Advanced Knowledge Based Systems
CS3411

Introductory Lecture:
Agents that Reason Logically

Enrico Franconi

Administrativia

- Class home page:
  `http://www.cs.man.ac.uk/~franconi/teaching/1999/3411/`
- Book: Russell and Norvig *Artificial Intelligence: A Modern Approach*
  Read Chapters 1 and 2 for this lecture's material

9.00	MT3111 M110	MT3241 M213	MT3031 M212	MT3031 M212	MT3031 M212
10.00	CS3451 1.5 MT3061 M214 waXMT3121M213	CS3421 LF14 BM3031 waXMT3111M110 wbXMT3241M213	CS3251 1.5 MT3241 M213	CS3321 1.5 MT3061 M214	CS3321 1.5 MT3211 M213
11.00	CS3251 1.5 MT3121 M218	CS3151 1.5 BM3031 MT3111 M110	CS3001 1.5 CS3121 1.3 MT3341 M213	CS3311 1.5 MT3121 M218	CS3411 LF14 EM2491 RosA waXMT3171M213 wbXMT3211M213
12.00	CS3071 1.1 wbXMT3061M214	MT3361 M213	CS3451 1.5 MT3171 M217	MT3211 M213	CS3411 LF14 EM2491 RosA MT3361 M213
1.00		CS3071 1.1	CS39x0 1.1 CS39x0 1.5	CS3131 1.3 CS3341 LF15	
2.00	CS3001 1.5 CS3121 1.3 wbXMT3061M214 MT3341 M213	CS3411 LF14 XEM2491 BM3010 B3 MT3101 M213 MT4321 MOH		CS3101 1.1 MT3101 M213 MT4321 MOH	CS3311 1.5 MT3171 M216
3.00	CS3131 1.3 CS3341 LF15 MT3071 M110	XEM2491 SBM3010 FSAB SBM3010 M107 SBM3010 S450 waXMT3101M213 XMT4321 MOH		CS3421 LF14 SBM3020	CS3041 1.5 BM3020 FeSm
4.00	CS3101 1.1 MT3071 M110	CS3041 1.5 SBM3010 FSAB SBM3010 M107 SBM3010 S450 waXMT3341M213 wbXMT3361MG08		CS3151 1.5 SBM3020	BM3020 FeSm

Some definitions of AI

"The exciting new effort to make computers think . . . *machines with minds*, in the full and literal sense" (Haugeland, 1985) "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . ." (Bellman, 1978)	"The study of mental faculties through the use of computational models" (Charniak and McDermott, 1985) "The study of the computations that make it possible to perceive, reason, and act" (Winston, 1992)
"The art of creating machines that perform functions that require intelligence when performed by people" (Kurzweil, 1990) "The study of how to make computers do things at which, at the moment, people are better" (Rich and Knight, 1991)	"A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes" (Schalkoff, 1990) "The branch of computer science that is concerned with the automation of intelligent behavior" (Luger and Stubblefield, 1993)

- Systems that think like humans
- Systems that act like humans
- Systems that think rationally
- Systems that act rationally

