Artificial Intelligence

Information Systems and Machine Learning Lab (ISMLL)
Tomáš Horváth

27th October, 2011

Philosophically about AI

What is Artificial Intelligence?

- Systems, that
  - THINK
    - Like humans
    - Rationally
  - ACT
    - Like humans
    - Rationally

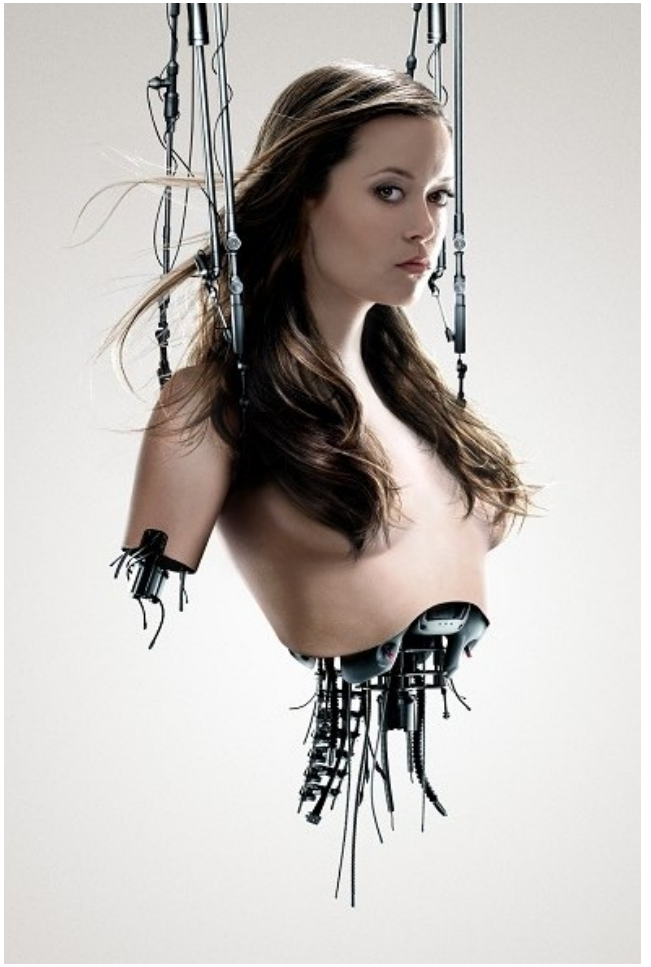

Thinking humanly

- If we would have a precise theory of the mind
  - We could express it as a computer program
- Cognitive science brings together
  - Computer models
  - Techniques from psychology

Acting humanly

- Turing test
  - a satisfactory operational definition of intelligence
    - (How long) Can a machine fool a human?
  - Instead of qualitative definition of requirements
    - Natural language processing, knowledge representation, automated reasoning, machine learning
  - www.jabberwacky.com

Thinking & Acting Rationally

- “Laws of thought” approach
  - Making correct, logical inferences
    - The world cannot be always described by correct logical notation
    - Difference between solving a problem “in principle” and doing it in practice
    - There are often situation with no correctly provable thing to do
    - Sometimes involving inference is not an advantage (hands on a hot stove)
- “Rational agent” approach
  - A rational agent is one that acts so as to achieve the best (or the best expected) outcome

Main disciplines contributing to AI

- Philosophy (428 B.C. – present)
  - Can formal rules be used to draw valid conclusions?
- Mathematics (800 – present)
  - What can be computed?
- Economics (1776 – present)
  - How should we make decisions to maximize payoff?
- Neuroscience (1861 – present)
  - How do brains process information?
- Psychology (1879 – present)
  - How do humans and animals think and act?
- Computer engineering (1940 – present)
  - How can we build an efficient computer?

