

CS4618 Artificial Intelligence I

Today: Introduction

Thomas Jansen

September 26th

Plans for Today

① Introduction

Overview

Formal Things

② Foundations of AI

'Definition' and Related Fields

History of AI

③ Intelligent Agents

Introduction and Task Environments

④ Summary

Summary & Take Home Message




Artificial Intelligence



Artificial Intelligence





Artificial Intelligence





Artificial Intelligence






Artificial Intelligence







Artificial Intelligence








Artificial Intelligence








Artificial Intelligence









Artificial Intelligence



Artificial Intelligence



Artificial Intelligence



Boring Formal Stuff

Boring Formal Stuff

Lecturer Thomas Jansen
Office WGB, 1-82
E-Mail t.jansen@cs.ucc.ie

Boring Formal Stuff

Lecturer Thomas Jansen
Office WGB, 1-82
E-Mail t.jansen@cs.ucc.ie

What and where

Lecture Wednesdays, 2pm–3pm WGB 15
Fridays, 11am–12pm WGB 09

Practicals

Boring Formal Stuff

Lecturer Thomas Jansen
Office WGB, 1-82
E-Mail t.jansen@cs.ucc.ie

What and where

Lecture	Wednesdays, 2pm–3pm	WGB 15
	Fridays, 11am–12pm	WGB 09
Practicals	to be announced	

Boring Formal Stuff

Lecturer Thomas Jansen
Office WGB, 1-82
E-Mail t.jansen@cs.ucc.ie

What and where

Lecture Wednesdays, 2pm–3pm WGB 15
Fridays, 11am–12pm WGB 09

Practicals to be announced

Assessment Weekly Assignments 10 marks

Boring Formal Stuff

Lecturer Thomas Jansen
Office WGB, 1-82
E-Mail t.jansen@cs.ucc.ie

What and where

Lecture	Wednesdays, 2pm–3pm	WGB 15
	Fridays, 11am–12pm	WGB 09
Practicals	to be announced	
Assessment	Weekly Assignments	10 marks
	In-Class Test	10 marks

Boring Formal Stuff

Lecturer Thomas Jansen
Office WGB, 1-82
E-Mail t.jansen@cs.ucc.ie

What and where

Lecture Wednesdays, 2pm–3pm WGB 15
 Fridays, 11am–12pm WGB 09
Practicals to be announced

Assessment Weekly Assignments 10 marks
 In-Class Test 10 marks
 End of Year Examination 80 marks
 1 × 1.5 hrs paper



Course Material


Course Material

Slides <http://www.cs.ucc.ie/~tj2/cs4618>

Course Material

Slides <http://www.cs.ucc.ie/~tj2/cs4618>

Books S. Russel, P. Norvig
Artificial Intelligence. A Modern Approach.
Third Edition. Pearson, 2010.




Course Material

Slides <http://www.cs.ucc.ie/~tj2/cs4618>

Books S. Russel, P. Norvig
Artificial Intelligence. A Modern Approach.
 Third Edition. Pearson, 2010.



T. Jansen
Analyzing Evolutionary Algorithms.
 Springer, 2013.





Course Outline

Course Outline

① Introduction

- 'Definition'
- (a bit of) History
- Intelligent Agents

Course Outline

① Introduction

- 'Definition'
- (a bit of) History
- Intelligent Agents

② 'Classical' AI

- Uninformed Search (BFS, DFS)
- Informed Search (A*)
- Game Play

Course Outline

① Introduction

- 'Definition'
- (a bit of) History
- Intelligent Agents

② 'Classical' AI

- Uninformed Search (BFS, DFS)
- Informed Search (A*)
- Game Play

③ Randomised Search Heuristics

- Local Search
- Simulated Annealing
- Evolutionary Algorithms



Point of View: Computer Science

Point of View: Computer Science

I think that it's extraordinarily important that we in computer science keep fun in computing. When it started out, it was an awful lot of fun. Of course, the paying customers got shafted every now and then, and after a while we began to take their complaints seriously. We began to feel as if we really were responsible for the successful, error-free perfect use of these machines. I don't think we are. I think we're responsible for stretching them, setting them off in new directions, and keeping fun in the house. I hope the field of computer science never loses its sense of fun.

Alan Perlis (1922–1990)

(taken from H. Abelson, G. J. Sussmann, J. Sussman (1985):
Structure and Interpretation of Computer Programs. MIT Press.)

What is Artificial Intelligence?

What is Artificial Intelligence?

thinking

What is Artificial Intelligence?

acting

thinking

What is Artificial Intelligence?

acting

thinking

rationally

What is Artificial Intelligence?

