

Sensor fusion in dynamical systems



Thomas Schön
Division of Automatic Control
Linköping University
Sweden

users.isy.liu.se/rt/schon

Joint work with (in alphabetical order): **Fredrik Gustafsson** (LiU), **Joel Hermansson** (Cybaero), **Jeroen Hol** (Xsens), **Johan Kihlberg** (Semcon), **Manon Kok** (LiU), **Henk Luinge** (Xsens), **Per-Johan Nordlund** (Saab), **Henrik Ohlsson** (Berkeley), **Simon Tegelid** (Xdin), **David Törnqvist** (LiU), **Niklas Wahlström** (LiU).



The sensor fusion problem



- Inertial sensors
- Camera
- Barometer



- Inertial sensors
- Radar
- Barometer
- Map



- Inertial sensors
- Cameras
- Radars
- Wheel speed sensors
- Steering wheel sensor



- Inertial sensors
- Ultra-wideband

How do we combine the information from the different sensors?

Might all seem to be very different problems at first sight. However, the same strategies can be used in dealing with all of these applications (and many more).



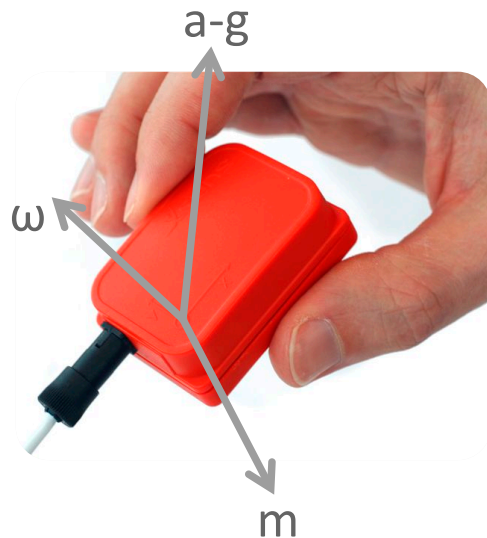
Introductory example (I/III)

Aim: Motion capture, find the motion (position, orientation, velocity and acceleration) of a person (or object) over time.

Industrial partner: Xsens Technologies.

Sensors used:

- 3D accelerometer (acceleration)
- 3D gyroscope (angular velocity)
- 3D magnetometer (magnetic field)

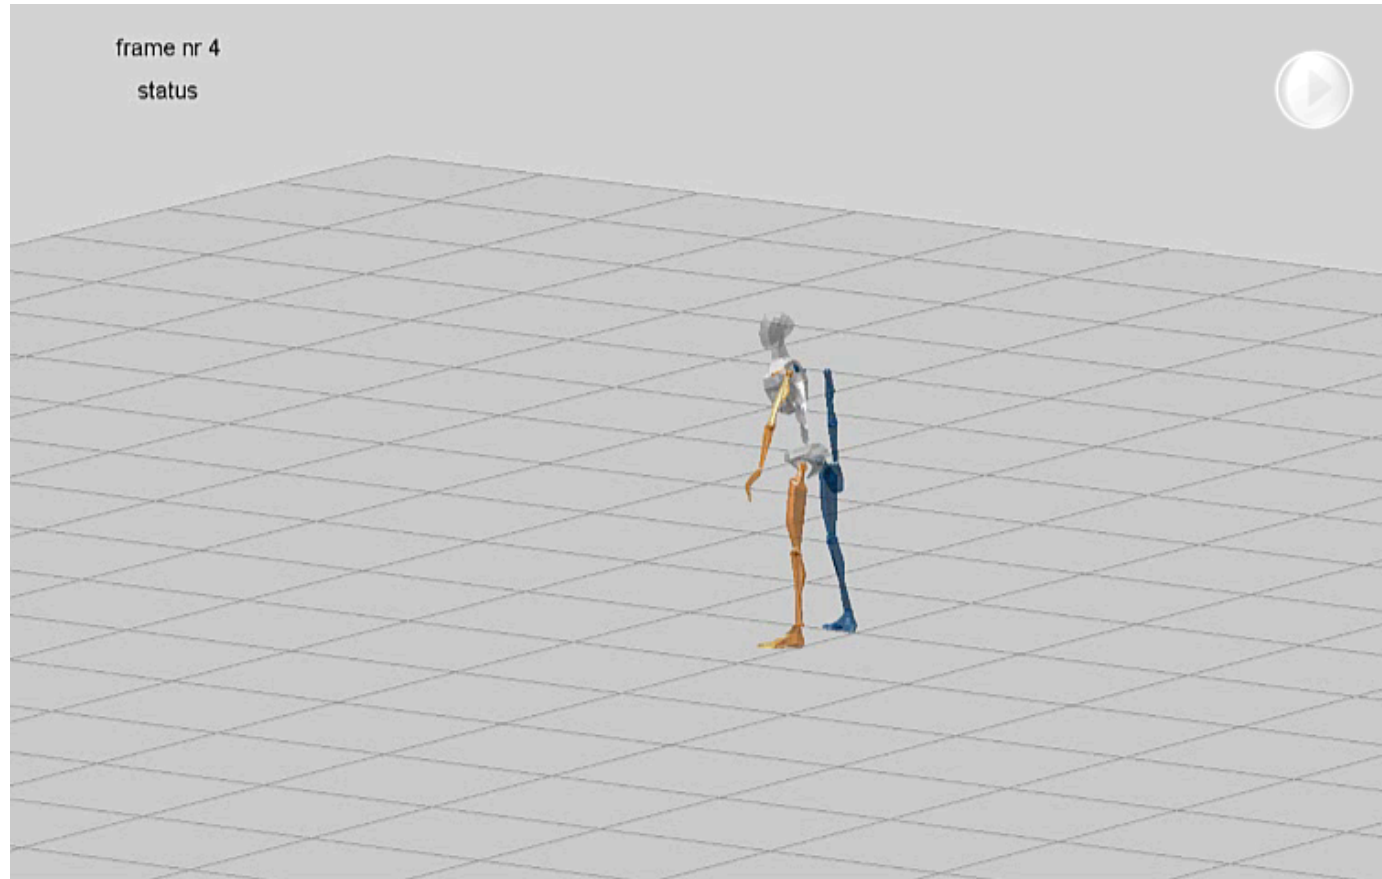


17 sensor units are mounted onto the body of the person.



Introductory example (II/III)

I. Only making use of the inertial information.

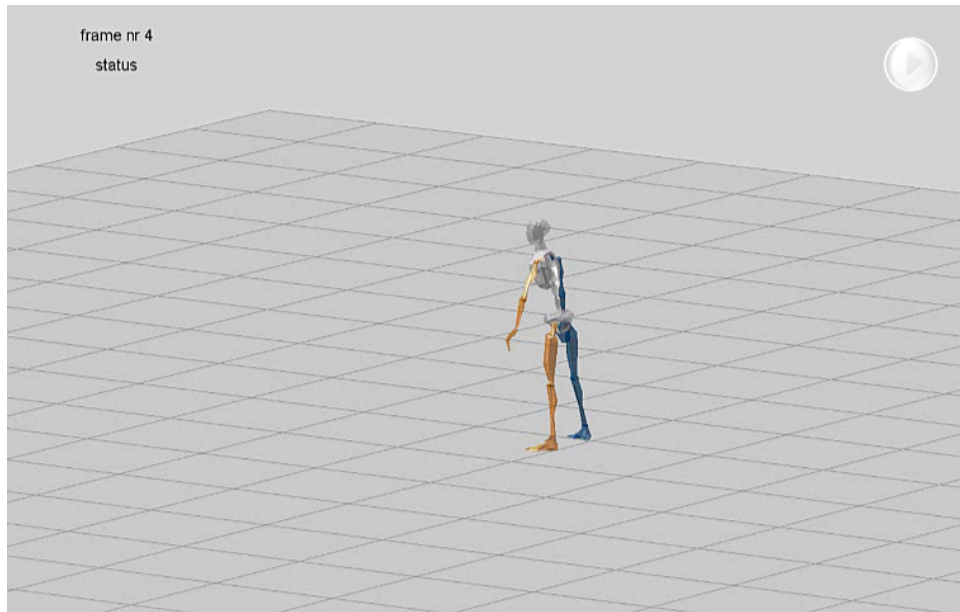


Movie courtesy of Daniel Roetenberg (Xsens)

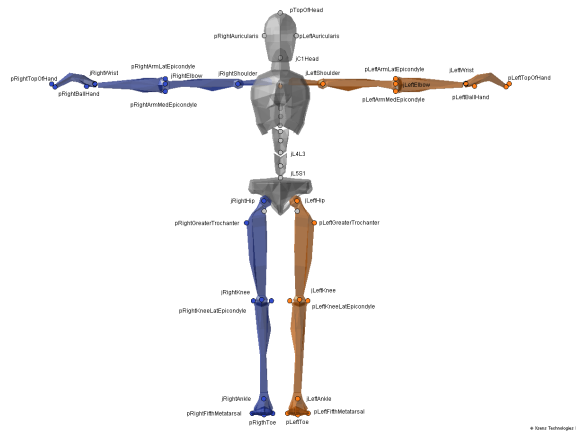


Introductory example (III/III)

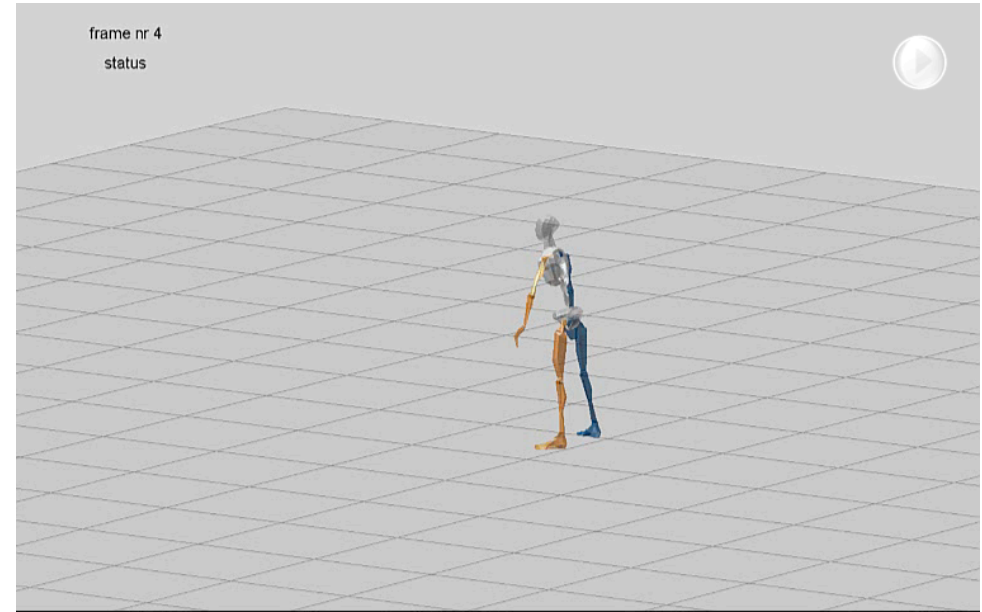
2. Inertial + biomechanical model



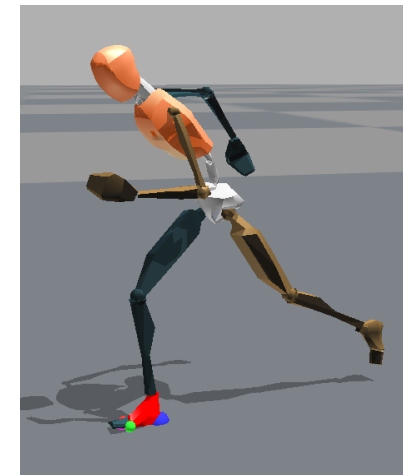
Movie courtesy of Daniel Roetenberg (Xsens)



3. Inertial + biomechanical model + world model



Movie courtesy of Daniel Roetenberg (Xsens)



These introductory examples leads to several questions, e.g.,

- Can we incorporate more sensors?
- Can we make use of more informative world models?
- How do we solve the inherent inference problem?
- Perhaps most importantly, can this be solved systematically?

