#### Intelligent Autonomous Agents Probabilistic Reasoning over Time (Hidden Markov Models, Kalman Filters)

#### Ralf Möller Hamburg University of Technology

#### Literature



Stuart Russell • Peter Norvig Prentice Hall Series in Artificial Intelligence

#### • Chapter 15

## **Dynamic Bayesian Networks**

- In addition to basic reasoning tasks, methods are needed for *learning* the transition and sensor models from observation.
- Learning can be done by inference, where inference provides an estimate of what transitions actually occurred and of what states generated the sensor readings. These estimates can be used to update the models.
- The updated models provide new estimates, and the process iterates to convergence.

### **Dynamic Bayesian Networks**

- *Learning* requires the full smoothing inference, rather than filtering, because it provides better estimates of the state of the process.
- Learning the parameters of a BN is done using Expectation – Maximization (EM) Algorithms. Iterative optimization method to estimate some unknowns parameters.

# **DBN – Special Cases**

#### • Hidden Markov Model (HMMs):

Temporal probabilistic model in which the state of the process is described by a single discrete random variable. (The simplest kind of DBN )

• Kalman Filter Models (KFMs):

Estimate the state of a physical system from noisy observations over time. Also known as linear dynamical systems (LDSs).

#### **Hidden Markov Models**

 $\mathbf{X}_t$  is a single, discrete variable (usually  $\mathbf{E}_t$  is too) Domain of  $X_t$  is  $\{1, \ldots, S\}$ 

Transition matrix 
$$\mathbf{T}_{ij} = P(X_t = j | X_{t-1} = i)$$
, e.g.,  $\begin{pmatrix} 0.7 & 0.3 \\ 0.3 & 0.7 \end{pmatrix}$ 

Sensor matrix  $O_t$  for each time step, diagonal elements  $P(e_t|X_t=i)$ e.g., with  $U_1 = true$ ,  $O_1 = \begin{pmatrix} 0.9 & 0 \\ 0 & 0.2 \end{pmatrix}$ 

Forward and backward messages as column vectors:

 $\mathbf{f}_{1:t+1} = \alpha \mathbf{O}_{t+1} \mathbf{T}^{\mathsf{T}} \mathbf{f}_{1:t}$  $\mathbf{b}_{k+1:t} = \mathbf{T} \mathbf{O}_{k+1} \mathbf{b}_{k+2:t}$ 

Forward-backward algorithm needs time  $O(S^2t)$  and space O(St)

Can avoid storing all forward messages in smoothing by running forward algorithm backwards:

$$\mathbf{f}_{1:t+1} = \alpha \mathbf{O}_{t+1} \mathbf{T}^{\mathsf{T}} \mathbf{f}_{1:t}$$
$$\mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \alpha \mathbf{T}^{\mathsf{T}} \mathbf{f}_{1:t}$$
$$\alpha'(\mathbf{T}^{\mathsf{T}})^{-1} \mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \mathbf{f}_{1:t}$$

Algorithm: forward pass computes  $\mathbf{f}_t$ , backward pass does  $\mathbf{f}_i$ ,  $\mathbf{b}_i$ 



Can avoid storing all forward messages in smoothing by running forward algorithm backwards:

$$\mathbf{f}_{1:t+1} = \alpha \mathbf{O}_{t+1} \mathbf{T}^{\mathsf{T}} \mathbf{f}_{1:t}$$
$$\mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \alpha \mathbf{T}^{\mathsf{T}} \mathbf{f}_{1:t}$$
$$\alpha'(\mathbf{T}^{\mathsf{T}})^{-1} \mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \mathbf{f}_{1:t}$$



Can avoid storing all forward messages in smoothing by running forward algorithm backwards:

$$\mathbf{f}_{1:t+1} = \alpha \mathbf{O}_{t+1} \mathbf{T}^{\top} \mathbf{f}_{1:t}$$
$$\mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \alpha \mathbf{T}^{\top} \mathbf{f}_{1:t}$$
$$\alpha'(\mathbf{T}^{\top})^{-1} \mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \mathbf{f}_{1:t}$$



Can avoid storing all forward messages in smoothing by running forward algorithm backwards:

$$\mathbf{f}_{1:t+1} = \alpha \mathbf{O}_{t+1} \mathbf{T}^{\top} \mathbf{f}_{1:t}$$
$$\mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \alpha \mathbf{T}^{\top} \mathbf{f}_{1:t}$$
$$\alpha'(\mathbf{T}^{\top})^{-1} \mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \mathbf{f}_{1:t}$$



Can avoid storing all forward messages in smoothing by running forward algorithm backwards:

$$\mathbf{f}_{1:t+1} = \alpha \mathbf{O}_{t+1} \mathbf{T}^{\mathsf{T}} \mathbf{f}_{1:t}$$
$$\mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \alpha \mathbf{T}^{\mathsf{T}} \mathbf{f}_{1:t}$$
$$\alpha'(\mathbf{T}^{\mathsf{T}})^{-1} \mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \mathbf{f}_{1:t}$$



