Learning to Control – an Agenda

Carl Edward Rasmussen
Computational and Biological Learning Lab
Department of Engineering
University of Cambridge

Adaptation and Learning in Autonomous Systems
Lund Center for Control of Complex Engineering Systems
April 21-23, 2010

Desiderata

Biological Learning is Fast and Flexible

- learning must be *model-based*
- models must *flexible*
- inference must be *efficient*

Probabilistic Models and Inference

Systems which rely on *experience* will always have some uncertainty associated with any prediction

- use probabilistic models

Probabilistic models capture *all types* of uncertainty

- inherent stochasticity
- measurement noise
- model uncertainty

To *reason efficiently* about past experience

- probabilistic, Bayesian inference
- principled framework
- forces you to be explicit about your assumptions
- exact computations may be intractable

Parametric vs Non-parametric Models

Different kinds of learning

- sometimes we're unsure about the *value* of some parameter
- more typically, we're unsure about *functional relationships*

Non-parametric models

- don't have a fixed parametric structure
- don't have a finite number of parameters
- automatically adapt their complexity to the observed data (Occam's Razor)

Gaussian Process Models

A Gaussian Process (GP) is a distribution over functions.

GPs are flexible, non-parametric, fully probabilistic Bayesian kernel machines where inference can be done in closed form.

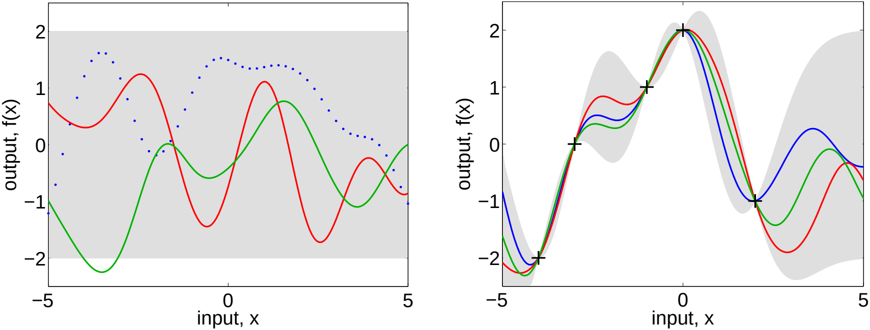

Understanding Control as Learning

Key idea: Bayesian inference provides as flexible and principled approach to control.

Traditional approaches first do *identification* of a dynamical system based on

1. simplifying assumptions
2. measurements

Then design a controller

My approach Don't make (parametric) assumptions. Use a stochastic model, taking into account uncertainties due to *noise* and *lack of knowledge*.
Controller based on predicted performance, *integrating* over all forms of uncertainty.

Short and Long Prediction Horizons

Typically, there is a dilemma concerning prediction horizons

- it is only feasible to learn short time dynamics
- good control requires the understanding of long term consequences

To resolve this we learning short time dynamics, then

- probabilistically, cascade many short term predictions to get long term consequences

Initially, this will typically cause rapidly rising uncertainties

The Learning Procedure

Repeatedly:

- Observe the behavior of the dynamical system, fit stochastic short term dynamics model
- *Probabilistically*, cascade short term predictions, to predict long term behavior
- Optimize the *simulated* behaviour wrt the controller.
- Apply the control law, record additional data

What priors did I use?

The prior information was

- short term dynamics are
  - *smooth*
  - *time invariant*
- a time scale: eigen frequency about 2 Hz
  - time discretisation 100 ms
  - horizon 2.5 s
- an error scale
  - 30 degrees is a 'moderate' error

Conclusions

- Learning is a powerful paradigm in
  - biology
  - can be exploited in engineering
- Fast learning from weak prior knowledge is possible and advantageous
  - avoid simplifying (parametric) assumptions
  - avoid deterministic ‘model identification’
- Probabilistic Inference and Stochastic models are ideally suited for learning
- Implications for
  - understanding biological systems
  - engineering control systems, eg robotics