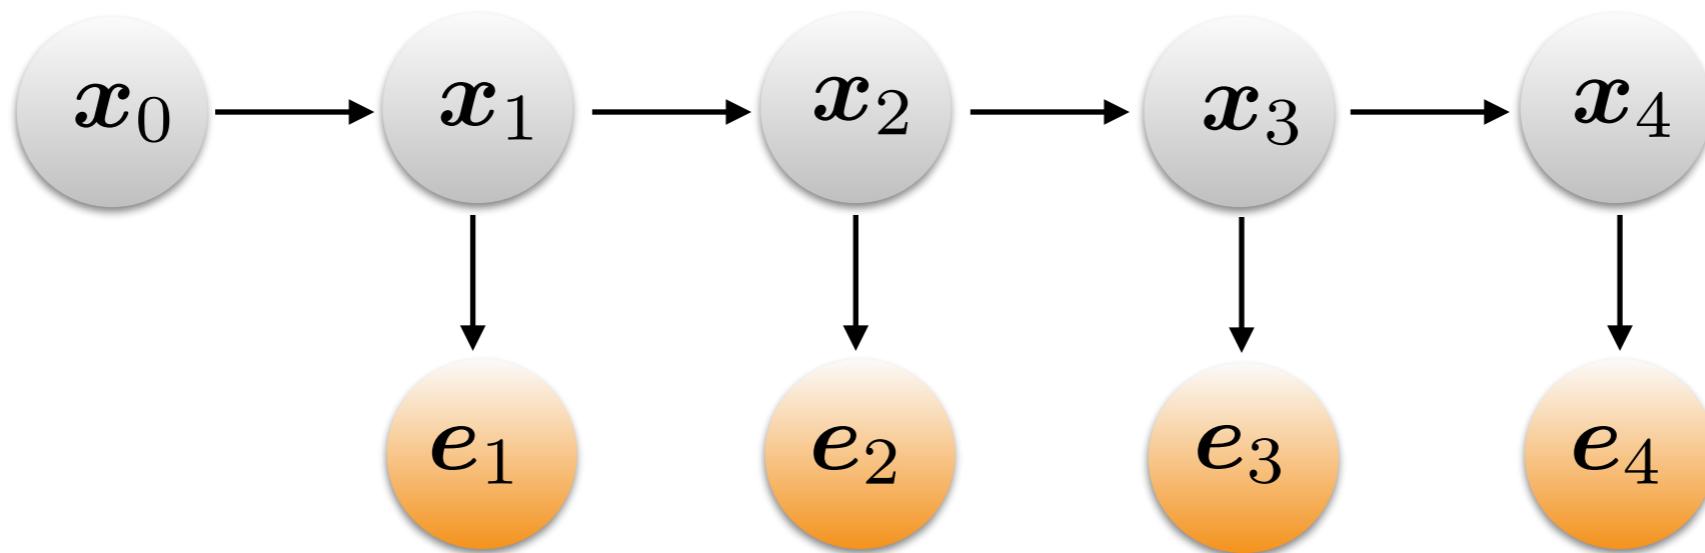


Kalman Filter

16-385 Computer Vision (Kris Kitani)
Carnegie Mellon University

Examples up to now have been **discrete** (binary) random variables

Kalman ‘filtering’ can be seen as a special case of a temporal inference with continuous random variables



Everything is continuous...

\boldsymbol{x} \boldsymbol{e} $P(\boldsymbol{x}_0)$ $P(\boldsymbol{e}|\boldsymbol{x})$ $P(\boldsymbol{x}_t|\boldsymbol{x}_{t-1})$

probability distributions are no longer tables but functions

Making the connection to the ‘filtering’ equations

(Discrete) Filtering

Tables

Tables

Tables

$$P(\boldsymbol{X}_{t+1} | \boldsymbol{e}_{1:t+1}) \propto P(\boldsymbol{e}_{t+1} | \boldsymbol{X}_{t+1}) \sum_{\boldsymbol{X}_t} P(\boldsymbol{X}_{t+1} | \boldsymbol{X}_t) P(\boldsymbol{X}_t | \boldsymbol{e}_{1:t})$$

Kalman Filtering

Gaussian

Gaussian

Gaussian

$$P(\mathbf{X}_{t+1} | \mathbf{e}_{1:t+1}) \propto P(\mathbf{e}_{t+1} | \mathbf{X}_{t+1}) \int_{\mathbf{x}_t} P(\mathbf{X}_{t+1} | \mathbf{x}_t) P(\mathbf{x}_t | \mathbf{e}_{1:t}) d\mathbf{x}_t$$

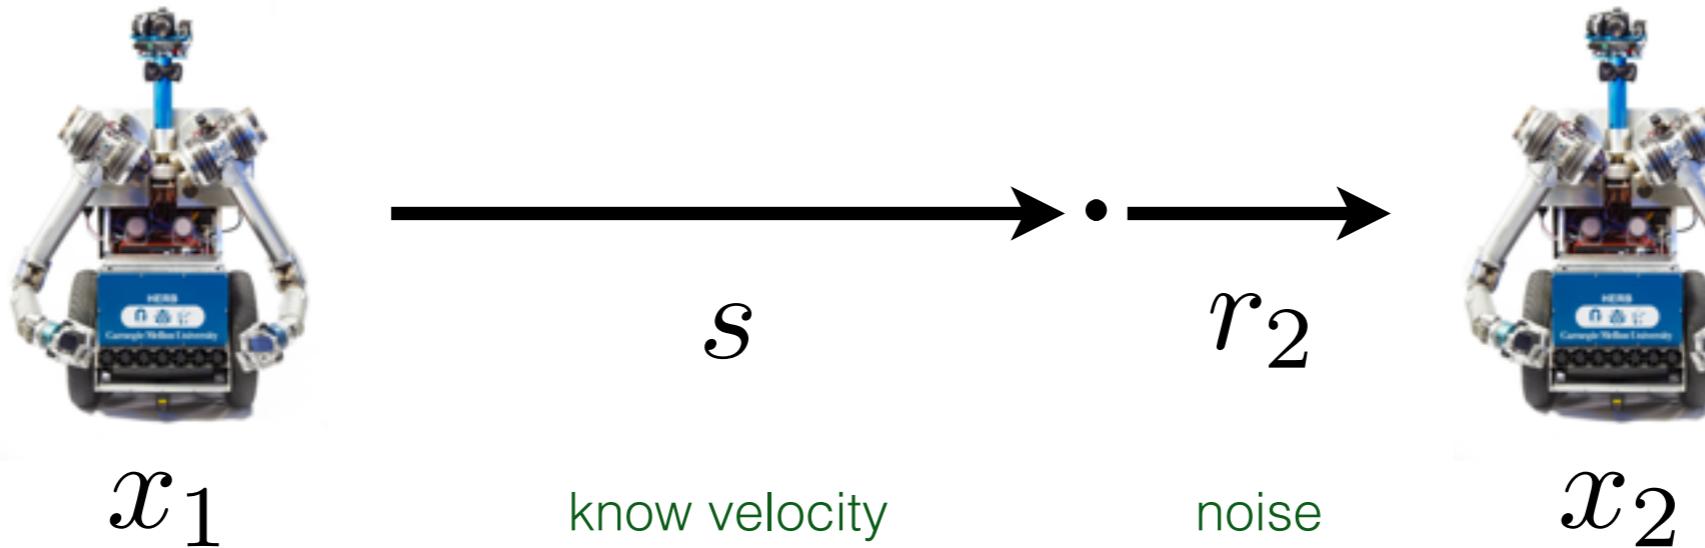
observation model $\int_{\mathbf{x}_t}$ motion model belief

integral because
continuous PDFs

Simple, 1D example...

x



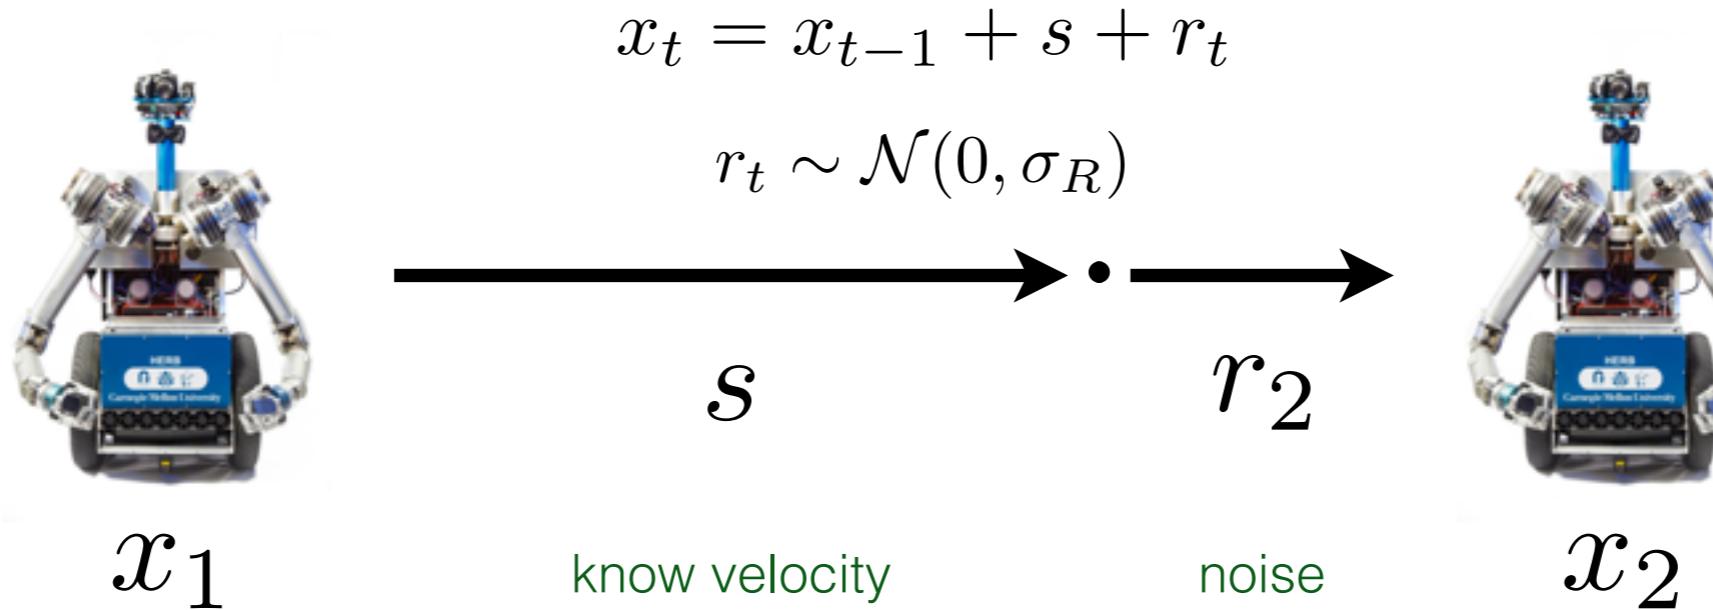


$$x_t = x_{t-1} + s + r_t$$

$$r_t \sim \mathcal{N}(0, \sigma_R)$$

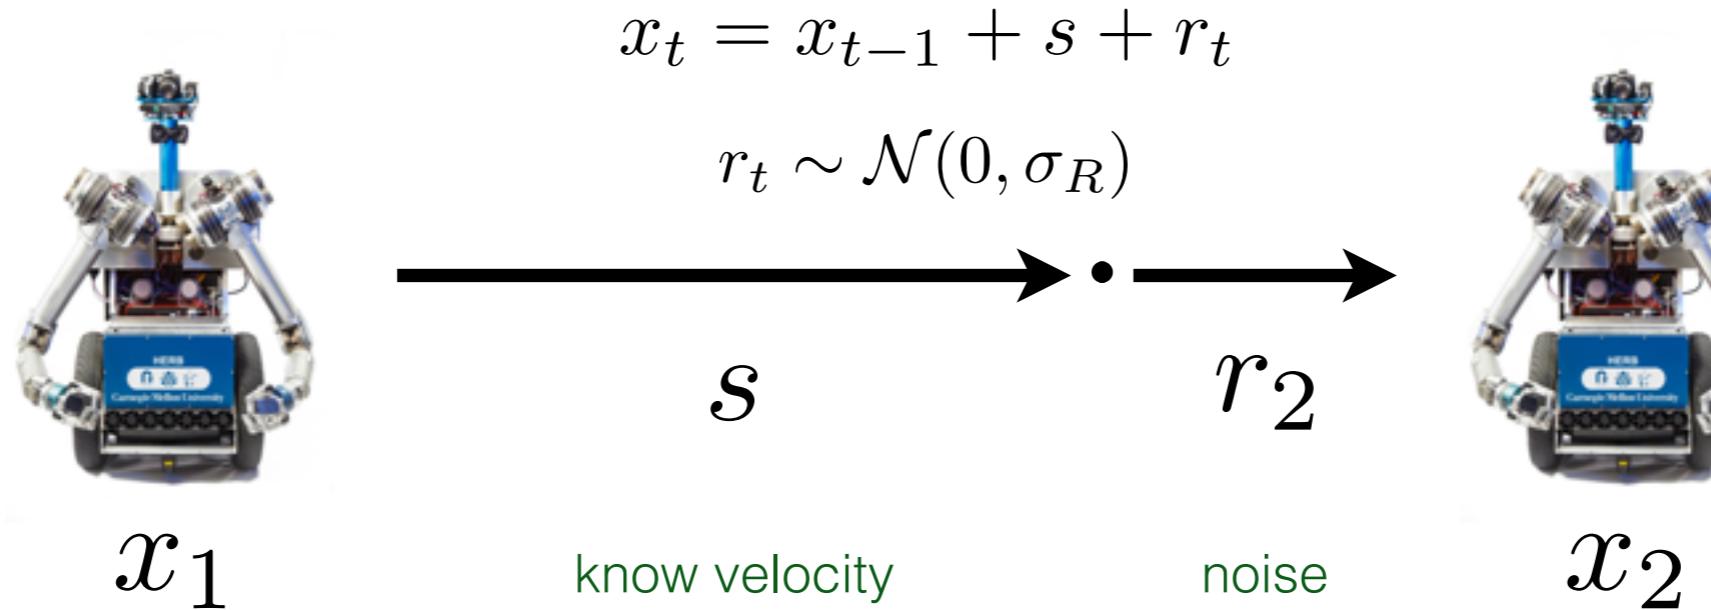
'sampled from'

System (motion) model



How do you represent the motion model?

$$P(x_t | x_{t-1})$$



How do you represent the motion model?

A linear Gaussian (continuous) transition model

$$P(x_t | x_{t-1}) = \mathcal{N}(x_t; x_{t-1} + s, \sigma_r)$$

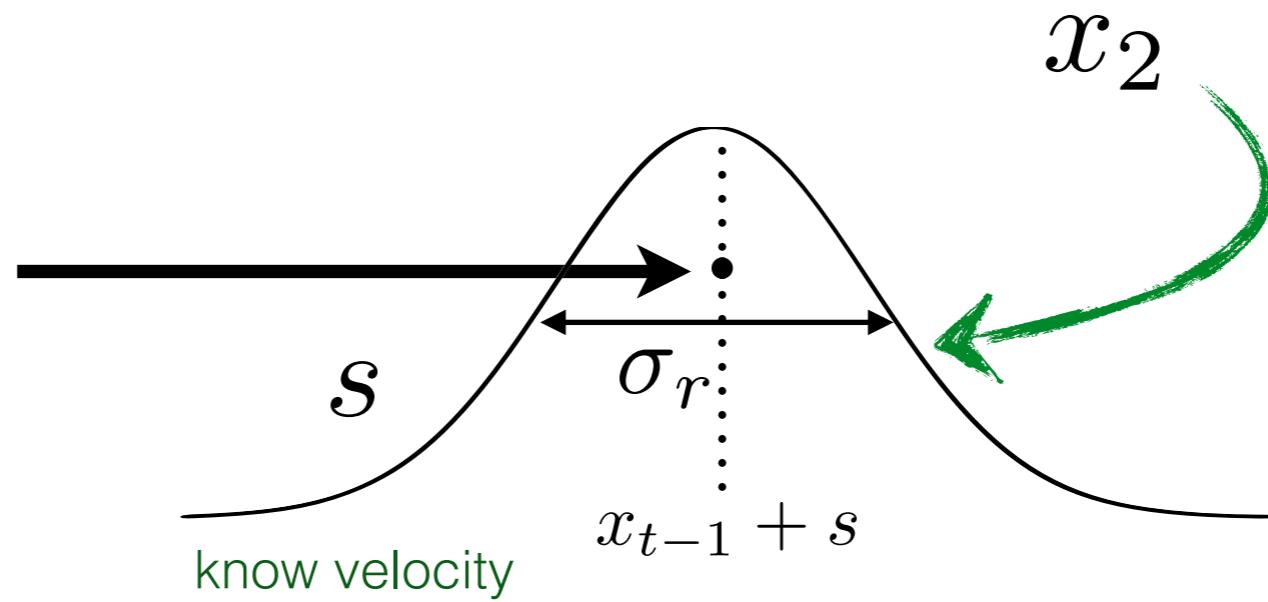
mean

standard deviation

How can you visualize this distribution?



