Intelligent Agents

Dr. Marcel Gehrke
Prof. Dr. Ralf Möller
PD Dr. Özgür Özçep

Universität zu Lübeck
Institut für Informationssysteme

Organization

- Module **Intelligent Agents (CS4514-KP12)**:
  - Agents, Mechanism, and Collaboration (lecture)
  - Agent Perception (Language and Vision) (lecture)
  - **Agent Planning, Reinforcement Learning, Relations, and Human-Awareness** (lecture)
  - Project
- Mode of all lectures: classically, on site, with integrated exercises
- Lecture slides partly available from last term, but new material produced on the run
- Today: overview of the general idea
- More details in Moodle (https://moodle.uni-luebeck.de/course/view.php?id=9399)

Organization

- Agents, Mechanisms, and Collaboration
  - Wednesdays, 14:15 - 15:45 Uhr, IFIS Seminarraum 2035,
  - Start: 18.10.2023
  - Lecturer: Özgür Özcep
- Agent Perception (Language and Vision)
  - Thursdays 14:15 - 15:45 Uhr, IFIS Seminarraum 2035
  - Start: 19.10.2023
  - Lecturer: Ralf Möller
- Agent Planning, Reinforcement Learning, Relations, and Human-Awareness
  - Fridays, 10:15 - 11:45 Uhr, IFIS Besprechungsraum 2032
  - Start: 20.10.2023
    Lecturer: Marcel Gehrke
- **Project**
  - Fridays, 12:15 - 13:45 Uhr, IFIS Seminarraum 2035
  - Start: 20.10.2023

A vision some time ago

The entertainment system was belting out the Beatles' "We Can Work It Out" when the phone rang. When Pete answered, his phone turned the sound down by sending a message to all the other *local* devices that had a *volume control*. His sister, Lucy, was on the line from the doctor's office: "Mom needs to see a specialist and then has to have a series of physical therapy sessions. Biweekly or something. I'm going to have my agent set up the appointments." Pete immediately agreed to share the chauffeuring. At the doctor's office, Lucy instructed her Semantic Web agent through her handheld Web browser. The agent promptly retrieved information about Mom's *prescribed treatment* from the doctor's agent, looked up several lists of *providers*, and checked for the ones *in-plan* for Mom's insurance within a *20-mile radius* of her *home* and with a *rating* of *excellent* or *very good* on trusted rating services. It then began trying to find a match between available *appointment times* (supplied by the agents of individual providers through their Web sites) and Pete's and Lucy's busy schedules. (The emphasized keywords indicate terms whose semantics, or meaning, were defined for the agent through the Semantic Web.)

In a few minutes the agent presented them with a plan. Pete didn't like it—University Hospital was all the way across town from Mom's place, and he'd be driving back in the middle of rush hour. He set his own agent to redo the search with stricter preferences about *location* and *time*. Lucy's agent, having *complete trust* in Pete's agent in the context of the present task, automatically assisted by supplying access certificates and shortcuts to the data it had already sorted through.

Almost instantly the new plan was presented: a much closer clinic and earlier times—but there were two warning notes. First, Pete would have to reschedule a couple of his *less important* appointments. He checked what they were—not a problem. The other was something about the insurance company's list failing to include this provider under *physical therapists*: "Service type and insurance plan status securely verified by other means," the agent reassured him. "(Details?)"

Lucy registered her assent at about the same moment Pete was muttering, "Spare me the details," and it was all set. (Of course, Pete couldn't resist the details and later that night had his agent explain how it had found that provider even though it wasn't on the proper list.)

T. Berners-Lee, J. Hendler, and O. Lassila: The semantic web.
Scientific American, 2001. (http://www.scientificamerican.com/)

A vision becomes true???

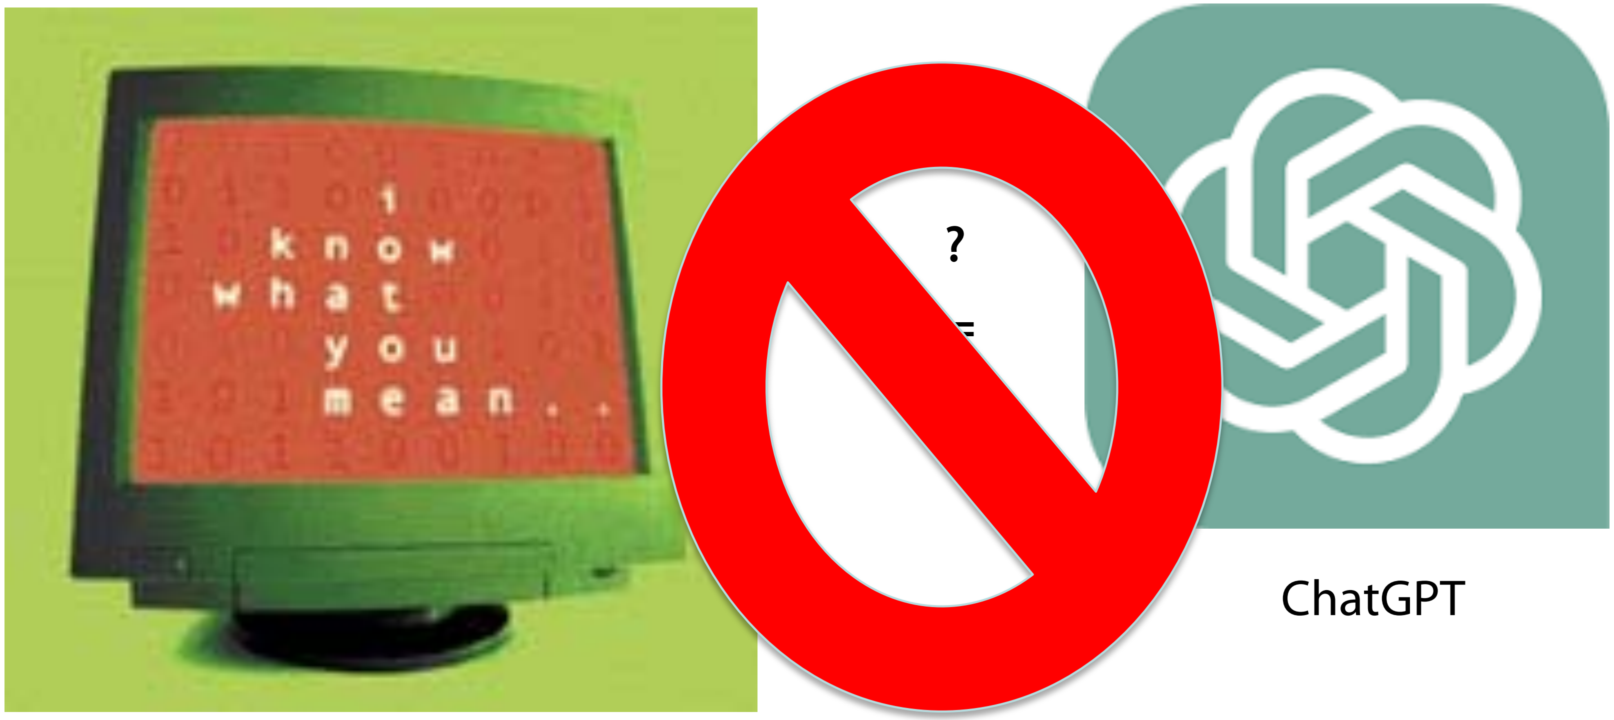

No meaning without perceiving and acting in the world!

Acting was already part of the original vision

“The real power of the Semantic Web will be realized when people create many programs that collect Web content from diverse sources, process the information and exchange the results with other programs. The effectiveness of such software agents will increase exponentially as more machine-readable Web content and automated services (including other agents) become available. The Semantic Web promotes this synergy: even agents that were not expressly designed to work together can transfer data among themselves when the data come with semantics“

T. Berners-Lee, J. Hendler, and O. Lassila: The semantic web. *Scientific American*, 2001. (http://www.scientificamerican.com/)

A vision becomes true!!? AgentGPT

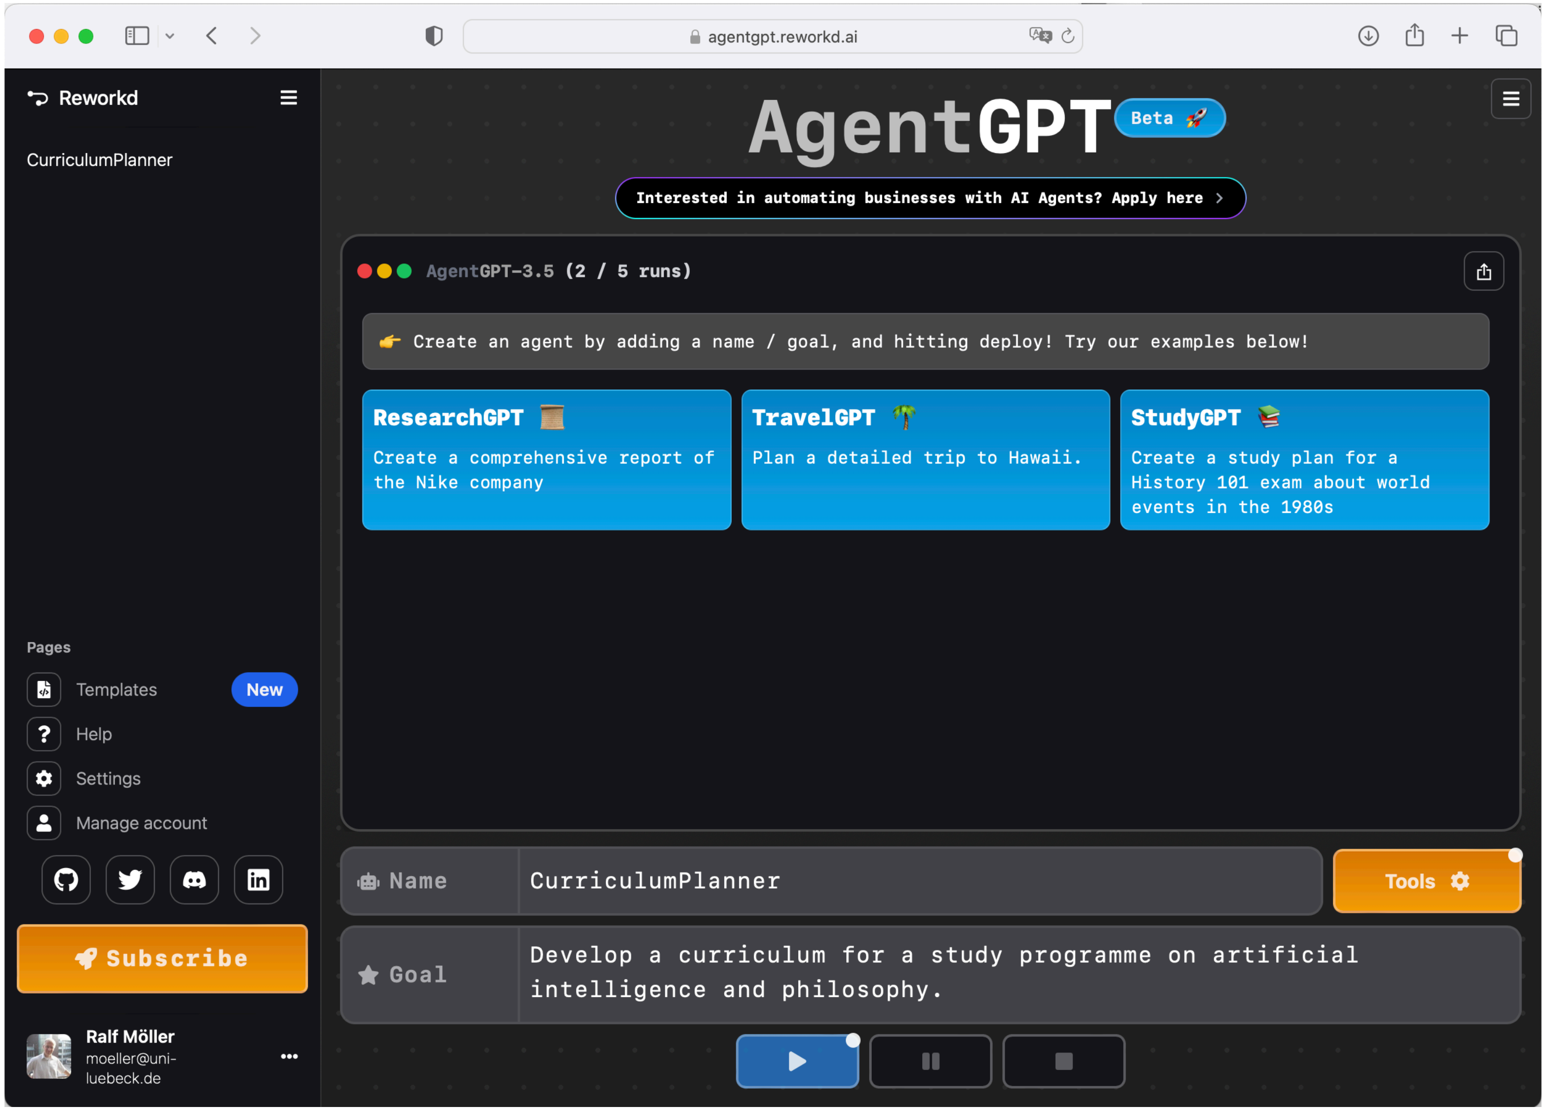

https://agentgpt.reworkd.ai/

Why still a question mark?

- Vision itself not realised in the intended form, only ist goals: The semantic web is a cemetery of ontologies (all died from annotationitis)

- On the other hand: can we „trust the semantics“ handled by AgentGPT?

