What Is Artificial Intelligence?

An Overview of Artificial Intelligence

Douglas H. Fisher

What Is Artificial Intelligence?

- "The study of the computations that make it possible to perceive, reason, and act" (p. 5, Winston, Patrick, (1993), *Artificial Intelligence*, 3rd Edition, Addison-Wesley, Retrieved from https://courses.csail.mit.edu/6.034f/ai3/rest.pdf)
- "Artificial intelligence (AI) may be defined as the branch of computer science that is concerned with the automation of intelligent behavior" (p. 1, Luger, George (2009), *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*, 6th Edition, Pearson Education, Inc. Retrieved from http://www.cs.fsu.edu/~cap5605/Luger_6th_ed.pdf)
- "Artificial intelligence, or AI, is the field that studies the synthesis and analysis of computational agents that act intelligently" (Poole, David, & Mackworth, Alan (2017), *Artificial Intelligence: Foundations of Computational Agents*, 2nd Edition, Cambridge University Press. Retrieved from https://artint.info/2e/html/ArtInt2e.Ch1.S1.html (cited page))
- "AI is 'the study and construction of agents that do the right thing,' where 'the right thing is defined by the objective that we provide to the agent'" (Russell, Stuart, & Norvig, Peter (2021), *Artificial Intelligence: A Modern Approach*, 4th Edition, Pearson Education, Inc.)

What Is Artificial Intelligence?

- The science and engineering of exploration, evaluation, and selection among alternatives by machine
- The science and engineering of approximating non-deterministic, intractable algorithms with "efficient" heuristic approaches

What Is Artificial Intelligence?

- The science and engineering of exploration, evaluation, and selection among alternatives by machine
- The science and engineering of approximating non-deterministic, intractable algorithms with "efficient" heuristic approaches
  - For example, a traditional characterization of an algorithm might look something like this:
    - Intractable algorithm X finds an optimal solution to some problem with worst case runtime performance of $\Omega(c^N)$ or $\Omega(N!)$
  - A formal characterization of an AI algorithm might look something like this:
    - Algorithm X' finds a solution within epsilon of the optimal solution 95% of the time, with average (or worst) case runtime performance of $\Omega(N^c)$

Examples of AI and Non-AI Algorithms

- Summing the individual assets in a stock portfolio to obtain the portfolio's worth is deterministic, no exploration, and is **not AI**.
- Designing a stock portfolio for a person under their constraints and preferences probably does explore and evaluate alternative designs, and so is **probably AI**.
- Precisely following, by machine or human, the best route from Nashville to Charlotte returned by a map directions application at the level that corresponds to the written directions is deterministic, involves no exploration, and is **not AI**.
- Following the best route… at a micro level is **probably AI**, with exploration of micro-alternatives such as braking, acceleration, lane changes, and interpreting the driving assistant's poorly pronounced street names.
- Finding the best route at any level of granularity **is AI** when alternative routes explored.
- One-shot use of a recommender system is probably **not AI, or is trivial AI**, if the recommendation results from a fully parameterized, deterministic process.
- In contrast, learning such a recommender system is **probably AI**, if alternative designs and parameterizations were explored to best fit user preferences.

What Is Artificial Intelligence?
From Narrow to General Artificial Intelligence

Narrow (aka task-specific, weak) AI		**General (aka autonomous, strong) AI**
	Integrative intelligence	
Music recommender system	Home assistant (Smart home)	Data (*Star Trek: The Next Generation*)
Retail facial identification	Intelligent vehicle	Exocomps (*Star Trek: The Next Generation*, episode 135, “The Quality of Life”)
Airport patrol scheduling		Elvex (Asimov, *Robot Dreams*)
Story telling		
Computing protein structures		

→ Increasing functional generality, variety, and integration

Why Study Artificial Intelligence?

An Overview of Artificial Intelligence

Douglas H. Fisher

Why Study Artificial Intelligence? Part I

- **Because it’s the hottest thing since sliced bread**
- Exposure and mastery of new computational tools—dealing with nondeterminism
- Societal implications, pro and con, including business, government, and literature
- Thinking about thinking (its own slides)

