

Artificial Intelligence


Artificial Intelligence and Intelligent Agents; Problem Solving

1131AI02

MBA, IM, NTPU (M5276) (Fall 2024)
Tue 2, 3, 4 (9:10-12:00) (B3F17)



<https://meet.google.com/paj-zhhj-mya>



Min-Yuh Day, Ph.D,
Professor

Institute of Information Management, National Taipei University

<https://web.ntpu.edu.tw/~myday>



Syllabus

Week Date Subject/Topics

1 2024/09/10 Introduction to Artificial Intelligence

2 2024/09/17 Mid-Autumn Festival (Day off)

3 2024/09/24 Artificial Intelligence and Intelligent Agents; Problem Solving

**4 2024/10/01 Knowledge, Reasoning and Knowledge Representation;
Uncertain Knowledge and Reasoning**

5 2024/10/08 Case Study on Artificial Intelligence I

6 2024/10/15 Machine Learning: Supervised and Unsupervised Learning

Syllabus

Week	Date	Subject/Topics
7	2024/10/22	The Theory of Learning and Ensemble Learning
8	2024/10/29	Midterm Project Report
9	2024/11/05	Self-Learning
10	2024/11/12	Deep Learning, Reinforcement Learning
11	2024/11/19	Case Study on Artificial Intelligence II
12	2024/11/26	Deep Learning for Natural Language Processing

Syllabus

Week Date Subject/Topics

13 2024/12/03 Computer Vision and Robotics

**14 2024/12/10 Generative AI,
Philosophy and Ethics of AI and the Future of AI**

15 2024/12/17 Final Project Report I

16 2024/12/24 Final Project Report II

Artificial Intelligence

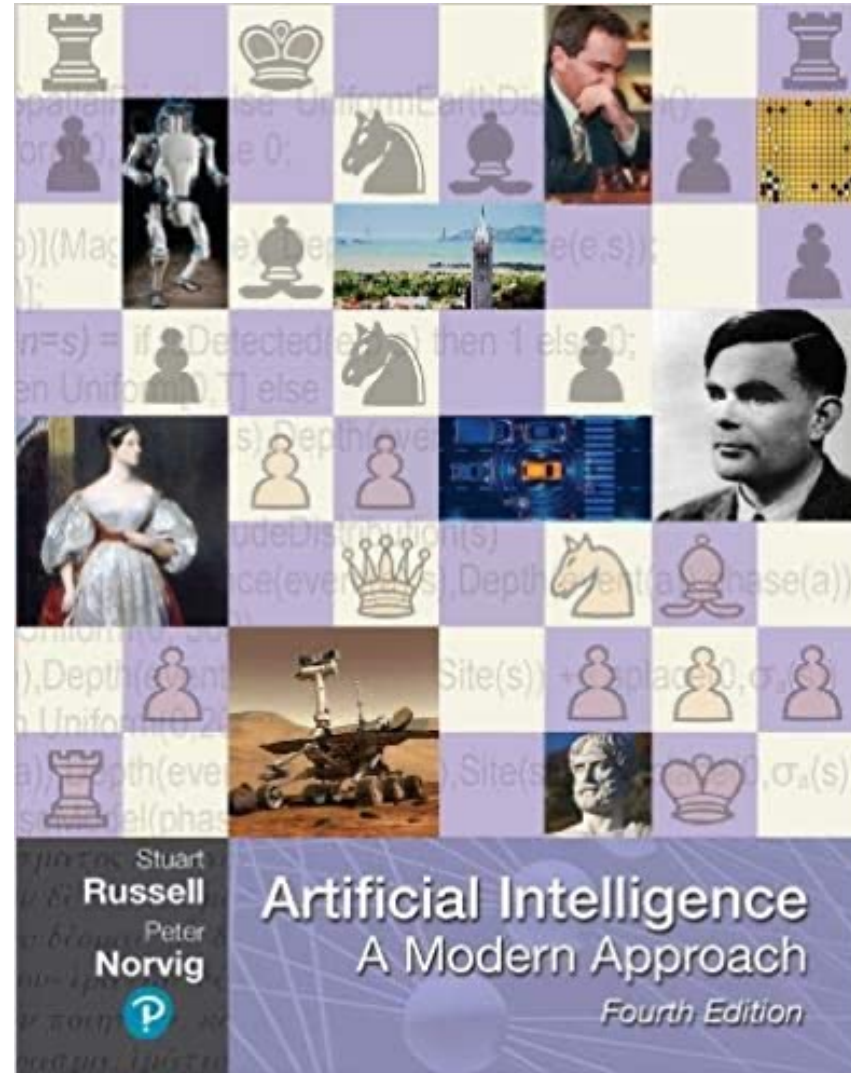
Intelligent Agents

Problem Solving

Outline

- **Artificial Intelligence**
- **Intelligent Agents**
- **Problem Solving**

Stuart Russell and Peter Norvig (2020),
Artificial Intelligence: A Modern Approach,
4th Edition, Pearson



Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

<https://www.amazon.com/Artificial-Intelligence-A-Modern-Approach/dp/0134610997/>

Artificial Intelligence: A Modern Approach

- 1. Artificial Intelligence**
- 2. Problem Solving**
- 3. Knowledge and Reasoning**
- 4. Uncertain Knowledge and Reasoning**
- 5. Machine Learning**
- 6. Communicating, Perceiving, and Acting**
- 7. Philosophy and Ethics of AI**

Artificial Intelligence: Intelligent Agents

Artificial Intelligence:

2. Problem Solving

- **Solving Problems by Searching**
- **Search in Complex Environments**
- **Adversarial Search and Games**
- **Constraint Satisfaction Problems**

Artificial Intelligence:

3. Knowledge and Reasoning

- **Logical Agents**
- **First-Order Logic**
- **Inference in First-Order Logic**
- **Knowledge Representation**
- **Automated Planning**

Artificial Intelligence:

4. Uncertain Knowledge and Reasoning

- **Quantifying Uncertainty**
- **Probabilistic Reasoning**
- **Probabilistic Reasoning over Time**
- **Probabilistic Programming**
- **Making Simple Decisions**
- **Making Complex Decisions**
- **Multiagent Decision Making**

Artificial Intelligence:

5. Machine Learning

- **Learning from Examples**
- **Learning Probabilistic Models**
- **Deep Learning**
- **Reinforcement Learning**

Artificial Intelligence:

6. Communicating, Perceiving, and Acting

- **Natural Language Processing**
- **Deep Learning for Natural Language Processing**
- **Computer Vision**
- **Robotics**


Artificial Intelligence:

Philosophy and Ethics of AI

The Future of AI

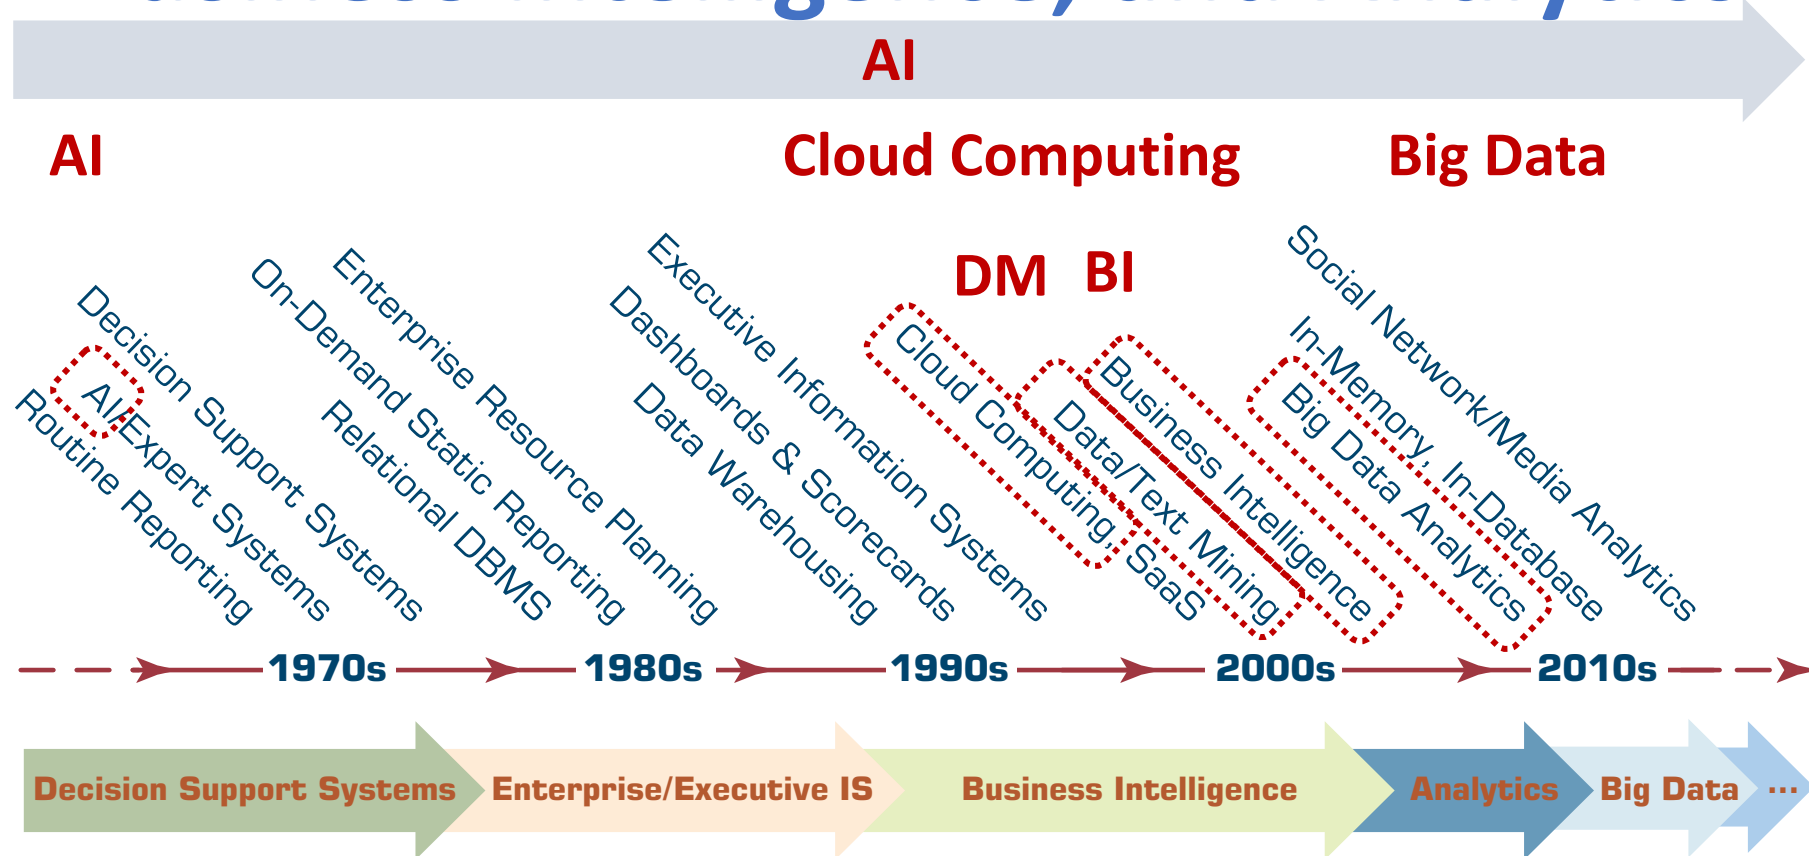
Artificial Intelligence (AI)

AI, ML, DL, Generative AI

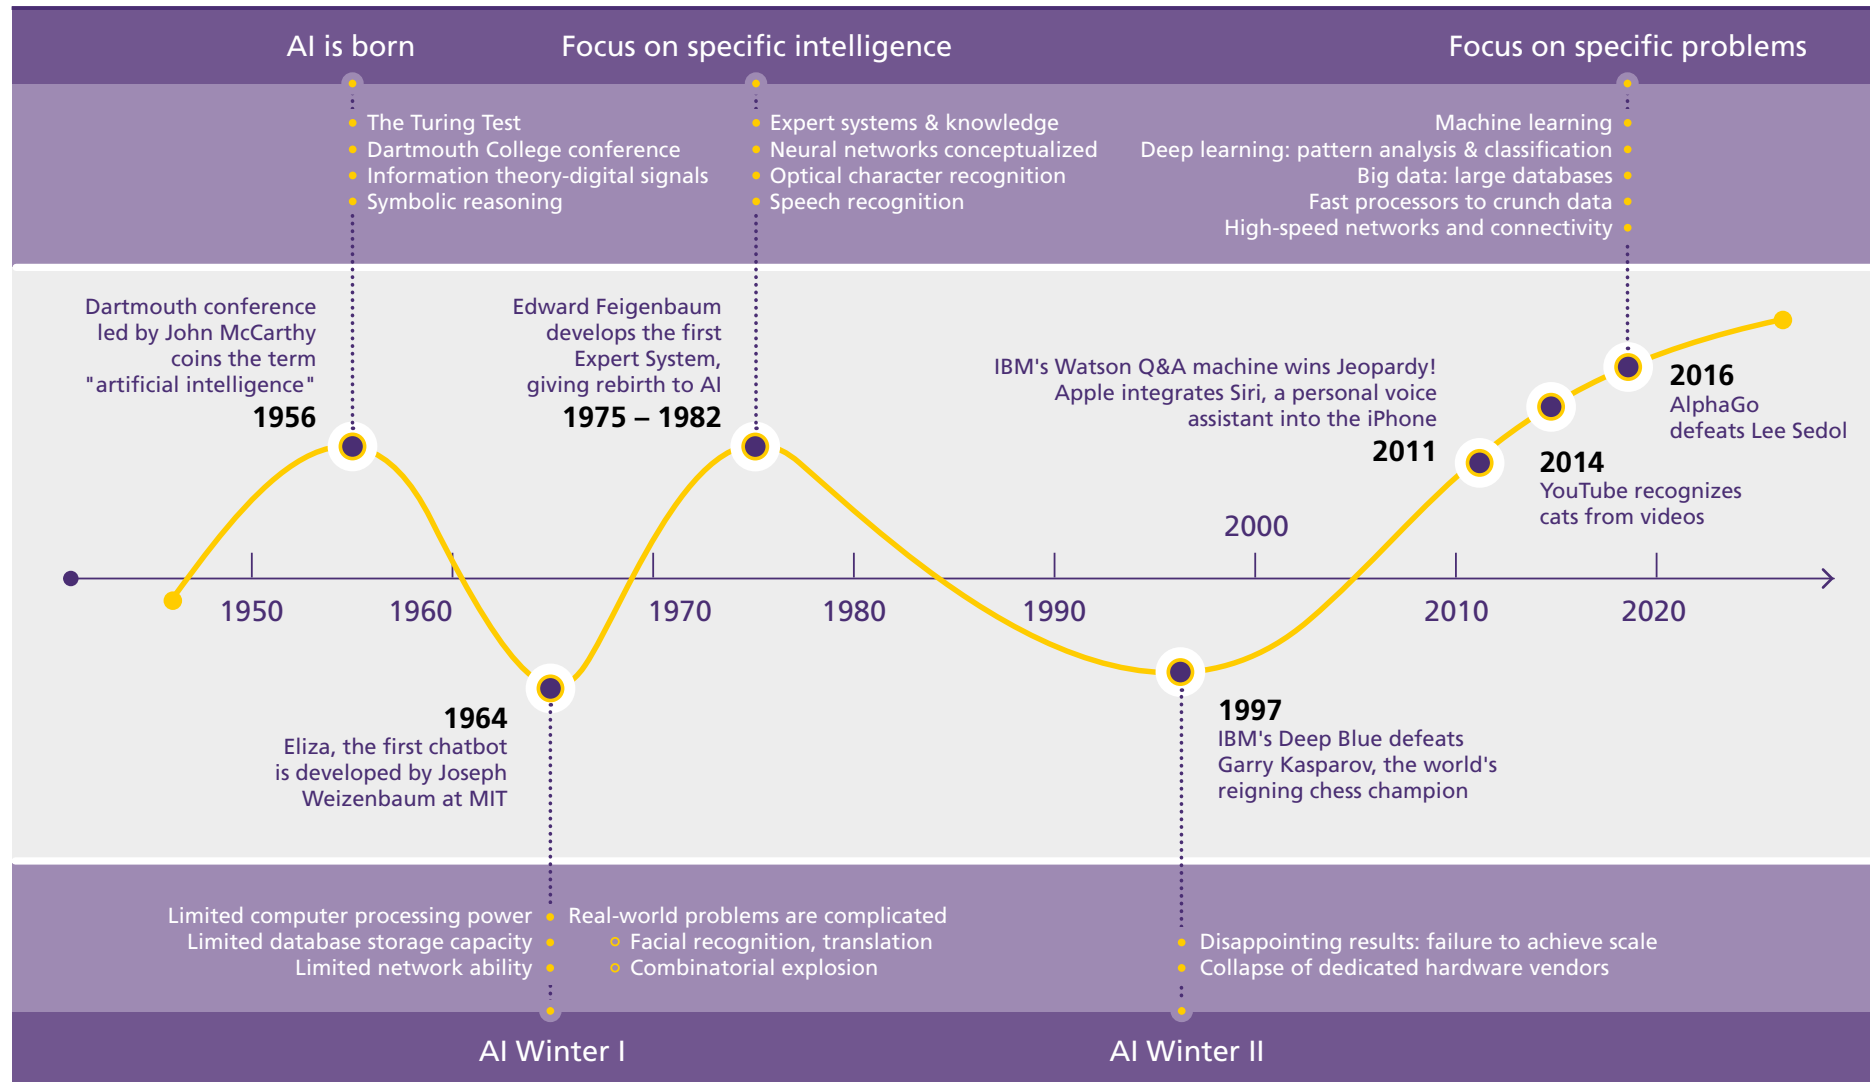


AI, Big Data, Cloud Computing

Evolution of Decision Support, Business Intelligence, and Analytics




The Rise of AI



The Development of LM-based Dialogue Systems

- 1) Early Stage (1966 - 2015)
- 2) The Independent Development of TOD and ODD (2015 - 2019)
- 3) Fusions of Dialogue Systems (2019 - 2022)
- 4) LLM-based DS (2022 - Now)



Task-oriented DS (TOD), Open-domain DS (ODD)

Definition of Artificial Intelligence (A.I.)

Artificial Intelligence

**“... the science and
engineering
of
making
intelligent machines”**

(John McCarthy, 1955)

Artificial Intelligence

**“... technology that
thinks and acts
like humans”**

Artificial Intelligence

**“... intelligence
exhibited by machines
or software”**

4 Approaches of AI

Thinking Humanly	Thinking Rationally
Acting Humanly	Acting Rationally

4 Approaches of AI

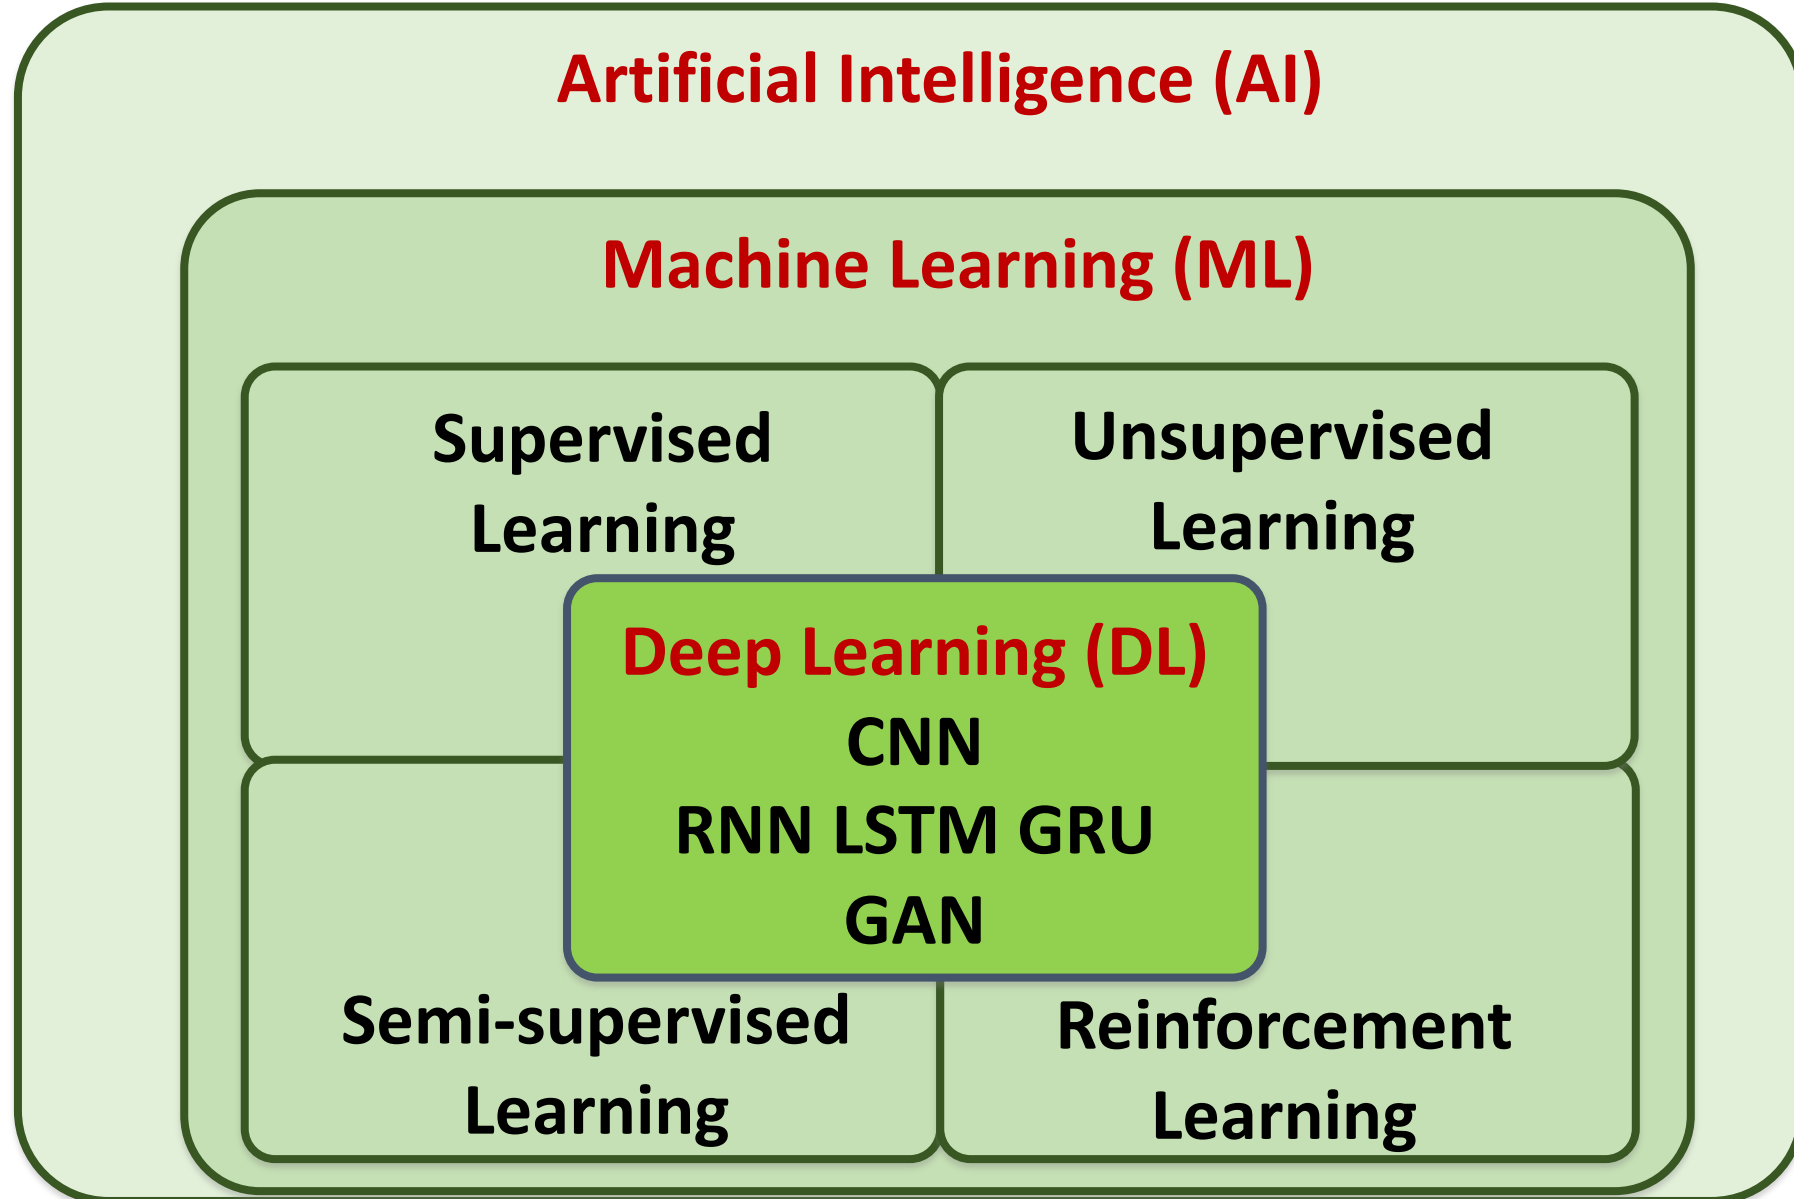
<p>2. Thinking Humanly: The Cognitive Modeling Approach</p>	<p>3. Thinking Rationally: The “Laws of Thought” Approach</p>
<p>1. Acting Humanly: The Turing Test Approach (1950)</p>	<p>4. Acting Rationally: The Rational Agent Approach</p>

AI Acting Humanly: The Turing Test Approach

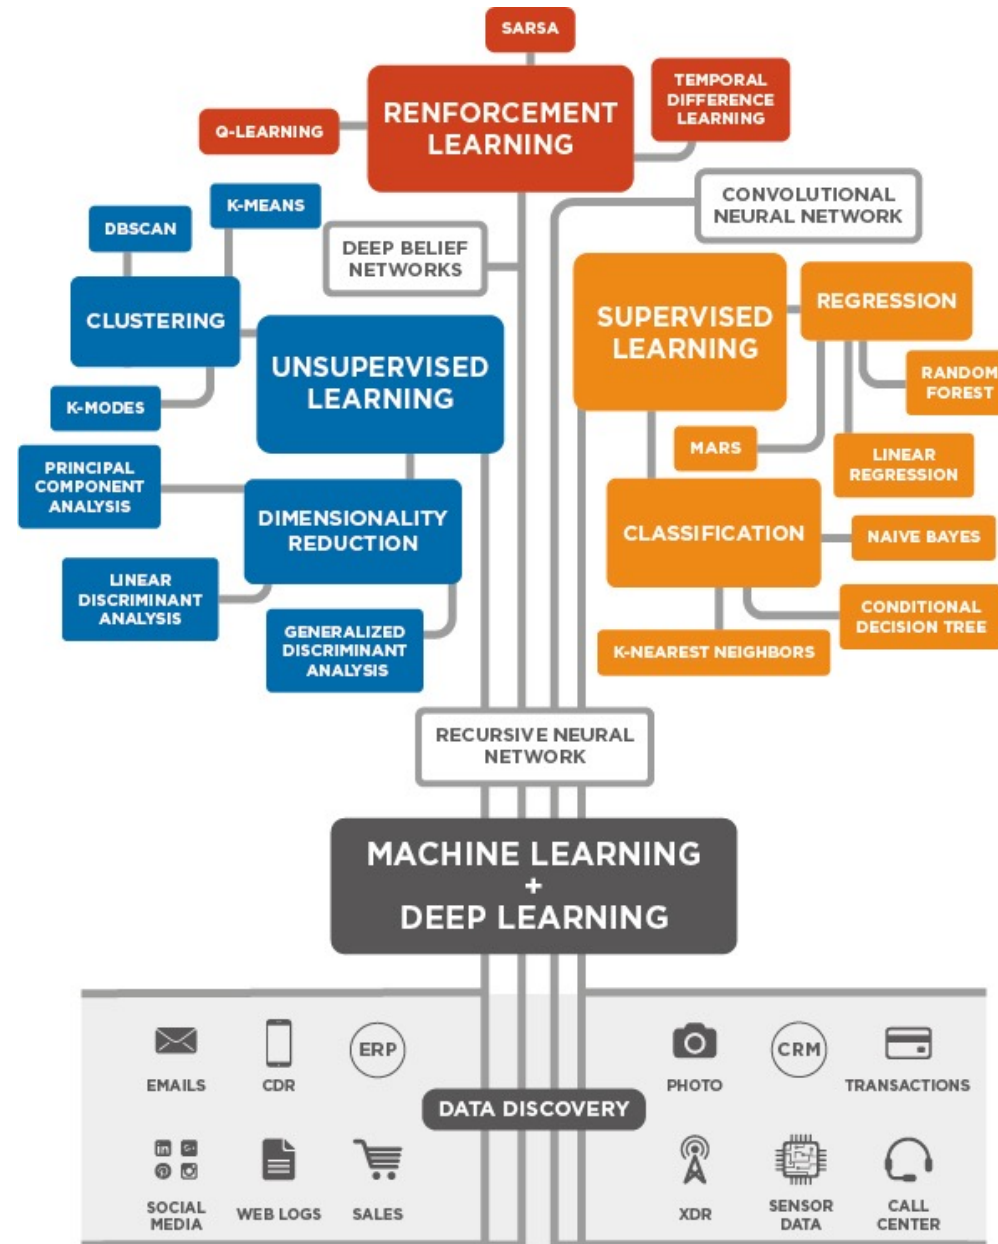
(Alan Turing, 1950)

- Knowledge Representation
- Automated Reasoning
- Machine Learning (ML)
 - Deep Learning (DL)
- Computer Vision (Image, Video)
- Natural Language Processing (NLP)
- Robotics


AI, ML, DL



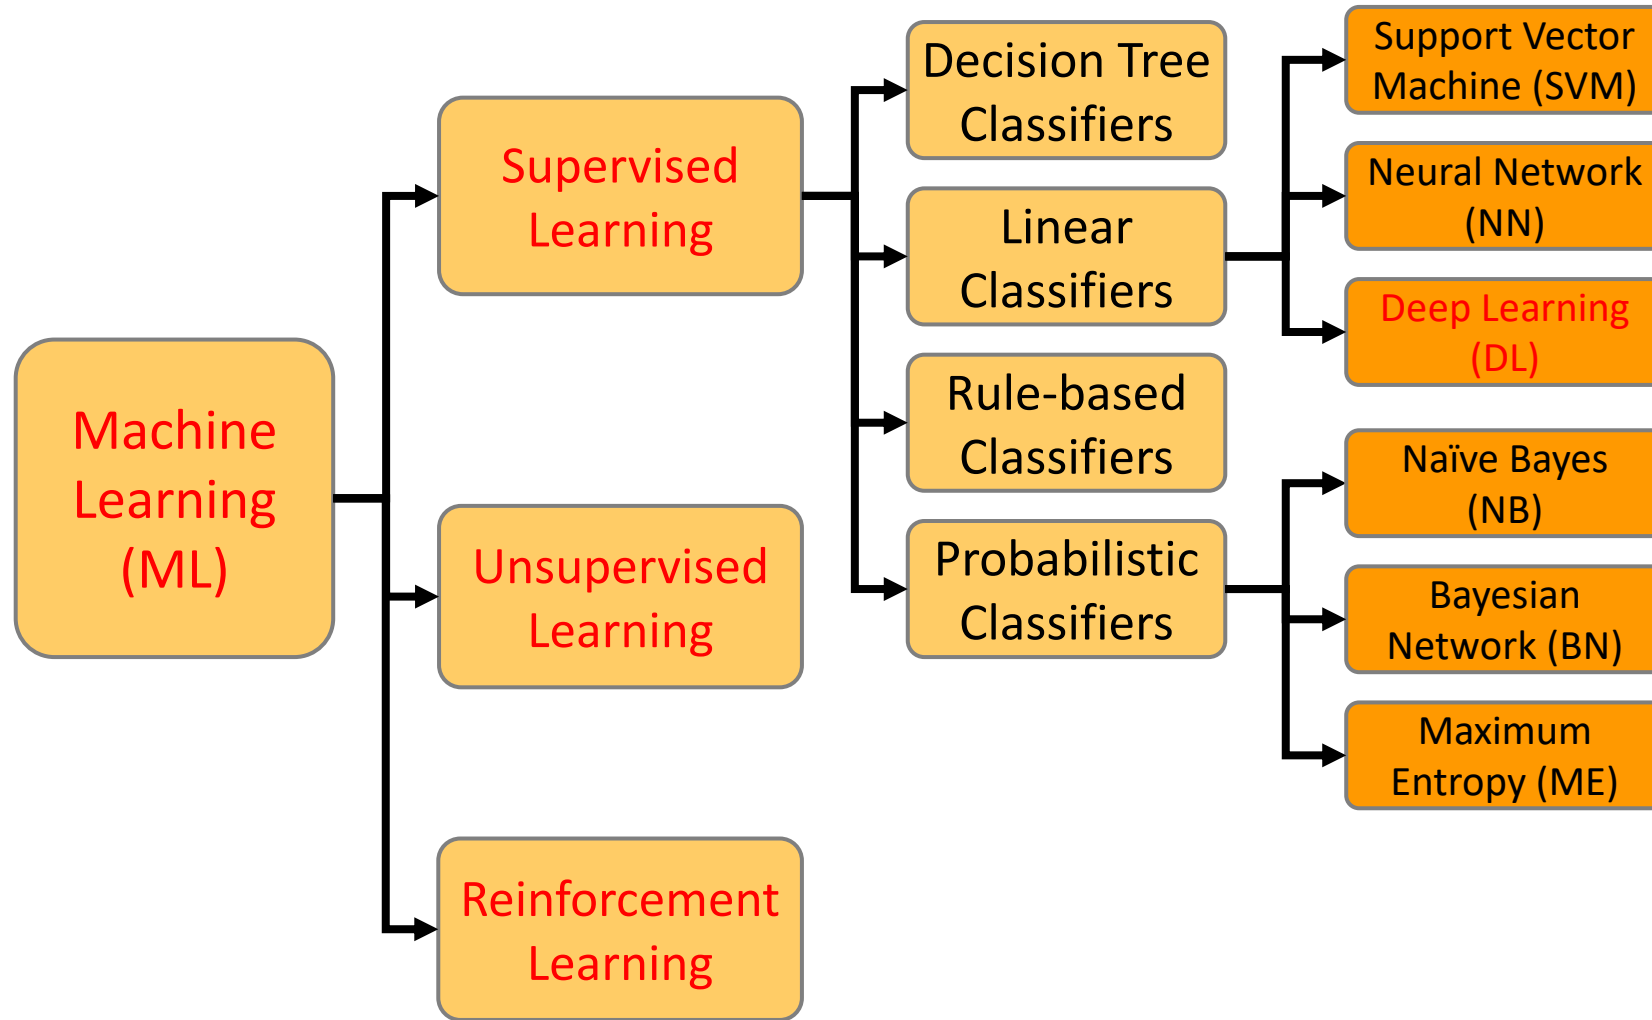
3 Machine Learning Algorithms




Machine Learning (ML)




Machine Learning (ML) / Deep Learning (DL)



Transformer Models



Large Language Models (LLMs)



Four Paradigms in NLP (LM)

Paradigm	Engineering	Task Relation
a. Fully Supervised Learning (Non-Neural Network)	Feature (e.g. word identity, part-of-speech, sentence length)	
b. Fully Supervised Learning (Neural Network)	Architecture (e.g. convolutional, recurrent, self-attentional)	
Transfer Learning: Pre-training, Fine-Tuning (FT)		
c. Pre-train, Fine-tune	Objective (e.g. masked language modeling, next sentence prediction)	
GAI: Pre-train, Prompt, and Predict (Prompting)		
d. Pre-train, Prompt, Predict	Prompt (e.g. cloze, prefix)	

Generative AI

**Text, Image, Video, Audio
Applications**

Comparison of Generative AI and Traditional AI

Feature	Generative AI	Traditional AI
Output type	New content	Classification/Prediction
Creativity	High	Low
Interactivity	Usually more natural	Limited

Generative AI

- **Generative AI: The Art of Creation**
- **Definition: AI systems capable of creating new content**
- **Characteristics: Creativity, interactivity**

Popular Generative AI

- **OpenAI ChatGPT (GPT-4o, GPT-4)**
- **Claude.ai (Claude 3.5)**
- **Google Gemini**
- **Chat.LMSys.org**
- **Perplexity.ai**
- **ChatPDF**
- **Stable Diffusion**
- **Video: D-ID, Synthesia**
- **Audio: Speechify**

OpenAI ChatGPT (GPT-4o, GPT-4)

ChatGPT 4 ▾


Model ⓘ


GPT-4o
Great for most tasks

o1-preview
Uses advanced reasoning


o1-mini
Faster at reasoning


More models GPT-4 >


 Temporary chat




Content calendar
for TikTok

 Quiz me on
world capitals

 Tell me the country
with the most
Olympic athletes

 Message ChatGPT



ChatGPT can make mistakes. Check important info.