Systems ⟺ Agents

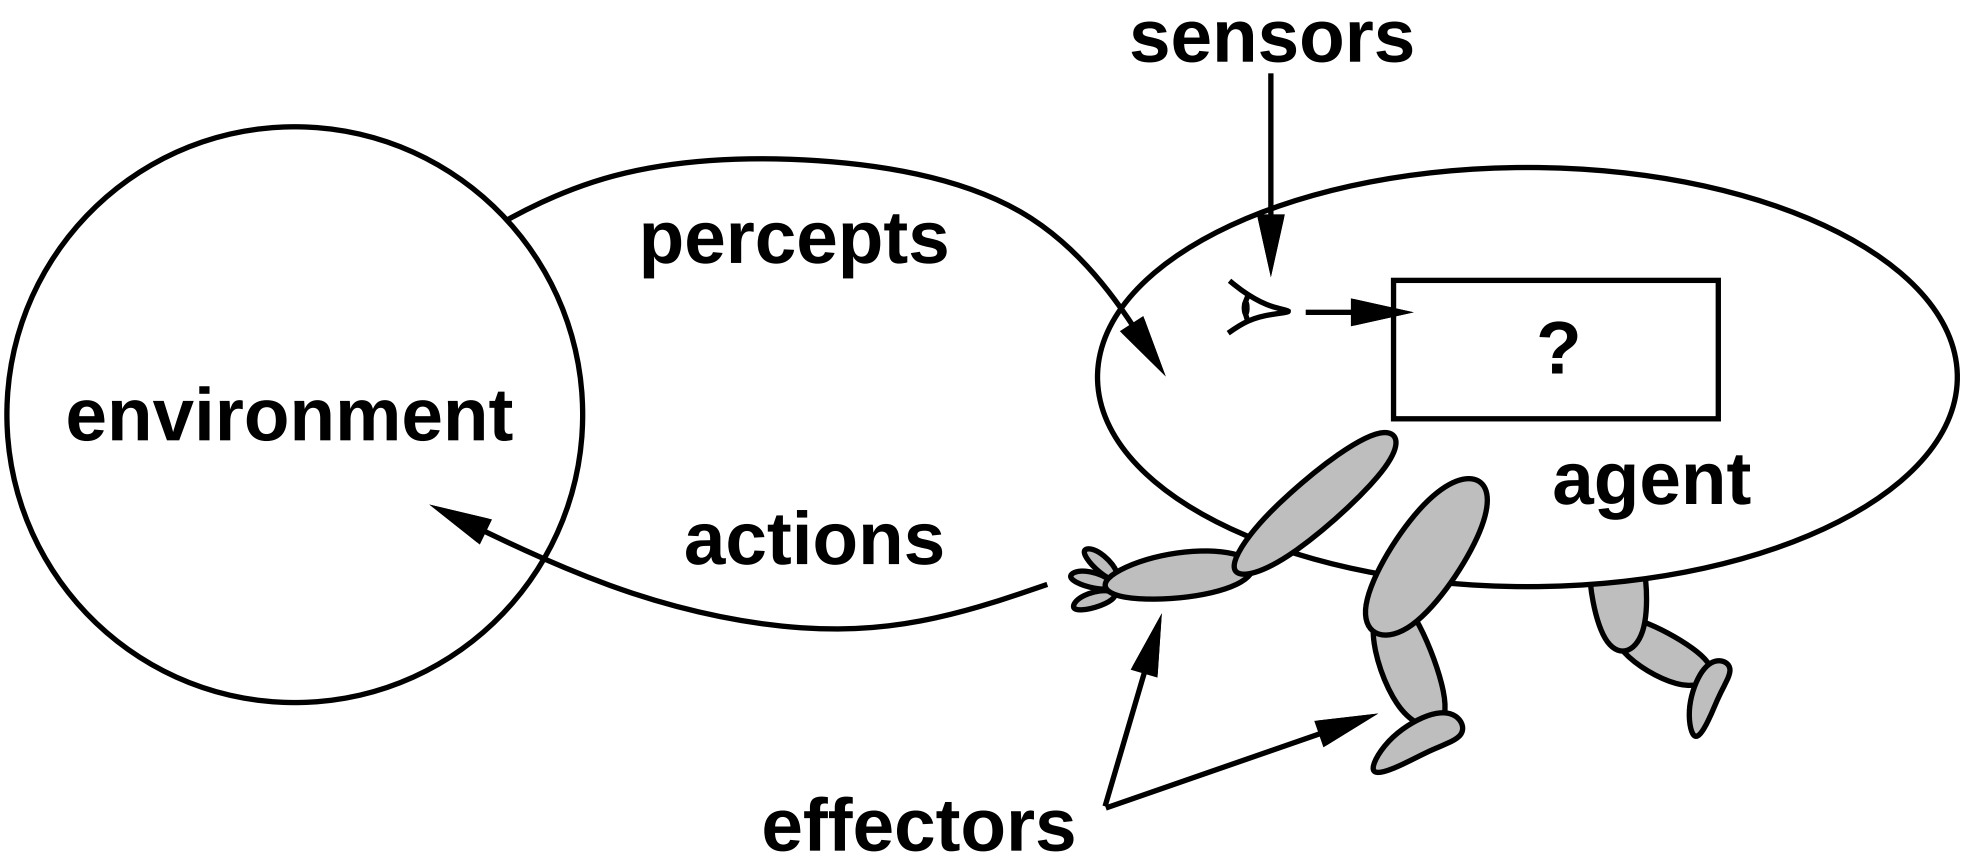

An agent

Consider, e.g., the task of designing an automated taxi:

Percepts: video, accelerometers, gauges, engine sensors, keyboard, GPS, . . .

Actions: steer, accelerate, brake, horn, speak/display, . . .

Goals: safety, reach destination, maximize profits, obey laws, passenger comfort, . . .

Environment: US urban streets, freeways, traffic, pedestrians, weather, customers, . . .

Internet shopping agent

Percepts:

Actions:

Goals:

Environment:

Rational Agents

An Agent as Reasoning module of a Rational Agent.