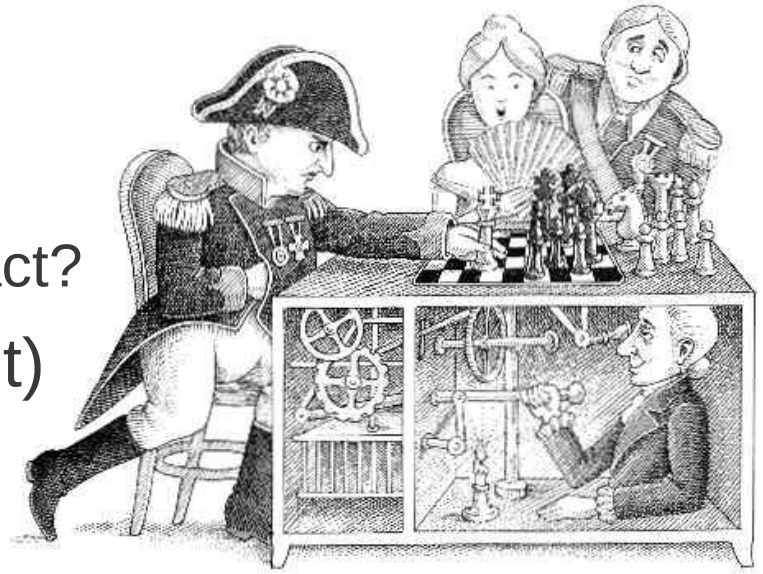

Overview of the lecture

Searching

- Find the shortest way from Arad to Bucharest

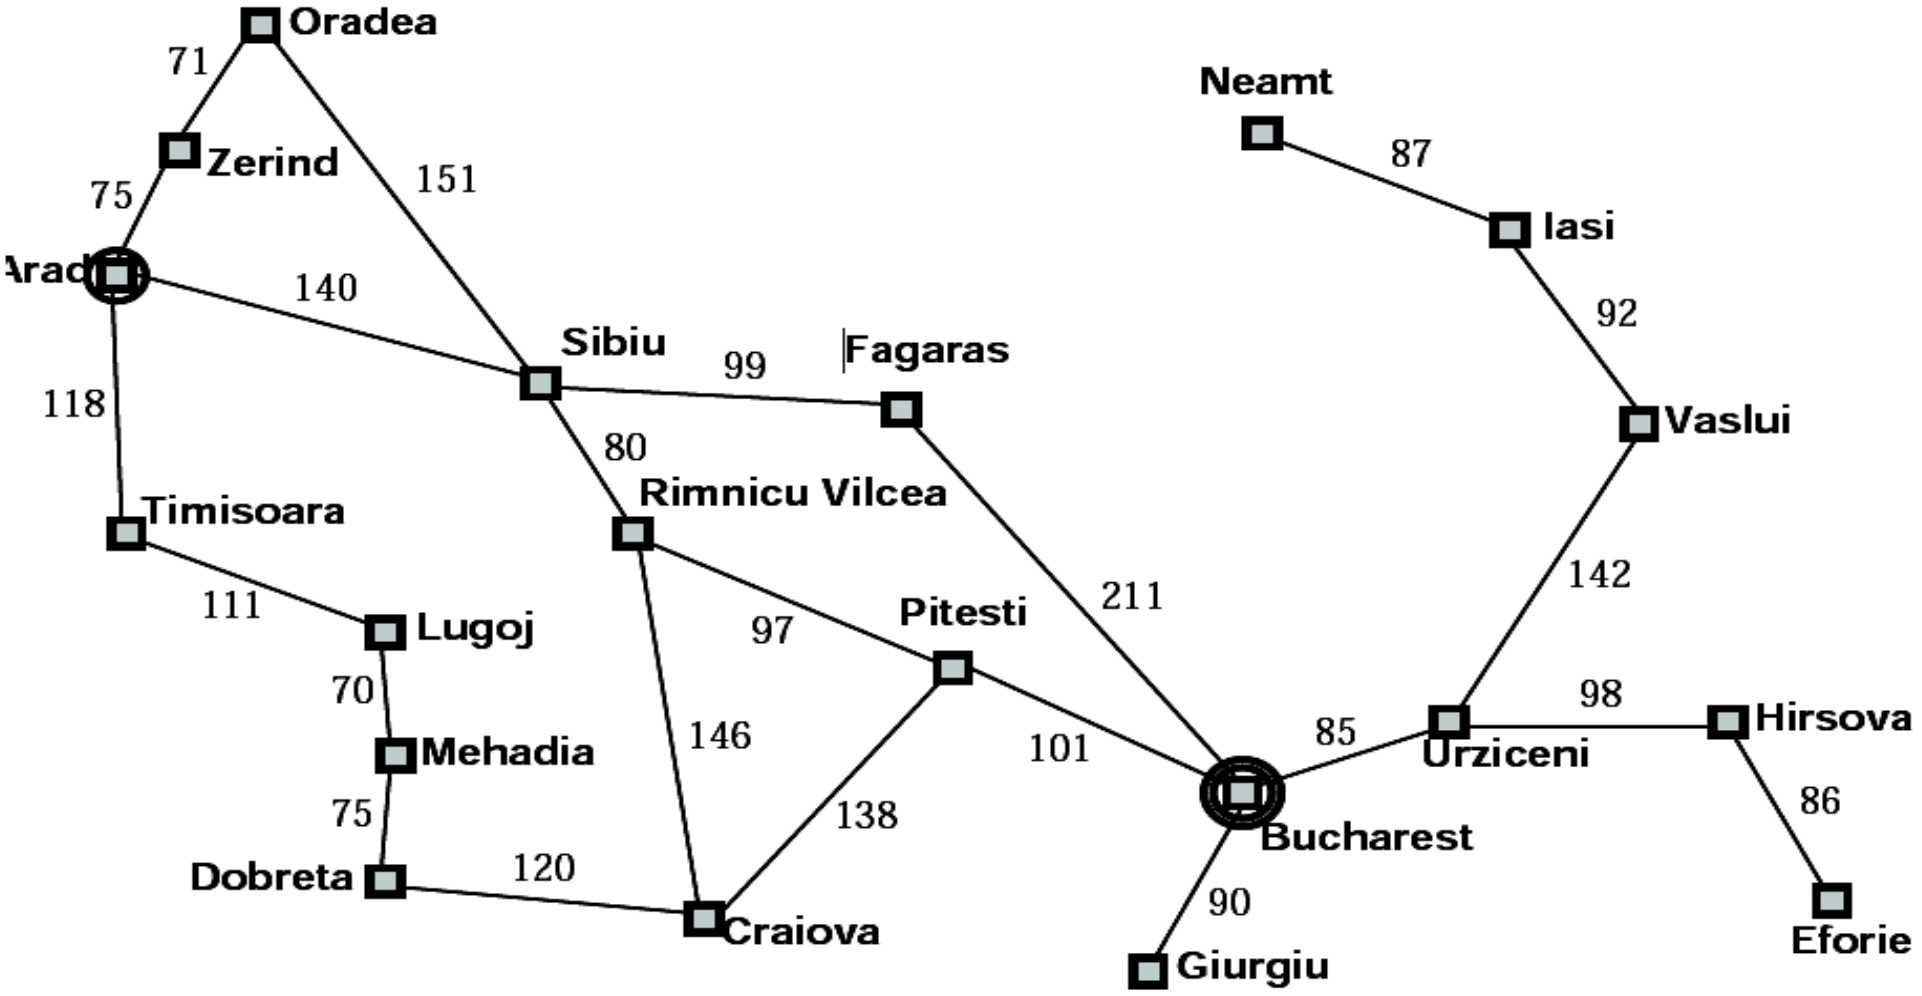

