

Slide from : cslu.ece.ogi.edu/publications/ps/UPF_CSLU_talk.pdf

Our reference : www.cavs.msstate.edu/~patil/particle_filter/UPF_CSLU_talk.pdf

The Unscented Particle Filter

Rudolph van der Merwe

(OGI)

Nando de Freitas

(UC Berkeley)

Arnaud Doucet

(Cambridge University)

Eric Wan

(OGI)

Outline

- **Optimal Estimation & Filtering**
- **Optimal Recursive Bayesian Solution**
- **Practical Solutions**
 - ◆ Gaussian approximations (EKF, UKF)
 - ◆ Sequential Monte Carlo methods (Particle Filters)
- **The Unscented Particle Filter**
 - ◆ The Unscented Transformation and UKF
 - ◆ Applications of UT/UKF to Particle Filters
- **Experimental Results**
- **Conclusions**

refer: Book "Sequential Monte Carlo Methods," by Arnaud Doucet, de Freitas, Gordon.

OUR ASSUMPTION will be:
MARKOVIAN, nonlinear, non-Gaussian State Space Model

Filtering

The unobserved signal [hidden state] is modelled as a Markov Process of initial distribution $p(x_0)$ and transition equation $p(x_t | x_{t-1})$.
The observations [y_t] are assumed to be conditionally independent given the process x_t and of marginal distribution $p(y_t | x_t)$.

● General problem statement

In Speech : y implies utterances, or speech data.
 x implies articulator positions, vocal tract parameters.

small y 's are instantaneous observations
capital Y 's are observation sequence.
Similarly,
for states denoted by x, X .

Observed

Unobserved

y_{k-2}

y_{k-1}

y_k

See the note at the bottom.

$$p(y_k | \mathbf{x}_k)$$

This is a nice equation - representing the probability of the current observation given the sequence. This representation is more appropriate as compared to $y = f(x)$. where in $f(x)$, would be linear or nonlinear function.

$p(y_k | x_k)$ -- emission probability. This is true as we know, y_k is a noisy observation.

Our aim is to Estimate recursively in time:
1. the posterior distribution $p(X_t | Y_{1_to_t})$
2. Its associated features [including the marginal distribution $p(x_t | Y_{1_to_t})$ - filtering distribution].
3. expectations for some function integrable w.r.t. posterior distribution.
The "some function" can be conditional mean, conditional covariance, or any other appropriate function

$p(x_k | x_{k-1})$ -- the forward transition probability.

$$p(\mathbf{x}_k | \mathbf{x}_{k-1})$$

This is a typical case, for in practice, the output is available or measurable and then we either try to estimate the input or try to model the system (estimate the transfer function of the system)

- ◆ Filtering is the problem of sequentially estimating the states (parameters or hidden variables) of a system as a set of observations become available on-line.

It is important to note that: A observation is related to the current state; in fact, the observation is not related to the previous sample or samples. It can be some how related to the previous sample(s) by the current state. Hence the equation : $y_k = f(x_k, \text{input, and measurement noise})$ on slide # 7.

Secondly, the current state is related to the previous state, input and process noise and NOT a function of observation per se. [Slide 7].

The word POSTERIOR implies measuring a certain current quantity [or sequence] after the observation(s) have been made. Herein,
 $p(\mathbf{X}_k | \mathbf{Y}_k) = p(x_k, x_{k-1}, x_{k-2}, \dots, x_0 | y_k, y_{k-1}, y_{k-2}, \dots, y_0)$
Thus, it is conditional joint probability of x_k given \mathbf{Y}_k .

Filtering

- **Solution of sequential estimation problem given by**

- ◆ Posterior density :

$$p(\mathbf{X}_k | \mathbf{Y}_k)$$

This is posterior density because we are estimating the current state after the current observation has been made. This is related to evidence and prior.

$$\mathbf{X}_k = \{ \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k \}$$

$$\mathbf{Y}_k = \{ \mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_k \}$$

- ◆ By recursively computing a marginal of the posterior, the *filtering density*,

$$p(\mathbf{x}_k | \mathbf{Y}_k)$$

Find about the last comment and have is it related to our discussion.

Posterior, Marginal. and not keeping the complete history of states.

Is it any way related to finding $P(x_{k=1} | Y_{k+1})$, and How?

one need not keep track of the complete history of the states.

Filtering

- Given the filtering density, a number of estimates of the system state can be calculated:

Still right now we are not sure of how $p(x_k|Y_k)$ is estimated.

- ◆ Mean (optimal MMSE estimate of state)

Need to understand more. left hand side should be $\text{mean}(x_k | Y_k)$. Check this!!!

$$\hat{\mathbf{x}}_k = E[x_k | \mathbf{Y}_k] = \int \mathbf{x}_k p(\mathbf{x}_k | \mathbf{Y}_k) d\mathbf{x}_k$$



Refer Slide # 3. This is conditional mean that the author is talking about. Expectation of any function - "some function".

- ◆ Mode

The mode is defined as the value with the maximum density function.

- ◆ Median

The median is defined as the value in the middle of the distribution if the values are arranged in ascending or descending order.

- ◆ Confidence intervals

The degree of certainty with which the value will lie with the threshold. This threshold is called as confidence intervals

- ◆ Kurtosis, etc.

Fourth order expectation of the value is called Kurtosis. This is not the formal definition, but i am not able to paste the formula here, but i still wanted to give some idea as to what Kurtosis is, hence the mention. For more information, refer any statistics book.

