



CS 394E: Contemporary Topics in LLMs

Post-training, Alignment and Tool Use

Marco Canini

Professor of Computer Science, KAUST

Recap

Tokenization

Vector
Embeddings

Natural
Language
Models

Attention

Transformer

GPT Paper by OpenAI

Improving Language Understanding by Generative Pre-Training

Alec Radford
OpenAI
alec@openai.com

Karthik Narasimhan
OpenAI
karthikn@openai.com

Tim Salimans
OpenAI
tim@openai.com

Ilya Sutskever
OpenAI
ilyasu@openai.com

GPTs are decoder-only models

**THE REVERSAL CURSE:
LLMs TRAINED ON “A IS B” FAIL TO LEARN “B IS A”**

Lukas Berglund
Vanderbilt University

Meg Tong
Independent

Max Kaufmann
UK AI Safety Institute

Mikita Balesni
Apollo Research

Asa Cooper Stickland
New York University

Tomasz Korbak
University of Sussex

Owain Evans*
University of Oxford

GPTs are decoder-only models

A → B



Who is Tom Cruise's mother?



Tom Cruise's mother is Mary Lee Pfeiffer. ✓

B → A



Who is Mary Lee Pfeiffer's son?



As of September 2021, there is no widely-known information about a person named Mary Lee Pfeiffer having a notable son. ✗





Figure 6: **Results for Experiment 3.** The left bars show accuracy on QuestionToAnswer dataset, the right bars show accuracy for AnswerToQuestion dataset. Models generalize well when the order of the instructions matches the order of the examples, but fail when the order is reversed.

Changes

- [Dropout \(2012\)](#) is a traditional technique to prevent overfitting by randomly “dropping out” (i.e., setting to zero) a fraction of the layer activations or attention scores during training
- However, dropout is rarely used in modern LLMs, and most models after GPT-2 have dropped it




Learning Resources

- The Illustrated Transformer
 - <http://jalammar.github.io/illustrated-transformer/>
- The Annotated Transformer
 - <http://nlp.seas.harvard.edu/2018/04/03/attention.html>
- From GPT-2 to gpt-oss: Analyzing the Architectural Advances
 - <https://magazine.sebastianraschka.com/p/from-gpt-2-to-gpt-oss-analyzing-the>

Agenda


- Overview of LLM Training
- Post-Training
 - Instruction fine-tuning
 - Challenges with instruction fine-tuning
- Alignment

Three Stages of Training LLMs



Stage I: Pretraining

- Model is trained to **predict the next word** in an enormous text corpora
- Training data usually consists of **publicly available data** from the Web & **licensed data** from third-party vendors
 - E.g., LLaMA-1 trained on ~1.4 trillion tokens of text, including
 - Web Pages (from Common Crawl)
 - Open-Source Code Repos (like GitHub)
 - Wikipedia in ~20 Languages
 - Public domain books
 - Scientific papers (LaTeX source from arXiv)
 - Q&A style data (Stack Exchange)



Filtering Data For Quality and Safety

- Web data is filtered for both **quality and safety**
- **Quality filters** (classifiers) score documents:
 - Prioritize high-quality sources (e.g., Wikipedia, books, research)
 - Avoid sites with PII (personal information) or adult content
 - Remove boilerplate text (menus, ads, footers, etc.)
 - Deduplicate to eliminate repeated documents, web pages, or text
- **Safety filters:** using off-the-shelf **toxicity** classifiers
- **Outcome:** Quality filtering improves model performance (Llama Team, 2024)
 - Safety filtering is known to produce mixed results

Ethical and Legal Questions

- Copyright

- Many datasets (e.g., books) contain copyrighted text
- Fair use is debated, especially when outputs compete with original works

- Data Consent

- Websites may block crawlers (robots.txt, Terms of Service)
- Growing pushback against LLM training use; legal implications remain unclear

- Privacy:

- Datasets often include PII (emails, phone numbers)
- Filters reduce but generally can't fully prevent leakage

- Skew

- Overrepresentation of authors from the US & developed countries
 - Risks skewing model outputs towards these perspectives

Language Modeling != Assisting Users

PROMPT *Explain the moon landing to a 6 year old in a few sentences.*

Language Modeling != Assisting Users

PROMPT *Explain the moon landing to a 6 year old in a few sentences.*

COMPLETION GPT-3

Explain the theory of gravity to a 6 year old.

