PROBABILISTIC MODELS

Gerhard Widmer

Institute for Computational Perception
Johannes Kepler University
Linz, Austria
gerhard.widmer@jku.at

JKU
JOHANNES KEPLER
UNIVERSITY LINZ

Department of
Computational
Perception

June 4, 2013

Example 1 – Automatic Diagnosis of Complex Systems

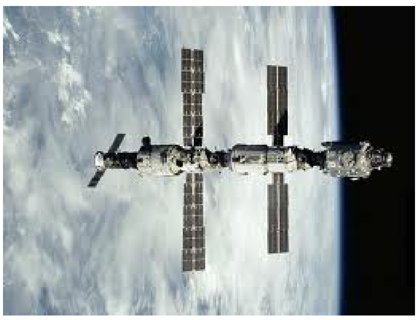

Given:

- A model of a spacecraft
- Sensor signals indicating certain errors and other conditions

Task:

- Identify most likely fault that could be the cause of this
- Evaluate what would happen if certain repair actions were taken
- Decide on measures to fix the problem.

Example 2 – Autonomous Driving

Given:

- A road map
- Streams of sensor signals (GPS, cameras, radar, infrared range sensors, ...)

Task:

- Autonomously drive a car through difficult terrain
- in the presence of other cars, obstacles, ...

Example 3 – Automatic Aircraft Tracking

Given:

- Radar or video signals (intermittent, sometimes occluded)

Task:

- Track aircraft over time
- Determine how many aircraft there are
- Keep track of which is which
- Predict where aircraft will be in 0.1 / 1 / 10 sec.

Example 4 – Speech Recognition and Voice Control

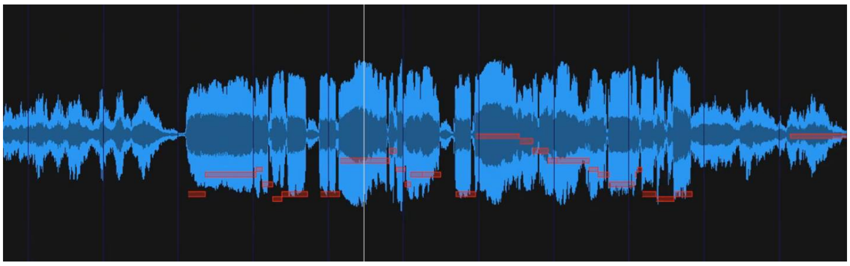

Given:

- Person speaking into a microphone (audio)

Task:

- Segment audio stream into separate words
- Understand / transcribe text that is spoken
- Use speech commands to control complex devices.

Example 5 – Real-time Music Following

Given:

- Live music performance on stage (audio stream through mike)
- Printed sheet music (score)

Task:

- Automatically track performers' position in score
- Synchronise live music with other events (e.g., page turning)

What These Scenarios Have in Common

An autonomous system (robot, car, repair system, tracker, ...)

- Receives information about its environment, at regular or irregular intervals
- Must identify objects, classify situations, make predictions, take decisions
- **Information may be incomplete, uncertain, partly wrong and contradictory**

☞ **Logic-based models or deterministic approaches not possible.**

What the Shown Systems Have in Common

They work with PROBABILISTIC (GRAPHICAL) MODELS.

Probabilistic Models

- Represent knowledge about the world, and about its uncertainty
- Support logical and probabilistic inference (decision-making)
- Can be learned from example observations
- Technically: Are graphical representations of factored probability distributions.