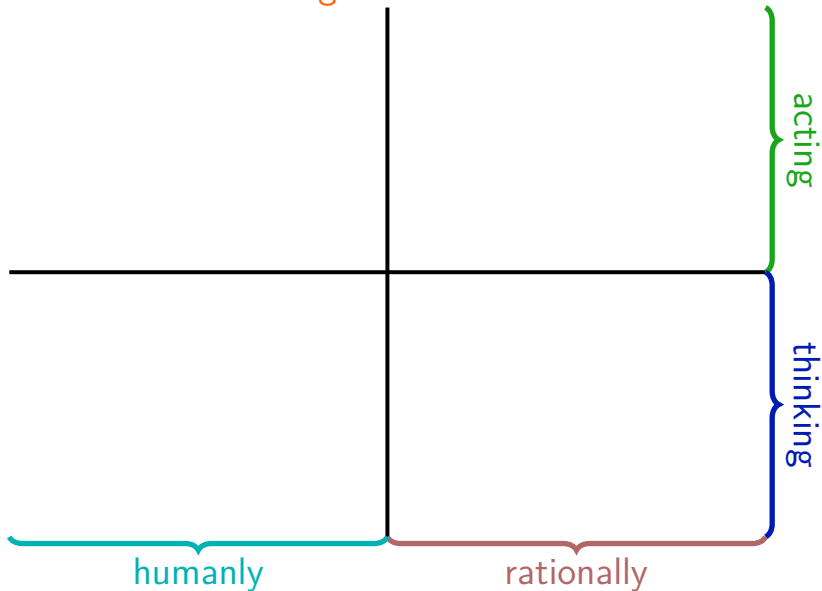
acting

thinking

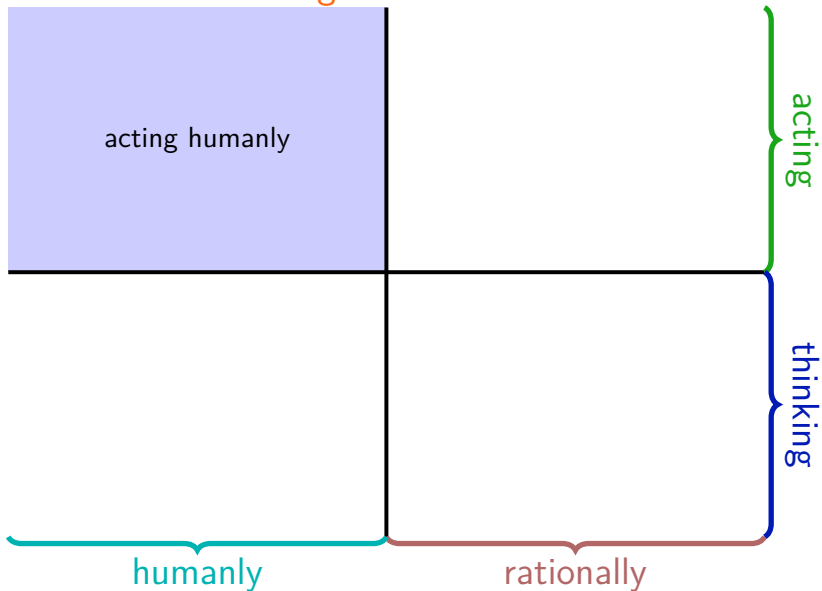
humanly

rationally

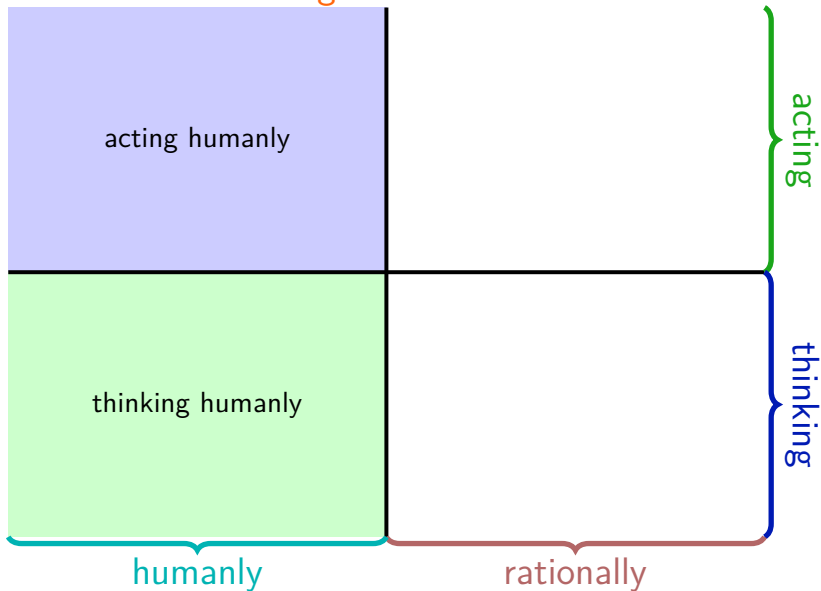
What is Artificial Intelligence?