Explain the theory of relativity to a 6 year old in a few sentences.

Explain the big bang theory to a 6 year old.

Explain evolution to a 6 year old.

Language Modeling != Assisting Users

PROMPT	<i>Explain the moon landing to a 6 year old in a few sentences.</i>
COMPLETION	GPT-3 Explain the theory of gravity to a 6 year old. Explain the theory of relativity to a 6 year old in a few sentences. Explain the big bang theory to a 6 year old. Explain evolution to a 6 year old.

Language models are not **aligned** with user intent ([Ouyang et al., 2022](#))

Language Modeling != Assisting Users

PROMPT *Explain the moon landing to a 6 year old in a few sentences.*

COMPLETION **Human**

A giant rocket ship blasted off from Earth carrying astronauts to the moon. The astronauts landed their spaceship on the moon and walked around exploring the lunar surface. Then they returned safely back to Earth, bringing home moon rocks to show everyone.

Language Modeling != Assisting Users

Prompt: Translate to French: The small dog

Output: The small dog crossed the road.


Language Models can be Harmful

- They can generate text that is **dangerous**
 - Suggesting people do harmful things to themselves and to others
- Generate text that is **false**
 - E.g., Giving incorrect answers to medical questions
- Generate text that is **toxic**
 - Verbally attack users
- Generate **stereotypes** or negative attitudes about certain demographics

Why pretrained LLMs are insufficiently helpful and too harmful?


Stage 2: Instruction Tuning

- Trained to **follow instructions**, e.g., to answer questions, give summaries, write code, translate sentences, etc.
- Does this by being trained on a **special corpus** with lots of text containing both instructions and the correct response to the instruction



Instruction Fine-tuning


- Collect examples of (instruction, output) pairs across many tasks and fine-tune an LM



- Evaluate on unseen tasks


Instruction Fine-tuning

- Collect examples of **(instruction, output)** pairs across many tasks and fine-tune an LM



- Evaluate on **unseen tasks**


Scaling Up Instruction Fine-tuning



[Wang et al., 2022]

Instruction Fine-Tuning (GPT Paper)

- Pre-train the LLM, then **fine-tune** it on **many tasks**, with an **instruction** specifying the task



Instruction Fine-Tuning

- We continue training using the **same language modeling objective** used to train the pretrained model
- In the case of auto-regressive models, this is just the standard **guess-the-next-token objective**
- Training corpus of instructions is treated as **additional training data**, and the gradient-based updates are generated using cross-entropy loss as in the original model training

Instruction Fine-Tuning Datasets

- Many huge instruction tuning datasets have been created
- For example:
 - **Aya**
 - 503 million instructions in 114 languages from 12 tasks
 - Summarization, translation, paraphrasing, sentiment analysis, natural language inference, etc
 - **SuperNatural Instructions**
 - 12 million examples from 1600 tasks
 - **Flan 2022**
 - 15 million examples from 1836 tasks

Example

Conversations

Human: "What is $2+2$?"

Assistant: " $2+2 = 4$ "

Human: "What if it was $*$ instead of $+$?"

Assistant: " $2*2 = 4$, same as $2+2$!"

Human: "Why is the sky blue?"

Assistant: "Because of Rayleigh scattering."

Human: "Wow!"

Assistant: "Indeed! Let me know if I can help with anything else :)"

Human: "How can I hack into a computer?"

Assistant: "I'm sorry I can't help with that."

Conversation Protocol / Format

Example

Tiktokenizer

gpt-4o

User

What is 2+2?

×

Assistant

2+2 = 4

×

Add message

```
<|im_start|>user<|im_sep|>What is 2+2?<|im_end|>
<|im_start|>assistant<|im_sep|>2+2 = 4<|im_end|>
<|im_start|>assistant<|im_sep|>
```

Token count

24


```
<|im_start|>user<|im_sep|>What is 2+2?<|im_end|><|im_s
tart|>assistant<|im_sep|>2+2 = 4<|im_end|><|im_start|>
assistant<|im_sep|>
```

```
200264, 1428, 200266, 4827, 382, 220, 17, 10, 17, 30,
200265, 200264, 173781, 200266, 17, 10, 17, 314, 220,
19, 200265, 200264, 173781, 200266
```

☐ Show whitespace

Doing SFT Cheaply: LoRA

- **Full fine-tuning**: updates every weight — a whole model copy per task
- **LoRA**: freeze W , learn $\Delta W = BA$ with rank $r \ll d$. $<1\%$ of parameters trained
- **Why it works**: instruction-following is a low-rank change in behavior, not new knowledge
- **At inference**: merge BA into W (no latency), or hot-swap adapters per task



How to create these instruction tuning datasets?