State Space Formulation of System

- General discrete-time nonlinear, non-Gaussian dynamic system

justify all the RED arrow words.

$$\mathbf{x}_k = \mathbf{f}(\mathbf{x}_{k-1}, \mathbf{u}_{k-1}, \mathbf{v}_{k-1})$$

$$\mathbf{y}_k = \mathbf{h}(\mathbf{x}_k, \mathbf{u}_k, \mathbf{n}_k)$$

state → \mathbf{x}_{k-1} in \mathbf{f}
 process noise → \mathbf{v}_{k-1} in \mathbf{f}
 known input → \mathbf{u}_k in \mathbf{h}
 measurement noise → \mathbf{n}_k in \mathbf{h}
 noisy observation → \mathbf{y}_k

The assumption is specific to the nonlinear state space models. for ref. www.cavs.msstate.edu/~patil/nonlinear_state_space_models.pdf, page 3 out of page 8, section 2.1 The Probabilistic Model

◆ Assumptions : Assumptions laid within the format of discussion. These are not the assumptions from the State Space Formulations.

1) States follow a first order Markov process

assumed that the current state is dependent only on the previous state.

$$p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{x}_{k-2}, \dots, \mathbf{x}_0) = p(\mathbf{x}_k | \mathbf{x}_{k-1})$$

The assumption validity needs to be checked for the case of nonlinear system. But, to start with the assumption offers simplicity in calculations.

2) Observations independent given the states

Where A is either one observation (any one) y_n or any set of previous observations. Y_n . Actually here n taking any value between 0 to k.

$$p(\mathbf{y}_k | \mathbf{x}_k, A) = p(\mathbf{y}_k | \mathbf{x}_k)$$

Assumption # 2 is used in finding the estimate of filtering density.

Recursive Bayesian Estimation

- Given this state space model, how do we recursively estimate the filtering density ?

Predict (update) stage

$$p(a | b) = p(b | a) p(a) / p(b)$$

$p(Y_k | x_k)$
 = $p(y_k, y_{k-1}, y_{k-2}, \dots, y_0 | x_k)$
 = club y_{k-1}, y_{k-2}, y_0 as Y_{k-1}
 = $p(y_k, Y_{k-1} | x_k)$

is $P(a,b) = P(a|b) P(b)$?

observations independent given the states. This assumption is used.

$$\begin{aligned}
 p(\mathbf{x}_k | \mathbf{Y}_k) &= \frac{p(\mathbf{Y}_k | \mathbf{x}_k) p(\mathbf{x}_k)}{p(\mathbf{Y}_k)} \\
 &= \frac{p(y_k, \mathbf{Y}_{k-1} | \mathbf{x}_k) p(\mathbf{x}_k)}{p(y_k, \mathbf{Y}_{k-1})} \\
 &= \frac{p(y_k | \mathbf{Y}_{k-1}, \mathbf{x}_k) p(\mathbf{Y}_{k-1} | \mathbf{x}_k) p(\mathbf{x}_k)}{p(y_k | \mathbf{Y}_{k-1}) p(\mathbf{Y}_{k-1})} \\
 &= \frac{p(y_k | \mathbf{Y}_{k-1}, \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{Y}_{k-1}) p(\mathbf{Y}_{k-1}) p(\mathbf{x}_k)}{p(y_k | \mathbf{Y}_{k-1}) p(\mathbf{Y}_{k-1}) p(\mathbf{x}_k)} \\
 &= \frac{p(y_k | \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{Y}_{k-1})}{p(y_k | \mathbf{Y}_{k-1})}
 \end{aligned}$$

First equation is got from the Basic Baye's Rule.

Joint Probability Distribution of current observation, observation sequence til previous observations given the current state. Similarly, the denominator consists of Joint Probability Distribution of current observation and the previous observation sequence.

Please also refer - www.cavs.msstate.edu/~patil/particle_filter/Tutorial_on_particlefilter.pdf

normalising factor or constant - depends on the likelihood function $p(y_k / x_k)$ defined in the measurement model

Recursive Bayesian Estimation

likelihood

prior



This parameter we want to estimate -

update stage :

$$p(\mathbf{x}_k | \mathbf{Y}_k) = \frac{p(\mathbf{y}_k | \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{Y}_{k-1})}{p(\mathbf{y}_k | \mathbf{Y}_{k-1})}$$

posterior

evidence

Prediction stage :

- ◆ **Prior :** $p(\mathbf{x}_k | \mathbf{Y}_{k-1}) = \int p(\mathbf{x}_k | \mathbf{x}_{k-1}) p(\mathbf{x}_{k-1} | \mathbf{Y}_{k-1}) d\mathbf{x}_{k-1}$

transition density given by *process model*



make a note!

(Propagation of past state into future before new observation is made.)

also refer eqn (3) of IEEE xactions on signal processing, feb. 02, pp 175

- ◆ **Likelihood :** defined in terms of *observation model*

- ◆ **Evidence :** $p(\mathbf{y}_k | \mathbf{Y}_{k-1}) = \int p(\mathbf{y}_k | \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{Y}_{k-1}) d\mathbf{x}_k$

Finding estimate or next observation value based on the previously available observation sequence is NOT based on MSE criterion. But, the estimate is the mostly likely [probable] value.

There can be more than one way to estimate the next observation, Evidence equation being one of them.

