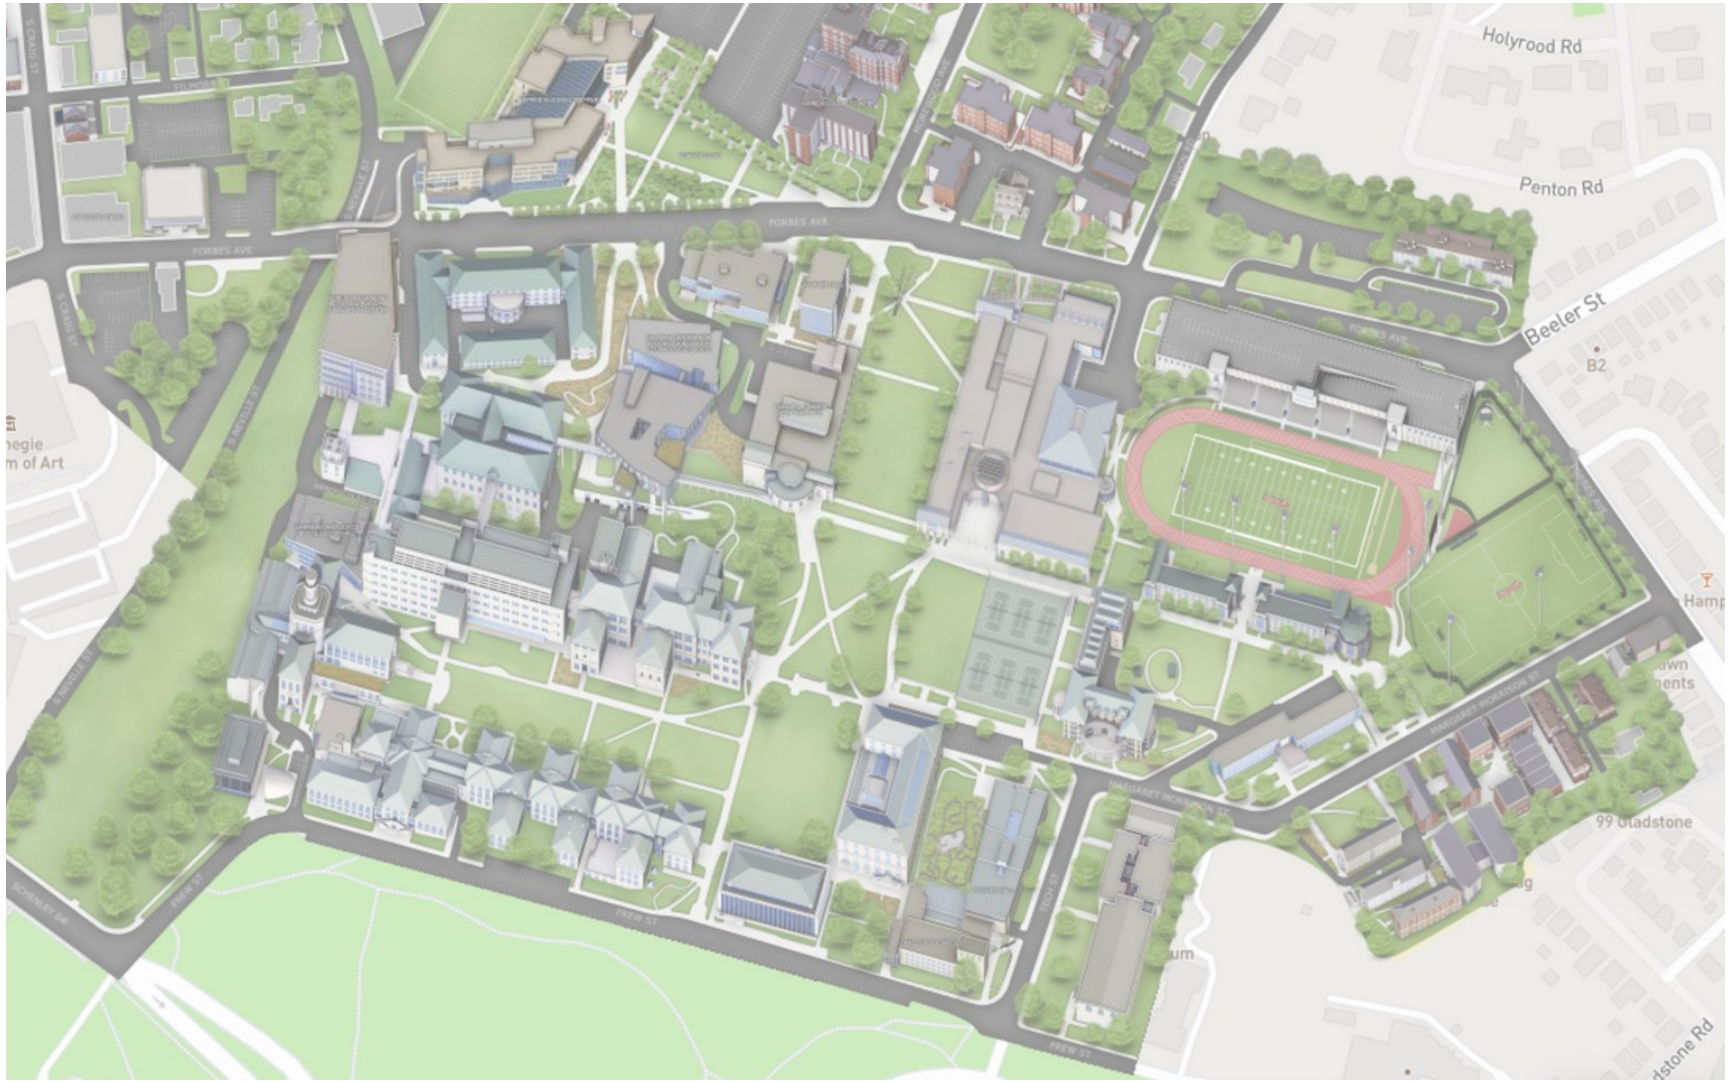

We need a new algorithm!