Searching

- Several strategies

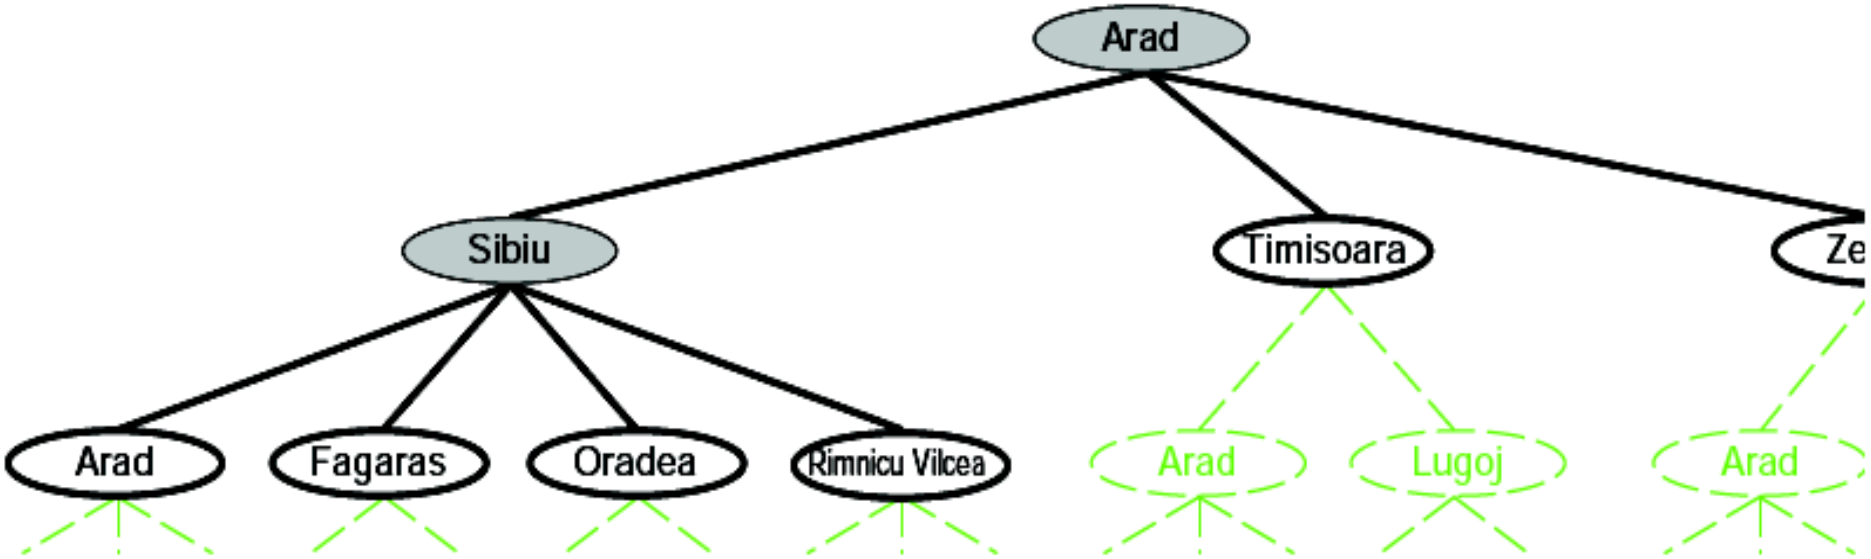

Searching

- Informed

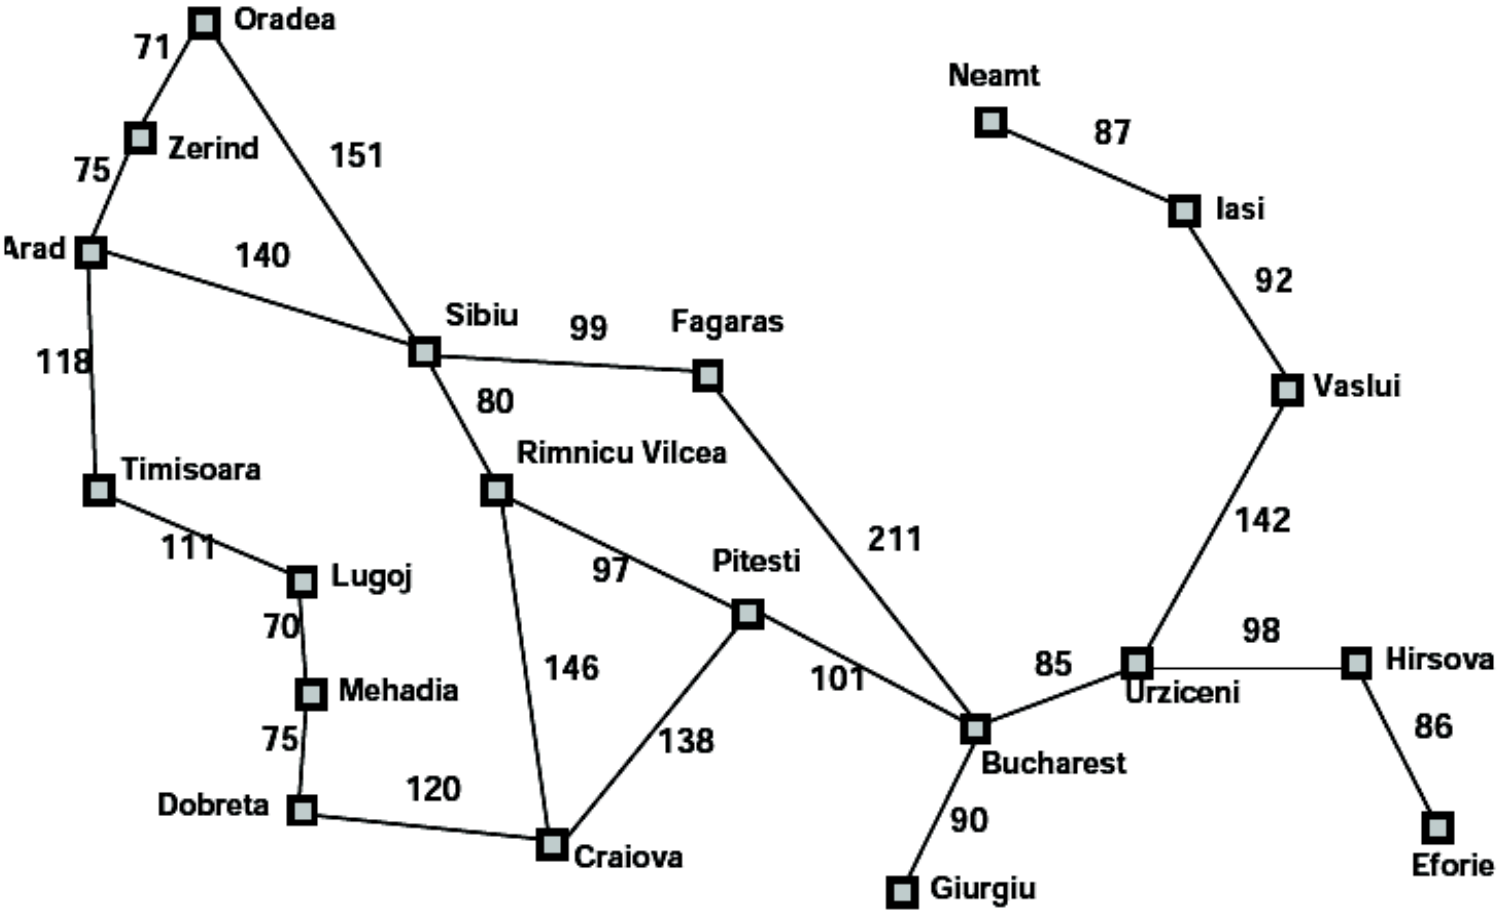

Straight–line distance to Bucharest	
Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	178
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	98
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