Called the normalising constant - also refer eqn (5) of IEEE xactions on signal processing, feb. 02, pp 175

Practical Solutions

- **Gaussian Approximations**

Kalman filters - Extended KF [EKF], Unscented KF [UKF],
Grid-based filters, Gaussian-sum filters

HMM filters are an application of approximate grid-based methods in a fixed-interval smoothing context and have been extensively used in speech processing. In HMM based tracking - Viterbi Algorithm is used to calculate the maximum posteriori estimate of the path through the trellis. Another approach is Baum-Welch Algorithm. The Viterbi and Baum-Welch Algorithms are frequently applied when the state space is approximated to be discrete. The algorithms are optimal IFF the underlying state space is truly discrete in nature.

- **Perfect Monte Carlo Simulation**

- **Sequential Monte Carlo Methods : “Particle Filters”**

Called by different name : Bootstrap Filtering, Condensation Algorithm, Particle Filtering, Interacting Particle Filtering, Survival of the fittest.
It is a technique for implementing a recursive Bayesian filter by MC simulations.

- ◆ **Bayesian Importance Sampling**

The weights are chosen using the principle of importance sampling.

- ◆ **Sampling-importance resampling (SIR)**

Gaussian Approximations

- Most common approach.
- Assume all RV statistics are *Gaussian*.
- Optimal recursive MMSE estimate is then given by



$$\hat{\mathbf{x}}_k = (\text{prediction of } \mathbf{x}_k) + \mathbf{k}_k [(\text{observation of } \mathbf{y}_k) - (\text{prediction of } \mathbf{y}_k)]$$

- Different implementations :

Basics of Kalman Filtering, Click here

- ◆ *Extended Kalman Filter* (EKF) : optimal quantities approximated via first order Taylor series expansion (linearization) of process and measurement models.
- ◆ *Unscented Kalman Filter* (UKF) : optimal quantities calculated using the Unscented Transformation (accurate to second order for any nonlinearity). Drastic improvement over EKF [Wan, van der Merwe, Nelson 2000].

for a link on use of UNSCENTED KALMAN FILTER for speech processing. CLICK here.

If the link does not work, go to www.cavs.msstate.edu/~patil/particle_filter/Ukf_final_speech.pdf

- **Problem : Gaussian approximation breaks down for most nonlinear real-world applications (multi-modal distributions, non-Gaussian noise sources, etc.)**

Find the details for the above statement!!!

Perfect Monte Carlo Simulation

- Allow for a complete representation of the posterior distribution.
- Map intractable integrals of optimal Bayesian solution to tractable discrete sums of weighted samples drawn from the posterior distribution.

$$\hat{p}(\mathbf{x}_k | \mathbf{Y}_k) = \frac{1}{N} \sum_{i=1}^N \mathbf{d}(\mathbf{x}_k - \mathbf{x}_k^{(i)})$$

independent, identically distributed random variable

$$\mathbf{x}_k^{(i)} \leftarrow \overset{I.I.D.}{p(\mathbf{x}_k | \mathbf{Y}_k)}$$

- So, any estimate of the form

$$E[f(\mathbf{x}_k)] = \int f(\mathbf{x}_k) p(\mathbf{x}_k | \mathbf{Y}_k) d\mathbf{x}_k$$

may be approximated by :

$$E[f(\mathbf{x}_k)] \approx \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_k^{(i)})$$

Particle Filters

● Bayesian Importance Sampling

Find out!

- ◆ It is often impossible to sample directly from the true posterior density.
- ◆ However, we can rather sample from a known, easy-to-sample, *proposal* distribution,

$$q(\mathbf{x}_k | \mathbf{Y}_k)$$

Is similar to importance density - $q(\mathbf{x}_k | \mathbf{Z}_k)$

and make use of the following substitution

$$\begin{aligned} E[f(\mathbf{x}_k)] &= \int f(\mathbf{x}_k) \frac{p(\mathbf{x}_k | \mathbf{Y}_k)}{q(\mathbf{x}_k | \mathbf{Y}_k)} q(\mathbf{x}_k | \mathbf{Y}_k) d\mathbf{x}_k \\ &= \int f(\mathbf{x}_k) \frac{p(\mathbf{Y}_k | \mathbf{x}_k) p(\mathbf{x}_k)}{p(\mathbf{Y}_k) q(\mathbf{x}_k | \mathbf{Y}_k)} q(\mathbf{x}_k | \mathbf{Y}_k) d\mathbf{x}_k \\ &= \int f(\mathbf{x}_k) \frac{w_k(\mathbf{x}_k)}{p(\mathbf{Y}_k)} q(\mathbf{x}_k | \mathbf{Y}_k) d\mathbf{x}_k \\ w_k(\mathbf{x}_k) &= \frac{p(\mathbf{Y}_k | \mathbf{x}_k) p(\mathbf{x}_k)}{q(\mathbf{x}_k | \mathbf{Y}_k)} \end{aligned}$$

Particle Filters

$$\begin{aligned} E[f(\mathbf{x}_k)] &= \frac{1}{p(\mathbf{Y}_k)} \int f(\mathbf{x}_k) w_k(\mathbf{x}_k) q(\mathbf{x}_k | \mathbf{Y}_k) d\mathbf{x}_k \\ &= \frac{\int f(\mathbf{x}_k) w_k(\mathbf{x}_k) q(\mathbf{x}_k | \mathbf{Y}_k) d\mathbf{x}_k}{\int p(\mathbf{Y}_k | \mathbf{x}_k) p(\mathbf{x}_k) \frac{q(\mathbf{x}_k | \mathbf{Y}_k)}{q(\mathbf{x}_k | \mathbf{Y}_k)} d\mathbf{x}_k} \\ &= \frac{\int f(\mathbf{x}_k) w_k(\mathbf{x}_k) q(\mathbf{x}_k | \mathbf{Y}_k) d\mathbf{x}_k}{\int w_k(\mathbf{x}_k) q(\mathbf{x}_k | \mathbf{Y}_k) d\mathbf{x}_k} \\ &= \frac{E_{q(\mathbf{x}_k | \mathbf{Y}_k)} [w_k(\mathbf{x}_k) f(\mathbf{x}_k)]}{E_{q(\mathbf{x}_k | \mathbf{Y}_k)} [w_k(\mathbf{x}_k)]} \end{aligned}$$