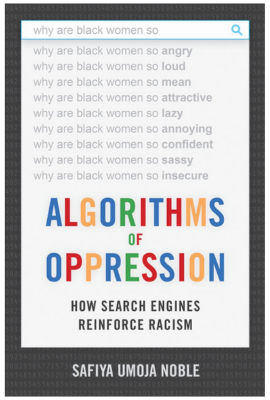

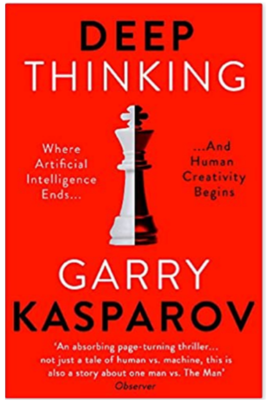

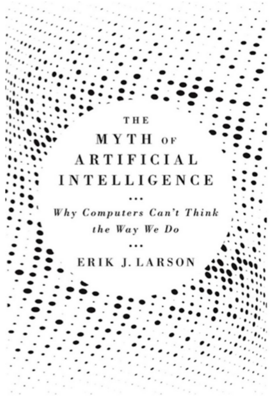

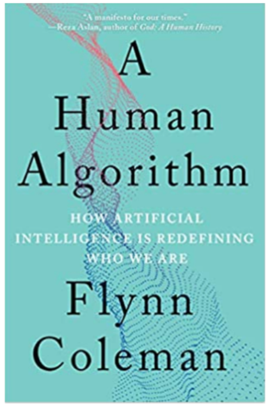

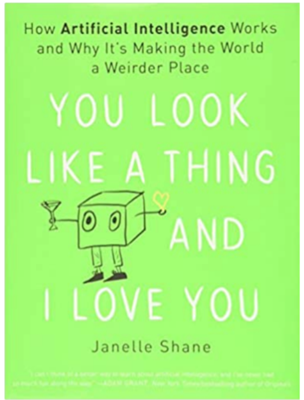

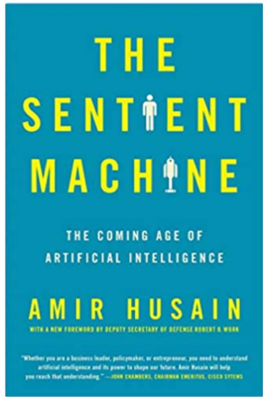

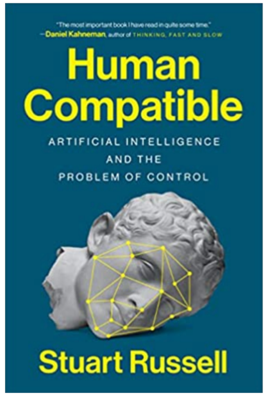

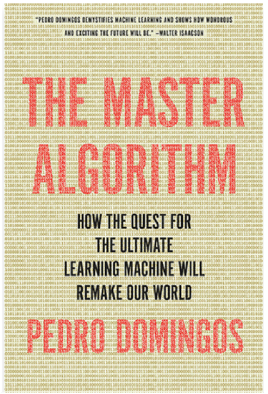

Douglas H. Fisher

Why Study Artificial Intelligence? Part II

- **Because it's the hottest thing since sliced bread**
- Exposure and mastery of new computational tools—dealing with nondeterminism
- Societal implications, pro and con (its own slides), including business, government, and literature
- Thinking about thinking (its own slides)

