Artificial Intelligence

Lars Schmidt-Thieme

Information Systems and Machine Learning Lab (ISMLL)
Institute of Economics and Information Systems
& Institute of Computer Science
University of Hildesheim
http://www.ismll.uni-hildesheim.de

1. What is Artificial Intelligence?

2. Overview

3. Organizational stuff

What is Artificial Intelligence?

Systems that ...

think like humans	**think rationally**
act like humans	**act rationally**

Acting humanly: Turing Test

Turing (1950) “Computing machinery and intelligence”:

- “Can machines think?”
  $\rightsquigarrow$ “Can machines behave intelligently?”

- Operational test for intelligent behavior:
  the Imitation Game.

- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes.

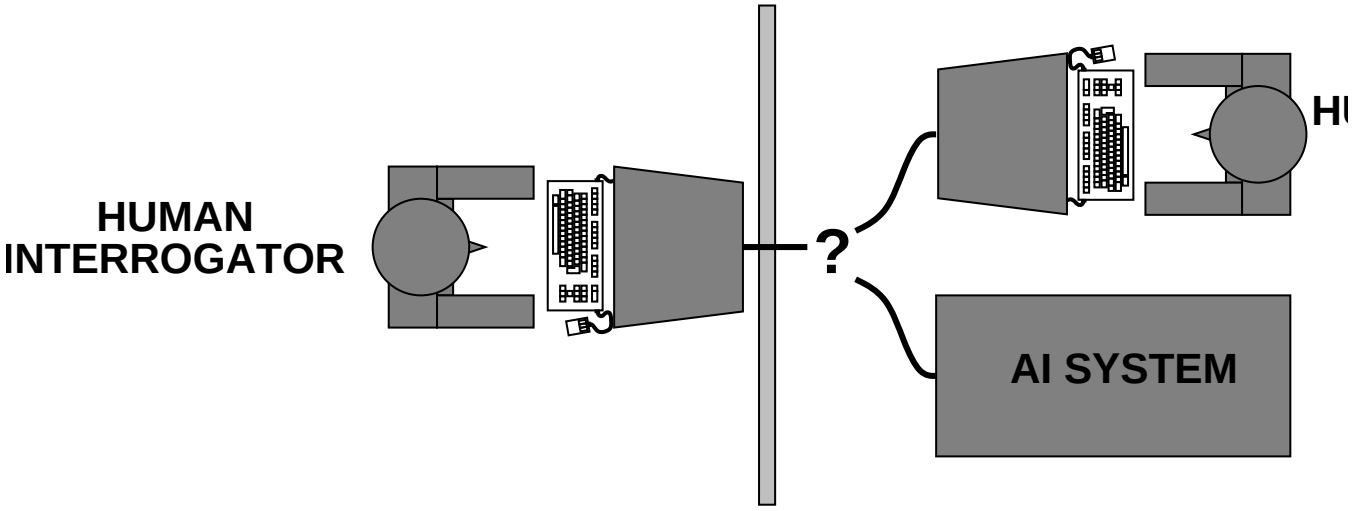

Turing Test / Loebner Price

Loebner Price:

- formal instantiation of the Turing Test.
- each year since 1991.
- $100.000 for the first computer whose responses were indistinguishable from a human's
- web chat interface.

Turing Test / Jabberwacky

An example conversation with Jabberwacky
(http://www.jabberwacky.com),
the Loebner Prize winner of 2005.

Jabberwacky: It is impossible to defend perfectly against the attack of those who want to die.

Lars: Helllo.

Jabberwacky: What is your name?

Lars: I am Lars. And who are you?

Jabberwacky: My name is Clint, nice to meet you.

Lars: Nice to meet you. What are you doing here, Clint?

Jabberwacky: I talk to you.

Lars: Well, yes.

Jabberwacky: How do you talk to people.

Lars: I meant what are you looking for in this chat room?

Jabberwacky: I am chatting at my computer.

A short history of AI (Russell/Norvig 2003)

1943	McCulloch & Pitts: Boolean circuit model of brain
1950	Turing's "Computing Machinery and Intelligence"
1952–69	Look, Ma, no hands!
1950s	Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
1956	Dartmouth meeting: "Artificial Intelligence" adopted
1965	Robinson's complete algorithm for logical reasoning
1966–74	AI discovers computational complexity Neural network research almost disappears
1969–79	Early development of knowledge-based systems
1980–88	Expert systems industry booms
1988–93	Expert systems industry busts: "AI Winter"
1985–95	Neural networks return to popularity
1988–	Resurgence of probability; general increase in technical depth "Nouvelle AI": ALife, GAs, soft computing
1995–	Agents, agents, everywhere . . .
2003–	Human-level AI back on the agenda

Examples of AI accomplishments

Autonomous Planning and Scheduling: NASA's Remote Agent Program (Jonsson et al. 2000).

Game Playing: IBM's Deep Blue bested Gary Kasparov (Goodman and Keene 1997).

Autonomous Control: DARPA Grand Challenge 2005: autonomous vehicle finds a 132 miles path over desert terrain.

Diagnosis: Lympn-node pathology diagnosis system corrects human expert (Heckerman 1991).

Logistics Planning: DART (Cross and Walker, 1994) planned the logistics in the First Persian Gulf War 1991.

Robotics: HipNav (DiGioia et al. 1996) uses computer vision to guide the insertion of a hip replacement prosthesis.

Language Understanding:
ProVerb (Littman et al. 1999) solves crossword puzzles.

1. What is Artificial Intelligence?
2. Overview
3. Organizational stuff

Russell & Norvig Textbook TOC

II. Problem-Solving
3 - Searching
4 - Informed Search/Exploration
5 - Constraint Satisfaction Problems
6 - Adversarial Search

II. Knowledge and Reasoning
7 - Propositional Logic
8/9 - First Order Logic
10 - Knowledge Representation

III. Planning
11 - Planning
12 - Planning and Acting in the Real World

IV. Uncertain Knowledge and Reasoning
13 - Uncertainty
14 - Probabilistic Reasoning
15 - Probabilistic Reasoning over Time
16 - Making Simple Decisions
17 - Making Complex Decisions

V. Learning
18 - Learning from Observations
19 - Knowledge in Learning
20 - Statistical Learning Methods
21 - Recinforcement Learning

VI. Example Applications
22 - Communication
23 - Probabilistic Language Processing
24 - Perception
25 - Robotics

Searching (1/2)

Find shortest way from Arad to Bucharest.