When $|X|$ is more than 10^6 or so (e.g., 3 ghosts in a 10x20 world), exact inference becomes infeasible

Likelihood weighting fails completely – number of samples needed grows ***exponentially*** with T

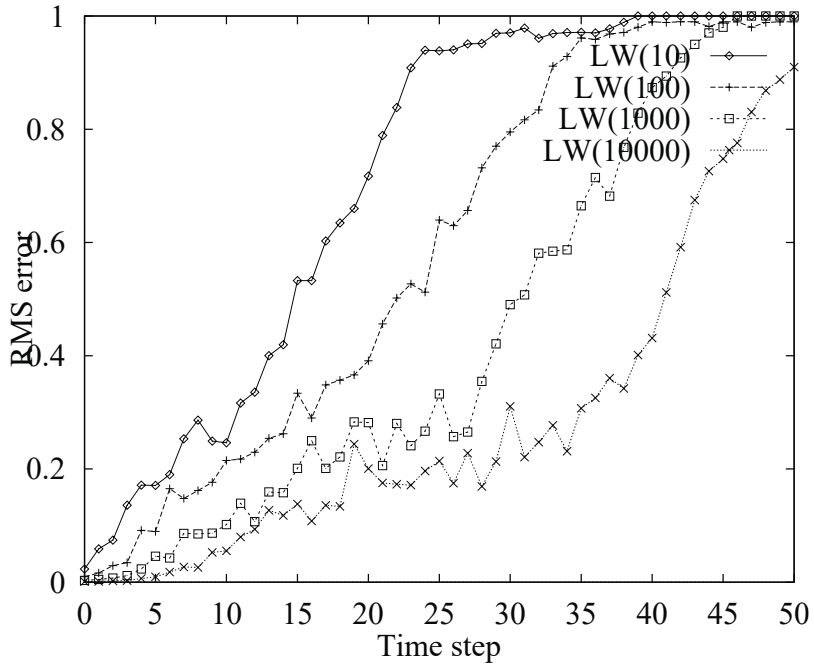

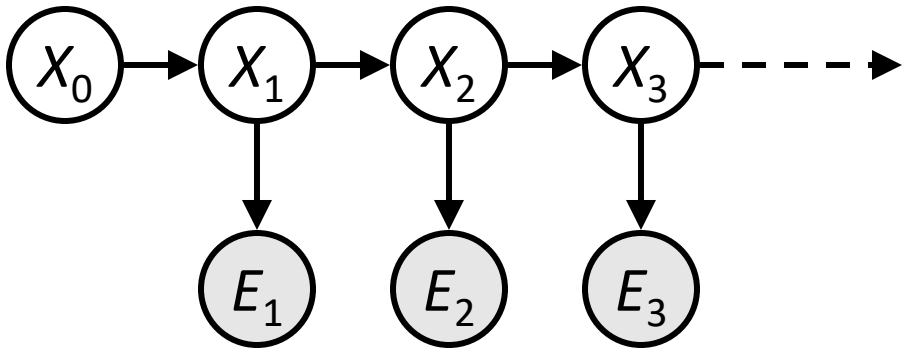

We need a new idea!

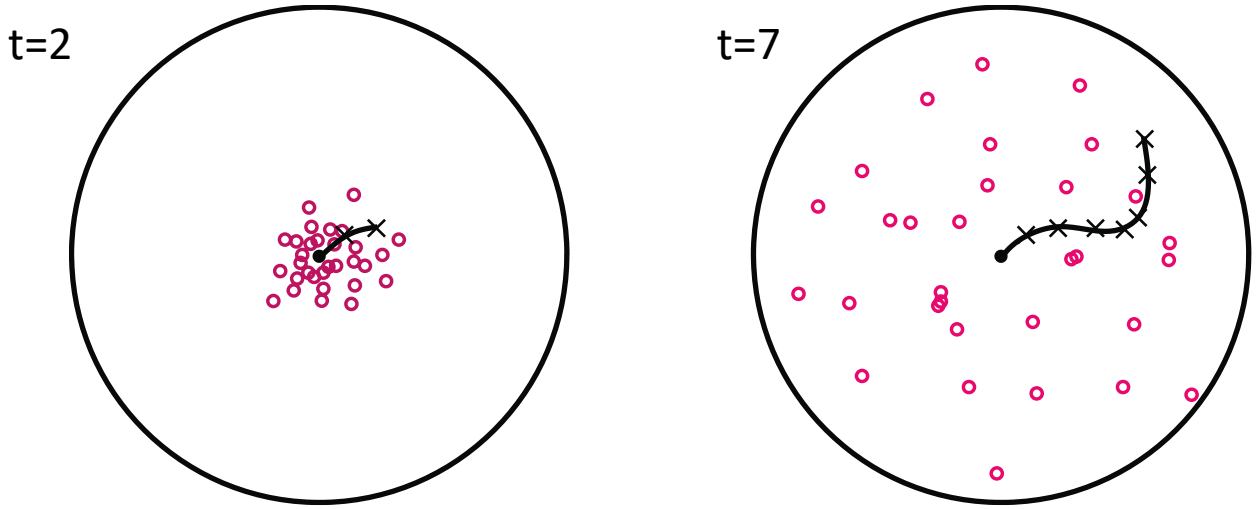

The problem: sample state trajectories go off into low-probability regions, ignoring the evidence; too few “reasonable” samples

Solution: kill the bad ones, make more of the good ones

This way the population of samples stays in the high-probability region

This is called resampling or survival of the fittest

Robot Localization

In robot localization:

- We know the map, but not the robot's position
- Observations may be vectors of range finder readings
- State space and readings are typically continuous (works basically like a very fine grid) and so we cannot store B(X)
- Particle filtering is a main technique

