Intelligent Systems: Reasoning and Recognition

James L. Crowley

Ensimag 2 — Second Semester 2019
Lecture 1 — 5 Feb 2019

Intelligence: Reasoning, Knowledge and Recognition

Outline

Class notes and exercises on the web:
http://www-prima.inrialpes.fr/Prima/Homepages/jlc/Courses/2018/ENSI2.SIRR/ENSI2.SIRR.html

Intelligence as a Description of Behavior

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

Turing defined intelligence as a description of behavior.
Intelligence is defined as human level performance at tasks requiring perception, action and communication.

The Turing Test: an imitation game

Alan Turing claimed that a machine would exhibit intelligence if it exhibited behavior that could not be distinguished from a person.

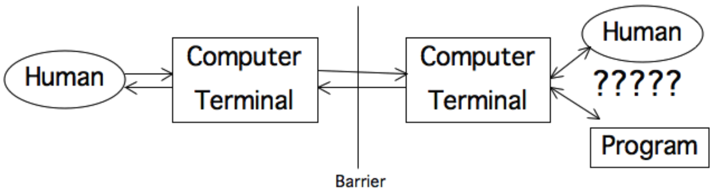

Turing proposed a test, in which a human observer interacts with an unknown agent over a teletype terminal. A machine (or program) is considered to be intelligent if the human observer cannot say whether it is a human or a machine.

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.
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.

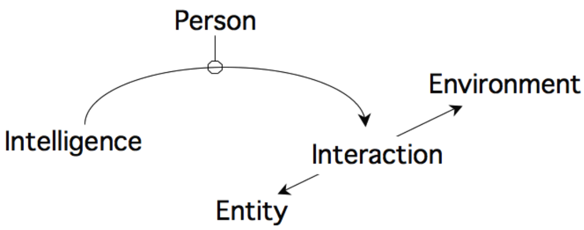

Intelligence describes the interaction of an entity with its environment.*
Intelligence is a description (a property assigned by an observer)
Intelligence describes an entity that interacts.

A modern definition defines intelligence as "Human level behavior at tasks requiring Perception, Action, and Interaction".

AI as a Scientific Discipline

The modern scientific domain emerged in the 1960s as a convergence of Cognitive Science, Logic, Planning, Pattern Recognition, Image Processing and other fields, driven by the emergence of Computer Science. The origin of the term "Artificial Intelligence" is generally credited to a Symposium at Dartmouth College in 1956, attended by the early AI pioneers, such as Arthur Samuel, John McCarthy, Marvin Minsky, Herb Simon, Allan Newell, etc.
The field has gone through several "epochs", each dominated by different paradigms.

Pre-1960: Automata and Pattern Recognition
1960-1985: Planning and Problem Solving
1980-1995: Expert Systems
1990-2005: Logic Programming and Behavioral Robotics
2000-2015: Bayesian Networks and the Semantic Web
2010-?: Deep Learning and Cognitive Computing

Paradigms are the problems and problem solutions that are shared by a scientific community. For AI, these are IJCAI, AAAI, and ECAI. In the 1980's these conferences focused on symbolic techniques for Machine Learning, Computer vision, robotics, as well as logic programming and Expert Systems.

In 1980, AI researchers predicted that Expert Systems would come to dominate all of computer science. This lead to enormous investment and huge conference attendance.

When the prediction failed to come true, funding and conference attendance dropped dramatically. Many schools stopped teaching AI.

From 1990 to 2010, steady progress was made in Computer Vision, Speech Recognition, Natural Language processing. The primary obstacle was machine learning.

In 2000 important progress was made with Support Vector machines and Bayesian techniques for machine learning. Since 2010 a major break through has occurred with the development of techniques based on Deep Learning. Deep Learning has been found to provide reliable solutions to longstanding problems perception and problem solving.

Deep learning requires High performance Computing and massive amounts of training data. Since 2010, this has been possible using Cloud Based data servers and the internet.

However, deep learning is only part of the solution. For human level performance at tasks requiring natural language requires explicit representation of knowledge. Natural language understanding offers access to the immense store of human knowledge coded in books and the internet, and to learn from and eventually teach humans.