Searching

- Informed

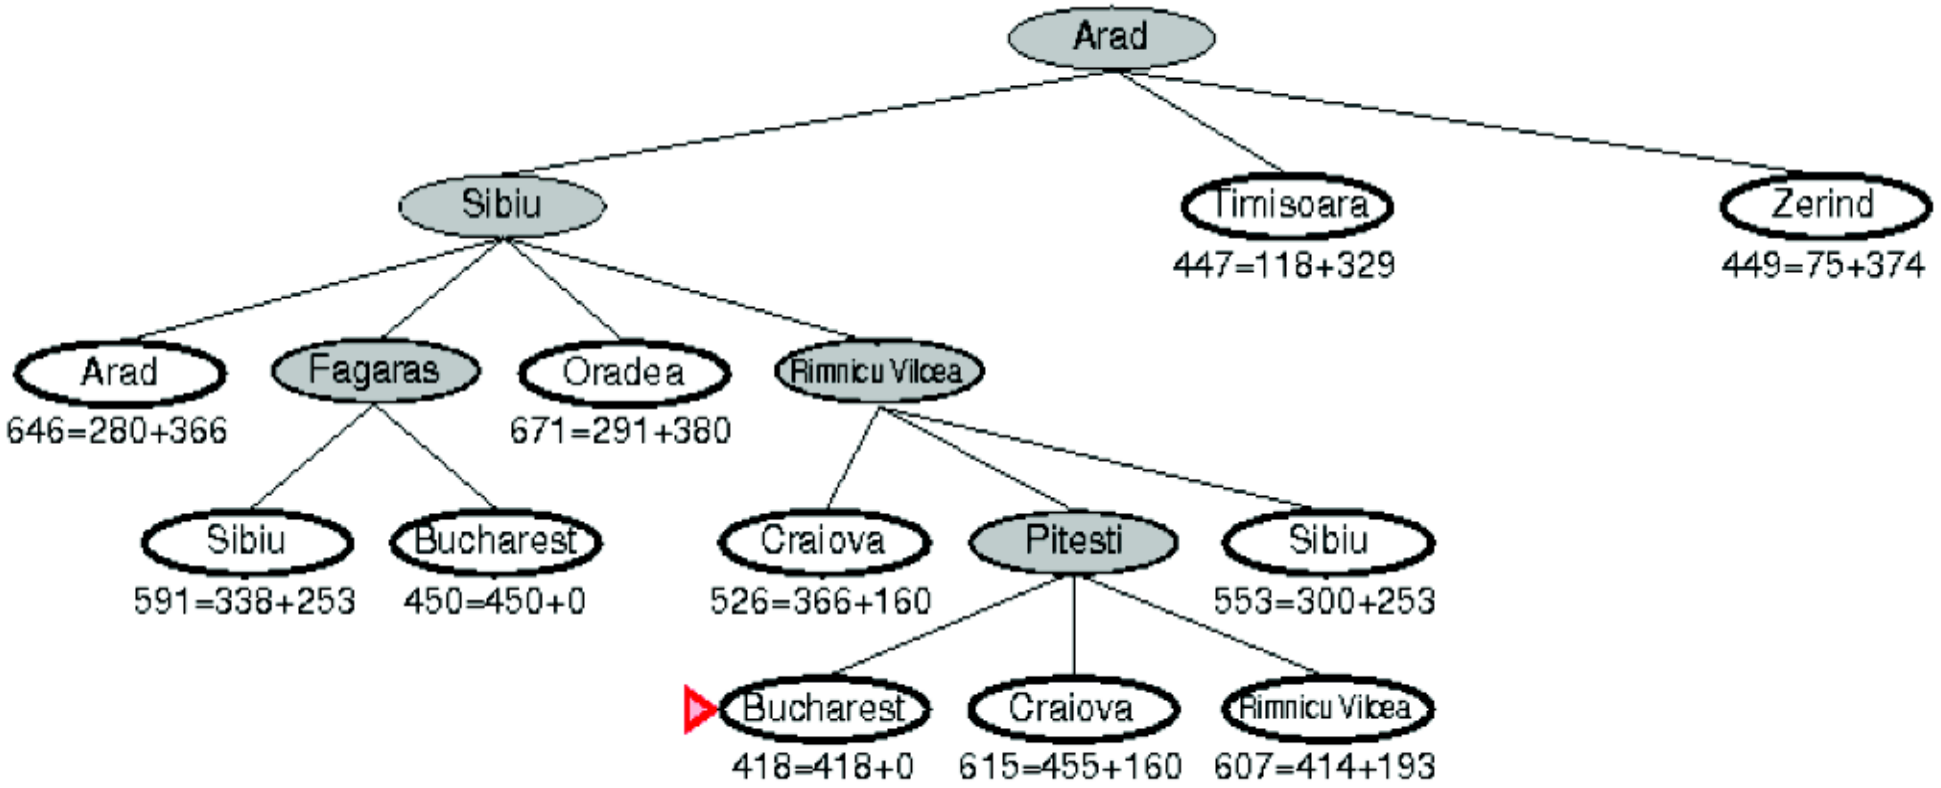

Adversarial Search

- Game against an opponent

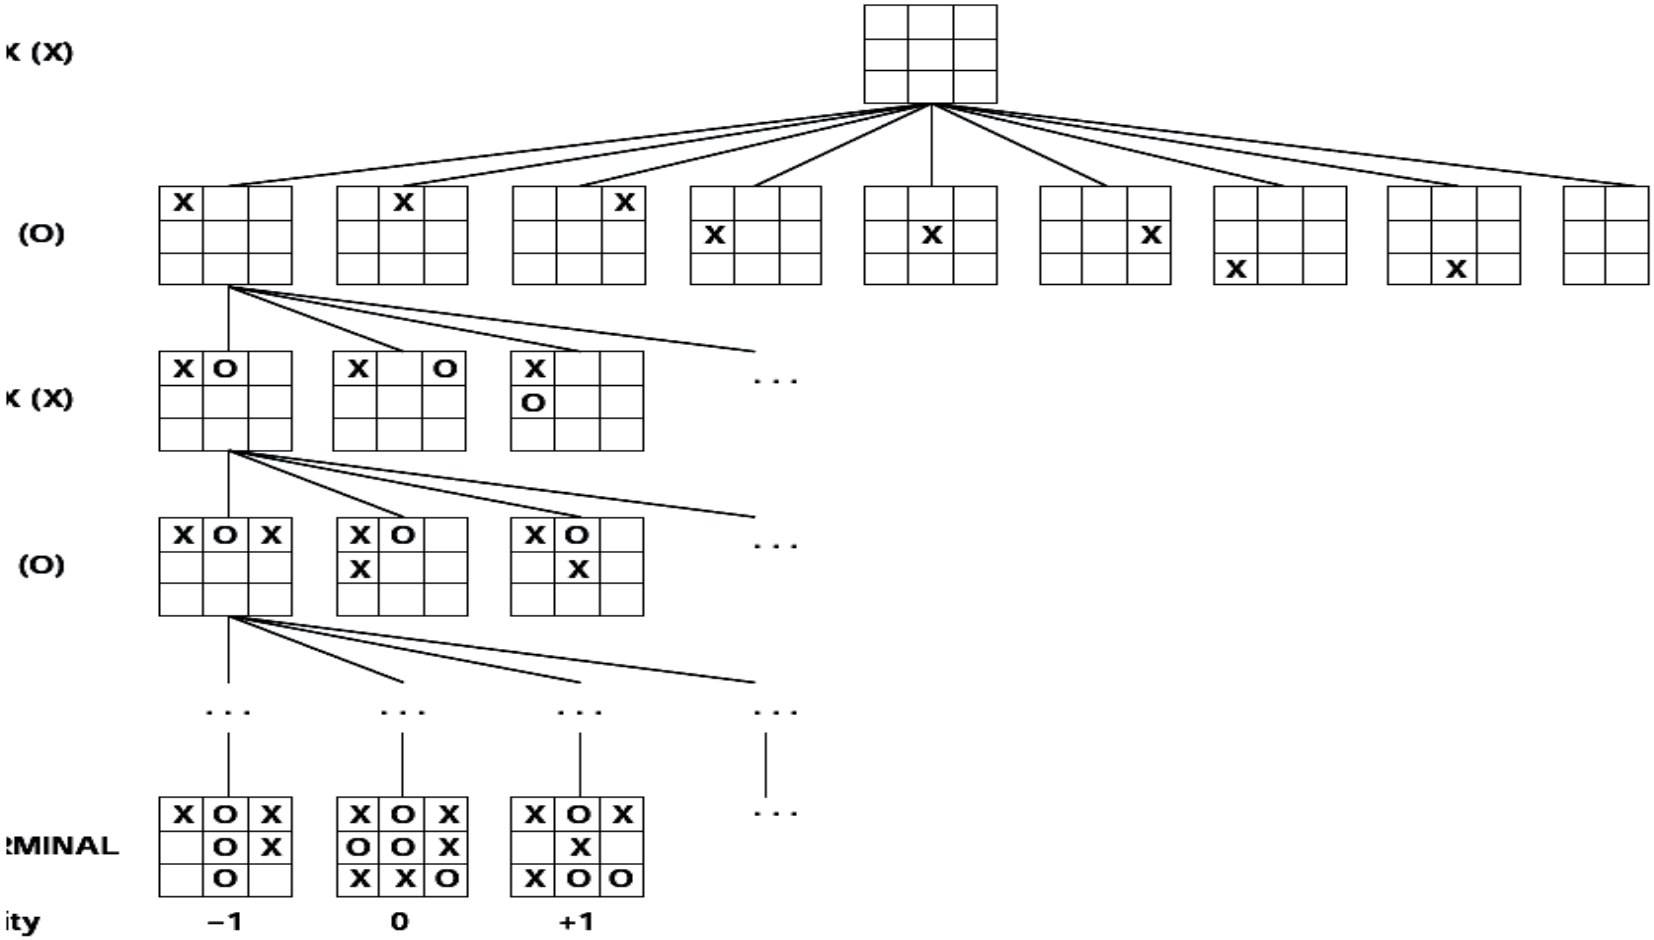

Propositional Logic

Propositional knowledge base:

$$P \implies Q$$
$$L \wedge M \implies P$$
$$B \wedge L \implies M$$
$$A \wedge P \implies L$$
$$A \wedge B \implies L$$
$$A$$
$$B$$

