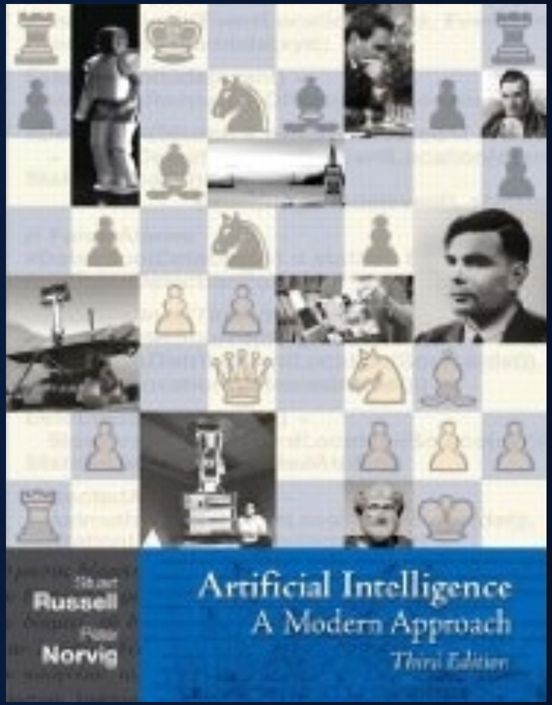

AI Fundamentals: an introduction

Maria Simi

AI Fundamentals: context

AI IN INDUSTRY AND SOCIETY

AI CURRICULUM

AI is taking over the world

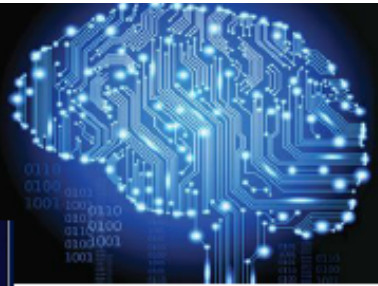

The deep learning tsunami

Deep Learning waves have lapped at the shores of computational linguistics for several years now, but 2015 seems like the year when the full force of the tsunami hit the major Natural Language Processing (NLP) conferences. [C. Manning]

Previous successes in the fields of image classification and speech …

Experts in the field (LeCun, Hinton, Bengio) agree on the fact that there will be important developments in text and video understanding, machine translation, question answering …

Google masters GO: *Deep-learning software defeats human professional for the first time*. AlphaGo. Nature 529, 445–446 (28 January 2016). In March, Lee Sedol defeated.

Is Deep Learning the answer to A.I.?

Andrew Ng

- Former director of the Stanford Artificial Intelligence Laboratory
- Now lead of Baidu's AI (1,200 people)
- Has directed many of the world's leading AI groups and built many AI products that are used by hundreds of millions of people

His answer:

1. *AI will transform many industries. But it's not magic*.
2. *Almost all of AI's recent progress is based on one type of AI, in which some input data (A) is used to quickly generate some simple response* (B) [A → B]

[https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now]

What Machine Learning Can Do

A simple way to think about supervised learning.

INPUT A	RESPONSE B	APPLICATION
Picture	Are there human faces? (0 or 1)	Photo tagging
Loan application	Will they repay the loan? (0 or 1)	Loan approvals
Ad plus user information	Will user click on ad? (0 or 1)	Targeted online ads
Audio clip	Transcript of audio clip	Speech recognition
English sentence	French sentence	Language translation
Sensors from hard disk, plane engine, etc.	Is it about to fail?	Preventive maintenance
Car camera and other sensors	Position of other cars	Self-driving cars

SOURCE ANDREW NG

What *this* AI can do

- Supervised learning Achilles' heel: it requires a huge amount of data.

  building a photo tagger requires anywhere from tens to hundreds of thousands of pictures (A) as well as labels or tags telling you if there are people in them (B)

- *AI work requires carefully choosing A and B and providing the necessary data to help the AI figure out the A→B relationship*.
- So what is the potential of implementing the mapping A→B?

  Ng's rule of tumb:

  If a typical person can do a mental task with less than one second of thought, we can probably automate it using AI either now or in the near future.

  Choosing A and B creatively has already revolutionized many industries. It is poised to revolutionize many more.

Issues for effective use of AI

Software. Not a problem: the community is quite open.

Data. *Among leading AI teams, many can likely replicate others' software in, at most, 1–2 years. But it is exceedingly difficult to get access to someone else's data. Thus data, rather than software, is the defensible barrier for many businesses*.

