

Intelligent Systems: Reasoning and Recognition

James L. Crowley

ENSIMAG 2 and MoSIG M1

Winter Semester 2017

Lecture 1

1 February 2017

Intelligence: Reasoning, Understanding and Knowledge

Outline

Intelligence, Knowledge and Reasoning	2
What do we mean by Intelligence?	2
Intelligence as the Ability to Solve Problems	2
Intelligence as a Description of Behaviour	3
The Turing Test: an imitation game	3
A Modern View of Intelligence	3
What is Knowledge?	5
Kinds of Knowledge	5
What is Reasoning?	6
Understanding	7
Symbolic Reasoning	8
What is a symbol?	8
Limitations of Physical Symbol Systems:	8
Learning	9
Categories of Machine Learning algorithms:	9
Course Overview	11

Class notes on the web :

<http://www-prima.inrialpes.fr/Prima/Homepages/jlc/Courses/2016/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.

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" of interaction.

However, the Turing test only tests Social intelligence. There are many other forms of intelligence.

A Modern View of Intelligence


Intelligence describes interaction with **People**, **Things** and **Self**.

People: Social interaction with humans

Things: Interaction with the physical environment (including robotics).

Self: Interaction with internal components (self-regulation, autonomy)

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) able to act.

Autonomous: Self-governing;
Have independent existence

Situated: Having behavior that is appropriate to the task and 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 is Knowledge?

What is knowledge? - Competence

Whatever enables the solution of problems.

Knowledge is defined by ability 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

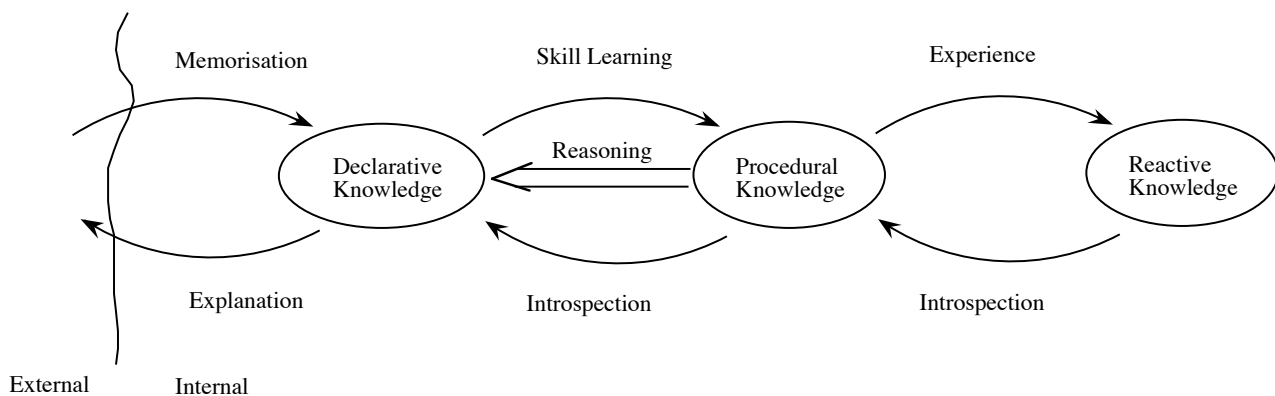
Declarative knowledge is used to communicate and to reason.

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.

Core Cognitive Abilities needed for Intelligence: perception, action, interaction

Meta-Cognitive abilities: Learning, Reasoning, Explanation, Anticipation, Imagination, Situation Awareness, Empathy, ...

Perception: (noun) 1: the ability become aware of something through the senses.
2: The process of recognizing and interpreting sensory stimulus.

Perception is more than sensing. Perception requires (1) Attention, (2) Sensing, (3) Recognition and (4) Assimilation.

Phenomena: Any pattern or event that can be perceived.

Attention: The human senses can be modeled as a hierarchy of filters. Attention is the process of tuning the filters to allow perception.

Sensing: Transformation of an external stimulus into an internal sensation. In biological systems, sensory signals are expressed as activations of neurons. In digital systems, sensory signals are expressed as symbols.

Recognition: Association of an external stimulus with a previously known or experienced phenomena.

