



# The Art of Artificial Reasoning for (Small) Language Models

Yejin Choi  
Stanford & NVIDIA ❤️

# Mission Impossible: Democratizing Generative AI by Transcending Scaling Laws

because (current) scaling laws demand  
extreme-scale compute that only a few can afford



Isn't that just impossible?



# A case in point...

Circa 2023

How can Indian startups create foundation models for India?

Rajan Anandan



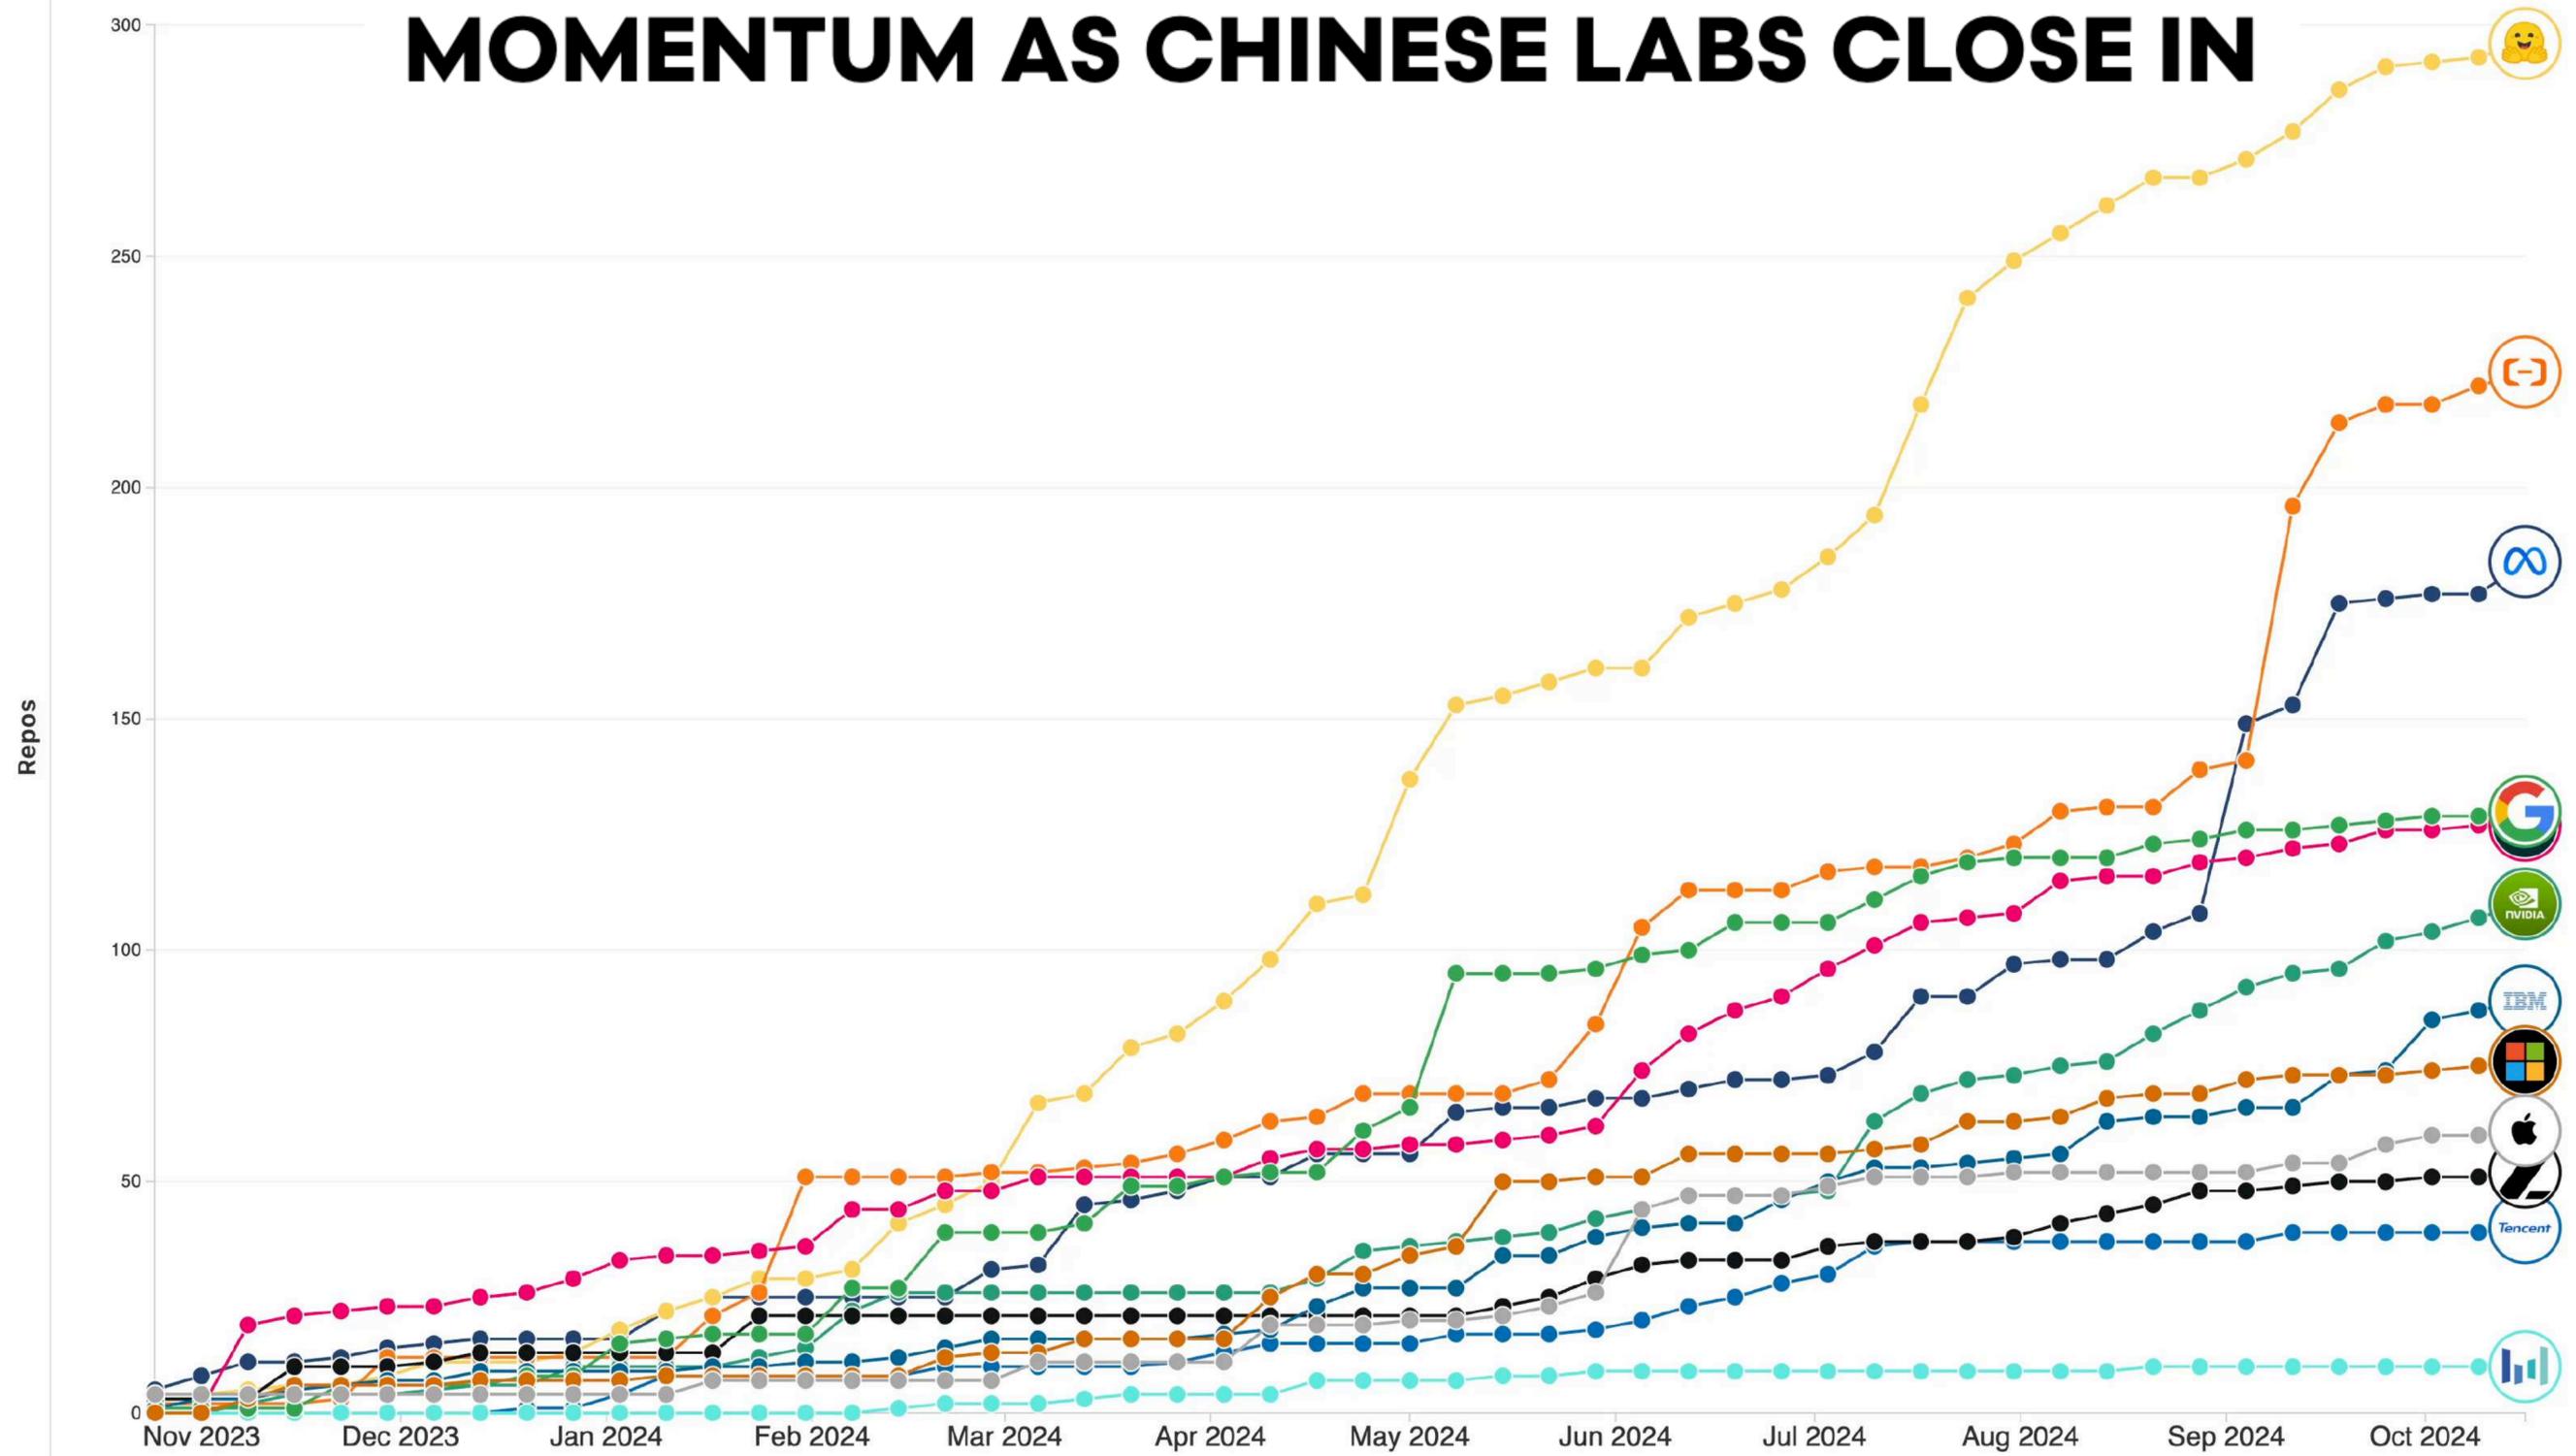
Sam Altman



It's hopeless to compete with OpenAI



# NVIDIA LEADS OPEN SOURCE AI MOMENTUM AS CHINESE LABS CLOSE IN



October 23, 2024

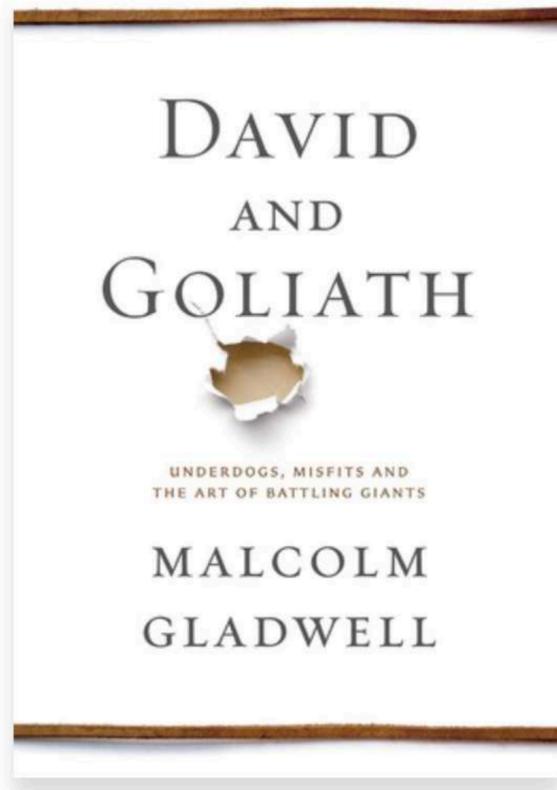
\* Considering repos above 10 space likes and 500 downloads



## NVIDIA LEADS OPEN SOURCE AI MOMENTUM AS CHINESE LABS CLOSE IN



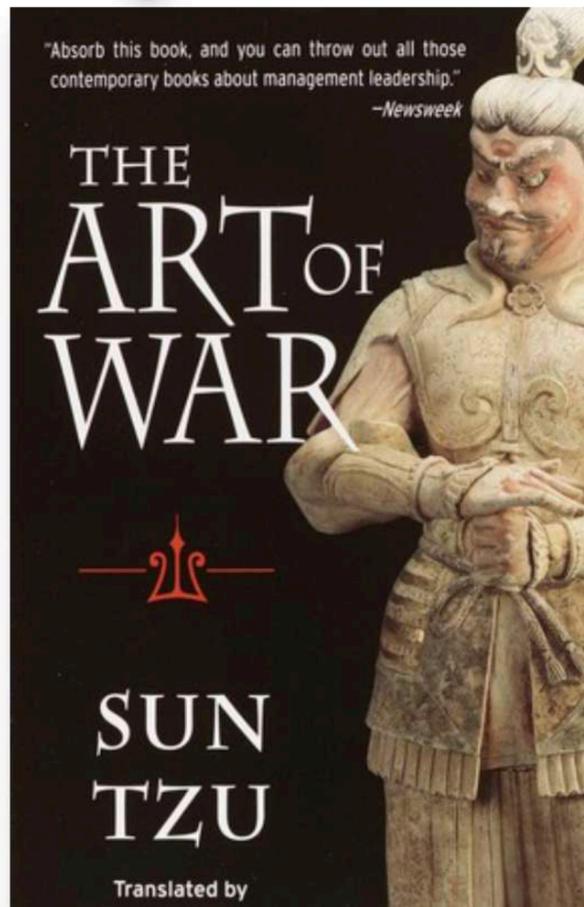
1. Small models are in high demand, with daily downloads hitting tens of millions (!!! 🔥) for top repositories
2. China's contribution to the open source ecosystem has skyrocketed, with Alibaba now surpassing Hugging Face 😲
3. NVIDIA has emerged as a rising star ★ (or a risen star? 😎), for open-source contributions ❤️



# David and Goliath: Underdogs, Misfits, and the Art of Battling Giants

by Malcolm Gladwell

*David vs. Goliath:  
the Art of Scaling Intelligence  
in the Era of Extreme-Scale Neural Models*



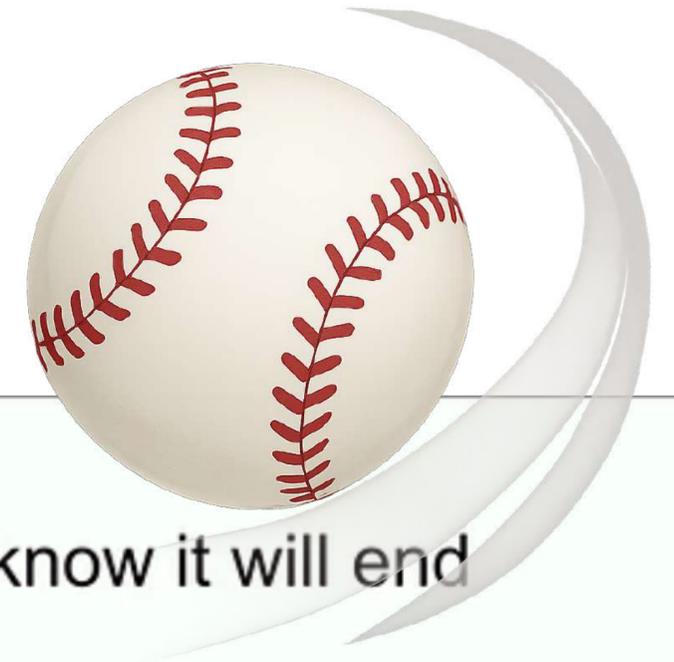
# The Art of War

by Sun Tzu, Thomas Cleary  
Pulat Otkan (Translator)

Three components to innovate:  
**Unconventional** data 🔥  
**Unconventional** algorithms 🚀  
**Unconventional** collaboration 🌍

# The Era of Brute-Force Scaling is Over

The Era of Smart Scaling Begins



Pre-training as we know it will end

Compute is growing:

- Better hardware
- Better algorithms
- Larger clusters

Data is not growing:

- We have but one internet
- **The fossil fuel of AI**

-- Ilya Sutskever from his test time award talk at NeurIPS 2024

# The Era of Brute-Force Scaling is Over

The Era of Smart Scaling Begins

stival.

THE NEW YORKER 100 Newsletter My Accc

re Fiction & Poetry Humor & Cartoons Magazine Puzzles & Games Video Podcasts G

OPEN QUESTIONS

## WHAT IF A.I. DOESN'T GET MUCH BETTER THAN THIS?

*GPT-5, a new release from OpenAI, is the latest product to suggest that progress on large language models has stalled.*

By Cal Newport  
August 12, 2025



Pre-training as we know it will end

Compute is growing:

- Better hardware
- Better algorithms
- Larger clusters

Data is not growing:

- We have but one internet
- The fossil fuel of AI

## NewScientist

Sign in

Enter search keywords

News Features Newsletters Podcasts Video Comment Culture Games | **This week's magazine**

Health Space Physics Technology Environment Mind Humans Life Mathematics Chemistry Earth Society

**Analysis and Technology**

## GPT-5's modest gains suggest AI progress is slowing down

OpenAI's latest large language model has achieved seemingly underwhelming improvements in performance, leading to questions about whether the AI industry can make significant advancements with its current designs

By [Alex Wilkins](#)

13 August 2025



# The Era of Brute-Force Scaling is Over

The Era of Smart Scaling Begins

Three broad ways to cope with data saturation:

## 1. Learn better & faster with limited data

- Alternative architectures
- Alternative training recipes

## 2. Synthesize new data

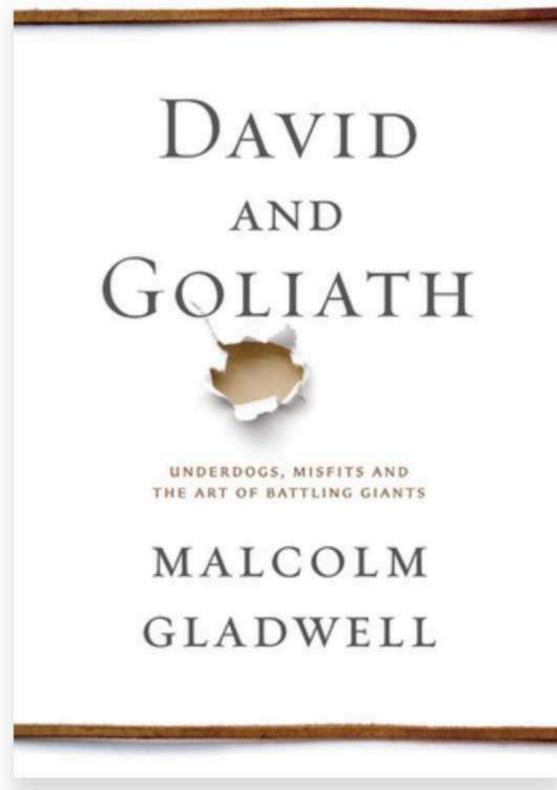
- Algorithms to generate the outer space of the internet data

## 3. Reason beyond what is in the data

- Test-time reasoning algorithms
- Test-time training algorithms

=> Scaling of “intelligence” continues

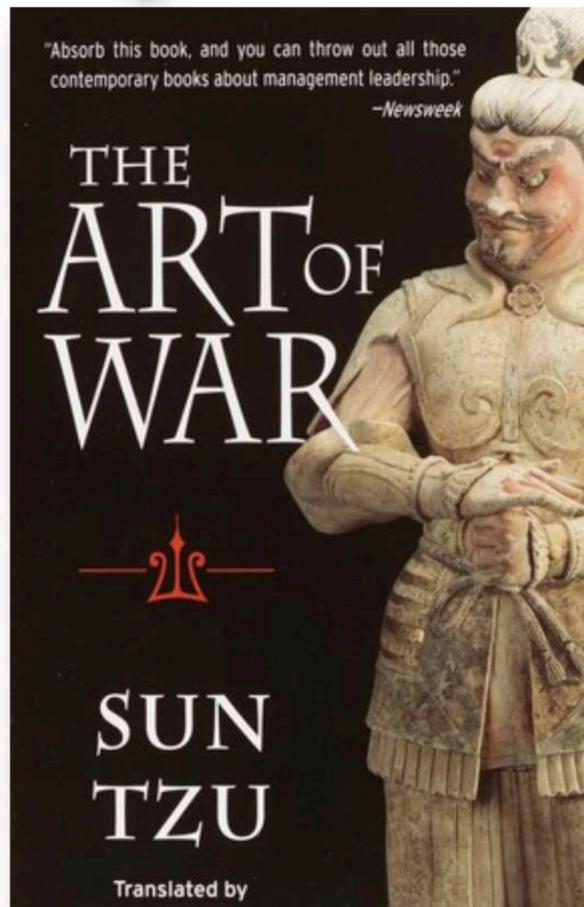




# David and Goliath: Underdogs, Misfits, and the Art of Battling Giants

by Malcolm Gladwell

*David vs. Goliath:  
the Art of Scaling Intelligence  
in the Era of Extreme-Scale Neural Models*



# The Art of War

by Sun Tzu, Thomas Cleary  
Pulat Otkan (Translator)

Three components to innovate:  
**Unconventional** data 🔥  
**Unconventional** algorithms 🚀  
**Unconventional** collaboration 🌍

# h: Underdogs, Art of Battling

In this talk:

===

ProRL: Prolonged RL

Prismatic Synthesis

RL as Pretraining

===

“Smaller but Better”

“Algorithms for the Win”



*David vs. Goliath:*

*the Art of Scaling Intelligence*

*in the Era of Extreme-Scale Neural Models*

early  
r)

Three components to innovate:

Unconventional data 🔥

Unconventional algorithms 🚀

Unconventional collaboration 🌍



— *NeurIPS 2025* —

# Prolonged Reinforcement Learning Expands Reasoning Boundaries in Large Language Models



**Mingjie Liu   Shizhe Diao   Ximing Lu   Jian Hu   Xin Dong**  
**Yejin Choi   Jan Kautz   Yi Dong**  
NVIDIA

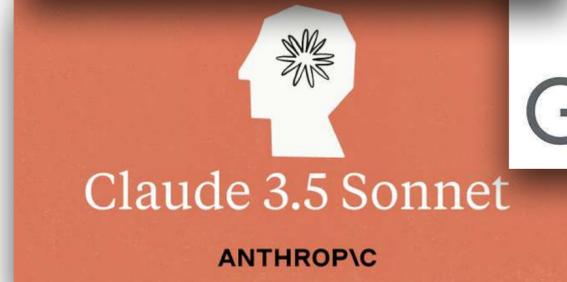
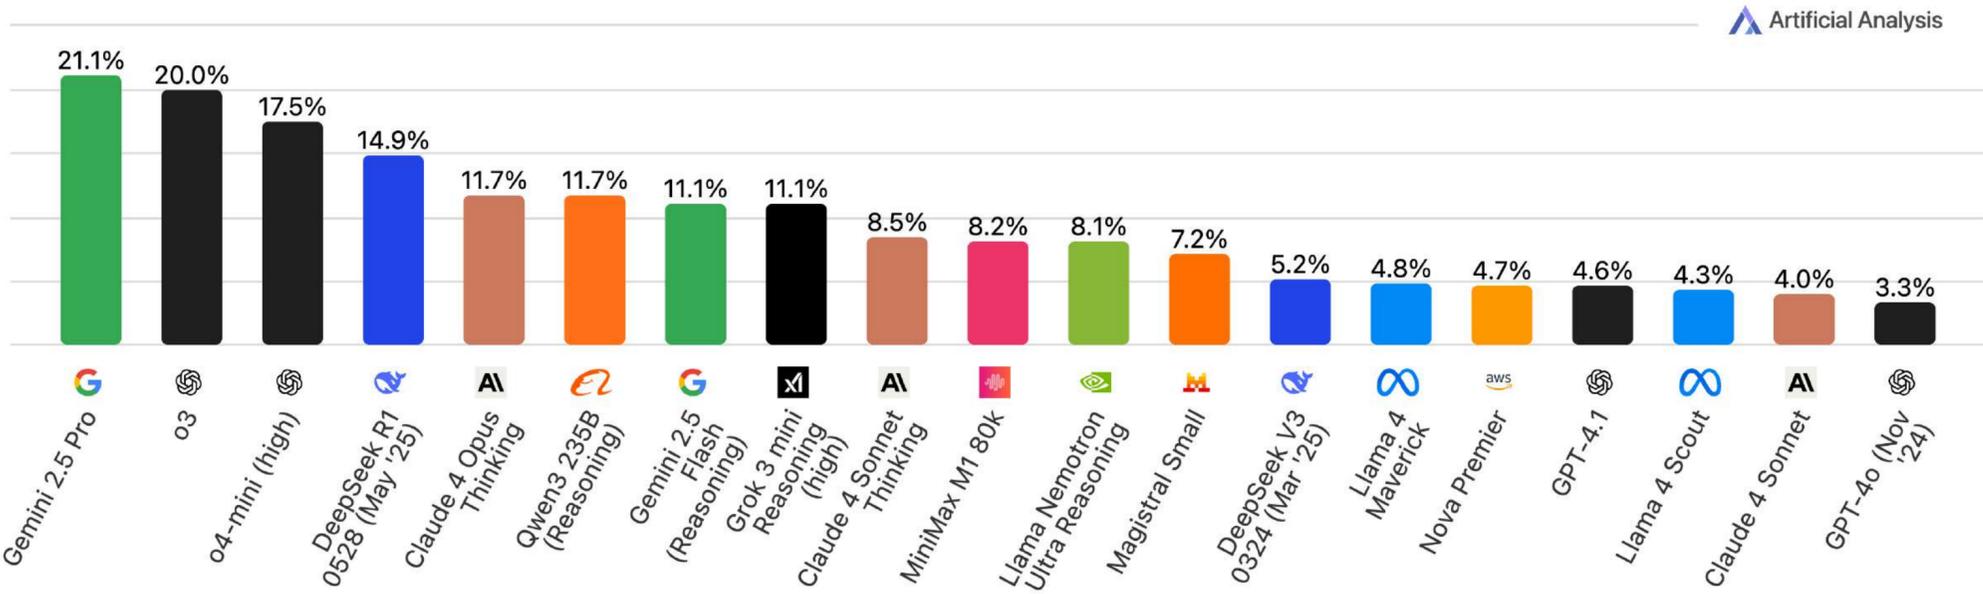
`{mingjiel, sdiao, ximingl, jianh, xind, yejinc, jkautz, yidong}@nvidia.com`

# 2025: The Rise of LRMs (as opposed to LLMs)



## Humanity's Last Exam Benchmark Leaderboard: Results

Independently conducted by Artificial Analysis



- **Long thought** (chain-of-thought)
- The power of **Reinforcement Learning**
- Shift from **Imitation Learning** to **Exploration Learning**

# 2025: The Rise of LRMs (as opposed to LLMs)

## Does Math Reasoning Improve General LLM Capabilities? Understanding Transferability of LLM Reasoning

Maggie Huan<sup>1,2,\*</sup>, Yuetai Li<sup>3,\*</sup>, Tuney Zheng<sup>4,\*</sup>, Xiaoyu Xu<sup>5</sup>, Seungone Kim<sup>1</sup>, Minxin Du<sup>5</sup>,  
Radha Poovendran<sup>3</sup>, Graham Neubig<sup>1</sup>, Xiang Yue<sup>1,†</sup>

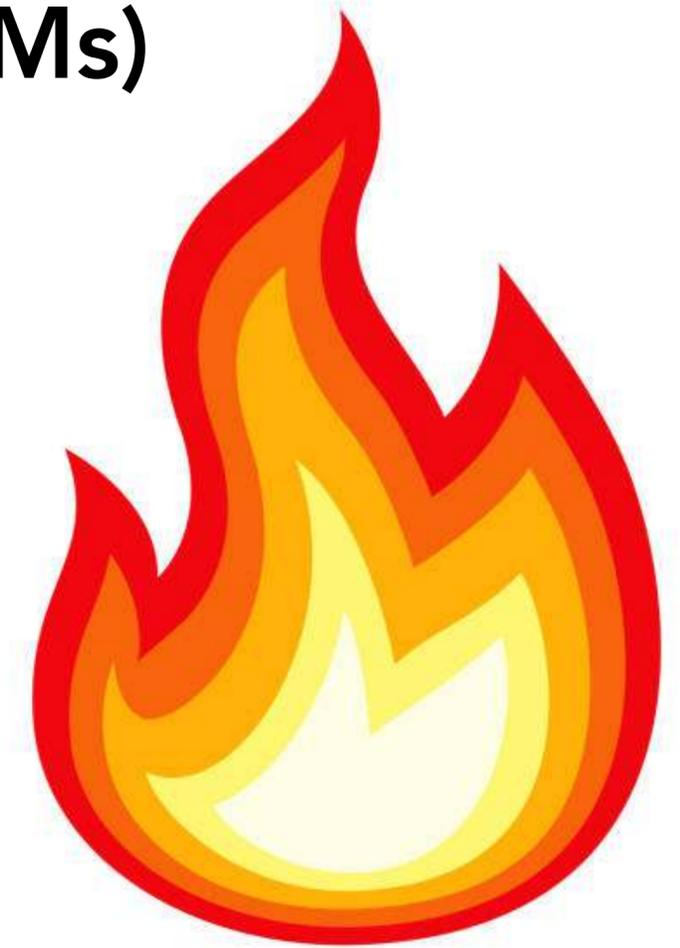
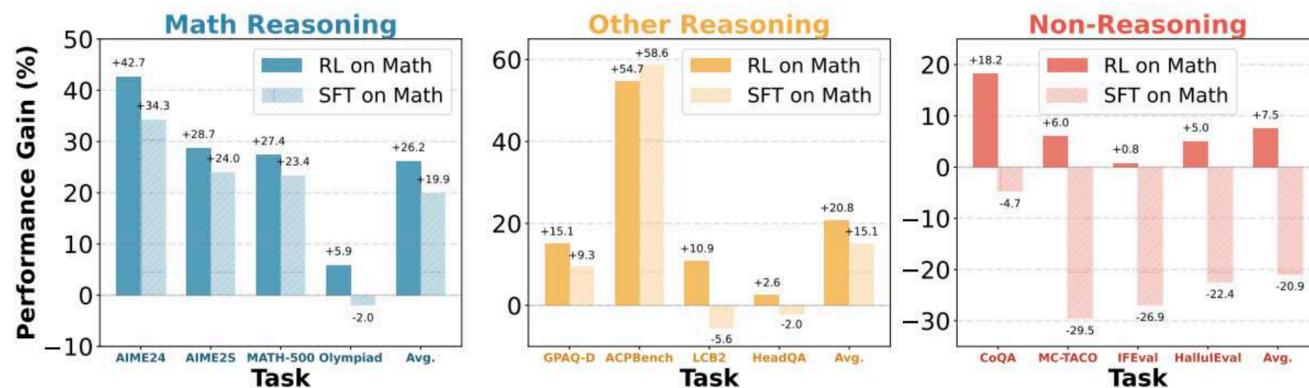
<sup>1</sup>Carnegie Mellon University <sup>2</sup>University of Pennsylvania <sup>3</sup>University of Washington  
<sup>4</sup>M-A-P <sup>5</sup>The Hong Kong Polytechnic University

ziyuh@seas.upenn.edu yuetaili@uw.edu xyue2@andrew.cmu.edu

**Abstract:** Math reasoning has become the poster child of progress in large language models (LLMs), with new models rapidly surpassing human-level performance on benchmarks like MATH and AIME. But as math leaderboards improve week by week, it is worth asking: *do these gains reflect broader problem-solving ability or just narrow overfitting?* To answer this question, we evaluate over 20 open-weight reasoning-tuned models across a broad suite of tasks, including math, scientific QA, agent planning, coding, and standard instruction-following. We surprisingly find that most models that succeed in math fail to transfer their gains to other domains. To rigorously study this phenomenon, we conduct controlled experiments on Qwen3-14B models using math-only data but different tuning methods. We find that reinforcement learning (RL)-tuned models generalize well across domains, while supervised fine-tuning (SFT)-tuned models often forget general capabilities. Latent-space representation and token-space distribution shift analyses reveal that SFT induces substantial representation and output drift, while RL preserves general-domain structure. Our results suggest a need to rethink standard post-training recipes, particularly the reliance on SFT-distilled data for advancing reasoning models.

