Intelligent Systems: Reasoning and Recognition

James L. Crowley

MoSIG M1 Winter Semester 2018

Lecture 1 1 February 2018

Intelligence: Reasoning, Understanding and Knowledge

Outline

Class notes on the web :
http://www-prima.inrialpes.fr/Prima/Homepages/jlc/Courses/2017/MOSIG.SIRR/MOSIG.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.

Roughly speaking, we will adopt the following definitions:

Intelligence: Competence. The ability to perform actions and reactions that are appropriate for the domain and goal.
Knowledge: Anything that enables competence.

Intelligence as a Description of Behaviour

But what do we mean by Intelligence? Alain Turing asked this question in 1936.
His approach remains valid today.

Turing proposed a “behavioralist” definition of intelligence:

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.

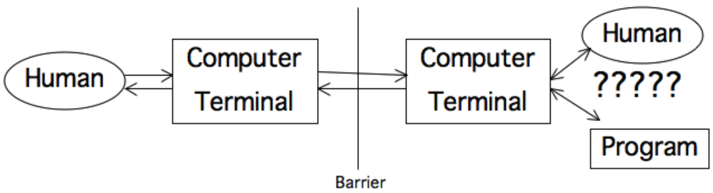

He proposed a test, in which a human observer interacts with an unknown agent over a teletype terminal. A machine (or program) is considered to be intelligent if the human observer cannot say whether it is a human or a machine.

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

The key idea is that Intelligence is a descriptive label not an intrinsic property.

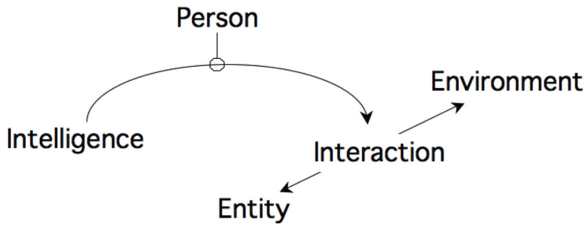

Intelligence describes the interaction of an entity with its environment.*
Intelligence is a description (a property assigned by an observer)
Intelligence describes an entity that interacts.

In the 1990's, research in robotics and perception, combined with insights from Cognitive science to bring about a view of intelligence as a description of interaction.

In this view, 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. The ability to act.
The ability to change the environment.

A "body" is a sensori-motor system for tightly coupled interaction with an environment.

Examples of Bodies:
- Natural: mammals, insects, bacteria, plants
- Artificial: Robots and other machines

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…

Autonomous: Self-governing; Have independent existence
Able to act to maintain self integrity. (the correct operation of the body).

Situated: Having behavior that is appropriate to the task and environment.
The ability to act and interact in a manner that is appropriate to the task and goal.

This leads to a view that there are many forms of intelligence, depending on the agent (the body), the domain of interaction (the environment), and the goals of the agent (autonomy, integrity).

Examples of common forms of human intelligence:
Social, mechanical, mathematical, financial, navigational, perceptual, etc
Animals also exhibit intelligence depending on their body and environment.

In conclusion: The Turing test posed the problem in terms of linguistic and social interaction, ignoring many other forms of intelligence.

Intelligence vs Rationality:
Rationality is the ability to choose actions to accomplish goals.

For many years, intelligence was defined as rationality.
This view dominated economics and psychology through the 1990's.

The problem is that humans are not rational.
Human behavior does not conform to mathematical models of risk and reward as measured by economics. Why?

Economic models assume explicit goals based on objective measures of value.
Human behavior is strongly influenced by a genetic heritage that has been shaped by evolution and social heritage that guides our interactions.

Knowledge

In the 1970's there was much debate about whether different forms of computer representation could be considered as "Knowledge".

In his 1980 Turing award lecture, Allen Newell ended the debate by defining knowledge as 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.

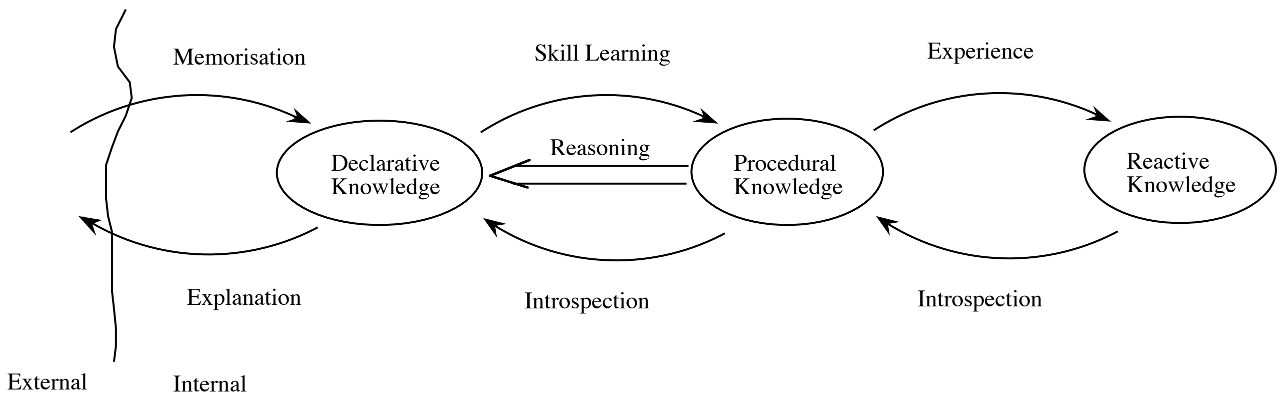

Newell proposes the distinction between "superficial" knowledge and "deep" knowledge.