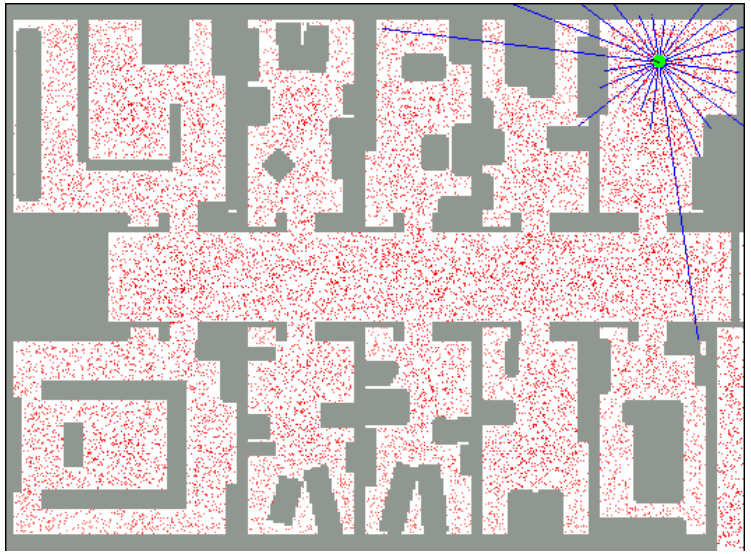

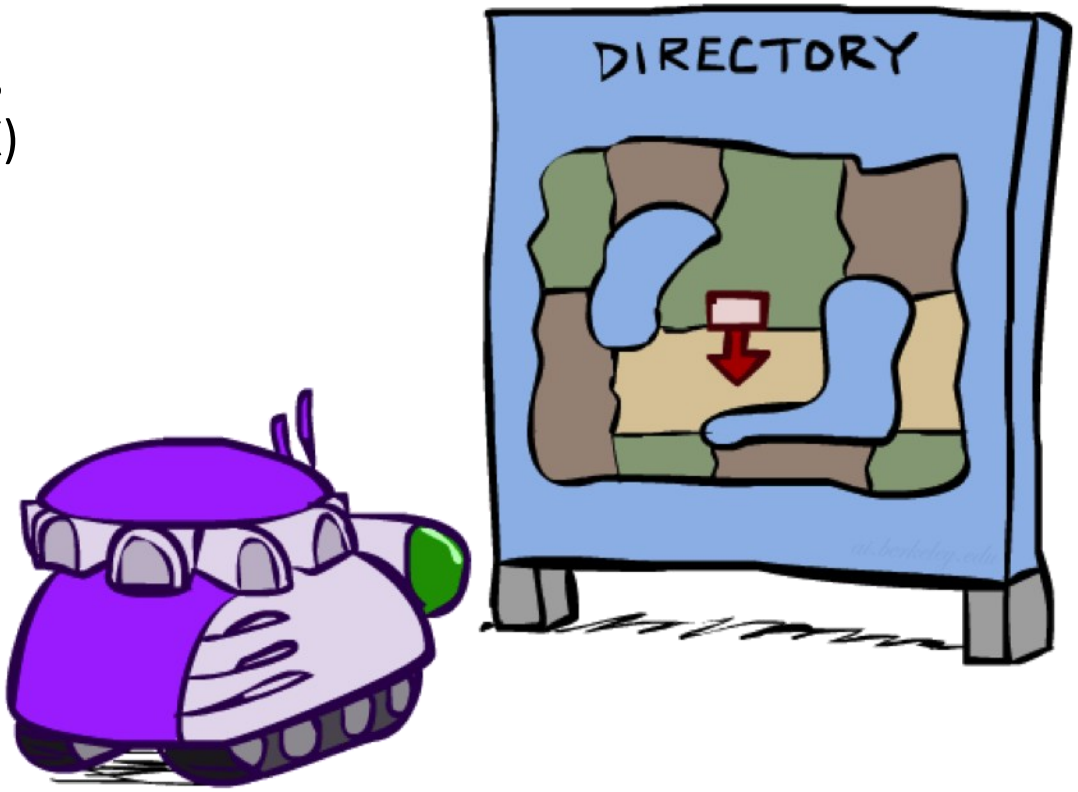

Particle Filter Localization (Sonar)

[Dieter Fox, et al.]

[Video: global-sonar-uw-annotated.avi]

Particle Filtering

- Represent belief state by a set of samples
  - Samples are called ***particles***
  - Time per step is linear in the number of samples
  - But: number needed may be large
- This is how robot localization works in practice

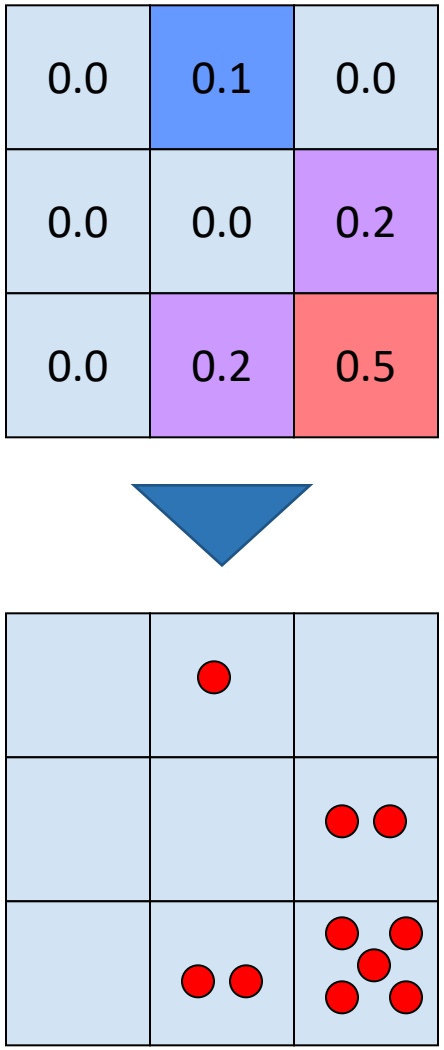

Representation: Particles

Our representation of P(X) is now a list of N particles (samples)

- Generally, N << |X|
- Storing map from X to counts would defeat the point

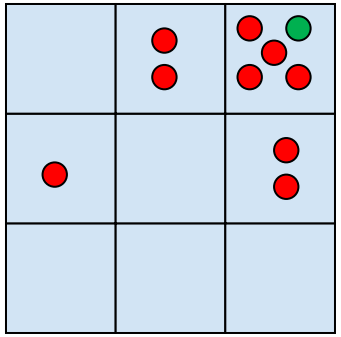

Particles:
(3,3)
(2,3)
(3,3)
(3,2)
(3,3)
(3,2)
(1,2)
(3,3)
(3,3)
(2,3)

P(x) approximated by number of particles with value x

- So, many x may have P(x) = 0!
- More particles, more accuracy
- Usually we want a low-dimensional marginal
  - E.g., “Where is ghost 1?” rather than “Are ghosts 1,2,3 in {2,6], [5,6], and [8,11]?”

For now, all particles have a weight of 1

Particle Filtering: Propagate forward

- A particle in state x_t is moved by sampling its next position directly from the transition model:
  - $x_{t+1} \sim P(X_{t+1} \mid x_t)$
  - Here, most samples move clockwise, but some move in another direction or stay in place
- This captures the passage of time
  - If enough samples, close to exact values before and after (consistent)

Particles:
(3,3)
(2,3)
(3,3)
(3,2)
(3,3)
(3,2)
(1,2)
(3,3)
(3,3)
(2,3)

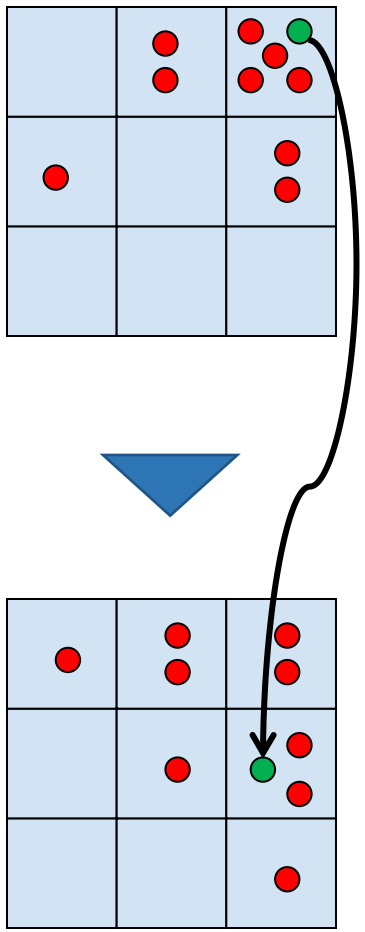

Particles:
(3,2)
(2,3)
(3,2)
(3,1)
(3,3)
(3,2)
(1,3)
(2,3)
(3,3)
(2,2)

Particle Filtering: Observe

- Slightly trickier:
  - Don't sample observation, fix it
  - Similar to likelihood weighting, weight samples based on the evidence
    - $W = P(e_t \mid x_t)$
  - Normalize the weights: particles that fit the data better get higher weights, others get lower weights

Particles:
(3,2)
(2,3)
(3,2)
(3,1)
(3,3)
(3,2)
(1,3)
(2,3)
(3,2)
(2,2)