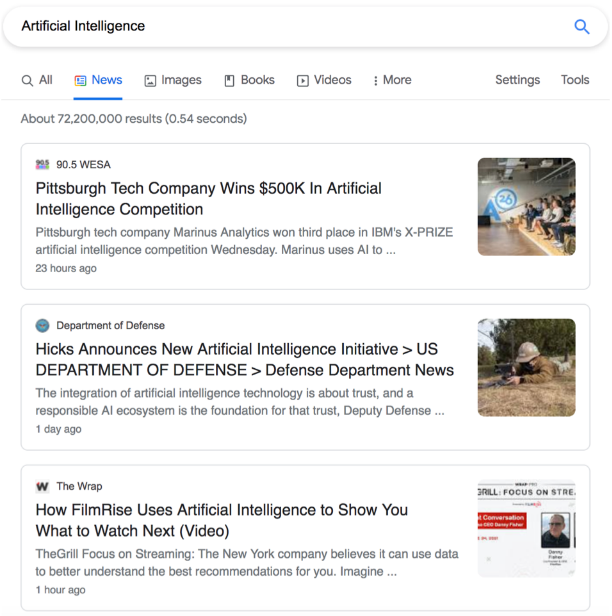

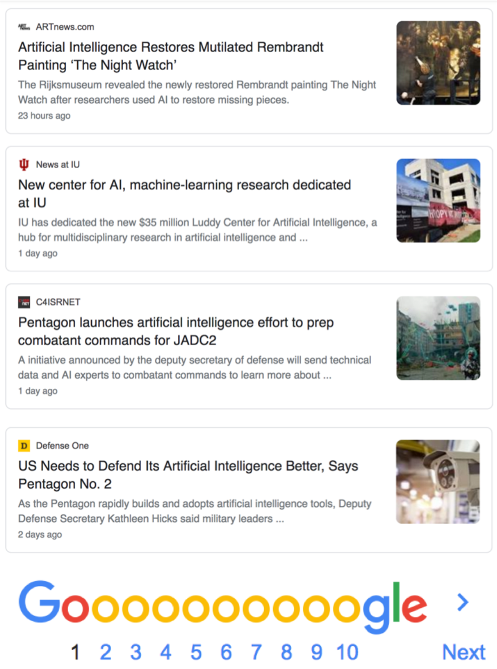

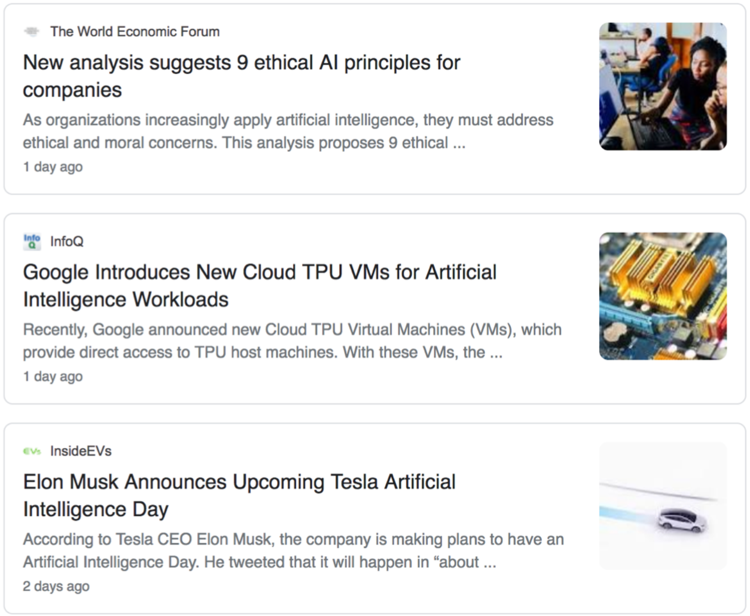

Douglas H. Fisher

Why Study Artificial Intelligence? Part III

- Because it's the hottest thing since sliced bread
- **Exposure and mastery of new computational tools—dealing with nondeterminism**
- Societal implications, pro and con, including business, government, and literature
- Thinking about thinking (its own slides)

- Agent architectures
- Uniformed and informed search
- Adversarial games
- Constraint-based reasoning
- Knowledge representation
- Deductive inference
- Planning without and with certainty
- Utility-driven decision-making
- Machine learning
- Natural language processing
- Computer vision
- Robotics

ABET Criteria

Students should have:

a. An ability to apply knowledge of mathematics, science, and engineering
b. An ability to design and conduct experiments, as well as to analyze and interpret data
c. **An ability to design a system, component, or process to meet desired needs within realistic constraints such as economic, environmental, social, political, ethical, health and safety, manufacturability, and sustainability**
d. An ability to function on multidisciplinary teams
e. An ability to identify, formulate, and solve engineering problems
f. An understanding of professional and ethical responsibility
g. An ability to communicate effectively
h. The broad education necessary to understand the impact of engineering solutions in a global, economic, environmental, and societal context
i. A recognition of the need for, and an ability to engage in life-long learning
j. A knowledge of contemporary issues
k. An ability to use the techniques, skills, and modern engineering tools necessary for engineering practice.

https://www.abet.org/accreditation/accreditation-criteria/criteria-for-accrediting-engineering-programs-2018-2019/

Why Study Artificial Intelligence?

- Because it's the hottest thing since sliced bread
- Exposure and mastery of new computational tools—dealing with nondeterminism
- **Societal implications, pro and con, including business, government, and literature**
- Thinking about thinking (its own slides)

- The Ethics and Religious Liberty Commission of the Southern Baptist Convention
- The Algorithmic Justice League

IEEE Code of Conduct

1. To hold paramount the safety, health, and welfare of the public, to strive to comply with ethical design and sustainable development practices, to protect the privacy of others, and to disclose promptly factors that might endanger the public or the environment;
2. **To improve the understanding by individuals and society of the capabilities and societal implications of conventional and emerging technologies, including intelligent systems;**
3. To avoid real or perceived conflicts of interest whenever possible, and to disclose them to affected parties when they do exist;
4. To avoid unlawful conduct in professional activities, and to reject bribery in all its forms;
5. To seek, accept, and offer honest criticism of technical work, to acknowledge and correct errors, to be honest and realistic in stating claims or estimates based on available data, and to credit properly the contributions of others;
6. **To maintain and improve our technical competence** and to undertake technological tasks for others only if qualified by training or experience, or after full disclosure of pertinent limitations;
7. To treat all persons fairly and with respect, and to not engage in discrimination based on characteristics such as race, religion, gender, disability, age, national origin, sexual orientation, gender identity, or gender expression;
8. To not engage in harassment of any kind, including sexual harassment or bullying behavior;
9. To avoid injuring others, their property, reputation, or employment by false or malicious actions, rumors or any other verbal or physical abuses;
10. To support colleagues and co-workers in following this code of ethics, to strive to ensure the code is upheld, and to not retaliate against individuals reporting a violation.

https://www.ieee.org/about/corporate/governance/p7-8.html

Thinking About Thinking

An Overview of Artificial Intelligence

Douglas H. Fisher

Why Study Artificial Intelligence?

