Gridspace

GRIDSPACE IAP 2024 LECTURE 1
CAN LLMS PLAN?

January 8, 2023

TODAY'S ROADMAP

- Course Overview
- Four Easy Tasks
- Algorithms
- Human Imagination & Planning
- DNN-Augmented Search & Games
- Recursive Prompting
- Search-Augmented Language Models

COURSE OVERVIEW

Introductions

gridspace.com

youtube.com/gridspaceinc

Where we're headed

- What are the limitations of Large Language Models (LLMs)?
- What are tasks that traditional algorithms or the human brain can tackle that a language model cannot?
- How can an LLM be augmented to recover this capability?

Objectives

- Show modern techniques for extending large language models to a wider range of speech and language tasks than next token prediction.
- Something for everyone; limited assumed knowledge
- Bi-Weekly challenge questions

Audience & Prerequisites

- Excitement about ML & speech
- Some familiarity with the mathematical language of ML and language models

Course Schedule

	MONDAY	TUESDAY	WEDNESDAY	THURSDAY	FRIDAY	SATURDAY	SUNDAY
PLANNING	Jan 8 Anthony: Can LLMs plan?	9	10	11 Wonkyum: Tools for LLM Planning	12	13	14
MEMORY	15 Martin Luther King Jr. Day	16 Nick: Can LLMs remember?	17	18 Lokman: Tools for LLM Memory	19	20	21
PERCEPTION	22 Jeremy: What do LLMs Perceive?	23	24	25 Phoebe: Tools for LLM Perception	26	27	28
LANGUAGE	29 Cole&Fulang: Philosophy of Generative Linguistics vs. LLMs	30	31	Feb 1 Cooper: Tools for LLM Conversations	2 Cole: Can LLMs do math proofs?	3	4

iap.gridspace.com

Staff & Admin

- Course lead: Fulang Chen (PhD MIT '23)
- Course support: iap@gridspace.com
- Video Releases
- Meeting Invites
- YouTube Recordings
- Remote versus in California

Structure

- Eight lectures over four units
- Bi-Weekly challenge questions
- Opportunity to present your work
- Wide span of topics. Call and response from fundamental science to practical applications.