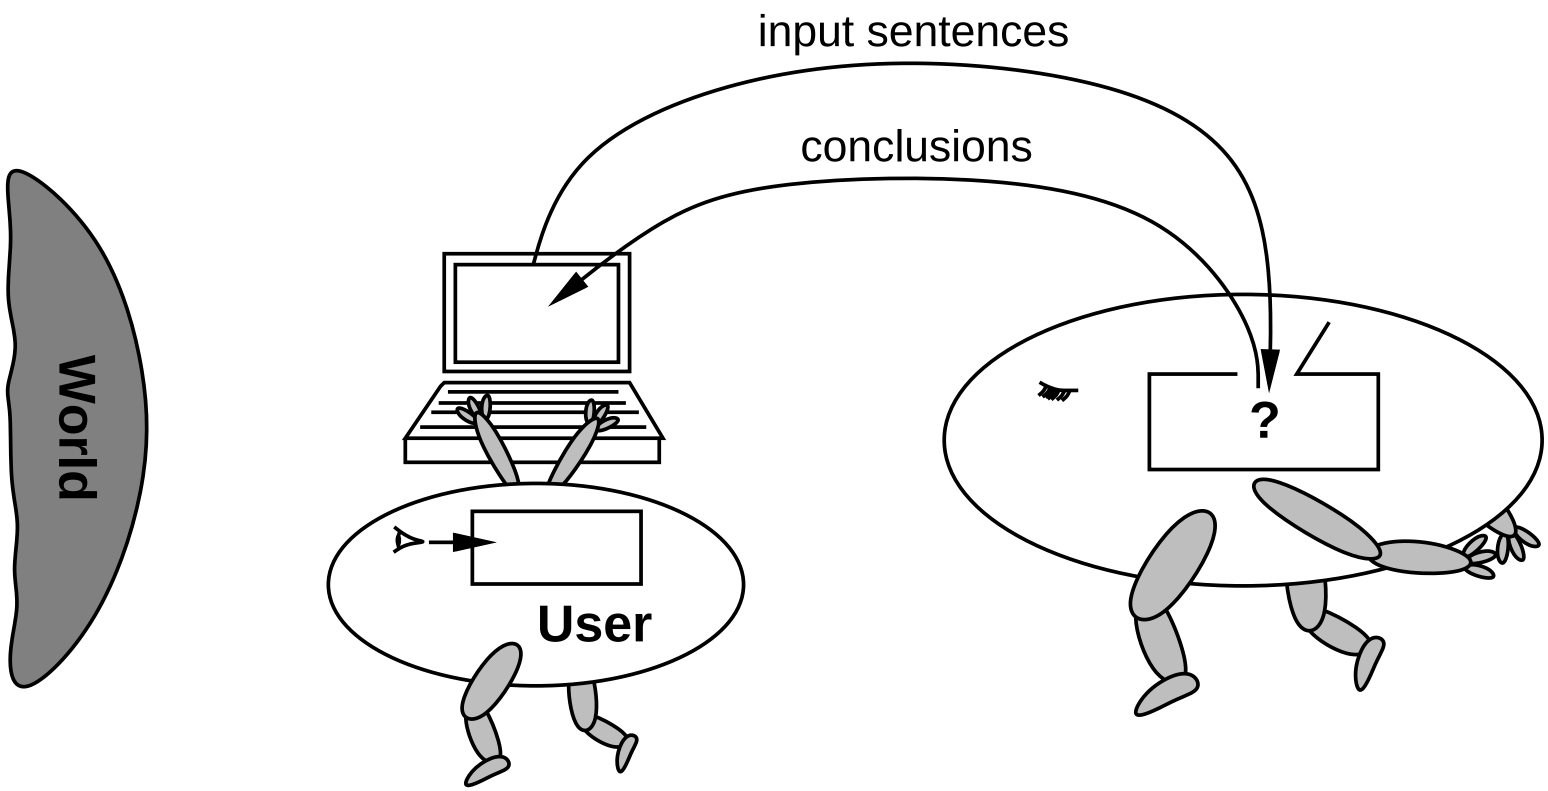

- An Agent as *Reasoning* module: a **logic**.
- *Logic*: a well formalized part of agent *knowledge* and *reasoning*.

Rational agents

- An *agent* is an entity that perceives and acts
- Abstractly, an agent is a function from percept histories and internal knowledge to actions:

$$f : \mathcal{P}^* \times \mathcal{K} \rightarrow \mathcal{A}$$

  For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveat: computational limitations make perfect rationality unachievable
  $\rightarrow$ design best *program* for given machine resources