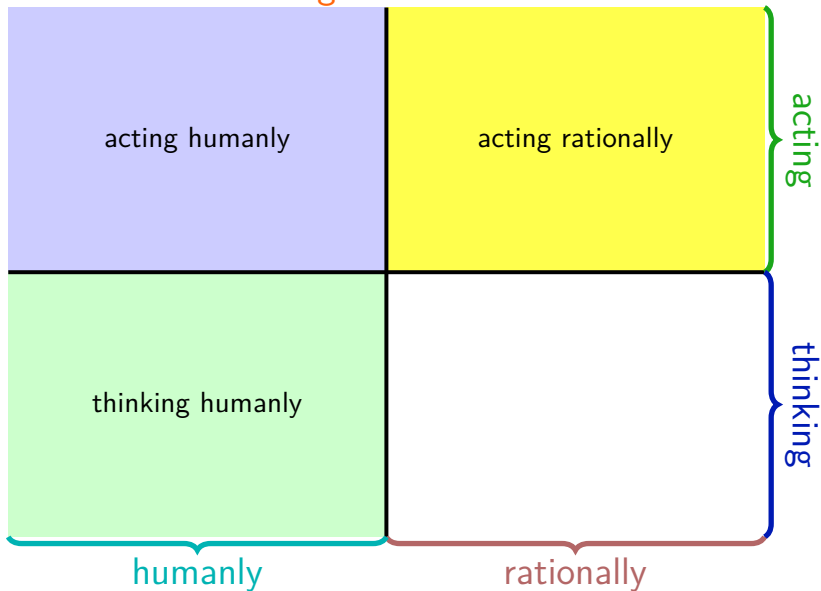
What is Artificial Intelligence?



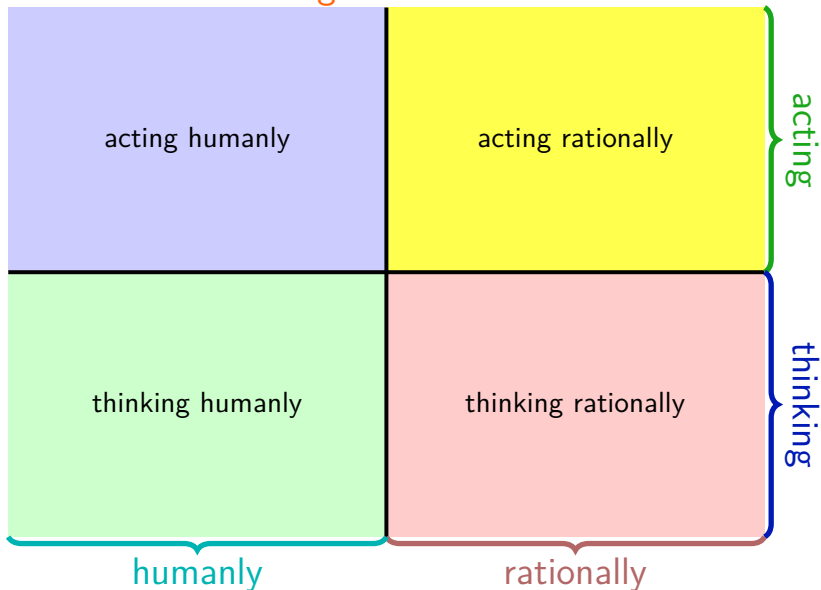
What is Artificial Intelligence?



What is Artificial Intelligence?



What is Artificial Intelligence?



Different Answers to 'What is Artificial Intelligence?'

Different Answers to 'What is Artificial Intelligence?'

- acting humanly

Different Answers to 'What is Artificial Intelligence?'

- acting humanly

Turing test

A. Turing (1950): Computing Machinery and Intelligence.
Mind 59(236):433–460.

Different Answers to 'What is Artificial Intelligence?'

- acting humanly

Turing test

A. Turing (1950): Computing Machinery and Intelligence.
Mind 59(236):433–460.

- thinking humanly

Different Answers to 'What is Artificial Intelligence?'

- acting humanly

Turing test

A. Turing (1950): Computing Machinery and Intelligence.
Mind 59(236):433–460.

- thinking humanly

cognitive modelling

Different Answers to 'What is Artificial Intelligence?'

- acting humanly

Turing test

A. Turing (1950): Computing Machinery and Intelligence.
Mind 59(236):433–460.

- thinking humanly
cognitive modelling
- thinking rationally

Different Answers to 'What is Artificial Intelligence?'

- acting humanly

Turing test

A. Turing (1950): Computing Machinery and Intelligence.
Mind 59(236):433–460.

- thinking humanly
cognitive modelling
- thinking rationally
logic

Different Answers to 'What is Artificial Intelligence?'

- acting humanly

Turing test

A. Turing (1950): Computing Machinery and Intelligence.
Mind 59(236):433–460.

- thinking humanly
cognitive modelling
- thinking rationally
logic
- acting rationally

Different Answers to 'What is Artificial Intelligence?'