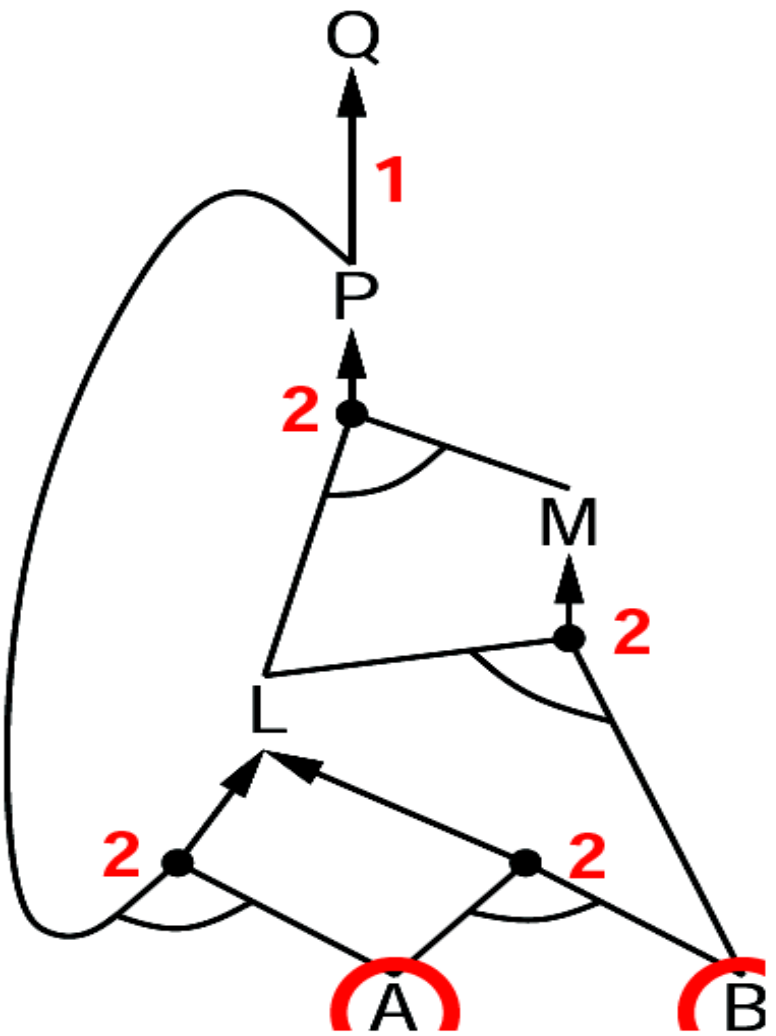

First-order Logic

FOL knowledge base:

$\text{American}(x) \wedge \text{Weapon}(y) \wedge \text{Sells}(x, y, z) \wedge \text{Hostile}(z) \implies \text{Criminal}(x)$

$\text{Owns}(\text{Nono}, M_1)$

$\text{Missile}(M_1)$

$\forall x \text{Missile}(x) \wedge \text{Owns}(\text{Nono}, x) \implies \text{Sells}(\text{West}, x, \text{Nono})$

