

The (IM)2 Newsletter
Status of the (IM)2 IP.
Program of the Kick-off meeting.
Details about the ICSI deal.

Machine Learning at IDIAP
See who is behind torch, the
ML library that will be used in
many (IM)2 developments.

Issue N° 01
March 2002
Editor: Jean-Albert Ferrez
www.im2.ch/nl nl@im2.ch



The (IM)2 Newsletter

Every month the (IM)2 Newsletter brings you the latest and hottest scientific and administrative news about the (IM)2 NCCR and related topics

Status of (IM)2

While the (IM)2 NCCR officially started on January 1st, its management has been busy on the preparation of the structures since long before. All of the Individual Projects (IP) have been redefined in order to have every partner bring its expertise to the common goal of (IM)2 while at the same time avoiding overlap and duplicate work. This redefinition of the IP was also necessary since almost two years have passed since the final project was submitted. The Management Board is still working on this, based on the input of the IP heads. It is expected that all Individual Projects will be accepted in their revised form by March, when the White Papers submissions are due.



The members of the (IM)2 Management Board in action: Thierry Pun (UniGE), Murat Kunt (EPFL), Rolf Ingold (UniFR), Horst Bunke (UniBE), Jean-Albert Ferrez (IDIAP), Hervé Bourlard (IDIAP) and Burkhardt Stiller (ETHZ, now replaced by Luc Van Gool).

The management of (IM)2 has also set up an agreement with ICSI and prepared some events, the first of which being the (IM)2 Kick-off meeting. The program of this initial get-together event is given in this issue of the (IM)2 Newsletter.

Like the (IM)2 web site, www.im2.ch, the Newsletter is an internal and external showcase where information about (IM)2 events, publications, scientific breakthrough, etc. can be announced. The editor relies on your input to make this Newsletter as useful as possible. We plan to publish one issue every month as long as we get enough content.

Dr Jean-Albert Ferrez
(IM)2 Program Manager

(IM)2 Kick-off meeting


To celebrate the launch of the (IM)2-NCCR on Interactive Multimodal Information Management and the creation of the Signal Processing Institute within EPFL, we have the pleasure of inviting you to a joint Kick-off meeting aiming at bringing together scientists from the various institutions involved.

This Kick-off meeting will take place on

**March 6, 2002 at 14h00
in room EL1, EPFL**

Program

- 14h00 Welcome
- 14h15 Prof. Hervé Bourlard,
Director of IDIAP and (IM)2:
"The (IM)2-NCCR on Interactive
Multimodal Information
Management"
- 14h45 Dr Jean-Albert Ferrez:
"A short presentation of
the 9 (IM)2 Individual Projects"
- 15h15 Dr Samy Bengio, IDIAP:
"The Torch machine-learning
library"
- 15h45 Break
- 16h00 Prof. Stefan Catsicas,
EPFL VP for Research
- 16h15 Prof. Murat Kunt, Director of ITS
"The Signal Processing Institute"
- 16h45 Dr Chuck Wooters, ICSI Berkeley:
"The International Computer
Science Institute in Berkeley
and the visitor agreement
signed with (IM)2 "
- 17h15 Discussion and questions
- 17h30 Cocktail



tswwww.epfl.ch

www.im2.ch

For further information, please contact
Jean-Albert.Ferrez@idiap.ch.

(IM)2 inks deal with ICSI

The (IM)2 Management Board has set up a young researcher exchange agreement with the International Computer Science Institute (ICSI) in Berkeley. With this agreement, 24 man-months worth of hosting are available for PhD students and postdocs working within (IM)2.

ICSI is a computer science research organization that is closely affiliated with the Computer Science Division of the Electrical Engineering and Computer Science (EECS) department at The University of California at Berkeley. ICSI is comprised of research scientists from all over the world working together with faculty, postdocs, graduate students, and support staff on a large assortment of projects.

ICSI and (IM)2 are working on similar projects, in particular a Meeting Recorder, and close collaboration will foster new developments. Furthermore, the perspective of spending a few month in Berkeley will certainly increase the attractiveness of positions offered within (IM)2.



www.icsi.berkeley.edu

Events

(IM)2 Inauguration 4.5.02

The official Inauguration of (IM)2 will take place on May 4 in Martigny. This event is synchronized with the (IM)2 Scientific and Industrial Advisory Boards meetings that take place on May 3. The program will be announced later.