- acting humanly

Turing test

A. Turing (1950): Computing Machinery and Intelligence.
Mind 59(236):433–460.

- thinking humanly

cognitive modelling

- thinking rationally

logic

- acting rationally

rational agents approach

Related and Relevant Fields

Related and Relevant Fields

- philosophy
logic, mind-body problem, knowledge and action, ...

Related and Relevant Fields

- **philosophy**
logic, mind-body problem, knowledge and action, ...
- **mathematics**
logic, computability, probability, ...

Related and Relevant Fields

- **philosophy**
logic, mind-body problem, knowledge and action, ...
- **mathematics**
logic, computability, probability, ...
- **economics**
decision making, planning, ...

Related and Relevant Fields

- **philosophy**
logic, mind-body problem, knowledge and action, ...
- **mathematics**
logic, computability, probability, ...
- **economics**
decision making, planning, ...
- **neuroscience**
natural information processing

Related and Relevant Fields

- **philosophy**
logic, mind-body problem, knowledge and action, ...
- **mathematics**
logic, computability, probability, ...
- **economics**
decision making, planning, ...
- **neuroscience**
natural information processing
- **psychology**
theory of thinking and acting

Related and Relevant Fields

- **philosophy**
logic, mind-body problem, knowledge and action, ...
- **mathematics**
logic, computability, probability, ...
- **economics**
decision making, planning, ...
- **neuroscience**
natural information processing
- **psychology**
theory of thinking and acting
- **computer engineering**
efficient hardware platforms

Related and Relevant Fields

- **philosophy**
logic, mind-body problem, knowledge and action, ...
- **mathematics**
logic, computability, probability, ...
- **economics**
decision making, planning, ...
- **neuroscience**
natural information processing
- **psychology**
theory of thinking and acting
- **computer engineering**
efficient hardware platforms
- **control theory and cybernetics**
theory of self-controlled systems

Related and Relevant Fields

- **philosophy**
logic, mind-body problem, knowledge and action, ...
- **mathematics**
logic, computability, probability, ...
- **economics**
decision making, planning, ...
- **neuroscience**
natural information processing
- **psychology**
theory of thinking and acting
- **computer engineering**
efficient hardware platforms
- **control theory and cybernetics**
theory of self-controlled systems
- **linguistics**
language and thought, formal grammars, ...

History of Artificial Intelligence

History of Artificial Intelligence

- 1943–1955: prior to AI

History of Artificial Intelligence

- 1943–1955: prior to AI
 - W. McCulloch, W. Pitts (1943): artificial neural nets

History of Artificial Intelligence

- 1943–1955: prior to AI
 - W. McCulloch, W. Pitts (1943): artificial neural nets
 - D. Hebb (1949): Hebbian learning

History of Artificial Intelligence

- 1943–1955: prior to AI
 - W. McCulloch, W. Pitts (1943): artificial neural nets
 - D. Hebb (1949): Hebbian learning
 - A. Turing (1950): Computing Machinery and Intelligence

History of Artificial Intelligence

- 1943–1955: prior to AI
 - W. McCulloch, W. Pitts (1943): artificial neural nets
 - D. Hebb (1949): Hebbian learning
 - A. Turing (1950): Computing Machinery and Intelligence
- 1956: birth of AI

History of Artificial Intelligence

- 1943–1955: prior to AI
 - W. McCulloch, W. Pitts (1943): artificial neural nets
 - D. Hebb (1949): Hebbian learning
 - A. Turing (1950): Computing Machinery and Intelligence
- 1956: birth of AI
 - Dartmouth workshop (Newell, Simon, McCarthy, Minsky, . . .)

History of Artificial Intelligence

- 1943–1955: prior to AI
 - W. McCulloch, W. Pitts (1943): artificial neural nets
 - D. Hebb (1949): Hebbian learning
 - A. Turing (1950): Computing Machinery and Intelligence
- 1956: birth of AI
 - Dartmouth workshop (Newell, Simon, McCarthy, Minsky, . . .)
- 1952–1969: early years

History of Artificial Intelligence

- 1943–1955: prior to AI
 - W. McCulloch, W. Pitts (1943): artificial neural nets
 - D. Hebb (1949): Hebbian learning
 - A. Turing (1950): Computing Machinery and Intelligence
- 1956: birth of AI
Dartmouth workshop (Newell, Simon, McCarthy, Minsky, . . .)
- 1952–1969: early years
 - A. Newell, H. Simon: General Problem Solver

History of Artificial Intelligence

- 1943–1955: prior to AI
 - W. McCulloch, W. Pitts (1943): artificial neural nets
 - D. Hebb (1949): Hebbian learning
 - A. Turing (1950): Computing Machinery and Intelligence
- 1956: birth of AI
 - Dartmouth workshop (Newell, Simon, McCarthy, Minsky, . . .)
- 1952–1969: early years
 - A. Newell, H. Simon: General Problem Solver
 - H. Gelernter: Geometry Theorem Prover