[github.com/ReasoningTransfer/Transferability-of-LLM-Reasoning](https://github.com/ReasoningTransfer/Transferability-of-LLM-Reasoning)

[huggingface.co/ReasoningTransferability](https://huggingface.co/ReasoningTransferability)



## SFT Memorizes, RL Generalizes: A Comparative Study of Foundation Model Post-training

Tianzhe Chu<sup>✳️</sup> Yuexiang Zhai<sup>♥️✳️</sup> Jihan Yang<sup>♦️</sup> Shengbang Tong<sup>♦️</sup>  
Saining Xie<sup>✳️♦️</sup> Dale Schuurmans<sup>✳️</sup> Quoc V. Le<sup>✳️</sup> Sergey Levine<sup>♥️</sup> Yi Ma<sup>✳️♥️</sup>

# 2025: The Rise of LRMs (as opposed to LLMs)

## Does Math Reasoning Improve General LLM Capabilities? Understanding Transferability of LLM Reasoning

Maggie Huan<sup>1,2,\*</sup>, Yuetai Li<sup>3,\*</sup>, Tuney Zheng<sup>4,\*</sup>, Xiaoyu Xu<sup>5</sup>, Seungone Kim<sup>1</sup>, Minxin Du<sup>5</sup>,  
Radha Poovendran<sup>3</sup>, Graham Neubig<sup>1</sup>, Xiang Yue<sup>1,†</sup>

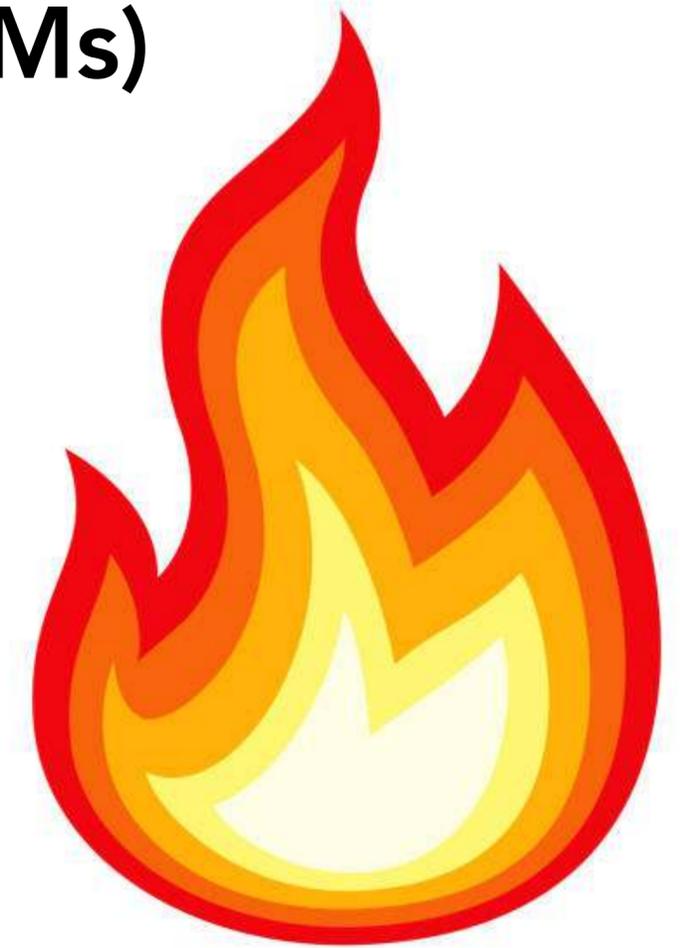
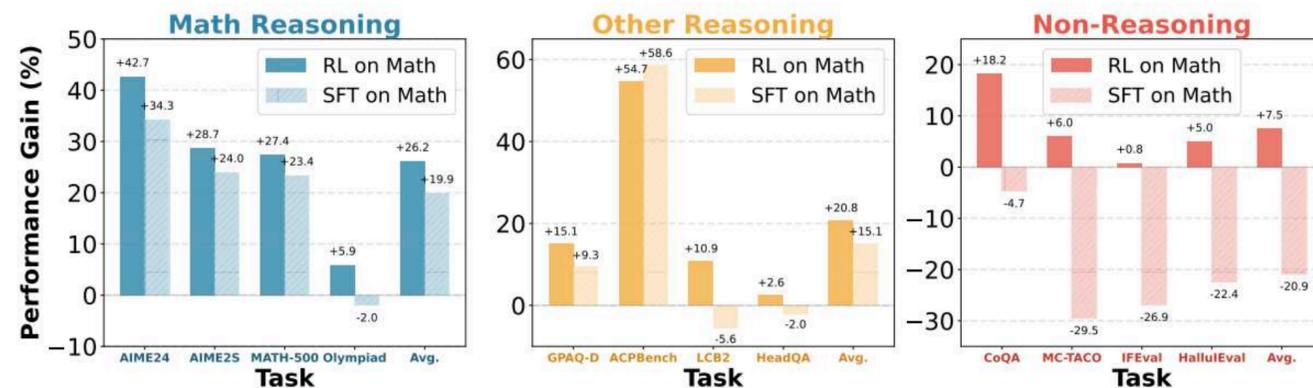
<sup>1</sup>Carnegie Mellon University <sup>2</sup>University of Pennsylvania <sup>3</sup>University of Washington  
<sup>4</sup>M-A-P <sup>5</sup>The Hong Kong Polytechnic University

ziyuh@seas.upenn.edu yuetaili@uw.edu xyue2@andrew.cmu.edu

**Abstract:** Math reasoning has become the poster child of progress in large language models (LLMs), with new models rapidly surpassing human-level performance on benchmarks like MATH and AIME. But as math leaderboards improve week by week, it is worth asking: *do these gains reflect broader problem-solving ability or just narrow overfitting?* To answer this question, we evaluate over 20 open-weight reasoning-tuned models across a broad suite of tasks, including math, scientific QA, agent planning, coding, and standard instruction-following. We surprisingly find that most models that succeed in math fail to transfer their gains to other domains. To rigorously study this phenomenon, we conduct controlled experiments on Qwen3-14B models using math-only data but different tuning methods. We find that reinforcement learning (RL)-tuned models generalize well across domains, while supervised fine-tuning (SFT)-tuned models often forget general capabilities. Latent-space representation and token-space distribution shift analyses reveal that SFT induces substantial representation and output drift, while RL preserves general-domain structure. Our results suggest a need to rethink standard post-training recipes, particularly the reliance on SFT-distilled data for advancing reasoning models.

[github.com/ReasoningTransfer/Transferability-of-LLM-Reasoning](https://github.com/ReasoningTransfer/Transferability-of-LLM-Reasoning)

[huggingface.co/ReasoningTransferability](https://huggingface.co/ReasoningTransferability)



## SFT Memorizes, RL Generalizes: A Comparative Study of Foundation Model Post-training

Tianzhe Chu<sup>✳\*</sup> Yuexiang Zhai<sup>♥✳\*</sup> Jihan Yang<sup>♦</sup> Shengbang Tong<sup>♦</sup>  
Saining Xie<sup>✳♦</sup> Dale Schuurmans<sup>✳\*</sup> Quoc V. Le<sup>✳</sup> Sergey Levine<sup>♥</sup> Yi Ma<sup>✳♥</sup>

# Striking findings against RL



May 19, 2025



## Does Reinforcement Learning Really Incentivize Reasoning Capacity in LLMs Beyond the Base Model?

Yang Yue<sup>1\*†</sup>, Zhiqi Chen<sup>1\*</sup>, Rui Lu<sup>1</sup>, Andrew Zhao<sup>1</sup>, Zhaokai Wang<sup>2</sup>, Yang Yue<sup>1</sup>, Shiji Song<sup>1</sup>, and Gao Huang<sup>1✉</sup>

<sup>1</sup> LeapLab, Tsinghua University    <sup>2</sup> Shanghai Jiao Tong University

\* Equal Contribution    † Project Lead    ✉ Corresponding Author

**Pass@1 better after RLVR**  
**Pass@K worse after RLVR**  
— compared to the base LLM —

# Striking findings against RL



May 19, 2025



## Does Reinforcement Learning Really Incentivize Reasoning Capacity in LLMs Beyond the Base Model?

Yang Yue<sup>1\*†</sup>, Zhiqi Chen<sup>1\*</sup>, Rui Lu<sup>1</sup>, Andrew Zhao<sup>1</sup>, Zhaokai Wang<sup>2</sup>, Yang Yue<sup>1</sup>, Shiji Song<sup>1</sup>, and Gao Huang<sup>1✉</sup>

<sup>1</sup> LeapLab, Tsinghua University    <sup>2</sup> Shanghai Jiao Tong University

\* Equal Contribution    † Project Lead    ✉ Corresponding Author

**Pass@1 better after RLVR**  
**Pass@K worse after RLVR**  
— compared to the base LLM —

# Cracking findings against RL



May 19, 2025

## Does Reinforcement Learning Increase Capacity in LLMs

Yang Yue<sup>1\*†</sup>, Zhiqi Chen<sup>1</sup>, Shiji Song<sup>1</sup>, and Gao Huang<sup>2</sup>

<sup>1</sup> LeapLab, Tsinghua University    <sup>2</sup> Shanghai Jiao Tong University

\* Equal Contribution    † Project Lead    ✉ Corresponding Author

...ing

Passage better after RLVR

Passage worse after RLVR

— compared to the base LLM —

## Echo Chamber: RL Post-training Amplifies Behaviors in Pretraining

**Rosie Zhao\***  
Harvard University  
Kempner Institute

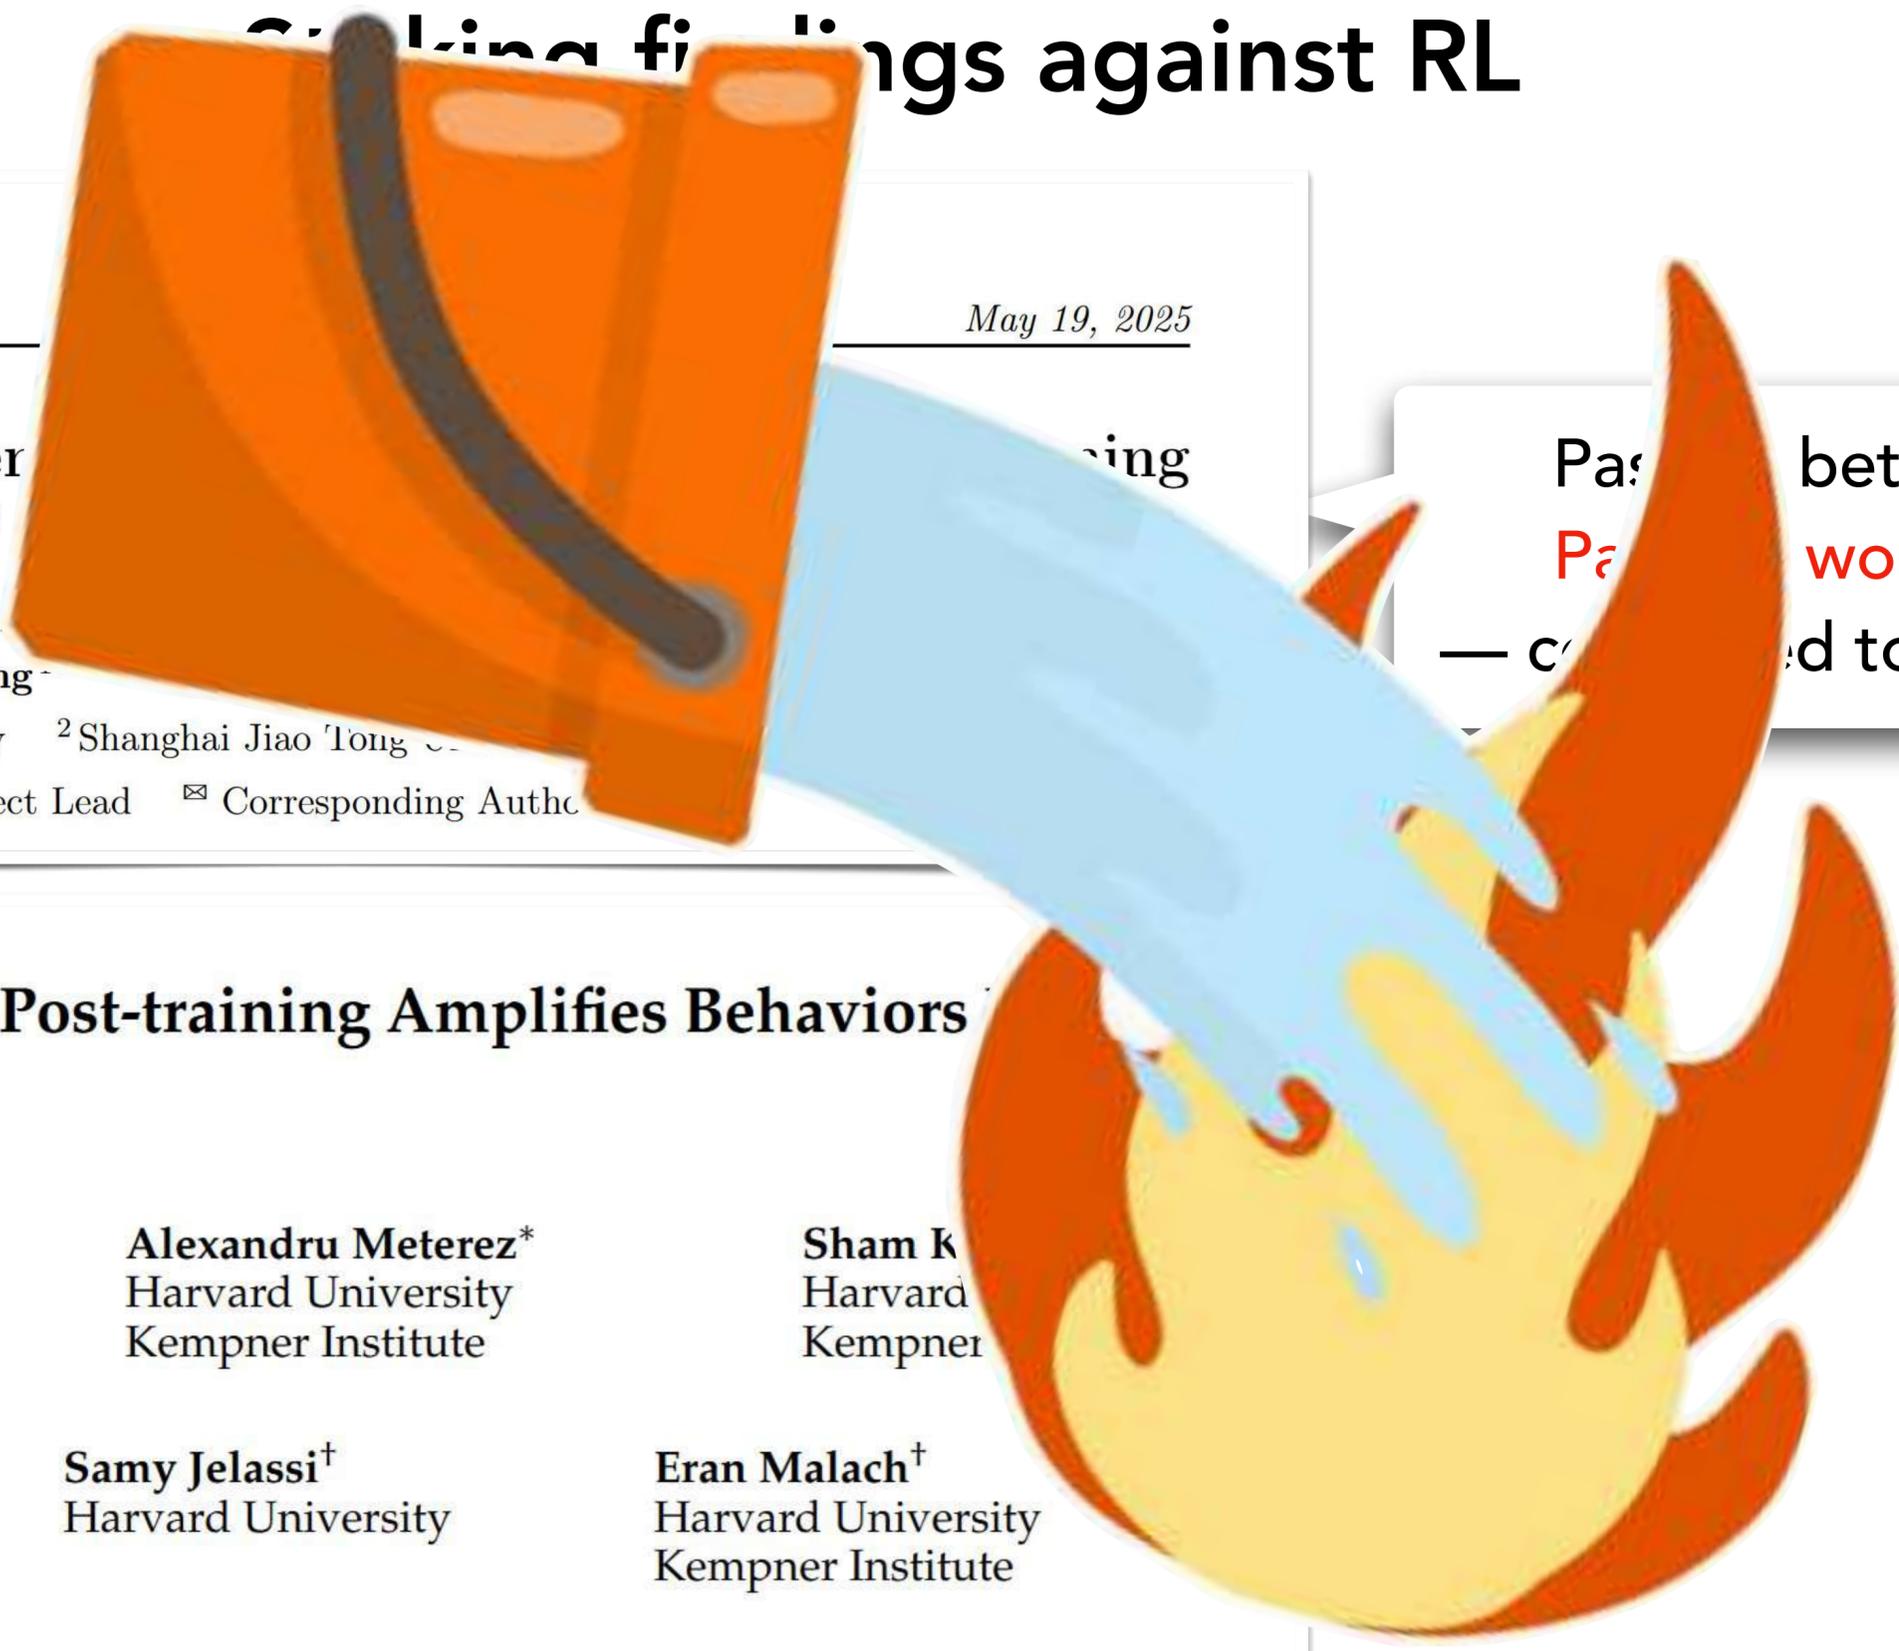
**Alexandru Meterez\***  
Harvard University  
Kempner Institute

**Sham K.**  
Harvard University  
Kempner Institute

**Cengiz Pehlevan**  
Harvard University  
Kempner Institute

**Samy Jelassi<sup>†</sup>**  
Harvard University

**Eran Malach<sup>†</sup>**  
Harvard University  
Kempner Institute



We cracked RLVR with... Random Rewards?!

Training Qwen2.5-Math-7B with our Spurious Rewards improved MATH-500 by:

- Random rewards: +21%
- Incorrect rewards: +25%
- (FYI) Ground-truth rewards: + 28.8%

How could this even work!?! Here's why:

Blogpost: [tinyurl.com/spurious-rewar...](https://tinyurl.com/spurious-rewar...)

# Striking findings against RL



## Spurious Rewards: Rethinking Training Signals in RLVR

Rulin Shao<sup>1\*</sup> Shuyue Stella Li<sup>1\*</sup> Rui Xin<sup>1\*</sup> Scott Geng<sup>1\*</sup> Yiping Wang<sup>1</sup>  
Sewoong Oh<sup>1</sup> Simon Shaolei Du<sup>1</sup> Nathan Lambert<sup>2</sup> Sewon Min<sup>3</sup> Ranjay Krishna<sup>1,2</sup>  
Yulia Tsvetkov<sup>1</sup> Hannaneh Hajishirzi<sup>1,2</sup> Pang Wei Koh<sup>1,2</sup> Luke Zettlemoyer<sup>1</sup>  
<sup>1</sup>University of Washington <sup>2</sup>Allen Institute for Artificial Intelligence  
<sup>3</sup>University of California, Berkeley  
{rulins, stelli, rx31, sgeng}@cs.washington.edu

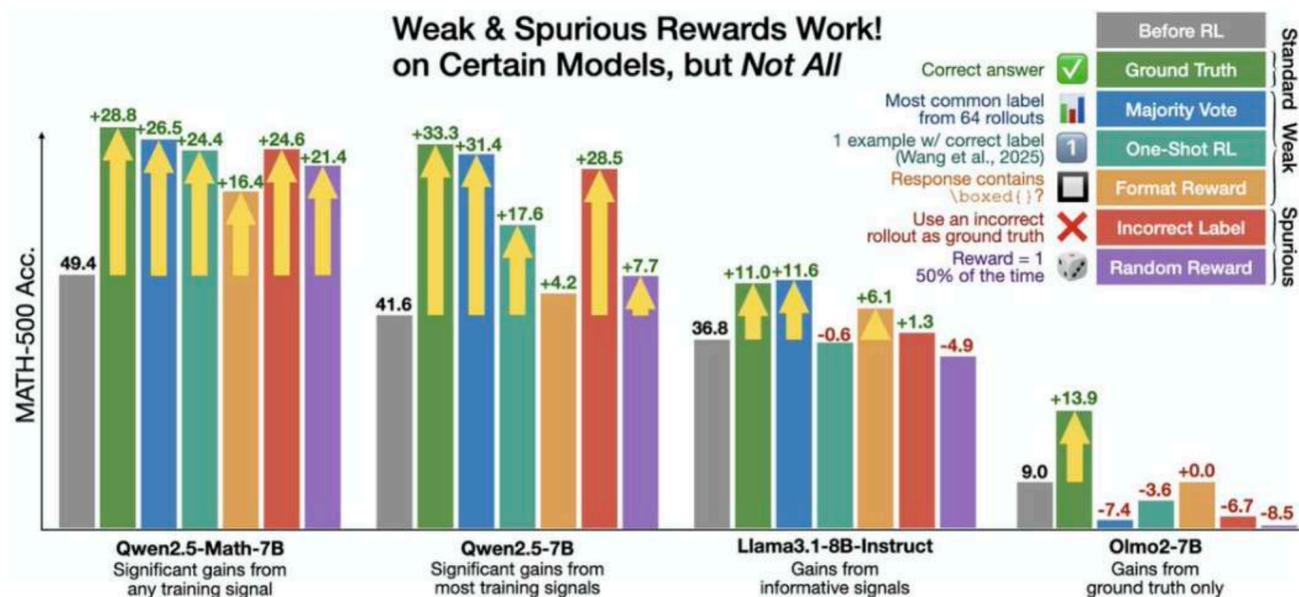


Figure 1: MATH-500 accuracy after 150 steps of RLVR on various training signals. We show that



May 19, 2025

## Does Reinforcement Learning Really Incentivize Reasoning Capacity in LLMs Beyond the Base Model?

Yang Yue<sup>1\*†</sup>, Zhiqi Chen<sup>1\*</sup>, Rui Lu<sup>1</sup>, Andrew Zhao<sup>1</sup>, Zhaokai Wang<sup>2</sup>, Yang Yue<sup>1</sup>, Shiji Song<sup>1</sup>, and Gao Huang<sup>1✉</sup>

<sup>1</sup>LeapLab, Tsinghua University <sup>2</sup>Shanghai Jiao Tong University

\* Equal Contribution † Project Lead ✉ Corresponding Author

## Echo Chamber: RL Post-training Amplifies Behaviors Learned in Pretraining

Rosie Zhao\*  
Harvard University  
Kempner Institute

Alexandru Meterez\*  
Harvard University  
Kempner Institute

Sham Kakade  
Harvard University  
Kempner Institute

Cengiz Pehlevan  
Harvard University  
Kempner Institute

Samy Jelassi†  
Harvard University

Eran Malach†  
Harvard University  
Kempner Institute



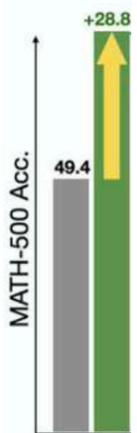
🤖 We cr  
Training C  
500 by:  
- Random  
- Incorre  
- (FYI) Gr  
How coul  
Blogpost:

Spuri

R  
Sewoong  
Yu



st RL



## Andrej Karpathy – “We’re summoning ghosts, not building animals”



**Dwarkesh Patel**  
1.11M subscribers

Subscribe

👍 23K



➦ Share

🌟 Ask

⬇️ Download



**Qwen2.5-Math-7B**  
Significant gains from any training signal

**Qwen2.5-7B**  
Significant gains from most training signals

**Llama3.1-8B-instruct**  
Gains from informative signals

**Olmoez-7B**  
Gains from ground truth only

May 19, 2025

Reasoning

2, Yang Yue<sup>1</sup>,

ors Learned

n Kakade  
ard University  
pner Institute

n<sup>+</sup>  
iversity  
stitute

Figure 1: MATH-500 accuracy after 150 steps of RLVR on various training signals. We show that



Stella Li → CogSci2025

@StellaLisy



🤖 We cracked RLVR with... Random Rewards?!

Training Qwen2.5-Math-7B with our Spurious Rewards improved MATH-500 by:

- Random rewards: +21%
- Incorrect rewards: +25%
- (FYI) Ground-truth rewards: + 28.8%

How could this even work!?! Here's why: 📖

Blogpost: [tinyurl.com/spurious-rewar...](https://tinyurl.com/spurious-rewar...)

# Striking findings against RL



Chemistry between the base LLM and RL matters  
 Conclusions from **effortless** RL  $\neq$  **effortful** RL

{trulins, stell1, rx31, sgeng}@cs.washington.edu

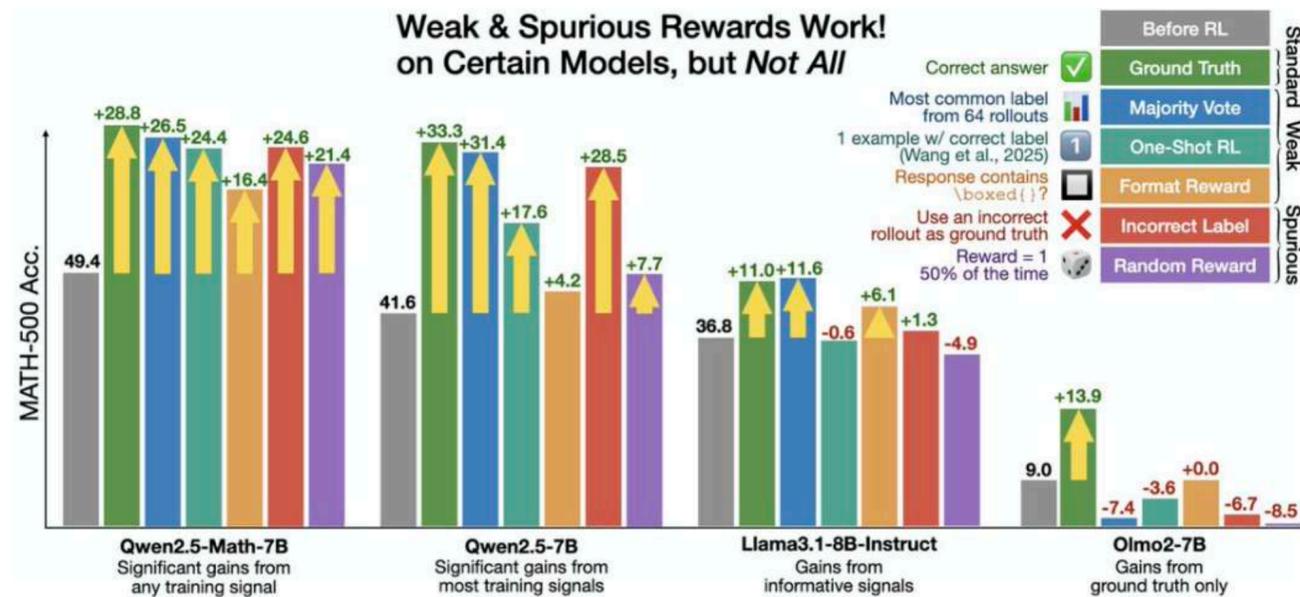


Figure 1: MATH-500 accuracy after 150 steps of RLVR on various training signals. We show that

<sup>1</sup> LeapLab, Tsinghua University    <sup>2</sup> Shanghai Jiao Tong University

\* Equal Contribution    † Project Lead    ✉ Corresponding Author

## Echo Chamber: RL Post-training Amplifies Behaviors Learned in Pretraining

Rosie Zhao\*  
Harvard University  
Kempner Institute

Alexandru Meterez\*  
Harvard University  
Kempner Institute

Sham Kakade  
Harvard University  
Kempner Institute

Cengiz Pehlevan  
Harvard University  
Kempner Institute

Samy Jelassi†  
Harvard University

Eran Malach†  
Harvard University  
Kempner Institute



“Rome wasn't built in a day”

What if we go for endurance with a small model (1.5B) rather than a sprint with a large model

— *NeurIPS 2025* —

# Prolonged Reinforcement Learning Expands Reasoning Boundaries in Large Language Models



**Mingjie Liu   Shizhe Diao   Ximing Lu   Jian Hu   Xin Dong**

**Yejin Choi   Jan Kautz   Yi Dong**

**NVIDIA**

`{mingjiel, sdiao, ximingl, jianh, xind, yejinc, jkautz, yidong}@nvidia.com`

ByteDance

# Prolonged Reinforcement Learning

Deepseek

ProRL is built on top of one of the best practices in RL—Decoupled Clip and Dynamic Sampling Policy Optimization (**DAPO**), a variation of Group Relative Policy Optimization (GRPO).

ByteDance

# Prolonged Reinforcement Learning

Deepseek

ProRL is built on top of one of the best practices in RL—Decoupled Clip and Dynamic Sampling Policy Optimization (**DAPO**), a variation of Group Relative Policy Optimization (GRPO).

$$\mathcal{L}_{\text{GRPO}}(\theta) = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[ \min \left( r_{\theta}(\tau)A(\tau), \text{clip}(r_{\theta}(\tau), 1 - \epsilon, 1 + \epsilon)A(\tau) \right) \right]$$

- **Dynamic Sampling** - targets mid-difficulty examples to maintain diverse learning signals
- **Decoupled Clipping** - 'clip-higher' uplifts low-probability tokens, encourages broader exploration.

$$\text{clip}(r_{\theta}(\tau), 1 - \epsilon_{\text{low}}, 1 + \epsilon_{\text{high}})$$

# Prolonged Reinforcement Learning

ProRL is built on top of one of the best practices in RL—Decoupled Clip and Dynamic Sampling Policy Optimization (DAPO), a variation of Group Relative Policy Optimization (GRPO).