There are many interesting problems that can be solved systematically, by addressing the following problem areas

Sensor fusion

1. Probabilistic models of dynamical systems
2. Probabilistic models of sensors and the world
3. Formulate and solve the state inference problem
4. Surrounding infrastructure

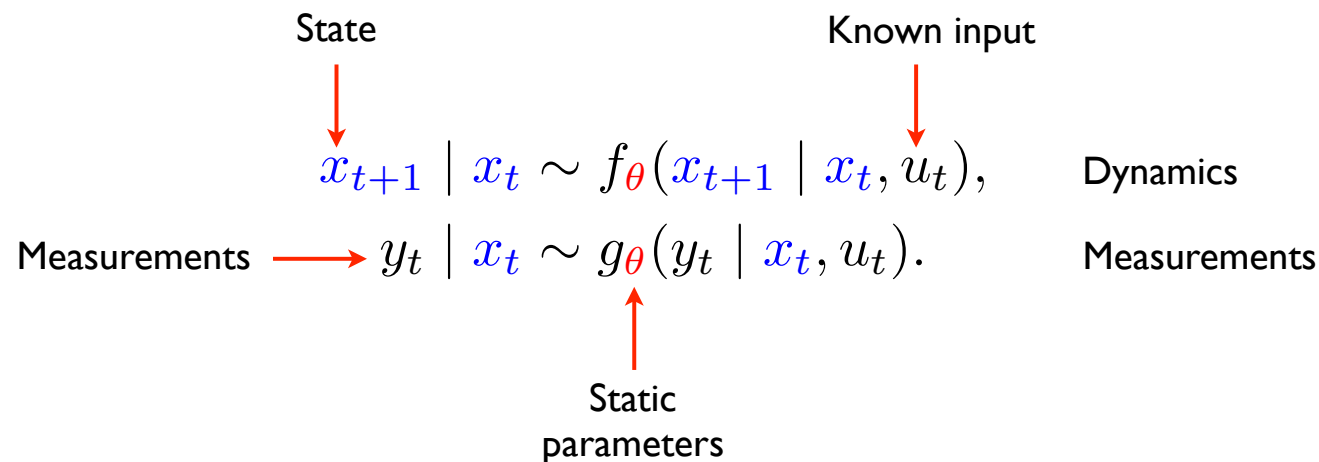


I. Probabilistic models of dynamical systems

Basic representation: Two discrete-time stochastic processes,

- $\{x_t\}_{t \geq 1}$ representing the state of the system
- $\{y_t\}_{t \geq 1}$ representing the measurements from the sensors

The probabilistic model is described using two (f and g) probability density functions (PDFs):



Model = PDF

This type of model is referred to as a **state space model (SSM)** or a **hidden Markov model (HMM)**.



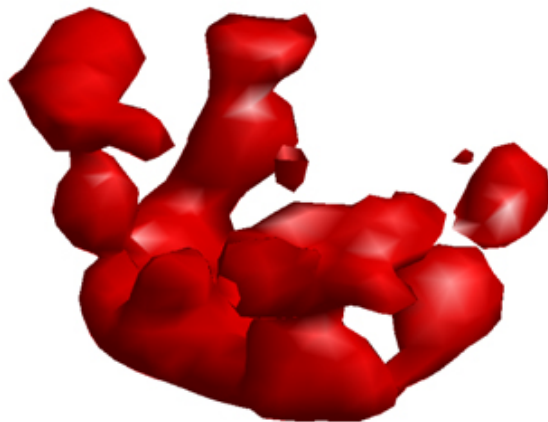
2. World model

The dynamical systems exist in a context.

This requires a **world model**.

Valuable (indeed often necessary) source of information in computing situational awareness.
There are more and more complex world models being built all the time.

An example is our new models of the magnetic contents in various objects, which opens up for interesting new possibilities....



(a) Estimated shape of table



(b) Real shape of table

Fig. 1: Estimated magnetic content in a table turned upside down.

Very much work in progress, we will present some initial results at the convention & exhibition center tomorrow.

Niklas Wahlström, Manon Kok, Thomas B. Schön and Fredrik Gustafsson. **Modeling magnetic fields using Gaussian processes**. Submitted to the 38th International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Vancouver, Canada, May 2013.

Manon Kok, Niklas Wahlström, Thomas B. Schön and Fredrik Gustafsson. **MEMS-based inertial navigation based on a magnetic field map**. Submitted to the 38th International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Vancouver, Canada, May 2013



3. Formulate and solve the inference problem

The inference problem amounts to **combining** the knowledge we have from the models (dynamic, world, sensor) and from the measurements.

The **aim** is to compute

$$p(x_{1:t}, \theta \mid y_{1:t})$$

and/or some of its marginal densities,

$$p(x_t \mid y_{1:t})$$

$$p(\theta \mid y_{1:t})$$

These densities are then commonly used to form point estimates, **maximum likelihood** or **Bayesian**.

-
- Everything we do rests on a firm foundation of probability theory and mathematical statistics.
 - If we have the wrong model, there is no estimation/learning algorithm that can help us.



3. Inference - the filtering problem

$$p(x_t | y_{1:t}) = \frac{\overbrace{p(y_t | x_t)}^{\text{sensor model}} \overbrace{p(x_t | y_{1:t-1})}^{\text{prediction density}}}{p(y_t | y_{1:t-1})}$$
$$p(x_{t+1} | y_{1:t}) = \int \underbrace{p(x_{t+1} | x_t)}_{\text{dynamical model}} \underbrace{p(x_t | y_{1:t})}_{\text{filtering density}} dx_t$$

In the application examples these equations are solved using particle filters (PF), Rao-Blackwellized particle filters (RBPF), extended Kalman filters (EKF) and various optimization based approaches.



4. The “surrounding infrastructure”

Besides models for dynamics, sensors and world, a successful sensor fusion solution heavily relies on a well functioning “surrounding infrastructure”.

This includes for example:

- Time synchronization of the measurements from the different sensors
- Mounting of the sensors and calibration
- Computer vision, radar processing
- Etc...

An example:



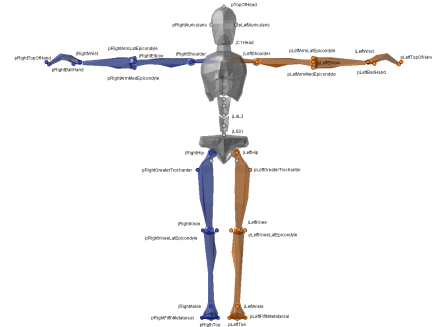
Relative pose calibration:

Compute the relative translation and rotation of the camera and the inertial sensors that are rigidly connected.

Jeroen D. Hol, Thomas B. Schön and Fredrik Gustafsson. **Modeling and Calibration of Inertial and Vision Sensors**. *International Journal of Robotics Research (IJRR)*, 29(2):231-244, February 2010.



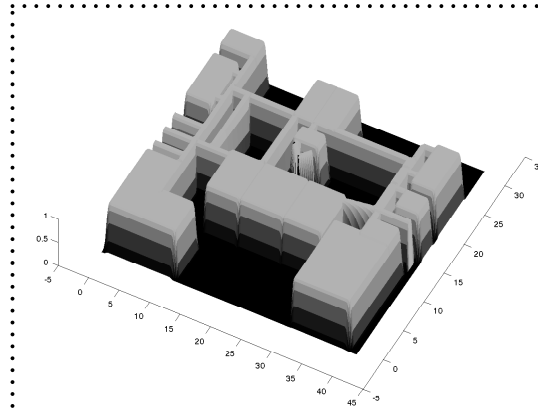
The story I am telling



$$\dot{x} = f(x, u, \theta)$$

1. We are dealing with dynamical systems

This requires a **dynamical model**.



2. The dynamical systems exist in a context.

This requires a **world model**.

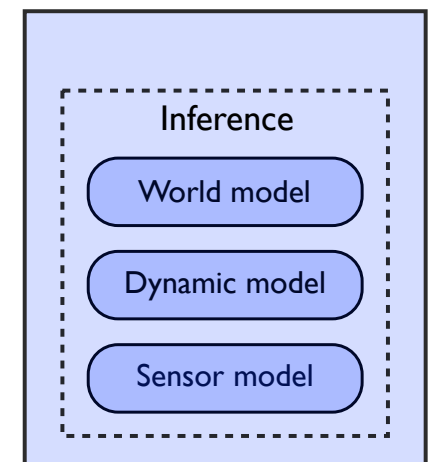
3. The dynamical systems must be able to perceive their own (and others') motion, as well as the surrounding world.

This requires sensors and **sensor models**.



4. We must be able to transform the measurements from the sensors into knowledge about the dynamical systems and their surrounding world.