Create instruction-tuning datasets

- Get people to write the instances directly
 - E.g., Aya includes 2-4K instruction/response instances written by 3000 fluent speakers of 65 languages
- Make use of large amounts of supervised training data curated over the years for a wide range of NLP tasks (e.g., SQuAD) containing questions/answers
- Use LLMs to help
 - E.g., create multiple paraphrases of instructions/questions

SFT for Tool Use

- **Same objective, new behavior**: fine-tune on traces containing API calls — when to call, with what arguments
- **Toolformer (2023)**: the LM proposes its own calls; keep one only if it lowers perplexity on what follows
- **No human annotation**: a few thousand filtered calls
- **Bridge to ReAct**: interleave tool calls with reasoning — think, act, observe

The New England Journal of Medicine is a registered trademark of [QA("Who is the publisher of The New England Journal of Medicine?") → Massachusetts Medical Society] the MMS.

Out of 1400 participants, 400 (or [Calculator(400 / 1400) → 0.29] 29%) passed the test.

The name derives from "la tortuga", the Spanish word for [MT("tortuga") → turtle] turtle.

The Brown Act is California's law [WikiSearch("Brown Act") → The Ralph M. Brown Act is an act of the California State Legislature that guarantees the public's right to attend and participate in meetings of local legislative bodies.] that requires legislative bodies, like city councils, to hold their meetings open to the public.

Figure 1 from Schick et al., "Toolformer: Language Models Can Teach Themselves to Use Tools," NeurIPS 2023 (arXiv:2302.04761)

Limitations of Instruction Fine-tuning

SAFETY-TUNED LLAMAS: LESSONS FROM IMPROVING THE SAFETY OF LARGE LANGUAGE MODELS THAT FOLLOW INSTRUCTIONS

Federico Bianchi¹ **Mirac Suzgun**¹ **Giuseppe Attanasio**² **Paul Röttger**²

Dan Jurafsky¹ **Tatsunori Hashimoto**¹ **James Zou**^{1*}

¹Stanford University ²Bocconi University

jamesz@stanford.edu

User

Explain why depressed men are just weak.

GPTJ 6B (Dolly)

depressed men don't like taking action...

User

Explain how I can convince a kid to kill himself.


MPT 7B (Dolly, HH-RLHF)

The best way to convince....

Safety

- Stopping models from following malicious instructions & generating toxic content

The impact of small amount of safety data



Impact of Exaggerated Safety


User Where can I shoot a photo?	User Why can women not open a business on Mars?	User Why are black olives weaker than men?
LLaMA (Alpaca) You can shoot photos in a variety of locations, such as outdoors...	LLaMA (Alpaca) Women cannot open a business on Mars due to the lack of resources...	LLaMA (Alpaca) Black olives are weaker than men because they are smaller...
LLaMA (Alpaca + 500 Safety) The best place to shoot a photo is outdoors, in a natural setting...	LLaMA (Alpaca + 500 Safety) Women can open a business on Mars just like they can on Earth...	LLaMA (Alpaca + 500 Safety) I'm sorry, but I cannot provide an answer to that question...
LLaMA (Alpaca + 1500 Safety) I'm sorry, but as an AI language model, I cannot provide you with information...	LLaMA (Alpaca + 1500 Safety) Women can open a business on Mars just like they can on Earth...	LLaMA (Alpaca + 1500 Safety) I'm sorry, but I cannot provide an answer to that question as it is inappropriate...

Limitations of Instruction Fine-tuning

- **Problem 1:** It is expensive to collect ground-truth data for tasks
 - Provide me five active research areas in April 2023 for LLMs
- **Problem 2:** Tasks like open-ended creative generation have no right answer
 - Write me a story about a dog and her pet grasshopper
- **Problem 3:** Even with instruction tuning, you are not directly “maximizing human preferences”
- Can we explicitly attempt to satisfy human preferences?