Particle Filters

- ◆ So, by drawing samples from $q(\mathbf{x}_k | \mathbf{Y}_k)$, we can approximate expectations of interest by the following:

$$\begin{aligned} E[f(\mathbf{x}_k)] &\approx \frac{\frac{1}{N} \sum_{i=1}^N w_k(\mathbf{x}_k^{(i)}) f(\mathbf{x}_k^{(i)})}{\frac{1}{N} \sum_{i=1}^N w_k(\mathbf{x}_k^{(i)})} \\ &\approx \sum_{i=1}^N \tilde{w}_k(\mathbf{x}_k^{(i)}) f(\mathbf{x}_k^{(i)}) \end{aligned}$$

- ◆ Where the normalized importance weights are given by

$$\tilde{w}_k(\mathbf{x}_k^{(i)}) = \frac{w_k(\mathbf{x}_k^{(i)})}{\sum_{j=1}^N w_k(\mathbf{x}_k^{(j)})}$$

Particle Filters

- ◆ Using the state space assumptions (1st order Markov / observational independence given state), the importance weights can be estimated recursively by [proof in De Freitas (2000)]

$$w_k = w_{k-1} \frac{p(\mathbf{y}_k | \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{x}_{k-1})}{q(\mathbf{x}_k | \mathbf{X}_{k-1}, \mathbf{Y}_k)}$$

- ◆ Problem with SIS is that the variance of the importance weights increase stochastically over time [Kong et al. (1994), Doucet et al. (1999)]
- ◆ To solve this, we need to resample the particles
 - keep / multiply particles with high importance weights
 - discard particles with low importance weights
- ◆ *Sampling-importance Resampling (SIR)*

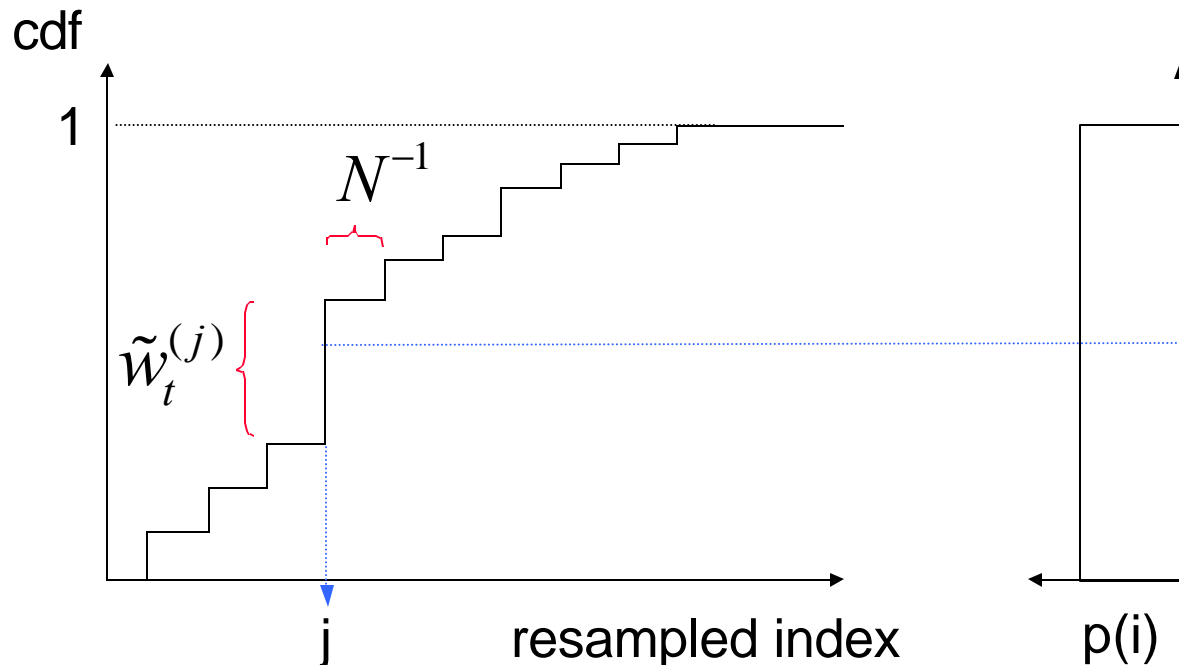
Particle Filters

- **Sampling-importance Resampling**

- ◆ Maps the N unequally weighted particles into a new set of N equally weighted samples.


$$\{\mathbf{x}_k^{(i)}, \tilde{w}_k^{(i)}\} \rightarrow \{\mathbf{x}_k^{(j)}, N^{-1}\}$$

- ◆ Method proposed by Gordon, Salmond & Smith (1993) and proven mathematically by Gordon (1994).