History of Artificial Intelligence

- 1943–1955: prior to AI
 - W. McCulloch, W. Pitts (1943): artificial neural nets
 - D. Hebb (1949): Hebbian learning
 - A. Turing (1950): Computing Machinery and Intelligence
- 1956: birth of AI
 - Dartmouth workshop (Newell, Simon, McCarthy, Minsky, . . .)
- 1952–1969: early years
 - A. Newell, H. Simon: General Problem Solver
 - H. Gelernter: Geometry Theorem Prover
 - A. Samuel: checkers

History of Artificial Intelligence

- 1943–1955: prior to AI
 - W. McCulloch, W. Pitts (1943): artificial neural nets
 - D. Hebb (1949): Hebbian learning
 - A. Turing (1950): Computing Machinery and Intelligence
- 1956: birth of AI
 - Dartmouth workshop (Newell, Simon, McCarthy, Minsky, . . .)
- 1952–1969: early years
 - A. Newell, H. Simon: General Problem Solver
 - H. Gelernter: Geometry Theorem Prover
 - A. Samuel: checkers
 - J. McCarthy: Lisp

History of Artificial Intelligence

- 1943–1955: prior to AI
 - W. McCulloch, W. Pitts (1943): artificial neural nets
 - D. Hebb (1949): Hebbian learning
 - A. Turing (1950): Computing Machinery and Intelligence
- 1956: birth of AI
 - Dartmouth workshop (Newell, Simon, McCarthy, Minsky, . . .)
- 1952–1969: early years
 - A. Newell, H. Simon: General Problem Solver
 - H. Gelernter: Geometry Theorem Prover
 - A. Samuel: checkers
 - J. McCarthy: Lisp
 - M. Minsky: AI Lab (several microworlds)

History of Artificial Intelligence

- 1943–1955: prior to AI
 - W. McCulloch, W. Pitts (1943): artificial neural nets
 - D. Hebb (1949): Hebbian learning
 - A. Turing (1950): Computing Machinery and Intelligence
- 1956: birth of AI
 - Dartmouth workshop (Newell, Simon, McCarthy, Minsky, . . .)
- 1952–1969: early years
 - A. Newell, H. Simon: General Problem Solver
 - H. Gelernter: Geometry Theorem Prover
 - A. Samuel: checkers
 - J. McCarthy: Lisp
 - M. Minsky: AI Lab (several microworlds)
 - F. Rosenblatt: perceptron

History of Artificial Intelligence (continued)

History of Artificial Intelligence (continued)

- 1966–1973: first crisis

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) not kept

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) **not kept**
 - **no general success** in machine translation

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) **not kept**
 - **no general success** in machine translation
“The spirit is willing but the flesh is weak.”

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) **not kept**
 - **no general success** in machine translation
“The spirit is willing but the flesh is weak.” (into Russian and back)

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) **not kept**
 - **no general success** in machine translation
“The spirit is willing but the flesh is weak.” (into Russian and back)
↪ “The vodka is good but the meat is rotten.”

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) **not kept**
 - **no general success** in machine translation
“The spirit is willing but the flesh is weak.” (into Russian and back)
↪ “The vodka is good but the meat is rotten.”
 - Lighthill report (1973): **failure** to come to grips with the combinatorial explosion

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) **not kept**
 - **no general success** in machine translation
“The spirit is willing but the flesh is weak.” (into Russian and back)
↪ “The vodka is good but the meat is rotten.”
 - Lighthill report (1973): **failure** to come to grips with the combinatorial explosion
 - M. Minsky, S. Papert (1969): Perceptrons. (formalizing what perceptrons **cannot** do)

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) **not kept**
 - **no general success** in machine translation
“The spirit is willing but the flesh is weak.” (into Russian and back)
↪ “The vodka is good but the meat is rotten.”
 - Lighthill report (1973): **failure** to come to grips with the combinatorial explosion
 - M. Minsky, S. Papert (1969): Perceptrons. (formalizing what perceptrons **cannot** do)
- 1969–1979: knowledge-based systems

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) **not kept**
 - **no general success** in machine translation
“The spirit is willing but the flesh is weak.” (into Russian and back)
↪ “The vodka is good but the meat is rotten.”
 - Lighthill report (1973): **failure** to come to grips with the combinatorial explosion
 - M. Minsky, S. Papert (1969): Perceptrons. (formalizing what perceptrons **cannot** do)
- 1969–1979: knowledge-based systems expert systems

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) **not kept**
 - **no general success** in machine translation
“The spirit is willing but the flesh is weak.” (into Russian and back)
↪ “The vodka is good but the meat is rotten.”
 - Lighthill report (1973): **failure** to come to grips with the combinatorial explosion
 - M. Minsky, S. Papert (1969): Perceptrons. (formalizing what perceptrons **cannot** do)
- 1969–1979: knowledge-based systems expert systems
- 1980 and ongoing: AI as an industry