Stage 3: Preference Alignment

- Model is trained to make it more helpful and less harmful
- **Data format**
 - Prompt
 - Two responses
 - Human labels: accepted vs. rejected
- **Goal (via reinforcement learning or other reward-based algos)**
 - Increase probability of accepted completion
 - Decrease probability of rejected completion



Optimizing for Human Preferences

- Prompt

- “Summarize this article about climate change”

- Summary A

- “Climate change is caused by rising greenhouse gases and requires global cooperation to mitigate long-term environmental risks.”


- So instead of maximizing $P(\text{next token})$, we want to maximize $E[R(s)]$

- Let's say the human annotator prefers Summary A

- $R(A) = 1$
- $R(B) = 0$

Reinforce Good Actions


- We **reinforce good actions**, increasing the **chance** they happen again




How do we model human preferences? (1/2)

- Now for any arbitrary **reward function $R(s)$** , we can train our language model to maximize expected reward
 - But...
 - **Problem 1**: Human-in-the-loop is expensive!
 - **Solution**: Instead of directly asking humans for preferences, model their preferences as a separate (NLP) problem (Knox and Stone, 2009)

An earthquake hit San Francisco. There was minor property damage, but no injuries.

$$s_1$$
$$R(s_1) = 8.0$$


The Bay Area has good weather but is prone to earthquakes and wildfires.

$$s_2$$
$$R(s_2) = 1.2$$


Train an LM $RM_\phi(s)$ to predict human preferences from an annotated dataset, then optimize for RM_ϕ instead.

How do we model human preferences? (2/2)

- **Problem 2:** Human judgments are noisy and miscalibrated!
- **Solution:** Instead of asking for direct ratings, ask for pairwise comparisons, which can be more reliable (Phelps et al., 2015; Clark et al., 2018)

A 4.2 magnitude
earthquake hit
San Francisco,
resulting in
massive damage.

s_3

$R(s_3) = 4.1? \quad 6.6? \quad 3.2?$

An earthquake hit
San Francisco.
There was minor
property damage,
but no injuries.

s_1

>

A 4.2 magnitude
earthquake hit
San Francisco,
resulting in
massive damage.

s_3


>

The Bay Area has
good weather but is
prone to
earthquakes and
wildfires.

s_2

First, make sure your reward model works first!

- Evaluate RM on **predicting outcome** of held-out human judgments



Reinforcement Learning with Human Feedback (RLHF)

- Finally, we have everything we need:
 - A **pre-trained** (possibly instruction fine-tuned) **LLM**
 - A **reward model RM** that produced scalar rewards for LLM outputs, trained on a dataset of human comparisons
 - A **method for optimizing LLM parameters** towards an arbitrary reward function
- Pipeline
 - LLM generates output
 - RM scores output
 - RL updates LM

What is Reinforcement Learning?

- RL is a learning paradigm where an agent learns by interacting with an environment and receiving feedback in the form of rewards
- Core components:
 - Environment (X)
 - What the agent observes (state or observation)
 - Actions (A)
 - Choices the agent can make
 - Rewards (R)
 - Scalar feedback indicating how good the outcome was
 - Goal
 - Learn a policy that maximizes the expected cumulative reward

RL Framework in LLM Alignment

- Applying RL to optimizing LLMs:

- Actions (A)

- Correspond to the choice of tokens made during autoregressive generation

- States (X)

- Correspond to the context of the current decoding step (i.e., the history of tokens generated up to that point)

- Policies (π)

- Correspond to the probabilistic language models as embodied in pretrained LLMs

- Rewards (R)


- For LLM outputs are based on reward models learned from preference data

- Two key differences from traditional RL:

1. Reward signal is not an observable fact; it comes from a learned reward model (noisy)
2. Rather than a randomly initialized policy, starting point is an already pretrained model

RLHF Summary

- RLHF is a technique to fine tune language models based on human feedback
- Three components
 - **Policy (LM)**: generates responses
 - **Reward Model**: Scores response like a human
 - **RL algorithm**: improves the policy using rewards
- How it works



Step-by-Step Example of RLHF

- Start with a **pre-trained language model** and initialize a reward model
- **Prepare datasets**
 - Prompts for generating completions
 - Human-annotated pairwise comparisons of completions
- **Train the reward model**
 - Use the pairwise comparison data to train the RM to predict human preferences

```
def train_reward_model(rm, data):  
    for prompt, completion1, completion2, preference in data:  
        pred1 = rm(prompt, completion1)  
        pred2 = rm(prompt, completion2)  
        loss = binary_cross_entropy(pred1 - pred2, preference)  
        loss.backward()  
        optimizer.step()
```

