Intelligent Systems: Reasoning and Recognition

James L. Crowley

ENSIMAG 2 and MoSIG M1   Winter Semester 2012

Lecture 1   1 February 2012

Intelligence, Knowledge, Reasoning and Recognition

Outline

Class notes on the web :

http://www-prima.inrialpes.fr/Prima/Homepages/jlc/Courses/2011/ENSI2.SIRR/ENSI2.SIRR.html

Intelligence, Knowledge and Reasoning

What do we mean by Intelligence?

INTELLIGENCE :

(Petit Robert) "La faculté de connaître et comprendre,
incluant la perception, l'apprentissage, l'intuition, le jugement et la conception."

(Dictionnaire American Heritage) "The ability to know and to reason"

In this course we are concerned with technologies for Knowledge, Reason and Understanding.

The term "Artificial Intelligence" emerged from a pioneering workshop at Dartmouth University in 1956. Pioneers attending this workshop included Alan Newell, Herb Simon, John McCarthy, Marvin Minsky, Nils Nilsson, and Ed Feigenbaum.

Intelligence as the Ability to Solve Problems

A. Newell and H. Simon defined Intelligence as the Application of Knowledge to Problem Solving"

Newell, A.; Shaw, J.C.; Simon, H.A. (1959). Report on a general problem-solving program. *Proceedings of the International Conference on Information Processing.* pp. 256-264.

Nilsson: STRIPS, A* GraphSearch
R. Fikes and N. Nilsson (1971), "STRIPS: A new approach to the application of theorem proving to problem solving", Artificial Intelligence 2: 189–208

This view allows us to define knowledge in terms of the ability to solve problems, and reasoning as the ability to generate knowledge.

What is Knowledge?

What is knowledge? - Competence
Whatever enables the solution of problems.

Knowledge is defined by function and not by representation.

Kinds of Knowledge
Cognitive Psychologists identify different categories of knowledge representation.

Declarative: A symbolic expression of competence.
                Declarative knowledge is abstract
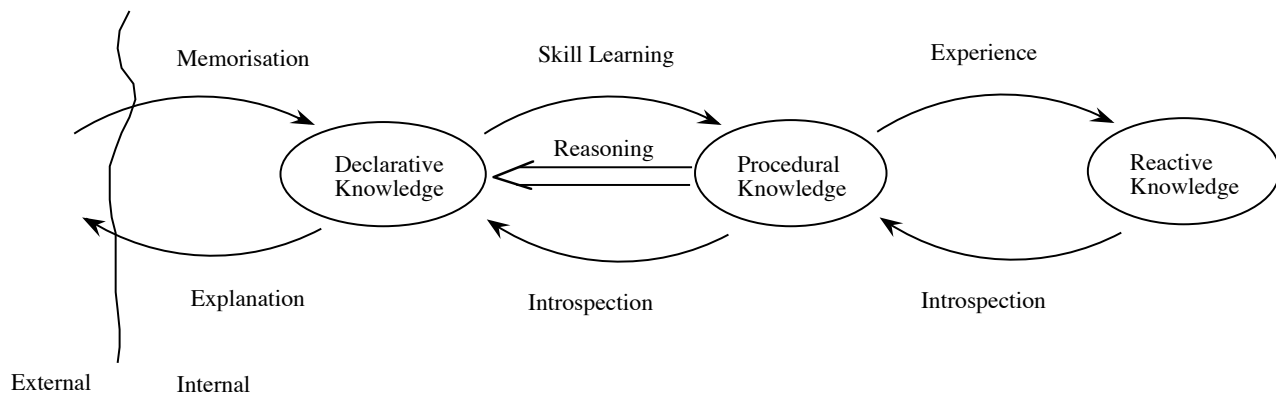   Declarative knowledge is used to communicate and to reason.
                Declarative knowledge must be interpreted to be used.

Procedural: A series of steps to solve a problem.
            A compiled expression of knowledge

Reactive: stimulus - response.