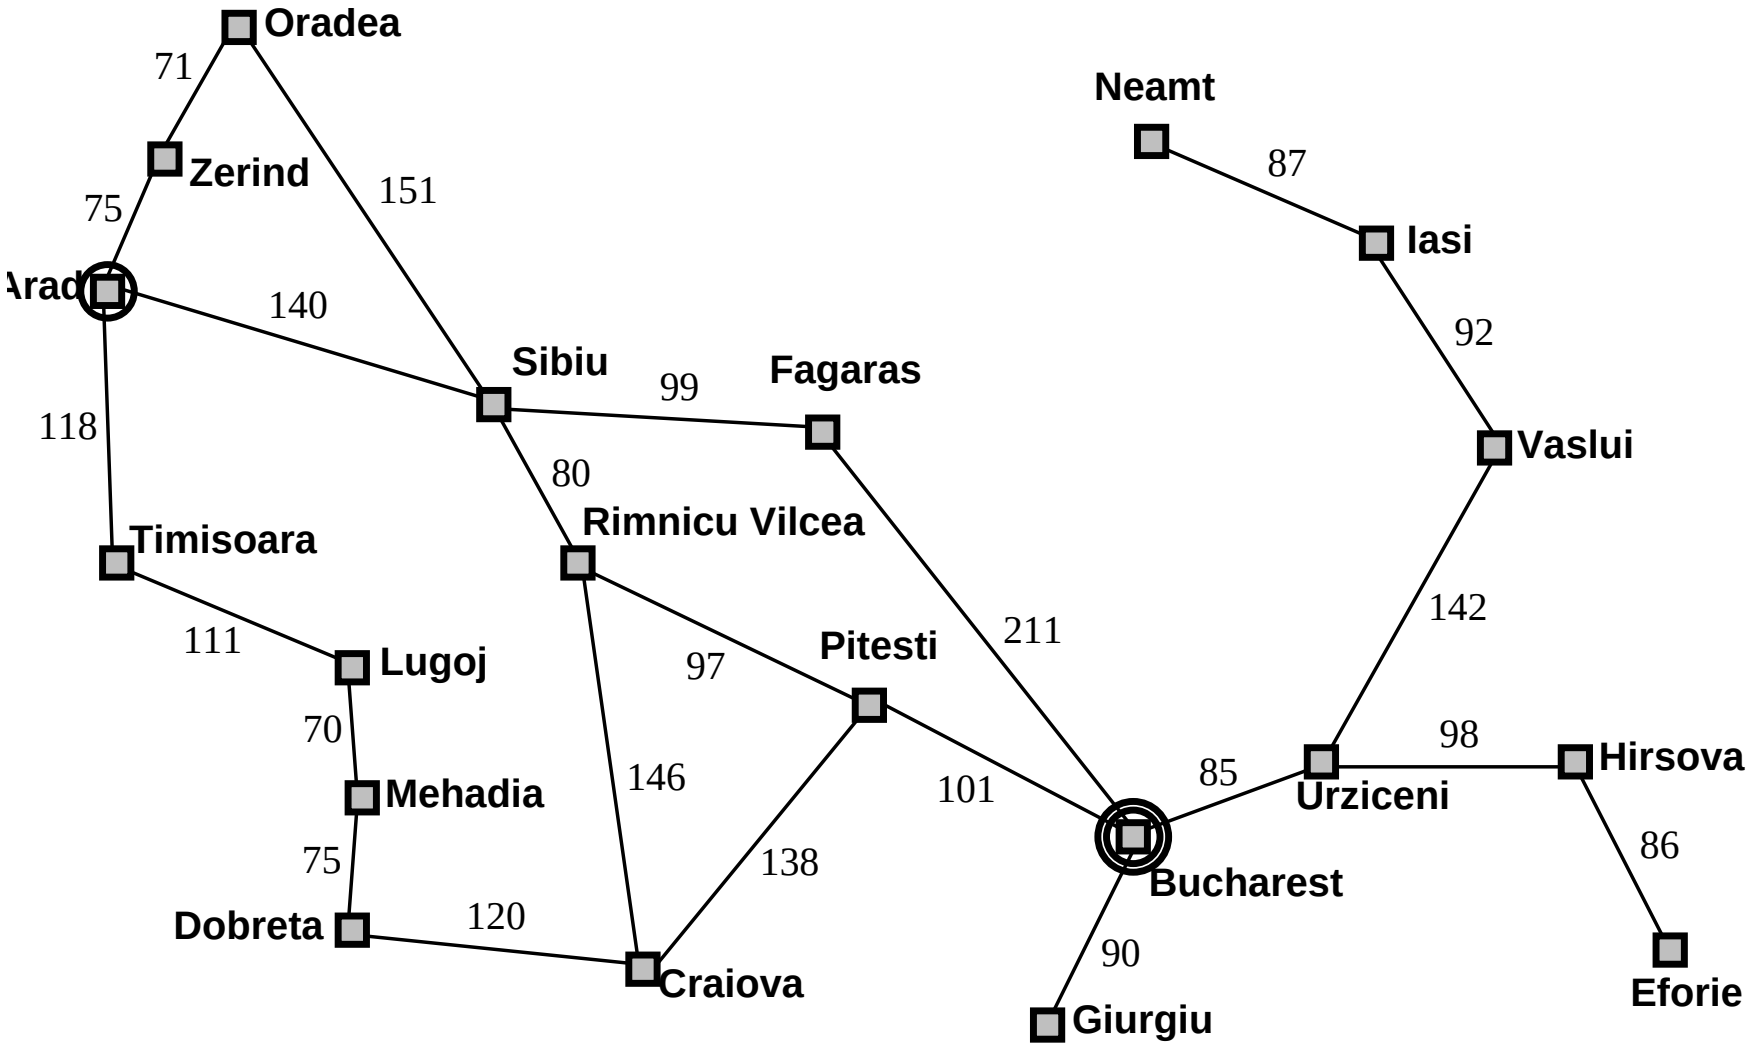

Searching (2/2)

Several strategies: breadth-first, depth-first, etc.