FOUR EASY TASKS

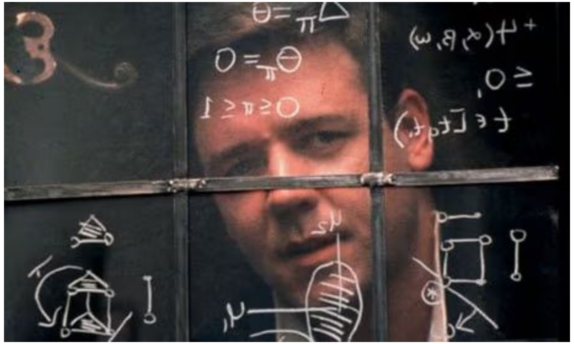

PLANNING

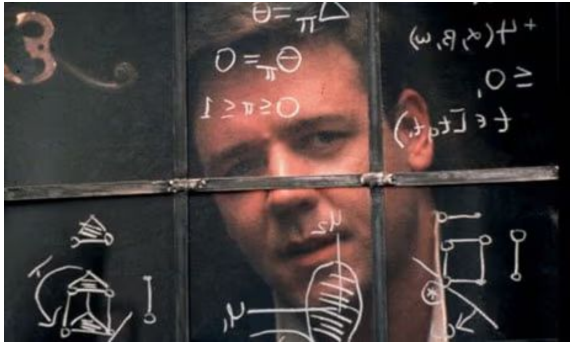

PLANNING

MEMORY

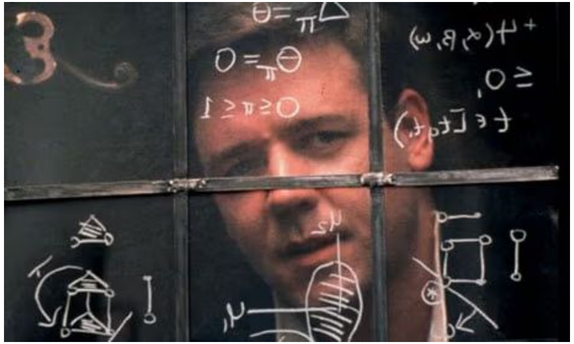

PLANNING

MEMORY

PERCEPTION

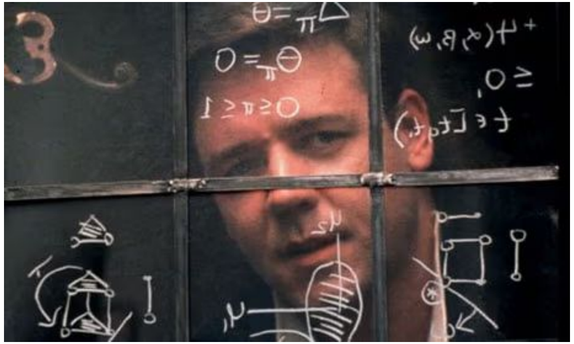

PLANNING

MEMORY

PERCEPTION

LANGUAGE &
SYMBOLIC
REASONING

ALGORITHMS


6
250
wikiHow to Do Long Division

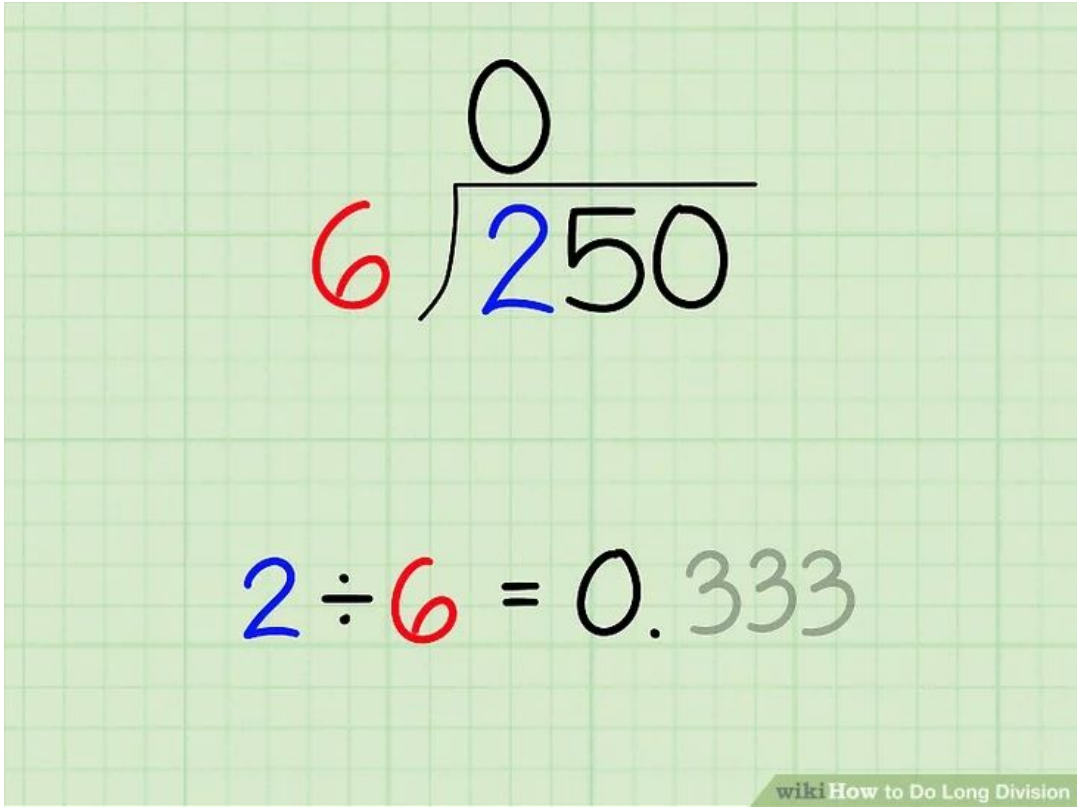
0
6) 250
2 ÷ 6 = 0.333
wikiHow to Do Long Division

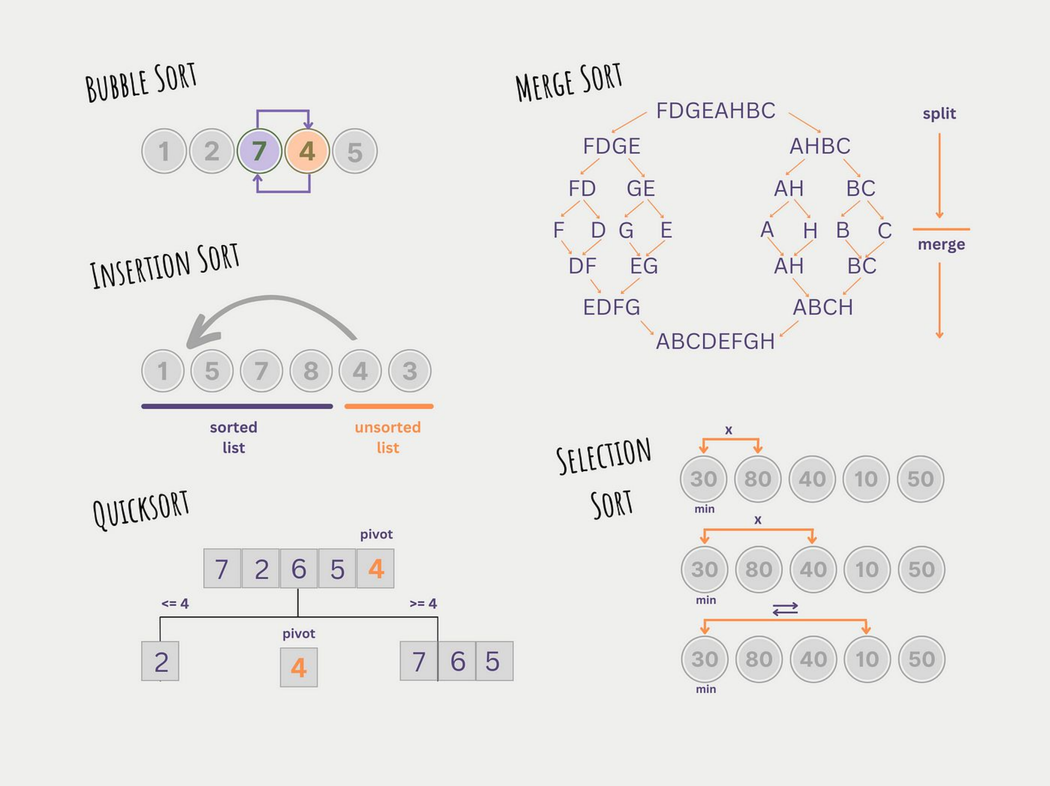
BUBBLE SORT
1 2 7 4 5
MERGE SORT
FDGEAHBC
FDGE AHBC
FD GE AH BC
F D G E A H B C
DF EG AH BC
EDFG ABCH
ABCDEFGH
split
merge
INSERTION SORT
1 5 7 8 4 3
sorted list
unsorted list
SELECTION SORT
x
30 80 40 10 50
min
QUICKSORT
pivot
7 2 6 5 4
<= 4
>= 4
2
4
7 6 5