- Because it's the hottest thing since sliced bread
- Exposure and mastery of new computational tools—dealing with nondeterminism
- Societal implications, pro and con (its own slides), including business, government, and literature
- **Thinking about thinking**
  - Computational thinking
  - Computing as a model of thought
  - AI will give you additional concepts on which to contemplate thought
  - Throughout the course, think about your thinking
    - Including the "mundane," like the examples to follow
    - AI may require that you formalize and implement your intuitions and insights

Why Study Artificial Intelligence?
Thinking About Thinking: Reminding Example

- In 2009, while working at the National Science Foundation (NSF), I ran into a colleague at the supermarket who had joined NSF about the same time that I had, but we had not seen each other since the initial orientation. We talked for a while about each staying a third year. I asked her what directorate at NSF she was in, and she said, “Biology, in the Division of Ecological Sciences,” and I told her that I would send her a link to a conference on societal and environmental sustainability, which included a lot of computer scientists who were addressing biodiversity issues.

Why Study Artificial Intelligence?

Thinking About Thinking: Reminding Example (cont.)

- The next day, in the early afternoon, having completely forgotten about my talk with her the previous evening, the elevator stopped on the sixth floor and I looked out and saw the following text (the blue square represents the elevator door and interior):

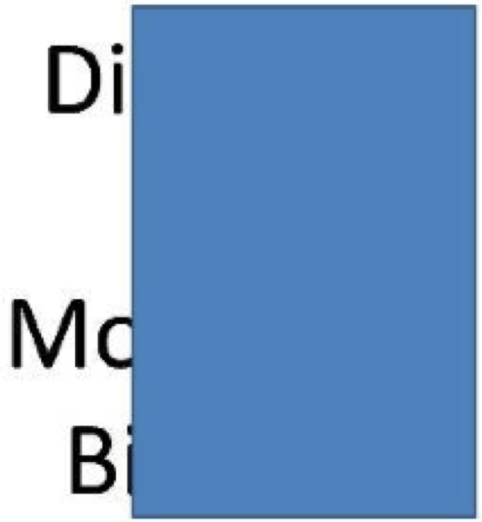

- I was reminded to email my colleague. Why might this have triggered a reminding? In your explanation, think about the assumptions you are making and why they are plausible, even probable. How could you confirm (or not) assumptions if you were pushed?

Why Study Artificial Intelligence?

Thinking About Thinking: Road Trip Example

- During the driving tour of the Midwest in July 2011, I took thousands of pictures with an old, non-GPS digital camera. After the trip, I posted some 750 of them. Without GPS coordinates, I used a variety of modalities available on Google—maps, satellite view, street view—to locate them. In locations with sufficient resolution I could place the photo right on the spot I was standing when I took the picture.
- Since I was placing the pictures a couple of weeks after the trip, there were different heuristics I used to place them—sometimes it was straightforward—a particular highway junction, or something otherwise named like a school or a cemetery or a mountain peak on Google maps. The order in which pictures were taken offered some constraints since having located one picture narrowed the possible locations for the next. All of these pieces of information are things that I would want an AI system, designed to do like activity, to represent and use.

Why Study Artificial Intelligence?
Thinking About Thinking: Road Trip Example

- In one case though, even with some known restricted area stemming from sequencing information, I was trying to locate a picture in the tiny town of Scribner, Nebraska, but I saw no way to identify the precise location of a building with a steeple.

Why Study Artificial Intelligence?

Thinking About Thinking: Road Trip Example, Part III

Can you identify where the building to the right is located?

Douglas H. Fisher

Why Study Artificial Intelligence?
Thinking About Thinking

As you introspect on your thoughts and actions, ask what information/knowledge might have been relevant.

Agents acting in an environment: inputs and output

What goals were you pursuing?

What environmental cues might have triggered access?

What prior knowledge was brought to bear?

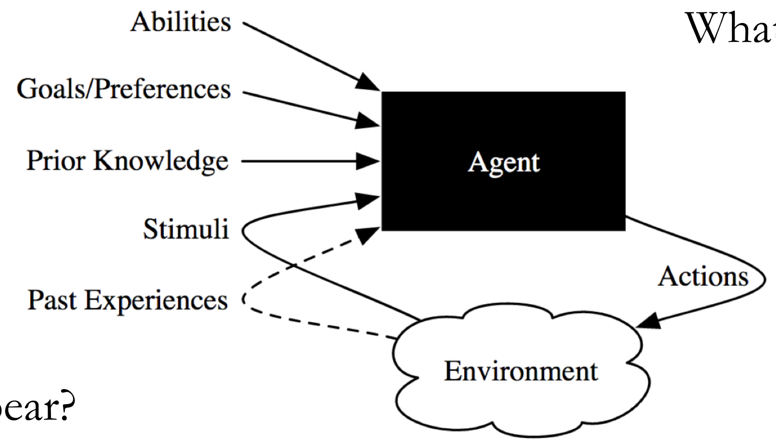

What processing might have occurred?