Simple reflex agents

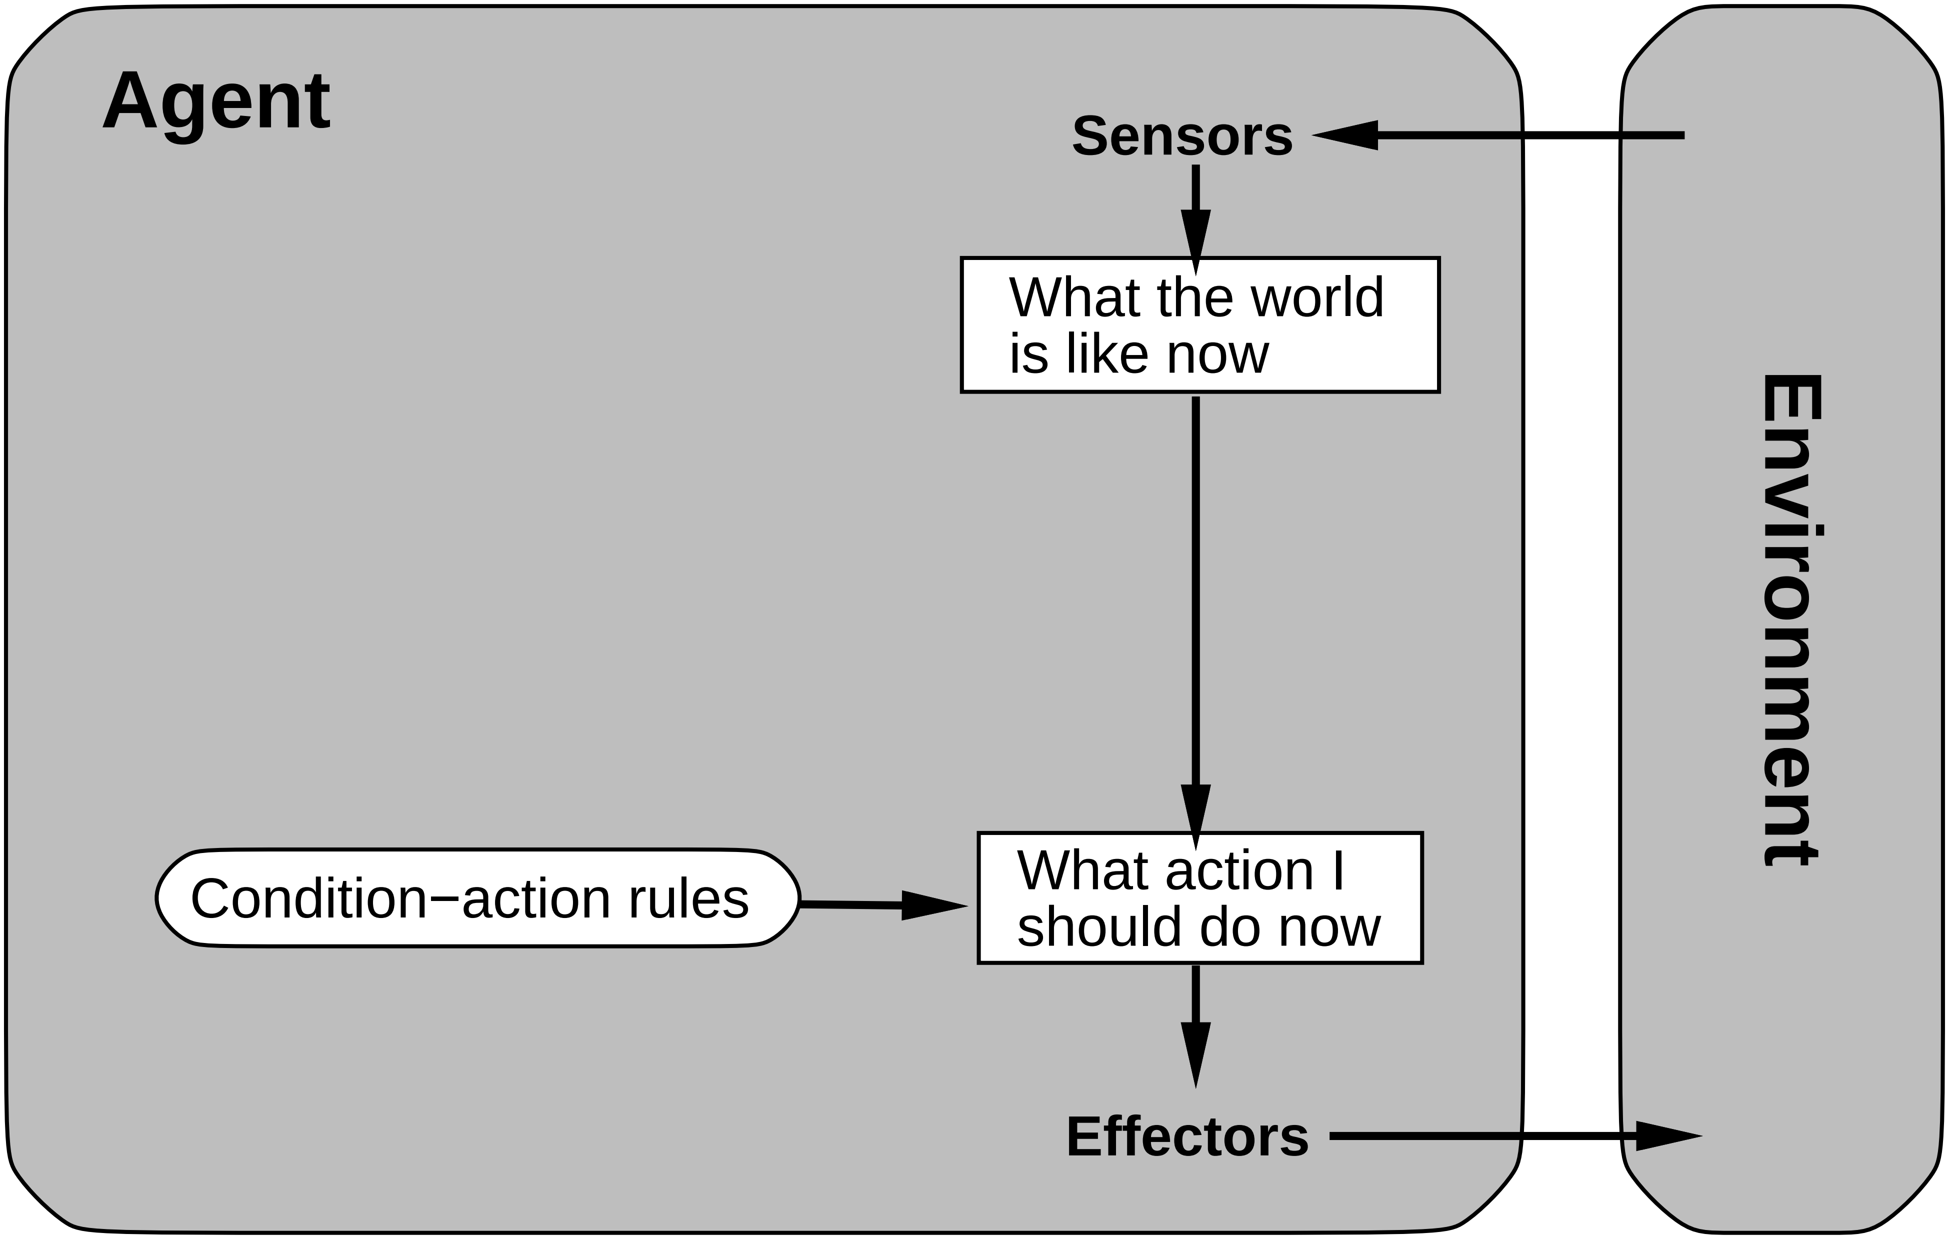

Knowledge Representation hypothesis

Any mechanically embodied intelligent process will be comprised of *structural* ingredients that

a) we as external observers naturally take to *represent* a propositional account of the knowledge the overall process exhibits, and

b) independent of such external semantical attribution, play a formal and essential role in *determining* the behavior that manifests that knowledge.

Declarative approach: logic as a *representation* and *inferencing* tool.