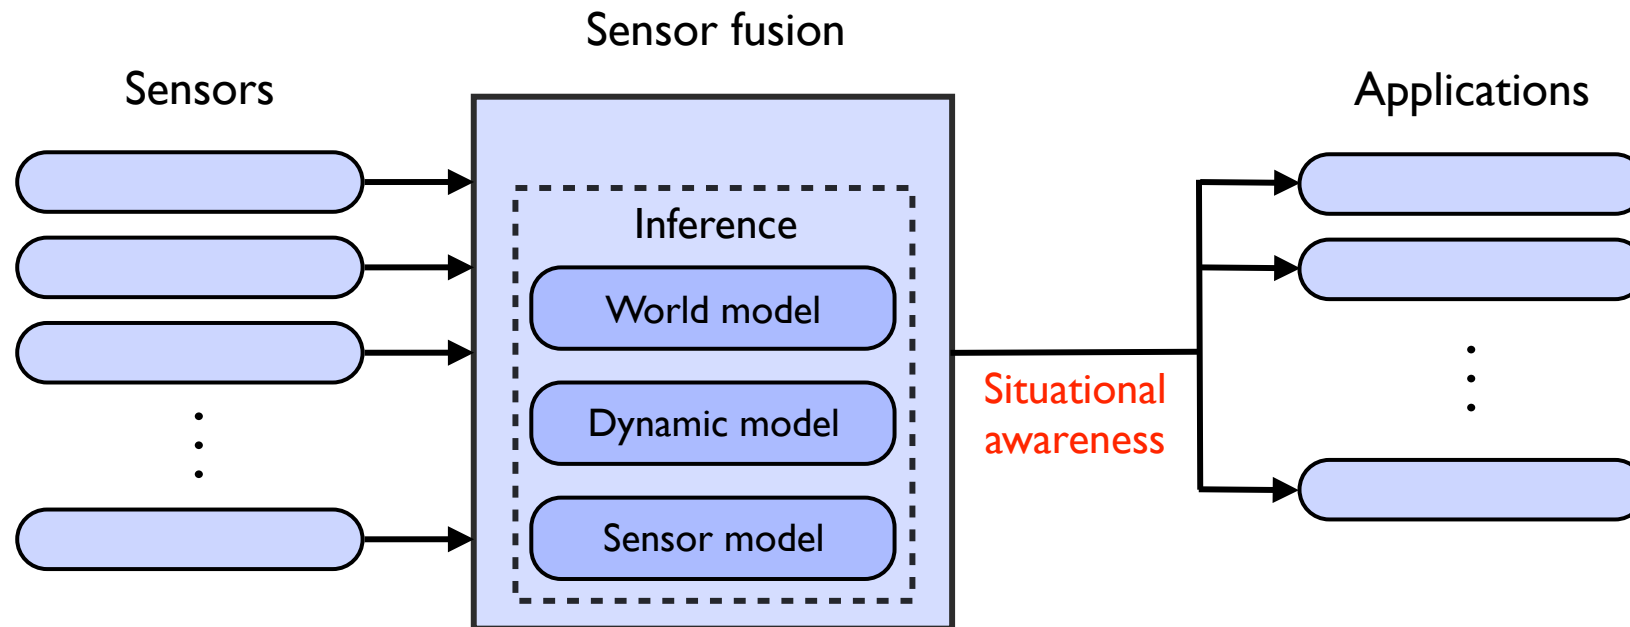
This requires **inference**.



Sensor fusion - definition

Definition (sensor fusion)

Sensor fusion is the process of using information from **several different** sensors to **infer** what is happening (this typically includes finding states of dynamical systems and various static parameters).



Sensor fusion

1. Probabilistic models of dynamical systems
2. Probabilistic models of sensors and the world
3. Formulate and solve the state inference problem
4. Surrounding infrastructure

A few words about the particle filter

Industrial application examples

1. Autonomous landing of a helicopter
2. Helicopter navigation
3. Fighter aircraft navigation
4. Indoor motion capture

Conclusions



State inference - simple special case

Consider the following special case (Linear Gaussian State Space (LGSS) model)

$$\begin{aligned}x_{t+1} &= Ax_t + Bu_t + v_t, & v_t &\sim \mathcal{N}(0, Q), \\y_t &= Cx_t + Du_t + e_t, & e_t &\sim \mathcal{N}(0, R).\end{aligned}$$

or, equivalently,

$$\begin{aligned}x_{t+1} \mid x_t &\sim f(x_{t+1} \mid x_t) = \mathcal{N}(x_{t+1} \mid Ax_t + Bu_t, Q), \\y_t \mid x_t &\sim g(y_t \mid x_t) = \mathcal{N}(y_t \mid Cx_t + Du_t, R).\end{aligned}$$

It is now straightforward to show that the solution to the time update and measurement update equations is given by the Kalman filter, resulting in

$$\begin{aligned}p(x_t \mid y_{1:t}) &= \mathcal{N}(x_t \mid \hat{x}_{t|t}, P_{t|t}), \\p(x_{t+1} \mid y_{1:t}) &= \mathcal{N}(x_{t+1} \mid \hat{x}_{t+1|t}, P_{t+1|t}).\end{aligned}$$



Obvious question: what do we do in an interesting case, for example when we have a nonlinear model including a world model in the form of a map?

- Need a general representation of the filtering PDF
- Try to solve the equations

$$p(x_t | y_{1:t}) = \frac{g(y_t | x_t)p(x_t | y_{1:t-1})}{p(y_t | y_{1:t-1})},$$

$$p(x_{t+1} | y_{1:t}) = \int f(x_{t+1} | x_t)p(x_t | y_{1:t})dx_t,$$

as accurately as possible.



State inference - the particle filter (I/II)

The particle filter provides an approximation of the filter PDF

$$p(\mathbf{x}_t \mid y_{1:t})$$

when the state evolves according to an SSM

$$\begin{aligned}\mathbf{x}_{t+1} \mid \mathbf{x}_t &\sim f(\mathbf{x}_{t+1} \mid \mathbf{x}_t, u_t), \\ y_t \mid \mathbf{x}_t &\sim h(y_t \mid \mathbf{x}_t, u_t), \\ \mathbf{x}_1 &\sim \mu(\mathbf{x}_1).\end{aligned}$$

The particle filter maintains an empirical distribution made up N samples (particles) and corresponding weights

$$\hat{p}(\mathbf{x}_t \mid y_{1:t}) = \sum_{i=1}^N w_t^i \delta_{\mathbf{x}_t^i}(\mathbf{x}_t)$$

“Think of each particle as one simulation of the system state. Only keep the good ones.”

This approximation converge to the true filter PDF,

Xiao-Li Hu, Thomas B. Schön and Lennart Ljung. **A Basic Convergence Result for Particle Filtering.** *IEEE Transactions on Signal Processing*, 56(4):1337-1348, April 2008.



State inference - the particle filter (II/II)

The weights and the particles in

$$\hat{p}(x_t | y_{1:t}) = \sum_{i=1}^N w_t^i \delta_{x_t^i}(x_t)$$

are updated as new measurements becomes available. This approximation can for example be used to compute an estimate of the mean value,

$$\hat{x}_{t|t} = \int x_t p(x_t | y_{1:t}) dx_t \approx \int x_t \sum_{i=1}^N w_t^i \delta_{x_t^i}(x_t) dx_t = \sum_{i=1}^N w_t^i x_t^i$$

The theory underlying the particle filter has been developed over the past two decades and the theory and its applications are still being developed at a very high speed. For a timely tutorial, see

A. Doucet and A. M. Johansen. **A tutorial on particle filtering and smoothing: fifteen years later**. In *Oxford Handbook of Nonlinear Filtering*, 2011, D. Crisan and B. Rozovsky (eds.). Oxford University Press.

or my new PhD course on computational inference in dynamical systems

users.isy.liu.se/rt/schon/course_CIDS.html



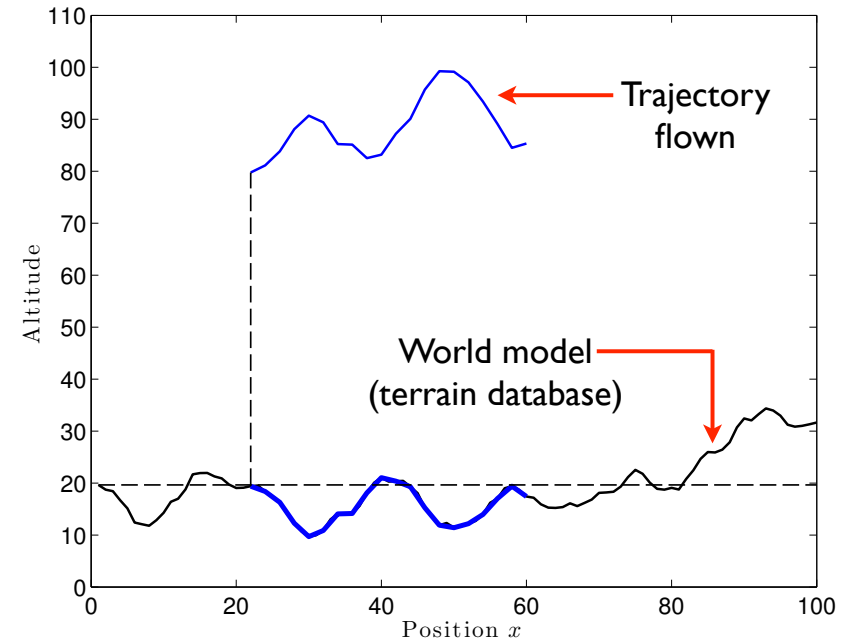
Using world models in solving state inference problems

Consider a 1D localization example.

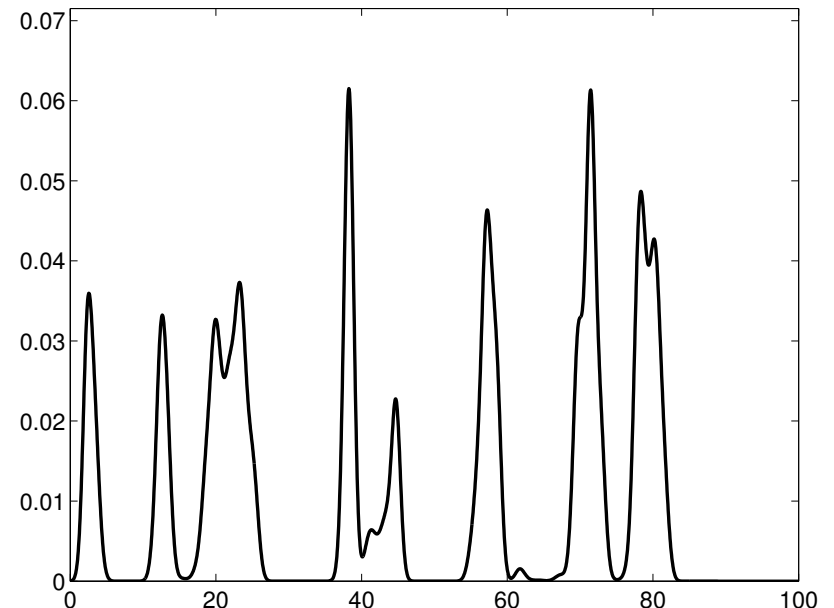
$$x_{t+1} = x_t + u_t + v_t,$$
$$y_t = h(x_t) + e_t.$$