©D.L. Poole and A.K. Mackworth 2010-2020 Artificial Intelligence, Lecture 1.1 5 / 15

Douglas H. Fisher

Tools, Agents, and Environments

An Overview of Artificial Intelligence

Douglas H. Fisher

Agent-Based Perspectives on AI

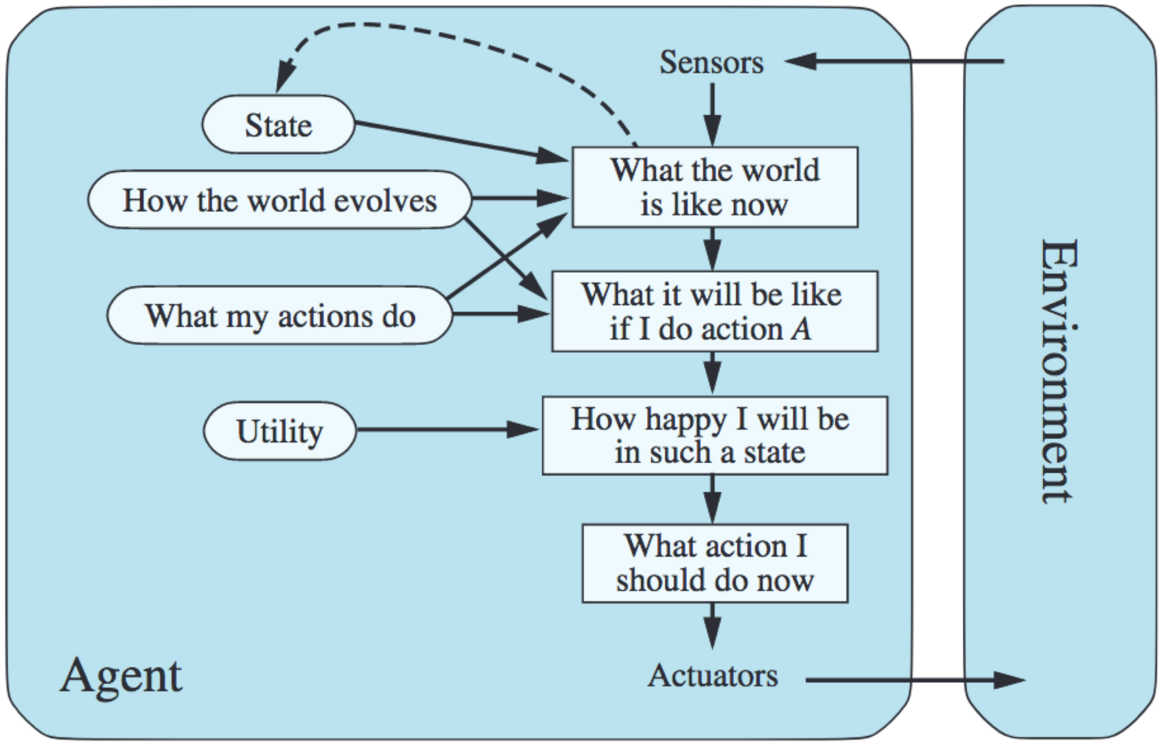

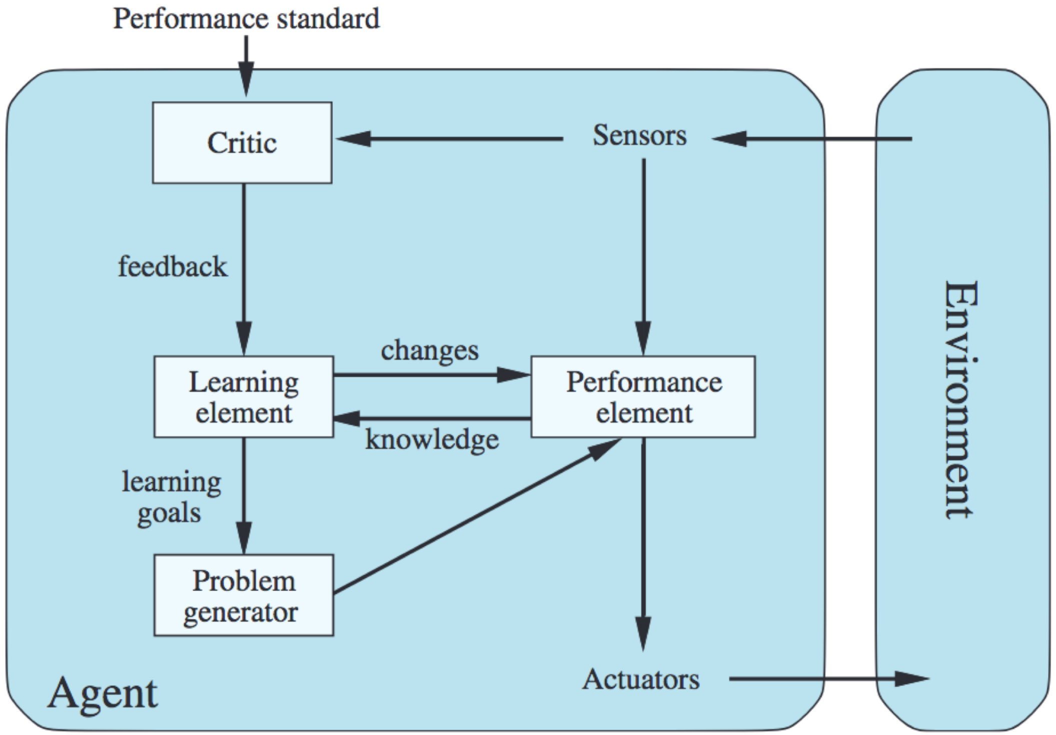

Figures 2.14 and 2.15 Utility Agent and Learning Agent from Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach*. Retrieved from http://aima.cs.berkeley.edu/figures.pdf (slides 8 and 9, respectively). Used with permission.

Agent-Based Perspectives on AI (cont.)