Example: A Taxonomy

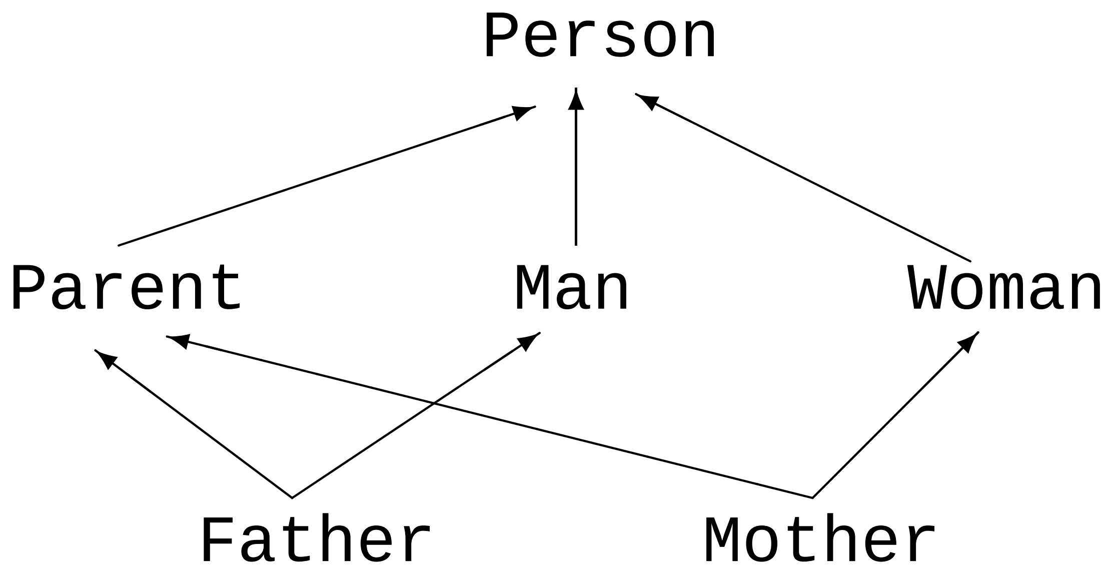

Question: Is X a subconcept of Y?

An Encoding of the Taxonomy

```
procedure subconcept-of(X,Y)
if      Y = 'Person'
then    return true
elseif  Y = 'Parent'
then    if     X = 'Mother' or X = 'Father'
        then   return true
        else   return false
elseif  Y = 'Woman'
then    if     X = 'Mother'
        then   return true
        else   return false
elseif  Y = 'Man'
then    if     X = 'Father'
        then   return true
        else   return false
else    return false
```

$-$ difficult to modify and to understand

$+$ efficient

The Taxonomy as Data Structure + Query

Superconcept	Subconcept
`‘Person’`	`‘Parent’`
`‘Person’`	`‘Woman’`
`‘Person’`	`‘Man’`
`‘Parent’`	`‘Mother’`
`‘Parent’`	`‘Father’`
`‘Woman’`	`‘Mother’`
`‘Man’`	`‘Father’`

& Query that computes `Subconcept-of`, i.e., the *transitive closure* of the Subconcept-Superconcept relationship.

$+$ easy to change

$-$ the semantics depends on the database $\neq$ taxonomy

$-$ intended semantics implicit in the query

Processing Knowledge = "Reasoning"

Representation alone is not useful.

We want to be able to access represented knowledge and to *process* it.

- access alone is, in general, insufficient
- *implicit* knowledge has to be made explicit

$\leadsto$ *deduction methods*

- – the results should only depend on the semantics ...
- – and not on accidental syntactic differences in representations