position \downarrow x_t velocity (measured input) \downarrow u_t

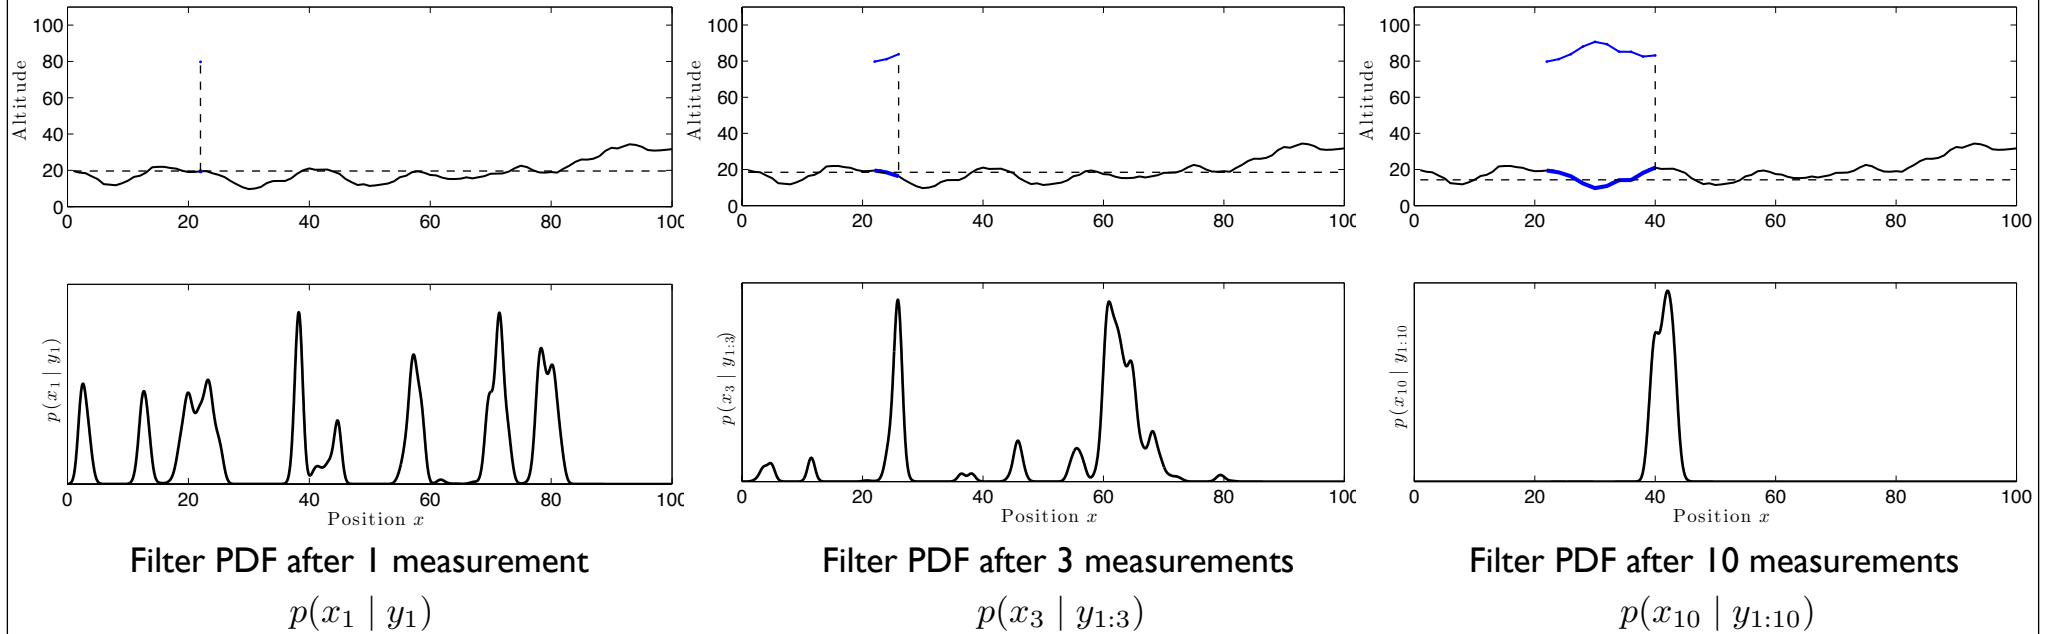
measurement (altitude) \uparrow y_t world model (terrain database) \uparrow $h(x_t)$



Filter PDF after 1 measurement $p(x_1 | y_1)$ \rightarrow



Using world models in solving state inference problems



Using world models in solving state inference problems

The simple ID localization example is an illustration of a problem involving a multimodal filter PDF

- **Straightforward** to represent and work with using a PF
- **Horrible** to work with using e.g. an extended Kalman filter

The example also highlights the **key capabilities** of the PF:

1. To automatically handle an unknown and dynamically changing number of hypotheses.

2. Work with nonlinear/non-Gaussian models

We have implemented a similar localization solution for this aircraft (Gripen).

Industrial partner: Saab



Sensor fusion

1. Probabilistic models of dynamical systems
2. Probabilistic models of sensors and the world
3. Formulate and solve the state inference problem
4. Surrounding infrastructure

A few words about the particle filter

Industrial application examples

1. Autonomous landing of a helicopter
2. Helicopter navigation
3. Fighter aircraft navigation
4. Indoor motion capture

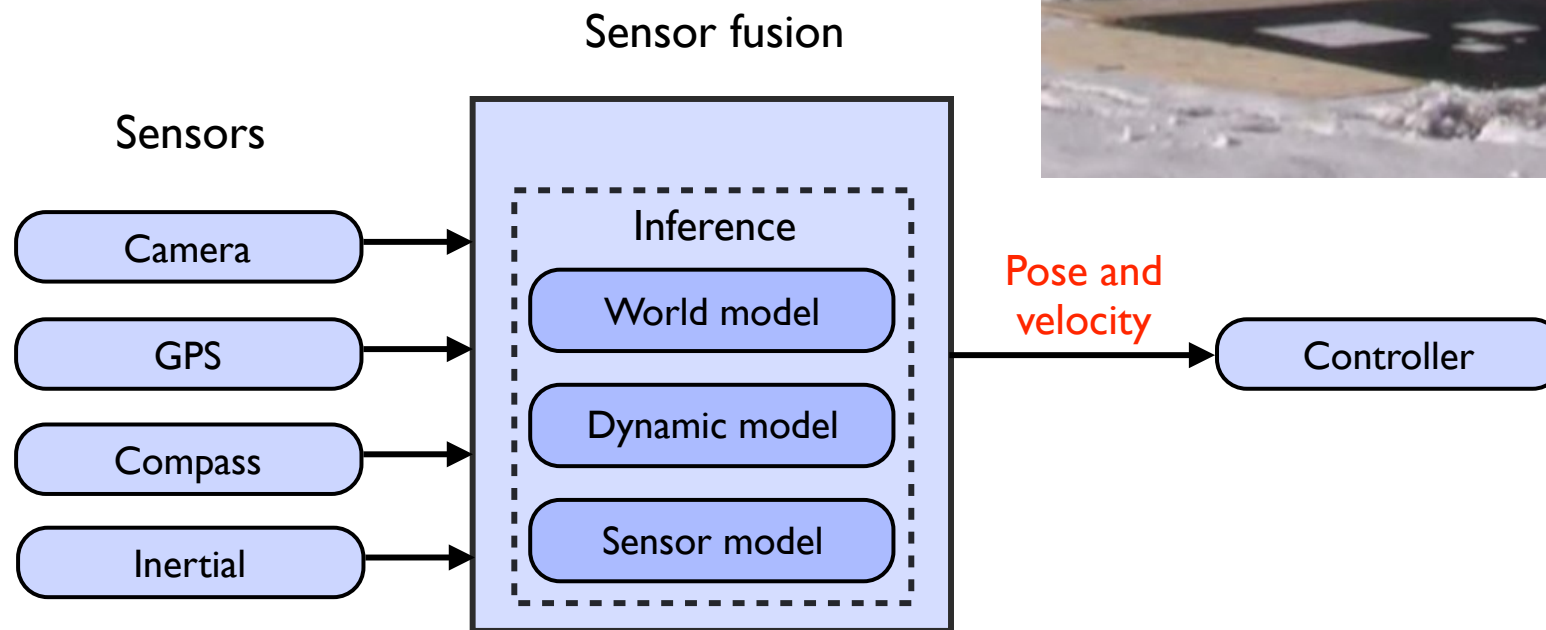
Conclusions



I. Autonomous helicopter landing (I/III)

Aim: Land a helicopter autonomously using information from a camera, GPS, compass and inertial sensors.

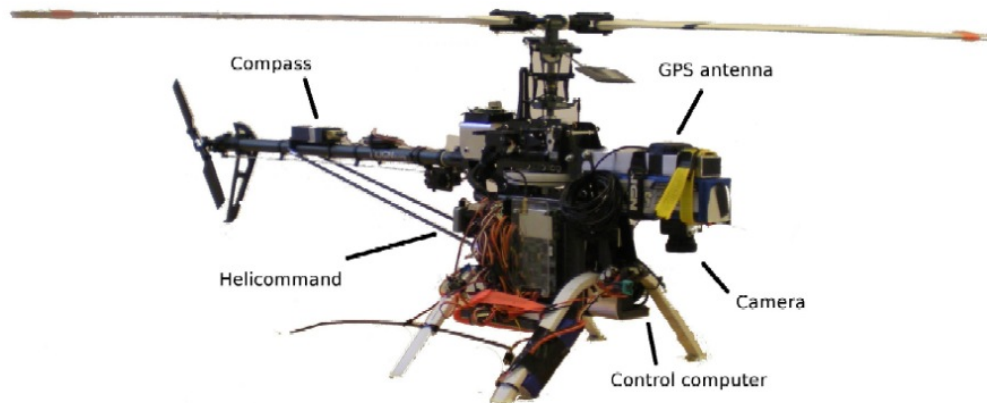
Industrial partner: Cybaero



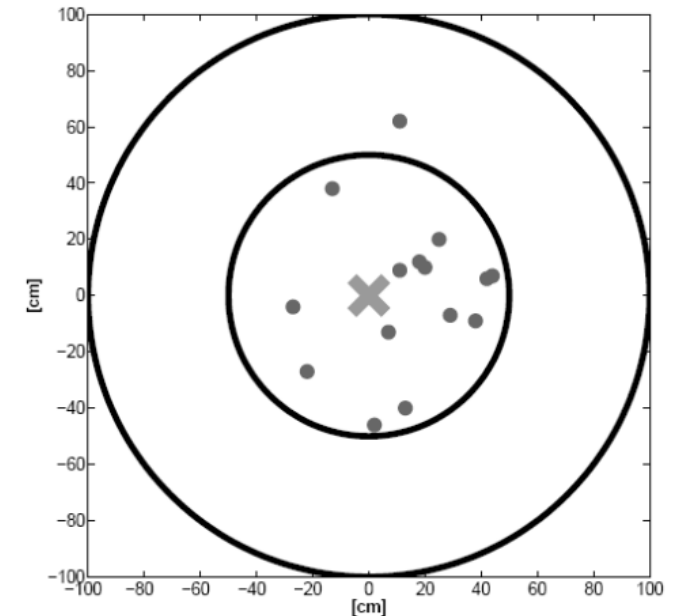
I. Autonomous helicopter landing (II/III)

Experimental helicopter

- Weight: 5kg
- Electric motor



Results from 15 landings



The two circles mark 0.5m and 1m landing error, respectively.

Dots = achieved landings
Cross = perfect landing

Joel Hermansson, Andreas Gising, Martin Skoglund and Thomas B. Schön. **Autonomous Landing of an Unmanned Aerial Vehicle.** *Reglermöte (Swedish Control Conference)*, Lund, Sweden, June 2010.