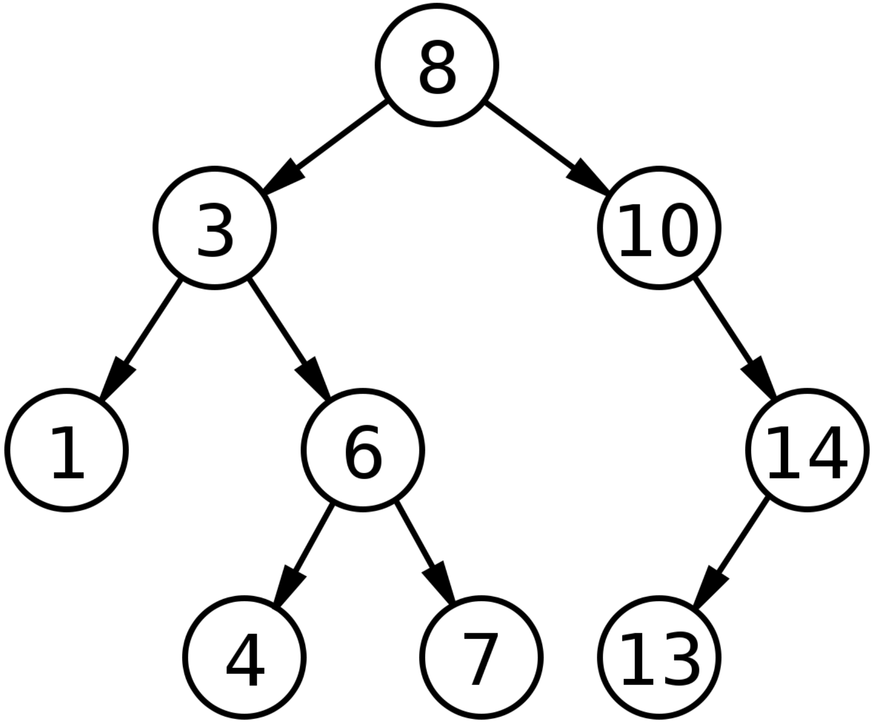
8
3
10
1
6
14
4
7
13

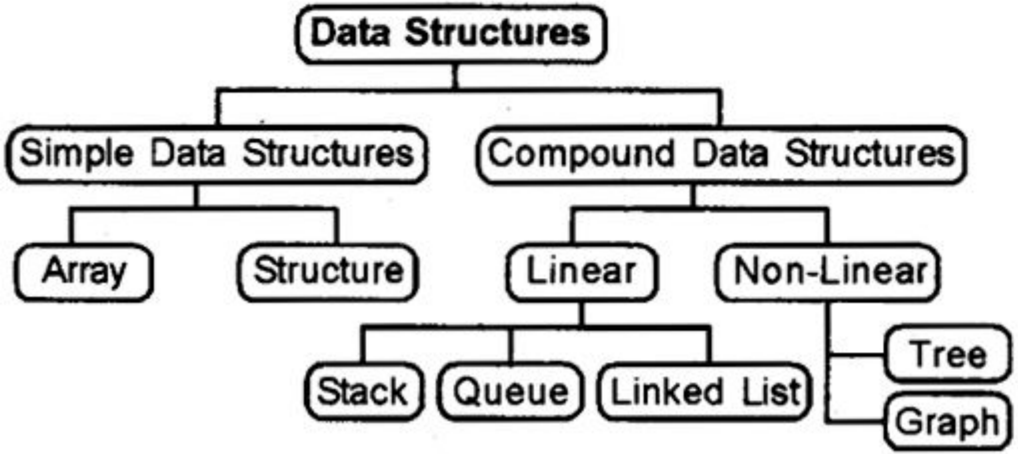
Data Structures
Simple Data Structures
Compound Data Structures
Array
Structure
Linear
Non-Linear
Stack
Queue
Linked List
Tree
Graph

Ex: $A = \begin{bmatrix} 1 & 3 & \\ 0 & 4 & \\ -1 & -2 & 2 \end{bmatrix}$

$$\left[\begin{array}{ccc|ccc} 1 & 3 & -2 & 1 & 0 & 0 \\ 0 & 4 & 1 & 0 & 1 & 0 \\ -1 & -2 & 2 & 0 & 0 & 1 \end{array}\right] \quad R_3 + R_1$$

$$\left[\begin{array}{ccc|ccc} 1 & 3 & -2 & 1 & 0 & 0 \\ 0 & 4 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \end{array}\right] \quad R_2 \longleftrightarrow R_3$$

$$\left[\begin{array}{ccc|ccc} 1 & 3 & -2 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 4 & 1 & 0 & 1 & 0 \end{array}\right] \quad R_3 - 4R_2$$

$$\left[\begin{array}{ccc|ccc} 1 & 3 & -2 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \end{array}\right] \quad R_1 + 2R_3$$

$$\underbrace{\begin{bmatrix} 1 & \\ 0 & \\ 0 & 0 \end{bmatrix}}_{= k - I}$$

$$\Rightarrow A^{-1} = \begin{bmatrix} -10 & \\ 1 & 0 \\ -4 & 1 \end{bmatrix}$$

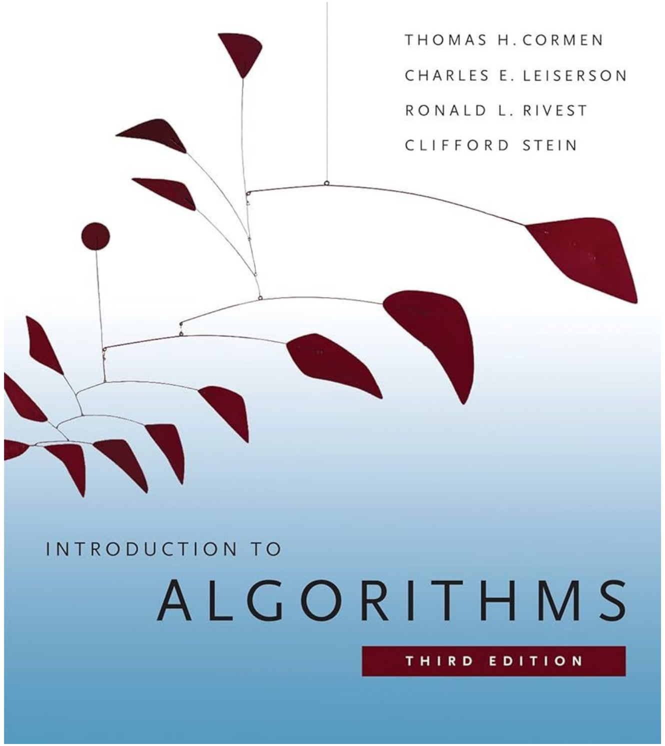
THOMAS H. CORMEN
CHARLES E. LEISERSON
RONALD L. RIVEST
CLIFFORD STEIN
INTRODUCTION TO
ALGORITHMS
THIRD EDITION

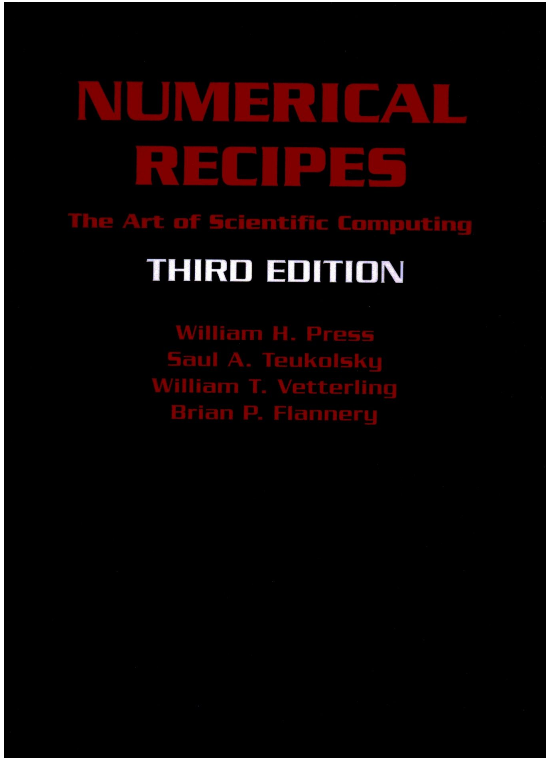
NUMERICAL RECIPES
The Art of Scientific Computing
THIRD EDITION
William H. Press
Saul A. Teukolsky
William T. Vetterling
Brian P. Flannery