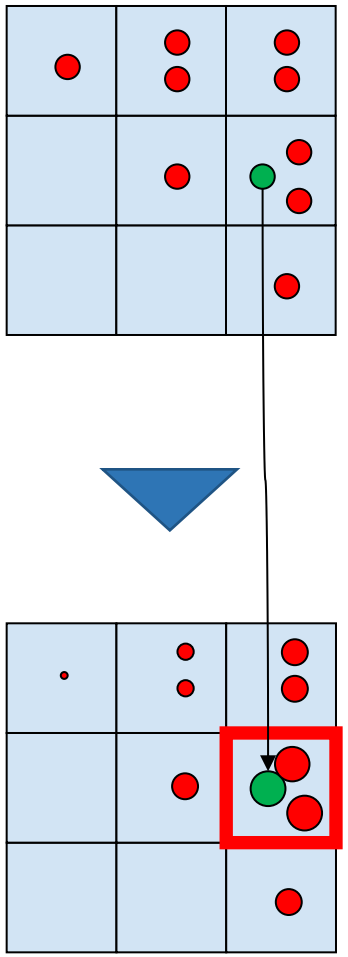

Particles:
(3,2) w=.9
(2,3) w=.2
(3,2) w=.9
(3,1) w=.4
(3,3) w=.4
(3,2) w=.9
(1,3) w=.1
(2,3) w=.2
(3,3) w=.4
(2,2) w=.4

Particle Filtering: Resample

Rather than tracking weighted samples, we ***resample***

We have an updated belief distribution based on the weighted particles

We sample N new particles from the weighted belief distributions

Now the update is complete for this time step, continue with the next one

Particles:
(3,2) w=.9
(2,3) w=.2
(3,2) w=.9
(3,1) w=.4
(3,3) w=.4
(3,2) w=.9
(1,3) w=.1
(2,3) w=.2
(3,3) w=.4
(2,2) w=.4

(New) Particles:
(3,2)
(2,2)
(3,2)
(2,3)
(3,3)
(3,2)
(1,3)
(2,3)
(3,2)
(3,2)

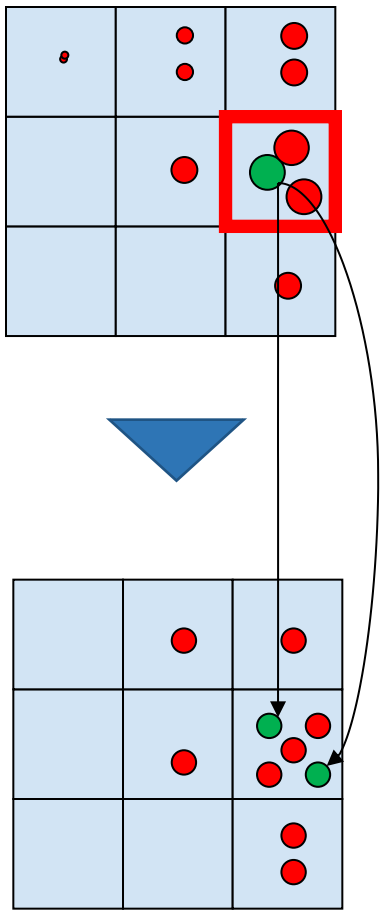

Summary: Particle Filtering

Particles: track samples of states rather than an explicit distribution

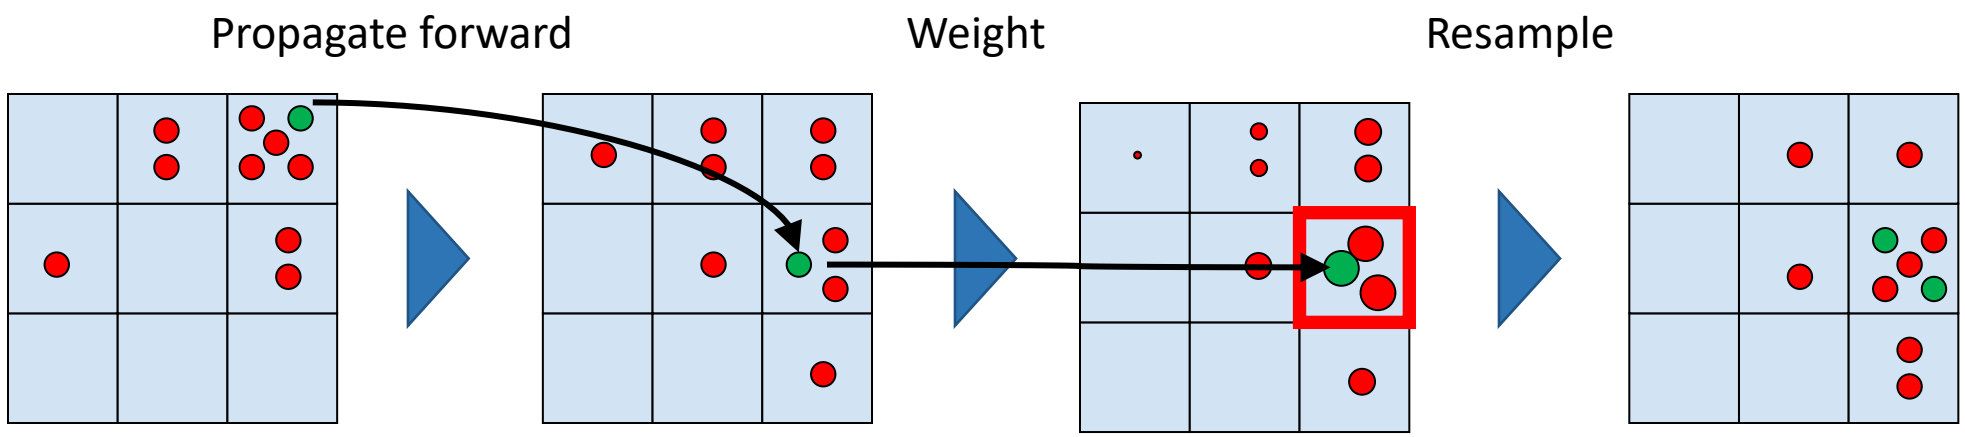

Particles:	Particles:	Particles:	(New) Particles:
(3,3)	(3,2)	(3,2) w=.9	(3,2)
(2,3)	(2,3)	(2,3) w=.2	(2,2)
(3,3)	(3,2)	(3,2) w=.9	(3,2)
(3,2)	(3,1)	(3,1) w=.4	(2,3)
(3,3)	(3,3)	(3,3) w=.4	(3,3)
(3,2)	(3,2)	(3,2) w=.9	(3,2)
(1,2)	(1,3)	(1,3) w=.1	(1,3)
(3,3)	(2,3)	(2,3) w=.2	(2,3)
(3,3)	(3,2)	(3,3) w=.4	(3,2)
(2,3)	(2,2)	(2,2) w=.4	(3,2)

Consistency: see proof in AIMA Ch. 14

[Demos: ghostbusters particle filtering (L15D3,4,5)]

Weighting and Resampling

How to compute a belief distribution given weighted particles

Weight

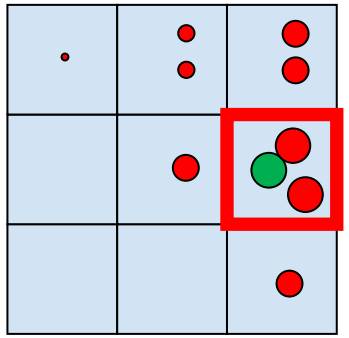

Particles:

(3,2) w=.9
(2,3) w=.2
(3,2) w=.9
(3,1) w=.4
(3,3) w=.4
(3,2) w=.9
(1,3) w=.1
(2,3) w=.2
(3,3) w=.4
(2,2) w=.4

Poll 2

If we only have one particle which of these steps are unnecessary?

