# Machine Learning @ NICTA

*Bob Williamson, 27 September 2011*

**Context**: Making sense of the ever increasing amounts of data gathered in all areas of human endeavor is an important contemporary challenge. Machine Learning (ML) technologies address that challenge. Traditionally "machine learning" might refer to symbolic computational approaches arising from the AI community, or statistical pattern recognition. However here it is used more broadly to include allied data-intensive areas such as information retrieval, and computational natural language processing. Machine learning researchers (ultimately and perhaps indirectly) construct software artifacts, and as such we need to be cognizant of aspects of software engineering and protocol design, as well as HCI design.

**Problem**: Currently ML is largely a technique-driven "black art". It is hard to compare different methods. Solutions are forced onto problems, rather than being tailored to fit. It is hard to reuse methods. It is often impossible for end users to exploit state of the art methods without the aid of an ML expert. Researchers within ML are often torn by an unresolved tension between further *developing* the underlying technology (which generates respect of peers) versus helping others solve the real problems by *using* ML technology. It is very hard to see ML as a field (as a whole) rather than a large bag of tricks.

**Opportunity**: NICTA can grasp the above tension and turn it into an opportunity and focus for the ML group. ML is an *enabling* technology – it can contribute significantly to solving challenging real problems, but it is almost never the solution solely in itself. *If* ML can be made more readily and widely usable, it will have great impact. Making it usable is *not* merely a matter of packaging existing solutions. ML needs to be viewed from a use-inspired perspective from the beginning. That means there needs to be a focus on problems (clearly articulating what the problem is being solved) cataloging techniques *against* problems, developing componentized software frameworks that allow the ready deployment of ML technologies and their integration with all of the other information and communication technologies needed to solve large scale and challenging problems. This *integration* is a common feature of mature engineering disciplines.

NICTA can embrace this opportunity whilst continuing to enable (in fact enhancing) the opportunities for creative contribution by its ML researchers.

## Vision

> *Turn Machine Learning into an Engineering Discipline*[1]

A variation on the above is

---

[1] Interestingly, the key ideas here are hardly new:

> "We can continue to pursue separate goals, invoke different methodologies, and develop disconnected theories, eventually leading to a wide array of subfields with few common concerns, concepts, or methods. This would not be a terrible fate, since progress would continue, though only in the narrow sense of that term. Alternatively, we can explore relations between various goals, attempt to combine methodologies, and search for integrated theories of learning that cross the paradigm boundaries which have formed in recent years. Not all learning researchers need devote full time to this endeavour, and in some cases, bridging the gap may involve little more than using new terminology or seeing methods in a new light. Some encouraging signs of cross-paradigmatic research have already started to emerge, but more remains to be done, and I invite experts and novices alike to join in the effort. This quest would result in a broader sort of progress, ultimately leading to a unified science of machine learning."
>
> – Pat Langley, "Towards a Unified Science of Machine Learning," *Machine Learning* **3**,253-259, 1989.

*Turn Machine Learning into an Engineering Discipline and an everyday tool of ICT developers*

The point of the more precise variant is that it is foolish to imagine that advanced ML technology will be directly used and assimilated by end users for a long time (except as a component part of a larger packaged solution). However we have seen in numerous projects within NICTA its appeal to other technologists. Focusing on developers as our customers will focus our efforts. It justifies carrying out work on languages and protocols. And it certainly justifies the focus on lightweight composable frameworks (which are generally more appealing to developers than large monolithic chunks of software). Longer term, of course it is desirable to make it even more accessible to average end-users, but realistically in the medium term that is only going to occur by having it embedded in overall solutions for end-use problems, rather than as a *tool* that a typical user might utilise (analogous, say, to a spreadsheet).

The vision can be viewed schematically (this is a first draft). At the centre is a map. This is the new and integrative component. It connects problems to methods, data representations to methods, and takes account of implementations technologies.