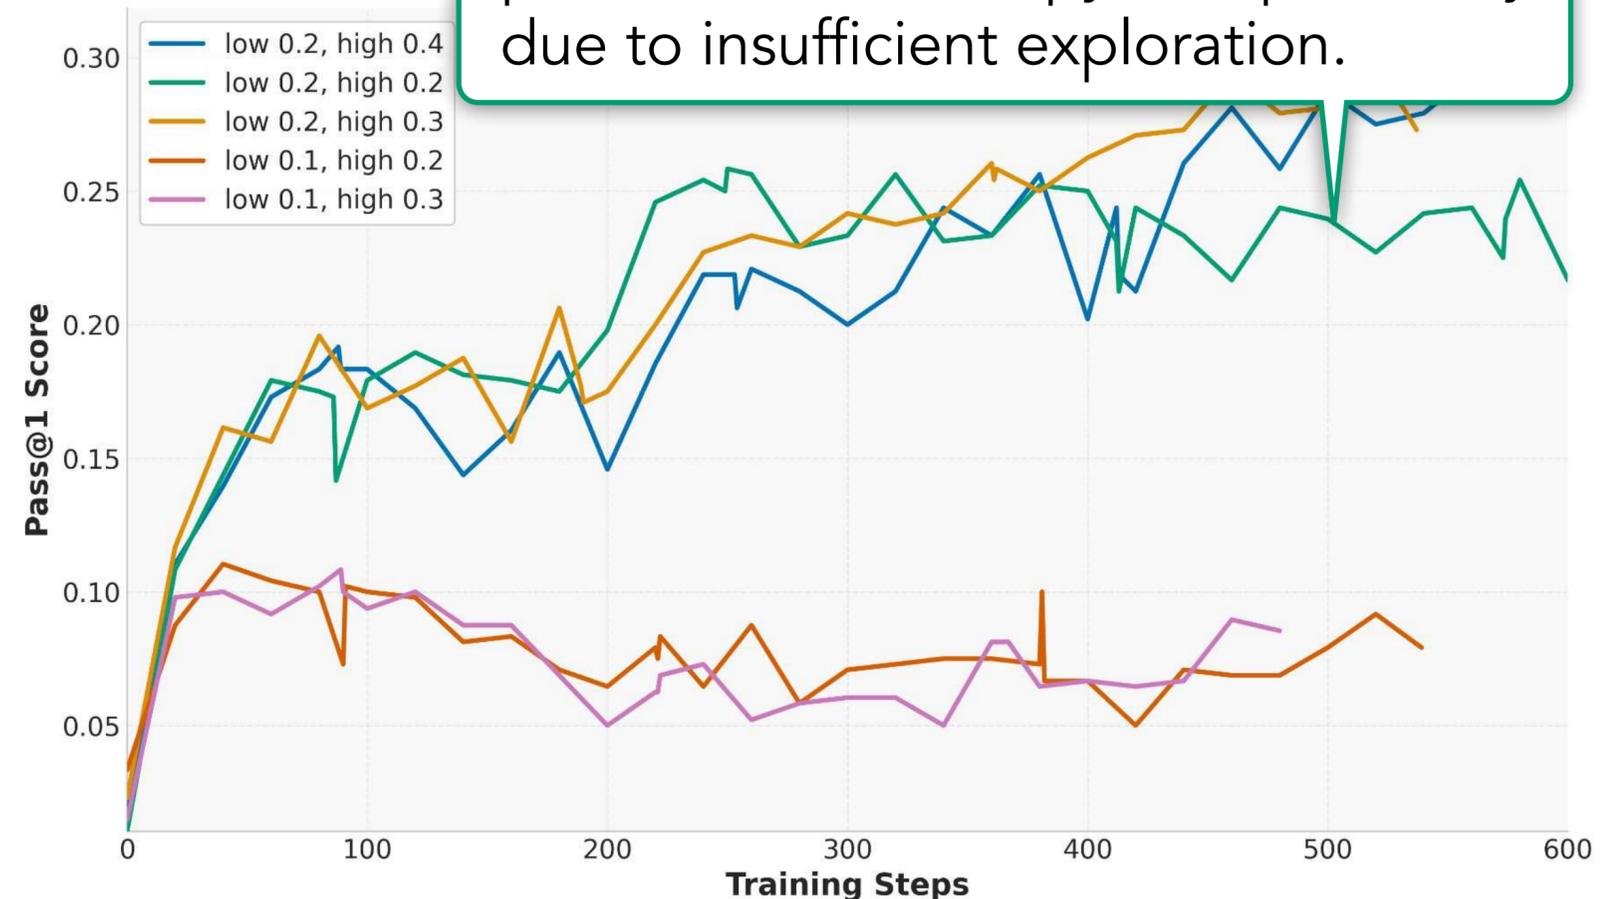
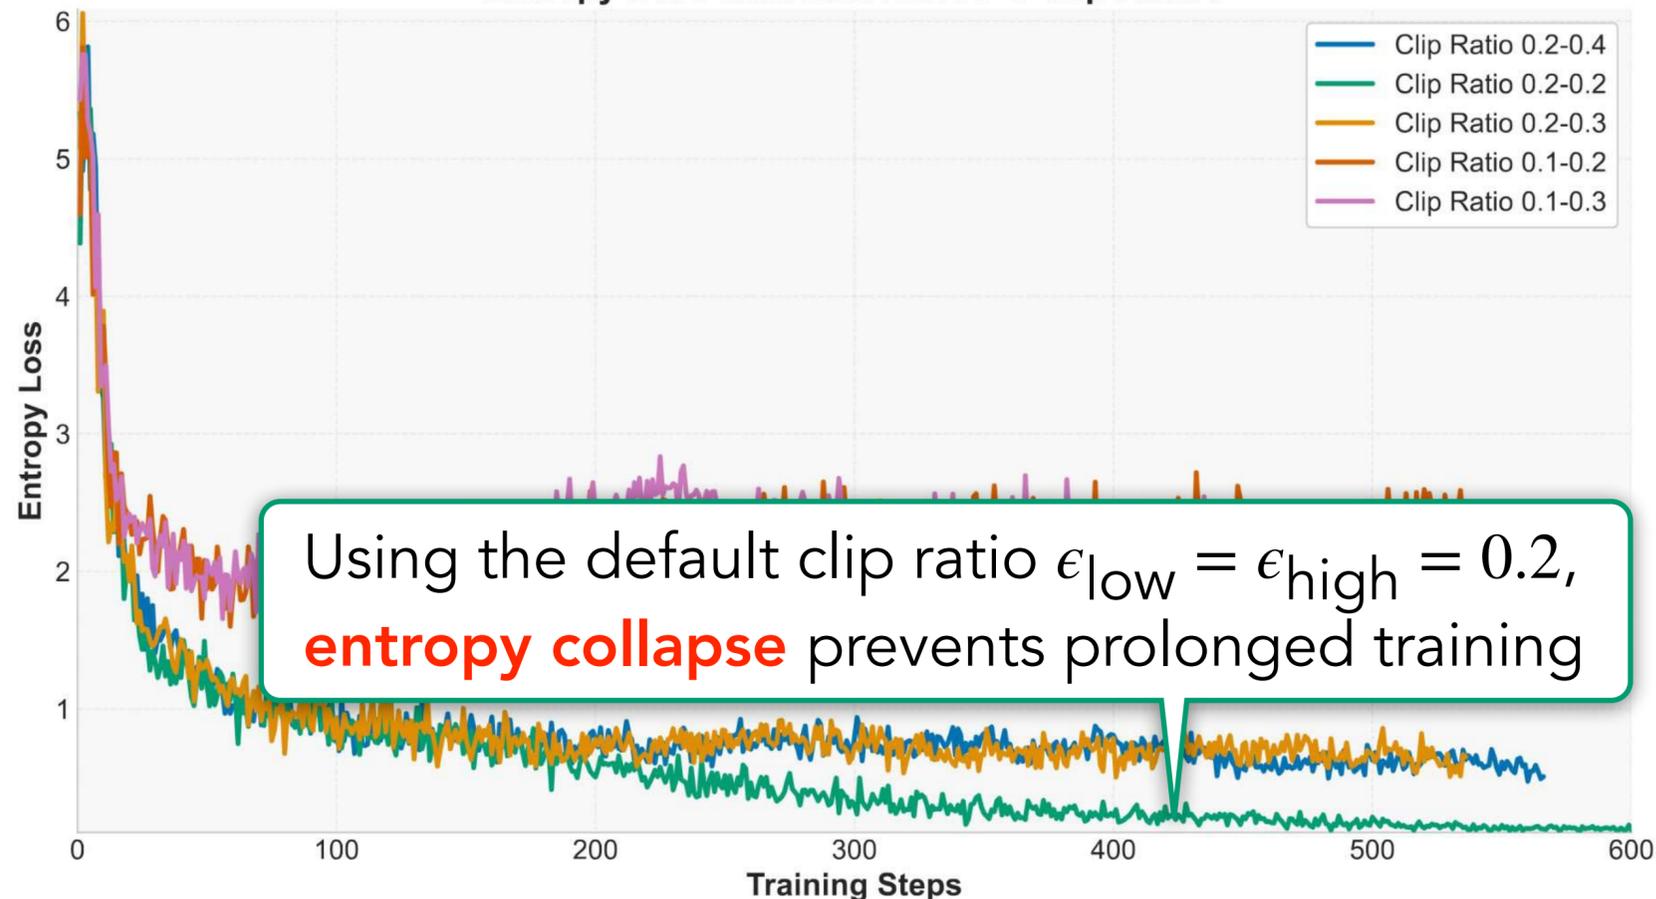


Lesson I: Maintain **sustainable entropy** via clip ratio, balance exploration and exploitation.

$$\text{clip}(r_{\theta}(\tau), 1 - \epsilon_{\text{low}}, 1 + \epsilon_{\text{high}})$$

Performance initially increases but plateaus after entropy collapse, likely due to insufficient exploration.

Entropy Loss with Different PPO Clip Ratios



# Prolonged Reinforcement Learning

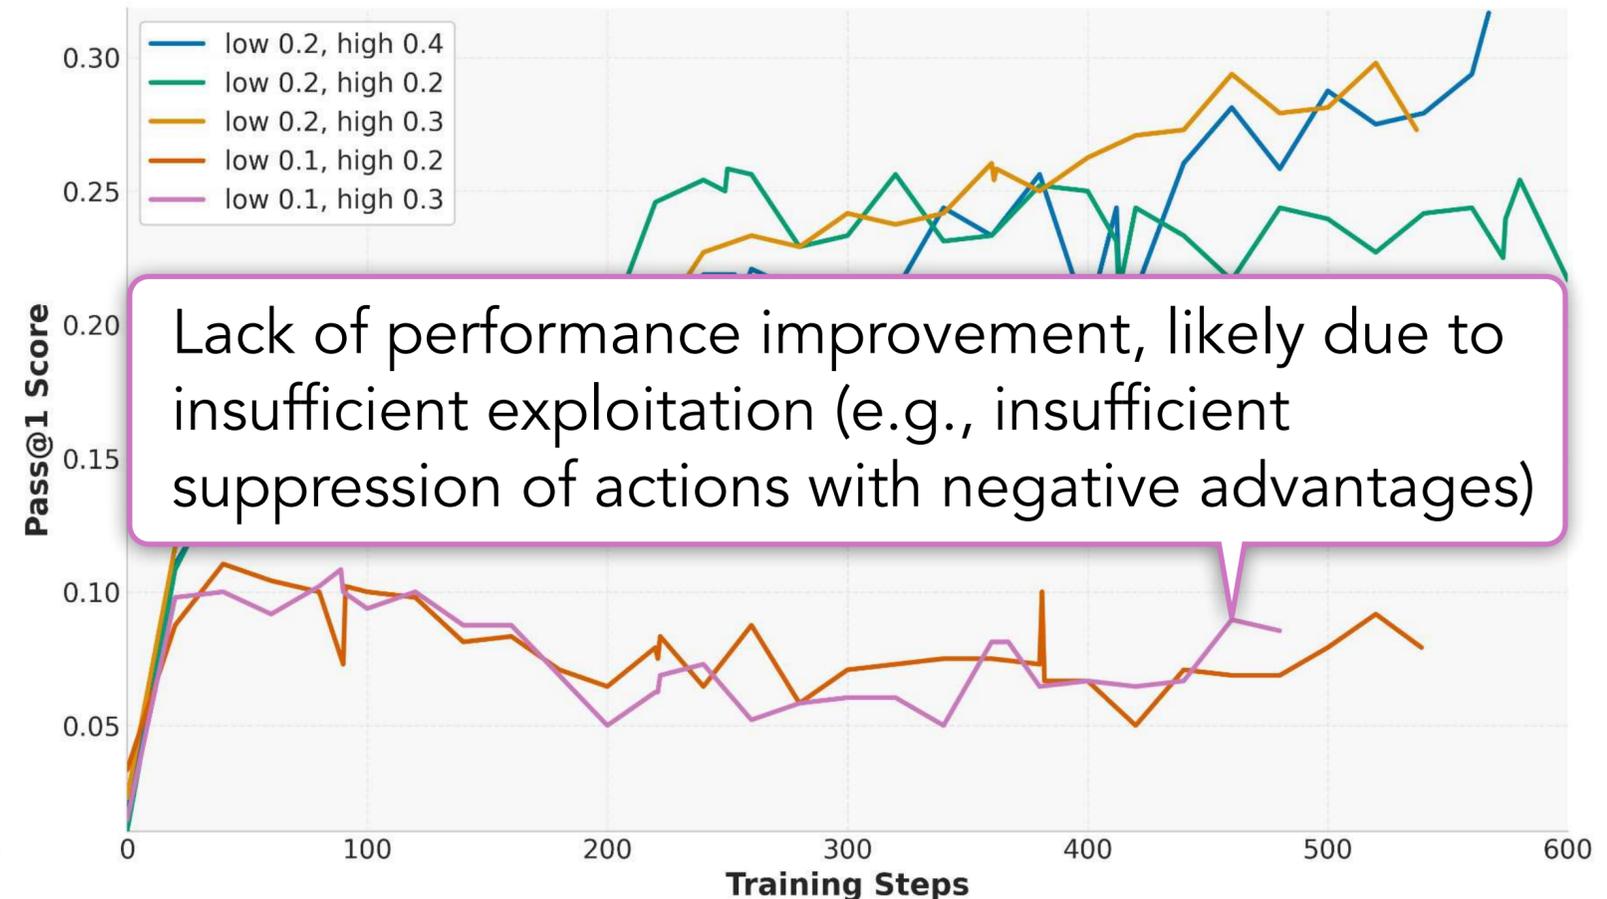
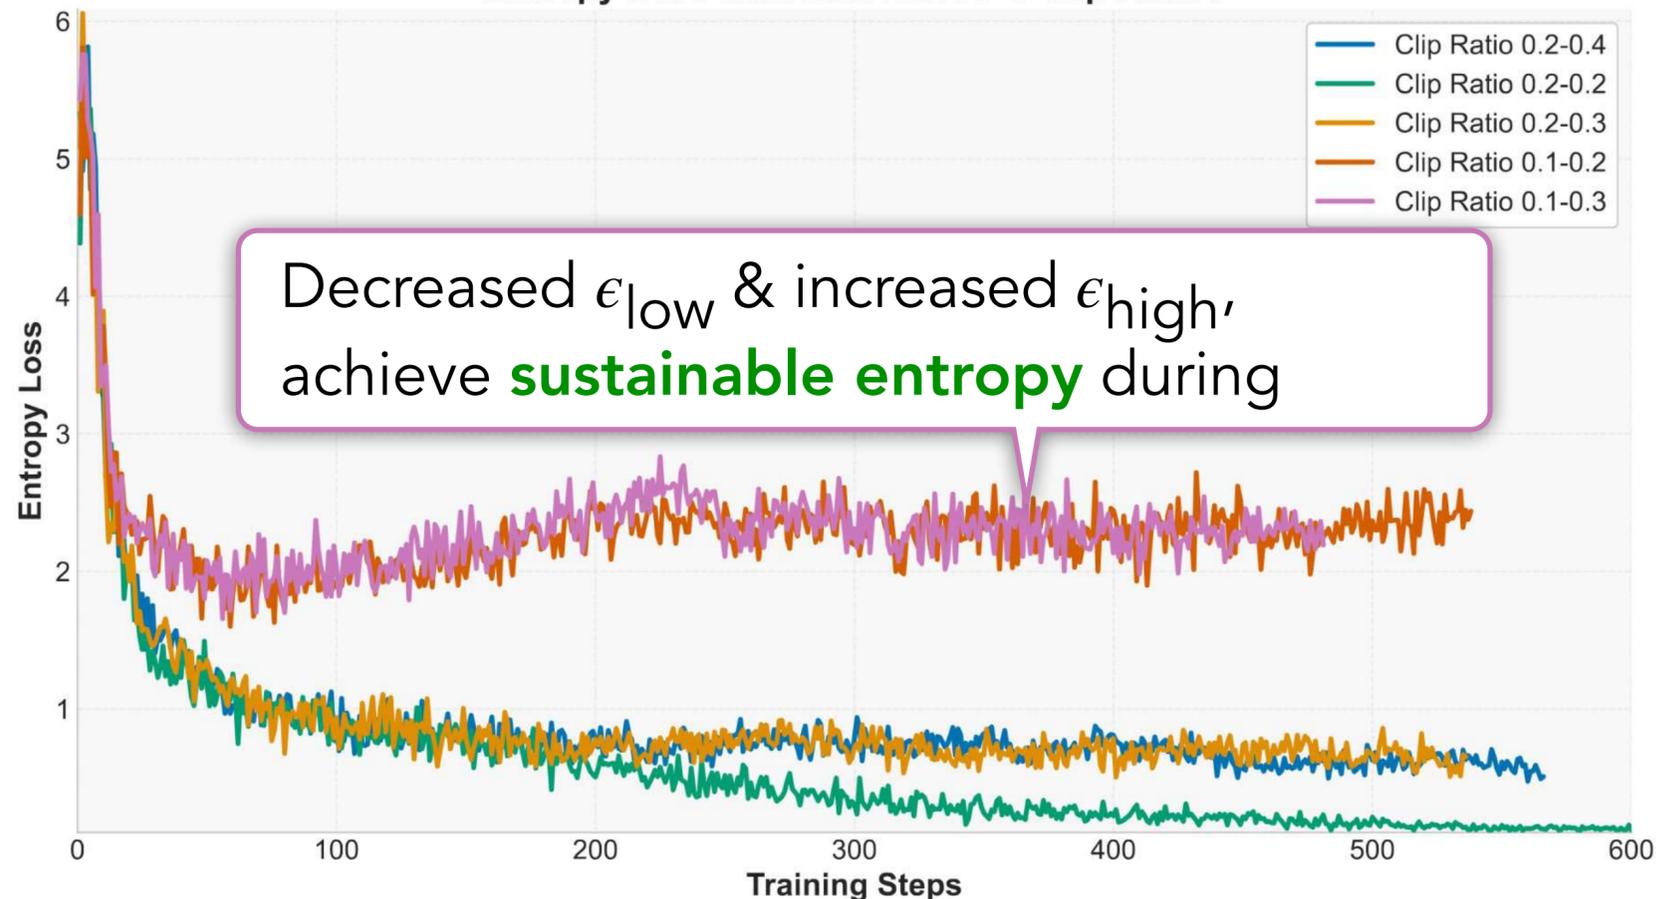
ProRL is built on top of one of the best practices in RL—Decoupled Clip and Dynamic Sampling Policy Optimization (DAPO), a variation of Group Relative Policy Optimization (GRPO).



Lesson I: Maintain **sustainable entropy** via clip ratio, balance exploration and exploitation.

$$\text{clip}(r_{\theta}(\tau), 1 - \epsilon_{\text{low}}, 1 + \epsilon_{\text{high}})$$

Entropy Loss with Different PPO Clip Ratios



# Prolonged Reinforcement Learning

ProRL is built on top of one of the best practices in RL—Decoupled Clip and Dynamic Sampling Policy Optimization (DAPO), a variation of Group Relative Policy Optimization (GRPO).

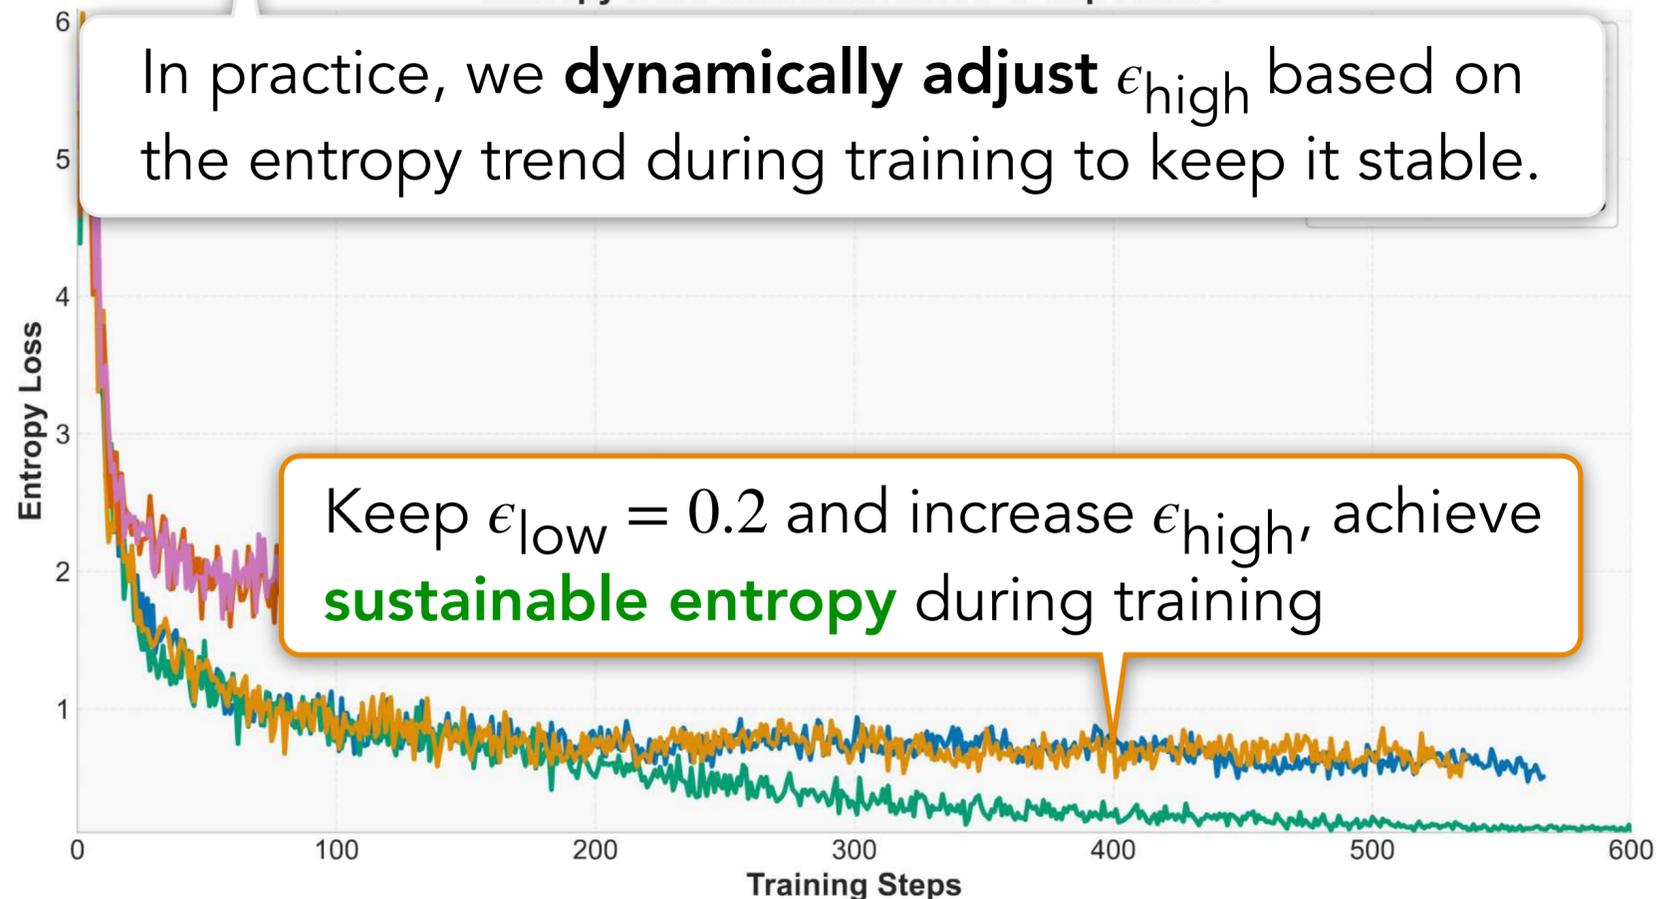


Lesson I: Maintain **sustainable entropy** via clip ratio, balance exploration and exploitation.

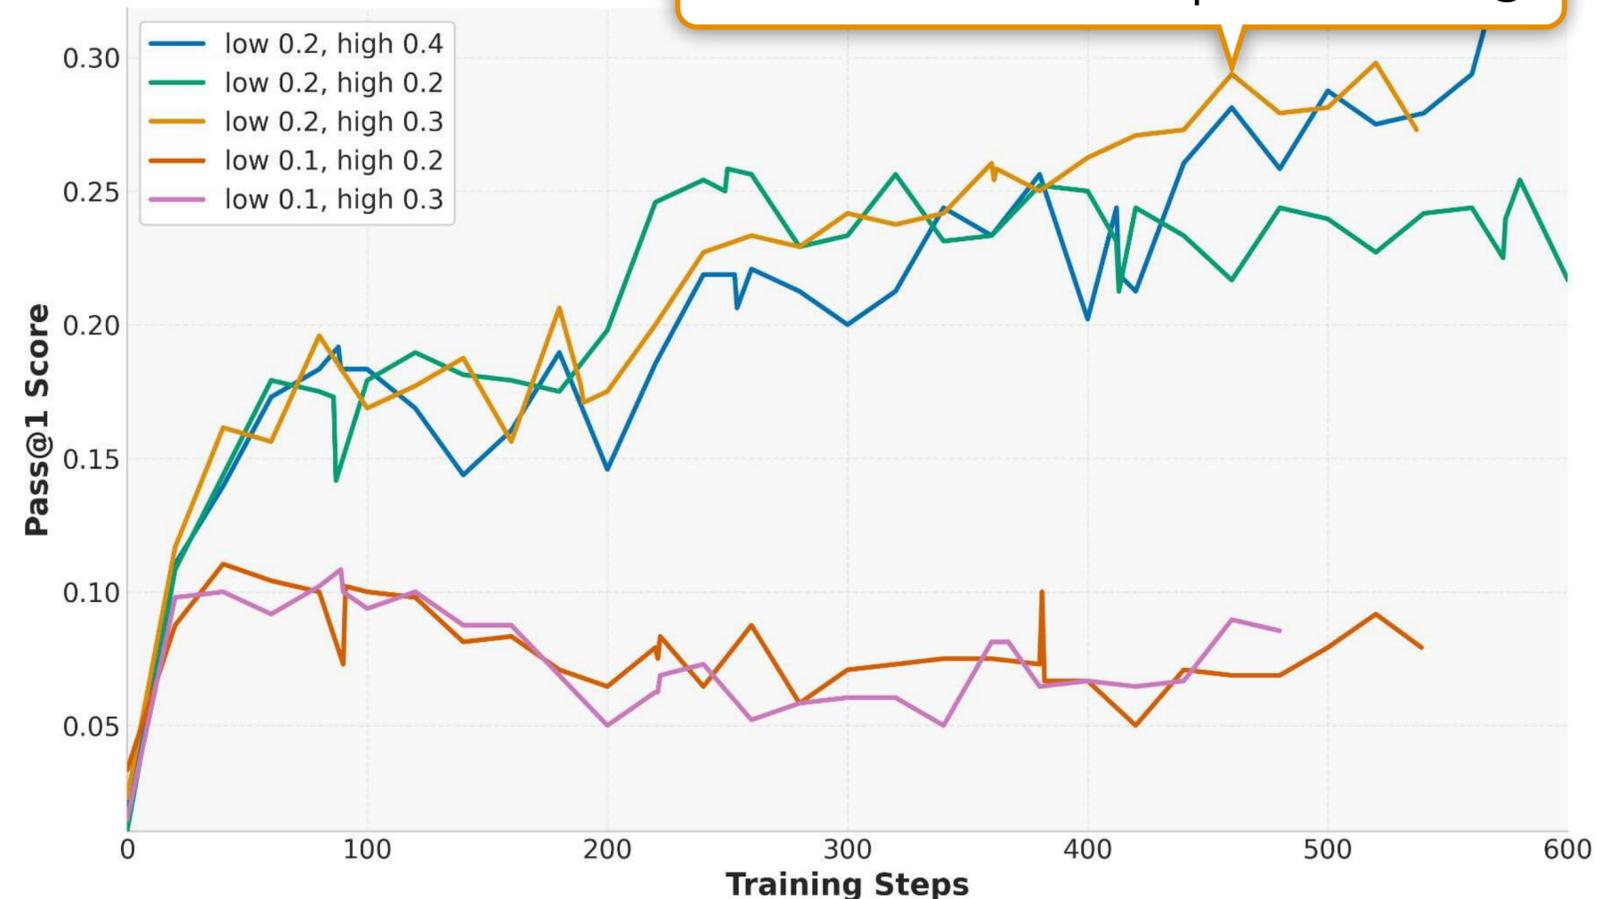
$$\text{clip}(r_{\theta}(\tau), 1 - \epsilon_{\text{low}}, 1 + \epsilon_{\text{high}})$$

Entropy Loss with Different PPO Clip Ratios

In practice, we **dynamically adjust**  $\epsilon_{\text{high}}$  based on the entropy trend during training to keep it stable.



Performance keep increasing

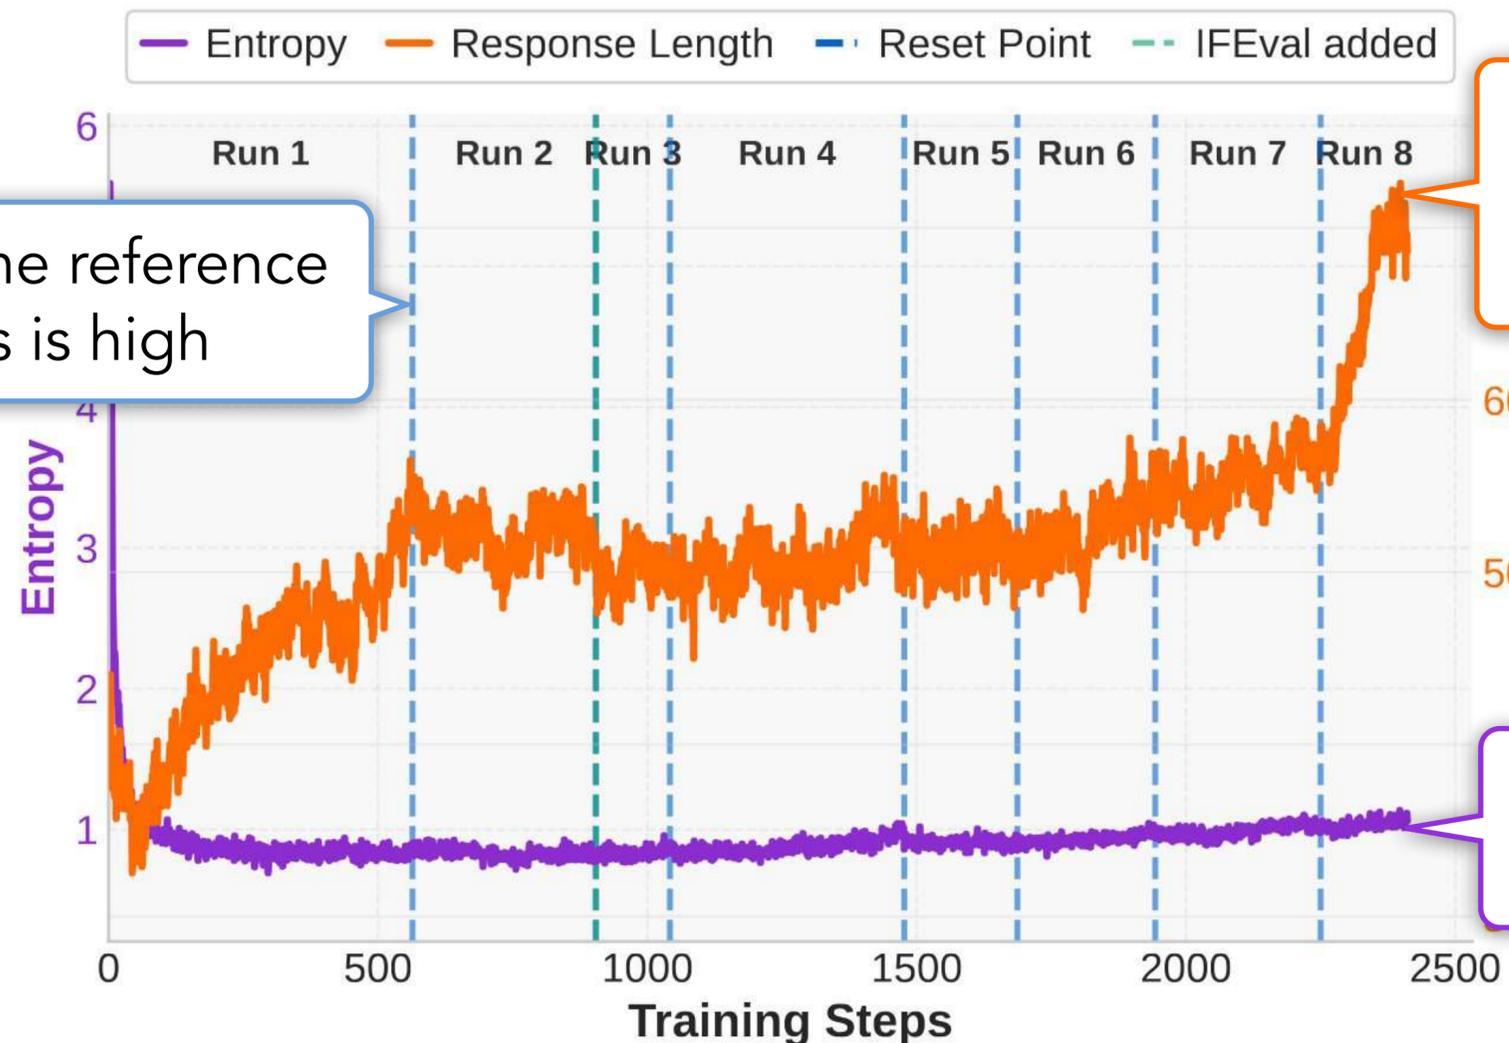


# Prolonged Reinforcement Learning

ProRL is built on top of one of the best practices in RL—Decoupled Clip and Dynamic Sampling Policy Optimization (DAPO), a variation of Group Relative Policy Optimization (GRPO).



