

Machine Learning Group

Who are we?

Our group is a team of over 20 researchers, working in all aspects of applying Machine Learning to real-life problems. We engage in cutting-edge research and deliver applications based on Machine Learning.

Some of our projects...

- **Parallel Machine Learning (PML) Toolbox:** A joint effort of our group and the Data Analytics department at the IBM Watson Lab is the first publically-available tool for execution of data mining and machine learning algorithms on multiple processor environments or on multiple threaded machines.
- **EuResist:** A pharmacogenomics project that integrates viral genomics with clinical data to predict responses to anti-HIV treatment. We are part of an EU consortium that is delivering a system to provide clinicians with a prediction of responses to antiretroviral treatment in HIV patients, thus helping them choose the best drugs and drug combinations for any given HIV genetic variant.
- **Coverage Directed Test Generation:** We developed a coverage-directed generation (CDG) methodology and technology, designed to automate the process of using feedback from coverage analysis for tuning generation stimuli towards areas not adequately verified. Our work uses Bayesian Networks to encode the complex joint input-output distribution space and ultimately infer generation directives to the stimuli generator
- **Identifying missing information:** Our work with the Information Retrieval team at the IBM Haifa Labs centers on assessing the amount of content in an information repository. We are developing methods to assess gaps in knowledge as well as identify content that users cannot find due to its format or it being masked by other information.
- **Machine Learning for Dynamic System Analysis (MeLoDy):** The Melody project aims to improve problem diagnosis processes by applying Machine Learning techniques to configuration data of operational computers. Our tools are installed in several IBM support centers around the world.
- **Virtual guest inspection, learning, and control:** We developed methods for identifying both extreme and mundane conditions of "guest" operating systems, as well as fault detection.
- **Improving chip manufacturing:** We work with the IBM chip fabrication facility to develop new classification and outlier detection methods for improving chip yields.

We are also active in academia

Publications and Patents

- Over 50 papers published during the last four years
- More than 25 patents filed
- IBM 2006 Best Paper award

Workshops

- SIGIR 2005 workshop on Query Difficulty Prediction
- DAC workshop on CSP-for-Verification full-day tutorial
- NIPS 2006 workshop on Revealing Hidden Elements of Dynamical Systems
- NIPS 2007 workshop on Learning with Large Data
- The annual Israeli Machine Learning workshop at HRL.

More information?

Shai Fine, fshai@il.ibm.com

Elad Yom-Tov, yomtov@il.ibm.com