<https://chatgpt.com/>

OpenAI ChatGPT (GPT-4) DALL·E 3

Vector graphic of a flowchart depicting the integration of generative AI in the education process, from content creation to virtual experiments, personalized learning, and innovative learning.



Claude 3.5 Sonnet



Claude 3.5, GPT-4o, Gemini 1.5 Pro

	Claude 3.5 Sonnet	Claude 3 Opus	GPT-4o	Gemini 1.5 Pro	Llama-400b (early snapshot)
Graduate level reasoning <i>GPQA, Diamond</i>	59.4%* 0-shot CoT	50.4% 0-shot CoT	53.6% 0-shot CoT	—	—
Undergraduate level knowledge <i>MMLU</i>	88.7%** 5-shot	86.8% 5-shot	—	85.9% 5-shot	86.1% 5-shot
	88.3% 0-shot CoT	85.7% 0-shot CoT	88.7% 0-shot CoT	—	—
Code <i>HumanEval</i>	92.0% 0-shot	84.9% 0-shot	90.2% 0-shot	84.1% 0-shot	84.1% 0-shot
Multilingual math <i>MGSM</i>	91.6% 0-shot CoT	90.7% 0-shot CoT	90.5% 0-shot CoT	87.5% 8-shot	—
Reasoning over text <i>DROP, F1 score</i>	87.1 3-shot	83.1 3-shot	83.4 3-shot	74.9 Variable shots	83.5 3-shot Pre-trained model
Mixed evaluations <i>BIG-Bench-Hard</i>	93.1% 3-shot CoT	86.8% 3-shot CoT	—	89.2% 3-shot CoT	85.3% 3-shot CoT Pre-trained model
Math problem-solving <i>MATH</i>	71.1% 0-shot CoT	60.1% 0-shot CoT	76.6% 0-shot CoT	67.7% 4-shot	57.8% 4-shot CoT
Grade school math <i>GSM8K</i>	96.4% 0-shot CoT	95.0% 0-shot CoT	—	90.8% 11-shot	94.1% 8-shot CoT

* Claude 3.5 Sonnet scores 67.2% on 5-shot CoT GPQA with maj@32

** Claude 3.5 Sonnet scores 90.4% on MMLU with 5-shot CoT prompting

Claude 3.5 Sonnet State-of-the-art vision

	Claude 3.5 Sonnet	Claude 3 Opus	GPT-4o	Gemini 1.5 Pro
Visual math reasoning <i>MathVista (testmini)</i>	67.7% 0-shot CoT	50.5% 0-shot CoT	63.8% 0-shot CoT	63.9% 0-shot CoT
Science diagrams <i>AI2D, test</i>	94.7% 0-shot	88.1% 0-shot	94.2% 0-shot	94.4% 0-shot
Visual question answering <i>MMMU (val)</i>	68.3% 0-shot CoT	59.4% 0-shot CoT	69.1% 0-shot CoT	62.2% 0-shot CoT
Chart Q&A <i>Relaxed accuracy (test)</i>	90.8% 0-shot CoT	80.8% 0-shot CoT	85.7% 0-shot CoT	87.2% 0-shot CoT
Document visual Q&A <i>ANLS score, test</i>	95.2% 0-shot	89.3% 0-shot	92.8% 0-shot	93.1% 0-shot

Google Gemini

Largest and most capable AI model
Making AI more helpful for everyone



LMSYS Chatbot Arena Leaderboard

GPT-4o

Claude 3.5

Rank* (UB)	Model	Arena Score	95% CI	Votes	Organization	License	Knowledge Cutoff
1	ChatGPT-4o-latest (2024-08-08)	1316	+4/-4	24358	OpenAI	Proprietary	2023/10
2	Gemini-1.5-Pro-Exp-0827	1301	+5/-5	19976	Google	Proprietary	2023/11
2	Gemini-1.5-Pro-Exp-0801	1298	+4/-3	25471	Google	Proprietary	2023/11
2	Grok-2-08-13	1295	+4/-6	10170	xAI	Proprietary	2024/3
5	GPT-4o-2024-05-13	1286	+3/-3	83181	OpenAI	Proprietary	2023/10
6	GPT-4o-mini-2024-07-18	1274	+4/-4	23318	OpenAI	Proprietary	2023/10
6	Gemini-1.5-Flash-Exp-0827	1270	+7/-6	6610	Google	Proprietary	2023/11
6	Claude 3.5 Sonnet	1270	+3/-3	53610	Anthropic	Proprietary	2024/4
6	Gemini Advanced App (2024-05-14)	1266	+3/-3	52225	Google	Proprietary	Online
6	Grok-2-Mini-08-13	1266	+6/-6	10939	xAI	Proprietary	2024/3
7	Meta-Llama-3.1-405b-Instruct	1266	+3/-4	24855	Meta	Llama 3.1 Community	2023/12
7	GPT-4o-2024-08-06	1262	+7/-5	14960	OpenAI	Proprietary	2023/10
10	Gemini-1.5-Pro-001	1259	+3/-3	74909	Google	Proprietary	2023/11

<https://lmarena.ai/>

Perplexity.ai



Anonymous Like

Share

New Thread

Home

Discover

Library

Login

Sign Up

Try Pro

Upgrade to Claude-2 or GPT-4, boost Copilot uses, and upload more files.

Learn More

Download

Large Language Models for Data Science

Sources

What Is a Large Language Model and What Can...

snowflake · 1

What are Large Language Models? | NVIDIA

nvidia · 2

Catch Up On Large Language Models | by...

towardsdatascien... · 3



View 3 more

Answer

Large Language Models (LLMs) are advanced AI systems designed to understand human language intricacies and generate intelligent, creative responses to queries ¹. They are deep learning algorithms that can recognize, summarize, translate, predict, and generate content using very large datasets ². LLMs are trained on enormous data sets typically measured in petabytes, sourced from books, articles, websites, and other text-based sources ¹. They are dynamic tools that are continuously updated and refined as new data becomes available ¹.

In the field of data science, LLMs can potentially transform how teams source, manage, and analyze data ¹. They are used in numerous data science applications, and their ability to process and interpret vast amounts of text data has made them an indispensable part of m

Ask follow-up...

Copilot




social media posts or customer reviews, to determine whether the overall sentiment is

<https://www.perplexity.ai/>

Transformer (Attention is All You Need)

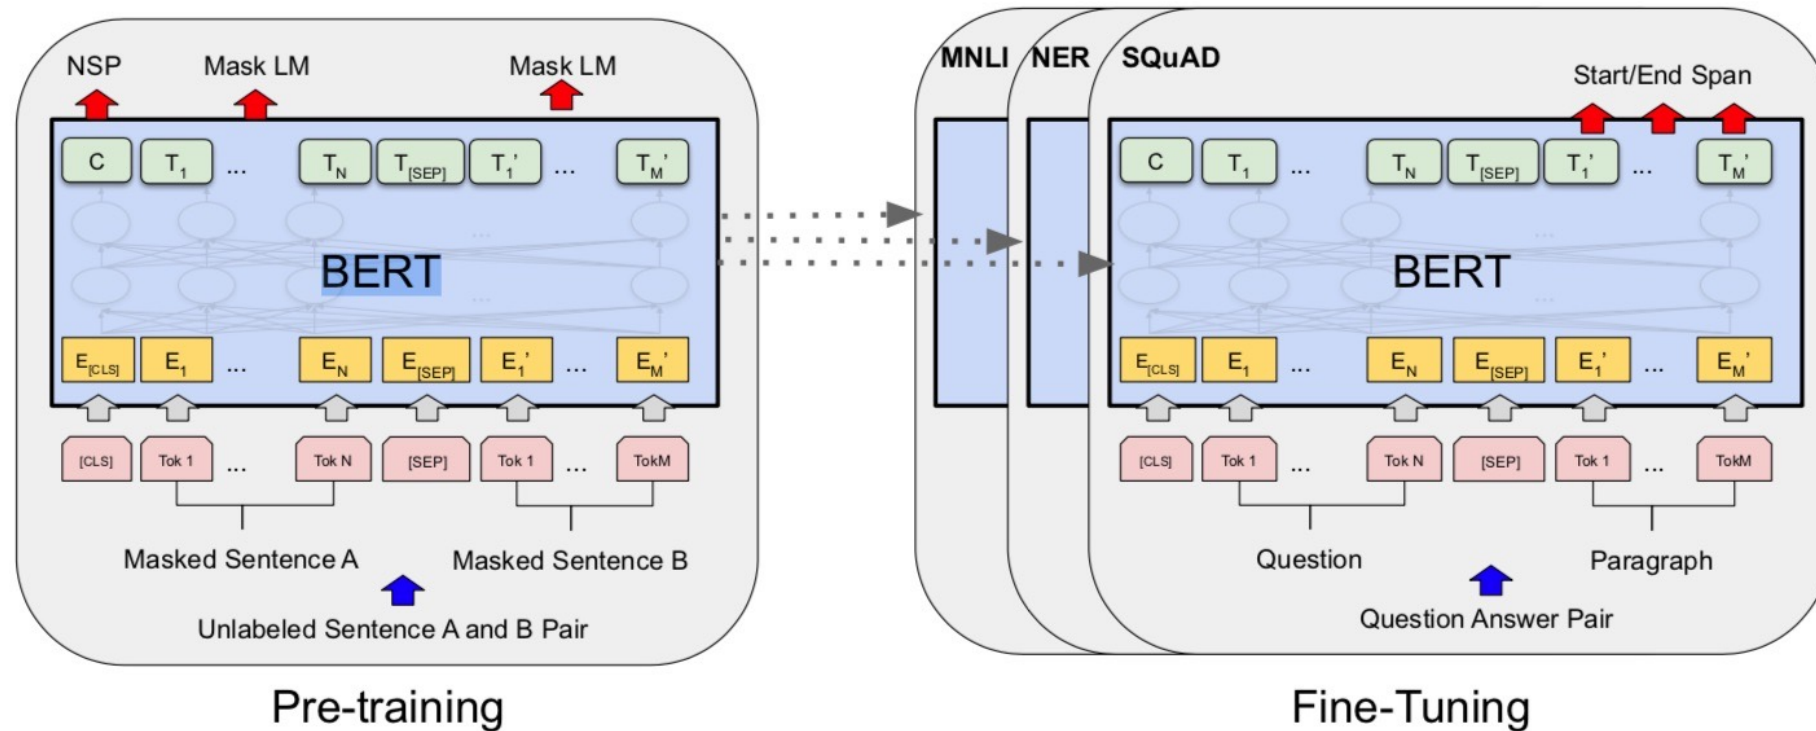
(Vaswani et al., 2017)



BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding

BERT (Bidirectional Encoder Representations from Transformers)

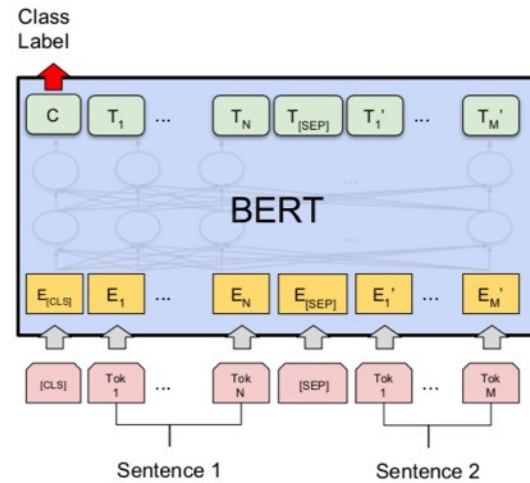
Overall pre-training and fine-tuning procedures for BERT




Source: Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2018).

"Bert: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint arXiv:1810.04805.


Fine-tuning BERT on Different Tasks



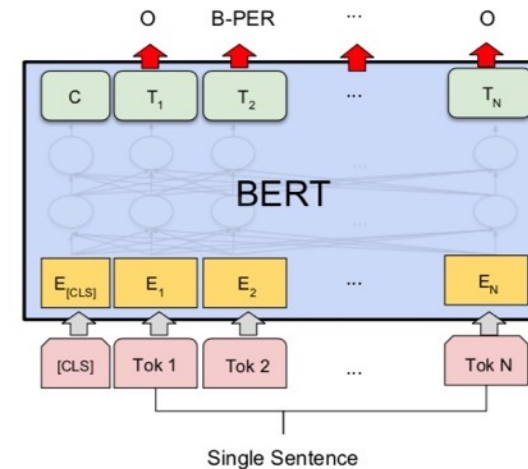
(a) Sentence Pair Classification Tasks:
MNLI, QQP, QNLI, STS-B, MRPC,
RTE, SWAG



(b) Single Sentence Classification Tasks:
SST-2, CoLA



(c) Question Answering Tasks:
SQuAD v1.1

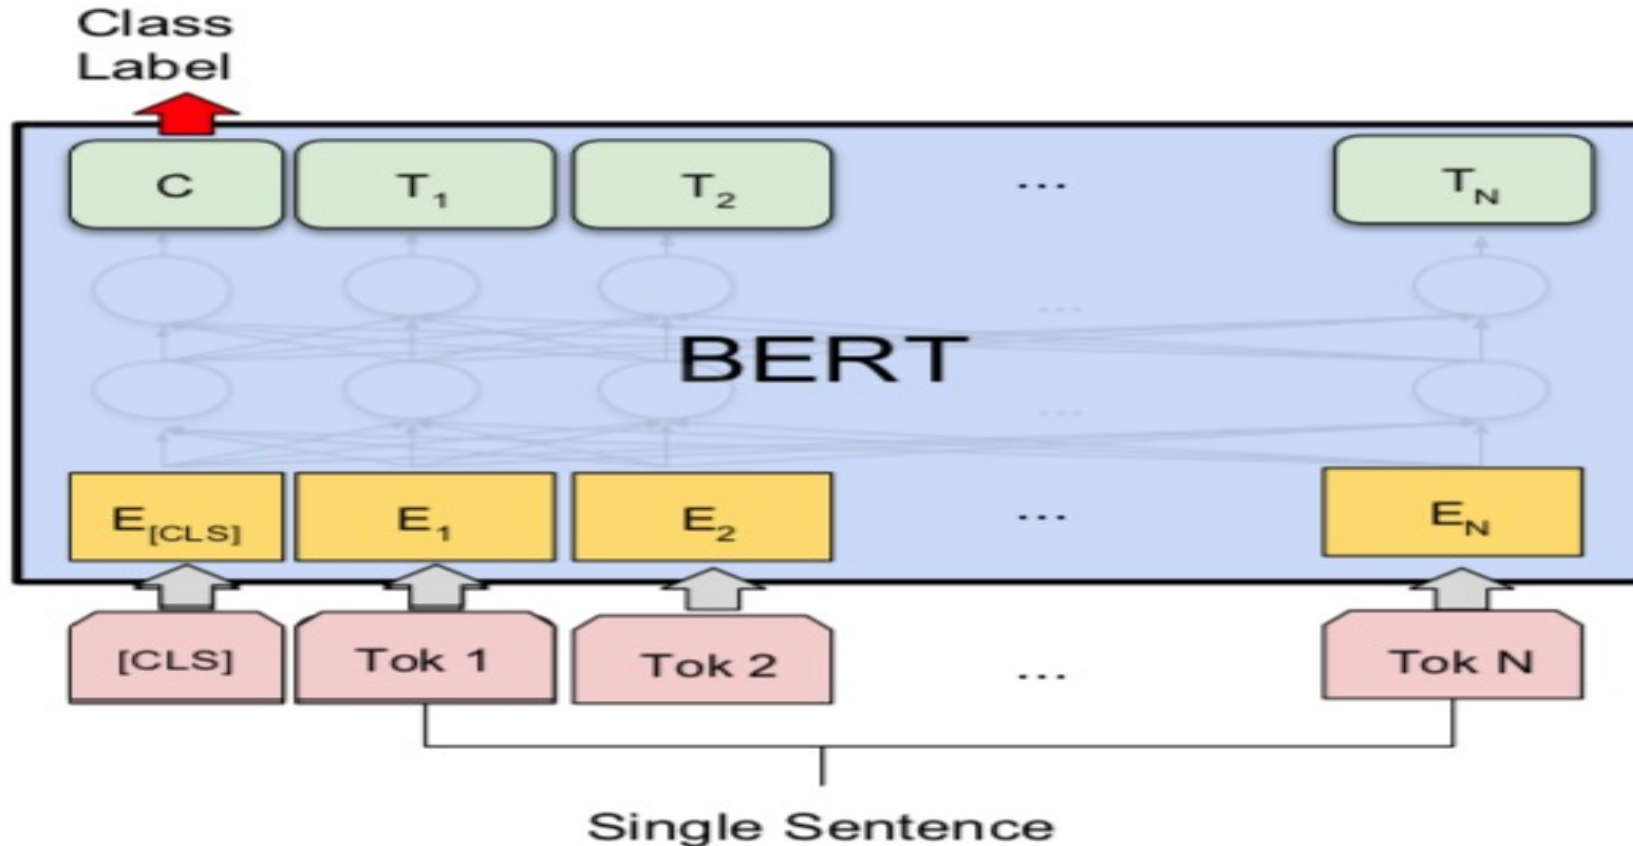


(d) Single Sentence Tagging Tasks:
CoNLL-2003 NER

Source: Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2018).


"Bert: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint arXiv:1810.04805.

Sentiment Analysis: Single Sentence Classification



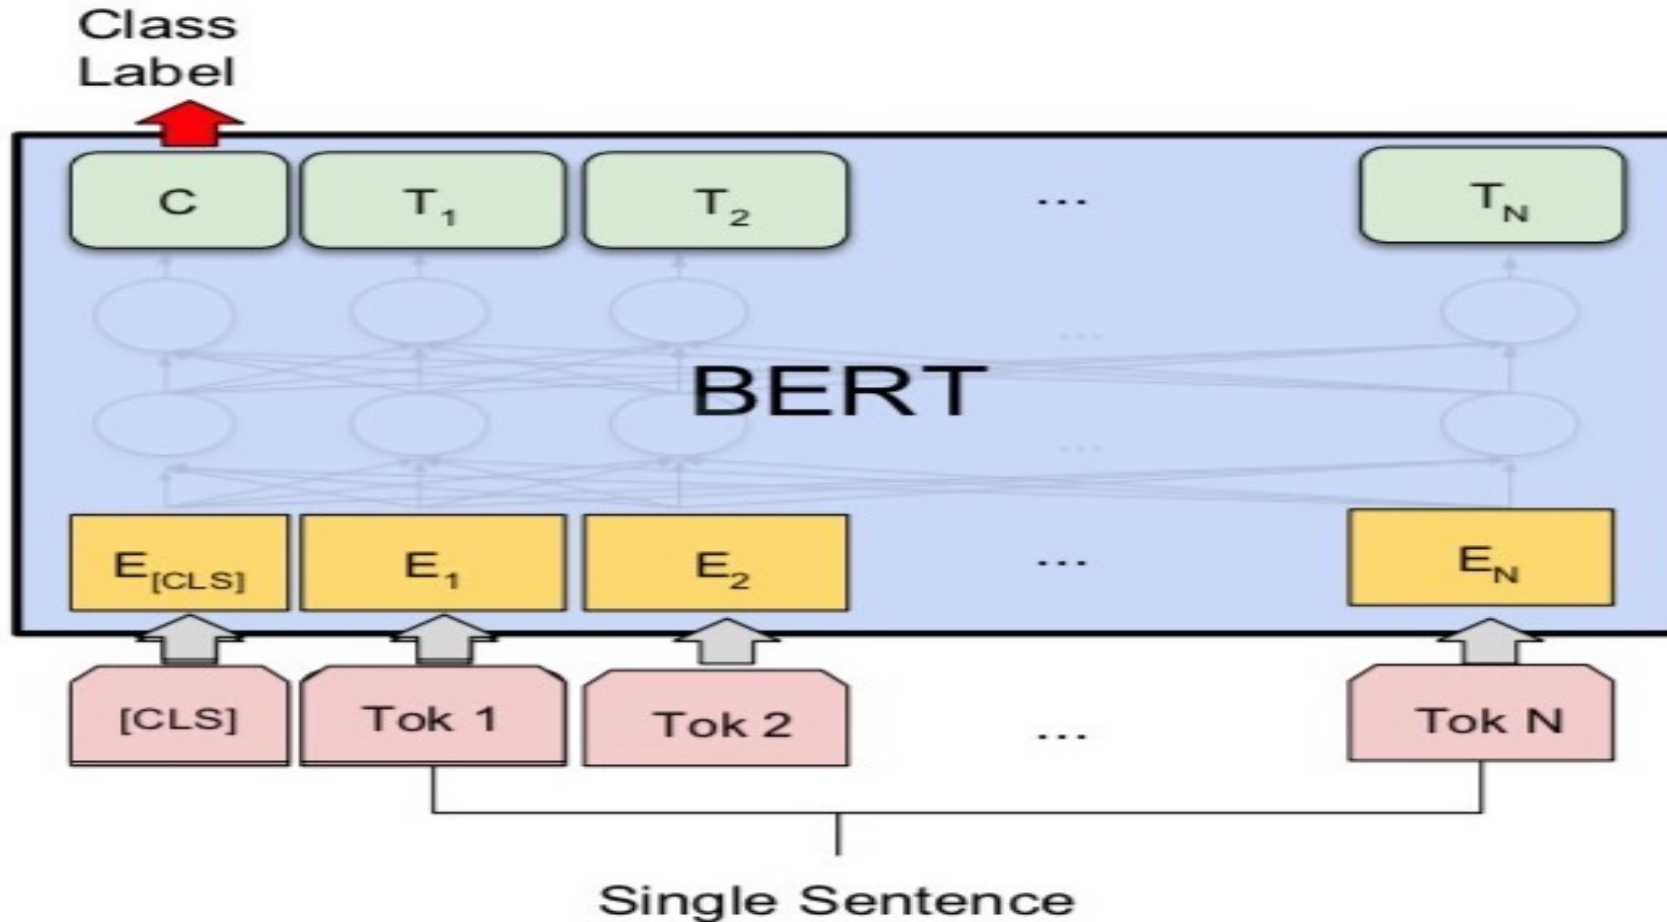
(b) Single Sentence Classification Tasks:
SST-2, CoLA

Fine-tuning BERT on Question Answering (QA)



(c) Question Answering Tasks:
SQuAD v1.1

Fine-tuning BERT on Dialogue Intent Detection (ID; Classification)



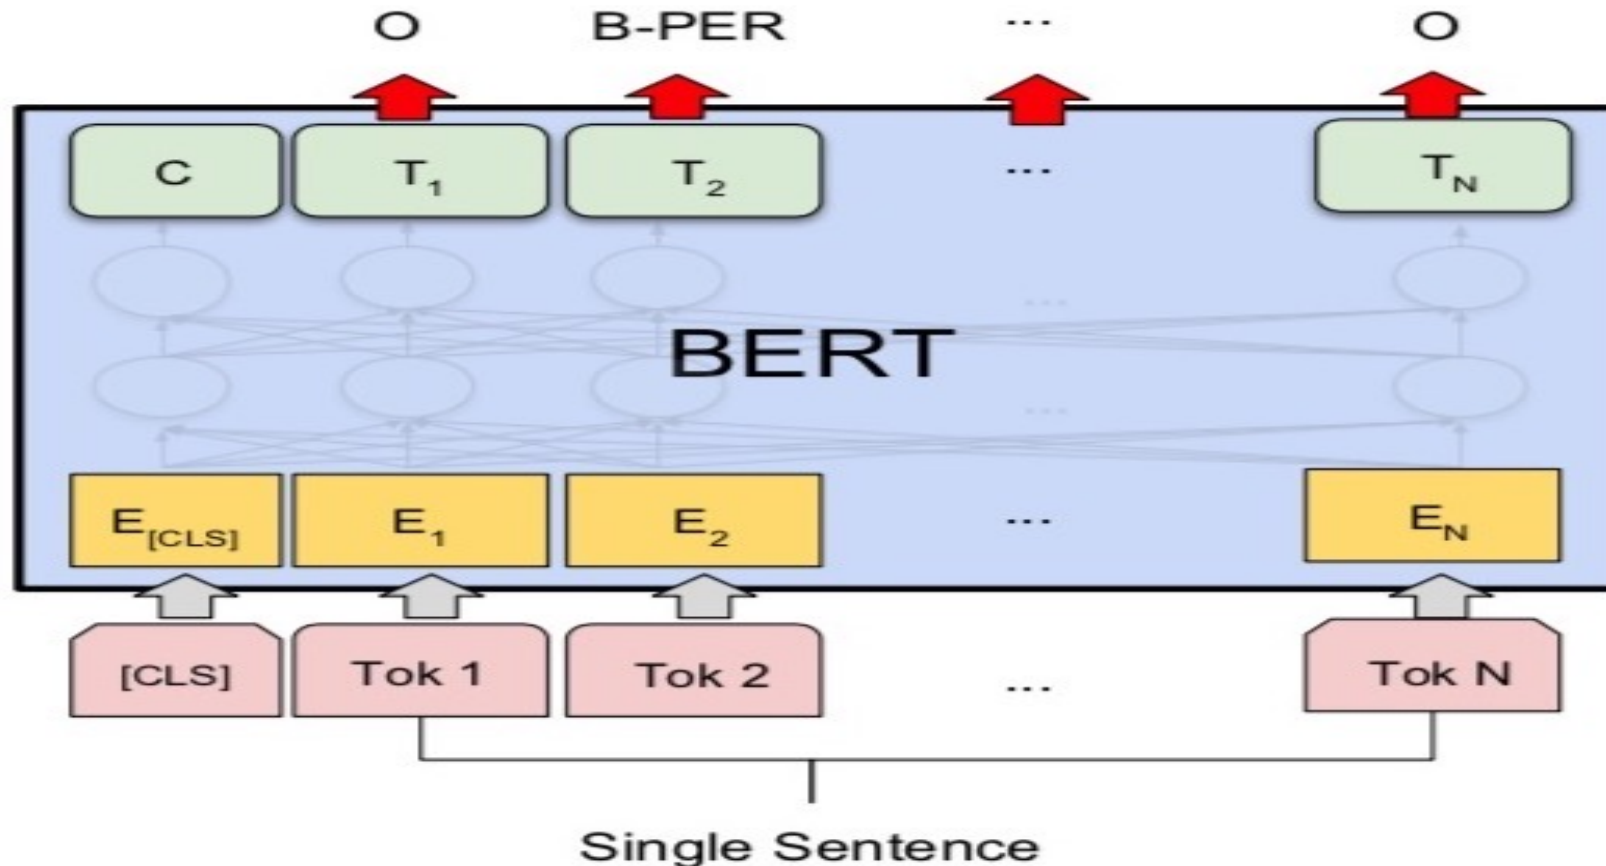
(b) Single Sentence Classification Tasks: SST-2, CoLA

Source: Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2018).

"Bert: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint arXiv:1810.04805.

Fine-tuning BERT on Dialogue

Slot Filling (SF)



(d) Single Sentence Tagging Tasks:
CoNLL-2003 NER


Source: Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2018).

"Bert: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint arXiv:1810.04805.

Task-Oriented Dialogue (ToD) System

Speech, Text, NLP

"Book me a cab to Russell Square"




"When do you want to leave?"

Source: Razumovskaia, Evgeniia, Goran Glavas, Olga Majewska, Edoardo M. Ponti, Anna Korhonen, and Ivan Vulic.

"Crossing the conversational chasm: A primer on natural language processing for multilingual task-oriented dialogue systems." *Journal of Artificial Intelligence Research* 74 (2022): 1351-1402.

Conversational AI

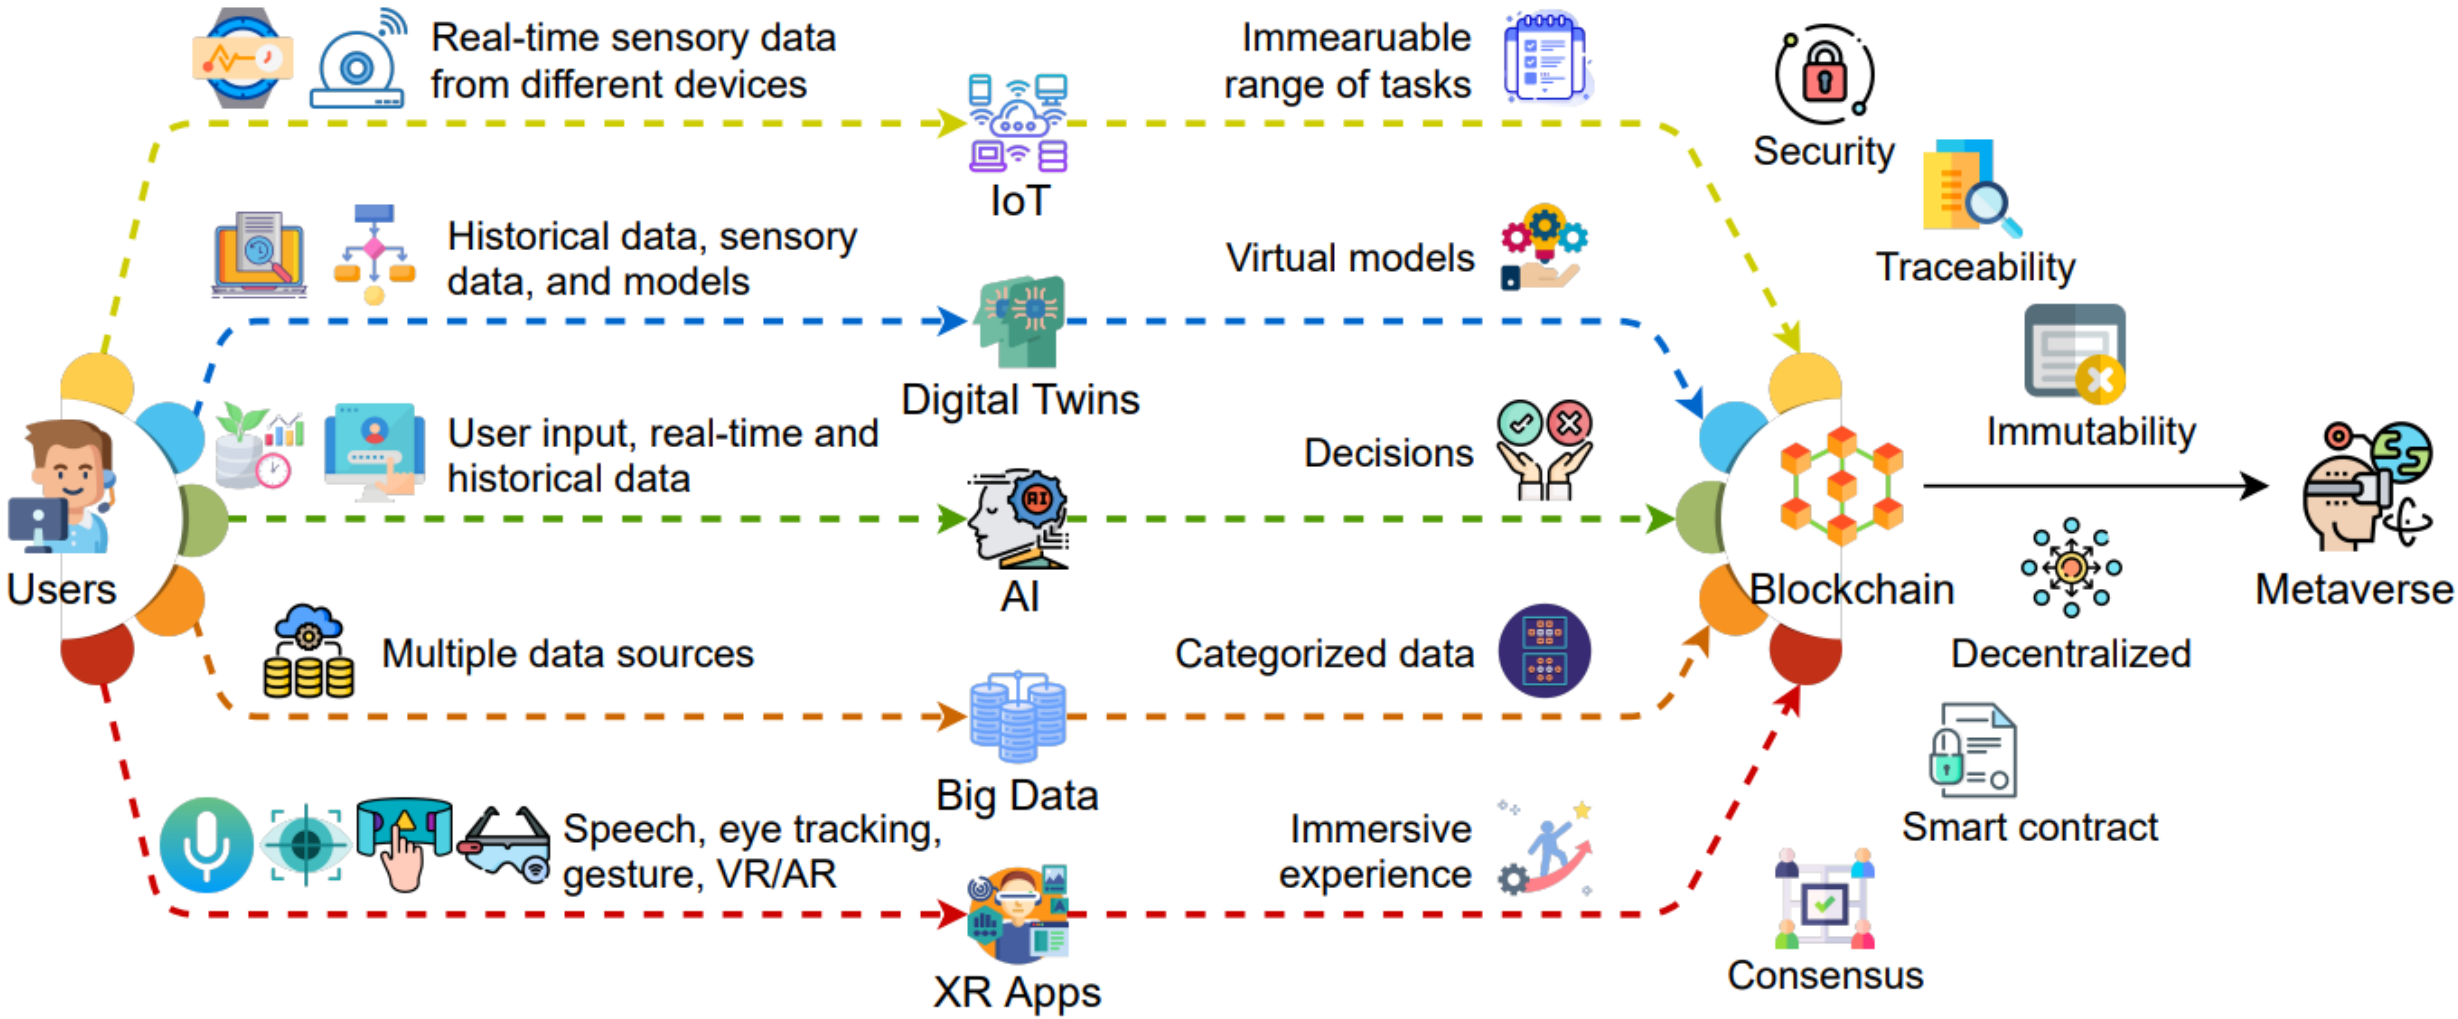
to deliver contextual and personal experience to users



Source: Huynh-The, Thien, Quoc-Viet Pham, Xuan-Quy Pham, Thanh Thi Nguyen, Zhu Han, and Dong-Seong Kim (2022). "Artificial Intelligence for the Metaverse: A Survey." arXiv preprint arXiv:2202.10336.