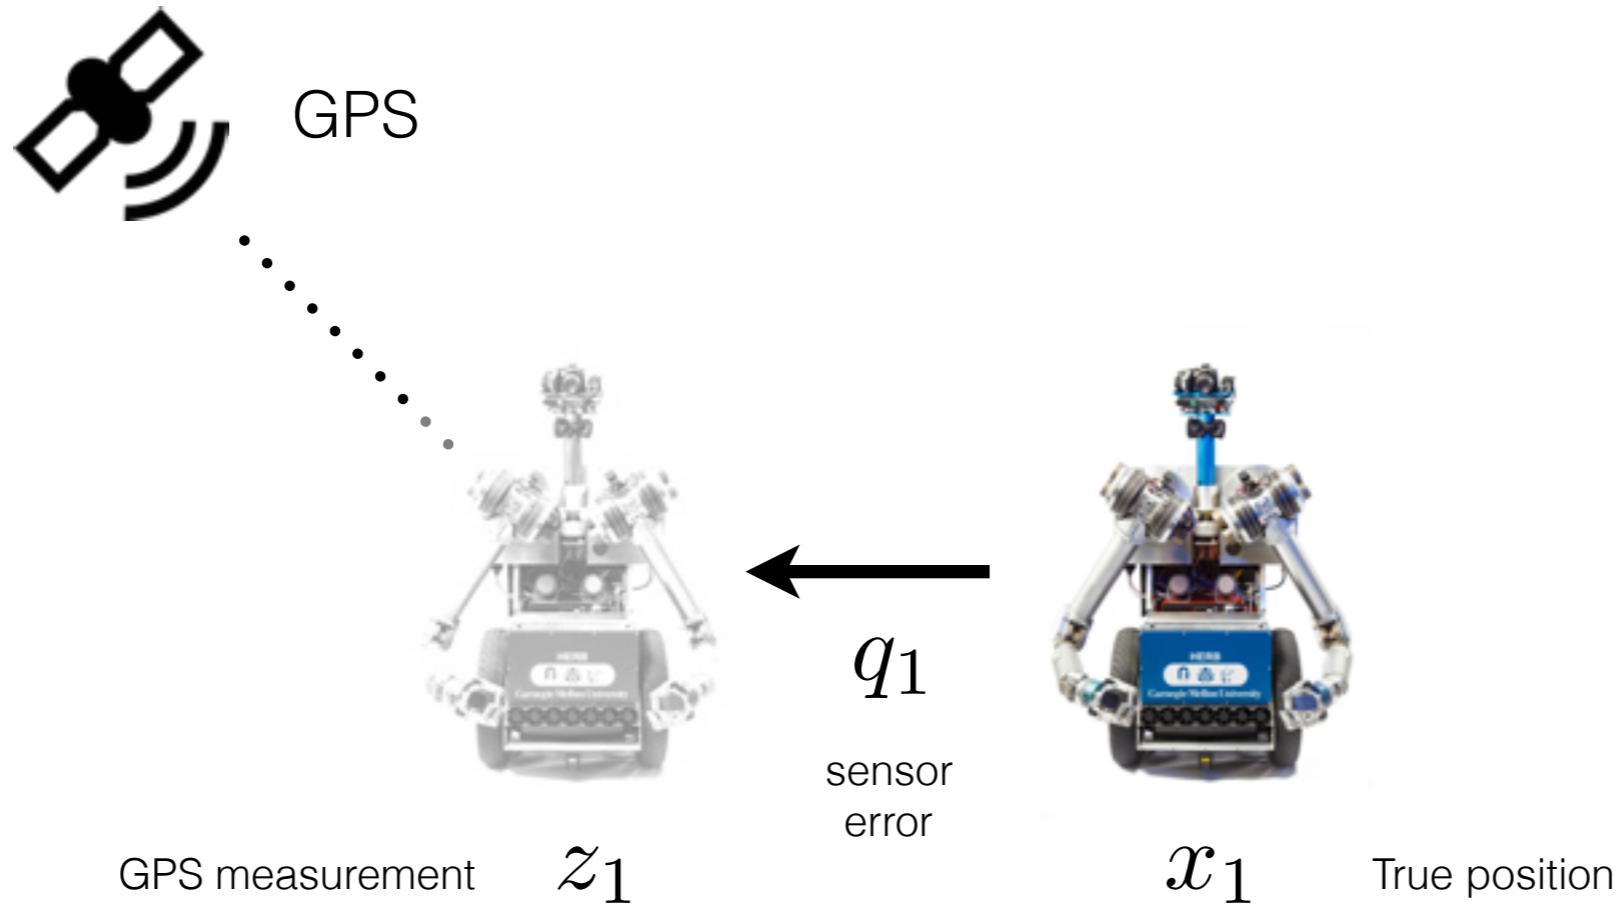
x_1



A linear Gaussian (continuous) transition model

$$P(x_t | x_{t-1}) = \mathcal{N}(x_t; x_{t-1} + s, \sigma_r)$$

Why don't we just use a table as before?

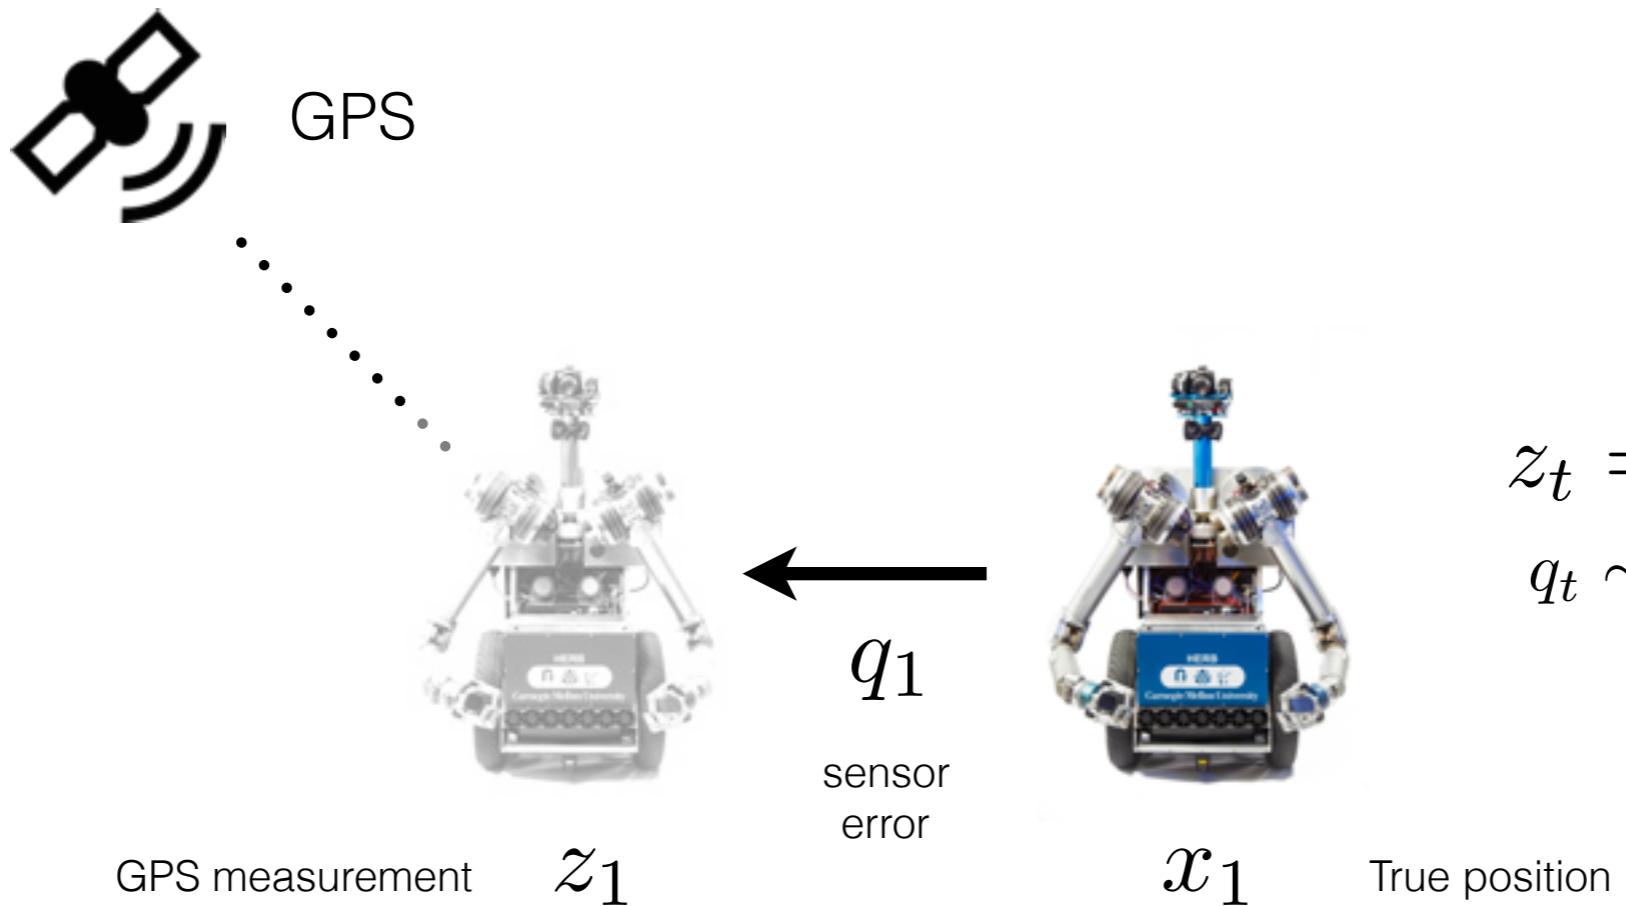


$$z_t = x_t + q_t$$

$$q_t \sim \mathcal{N}(0, \sigma_Q)$$

sampled from a Gaussian

Observation (measurement) model



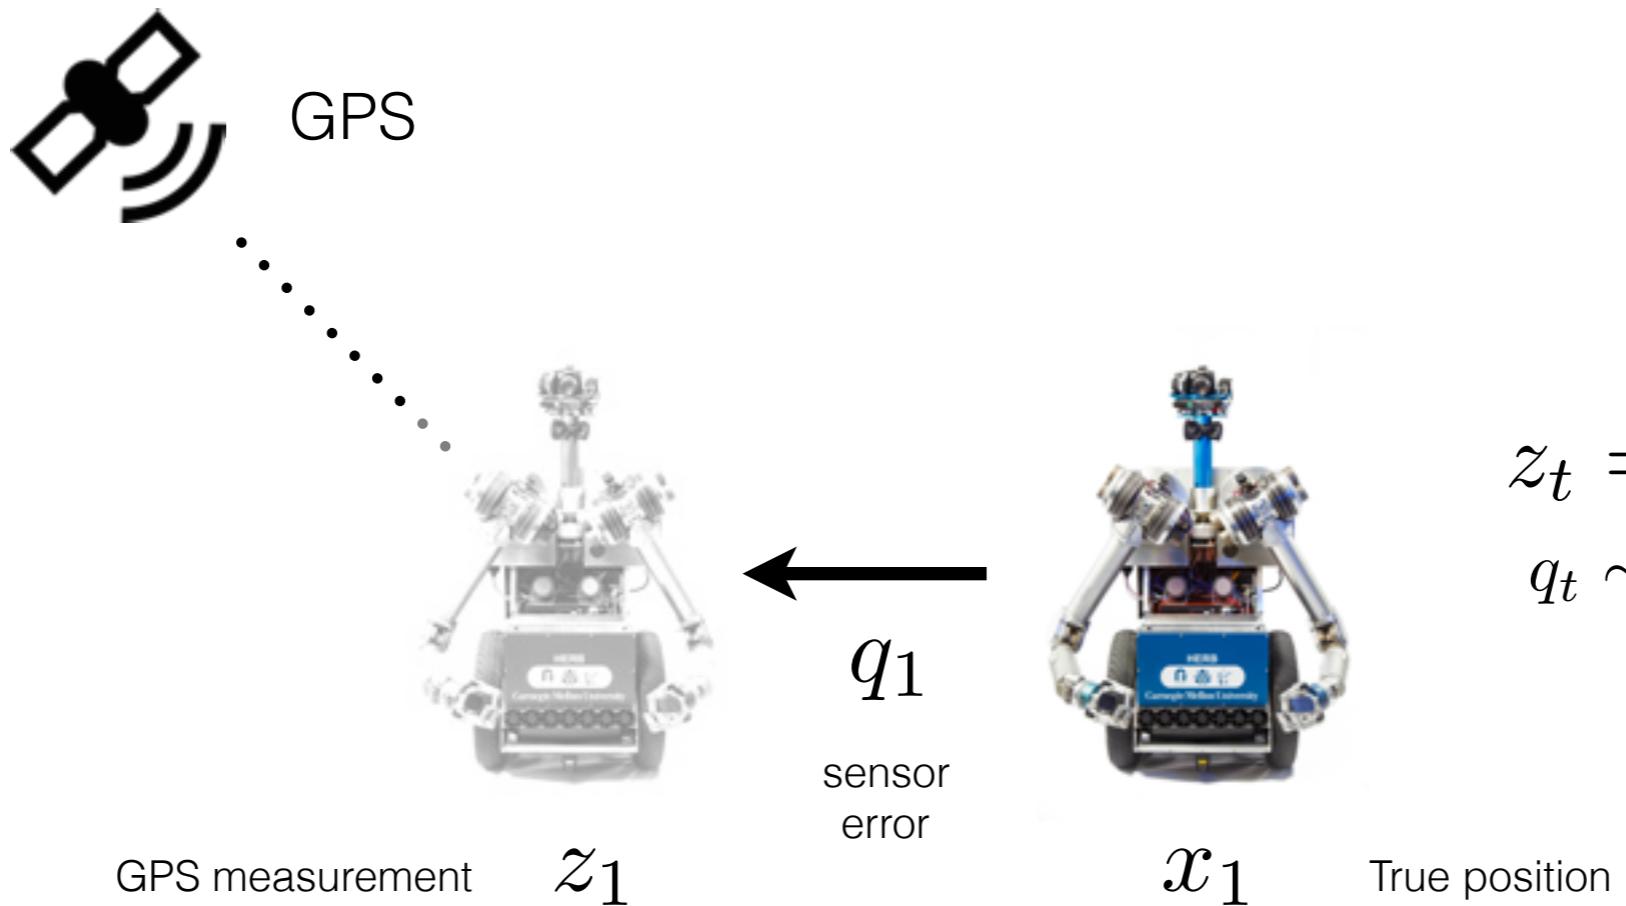
$$z_t = x_t + q_t$$

$$q_t \sim \mathcal{N}(0, \sigma_Q)$$

How do you represent the observation (measurement) model?

$$P(e|x)$$

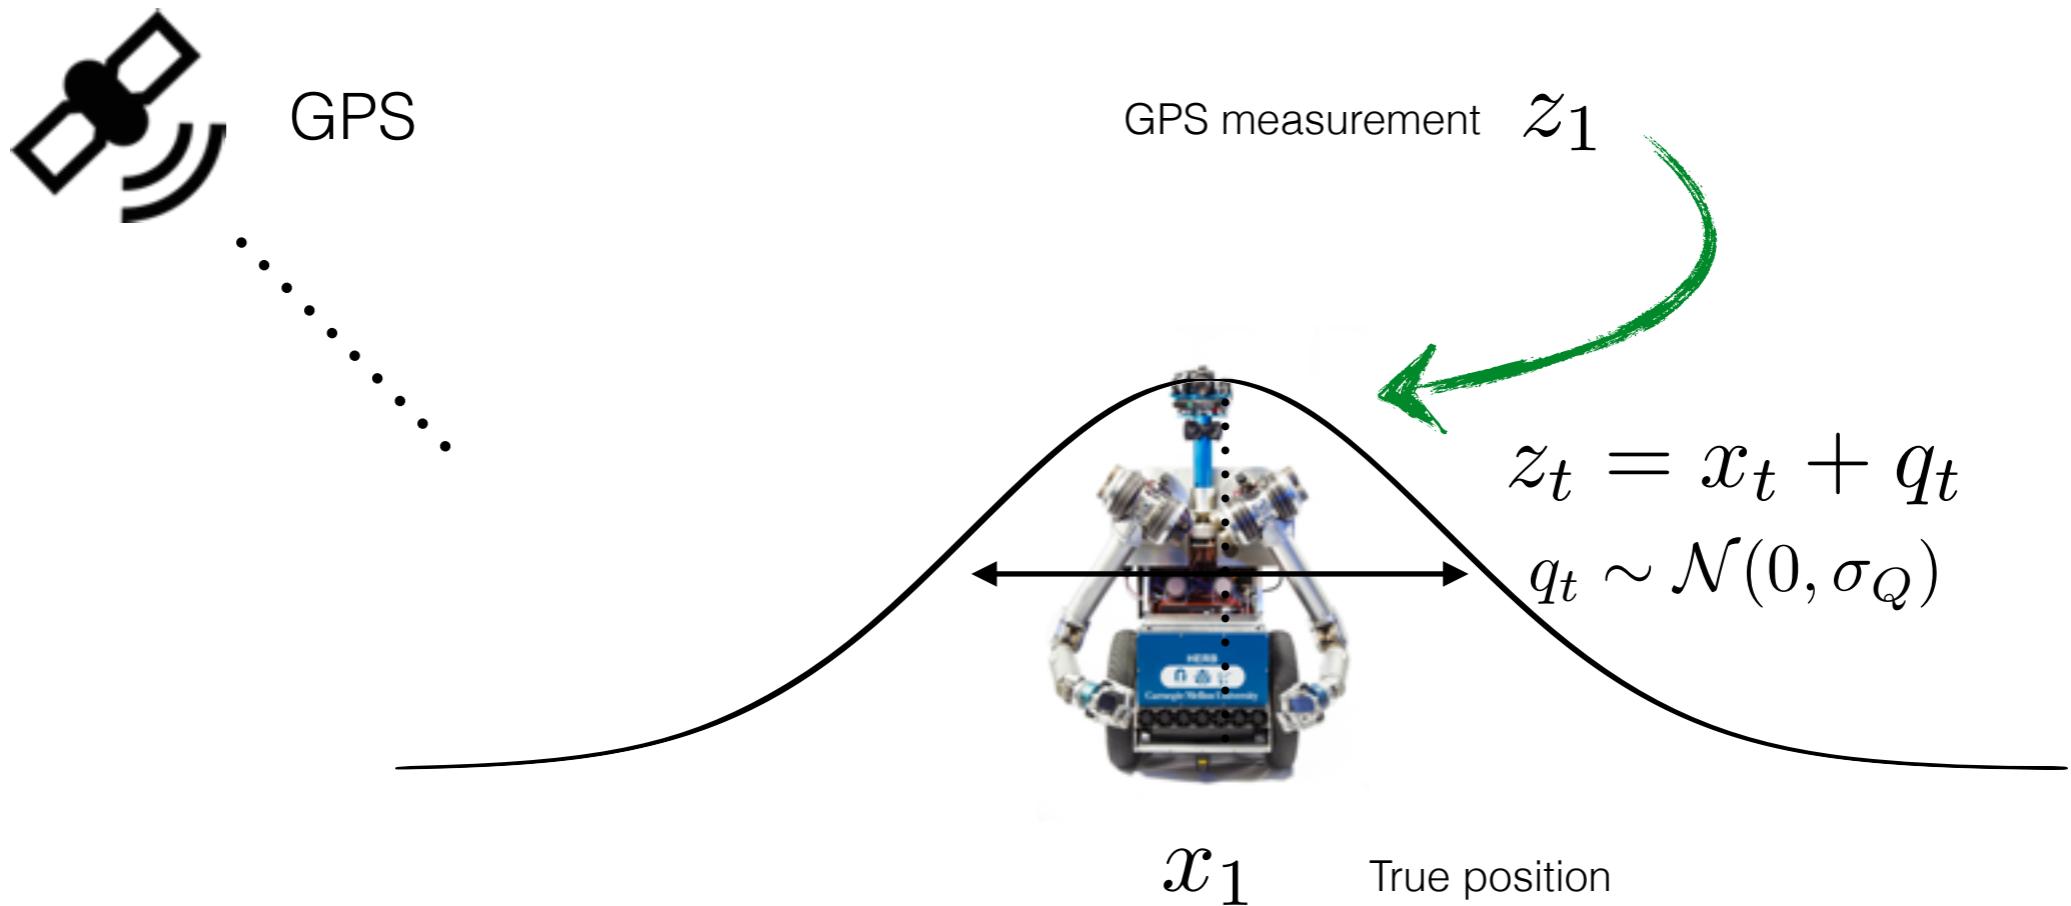
e represents z



How do you represent the observation (measurement) model?