- Revival of ontologies (?) -> Neurosymbolic AI NeSy ... though not really neural networks play a role

- Even more: we are in an era of Postneural AI

Postneural AI

- It‘s about agents
- It‘s about humans
- It‘s about regulating interaction within the agents/humans
  - towards overall goal of beneficiary of humankind

- The right global architecture is not that of the semantic web, rather it‘s
  - Agents (treated in this course)
  - Formal Ethics (descriptive part treated in this course)
  - and Mechanism Design (treated in this course)

Artificial Intelligence and Intelligent Agents

- Artificial intelligence (AI) is the science of systematic synthesis and analysis of computational agents that act intelligently
  - Agents are central to AI (and vice versa)
  - Intelligent agent = computational agent that acts intelligently
  - Talking about AI w/o talking about agents misses the point (and vice versa)
- Need to technically define the notion of "acting intelligently"
- Systems are called computational agents in AI, or agents for short

Literature

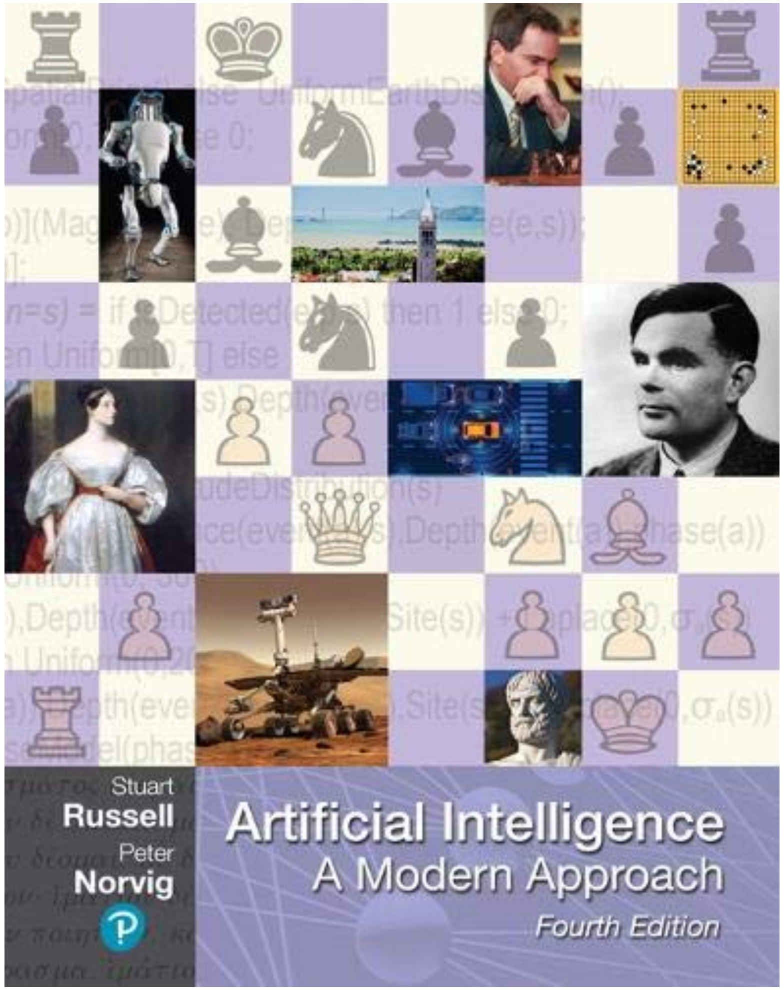

http://aima.cs.berkeley.edu
(AIMA, 1st edition 1995)

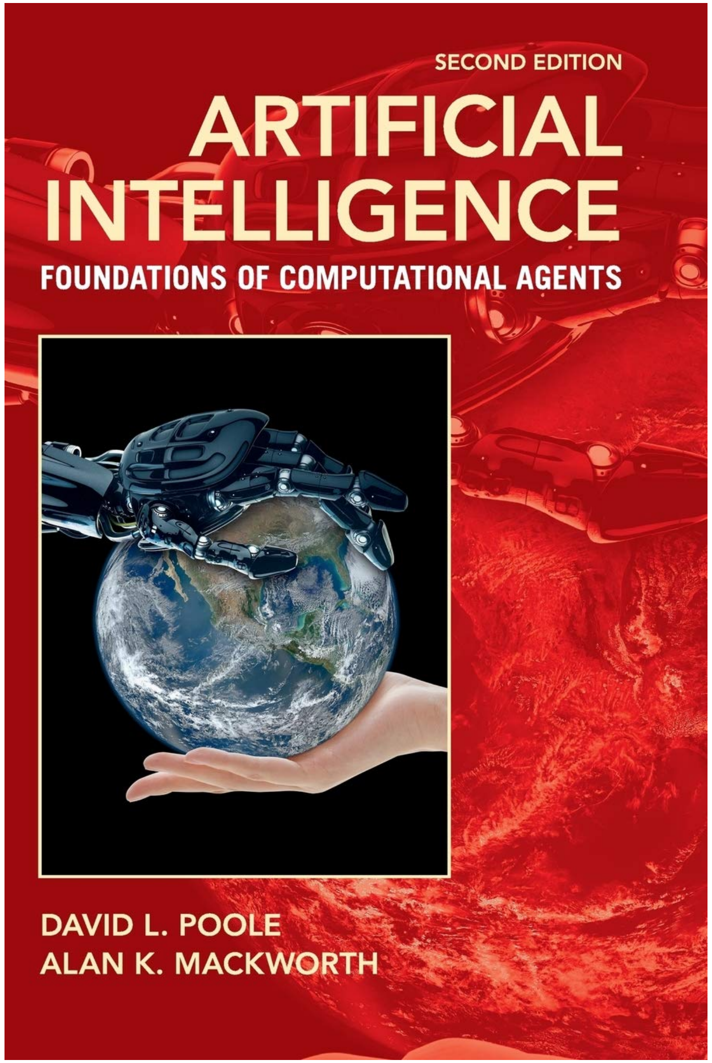

http://artint.info
(AIFCA, 1st edition 2010)

Literature

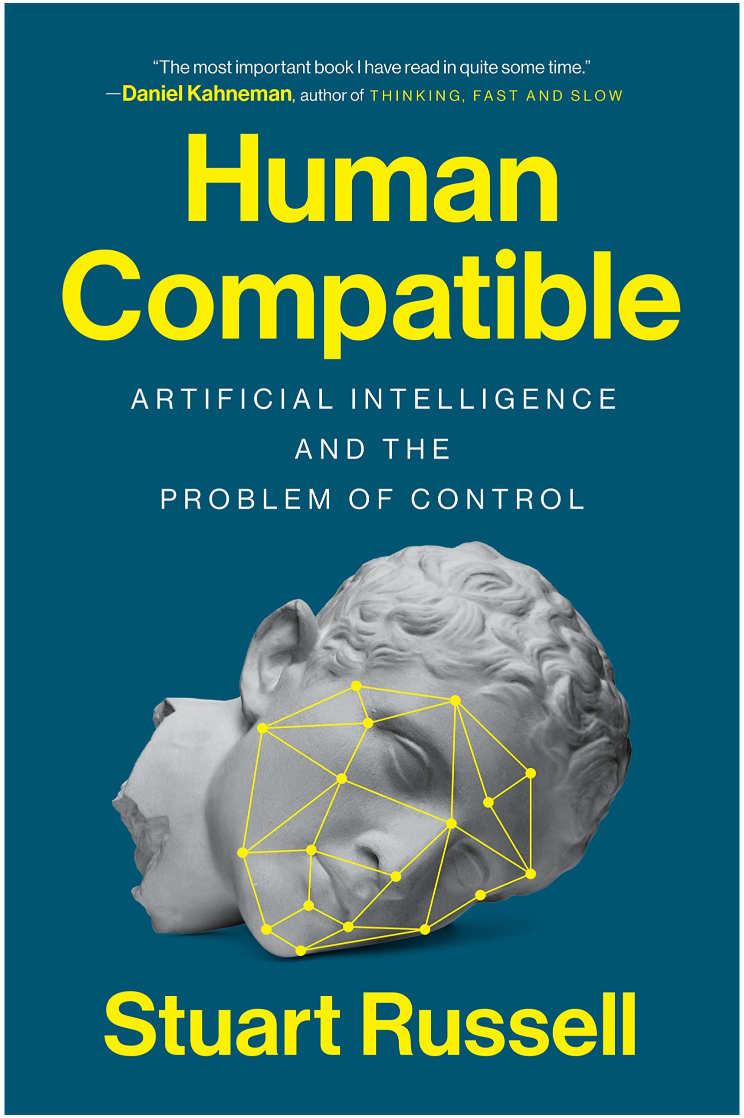

"The most important book I have read in quite some time."
—Daniel Kahneman, author of THINKING, FAST AND SLOW
Human Compatible
ARTIFICIAL INTELLIGENCE AND THE PROBLEM OF CONTROL
Stuart Russell

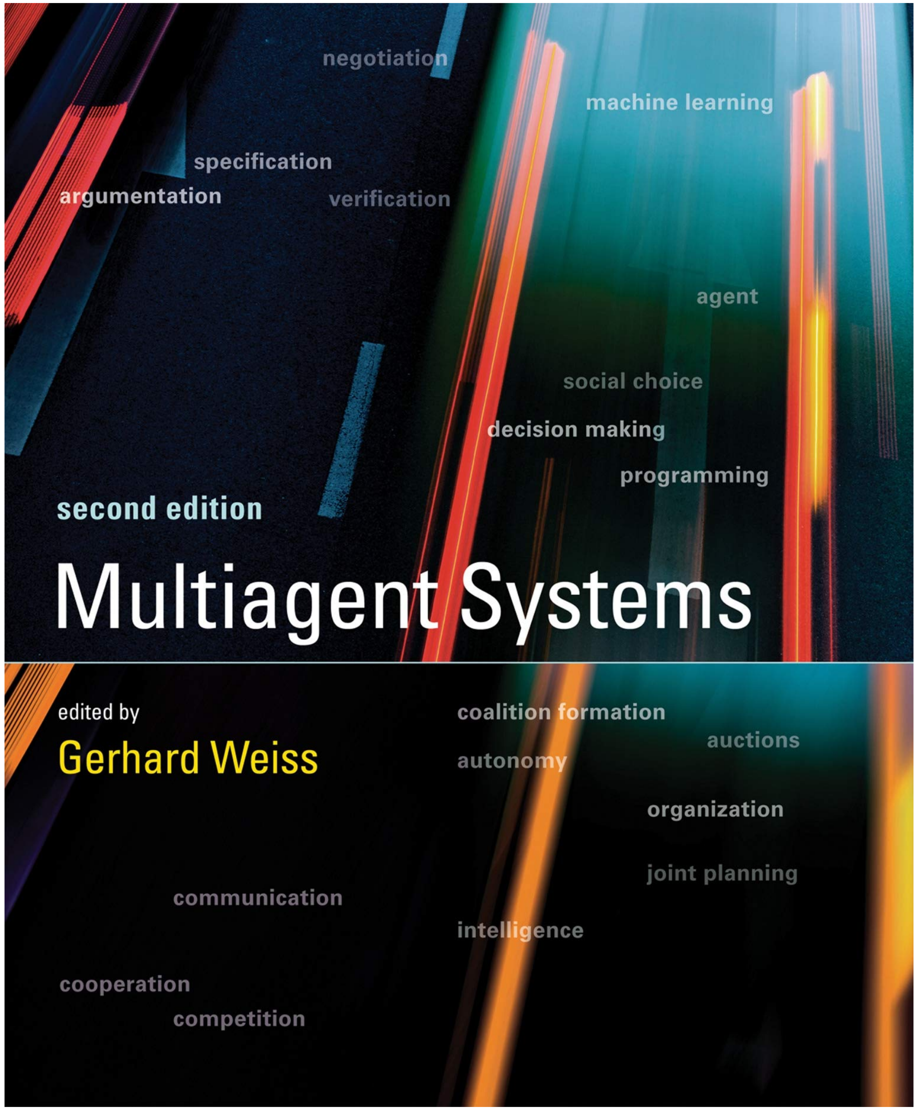

negotiation
machine learning
specification
argumentation
verification
agent
social choice
decision making
programming
second edition
Multiagent Systems
edited by
Gerhard Weiss
coalition formation
auctions
autonomy
organization
joint planning
communication
intelligence
cooperation
competition

Literature

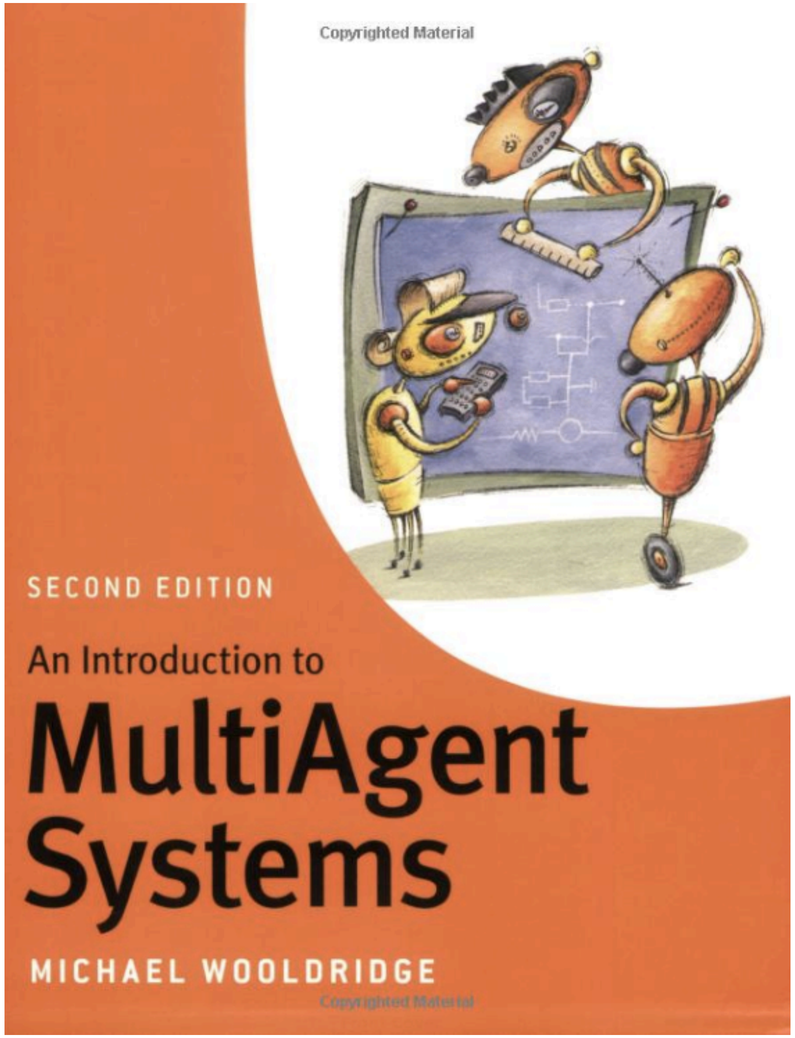
Copyrighted Material
SECOND EDITION
An Introduction to
MultiAgent
Systems
MICHAEL WOOLDRIDGE

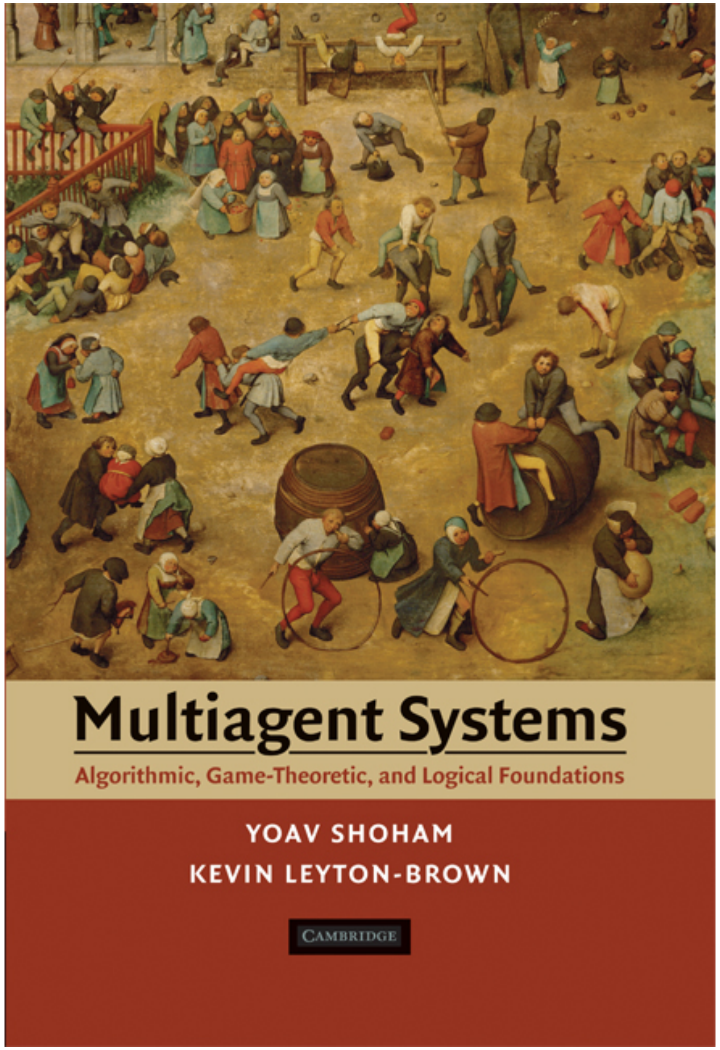
Multiagent Systems
Algorithmic, Game-Theoretic, and Logical Foundations
YOAV SHOHAM
KEVIN LEYTON-BROWN
CAMBRIDGE

Literature


Ronald Fagin
Joseph Y. Halpern
Yoram Moses
Moshe Y. Vardi
Reasoning About Knowledge


SYNTHESE LIBRARY
Hans van Ditmarsch
Wiebe van der Hoek
Barteld Kooi
Dynamic Epistemic Logic
Springer

What is an Agent?

- Anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators
[AIMA-Def]

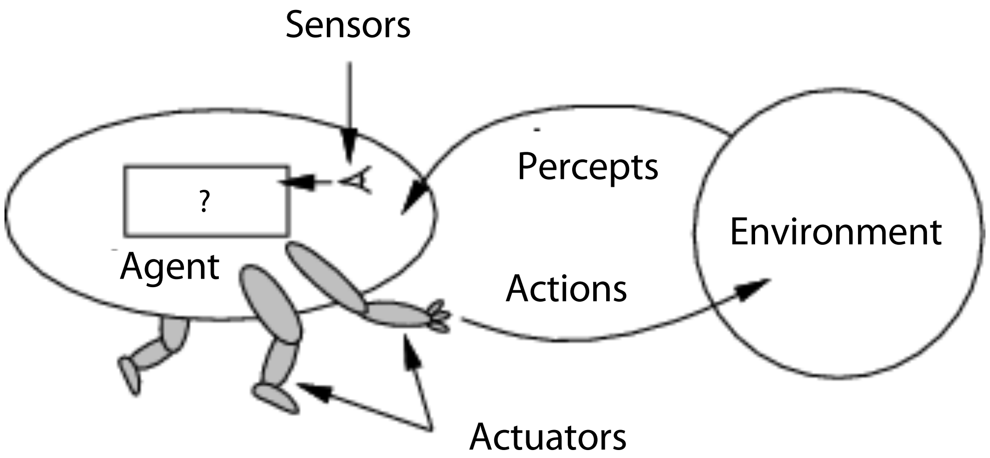

  - Human agent
eyes, ears, and other
organs for sensors; hands, legs, mouth, and other body parts for actuators
  - Robotic agent
cameras and infrared range finders for sensors; various motors for actuators
  - Software agent
interfaces, data integration, interpretation, data manipulation/output

Abstractions: Agents and Environments

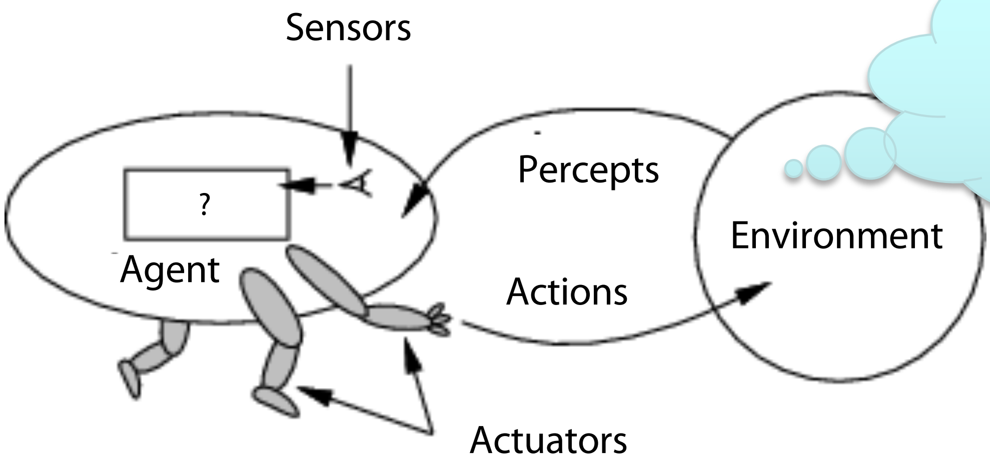

Environment contains other agents -> Mechanism Design

- The agent function maps from percept histories to actions:

  [f: P* → A]

- The agent program runs on a physical architecture to produce f
- Agent = architecture + program

Really insist on functional behavior?

Reactive vs. Goal-based Agents

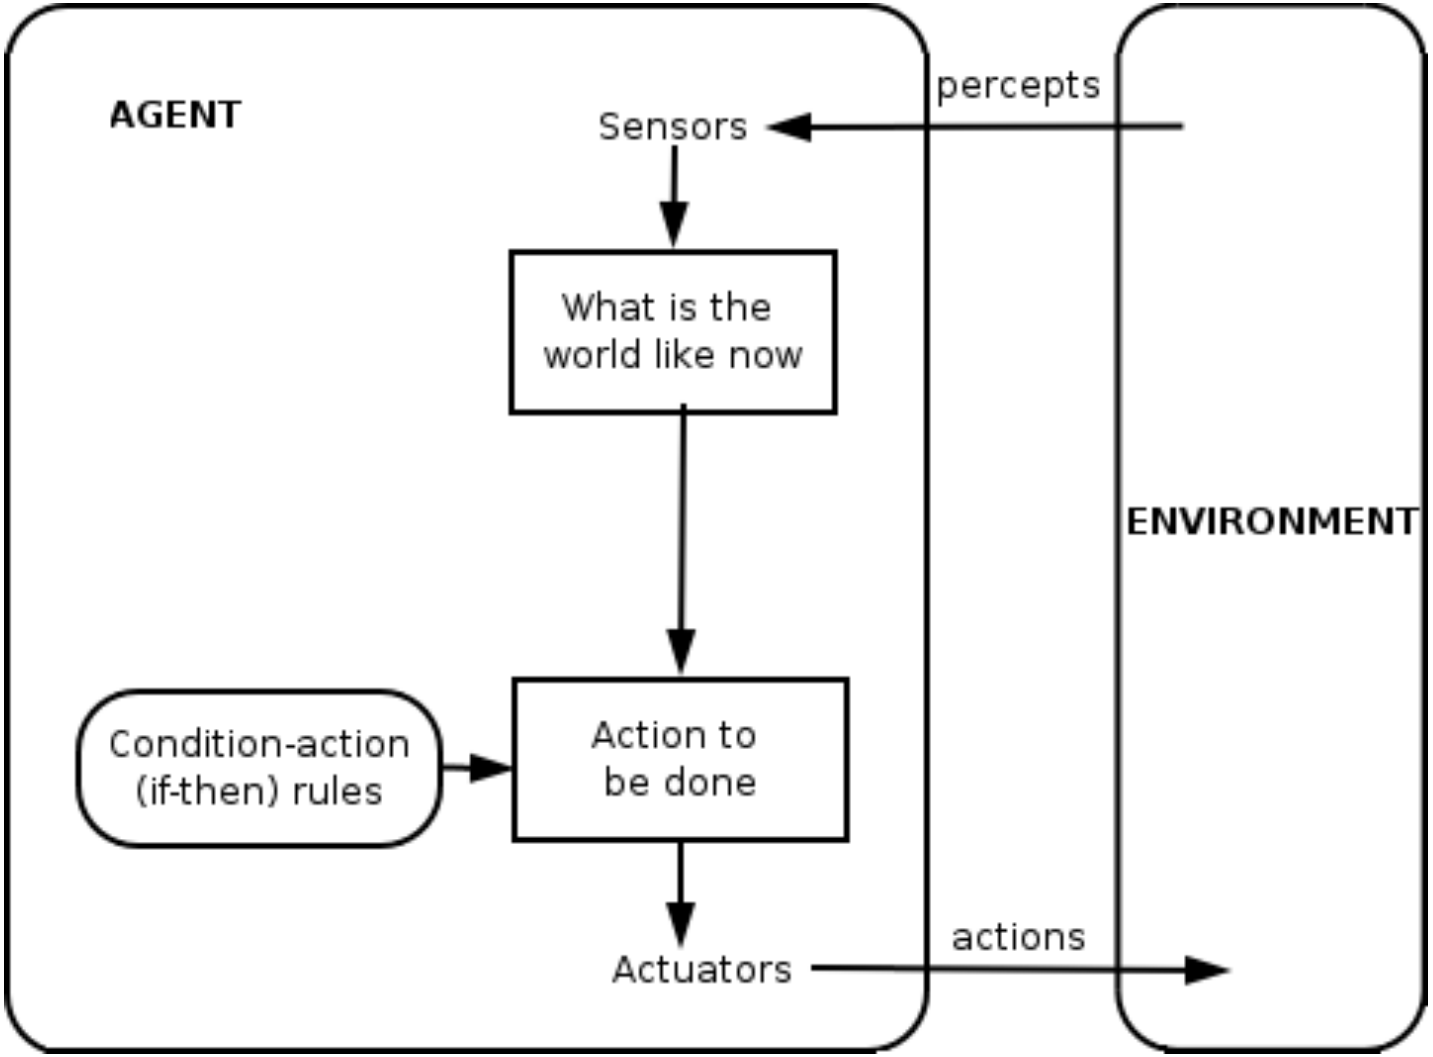

Reactive vs. *Goal-based* Agents

Agent Perception (Language and Vision)

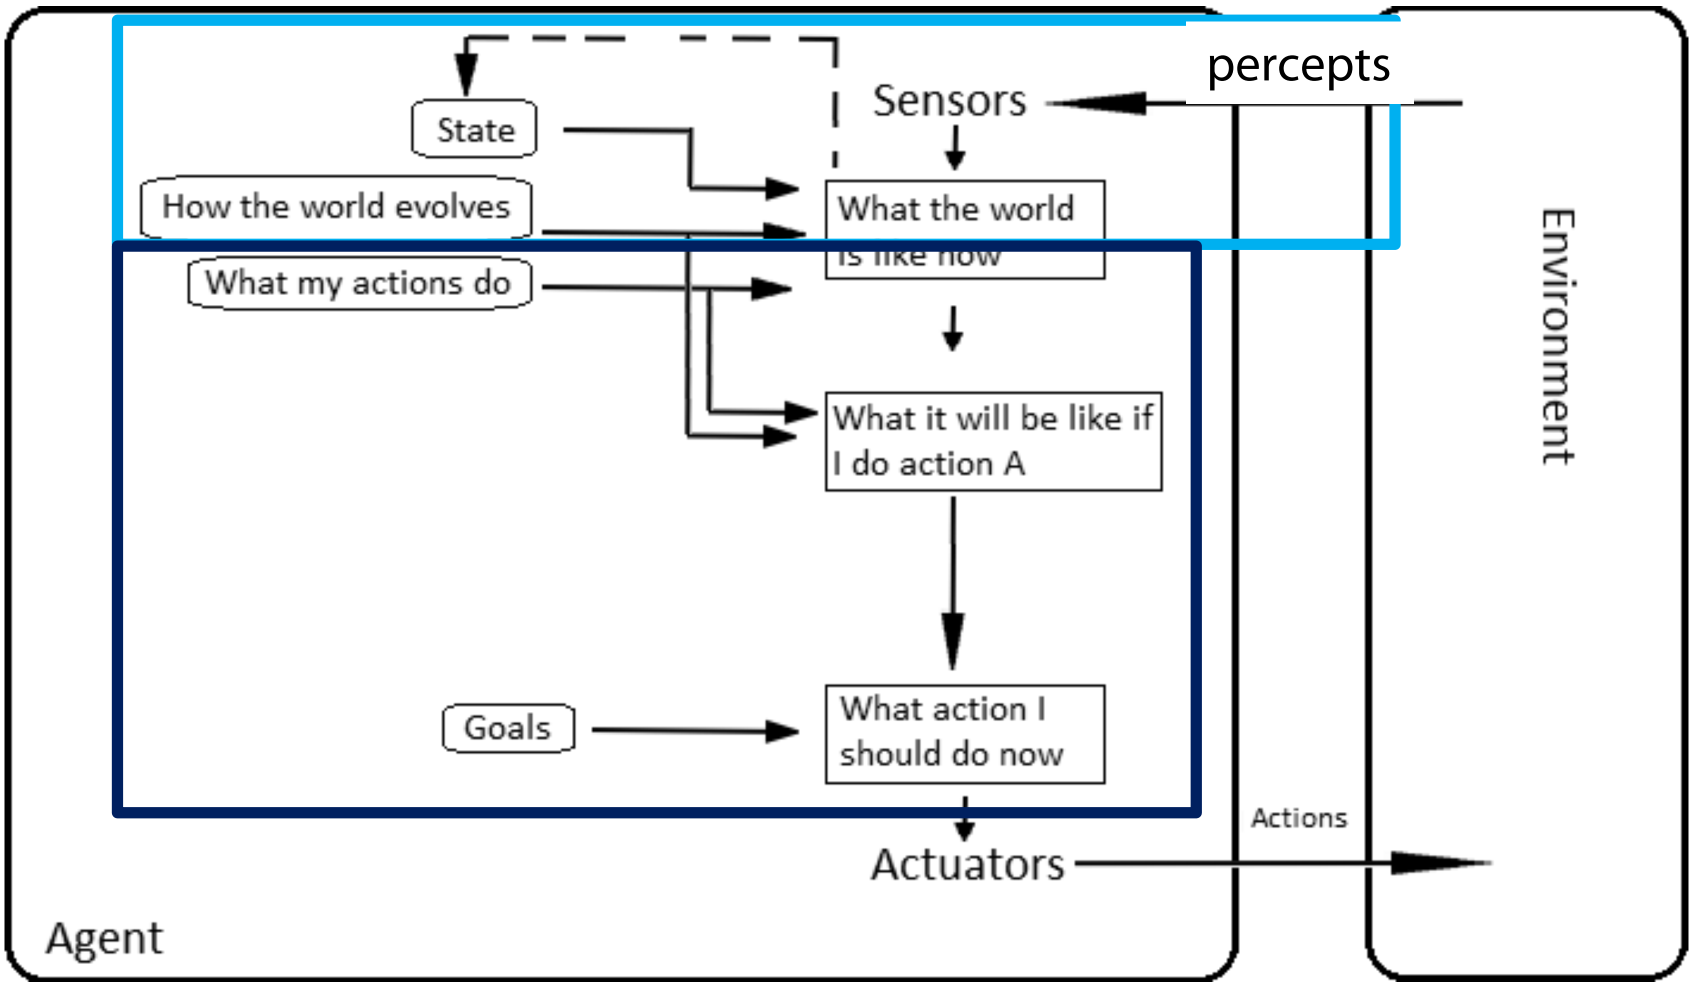

Agent Planning, Reinforcement Learning, Relation and Human-Awareness

Balancing Reactive and Goal-Oriented Behavior

- We want our agents to be reactive, responding to changing conditions in an appropriate fashion (e.g., timely)
- We want our agents to systematically work towards long-term goals
- These two considerations can be at odds with one another
  - Designing an agent that can balance the two remains an open research problem
  - Achieve maximum freedom of action if there is no specific short-term goal (e.g., keep batteries charged)