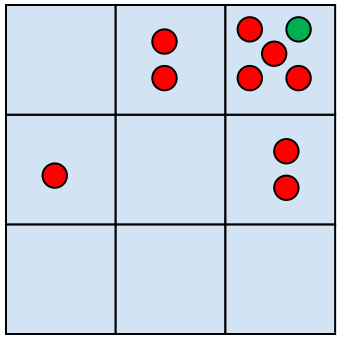

Select all that are unnecessary.

A. Propagate forward
B. Weight
C. Resample
D. None of the above

[Demos: ghostbusters particle filtering (L15D3,4,5)]

Poll 2

If we only have one particle which of these steps are unnecessary?

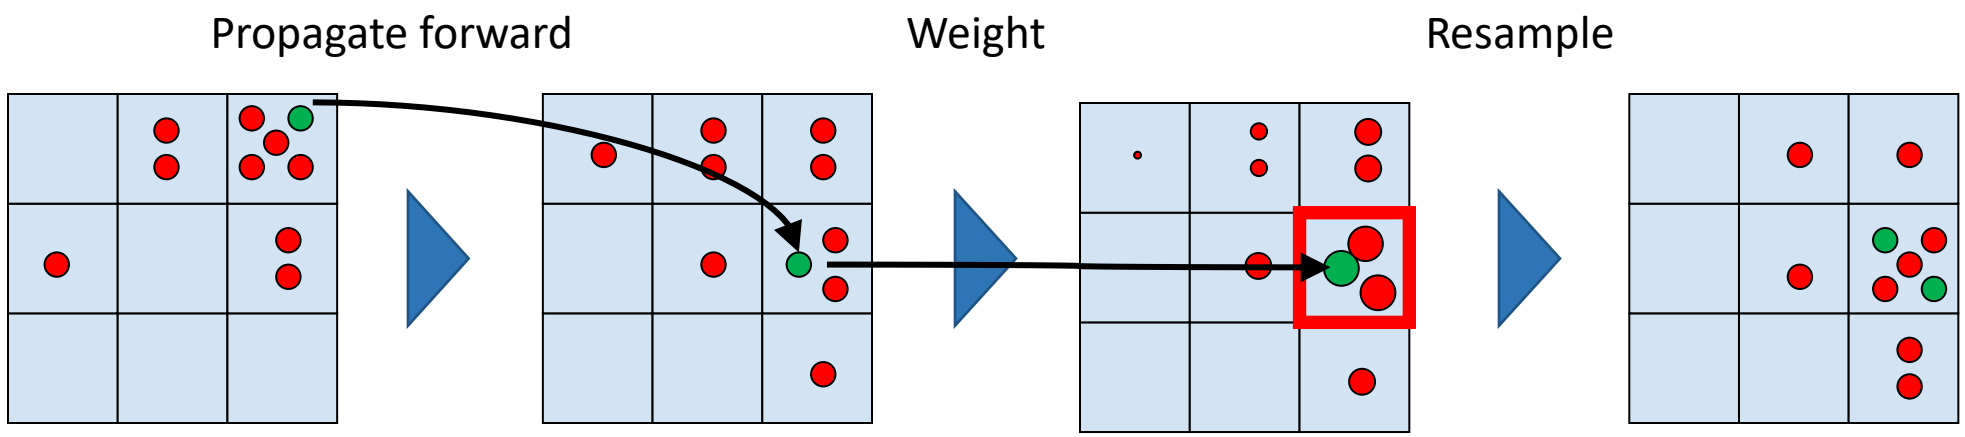

Select all that are unnecessary.

A. Propagate forward

B. Weight

C. Resample

D. None of the above

Unless the weight is zero, in which case, you'll want to resample from the beginning ☹

[Demos: ghostbusters particle filtering (L15D3,4,5)]

Particle Filter Localization (Laser)