Talent. *Simply downloading and "applying" open-source software to your data won't work. AI needs to be customized to your business context and data. This is why there is currently a war for the scarce AI talent that can do this work*.

Deep learning and AI

- *Deep learning* is only one approach inside the much wider field of machine learning and … *Machine learning* is one approach within the wider field of AI
- *Many researchers are exploring other forms of AI, some of which have proved useful in limited contexts; there may well be a breakthrough that makes higher levels of intelligence possible, but there is still no clear path yet to this goal* [Ng]
- The goal of building AI is far from being solved and is still quite challenging in its own.

Limits of machine learning

Main arguments against a **ML-only** approach to A.I.

1. Explanation and accountability: ML systems are not (yet?) able to justify in human terms their results. For some application it is essential.
   - Knowledge must be meaningful to humans to be able to generate explanations
   - Some regulations requires the **right to an explanation** in decision-making, and try to **prevent discrimination** based on race, opinions, health, sex … (e.g. GDPR)
   - AI for decision support for humans? Yes, but …
2. ML systems learn what’s in the data, *without understanding what's true or false, real or imaginary, fair or unfair*
   - Most popular opinion in the training data; possible to develop **bad/unfair** models
   - People are generally more critical about information

Building complex AI systems requires the combination of several techniques and approaches, not only ML.

AI fundamentals

AI fundamentals has the role, within the Ai curriculum) of teaching you about "Good Old-Fashioned Artificial Intelligence” (GOFAI) or “symbolic AI”.

High-level "symbolic" (human-readable) representations of problems, logic and search. Symbolic AI was the dominant paradigm of AI research from the mid-1950s until the late 1980s

The approach is based on the assumption that many aspects of intelligence can be achieved by the manipulation of symbols (the **physical symbol system hypothesis)**:

> *“A physical symbol system has the necessary and sufficient means for general intelligent action*” [Allen Newell, Herbert A. Simon]

AI curriculum: structure

Curriculum specific courses (caratterizzanti)	**CFU**
Artificial Intelligence Fundamentals (sem 1)	6
Machine Learning (sem 1)	9
Computational mathematics for learning and data analysis (sem 1)	9
Natural Languages Technologies (sem 2)	9
Distributed systems: paradigms and models (sem 2)	9
Intelligent Systems for Pattern Recognition (sem 3)	6
Smart Applications (sem 3)	9
Electives	30
Free choice	9
Thesis	24
Total	**120**

AI curriculum: electives

Curriculum electives	**CFU**
Algorithm engineering (KD)	9
Data mining (KD)	9
Mobile and cyber-physical systems (ICT)	9
Information retrieval (KD)	6
Computational neuroscience (ING)	6
Social and ethical issues in computer technology	6
Robotics (S.Anna)	6
Semantic web (CNR)	6
Free choice	9
Total	**39**

Suggested combination: 2*9 + 2*6 +1*9 (free)   Other?

AI Fundamentals: at a glance

THE STRUCTURE OF THE COURSE

TEACHING METHODOLOGY

The main topics

*What are the contents … and **why is different from IIA**?*

The course assumes the first part of IIA (or similar topics) as a prerequisite.

Introduction (1)

I - Constraint satisfaction (3+1)

II - Knowledge representation and reasoning (6+1)

III - Reasoning under uncertainty (3+1)

IV - Planning (3+1)

V - Rule based systems (3+1)

Total: 24

I - Constraint Satisfaction Problems

1. Problem formulation
2. Problem reduction, consistency checking techniques
3. Heuristic and efficient search, local repair methods; problem structure.
4. *I: exercise review and student's presentations*

Prerequisites: building on problem solving as search

- problem formulation
- search algorithms in a state space
- heuristic search, local search

II – Knowledge representation and reasoning

1. The KR&R hypothesis. Review of reasoning in classical logic.
2. Non-monotonic reasoning
3. Reasoning about knowledge and belief
4. Reasoning about change; situation calculus and the frame problem. Temporal reasoning.
5. Semantic networks and frames
6. Reasoning about ontologies and description logics (the basics)
7. *II: exercise review and student's presentations of seminal papers*

III - Reasoning under uncertainty

1. Representing uncertain knowledge and probabilistic reasoning.
2. Belief networks and inference
3. Reasoning over time
4. *III: exercise review and student’s presentations*