A Decomposition of the Delivery Robot

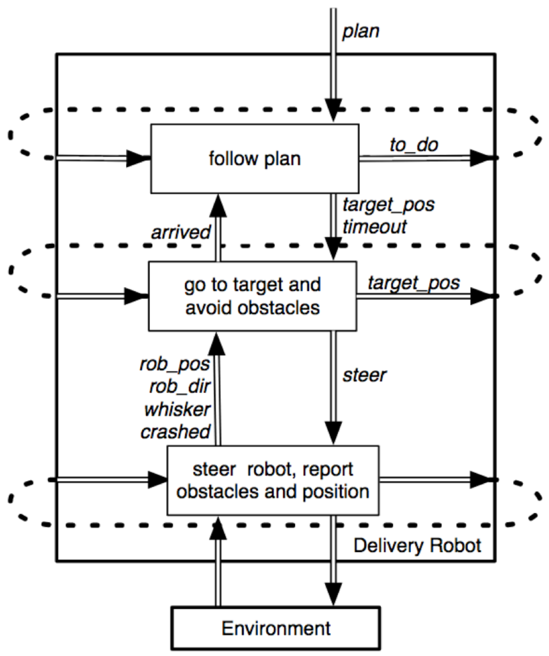

©D. Poole and A. Mackworth 2019 Artificial Intelligence, Lecture 2.2 7 / 14

Douglas H. Fisher

Cognitive Architectures for Integrative Intelligence

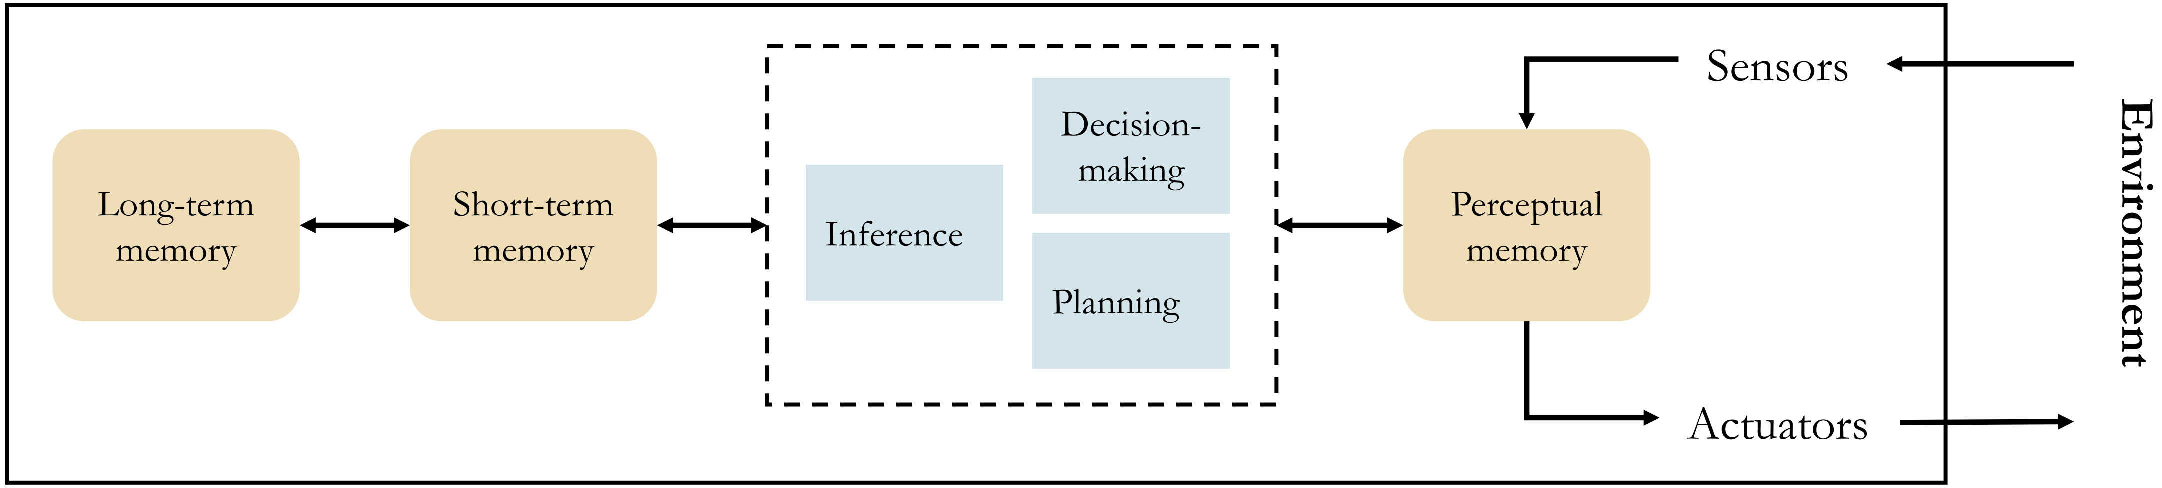

Various forms of memory

- Concept memory
- Scene memory
- Procedural memory
- Skill memory
- Episodic memory
- Sequential memory
- Relational memory

Douglas H. Fisher

Agent Types

Agent Type	Performance Measure	Environment	Actuators	Sensors
Medical diagnosis system	Healthy patient, reduced costs	Patient, hospital, staff	Display of questions, tests, diagnoses, treatments	Touchscreen/voice entry of symptoms and findings
Satellite image analysis system	Correct categorization of objects, terrain	Orbiting satellite, downlink, weather	Display of scene categorization	High-resolution digital camera
Part-picking robot	Percentage of parts in correct bins	Conveyor belt with parts; bins	Jointed arm and hand	Camera, tactile and joint angle sensors
Refinery controller	Purity, yield, safety	Refinery, raw materials, operators	Valves, pumps, heaters, stirrers, displays	Temperature, pressure, flow, chemical sensors
Interactive English tutor	Student's score on test	Set of students, testing agency	Display of exercises, feedback, speech	Keyboard entry, voice

Figures 2.5 Examples of agent types and their PEAS description from Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach*. Retrieved from http://aima.cs.berkeley.edu/figures.pdf (slide 5). Used with permission.

Task Environments

Task Environment	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword puzzle	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with a clock	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Backgammon	Fully	Multi	Stochastic	Sequential	Static	Discrete
Taxi driving	Partially	Multi	Stochastic	Sequential	Dynamic	Continuous
Medical diagnosis	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Image analysis	Fully	Single	Deterministic	Episodic	Semi	Continuous