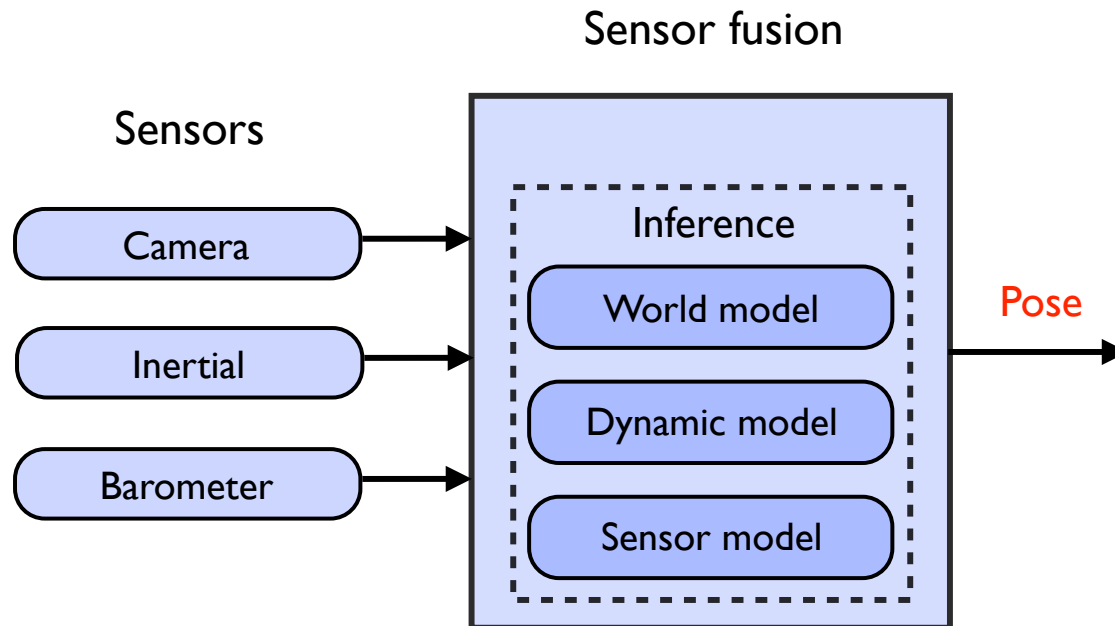


I. Autonomous helicopter landing (III/III)



2. Helicopter pose estimation using a map (I/III)

Aim: Compute the position and orientation of a helicopter by exploiting the information present in Google maps images of the operational area.



2. Helicopter pose estimation using a map (II/III)



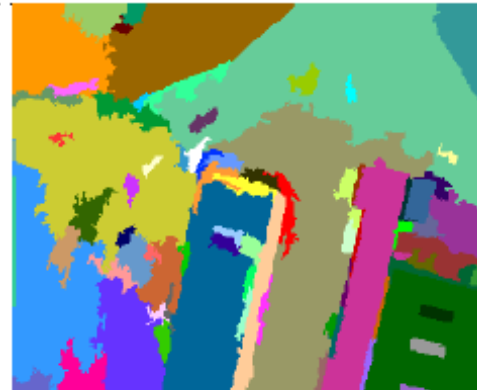
Map over the operational environment obtained from Google Earth.



Manually classified map with grass, asphalt and houses as pre-specified classes.



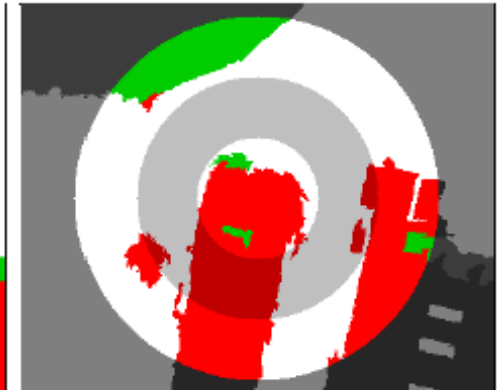
Image from on-board camera



Extracted superpixels



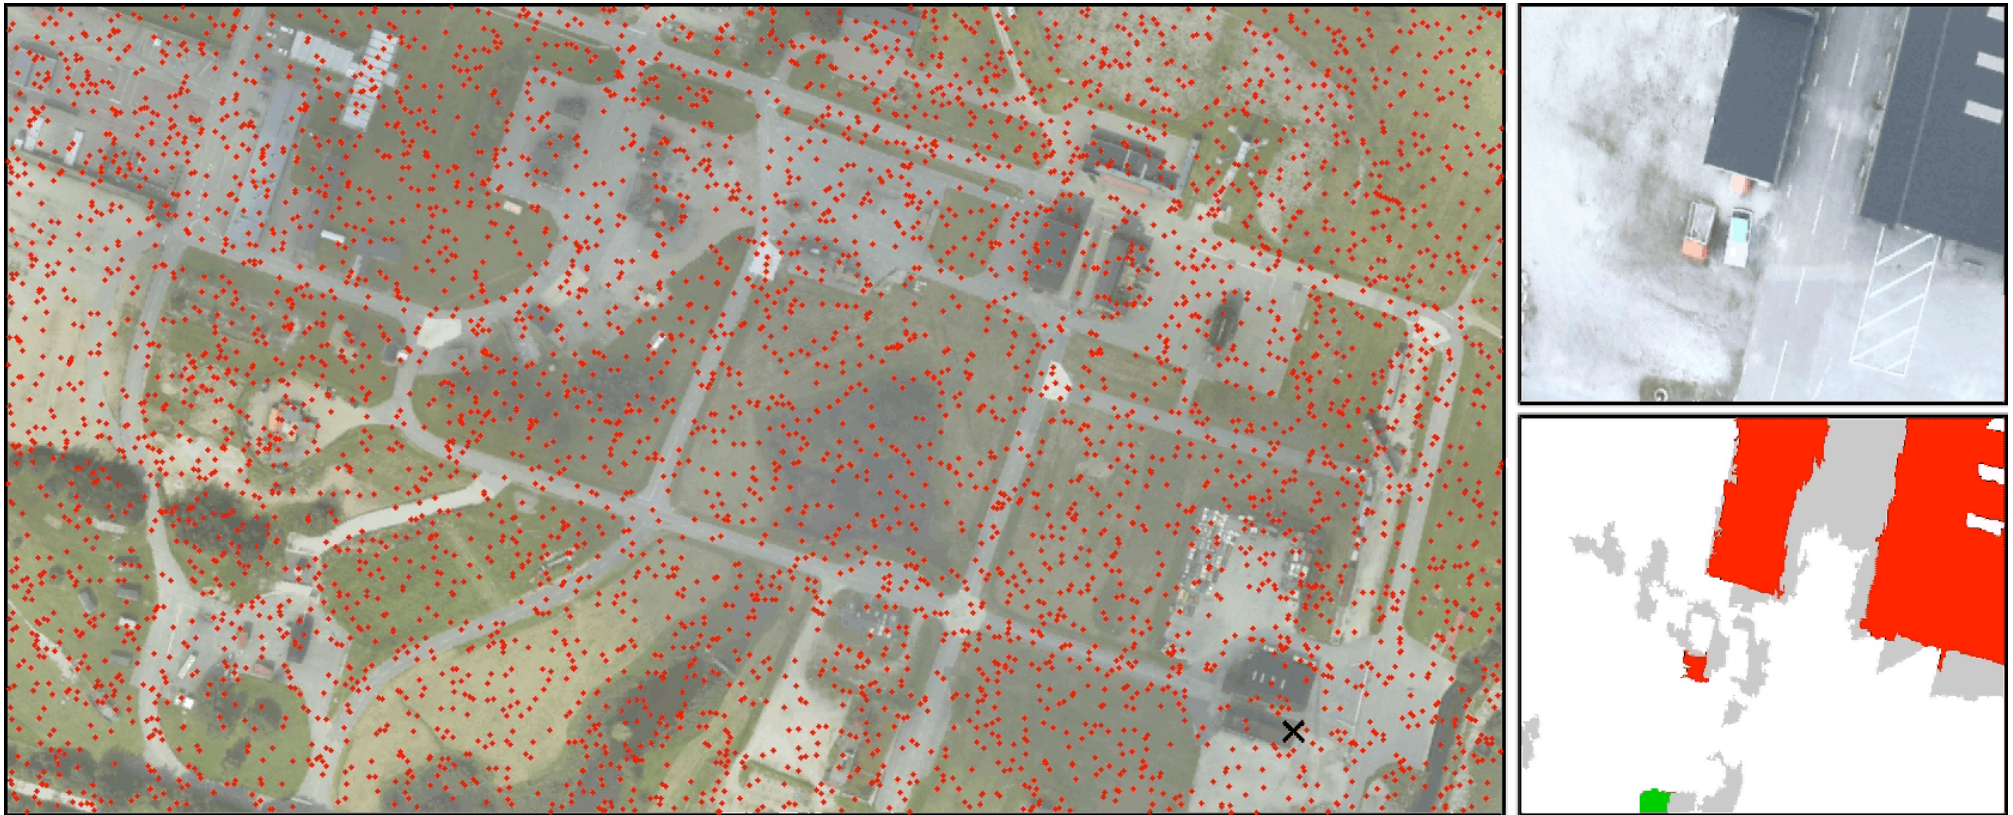
Superpixels classified as grass, asphalt or house



Three circular regions used for computing class histograms



2. Helicopter pose estimation using a map (III/III)



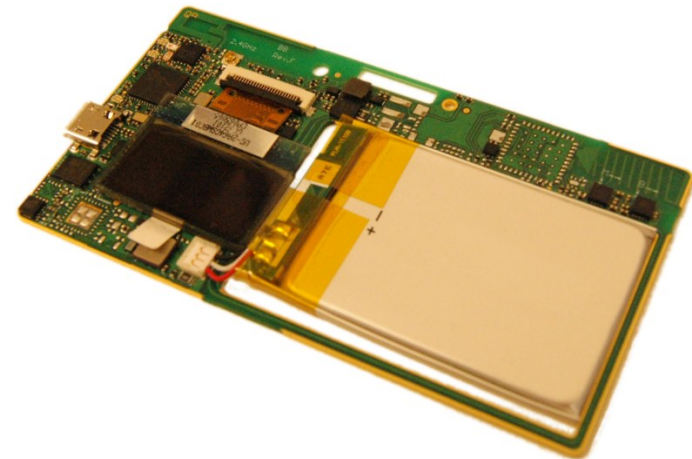
“Think of each particle as one simulation of the system state (in the movie, only the horizontal position is visualized). Only keep the good ones.”

Fredrik Lindsten, Jonas Callmer, Henrik Ohlsson, David Törnqvist, Thomas B. Schön, Fredrik Gustafsson, **Geo-referencing for UAV Navigation using Environmental Classification**. *Proceedings of the International Conference on Robotics and Automation (ICRA)*, Anchorage, Alaska, USA, May 2010.