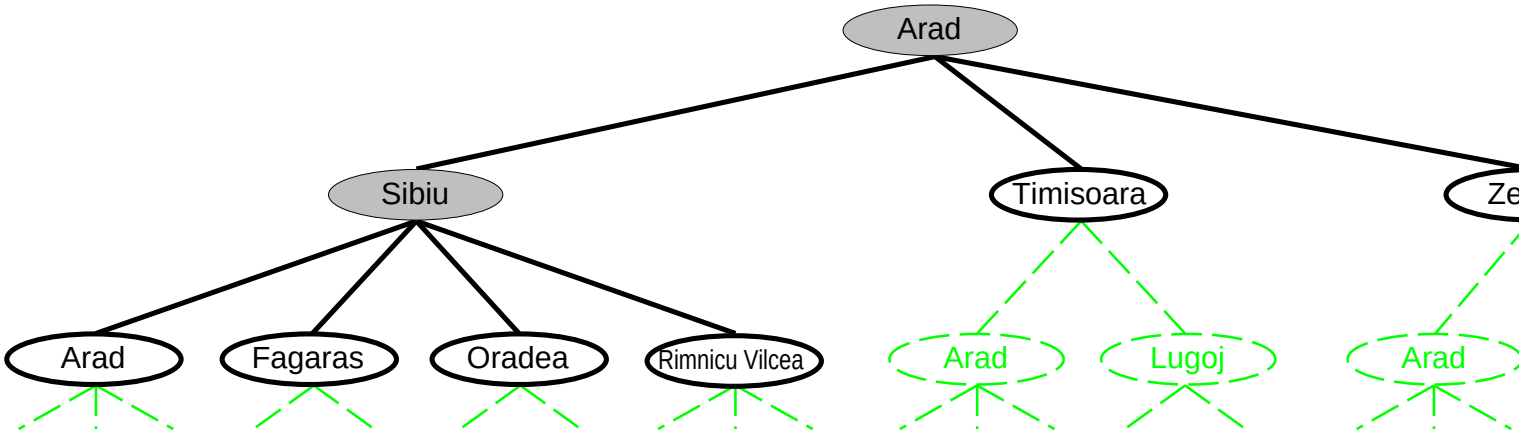

Informed Search (1/2)

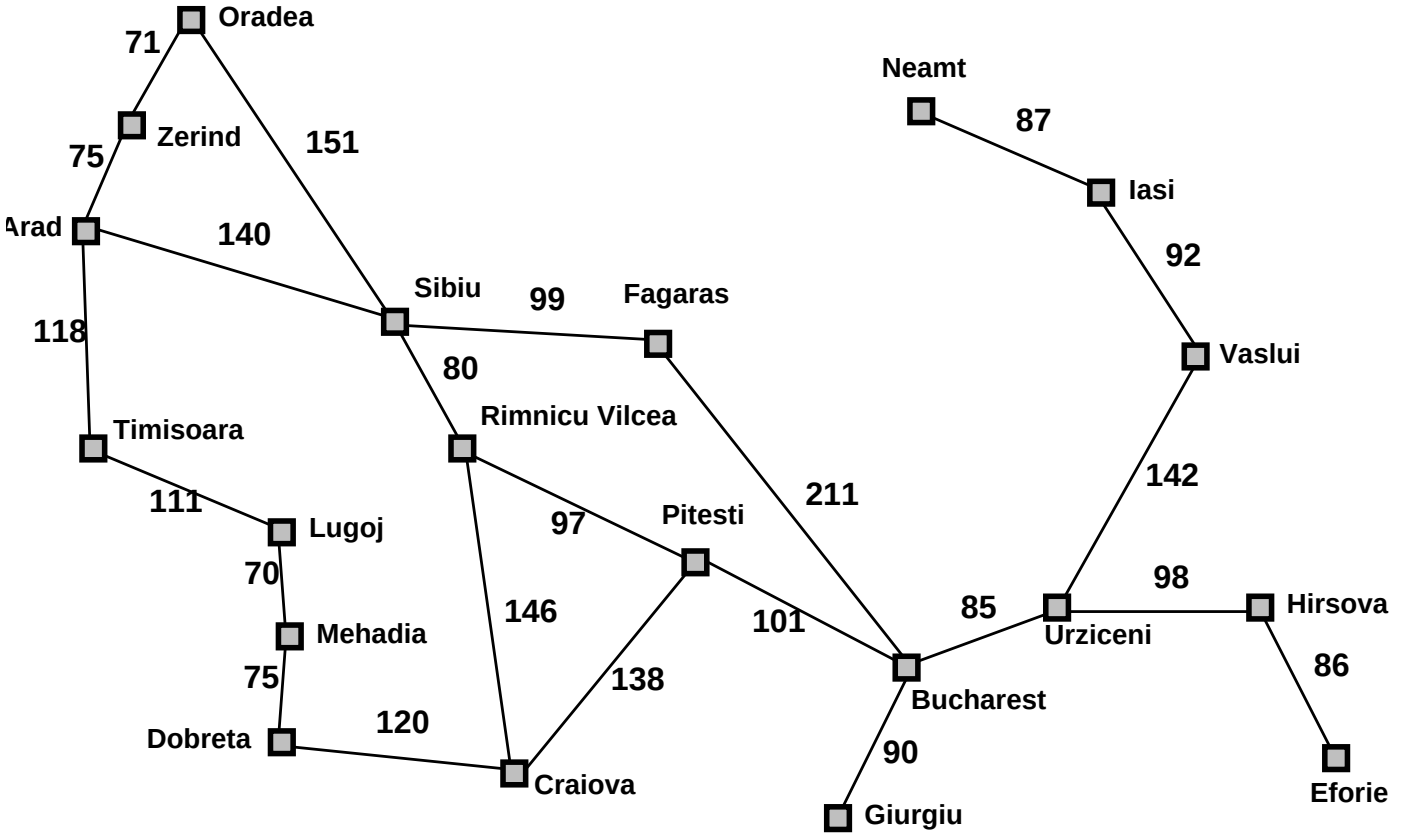

Straight−line distance to Bucharest	
Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	178
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	98
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