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) **not kept**
 - **no general success** in machine translation
“The spirit is willing but the flesh is weak.” (into Russian and back)
↪ “The vodka is good but the meat is rotten.”
 - Lighthill report (1973): **failure** to come to grips with the combinatorial explosion
 - M. Minsky, S. Papert (1969): Perceptrons. (formalizing what perceptrons **cannot** do)
- 1969–1979: knowledge-based systems expert systems
- 1980 and ongoing: AI as an industry
- 1986 and ongoing: return of artificial neural networks

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) **not kept**
 - **no general success** in machine translation
“The spirit is willing but the flesh is weak.” (into Russian and back)
↪ “The vodka is good but the meat is rotten.”
 - Lighthill report (1973): **failure** to come to grips with the combinatorial explosion
 - M. Minsky, S. Papert (1969): Perceptrons. (formalizing what perceptrons **cannot** do)
- 1969–1979: knowledge-based systems expert systems
- 1980 and ongoing: AI as an industry
- 1986 and ongoing: return of artificial neural networks
- 1995 and ongoing: intelligent agents

History of Artificial Intelligence (continued)

- 1966–1973: first crisis
 - early promises (H. Simon (1957): *In a visible future, the range of problems they (intelligent machines) can handle will be coextensive with the range to which the human mind has been applied.*) **not kept**
 - **no general success** in machine translation
“The spirit is willing but the flesh is weak.” (into Russian and back)
↪ “The vodka is good but the meat is rotten.”
 - Lighthill report (1973): **failure** to come to grips with the combinatorial explosion
 - M. Minsky, S. Papert (1969): Perceptrons. (formalizing what perceptrons **cannot** do)
- 1969–1979: knowledge-based systems expert systems
- 1980 and ongoing: AI as an industry
- 1986 and ongoing: return of artificial neural networks
- 1995 and ongoing: intelligent agents
- 2001 and ongoing: availability of very large data sets

Current Applications

Current Applications

- robotics

Current Applications

- robotics
- robotic vehicles

Current Applications

- robotics
- robotic vehicles
- logistics planning

Current Applications

- robotics
- robotic vehicles
- logistics planning
- autonomous planning and scheduling

Current Applications

- robotics
- robotic vehicles
- logistics planning
- autonomous planning and scheduling
- speech recognition

Current Applications

- robotics
- robotic vehicles
- logistics planning
- autonomous planning and scheduling
- speech recognition
- machine translation

Current Applications

- robotics
- robotic vehicles
- logistics planning
- autonomous planning and scheduling
- speech recognition
- machine translation
- game playing

Current Applications

- robotics
- robotic vehicles
- logistics planning
- autonomous planning and scheduling
- speech recognition
- machine translation
- game playing
- spam fighting

Current Applications

- robotics
- robotic vehicles
- logistics planning
- autonomous planning and scheduling
- speech recognition
- machine translation
- game playing
- spam fighting
- ...

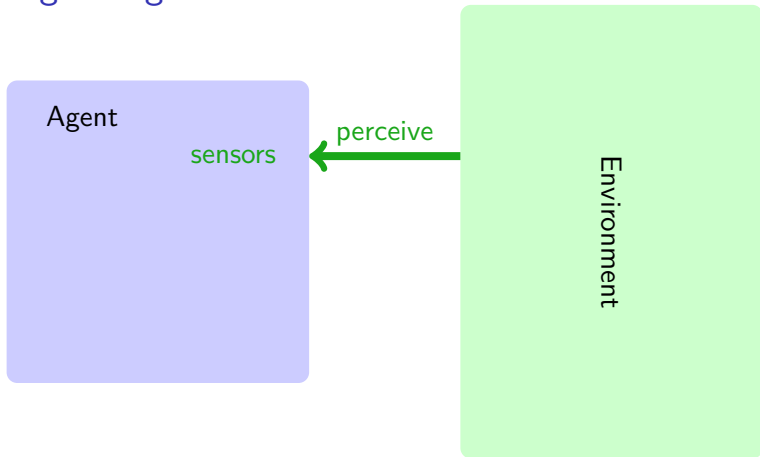
Intelligent Agents and Their Environment