Also a linear Gaussian model

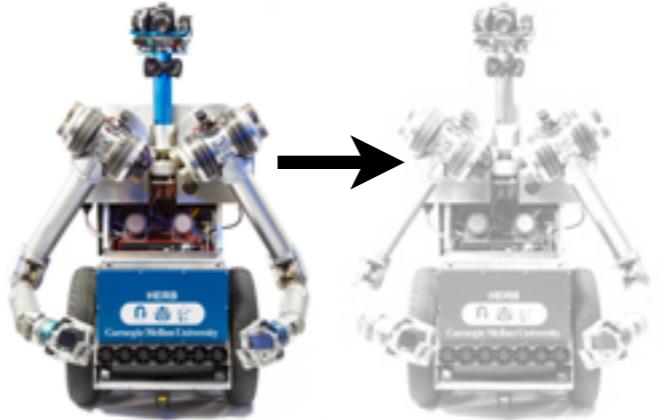
$$P(z_t | x_t) = \mathcal{N}(z_t; x_t, \sigma_Q)$$



How do you represent the observation (measurement) model?

Also a linear Gaussian model

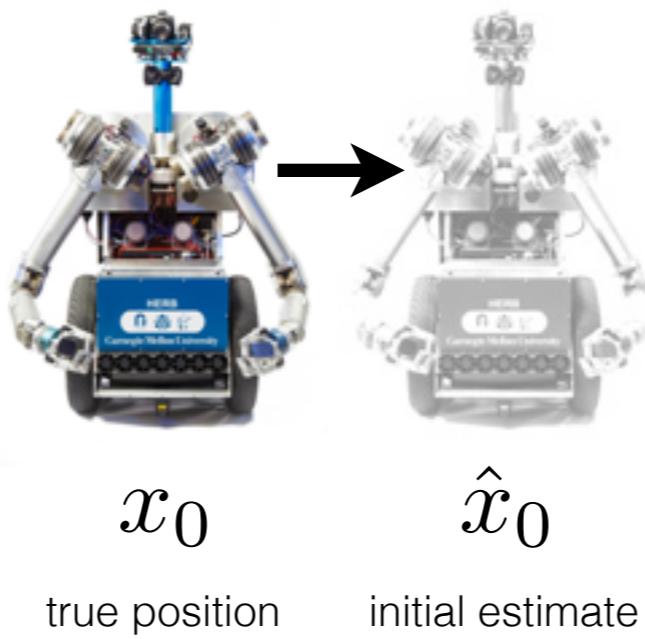
$$P(z_t | x_t) = \mathcal{N}(z_t; x_t, \sigma_Q)$$



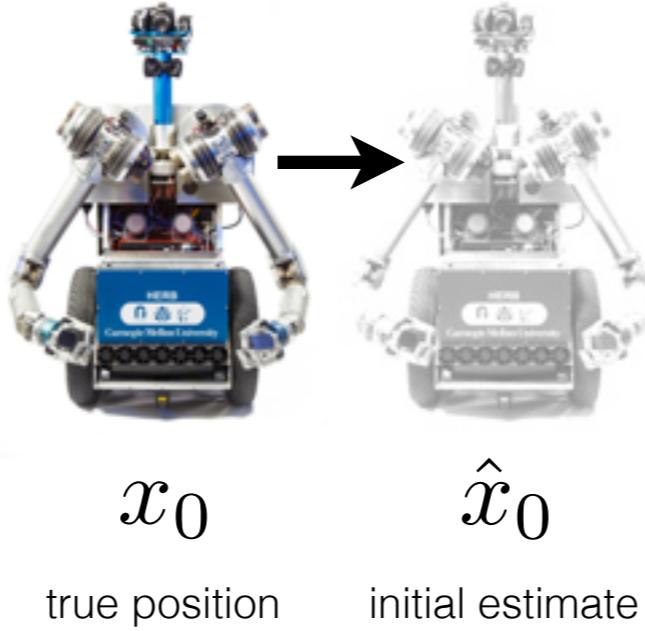
x_0 \hat{x}_0
true position initial estimate

initial estimate uncertainty σ_0

Prior (initial) State



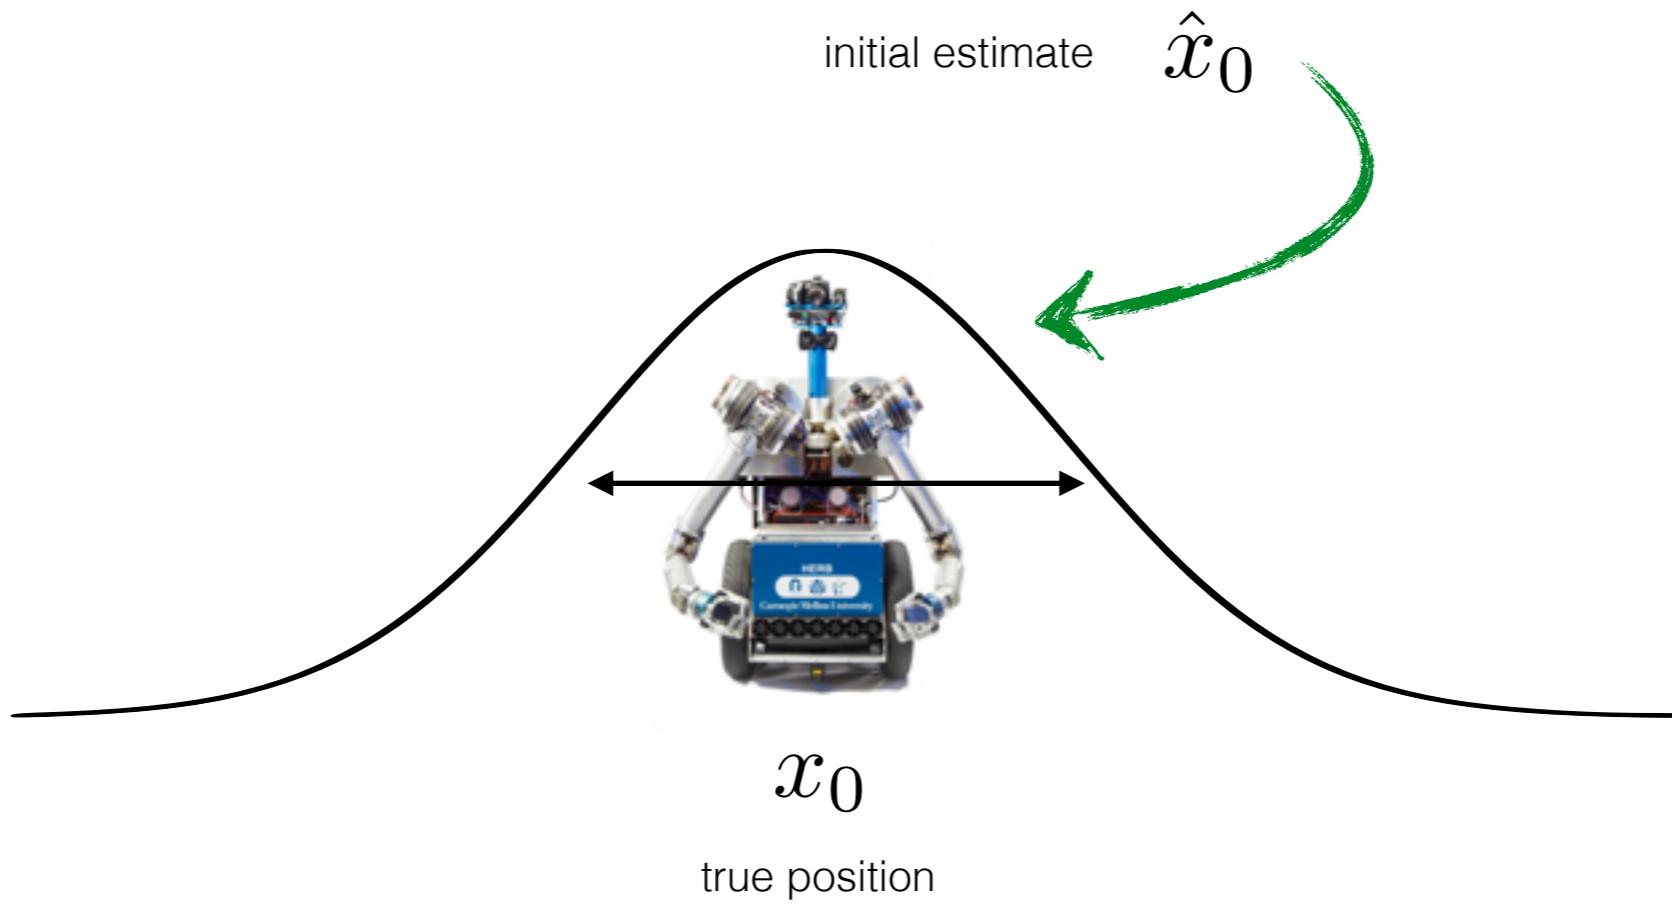
How do you represent the prior state probability?



How do you represent the prior state probability?

Also a linear Gaussian model!

$$P(\hat{x}_0) = \mathcal{N}(\hat{x}_0; x_0, \sigma_0)$$



How do you represent the prior state probability?

Also a linear Gaussian model

$$P(\hat{x}_0) = \mathcal{N}(\hat{x}_0; x_0, \sigma_0)$$

Inference

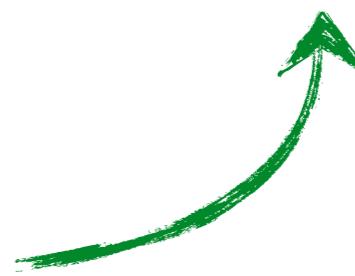
So how do you do temporal filtering with the KL?

Recall: the first step of filtering was the ‘prediction step’

$$P(\mathbf{X}_{t+1}|\mathbf{e}_{1:t+1}) \propto P(\mathbf{e}_{t+1}|\mathbf{X}_{t+1}) \int_{\mathbf{x}_t} P(\mathbf{X}_{t+1}|\mathbf{x}_t) P(\mathbf{x}_t|\mathbf{e}_{1:t}) d\mathbf{x}_t$$

prediction step
motion model belief

compute this!
It's just another Gaussian



need to compute the ‘prediction’ mean and variance...

Prediction

(Using the motion model)

How would you predict \hat{x}_1 given \hat{x}_0 ?

using this ‘cap’ notation to
denote ‘estimate’

$$\hat{x}_1 = \hat{x}_0 + s \quad (\text{This is the mean})$$

$$\sigma_1^2 = \sigma_0^2 + \sigma_r^2 \quad (\text{This is the variance})$$

prediction step

$$P(\mathbf{X}_{t+1} | \mathbf{e}_{1:t+1}) \propto P(\mathbf{e}_{t+1} | \mathbf{X}_{t+1}) \int_{\mathbf{x}_t} P(\mathbf{X}_{t+1} | \mathbf{x}_t) P(\mathbf{x}_t | \mathbf{e}_{1:t}) d\mathbf{x}_t$$

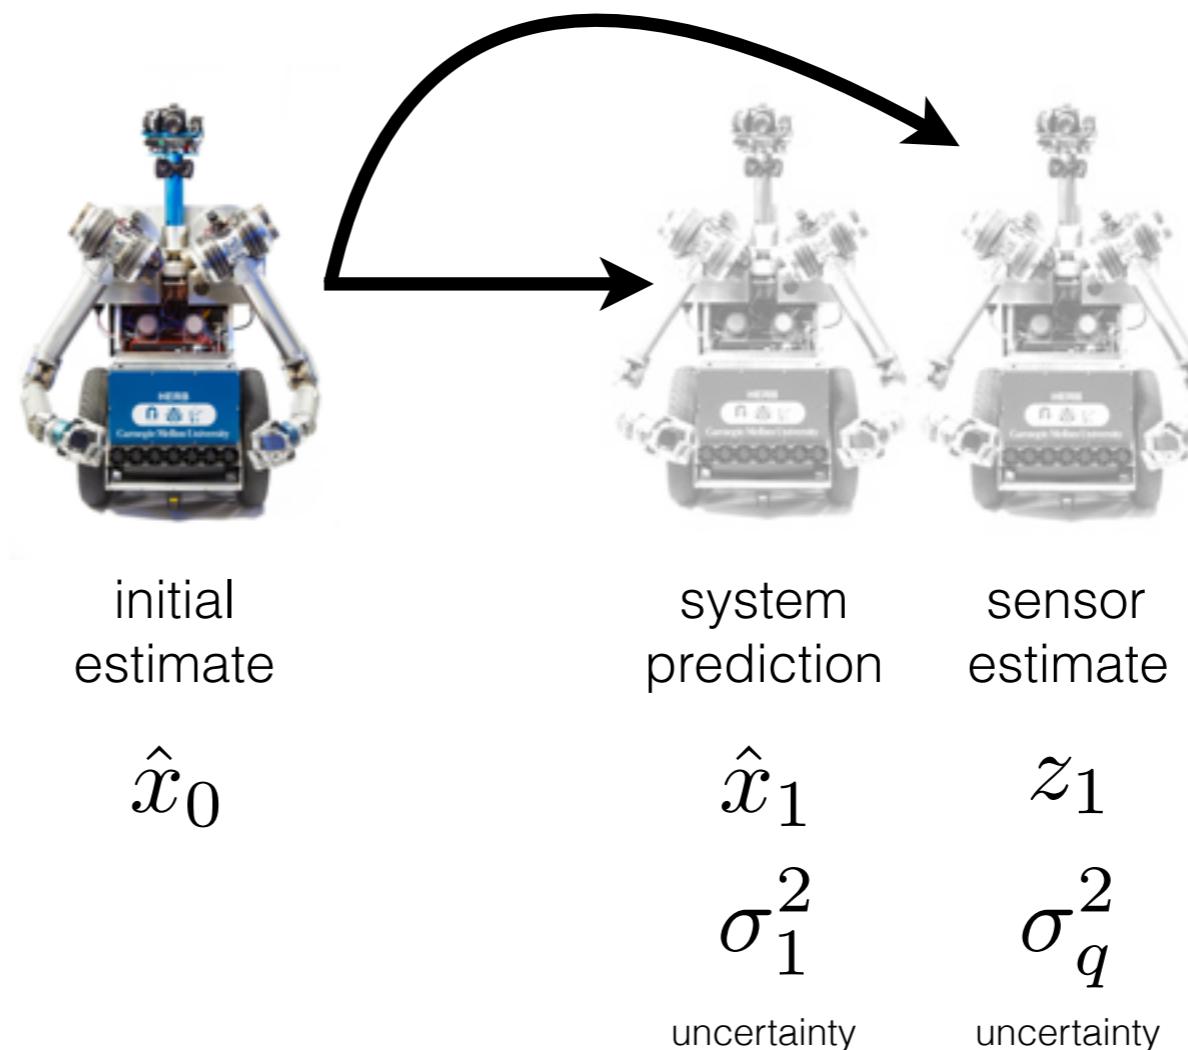
motion model
belief

the second step after prediction is ...

... update step!

compute this ✓
(using results of the prediction step)

In the **update step**, the **sensor measurement corrects** the system **prediction**



Which estimate is correct? Is there a way to know?

Is there a way to merge this information?

Intuitively, the smaller variance mean less uncertainty.



system
prediction σ_1^2



sensor
estimate σ_q^2

So we want a weighted state estimate correction

something
like this...

$$\hat{x}_1^+ = \frac{\sigma_q^2}{\sigma_1^2 + \sigma_q^2} \hat{x}_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_q^2} z_1$$

This happens naturally in the Bayesian filtering (with Gaussians) framework!

Recall the filtering equation:

What is the product of two Gaussians?

Recall ...

When we multiply the prediction (Gaussian) with the observation model (Gaussian) we get ...