Can avoid storing all forward messages in smoothing by running forward algorithm backwards:

$$\mathbf{f}_{1:t+1} = \alpha \mathbf{O}_{t+1} \mathbf{T}^{\top} \mathbf{f}_{1:t}$$
$$\mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \alpha \mathbf{T}^{\top} \mathbf{f}_{1:t}$$
$$\alpha'(\mathbf{T}^{\top})^{-1} \mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \mathbf{f}_{1:t}$$



Can avoid storing all forward messages in smoothing by running forward algorithm backwards:

$$\mathbf{f}_{1:t+1} = \alpha \mathbf{O}_{t+1} \mathbf{T}^{\top} \mathbf{f}_{1:t}$$
$$\mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \alpha \mathbf{T}^{\top} \mathbf{f}_{1:t}$$
$$\alpha'(\mathbf{T}^{\top})^{-1} \mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \mathbf{f}_{1:t}$$



Can avoid storing all forward messages in smoothing by running forward algorithm backwards:

$$\mathbf{f}_{1:t+1} = \alpha \mathbf{O}_{t+1} \mathbf{T}^{\top} \mathbf{f}_{1:t}$$
$$\mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \alpha \mathbf{T}^{\top} \mathbf{f}_{1:t}$$
$$\alpha'(\mathbf{T}^{\top})^{-1} \mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \mathbf{f}_{1:t}$$



Can avoid storing all forward messages in smoothing by running forward algorithm backwards:

$$\mathbf{f}_{1:t+1} = \alpha \mathbf{O}_{t+1} \mathbf{T}^{\mathsf{T}} \mathbf{f}_{1:t}$$
$$\mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \alpha \mathbf{T}^{\mathsf{T}} \mathbf{f}_{1:t}$$
$$\alpha'(\mathbf{T}^{\mathsf{T}})^{-1} \mathbf{O}_{t+1}^{-1} \mathbf{f}_{1:t+1} = \mathbf{f}_{1:t}$$



#### **Kalman Filters**

Modelling systems described by a set of continuous variables,

e.g., tracking a bird flying— $\mathbf{X}_t = X, Y, Z, \dot{X}, \dot{Y}, \dot{Z}$ .

Airplanes, robots, ecosystems, economies, chemical plants, planets, ...



Gaussian prior, linear Gaussian transition model and sensor model

#### **Updating Gaussian Distributions**

Prediction step: if  $P(\mathbf{X}_t | \mathbf{e}_{1:t})$  is Gaussian, then prediction

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

is Gaussian. If  $\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t})$  is Gaussian, then the updated distribution

 $\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t+1}) = \alpha \mathbf{P}(\mathbf{e}_{t+1}|\mathbf{X}_{t+1})\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t})$ 

is Gaussian

Hence  $\mathbf{P}(\mathbf{X}_t | \mathbf{e}_{1:t})$  is multivariate Gaussian  $N(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$  for all t

General (nonlinear, non-Gaussian) process: description of posterior grows **unboundedly** as  $t \to \infty$ 

#### Simple 1–D Example

Gaussian random walk on X-axis, s.d.  $\sigma_x$ , sensor s.d.  $\sigma_z$ 





s.d. = sacrum diis

#### **General Kalman Update**

Transition and sensor models:

 $P(\mathbf{x}_{t+1}|\mathbf{x}_t) = N(\mathbf{F}\mathbf{x}_t, \mathbf{\Sigma}_x)(\mathbf{x}_{t+1})$  $P(\mathbf{z}_t|\mathbf{x}_t) = N(\mathbf{H}\mathbf{x}_t, \mathbf{\Sigma}_z)(\mathbf{z}_t)$  *Left for further studies* 

F is the matrix for the transition;  $\Sigma_x$  the transition noise covariance H is the matrix for the sensors;  $\Sigma_z$  the sensor noise covariance

Filter computes the following update:

 $\boldsymbol{\mu}_{t+1} = \mathbf{F}\boldsymbol{\mu}_t + \mathbf{K}_{t+1}(\mathbf{z}_{t+1} - \mathbf{H}\mathbf{F}\boldsymbol{\mu}_t) \\ \boldsymbol{\Sigma}_{t+1} = (\mathbf{I} - \mathbf{K}_{t+1})(\mathbf{F}\boldsymbol{\Sigma}_t\mathbf{F}^\top + \boldsymbol{\Sigma}_x)$ 

where  $\mathbf{K}_{t+1} = (\mathbf{F} \boldsymbol{\Sigma}_t \mathbf{F}^\top + \boldsymbol{\Sigma}_x) \mathbf{H}^\top (\mathbf{H} (\mathbf{F} \boldsymbol{\Sigma}_t \mathbf{F}^\top + \boldsymbol{\Sigma}_x) \mathbf{H}^\top + \boldsymbol{\Sigma}_z)^{-1}$ is the Kalman gain matrix