Intelligence as Embodied, Situated Behaviour

The scientific community working on Robotics have developed a different definition. The common definition used in robotics and related field is that 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 Able to act to maintain ones integrity.
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. Much of the modern debate concerns whether a given system can be considered "embodied". May scientists take a hard-line view that a system must have a "body" that resembles a human. An alternative is that embodiment means having the ability to change the state of the external environment.

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…

In biology and in robotics, one distinguishes the “internal” and “external” environment. The is debate on whether acting on one’s internal environment (as with the human autonomic system) is sufficient for intelligence.

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

With this view, a remotely operated robot or vehicle is not autonomous.

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.

Knowledge and Reasoning and Understanding

A common dictionary definition of intelligence is "The ability to know and to reason" In this course we will be concerned with technologies for Knowledge, Reason and Understanding.

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. Anything that enabled the solution of problems.

The problem is that this lead to formulating intelligence as rationality. The problem is that human behavior is always rational.

Intelligence vs Rationality:

Rationality is the ability to choose actions to accomplish goals. For many years, the standard for intelligence was rationality. This view dominated economics and psychology through the 1990's. The problem is that humans are not individually rational. Human behavior does not conform to mathematical models of risk and reward. Why?

Human behavior is strongly influenced by a genetic heritage that has been shaped by evolution and social heritage that guides our interactions.

Kinds of Knowledge

Cognitive Psychologists identify different categories of knowledge representation.
A common cognitive model organized Declarative, Procedural and Reactive knowledge as a hierarchy.

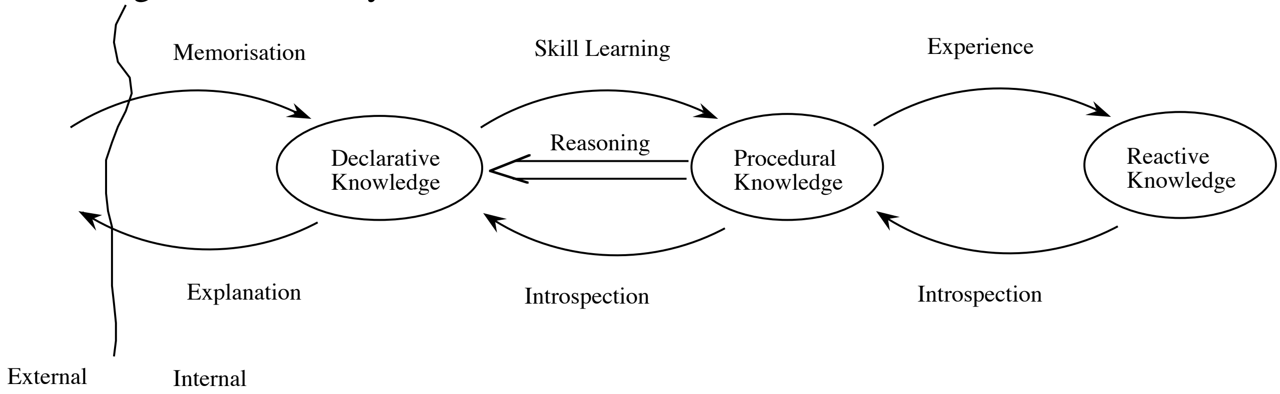

Declarative: A symbolic expression of competence.
Declarative knowledge is abstract
Declarative knowledge is used to communicate and to reason.
Declarative knowledge must be interpreted to be used.

Procedural: A series of steps to solve a problem.

Reactive: Automatic response to stimulus; conditioned reflex

Concepts, Working Memory and Cognition

Modern cognitive scientists study intelligence as interaction different forms of memories. Most models posit a cognitive architecture composed of a limited working memory interacting with perception and action based on episodic, procedural and conceptual memories.

Symbolic descriptions used for such models are like the mathematics used in other sciences. They describe functionality for interacting components. Actual systems may be implemented with a variety of programming tools, including Neural Networks.

A common cognitive model used for study of comprehension centers reasoning on processing of concepts in working memory.