... a product of two Gaussians

$$\mu = \frac{\mu_1 \sigma_2^2 + \mu_2 \sigma_1^2}{\sigma_2^2 + \sigma_1^2} \quad \sigma = \frac{\sigma_1^2 \sigma_2^2}{\sigma_1^2 + \sigma_2^2}$$

applied to the filtering equation...

$$P(\mathbf{X}_{t+1} | \mathbf{e}_{1:t+1}) \propto P(\mathbf{e}_{t+1} | \mathbf{X}_{t+1}) \int_{\mathbf{x}_t} P(\mathbf{X}_{t+1} | \mathbf{x}_t) P(\mathbf{x}_t | \mathbf{e}_{1:t}) d\mathbf{x}_t$$

mean: z_1

variance: σ_q

mean: \hat{x}_1

variance: σ_1

new mean:

$$\hat{x}_1^+ = \frac{\hat{x}_1 \sigma_q^2 + z_1 \sigma_1^2}{\sigma_q^2 + \sigma_1^2}$$

new variance:

$$\hat{\sigma}_1^{2+} = \frac{\sigma_q^2 \sigma_1^2}{\sigma_q^2 + \sigma_1^2}$$

'plus' sign means post
'update' estimate



system
prediction σ_1^2



sensor
estimate σ_q^2

With a little algebra...

$$\hat{x}_1^+ = \frac{\hat{x}_1 \sigma_q^2 + z_1 \sigma_1^2}{\sigma_q^2 + \sigma_1^2} = \hat{x}_1 \frac{\sigma_q^2}{\sigma_q^2 + \sigma_1^2} + z_1 \frac{\sigma_1^2}{\sigma_q^2 + \sigma_1^2}$$

We get a weighted state estimate correction!

Kalman gain notation

With a little algebra...

$$\hat{x}_1^+ = \hat{x}_1 + \frac{\sigma_1^2}{\sigma_q^2 + \sigma_1^2} (z_1 - \hat{x}_1) = \hat{x}_1 + K(z_1 - \hat{x}_1)$$


‘Kalman gain’ ‘Innovation’

With a little algebra...

$$\sigma_1^+ = \frac{\sigma_1^2 \sigma_q^2}{\sigma_1^2 + \sigma_q^2} = \left(1 - \frac{\sigma_1^2}{\sigma_1^2 + \sigma_q^2} \right) \sigma_1^2 = (1 - K) \sigma_1^2$$

Summary (1D Kalman Filtering)

To solve this...

$$P(\mathbf{X}_{t+1} | \mathbf{e}_{1:t+1}) \propto P(\mathbf{e}_{t+1} | \mathbf{X}_{t+1}) \int_{\mathbf{x}_t} P(\mathbf{X}_{t+1} | \mathbf{x}_t) P(\mathbf{x}_t | \mathbf{e}_{1:t}) d\mathbf{x}_t$$

Compute this...

$$\hat{x}_1^+ = \hat{x}_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_q^2} (z_1 - \hat{x}_1) \quad \sigma_1^{2+} = \sigma_1^2 - \frac{\sigma_1^2}{\sigma_1^2 + \sigma_q^2} \sigma_1^2$$

$$K = \frac{\sigma_1^2}{\sigma_1^2 + \sigma_q^2}$$

‘Kalman gain’

$$\hat{x}_1^+ = \hat{x}_1 + K(z_1 - \hat{x}_1)$$

mean of the new Gaussian

$$\sigma_1^{2+} = \sigma_1^2 - K\sigma_1^2$$

variance of the new Gaussian

Simple 1D Implementation

```
[x p] = KF(x, v, z)
```

```
x = x + s;  
v = v + q;
```

```
K = v / (v + r);
```

```
x = x + K * (z - x);  
p = v - K * v;
```

Just 5 lines of code!

or just 2 lines

$$\begin{aligned}[x \ P] &= KF(x, v, z) \\ x &= (x+s) + (v+q) / ((v+q)+r) * (z - (x+s)) ; \\ p &= (v+q) - (v+q) / ((v+q)+r) * v ;\end{aligned}$$

Bare computations (algorithm) of Bayesian filtering:

$\text{KalmanFilter}(\mu_{t-1}, \Sigma_{t-1}, u_t, z_t)$

$\bar{\mu}_t = A_t \mu_{t-1} + B u_t$

prediction mean

motion control 'old' mean

$\bar{\Sigma}_t = A_t \Sigma_{t-1} A_t^\top + R$

prediction covariance

'old' covariance Gaussian noise

$K_t = \bar{\Sigma}_t C_t^\top (C_t \bar{\Sigma}_t C_t^\top + Q_t)^{-1}$

Gain

$\mu_t = \bar{\mu}_t + K_t (z_t - C_t \bar{\mu}_t)$

Prediction

$\Sigma_t = (I - K_t C_t) \bar{\Sigma}_t$

Update

Simple Multi-dimensional Implementation (also 5 lines of code!)

```
[x P] = KF(x, P, z)
```

```
x = A*x;  
P = A*P*A' + Q;
```

```
K = P*C' / (C*P*C' + R);
```

```
x = x + K*(z - C*x);  
P = (eye(size(K, 1)) - K*C)*P;
```

2D Example



x

y

state

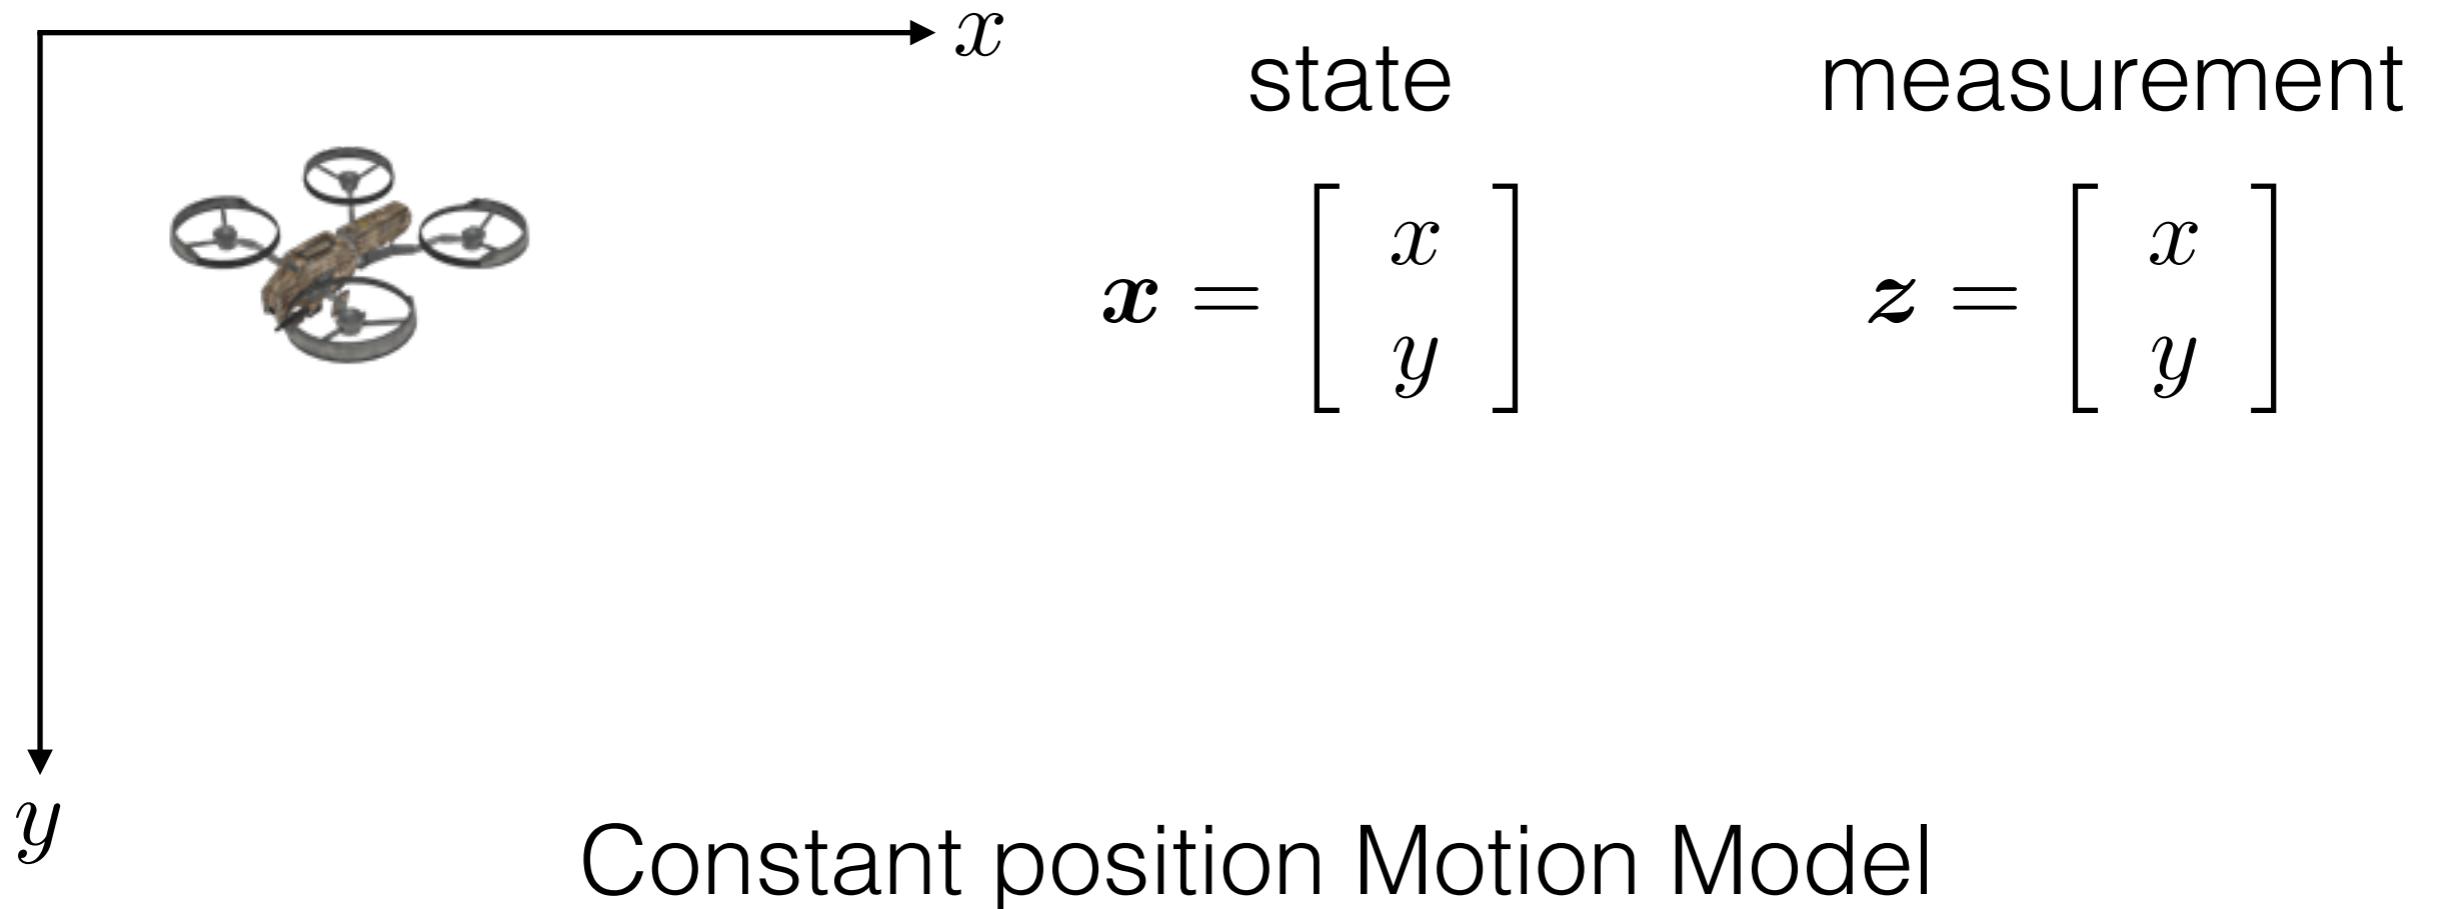
$$\boldsymbol{x} = \begin{bmatrix} x \\ y \end{bmatrix}$$

measurement

$$\boldsymbol{z} = \begin{bmatrix} x \\ y \end{bmatrix}$$

Constant position Motion Model

$$\boldsymbol{x}_t = A\boldsymbol{x}_{t-1} + B\boldsymbol{u}_t + \boldsymbol{\epsilon}_t$$



$$\boldsymbol{x}_t = A\boldsymbol{x}_{t-1} + B\boldsymbol{u}_t + \epsilon_t$$

Constant position

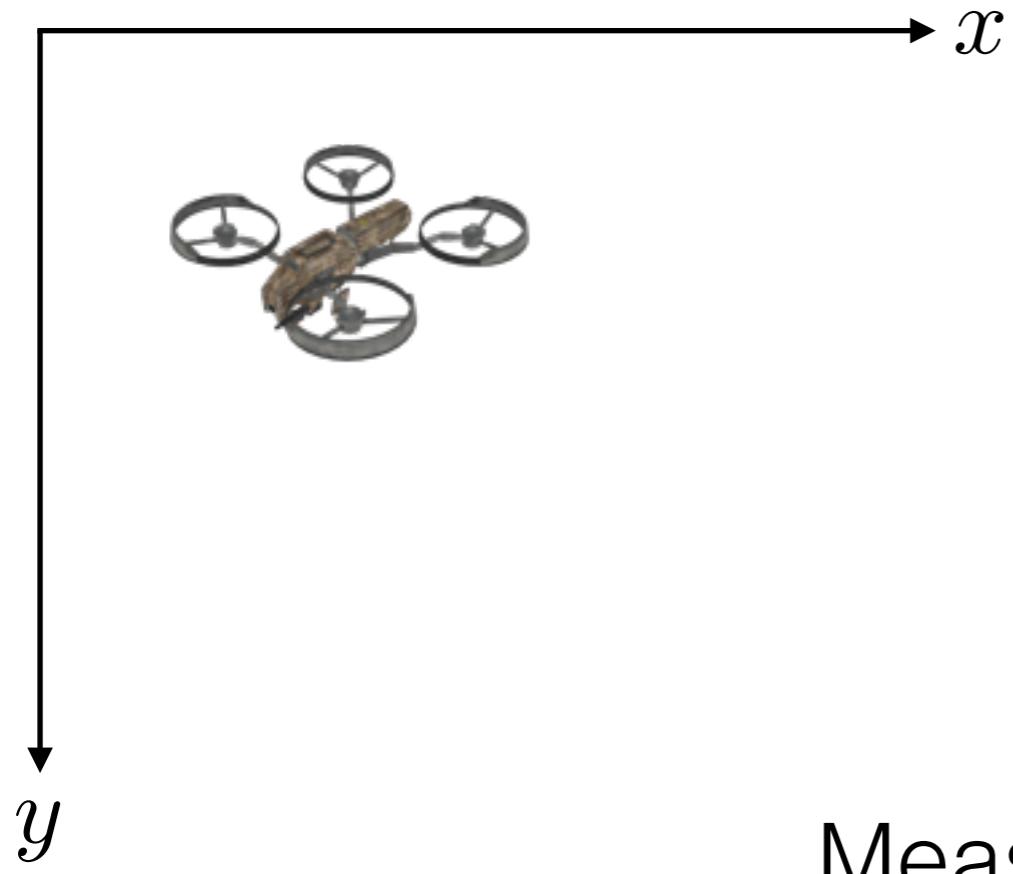
$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