Social Ability

- The real world is a multi-agent environment: we cannot try achieving goals without taking others into account
- Some goals can only be achieved with the cooperation of others
- Social ability in agents is the ability to interact with other agents (and possibly humans) via some kind of agent-communication language …
  -> Agent Perception (Language and Vision)
- ... with the goal to let other agents to make commitments (of others) or reinforcements (about its own behavior)
- Need to represent and reason about beliefs about other agents
  -> Agents, Mechanism, and Collaboration

Rational Agents

- Rational Agent: For each possible percept sequence, a rational agent
  - should select an action
  - that is expected to maximize its local performance measure,
  - given the evidence provided by the percept sequence and
  - whatever built-in knowledge the agent has
- Rational = Intelligent ?
  - There is more to intelligence than to meet rationality

Autonomous Agents

- Rationality is distinct from omniscience (all-knowing with infinite knowledge)
- Computing the best action usually intractable
- Rationality is bounded
- Agents can perform actions in order to modify future percepts so as to obtain useful information
- An agent is autonomous if its behavior is determined by its own "experience" (with ability to learn and adapt)
  - What matters for the "experience" is the
    - percept sequence (which the agents can determine), the
    - state representation, and the
    - "computational success" of computing the best action as well as learning and adapting for the future

Autonomy ≠ No influence possible

- Clearly: we do not prescribe concretely the agent function

- But one wants and can have influence on the meta-level
(-> see revival of meta-reasoning)

- In other words: Have influence on the rules of the interaction game
- Agents, autonomously following them, will act towards an environment in a desired, beneficial state

- This is part of
mechanism design = game theory + social choice theory

Human-compatible Behavior

- Agents act on behalf of humans, whose goals the agents have to fufill (this is the single goal of agents)
- Agent should consider its initial goals to be uncertain
- Agent should be able to prove their behavior is beneficial to humans
- Artificial intelligence, agents, and ethics
  - Agents (and their designer) must act in an ethical way
    Developers should be able to prove …
    … that agents are able to prove
    … that they (the agents) act in an ethical way
  - Simple technology assessment is not enough
  - And yes, there are formal ethics, there is deontic logic, ...

Human-aware Behavior

- Agents interact with humans
- Selected actions must match human expectations
  - Maybe the presumably expected action might not be the best (for the human or the agent, or both)
- Selected actions that are assumed to not match human expectations must be explained

Agent Model vs. Human Model

Agent Planning, Reinforcement Learning, Relations, and Human-Awareness

Agents, Mechanism, and Collaboration (lecture)

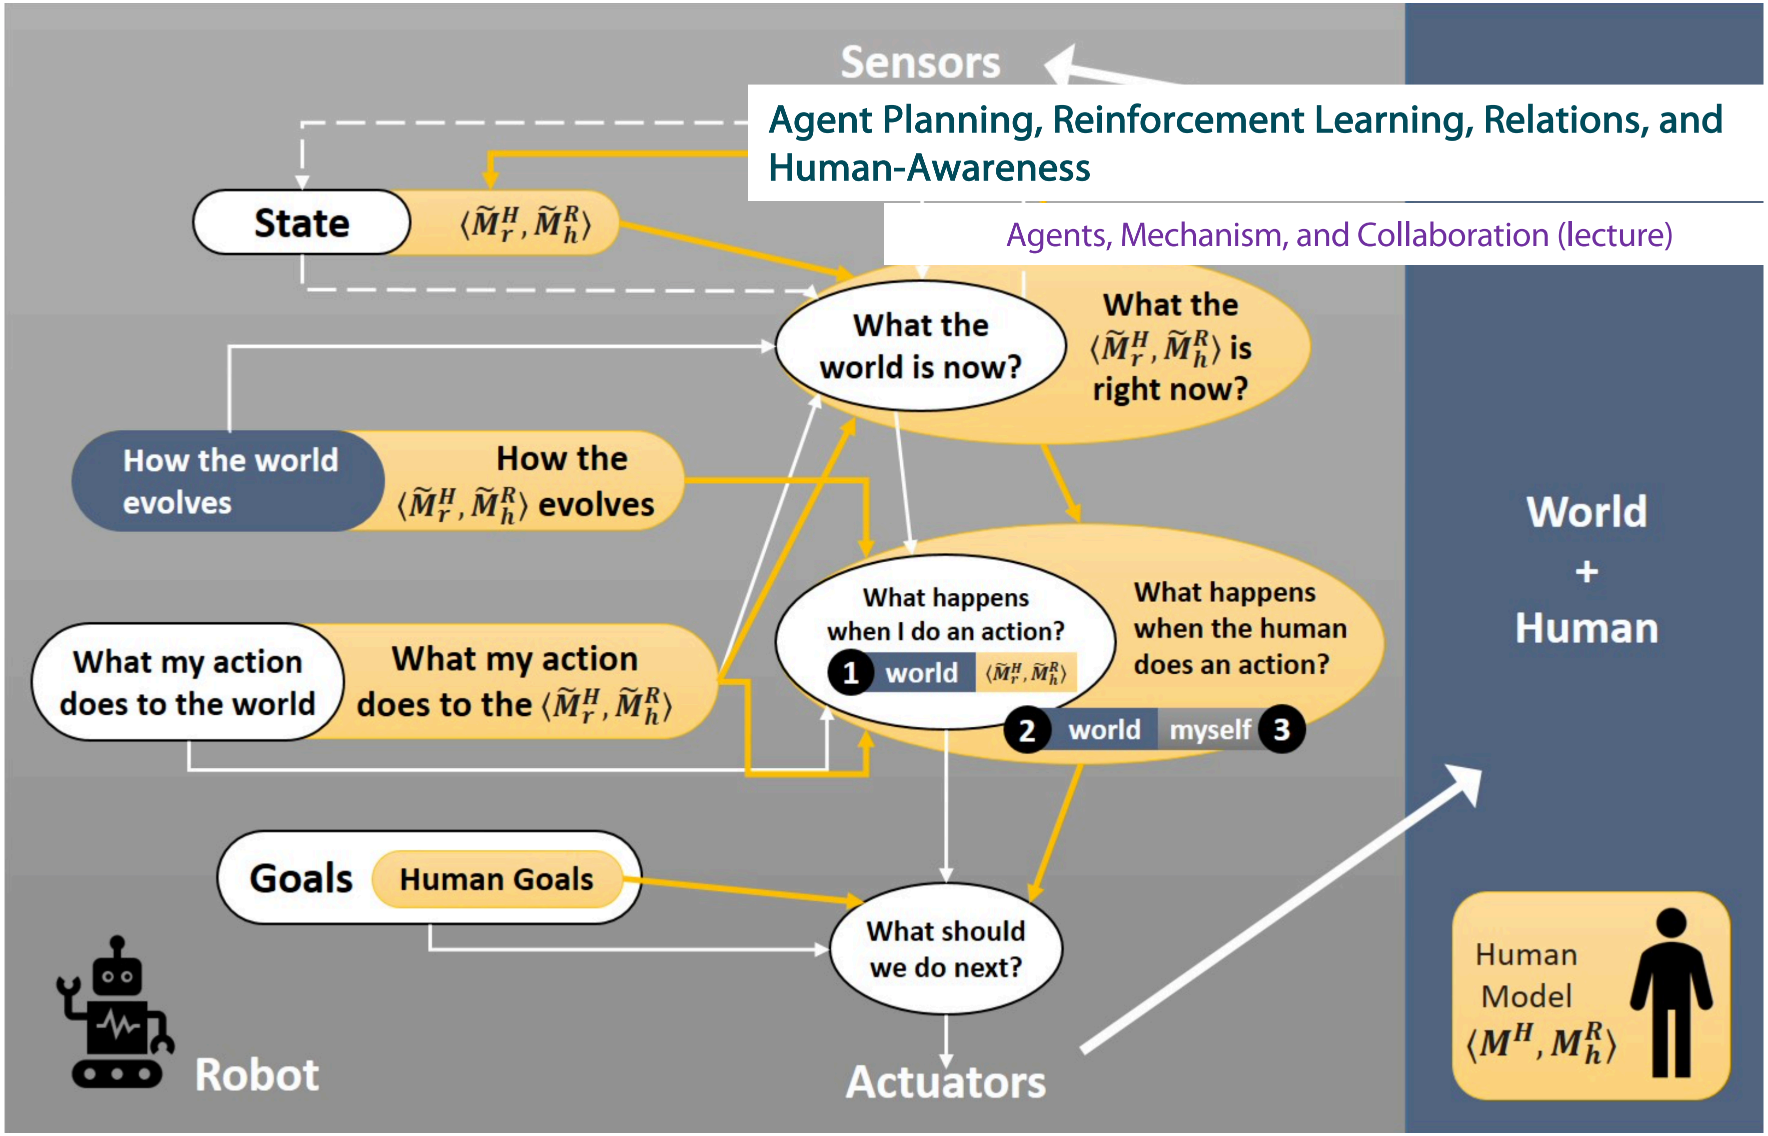

Learning Agents (Online)

- Ever extended percept sequence (incl. more or less explicitly encoded reinforcement feedback or rewards) is …
  - … sparse (no big data), but gives rise to model updates
  - … with the aim to better (faster) achieve goals
- We say: Agents learn (and we mean: while acting, or online)
  - Optimize a performance measure
- Setting up agents' online learning engines
  - Dedicated knowledge about online learning required
- Setting up an agent's initial model by exploiting data:
  - Dedicated knowledge of machine learning required
  - Also basically optimizing a performance measure

Reactive vs. *Goal-based* Agents

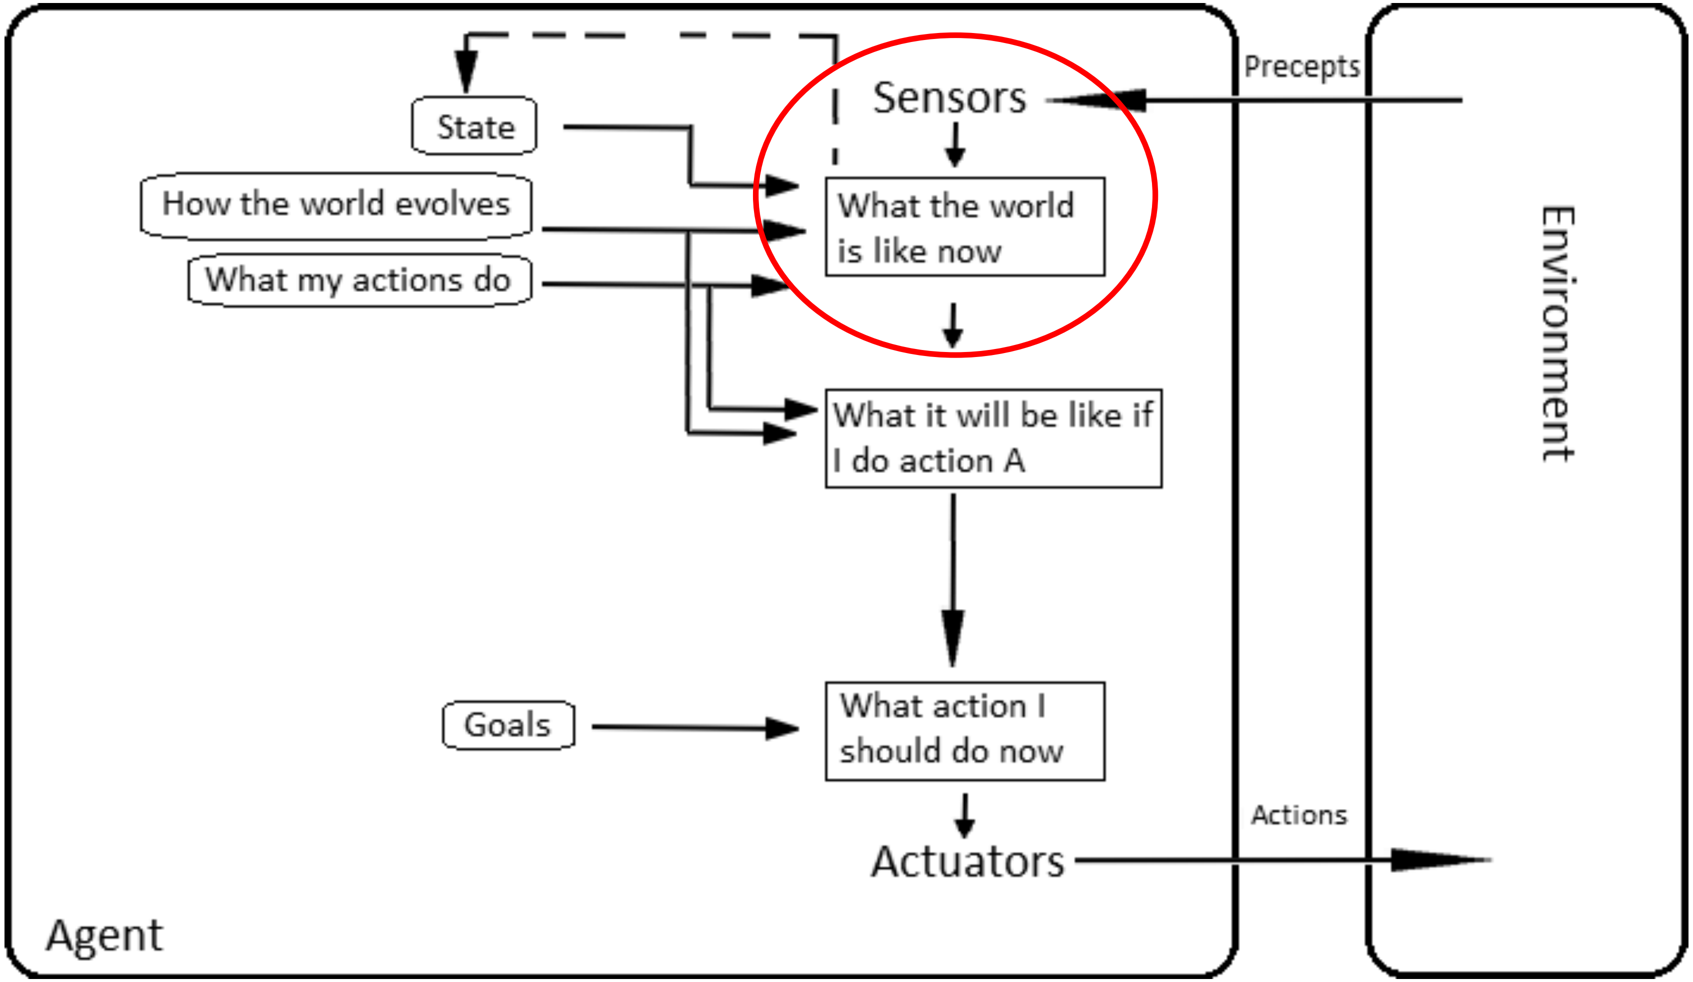

Misunderstandings

- Applying ML to implement a function f some people say: "I have used ML technique X to create an AI"
- Unconsciously, AI is used as a synonym for agent, but … … mostly a very simple one
  - f: P → A
- Claiming that f is "an AI" is an indication of lack of understanding …
- … even if the last n percepts are considered
  - f: P × … × P → A
- One is lost w/o an understanding of intelligent agents
  - f: P* → A