Lesson III: Dynamically adjust hyper-parameters to **keep one hero run alive**



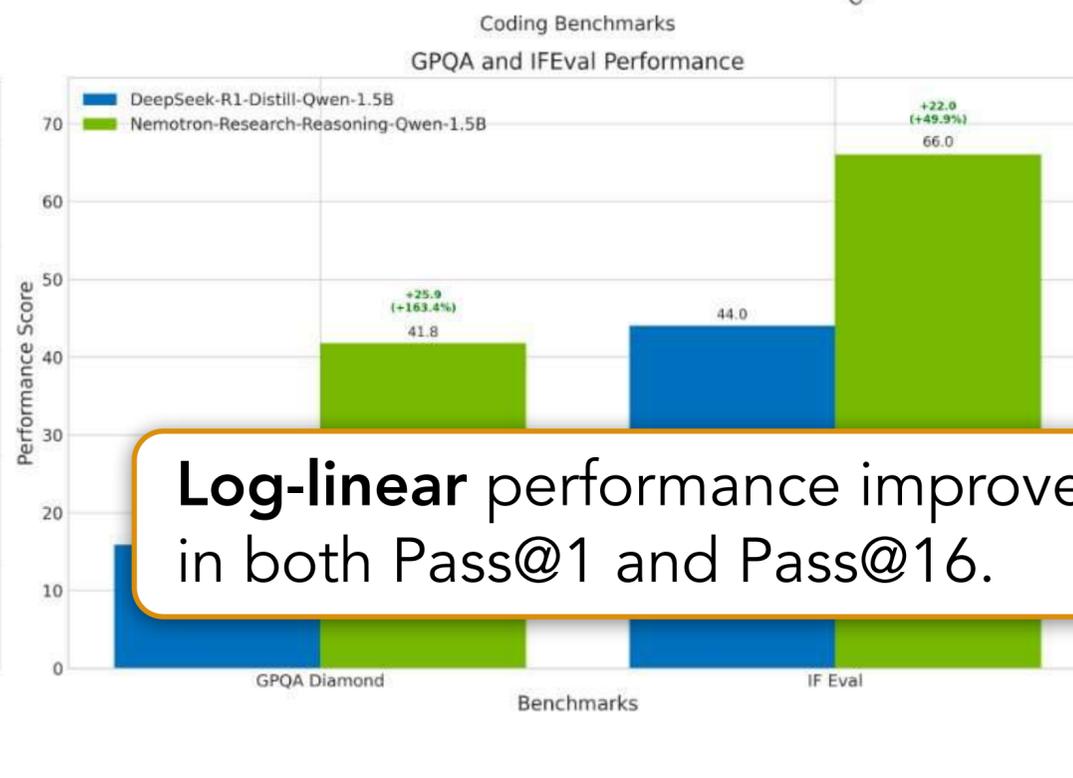
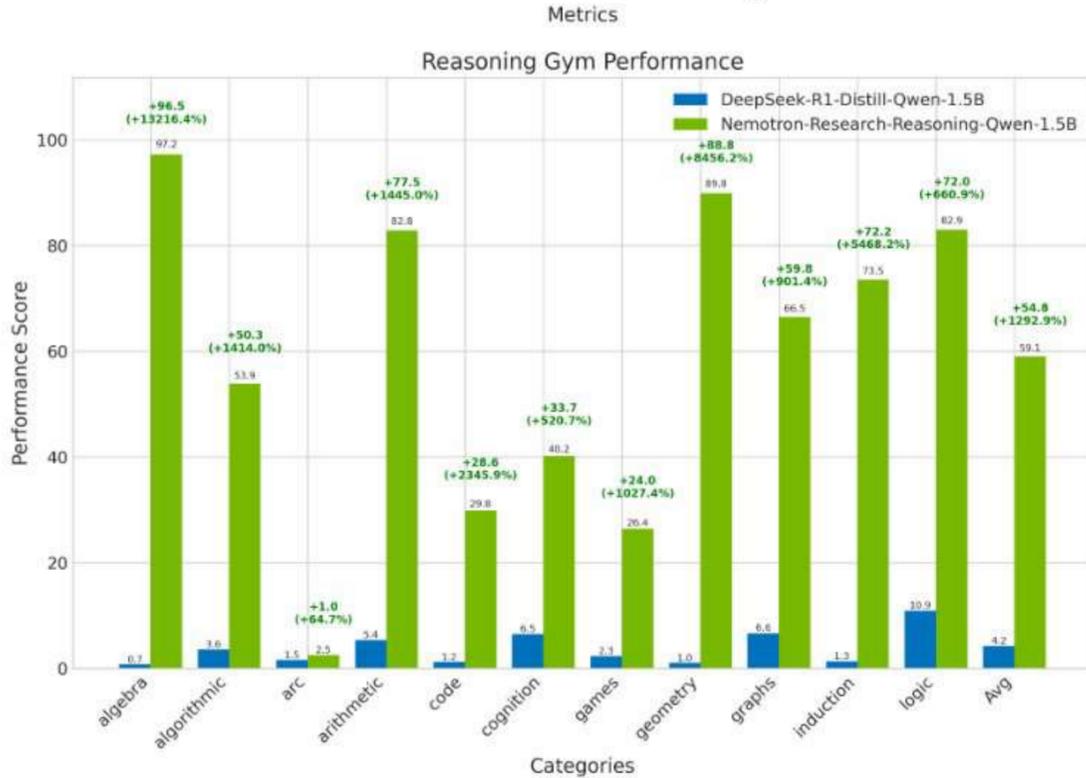
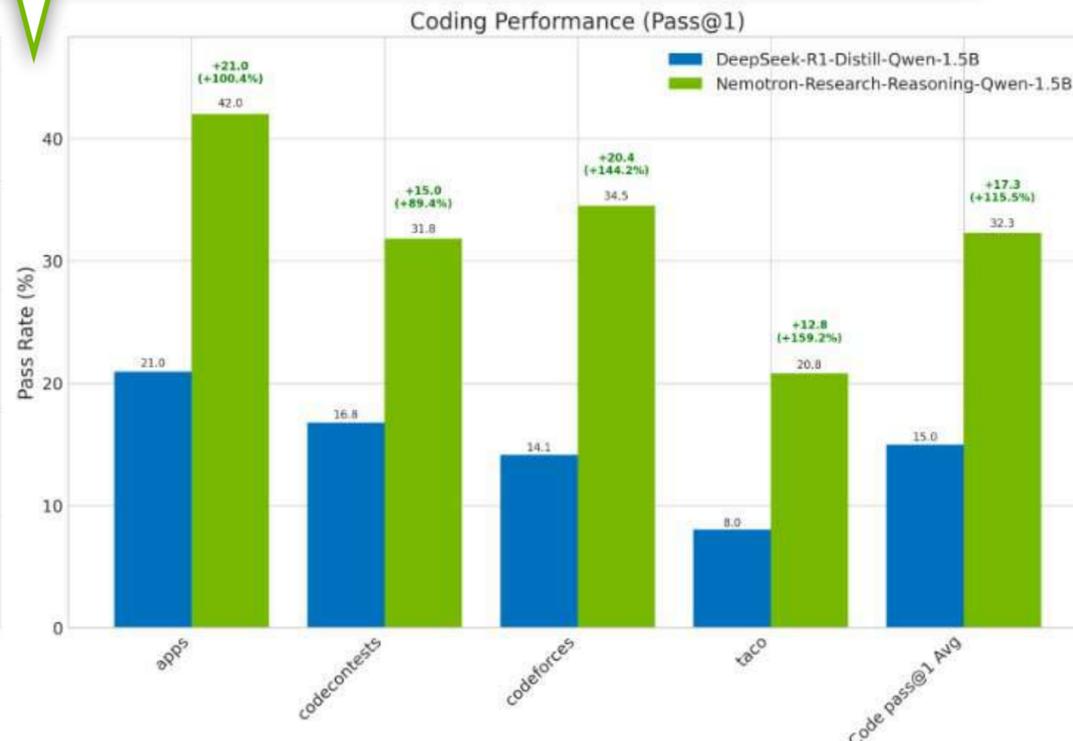
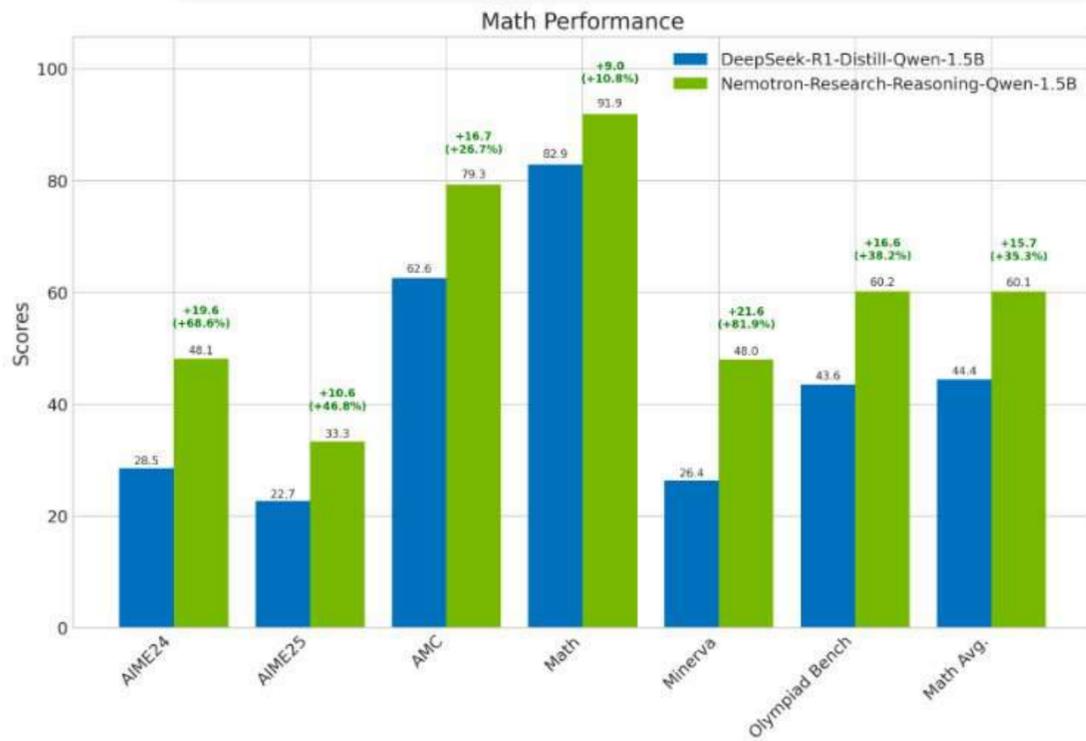
Periodically reset the reference policy when KL loss is high

Keep response length under 8k for training efficiency; extend context length in the final run.

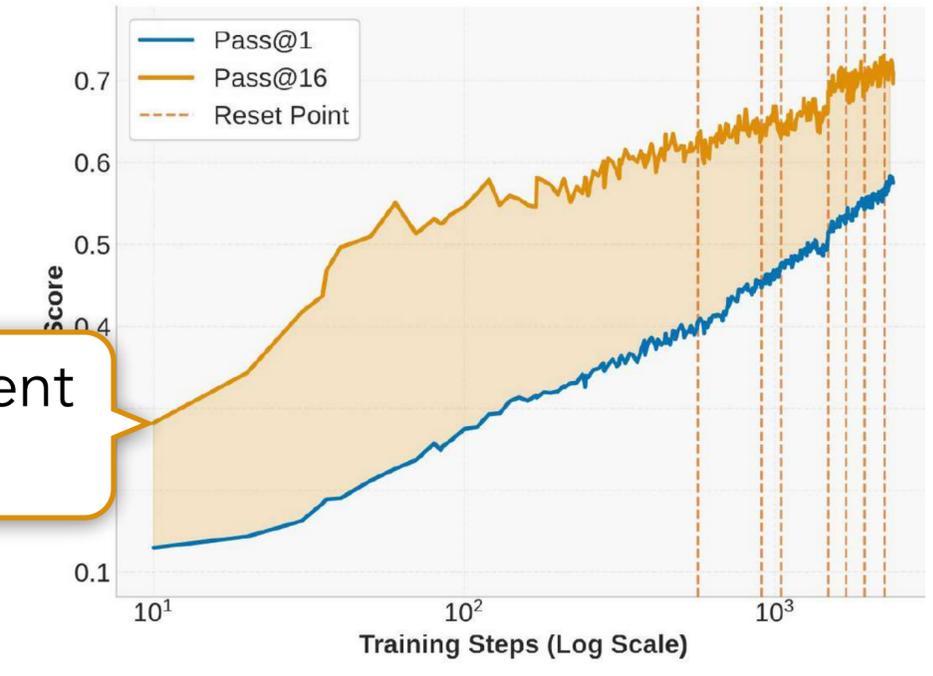
Dynamically adjust  $\epsilon_{\text{high}}$  to maintain **sustainable entropy**

Evaluated across diverse domains—math, code, STEM, reasoning gym, and instruction following—**ProRL** results in significant gains over the base model, **Distill-R1-1.5B**, in all cases.

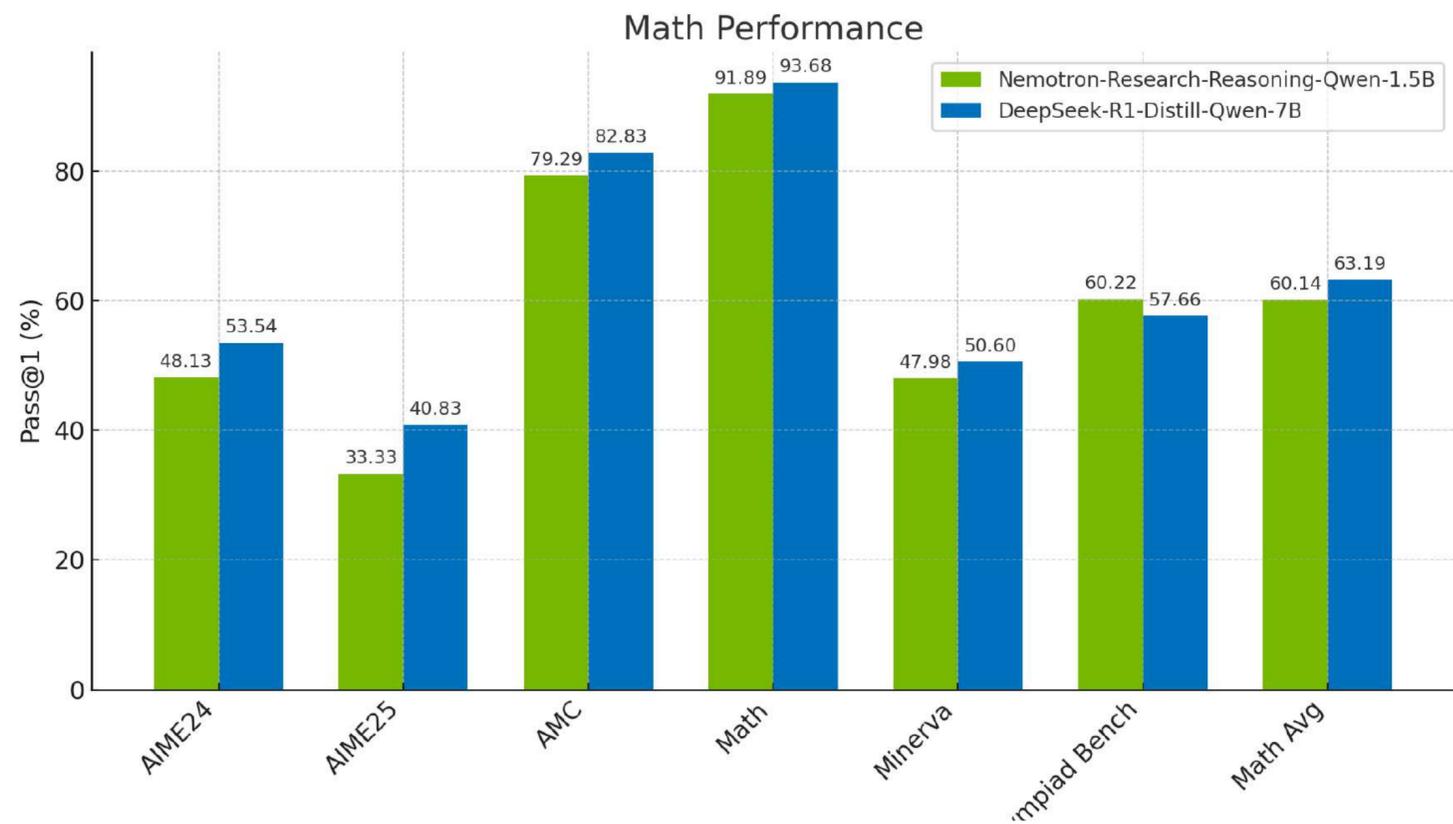
The gain is particularly large in out-of-domain tasks and reasoning gym where the base model struggles.



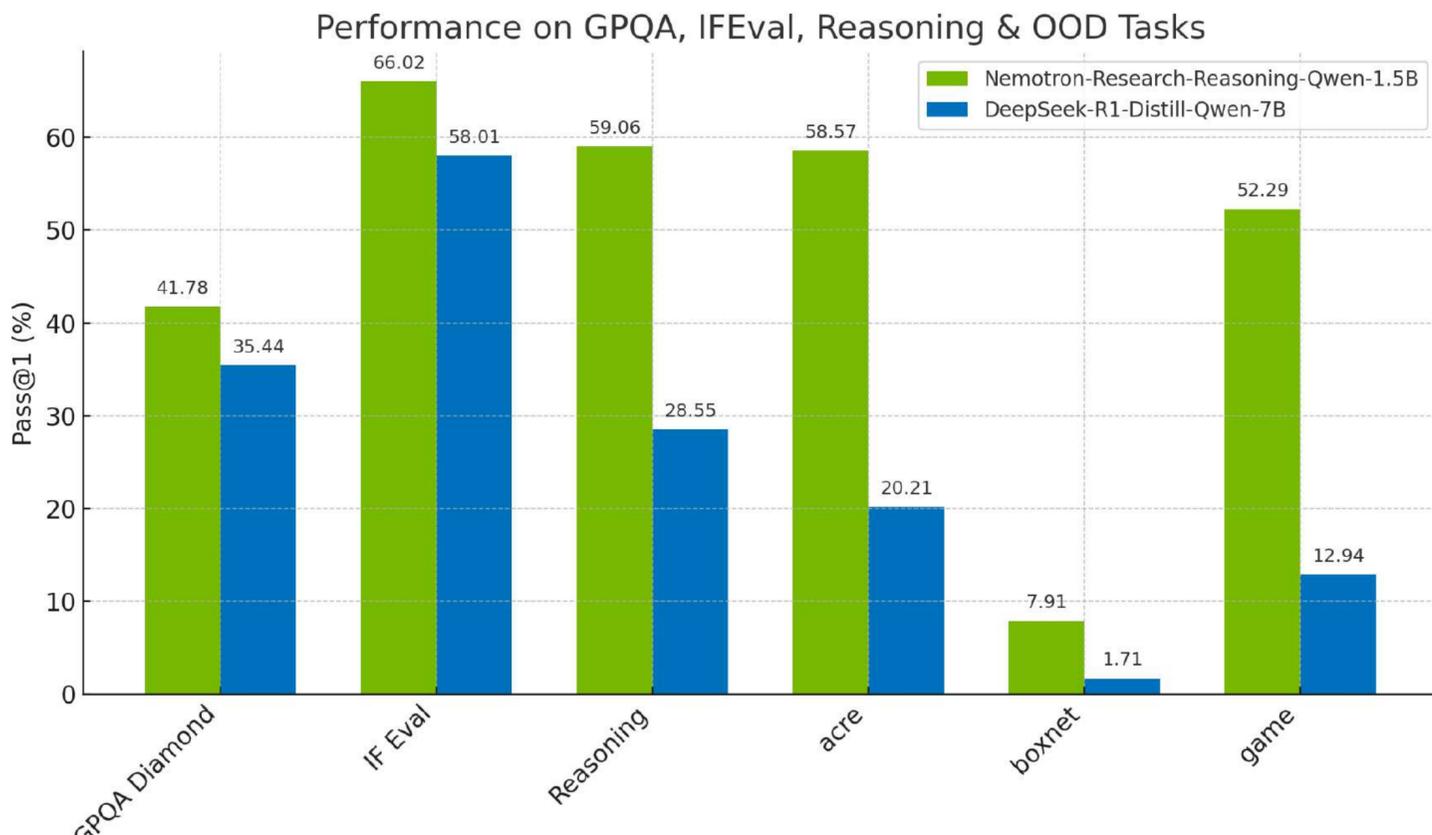
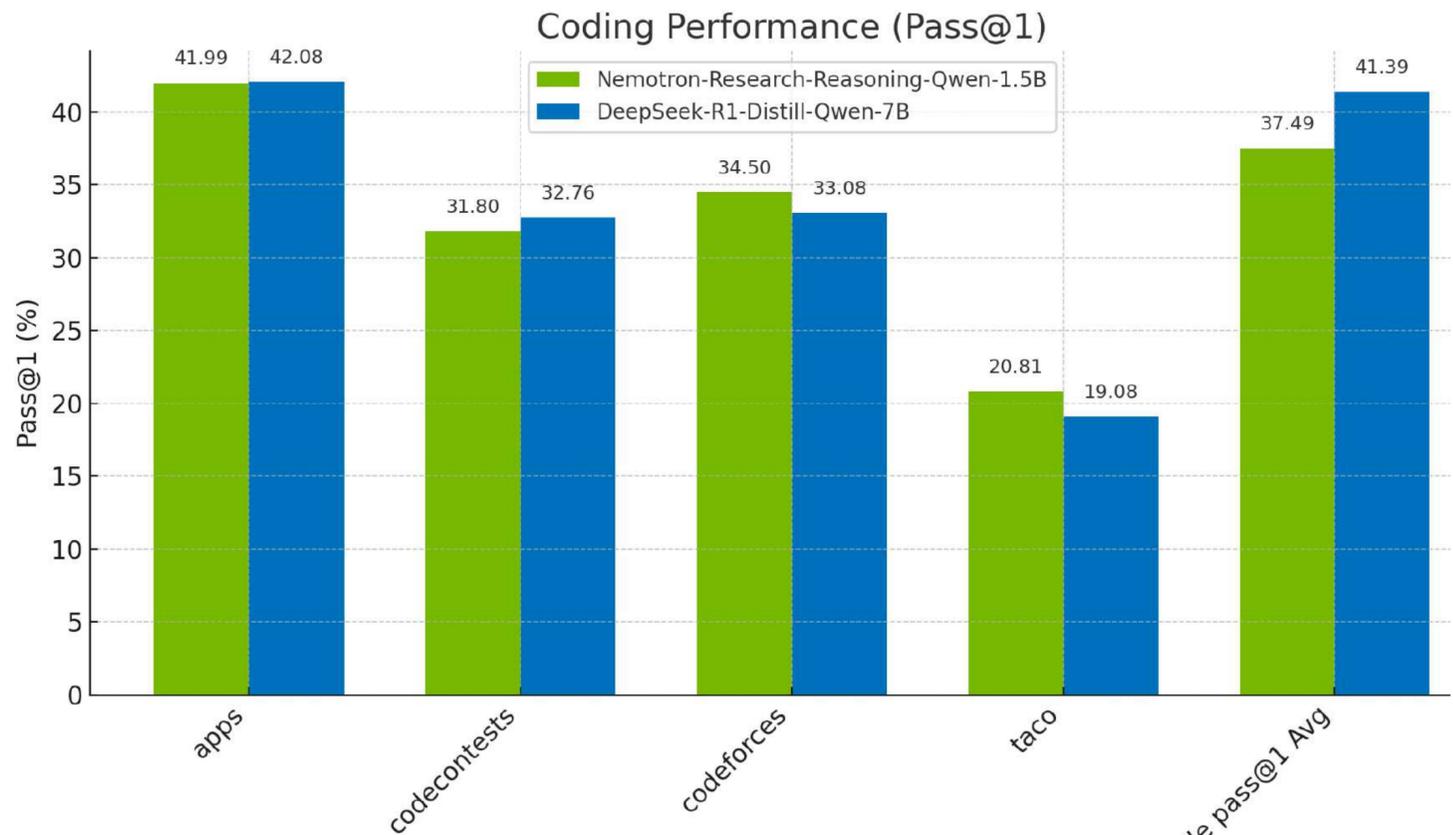
**Log-linear** performance improvement in both Pass@1 and Pass@16.



# The SOTA 1.5B Reasoning Model



Our ProRL-ed model, **Nemotron-Reasoning -1.5B**, achieves comparable or better performance than the 4.5× larger model, **DeepSeek-R1-7B**.





Stella Li → CogSci2025

@StellaLisy



🤖 We cracked RLVR with... Random Rewards?!

Training Qwen2.5-Math-7B with our Spurious Rewards improved MATH-500 by:

- Random rewards: +21%
- Incorrect rewards: +25%
- (FYI) Ground-truth rewards: + 28.8%

How could this even work!?! Here's why: 📖

Blogpost: [tinyurl.com/spurious-rewar...](https://tinyurl.com/spurious-rewar...)

# Striking findings against RL



Chemistry between the base LLM and RL matters  
 Conclusions from **effortless** RL  $\neq$  **effortful** RL

{trulins, stell1, rx31, sgeng}@cs.washington.edu

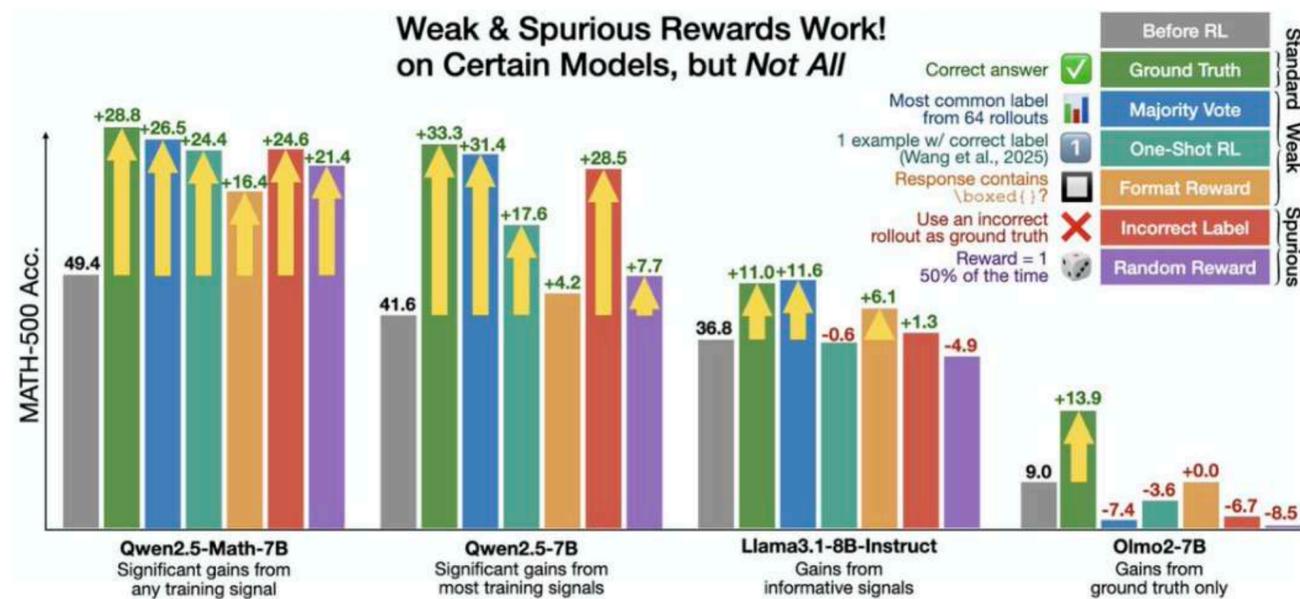


Figure 1: MATH-500 accuracy after 150 steps of RLVR on various training signals. We show that

<sup>1</sup> LeapLab, Tsinghua University    <sup>2</sup> Shanghai Jiao Tong University

\* Equal Contribution    † Project Lead    ✉ Corresponding Author

## Echo Chamber: RL Post-training Amplifies Behaviors Learned in Pretraining

Rosie Zhao\*  
Harvard University  
Kempner Institute

Alexandru Meterez\*  
Harvard University  
Kempner Institute

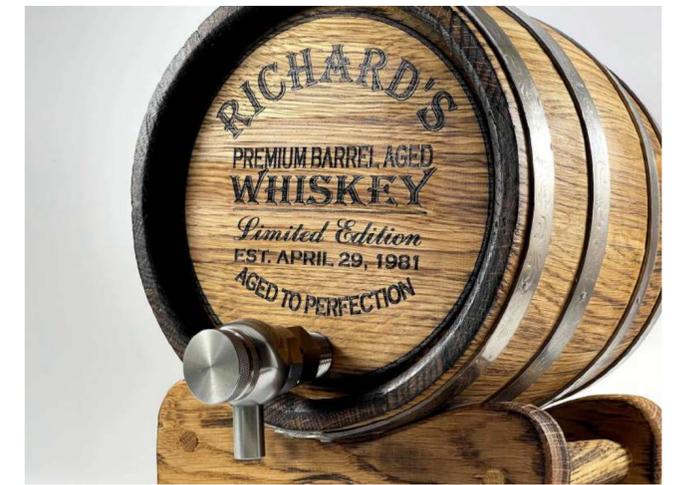
Sham Kakade  
Harvard University  
Kempner Institute

Cengiz Pehlevan  
Harvard University  
Kempner Institute

Samy Jelassi†  
Harvard University

Eran Malach†  
Harvard University  
Kempner Institute

# Conclusion (🌹 vs 💀)



- 🌹 Conclusions from effortless RL  $\neq$  effortful RL
- 🌹 An effortful, proloooooonged RL on 1.5B model can go faaaaaaaarrrr
- 🌹 ProRL shows that a dynamic control of entropy is critical

## BRORL: SCALING REINFORCEMENT LEARNING VIA BROADENED EXPLORATION

Jian Hu<sup>1</sup> Mingjie Liu<sup>1</sup> Ximing Lu<sup>1</sup> Fang Wu<sup>2</sup> Zaid Harchaoui<sup>3</sup> Shizhe Diao<sup>1</sup>  
Yejin Choi<sup>1</sup> Pavlo Molchanov<sup>1</sup> Jun Yang<sup>1</sup> Jan Kautz<sup>1</sup> Yi Dong<sup>1</sup>  
<sup>1</sup>NVIDIA <sup>2</sup>Stanford University <sup>3</sup>University of Washington

💀 Can effortful RL succeed on say, GPT2?

# The **Entropy** Mechanism of Reinforcement Learning for Reasoning Language Models

Ganqu Cui<sup>1\*</sup>, Yuchen Zhang<sup>1,4\*</sup>, Jiacheng Chen<sup>1\*</sup>, Lifan Yuan<sup>3</sup>, Zhi Wang<sup>5</sup>, Yuxin Zuo<sup>2</sup>, Haozhan  
Yuchen Fan<sup>1</sup>, Huayu Chen<sup>2</sup>, Weize Chen<sup>2</sup>, Zhiyuan Liu<sup>2</sup>, Hao Peng<sup>3</sup>, Lei Bai<sup>1</sup>, Wanli Ouyang,  
Yu Cheng<sup>1,6†</sup>, Bowen Zhou<sup>1,2†</sup>, Ning Ding<sup>2,1†</sup>

<sup>1</sup> Shanghai AI Laboratory <sup>2</sup> Tsinghua University <sup>3</sup> UIUC <sup>4</sup> Peking University <sup>5</sup> Nanjing University <sup>6</sup> CUHK

# The Unreasonable Effectiveness of **Entropy** Minimization in LLM Reasoning

Shivam Agarwal, Zimin Zhang, Lifan Yuan, Jiawei Han, Hao Peng  
University of Illinois Urbana-Champaign

# Beyond the 80/20 Rule: High-**Entropy** Minority Tokens Drive Effective Reinforcement Learning for LLM Reasoning

Shenzhi Wang<sup>1,2</sup>, Le Yu<sup>1</sup>, Chang Gao<sup>1</sup>, Chujie Zheng<sup>1</sup>, Shixuan Liu<sup>1</sup>, Rui Lu<sup>2</sup>,  
Dang<sup>1</sup>, Xiong-Hui Chen<sup>1</sup>, Jianxin Yang<sup>1</sup>, Zhenru Zhang<sup>1</sup>, Yuqiong Liu<sup>1</sup>, An Yang<sup>1</sup>,  
Andrew Zhao<sup>2</sup>, Yang Yue<sup>2</sup>, Shiji Song<sup>2</sup>, Bowen Yu<sup>1,⊗,†</sup>, Gao Huang<sup>2,⊗</sup>, Junyang Lin<sup>1</sup>  
<sup>1</sup> Qwen Team, Alibaba Inc. <sup>2</sup> LeapLab, Tsinghua University

EPO 🤗💧

EPO 🤗💧

## **Evolutionary** Policy Optimization

Jianren Wang<sup>\*1</sup> Yifan Su<sup>\*1</sup> Abhinav Gupta<sup>1</sup> Deepak Pathak<sup>1</sup>

## EPO: **ENTROPY**-REGULARIZED POLICY OPTIMIZATION FOR LLM AGENTS REINFORCEMENT LEARNING

Wujiang Xu<sup>1\*</sup>, Wentian Zhao<sup>2</sup>, Zhenting Wang<sup>1</sup>, Yu-Jhe Li<sup>2</sup>,  
Can Jin<sup>1</sup>, Mingyu Jin<sup>1</sup>, Kai Mei<sup>1</sup>, Kun Wan<sup>2</sup>, Dimitris N. Metaxas<sup>1</sup>  
<sup>1</sup> Rutgers University <sup>2</sup> Adobe Inc.



Stella Li → CogSci2025

@StellaLisy



🤖 We cracked RLVR with... Random Rewards?!

Training Qwen2.5-Math-7B with our Spurious Rewards improved MATH-500 by:

- Random rewards: +21%
- Incorrect rewards: +25%
- (FYI) Ground-truth rewards: + 28.8%

How could this even work!?! Here's why: 📖

Blogpost: [tinyurl.com/spurious-rewar...](https://tinyurl.com/spurious-rewar...)