$$B\boldsymbol{u} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

system noise

$$\epsilon_t \sim \mathcal{N}(\mathbf{0}, R)$$

$$R = \begin{bmatrix} \sigma_r^2 & 0 \\ 0 & \sigma_r^2 \end{bmatrix}$$



state

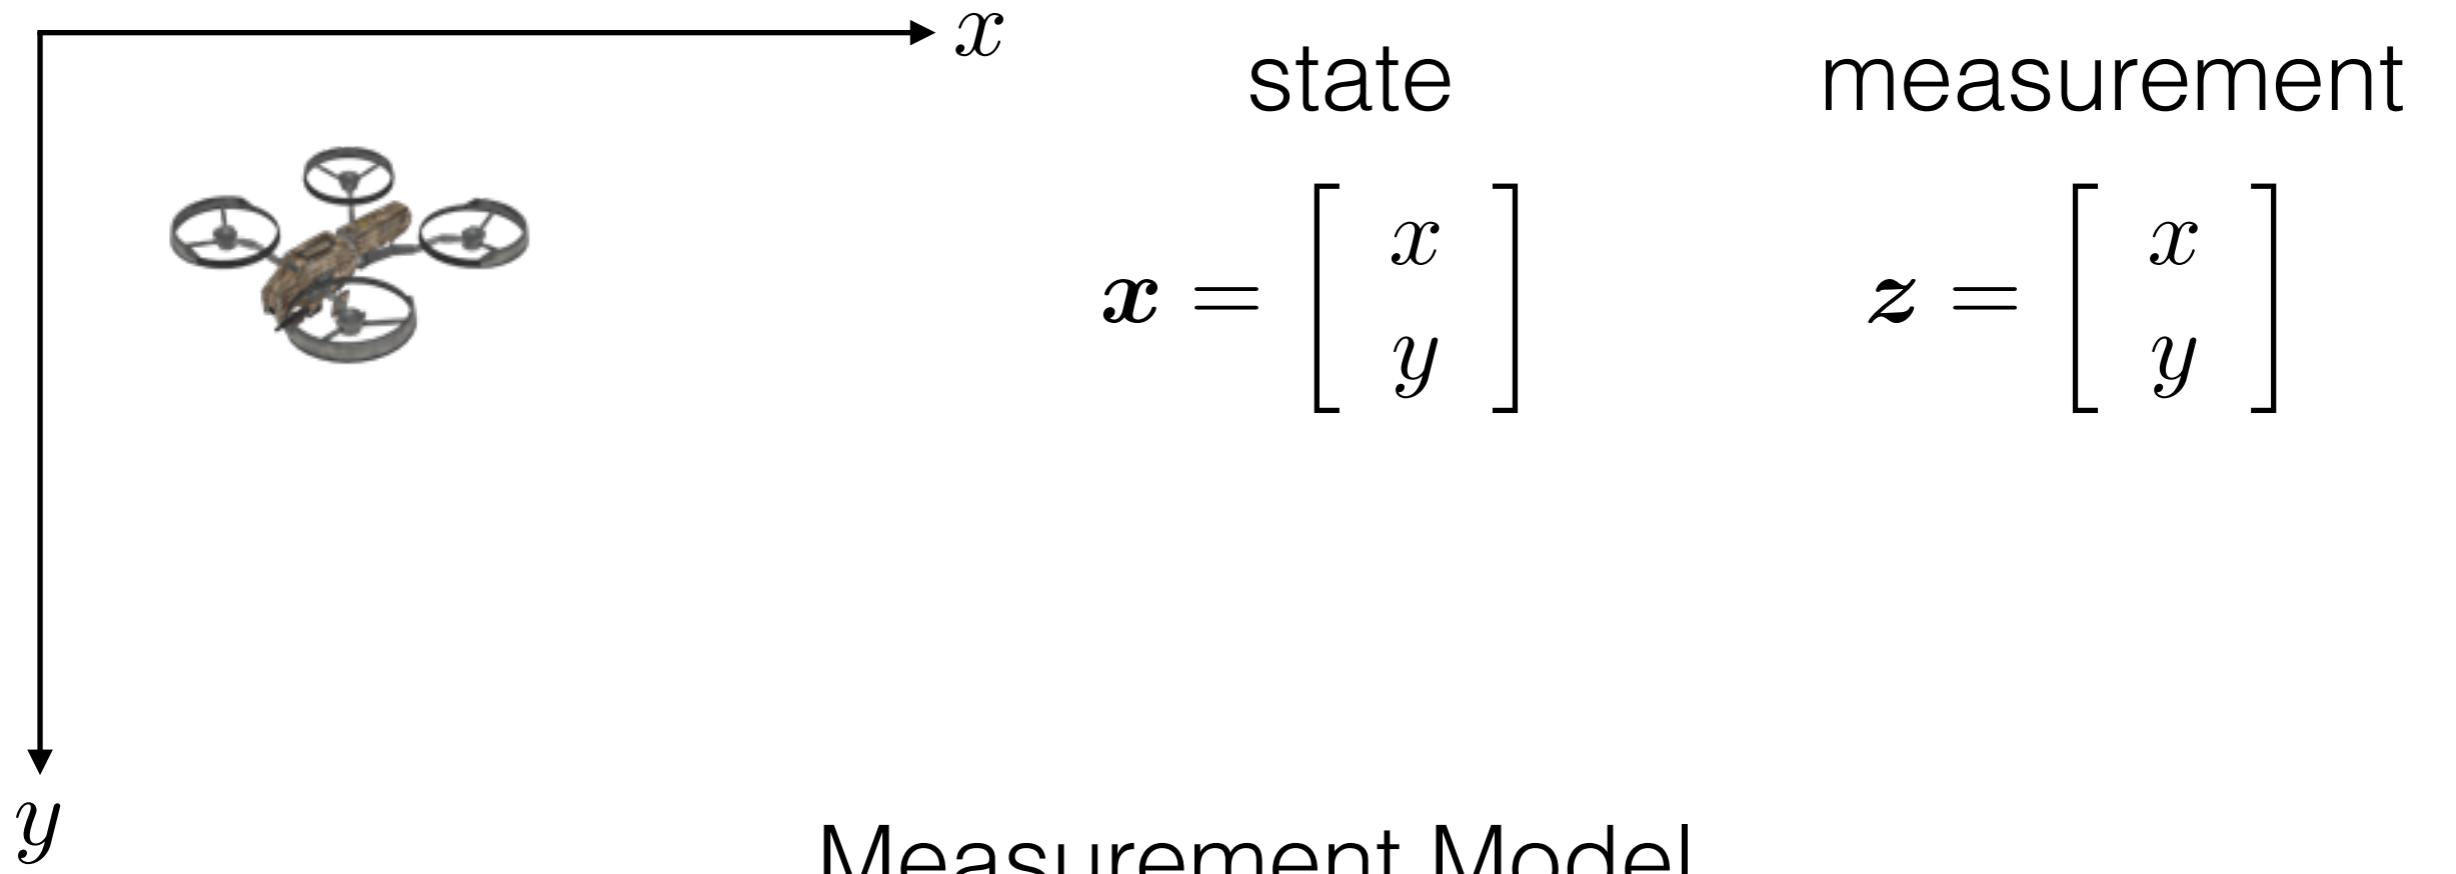
$$\boldsymbol{x} = \begin{bmatrix} x \\ y \end{bmatrix}$$

measurement

$$\boldsymbol{z} = \begin{bmatrix} x \\ y \end{bmatrix}$$

Measurement Model

$$\boldsymbol{z}_t = C_t \boldsymbol{x}_t + \boldsymbol{\delta}_t$$



Measurement Model

$$\boldsymbol{z}_t = C_t \boldsymbol{x}_t + \delta_t$$

zero-mean measurement noise

$$C = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \delta_t \sim \mathcal{N}(\mathbf{0}, Q) \quad Q = \begin{bmatrix} \sigma_q^2 & 0 \\ 0 & \sigma_q^2 \end{bmatrix}$$

Algorithm for the 2D object tracking example



$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

motion model

$$C = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

observation model

General Case

$$\bar{\mu}_t = A_t \mu_{t-1} + B u_t$$

$$\bar{\Sigma}_t = A_t \Sigma_{t-1} A_t^\top + R$$

$$K_t = \bar{\Sigma}_t C_t^\top (C_t \bar{\Sigma}_t C_t^\top + Q_t)^{-1}$$

$$\mu_t = \bar{\mu}_t + K_t(z_t - C_t \bar{\mu}_t)$$

$$\Sigma_t = (I - K_t C_t) \bar{\Sigma}_t$$

Constant position Model

$$\bar{x}_t = x_{t-1}$$

$$\bar{\Sigma}_t = \Sigma_{t-1} + R$$

$$K_t = \bar{\Sigma}_t (\bar{\Sigma}_t + Q)^{-1}$$

$$x_t = \bar{x}_t + K_t(z_t - \bar{x}_t)$$

$$\Sigma_t = (I - K_t)^2 \bar{\Sigma}_t$$

Just 4 lines of code

```
[x P] = KF_constPos(x, P, z)  
P = P + Q;  
K = P / (P + R);  
x = x + K * (z - x);  
P = (eye(size(K, 1)) - K) * P;
```

Where did the 5th line go?

General Case

$$\bar{\mu}_t = A_t \mu_{t-1} + Bu_t$$

$$\bar{\Sigma}_t = A_t \Sigma_{t-1} A_t^\top + R$$

$$K_t = \bar{\Sigma}_t C_t^\top (C_t \bar{\Sigma}_t C_t^\top + Q_t)^{-1}$$

$$\mu_t = \bar{\mu}_t + K_t(z_t - C_t \bar{\mu}_t)$$

$$\Sigma_t = (I - K_t C_t) \bar{\Sigma}_t$$

Constant position Model

$$\bar{x}_t = x_{t-1}$$

$$\bar{\Sigma}_t = \Sigma_{t-1} + R$$

$$K_t = \bar{\Sigma}_t (\bar{\Sigma}_t + Q)^{-1}$$

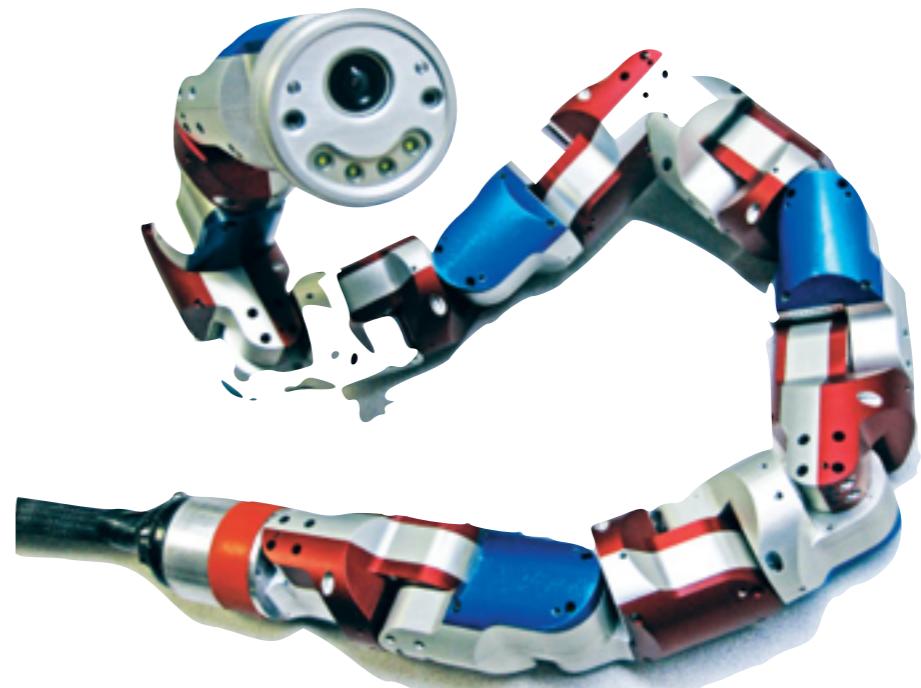
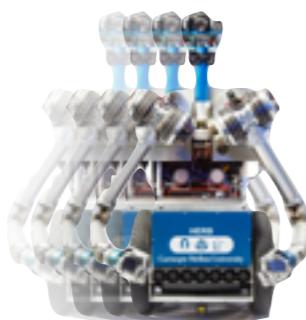
$$x_t = \bar{x}_t + K_t(z_t - \bar{x}_t)$$

$$\Sigma_t = (I - K_t) \bar{\Sigma}_t$$

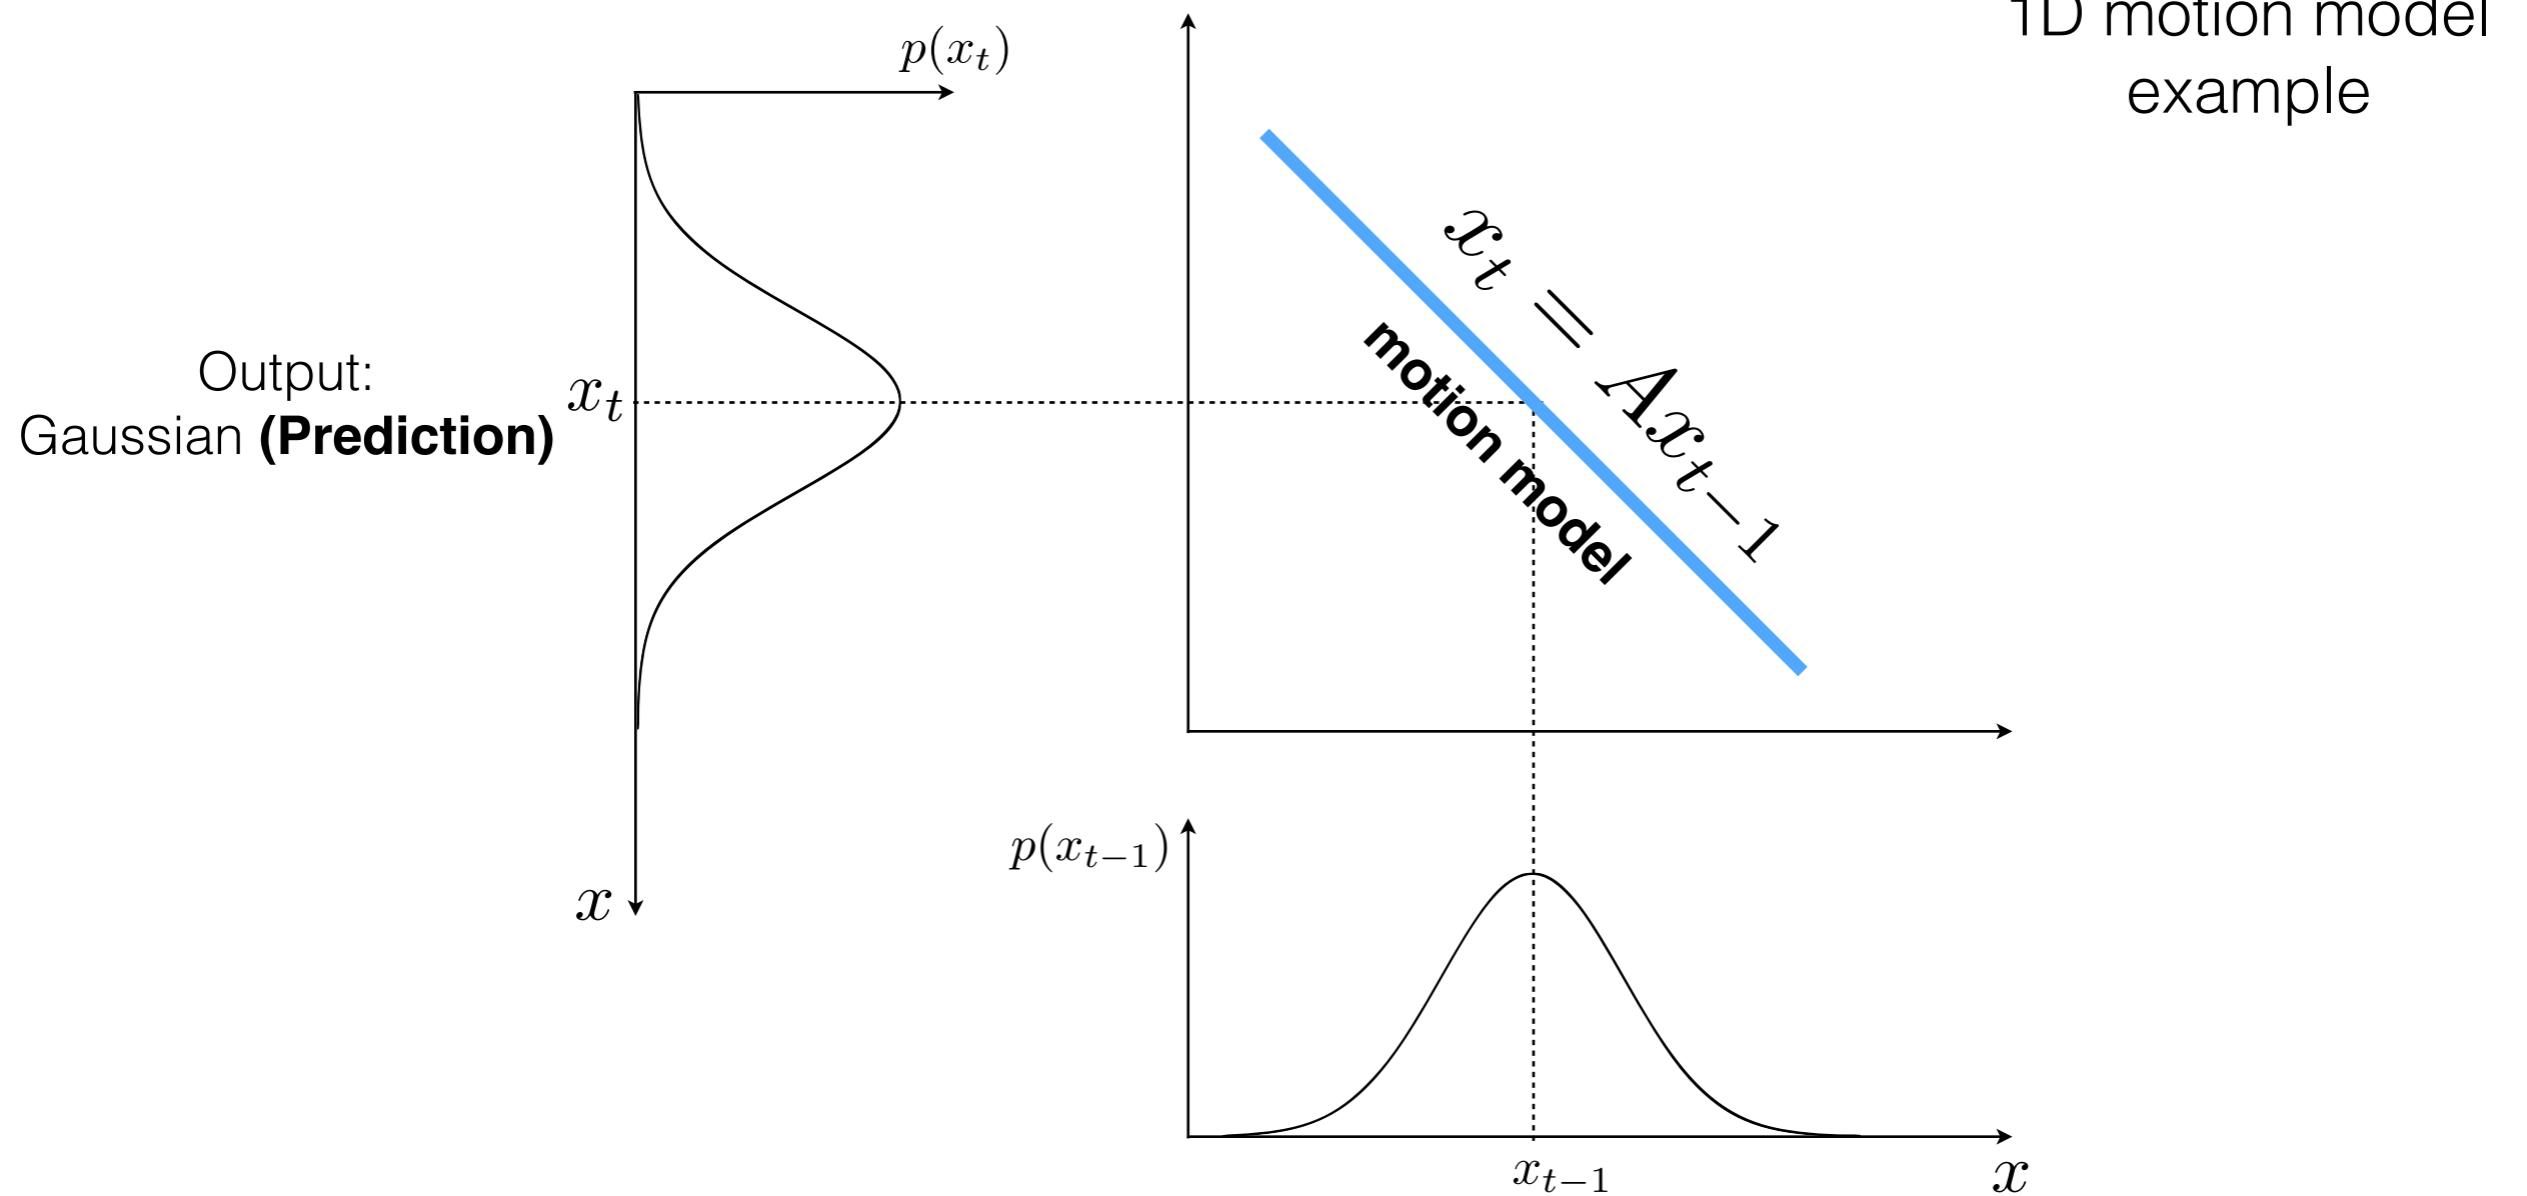
Motion model of the Kalman filter is linear