 $\Sigma_t$  and  $\mathbf{K}_t$  are independent of observation sequence, so compute offline

#### 2-D Tracking: Filtering



#### 2-D Tracking: Smoothing



### Where it breaks

Cannot be applied if the transition model is nonlinear

Extended Kalman Filter models transition as locally linear around  $\mathbf{x}_t = \boldsymbol{\mu}_t$ Fails if systems is locally unsmooth



# Dynamic Bayesian Networks

 $X_t$ ,  $E_t$  contain arbitrarily many variables in a replicated Bayes net



#### **DBNs vs. HMMs**

Every HMM is a single-variable DBN; every discrete DBN is an HMM





Consider the transition model

Z

Sparse dependencies  $\Rightarrow$  exponentially fewer parameters;

e.g., 20 state variables, three parents each DBN has  $20 \times 2^3 = 160$  parameters, HMM has  $2^{20} \times 2^{20} \approx 10^{12}$ 

#### **DBNs vs. Kalman Filters**

Every Kalman filter model is a DBN, but few DBNs are KFs; real world requires non-Gaussian posteriors

E.g., where are bin Laden and my keys? What's the battery charge?



#### **Exact Inference in DBNs**

Naive method: unroll the network and run any exact algorithm



Problem: inference cost for each update grows with t

Rollup filtering: add slice t + 1, "sum out" slice t using variable elimination

# Likelihood Weighting

Set of weighted samples approximates the belief state



LW samples pay no attention to the evidence!

- $\Rightarrow$  fraction "agreeing" falls exponentially with t
- $\Rightarrow$  number of samples required grows exponentially with t



## **Particle Filtering**

Basic idea: ensure that the population of samples ("particles") tracks the high-likelihood regions of the state-space

Replicate particles proportional to likelihood for  $\mathbf{e}_t$ 



Widely used for tracking nonlinear systems, esp. in vision

Also used for simultaneous localization and mapping in mobile robots  $10^5$ -dimensional state space

# Particle Filtering (cntd.)

Assume consistent at time t:  $N(\mathbf{x}_t | \mathbf{e}_{1:t}) / N = P(\mathbf{x}_t | \mathbf{e}_{1:t})$ 

Propagate forward: populations of  $\mathbf{x}_{t+1}$  are

 $N(\mathbf{x}_{t+1}|\mathbf{e}_{1:t}) = \sum_{\mathbf{x}_t} P(\mathbf{x}_{t+1}|\mathbf{x}_t) N(\mathbf{x}_t|\mathbf{e}_{1:t})$ 

Weight samples by their likelihood for  $e_{t+1}$ :

 $W(\mathbf{x}_{t+1}|\mathbf{e}_{1:t+1}) = P(\mathbf{e}_{t+1}|\mathbf{x}_{t+1})N(\mathbf{x}_{t+1}|\mathbf{e}_{1:t})$ 

Resample to obtain populations proportional to W:

$$N(\mathbf{x}_{t+1}|\mathbf{e}_{1:t+1})/N = \alpha W(\mathbf{x}_{t+1}|\mathbf{e}_{1:t+1}) = \alpha P(\mathbf{e}_{t+1}|\mathbf{x}_{t+1})N(\mathbf{x}_{t+1}|\mathbf{e}_{1:t})$$
  
$$= \alpha P(\mathbf{e}_{t+1}|\mathbf{x}_{t+1})\Sigma_{\mathbf{x}_t}P(\mathbf{x}_{t+1}|\mathbf{x}_t)N(\mathbf{x}_t|\mathbf{e}_{1:t})$$
  
$$= \alpha' P(\mathbf{e}_{t+1}|\mathbf{x}_{t+1})\Sigma_{\mathbf{x}_t}P(\mathbf{x}_{t+1}|\mathbf{x}_t)P(\mathbf{x}_t|\mathbf{e}_{1:t})$$
  
$$= P(\mathbf{x}_{t+1}|\mathbf{e}_{1:t+1})$$

#### **Particle Filtering: Performance**

Approximation error of particle filtering remains bounded over time, at least empirically—theoretical analysis is difficult



#### **Summary**

Temporal models use state and sensor variables replicated over time

Markov assumptions and stationarity assumption, so we need

- transition model  $\mathbf{P}(\mathbf{X}_t | \mathbf{X}_{t-1})$
- sensor model  $\mathbf{P}(\mathbf{E}_t|\mathbf{X}_t)$

Tasks are filtering, prediction, smoothing, most likely sequence; all done recursively with constant cost per time step

Hidden Markov models have a single discrete state variable; used for speech recognition

Kalman filters allow n state variables, linear Gaussian,  ${\cal O}(n^3)$  update

Dynamic Bayes nets subsume HMMs, Kalman filters; exact update intractable

Particle filtering is a good approximate filtering algorithm for DBNs