HUMAN IMAGINATION & PLANNING

https://www.chess.com/daily-chess-puzzle


24 FEB 91 (PM)
U.S. ARMY

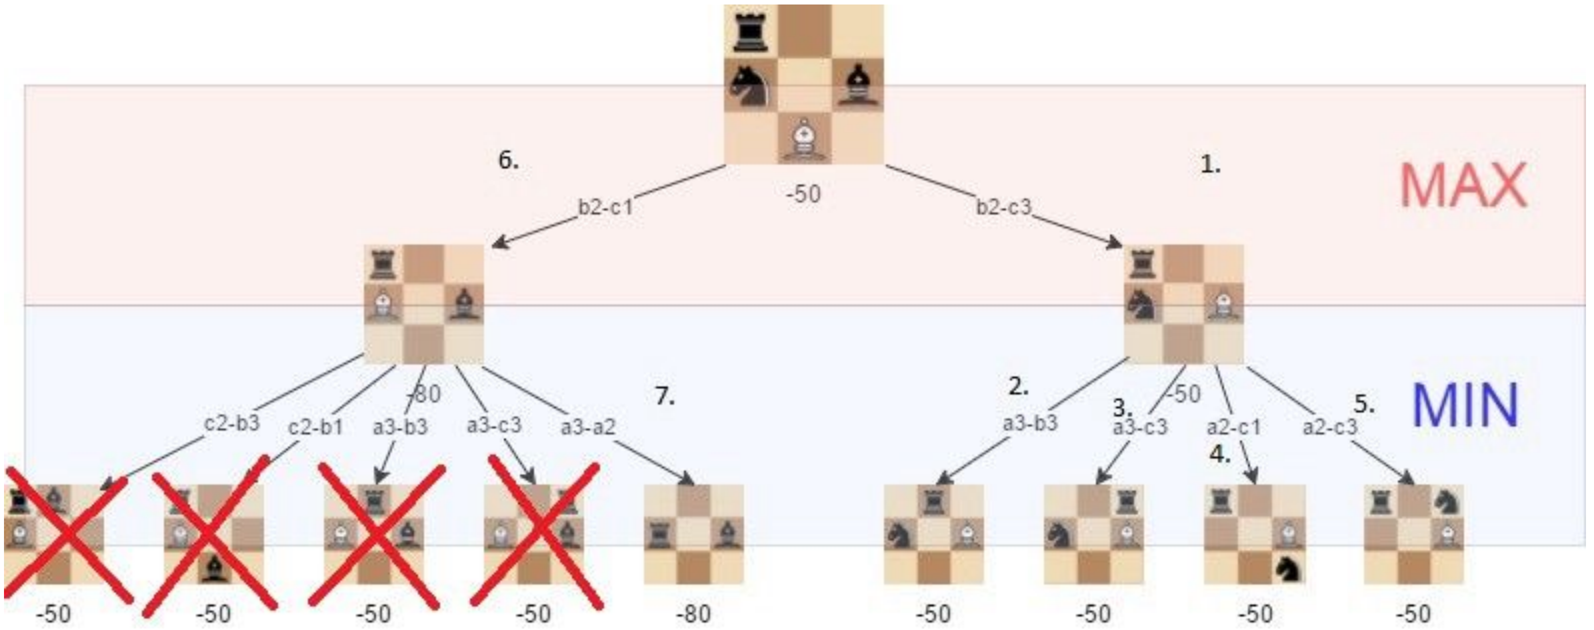
MAX
MIN
-50
6.
1.
b2-c1
b2-c3
-80
-50
7.
2.
3.
4.
5.
c2-b3
c2-b1
a3-b3
a3-c3
a3-a2
a3-b3
a3-c3
a2-c1
a2-c3
-50
-50
-50
-50
-80
-50
-50
-50
-50

DNN-AUGMENTED SEARCH & GAMES

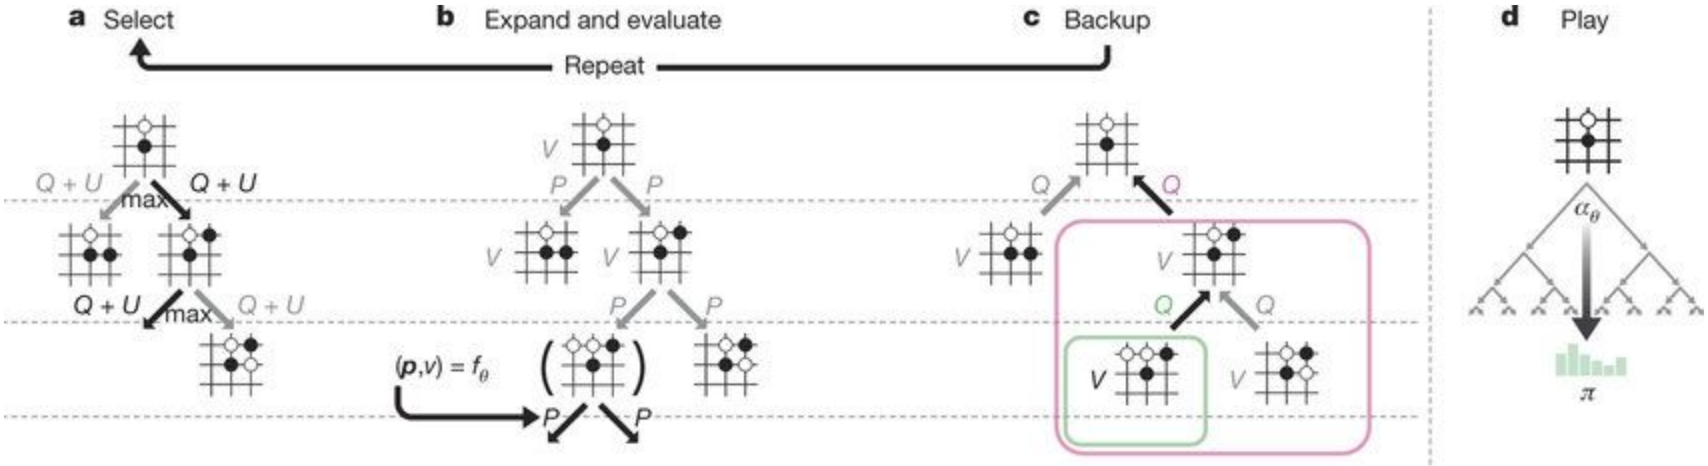

Silver, David, et al. "Mastering the game of go without human knowledge." nature 550.7676 (2017): 354-359.

Silver, David, et al. "Mastering chess and shogi by self-play with a general reinforcement learning algorithm." arXiv preprint arXiv:1712.01815 (2017).