# Striking findings against RL



Chemistry between the base LLM and RL matters  
 Conclusions from **effortless** RL  $\neq$  **effortful** RL

{trulins, stell1, rx31, sgeng}@cs.washington.edu

<sup>1</sup> LeapLab, Tsinghua University <sup>2</sup> Shanghai Jiao Tong University

Conclusions from **effortless** SFT  $\neq$  **effortful** SFT

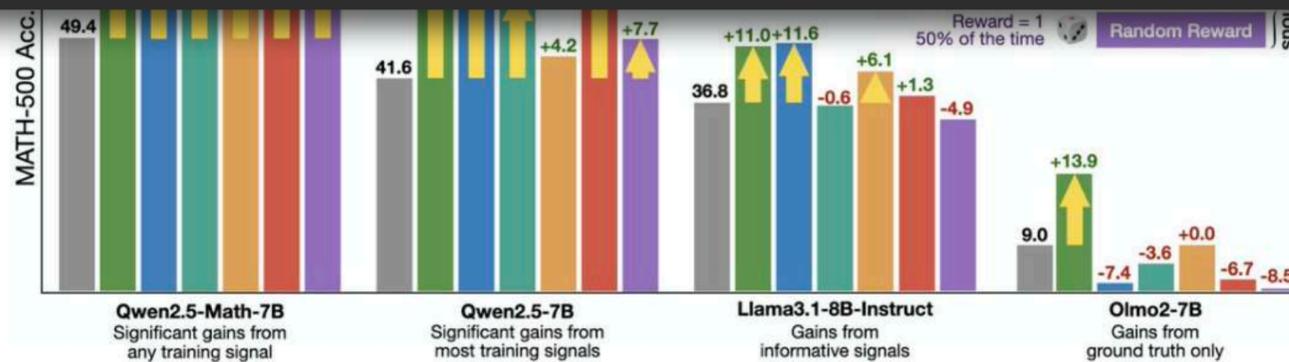


Figure 1: MATH-500 accuracy after 150 steps of RLVR on various training signals. We show that

Rosie Zhao\*  
Harvard University  
Kempner Institute

Alexandru Meterez\*  
Harvard University  
Kempner Institute

Sham Kakade  
Harvard University  
Kempner Institute

Cengiz Pehlevan  
Harvard University  
Kempner Institute

Samy Jelassi<sup>†</sup>  
Harvard University

Eran Malach<sup>†</sup>  
Harvard University  
Kempner Institute

🤖 We cracked RLVR with... Random Rewards?!  
Training Qwen2.5-Math-7B with our Spurious Rewards improved MATH-500 by:  
Random rewards: +21%

# Striking findings against RL

Conclusions from **effortless** SFT  $\neq$  **effortful** SFT

## Spurious Rewards: Rethinking Training Signals in RLVR

Rulin Shao<sup>1\*</sup> Shuyue Stella Li<sup>1\*</sup> Rui Xin<sup>1\*</sup> Scott Geng<sup>1\*</sup> Yiping Wang<sup>1</sup>  
Sewoong Oh<sup>1</sup> Simon Shaolei Du<sup>1</sup> Nathan Lambert<sup>2</sup> Sewon Min<sup>3</sup> Ranjay Krishna<sup>1,2</sup>  
Yulia Tsvetkov<sup>1</sup> Hannaneh Hajishirzi<sup>1,2</sup> Pang Wei Koh<sup>1,2</sup> Luke Zettlemoyer<sup>1</sup>  
<sup>1</sup>University of Washington <sup>2</sup>Allen Institute for Artificial Intelligence  
<sup>3</sup>University of California, Berkeley  
{rulins, stelli, rx31, sgeng}@cs.washington.edu

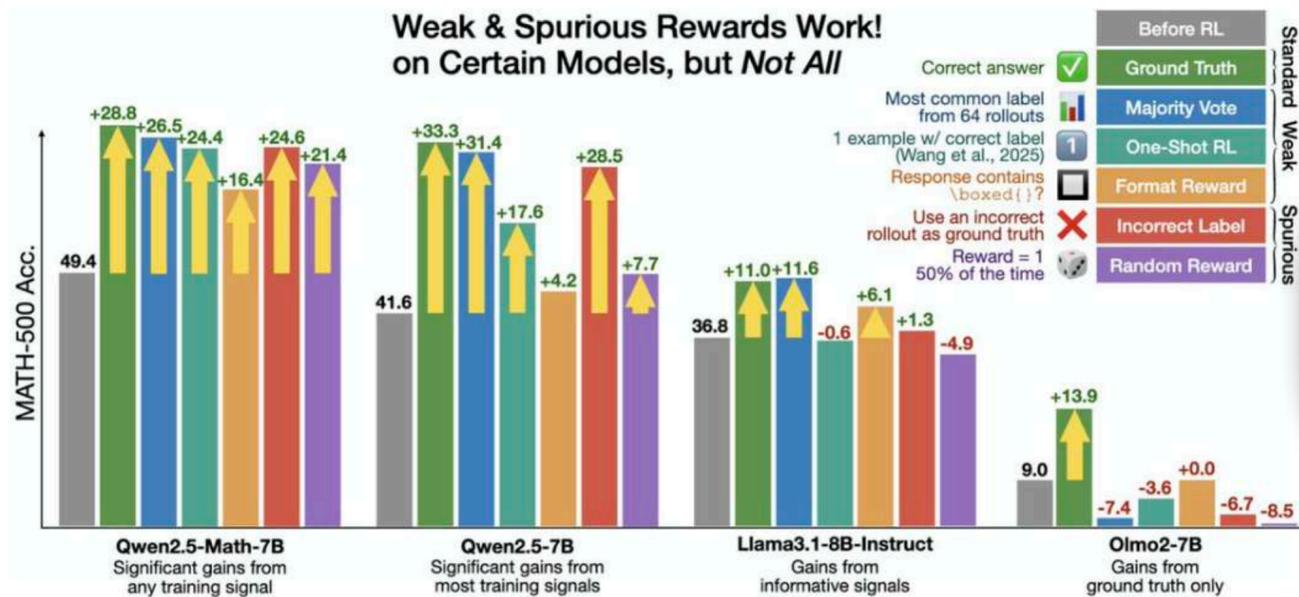


Figure 1: MATH-500 accuracy after 150 steps of RLVR on various training signals. We show that



May 19, 2025

## Does Reinforcement Learning Really Incentivize Reasoning Capacity in LLMs Beyond the Base Model?

Yang Yue<sup>1\*†</sup>, Zhiqi Chen<sup>1\*</sup>, Rui Lu<sup>1</sup>, Andrew Zhao<sup>1</sup>, Zhaokai Wang<sup>2</sup>, Yang Yue<sup>1</sup>, Shiji Song<sup>1</sup>, and Gao Huang<sup>1✉</sup>

<sup>1</sup>LeapLab, Tsinghua University <sup>2</sup>Shanghai Jiao Tong University

\* Equal Contribution † Project Lead ✉ Corresponding Author

Why is Qwen's chemistry so good with RL?  
SFT or SFT-style "mid-training" during pre-training

# 2025: The Rise of LRMs (as opposed to LLMs)

## Does Math Reasoning Improve General LLM Capabilities? Understanding Transferability of LLM Reasoning

Maggie Huan<sup>1,2,\*</sup>, Yuetai Li<sup>3,\*</sup>, Tuney Zheng<sup>4,\*</sup>, Xiaoyu Xu<sup>5</sup>, Seungone Kim<sup>1</sup>, Minxin Du<sup>5</sup>,  
Radha Poovendran<sup>3</sup>, Graham Neubig<sup>1</sup>, Xiang Yue<sup>1,†</sup>

<sup>1</sup>Carnegie Mellon University <sup>2</sup>University of Pennsylvania <sup>3</sup>University of Washington  
<sup>4</sup>M-A-P <sup>5</sup>The Hong Kong Polytechnic University

ziyuh@seas.upenn.edu yuetaili@uw.edu xyue2@andrew.cmu.edu

**Abstract:** Math reasoning has become the poster child of progress in large language models (LLMs), with new models rapidly surpassing human-level performance on benchmarks like MATH and AIME. But as math leaderboards improve week by week, it is worth asking: *do these gains reflect broader problem-solving ability or just narrow overfitting?* To answer this question, we evaluate over 20 open-weight reasoning-tuned models across a broad suite of tasks, including math, scientific QA, agent planning, coding, and standard instruction-following. We surprisingly find that most models that succeed in math fail to transfer their gains to other domains. To rigorously study this phenomenon, we conduct controlled experiments on Qwen3-14B models using math-only data but different tuning methods. We find that reinforcement learning (RL)-tuned models generalize well across domains, while supervised fine-tuning (SFT)-tuned models often forget general capabilities. Latent-space representation and token-space distribution shift analyses reveal that SFT induces substantial representation and output drift, while RL preserves general-domain structure. Our results suggest a need to rethink standard post-training recipes, particularly the reliance on SFT-distilled data for advancing reasoning models.

[github.com/ReasoningTransfer/Transferability-of-LLM-Reasoning](https://github.com/ReasoningTransfer/Transferability-of-LLM-Reasoning)

[huggingface.co/ReasoningTransferability](https://huggingface.co/ReasoningTransferability)



Then, SFT cannot generalize?  
It depends...



## SFT Memorizes, RL Generalizes: A Comparative Study of Foundation Model Post-training

Tianzhe Chu<sup>✳</sup> Yuexiang Zhai<sup>♥✳</sup> Jihan Yang<sup>♦</sup> Shengbang Tong<sup>♦</sup>  
Saining Xie<sup>✳♦</sup> Dale Schuurmans<sup>✳</sup> Quoc V. Le<sup>✳</sup> Sergey Levine<sup>♥</sup> Yi Ma<sup>✳♥</sup>

Conclusions from **effortless** SFT  $\neq$  **effortful** SFT

# h: Underdogs, Art of Battling

In this talk:

===

ProRL: Prolonged RL

Prismatic Synthesis

RL as Pretraining

===

“Smaller but Better”

“Algorithms for the Win”



*David vs. Goliath:*

*the Art of Scaling Intelligence*

*in the Era of Extreme-Scale Neural Models*

Three components to innovate:

Unconventional data 

Unconventional algorithms 

Unconventional collaboration 

early  
r)



# Prismatic Synthesis:

## Gradient-based Data Diversification

### Boosts Generalization in LLM Reasoning

— *NeurIPS 2025* —

Jaehun Jung



Seungju Han



Ximing Lu



Skyler Hallinan



David Acuna



Shrimai Prabhunoye



Mostafa Patwary



Mohammad Shoeybi



Bryan Catanzaro



Yejin Choi



# The Era of Brute-Force Scaling is Over

## The Era of Smart Scaling Begins



### Pre-training as we know it will end

Compute is growing:

- Better hardware
- Better algorithms
- Larger clusters

Data is not growing:

- We have but one internet
- **The fossil fuel of AI**

-- Ilya Sutskever from his test time award talk at NeurIPS 2024

# Synthetic Data to the Rescue!

Already used in both mid / post-training:



AI



Microsoft  
**Phi-4**

Sometimes, for the entirety of LLM training...

Common practice & assumption

(1) rely on the largest / strongest teacher models

# Obvious Concerns: Mode Collapse

nature

Explore content ▾ About the journal ▾ Publish with us ▾

[nature](#) > [articles](#) > article

PRICE \$8.99

THE

NOV. 20, 2023

NEW YORKER

Conclusions from *effortless* synthetic data  
!= *effortful* synthetic data

## Abstract

Stable diffusion re  
(ref. [2](#)) and GPT-4 (

## THE CURSE OF RECURSION: TRAINING ON GENERATED DATA MAKES MODELS FORGET

**Ilya Shumailov\***  
University of Oxford

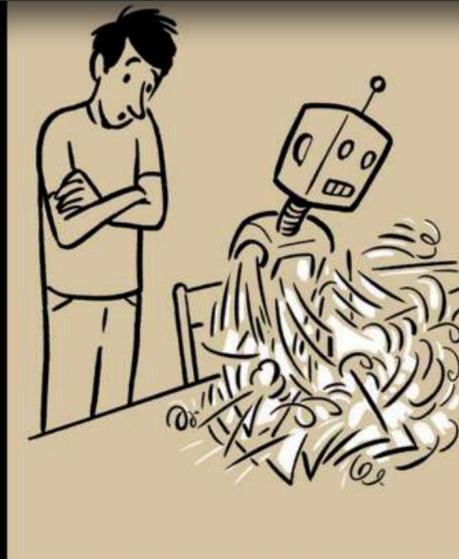
**Zakhar Shumaylov\***  
University of Cambridge

**Yiren Zhao**  
Imperial College London

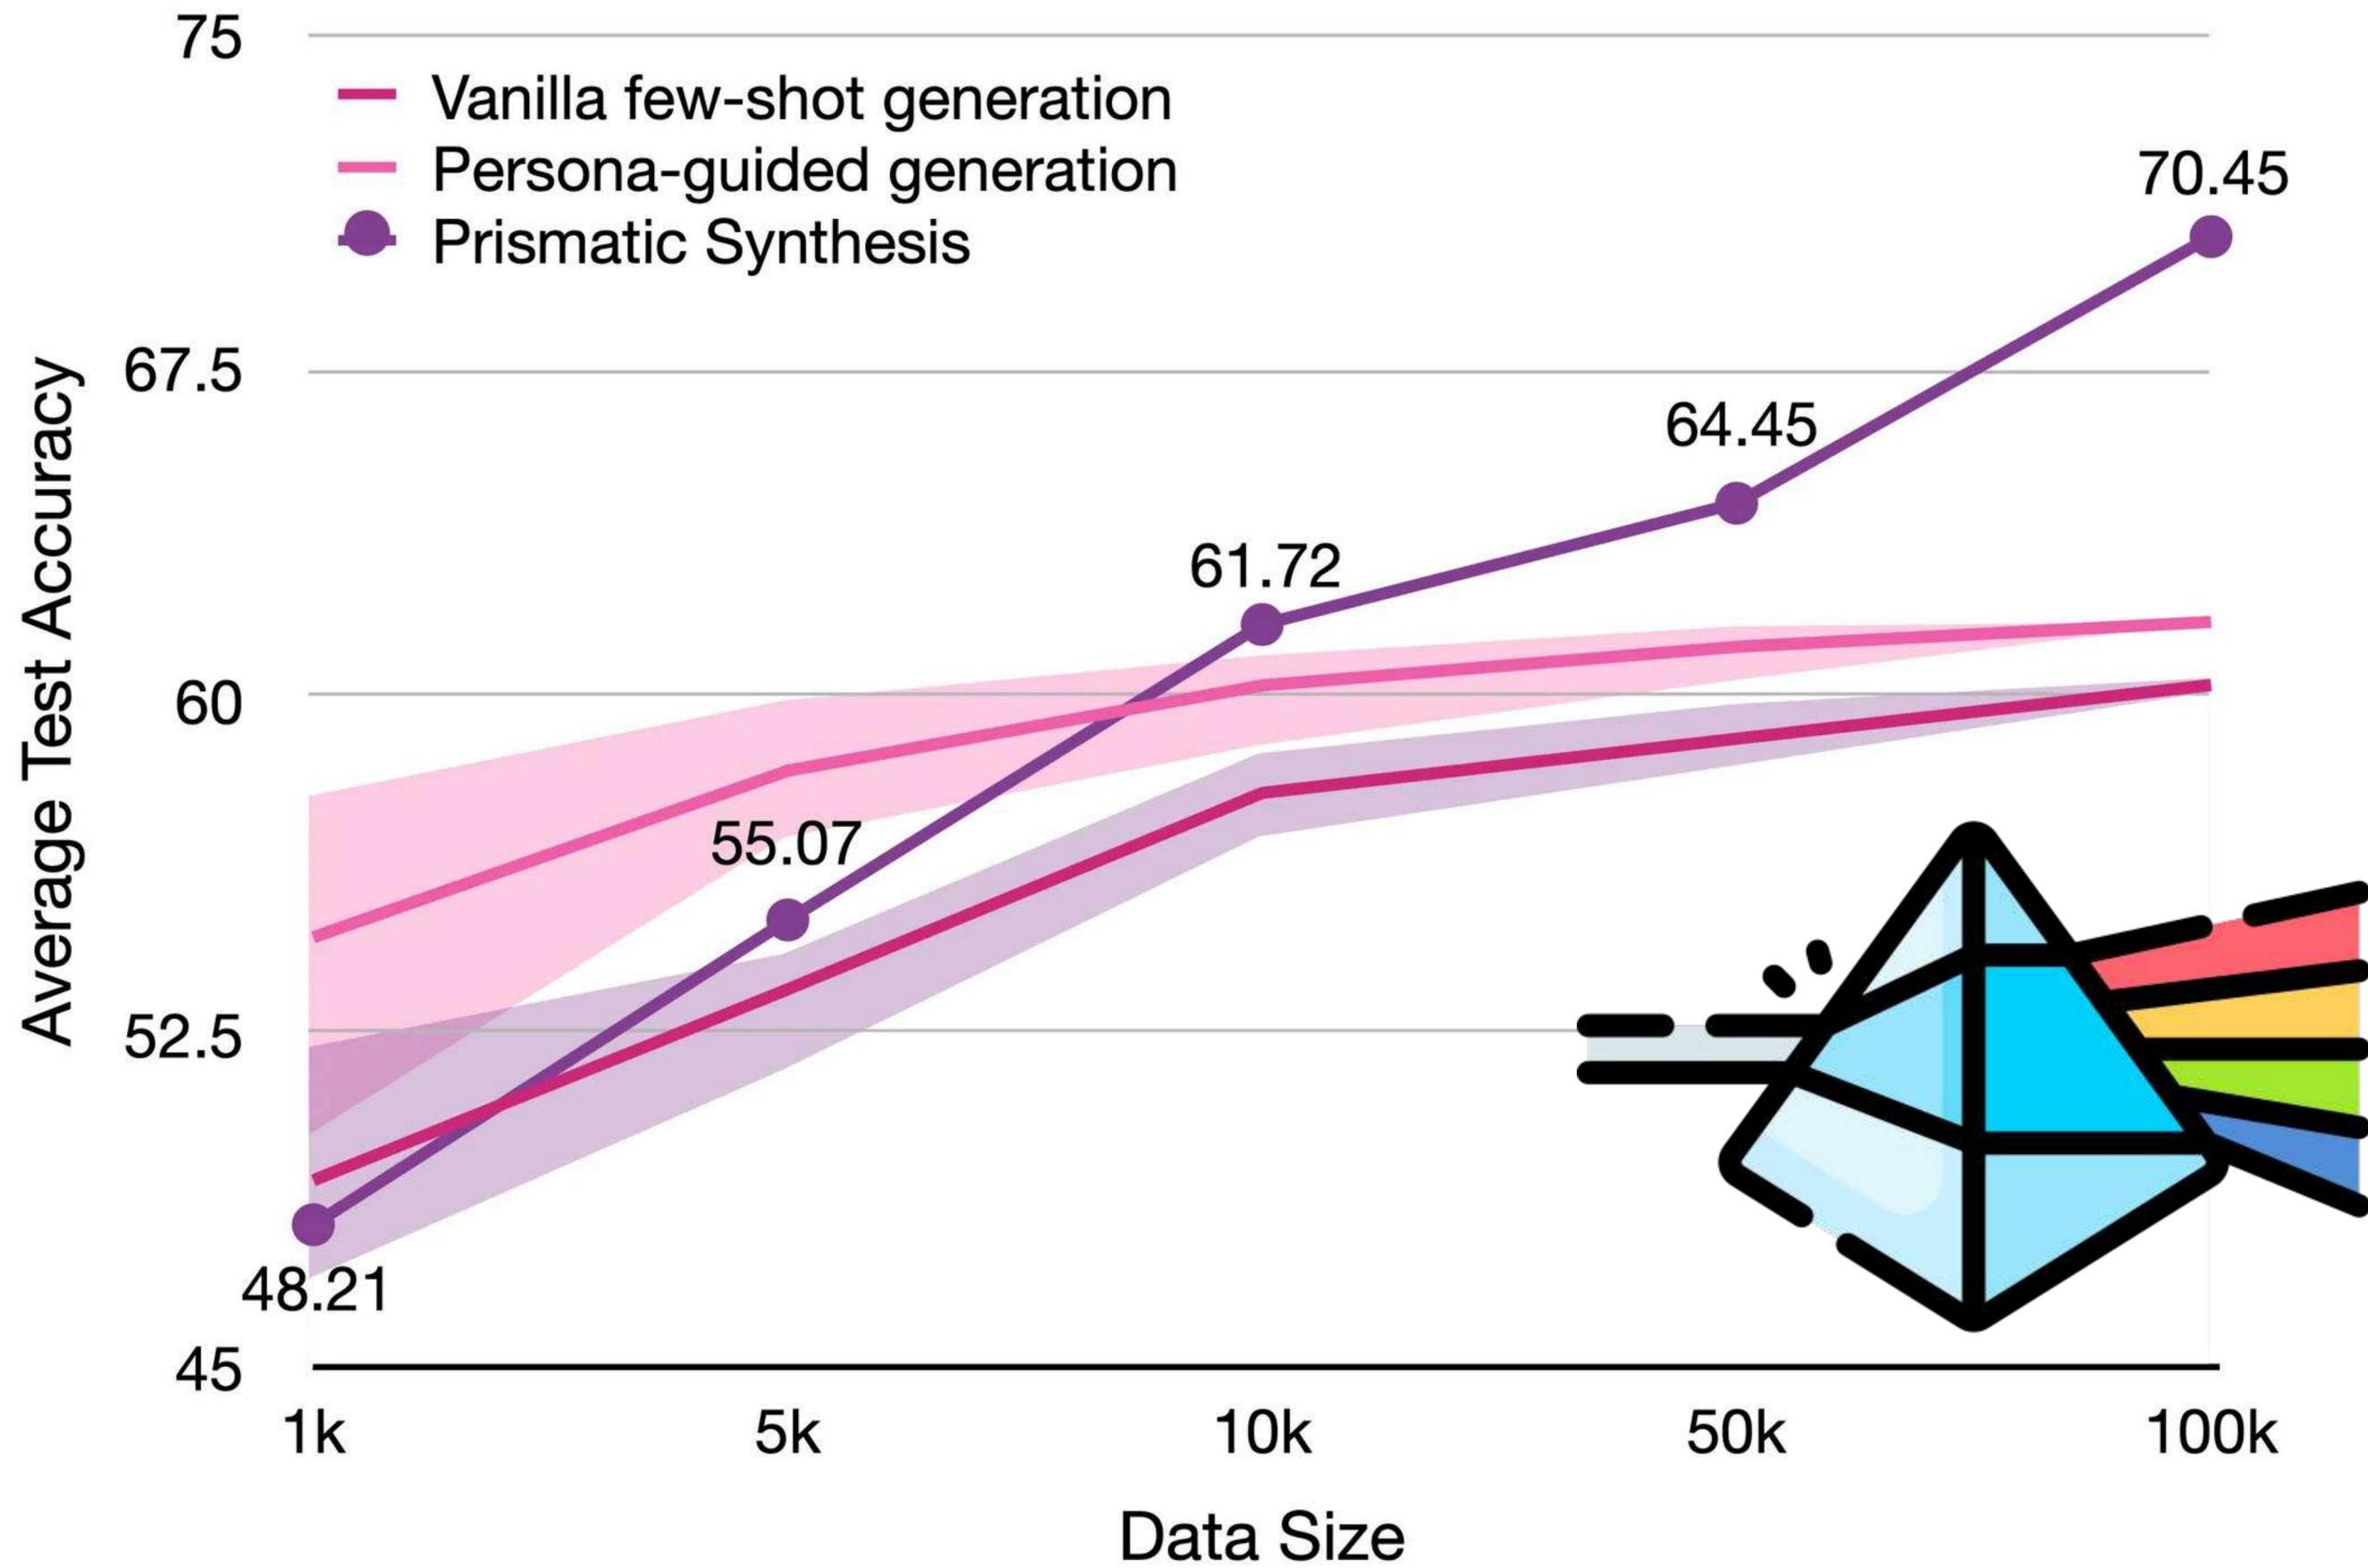
**Yarin Gal**  
University of Oxford

**Nicolas Papernot**  
University of Toronto & Vector Institute

**Ross Anderson**  
University of Cambridge & University of Edinburgh



CN23



# Gradient as Data Representation

## Input $\mathbf{x}$

Bill walks 0.5 mile south, then 0.75 mile east, and finally 0.5 mile south. How many miles is he, in a direct line, from his starting point?

## Output $\mathbf{y}$

Bill walked  $0.5+0.5 = 1.0$  miles south total and walked 0.75 miles east in total. Therefore, using Pythagorean... the answer is 1.25 miles.

Use a small reference model to  
compute  $\nabla_{\theta} \mathbf{P}(\mathbf{y} | \mathbf{x})$

Rademacher Projection for  
dimensionality reduction



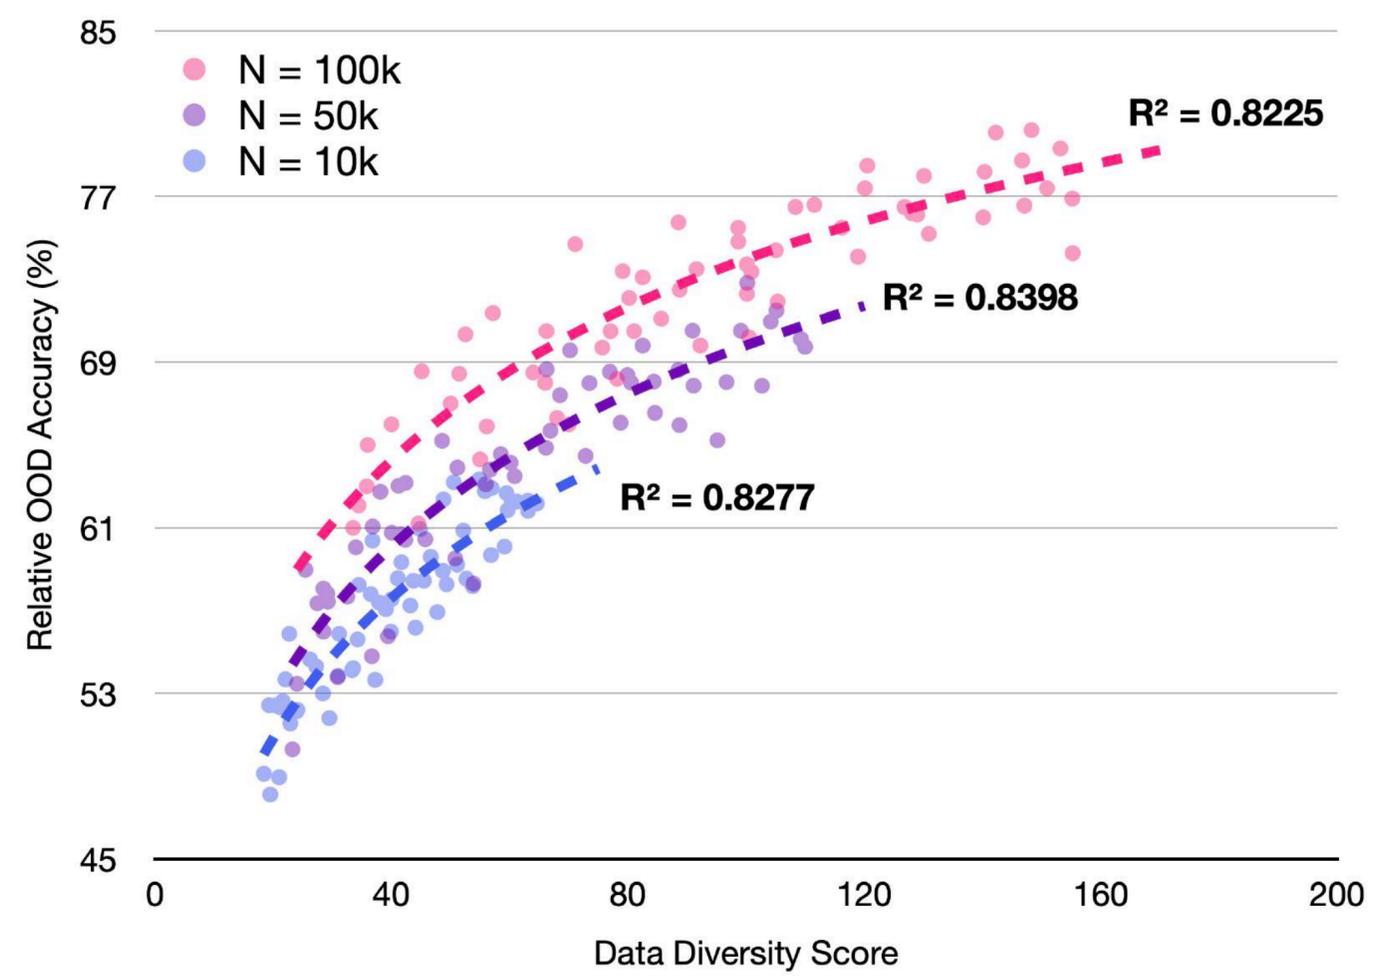
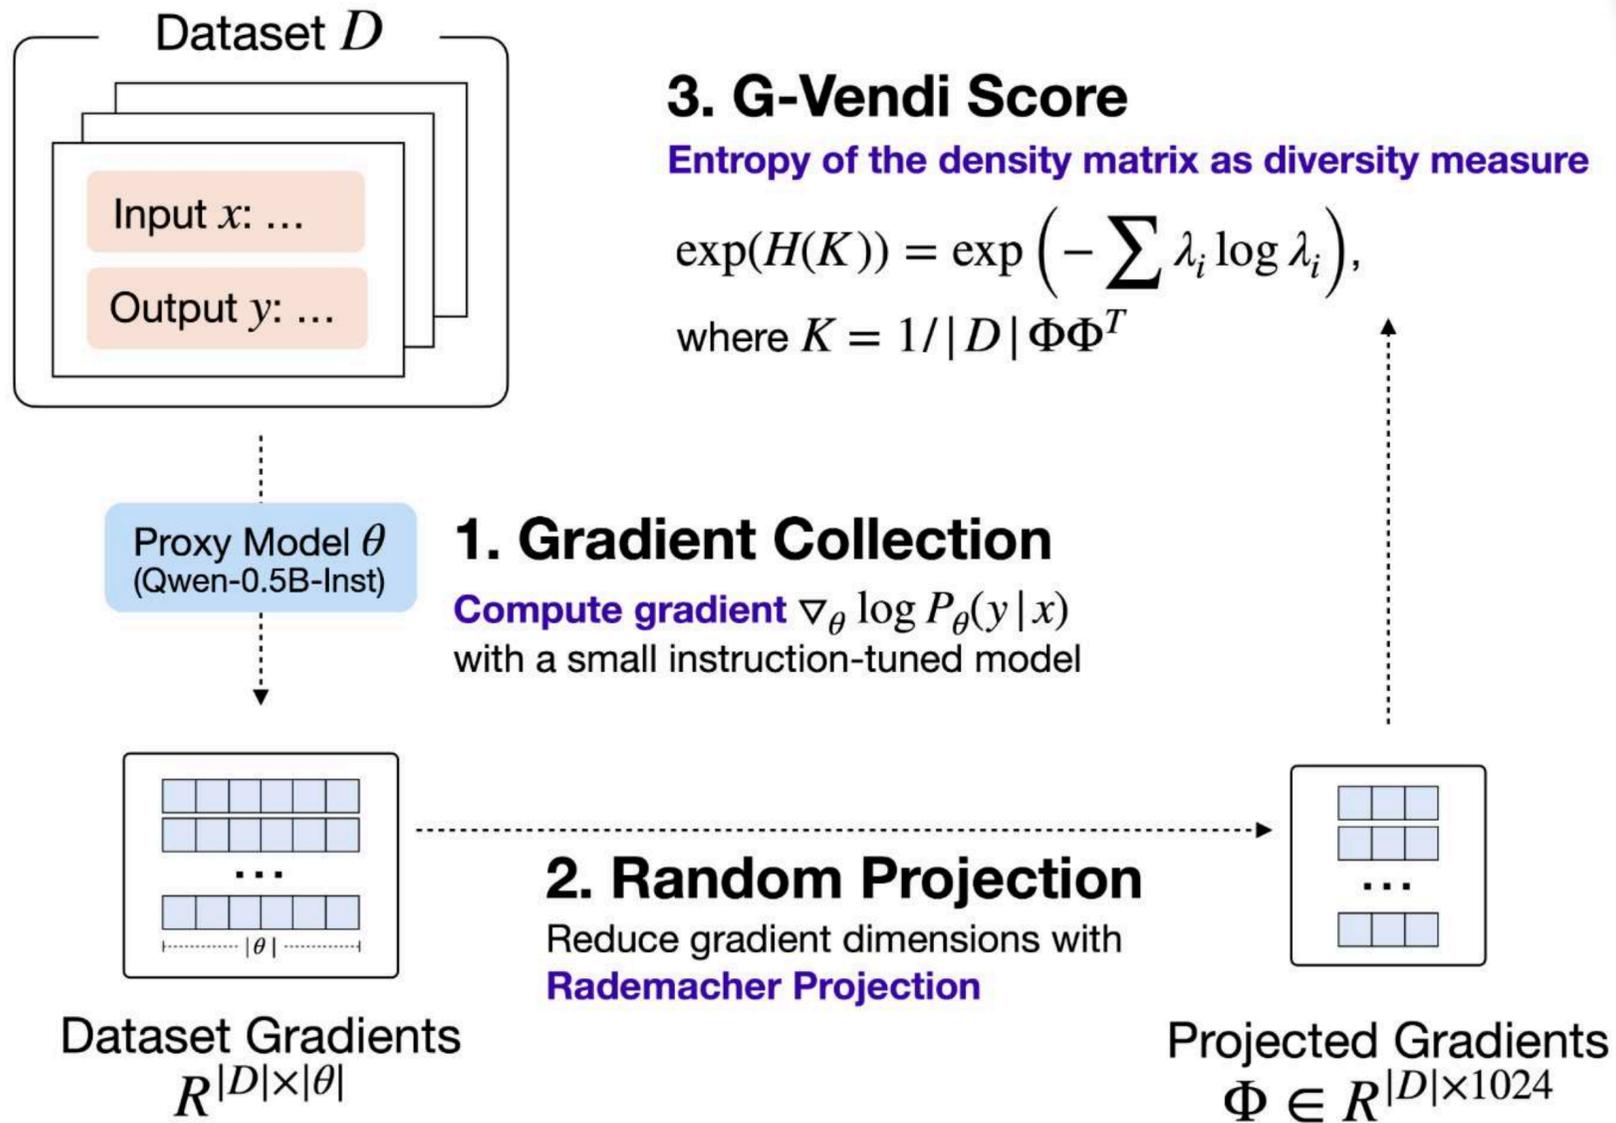
Sample A