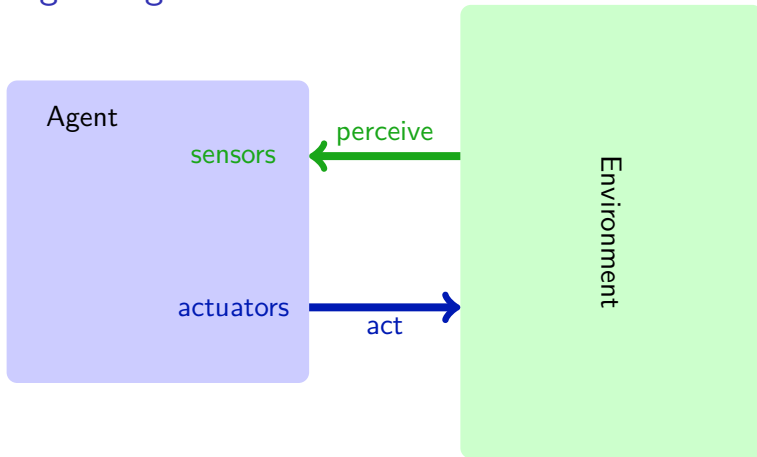
Agent

Environment


Intelligent Agents and Their Environment




Intelligent Agents and Their Environment



Intelligent Agents and Their Environment




Intelligent Agents and Their Environment



Classify systems according to

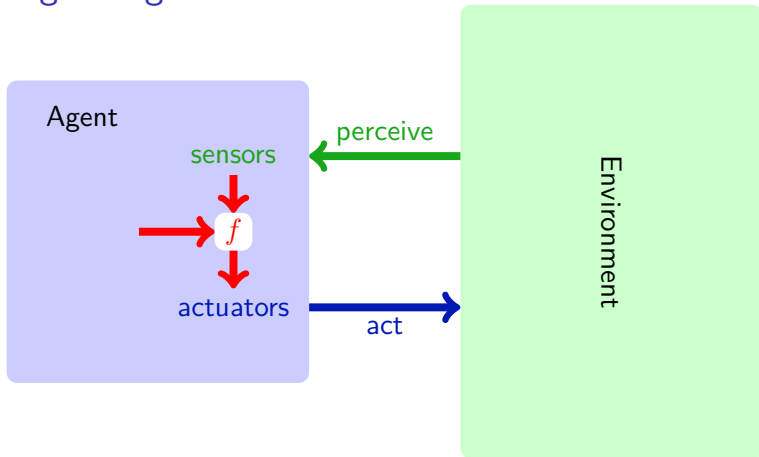
Intelligent Agents and Their Environment



Classify systems according to

- properties of the environment


Intelligent Agents and Their Environment



Classify systems according to

- properties of the environment
- properties of the agent

Intelligent Agents and Their Environment



Classify systems according to

- properties of the environment
- properties of the agent
- properties of the interplay

Rational Agents

Rational Agents

Roughly speaking rational $\hat{=}$ doing the right thing

Rational Agents

Roughly speaking rational $\hat{=}$ doing the right thing

What defines what is the right thing?

Rational Agents

Roughly speaking rational $\hat{=}$ doing the right thing

What defines what is the right thing?

Crucial performance measure

maps sequence of environment states to performance

Rational Agents

Roughly speaking rational $\hat{=}$ doing the right thing

What defines what is the right thing?

Crucial performance measure

maps sequence of environment states to performance

Definition (Rational Agent)

For each possible percept sequence, a rational agent should select an action that is expected to maximise the performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has.

Rational Agents

Roughly speaking rational $\hat{=}$ doing the right thing

What defines what is the right thing?

Crucial performance measure

maps sequence of environment states to performance

Definition (Rational Agent)

For each possible percept sequence, a rational agent should select an action that is expected to maximise the performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has.

Note exploration, information gathering, learning
crucial in unknown (or partially known) environments

Task Environments

Problem statement for rational agent(s) $\hat{=}$ task environment

Task Environments

Problem statement for rational agent(s) $\hat{=}$ task environment

Components

Task Environments

Problem statement for rational agent(s) $\hat{=}$ task environment

Components

- performance measure

Task Environments

Problem statement for rational agent(s) $\hat{=}$ task environment

Components

- performance measure
- environment

Task Environments

Problem statement for rational agent(s) $\hat{=}$ task environment

Components

- performance measure
- environment
- actuators

Task Environments

Problem statement for rational agent(s) $\hat{=}$ task environment

Components

- performance measure
- environment
- actuators
- sensors

Task Environments

Problem statement for rational agent(s) $\hat{=}$ task environment

Components

- performance measure
- environment
- actuators
- sensors

Helpful categorisation of task environments according to properties ('dimensions')

Dimensions of Task Environments

Remember task environment encompasses
performance measure, environment, actuators, sensors

Dimensions of Task Environments

Remember task environment encompasses
 performance measure, environment, actuators, sensors

- fully observable vs. partially observable

Dimensions of Task Environments

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
property of the sensors

Dimensions of Task Environments

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
property of the sensors
- single agent vs. multiagent

Dimensions of Task Environments

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
property of the sensors
- single agent vs. multiagent
sub-dimension competitive vs. cooperative

Dimensions of Task Environments

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
property of the sensors
- single agent vs. multiagent
sub-dimension competitive vs. cooperative
communication

Dimensions of Task Environments

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
property of the sensors
- single agent vs. multiagent
sub-dimension competitive vs. cooperative
communication
- deterministic vs. stochastic

Dimensions of Task Environments

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
property of the sensors
- single agent vs. multiagent
sub-dimension competitive vs. cooperative
communication
- deterministic vs. stochastic
property of the environment