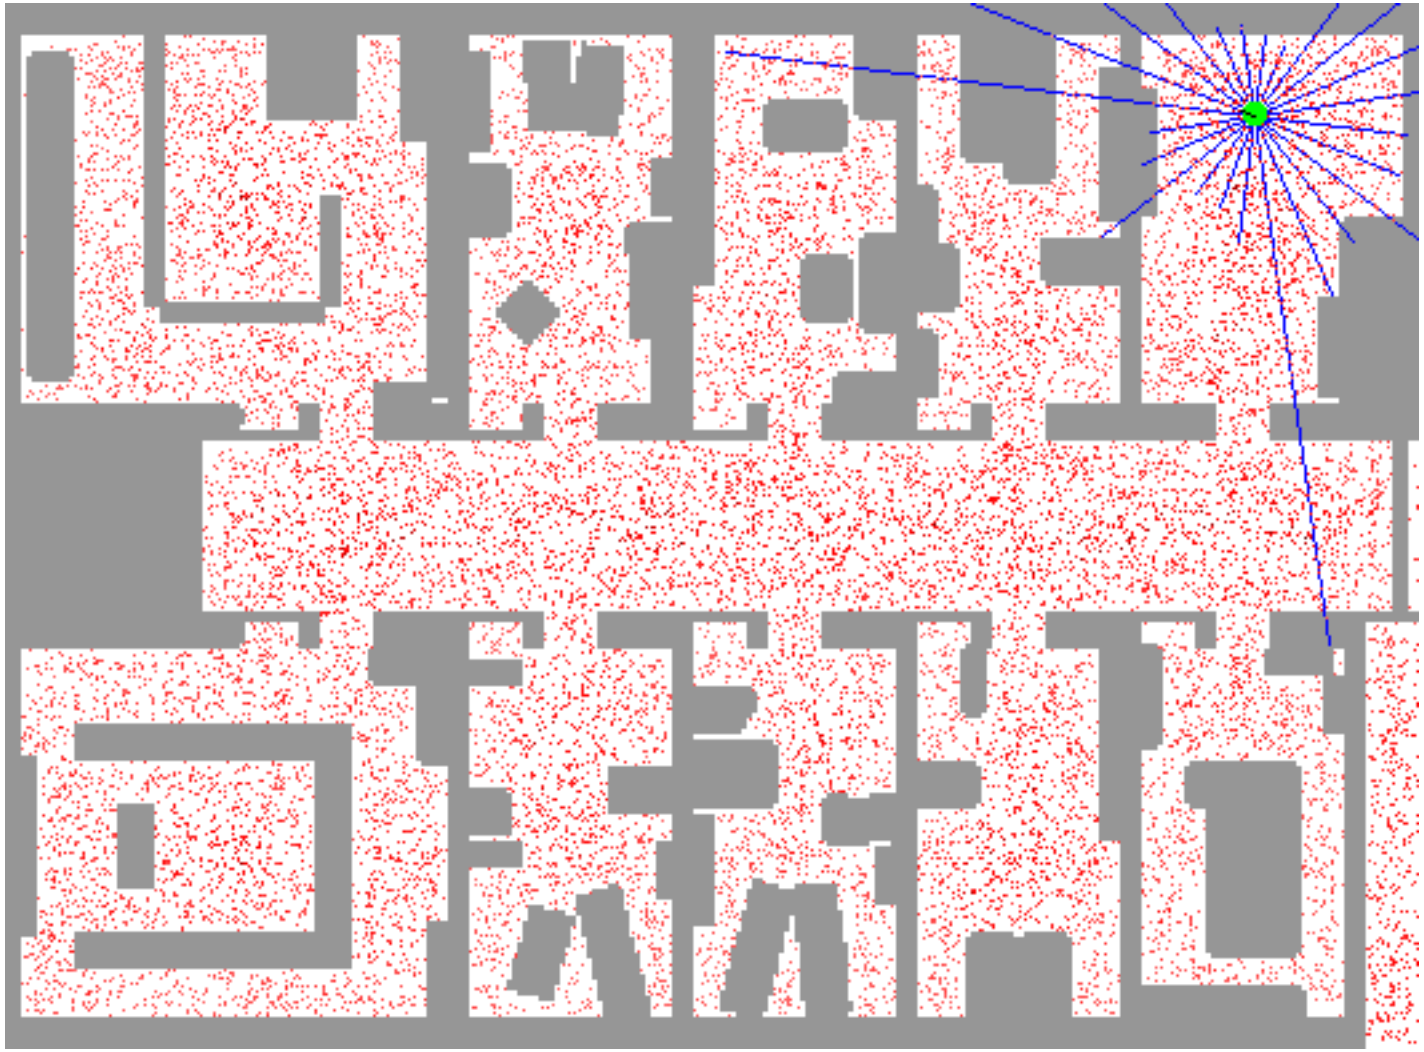

[Dieter Fox, et al.]

[Video: global-floor.gif]

Robot Mapping

SLAM: Simultaneous Localization And Mapping

- We do not know the map or our location
- State consists of position AND map!
- Main techniques: Kalman filtering (Gaussian HMMs) and particle methods

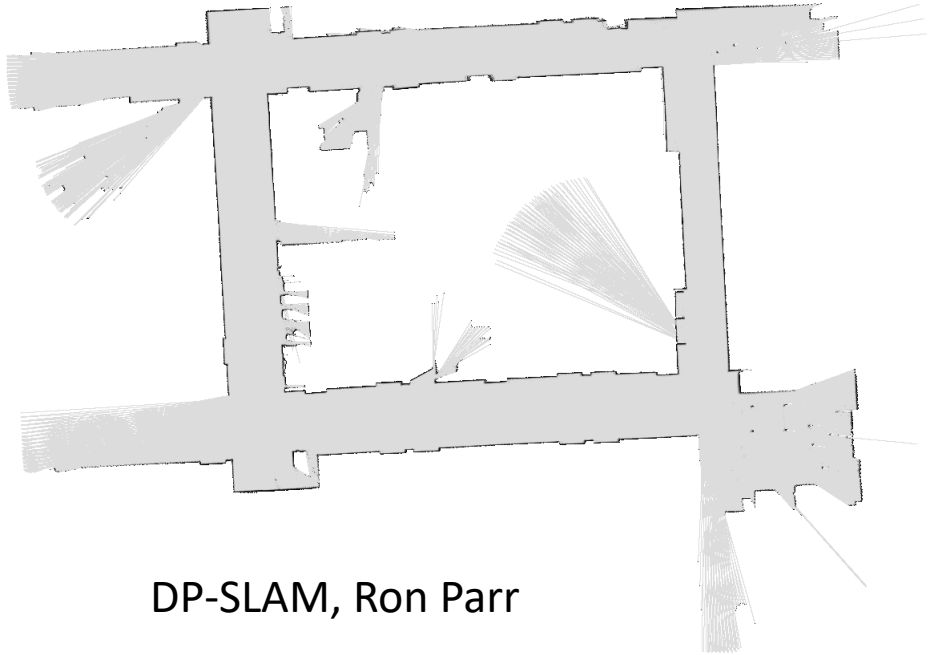

DP-SLAM, Ron Parr

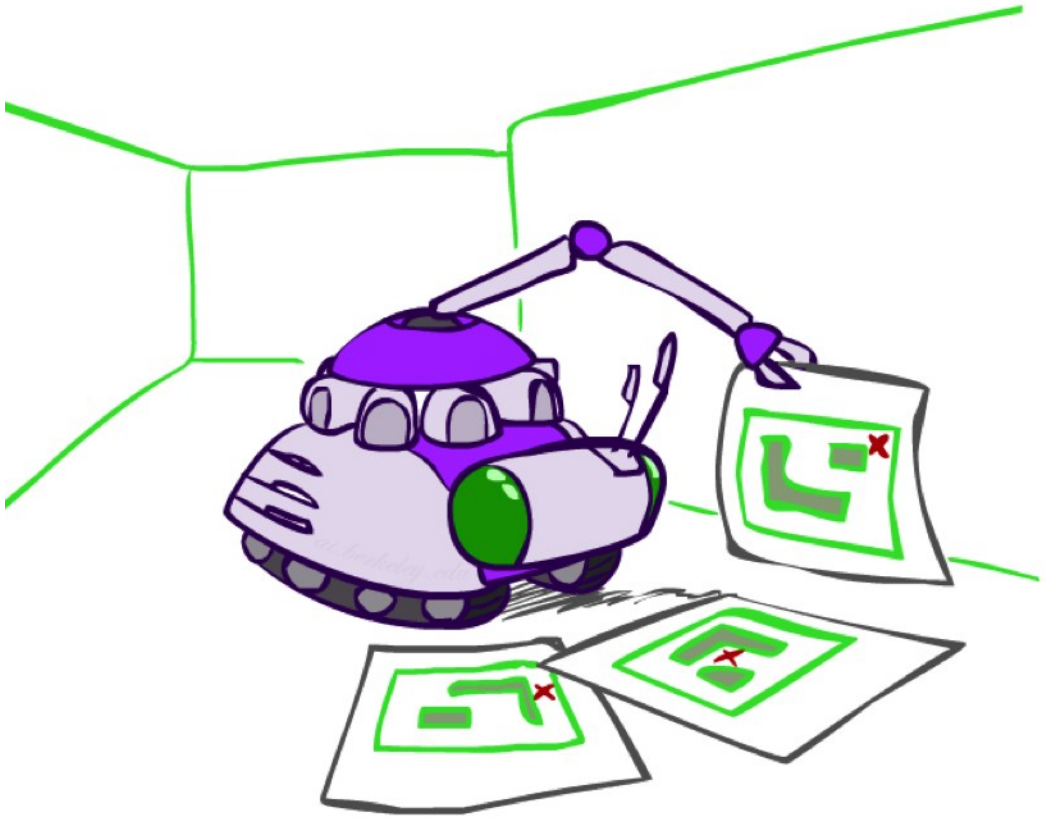

[Demo: PARTICLES-SLAM-mapping1-new.avi]