(IM)2 Summer Institute 3-8.10.02

From October 3 to October 8, Verbier will host the first (IM)2 internal workshop. This will be the perfect opportunity to look back on the first months of the NCCR, to identify gaps and overlaps in the various research plans.

The Machine Learning Group at

The Machine Learning group at IDIAP is mainly interested in statistical machine learning, a research domain mostly related to statistical inference, artificial intelligence, and optimization. Its aim is to construct systems able to learn to solve tasks given a set of examples that were drawn from an unknown probability distribution, eventually given some prior knowledge of the task. Another important goal of statistical machine learning is to measure the expected performance of these systems on new examples drawn from the same probability distribution.

People

The group is lead by Dr. Samy Bengio, who has more than 12 years of research experience in the field. He currently supervises 5 PhD students who will be awarded degrees from either EPFL, UniLausanne or UniBern, as well as 1 senior researcher, 2 development engineers and 2 short-term visitors.



Part of the Machine Learning Group at IDIAP: Front: Quan Le, Nicolas Gilardi, Johnny Mariéthoz. Back: Samy Bengio, Silvia Chiappa, Michael Kanevski, Christos Dimitrakakis, Christine Marcel.

Research Themes

Identity Verification: based on the voice of the client as well as an image or a video of him, we develop multimodal fusion algorithms for authentication.

Multimodal Input Processing: given multiple sequences conveying possibly related information at different times and time scales, the goal is to merge these sequences in order to take a unique and robust decision using efficient decoding algorithms.

Large scale data analysis: most actual powerful machine learning algorithms have been used for medium scale datasets: less than one hundred features describing one example and less than ten thousand examples in the dataset. Decomposition of the problem into sub-problems can lead to efficient solutions.

Ensemble models: One way to enhance generalization performance of machine learning algorithms is to combine the output of many algorithms instead of relying on only one algorithm. Many such methods are already known, such as AdaBoost, Bagging, Mixture of Experts.

Feature selection: Another way to enhance generalization performance of machine learning algorithms is to select and use only the input features that are well suited to solve a given problem.

Fusion of generative and discriminative models: two classes of machine learning algorithms are known and they have different advantages and disadvantages, depending on the problem to solve. We are interested in new algorithms that take advantages of both approaches.

Generalization performance analysis: It is important to analyze and understand machine learning algorithms in order to be able to compare them to other state-of-the-art algorithms.

Sequence modeling: most recent machine learning have been tailored for static problems. Given IDIAP's interest in speech processing, our group is also interested in developing and analyzing specific machine learning algorithm for sequence processing, including time series prediction and biological sequence analysis.

Spatial data analysis: We are specifically interested in building machine learning algorithms that would take into account spatial correlation between the input features and the target output in order to simultaneously enhance the prediction performance while preserving the spatial distribution of the dataset.

Multi-class classification: Many machine learning algorithms are in fact classification problems with multiple classes. One such problem in speech is the prediction of the phoneme (one out of 40 different phonemes) given the input features, at every time step.

Application Examples

The applications of Statistical Machine Learning are quite diverse. The *natural* target applications of our group are related to speech and vision. On top of them, here is a sample of other interesting application domains:

- Data Mining: how to extract interesting information from huge database warehouses (for instance, churn detection, client modeling and prediction).
- Environment: geostatistical analysis, specific 2D and 3D modelling.
- Finance and Economy: financial portfolio management, asset prediction, portfolio selection, auction analysis.
- Pattern Recognition: handwritten character recognition, speech recognition, face detection, speaker and face verification and identification.
- Biological Sequence Analysis: classification of DNA or RNA sequences.

www.torch.ch



The Machine Learning group has recently designed and developed a new library written in C++ and distributed under the GPL license, which should contain most of the state-of-the-art machine learning algorithms. It already contains

- all kinds of neural networks including multi-layer perceptrons, mixtures of experts, radial basis functions,
- support vector machines, for classification and regression,
- kmeans, gaussian mixture models, hidden markov models,
- non-parametric models such as k-nearest neighbors or parzen windows,
- ensemble models such as bagging and adaboost,
- embedded speech recognition for small vocabularies, speaker verification,
- general purpose matrix package and random generator functions,
- automatic documentation generator.

It is based on a modular design where efficiency was the core objective. Being Open Source (GPL), everybody can have access to it through the web page www.torch.ch and even contribute to it.

sb