Part-picking robot	Partially	Single	Stochastic	Episodic	Dynamic	Continuous
Refinery controller	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
English tutor	Partially	Multi	Stochastic	Sequential	Dynamic	Discrete

Figure 2.6 Examples of task environments and their characteristics.

Figure 2.6 Examples of task environments and their characteristics from Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach*. Retrieved from http://aima.cs.berkeley.edu/figures.pdf (slide 5). Used with permission.

Capabilities of Cognitive Architectures

- Recognition and categorization
  - Represent patterns and situations in memory
  - Learn these patterns
- Decision-making and choice (one step plans?)
  - Allowable alternatives
  - Desirability of alternatives
    - Goals, objectives, and utilities
  - Learning allowability/desirability/effectiveness
- Perception and situation assessment
  - Compose large-scale environment models from percepts
  - Relies recognition and categorization of patterns in the environment
  - Relies on inferential mechanisms
- Prediction and monitoring
  - Model of the environment
  - Effects of actions
- Problem solving and planning
  - Goals, objective, and utilities
  - Partially ordered actions
  - Enabling conditions
  - Predicted effects
  - Learning to reduce effective breadth and depth of search
- Reasoning and belief maintenance
  - Deductive reasoning
  - Abductive reasoning
  - Inductive reasoning
  - Incremental or online learning
- Execution and action
  - Actuators in environment
  - Primitive actions
  - Composite actions
- Interaction and communication
  - Translating knowledge for other agents
  - Question asking and answering
- Remembering, reflection, and learning
  - Cognitive structures formed during external or cognitive activities
  - Explanation/justification
  - Meta-reasoning
- Learning is pervasive and, in human instantiations, perhaps emotional awareness and response is too

Langley, P., Laird, J., & Rogers, S. (2009). Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*, *10*, 141-160.