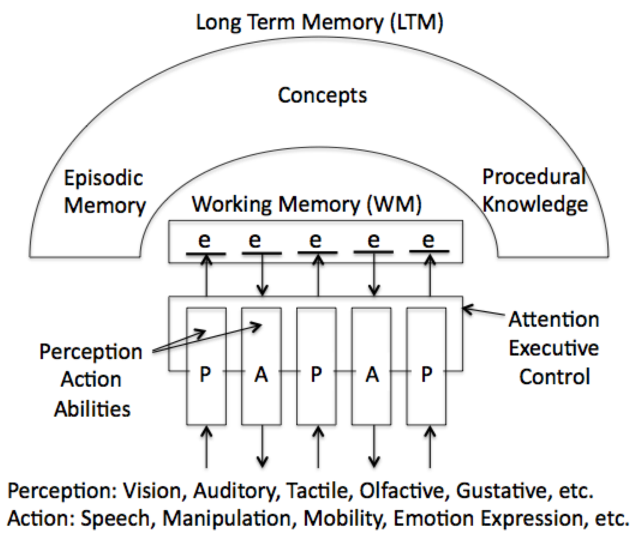

(inspired by Rasmussen, Card-Moran-Newell, Anderson, Kintsch, many others)

Most models of human cognitive models share a number of common elements:

- Perception: Transforms and combines sensory stimuli to Phenomena
- Short Term Perceptual Memory: Temporary buffer holding recent stimuli
- Action: Activation patterns for muscle groups.
- Working Memory: 7+/-2 memory slots (perceived or remembered)
- Long Term Memory

Long-term memory (LTM) refers to memory structures used in several different cognitive abilities:

- Episodic Memories: recordings of significant sensory experiences
- Semantic Memory: Abstract representations for sensory experiences
- Procedural Memory: Sequences of operations to accomplish goals
- Spatial memory (e.g. network of places in the hippocampus)

Spreading Activation

Most theories posit some form of "spreading activation" (Anderson 83) in which activation energy propagates through a network of cognitive "units".

Spreading activation is mechanism for associating cognitive units and controlling the contents of the limited Short-term memory.

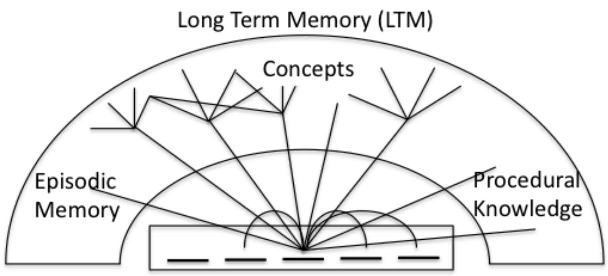

Activation energy spreads from short-term memory to other elements of short-term memory and to long-term memory including concept memory, episodic memory, procedural knowledge, etc. Activated units then spread their energy to other units where it can arrive from multiple paths and accumulate. At the same time the energy decays with time.

Units that receive energy from several other units can become "activated" and can replace one of the 7+/- 2 active units in short term memory. (Miller 56)

Reasoning

Reasoning is the capacity to give meaning to phenomena, where a phenomena is anything that can be perceived or remembered. Reasoning can apply to perception, or episodic, procedural or conceptual memory.

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

Understanding

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 associate it with a model that allows identification of components, explanation of causal relations and predictions of future 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 learn to recognize and reason.

Course Overview

Part 1 – Machine Learning for Recognition
1) Supervised learning and Performance Evaluation
2) Bayesian Learning, non-parametric methods.
3) Non-supervised learning with EM and K-Means
4) Support Vector Machines
5) Artificial Neural Networks, Back Propagation, and Architectures.

Part 2 – Reasoning
1) Planning and problem solving
2) Temporal and Spatial Reasoning
3) Concepts, Frames and Structured Knowledge Representations
4) Rule Based Reasoning
5) Causal Reasoning

Every week I will hand out exercise problems.
Exercises will NOT be graded.
However, the exam will be composed of (modified) versions of the exercises!
Do the exercises and the exam will be easy.
Ignore the exercises and the exam will be very hard.

Exercises may be done individually or in a group.
Exercises should be done within 2 weeks of assignment.

Completed exercises should be COPIED INTO AN EMAIL including the names of all persons who contributed to the solution. It is acceptable to work the exercise on paper and send a photo by email. The photo must be easily readable.

Feedback will be returned by email. Please allow at least 2 weeks for feedback.