$$x_t = Ax_{t-1} + Bu_t + \epsilon_t$$

but motion is not always linear



Visualizing linear models

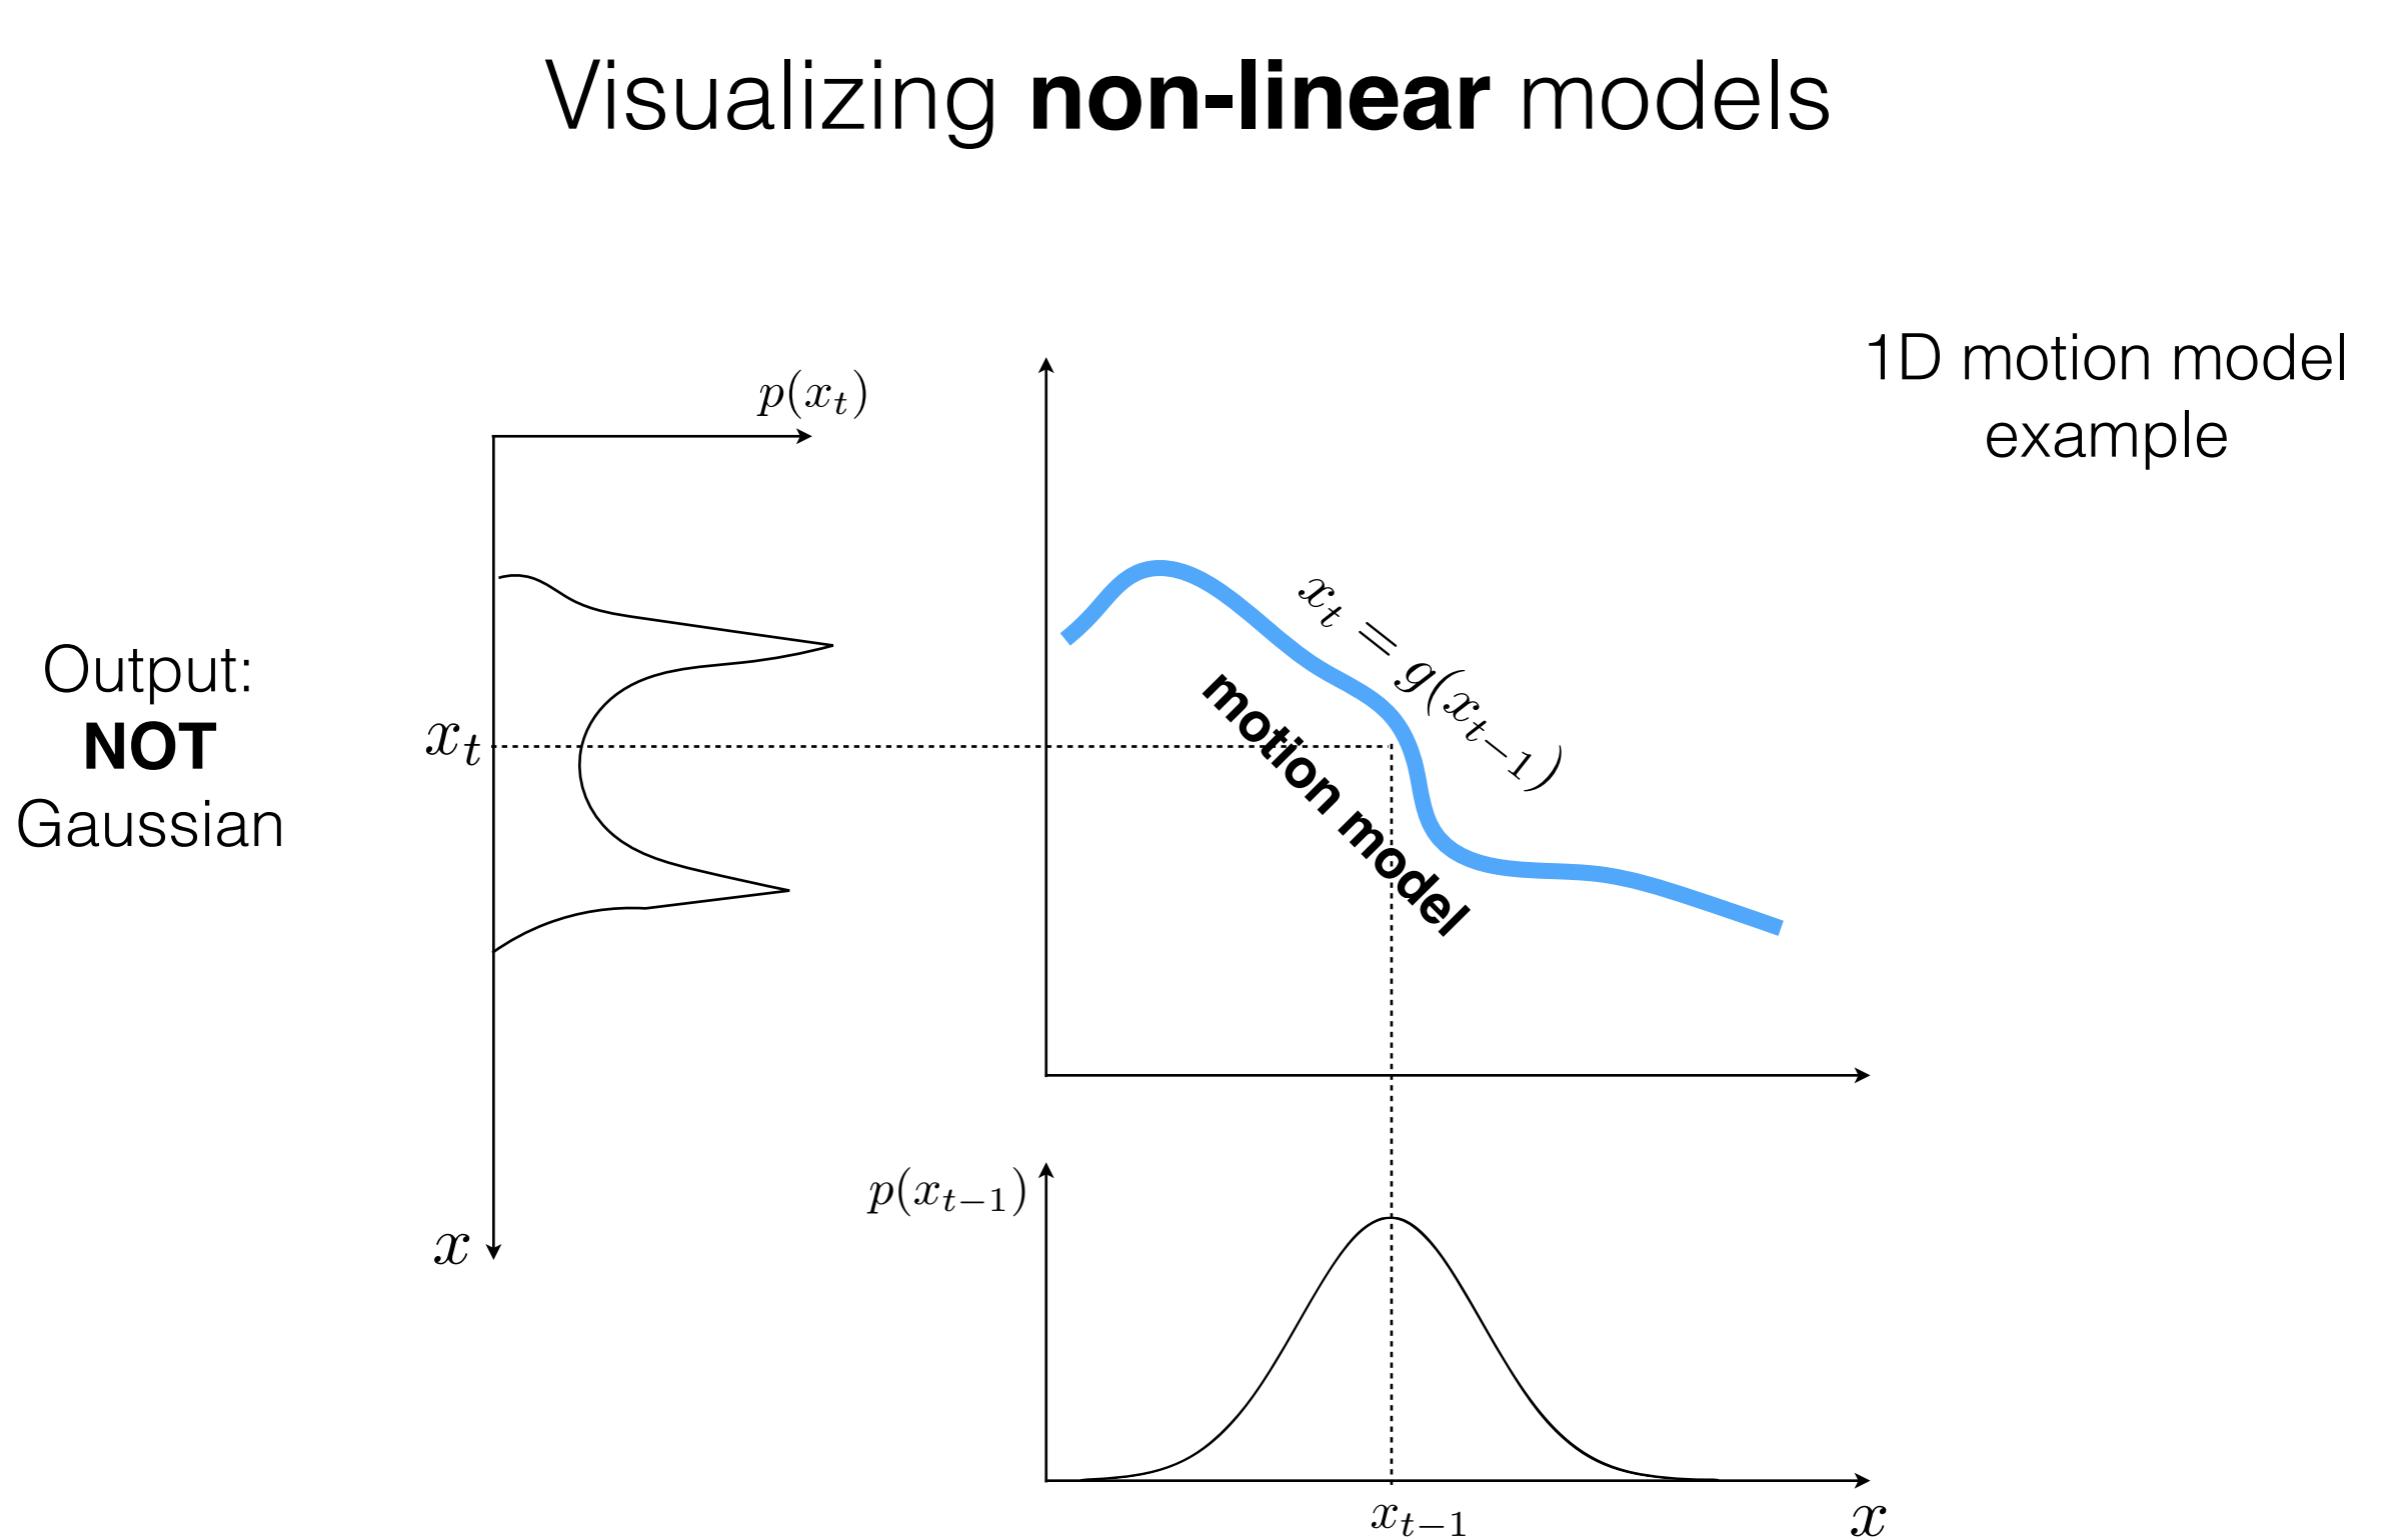


Can we use the Kalman
Filter?

(motion model and observation model are linear)

Input:
Gaussian (**Belief**)

