Intelligent Systems: Reasoning and Recognition

James L. Crowley

ENSIMAG 2 and MoSIG M1 Winter Semester 2013

Lecture 1 6 February 2013

Intelligence, Knowledge, Reasoning and Recognition

<u>Outline</u>

Class notes on the web :
http://www-prima.inrialpes.fr/Prima/Homepages/jlc/Courses/2012/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.

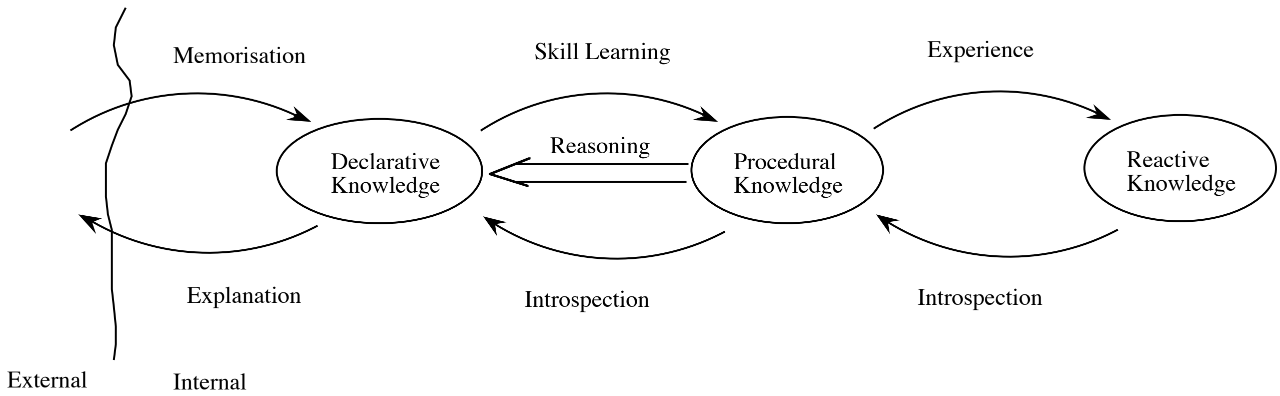

Newell proposes the distinction between "superficial" knowledge and "deep" knowledge.

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

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

Power of Symbolic Reasoning: Generalisation!

Illustration from 1st order predicate calculus:

A Gentle Introduction to 1st Order Logic

First-order logic is the standard formal logic for axiomatic systems.
Bertrand Russel (1909) showed that all arithmetic could be reduced to 1st order logic.

1st order symbolic Logic: Symbol Set, Logic operators, Predicates, Quantifiers.

Basic Elements:
1) Symbols: {S} (the domain of discours).
The symbol set is predefined and fixed.

2) logical operators: AND, OR, NOT (note: any 2 are sufficient).
Different Symbols are used. For example

AND	$\wedge$
OR	$\vee$
NOT	$\neg$
Implies	$\rightarrow$
Equivalence	$\leftrightarrow$
Meta Implication	$\Rightarrow$
Meta Equvalence	$\Leftrightarrow$

With {S} and the logical operators we can state logical propositions.

The meaning (semantics) of logical operators are defined by truth tables
(L. Wittgenstein, Tractatus, 1921).

3) Predicate Functions: Truth functions.

Predicates are functions that map the symbols to one of two values: TRUE or FALSE

$$P(A) \rightarrow P(B) \wedge \neg Q(C)$$

Generalization requires the addition of variables.
Variables are defined by Quantifiers

For All x $\forall x$:
For Each x $\exists x$:

Note these are redundant: $\forall x: \neg P(x) \Leftrightarrow \neg \exists x: P(x)$

The expressive power of 1st order results from the use of symbols.

Problems:
1) 1st order logic is restricted to closed domains. The set of symbols are predefined and not extensible. The real world is open

2) 1st order logic is a purely syntactic (shallow) reasoning system.

3) Knowledge acquisition (learning) is very expensive and difficult.

Limitations of Physical Symbol Systems:

There are several problems with Newell's hypothesis

1) It restricts intelligence to symbol manipulation. Intelligence is more general.

Newell claimed that symbol manipulation was necessary and sufficient for intelligence. We now know that this is not show.

We can show limited forms of intelligence without symbols (or generalisation).
Thus symbol systems are not necessary.

In addition, intelligence requires an open universe. 1st order logic is not sufficient. No adequate higher order logics have been found, in general because of the very high cost of knowledge acquisition.

More broadly, I would argue that Newel confused "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
Edinburgh and Grenoble => Logic Programming (Prolog)
Stanford => Expert System technologies.
MIT => Behavioral robotics, Frames, Schema systems.
CMU => Rule based production systems and Cognitive Models and

Expert Systems

From the 1950s to the 1990s, Artificial Intelligence concentrated on symbolic reasoning techniques. In the 1980's a technology for hand crafting "expert system" developded. 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.

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.

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.

In the 1990s, numerical pattern recognition was shown to provide simple and effective solutions to many hard AI problems. This lead to emergence of a new approach to AI.

Bayesian Machine Learning

From the 1980's Bayesian reasoning and statistics have provided a rich alternative to physical symbols systems.

History:

In the early 1960's, Frank Rosenblatt demonstrated a numerical learning algorithm named the "perceptron". [Rosenblat 62], using an analog computer.

This led to study of Probabilistic Pattern Recognition, summarized by textbooks by [Nilsson 65] and [Duda-Hart 73]. These methods were based on the mathematical field of Probability and statistics and Bayes rule.

AI researchers rejected this approach until the 90's when simple solutions to many hard AI problems were shown with a two layer perceptron (Hinton 86).

This lead to an explosion of techniques using Bayesian reasoning.

Two branches:

1) Statistical machine learning - use Baye's rule and statistics to learn functions to detect, recognize and discriminate patterns.

2) Bayesian Reasoning - A powerful method for accumulating evidence under uncertainty.

We will look at these in the second half of the course.

The goal is to combine the learning power and open universe of Bayesian methods with the expressive power and generalisation of symbolic logic.

Intelligence as a Description of Behaviour

But what do we mean by Intelligence? Alain Turing asked this question in 1936

The Turing Test: an imitation game

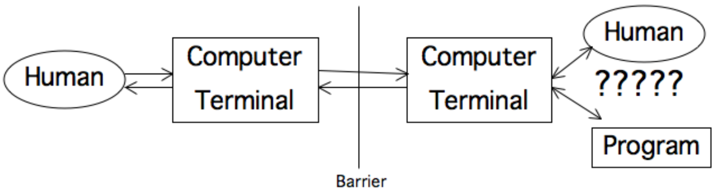

In 1936, Alan Turing claimed that a machine would exhibit intelligence if it exhibited behaviour that could not be distinguished from a person.

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.

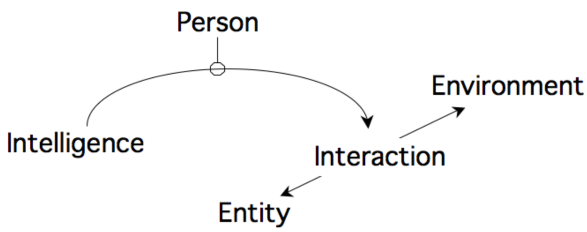

Intelligence describes the interaction of an entity with its environment.*
Intelligence is a description (an ascribed property)
Intelligence describes an entity that interacts.

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.