Dimensions of Task Environments

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
property of the sensors
- single agent vs. multiagent
sub-dimension competitive vs. cooperative
communication
- deterministic vs. stochastic
property of the environment
note other agents not taken into account

Dimensions of Task Environments

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
property of the sensors
- single agent vs. multiagent
sub-dimension competitive vs. cooperative
communication
- deterministic vs. stochastic
property of the environment
note other agents not taken into account
- episodic vs. sequential

Dimensions of Task Environments

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
property of the sensors
- single agent vs. multiagent
sub-dimension competitive vs. cooperative
communication
- deterministic vs. stochastic
property of the environment
note other agents not taken into account
- episodic vs. sequential
- static vs. dynamic

Dimensions of Task Environments

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
property of the sensors
- single agent vs. multiagent
sub-dimension competitive vs. cooperative
communication
- deterministic vs. stochastic
property of the environment
note other agents not taken into account
- episodic vs. sequential
- static vs. dynamic
special case semidynamic: only performance measure changes
over time

Dimensions of Task Environments (continued)

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
- single agent vs. multiagent
- deterministic vs. stochastic
- episodic vs. sequential
- static vs. dynamic

Dimensions of Task Environments (continued)

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
- single agent vs. multiagent
- deterministic vs. stochastic
- episodic vs. sequential
- static vs. dynamic
- discrete vs. continuous

Dimensions of Task Environments (continued)

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
- single agent vs. multiagent
- deterministic vs. stochastic
- episodic vs. sequential
- static vs. dynamic
- discrete vs. continuous
 - applies to
 - environment, in particular time
 - percepts
 - actions

Dimensions of Task Environments (continued)

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
- single agent vs. multiagent
- deterministic vs. stochastic
- episodic vs. sequential
- static vs. dynamic
- discrete vs. continuous
 - applies to
 - environment, in particular time
 - percepts
 - actions
- known vs. unknown

Dimensions of Task Environments (continued)

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
- single agent vs. multiagent
- deterministic vs. stochastic
- episodic vs. sequential
- static vs. dynamic
- discrete vs. continuous
 - applies to
 - environment, in particular time
 - percepts
 - actions
- known vs. unknown
 - note different from observable

Looking at 'Dimensions of Task Environments'

Remember task environment encompasses
performance measure, environment, actuators, sensors

- fully observable vs. partially observable
- single agent vs. multiagent
- deterministic vs. stochastic
- episodic vs. sequential
- static vs. dynamic
- discrete vs. continuous
- known vs. unknown

Summary & Take Home Message

Things to remember

Summary & Take Home Message

Things to remember

- AI in this course encompasses 'classical' AI algorithms and randomised search heuristics, concentrates on search

Summary & Take Home Message

Things to remember

- AI in this course encompasses 'classical' AI algorithms and randomised search heuristics, concentrates on search
- AI between thinking ↔ acting, rationally ↔ humanly

Summary & Take Home Message

Things to remember

- AI **in this course** encompasses 'classical' AI algorithms and randomised search heuristics, concentrates on search
- AI between thinking ↔ acting, rationally ↔ humanly
- many related and relevant fields

Summary & Take Home Message

Things to remember

- AI **in this course** encompasses 'classical' AI algorithms and randomised search heuristics, concentrates on search
- AI between thinking↔acting, rationally↔humanly
- many related and relevant fields
- history with early enthusiasm, crisis, and continuing efforts and successful applications

Summary & Take Home Message

Things to remember

- AI **in this course** encompasses 'classical' AI algorithms and randomised search heuristics, concentrates on search
- AI between thinking↔acting, rationally↔humanly
- many related and relevant fields
- history with early enthusiasm, crisis, and continuing efforts and successful applications
- intelligent agents and environments: task environments

Summary & Take Home Message

Things to remember

- AI **in this course** encompasses 'classical' AI algorithms and randomised search heuristics, concentrates on search
- AI between thinking↔acting, rationally↔humanly
- many related and relevant fields
- history with early enthusiasm, crisis, and continuing efforts and successful applications
- intelligent agents and environments: task environments

Take Home Message

Summary & Take Home Message

Things to remember

- AI **in this course** encompasses 'classical' AI algorithms and randomised search heuristics, concentrates on search
- AI between thinking↔acting, rationally↔humanly
- many related and relevant fields
- history with early enthusiasm, crisis, and continuing efforts and successful applications
- intelligent agents and environments: task environments

Take Home Message

- AI is a broad, diverse, rich, interesting, and important field.

Summary & Take Home Message

Things to remember

- AI **in this course** encompasses 'classical' AI algorithms and randomised search heuristics, concentrates on search
- AI between thinking↔acting, rationally↔humanly
- many related and relevant fields
- history with early enthusiasm, crisis, and continuing efforts and successful applications
- intelligent agents and environments: task environments

Take Home Message

- AI is a broad, diverse, rich, interesting, and important field.
- Intelligent agents and their environments are a flexible and general framework for AI (applications).