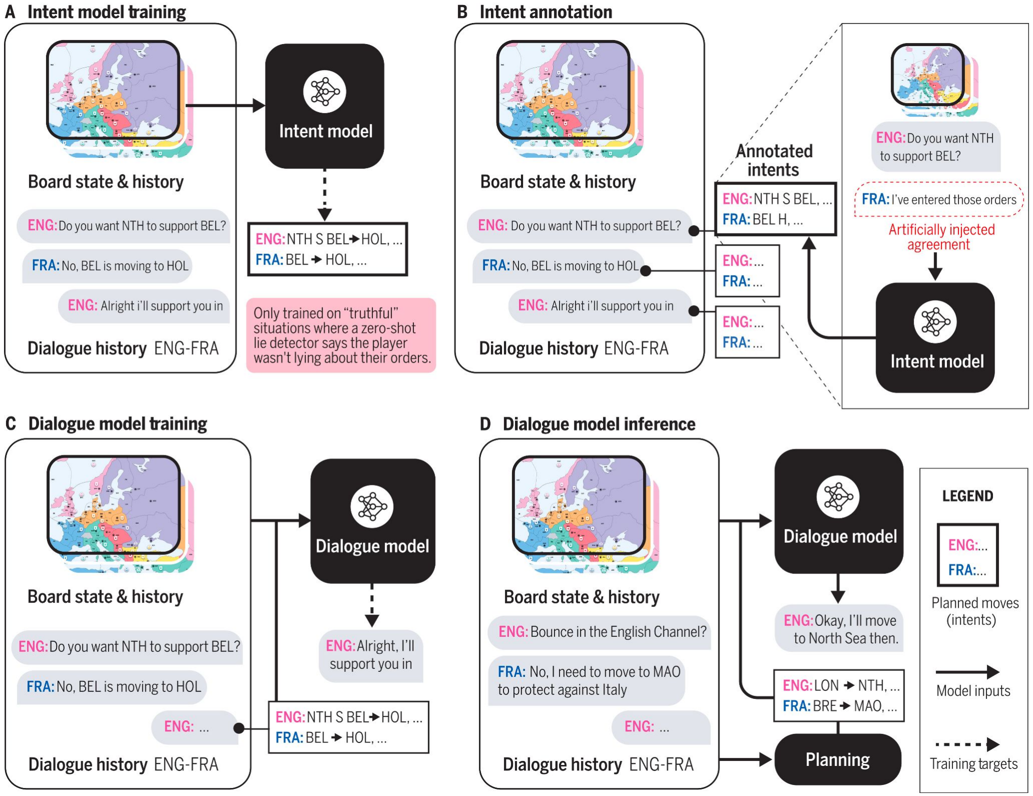

Meta Fundamental AI Research Diplomacy Team (FAIR)†, et al. "Human-level play in the game of Diplomacy by combining language models with strategic reasoning." Science 378.6624 (2022): 1067-1074.

RECURSIVE PROMPTING

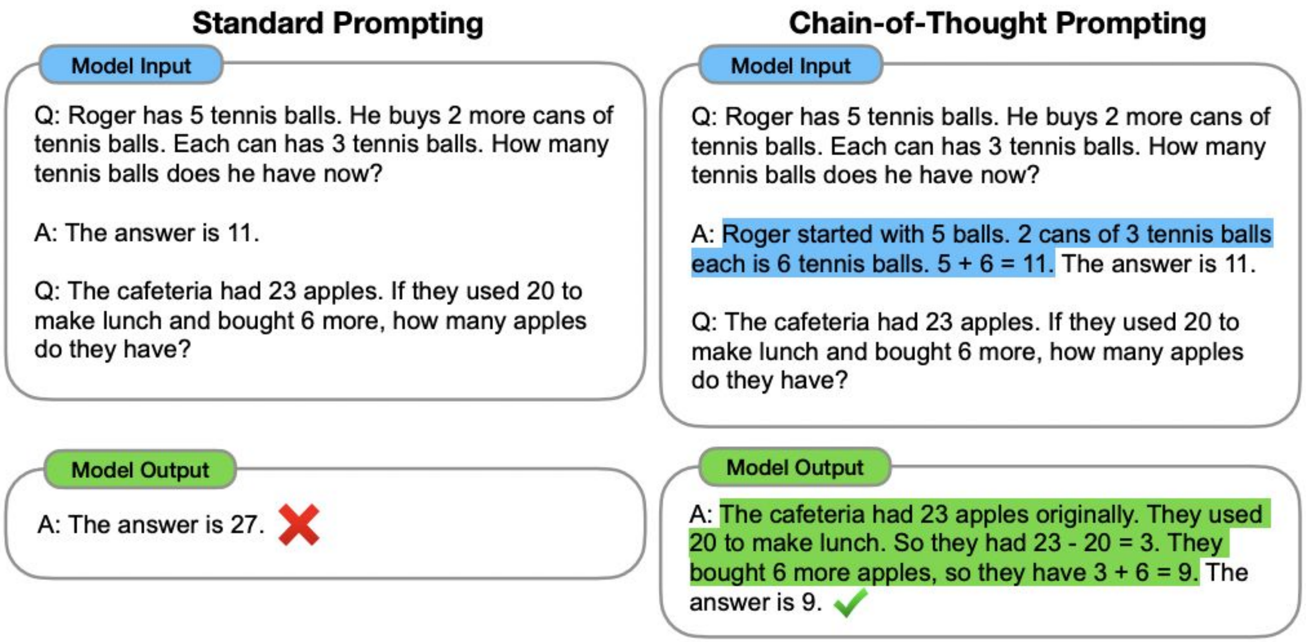

Figure 1: Chain-of-thought prompting enables large language models to tackle complex arithmetic, commonsense, and symbolic reasoning tasks. Chain-of-thought reasoning processes are highlighted.

Wei, Jason, et al. "Chain-of-thought prompting elicits reasoning in large language models." Advances in Neural Information Processing Systems 35 (2022): 24824-24837.