Particle Filters

$i=1, \dots, N=10$ particles



Particle Filters

- **Choice of Proposal Distribution**

$$w_k = w_{k-1} \frac{p(\mathbf{y}_k | \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{x}_{k-1})}{q(\mathbf{x}_k | \mathbf{X}_{k-1}, \mathbf{Y}_k)}$$

Absolutely Critical Design Step

→ critical design issue for successful particle filter

- samples/particles are drawn from this distribution
- used to evaluate importance weights

- ◆ **Requirements**

- 1) Support of proposal distribution must include support of true posterior distribution, i.e. *heavy-tailed* distributions are preferable.
- 2) Must include most recent observations.

Particle Filters

- ◆ Most popular choice of proposal distribution does not satisfy these requirements though:

$$q(\mathbf{x}_k | \mathbf{X}_{k-1}, \mathbf{Y}_k) = p(\mathbf{x}_k | \mathbf{x}_{k-1})$$

[Isard and Blake 96, Kitagawa 96, Gordon et al. 93, Beadle and Djuric 97, Avitzour 95]

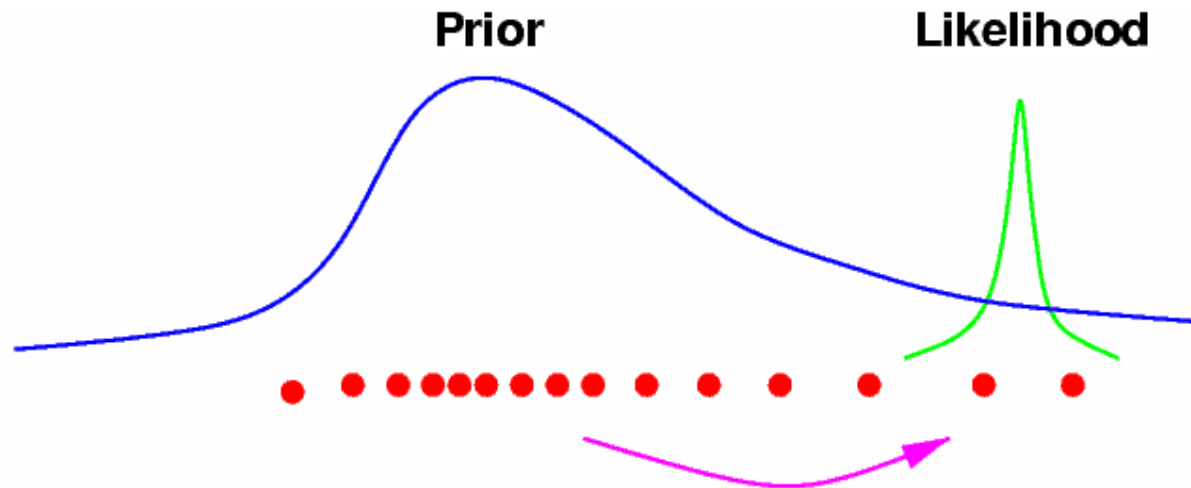
- ◆ Easy to implement :

$$\begin{aligned} w_k &= w_{k-1} \frac{p(\mathbf{y}_k | \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{x}_{k-1})}{p(\mathbf{x}_k | \mathbf{x}_{k-1})} \\ &= w_{k-1} p(\mathbf{y}_k | \mathbf{x}_k) \end{aligned}$$

- ◆ Does not incorporate most recent observation though !

Improving Particle Filters

- Incorporate New Observations into Proposal



- ◆ Use Gaussian approximation (i.e. Kalman filter) to generate proposal by combining new observation with prior

$$\begin{aligned} q(\mathbf{x}_k | \mathbf{X}_{k-1}, \mathbf{Y}_k) &= p_G(\mathbf{x}_k | \mathbf{X}_{k-1}, \mathbf{Y}_k) \\ &= \mathcal{N}(\hat{\mathbf{x}}_k, \text{cov}[\mathbf{x}_k]) \end{aligned}$$


Improving Particle Filters

- ***Extended Kalman Filter Proposal Generation***
 - ◆ De Freitas (1998), Doucet (1998), Pitt & Shephard (1999).
 - ◆ Greatly improved performance compared to standard particle filter in problems with very accurate measurements, i.e. likelihood very peaked in comparison to prior.
 - ◆ In highly nonlinear problems, the EKF tends to be very inaccurate and *underestimates* the true covariance of the state. This violates the distribution support requirement for the proposal distribution and can lead to poor performance and filter divergence.
- We propose the use of the ***Unscented Kalman Filter*** for proposal generation to address these problems !