Visualizing **non-linear** models

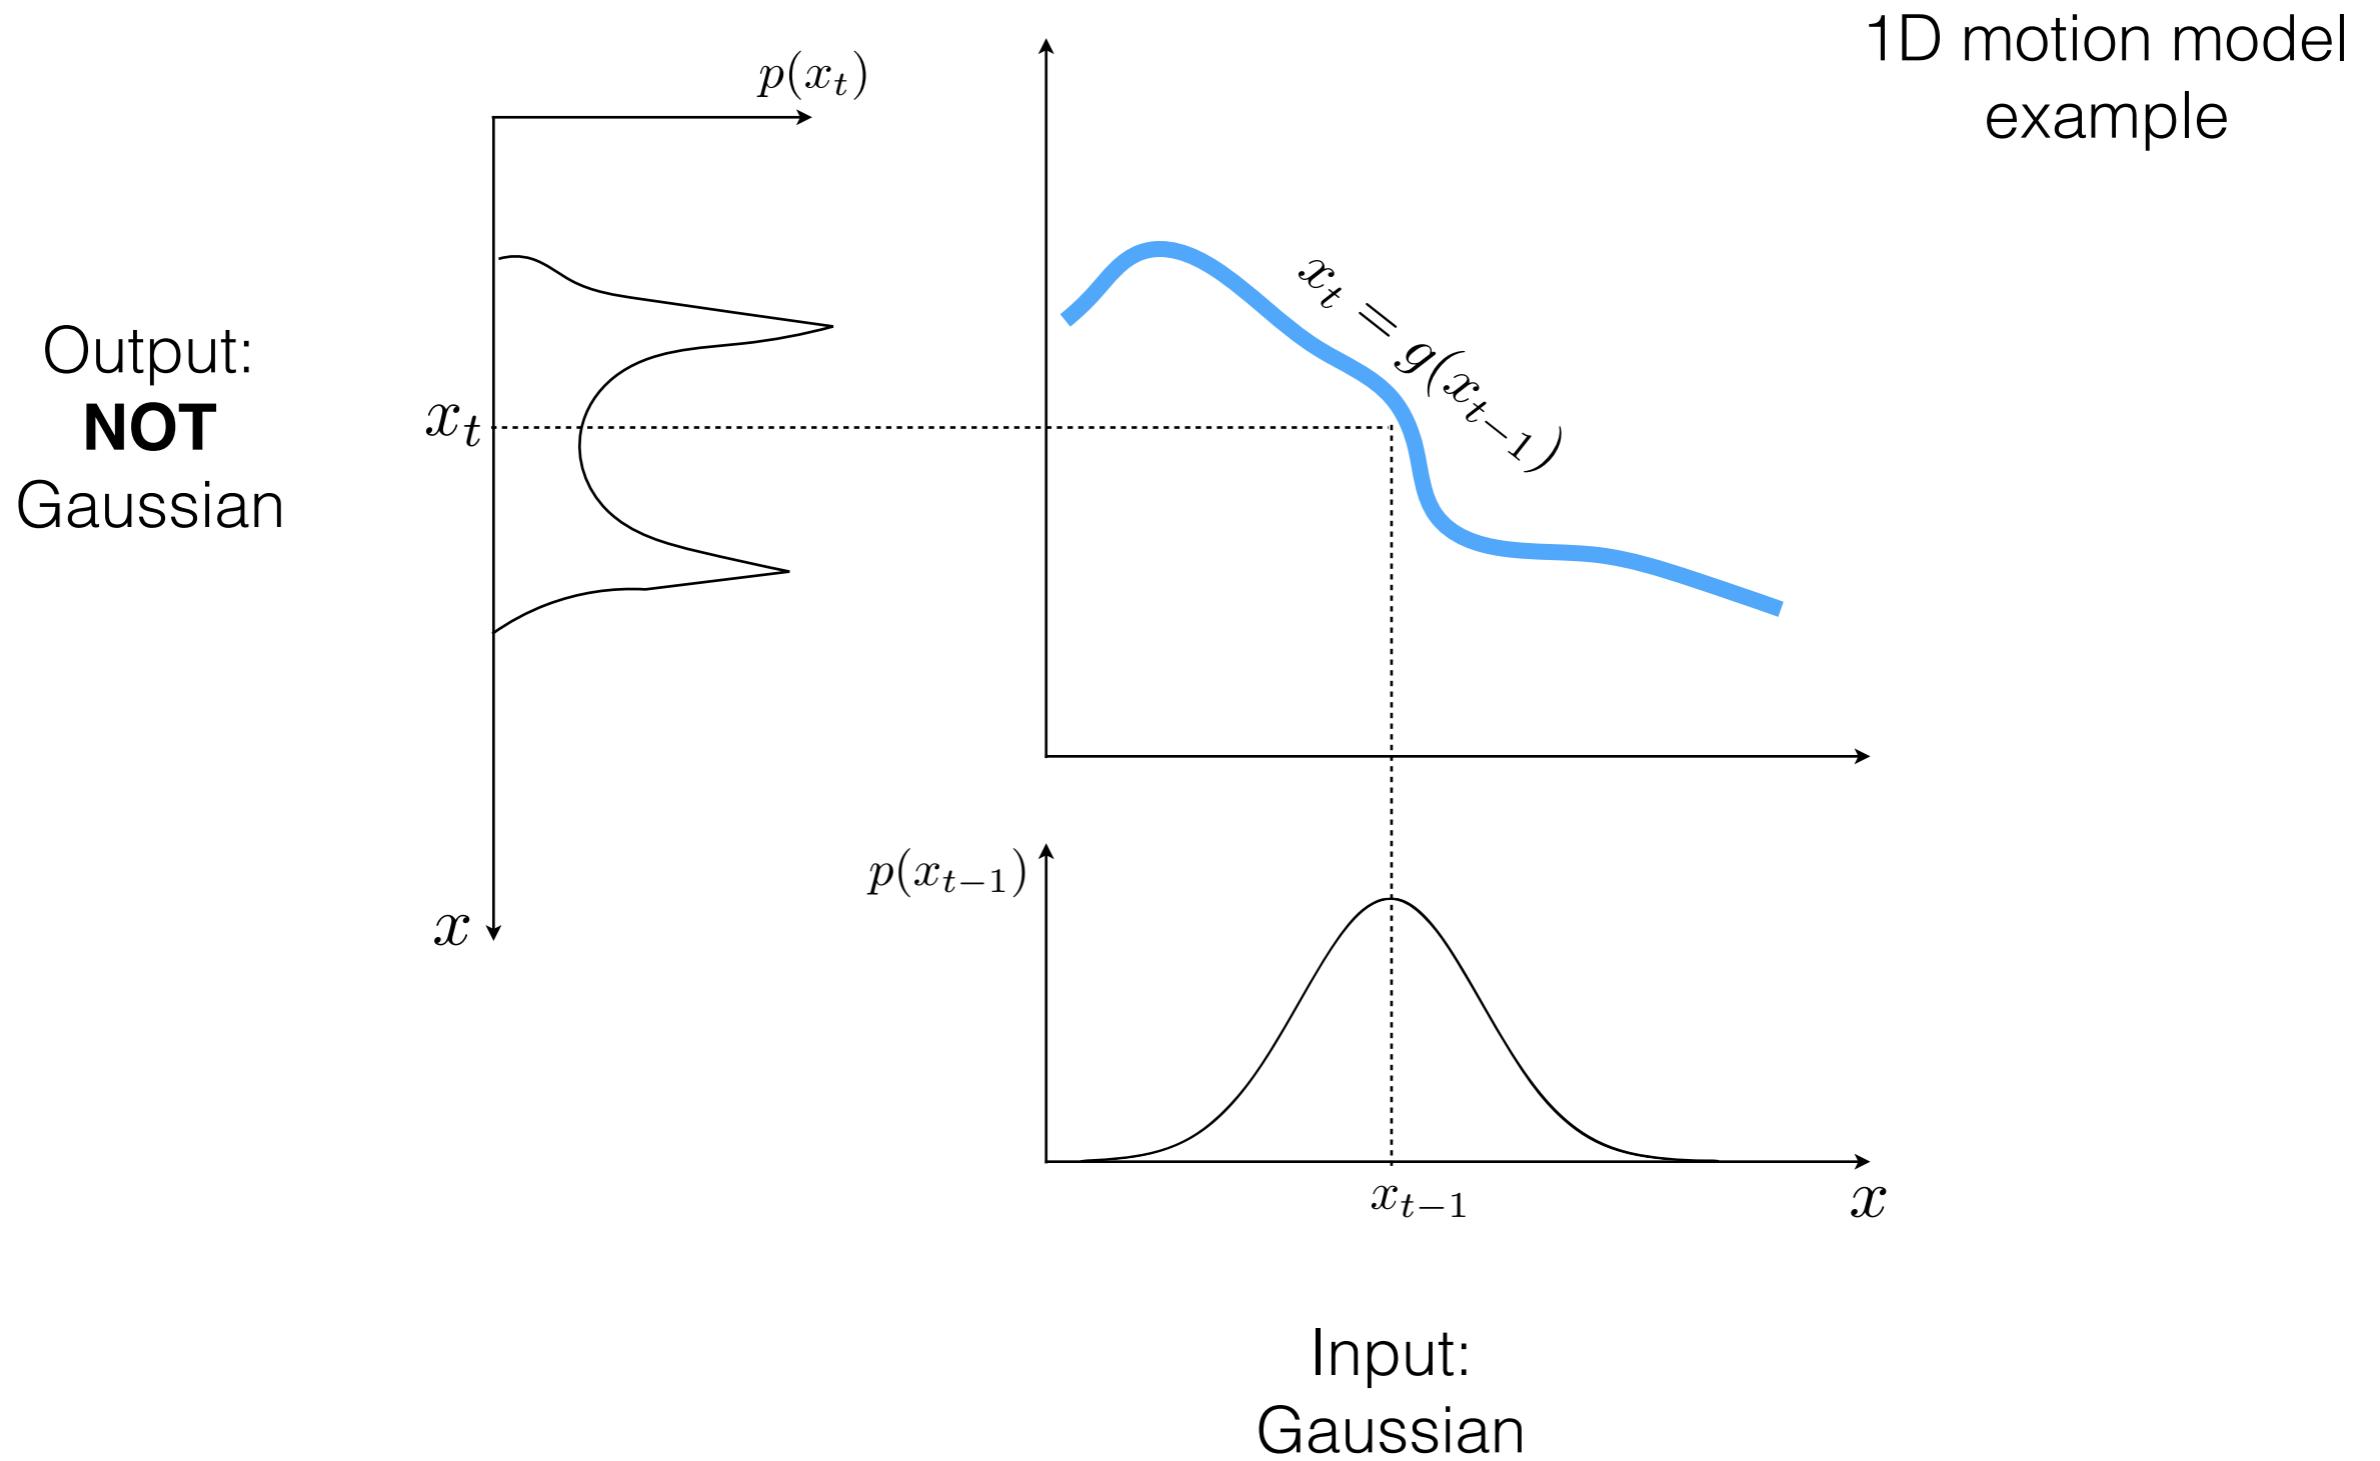


*Can we use the Kalman
Filter?*

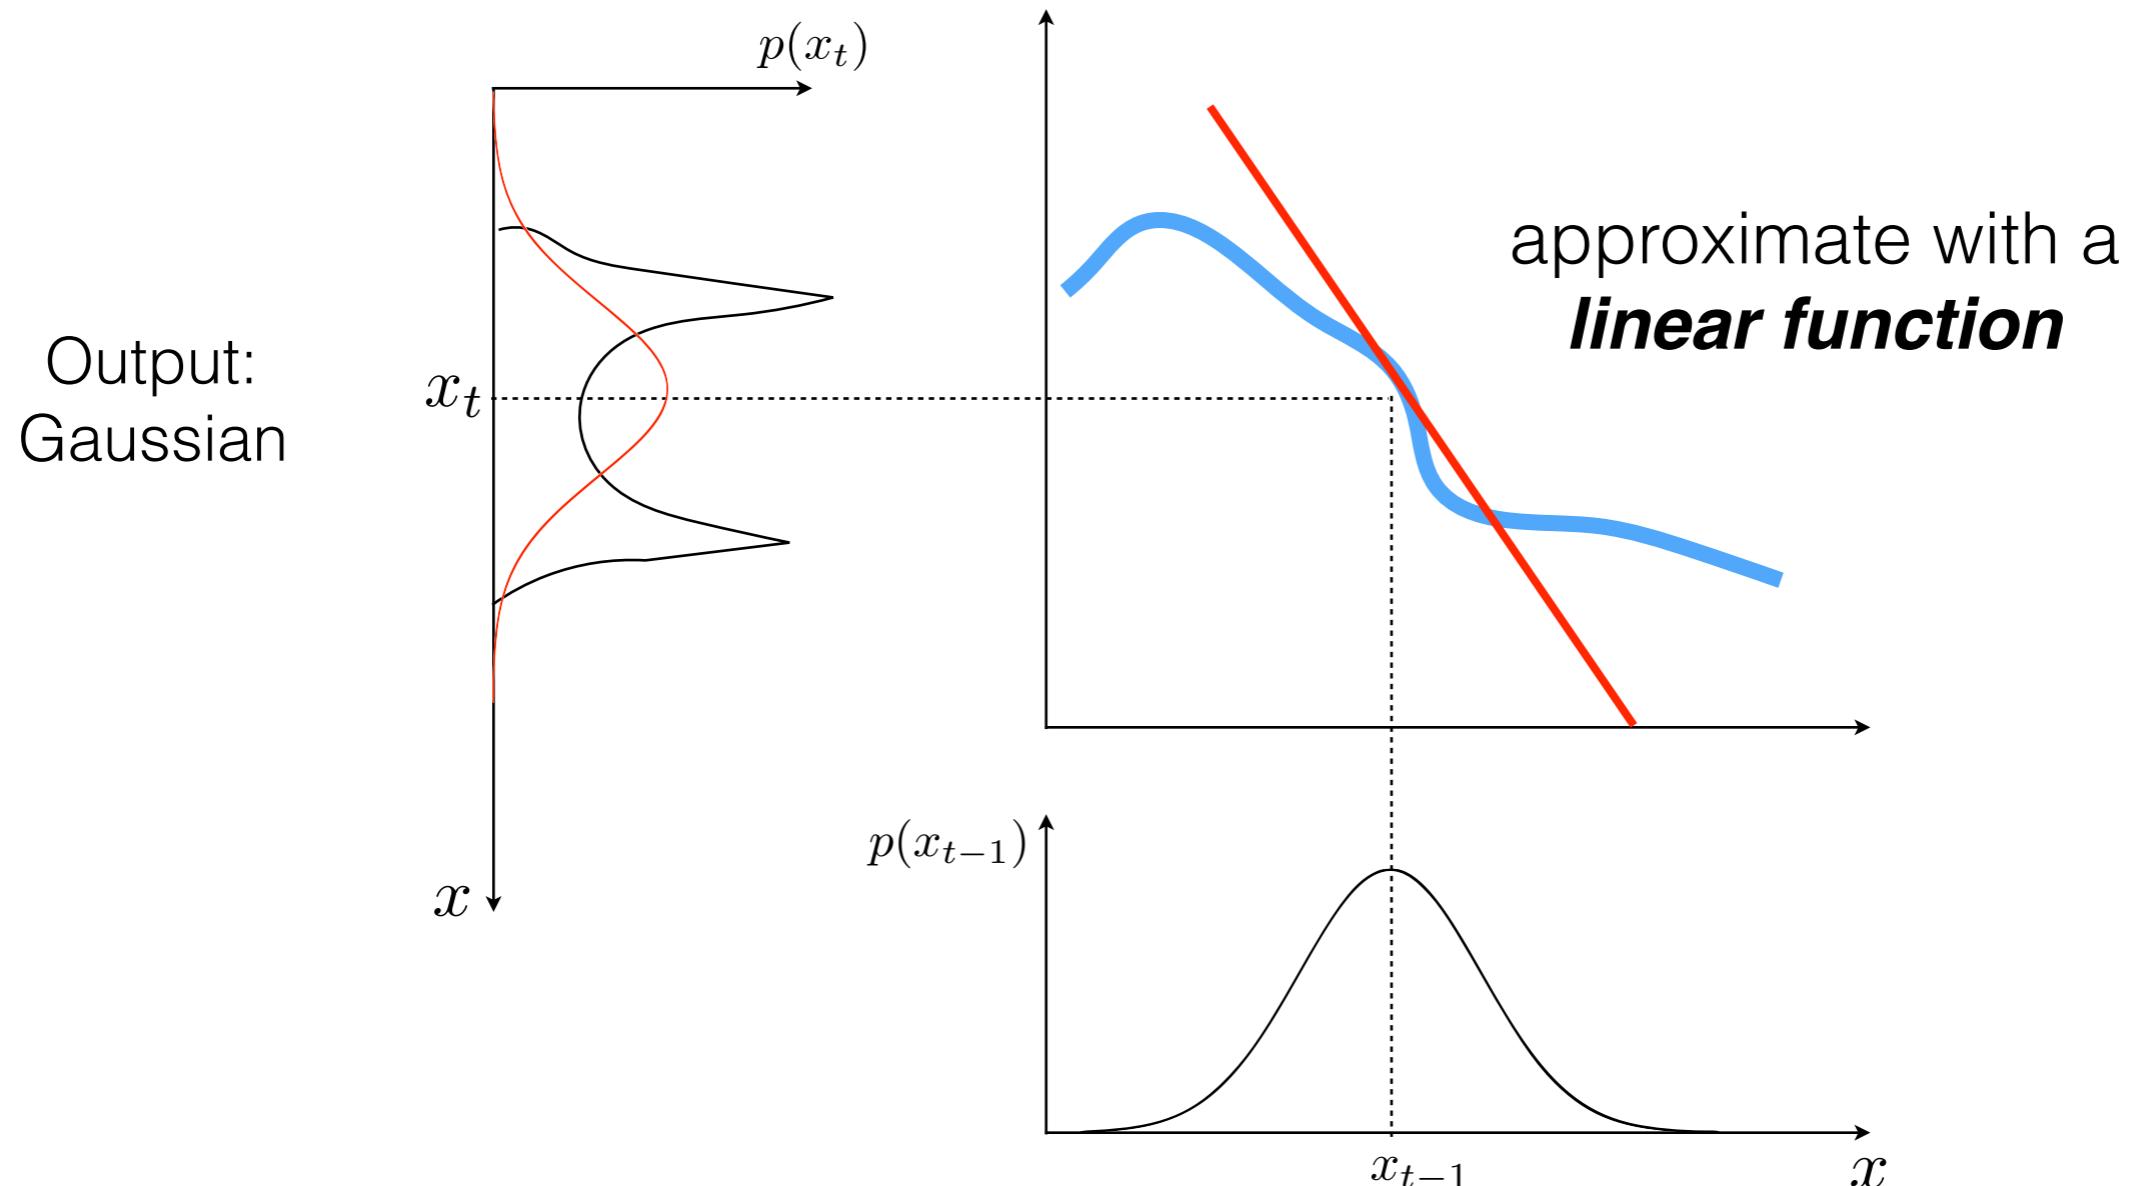
(motion model is not linear)

Input:
Gaussian (**Belief**)

How do you deal with non-linear models?



How do you deal with non-linear models?



When does this trick work?

Input:
Gaussian

Extended Kalman Filter

- Does not assume linear Gaussian models
- Assumes Gaussian noise
- Uses local linear approximations of model to keep the efficiency of the KF framework

Kalman Filter

linear motion model

$$x_t = Ax_{t-1} + Bu_t + \epsilon_t$$

linear sensor model

$$z_t = C_t x_t + \delta_t$$

Extended Kalman Filter

non-linear motion model

$$x_t = g(x_{t-1}, u_t) + \epsilon_t$$

non-linear sensor model

$$z_t = H(x_t) + \delta_t$$

Motion model linearization

$$g(x_{t-1}, u_t) \approx g(\mu_{t-1}, u_t) + \frac{\partial g(\mu_{t-1}, u_t)}{\partial x_{t-1}}(x_{t-1} - \mu_{t-1})$$

Taylor series expansion

Motion model linearization

$$\begin{aligned} g(x_{t-1}, u_t) &\approx g(\mu_{t-1}, u_t) + \frac{\partial g(\mu_{t-1}, u_t)}{\partial x_{t-1}} (x_{t-1} - \mu_{t-1}) \\ &\approx g(\mu_{t-1}, u_t) + G_t (x_{t-1} - \mu_{t-1}) \end{aligned}$$


What's this called?

Motion model linearization

$$g(x_{t-1}, u_t) \approx g(\mu_{t-1}, u_t) + \frac{\partial g(\mu_{t-1}, u_t)}{\partial x_{t-1}} (x_{t-1} - \mu_{t-1})$$
$$\approx g(\mu_{t-1}, u_t) + G_t (x_{t-1} - \mu_{t-1})$$



What's this called?