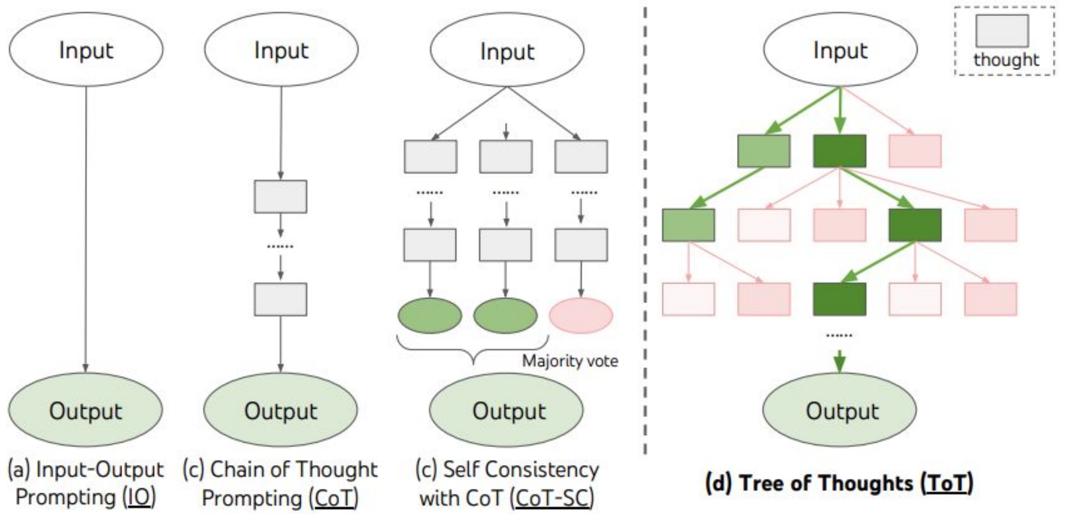

Figure 1: Schematic illustrating various approaches to problem solving with LLMs. Each rectangle box represents a *thought*, which is a coherent language sequence that serves as an intermediate step toward problem solving. See concrete examples of how thoughts are generated, evaluated, and searched in Figures 2,4,6.

Yao, Shunyu, et al. "Tree of thoughts: Deliberate problem solving with large language models." arXiv preprint arXiv:2305.10601 (2023).

Question: A needle 35 mm long rests on a water surface at $20 \circ C$. What force over and above the needle's weight is required to lift the needle from contact with the water surface? $\sigma = 0.0728m$.
`<work>`

$$\sigma = 0.0728N/m$$
$$\sigma = F/L$$
$$0.0728 = F/(2 \times 0.035)$$
$$F = 0.0728(2 \times 0.035)$$

```
calculate.py
"‘
f = 0.0728*(2*0.035)
with open("output.txt", "w") as file:
file.write(str(round(f, 5)))
"’
```

«run: `calculate.py`»

«read: `output.txt`»

0.0051

`</work>`

Answer: $F = 0.0051N$

Figure 4: Working memory example from Taylor et al. (2022). This prompt and its output are seen during LM pre-training.

Mialon, Grégoire, et al. "Augmented language models: a survey." arXiv preprint arXiv:2302.07842 (2023).

SEARCH-AUGMENTED LANGUAGE MODELS

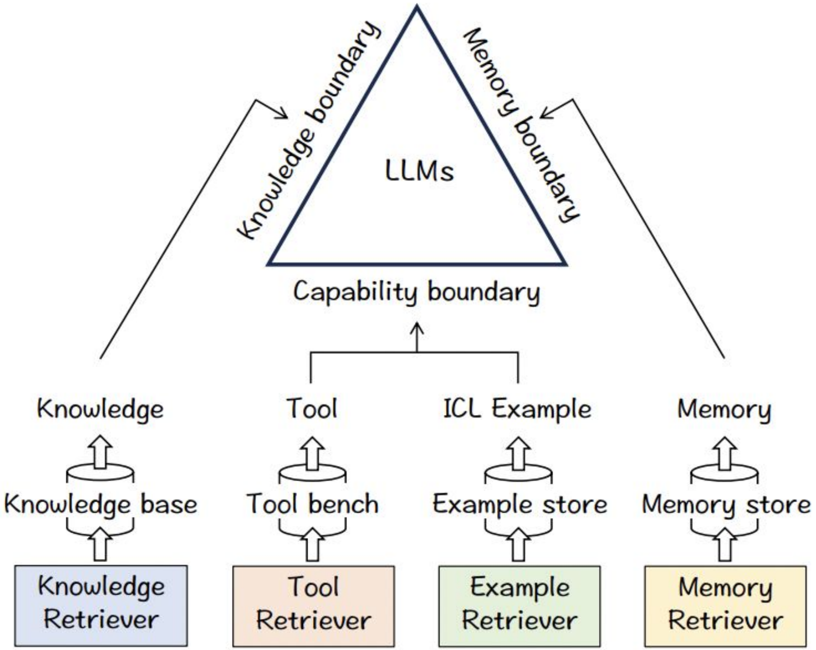

Figure 1: Confront the threefold inherent boundaries of LLMs on top of retrieval augmentation.

Zhang, Peitian, et al. "Retrieve anything to augment large language models." arXiv preprint arXiv:2310.07554 (2023).

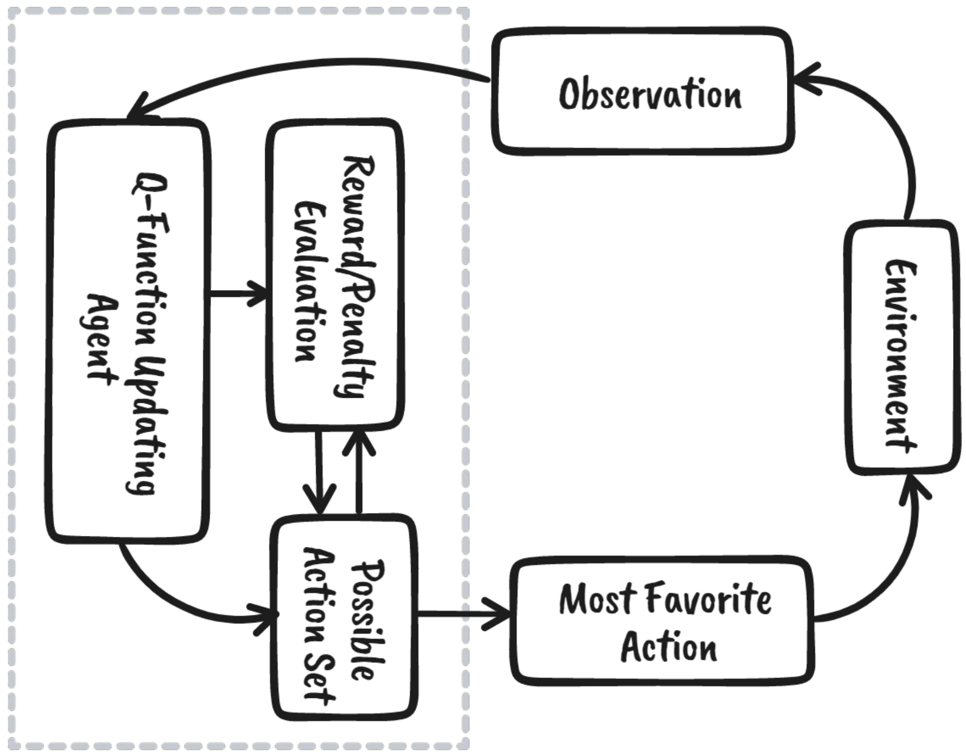
Observation
Q-Function Updating Agent
Reward/Penalty Evaluation
Possible Action Set
Most Favorite Action
Environment