Frame Agents

- Assume that machine learning techniques are used to build models at agent setup time
- Runtime behavior of agent always depends on last n elements of percept sequence only $f: P \times \ldots \times P \rightarrow A$
- No interaction w/ environment, no feedback
- Agent is fake (simply a frame around standard SW/HW)
  - Also holds when setup training data is camouflaged as initial percepts (but no actions towards goals are computed until training completed)
- Maybe even enlightening for practical applications, but agent idea …
- … does not show its full potential

Learning-based Software Development

- There is no need to deliberately conflate machine learning with agents and AI !
- No need to invent frame agents !
- Can build extremely cool SW/HW w/ machine learning techniques (e.g., for industrial image processing applications)

- → Probabilistic Differential Programming (CS5071-KP04)
- → Deep Learning Lab (CS5071-KP04)

Back to the Future: Human-guided Learning

- Develop machine learning techniques that achieve good performace w/o too much training material
- Exploit human capabilities
- Artificial agents and human agents cooperate
- Machine learning becomes agent online learning
  - Motivation for studying agents!
  - Machine learning cannot go w/o agents in the future

- Agents allow for more or less learning (incl. no learning)
- Next: Proper agent with no learning

Proper Agent: An Example

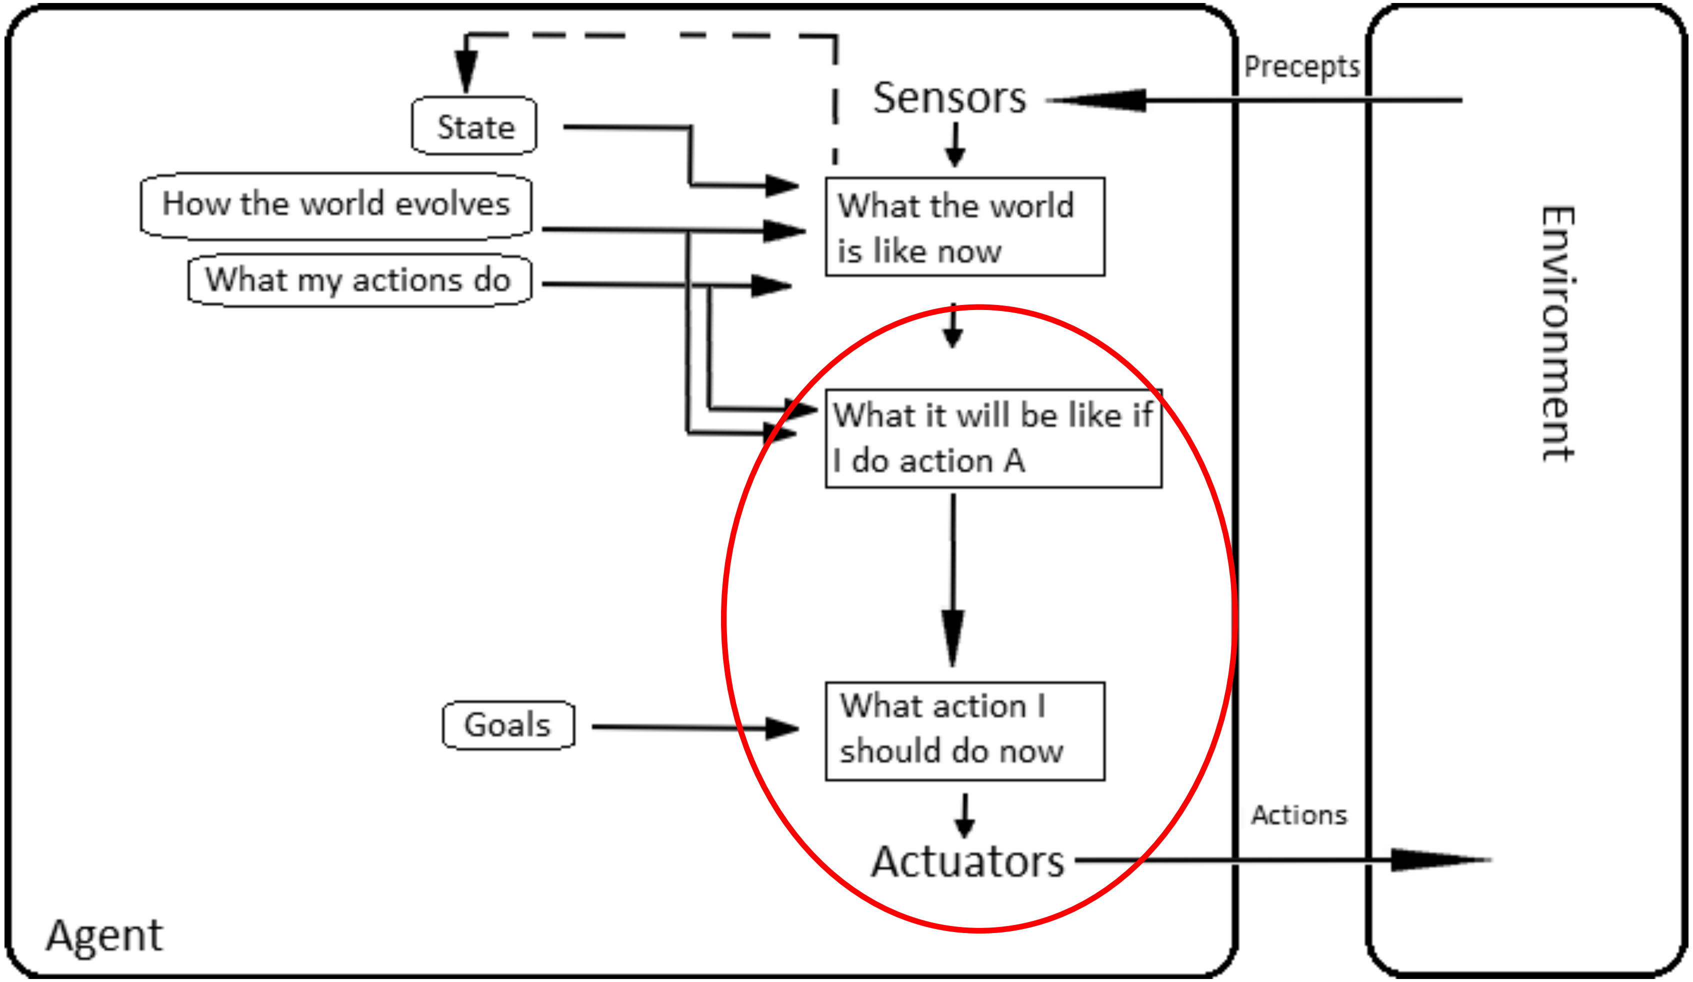

Proper Agent: An Example

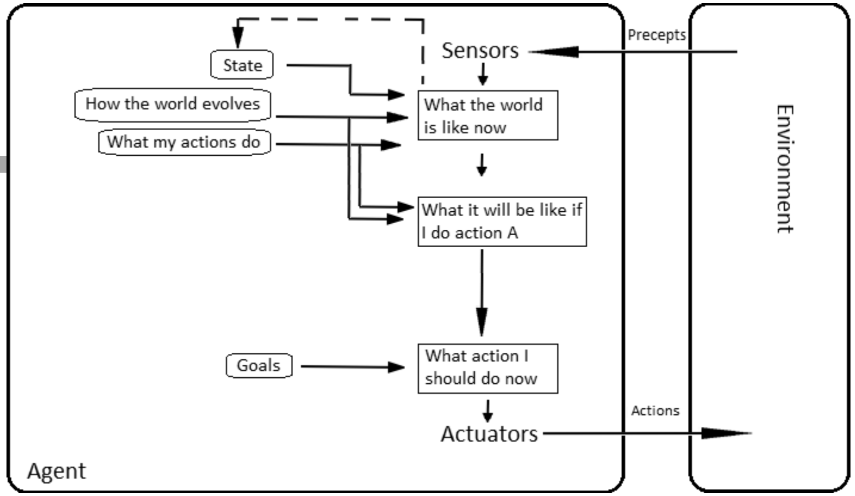

Given:

- Current state of the environment
- Description of goal state
- Set of action descriptions

→ Find sequence of actions (a plan) for transforming current state into goal state

→ Select first action, and hope that plan can be completed

STRIPS Formalism

- States modeled as set of ground atoms (database)
  - Current state as well as goal state
  - Example: Blocks World
    - On_Table(A), On_Table(B), On_Table(C)
    - On_Block(C, B), On_Block(B, A)
- Historical & name convention note
  - Nowadays STRIPS = representation language FOL
  - Part of PDDL (planning domain definition language)
  - Originally, STRIPS also denoted (linear, non-complete) hill-climbing style algorithm
  - Here we consider a different algorithm based on partial order plans

STRIPS = Stanford Research Institute Problem Solver (1971)

STRIPS Planning Operators

$$
\begin{aligned}
&Op(\text{Action: } Go(there),\\
&\quad \text{Precond: } At(here) \land Path(here, there),\\
&\quad \text{Effect: } At(there) \land \neg At(here))
\end{aligned}
$$

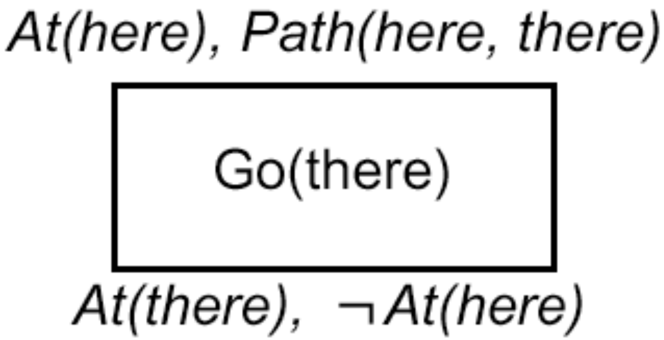

Initial Plan

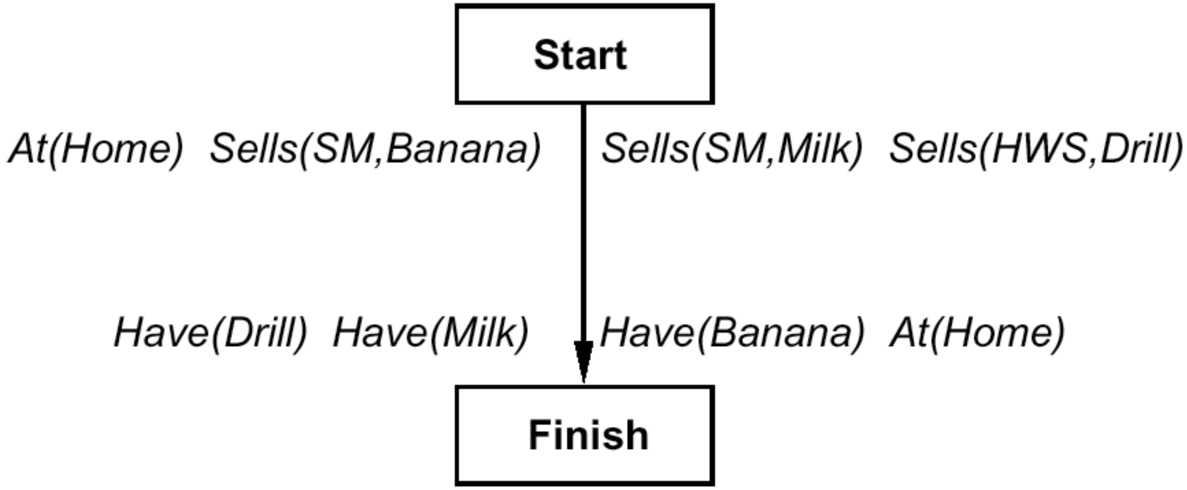

Action(Type)s:

Op (Action: *go(there)*,
Precond: *At(here)*,
Effect: *At(there)* $\wedge$ $\neg$*At(here)*)