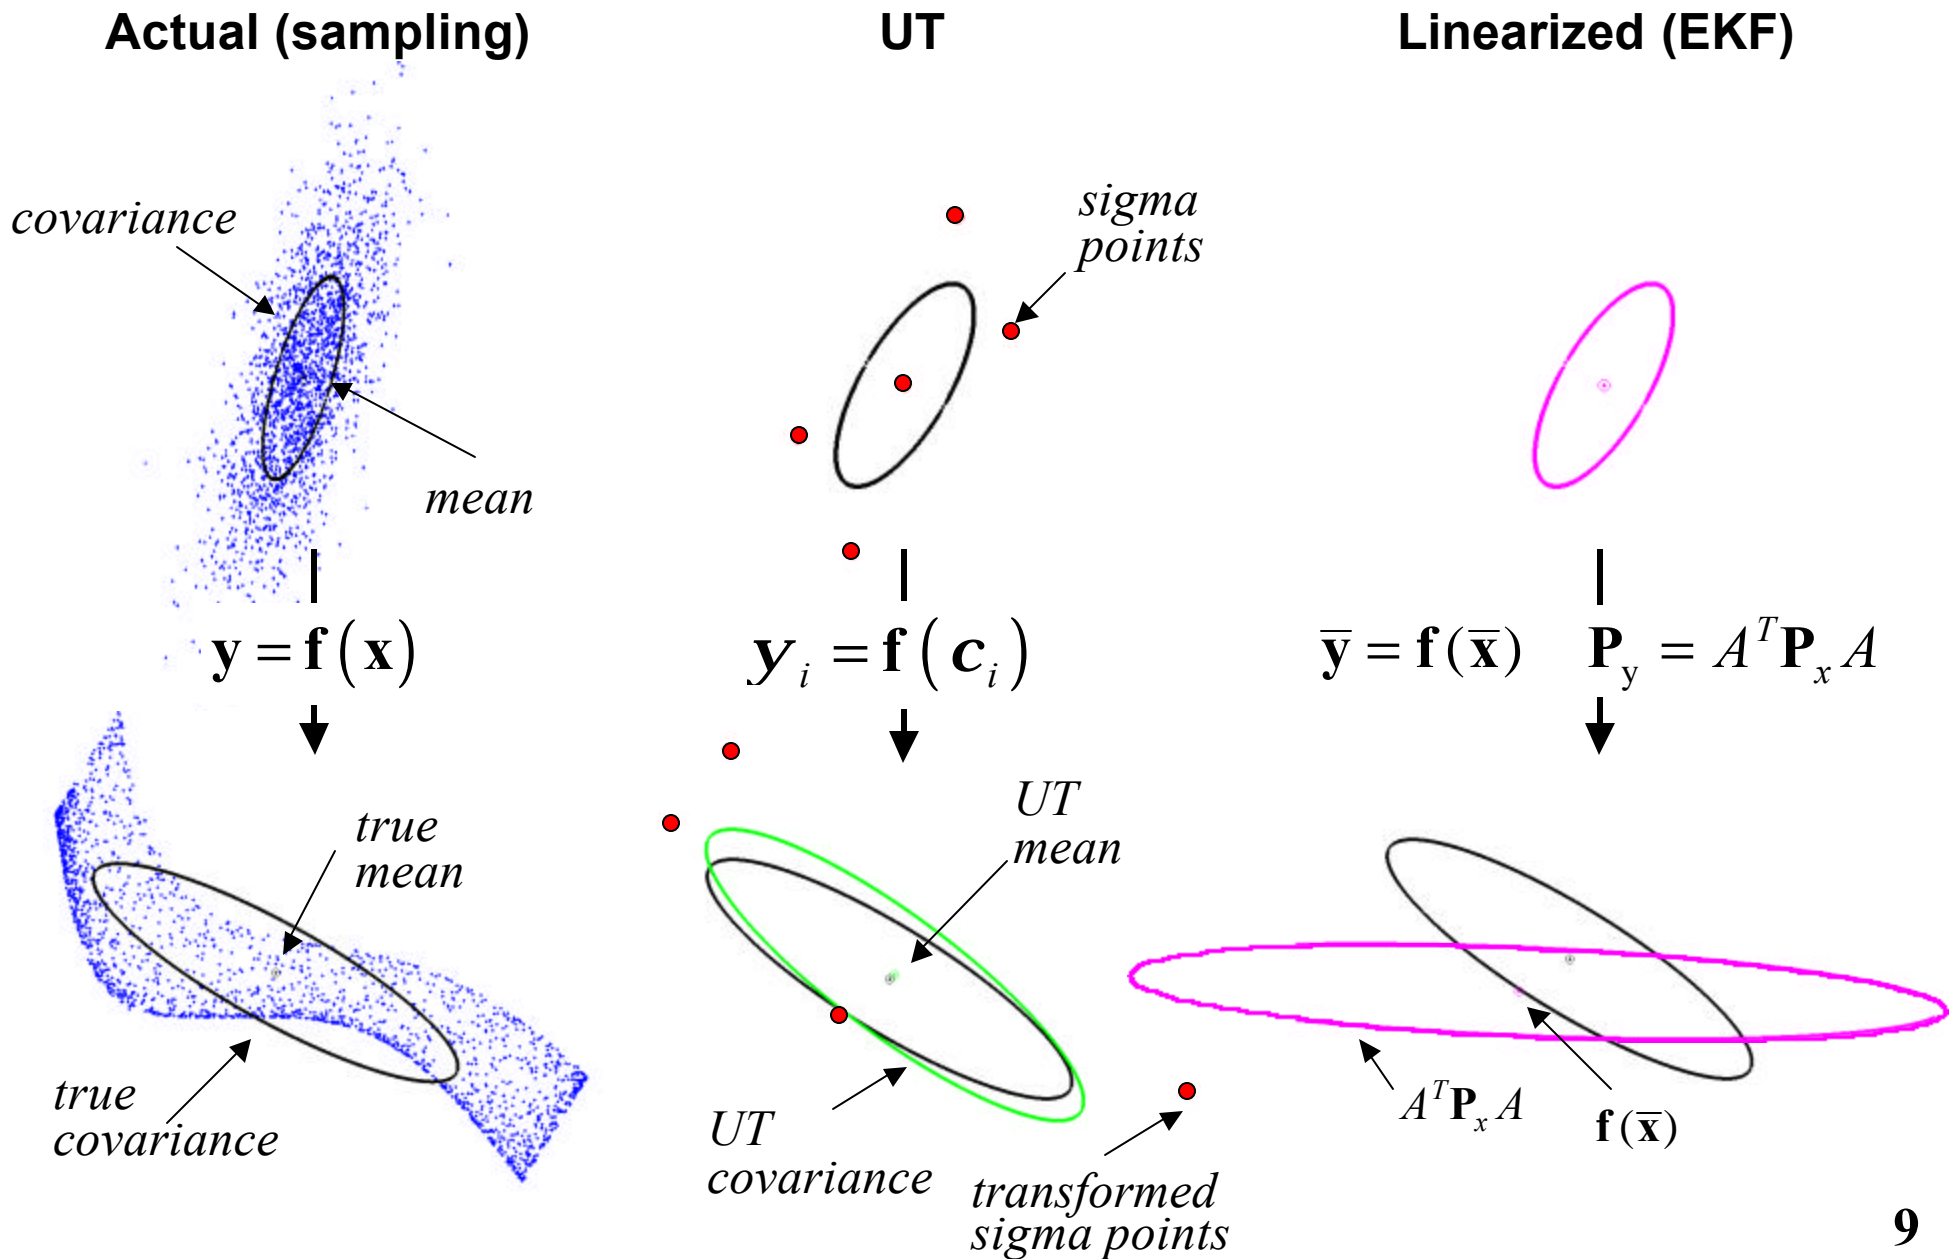
Improving Particle Filters

- ***Unscented Kalman Filter Proposal Generation***
 - ◆ UKF is a recursive MMSE estimator based on the Unscented Transformation (UT).
 - ◆ UT : Method for calculating the statistics of a RV that undergoes a nonlinear transformation (Julier and Uhlmann 1997)
 - ◆ UT/UKF : - accurate to 3rd order for Gaussians
 - higher order errors scaled by choice of transform parameters.