1024-dim representation of data point

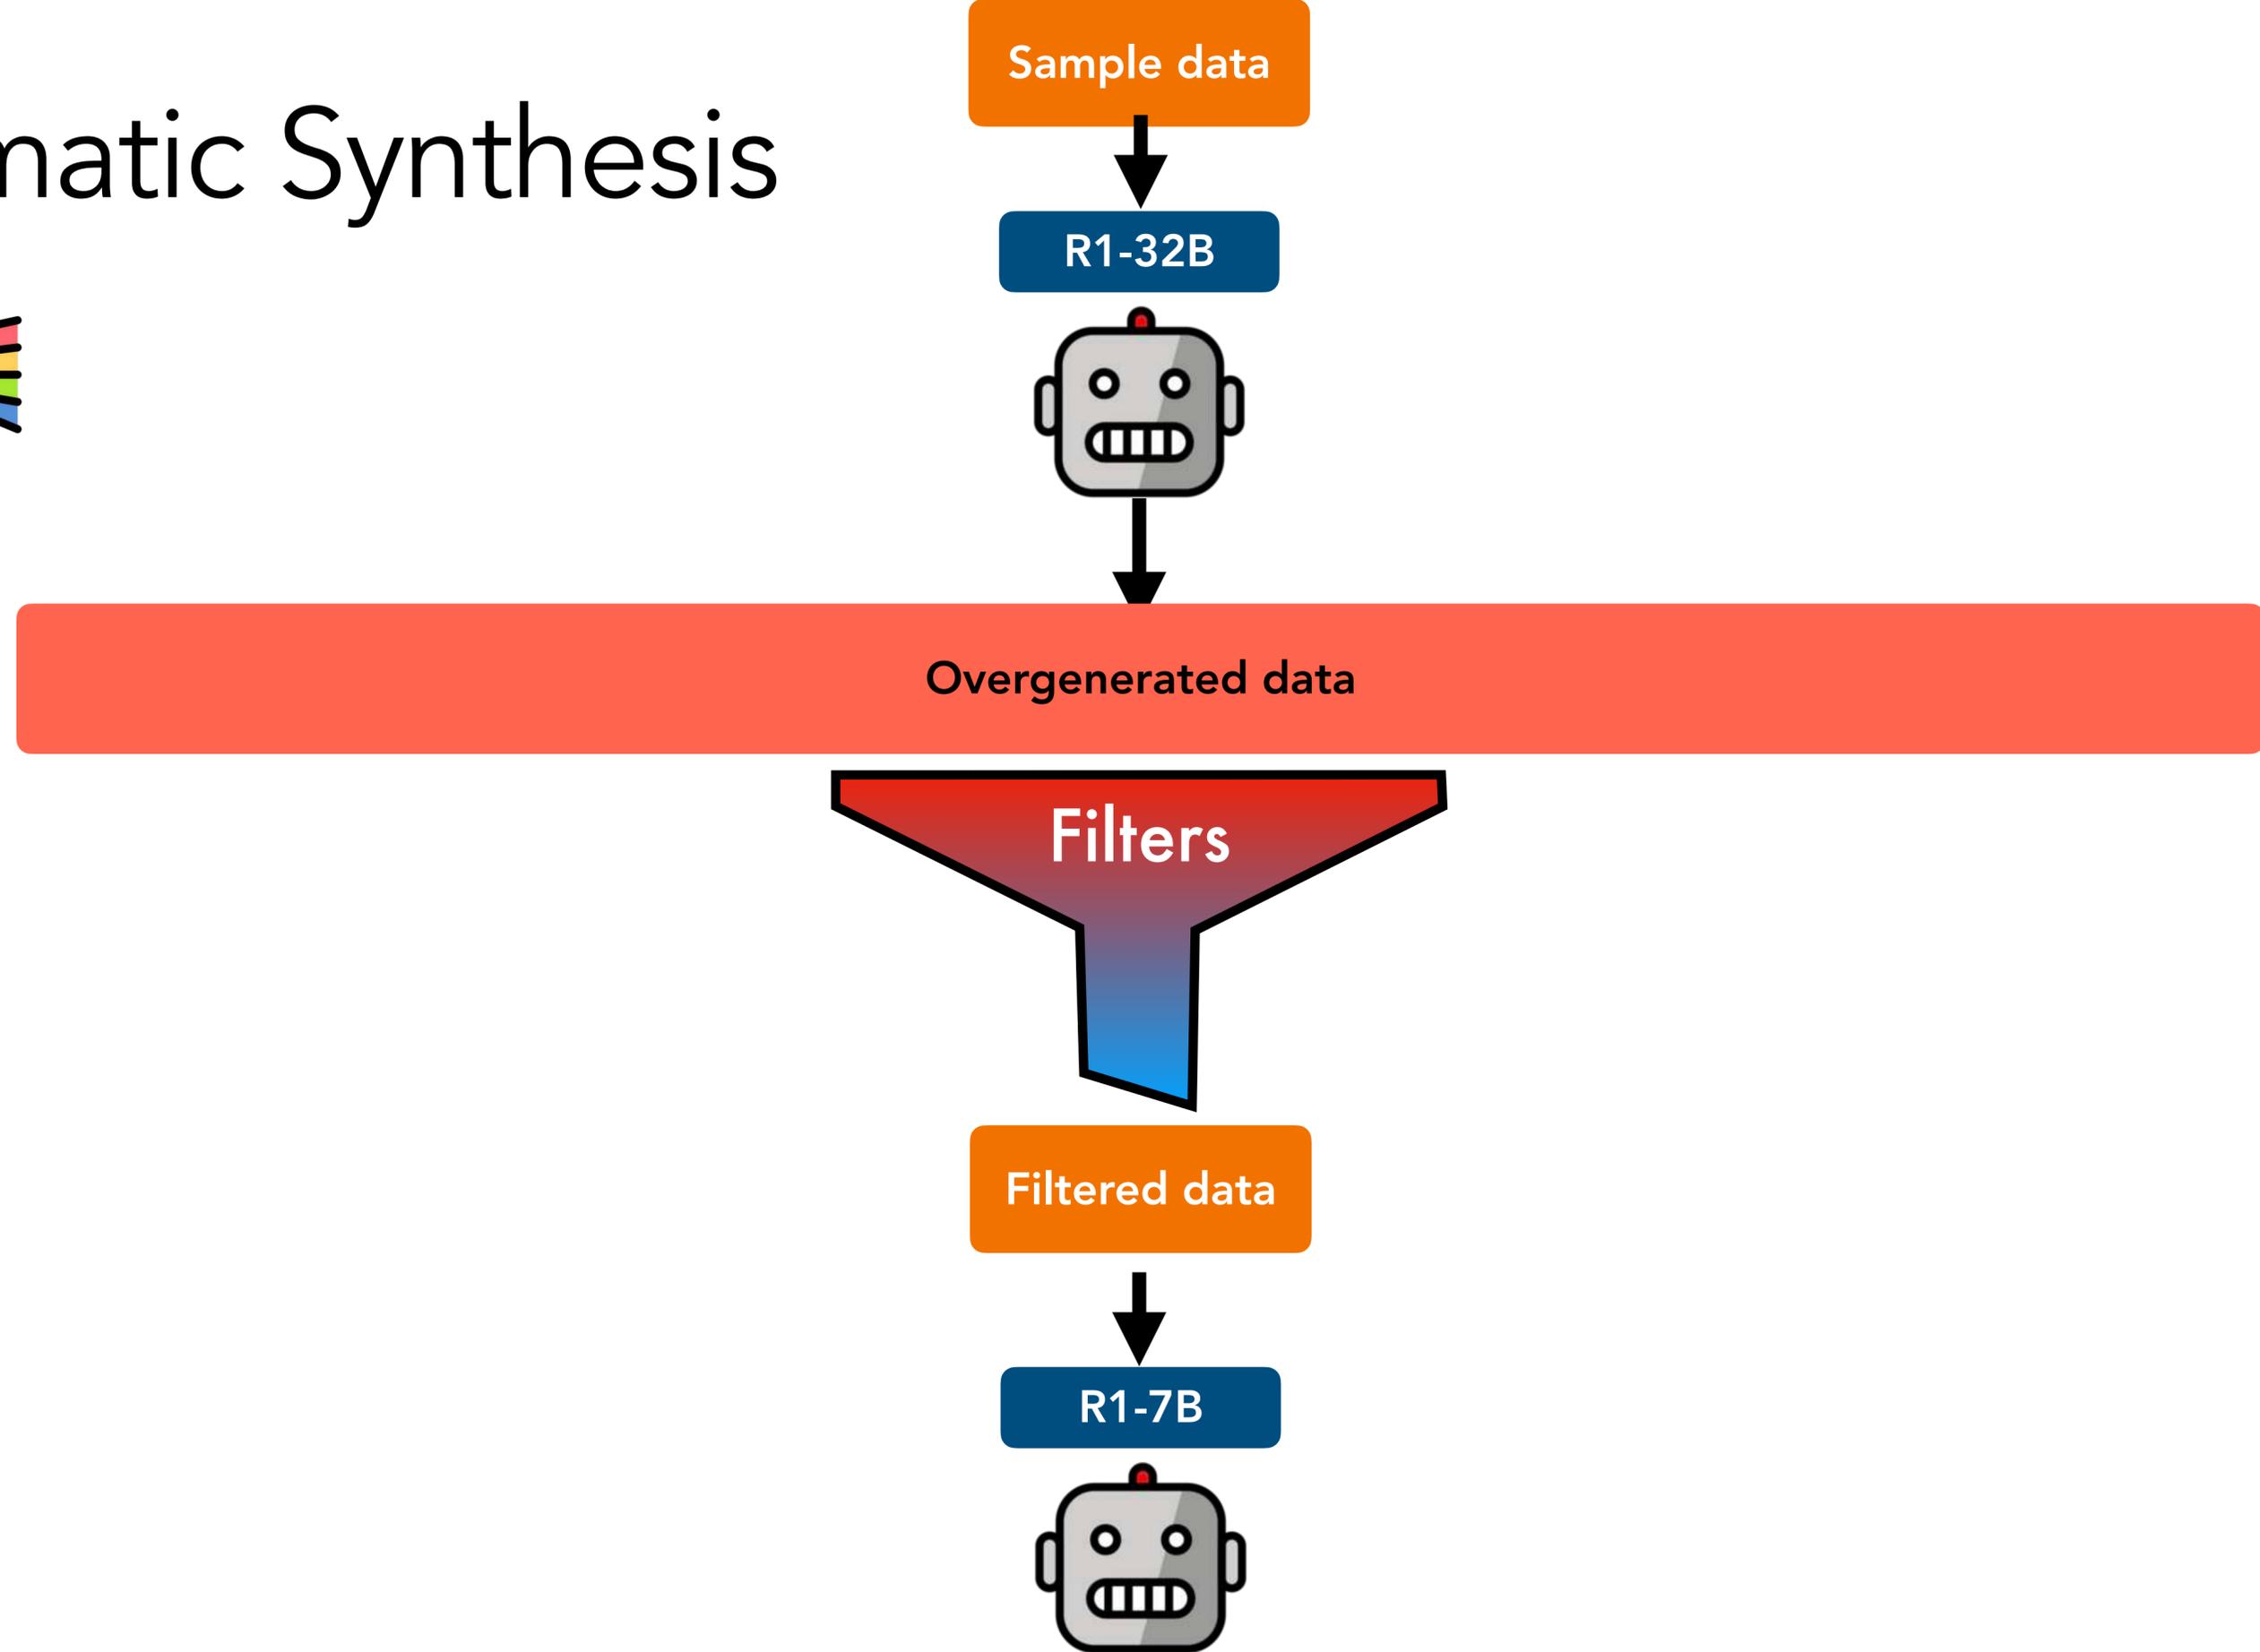
Gradient encodes the mapping from  
 $\mathbf{x}$  to  $\mathbf{y}$ ;  
which naturally represents  
“reasoning”  
between input and output!

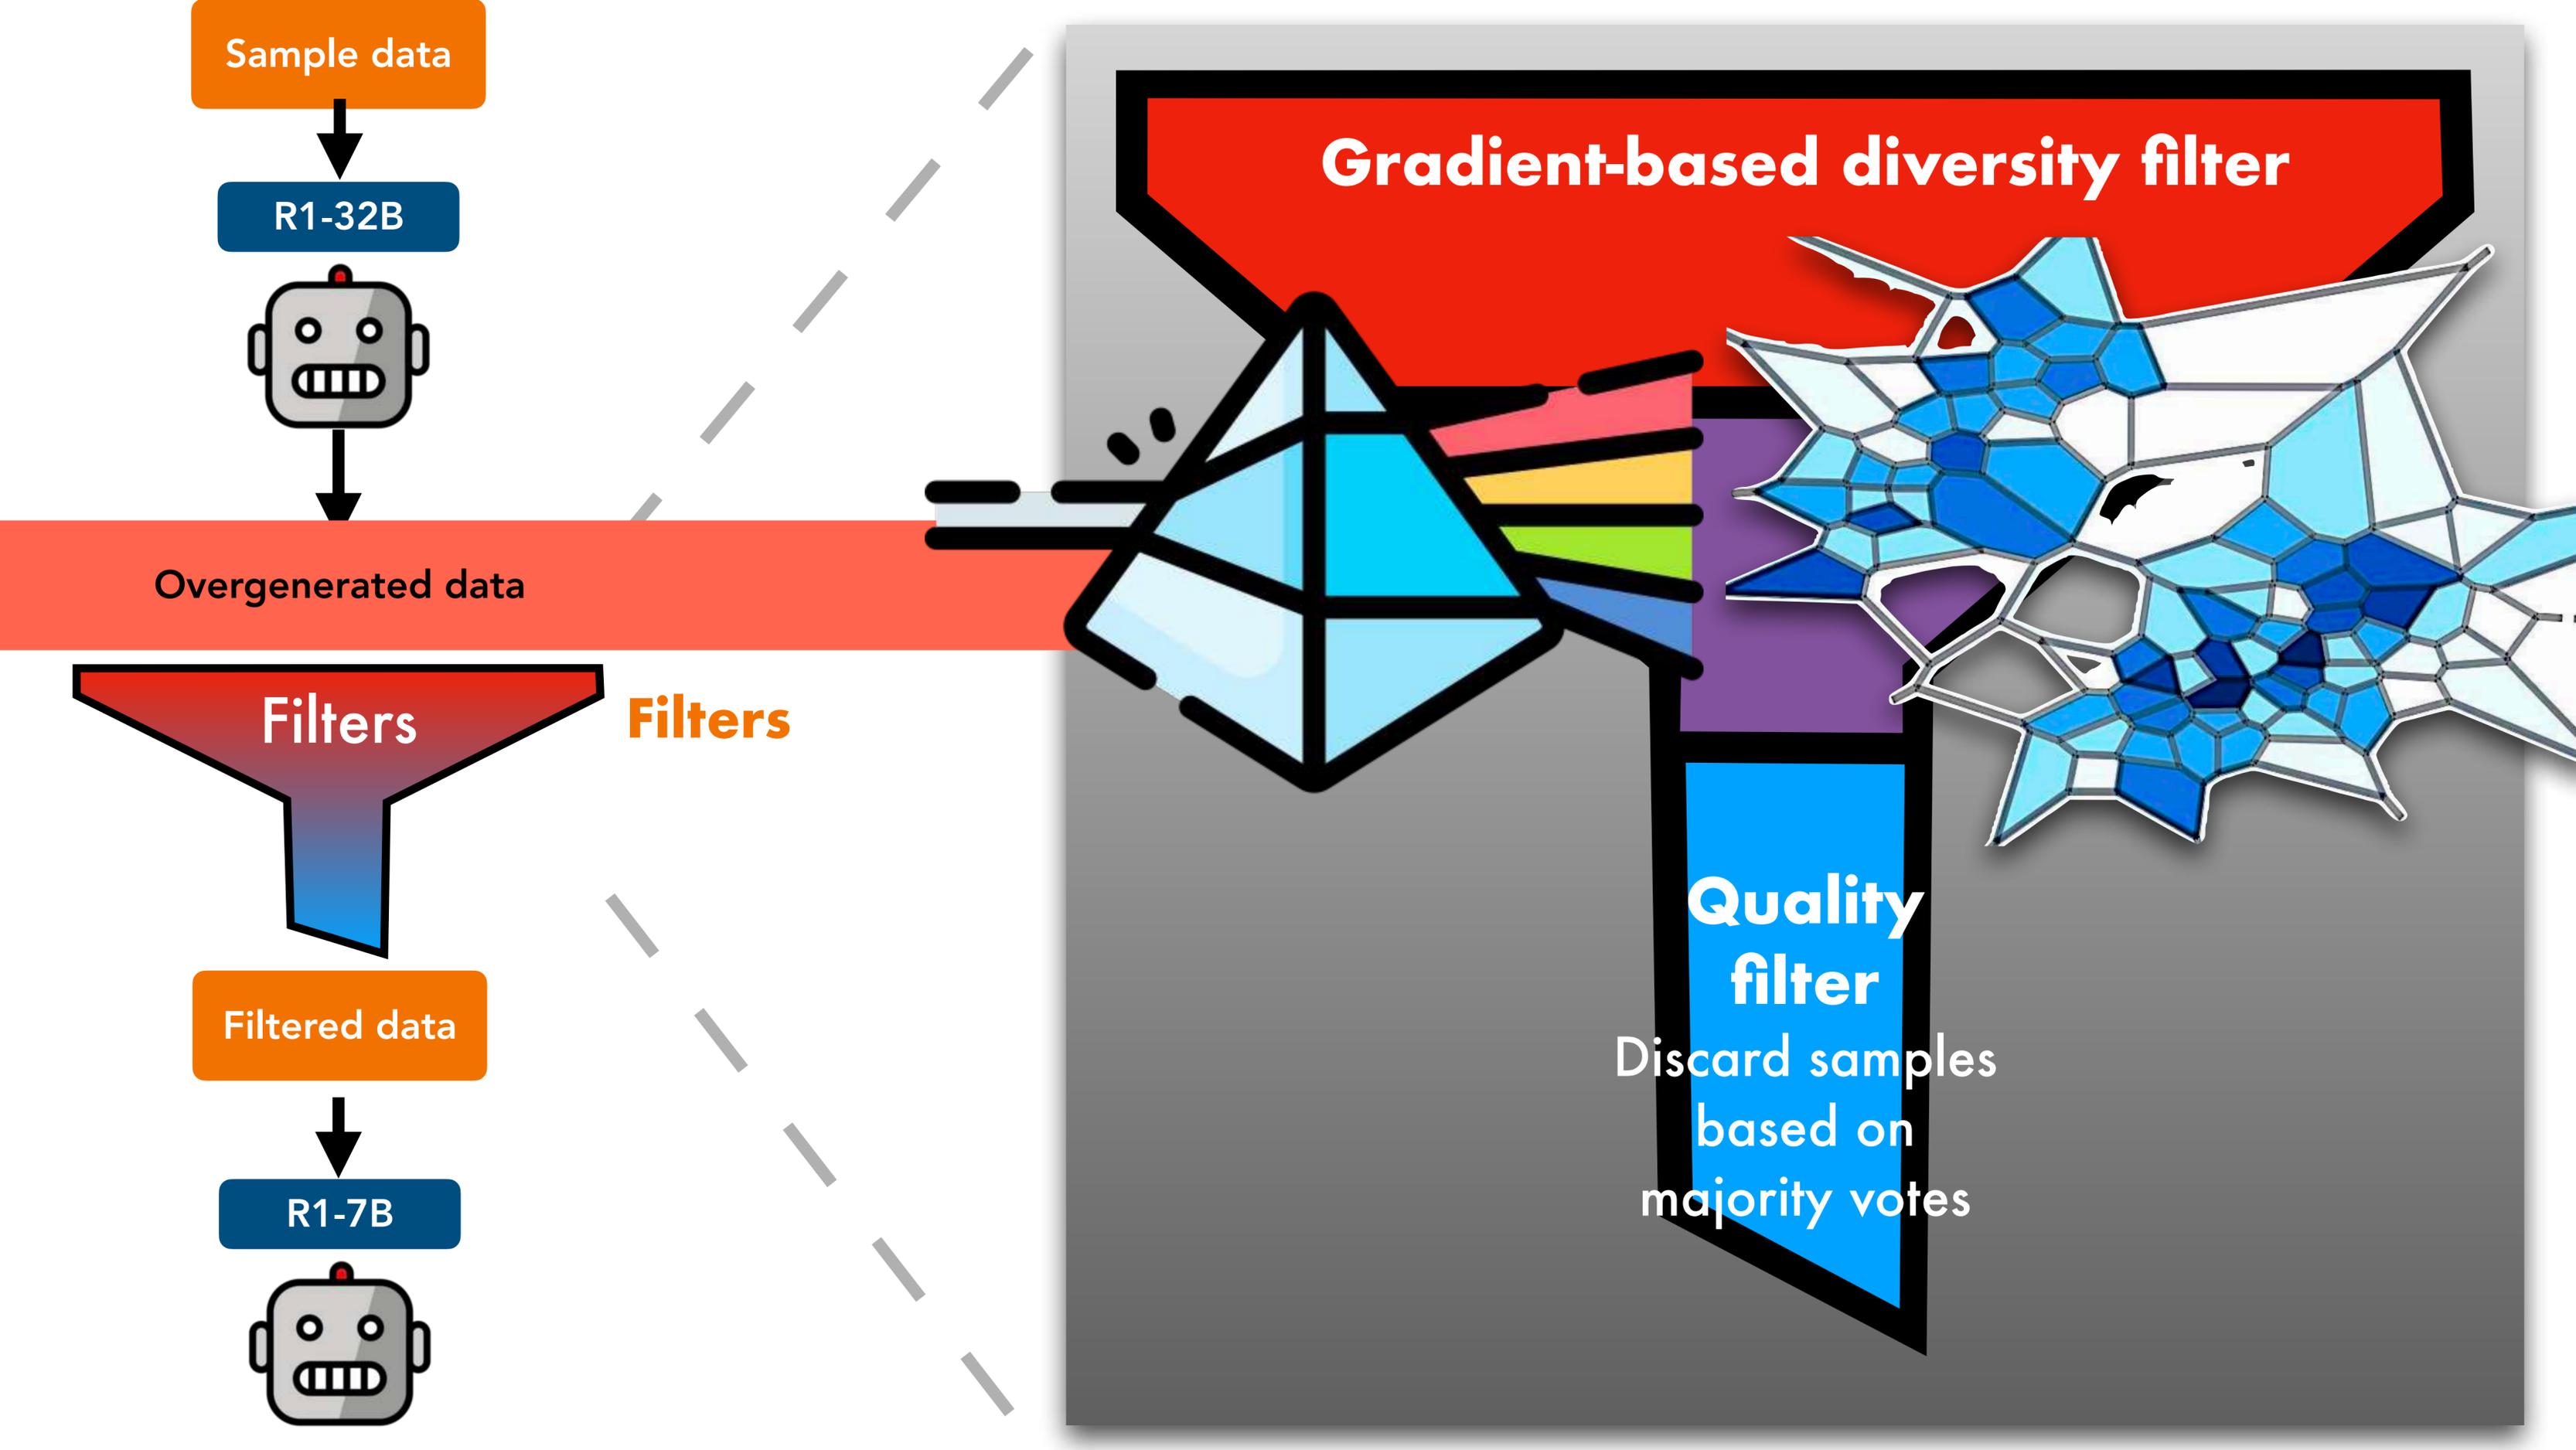
# G-Vendi Score

G-vendi score predicts generalization (i.e., *out-of-distribution* generalization)!