AI and Blockchain

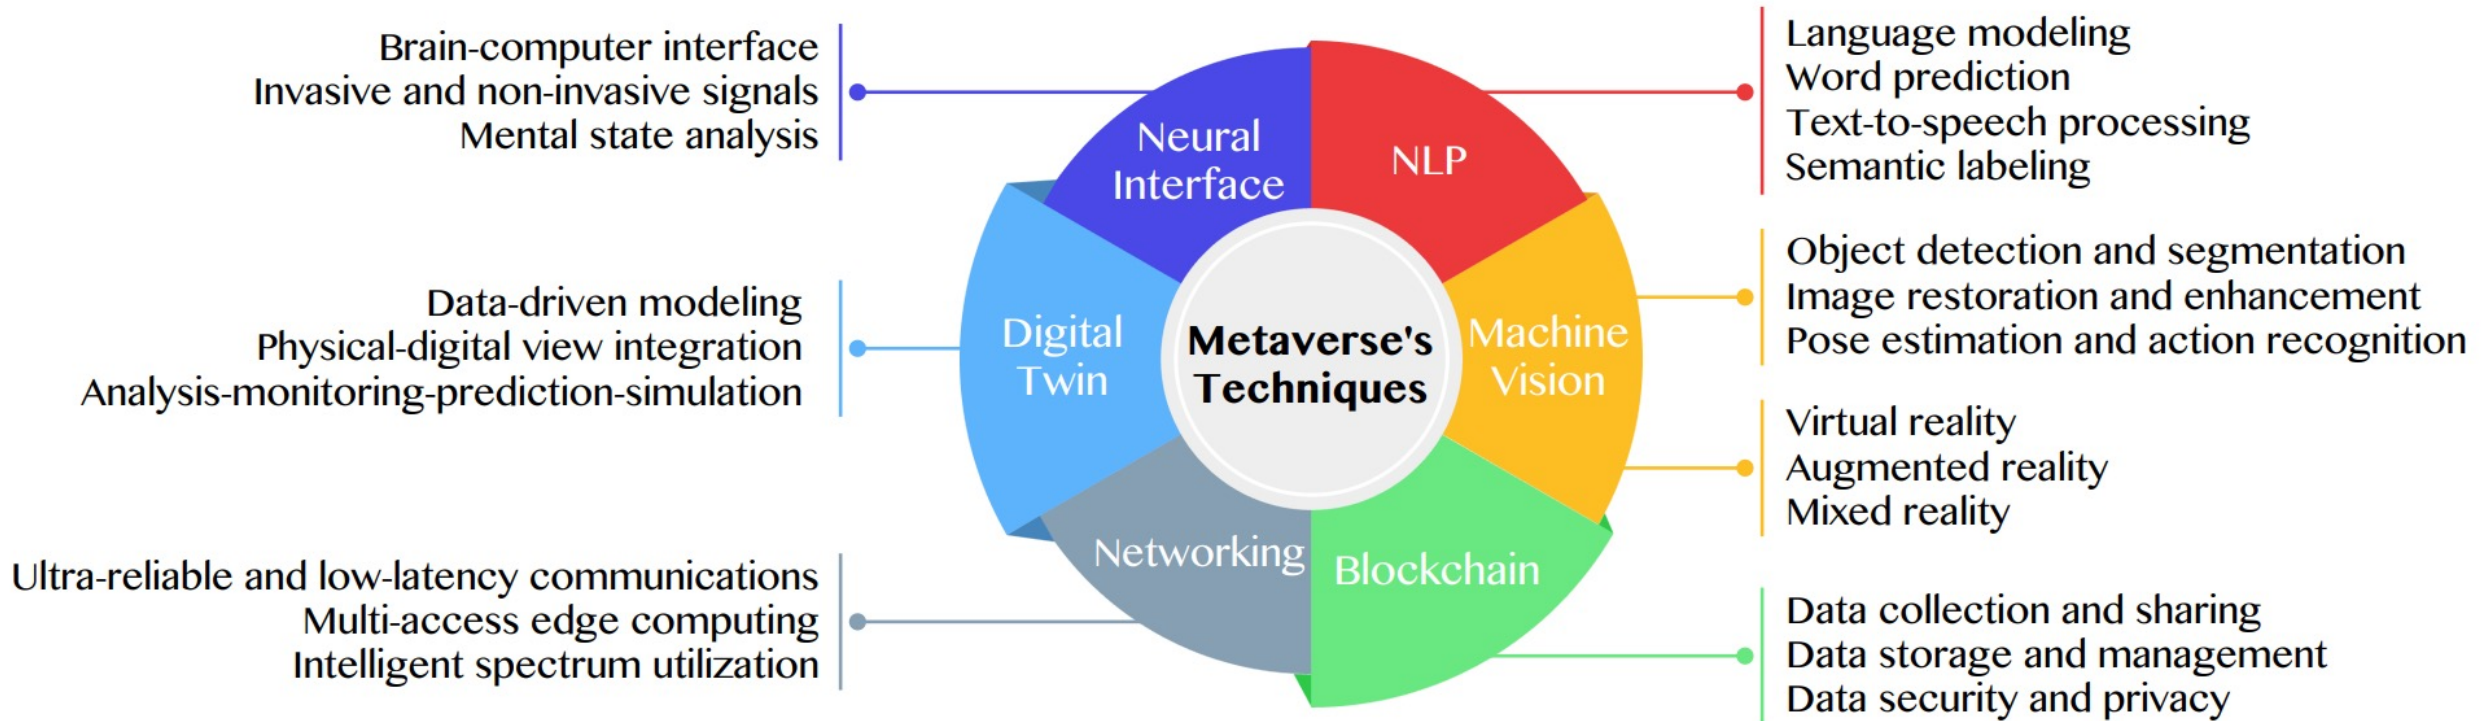
Key Enabling Technologies of the Metaverse



Source: Gadekallu, Thippa Reddy, Thien Huynh-The, Weizheng Wang, Gokul Yenduri, Pasika Ranaweera, Quoc-Viet Pham, Daniel Benevides da Costa, and Madhusanka Liyanage (2022). "Blockchain for the Metaverse: A Review." arXiv preprint arXiv:2203.09738..

Primary Technical Aspects in the Metaverse

AI with ML algorithms and DL architectures is advancing the user experience in the virtual world




Source: Huynh-The, Thien, Quoc-Viet Pham, Xuan-Quy Pham, Thanh Thi Nguyen, Zhu Han, and Dong-Seong Kim (2022).

"Artificial Intelligence for the Metaverse: A Survey." arXiv preprint arXiv:2202.10336.

AI for the Metaverse in the Application Aspects

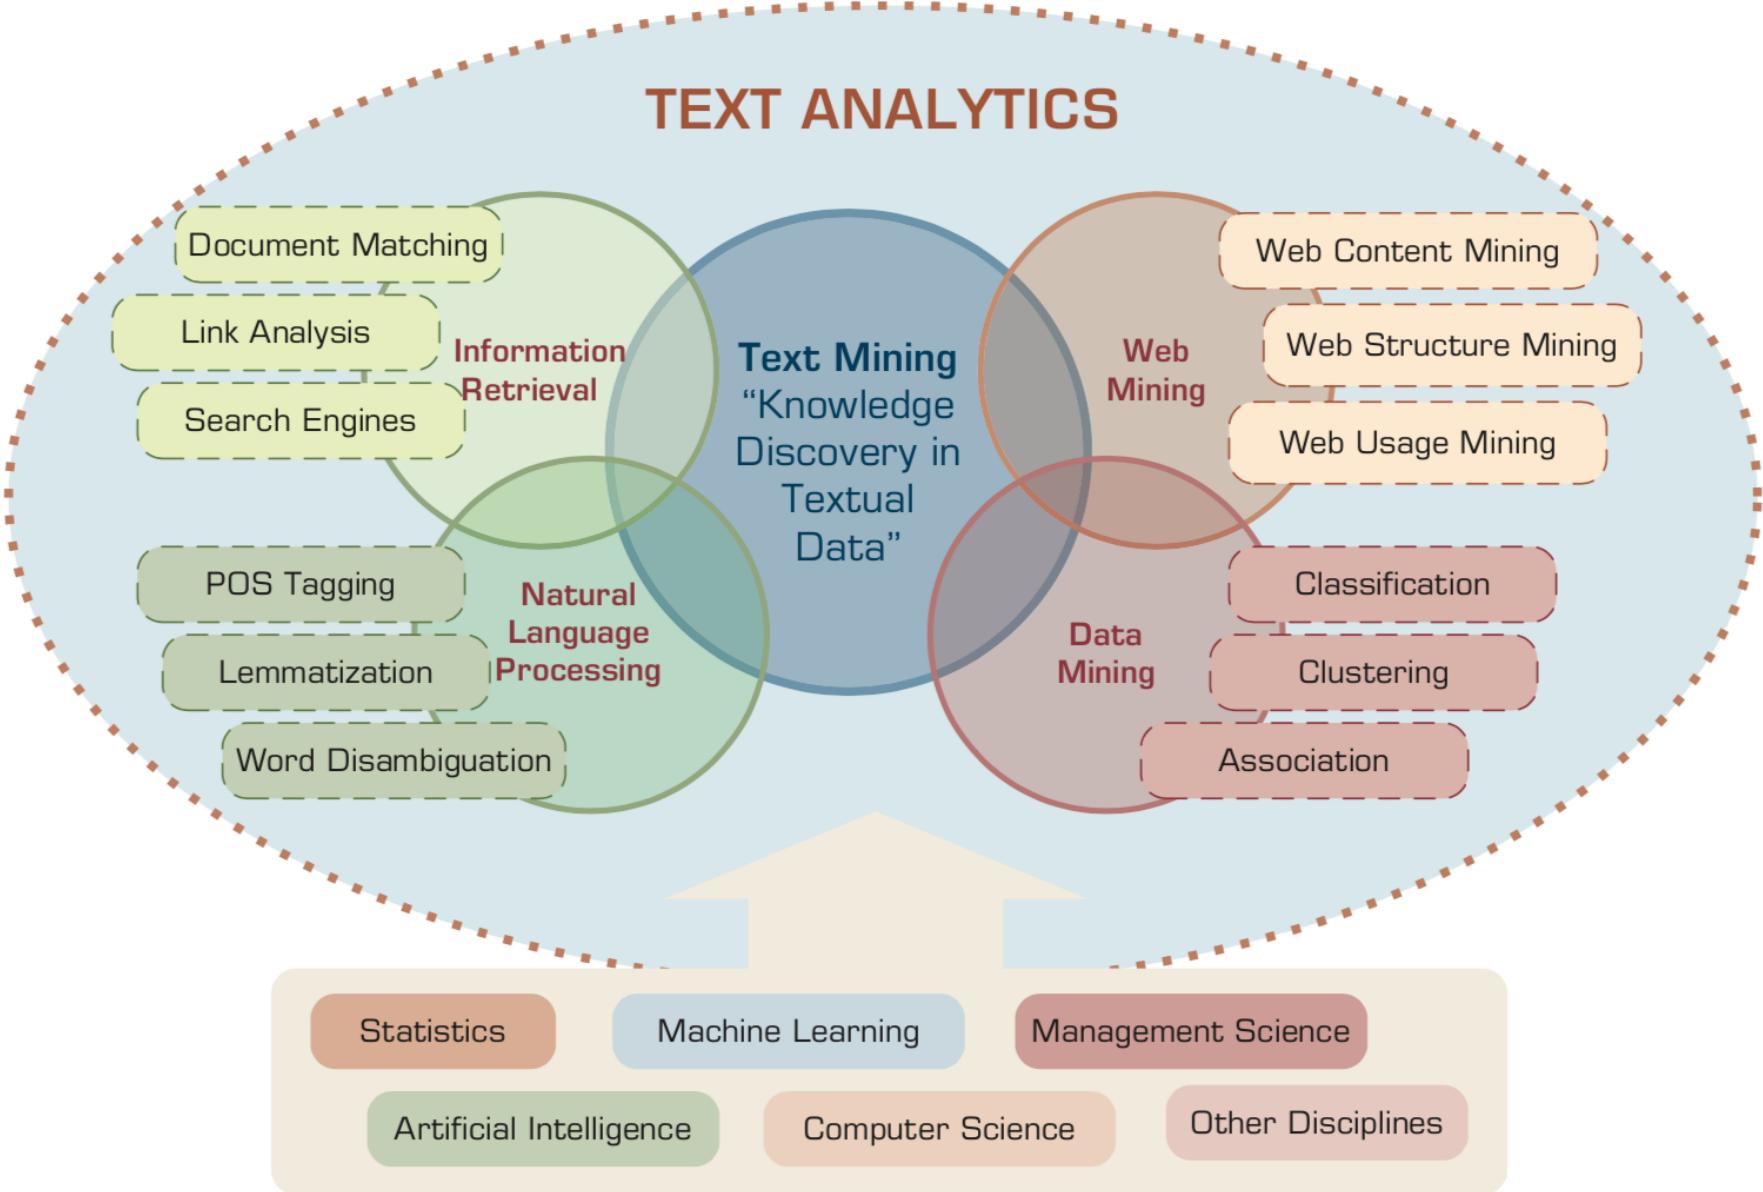
healthcare, manufacturing, smart cities, gaming
E-commerce, human resources, real estate, and DeFi



Source: Huynh-The, Thien, Quoc-Viet Pham, Xuan-Quy Pham, Thanh Thi Nguyen, Zhu Han, and Dong-Seong Kim (2022).

"Artificial Intelligence for the Metaverse: A Survey." arXiv preprint arXiv:2202.10336.

AI for Text Analytics



Source: Ramesh Sharda, Dursun Delen, and Efraim Turban (2017), Business Intelligence, Analytics, and Data Science: A Managerial Perspective, 4th Edition, Pearson

4 Approaches of AI


<p>2. Thinking Humanly: The Cognitive Modeling Approach</p>	<p>3. Thinking Rationally: The “Laws of Thought” Approach</p>
<p>1. Acting Humanly: The Turing Test Approach (1950)</p>	<p>4. Acting Rationally: The Rational Agent Approach</p>

Acting Rationally: The Rational Agent Approach

- AI has focused on the study and construction of agents that **do the right thing.**
- **Standard model**

Neuroscience

The parts of a **nerve cell** or **neuron**



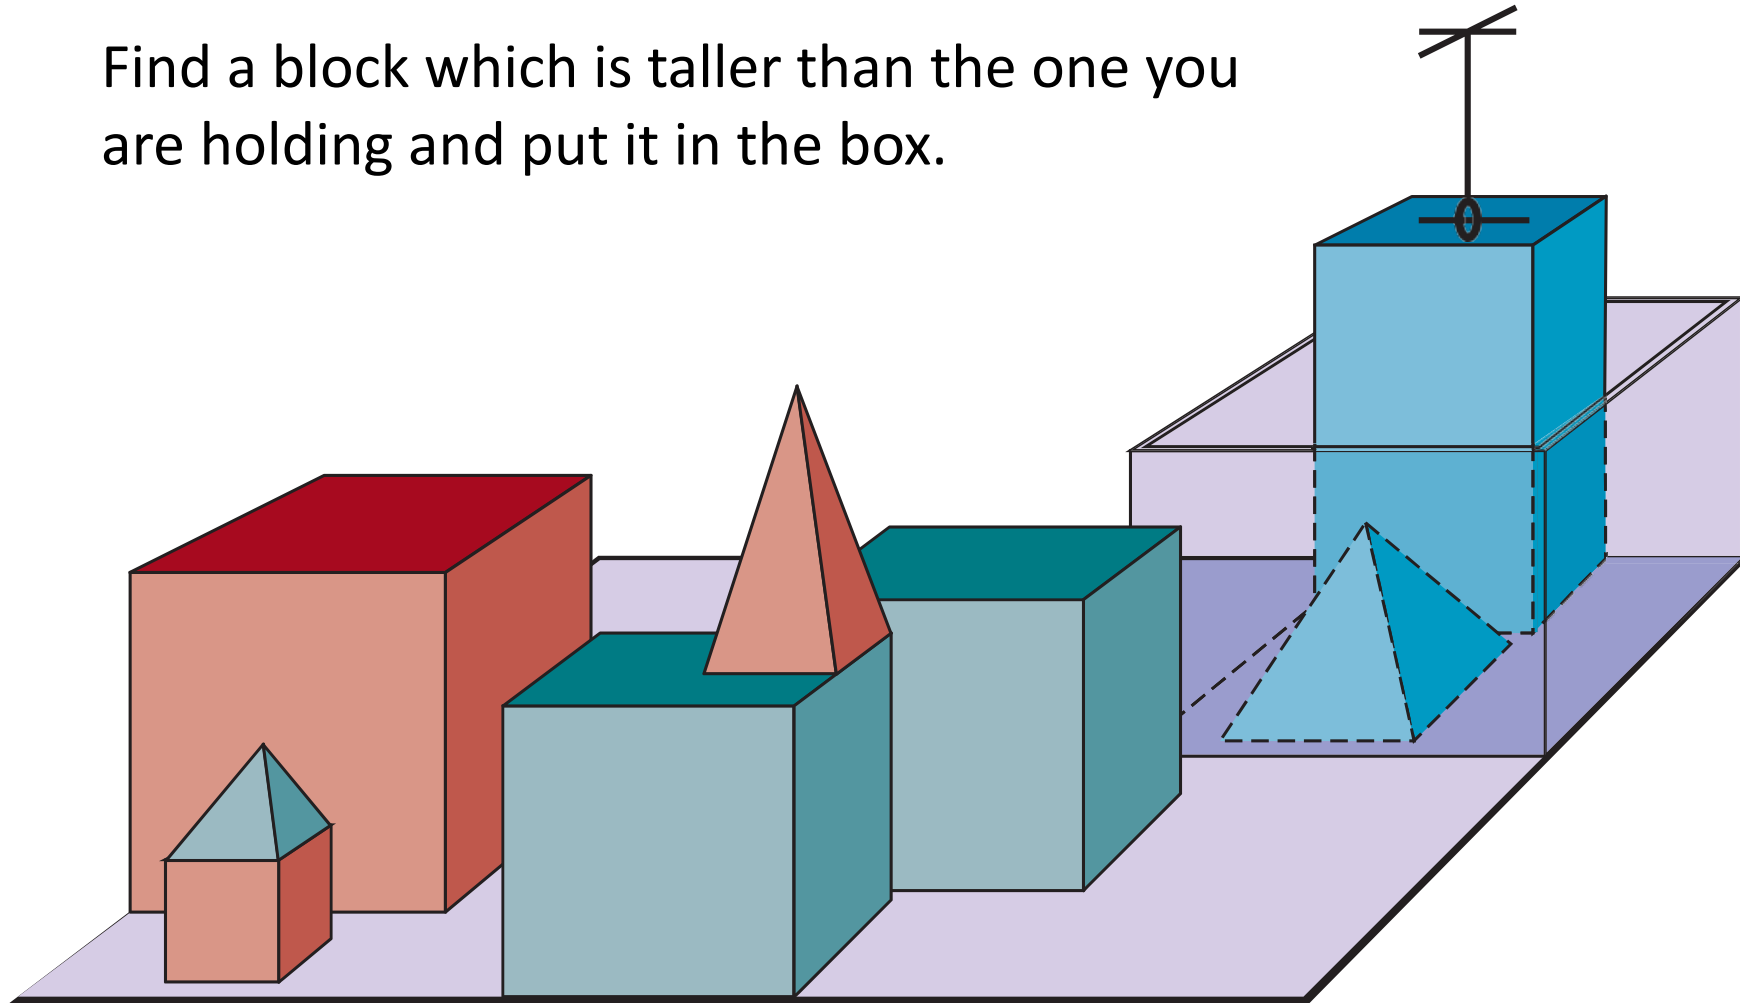
Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Comparison of Computer and Human Brain

	Supercomputer	Personal Computer	Human Brain
Computational units	10^6 GPUs + CPUs	8 CPU cores	10^6 columns
	10^{15} transistors	10^{10} transistors	10^{11} neurons
Storage units	10^{16} bytes RAM	10^{10} bytes RAM	10^{11} neurons
	10^{17} bytes disk	10^{12} bytes disk	10^{14} synapses
Cycle time	10^{-9} sec	10^{-9} sec	10^{-3} sec
Operations/sec	10^{18}	10^{10}	10^{17}

A scene from the blocks world

Find a block which is taller than the one you are holding and put it in the box.

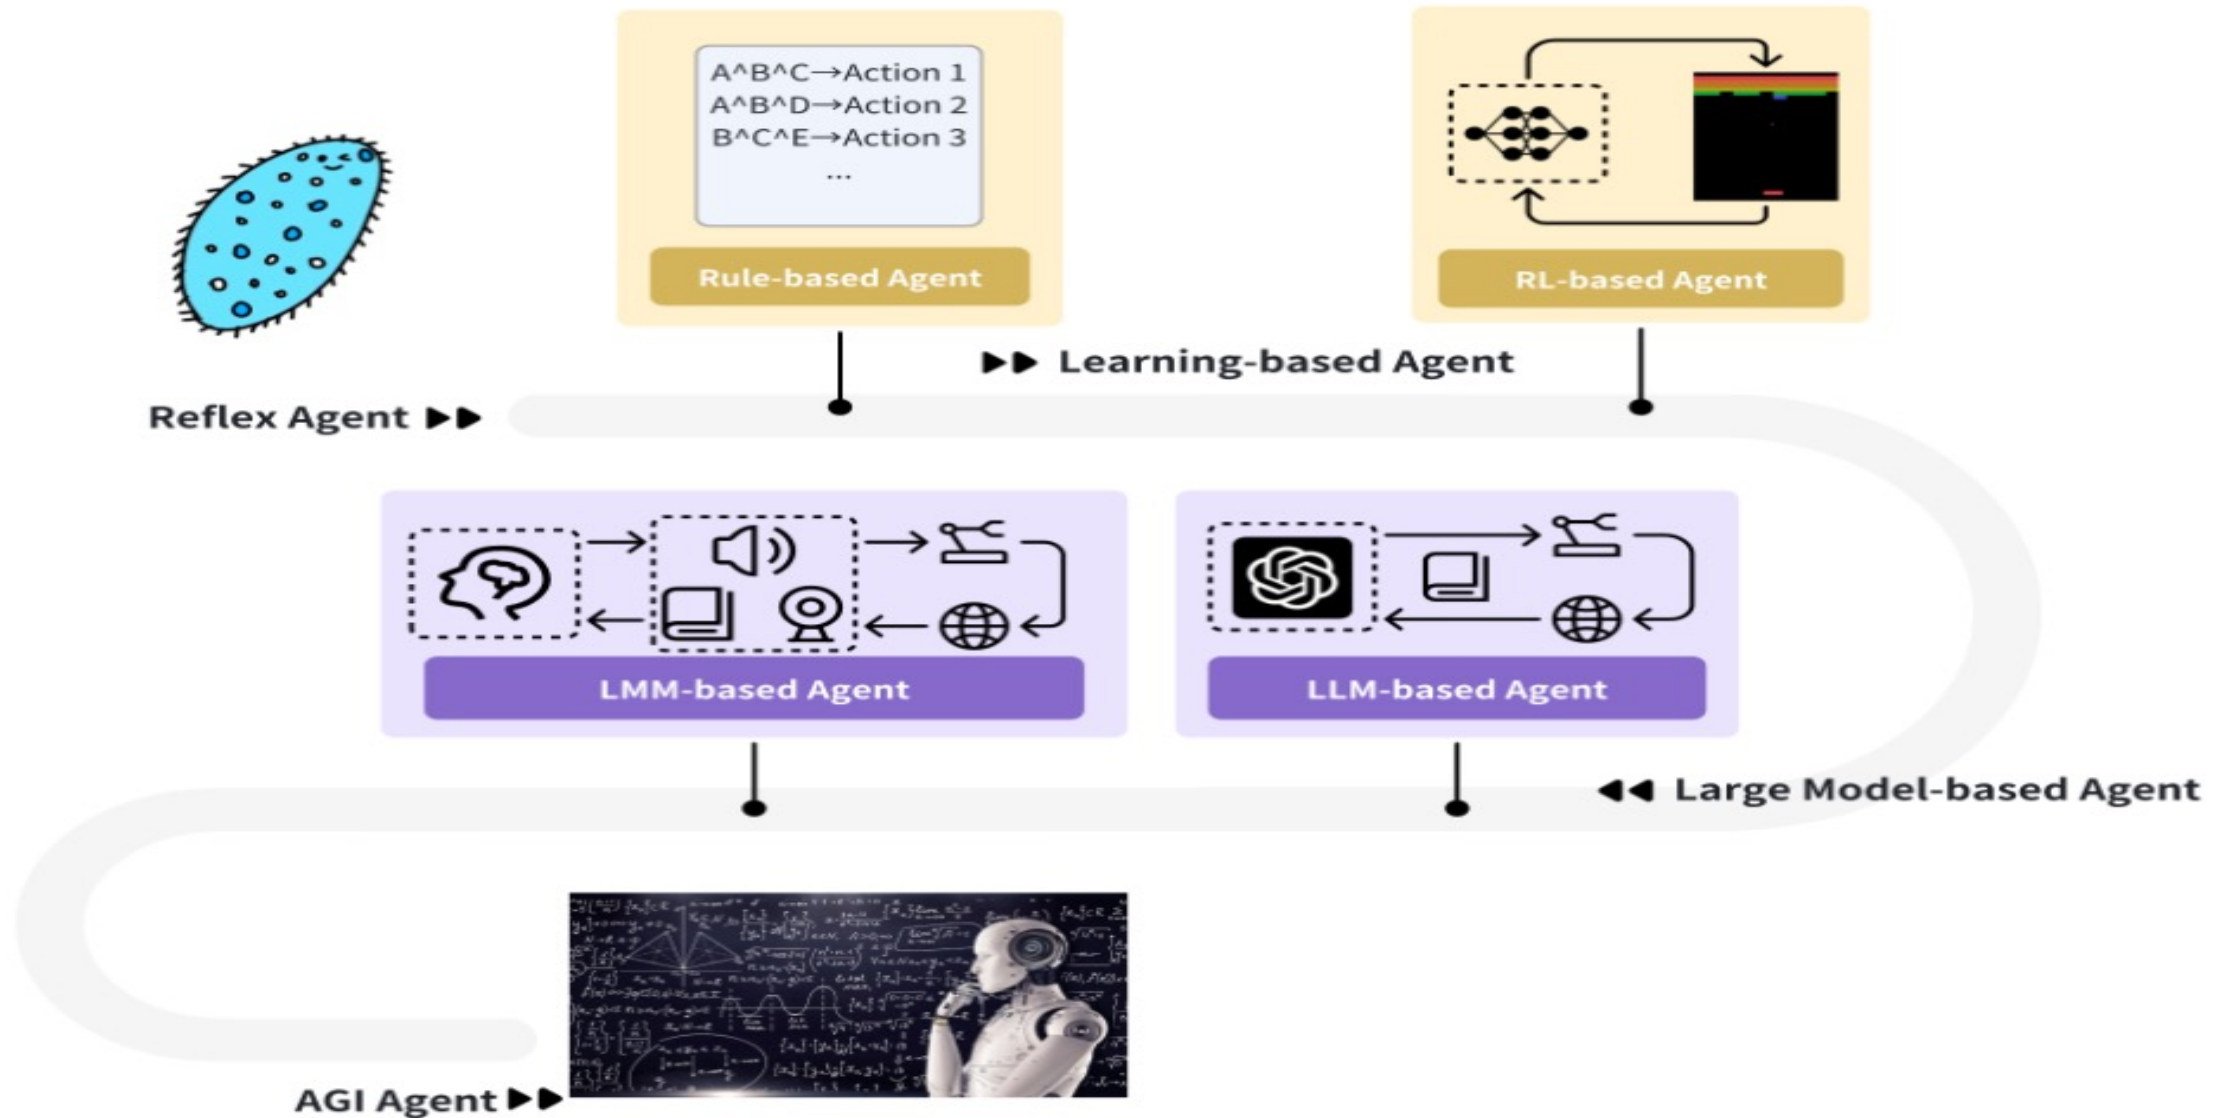


Intelligent Agents

4 Approaches of AI

<p>2. Thinking Humanly: The Cognitive Modeling Approach</p>	<p>3. Thinking Rationally: The “Laws of Thought” Approach</p>
<p>1. Acting Humanly: The Turing Test Approach <small>(1950)</small></p>	<p>4. Acting Rationally: The Rational Agent Approach</p>

Intelligent Agents Roadmap



AI Agents


- **Traditional AI Agents**

- **Simple reflex agents**
- **Model-based reflex agents**
- **Goal-based agents**
- **Utility-based agents**
- **Learning agents**


- **Evolution of AI Agents**

- **LLM-based Agents**
- **Multi-modal agents**
- **Embodied AI agents in virtual environments**
- **Collaborative AI agents**

Agents interact with environments through sensors and actuators



AI Agents



- **Definition:** An **AI agent** is an **entity** that **perceives** its **environment** and takes **actions** to achieve **goals**
- **Components:**
 - 1. Sensors:** Perceive the environment
 - 2. Actuators:** Act upon the environment
 - 3. Decision-making** mechanism: Process inputs and decide on actions

Reinforcement Learning (DL)


Agent

Environment


Reinforcement Learning (DL)



Reinforcement Learning (DL)



A vacuum-cleaner world with just two locations



Partial tabulation of a simple agent function for the vacuum-cleaner world

Percept sequence	Action
<i>[A, Clean]</i>	<i>Right</i>
<i>[A, Dirty]</i>	<i>Suck</i>
<i>[B, Clean]</i>	<i>Left</i>
<i>[B, Dirty]</i>	<i>Suck</i>
<i>[A, Clean], [A, Clean]</i>	<i>Right</i>
<i>[A, Clean], [A, Dirty]</i>	<i>Suck</i>
<i>⋮</i>	<i>⋮</i>
<i>[A, Clean], [A, Clean], [A, Clean]</i>	<i>Right</i>
<i>[A, Clean], [A, Clean], [A, Dirty]</i>	<i>Suck</i>
<i>⋮</i>	<i>⋮</i>

PEAS description of the task environment for an automated taxi driver

Agent Type	Performance Measure	Environment	Actuators	Sensors
Taxi driver	Safe, fast, legal, comfortable trip, maximize profits, minimize impact on other road users	Roads, other traffic, police, pedestrians, customers, weather	Steering, accelerator, brake, signal, horn, display, speech	Cameras, radar, speedometer, GPS, engine sensors, accelerometer, microphones, touchscreen

Examples of Agent Types and their PEAS descriptions

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

Examples of Task Environments and their Characteristics

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

The TABLE-DRIVEN-AGENT program is invoked for each new percept and returns an action each time.