Newell proposes the distinction between "superficial" knowledge and "deep" knowledge.

Superficial knowledge provides reasoning without understanding. A common example of **superficial** reasoning is reasoning by symbol manipulation, without regard to the meaning of the symbols.

Deep knowledge requires the ability to predict and explain, and requires some form of model.

What is Reasoning?
Generation of new knowledge by inference.
Examples of types of inference:

        Deduction : $(p \wedge (p \rightarrow q)) \Rightarrow (q)$
        Abduction : $(q \wedge (p \rightarrow q)) \Rightarrow \text{Maybe}(p)$

$$\text{Induction: } p(A) \rightarrow q, p(B) \rightarrow q, ... \Rightarrow \forall x \, (\, p(x) \rightarrow q)$$

The Physical Symbol System hypothesis

In 1980, Alan Newell proposed that Intelligence REQUIRED the manipulation of symbols by a physical system. Newell's was based on a linguistic view of intelligence dating back to the 19th century.

What is a symbol?
A symbol is a 3rd order relation between
   A sign
   A thing
   An interpreter

There are two problems with Newell's hypothesis
1) It restricts intelligence to symbol manipulation. Intelligence is more general.
Newell claimed that only symbol manipulation system could be intelligent.
2) It confuses "What intelligence is" with "How intelligence is achieved".

Throughout the 1960s and 1970's research was limited to four or five universities:

MIT : Approach was to invent a powerful universal reasoning machine
Stanford: Approach was Lots of symbolic knowledge and a little bit of reasoning
CMU: Build a computing science based on cognitive models of biological systems.
Edinburgh (and Grenoble): Logic Programming as a universal reasoning mechanism.

Results
MIT => Behavioral robotics, Frames, Schema systems.
CMU => Rule based production systems and Cognitive Models and
Edinburgh and Grenoble => Logic Programming (Prolog)
Stanford => Expert System technologies.

Expert Systems

In 1980, the commercial successes of expert systems brought a great interest in Artificial Intelligence. In the early 1980, AI was presented as the future of informatics.

For expert systems technology, reasoning is accomplished by manipulation symbols.

The technologies for "Expert Systems" are based on symbol manipulation, without regard for the meaning of the symbols. Thus Expert Systems are based on superficial knowledge.

The knowledge acquisition barrier

While Expert Systems technologies provide useful solutions for many applications, there is a fundamental problem: Hand-crafting a knowledge base is generally a very expensive and difficult process.

By the late 1980's this led to a strong growth in interest in machine learning techniques for knowledge acquisition.

Machine Learning for Recognition

Historical Roots : Symbolic Learning vs Statistical Learning

In the 1990s, there were TWO machine learning communities.
Symbolic learning and Statistical learning.

Symbolic Learning
Symbolic learning was primarily motivated with solving the knowledge acquisition problem. The primary approach was to manipulate symbolic representations of declarative knowledge to infer new knowledge. Progress in this area has been slow.

Statistical Machine Learning.
Pattern analysis has long been a part of probability and statistics.

Statistical machine learning has its roots in Bayesian Pattern Recognition as well as Probability and Statistics.

In the 1950's Pattern Recognition was engineering discipline concerned with hand-crafting recognition machines.

In the early 1960's, Rosenblatt demonstrated the Perceptron learning machine [Rosenblat 62], with an implementation as a large analog computer.

Nilsson, among others, published a treatise on learning machines [Nilsson 65] grounded on Bayesian probability theory.

The textbook by Duda and Hart summarized the state of the art [Duda-Hart 73] in Bayesian pattern recognition in a remarkably modern and accessible form.

However an important barrier was the requirement for large amounts of data.

The availability of programmable computers made possible automatic algorithms for learning for recognition. The internet and digital sensing have brought about easy access to large volumes of data, making this approach very practical.

Since the 1990's we have seen an explosion in techniques for statistical machine learning. We will devote the second half of the class to such methods.