Informed Search (2/2)

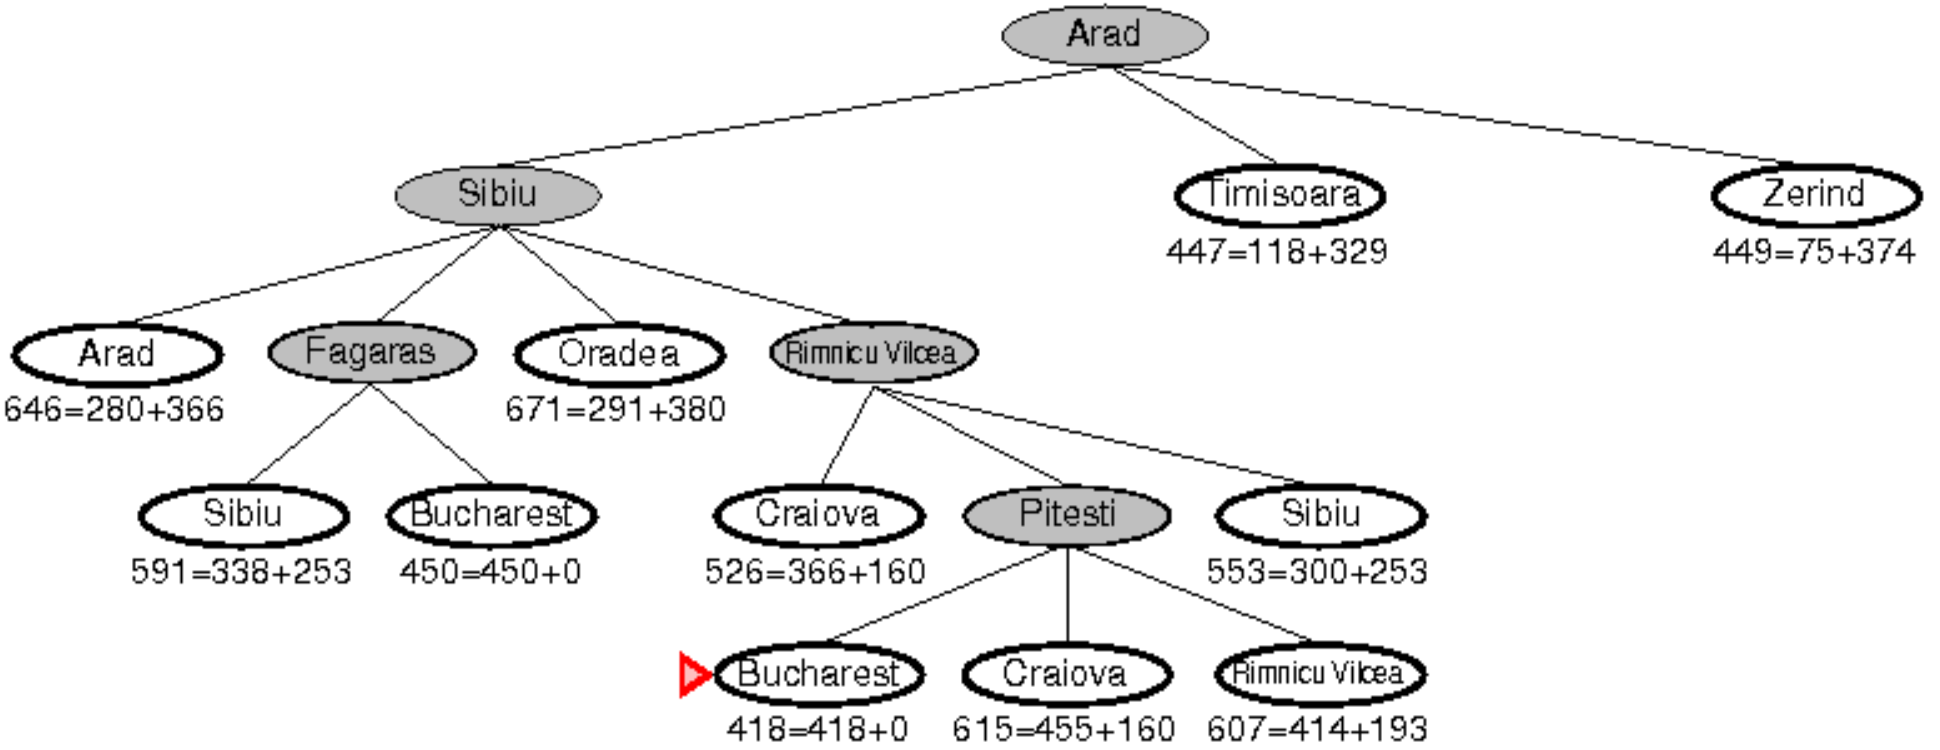

Constraint Satisfaction Problems

Color a map with 3 colors s.t. no two adjacent regions have the same color.