It retains the complete percept sequence in memory.

function TABLE-DRIVEN-AGENT(*percept*) **returns** an action

persistent: *percepts*, a sequence, initially empty

table, a table of actions, indexed by percept sequences, initially fully specified

append *percept* to the end of *percepts*

action \leftarrow LOOKUP(*percepts*, *table*)

return *action*


The agent program for a simple reflex agent in the two-location vacuum environment.

function REFLEX-VACUUM-AGENT(*[location, status]*) **returns** an action

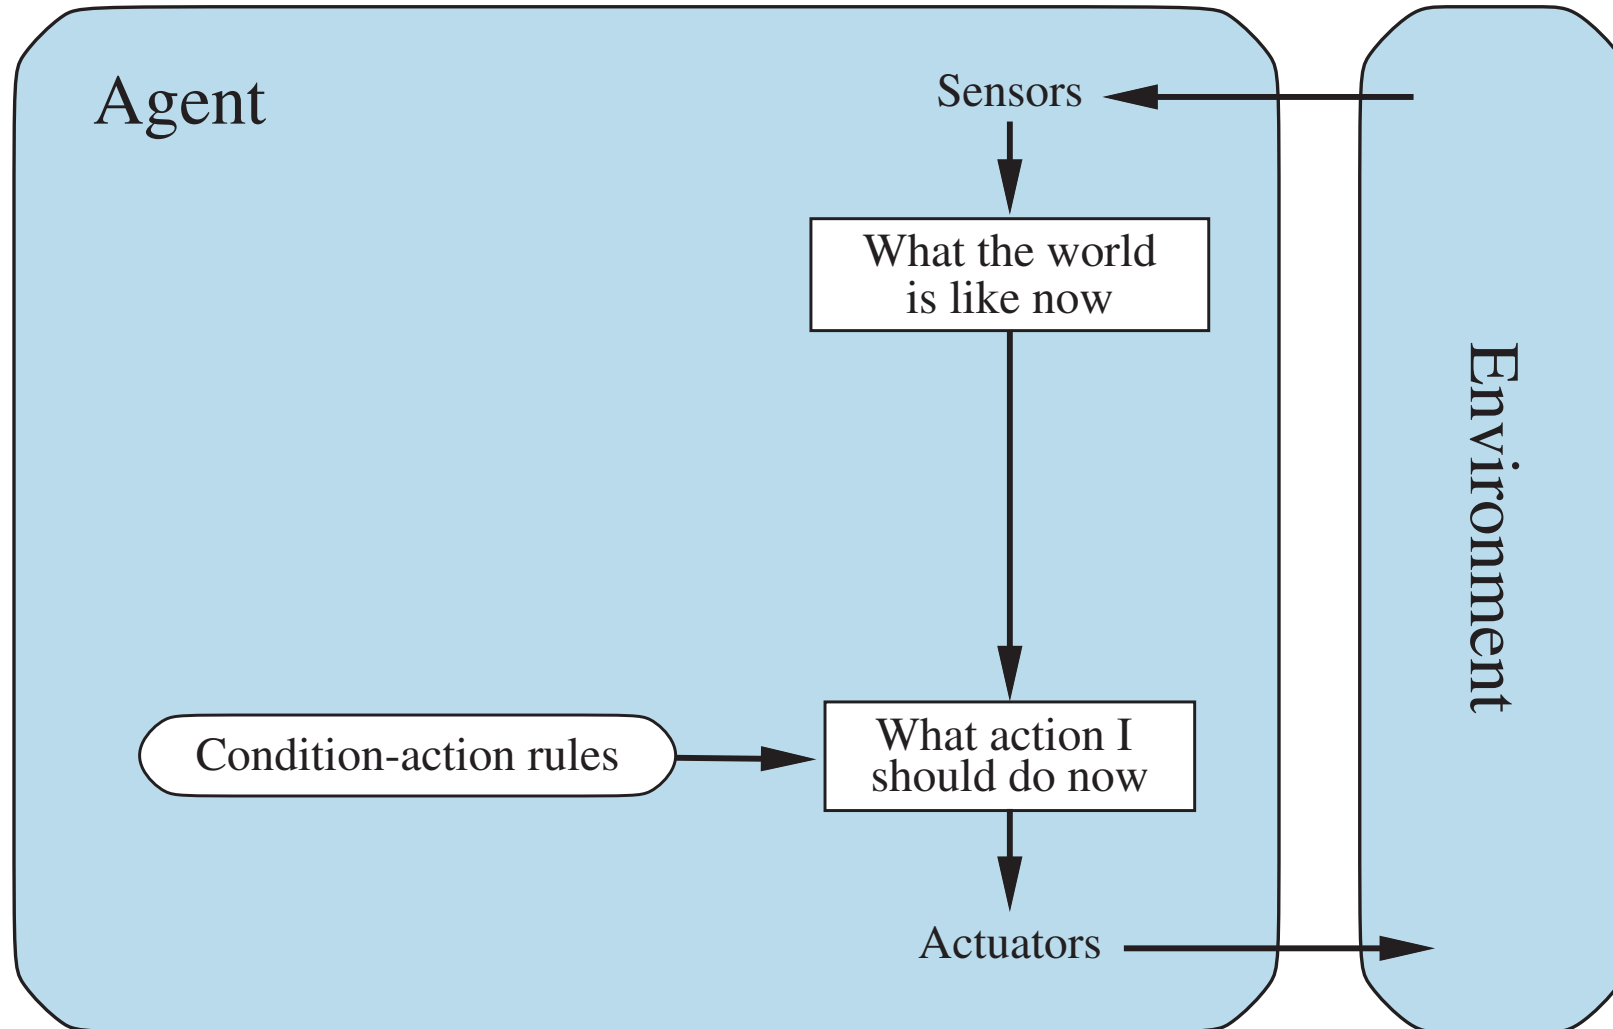
if *status = Dirty* **then return** *Suck*

else if *location = A* **then return** *Right*

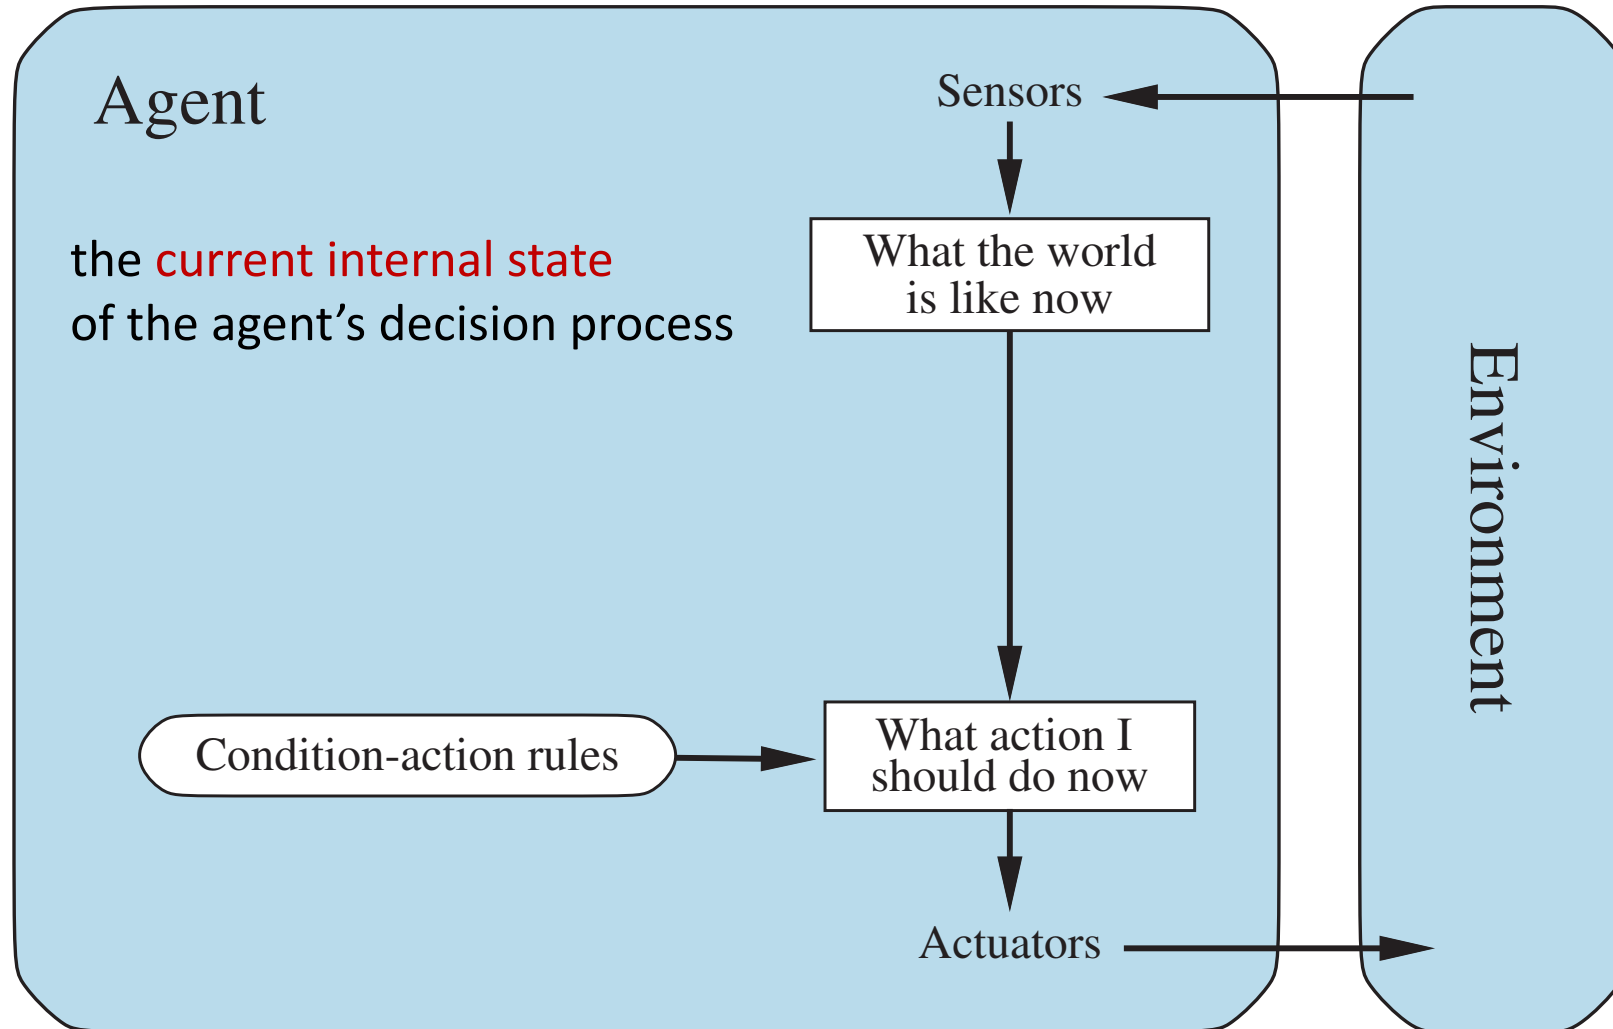
else if *location = B* **then return** *Left*



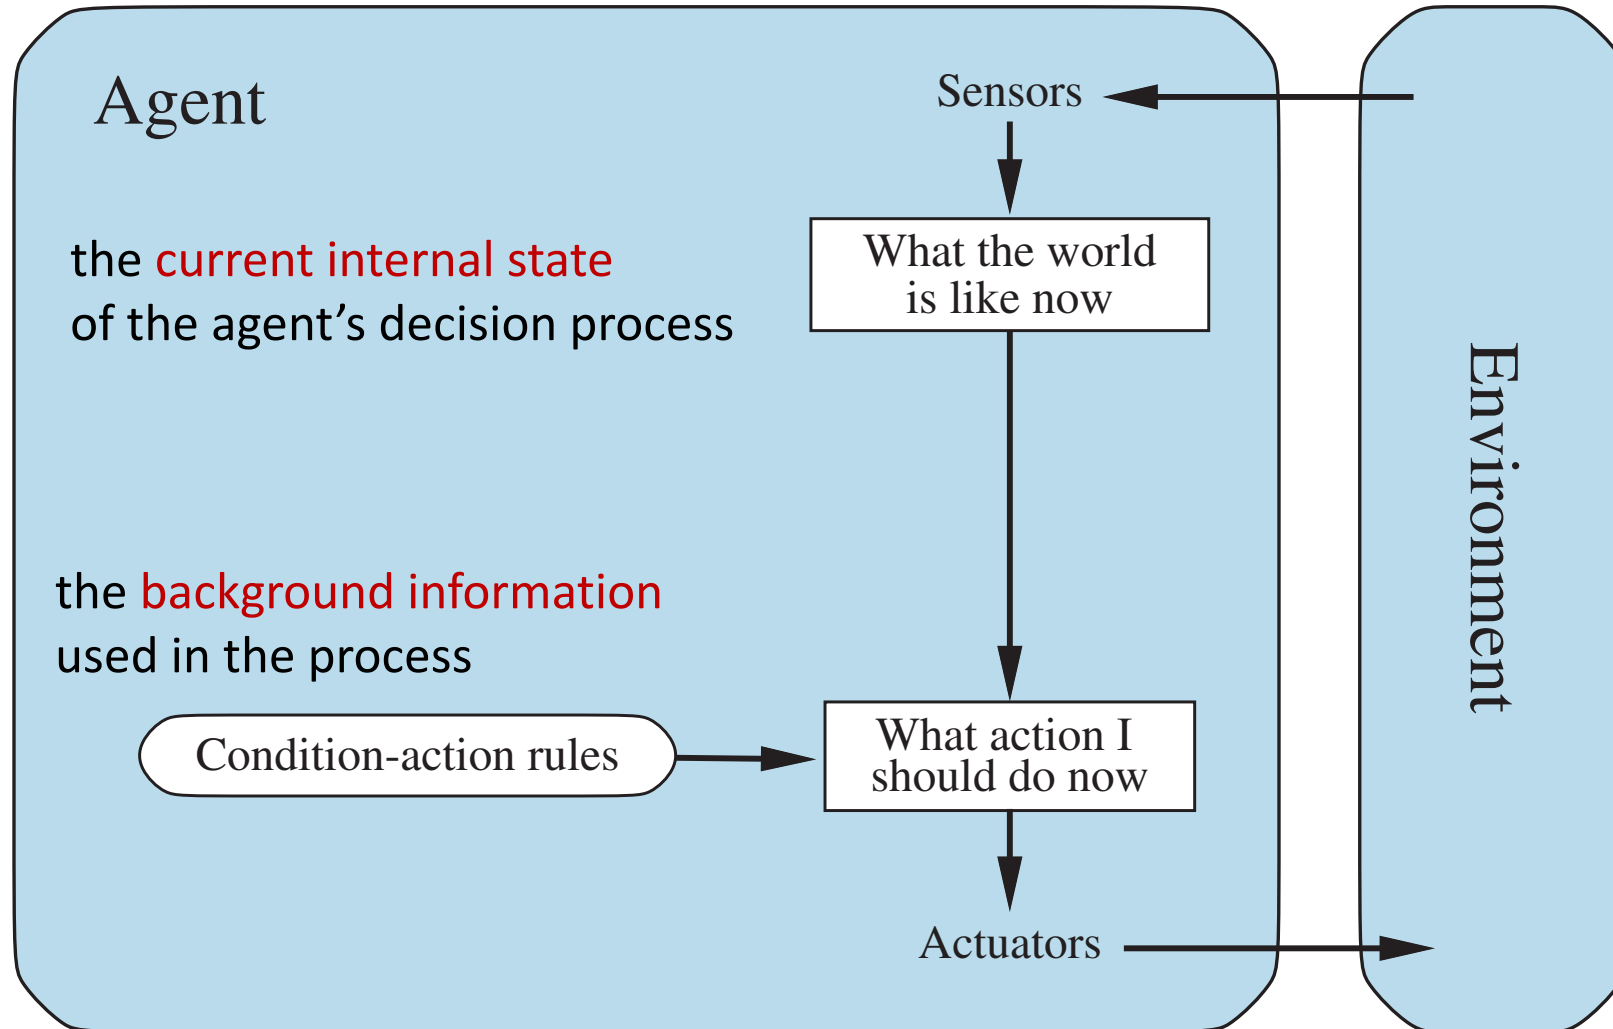
Schematic Diagram of a Simple Reflex Agent



Schematic diagram of a simple reflex agent



Schematic diagram of a simple reflex agent



A Simple Reflex Agent

It acts according to a rule whose condition matches the current state, as defined by the percept.


function SIMPLE-REFLEX-AGENT(*percept*) **returns** an action
persistent: *rules*, a set of condition–action rules

state \leftarrow INTERPRET-INPUT(*percept*)


rule \leftarrow RULE-MATCH(*state*, *rules*)

action \leftarrow *rule*.ACTION

return *action*



A Model-based Reflex Agent



A model-based reflex agent

It keeps track of the current state of the world,
using an internal model.

It then chooses an action in the same way as the reflex agent.


function MODEL-BASED-REFLEX-AGENT(*percept*) **returns** an action
persistent: *state*, the agent's current conception of the world state
transition_model, a description of how the next state depends on
the current state and action
sensor_model, a description of how the current world state is reflected
in the agent's percepts
rules, a set of condition–action rules
action, the most recent action, initially none

state ← UPDATE-STATE(*state*, *action*, *percept*, *transition_model*, *sensor_model*)

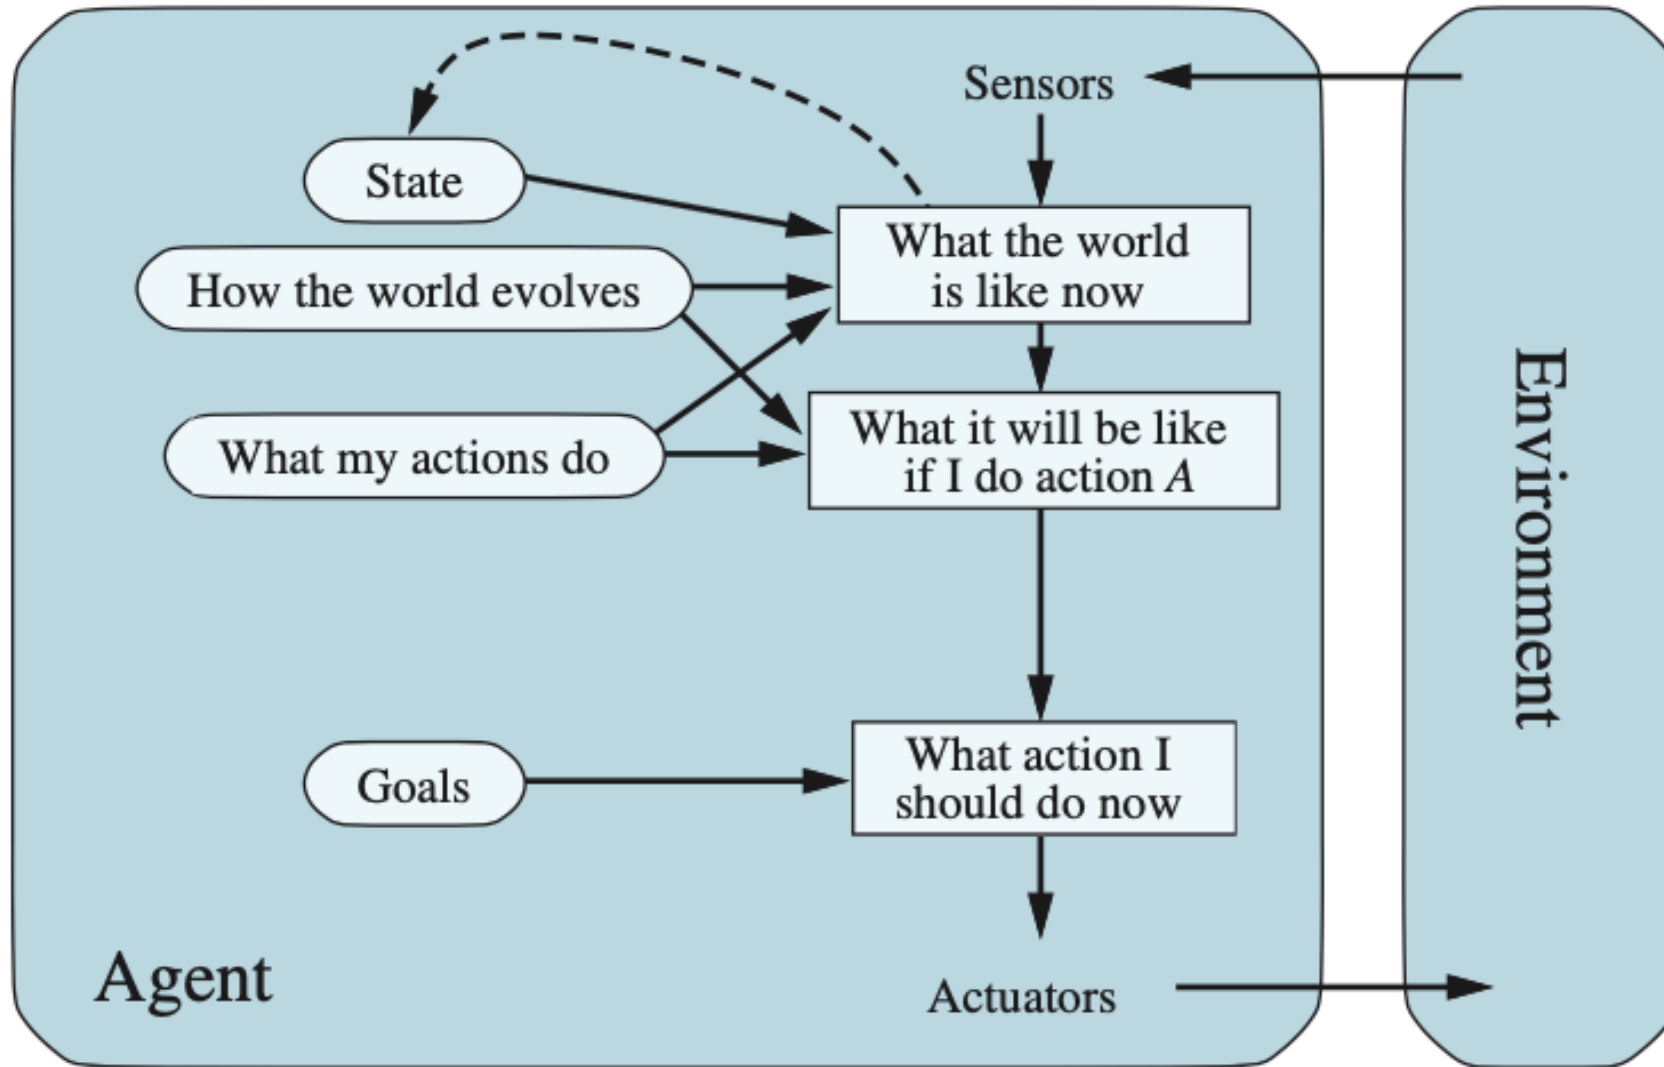
rule ← RULE-MATCH(*state*, *rules*)

action ← *rule*.ACTION

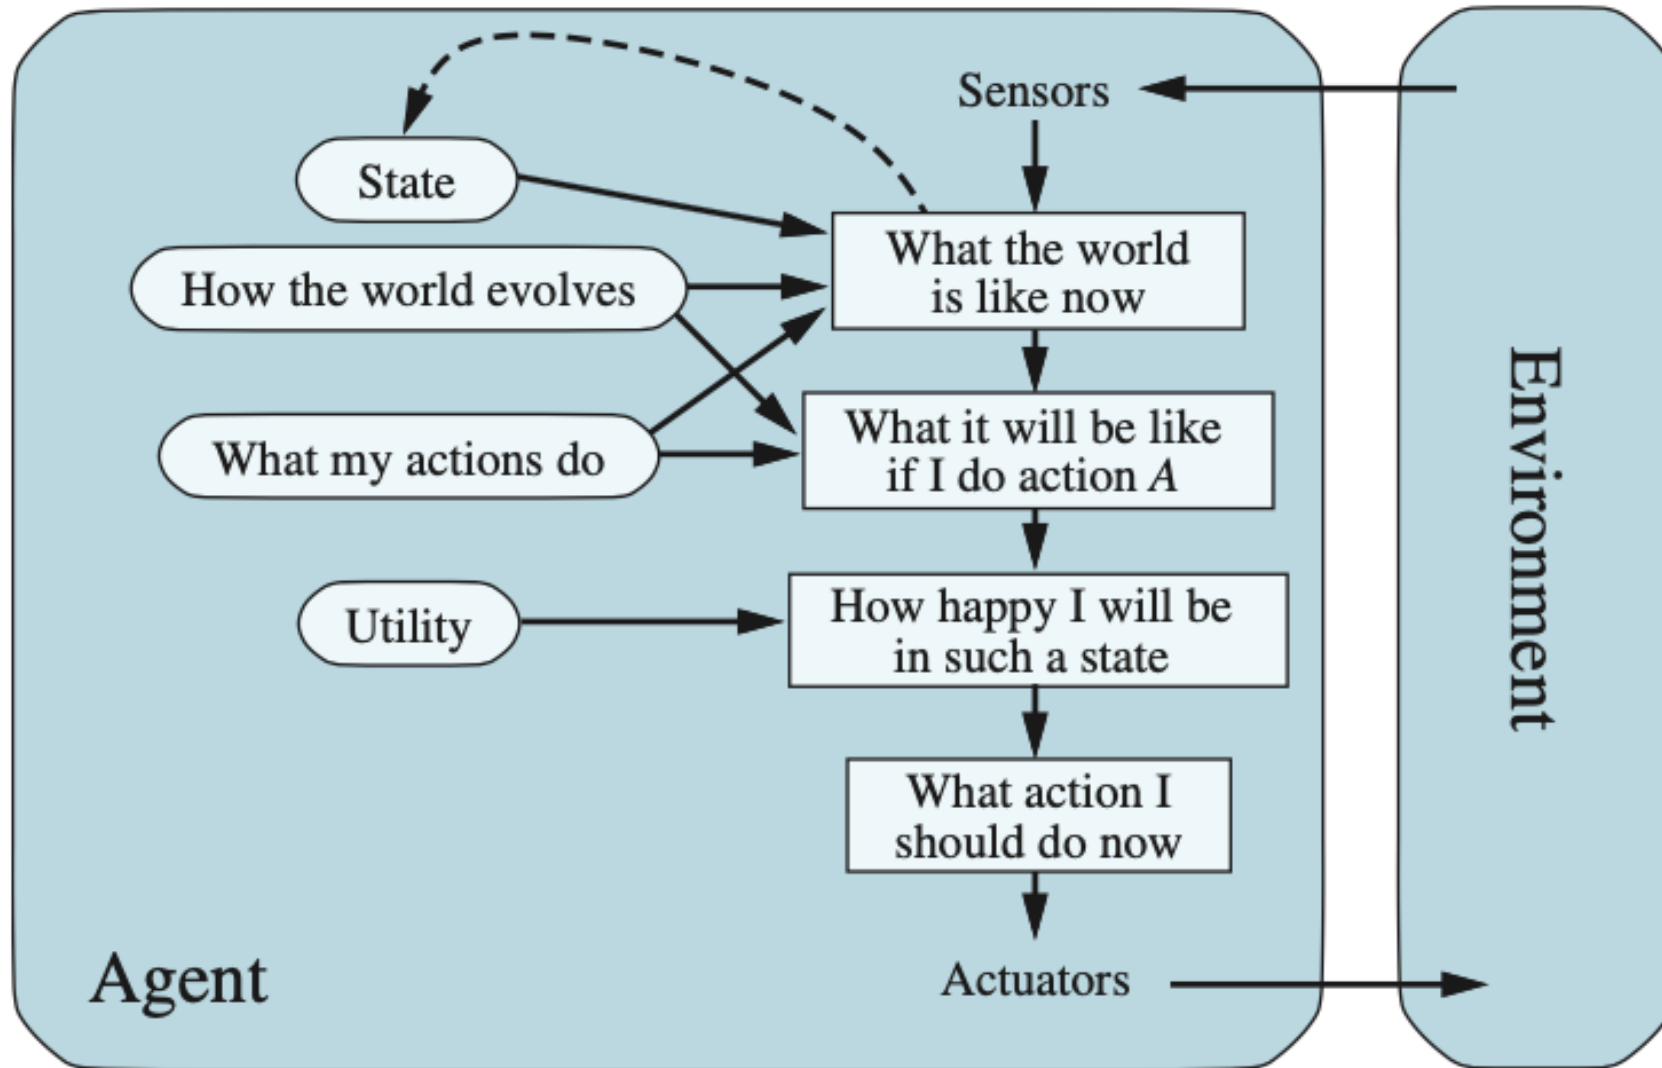
return *action*



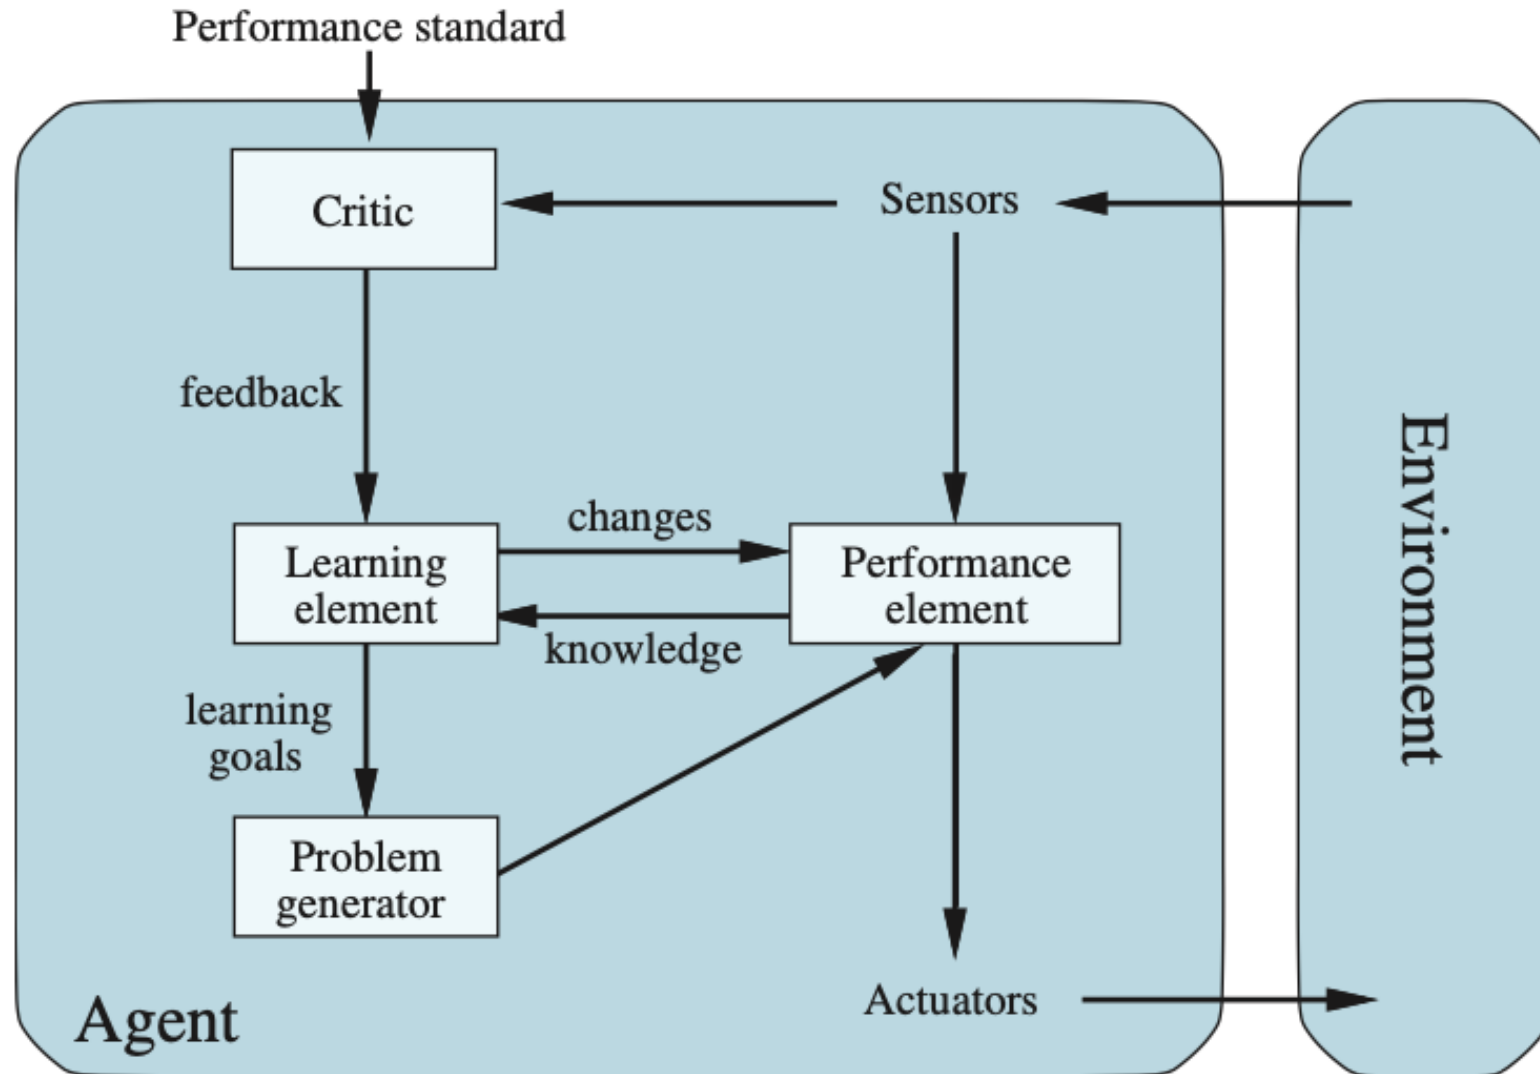
A model-based, goal-based agent




A model-based, utility-based agent




A general learning agent




Three ways to represent **states** and the **transitions** between them



(a) Atomic




(b) Factored



(c) Structured

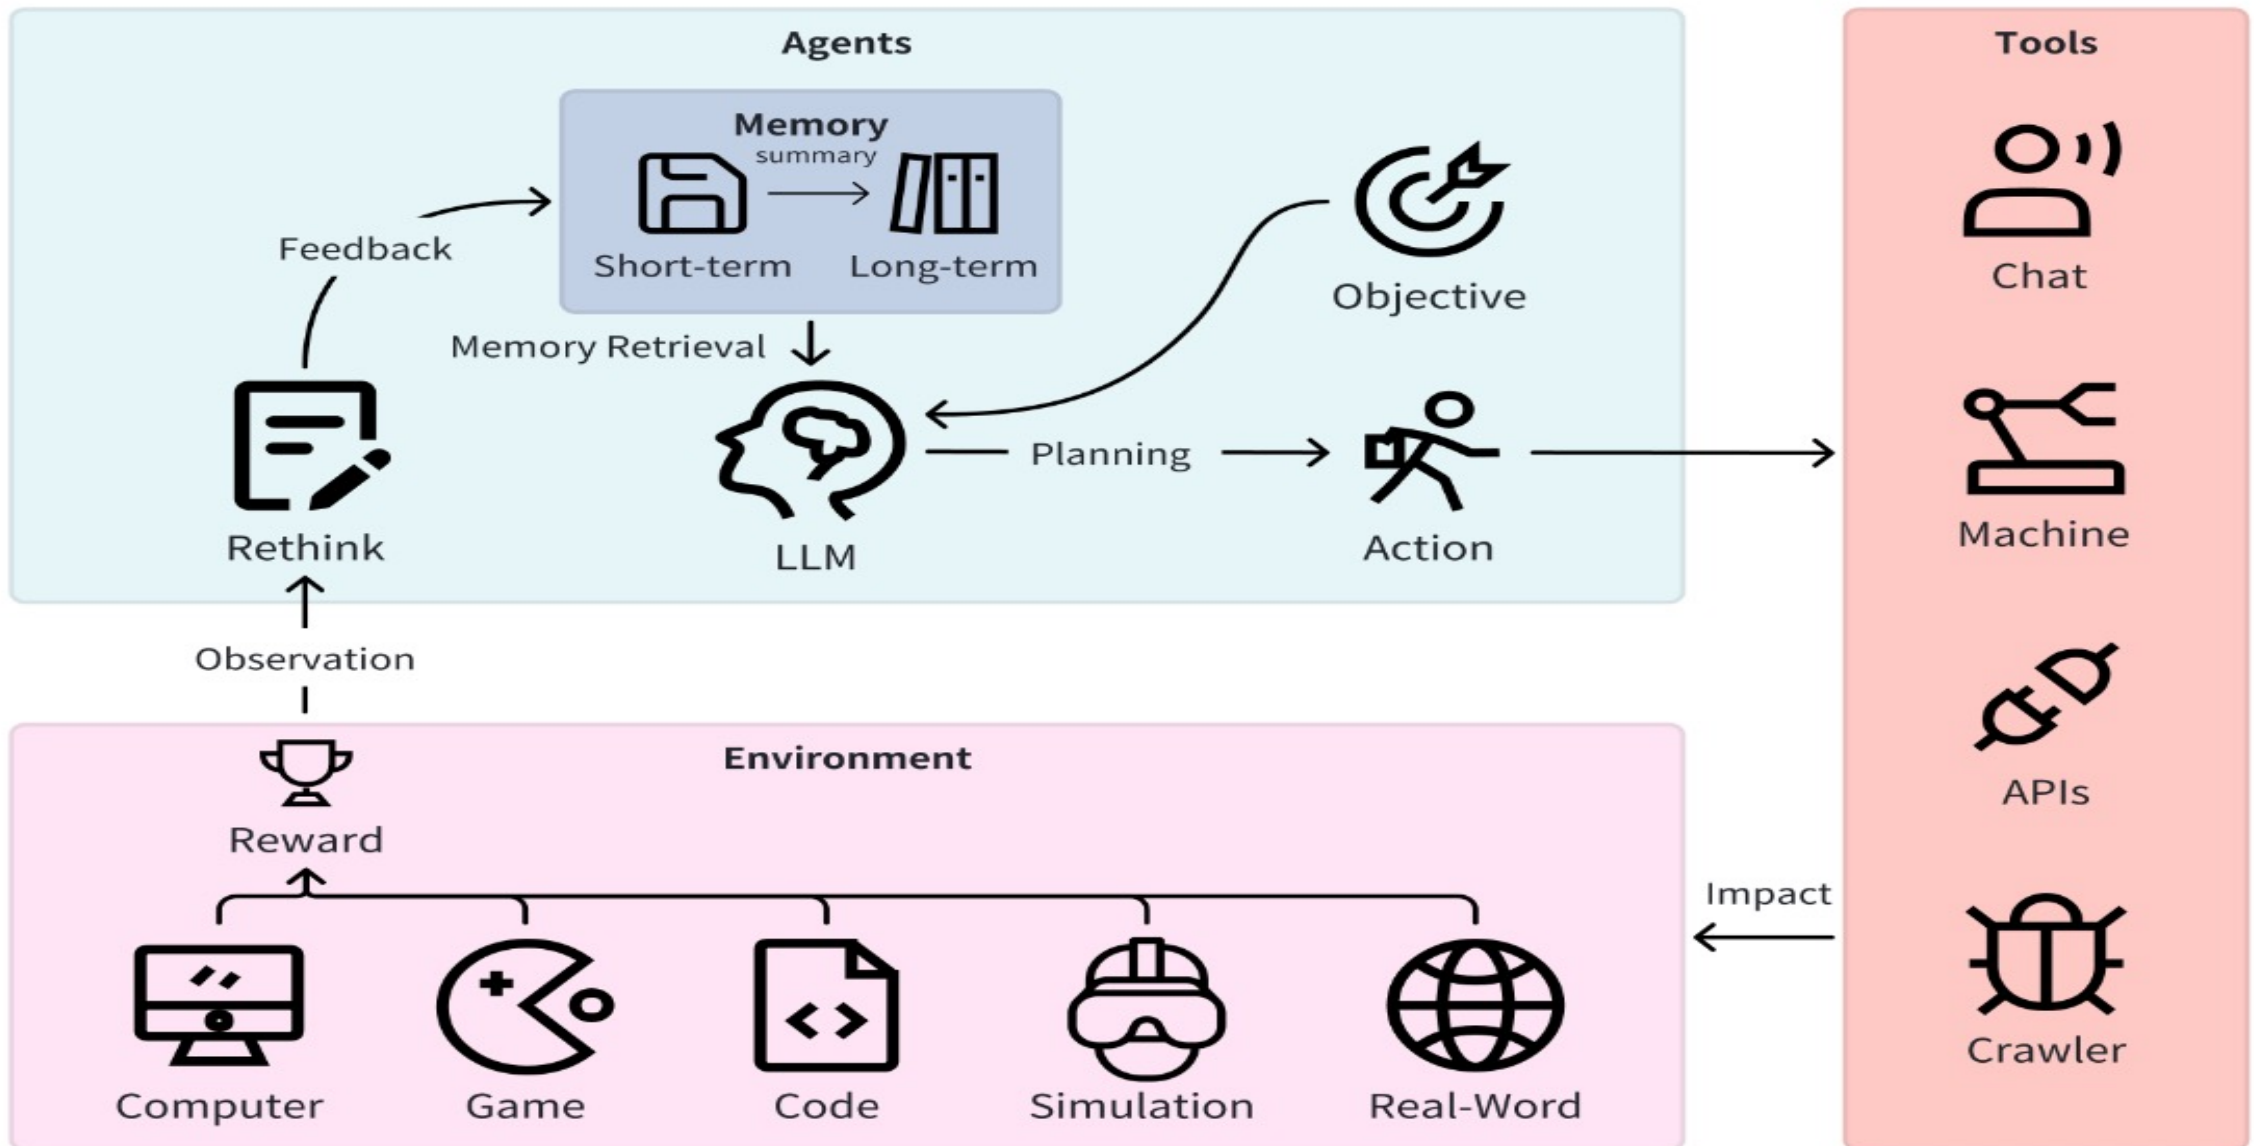
Large Language Model (LLM) based Agents




LLM-based Agents

- **Definition:** **AI agents** that use **Large Language Models** as their **core decision-making** mechanism
- **Key Features:**
 - **Natural language interface**
 - **Vast knowledge base**
 - **Ability to understand context and nuance**
 - **Generalize to new tasks with minimal additional training**