IV - Planning

1. The planning problem, representation for actions. Planning as state-space search, regressive planning
2. Partial order planning, planning graphs
3. Planning in the real world: dealing with temporal and resource constraints, hierarchical planning, planning in non-deterministic domains, multi-agent planning.
4. IV: *exercise review and student's presentations of seminal papers.*

V - Rule based systems

1. Logic programming and rule based production systems.
2. Uncertainty in rule based systems; efficient implementation.
3. Constraint logic programming / abductive logic programming
4. V: *exercise review and student's presentations of seminal papers*

AI Fundamentals: you, the audience

STUDENTS

Present yourself

Your name

Where do you come from (undergraduate studies, town, country …)

Enrolled in AI curriculum/chances you will enroll

Your motivations, expectations and learning goals

Register into the Moodle platform, leave a statement:

https://elearning.di.unipi.it/course/view.php?id=96

AI Fundamentals: necessary background

WHAT YOU NEED TO KNOW

HOW TO FILL THE GAPS

Prerequisites

- Some background from computer science
  - ✓ Algorithms and complexity
  - ✓ Formal logic
  - ✓ Computability
  - ✓ Elements of probability calculus
- A basic course in Artificial Intelligence
  - ✓ Problem solving as search
  - ✓ Representation and reasoning in classical logic

You turn

1. Make sure you know the following:
   - ✓ Problem formulation according to the paradigm of *problem solving as search*. AIMA cap. 3 – Lecture slides IIA-2017, Lecture 3
   - ✓ Searching for solutions: basic search algorithms, heuristic search algorithms (A* and variants) AIMA cap. 3. Lecture slides IIA-2017, Lecture 3-4.
   - ✓ Local search algorithms and optimization problems: AIMA cap. 4.1, Lecture slides IIA-2017, Lecture 5.

AI as building intelligent computational agents

ONLINE BOOK:

“ARTIFICIAL INTELLIGENCE: FOUNDATIONS OF COMPUTATIONAL AGENTS”

BY POOLE AND MACWORTH

AI means building computational agents

Artificial intelligence, or **AI**, is the field that studies the synthesis and analysis of **computational agents** that act intelligently.

An **agent** is something that acts in an environment/it does something.

We are interested in what an agent does; that is, how it **acts**. We judge an agent by its actions.

An agent acts **intelligently** when

- ✓ what it does is appropriate for its circumstances and its goals
- ✓ it is flexible to changing environments and changing goals
- ✓ it learns from experience, and
- ✓ it makes appropriate choices given its perceptual and computational limitations.

Computational agents

A **computational** agent is an agent whose decisions about its actions can be explained in terms of computation (the decision can be broken down into primitive operation that can be implemented in a physical device).

The central **scientific goal** of AI is to understand the principles that make intelligent behavior possible in natural or artificial systems.

The central **engineering goal** of AI is the **design** and **synthesis** of useful, intelligent artefacts, agents, that are useful in many applications.

This is done by

- ✓ the **analysis** of natural and artificial agents;
- ✓ formulating and testing hypotheses about what it takes to construct intelligent agents;
- ✓ designing, building, and experimenting with computational systems that perform tasks commonly viewed as requiring intelligence.

The term "Artificial Intelligence"

- Artificial vs real/natural intelligence

  Artificial Intelligence **is not the opposite** of real Intelligence (not fake vs real).

  Intelligence cannot be *fake*. If an artificial agent behaves intelligently, it is intelligent. It is only the external behavior that defines intelligence (like in the Turing test).

  Artificial intelligence is real intelligence created artificially.

- Winograd schemas as a test of intelligence:
  - *The city councilmen refused the demonstrators a permit because they feared violence*.
    Who feared violence?
  - *The city councilmen refused the demonstrators a permit because they advocated violence*.
    Who advocated violence?

Human intelligence

The obvious naturally intelligent agent is the human being. Human intelligence comes from three main sources:

1. **biology**: Humans have evolved into adaptable animals that can survive in various habitats.
2. **culture**: Culture provides not only language, but also useful tools, useful concepts, and the wisdom that is passed from parents and teachers to children. Language, which is part of culture, provides distinctions in the world that should be noticed for learning.
3. **life-long learning (experience)**: Humans learn throughout their life and accumulate knowledge and skills.

Another form of intelligence is social intelligence, the one exhibited by communities and organizations.