 - ◆ More accurate estimates than EKF (Wan, van der Merwe, Nelson 2000)
 - ◆ Have some control over higher order moments, i.e. kurtosis, etc. → heavy tailed distributions !

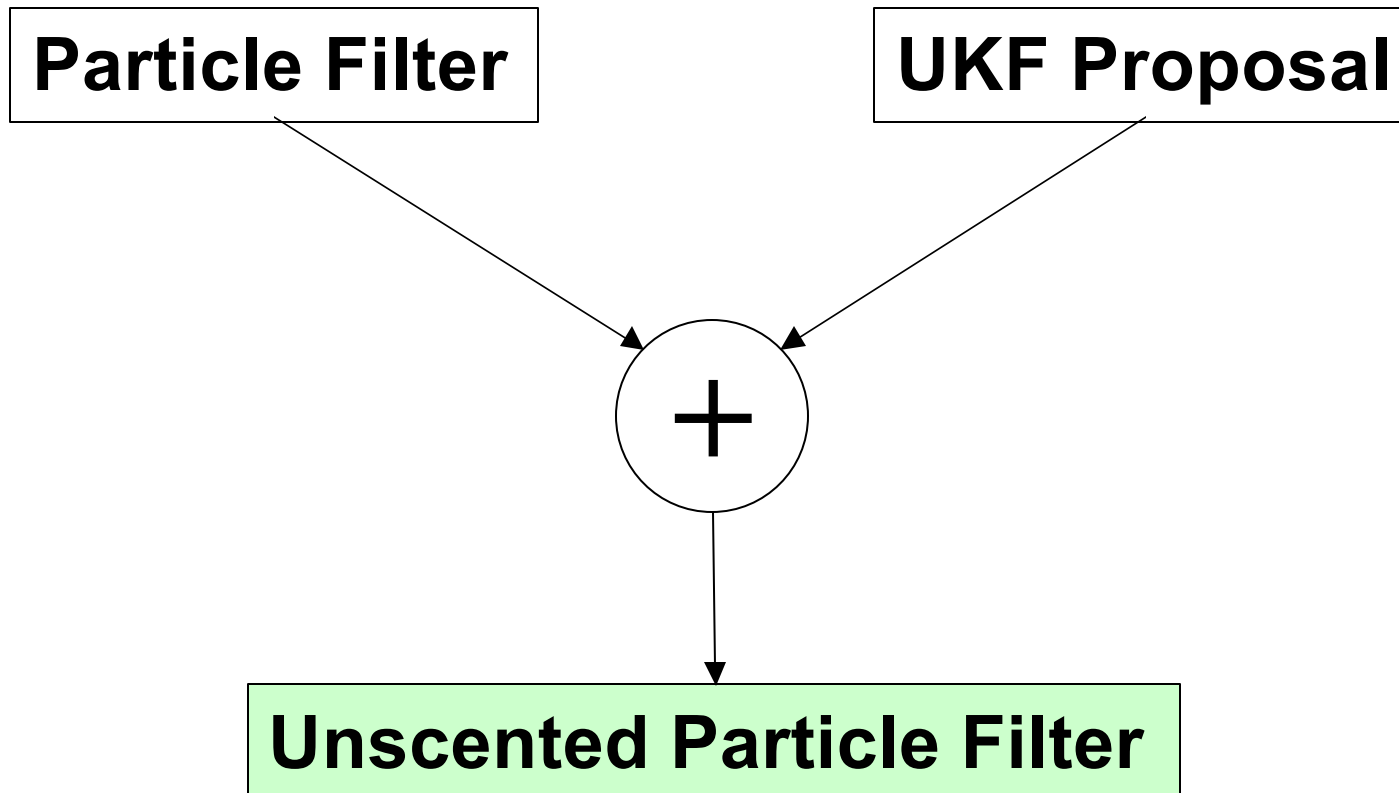
Unscented Transformation



The Unscented Transformation



Unscented Particle Filter



Experimental Results

- Synthetic Experiment

- ◆ Time-series

- process model :

$$x_{k+1} = 1 + \sin(\mathbf{w}p k) + \mathbf{f}x_k + v_k$$

process noise (Gamma) 