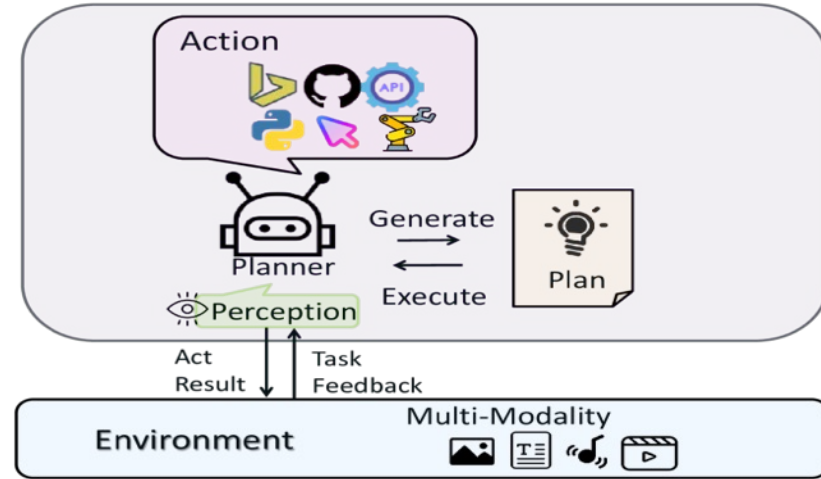
LLM-based Agents




Large Multimodal Agents (LMA)




(a)



(b)





(c)



(d)

Large Multimodal Agents (LMA)



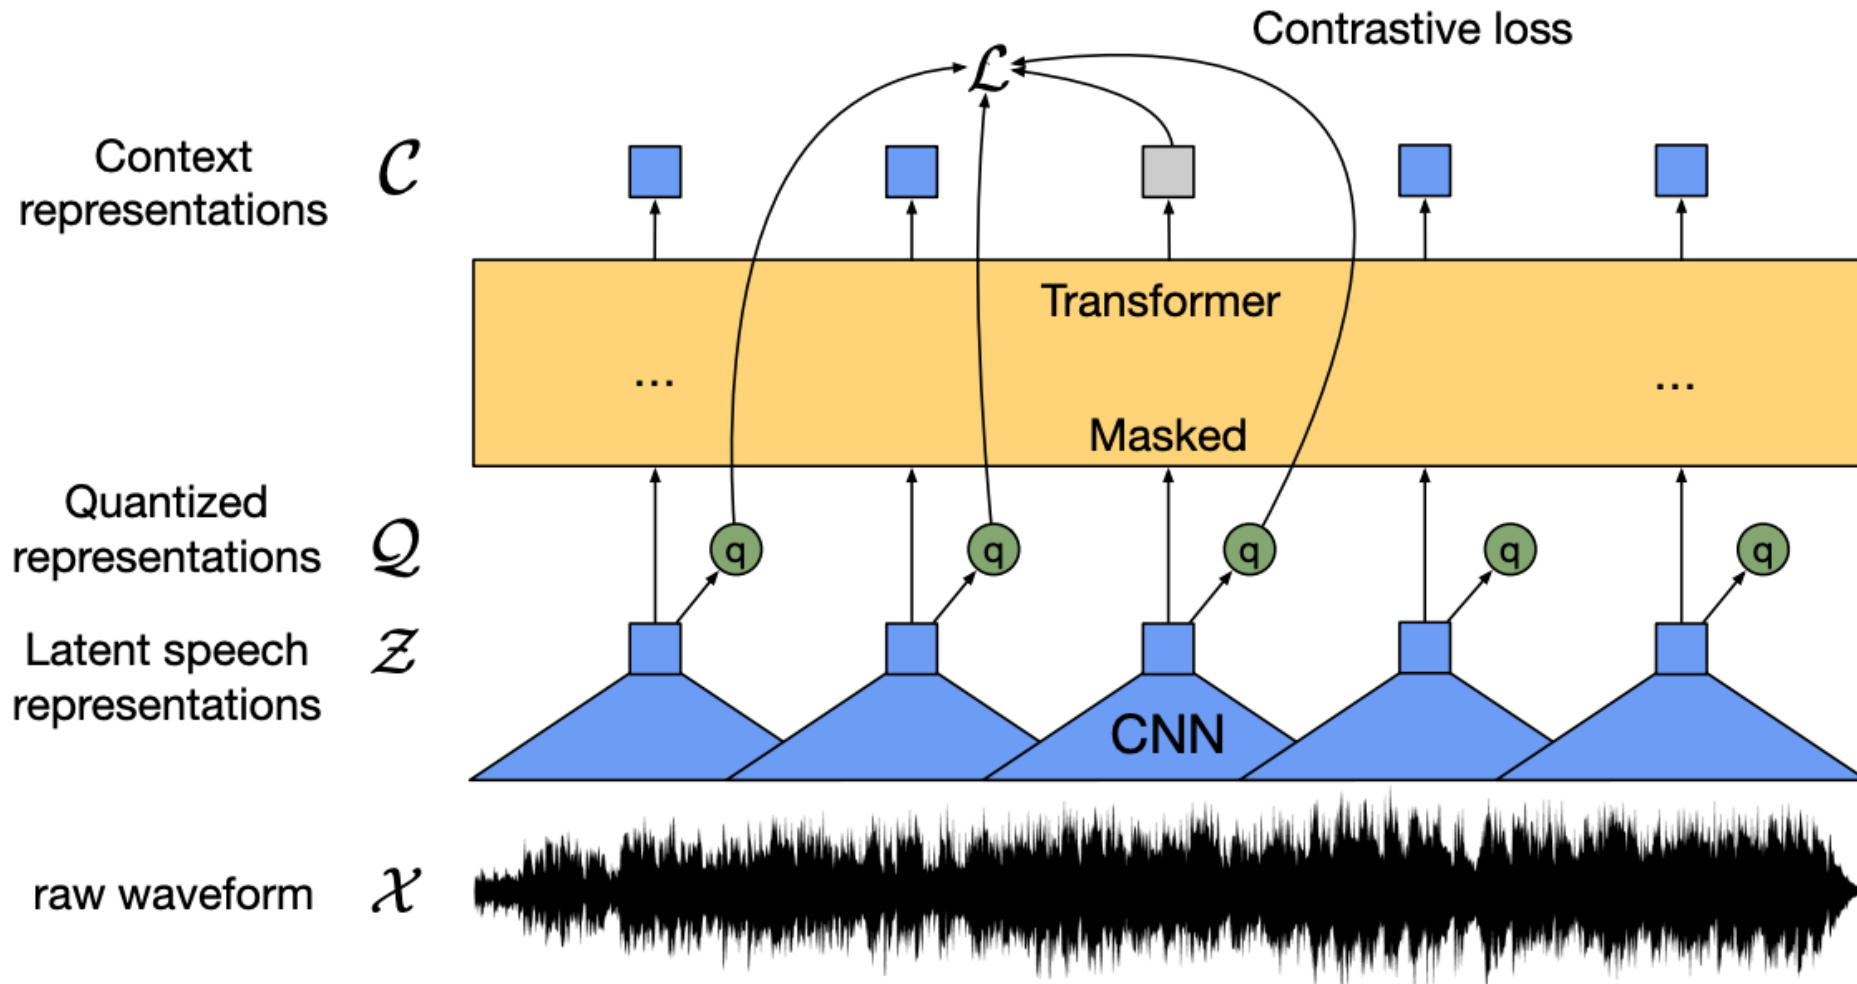
Artificial Intelligence:

6. Communicating, Perceiving, and Acting

- **Natural Language Processing**
- **Deep Learning for Natural Language Processing**
- **Computer Vision**
- **Robotics**

wav2vec 2.0:

A framework for self-supervised learning of speech representations



Source: Baevski, Alexei, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli.

"wav2vec 2.0: A framework for self-supervised learning of speech representations." Advances in Neural Information Processing Systems 33 (2020): 12449-12460.

Computer Vision: Image Classification, Object Detection, Object Instance Segmentation

Classification




CAT

Classification
+ Localization



CAT

Object
Detection



CAT, DOG, DUCK

Instance
Segmentation



CAT, DOG, DUCK


Single Objects

Multiple Objects


Computer Vision: Object Detection




(a) Object Classification



(b) Generic Object Detection (Bounding Box)



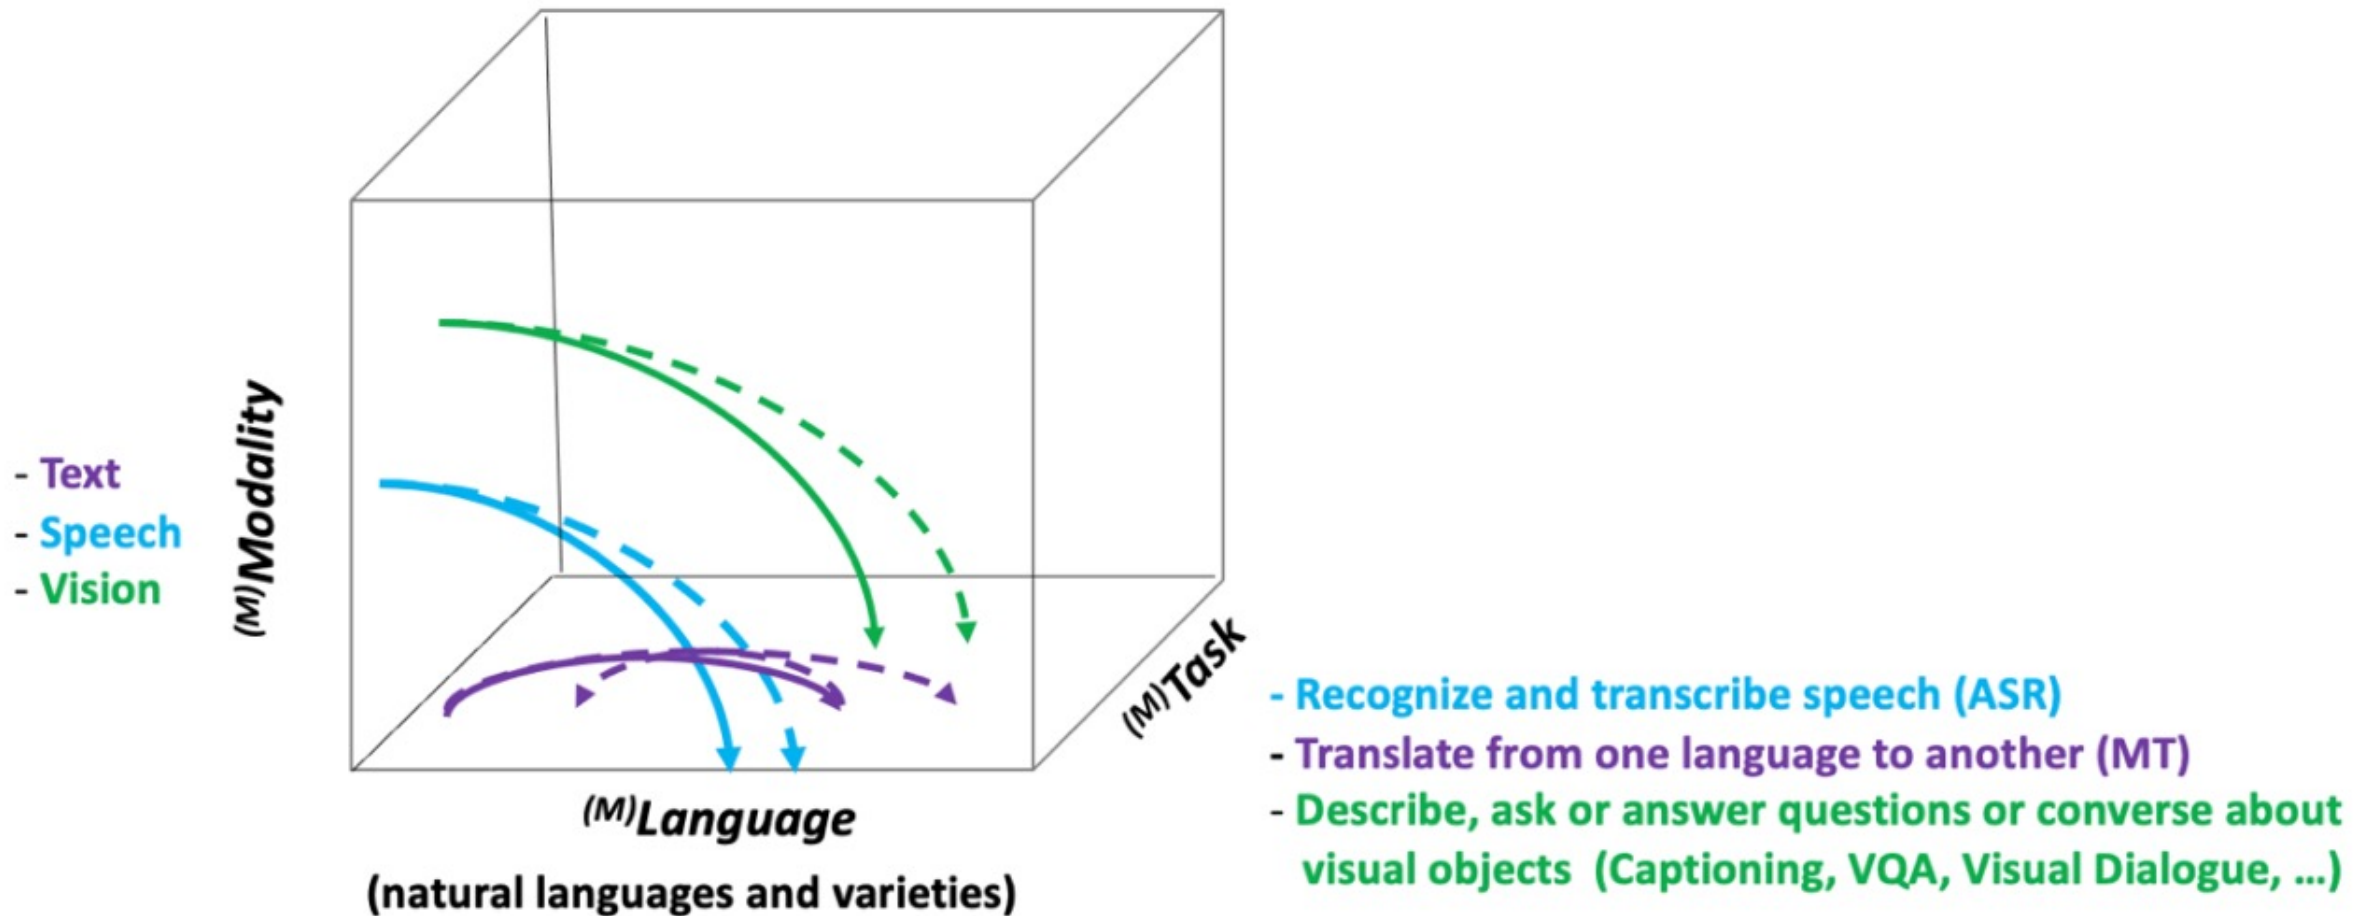
(c) Semantic Segmentation



(d) Object Instance Segmentation

NLG from a Multilingual, Multimodal and Multi-task perspective


Multi³(Natural Language) Generation



Source: Erdem, Erkut, Menekse Kuyu, Semih Yagcioglu, Anette Frank, Letitia Parcalabescu, Barbara Plank, Andrii Babii et al.

"Neural Natural Language Generation: A Survey on Multilinguality, Multimodality, Controllability and Learning." Journal of Artificial Intelligence Research 73 (2022): 1131-1207.

Text-and-Video Dialog Generation Models with Hierarchical Attention




Source: Erdem, Erkut, Menekse Kuyu, Semih Yagcioglu, Anette Frank, Letitia Parcalabescu, Barbara Plank, Andrii Babii et al.

"Neural Natural Language Generation: A Survey on Multilinguality, Multimodality, Controllability and Learning." Journal of Artificial Intelligence Research 73 (2022): 1131-1207.

FinBrain: when Finance meets AI 2.0

(Zheng et al., 2019)



Source: Xiao-lin Zheng, Meng-ying Zhu, Qi-bing Li, Chao-chao Chen, and Yan-chao Tan (2019), "Finbrain: When finance meets AI 2.0." Frontiers of Information Technology & Electronic Engineering 20, no. 7, pp. 914-924

Technology-driven Financial Industry Development

Development stage	Driving technology	Main landscape	Inclusive finance	Relationship between technology and finance
Fintech 1.0 (financial IT)	Computer	Credit card, ATM, and CRMS	Low	Technology as a tool
Fintech 2.0 (Internet finance)	Mobile Internet	Marketplace lending, third-party payment, crowdfunding, and Internet insurance	Medium	Technology- driven change
Fintech 3.0 (financial intelligence)	AI, Big Data, Cloud Computing, Blockchain	Intelligent finance	High	Deep fusion

Deep learning for financial applications: A survey

Applied Soft Computing (2020)

Source:


Ahmet Murat Ozbayoglu, Mehmet Ugur Gudelek, and Omer Berat Sezer (2020). "Deep learning for financial applications: A survey." Applied Soft Computing (2020): 106384.

**Financial
time series forecasting with
deep learning:
A systematic literature review:
2005–2019
Applied Soft Computing (2020)**

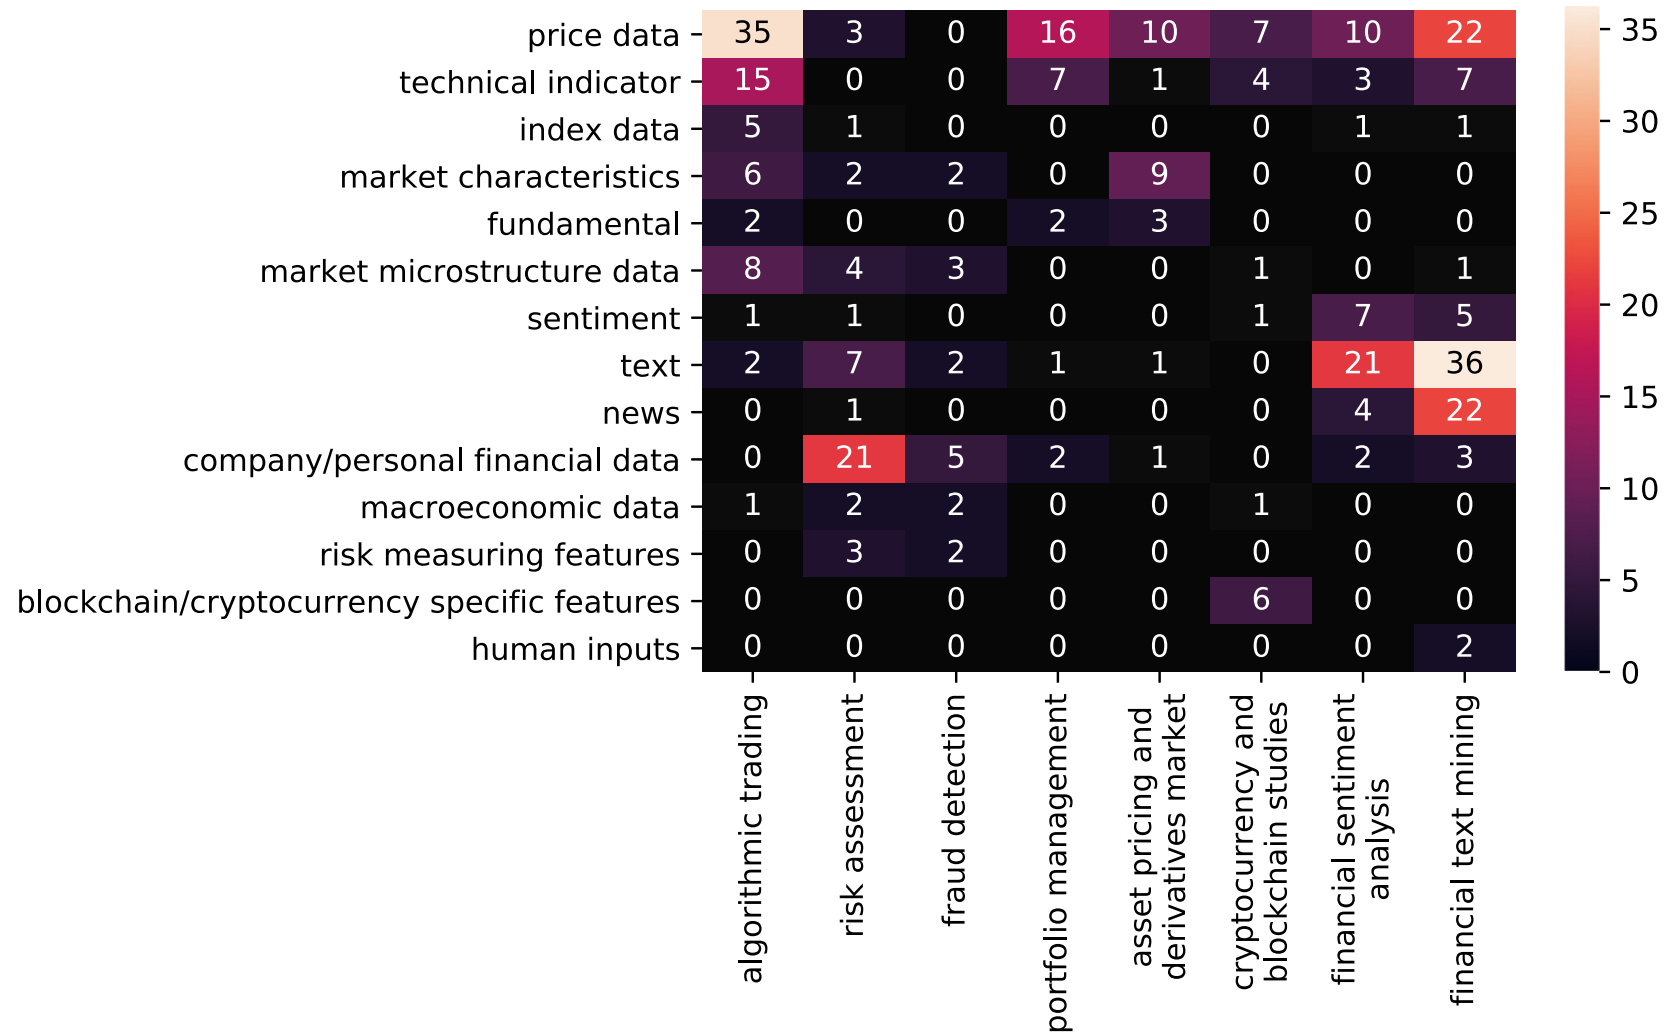
Source:

Omer Berat Sezer, Mehmet Ugur Gudelek, and Ahmet Murat Ozbayoglu (2020),
"Financial time series forecasting with deep learning: A systematic literature review:
2005–2019." *Applied Soft Computing* 90 (2020): 106181.

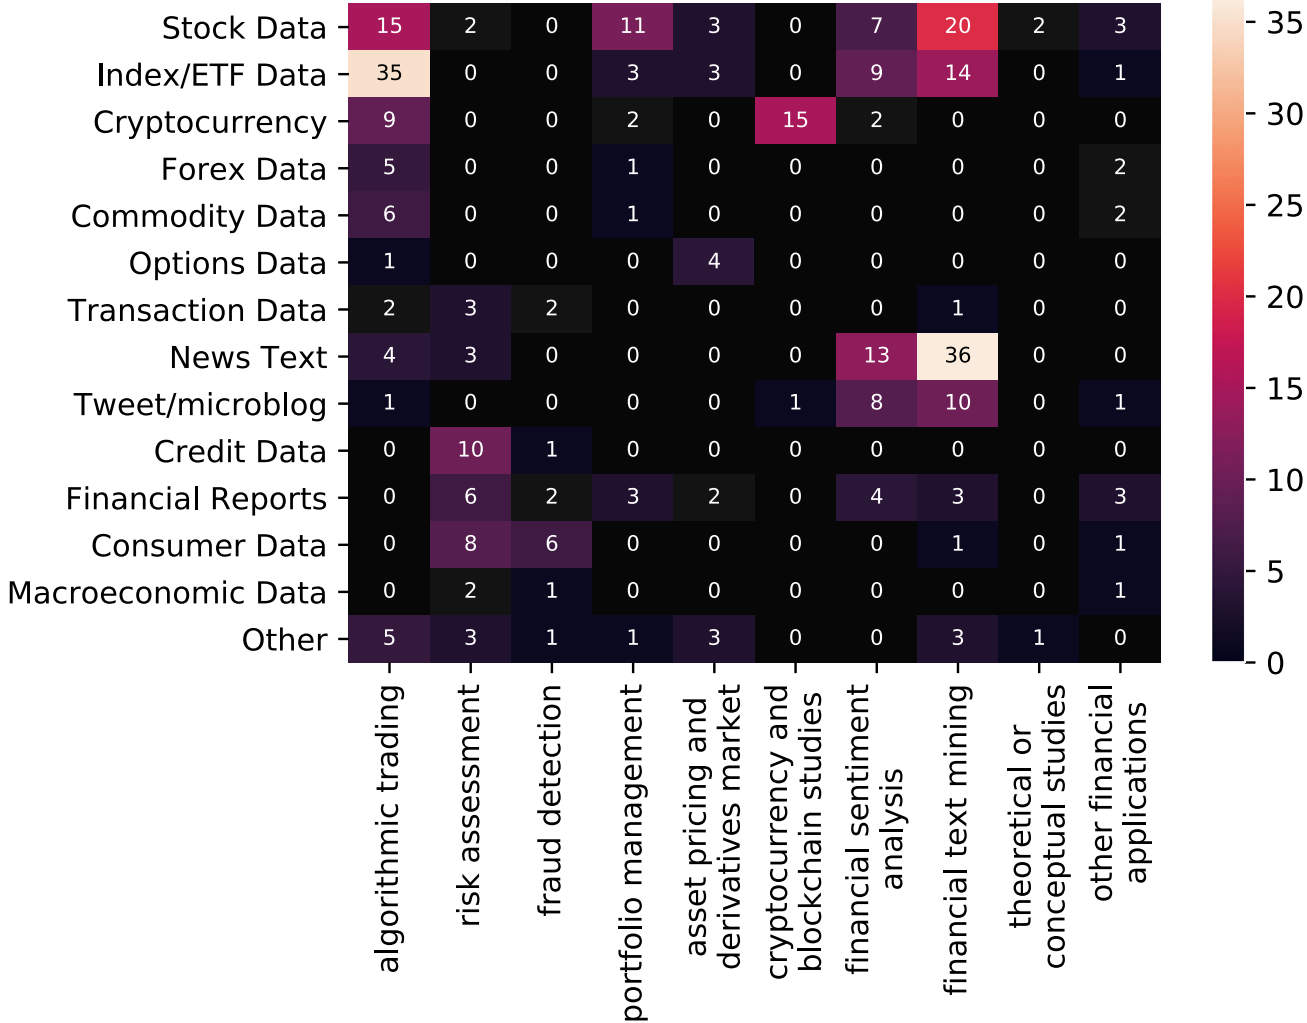
Deep learning for financial applications: Topic-Model Heatmap



Deep learning for financial applications: Topic-Feature Heatmap




Deep learning for financial applications: Topic-Dataset Heatmap



Source: Ahmet Murat Ozbayoglu, Mehmet Ugur Gudelek, and Omer Berat Sezer (2020). "Deep learning for financial applications: A survey." Applied Soft Computing (2020): 106384.

Financial time series forecasting with deep learning: Topic-model heatmap



Artificial Intelligence Problem Solving

Artificial Intelligence: A Modern Approach

1. Artificial Intelligence
2. Problem Solving
3. Knowledge and Reasoning
4. Uncertain Knowledge and Reasoning
5. Machine Learning
6. Communicating, Perceiving, and Acting
7. Philosophy and Ethics of AI

Artificial Intelligence:

2. Problem Solving

- Solving Problems by Searching
- Search in Complex Environments
- Adversarial Search and Games
- Constraint Satisfaction Problems

LLM-enhanced Problem Solving

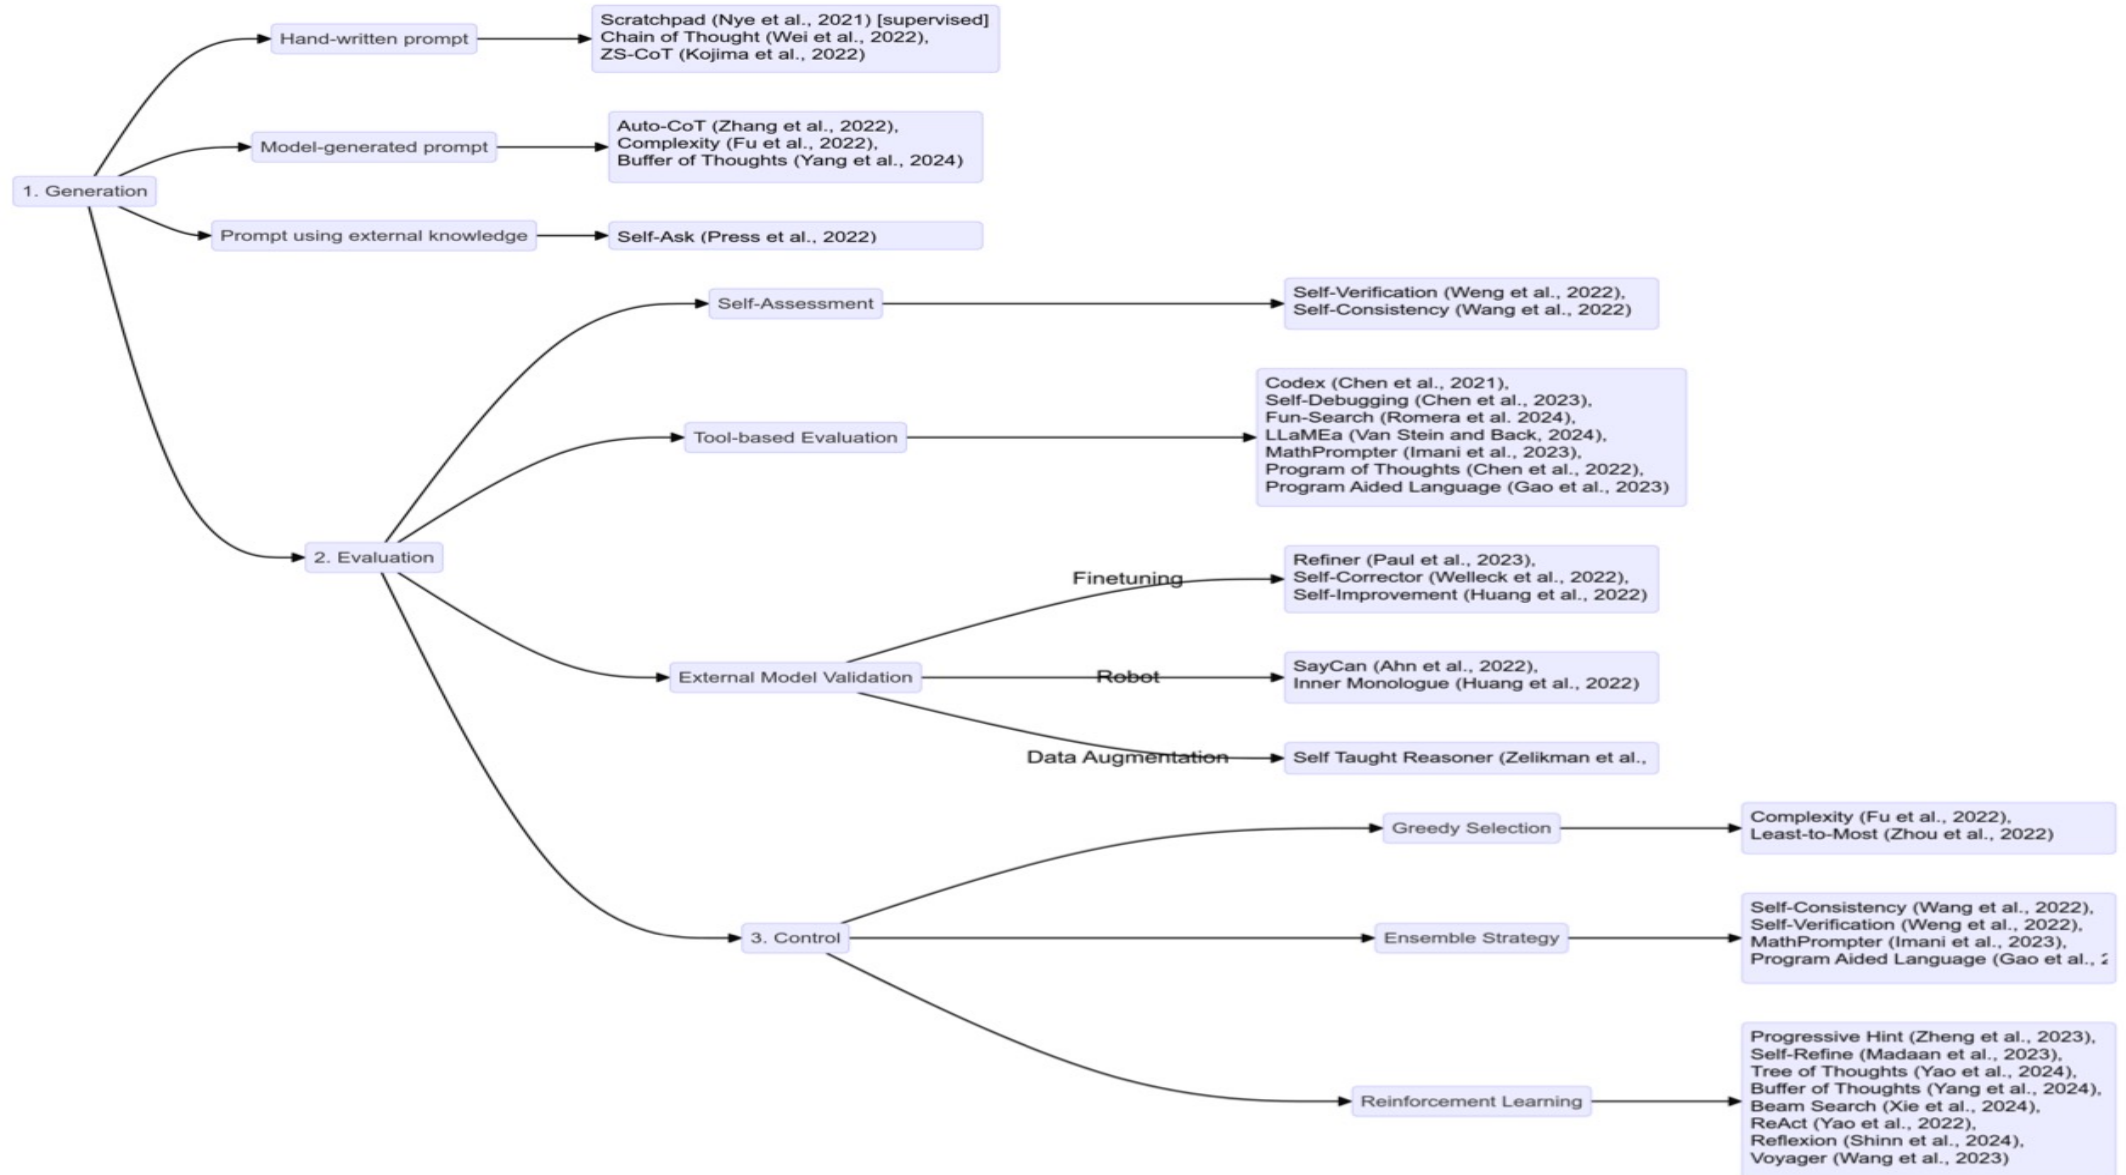
Problem Solving and LLM-enhanced Techniques