Logic

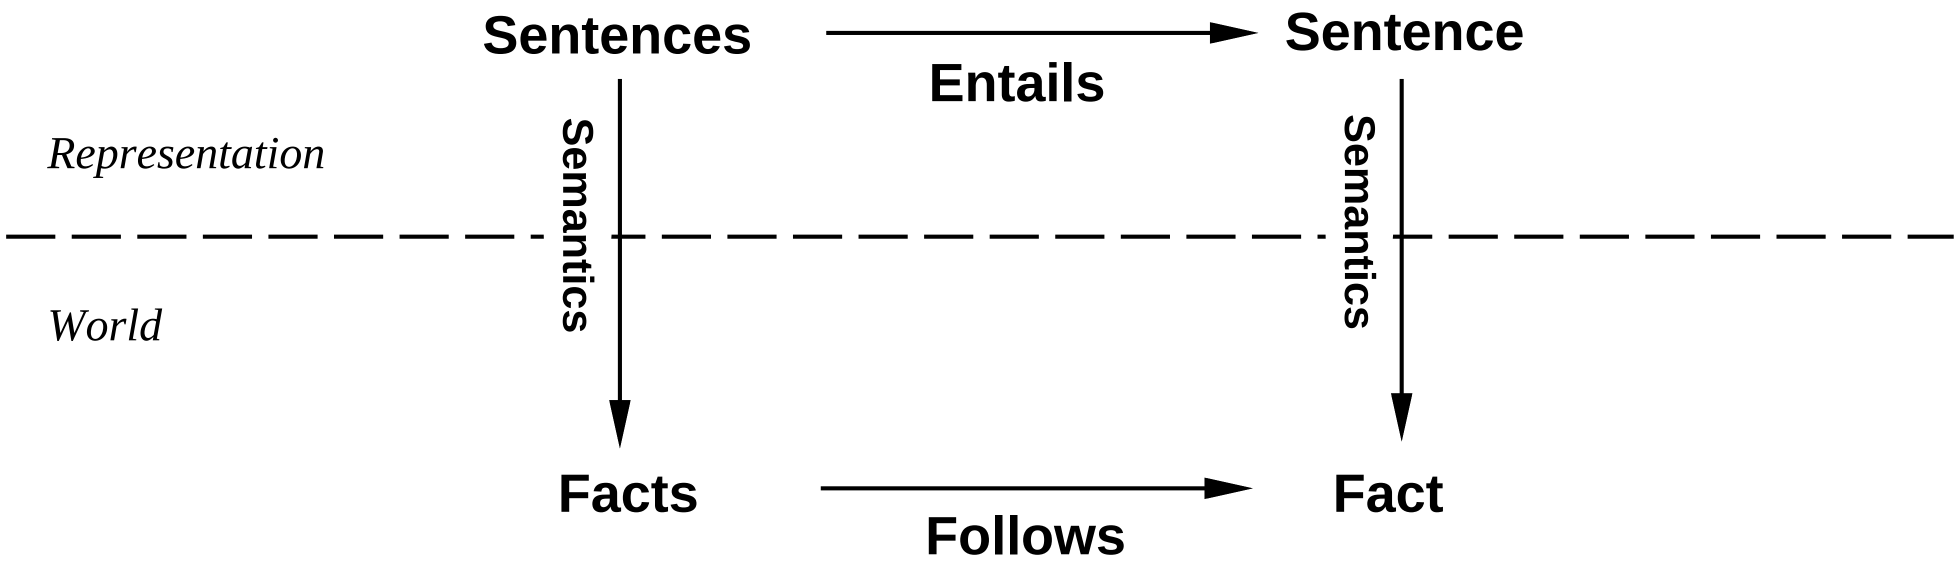

A logic allows the axiomatization of the domain information, and the drawing of conclusions from that information.

- Syntax
- Semantics
- Logical inference = *reasoning*

Important Questions

- **Expressive Power** of representation language

$\leadsto$ able to *represent* the problem

- **Correctness** of entailment procedure

$\leadsto$ *no false* conclusions are drawn

- **Completeness** of entailment procedure

$\leadsto$ *all correct* conclusions are drawn

- **Decidability** of entailment problem

$\leadsto$ there exists a (terminating) algorithm to compute entailment

- **Complexity**

$\leadsto$ resources needed for computing the solution

What is a Logic

Clearly distinguish the definitions of:

- the *formal language*
  - Syntax
  - Semantics
  - Expressive Power
- the *reasoning problem* (e.g., entailment)
  - Decidability
  - Computational Complexity
- the *problem solving procedure*
  - Soundness and Completeness
  - (Asymptotic) Complexity

The ideal Logic

- Expressive
- With decidable reasoning problems
- With sound and complete reasoning procedures
- With efficient reasoning procedures – possibly sub-optimal

Goals of research in the field

- Study how **declarative knowledge** can be *formally defined* using a logic-based approach.
- Give a *computational* account to it, in order to reproduce it in a computing device.

Main topics of the course

- Classical Logic
- Building a Knowledge Base
- Logical Reasoning Systems
- Description Logics

Conclusions

- A warning
  - Rigorous and formal course

Conclusions

- A warning
  - Rigorous and formal course
- Two promises
  - Many examples
  - Only few main important topics