- nonstationary observation model :

$$y_k = \begin{cases} \mathbf{f}x_k^2 + n_k & k \leq 30 \\ \mathbf{f}x_k - 2 + n_k & k > 30 \end{cases}$$

 measurement noise (Gaussian)


Experimental Results

- **Synthetic Experiment : (100 independent runs)**

Filter	MSE	
	mean	variance
Extended Kalman Filter (EKF)	0.374	0.015
Unscented Kalman Filter (UKF)	0.280	0.012
Particle Filter : generic	0.424	0.053
Particle Filter : EKF proposal	0.310	0.016
<i>Unscented Particle Filter</i>	<i>0.070</i>	<i>0.006</i>

Experimental Results

- Synthetic Experiment



Experimental Results

● Pricing Financial Options

- ◆ Options : financial derivative that gives the holder the right (but not obligation) to do something in the future.
 - Call option : - allow holder to *buy* an underlying cash product
 - at a *specified future date* (“maturity time”)
 - for a *predetermined price* (“strike price”)
 - Put option : - allow holder to *sell* an underlying cash product
- ◆ *Black Scholes* partial differential equation
 - Main industry standard for pricing options

$$\frac{\partial f}{\partial t} + rS \frac{\partial f}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} = rf$$

risk-free interest rate

option value

volatility of cash product

value of underlying cash product

The diagram shows the Black-Scholes partial differential equation: $\frac{\partial f}{\partial t} + rS \frac{\partial f}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} = rf$. The equation is enclosed in a light green rectangular box. Red arrows point from text labels to specific parts of the equation: 'risk-free interest rate' points to the 'r' in the first term; 'option value' points to the 'f' in the second term; 'volatility of cash product' points to the 'σ' in the third term; and 'value of underlying cash product' points to the 'S' in the third term.

Experimental Results

- **Pricing Financial Options**

- ◆ Black & Scholes (1973) derived the following pricing solution:

$$C = S\mathcal{N}_c(d_1) - Xe^{-rt_m}\mathcal{N}_c(d_2)$$

$$P = -S\mathcal{N}_c(-d_1) + Xe^{-rt_m}\mathcal{N}_c(-d_2)$$

$$d_1 = \frac{\ln(S/X) + (r + \frac{1}{2}\mathbf{S}^2)t_m}{\mathbf{S}\sqrt{t_m}}$$

$$d_2 = d_1 - \mathbf{S}\sqrt{t_m}$$

$\mathcal{N}_c(\cdot)$ = cumulative normal distribution

Experimental Results

- **Pricing Financial Options**

- ◆ State-space representation to model system for particle filters

- Hidden states : r , S
- Output observations: C , P
- Known control signals: t_m , S

- ◆ Estimate call and put prices over a 204 day period on the FTSE-100 index.

- Performance : normalized square error for one-step-ahead predictions

$$NSE = \sqrt{\sum_k (y_k - \hat{y}_k)^2}$$



Experimental Results

- Options Pricing Experiment : (100 independent runs)

Option Type	Algorithm	NSE	
		mean	var
Call	Trivial	0.078	0.000
	Extended Kalman Filter (EKF)	0.037	0.000
	Unscented Kalman Filter (UKF)	0.037	0.000
	Particle Filter : generic	0.037	0.000
	Particle Filter : EKF proposal	0.092	0.508
	<i>Unscented Particle Filter</i>	<i>0.009</i>	<i>0.000</i>
Put	Trivial	0.035	0.000
	Extended Kalman Filter (EKF)	0.023	0.000
	Unscented Kalman Filter (UKF)	0.023	0.000
	Particle Filter : generic	0.023	0.000
	Particle Filter : EKF proposal	0.024	0.007
	<i>Unscented Particle Filter</i>	<i>0.008</i>	<i>0.000</i>


Experimental Results

- Options Pricing Experiment : UPF one-step-ahead predictions





Experimental Results

- Options Pricing Experiment : Estimated interest rate and volatility



Experimental Results

- Options Pricing Experiment : Probability distributions of implied interest rate and volatility



Particle Filter Demos

- *Visual Dynamics Group, Oxford. (Andrew Blake)*


Tracking agile motion



Tracking motion against camouflage



Mixed state tracking



Conclusions

- Particle filters allow for a practical but complete representation of posterior probability distribution.
- Applicable to general nonlinear, non-Gaussian estimation problems where standard Gaussian approximations fail.
- Particle filters rely on importance sampling, so the proper choice of proposal distribution is very important:
 - ◆ Standard approach (i.e. transition prior proposal) fails when likelihood of new data is very peaked (accurate sensors) or for heavy-tailed noise sources.
 - ◆ EKF proposal : Incorporates new observations, but can diverge due to inaccurate and inconsistent state estimates.
 - ◆ *Unscented Particle Filter* : UKF proposal
 - More consistent and accurate state estimates.
 - Larger support overlap, can have heavier tailed distributions.
 - Theory predicts and experiments prove **significantly better** performance.

The End