What You Will Learn

SEMANTICS: What types of models there are, and what they represent

- Bayesian Networks
- Temporal Models: Dynamic Bayes Nets, HMMs, Kalman Filters
- Briefly: Undirected and Semi-directed Models

INFERENCE: How to use models to answer questions

- Deterministic algorithms
- Probabilistic algorithms (stochastic sampling)

LEARNING: How to learn models from data

- Parameter learning
- Structure learning

What the Math Will Look Like

Standard Inference Algorithm for BNs

1. **Compute joint probability** $P(\boldsymbol{X}, \boldsymbol{e})$:

$$\boxed{P(\boldsymbol{X}, \boldsymbol{e}) = \sum_{\boldsymbol{z} \in Val(\boldsymbol{Z})} P(\boldsymbol{X}, \boldsymbol{e}, \boldsymbol{z})}$$

where $\boldsymbol{Z} = \mathcal{X} - \boldsymbol{X} - \boldsymbol{E}$
and each $P(\boldsymbol{x}, \boldsymbol{e}, \boldsymbol{z})$ is computed directly from the BN:

$$\boxed{P(\boldsymbol{x}, \boldsymbol{e}, \boldsymbol{z}) = \prod_{x_i \in \boldsymbol{x} \cup \boldsymbol{e} \cup \boldsymbol{z}} P(x_i | \mathrm{pa}_{X_i}).}$$

2. **Conditioning via Renormalisation:**

$$\boxed{P(\boldsymbol{X} \mid \boldsymbol{e}) = \frac{1}{Z} P(\boldsymbol{X}, \boldsymbol{e})} \quad \text{where } Z = P(\boldsymbol{e}) = \sum_{\boldsymbol{x}} P(\boldsymbol{x}, \boldsymbol{e})$$

... and sometimes we will meet our good old friend

... the Gaussian probability density function:

$$\boxed{p(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}}}$$

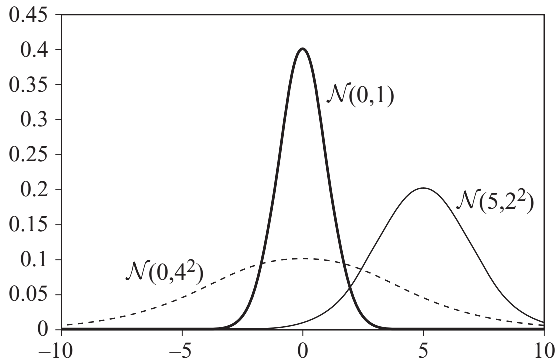

What You Will Know After This Class

You will understand the computational principles behind our on-line, probabilistic, particle-filtering music tracker:

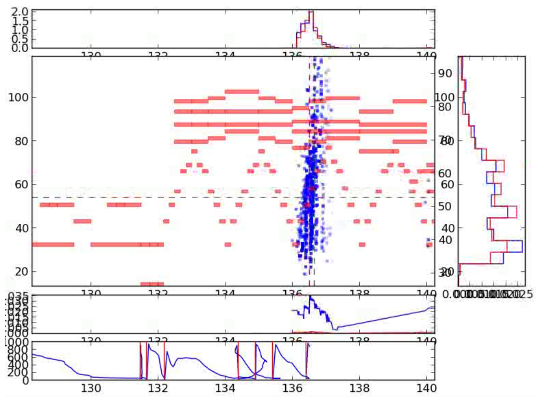

... and generally how to model and reason about complex systems.

Literature

Koller, Daphne and Friedman, Nir (2009).
Probabilistic Graphical Models: Principles and Techniques. Cambridge, MA: MIT Press.

Russell, Stuart J. and Norvig, Peter (2003).
Artificial Intelligence: A Modern Approach. Englewood Cliffs, NJ: Prentice Hall.

Doucet, A. and Johansen, A.M. (2008).
A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later. In: D. Crisan & B. Rozovskii (eds), *The Oxford Handbook of Nonlinear Filtering.* Oxford University Press.

Korzeniowski, F., Arzt, A., Krebs, F. and Widmer, G. (2013).
Tracking Rests and Tempo Changes: Improved Score Following with Particle Filters. In *Proceedings of the 40th International Computer Music Conference (ICMC 2013)*, Perth, Australia.

Rabiner, Lawrence E. (1989).
A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proceedings of the IEEE* 77(2), 257 - 286.

van der Merwe, R., Doucet, A., de Freitas, N, and Wan, E. (2000).
The Unscented Particle Filter. In *Proceedings of the 14th International Conference on Neural Information Processing Systems (NIPS).*