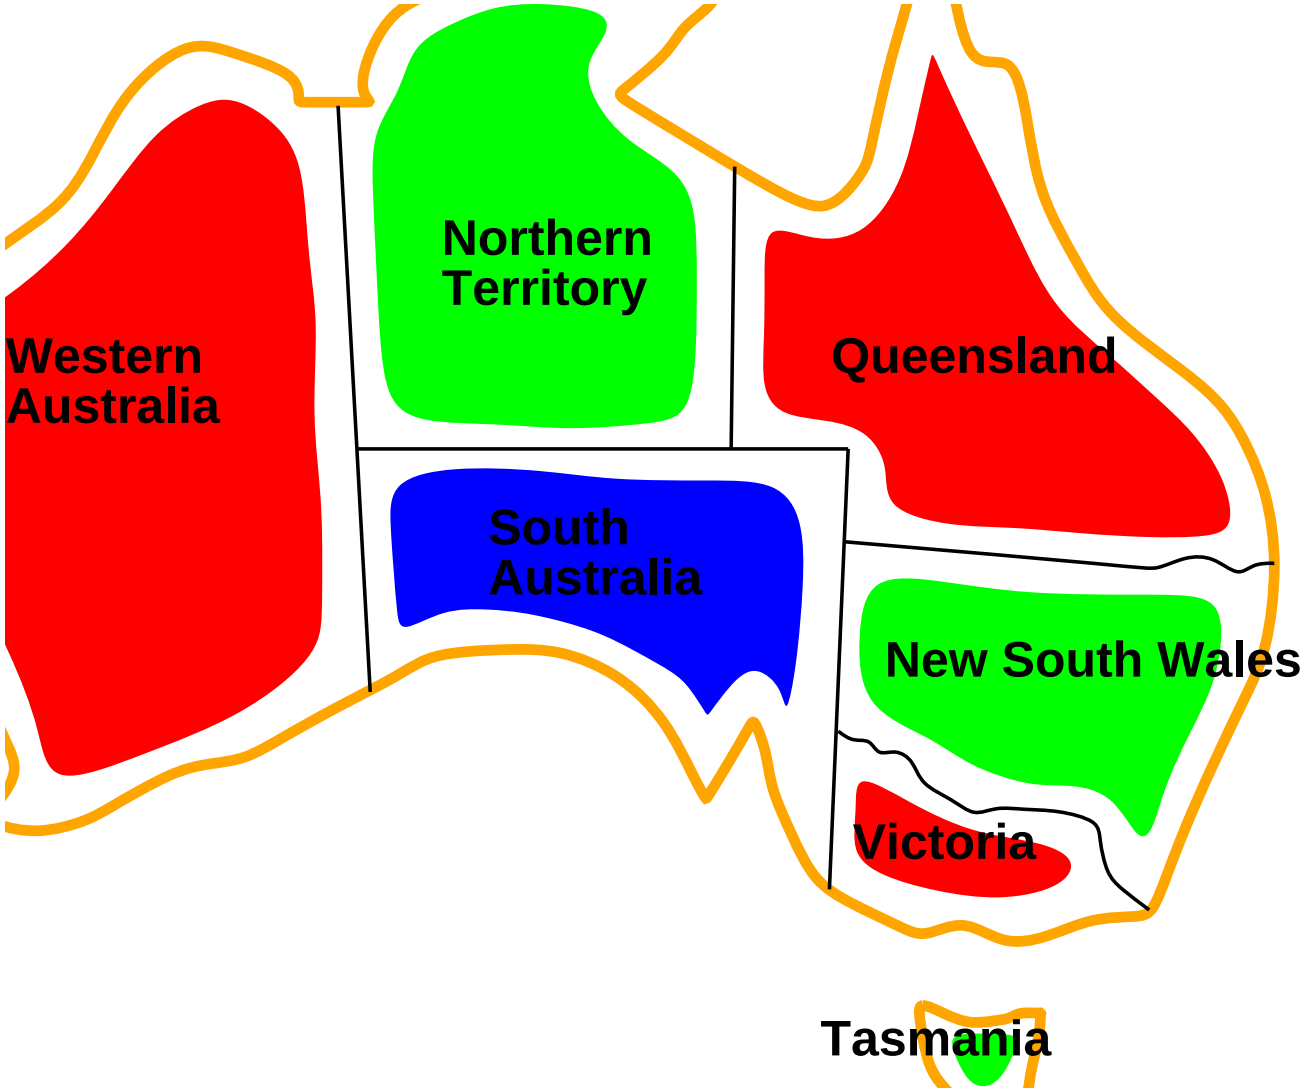

Adversarial Search

Game against an opponent: specify an action for every possible reply.