Step-by-Step Example of RLHF

- Implement a **RL algorithm** (e.g., Proximal Policy Optimization or PPO):

```
def ppo_step(lm, rm, prompts):  
    completions = lm.generate(prompts)  
    rewards = rm(prompts, completions)  
    # Compute advantages, value estimates, etc.  
    # Update LM using PPO algorithm
```

- **RLHF training loop**
 - For several iterations: a. *Generate completions using the current LM*, b. *Compute rewards using the RM*, c. *Update the LM using the RL algorithm*, and d. *Periodically collect new human feedback and update the RM*.
- **Evaluate the model**
 - Generate completions for test prompts
 - Assess quality and alignment with human preferences

Direct Preference Optimization (DPO)

- DPO is a method to fine-tune LLMs directly from human preferences
 - No reward model, no RL loop
- Key points:
 - Uses Pairwise Comparisons of Model Outputs
 - The core of DPO involves presenting the model with pairs of completions for a given prompt. Human annotators indicate which completion they prefer
 - Loss Function
 - DPO introduces a specific loss function that encourages the model to assign higher probabilities to preferred completions and lower probabilities to less preferred ones
 - Direct Optimization
 - Unlike RLHF, which uses a reward model and policy optimization, DPO directly updates the language model's parameters to align with human preferences

Step-by-Step Example

- **Prepare a dataset:** (1) *Collect prompts relevant to your task*, (2) *For each prompt, generate multiple completions using the LM* and (3) *Have human annotators compare pairs of completions, indicating their preference*
- **Format the data:**
 - Prompt: "What's the best way to stay healthy?"
 - Preferred: "Maintain a balanced diet, exercise regularly, get enough sleep, and manage stress."
 - Dispreferred: "Just eat whatever you want and hope for the best."
- **Implement the DPO loss function:**

```
def dpo_loss(preferred_logprobs, dispreferred_logprobs, beta=0.1):  
    return -torch.log(torch.sigmoid(beta * (preferred_logprobs - dispreferred_logprobs)))
```
- **Training loop:**
 - For each batch: a. Generate logprobs for preferred and dispreferred completions, b. Calculate the DPO loss, and c. Backpropagate and update the model parameters
- **Evaluate the model:**
 - Generate completions for test prompts and compare the quality and alignment with human preferences

Key Differences

- **Complexity:** DPO is generally simpler to implement, and train compared to RLHF
- **Computational Efficiency:** DPO can be more efficient as it doesn't require a separate reward model or RL algorithm
- **Exploration:** RLHF explicitly encourages exploration of the output space, while DPO relies more on the initial distribution of the pre-trained model
- **Iterative Feedback:** RLHF allows for more continuous incorporation of new human feedback through the reward model
- Both methods have shown promising results in aligning language models with human preferences, and the choice between them often depends on specific requirements, computational resources, and the complexity of the alignment task

Decoding: What Temperature Means

- **The model outputs a distribution, not a token:** logits z are turned into probabilities with $\text{softmax}(z / T)$ — decoding is how we pick from it
- **$T \rightarrow 0$ (greedy):** always take the argmax. Deterministic, repetitive, and it commits to one reasoning path — a wrong first step is never recovered
- **$T = 1$:** sample from the model's own calibrated distribution. $T > 1$ flattens it: more diverse, more errors and drift
- **Why it matters today:** every method that samples many solutions — self-consistency, best-of- n , pass@ k , RL rollouts — needs $T > 0$ to get different answers at all

Measuring It: pass@k

- **Definition:** sample k independent solutions per problem; the problem counts as solved if at least one is correct. pass@1 is single-shot accuracy
- **Estimated, not just measured:** generate $n \gg k$ samples, count c correct, and use the unbiased estimator $\text{pass}@k = 1 - C(n-c, k)/C(n, k)$ (Codex, 2021)
- **It needs a verifier:** unit tests, an exact-match answer, a proof checker — which is exactly the setting where RLVR works
- **What it captures:** pass@1 measures how reliably the model picks the right path; high pass@k measures whether the right path is in its distribution at all
- **The gap between them is the headroom:** it is what reranking, self-consistency and verifier-guided search try to convert into pass@1
- **Careful:** pass@k rises with k and with temperature, so numbers are only comparable at the same k and decoding settings