Particle Filter SLAM – Video 1

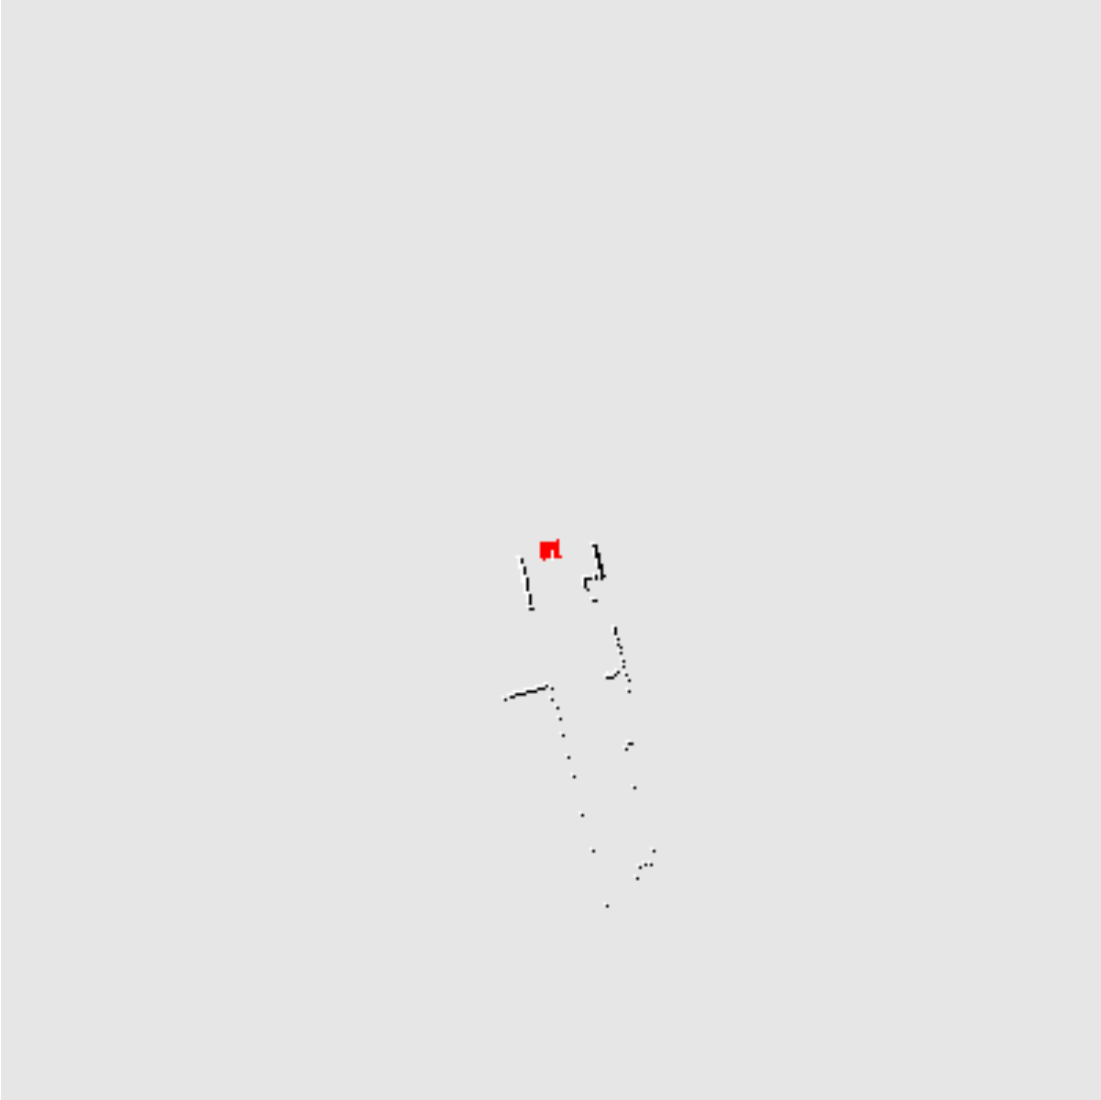

[Sebastian Thrun, et al.]

[Demo: PARTICLES-SLAM-mapping1-new.avi]