# Prismatic Synthesis





Sample data

R1-32B

Gradient-based diversity filter

Overgenerated data

Filters

Filters

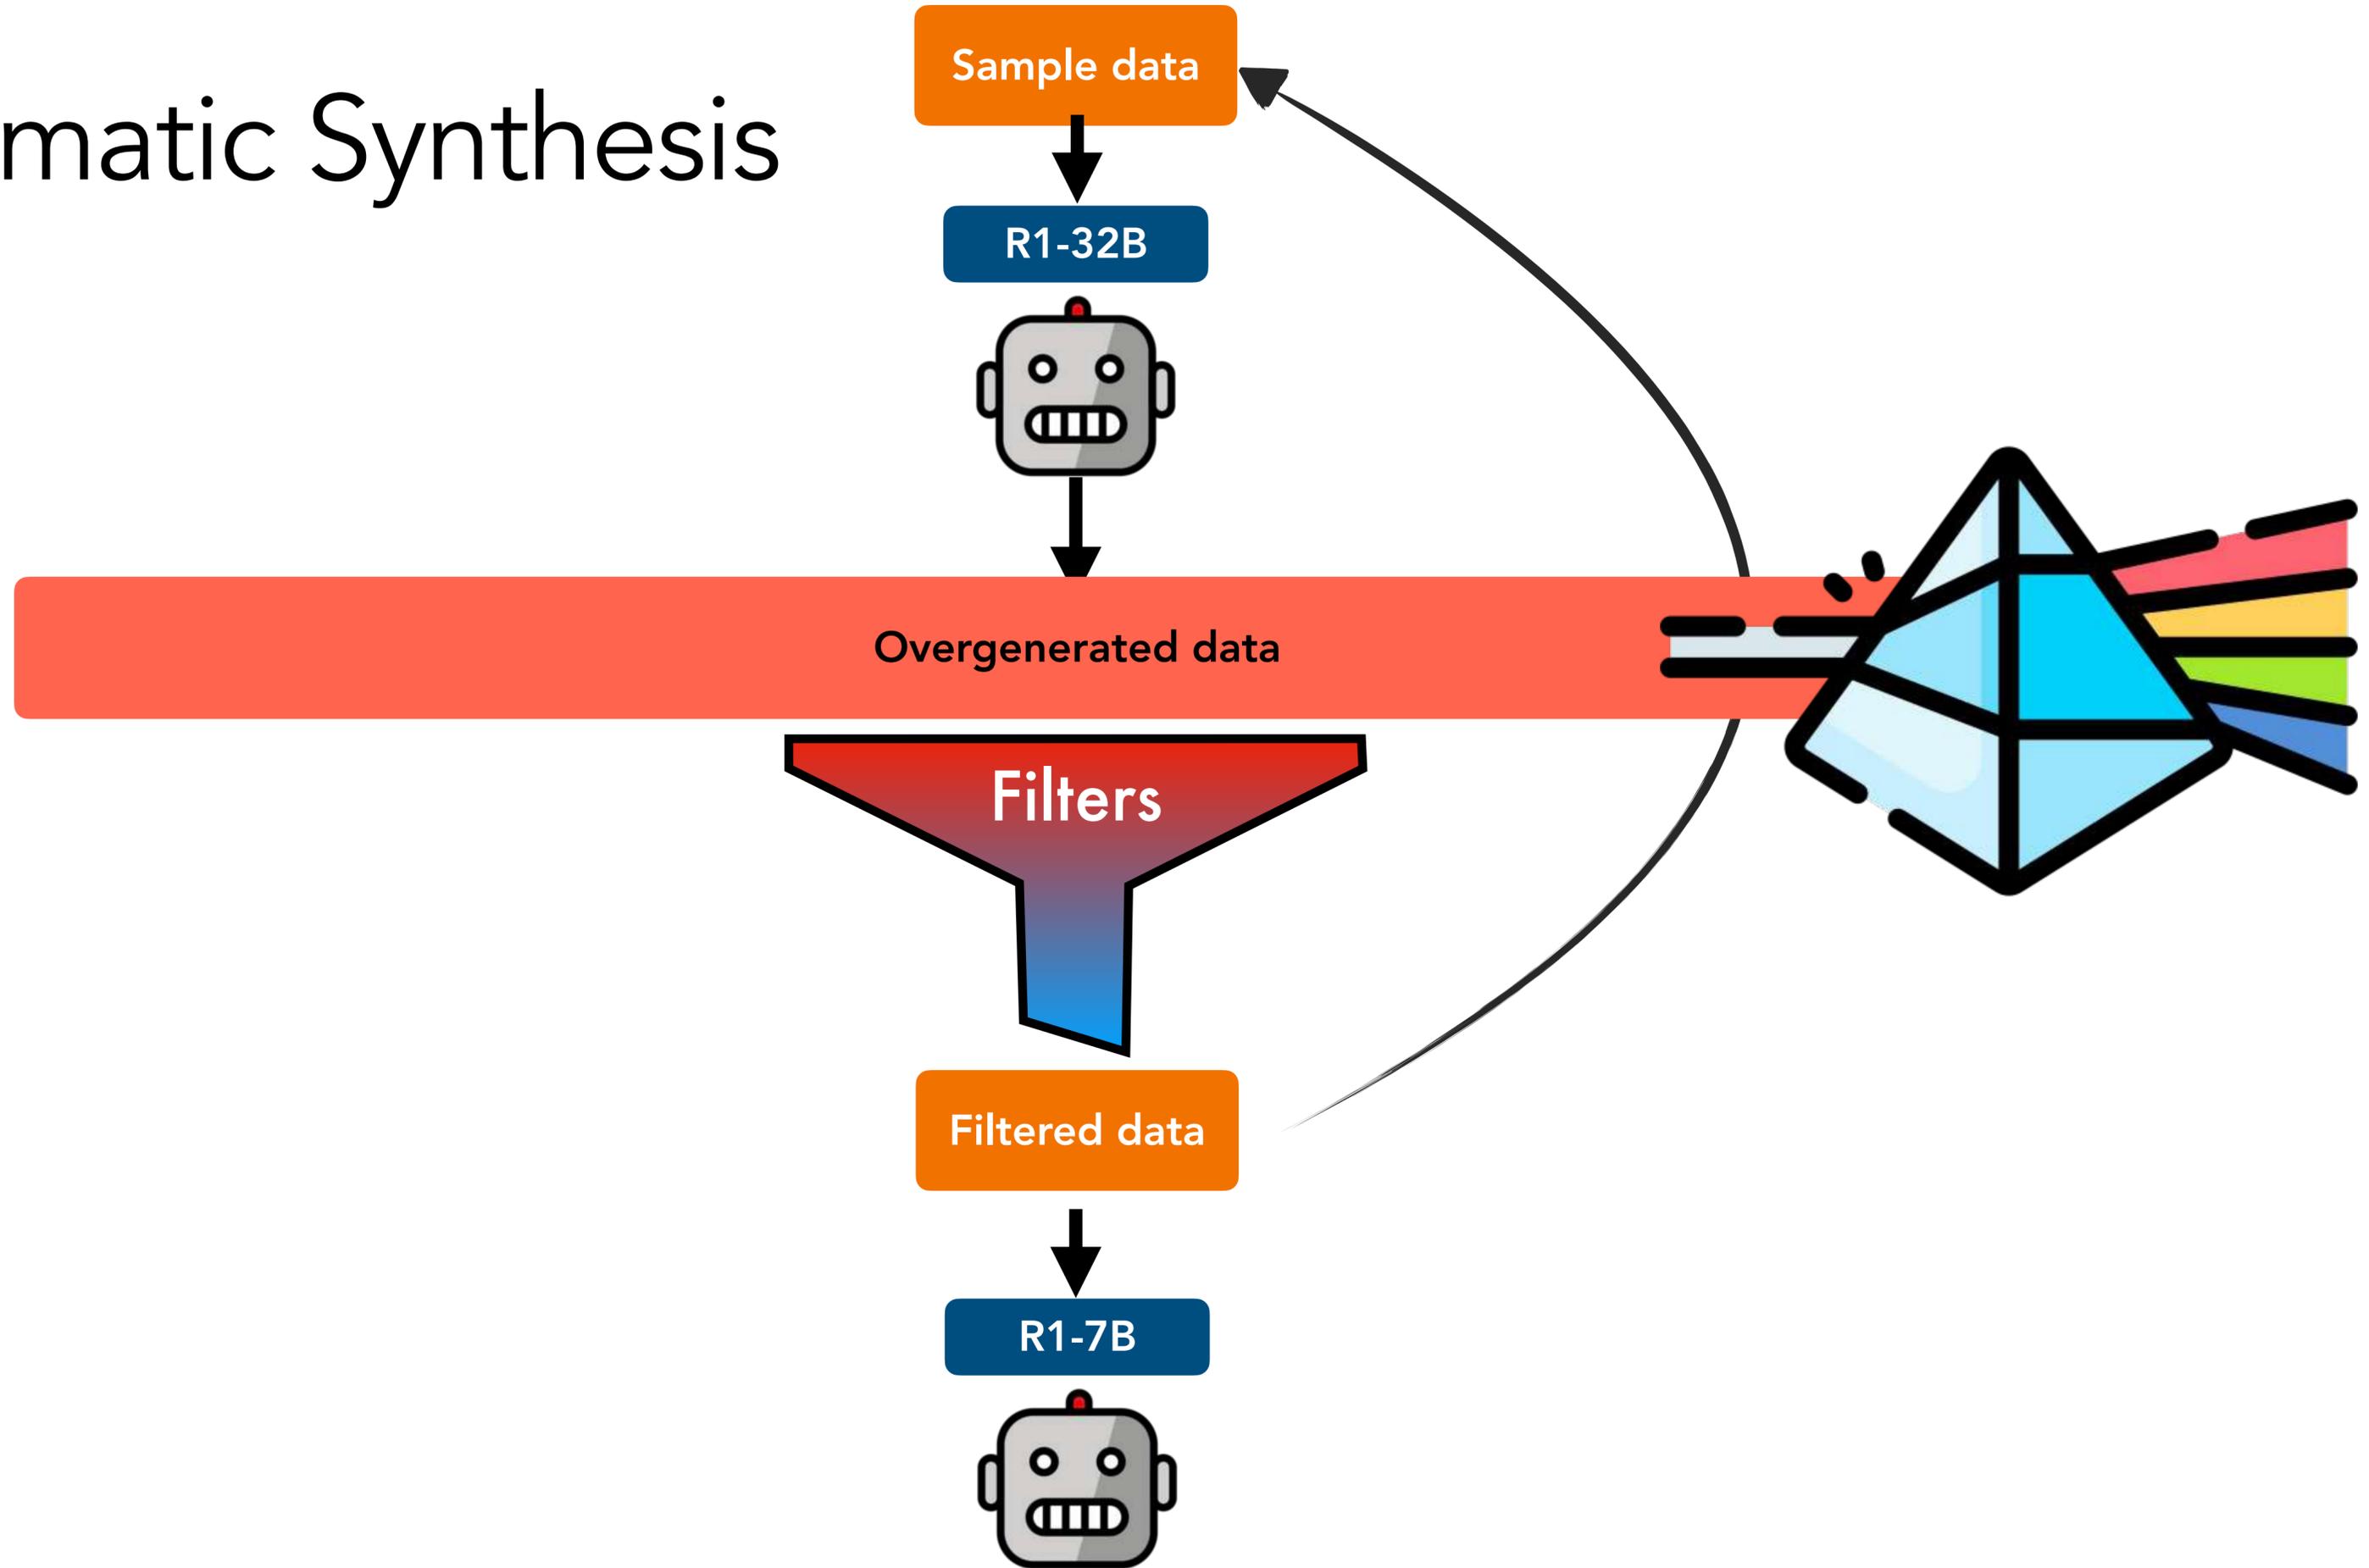
Quality filter

Discard samples based on majority votes

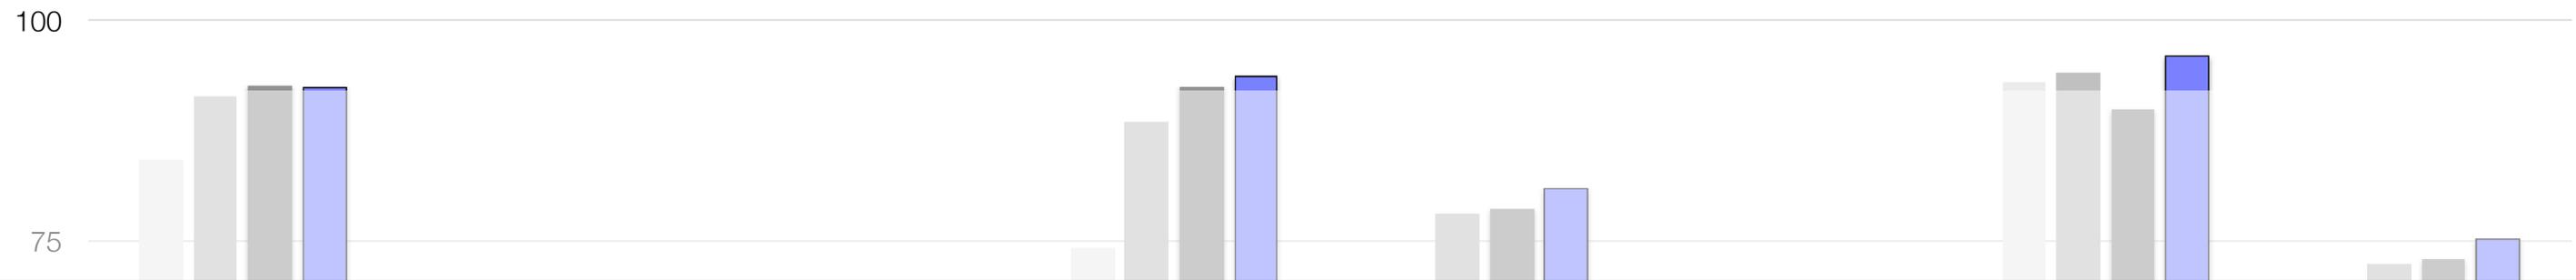
Filtered data

R1-7B

# Prismatic Synthesis



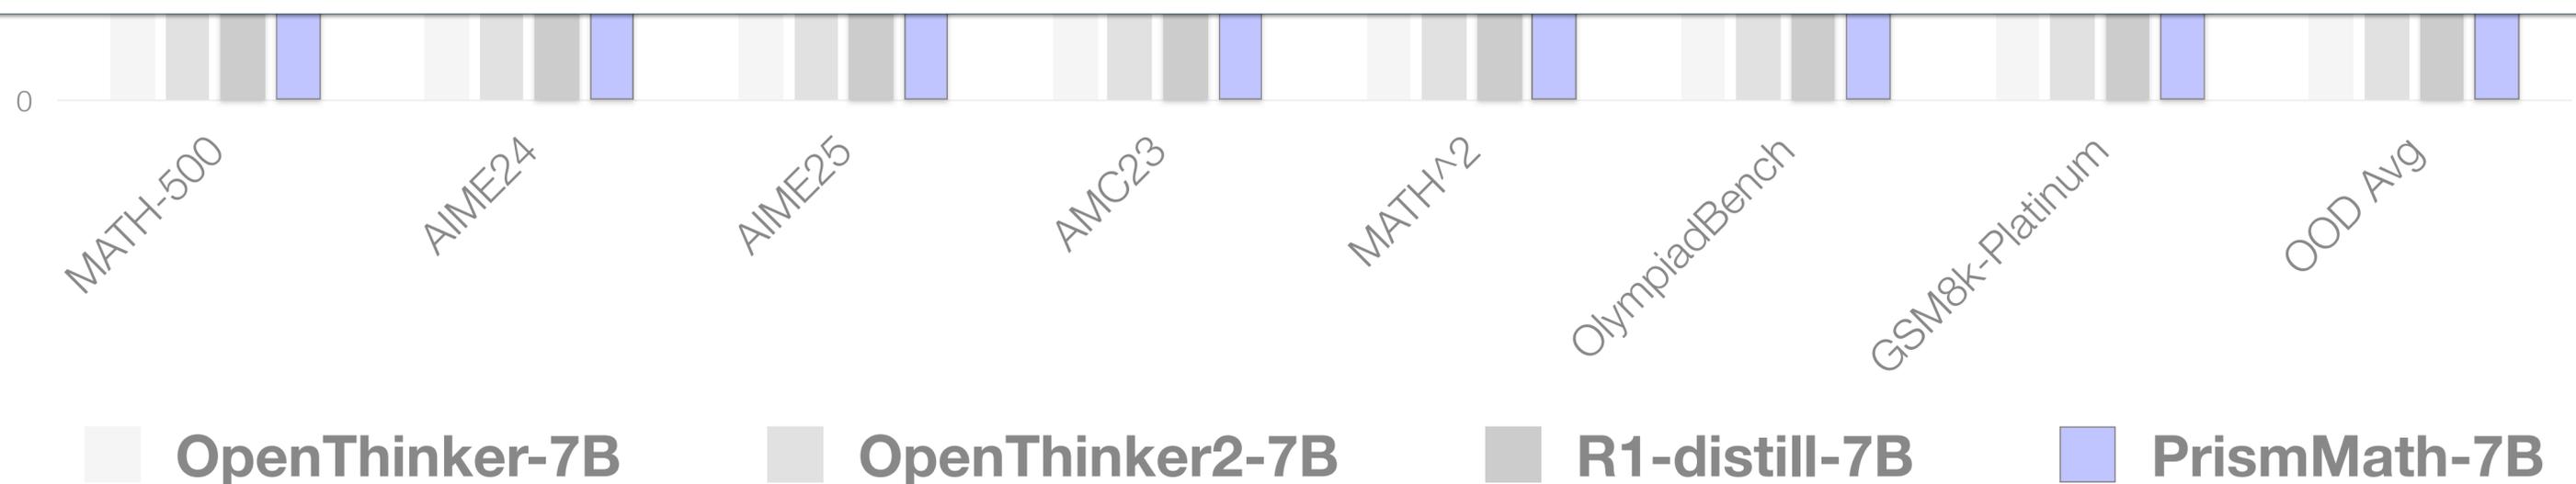
# Prismatic Synthesis improves long-CoT reasoning



Outperforms baselines, while

(1) using **20x smaller R1-32B** as data generator,

(2) relying on **zero human-labeled answers**, i.e., entirely model generated!



# Concluding Thoughts

---

- Reasoning requires data that transcends the internet data
- Synthetic data to the rescue
- RL can be viewed as a form of (implicit) synthetic data generation via exploration
- RL or not, existing methods often lack the bird eye view on the overall diversity
- Systematic diversification (+ quality check) can make a big difference, even to overcome 20x size difference of the R1 teacher models

# h: Underdogs, Art of Battling

In this talk:

===

ProRL: Prolonged RL

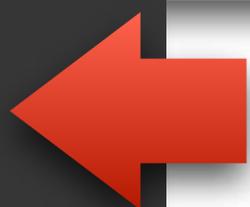
Prismatic Synthesis

RL as Pretraining

===

“Smaller but Better”

“Algorithms for the Win”



*David vs. Goliath:*

*the Art of Scaling Intelligence*

*in the Era of Extreme-Scale Neural Models*

Three components to innovate:

Unconventional data 🔥

Unconventional algorithms 🚀

Unconventional collaboration 🌍

# RLP

## Reinforcement as a Pretraining Objective

— ICLR 2026 —

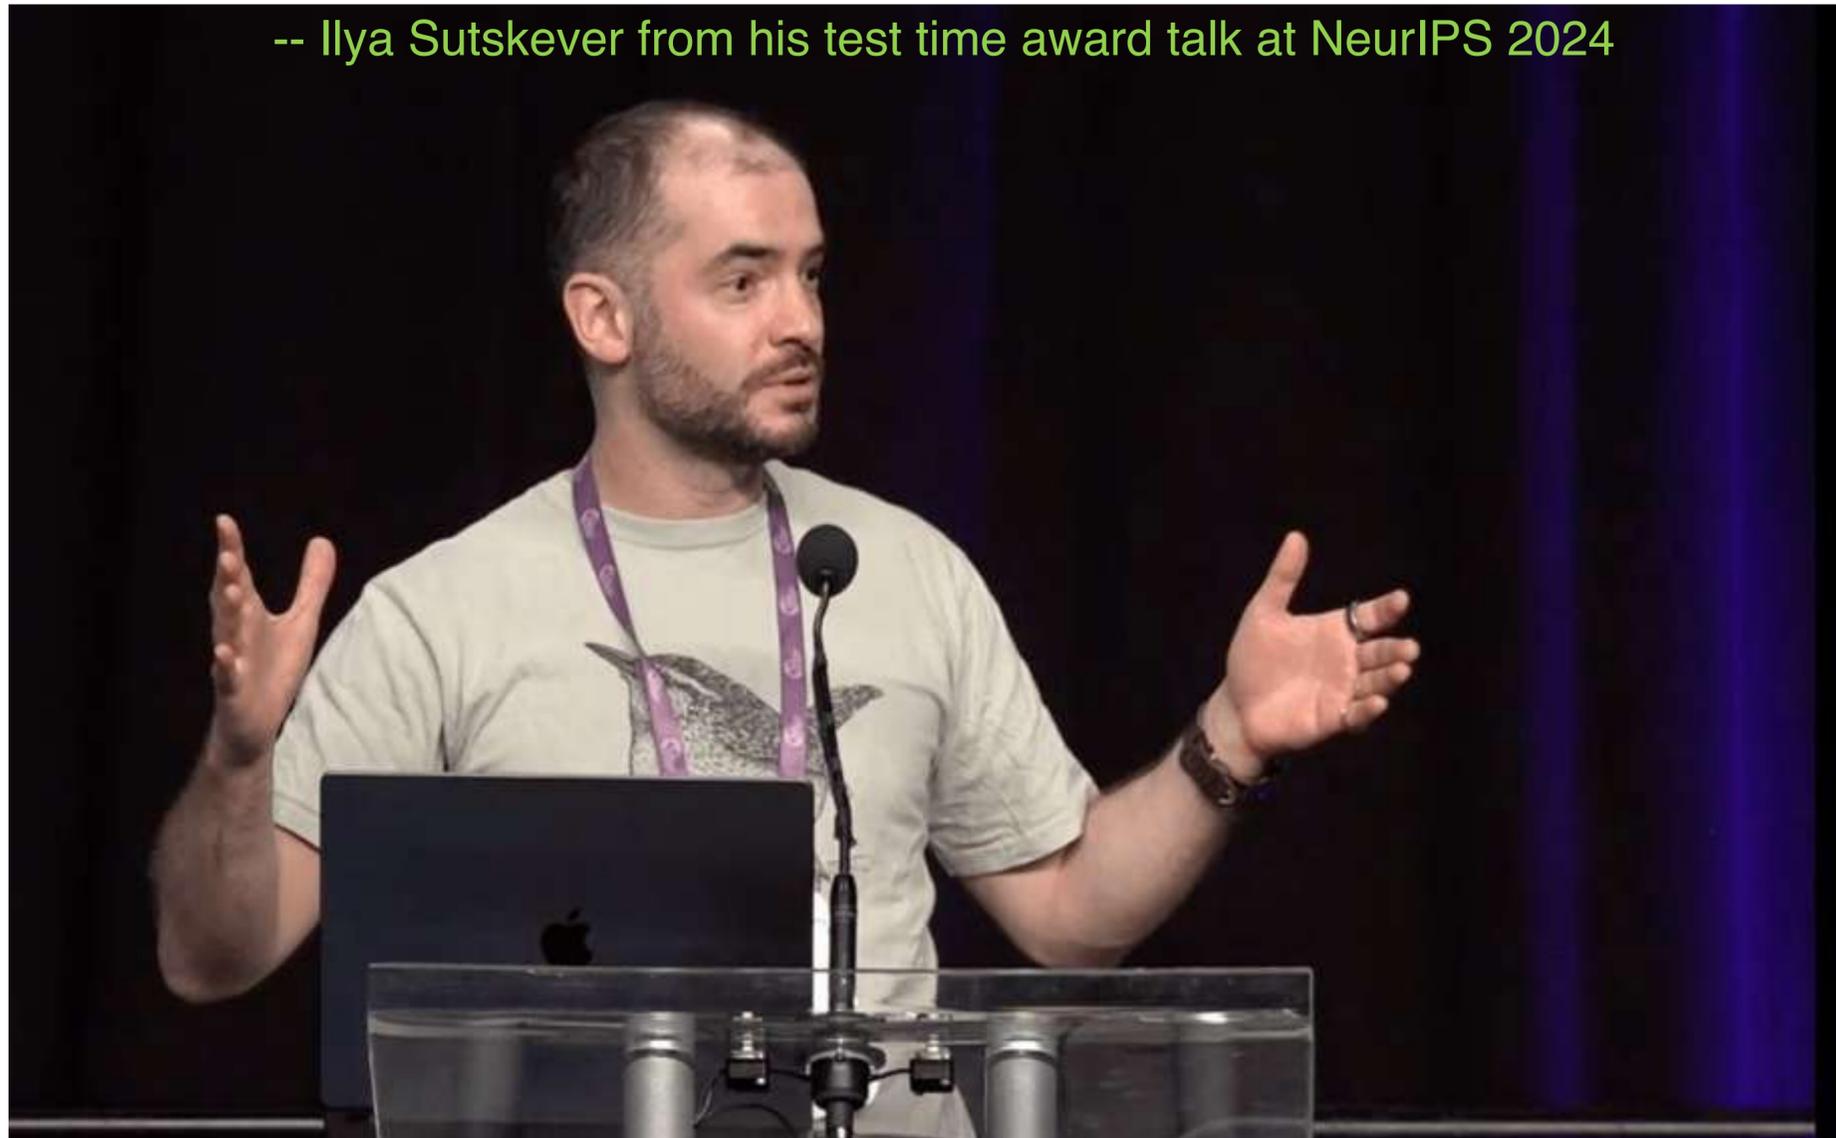


**Ali Hatamizadeh<sup>†1</sup>, Syeda Nahida Akter<sup>†2\*</sup>, Shrimai Prabhume<sup>†1,3</sup>, Jan Kautz<sup>1</sup>,  
Mostofa Patwary<sup>1</sup>, Mohammad Shoeybi<sup>1</sup>, Bryan Catanzaro<sup>1</sup>, Yejin Choi<sup>1,4</sup>**  
NVIDIA<sup>1</sup>, Carnegie Mellon University<sup>2</sup>, Boston University<sup>3</sup>, Stanford University<sup>4</sup>  
ahatamizadeh@nvidia.com, sprabhume@nvidia.com

# The End of Ever More Pretraining

- Scaling compute is no longer the bottleneck
- High-quality data is finite and increasingly exhausted
- The era shifts from scaling data to extracting more value per token

-- Ilya Sutskever from his test time award talk at NeurIPS 2024



## Pre-training as we know it will end

Compute is growing:

- Better hardware
- Better algorithms
- Larger clusters

Data is not growing:

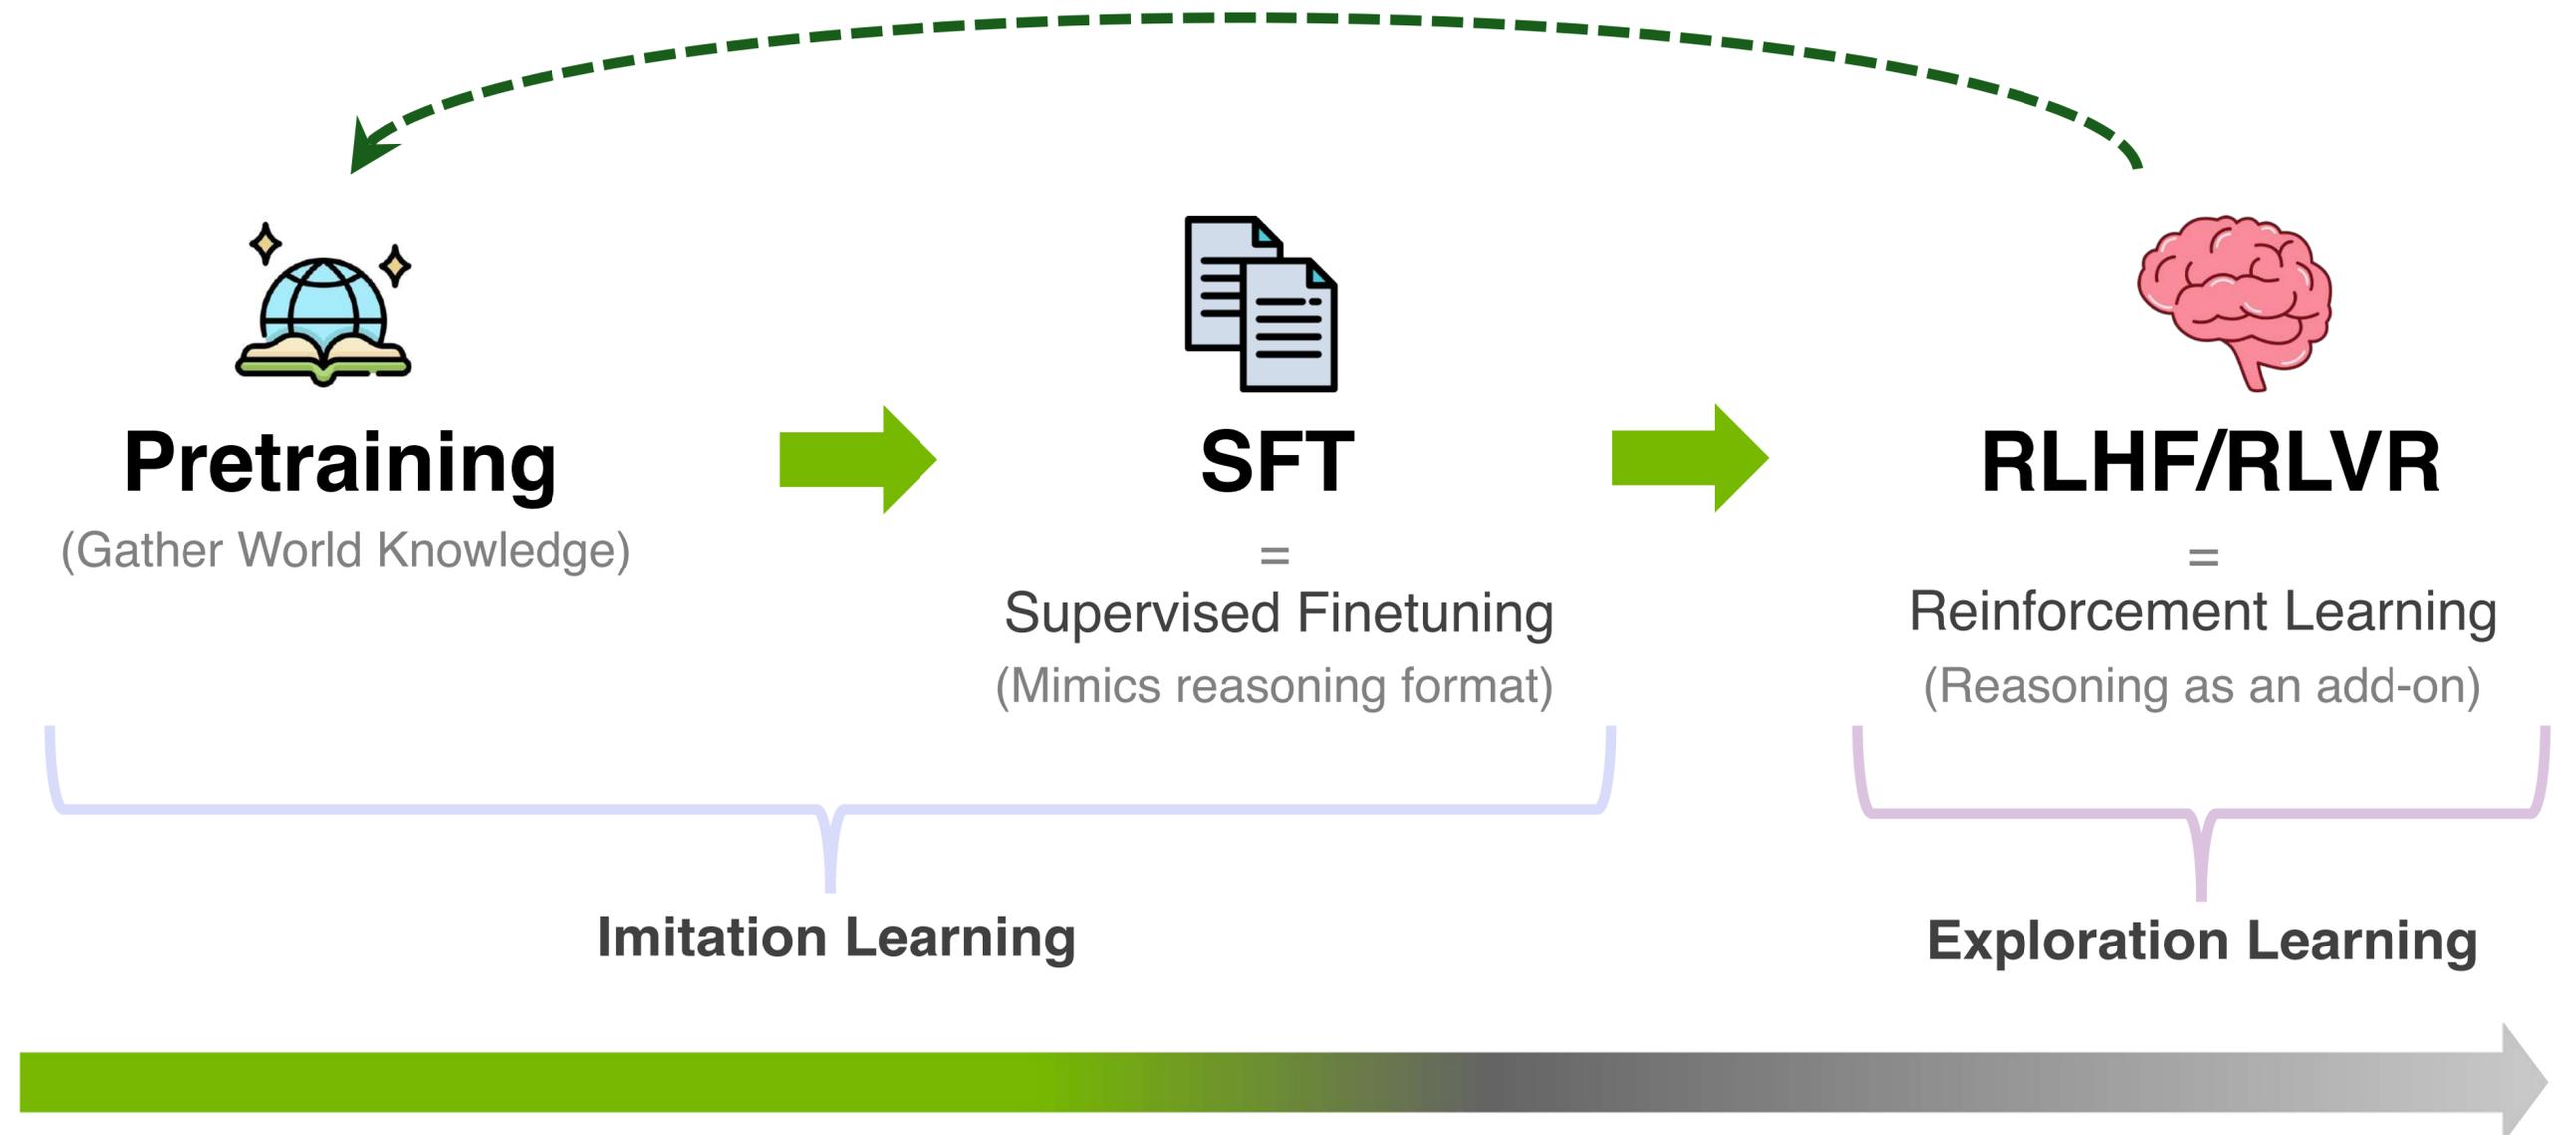
- We have but one internet
- **The fossil fuel of AI**



# The Problem with Standard LLM Training

Reasoning is an afterthought — we can do better

**Q1: Can reasoning be baked in earlier during pretraining — not just added post-hoc?**



**Q2: Do gains from early reasoning exposure persist through post-training — or get washed out?**

# Vanilla Pretraining vs RLP Pretraining

Same context — but RLP induces reasoning.



Photosynthesis is the process plants, algae and some bacteria use to make their own food using \_\_\_\_\_

**Vanilla Pretraining  
(Next Token Prediction)**

**P(next token | context)**  
(Pattern Completion)

**RLP Pretraining**

**<think>**Photosynthesis relies on solar energy.  
Hence the next token must be sunlight. **</think>**

**P(next token | context, thought)**  
(Reasoning driven prediction)

**sunlight**

**Key difference:** RLP produces an explicit reasoning trace before predicting the token — making the “why” visible and trainable, not just the final answer.

# From NTP to RLP: Teaching Models to Think First

Treat CoT as an action; reward by information gain; update only thought tokens.

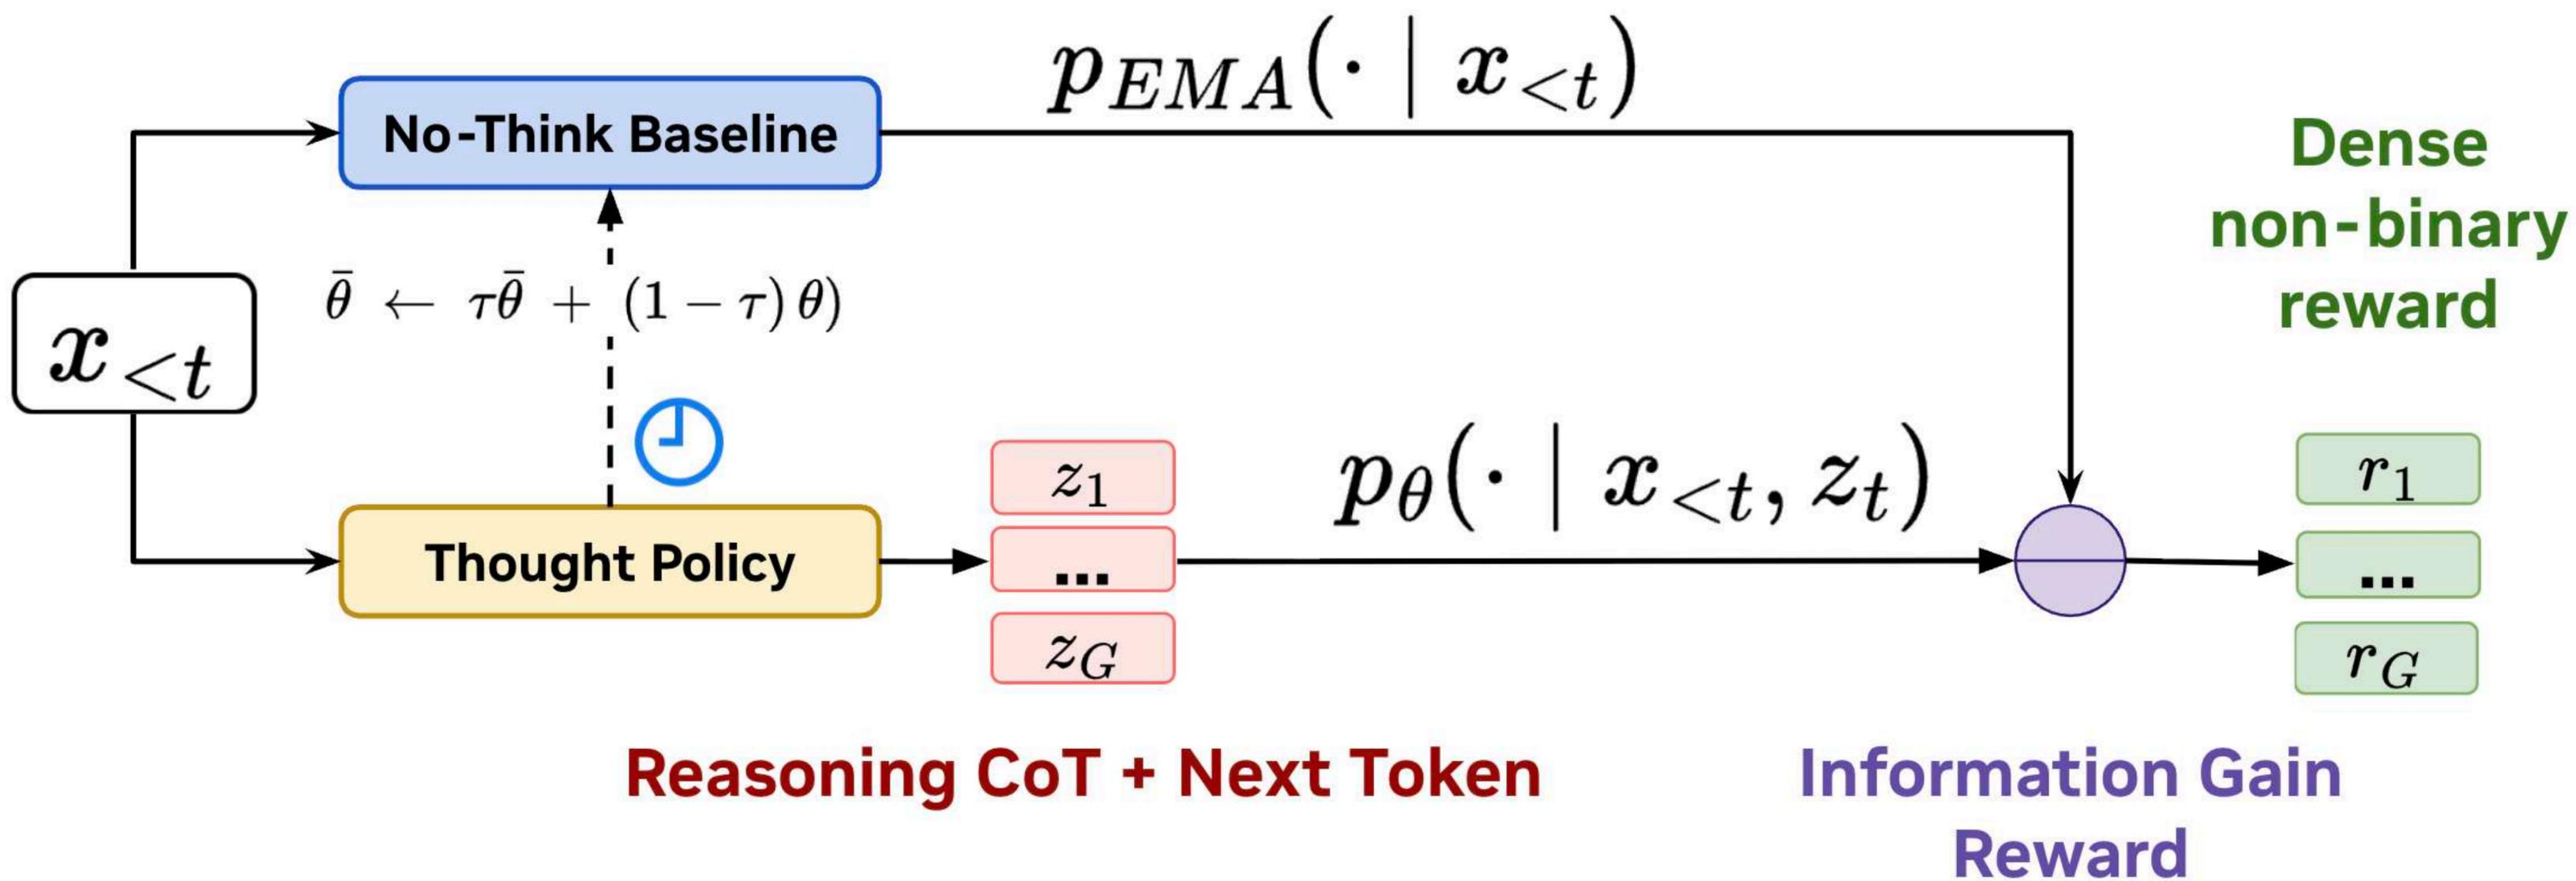
## How RLP works

- Sample a thought before predicting the next token.
- Compare likelihoods with vs. without the thought

$$r(c_t) = \log p_{\theta}(x_t | x_{<t}, c_t) - \log \bar{p}_{\phi}(x_t | x_{<t})$$

- Verifier-free, position-wise credit at every step.

# RLP Framework



# Training Recipe

- **Single network** acts as policy + reasoned predictor.
- **EMA “no-think” teacher** with slow lag  $\phi \leftarrow \tau \phi + (1 - \tau) \theta$ ,  $\tau = 0.999$
- **Group-relative advantages** with inclusive-mean correction

$$\bar{r} = \frac{1}{G} \sum_{j=1}^G r(c_t^{(j)}), \quad A^{(i)} = \frac{G}{G-1} \left( r(c_t^{(i)}) - \bar{r} \right)$$

- **Clipped, per-token surrogate** (GRPO-style) on **thought tokens** only

$$\mathcal{L}_{\text{clip}} = - \mathbb{E} \left[ \frac{1}{|c_t^{(i)}|} \sum_u \min \left( \rho_u^{(i)} A^{(i)}, \text{clip}(\rho_u^{(i)}; 1 - \epsilon_\ell, 1 + \epsilon_h) A^{(i)} \right) \right]$$

with per-token importance ratios

$$\rho_u^{(i)} = \exp \left( \log \pi_\theta(\ell_u^{(i)} \mid x_{<t}, \ell_{1:u-1}^{(i)}) - \log \pi_{\theta_{\text{old}}}(\ell_u^{(i)} \mid x_{<t}, \ell_{1:u-1}^{(i)}) \right)$$

# Leverage Early Reasoning with "RLP"

Q1

Can RLP improve the reasoning ability of a base model without any prior task-specific tuning? (And better than Base + CPT?)

Does the gain of RLP survive even after SFT+RLVR?

Q2

Can RLP, when applied to an earlier pre-training checkpoint with 200B less tokens (to match the pretraining FLOPs), match the performance of a fully trained base model?