$\text{Missile}(x) \Rightarrow \text{Weapon}(x)$

$\text{Enemy}(x, \text{America}) \implies \text{Hostile}(x)$

$\text{American}(\text{West})$

$\text{Enemy}(\text{Nono}, \text{America})$

First-order Logic

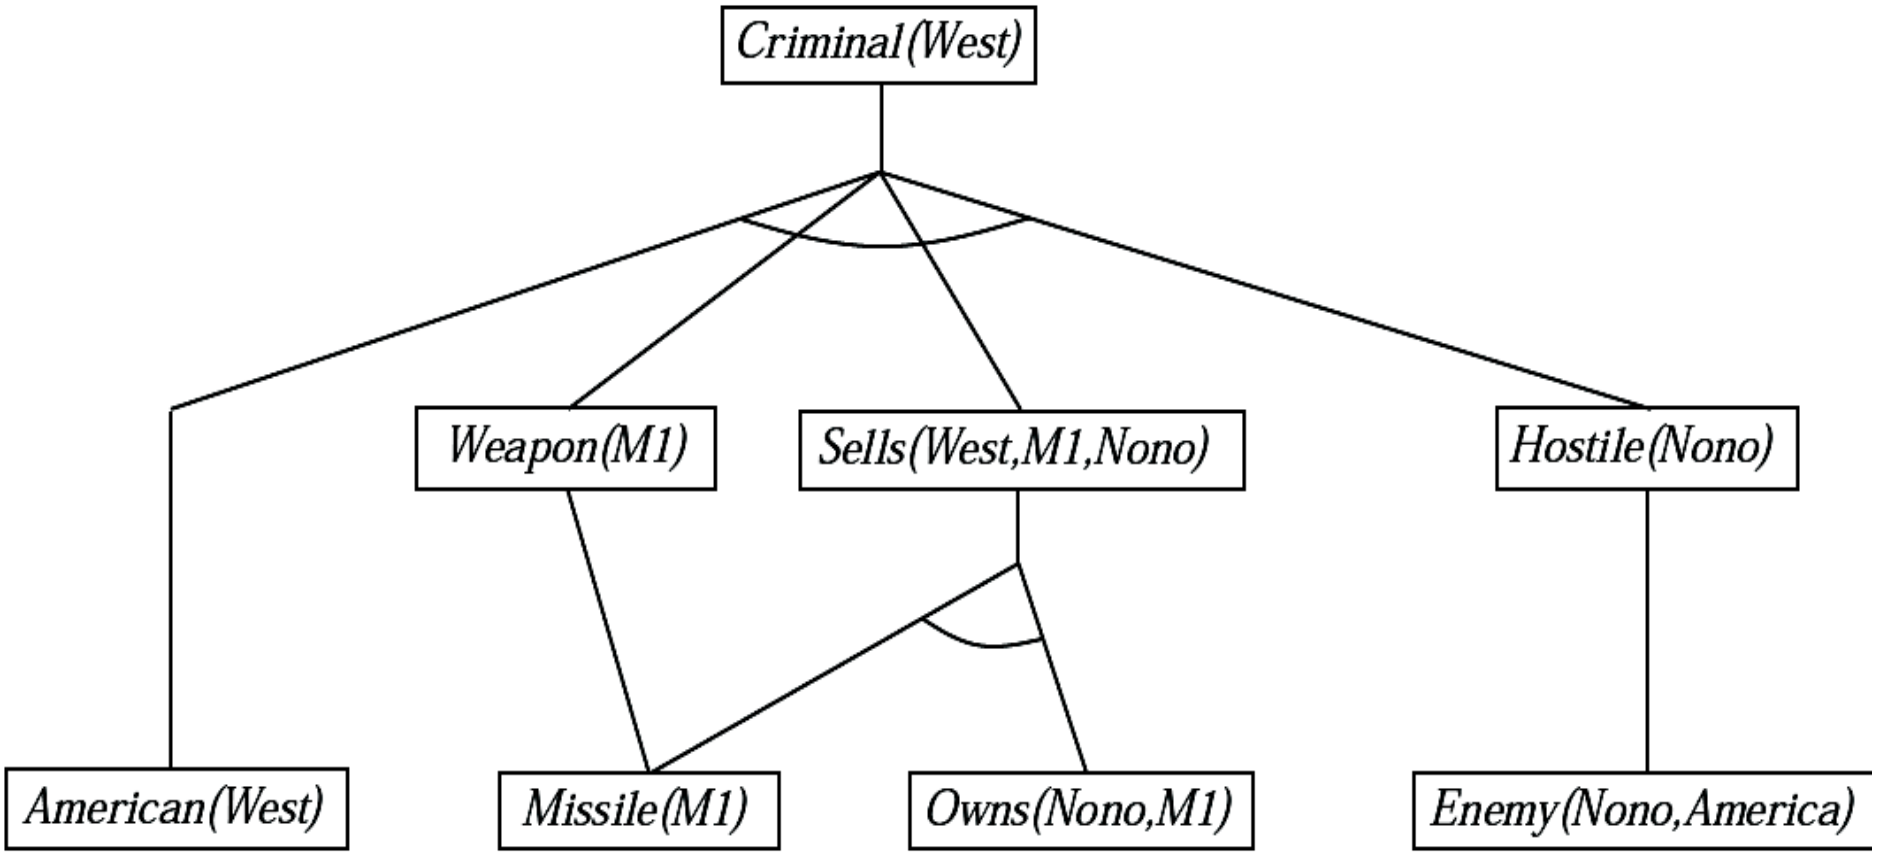

Logic Programming (Prolog)