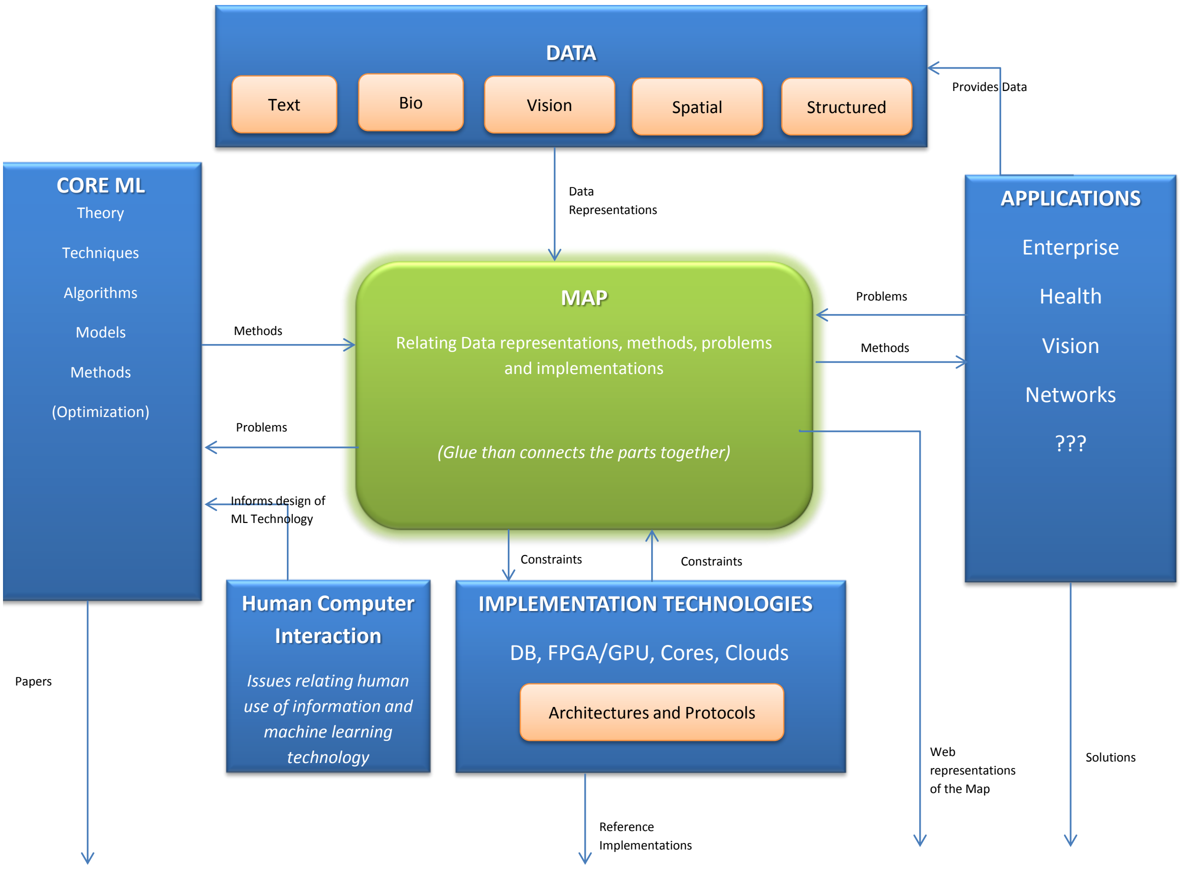


**Figure 1 Schematic of Machine Learning vision**

Key aspects of the vision are:

- *Embrace problem orientation* – break the tension between developing core new technology (and then looking for ways to apply it) versus engaging with end-use (and as a consequence not extending and developing the discipline).

- *Embrace pluralism* – It would be foolish to be prescriptive regarding the flavor of machine learning done (whether that is in terms of style (theory, experiment) or method (Bayes nets, decision theoretic etc.). A valuable consequence of our size is the diversity of approaches that we can bring to bear on problems.
- *Connect with software architecture and protocol design* and build the means by which ML technology can solve a range of end-use problems.
- *Develop composable technologies* that are readily applicable to end-use problems by cataloging the problems and hence solutions. Develop theories and methods by which ML effectively interfaces with knowledge representation and reasoning with uncertainty.
- *Avoid building monolithic solutions* in search of a problem; develop componentized, light-weight software stacks that allow the ready composition with other technologies. Avoid developing "me-too" solutions that exist in other forms.
- *Focus these above efforts* within a small number of projects in diverse domains (bioinformatics, enterprise documents, computer vision, and xxx (need to have discussions with business areas)). Ensure we have real expertise within these domains (this can arise from deep engagement with researchers and engineers from the domain, or our own ML researchers developing the expertise themselves).
- *Embrace the cloud, embedded systems and social computing.* These novel but pervasive computing paradigms hold many opportunities for the widespread deployment of Machine Learning technologies.

## Objectives

The key objectives I would like to see achieved are:

- Lead a systematic effort to relate problems to techniques, classify and categorise the problems, relate them to each other, and build languages that make it easier for a developer to make use of ML technology. This is the key new (grand) challenge;
- The majority of work should remain on developing new machine learning methods, solutions, theories and systems. My point is that the above "glue" is central to the vision, but, to speak metaphorically, a chair needs glue, but needs to be primarily wood. We still need to build the wood. But it needs to be shaped by and support the overarching vision.
- The key success criterion is that our research gets used. This does not merely mean software that we write gets utilised. It can also mean that new theories we develop become standard.