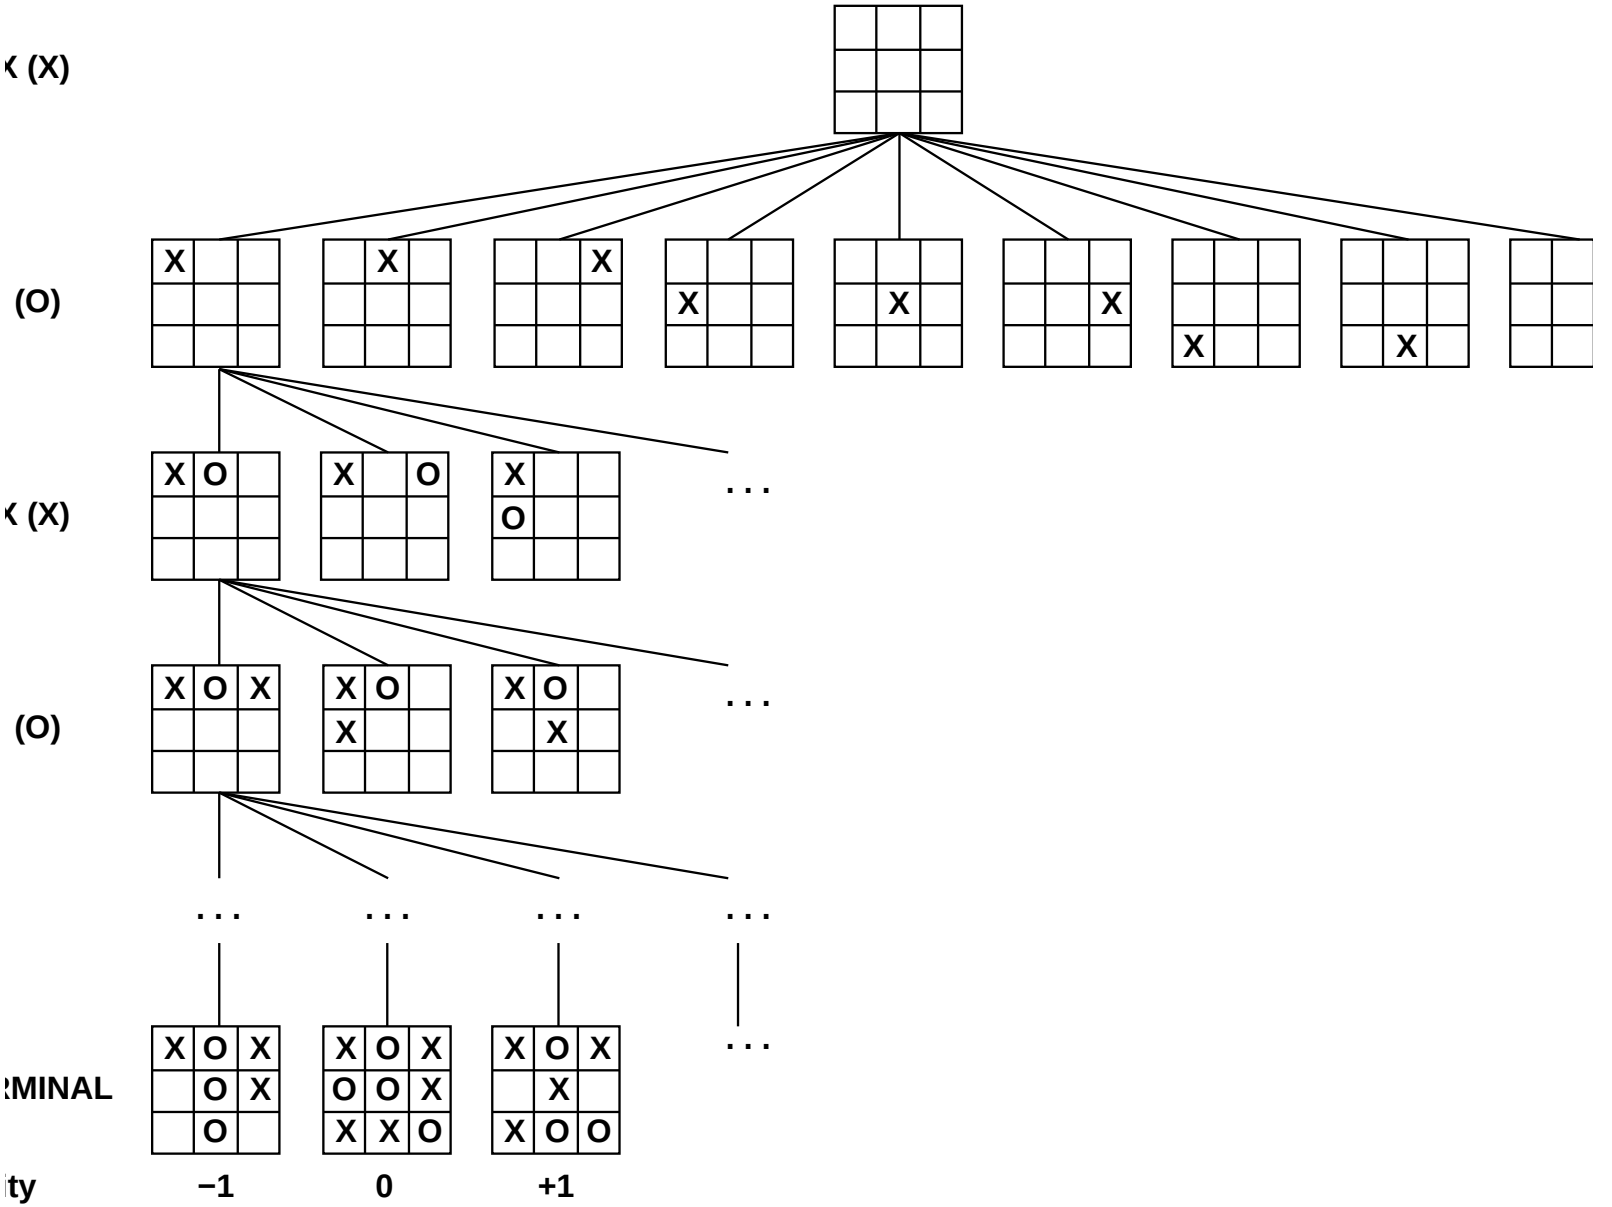

Propositional Logic

Propositional knowledge base:

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

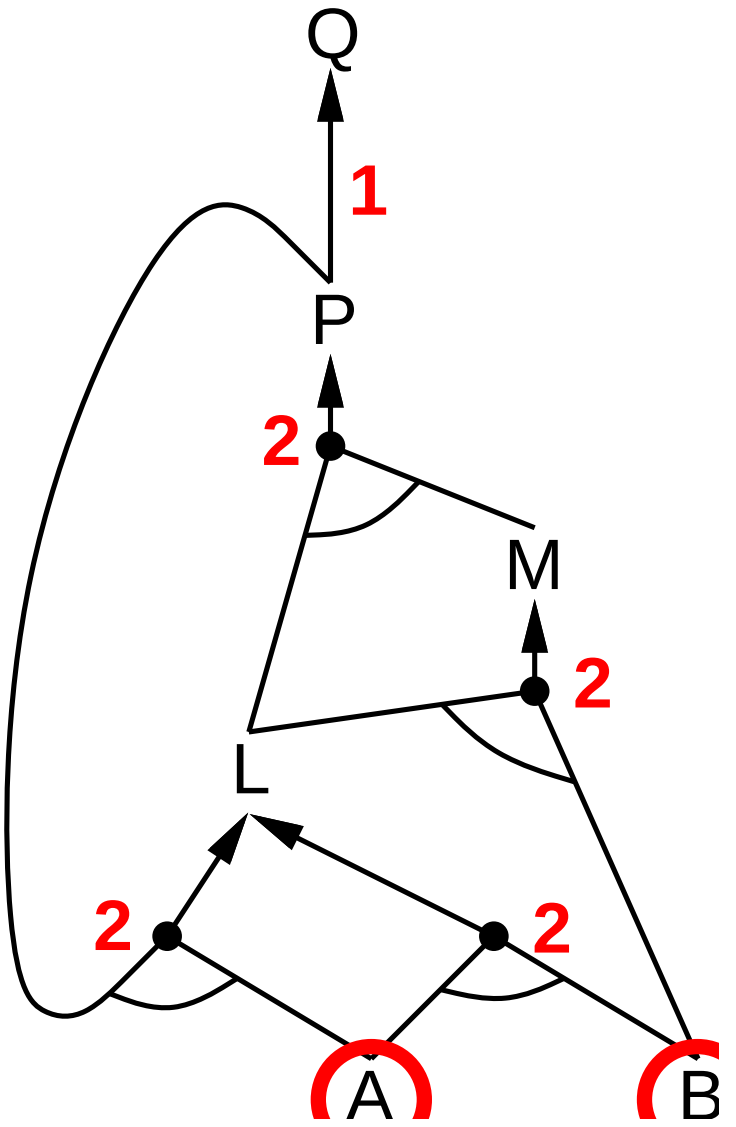

First Order Logic (1/2)