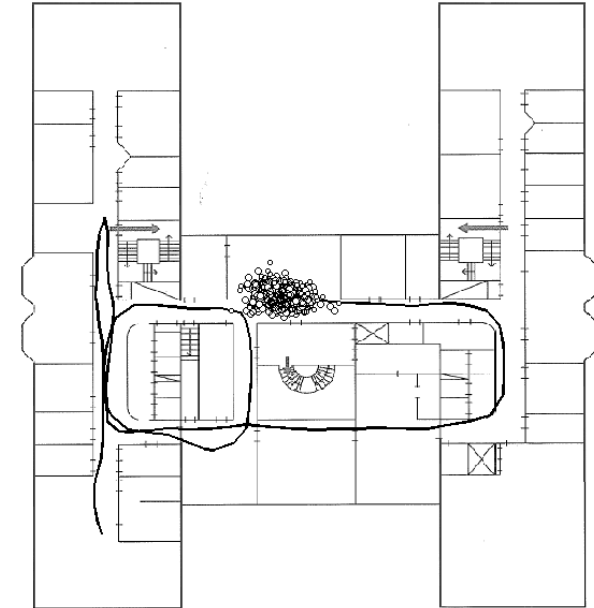
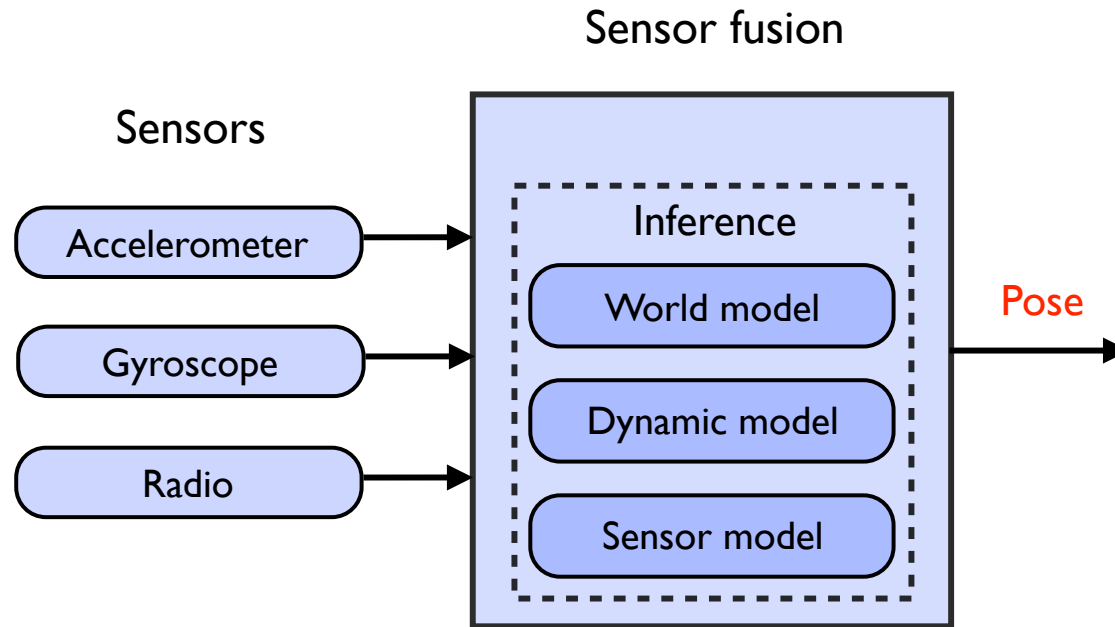
3. Indoor localization (I/III)

Aim: Compute the position of a person moving around indoors using sensors (inertial, magnetometer and radio) located in an ID badge and a map.

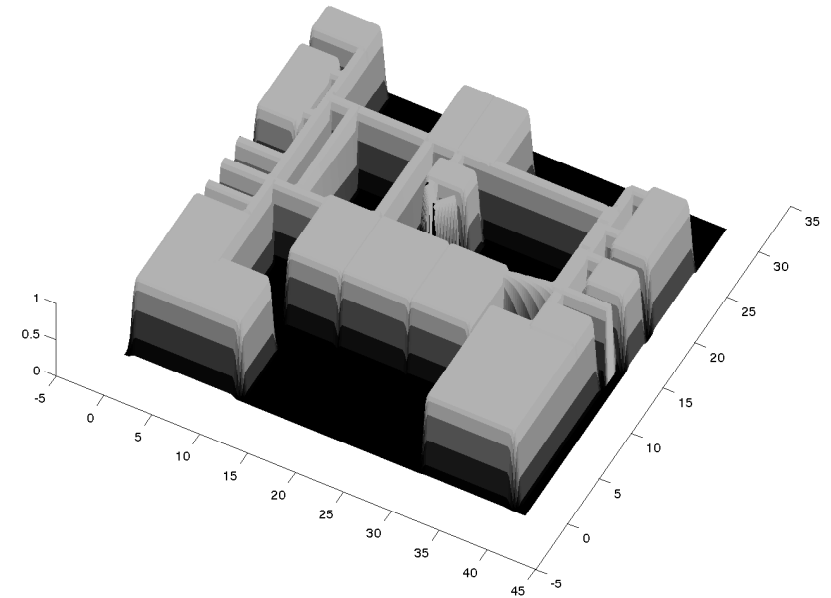
Industrial partner: Xdin



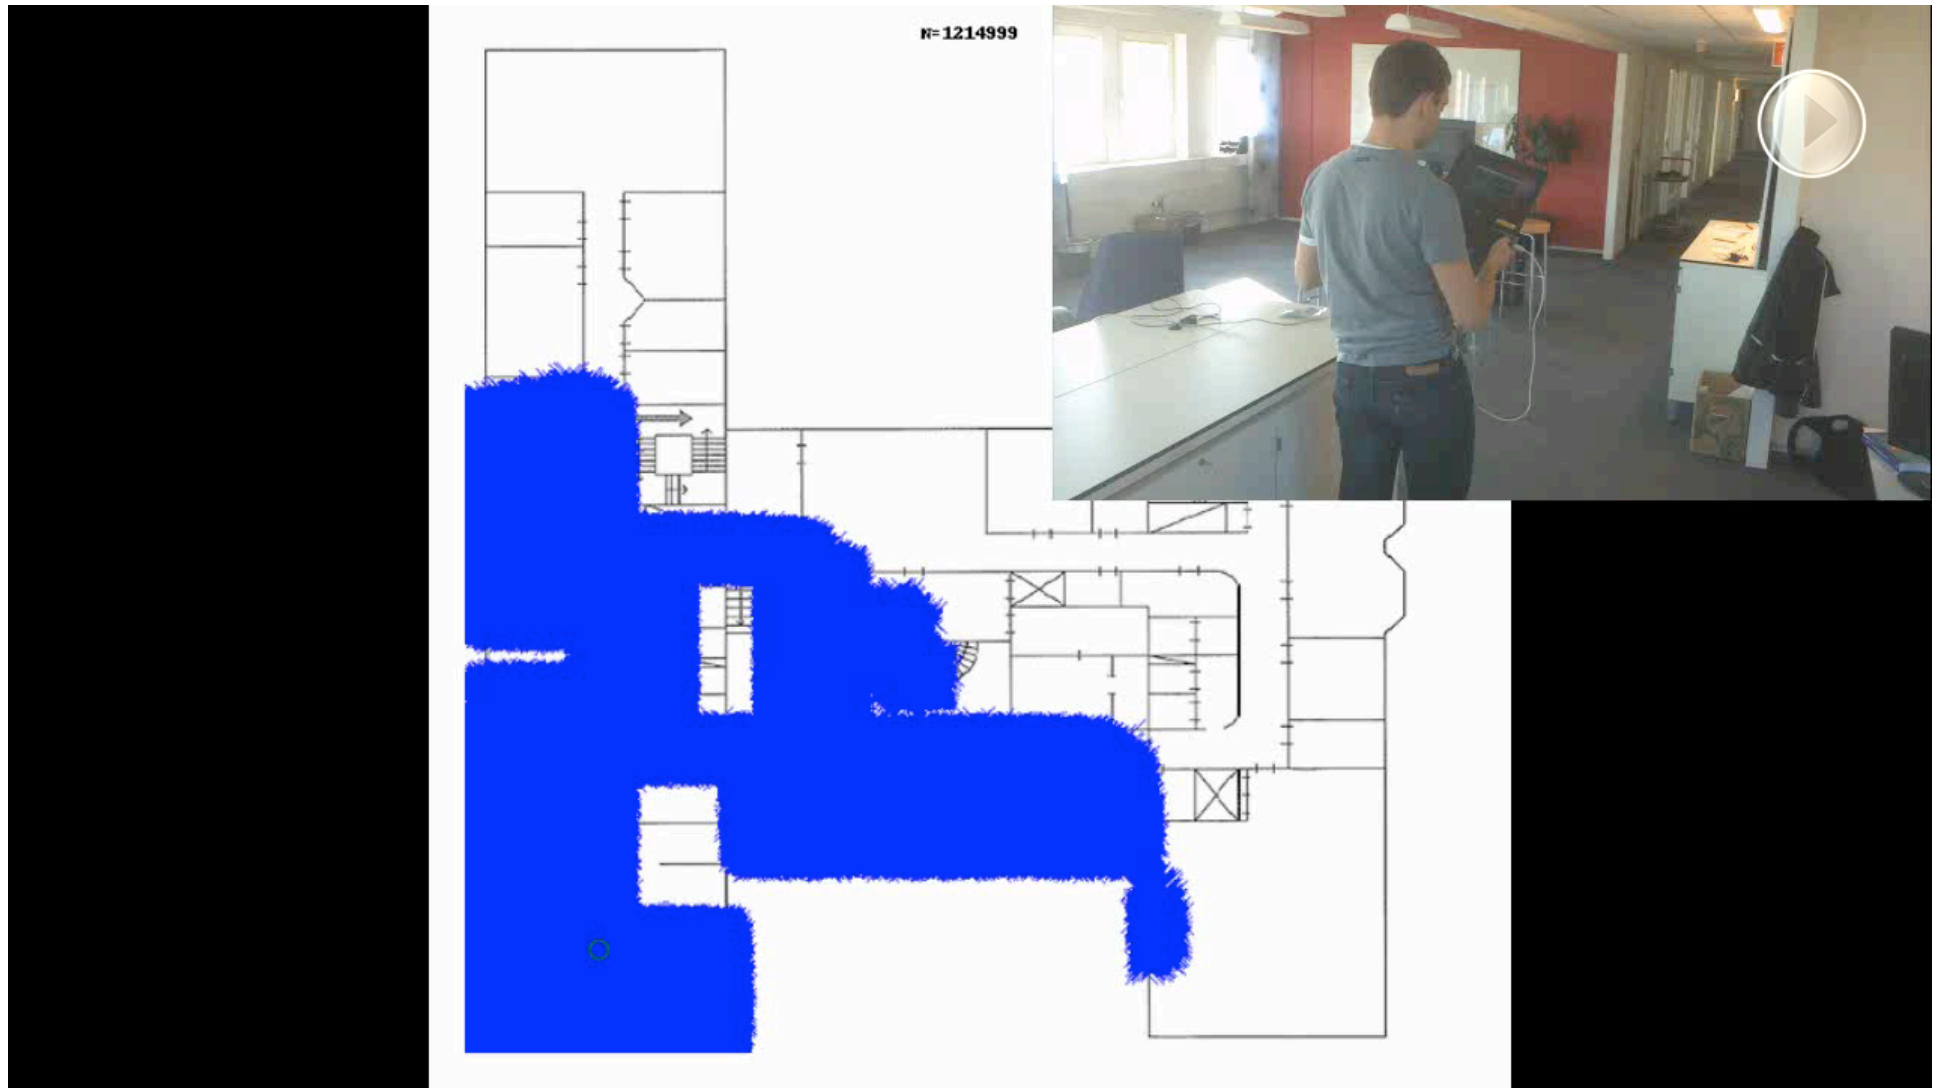
3. Indoor localization (II/III)



PDF of an office environment, the bright areas are rooms and corridors (i.e., walkable space).