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, manipulation, navigation, Anticipation, Empathy, …

Meta-Cognitive abilities: Learning, Reasoning, Explanation, Imagination
(abilities that create cognitive abilities)

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.

Sensation: Transformation of an external stimulus into an internal representation (phenomena). In biological systems, sensory signals are expressed as activations of neurons.

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

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

Humans learn to recognize and reason through experience.

Reasoning

Reasoning is the capacity to give meaning to phenomena.

Reasoning generally refers to associating a phenomena to some form of knowledge making it possible to explain and predict. In this common use, reasoning leads to understanding (comprehension).

Reasoning has historically been associated with symbolic logic.

Symbolic Reasoning

The term "Artificial Intelligence" emerged from a pioneering workshop at Dartmouth University in 1956. This workshop led to the emergence of a large scientific community identified as "Artificial Intelligence". Primary conferences included IJCAI, AAAI, and ECAI.

In the 1980's these conferences included work on Machine Learning, Computer vision, robotics, as well as logic programming and Expert Systems. In 1980, the AI conferences predicted that Expert Systems would come to dominate all of computer science. This lead to enormous investment and huge conference attendance.

The Physical Symbol System Hypothesis.

In the mid-1980s, the AI community adopted the view that Symbolic Reasoning was Necessary and Sufficient for intelligence.
This was referred to as the Physcial Symbol System hypothesis (PSS)

PSS was based on a linguistic view of intelligence dating back to the 19th century.
As AI adopted the PSS, was on neural networks, statistical machine learning, robotics, speech recognition and computer vision were no longer considered as relevant to AI. Scientific research on AI was dominated by Symbolic logic.

Symbolic Logic was the defined and the mathematics of intelligent systems.
In particular, first order predicate calculus.

Examples of types of reasoning using logic included:

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

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

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

Such techniques are examples of superficial reasoning.

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.

Weaknesses: Learning, Errors and Incompleteness.

Limitations of Physical Symbol Systems.

Fundamental problems soon emerged and the AI "Bubble" led to the AI winter with a drop in research funding and community size.

There are several problems with the "Physical Symbol System" (PSS) hypothesis.

1) Contra-evidence: PSS claimed that intelligence required symbolic reasoning.
We can show forms of intelligence without explicit symbols. Thus symbol systems are not necessary.

2) Intolerance of errors: In a logic system, the errors lead to contradictions. The result is a very fragile reasoning system. However, perception is uncertain and any system that reasons about perception must be robust to errors.

3) Knowledge Acquisition: Hand crafting is expensive and difficult.
Learning for symbolic knowledge representation remains an open (unsolved) problem.

4) Closed world hypothesis: Most current techniques in symbolic reasoning require a pre-defined (closed) system of symbols. Intelligence requires an open universe.

3) Learning Symbolic Representations for intelligence is VERY hard.
Very few stable techniques have been found. Most attempts to build symbolic learning systems have failed. No general technique has been found.

Understanding.

What does it mean to Understand?

In common usage, understanding has many meanings.
Researchers in Cognitive science prefer the term "comprehension".
We will use the two terms interchangeably.

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.

We "understand" a phenomena when we can relate it to a model that allows identification of components, explanation of causal relations and predictions of future phenomena.

Narratives are a common form of causal model. Narratives are declarative representation of causal sequences of events. Associating phenomena to a narrative makes it possible identify unseen causes and predict future events.

Other forms of models decompose a perception into interacting components. Decomposing the model into components and interactions between components provides a means to explain a process or phenomena and to develop causal models.

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.

In the second half of the class we will look at techniques for structured knowledge representation, situation modeling and narrative theory as models for reasoning and understanding.

Machine 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: 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 techniques for aquiring any form of cognitive ability, including Skill Acquisition, Social Interaction, spatial reasoning, 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.
Semi-supervised learning: Learning from unlabeled data with some form of error correction.
Reinforcement learning: Learning to perform actions from experience
Clustering: Discovering groups (categories) of data based on similarity of some property or measure.

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 linear classifier followed by a decision. Perceptrons became unpopular when Minsky and Papert showed that many simple patterns could not be recognized by a perceptron.

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.

The core hard problem was acquiring symbolic knowledge.

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 and displaced Bayesian recognition techniques.

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.
For example, it is not currently possible for a network to explain its reasoning.

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) Bayesian Learning, non-parametric methods.
3) Clustering and non-supervised learning with EM and K-Means
4) Support Vector Machines
5) Artificial Neural Networks (programming exercises with Keras)

Part 2 – Reasoning

1) Symbolic Reasoning, Expert Systems
2) Situation Models and Planning
3) Structured Knowledge Representations
4) Temporal and Spatial Reasoning
5) Causal Reasoning
6) Rule Based Systems (programming exercises with CLIPS)

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.