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 (2/2)

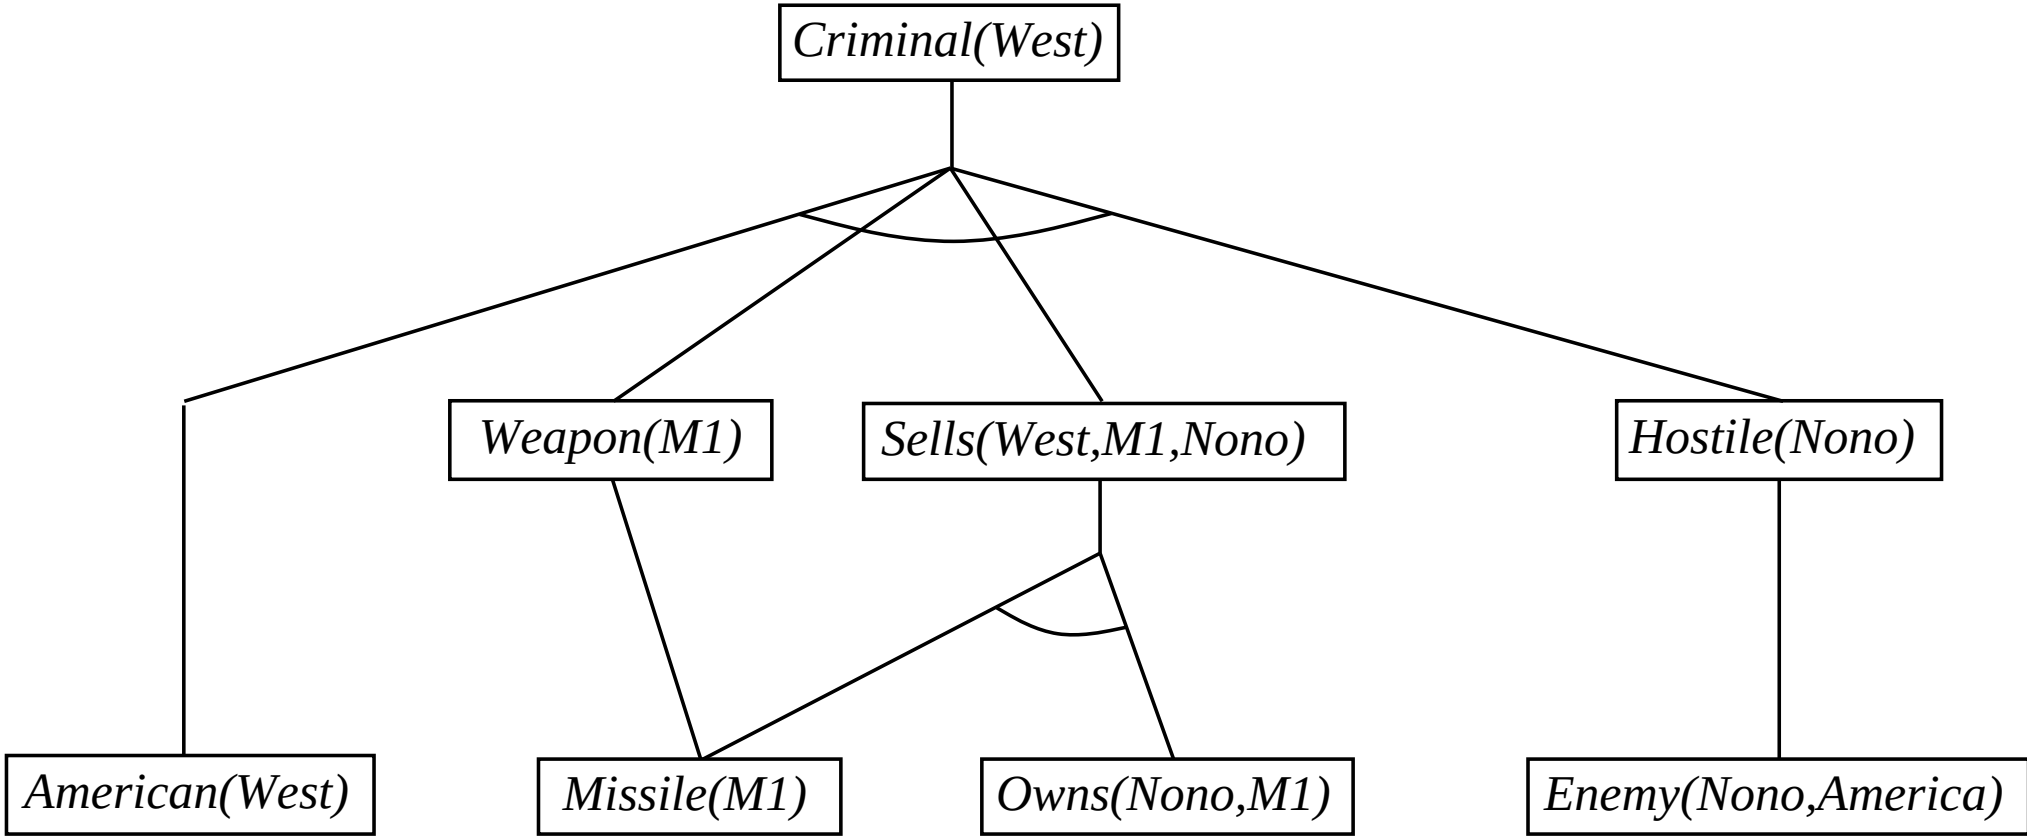

Planning

Blocks World: move one block at a time s.t. a given goal configuration is reached.