Agents Situated in Environments

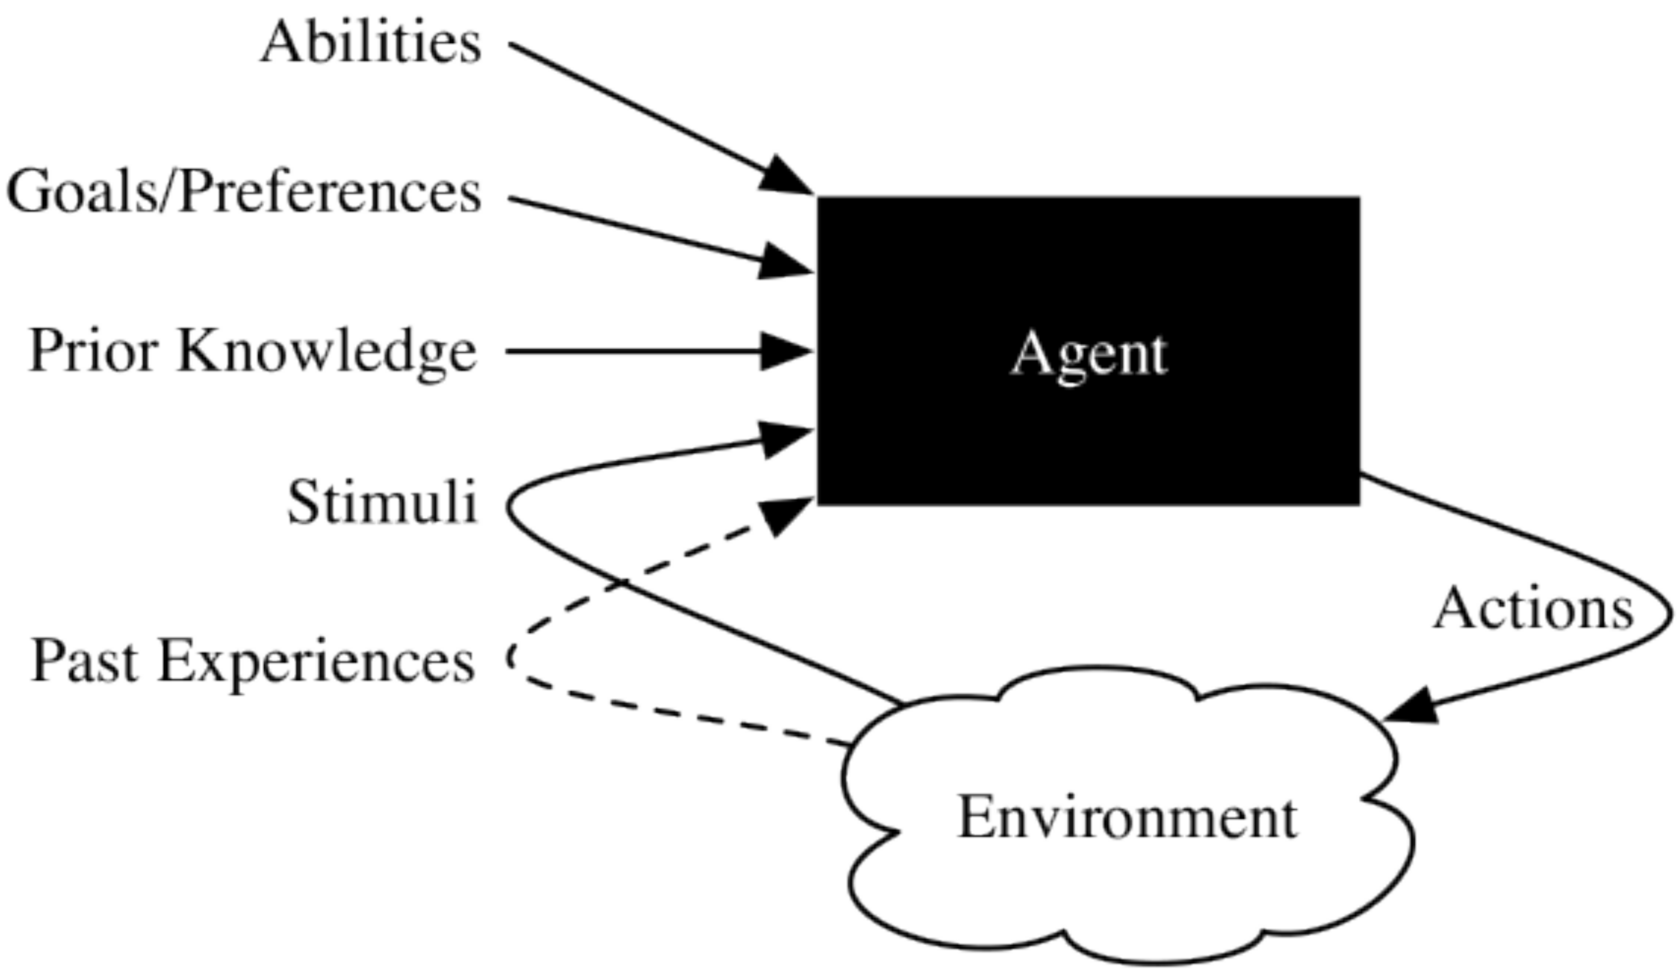

Design process for agents

Three aspects of computation that must be distinguished:

1. **Design time computation**, that goes into the design of the agent
2. **Offline computation**, that the agent can do before acting in the world
3. **Online computation,** the computation that is done by the agent as it is acting.

If most knowledge is given either at design time or offline, the agent is highly specialized and simple but may lack flexibility. However, designing an agent that can adapt to complex environments and changing goals is a major challenge.

To reach this goal, we have two strategies:

- simplify environments and build complex reasoning systems for these simple environments;
- build simple agents for natural environments, simplifying the tasks

Steps in the design process

1. Define the **task**: specify *what* needs to be computed
2. Define what constitutes a **solution:** optimal solution, satisficing solution, approximately optimal solution, probable solution.
3. Choose a **formal representation** for the task; this means choosing how to **represent knowledge** for the task and building a **knowledge base**
4. Compute an **output**
5. **Interpret** output as a solution

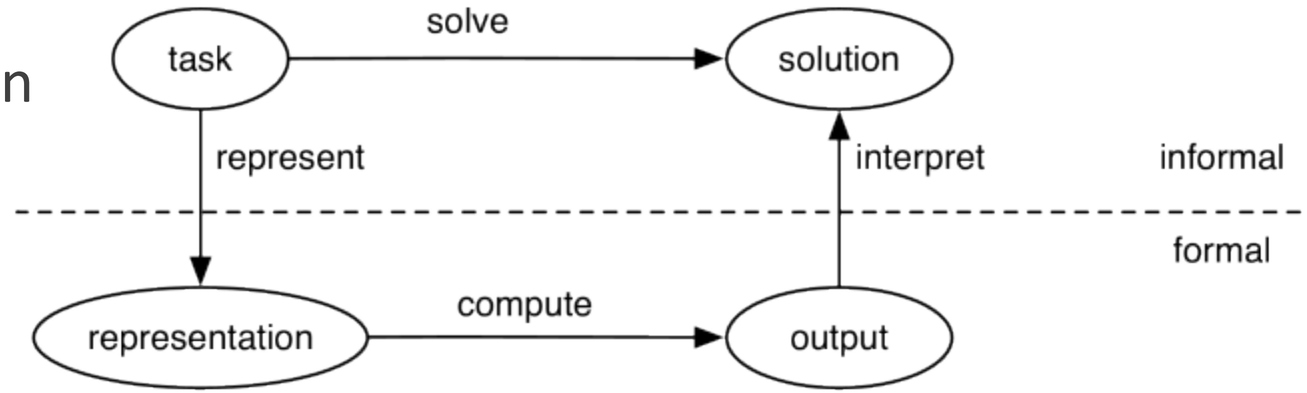

Agent design space: dimensions of complexity

1. Modularity
2. Planning Horizon: *how far ahead to plan*
3. Representation: *how to describe the world*
4. Computational limits: *real agents have limited computational resources*
5. Learning: *how to learn from experience*
6. Uncertainty, in both perception and the effects of actions
7. Preference: *the structure of goals or preferences*
8. Number of Agents
9. Interaction

Modularity

Modularity is the extent to which a system can be decomposed into interacting modules and it is a key factor for reducing complexity.

In the modularity dimension, an agent's structure is one of the following:

- **flat** – there is no organizational structure
- **modular** – the system is decomposed into interacting modules that can be understood on their own
- **hierarchical** – the system is modular, and the modules themselves are decomposed into simpler modules, each of which are hierarchical systems or simple components. The agent reasons at **multiple levels of abstraction**.

Planning Horizon

The planning horizon dimension is how far ahead in time the agent plans.

In this dimension an agent is one of the following:

- **Non-planning agent:** does not look at the future.
- **Finite horizon planner**: agent that looks for a fixed finite number of stages. Greedy if only looks one time step ahead.
- **Indefinite horizon planner** is an agent that looks ahead some finite, but not predetermined, number of stages.
- **Infinite horizon planner** is an agent that keeps planning forever. Ex. stabilization module of a legged robot

Representation

The representation dimension concerns how the state of the world is described. A state of the world specifies the agent's internal state (its belief state) and the environment state.