Jacobian Matrix

'the rate of change in x'
'slope of the function'

Motion model linearization

$$\begin{aligned} g(x_{t-1}, u_t) &\approx g(\mu_{t-1}, u_t) + \frac{\partial g(\mu_{t-1}, u_t)}{\partial x_{t-1}}(x_{t-1} - \mu_{t-1}) \\ &\approx g(\mu_{t-1}, u_t) + G_t(x_{t-1} - \mu_{t-1}) \end{aligned}$$

Jacobian Matrix

'the rate of change in x'
'slope of the function'

Sensor model linearization

$$\begin{aligned} h(x_t) &\approx h(\bar{\mu}_t) + \frac{\partial h(\bar{\mu}_t)}{\partial x_t}(x_t - \bar{\mu}_t) \\ &\approx h(\bar{\mu}_t) + H_t(x_t - \bar{\mu}_t) \end{aligned}$$

New EKF Algorithm

(pretty much the same)

Kalman Filter

$$\bar{\mu}_t = A_t \mu_{t-1} + B u_t$$

$$\bar{\Sigma}_t = A_t \Sigma_{t-1} A_t^\top + R$$

$$K_t = \bar{\Sigma}_t C_t^\top (C_t \bar{\Sigma}_t C_t^\top + Q_t)^{-1}$$

$$\mu_t = \bar{\mu}_t + K_t (z_t - C_t \bar{\mu}_t)$$

$$\Sigma_t = (I - K_t C_t) \bar{\Sigma}_t$$

Extended KF

$$\bar{\mu}_t = g(\mu_{t-1}, u_t)$$

$$\bar{\Sigma}_t = G_t \bar{\Sigma}_{t-1} G_t^\top + R$$

$$K_t = \bar{\Sigma}_t H_t^\top (H_t \bar{\Sigma}_t H_t^\top + Q)^{-1}$$

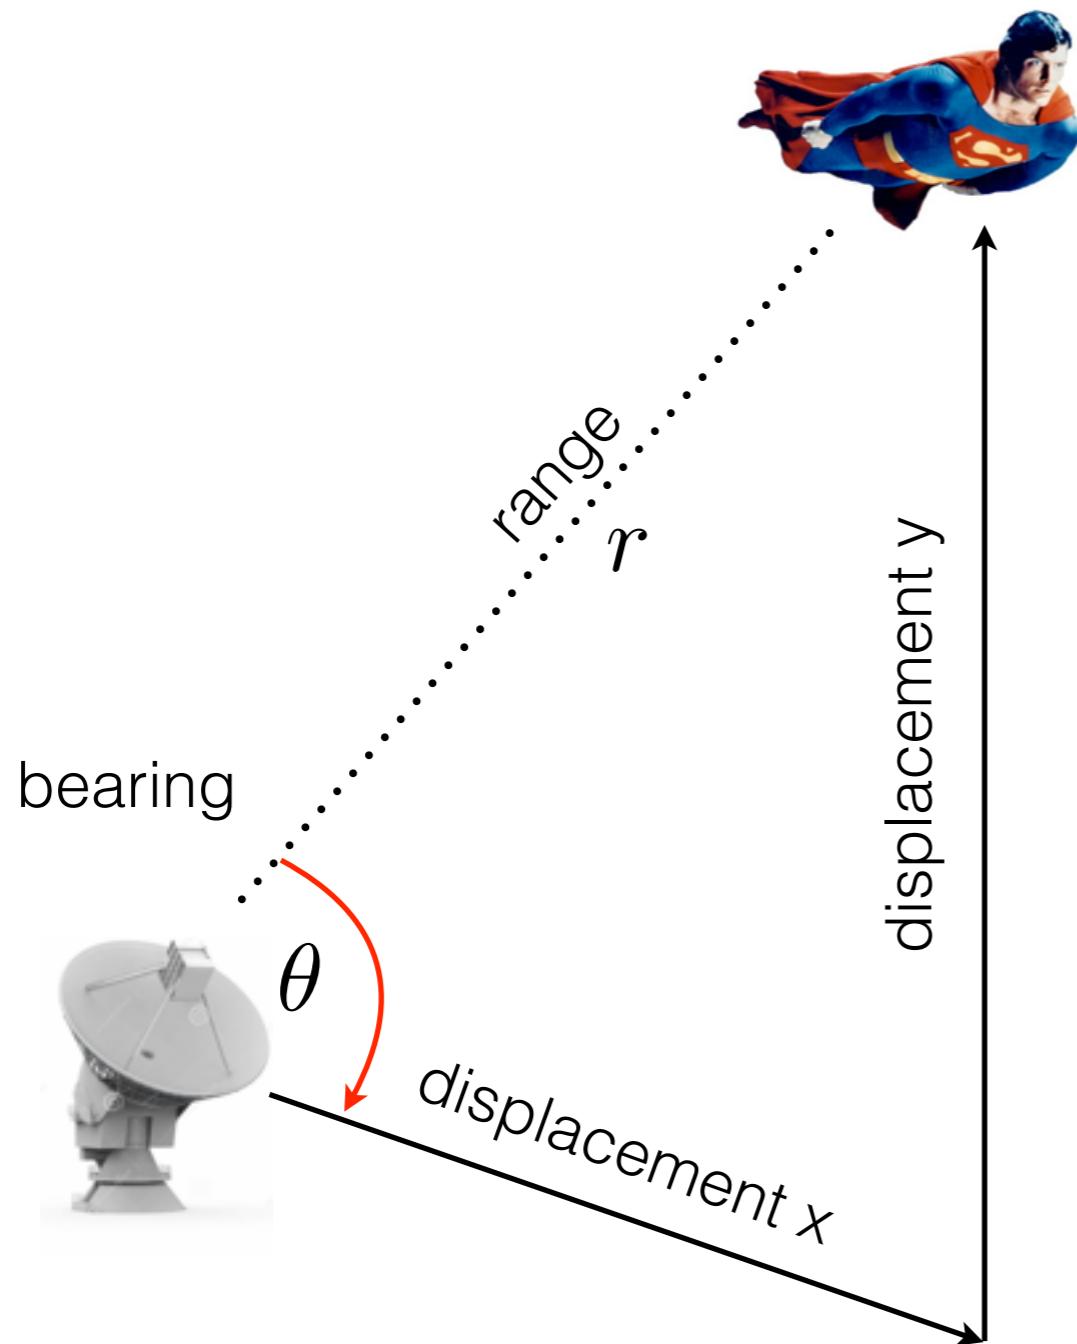
$$\mu_t = \bar{\mu}_t + K_t (z_t - h(\bar{\mu}_t))$$

$$\Sigma_t = (I - K_t H_t) \bar{\Sigma}_t$$

2D example



state: position-velocity



$$\mathbf{x} = \begin{bmatrix} x \\ \dot{x} \\ y \\ \dot{y} \end{bmatrix} \quad \begin{array}{l} \text{position} \\ \text{velocity} \\ \text{position} \\ \text{velocity} \end{array}$$

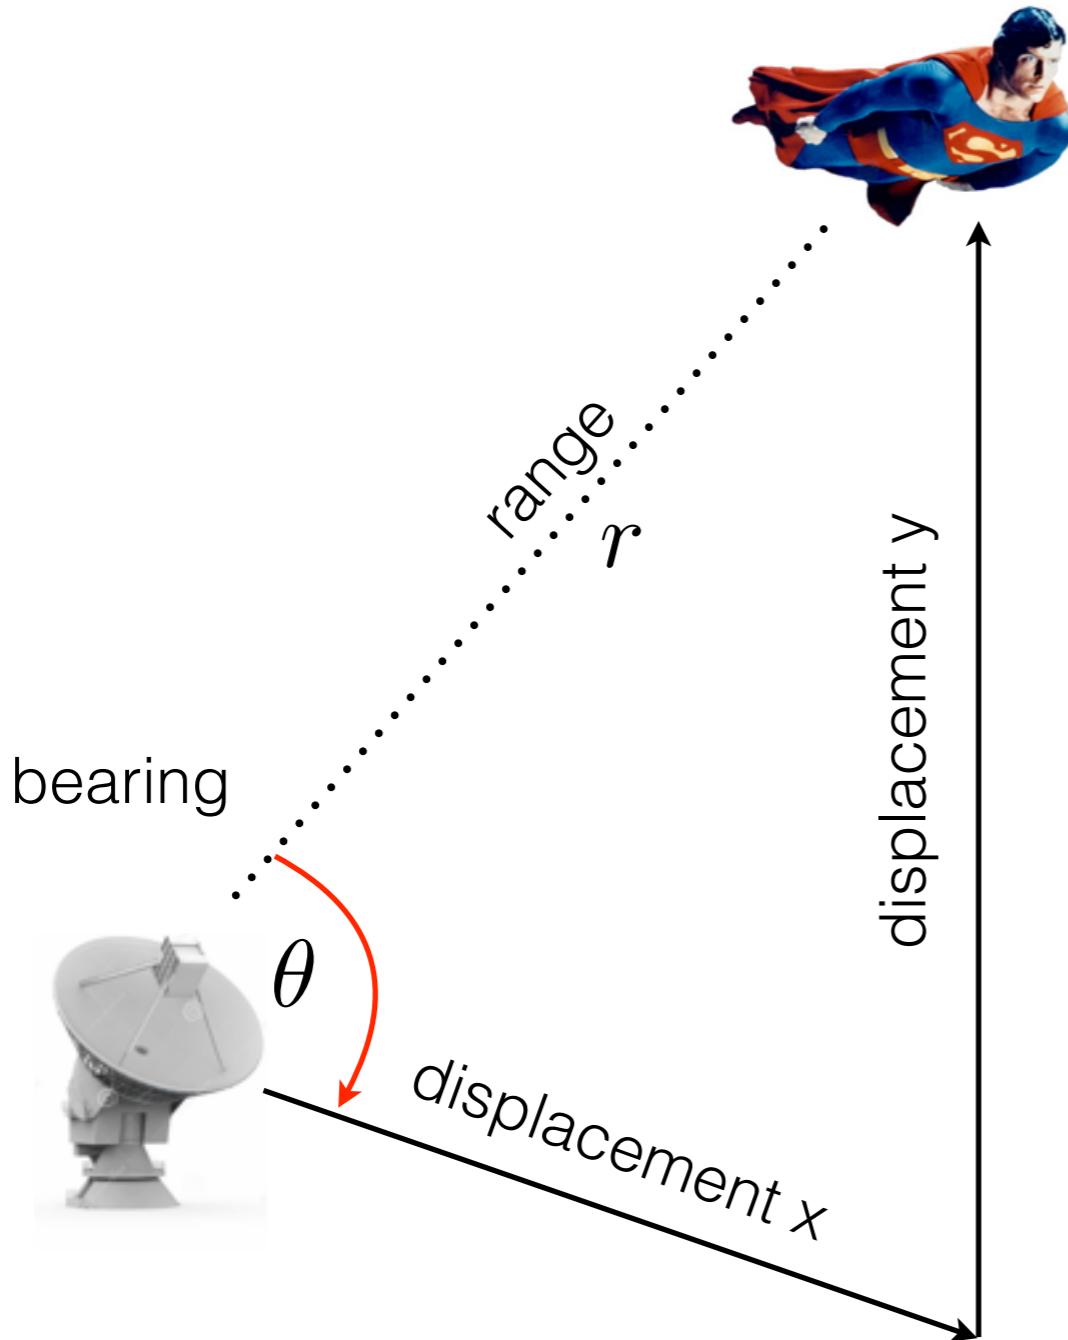
constant velocity motion model

$$A = \begin{bmatrix} 1 & \Delta t & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & \Delta t \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

with additive Gaussian noise

Motion model is linear but ...

measurement: range-bearing



$$\begin{aligned} z &= \begin{bmatrix} r \\ \theta \end{bmatrix} \\ &= \begin{bmatrix} \sqrt{x^2 + y^2} \\ \tan^{-1}(y/x) \end{bmatrix} \end{aligned}$$

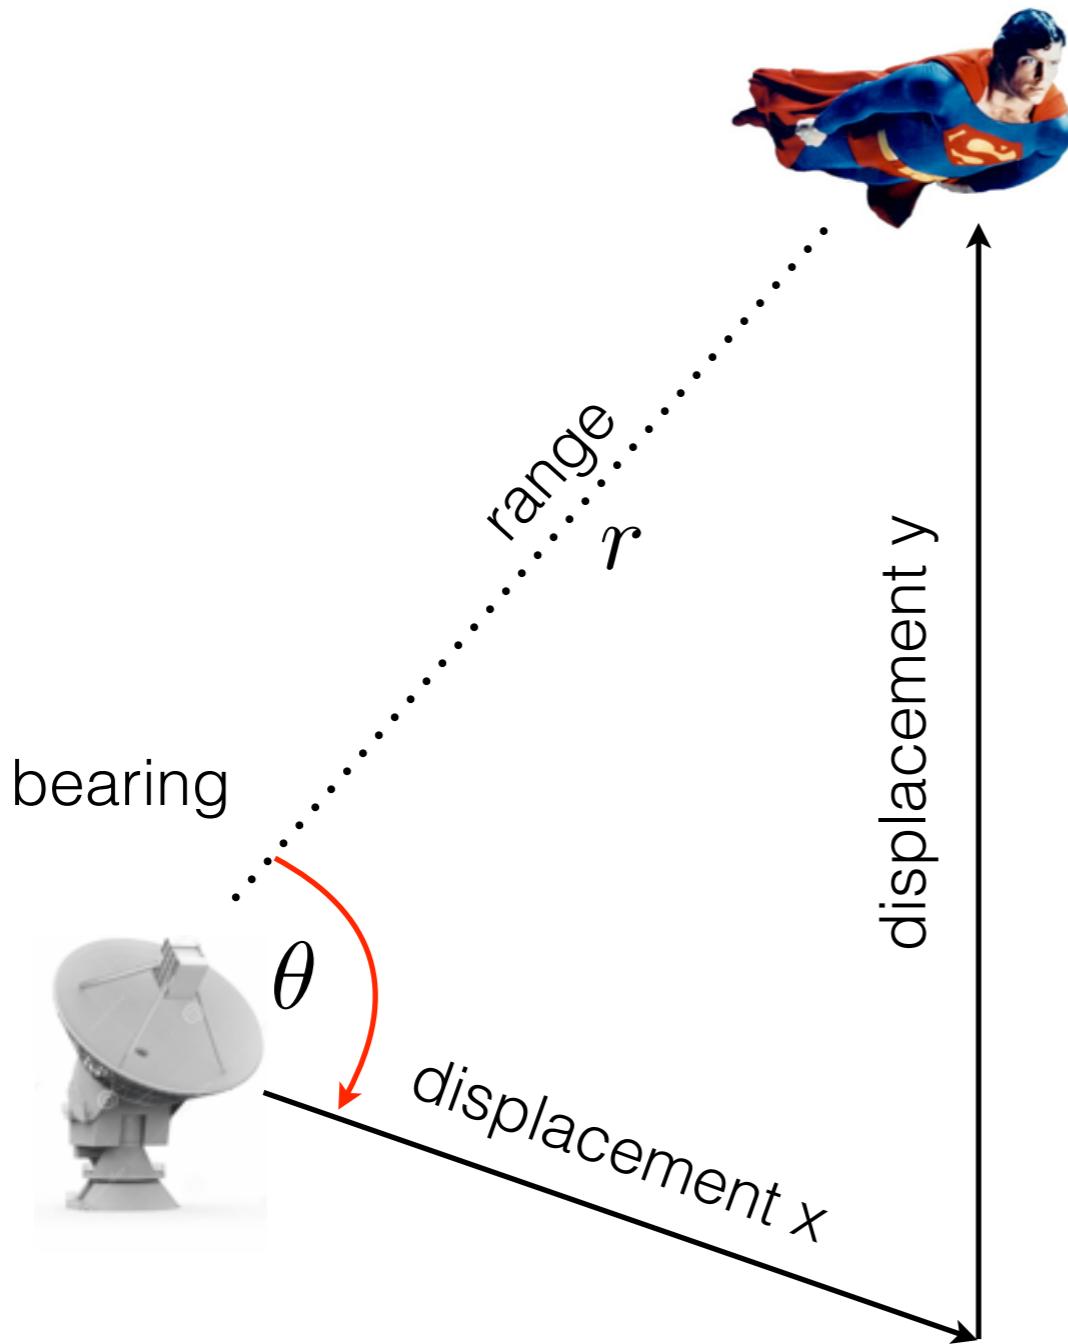
measurement model

Is the measurement model linear?

$$z = h(r, \theta)$$

with additive Gaussian noise

measurement: range-bearing



$$\begin{aligned} z &= \begin{bmatrix} r \\ \theta \end{bmatrix} \\ &= \begin{bmatrix} \sqrt{x^2 + y^2} \\ \tan^{-1}(y/x) \end{bmatrix} \end{aligned}$$

measurement model

Is the measurement model linear?

$$z = h(r, \theta)$$

with additive Gaussian noise

non-linear!

What should we do?

linearize the observation/measurement model!

$$\begin{aligned} z &= \begin{bmatrix} r \\ \theta \end{bmatrix} \\ &= \begin{bmatrix} \sqrt{x^2 + y^2} \\ \tan^{-1}(y/x) \end{bmatrix} \end{aligned}$$

$$H = \frac{\partial z}{\partial x} = ?$$

What is the Jacobian?

$$H = \begin{bmatrix} \frac{\partial r}{\partial x} & \frac{\partial r}{\partial \dot{x}} & \frac{\partial r}{\partial y} & \frac{\partial r}{\partial \dot{y}} \\ \frac{\partial \theta}{\partial x} & \frac{\partial \theta}{\partial \dot{x}} & \frac{\partial \theta}{\partial y} & \frac{\partial \theta}{\partial \dot{y}} \end{bmatrix} =$$

$$\begin{aligned} \mathbf{z} &= \begin{bmatrix} r \\ \theta \end{bmatrix} \\ &= \begin{bmatrix} \sqrt{x^2 + y^2} \\ \tan^{-1}(y/x) \end{bmatrix} \end{aligned}$$

$$H = \frac{\partial \mathbf{z}}{\partial \mathbf{x}} = ?$$

What is the Jacobian?

Jacobian used in the Taylor series expansion looks like ...

$$H = \begin{bmatrix} \frac{\partial r}{\partial x} & \frac{\partial r}{\partial \dot{x}} & \frac{\partial r}{\partial y} & \frac{\partial r}{\partial \dot{y}} \\ \frac{\partial \theta}{\partial x} & \frac{\partial \theta}{\partial \dot{x}} & \frac{\partial \theta}{\partial y} & \frac{\partial \theta}{\partial \dot{y}} \end{bmatrix} = \begin{bmatrix} \cos(\theta) & 0 & \sin(\theta) & 0 \\ -\sin(\theta)/r & 0 & \cos(\theta)/r & 0 \end{bmatrix}$$

```
[x P] = EKF(x, P, z, dt)
```

```
r = sqrt (x(1)^2+x(3)^2);  
b = atan2(x(3),x(1));  
y = [r; b];
```

```
H = [ cos(b) 0 sin(b) 0;  
      -sin(b)/r 0 cos(b)/r 0];
```

```
x = F*x;  
P = F*P*F' + Q;
```

```
K = P*H'/ (H*P*H' + R);
```

```
x = x + K*(z - y);  
P = (eye(size(K,1))-K*H)*P;
```

Parameters:

```
Q = diag([0 .1 0 .1]);  
R = diag([50^2 0.005^2]);  
F = [ 1 dt 0 0;  
      0 1 0 0;  
      0 0 1 dt;  
      0 0 0 1];
```

extra computation for
the EKF measurement
model Jacobian

Problems with EKFs

Taylor series expansion = poor approximation of non-linear functions
success of linearization depends on limited uncertainty and amount
of local non-linearity

Computing partial derivatives is a pain

Drifts when linearization is a bad approximation

Cannot handle multi-modal (multi-hypothesis) distributions