- **Traditional Search Algorithms**
 - **Breadth-first Search (BFS)**
 - **Depth-first Search (DFS)**
 - **A* Search**
- **LLM-enhanced Problem Solving**
 - **Chain-of-thought (CoT) prompting**
 - **Few-shot learning for problem decomposition**
 - **Integration with external tools and APIs**

LLM-based Agents for Complex Problem Solving

- **ReAct: Reasoning and Acting in Language Models**
- **MRKL (Modular Reasoning, Knowledge and Language) systems**
- **LLM-powered planning and decision making**
- **Chain-of-thought Prompting**
 - **Solving a complex problem using chain-of-thought prompting**

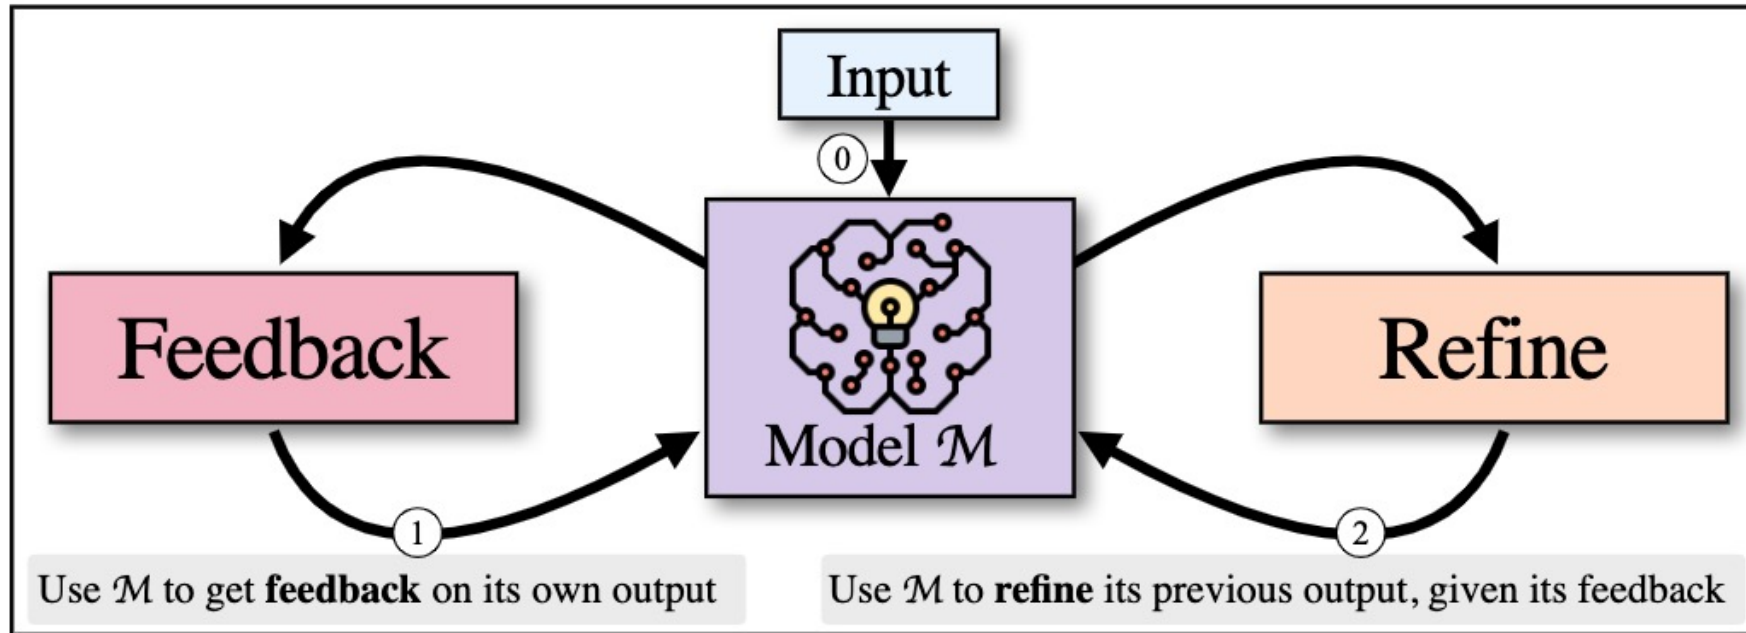
LLM-Reasoning Approaches: Prompt Generation, Evaluation, and Control



LLM-Reasoning Approaches: Prompt Generation, Evaluation, and Control

Approach	Domain	Step generation	Step evaluation	Step control
Scratchpad [Nye et al., 2021]	math word	hand-wr/supervised	-	greedy/1prompt
Chain-of-thought [Wei et al., 2022b]	math word	hand-written	-	greedy/1prompt
ZS-CoT [Kojima et al., 2022]	math word	hand-written	-	greedy/1prompt
Auto-CoT [Zhang et al., 2022]	math word	model-generated	-	clustering
Complexity [Fu et al., 2022]	math word	hand-written	self-consistency	greedy/1prompt
Self-ask [Press et al., 2022]	math word	external knowledge	LLM	multi-hop questions
Self-verification [Weng et al., 2022]	math word	hand-written	back-verify	ensemble
Self-consistency [Wang et al., 2022b]	math word	hand-written	majority	ensemble
Codex [Chen et al., 2021]	code	-	tool-based	-
Self-debugging [Chen et al., 2023]	code	hand-written	tool-based	greedy
Fun-search [Romera-Paredes et al., 2024]	code	hand-written	tool-based	evolutionary algorithm
LLaMEa [van Stein and Bäck, 2024]	code	hand-written	tool-based	evolutionary algorithm
MathPrompter [Imani et al., 2023]	math	hand-written	tool-based	ensemble
Program-of-thoughts [Chen et al., 2022]	math word	hand-written, Codex	Python+Consist.	decouple reason/compute
Program-aided-language [Gao et al., 2023]	math word	hand-written, Codex	NLP/Python	ensemble
Refiner [Paul et al., 2023]	math word	finetune	critic model	gen/crit feedback
Self-corrector [Welleck et al., 2022]	math word	finetune	corrector model	gen/corr feedback
Self-improvement [Huang et al., 2022a]	math word	finetune	self-assessment	CoT/consistency
Say-can [Ahn et al., 2022]	robot	model-generated	external model	greedy
Inner-monologue [Huang et al., 2022b]	robot	hand-written	various	greedy
Self-taught-reasoner [Zelikman et al., 2022]	math word	finetune	augmentation	greedy/feedback
Least-to-most [Zhou et al., 2022]	math word	hand-written	self-assessment	curriculum
Progressive-hint [Zheng et al., 2023]	math word	model-generated	self-assessment	stable prompt
Self-refine [Madaan et al., 2023]	math word	model-generated	self-assessment	greedy/feedback
Tree-of-thoughts [Yao et al., 2024]	puzzles	model-generated	self-assessment	BFS/DFS
Buffer-of-thoughts [Yang et al., 2024]	math word	thought template	self-assessment	buffer manager
Beam-search [Xie et al., 2024]	math word	model-generated	self-assessment	Beam Search
ReAct [Yao et al., 2022]	action	external knowledge	self-assessment	reinforcement learning
Reflexion [Shinn et al., 2024]	decision	model-generated	ext model	reinforcement learning
Voyager [Wang et al., 2023]	Minecraft	model-generated	Minecraft	reinforcement learning

Self-Refine: Iterative refinement with self-feedback



(a) **Dialogue:** x, y_t

(b) **FEEDBACK** fb

(c) **REFINE** y_{t+1}

User: I am interested in playing Table tennis.

Response: I'm sure it's a great way to socialize, stay active

Engaging: Provides no information about table tennis or how to play it.

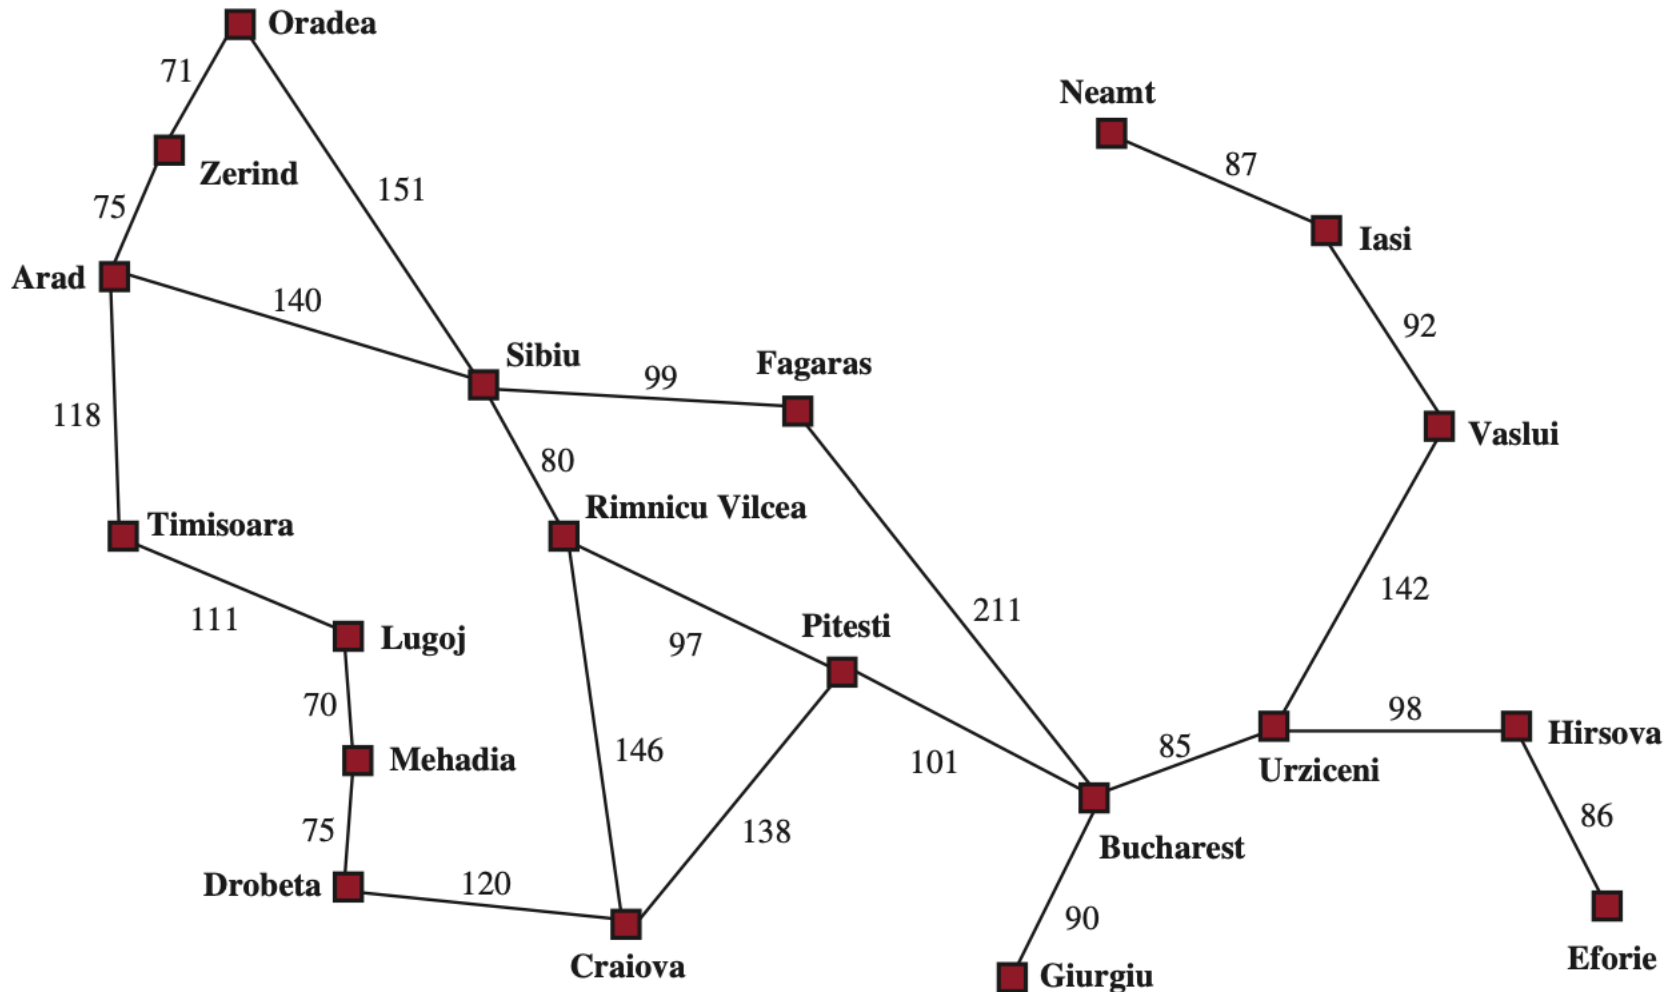
User understanding: Lacks understanding of user's needs and state of mind.

Response (refined): That's great to hear (...) ! It's a fun sport requiring quick reflexes and good hand-eye coordination. Have you played before, or are you looking to learn?

Solving Problems by Searching


AI: Solving Problems by Searching

A simplified road map of part of Romania, with road distances in miles.



The state-space graph for the two-cell vacuum world

There are 8 states and three actions for each state:
L = Left, R = Right, S = Suck.



A typical instance of the 8-puzzle

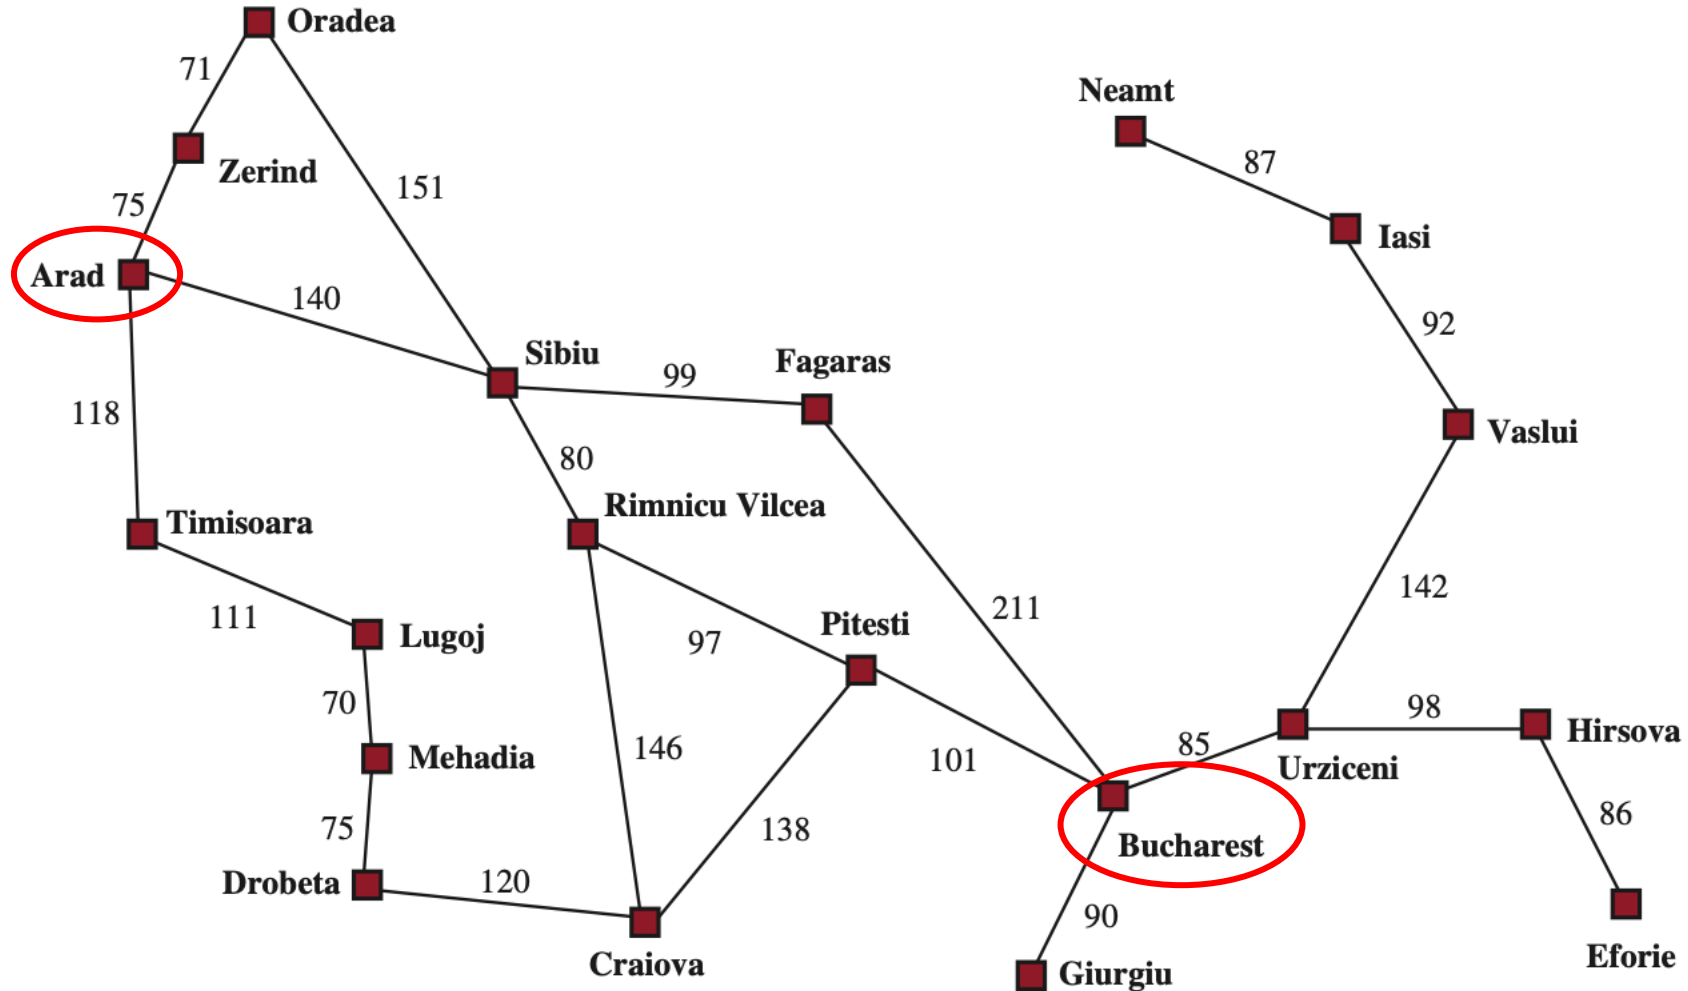
7	2	4
5		6
8	3	1

Start State

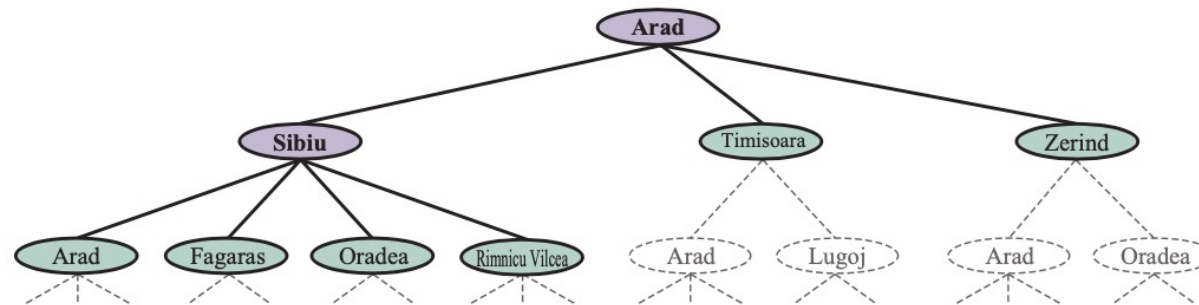
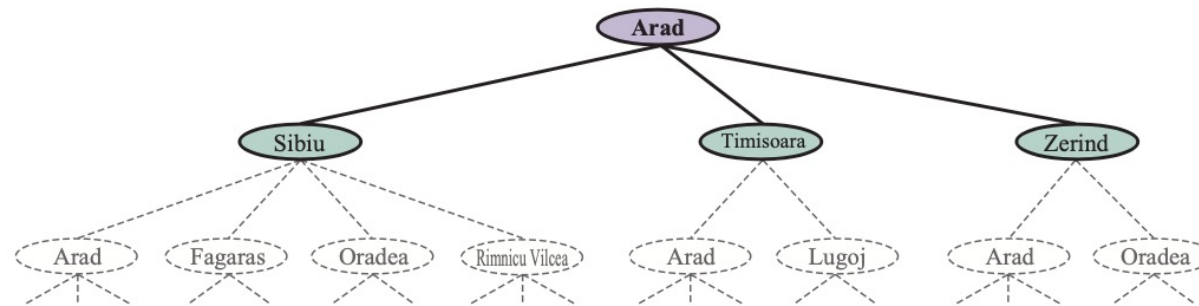
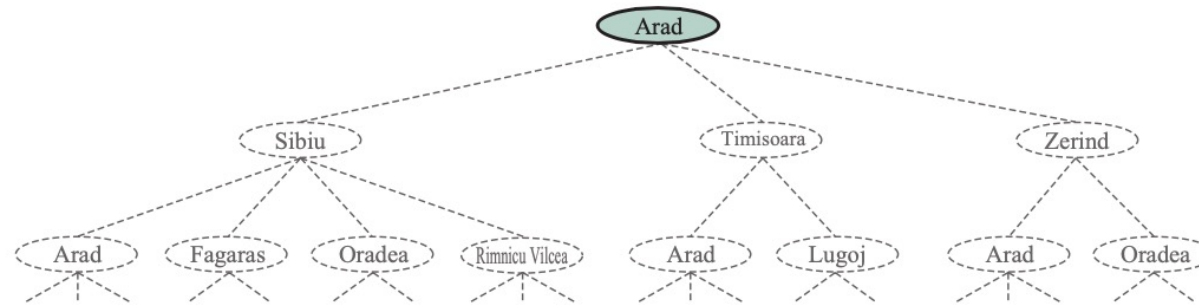
	1	2
3	4	5
6	7	8

Goal State

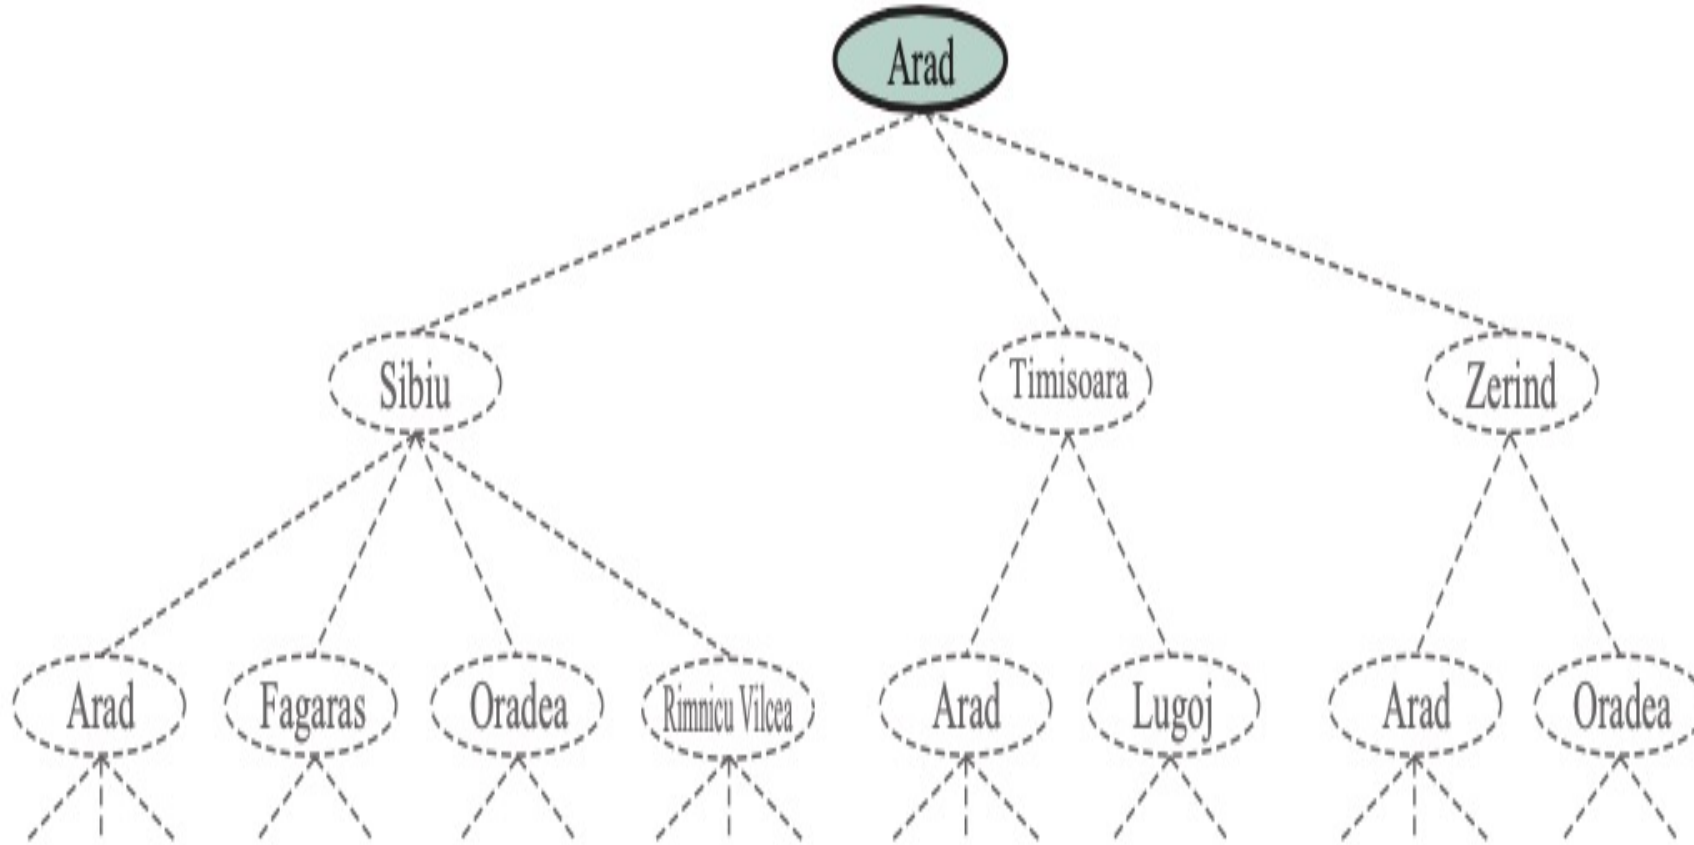
Arad to Bucharest



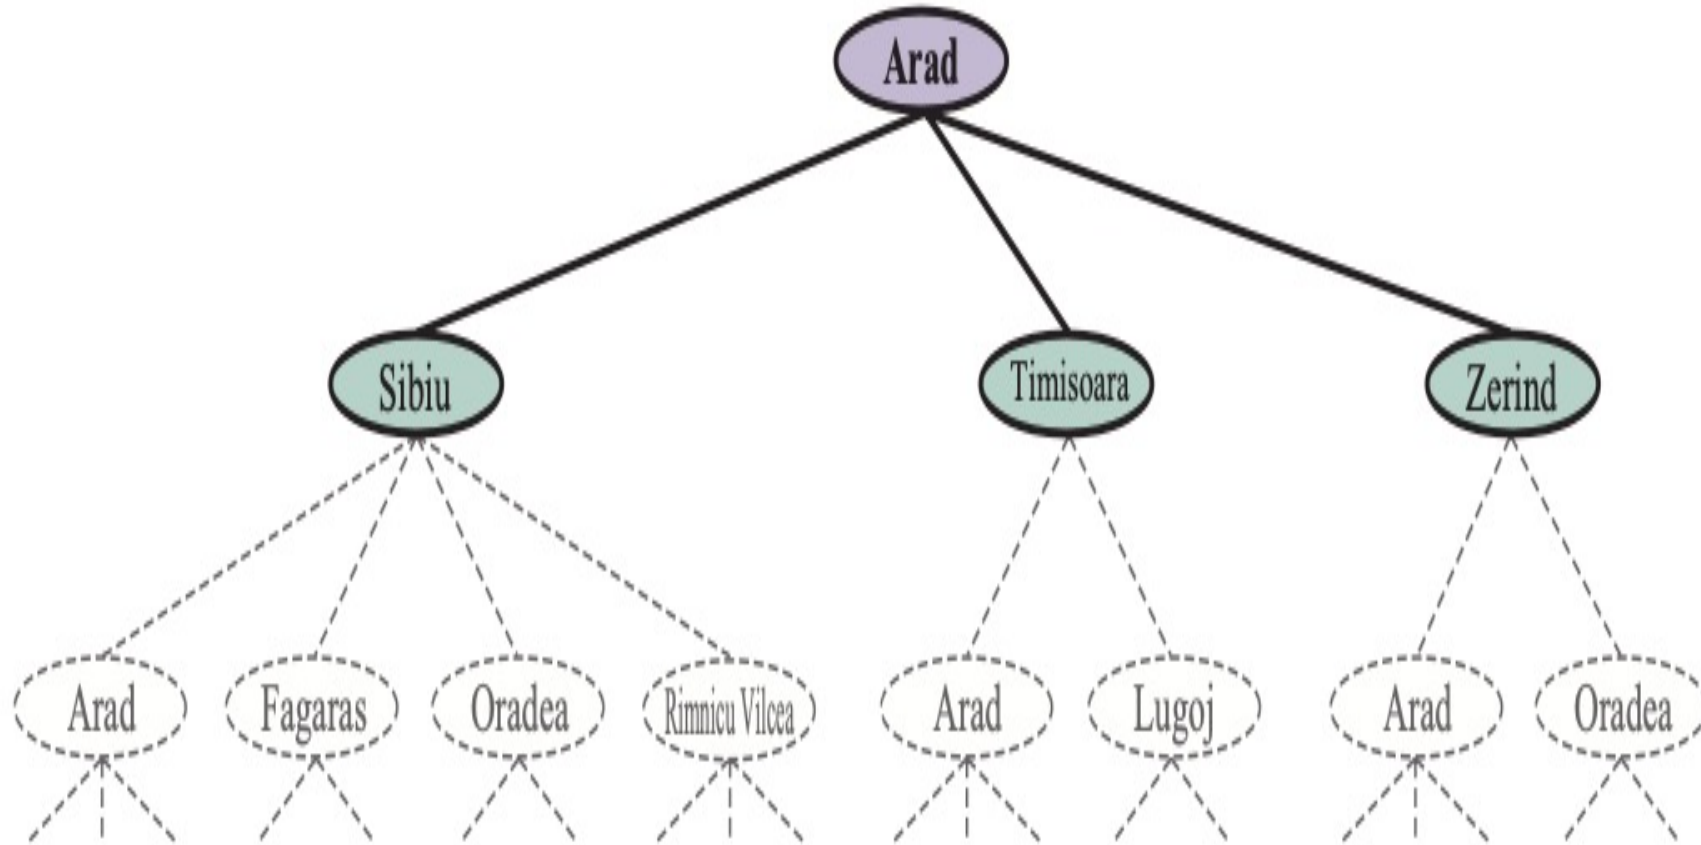
Three partial search trees for finding a route from Arad to Bucharest




Three partial search trees for finding a route from Arad to Bucharest




Three partial search trees for finding a route from Arad to Bucharest



Three partial search trees for finding a route from Arad to Bucharest




A sequence of search trees generated by a graph search on the Romania problem




The Separation Property of Graph Search

illustrated on a rectangular-grid problem


The frontier (green) separates the interior (lavender) from the exterior (faint dashed)



(a)



(b)



(c)





The Best-First Search (BFS) Algorithm

```
function BEST-FIRST-SEARCH(problem, f) returns a solution node or failure  
  node  $\leftarrow$  NODE(STATE=problem.INITIAL)  
  frontier  $\leftarrow$  a priority queue ordered by f, with node as an element  
  reached  $\leftarrow$  a lookup table, with one entry with key problem.INITIAL and value node  
  while not IS-EMPTY(frontier) do  
    node  $\leftarrow$  POP(frontier)  
    if problem.IS-GOAL(node.STATE) then return node  
    for each child in EXPAND(problem, node) do  
      s  $\leftarrow$  child.STATE  
      if s is not in reached or child.PATH-COST < reached[s].PATH-COST then  
        reached[s]  $\leftarrow$  child  
        add child to frontier  
  return failure
```

```
function EXPAND(problem, node) yields nodes  
  s  $\leftarrow$  node.STATE  
  for each action in problem.ACTIONS(s) do  
    s'  $\leftarrow$  problem.RESULT(s, action)  
    cost  $\leftarrow$  node.PATH-COST + problem.ACTION-COST(s, action, s')  
    yield NODE(STATE=s', PARENT=node, ACTION=action, PATH-COST=cost)
```


Breadth-First Search on a Simple Binary Tree

Bread-First Search (BFS)




Breadth-First Search on a Simple Binary Tree

**Bread-First
Search
(BFS)**




Breadth-First Search on a Simple Binary Tree

**Bread-First
Search
(BFS)**




Breadth-First Search on a Simple Binary Tree

**Bread-First
Search
(BFS)**



Breadth-First Search on a Simple Binary Tree

**Bread-First
Search
(BFS)**




Breadth-First Search and Uniform-Cost Search Algorithms

function BREADTH-FIRST-SEARCH(*problem*) **returns** a solution node or *failure*
 node ← NODE(*problem*.INITIAL)
 if *problem*.IS-GOAL(*node*.STATE) **then return** *node*
 frontier ← a FIFO queue, with *node* as an element
 reached ← {*problem*.INITIAL}
 while not IS-EMPTY(*frontier*) **do**
 node ← POP(*frontier*)
 for each *child* **in** EXPAND(*problem*, *node*) **do**
 s ← *child*.STATE
 if *problem*.IS-GOAL(*s*) **then return** *child*
 if *s* is not in *reached* **then**
 add *s* to *reached*
 add *child* to *frontier*
 return *failure*

function UNIFORM-COST-SEARCH(*problem*) **returns** a solution node, or *failure*
 return BEST-FIRST-SEARCH(*problem*, PATH-COST)