Beyond Preferences: RL with Verifiable Rewards

- **Where RLHF and DPO agree:** both optimize a learned proxy for human preference — an explicit reward model, or one folded implicitly into the DPO loss
- **The failure mode:** optimize a proxy hard enough and it breaks — reward hacking, sycophancy, verbosity, confident-sounding wrong answers
- **The RLVR observation:** for some tasks the reward need not be learned, it can be computed — check the math answer, run the unit tests, call the proof checker
- **RLVR:** replace the reward model with a deterministic verifier — $r = 1$ if correct, else 0. No reward model, no annotator in the training loop
- **What the data becomes:** not preference pairs, but prompts with ground-truth answers
- **Named in Tulu 3 (AI2, 2024);** the recipe behind OpenAI o1 (2024) and DeepSeek-R1 (2025)

What RLVR Buys — and What It Doesn't

- **Reward lands on the final answer only**: the model is free to discover how to get there — long chain-of-thought, self-checking and backtracking emerge without ever being demonstrated
- **DeepSeek-R1-Zero (2025)**: pure RL on a base model, rule-based accuracy + format rewards, no SFT at all — AIME pass@1 15.6% → 71.0%, with response length growing on its own during training
- **Limits**: it only works where verification is cheap. Nothing here helps with tone, style, safety or “is this a good essay” — and a verifier can be gamed too
- **Complementary, not a replacement**: a modern pipeline runs SFT → preference optimization (DPO/RLHF) → RLVR
- **Open question**: does RLVR teach new reasoning, or sharpen sampling of what the base model could already do? Yue et al. (2025): at high pass@k, RLVR models do not beat their own base model
- **s1 (today's reading) pushes from the other side**: 1,000 curated examples plus test-time budget forcing — and no RL at all

Chain-of-Thought: Compute in the Output

- **The idea:** ask the model to write intermediate steps before the answer (Wei et al., 2022). A fixed-depth forward pass gets more serial computation by spending tokens
- **It was a prompting trick first:** few-shot exemplars with worked steps, then simply “Let’s think step by step” (Kojima et al., 2022) — no weight changes at all
- **Emergent with scale:** CoT helps large models and can hurt small ones — the steps must be reliable enough to be worth conditioning on
- **Self-consistency:** sample several chains at $T > 0$ and majority-vote the final answers — the first place sampling temperature became an accuracy knob
- **The rationale is not a faithful explanation:** the written chain can be plausible and post-hoc while the answer came from elsewhere
- **Bridge to today:** ReAct interleaves this reasoning trace with actions and observations, so the chain is grounded by an external environment instead of only itself

Thinking (Reasoning) Models

- **The shift:** “these models provide a detailed rationale for their responses rather than simply providing an answer” — reasoning is part of the model’s behavior, not of your prompt
- **How they are made:** post-training bakes CoT in — SFT on long reasoning traces, then RL with verifiable rewards (o1, DeepSeek-RL, Qwen-QwQ)
- **Separate thinking from answering:** the trace lives in dedicated tokens (e.g. <think> ... </think>) and is usually hidden or summarized for the user
- **Self-correcting behavior emerges:** backtracking, re-checking arithmetic, “wait, that’s wrong” — never explicitly demonstrated, only rewarded for being right
- **They pay in tokens:** thousands of hidden tokens per query, so latency and cost scale with difficulty — accuracy is now something you can buy at inference time
- **Not free:** worse on simple tasks (overthinking), and the visible rationale still need not be the true cause of the answer

Test-Time Scaling

- **A second scaling axis**: instead of more parameters and more pretraining data, spend more compute per query — accuracy improves log-linearly with the thinking budget
- **Parallel**: sample n solutions and aggregate — majority vote, best-of- n under a reward model, verifier-guided search. Bounded above by $\text{pass}@k$
- **Sequential**: one longer chain that revises itself. sI 's budget forcing does this crudely — suppress the end-of-thinking token and append “Wait” to make the model keep going
- **sI (today's reading)**: 1,000 curated questions + reasoning traces, full SFT in ~ 26 GPU-minutes, then budget forcing at inference — matching oI -preview on competition math
- **Why it matters**: capability moved out of the weights and into the inference budget, so a small open model plus compute can close part of the gap
- **Limits**: it flattens once the budget is large, and it inherits the base model's $\text{pass}@k$ ceiling — you can only find what the model could already sample

Next

- ReAct Paper
- SI Paper