Op (Action: *buy(x)*,
Precond: *At(store)* $\wedge$ *Sells(store,x)*,
Effect: *Have(x)*)

there, here, x, store are variables

Planning as („Lifted“) Search

- Forward
- Backward

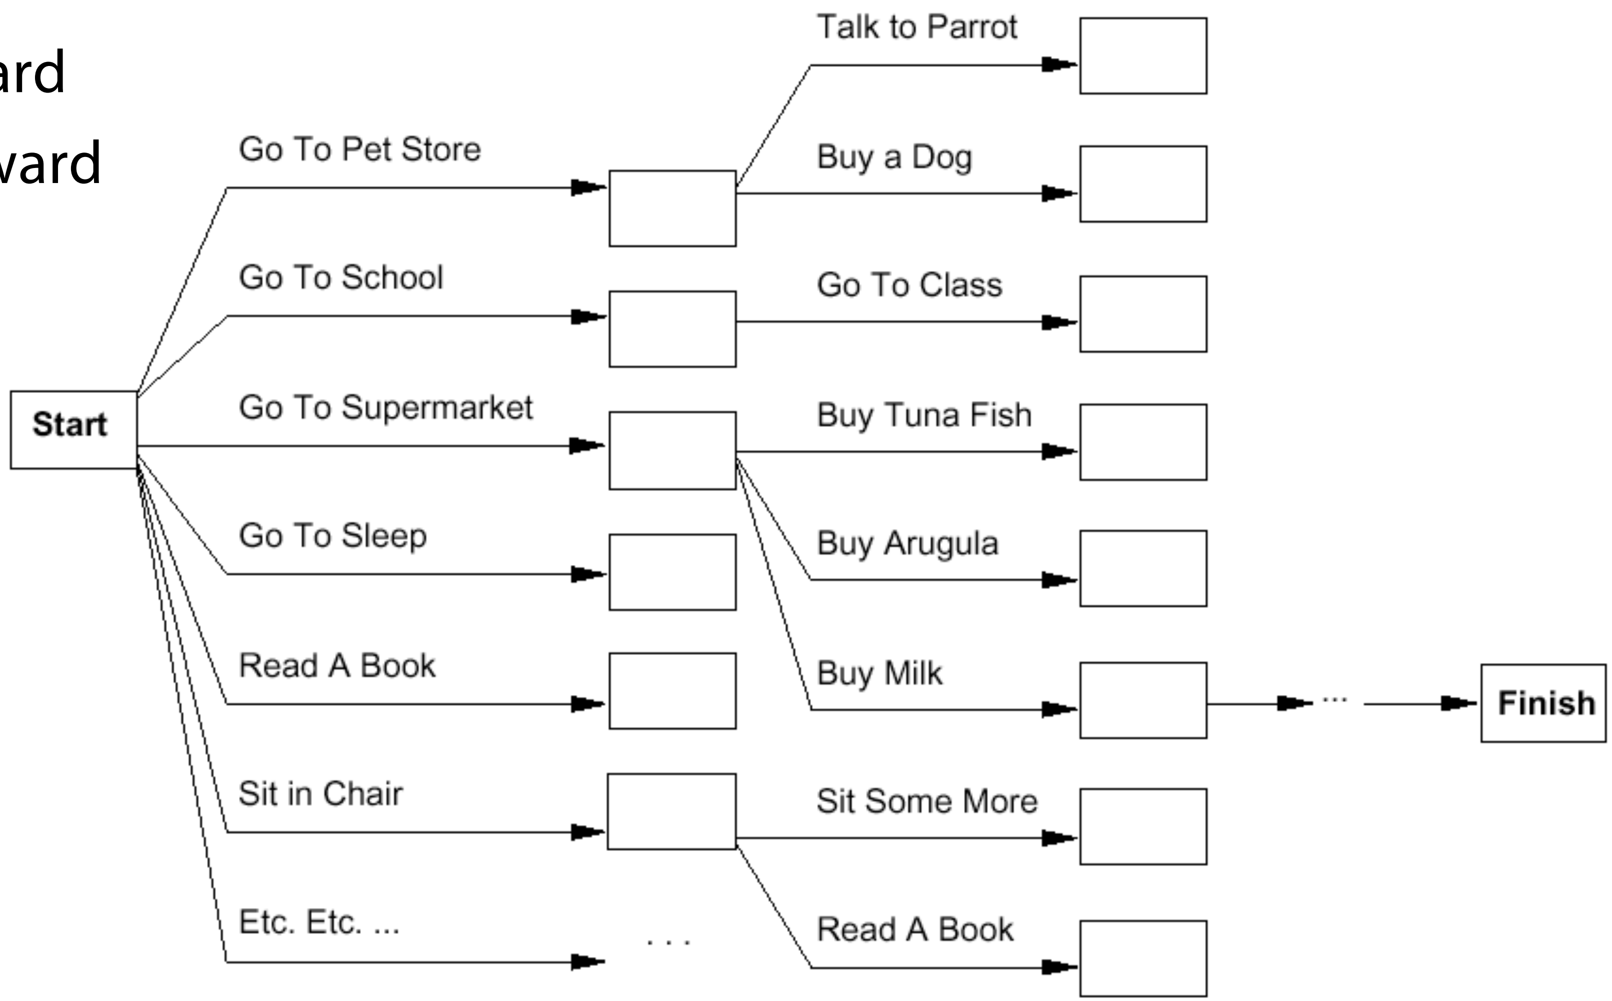

Classical Search vs. Planning

Problem solving by search

- Problem = single state space
- Program search algorithm

Problem solving by planning

- Specify problem declaratively
- Solve by general planning algorithm
- E.g.: instead of go-to-SM1, go-to-HW
  rather go(place)

Plan = (Linear) sequence of Actions?

Apply principle of
Least Commitment

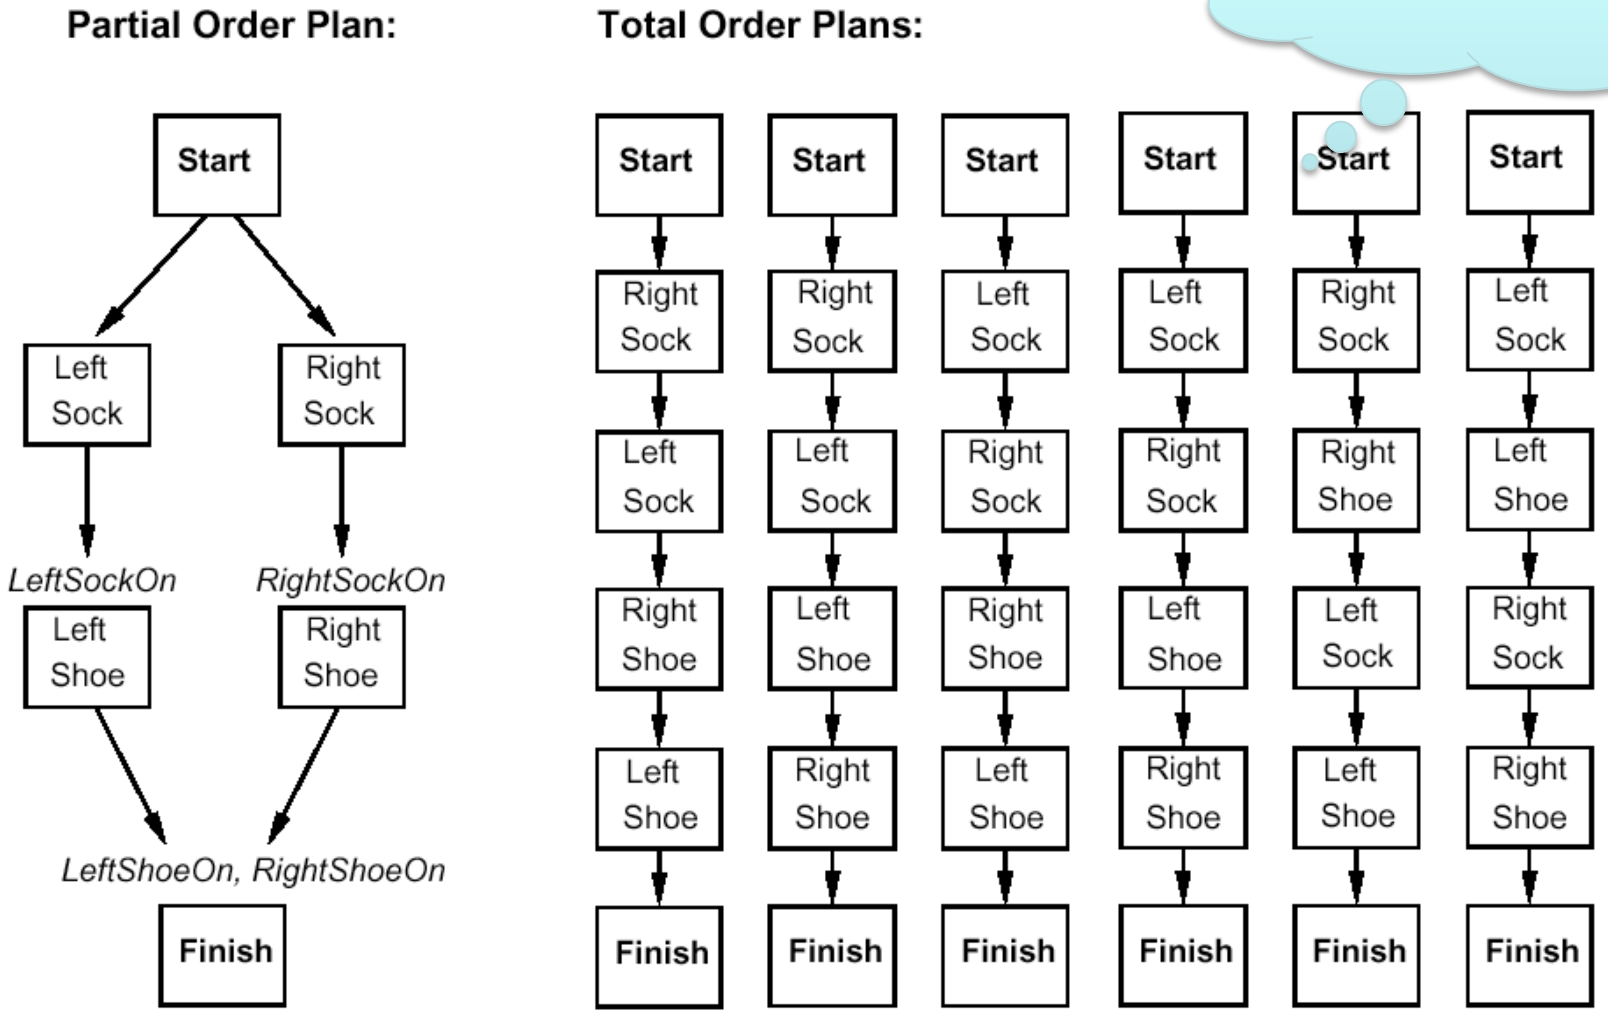

Representation of Partial-Order Plans

- Plan step = STRIPS Operator
- Plan consists of
  - Plan steps with partial order (<),
    where $S_i < S_j$ iff S_i is to be executed before S_j
  - Set of variable assignments $x = t$,
    where x is a variable and t is a constant or variable
  - Set of causal relations:
    $S_i \rightarrow^c S_j$ means S_i creates the precondition c of S_j (entails $S_i < S_j$)

- Solutions to planning problems …
  … must satisfy certain conditions

STRIPS = Stanford Research Institute Problem Solver (1971)

Completeness and Consistency

- Complete plan
  - Every precondition of a step is fulfilled
  - $\forall S_j$ with $c \in Precond(S_j)$,
    - $\exists S_i$ s.t. $S_i < S_j$ and $c \in Effects(S_i)$, and
    - for every linearization it holds that:
      - $\forall S_k$ with $S_i < S_k < S_j$, $\neg c \notin Effects(S_k)$
- Consistent plan
  - If $S_i < S_j$, then $S_j \not< S_i$ and
  - If $x = A$, then $x \neq B$ for different A and B for variable x (Unique Names Assumption)
- Solution of the planning problem: complete and consistent plan

Plan Refinement (1)

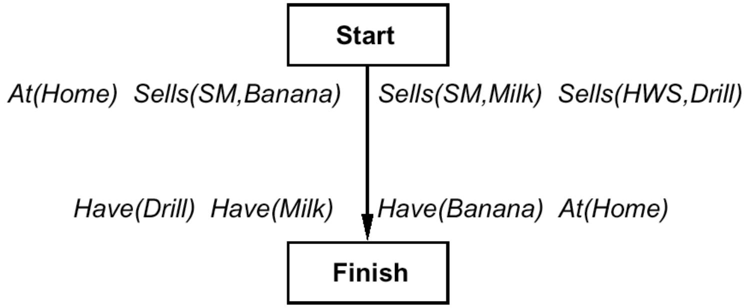

- Backward planning

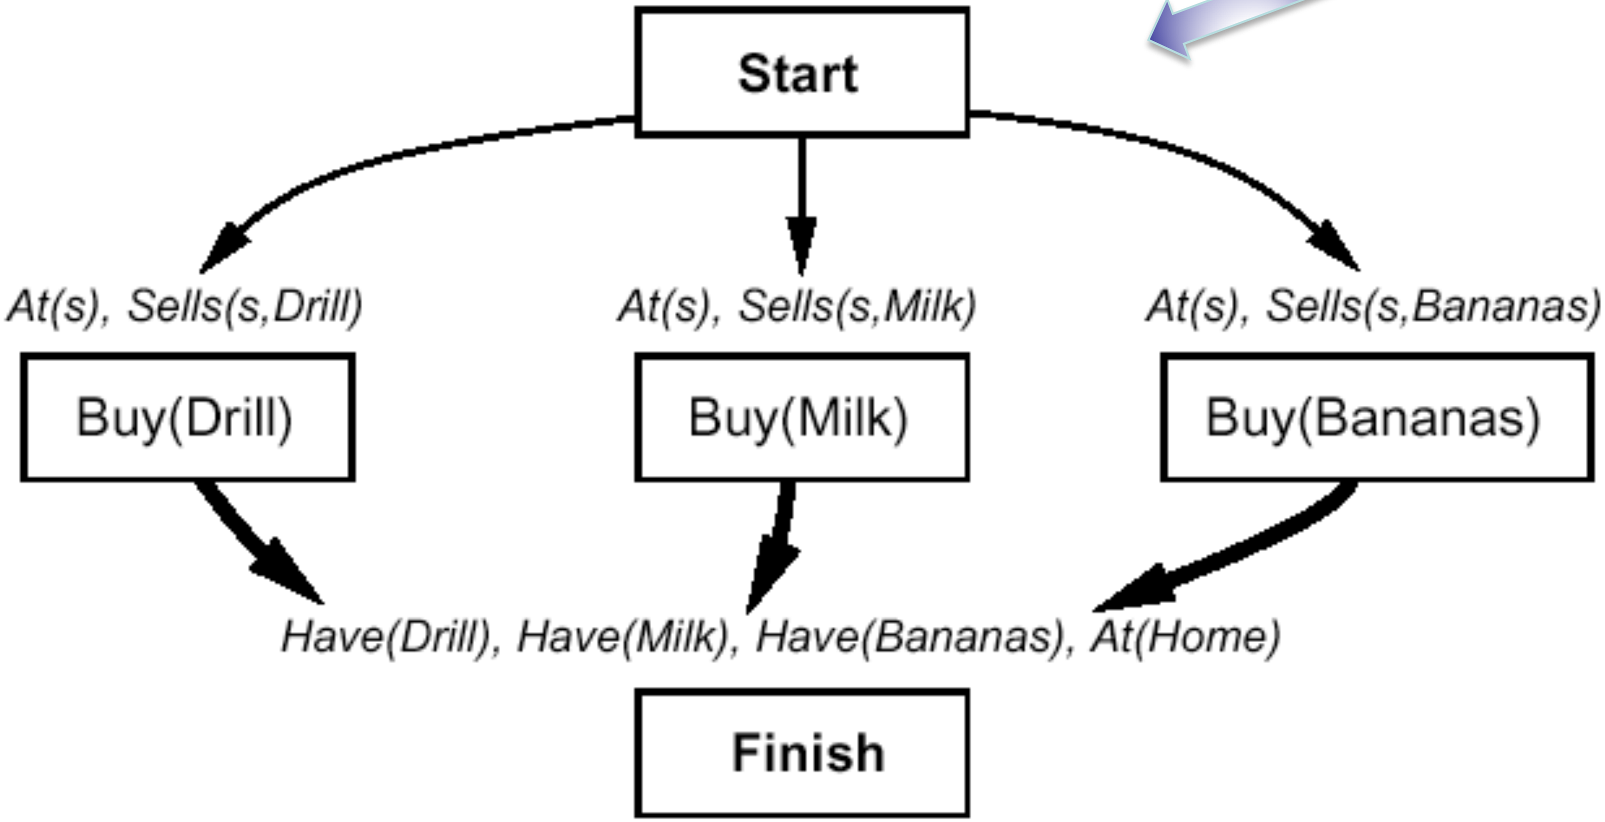

Thin arrows = <

Fat arrows = causal relation + <

Plan Refinement (1)