From simple to complex:

- **Atomic states**, as in problem solving.
- **Feature-based representation:** set of propositions that are true or false of the state, properties with a set of possible values. (PROP, CSP, most machine learning).
- **Individuals and relations** (often called relational representations). Representations at the expressive level of FOL (or contractions)

Computational limits

An agent must decide on its best action within time constraints or other constraints in computational resources (memory, precision, …)

The computational limits dimension determines whether an agent has

- **perfect rationality**, where an agent is able to reasons about the best action without constraints
- **bounded rationality**, where an agent decides on the best action that it can find given its computational limitations.

An **anytime algorithm** is an algorithm where the solution quality improves with time. To take into account bounded rationality, an agent must decide whether it should act or reason for longer.

Learning

Learning is necessary when the designer does not have a good model.

The **learning dimension** determines whether

- knowledge is given, or
- knowledge is learned (from data or past experience).

Learning typically means finding the best model that fits the data and produces a good predictive model.

Modelling formalisms and approaches are dealt in this course. All the issues concerned with learning are dealt in the **Machine Learning** course.

Uncertainty

Uncertainty is divided into two dimensions:

- uncertainty from sensing/perception (**fully observable**, **partially observable** states)
- uncertainty about the effects of actions (**deterministic**, **stochastic**).

  When the effect is stochastic, there is only a probability distribution over the resulting states.

We will deal with uncertainty in Section III.

Preference

The preference dimension considers whether the agent has goals or richer preferences:

- A **goal** is either an achievement goal, which is a proposition to be true in some final state, or a maintenance goal, a proposition that must be true in all visited states.
- **Complex preferences** involve trade-offs among the desirability of various outcomes, perhaps at different times.
  - An **ordinal preference** is where only the ordering of the preferences is important.
  - A **cardinal preference** is where the magnitude of the values matters. States are evaluated by **utility functions**.

Number of agents

The **number of agents** dimension considers whether the agent explicitly considers other agents:

- **Single agent reasoning** means the agent assumes that there are no other agents in the environment or that all other agents are "part of nature", and so are **non-purposive**.
- **Multiple agent reasoning** (or multiagent reasoning) means the agent takes the reasoning of other agents into account. This occurs when there are other intelligent agents whose goals or preferences depend, in part, on what the agent does or if the agent must communicate with other agents.

Interaction

The interaction dimension considers whether the agent does

- **offline reasoning** where the agent determines what to do before interacting with the environment, or
- **online reasoning** where the agent must determine what action to do while interacting in the environment, and needs to make timely decisions.

More sophisticated agents reason while acting; this includes long-range strategic reasoning as well as reasoning for reacting in a timely manner to the environment.

Summary

Dimension	Values
Modularity	flat, modular, hierarchical
Planning horizon	non-planning, finite stage, indefinite stage, infinite stage
Representation	states, features, relations
Computational limits	perfect rationality, bounded rationality
Learning	knowledge is given, knowledge is learned
Sensing uncertainty	fully observable, partially observable
Effect uncertainty	deterministic, stochastic
Preference	goals, complex preferences
Number of agents	single agent, multiple agents
Interaction	offline, online

Your turn

2. Read the introduction of the online book "**Artificial Intelligence: Foundations of Computational Agents"** by Poole and Macworth http://artint.info/2e/html/ArtInt2e.html
3. Study one of the examples in the introduction (1.6 Prototypical Applications) and be prepared to report about it.

No classes on Thursday 21st!!!

Main references for the course

[AI-FCA] David L. Poole, Alan K. Mackworth. *Artificial Intelligence: foundations of computational agents*, Cambridge University Press, Apr 19, 2010 – Computers. http://artint.info/2e/html/ArtInt2e.html

[AIMA] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach* (3rd edition). Pearson Education 2010.

[KR&R] Ronald Brachman and Hector Levesque. *Knowledge Representation and Reasoning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. 2004.

[AI-LF] Genesereth, M., and Nilsson, N., Logical Foundations of Artificial Intelligence, San Francisco: Morgan Kaufmann, 1987.

[AI-NS] Nils Nilsson, N., *Artificial Intelligence: A New Synthesis*, San Francisco: Morgan Kaufmann, 1998.