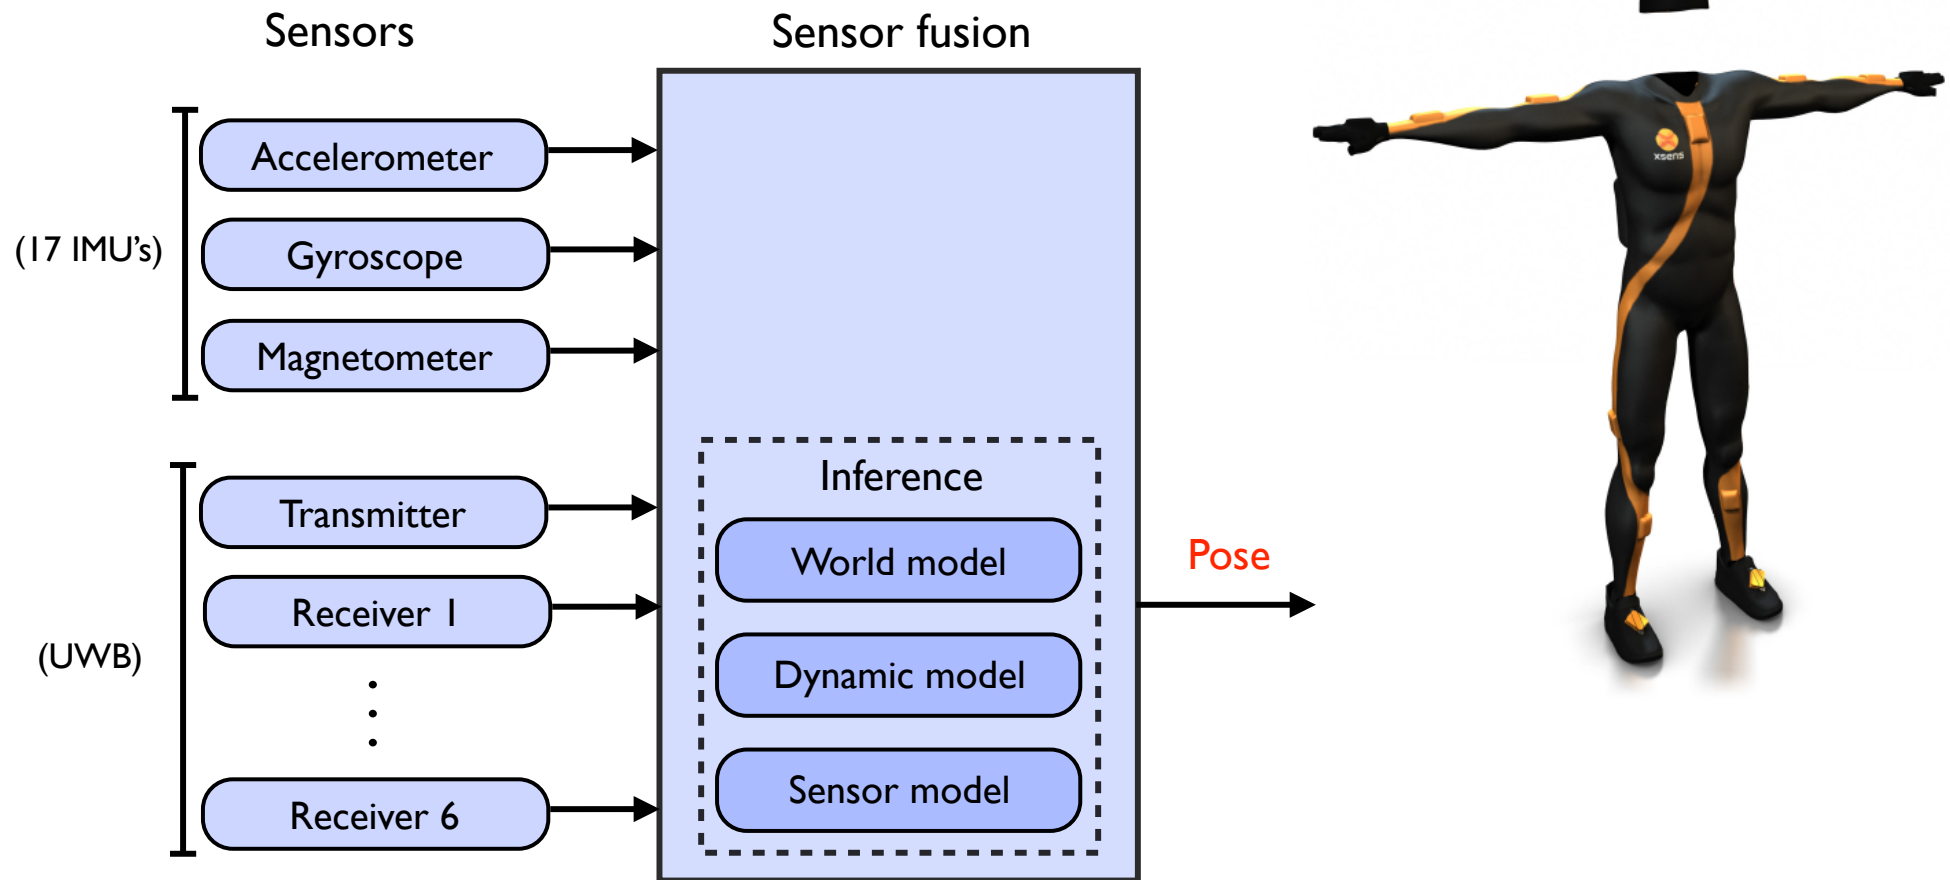
3. Indoor localization (III/III)



4. Indoor human motion estimation (I/M)

Aim: Estimate the position and orientation of a human (i.e. human motion) using measurements from inertial sensors and ultra-wideband (UWB).

Industrial partner: Xsens Technologies



4. Indoor human motion estimation (II/V)



Sensor unit integrating an IMU and a UWB transmitter into a single housing.



UWB - impulse radio using very short pulses (\sim 1ns)

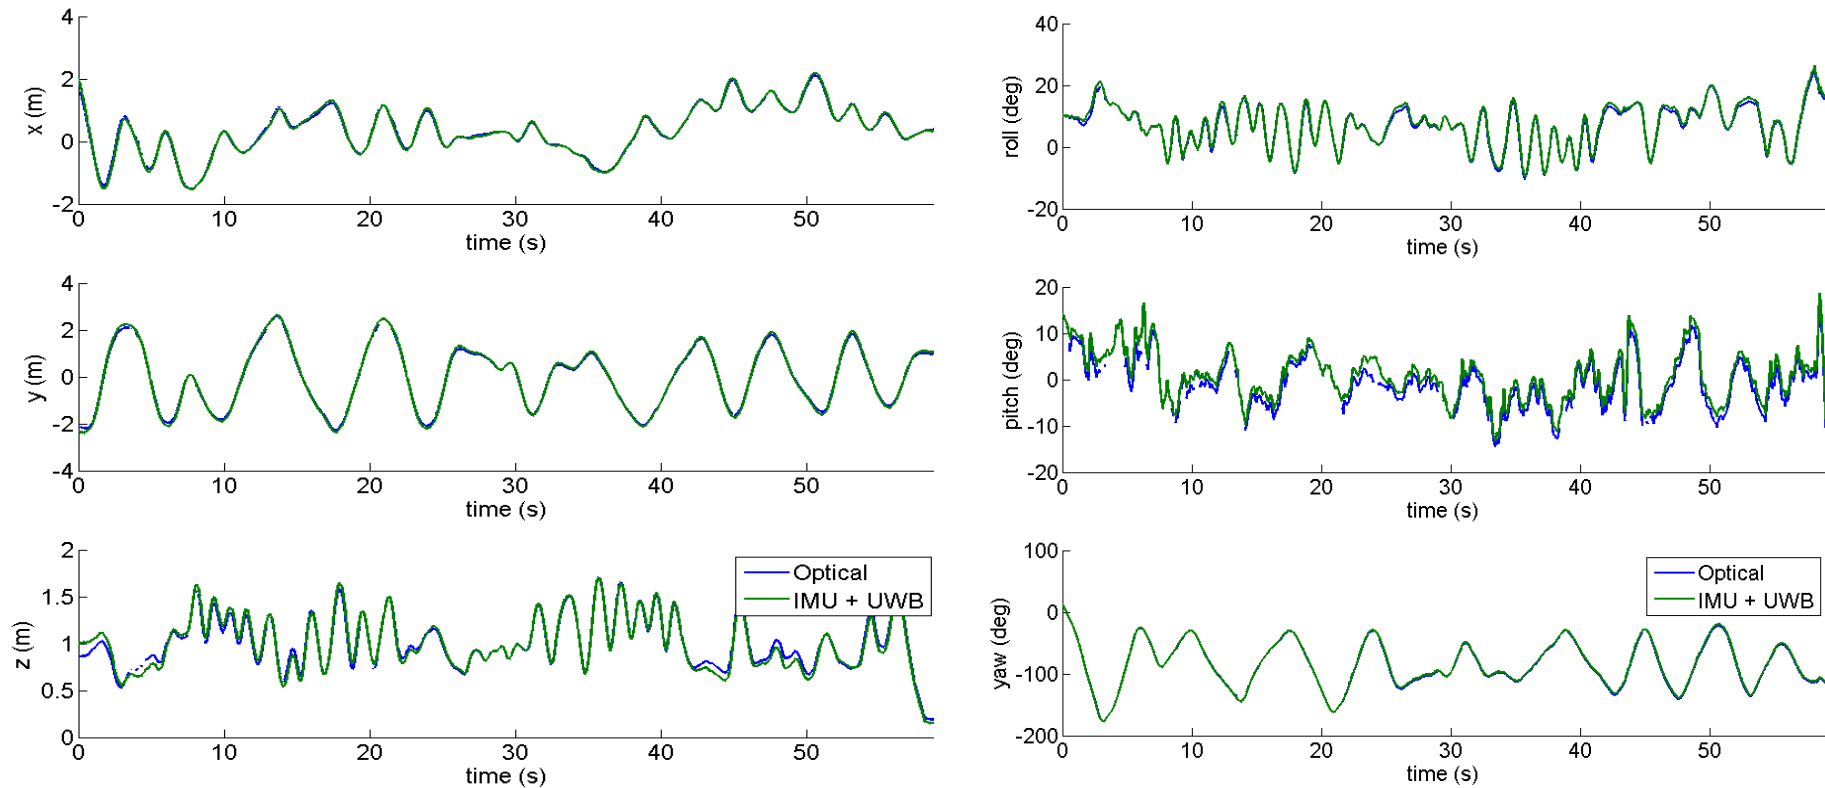
- Low energy over a wide frequency band
- High spatial resolution
- Time-of-arrival (TOA) measurements
- Mobile transmitter and 6 stationary, synchronized receivers at known positions.