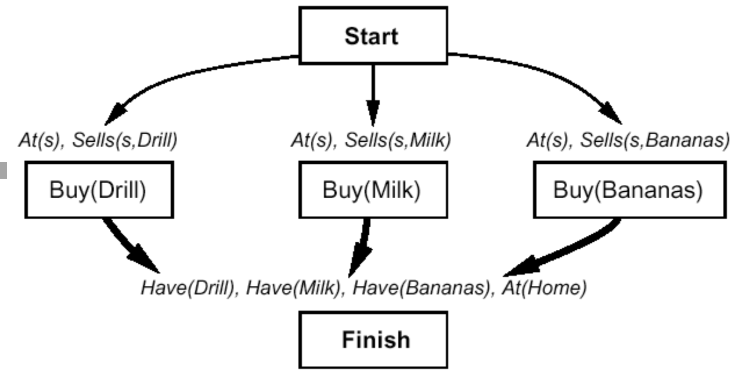

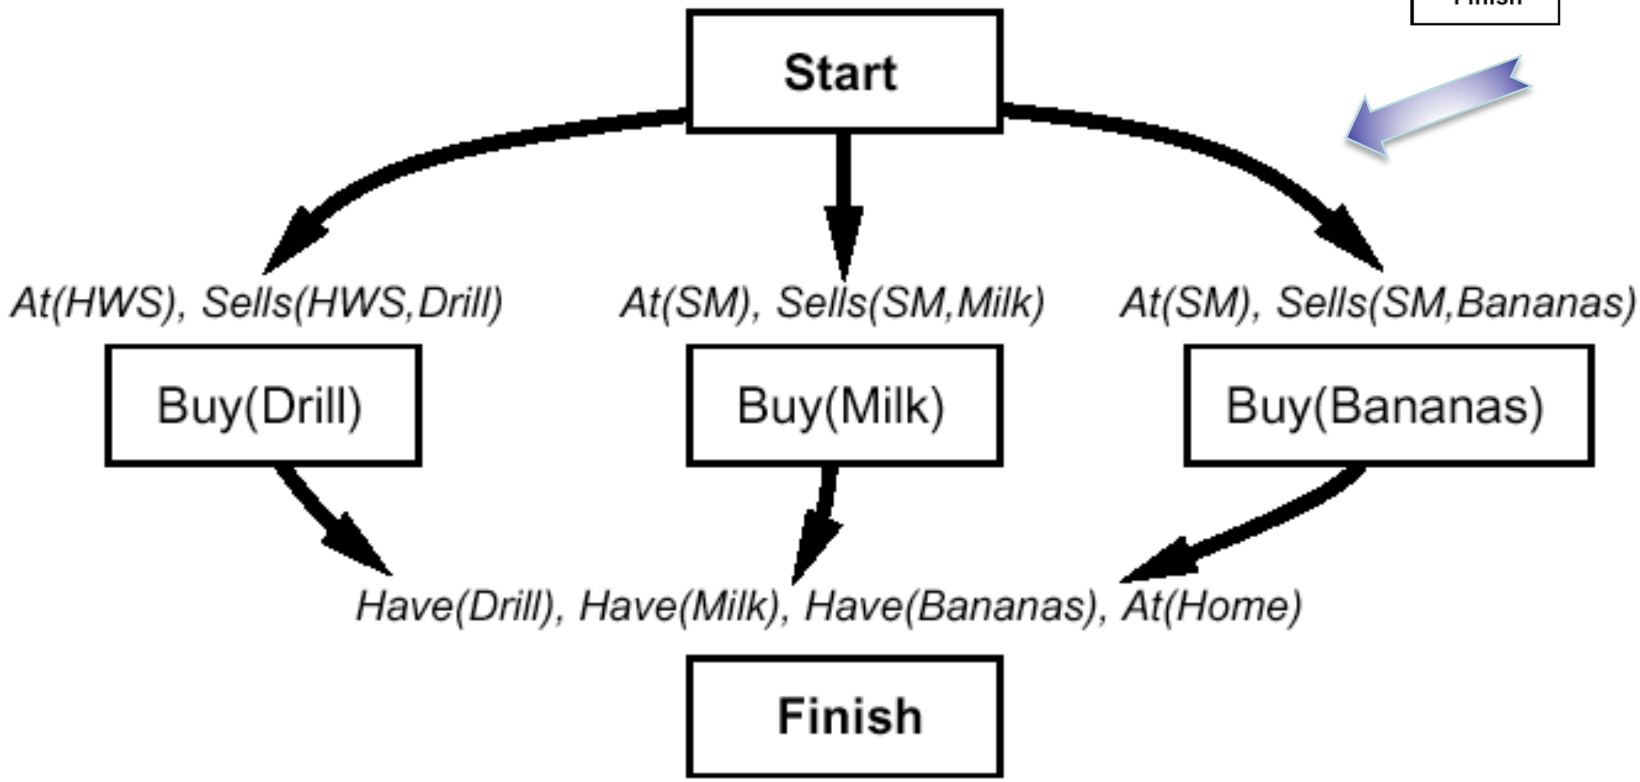

. . . after variable instantiation

Plan Refinement (2)

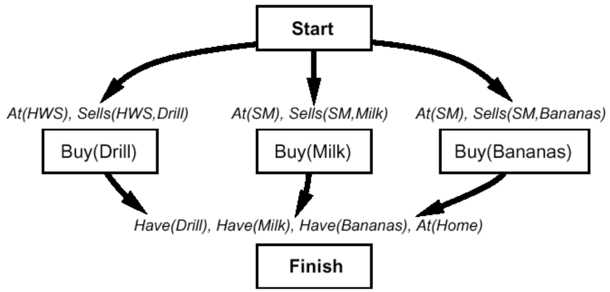

- ... buy at the right store

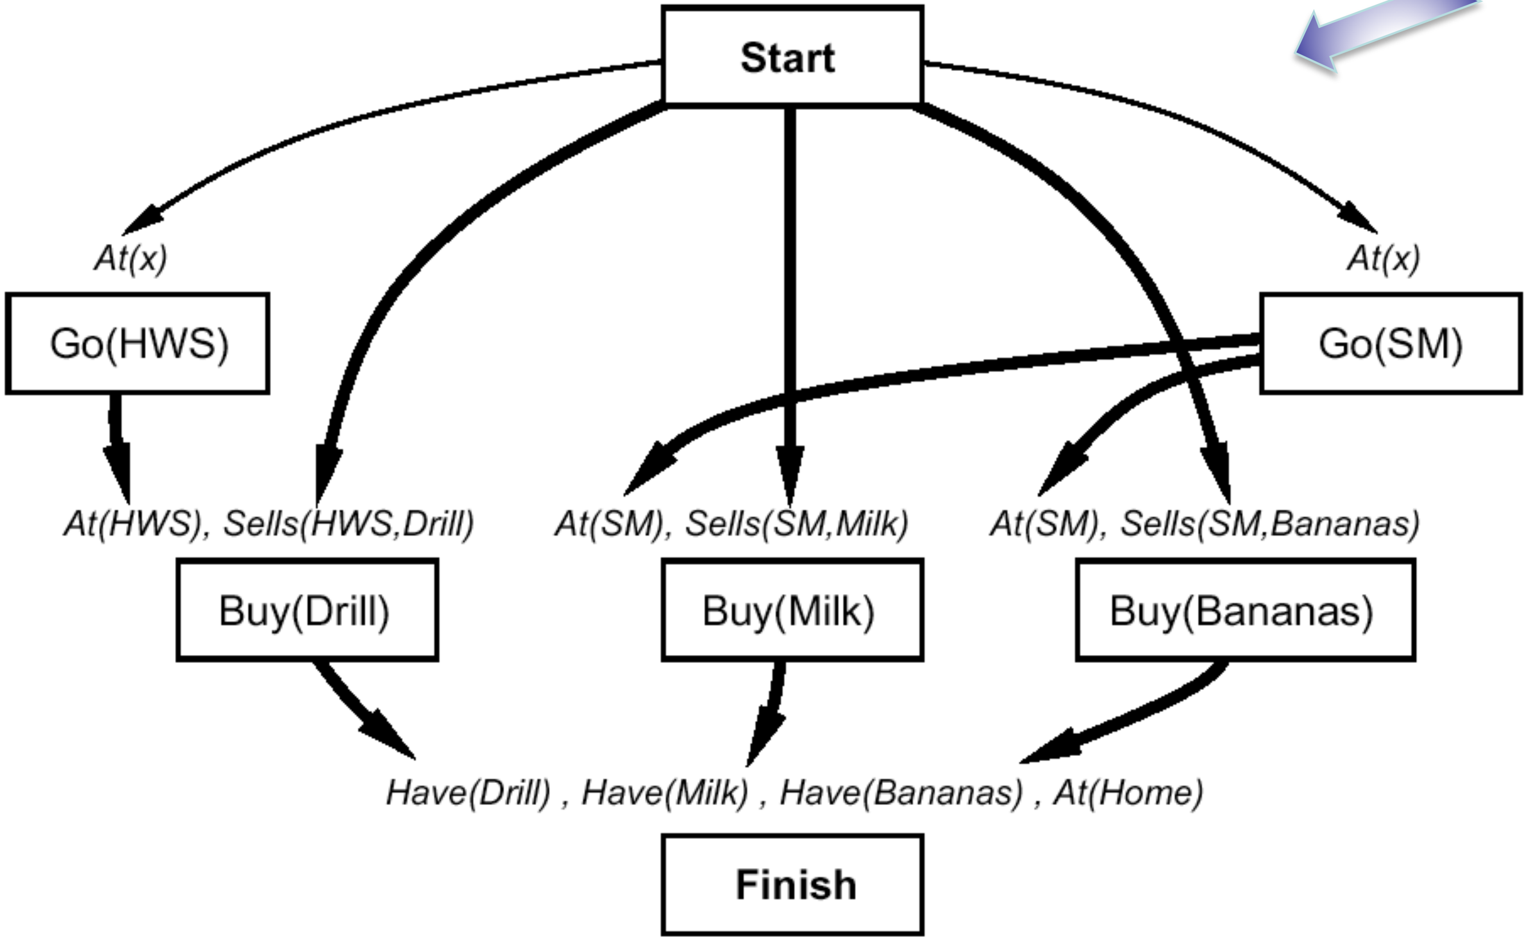

Plan Refinement (3)

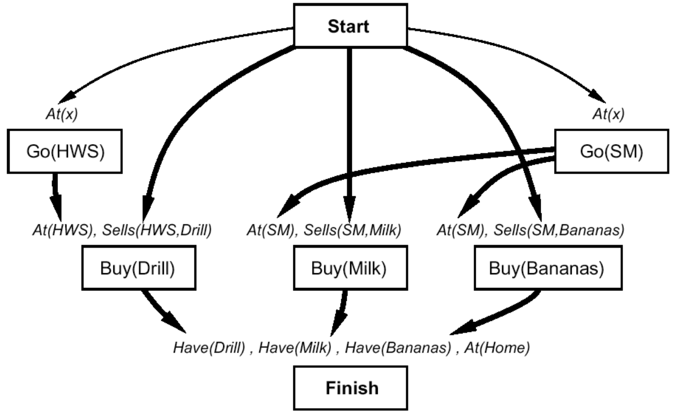

- ... but you must get there

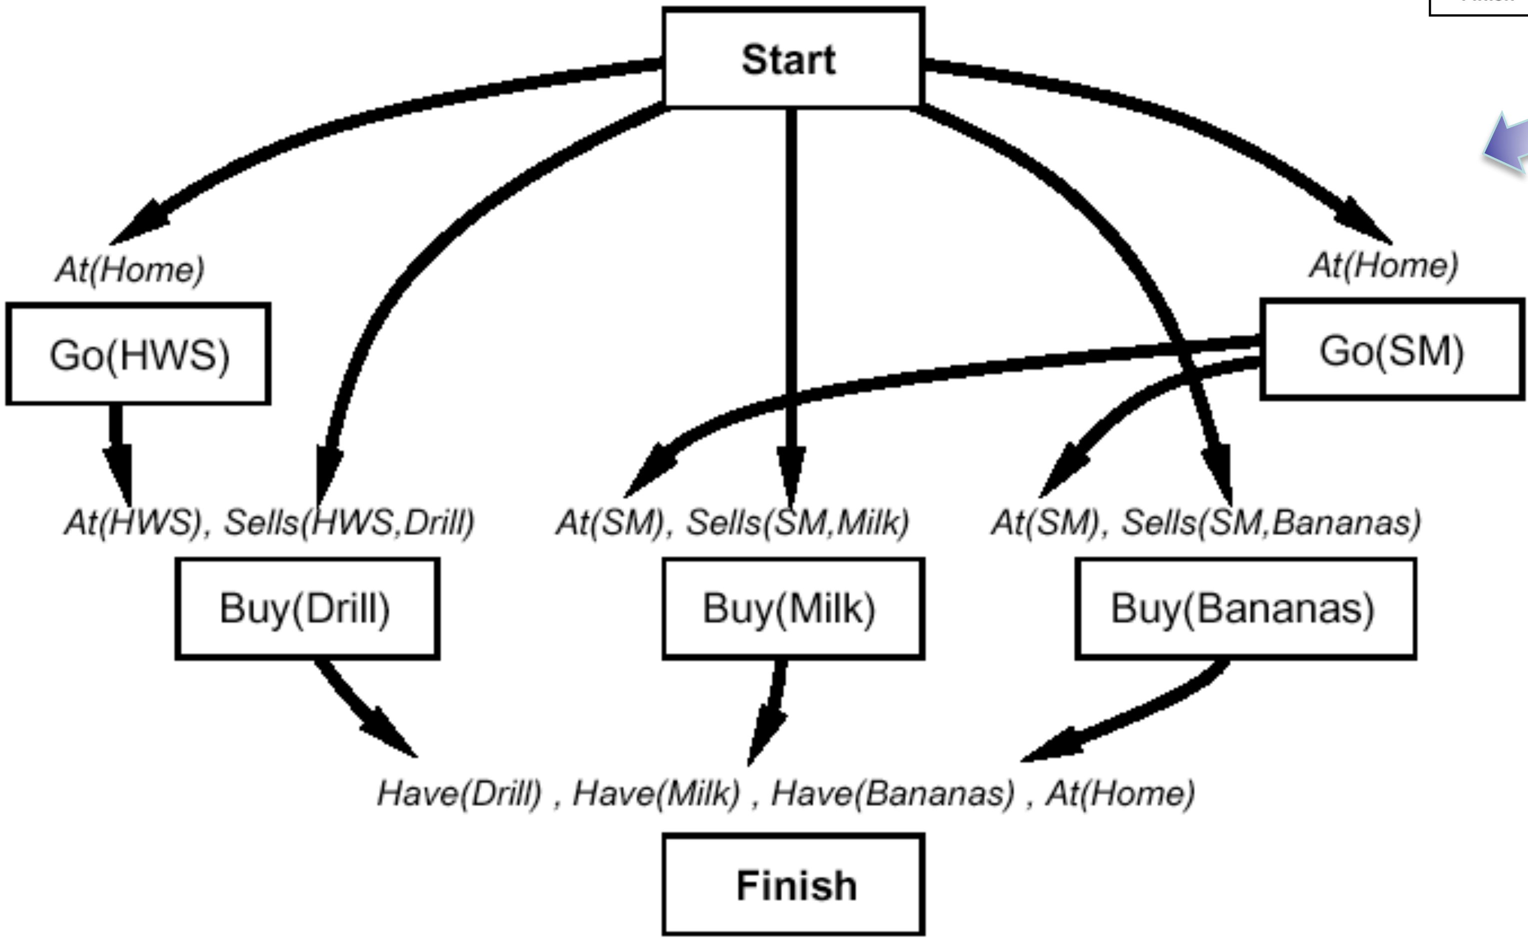

Plan Refinement (3)