- Programming in
  - Procedural languages (Java, C++, …)
    - "tell" the computer "what to do", i.e. what steps to provide, to achieve a desired result

```
for (Student student : students)
    if (student.age > 21 && student.sex == "f")
        System.out.println(student.id + "\t" + student.name);
```

  - Declarative languages (Prolog, SQL, …)
    - "tell" the computer "what information we want to get" without telling the exact steps to get it.

```
provide(X,Y) :- student(X,Y,V,Z), V > 21, Z = f.
```

```
SELECT id, name FROM students WHERE age > 21 AND sex = 'f'
```

Machine Learning (basics only)

- ML is a research area “combining”
  - statistics
  - optimization theory
- how can we derive some patterns (rules) from the observed data?
  - what is a pattern?
  - how to measure the accuracy?
  - how those patterns ease the decision process?

Inductive Logic Programming

Learning `daughter/2`

INPUT

Training examples		*Background knowledge*
$daughter(mary, ann).$	$\oplus$	$mother(ann, mary).$ $female(ann).$
$daughter(eve, tom).$	$\oplus$	$mother(ann, tom).$ $female(mary).$
$daughter(tom, ann).$	$\ominus$	$father(tom, eve).$ $female(eve).$
$daughter(eve, ann).$	$\ominus$	$father(tom, ian).$
		$parent(X, Y) \leftarrow mother(X, Y)$
		$parent(X, Y) \leftarrow father(X, Y)$

OUTPUT

$$daughter(X, Y) \leftarrow female(X), parent(Y, X)$$

Textbook

- Stuart Russell and Peter Norvig, Artificial Intelligence – A Modern Approach, Prentice Hall 2003.

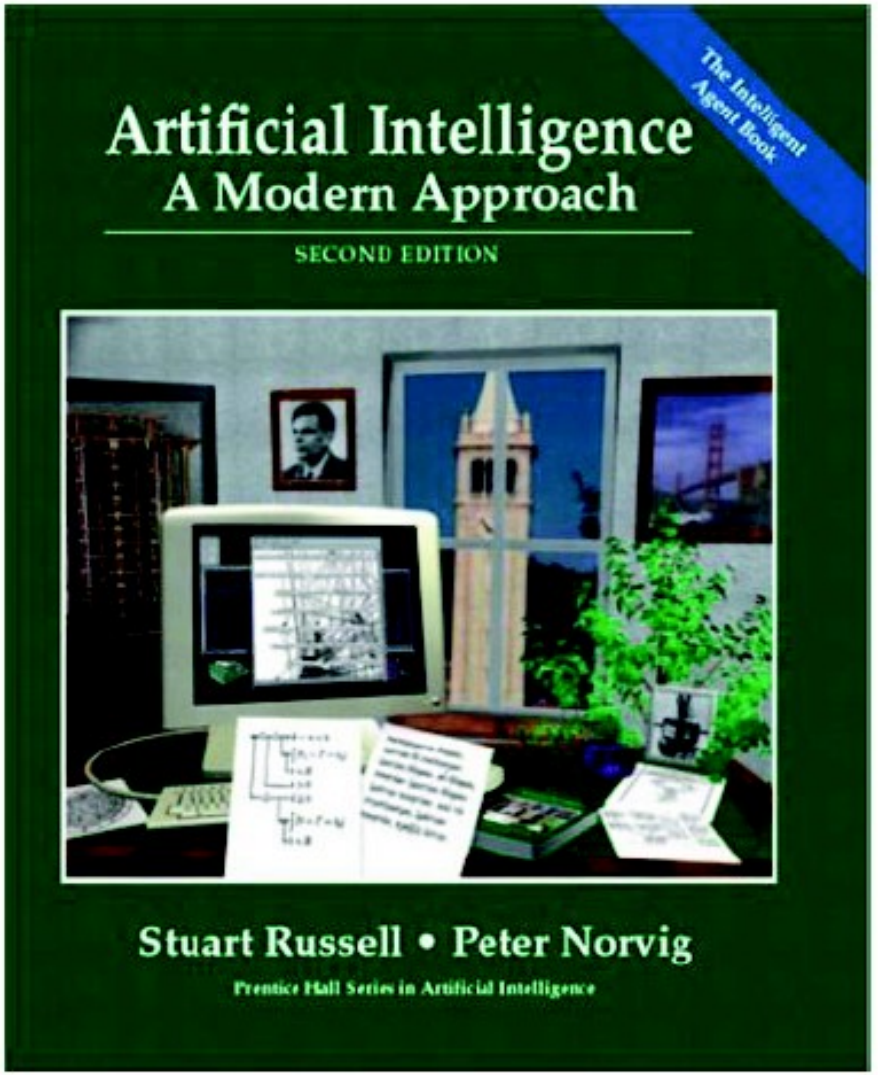

Administrative things

Lectures

- Lectures
  - Wednesday 8:00 - 10:00
  - Thursday 10:00 - 12:00 (bi-weekly)
  - Final examination (test) grade A

- Lectures in english

Tutorials

- Tutorials
  - Monday 16:00 - 18:00 (bi-weekly)
  - Final grade for the exercises B

- Tutorials in english

Final Grade

- Final grade from the course will be computed as (9.A + B) / 10
  - successful, but unsatised students will be recommended to make an additional work, e.g. a project, paper presentation, . . .
  - unsuccessful students will be recommended to study better in the future

Thanks for Your attention!

Questions?