Particle Filter SLAM – Video 2

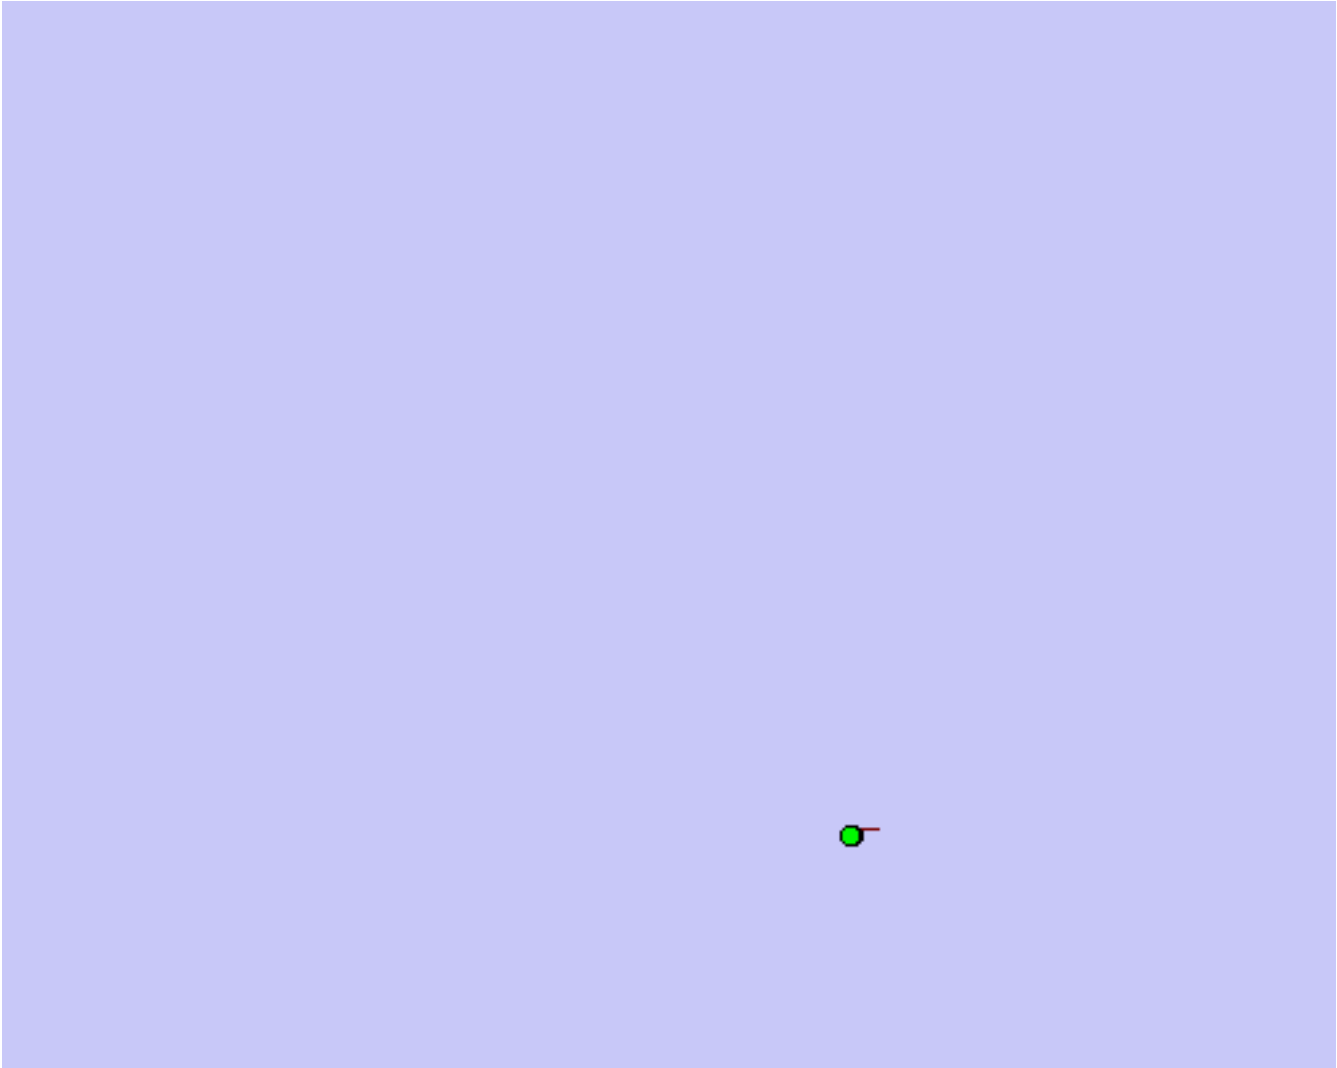

[Dirk Haehnel, et al.]

[Demo: PARTICLES-SLAM-fastslam.avi]

SLAM

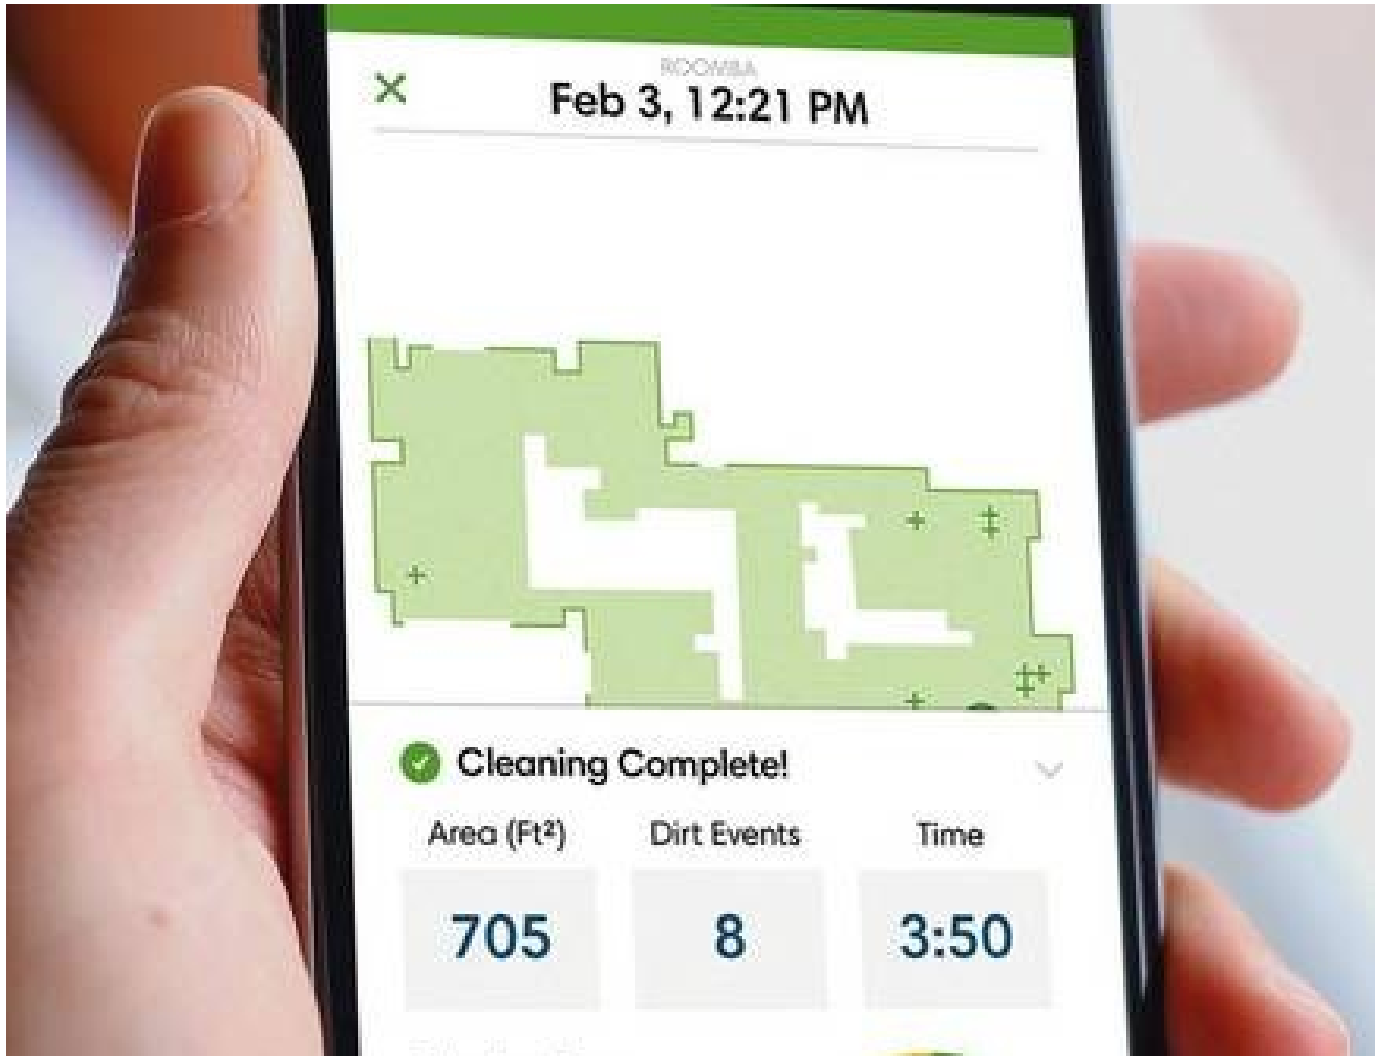

https://www.irobot.com/