Recognition (noun) The identification of something as having been previously perceived, experienced or known.

Humans learn to recognize and reason through experience.

What is Reasoning?

Generation of new knowledge by inference.

Examples of types of inference using Symbolic Logic

Deduction : $(p \wedge (p \rightarrow q)) \Rightarrow (q)$

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

Induction: $p(A) \rightarrow q, p(B) \rightarrow q, \dots \Rightarrow \forall x (p(x) \rightarrow q)$

Understanding

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.

The ability to predict phenomena is fundamental to survival.

The ability to explain phenomena is key to learning.

Our goal in this course is to learn the theories and models for building systems that can recognize and reason.

Symbolic Reasoning

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: abstraction! We can use symbols to designate sets of individuals, defined by intention (properties) or extension (list of members).

We can combine the symbols hierarchically.

Weakness: Cost of knowledge acquisition. Hand crafting is expensive and difficult.

Learning for symbolic knowledge representation remains an open (unsolved) problem.

Limitations of Physical Symbol Systems:

There are several problems with Newell's hypothesis

1) Newell reduced intelligence to symbol manipulation.

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

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

2) Intelligence requires an open universe. Most current techniques in symbolic reasoning require a pre-defined (closed) system of symbols.

More broadly, I would argue that Newell confused "What intelligence is" with "How intelligence is achieved".

Learning

Machine Learning: the construction of algorithms that can make predictions from data.

Definition proposed by Tom Mitchell (CMU): A computer program is said to **learn** from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T , as measured by P , improves with experience E .

In short: $\text{Learn}(P, T, E) : P(T)$ improves with E .

Most popular machine learning technique concerns the task of Recognition (also called Classification).

However, Machine Learning concerns other tasks such as Function Learning, Skill Acquisition, Clustering, Data mining, Concept Formation and Reasoning.

Categories of Machine Learning algorithms:

Supervised Learning: Learning a function from a labeled data set

Unsupervised Learning: Learning a function from unlabeled data.

Clustering: Discovering groups (categories) of data based on similarity of some property or measure.

Reinforcement learning: Learning to perform actions from experience

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

The dominant paradigm was Probabilistic Pattern Recognition, summarized in textbooks by [Nilsson 65] and [Duda-Hart 73], using methods based on Probability, Statistics and Bayes rule.

In the 1970s and 1980s AI researchers rejected Probabilistic Pattern Recognition in favor of hand coded symbolic reasoning systems. The dominant paradigm was "Expert systems", by programming problem solving techniques used by domain experts.

In the 1990's when it was shown that many intractable AI problems could be easily solved with a two-layer perceptron, also called an Artificial Neural Networks (Rummelhard and Hinton 89).

Within a few years, Bayesian machine learning techniques were shown to outperform Artificial Neural Networks for learning functions to detect and recognize patterns.

In the early 2000's discriminative techniques such as boosted learning, support Vector Machines and Kernel Methods provided improved performance for recognition.

Since 2010 a new revolution in recognition has occurred as massive amounts of data and computing power, made possible by the Internet and grid computing, have enabled techniques for learning deep Multi-layer Neural Networks with up to 20 or so layers (Deep Learning). Performance gains of 20% to 40% on classic problems in Computer Vision and Speech recognition were observed.

It is hard to say where the next revolution will come from.

A current challenge is to combine the learning power and open universe of Machine Learning with the expressive power and generalization of symbolic logic.

Rather than learning this year's hot technique, in this class we will concentrate on the fundamentals of intelligent systems.

Course Overview

Part 1 – Recognition

- 1) Supervised learning and Performance Evaluation
- 2) Perceptrons, Support Vector Machines
- 3) Artificial Neural Networks
- 4) Bayesian Learning, non-parametric methods.
- 5) Clustering and non-supervised learning with EM and K-Means

Part 2 – Reasoning

- 1) Symbolic Reasoning, Expert Systems
- 2) Rule Based Systems
- 3) Temporal Reasoning
- 4) Structured Knowledge Representations
- 5) Situation Models and Planning

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.

Some exercises in Part 1 may be performed in Python

Programming exercise in part 2 will use CLIPS. – C Language Integrated Production Systems.

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