- Note:
  Up to now no search, but simple „backward chaining“
- Now:
  Conflict! After go(HWS) is executed, At(Home) no longer holds (similarly for go(SM))

Protection of Causal Relations

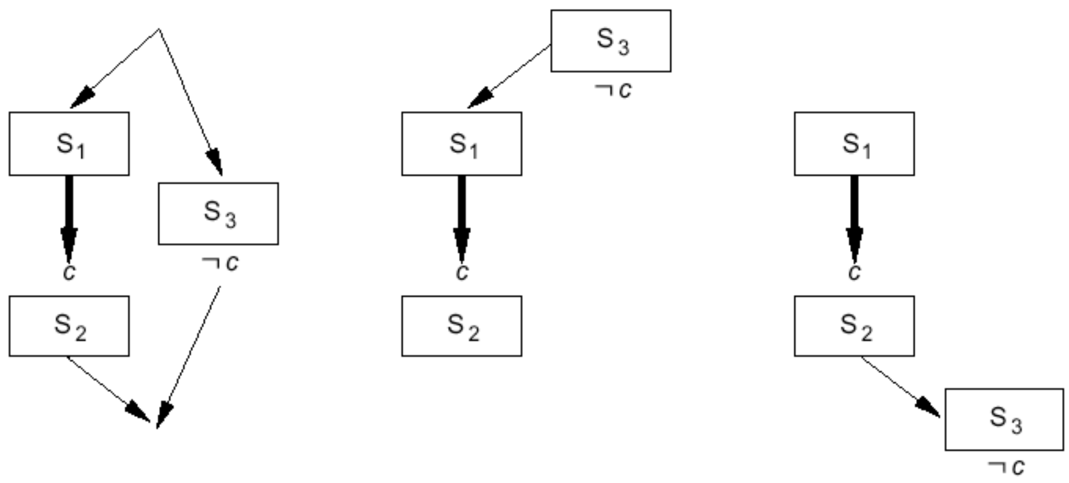

(a) (b) (c)

- Conflict:
  - S3 "threatens" causal relation between S1 and S2
- Conflict resolution:
  - Promotion: Put threat before causal relation (b)
  - Demotion: Put threat after causal relation (c)

Another Plan Refinement …

- Assumption: Cannot resolve conflict by protection
- Made a wrong step during plan refinement
- Alternative
  - Select x = HWS (with causal relation) while instantiating At(x) in go(SM)

Another Plan Refinement …

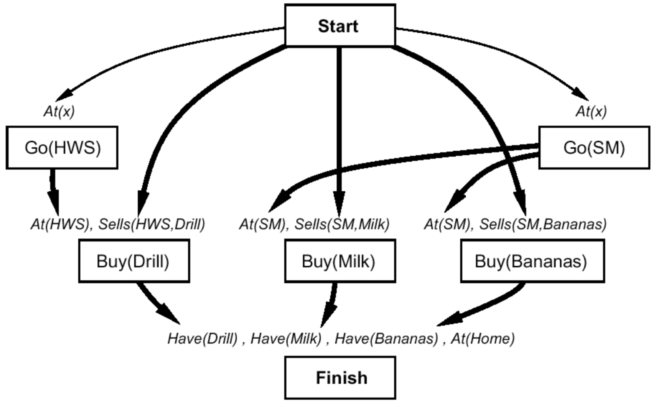

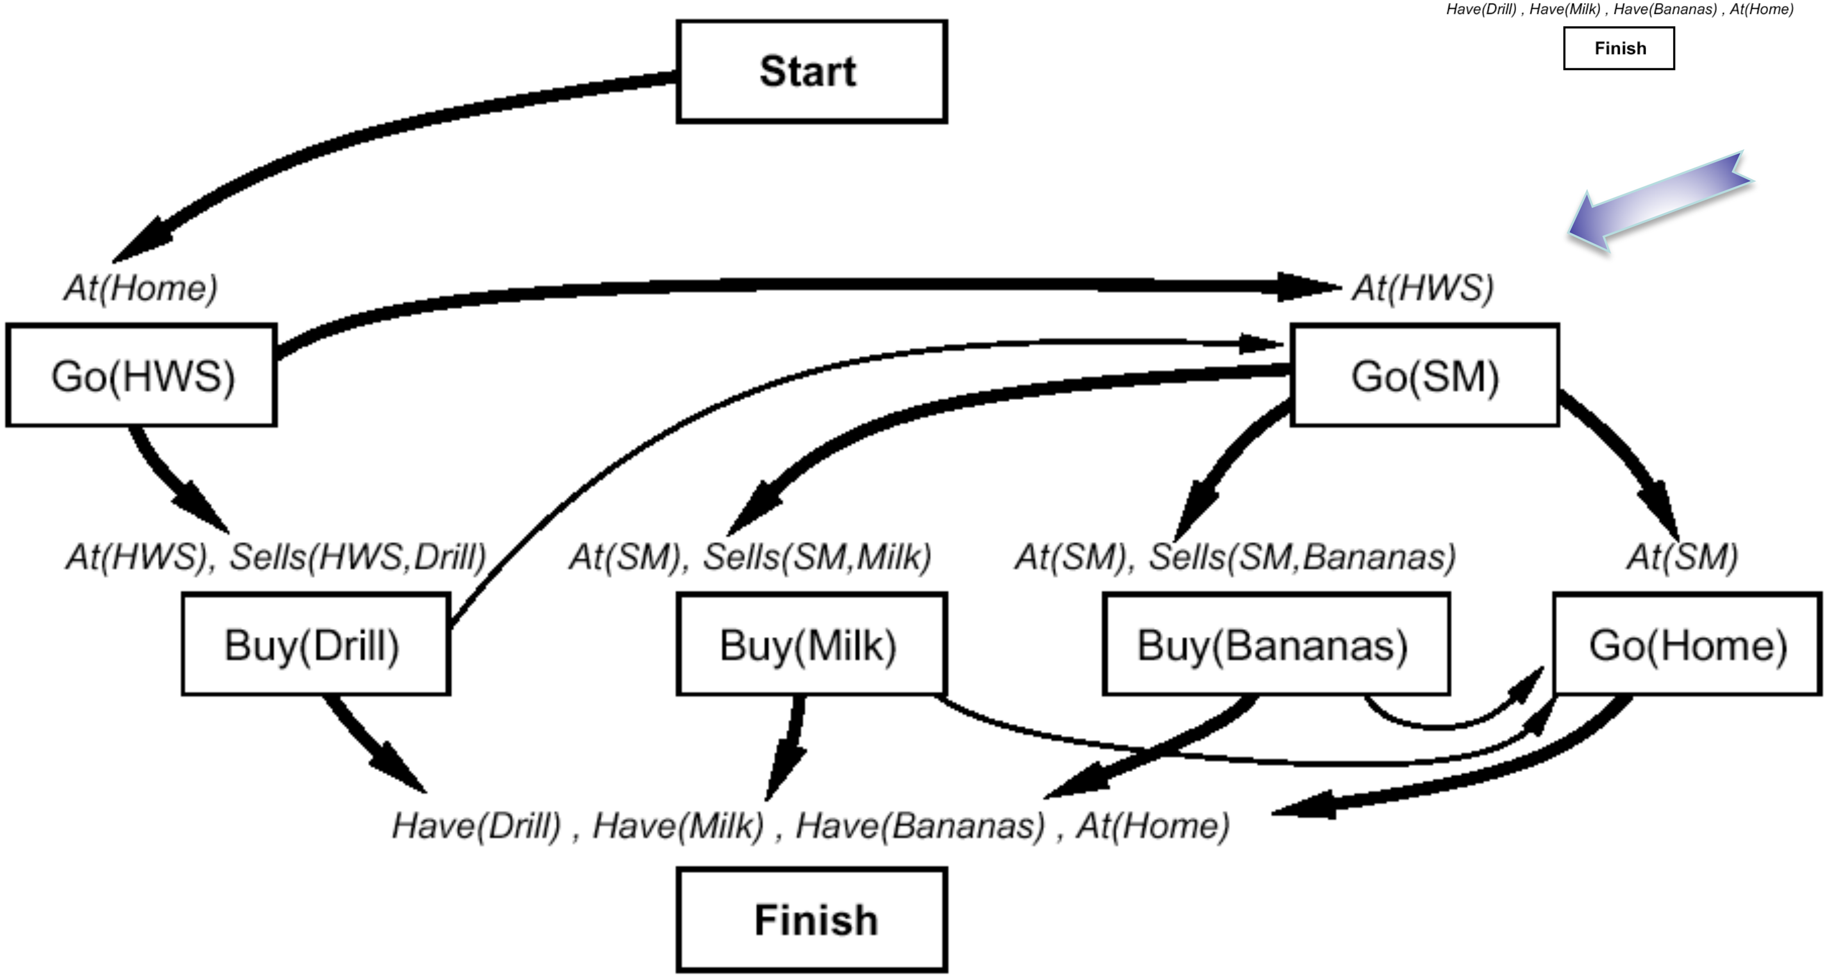

The Complete Solution …

- … with all links
- Computation by so-called POP Algorithm
  - Complete
  - … and correct
- Additionally, not considered here, correct treatment of variables

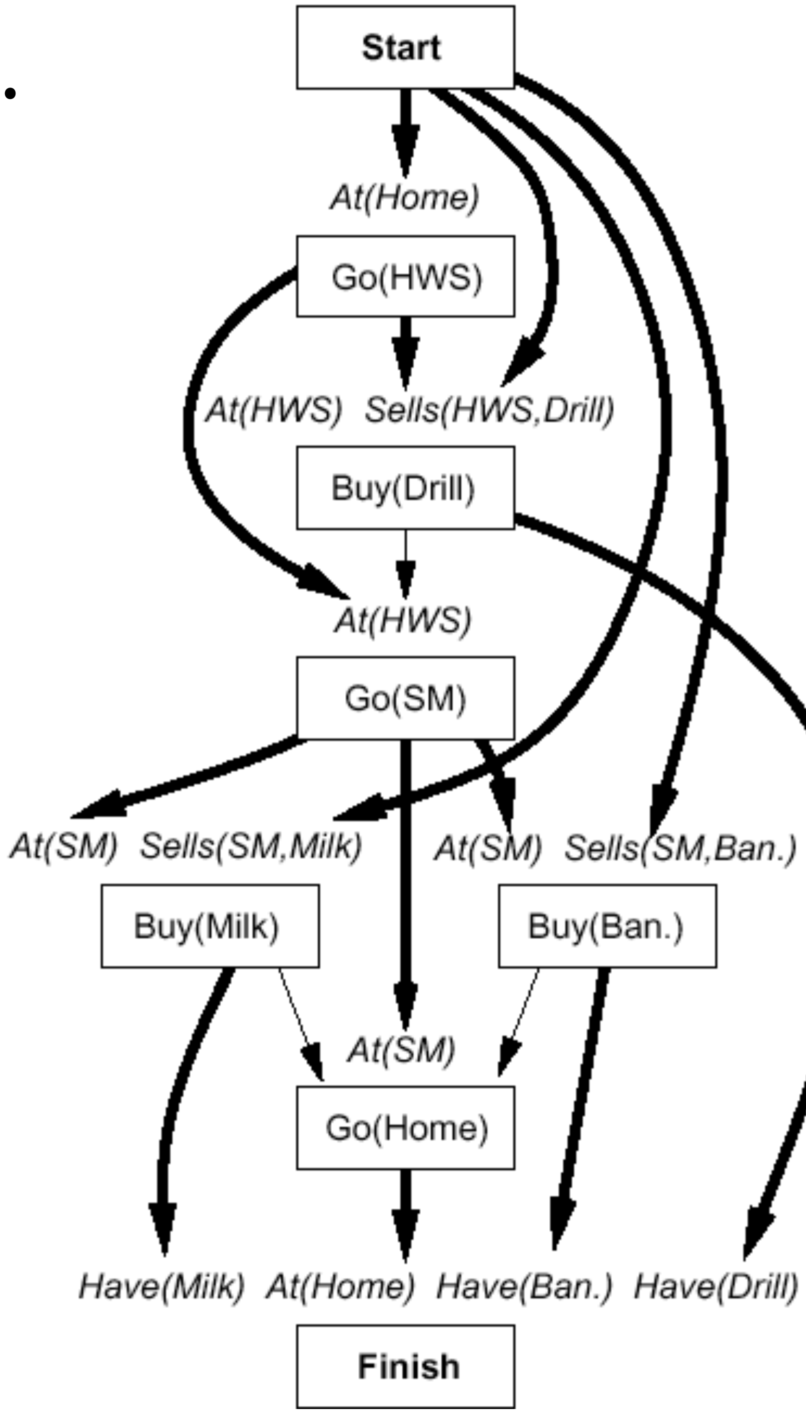

Last Century Planning Systems (Last Decade!)

- UCPOP (Weld, UW)
  (http://www.cs.washington.edu/ai/ucpop.html)
- Sensory Graphplan (Weld, Blum, and Furst: UW)
  (http://aiweb.cs.washington.edu/ai/sgp.html)
- IPP (Köhler and Nebel: Univ. Freiburg)
  (https://idw-online.de/de/news5468)
- Prodigy: Planning and Learning (Veloso: CMU)
  (http://www-2.cs.cmu.edu/afs/cs.cmu.edu/project/prodigy/Web/prodigy-home.html)

- All systems have found interesting applications

Planning is an Active Field of Research

- More powerful successors
  - Systems learn how to plan fast for specific problem instances
  - Can deal with uncertainty
    - About state estimation
    - About effects of actions
- Very powerful problem solvers can be set up …
  - w/ less effort/knowledge than with mathematical optimization theory and respective tools

→ Lecture **Agent Planning, Reinforcement Learning, Relations, and Human-Awarenes** of this module

(→ Automated Planning and Acting (CS5072-KP04)

(https://www.ifis.uni-luebeck.de/index.php?id=dski-aktuell-ss20&L=2))

Round-Up

- Rational agents that:
  - Act autonomously and are persistent;
    nonetheless can be influenced via mechanism design
  - Achieve goals surprisingly fast
    (despite bounded rationality)
  - Learn how to behave in a clever way
    (even learn computational strategies)
- Can adapt their goals
  to anticipate humans needs and expectations
  - Human compatibility, human awareness
- Can learn new models online to
  - Keep high performance over time
  - Support human-guided machine learning