Algorithm 1 Deep Q-learning with Experience Replay

Initialize replay memory $\mathcal{D}$ to capacity N
Initialize action-value function Q with random weights
for episode $= 1, M$ **do**
  Initialise sequence $s_1 = \{x_1\}$ and preprocessed sequenced $\phi_1 = \phi(s_1)$
  for $t = 1, T$ **do**
    With probability ϵ select a random action a_t
    otherwise select $a_t = \max_a Q^*(\phi(s_t), a; \theta)$
    Execute action a_t in emulator and observe reward r_t and image x_{t+1}
    Set $s_{t+1} = s_t, a_t, x_{t+1}$ and preprocess $\phi_{t+1} = \phi(s_{t+1})$
    Store transition $(\phi_t, a_t, r_t, \phi_{t+1})$ in $\mathcal{D}$
    Sample random minibatch of transitions $(\phi_j, a_j, r_j, \phi_{j+1})$ from $\mathcal{D}$
    Set $y_j = \begin{cases} r_j & \text{for terminal } \phi_{j+1} \\ r_j + \gamma \max_{a'} Q(\phi_{j+1}, a'; \theta) & \text{for non-terminal } \phi_{j+1} \end{cases}$
    Perform a gradient descent step on $(y_j - Q(\phi_j, a_j; \theta))^2$ according to equation 3
  end for
end for

Mnih, Volodymyr, et al. "Playing atari with deep reinforcement learning." arXiv preprint arXiv:1312.5602 (2013).

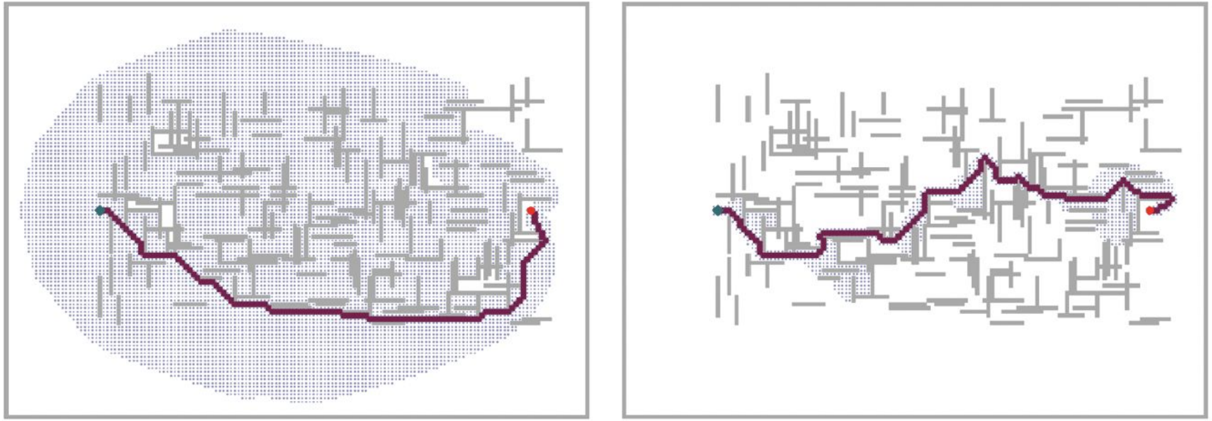

Two searches on the same grid: (a) an A* search and (b) a weighted A* search with weight $W = 2$. The gray bars are obstacles, the purple line is the path from the green start to red goal, and the small dots are states that were reached by each search. On this particular problem, weighted A* explores 7 times fewer states and finds a path that is 5% more costly.

Russell, Stuart J., and Peter Norvig. Artificial intelligence a modern approach. London, 2010.

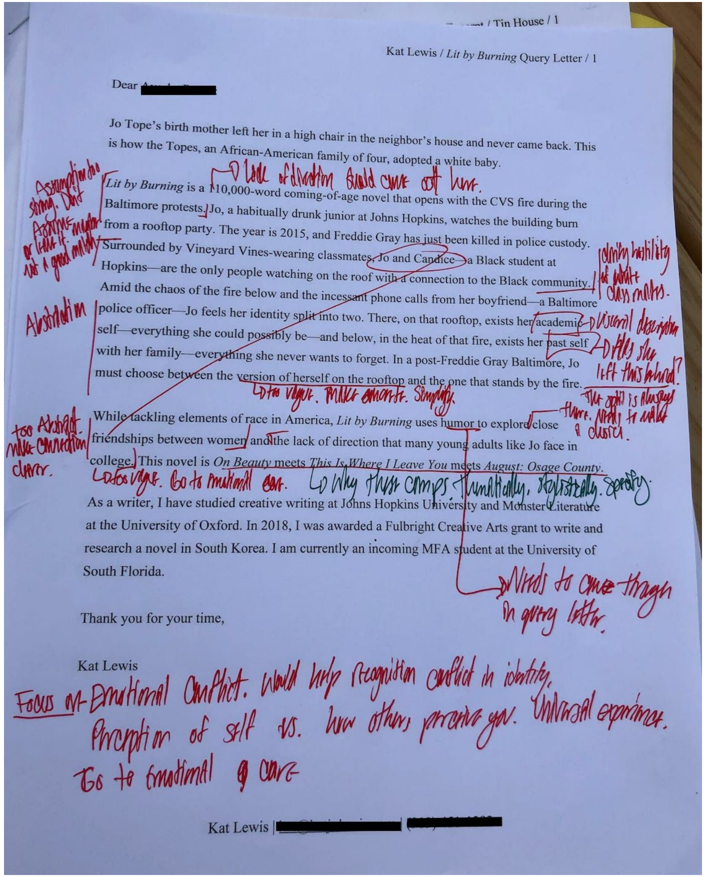

Dear [redacted]

Jo Tope's birth mother left her in a high chair in the neighbor's house and never came back. This is how the Topes, an African-American family of four, adopted a white baby.

Lack of direction should come out here.

Assumption too strong. Don't assume murder or leave it. Not a good match.