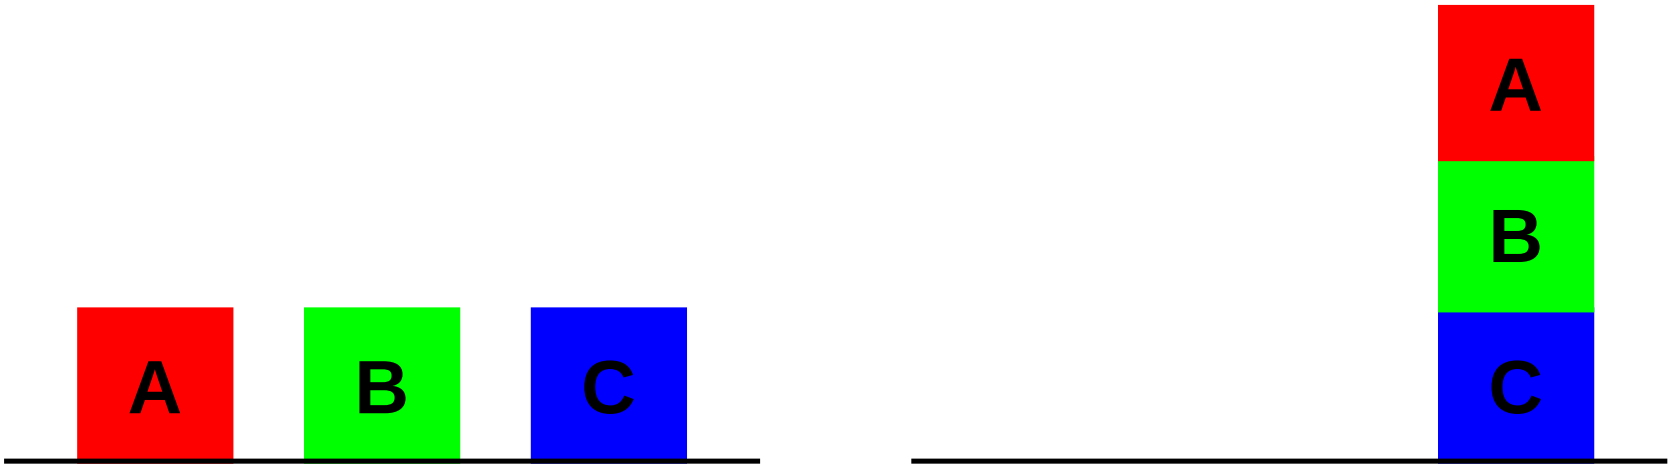

Inductive Logic Programming

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

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

[Jamens Cussens 2005]

1. What is Artificial Intelligence?
2. Overview
3. Organizational stuff

Textbook

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

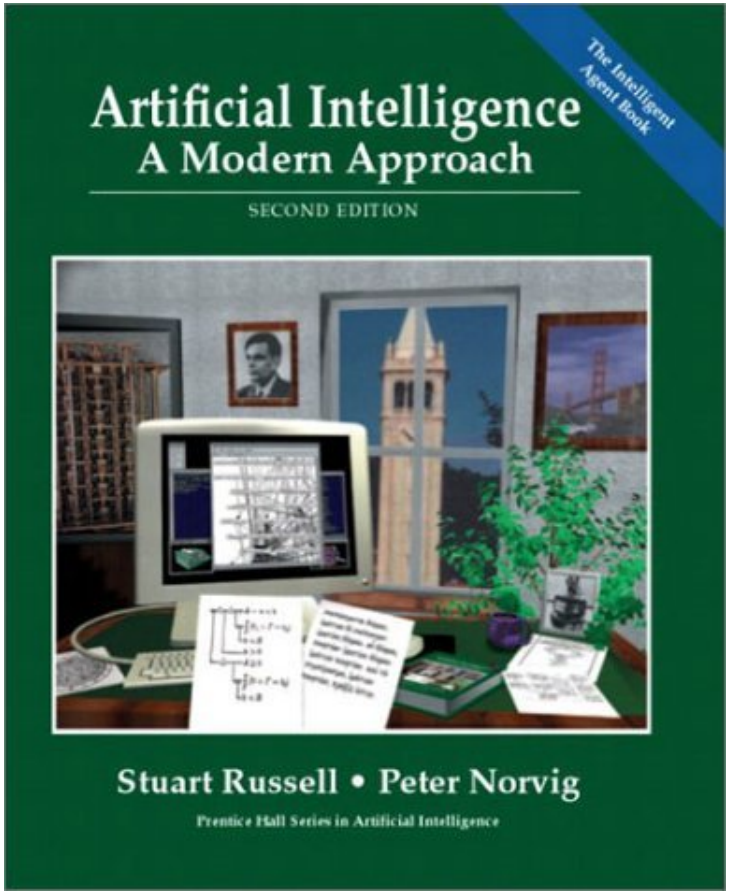

Exercises and Tutorials

- There will be a weekly sheet with two exercises handed out **each Wednesday** in the lecture.
  1st sheet will be handed out this Wed. 4.11.

- Solutions to the exercises can be submitted until **every next Wednesday before the lecture** in the letter box
  1st sheet is due Wed. 11.11.

- Exercises will be corrected by your tutor.

- Tutorials **each Monday 14-16**,
  1st tutorial at Mon. 9.11.

Exam and credit points

- There will be an exam at end of term (2h, 4 problems).
- You can get up to 10% of the points as bonus points from the tutorial.
- The course gives 8 ECTS.
- This is a Bachelor course that can be used in
  - BSc Wirtschaftsinformatik / Gebiet KI & ML
  - BSc IMIT (neu) / Gebiet KI & ML
  - BSc IMIT (alt) / Modul IT3 Machine Learning.
  - but not in any MSc program!

References

[RN03] Stuart Russell and Peter Norvig. *Artificial Intelligence – A Modern Approach*. Prentice Hall, 2003.