Representing the Environment

Often Requires a Combination of Atomic, Factored, and Structural Representations

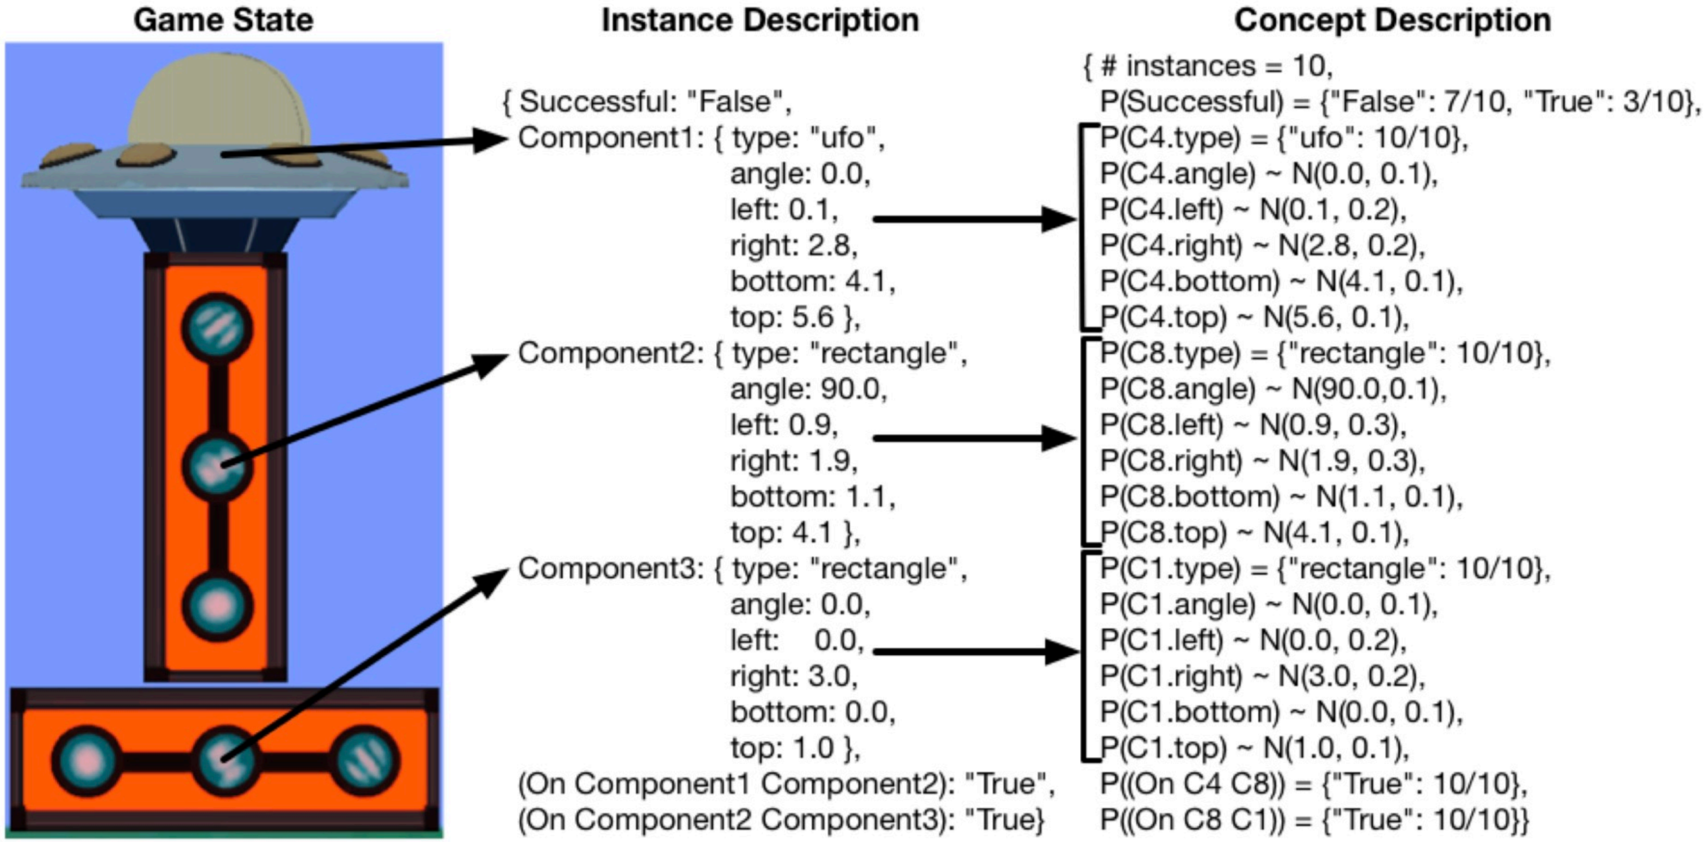

From MacLellan, C. J., Harpstead, E., Aleven, V., Koedinger, K. R. (2016). [*TRESTLE: A Model of Concept Formation in Structured Domains. Advances in Cognitive Systems*](http://www.cogsys.org/journal/volume4/article-4-10.pdf), *4*, 131–150. Retrieved from http://www.cogsys.org/journal/volume4/article-4-10.pdf (Used with permission)

Tools, Agents, and Environments

The End