Lit by Burning is a 110,000-word coming-of-age novel that opens with the CVS fire during the Baltimore protests. Jo, a habitually drunk junior at Johns Hopkins, watches the building burn from a rooftop party. The year is 2015, and Freddie Gray has just been killed in police custody. Surrounded by Vineyard Vines-wearing classmates, Jo and Candice—a Black student at Hopkins—are the only people watching on the roof with a connection to the Black community. Amid the chaos of the fire below and the incessant phone calls from her boyfriend—a Baltimore police officer—Jo feels her identity split into two. There, on that rooftop, exists her academic self—everything she could possibly be—and below, in the heat of that fire, exists her past self with her family—everything she never wants to forget. In a post-Freddie Gray Baltimore, Jo must choose between the version of herself on the rooftop and the one that stands by the fire.

Abstraction

Clarity hostility of white classmates.

Visceral description

Did she left this behind?

Too vague. Make concrete. Simplify.

The split is always there. Needs to make a choice.

Too abstract. Make connection clearer.

While tackling elements of race in America, *Lit by Burning* uses humor to explore close friendships between women and the lack of direction that many young adults like Jo face in college. This novel is *On Beauty* meets *This Is Where I Leave You* meets *August: Osage County*.

Too vague. Go to emotional core.

Why these comps? Thematically, stylistically. Specify.

As a writer, I have studied creative writing at Johns Hopkins University and Monster Literature at the University of Oxford. In 2018, I was awarded a Fulbright Creative Arts grant to write and research a novel in South Korea. I am currently an incoming MFA student at the University of South Florida.

Needs to come through in query letter.

Thank you for your time,

Kat Lewis

Focus on Emotional Conflict. Would help recognition conflict in identity. Perception of self vs. how others perceive you. Universal experience.

Go to emotional core

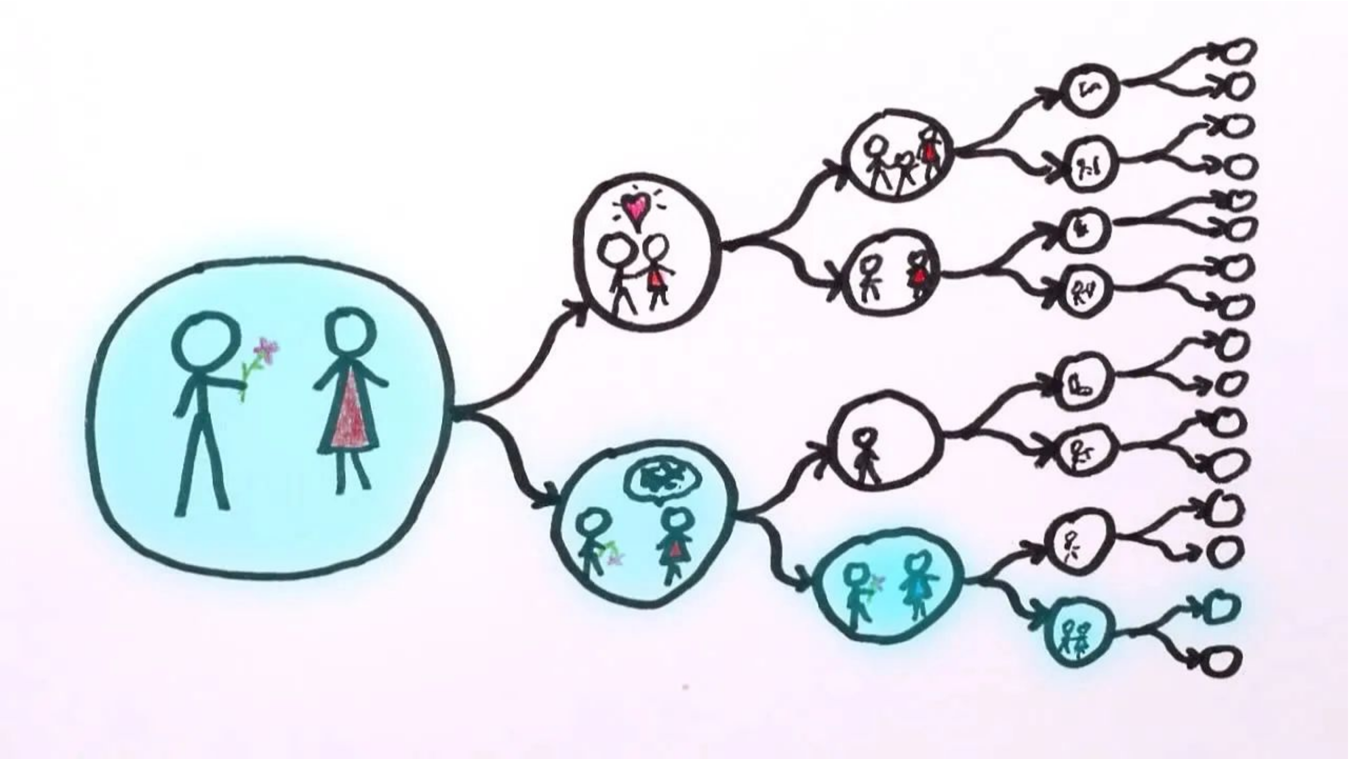

Exercises

- Hallucinations are a major problem with LLM generations. What are strategies for maximizing truthfulness of model outputs?

- How might you align an LLM to play a novel abstract strategy game using only a description of the rules?

- Stockfish is currently the most powerful chess engine by rating. Based on your own research, how much of this power can be attributed to machine learning versus non-ML algorithms.

REFERENCES

- Meta Fundamental AI Research Diplomacy Team (FAIR)†, et al. "Human-level play in the game of Diplomacy by combining language models with strategic reasoning." Science 378.6624 (2022): 1067-1074.
- Wei, Jason, et al. "Chain-of-thought prompting elicits reasoning in large language models." Advances in Neural Information Processing Systems 35 (2022): 24824-24837.
- Yao, Shunyu, et al. "Tree of thoughts: Deliberate problem solving with large language models." arXiv preprint arXiv:2305.10601 (2023).
- Mialon, Grégoire, et al. "Augmented language models: a survey." arXiv preprint arXiv:2302.07842 (2023).
- Silver, David, et al. "Mastering the game of go without human knowledge." nature 550.7676 (2017): 354-359.
- Silver, David, et al. "Mastering chess and shogi by self-play with a general reinforcement learning algorithm." arXiv preprint arXiv:1712.01815 (2017).
- Liu, Jiongnan, et al. "RETA-LLM: A Retrieval-Augmented Large Language Model Toolkit." arXiv preprint arXiv:2306.05212 (2023).
- Mnih, Volodymyr, et al. "Playing atari with deep reinforcement learning." arXiv preprint arXiv:1312.5602 (2013).
- Russell, Stuart J., and Peter Norvig. Artificial intelligence a modern approach. London, 2010.