

Research Programme on Machine Learning

Chris Williams
School of Informatics

October 2008

What is Machine Learning?


- ▶ *Supervised learning*: predict y from \mathbf{x} , or model $p(y|\mathbf{x})$.
Examples:
 - ▶ Predict speedup of a program y given information on compiler transformation sequence \mathbf{x} (O'Boyle)
 - ▶ Predict German sentence \mathbf{y} corresponding to English sentence \mathbf{x} (Osborne/Koehn)
- ▶ *Reinforcement Learning*: System receives sequence of inputs representing the state of the world. System is provided with occasional payoffs that define the success of the action sequence. Goal is to maximize payoff.
Examples:
 - ▶ RL for optimizing human-computer interaction (Lemon)
 - ▶ RL for robotic control (Ramamoorthy)

- ▶ *Unsupervised learning*: no targets. Model $p(\mathbf{x})$ using hidden variables. Simple example: clustering
- ▶ More complex example

Data



Results



- ▶ *Other tasks*, e.g. circuit inference

Why have a Research Programme on Machine Learning?

- ▶ *Internal:* Machine learning methods are widely researched and used in many places in the School and elsewhere. With the co-location of the School in the Forum there is an opportunity to develop the links between these various groups, so as to enhance mutual understanding, and likely lead to new research opportunities.
- ▶ First goal understand the nature of the work carried out in the different groups.
- ▶ Possibility to have e.g. an ongoing seminar series
- ▶ *External:* Could perhaps create an outward looking face to the world (cf Language at Edinburgh)
- ▶ Please think of useful things you would like to see the RPML do


- ▶ EU FP7 Network of Excellence, ~ 50 partners, <http://www.pascal-network.org/>. UoE is a member, and I am the site manager
- ▶ Emphasis in PASCAL2 is on *integrative* activities, e.g. challenges organized, workshops organized, long visits made to other PASCAL sites
- ▶ Multi-Component Learning Thematic Programme (Oct 08 - Feb 09)
- ▶ Partial or Delayed Feedback Thematic Programme (Mar 09 - Sep 09; semi-supervised/unsupervised learning and RL)
- ▶ Cognitive Architecture & Representation Thematic Programme (Oct 09 - Feb 2010)
- ▶ Other activities outwith thematic programmes are also ongoing

My Research Interests


- ▶ Particularly interested in unsupervised learning with structured probabilistic models
- ▶ Models for understanding time-series/condition monitoring
- ▶ Object recognition in images
- ▶ Gaussian processes for prediction
- ▶ Unsupervised learning to provide hypotheses for how the brain should work (!)

Factorial Switching Kalman Filter

Modelling artifactual and physiological problems in neonatal baby monitoring time series



Object recognition



Tasks: (for each of 20 classes)


- ▶ Classification: Is there a dog in this image?
- ▶ Localization: For each dog in this image, draw its bounding box
- ▶ Segmentation: which pixels belong to each dog?
- ▶ <http://pascallin.ecs.soton.ac.uk/challenges/VOC/>

Gaussian process prediction

Carl Edward Rasmussen and
Chris Williams
MIT Press, 2006

www.GaussianProcess.org/gpml

Now available free on the inter-
net



Oct 15 2008

- ▶ 1545 Mike O'Boyle: Machine Learning for automating compiler/architecture co-design
- ▶ 1600 Gillian Hayes: ML for robotic and software agents
- ▶ 1615 Amos Storkey: Modelling time sequences
- ▶ 1630 Frank Keller: ML for language learning and processing
- ▶ 1645 Discussion time
- ▶ 1700 Further discussions & food in G02

See <http://www.dai.ed.ac.uk/homes/ckiw/rpml/> for full listings

Oct 29

- ▶ 1530 Subramanian Ramamoorthy: ML for planning and control problems
- ▶ 1545 Jano van Hemert: Advanced Data Mining and Integration Research for Europe (ADMIRE)
- ▶ 1600 Steve Renals: Challenges arising from the instrumented meeting room
- ▶ 1615 Sethu Vijayakumar: Machine Learning for Sensorimotor control
- ▶ 1630 Bjoern Franke: An overview of the PASTA project
- ▶ 1645 Discussion time
- ▶ 1700 Further discussions & food in 2.32 (level 2 mini forum)

Nov 12

- ▶ 1530 Michael Rovatsos: Collaborative and strategic multiagent learning
- ▶ 1545 Simon King: Hidden Markov Model-based speech synthesis
- ▶ 1600 Mike Dewar: iBehave: Toward the Sequencing of Behaviour
- ▶ 1615 Mike Davies (SEE): Sparse representations and compressed sensing
- ▶ 1630 Oliver Lemon: Learning to optimize human-computer interaction: training multimodal dialogue systems with Reinforcement Learning
- ▶ 1645 Discussion time
- ▶ 1700 Further discussions & food in 4.40 (level 4 mini forum)

Nov 26

- ▶ 1530 Philipp Koehn: Statistical Machine Translation
- ▶ 1545 Leena Patel (SEE): Biologically Inspired Systems in Neuromorphics
- ▶ 1600 Hiroshi Shimodaira: Speech-driven facial animation: how to learn a stream to stream mapping?
- ▶ 1615 Steve McLaughlin (SEE): Signals, Sampling and Spectrum
- ▶ 1630 Jan Wildenhain (Cell Biology): TBA
- ▶ 1645 Discussion time
- ▶ 1700 Further discussions & food in 4.40 (level 4 mini forum)