Part of the Romania State Space

Uniform-Cost Search



Depth-First Search (DFS)

Depth-First Search (DFS)



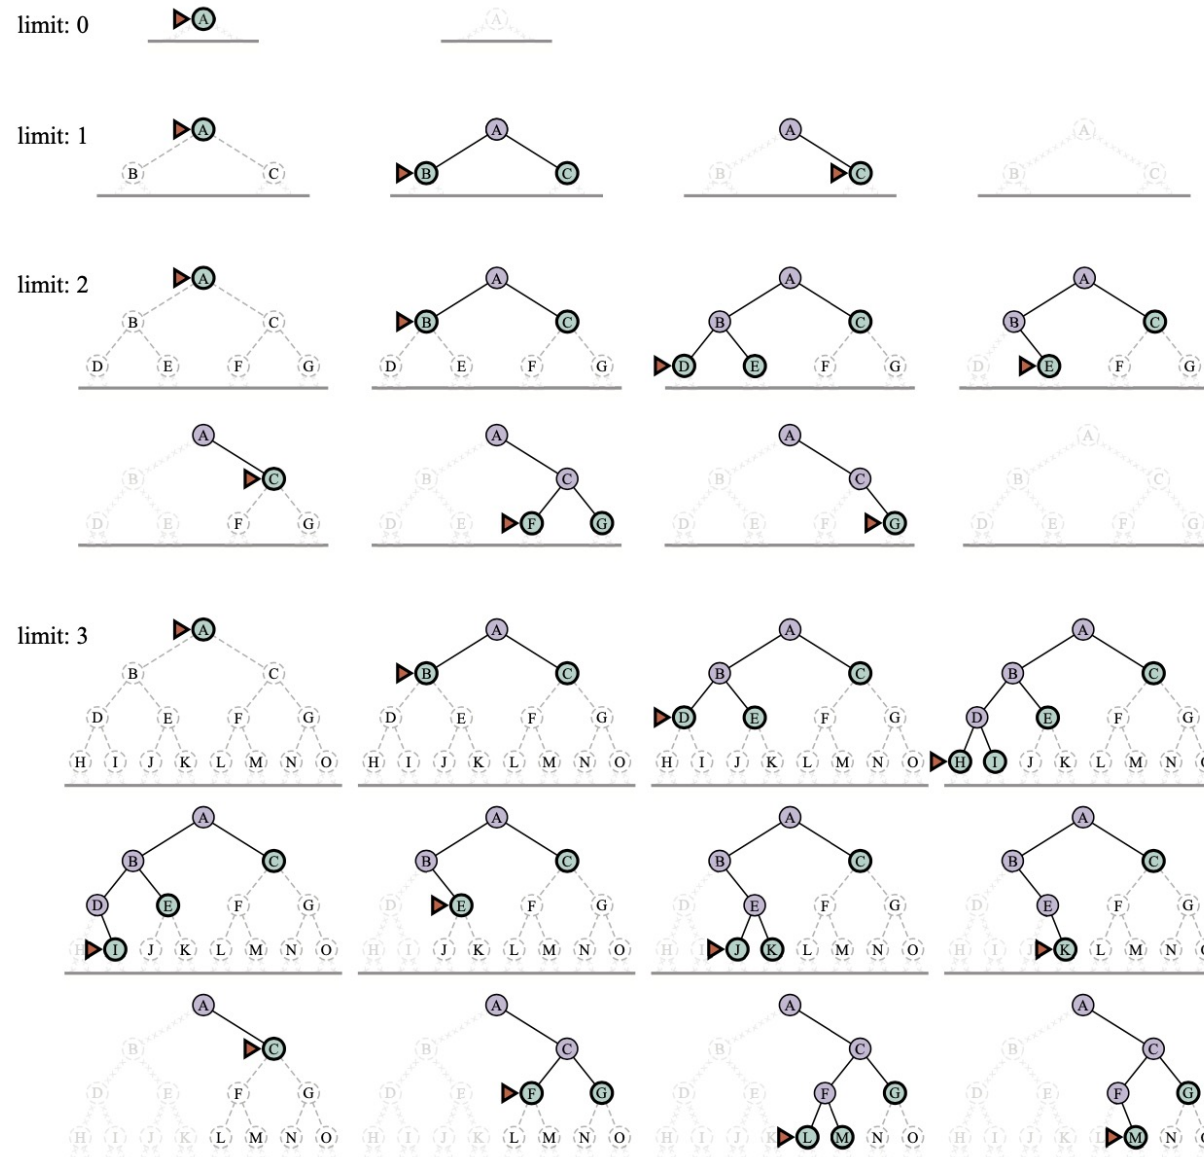
Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Iterative deepening and depth-limited tree-like search

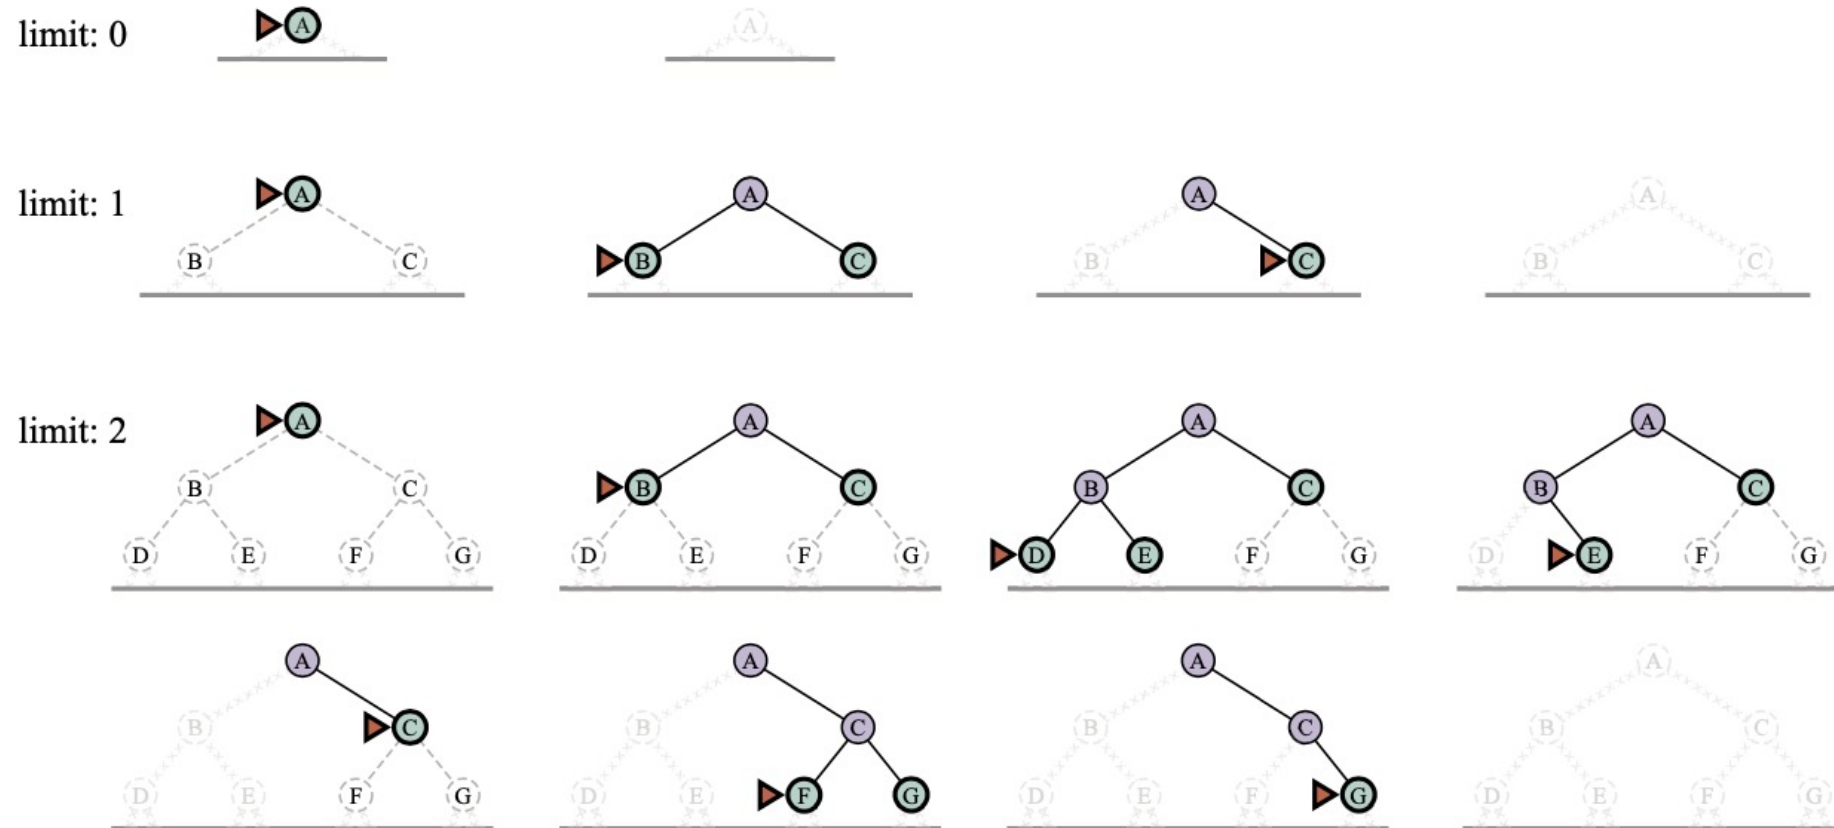
function ITERATIVE-DEEPENING-SEARCH(*problem*) **returns** a solution node or *failure*
 for *depth* = 0 **to** ∞ **do**
 result \leftarrow DEPTH-LIMITED-SEARCH(*problem*, *depth*)
 if *result* \neq *cutoff* **then return** *result*

function DEPTH-LIMITED-SEARCH(*problem*, ℓ) **returns** a node or *failure* or *cutoff*
 frontier \leftarrow a LIFO queue (stack) with NODE(*problem*.INITIAL) as an element
 result \leftarrow *failure*
 while not IS-EMPTY(*frontier*) **do**
 node \leftarrow POP(*frontier*)
 if *problem*.IS-GOAL(*node*.STATE) **then return** *node*
 if DEPTH(*node*) > ℓ **then**
 result \leftarrow *cutoff*
 else if not IS-CYCLE(*node*) **do**
 for each *child* **in** EXPAND(*problem*, *node*) **do**
 add *child* to *frontier*
 return *result*

Four iterations of iterative deepening search




Four iterations of iterative deepening search



Four iterations of iterative deepening search

limit: 3



Bidirectional Best-First Search

keeps two frontiers and two tables of reached states

```
function BIBF-SEARCH( $problem_F, f_F, problem_B, f_B$ ) returns a solution node, or failure  
   $node_F \leftarrow$  NODE( $problem_F$ .INITIAL) // Node for a start state  
   $node_B \leftarrow$  NODE( $problem_B$ .INITIAL) // Node for a goal state  
   $frontier_F \leftarrow$  a priority queue ordered by  $f_F$ , with  $node_F$  as an element  
   $frontier_B \leftarrow$  a priority queue ordered by  $f_B$ , with  $node_B$  as an element  
   $reached_F \leftarrow$  a lookup table, with one key  $node_F$ .STATE and value  $node_F$   
   $reached_B \leftarrow$  a lookup table, with one key  $node_B$ .STATE and value  $node_B$   
   $solution \leftarrow failure$   
  while not TERMINATED( $solution, frontier_F, frontier_B$ ) do  
    if  $f_F$ (TOP( $frontier_F$ )) <  $f_B$ (TOP( $frontier_B$ )) then  
       $solution \leftarrow$  PROCEED( $F, problem_F, frontier_F, reached_F, reached_B, solution$ )  
    else  $solution \leftarrow$  PROCEED( $B, problem_B, frontier_B, reached_B, reached_F, solution$ )  
  return  $solution$ 
```

Bidirectional Best-First Search

keeps two frontiers and two tables of reached states

```
function PROCEED(dir, problem, frontier, reached, reached2, solution) returns a solution
    // Expand node on frontier; check against the other frontier in reached2.
    // The variable “dir” is the direction: either F for forward or B for backward.
    node ← POP(frontier)
    for each child in EXPAND(problem, node) do
        s ← child.STATE
        if s not in reached or PATH-COST(child) < PATH-COST(reached[s]) then
            reached[s] ← child
            add child to frontier
            if s is in reached2 then
                solution2 ← JOIN-NODES(dir, child, reached2[s])
                if PATH-COST(solution2) < PATH-COST(solution) then
                    solution ← solution2
    return solution
```

Evaluation of search algorithms

Criterion	Breadth- First	Uniform- Cost	Depth- First	Depth- Limited	Iterative Deepening	Bidirectional (if applicable)
Complete?	Yes ¹	Yes ^{1,2}	No	No	Yes ¹	Yes ^{1,4}
Optimal cost?	Yes ³	Yes	No	No	Yes ³	Yes ^{3,4}
Time	$O(b^d)$	$O(b^{1+\lceil C^*/\epsilon \rceil})$	$O(b^m)$	$O(b^\ell)$	$O(b^d)$	$O(b^{d/2})$
Space	$O(b^d)$	$O(b^{1+\lceil C^*/\epsilon \rceil})$	$O(bm)$	$O(b\ell)$	$O(bd)$	$O(b^{d/2})$

b is the branching factor; m is the maximum depth of the search tree;
 d is the depth of the shallowest solution, or is m when there is no solution;
 ℓ is the depth limit

Values of *hSLD*

—straight-line distances to Bucharest.


Arad	366	Mehadia	241
Bucharest	0	Neamt	234
Craiova	160	Oradea	380
Drobeta	242	Pitesti	100
Eforie	161	Rimnicu Vilcea	193
Fagaras	176	Sibiu	253
Giurgiu	77	Timisoara	329
Hirsova	151	Urziceni	80
Iasi	226	Vaslui	199
Lugoj	244	Zerind	374

A* search

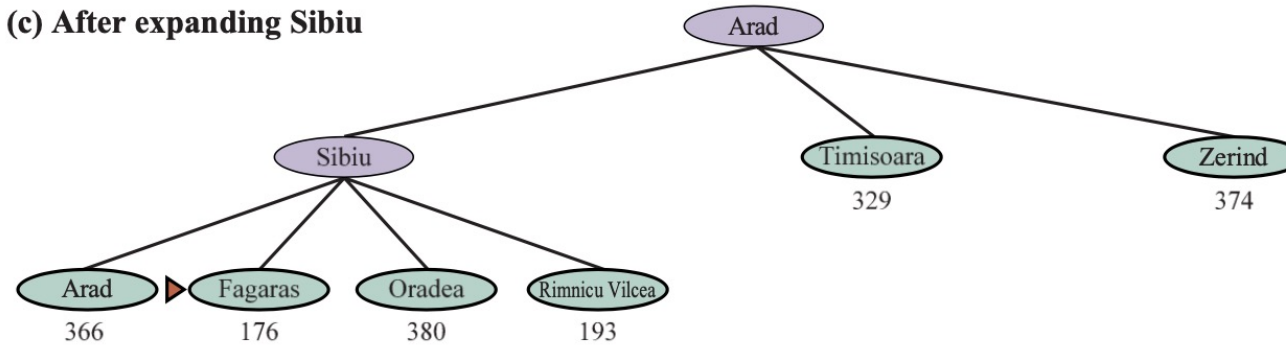
(a) The initial state



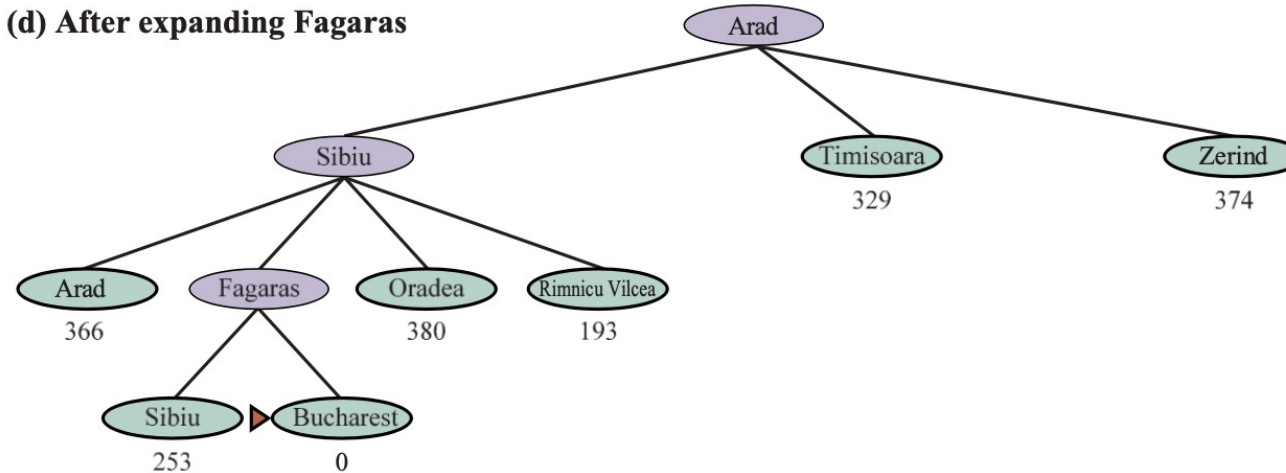
(b) After expanding Arad



(c) After expanding Sibiu



(d) After expanding Fagaras




A* search


Nodes are labeled with $f = g + h$.

The h values are the Straight-Line Distances heuristic h_{SLD}

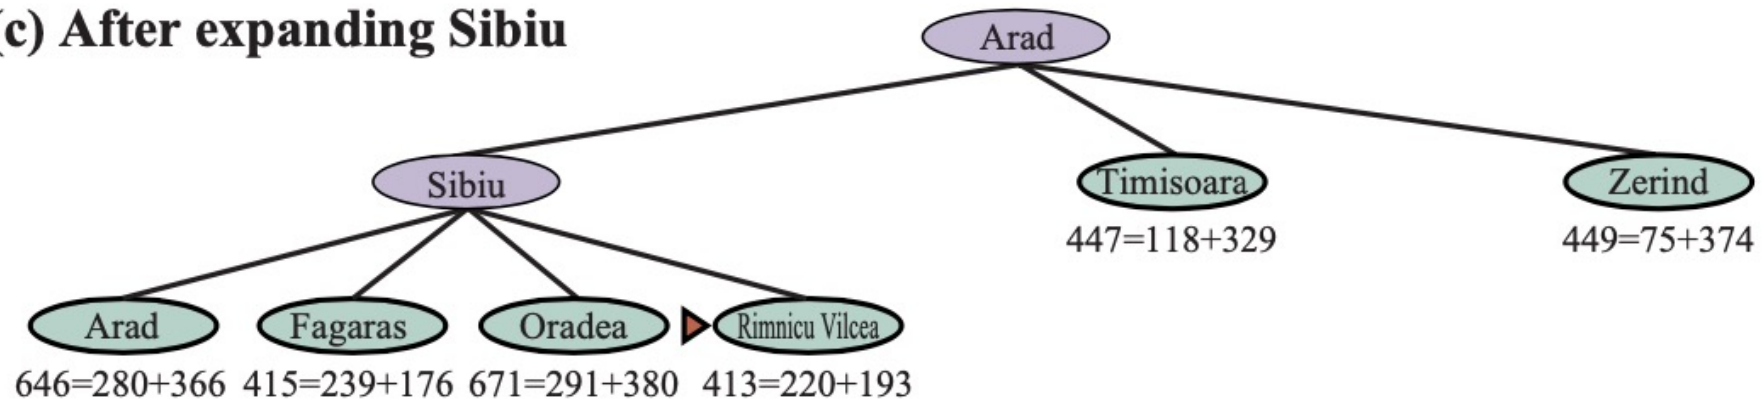
(a) The initial state



(b) After expanding Arad



(c) After expanding Sibiu

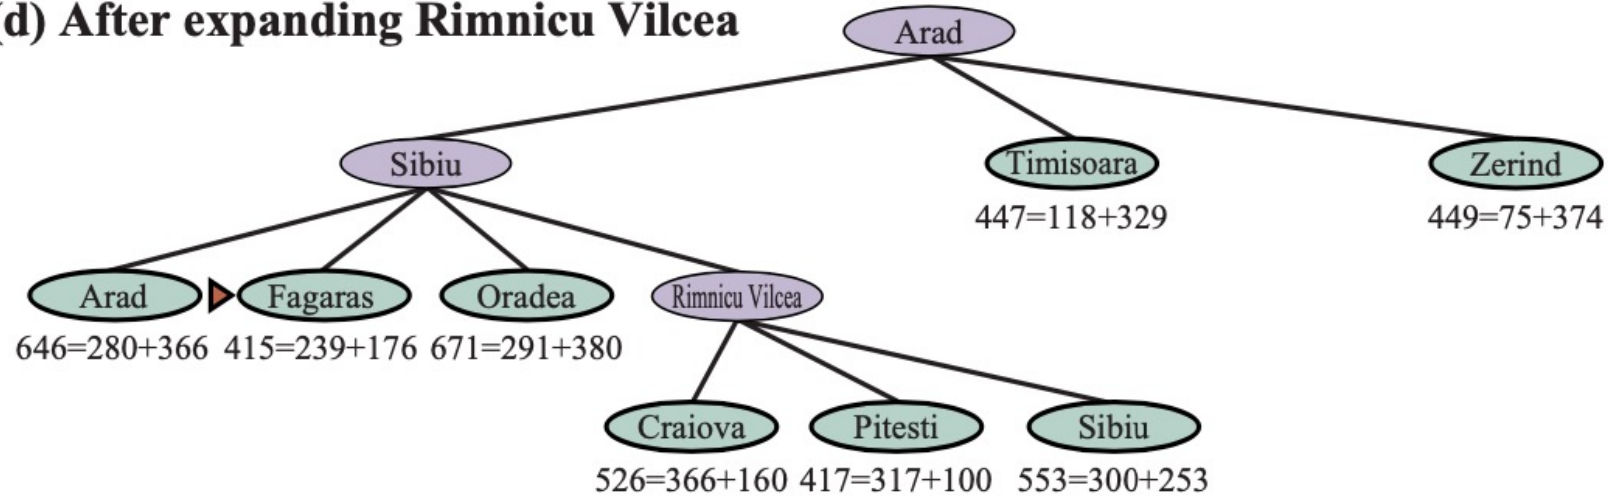


A* search

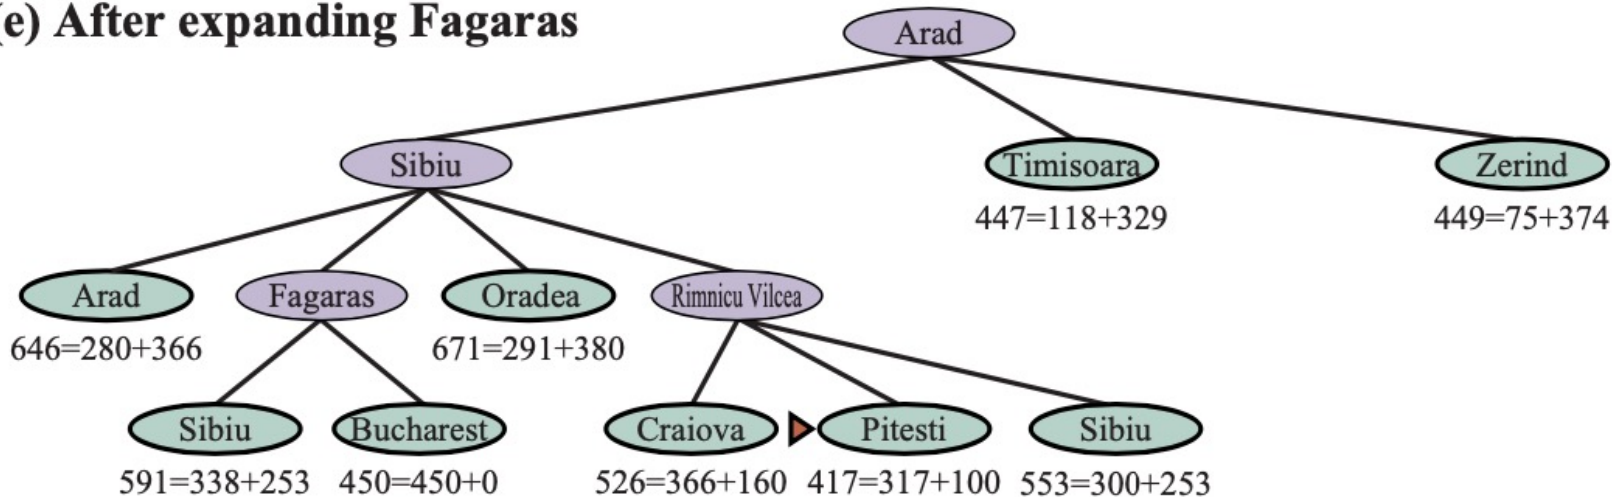
Nodes are labeled with $f = g + h$.

The h values are the Straight-Line Distances heuristic h_{SLD}

(d) After expanding Rimnicu Vilcea



(e) After expanding Fagaras

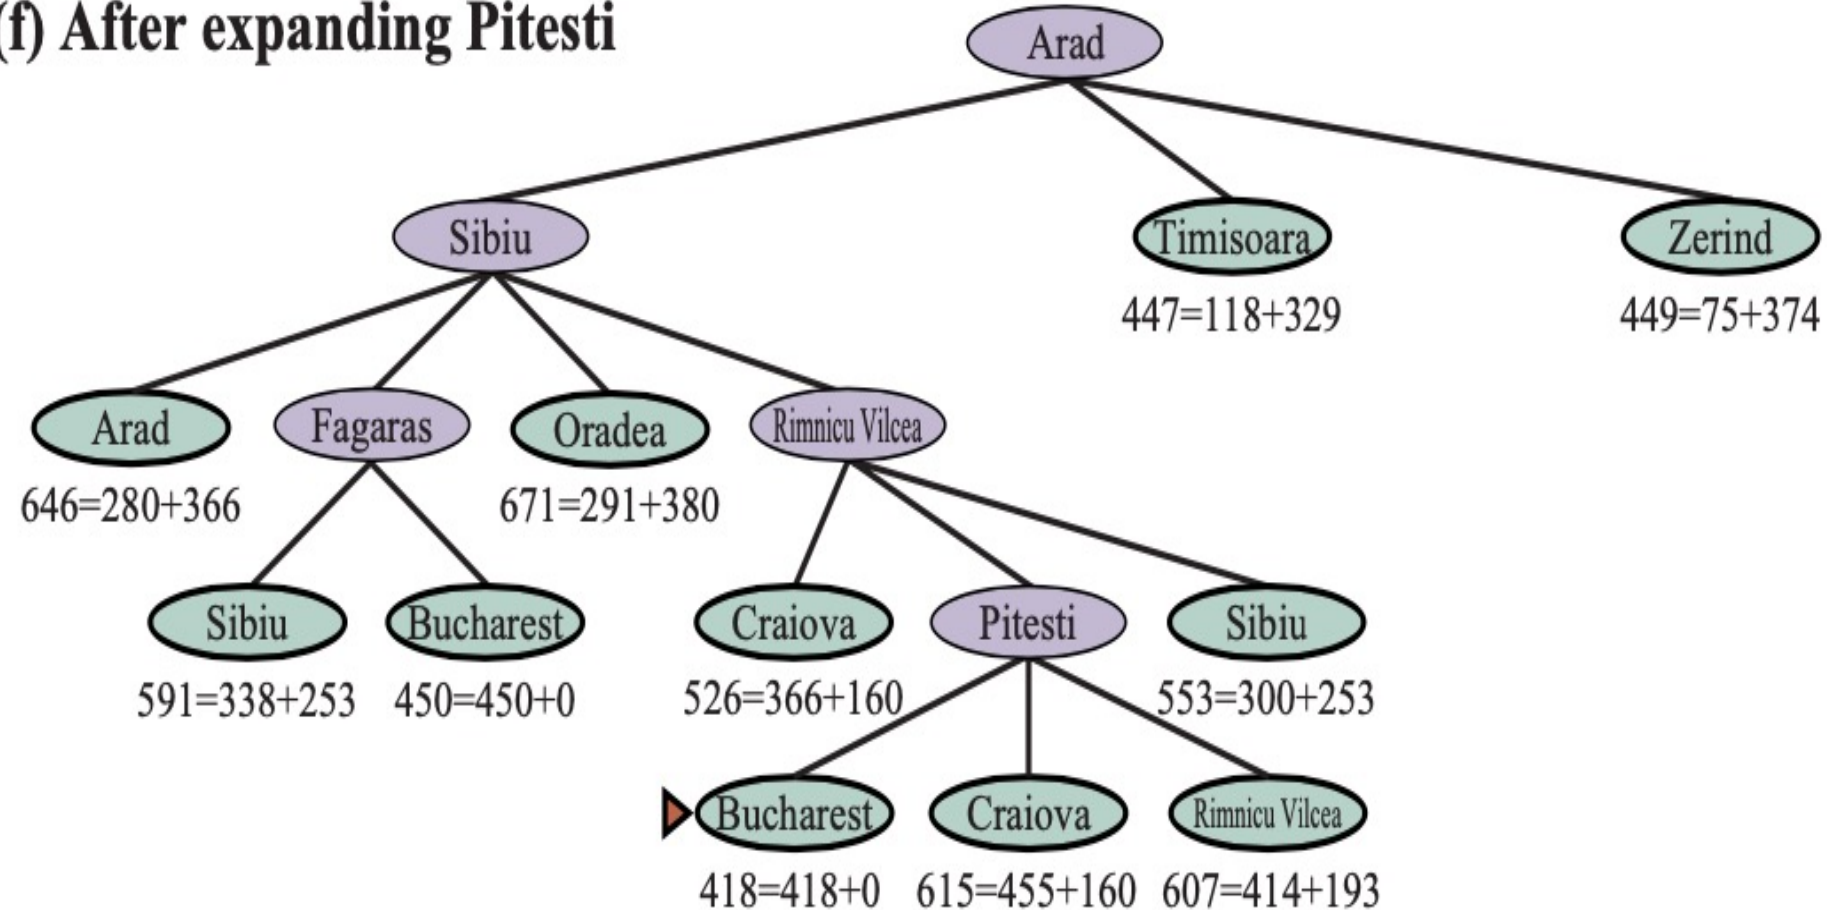


A* search


Nodes are labeled with $f = g + h$.

The h values are the Straight-Line Distances heuristic h_{SLD}

(f) After expanding Pitesti




Triangle Inequality




If the heuristic h is consistent, then the single number $h(n)$ will be less than the sum of the cost $c(n, a, a')$ of the action from n to n' plus the heuristic estimate $h(n')$.

Map of Romania showing contours at $f = 380$, $f = 400$, and $f = 420$, with Arad as the start state




(a) A* Search

(b) Weighted A* Search



(a)



(b)

The gray bars are obstacles, the purple line is the path from the green start to red goal, and the small dots are states that were reached by each search.

On this particular problem, weighted A* explores 7 times fewer states and finds a path that is 5% more costly.