Intelligence as a Description of Behaviour

The Turing Test: an imitation game
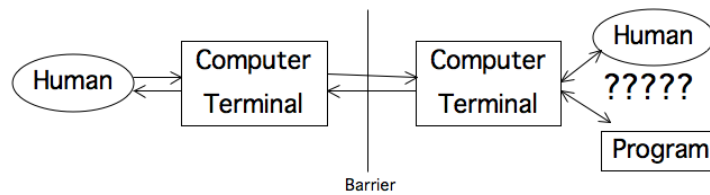
In 1936, Alan Turing claimed that a machine would exhibit intelligence if it exhibited behaviour that could not be distinguished from a person.

Limits: 1) Assumes that only humans are intelligent
          2) Reduced intelligence to human linguistic interaction

Turing posed the problem in terms of linguistic and social interaction, ignoring many other forms of intelligence.

However, Turing gave an important insight: Intelligence is NOT an intrinsic property of an agent. Intelligence is a "DESCRIPTION".

A Modern View of Intelligence

In the 1990's, research in robotics and perception, combined with insights from Cognitive science to bring about a new view of intelligence as a description of interaction. The key idea is that Intelligence is a descriptive label not an intrinsic property.
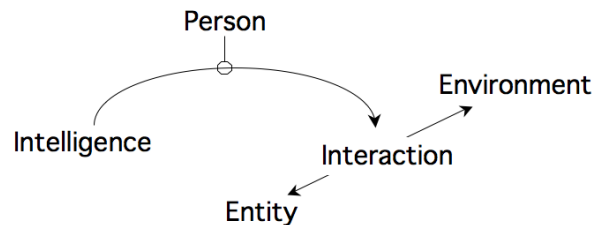
Intelligence <u>describes</u> the <u>interaction</u> of an <u>entity</u> with its <u>environment</u>.*
Intelligence is a <u>description</u> (an ascribed property)
Intelligence describes an <u>entity</u> that <u>interacts</u>.

To be considered "intelligent", a system must be embodied, autonomous, and situated [Breazeal 02], [Brooks 94].

Embodied:	Possessing a body (sensory/motor components)
Autonomous:	Self-governing;
	Have independent existence
Situated:	Behaviour determined by the environment

[Breazeal 02] C. Breazeal, Designing Sociable Robots, MIT Press, 2002.
[Steels and Brooks 94] L. Steels, and R. Brooks,   The artificial life route to artificial intelligence: Building Situated Embodied Agents. New Haven: Lawrence Erlbaum Ass., 1994.

Embodied: Incarnated. Possessing a body.

Body: A sensori-motor system for tightly coupled interaction with an environment.

Examples of Bodies:
     Natural: Human, mammal, insects, bacteria, plants,
     Artificial: Humanoid Robot, AIBO, mobile robots, roomba?

Environment: A system composed of multiple interacting entities.

Examples of Environments:

     Natural: Jungle, desert, sea floor….

     Artificial: Office, home, family, social network, computer games…

     Abstract: Chess, mathematics, any academic discipline…

What does it mean to Understand?

Understanding can be described as the ability to predict and explain. Understanding typically relies on some form of model that can be used to predict the outcome of a process or phenomena. Decomposing the model into components and interactions between components provides a means to explain a process or phenomena.

Course Overview

Part 1 – Reasoning
   1) Expert Systems
   2) Planning
   3) Rule Based Systems
   4) Structured Knowledge Representations

Part 2 – Recognition
   1) Bayes Rule
   2) Non-parametric methods for Bayesian Classification.
   3) Generative Methods: Gaussian Mixtures Models and EM.
   4) Discriminative Methods: Linear Classifiers and Boosting.

Programming exercise: Using CLIPS. – C Language Integrated Production Systems.

Exercises will NOT be graded.
Feedback and corrections will be provided for COMPLETED exercises.
Completed exercises should be COPIED INTO AN EMAIL including the names of all persons who contributed to the solution. Feedback will be returned by email. Please allow at least 2 weeks for feedback.

Note: DO NOT Send a file named file.clips or Exercise.clips, etc.