- Inertial measurements @ 200 Hz
- UWB measurements @ 50 Hz

Excellent for indoor positioning



4. Indoor human motion estimation (III/V)



Performance evaluation using a camera-based reference system (Vicon).

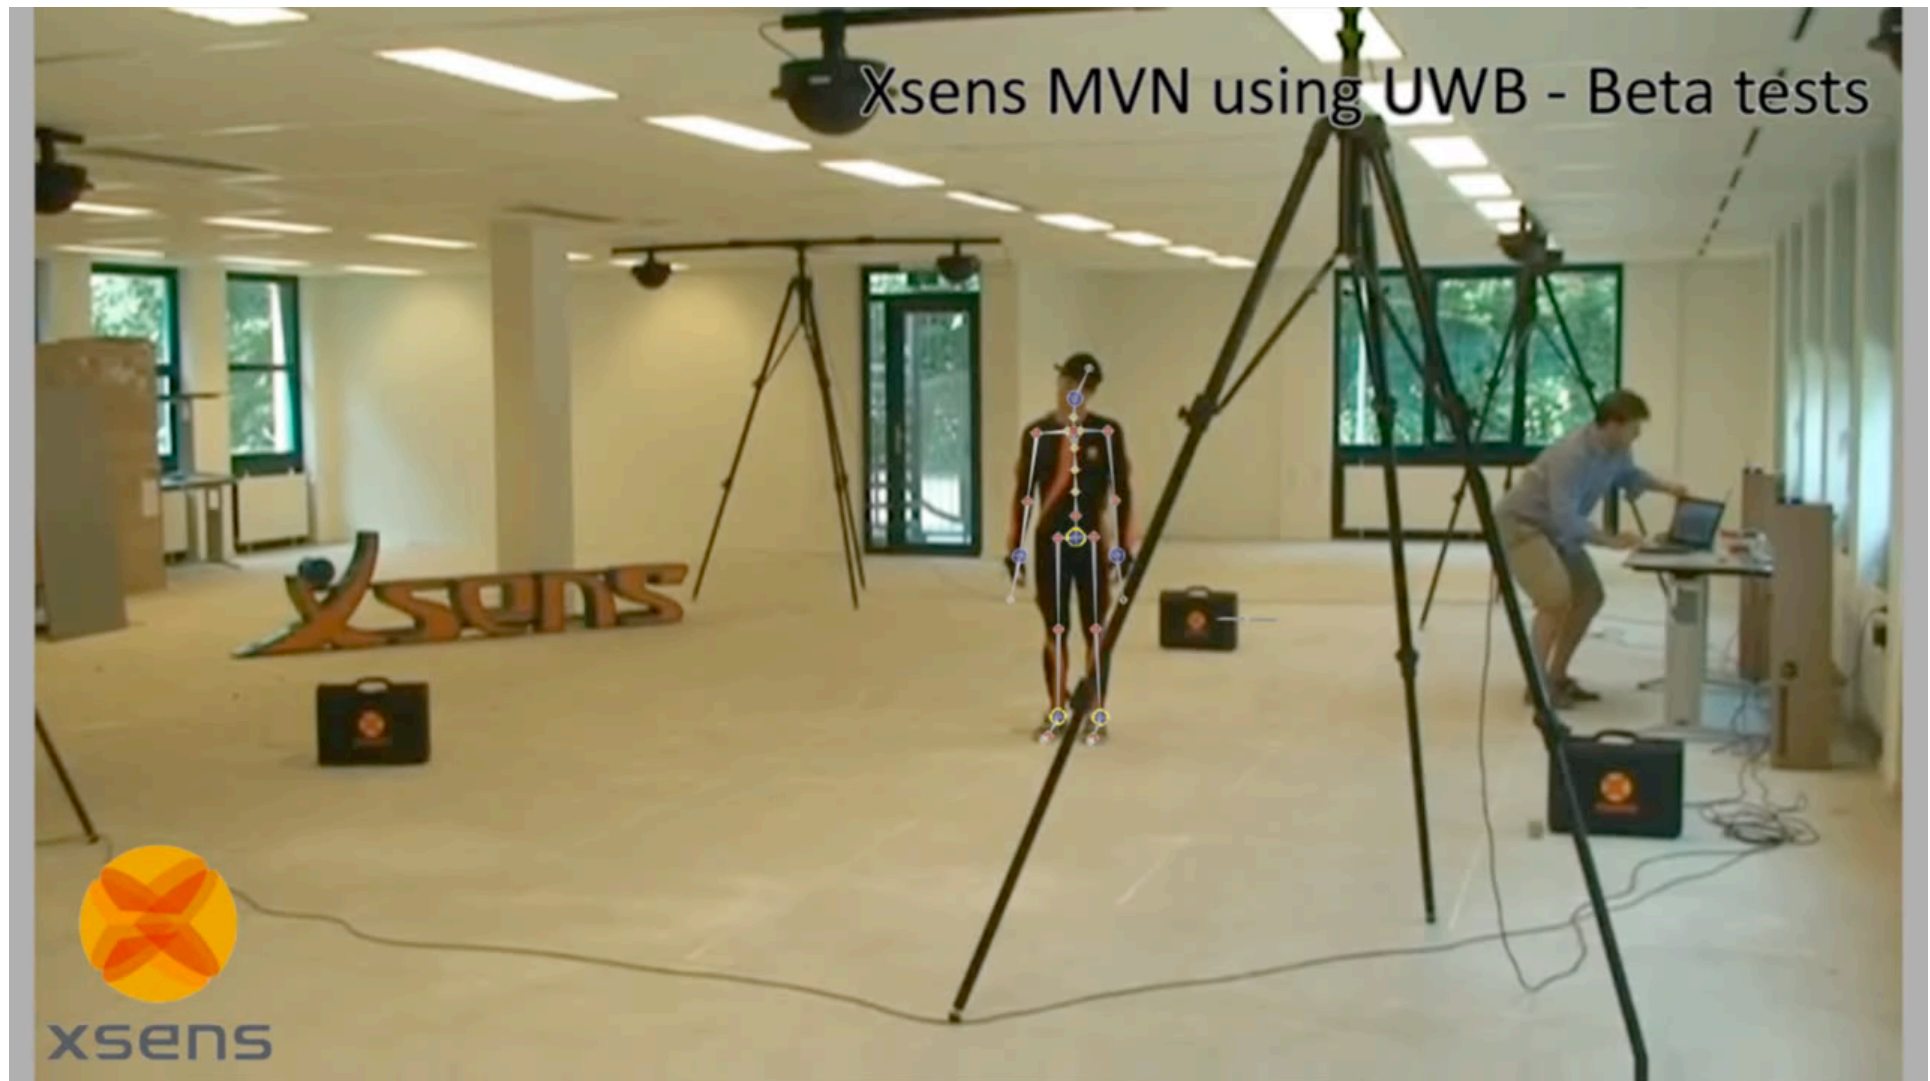
RMSE: 0.6 deg. in orientation and 5 cm in position.

Jeroen Hol, Thomas B. Schön and Fredrik Gustafsson, **Ultra-Wideband Calibration for Indoor Positioning**. *Proceedings of the IEEE International Conference on Ultra-Wideband (ICUWB)*, Nanjing, China, September 2010.

Jeroen Hol, Fred Dijkstra, Henk Luinge and Thomas B. Schön, **Tightly Coupled UWB/IMU Pose Estimation**. *Proceedings of the IEEE International Conference on Ultra-Wideband (ICUWB)*, Vancouver, Canada, September 2009.



4. Indoor human motion estimation (IV/V)



4. Indoor human motion estimation (V/V)



Quite a few different applications from different areas, all solved using the **same underlying sensor fusion strategy**

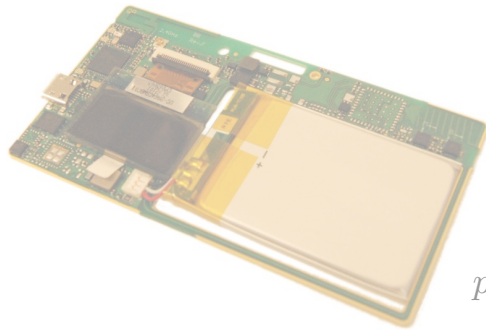
- **Model** the dynamics
- **Model** the sensors
- **Model** the world
- Solve the resulting **inference** problem

and, do not underestimate the “surrounding infrastructure”!

- There is a lot of **interesting research** that remains to be done!
- The number of available sensors is currently skyrocketing
- The **industrial utility** of this technology is **growing** as we speak!

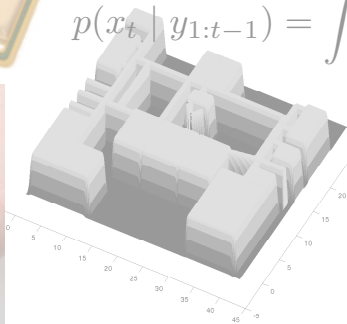


Thank you for your attention!!

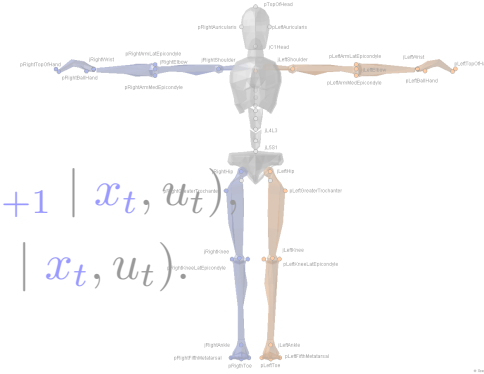


$$p(x_t | y_{1:t}) = \frac{h(y_t | x_t)p(x_t | y_{1:t-1})}{p(y_t | y_{1:t-1})},$$

$$p(x_t | y_{1:t-1}) = \int f(x_t | x_{t-1})p(x_{t-1} | y_1)$$



$$x_{t+1} | x_t \sim f_{\theta}(x_{t+1} | x_t, u_t)$$
$$y_t | x_t \sim h_{\theta}(y_t | x_t, u_t).$$



Joint work with (in alphabetical order): **Fredrik Gustafsson** (LiU), **Joel Hermansson** (Cybaero), **Jeroen Hol** (Xsens), **Johan Kihlberg** (Semcon), **Manon Kok** (LiU), **Henk Luinge** (Xsens), **Per-Johan Nordlund** (Saab), **Henrik Ohlsson** (Berkeley), **Simon Tegelid** (Xdin), **David Törnqvist** (LiU), **Niklas Wahlström** (LiU).