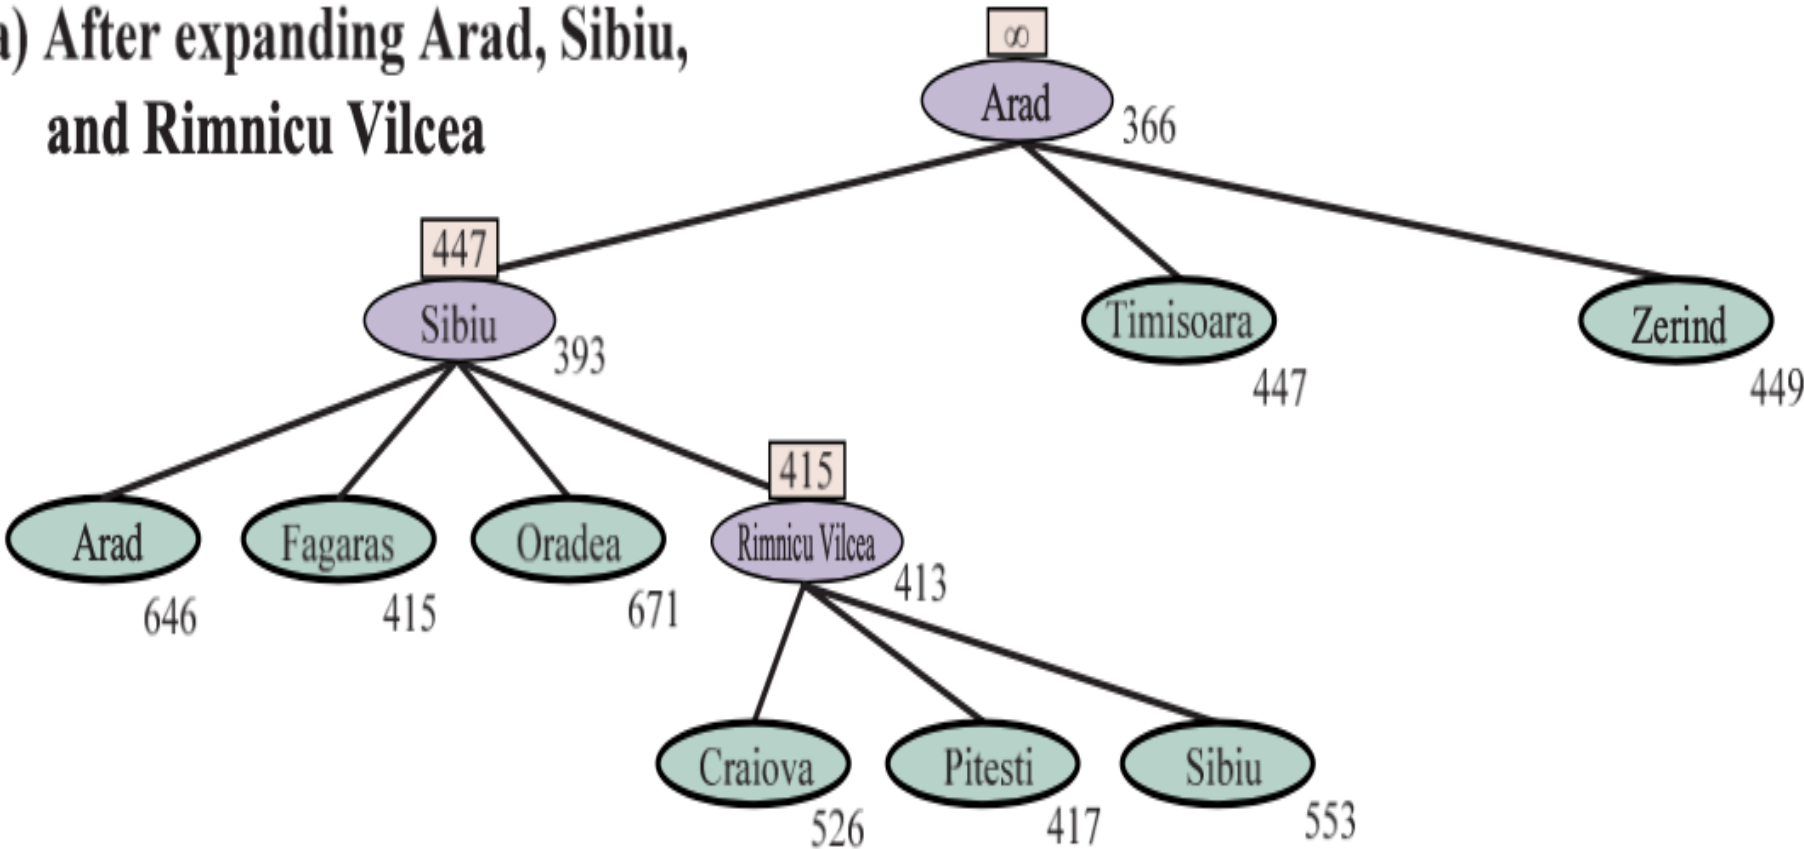
Recursive Best-First Search (RBFS) Algorithm

function RECURSIVE-BEST-FIRST-SEARCH(*problem*) **returns** a solution or *failure*
 solution, fvalue \leftarrow RBFS(*problem*, NODE(*problem*.INITIAL), ∞)
return *solution*

function RBFS(*problem, node, f_limit*) **returns** a solution or *failure*, and a new *f*-cost limit
 if *problem*.IS-GOAL(*node*.STATE) **then return** *node*
 successors \leftarrow LIST(EXPAND(*node*))
 if *successors* is empty **then return** *failure, ∞*
 for each *s* **in** *successors* **do** // update *f* with value from previous search
 s.f \leftarrow max(*s*.PATH-COST + *h*(*s*), *node.f*)
 while true do
 best \leftarrow the node in *successors* with lowest *f*-value
 if *best.f* > *f_limit* **then return** *failure, best.f*
 alternative \leftarrow the second-lowest *f*-value among *successors*
 result, best.f \leftarrow RBFS(*problem, best, min(f_limit, alternative)*)
 if *result* \neq *failure* **then return** *result, best.f*

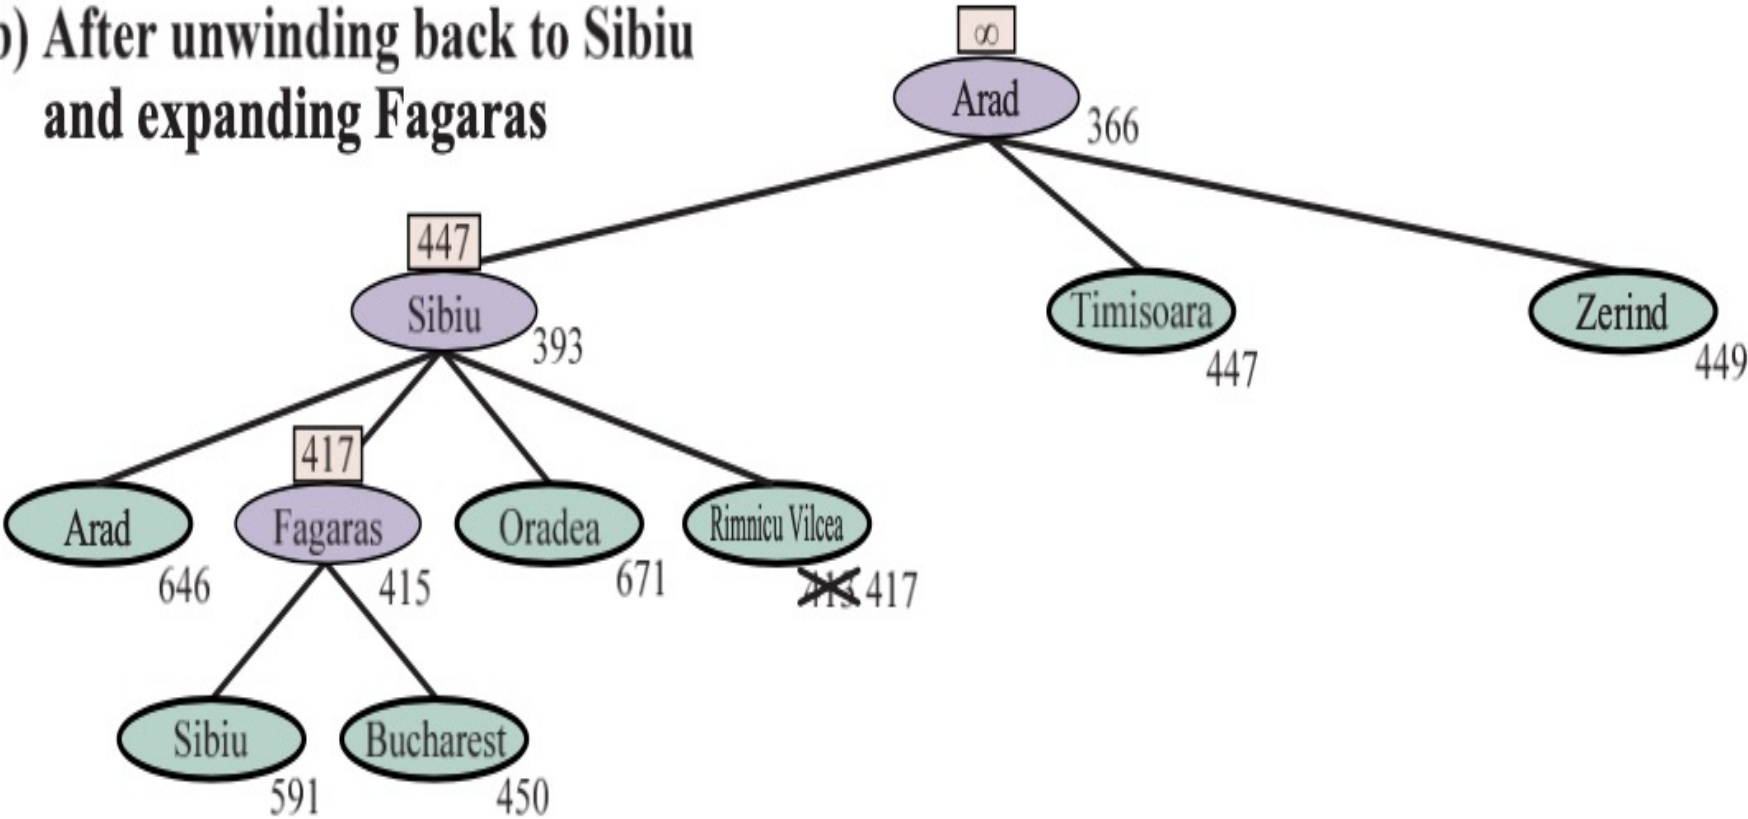
Recursive Best-First Search (RBFS)

(a) After expanding Arad, Sibiu, and Rimnicu Vilcea



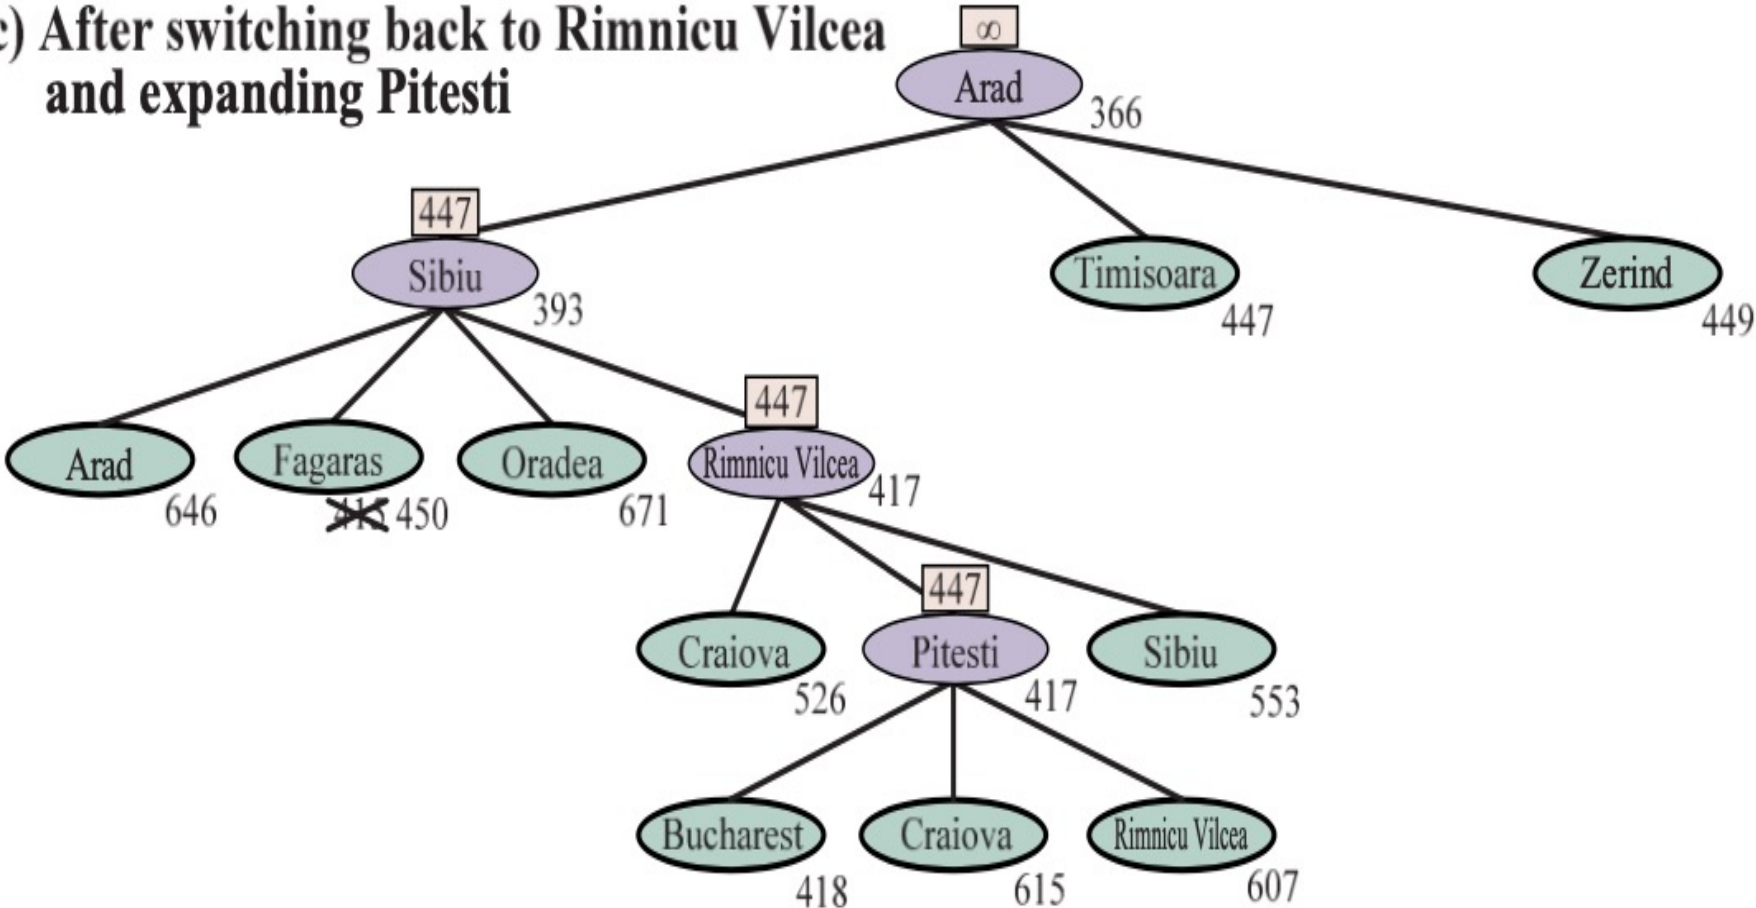
Recursive Best-First Search (RBFS)

(b) After unwinding back to Sibiu and expanding Fagaras




Recursive Best-First Search (RBFS)

(c) After switching back to Rimnicu Vilcea and expanding Pitesti



Bidirectional Search maintains two frontiers



On the left, nodes A and B are successors of Start;
on the right, node F is an inverse successor of Goal

A typical instance of the 8-puzzle

The shortest solution is 26 actions long

7	2	4
5		6
8	3	1

Start State

	1	2
3	4	5
6	7	8

Goal State

Comparison of the search costs and effective branching factors for 8-puzzle problems

d	Search Cost (nodes generated)			Effective Branching Factor		
	BFS	$A^*(h_1)$	$A^*(h_2)$	BFS	$A^*(h_1)$	$A^*(h_2)$
6	128	24	19	2.01	1.42	1.34
8	368	48	31	1.91	1.40	1.30
10	1033	116	48	1.85	1.43	1.27
12	2672	279	84	1.80	1.45	1.28
14	6783	678	174	1.77	1.47	1.31
16	17270	1683	364	1.74	1.48	1.32
18	41558	4102	751	1.72	1.49	1.34
20	91493	9905	1318	1.69	1.50	1.34
22	175921	22955	2548	1.66	1.50	1.34
24	290082	53039	5733	1.62	1.50	1.36
26	395355	110372	10080	1.58	1.50	1.35
28	463234	202565	22055	1.53	1.49	1.36

A subproblem of the 8-puzzle

*	2	4
*		*
*	3	1

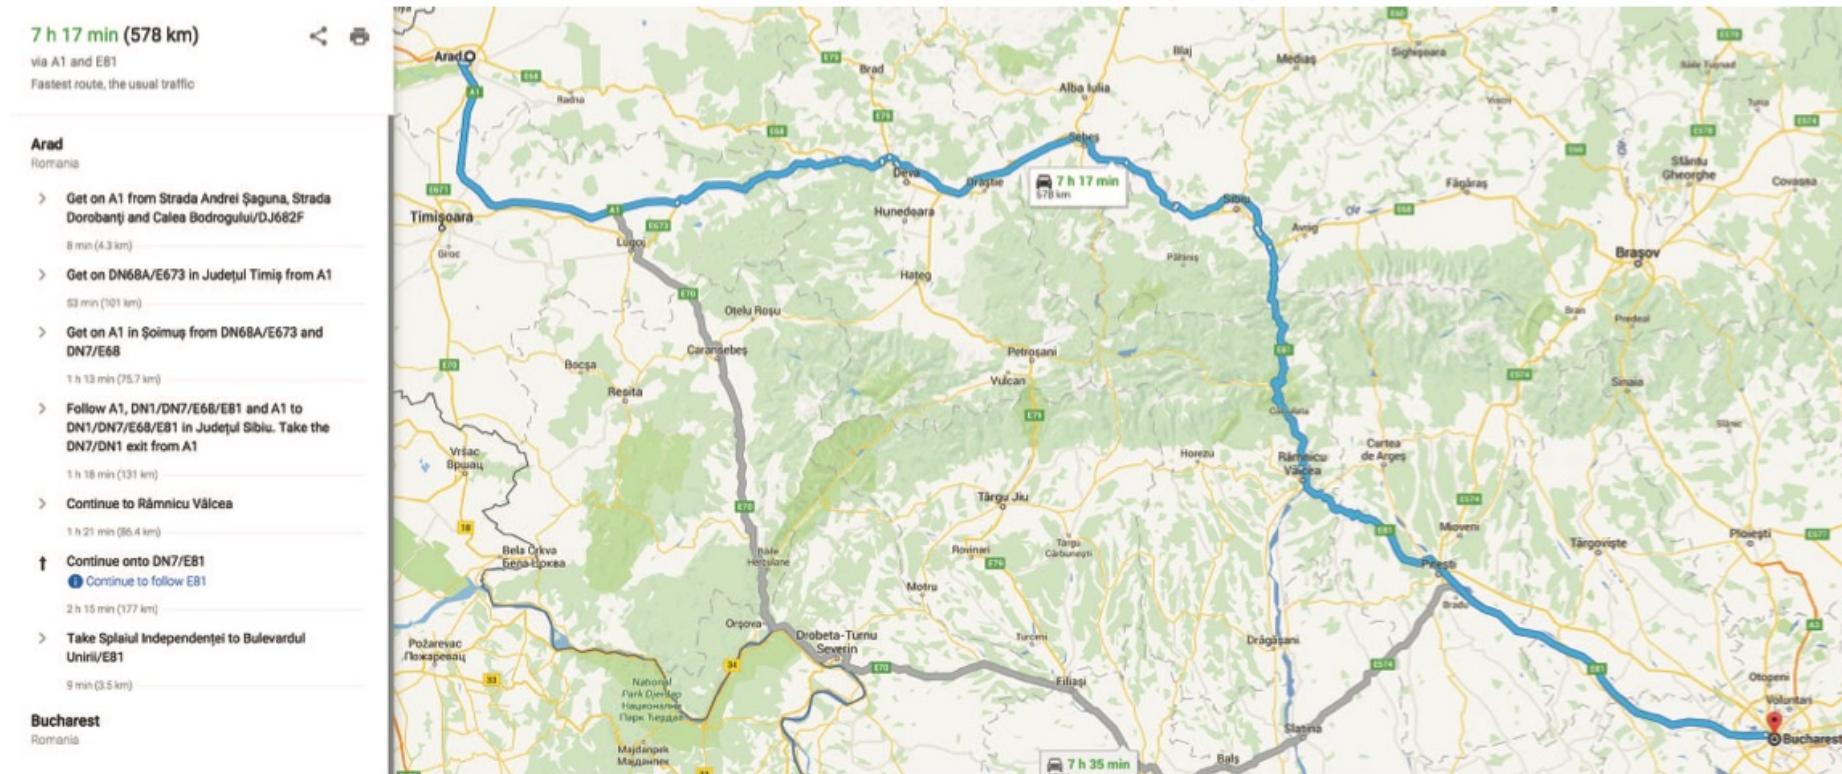
Start State

	1	2
3	4	*
*	*	*

Goal State


The task is to get tiles 1, 2, 3, 4, and the blank into their correct positions, without worrying about what happens to the other tiles

A Web service providing driving directions, computed by a search algorithm.




Search in Complex Environments

A one-dimensional state-space landscape



Adversarial Search and Games

Game Tree for the Game of Tic-tac-toe




Constraint Satisfaction Problems

The Map-Coloring Problem Represented as a Constraint Graph




(a)



(b)


A Tree Decomposition of the Constraint Graph



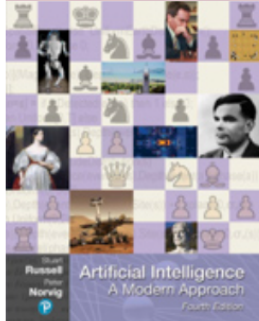
Artificial Intelligence: A Modern Approach (AIMA)

- **Artificial Intelligence: A Modern Approach (AIMA)**
 - <http://aima.cs.berkeley.edu/>
- **AIMA Python**
 - <http://aima.cs.berkeley.edu/python/readme.html>
 - <https://github.com/aimacode/aima-python>
- **Search**
 - <http://aima.cs.berkeley.edu/python/search.html>
- **Games: Adversarial Search**
 - <http://aima.cs.berkeley.edu/python/games.html>
- **CSP (Constraint Satisfaction Problems)**
 - <http://aima.cs.berkeley.edu/python/csp.html>

Artificial Intelligence: A Modern Approach (AIMA)



- △ US Edition
- △ Global Edition
- Acknowledgements
- Code
- Courses
- Editions
- Errata
- Exercises
- Figures
- Instructors Page
- Pseudocode
- Reviews



Artificial Intelligence: A Modern Approach, 4th US ed.

by Stuart Russell and Peter Norvig

The authoritative, most-used AI textbook, adopted by over 1500 schools.

Table of Contents for the US Edition (or see the Global Edition)

[Preface \(pdf\)](#); [Contents with subsections](#)

I Artificial Intelligence

- 1 Introduction ... 1
- 2 Intelligent Agents ... 36

II Problem-solving

- 3 Solving Problems by Searching ... 63
- 4 Search in Complex Environments ... 110
- 5 Adversarial Search and Games ... 146
- 6 Constraint Satisfaction Problems ... 180

III Knowledge, reasoning, and planning

- 7 Logical Agents ... 208
- 8 First-Order Logic ... 251
- 9 Inference in First-Order Logic ... 280
- 10 Knowledge Representation ... 314
- 11 Automated Planning ... 344

IV Uncertain knowledge and reasoning

- 12 Quantifying Uncertainty ... 385
- 13 Probabilistic Reasoning ... 412
- 14 Probabilistic Reasoning over Time ... 461
- 15 Probabilistic Programming ... 500
- 16 Making Simple Decisions ... 528
- 17 Making Complex Decisions ... 562
- 18 Multiagent Decision Making ... 599

V Machine Learning

- 19 Learning from Examples ... 651
- 20 Learning Probabilistic Models ... 721
- 21 Deep Learning ... 750
- 22 Reinforcement Learning ... 789

VI Communicating, perceiving, and acting

- 23 Natural Language Processing ... 823
- 24 Deep Learning for Natural Language Processing ... 856
- 25 Computer Vision ... 881
- 26 Robotics ... 925

VII Conclusions

- 27 Philosophy, Ethics, and Safety of AI ... 981
- 28 The Future of AI ... 1012
- Appendix A: Mathematical Background ... 1023
- Appendix B: Notes on Languages and Algorithms ... 1030
- Bibliography ... 1033 ([pdf](#) and [LaTeX .bib file](#) and [bib data](#))
- Index ... 1069 ([pdf](#))


[Exercises \(website\)](#)

[Figures \(pdf\)](#)

[Code \(website\)](#); [Pseudocode \(pdf\)](#)

Covers: [US](#), [Global](#)

AIMA Code



aimacode

Code for the book "Artificial Intelligence: A Modern Approach"

358 followers Berkeley, CA <http://aima.cs.berkeley.edu> peter@norvig.com

Follow

Overview Repositories 13 Projects Packages People

Popular repositories

aima-python

Public

Python implementation of algorithms from Russell And Norvig's "Artificial Intelligence - A Modern Approach"

Jupyter Notebook ☆ 6.6k 🍴 3.2k

aima-java

Public

Java implementation of algorithms from Russell And Norvig's "Artificial Intelligence - A Modern Approach"

Java ☆ 1.4k 🍴 767

aima-pseudocode

Public

Pseudocode descriptions of the algorithms from Russell And Norvig's "Artificial Intelligence - A Modern Approach"

☆ 740 🍴 386

aima-exercises

Public

Exercises for the book Artificial Intelligence: A Modern Approach

HTML ☆ 611 🍴 353

aima-javascript

Public

Javascript visualization of algorithms from Russell And Norvig's "Artificial Intelligence - A Modern Approach"

JavaScript ☆ 495 🍴 208

aima-lisp

Public

Common Lisp implementation of algorithms from Russell And Norvig's "Artificial Intelligence - A Modern Approach"

Common Lisp ☆ 342 🍴 95

You can now follow organizations

Organization activity like new discussions, sponsorships, and repositories will appear in [your dashboard feed](#).

OK, got it!

members. You must be a member to see who's a part of this organization.

Top languages

JavaScript Python Java
Common Lisp Scala

Report abuse

<https://github.com/aimacode>

AIMA Python

aimacode / aima-python Public

Watch 337

Fork 3.2k

Star 6.6k

Code Issues 120 Pull requests 79 Actions Projects Wiki Security Insights

master

1 branch 0 tags

Go to file

Add file

Code



mcventur Fixed bug in treatment of repeated nodes in frontier... 61d695b on Dec 5, 2021 1,190 commits



aima-data @ f6cbea6 updating submodule (#994) 4 years ago



gui fixed tests (#1191) 2 years ago



images add perception and tests (#1091) 3 years ago



js Added TicTacToe to notebook (#213) 7 years ago



notebooks Image Rendering problem resolved (#1178) 3 years ago



tests fixed tests (#1191) 2 years ago



.coveragerc Added coverage report generation to Travis (#1058) 3 years ago



.flake8 Fix flake8 warnings (#508) 5 years ago



.gitignore Reworked PriorityQueue and Added Tests (#1025) 4 years ago



.gitmodules Updating Submodule (#647) 5 years ago



.travis.yml fixed svm for not posdef kernel matrix, updated .travis.yml wi... 2 years ago

About

Python implementation of algorithms from Russell And Norvig's "Artificial Intelligence - A Modern Approach"

Readme

MIT license

6.6k stars

337 watching

3.2k forks

Releases

No releases published

Packages

No packages published

Papers with Code State-of-the-Art (SOTA)



Search for papers, code and tasks



[Browse State-of-the-Art](#)

[Follow](#)

[Discuss](#)

[Trends](#)

[About](#)

[Log In/Register](#)

Browse State-of-the-Art

1509 leaderboards • 1327 tasks • 1347 datasets • 17810 papers with code

Follow on [Twitter](#) for updates

Computer Vision



Semantic Segmentation

33 leaderboards
667 papers with code



Image Classification

52 leaderboards
564 papers with code



Object Detection

54 leaderboards
467 papers with code



Image Generation

51 leaderboards
231 papers with code



Pose Estimation

40 leaderboards
231 papers with code

[See all 707 tasks](#)

Natural Language Processing



Machine Translation



Language Modelling



Question Answering



Sentiment Analysis

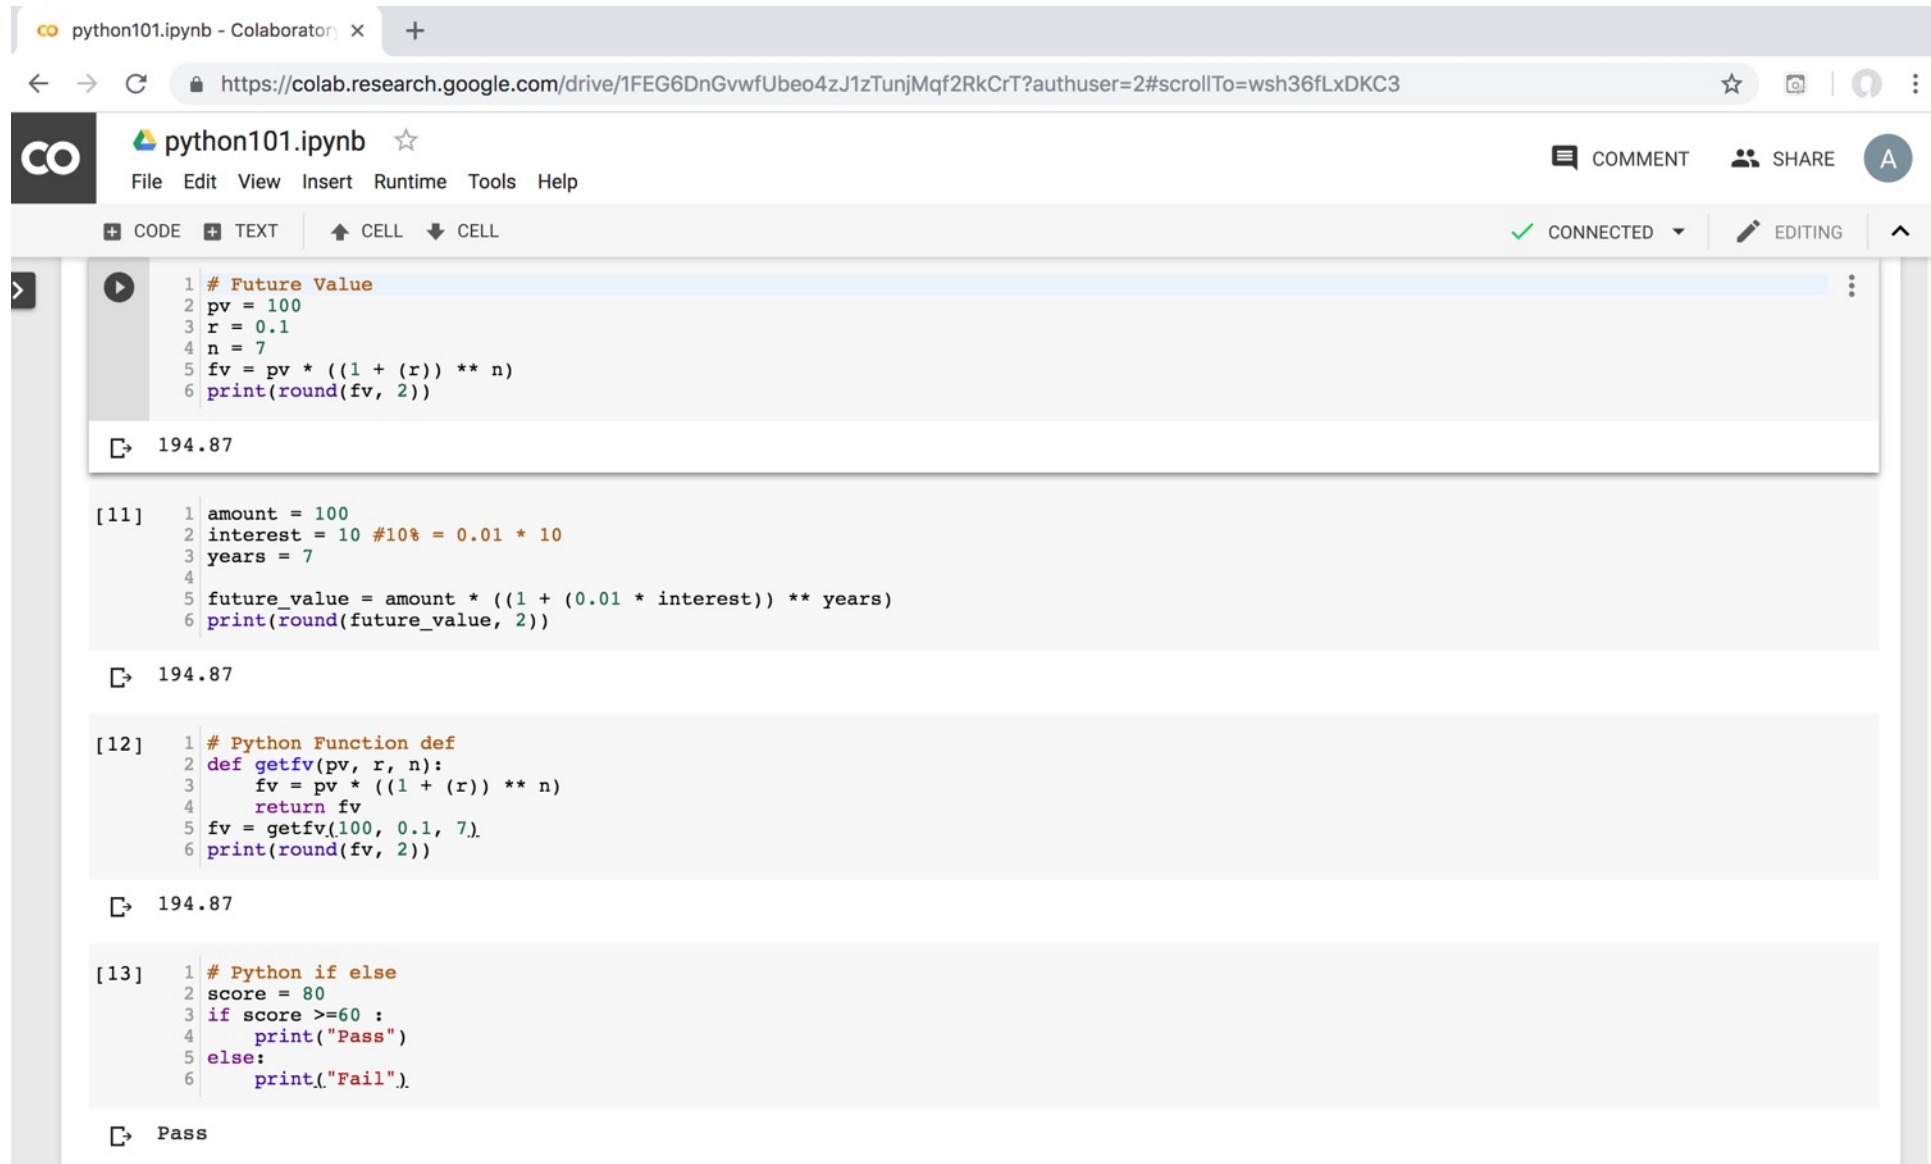


Text Generation

<https://paperswithcode.com/sota>

Python in Google Colab (Python101)

<https://colab.research.google.com/drive/1FEG6DnGvwfUbeo4zJ1zTunjMqf2RkCrT>



The screenshot shows a Google Colab notebook titled "python101.ipynb". The interface includes a browser address bar with the URL <https://colab.research.google.com/drive/1FEG6DnGvwfUbeo4zJ1zTunjMqf2RkCrT?authuser=2#scrollTo=wsh36fLxDKC3>. The notebook has a menu bar with "File", "Edit", "View", "Insert", "Runtime", "Tools", and "Help". The top right shows "COMMENT", "SHARE", and a user profile icon. The notebook is in "EDITING" mode and is "CONNECTED".

The notebook contains four code cells:

- Cell 1:** A code cell with the following Python code:

```
1 # Future Value
2 pv = 100
3 r = 0.1
4 n = 7
5 fv = pv * ((1 + (r)) ** n)
6 print(round(fv, 2))
```

The output is "194.87".
- Cell [11]:** A code cell with the following Python code:

```
1 amount = 100
2 interest = 10 #10% = 0.01 * 10
3 years = 7
4
5 future_value = amount * ((1 + (0.01 * interest)) ** years)
6 print(round(future_value, 2))
```

The output is "194.87".
- Cell [12]:** A code cell with the following Python code:

```
1 # Python Function def
2 def getfv(pv, r, n):
3     fv = pv * ((1 + (r)) ** n)
4     return fv
5 fv = getfv(100, 0.1, 7).
6 print(round(fv, 2))
```

The output is "194.87".
- Cell [13]:** A code cell with the following Python code:

```
1 # Python if else
2 score = 80
3 if score >=60 :
4     print("Pass")
5 else:
6     print("Fail").
```

The output is "Pass".

<https://tinyurl.com/aintpupython101>

Summary

- **Artificial Intelligence**
- **Intelligent Agents**
- **Problem Solving**

References

- Stuart Russell and Peter Norvig (2020), *Artificial Intelligence: A Modern Approach*, 4th Edition, Pearson.
- Numa Dhamani and Maggie Engler (2024), *Introduction to Generative AI*, Manning
- Denis Rothman (2024), *Transformers for Natural Language Processing and Computer Vision - Third Edition: Explore Generative AI and Large Language Models with Hugging Face, ChatGPT, GPT-4V, and DALL-E 3*, 3rd ed. Edition, Packt Publishing
- Ben Auffarth (2023), *Generative AI with LangChain: Build large language model (LLM) apps with Python, ChatGPT and other LLMs*, Packt Publishing.
- Aurélien Géron (2022), *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*, 3rd Edition, O'Reilly Media.
- Steven D'Ascoli (2022), *Artificial Intelligence and Deep Learning with Python: Every Line of Code Explained For Readers New to AI and New to Python*, Independently published.
- Nithin Buduma, Nikhil Buduma, Joe Papa (2022), *Fundamentals of Deep Learning: Designing Next-Generation Machine Intelligence Algorithms*, 2nd Edition, O'Reilly Media.
- Khaled Bayouhd, Raja Knani, Fayçal Hamdaoui, and Abdellatif Mtibaa (2022). "A survey on deep multimodal learning for computer vision: advances, trends, applications, and datasets." *The Visual Computer* 38, no. 8: 2939-2970.
- Abdelrahman Mohamed, Hung-yi Lee, Lasse Borgholt, Jakob D. Havtorn, Joakim Edin, Christian Igel, Katrin Kirchhoff et al. (2022). "Self-Supervised Speech Representation Learning: A Review." *arXiv preprint arXiv:2205.10643*.
- Longbing Cao (2022). "Decentralized ai: Edge intelligence and smart blockchain, metaverse, web3, and desc." *IEEE Intelligent Systems* 37, no. 3: 6-19.
- Thien Huynh-The, Quoc-Viet Pham, Xuan-Quy Pham, Thanh Thi Nguyen, Zhu Han, and Dong-Seong Kim (2022). "Artificial Intelligence for the Metaverse: A Survey." *arXiv preprint arXiv:2202.10336*.
- Thippa Reddy Gadekallu, Thien Huynh-The, Weizheng Wang, Gokul Yenduri, Pasika Ranaweera, Quoc-Viet Pham, Daniel Benevides da Costa, and Madhusanka Liyanage (2022). "Blockchain for the Metaverse: A Review." *arXiv preprint arXiv:2203.09738*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry et al. (2021) "Learning transferable visual models from natural language supervision." In *International Conference on Machine Learning*, pp. 8748-8763. PMLR.
- Aske Laat, Annie Wong, Suzan Verberne, Joost Broekens, Niki van Stein, and Thomas Back. (2024) "Reasoning with Large Language Models, a Survey." *arXiv preprint arXiv:2407.11511*.
- Madaan, Aman, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon et al. (2024) "Self-refine: Iterative refinement with self-feedback." *Advances in Neural Information Processing Systems* 36.