## Training Ingredients

**Model:** Qwen3-1.7B-Base

**Checkpoint:** Final model

**Token Budget:** 1B Tokens

**CPT Data (D):**

- General Pretraining Corpora (Pretraining-Dataset)

## Comparison with:

- Base model
  - **BASE**
  - **BASE+POST**
- **Token Matched:** NTP-based CPT on base model
  - **BASE+CPT**
  - **BASE+CPT+POST**

## Training Ingredients

**Model:** NVIDIA-Nemotron-Nano-12B-v2-Base

**Checkpoint:** Model trained till 19.8T tokens

**Token Budget:** 250M Tokens

**CPT Data (D):**

- General Pretraining Corpora (Nemotron-Pretraining-Dataset)

## Comparison with:

- FLOP Matched: Base model (20T)
  - **BASE**
  - **BASE+POST**

**RLP with 200B less tokens!**

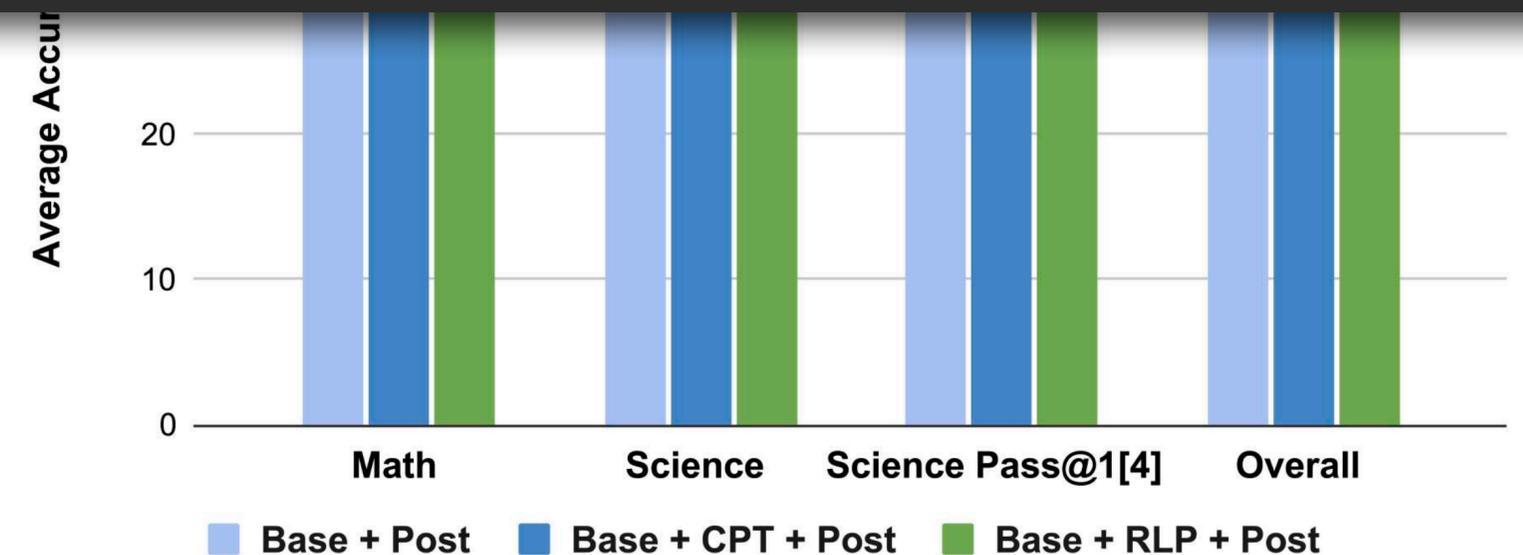
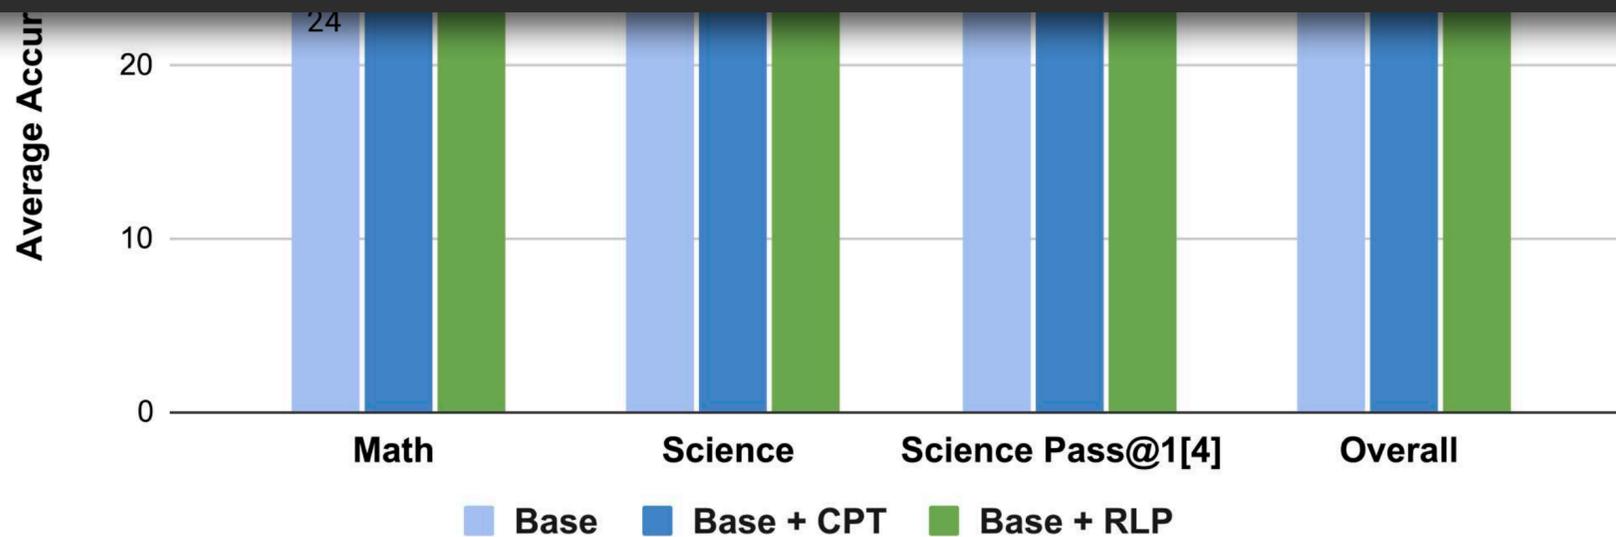
# RLP shows a significant improvement on Qwen3-1.7B-Base

**RLP** outperforms **BASE** by **+19%** and **CPT** by **+17%** on average across math and science benchmarks.

After identical SFT + RLVR, **RLP** compounds its **advantage: +8%** relative over **BASE+Post**; **+7%** relative over **CPT+Post**

Unlike NPT or continuous pretraining, **RLP** instills reasoning during pretraining itself

The early advantages **compound** through post training, giving models stronger and more robust reasoning capabilities.



# RLP shows a significant improvement on Qwen3-1.7B-Base



Does the improvement sustain under compute equivalent baselines?

$$T_{flop} = (n \times l_{seq} \times bs \times iters) + T_{inp}$$

$n$  = rollouts

$bs$  = batch size

$l_{seq}$  = sequence length

## Training Ingredients

**Model:** Qwen3-1.7B-Base

**Checkpoint:** Final model

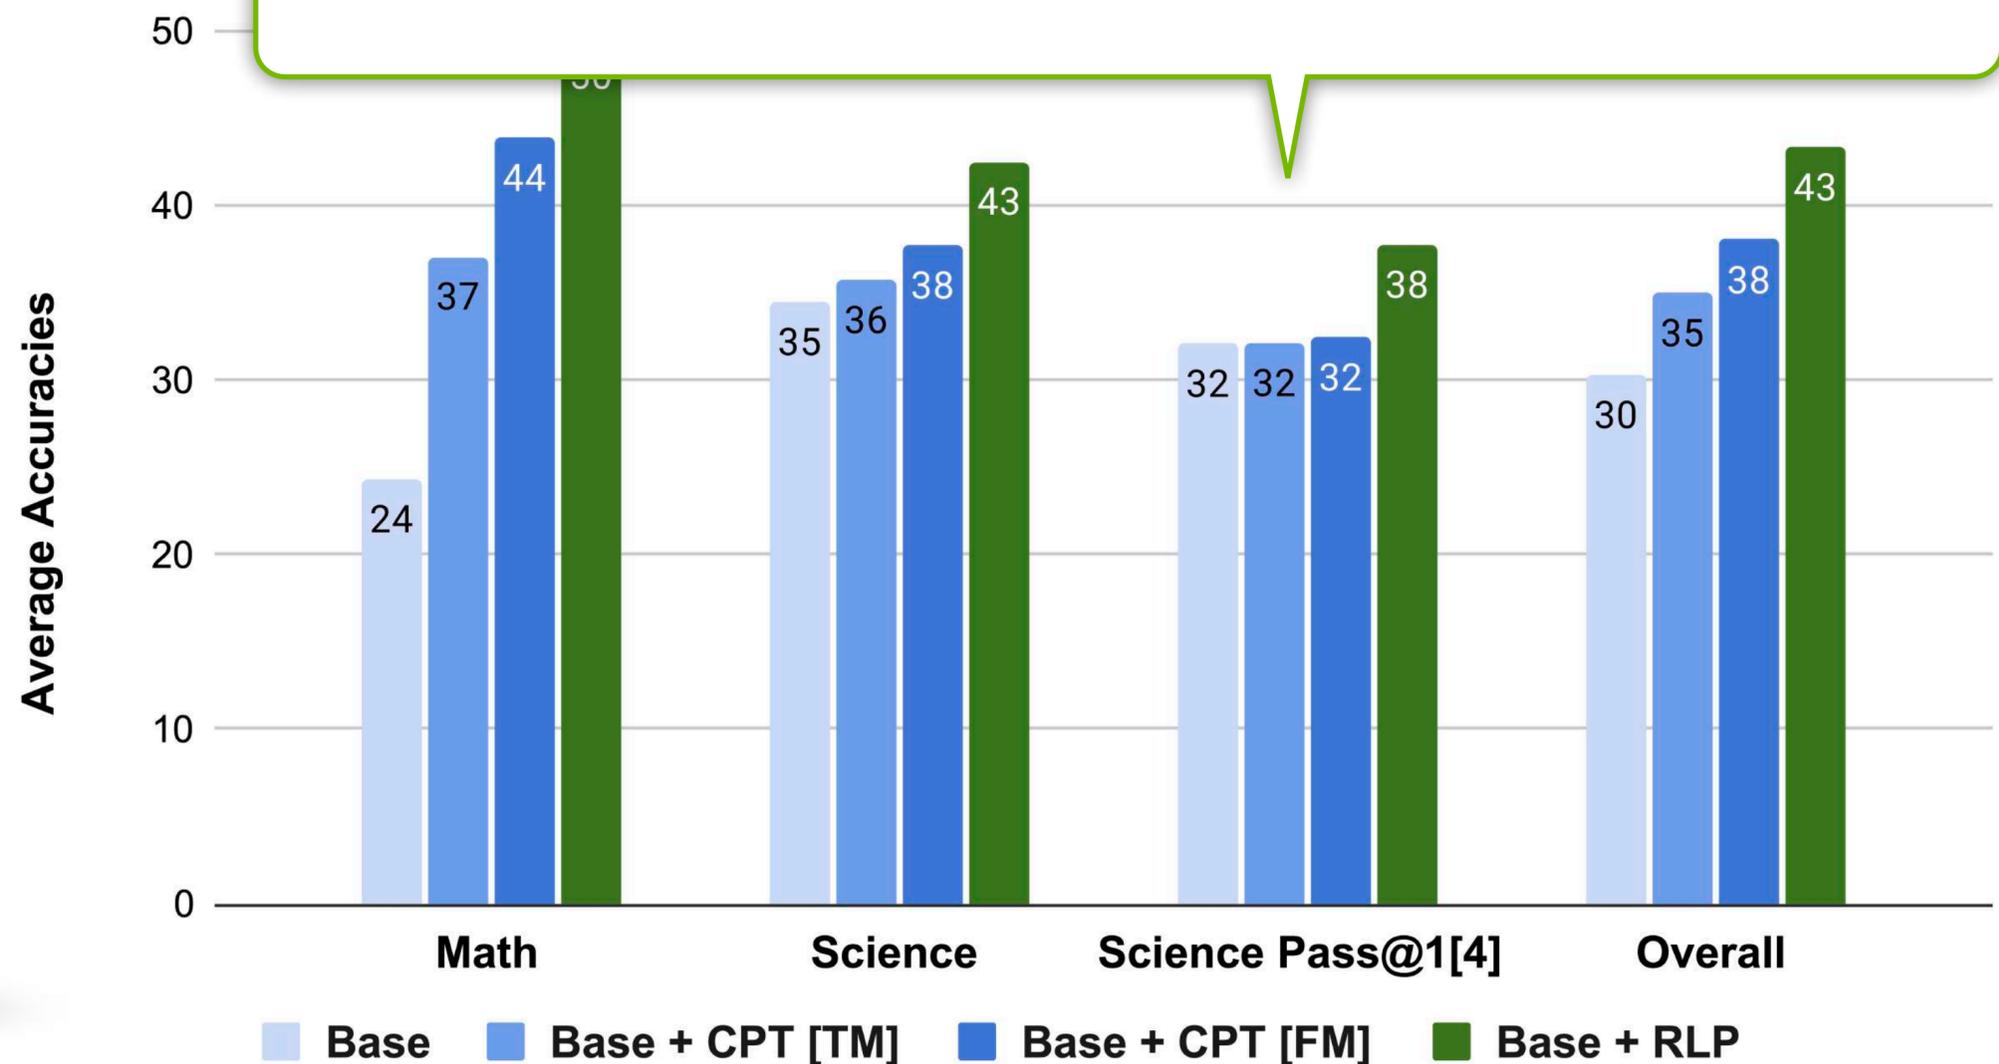
**Token Matching ( $T_{inp}$ ):** 170M Tokens

**Flop Matching ( $T_{flop}$ ):** 6B Tokens

**CPT Data ( $\mathcal{D}$ ):**

- Nemotron-CrossThink

Even with flop matching training, **RLP outperforms CPT** (exposed to **35X more data**) by **+14%** on average!



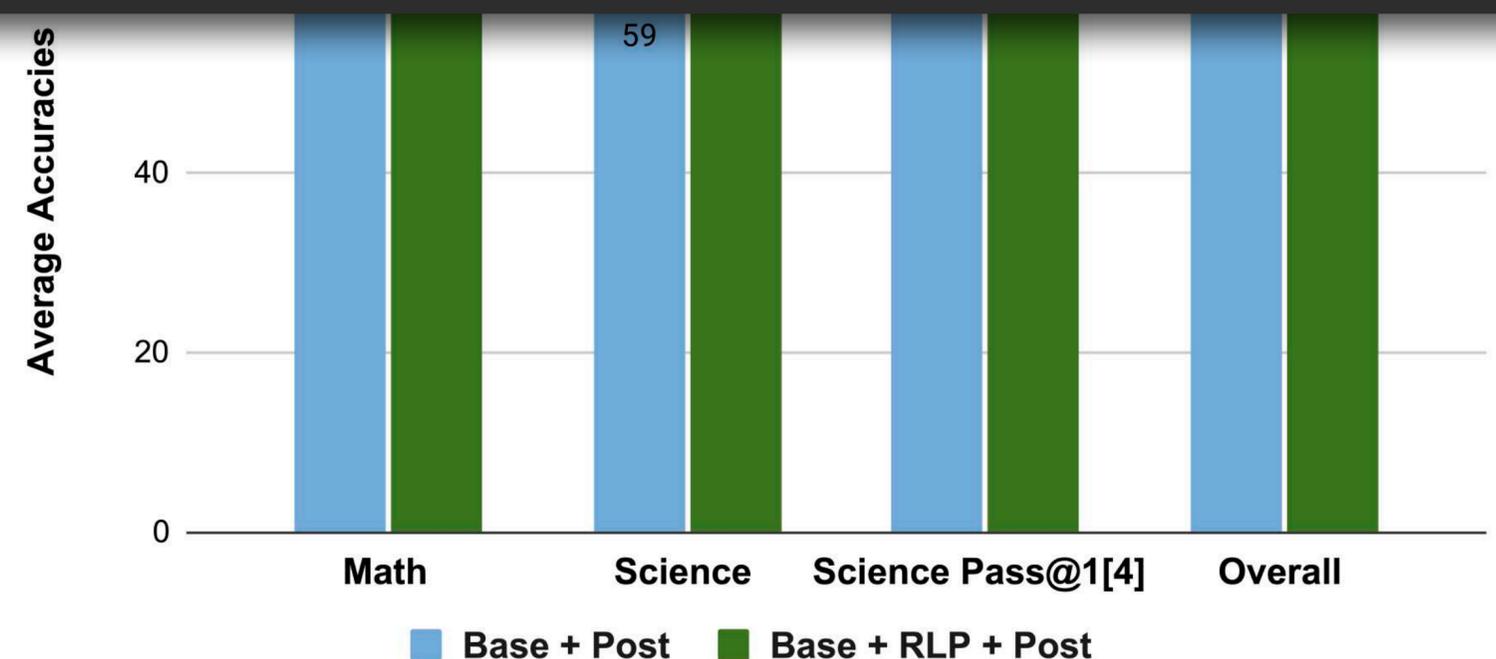
# RLP scales with LLM size (12B) and frameworks

**RLP**, trained on **~200B fewer** tokens, achieves a **35% average gain**, with the largest boost in science (**+23% absolute**), showing robust cross-domain benefits at

After identical post-training, **RLP** outperforms **BASE** by **+2% absolute margin**, which has been seen **~200B** more tokens during pretraining.

The benefits of **RLP** *persist and even amplify* when scaling to larger models and across architectures.

*Efficiency meets intelligence*—proof that reasoning, not sheer scale, drives progress.



---

# Quiet-STaR: Language Models Can Teach Themselves to Think Before Speaking

**Eric Zelikman**  
Stanford University

**Georges Harik**  
Notbad AI Inc

**Yijia Shao**  
Stanford University

**Varuna Jayasiri**  
Notbad AI Inc

**Nick Haber**  
Stanford University

**Noah D. Goodman**  
Stanford University



2025-09-26

---

## Reinforcement Pre-Training

---

**Yao Tang**<sup>†</sup>   **Qingxiu Dong**<sup>\*†‡</sup>   **Li Dong**<sup>\*†</sup>  
**Tianzhu Ye**<sup>†§</sup>   **Yutao Sun**<sup>†§</sup>   **Zhifang Sui**<sup>‡</sup>   **Furu Wei**<sup>†◇</sup>

<sup>†</sup> Microsoft Research

<sup>‡</sup> Peking University

<sup>§</sup> Tsinghua University

<https://aka.ms/GeneralAI>

## Reinforcement Learning on Pre-Training Data

Siheng Li<sup>1,3,\*†</sup>, Kejiao Li<sup>1,†</sup>, Zenan Xu<sup>1,†</sup>, Guanhua Huang<sup>1</sup>, Evander Yang<sup>1</sup>, Kun Li<sup>1,3,\*</sup>, Haoyuan Wu<sup>1</sup>, Jiajia Wu<sup>1</sup>, Zihao Zheng<sup>1</sup>, Chenchen Zhang<sup>1</sup>, Kun Shi<sup>1</sup>, Kyrierl Deng<sup>1</sup>, Qi Yi<sup>1</sup>, Ruibin Xiong<sup>1</sup>, Tingqiang Xu<sup>1,\*</sup>, Yuhao Jiang<sup>1</sup>, Jianfeng Yan<sup>1</sup>, Yuyuan Zeng<sup>1</sup>, Guanghui Xu<sup>1</sup>, Jinbao Xue<sup>2</sup>, Zhijiang Xu<sup>2</sup>, Zheng Fang<sup>2</sup>, Shuai Li<sup>2</sup>, Qibin Liu<sup>2</sup>, Xiaoxue Li<sup>2</sup>, Zhuoyu Li<sup>2</sup>, Yangyu Tao<sup>2</sup>, Fei Gao<sup>2</sup>, Cheng Jiang<sup>2</sup>, Bo Chao Wang<sup>2</sup>, Kai Liu<sup>2</sup>, Jianchen Zhu<sup>2</sup>, Wai Lam<sup>3</sup>, Bo Zhou<sup>1,‡</sup>, Di Wang<sup>1</sup>

<sup>1</sup>LLM Department, Tencent   <sup>2</sup>HunYuan Infra Team

<sup>3</sup>The Chinese University of Hong Kong

✉ [chaysezhou@tencent.com](mailto:chaysezhou@tencent.com)

# RLP: The Core Idea

Chain-of-thought as an exploratory action during pretraining

$$\text{Reward}(\text{thought}) = \log P(\text{next token} \mid \text{context} + \text{thought}) - \log P(\text{next token} \mid \text{context alone})$$



## Verifier-Free

No external judge needed. The reward signal comes directly from information gain on the next token — applicable to any ordinary pretraining text.



## Dense Reward

Position-wise credit assigned at every token where reasoning helps. Unlike sparse post-training rewards, RLP provides continuous feedback throughout training.



## Scalable

Works on the same massive pretraining corpora used for standard training. No specialized datasets required — reasoning emerges from plain text.

RLP bridges next-token prediction and the emergence of useful chain-of-thought reasoning — all during pretraining.

# Front-Loading Reasoning

## The Synergy between Pretraining and Post-Training Data

— ICLR 2026 —



**Syeda Nahida Akter<sup>2\*</sup>, Shrimai Prabhumoye<sup>1,3</sup>, Eric Nyberg<sup>2</sup>, Mostofa Patwary<sup>1</sup>,  
Mohammad Shoeybi<sup>1</sup>, Yejin Choi<sup>1,4</sup>, Bryan Catanzaro<sup>1</sup>**

NVIDIA<sup>1</sup>, Carnegie Mellon University<sup>2</sup>, Boston University<sup>3</sup>, Stanford University<sup>4</sup>  
sakter@andrew.cmu.edu, sprabhumoye@nvidia.com

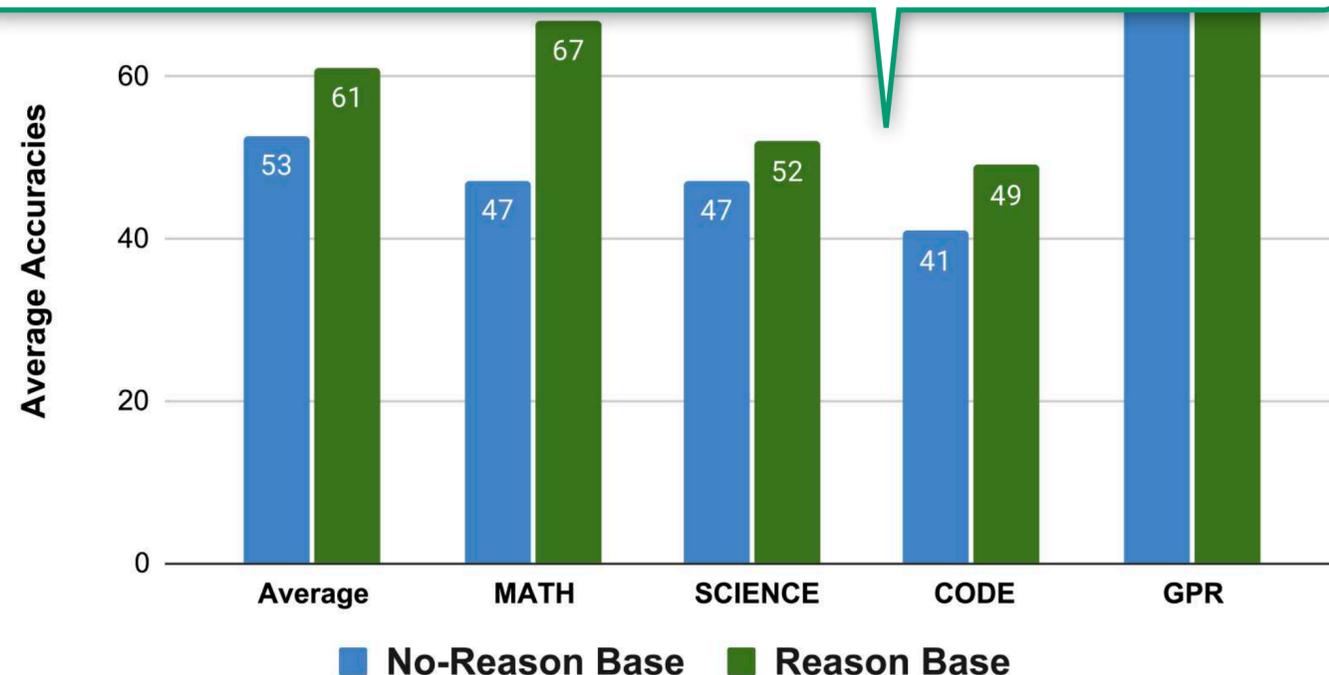
# Front Loading Reasoning

FLR systematically injects reasoning-style data at different phases of training—pretraining, SFT, RL—while varying its diversity, quantity, and quality.

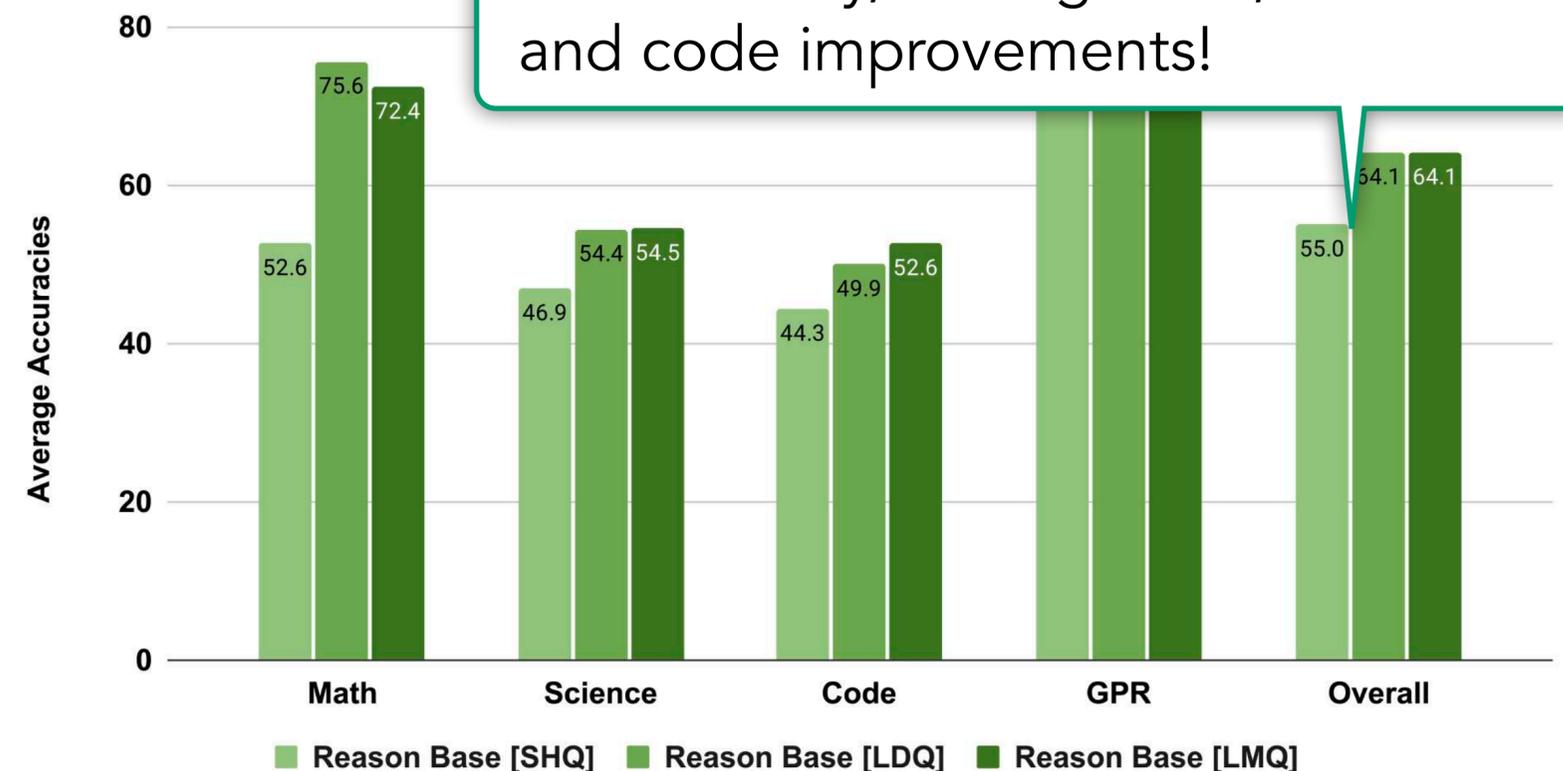
FLR

RQ1: Is the inclusion of reasoning data in pretraining beneficial for the base model?  
Lesson I: Yes, **highly beneficial**, and the benefit scales with **diversity and quantity**!

Using reasoning data in pretraining yields an absolute **+16% average improvement** over the no-reason baseline.



Pretraining gains improve with scale and diversity, driving math, science and code improvements!

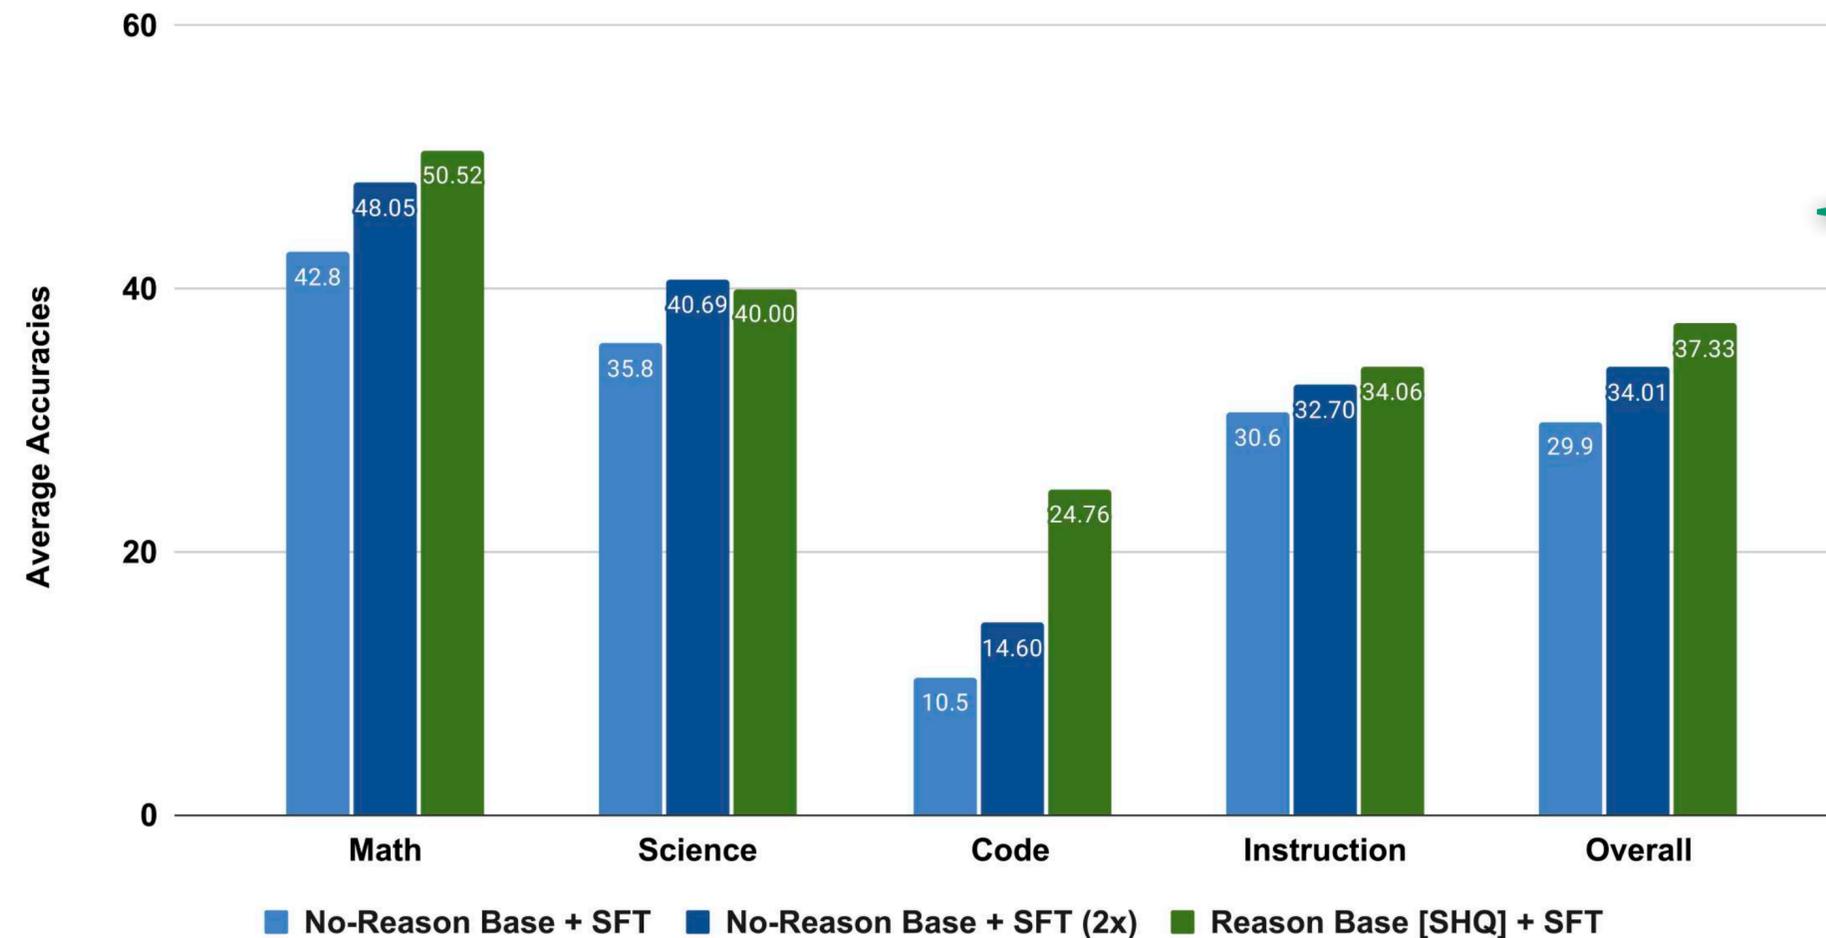


# Front Loading Reasoning

FLR systematically injects reasoning-style data at different phases of training—pretraining, SFT, RL—while varying its diversity, quantity, and quality.



RQ3: Can a model with no-reason pretraining "catch up" by getting more SFT compute?  
Lesson IV: **No!** Early reasoning builds an **irreplaceable** foundation that SFT can't replicate



Doubling SFT data for the baseline improves performance by **+4.09%**, but it still falls short of even the weakest reasoning-pretrained model (**+3.32%**).

# Front Loading Reasoning

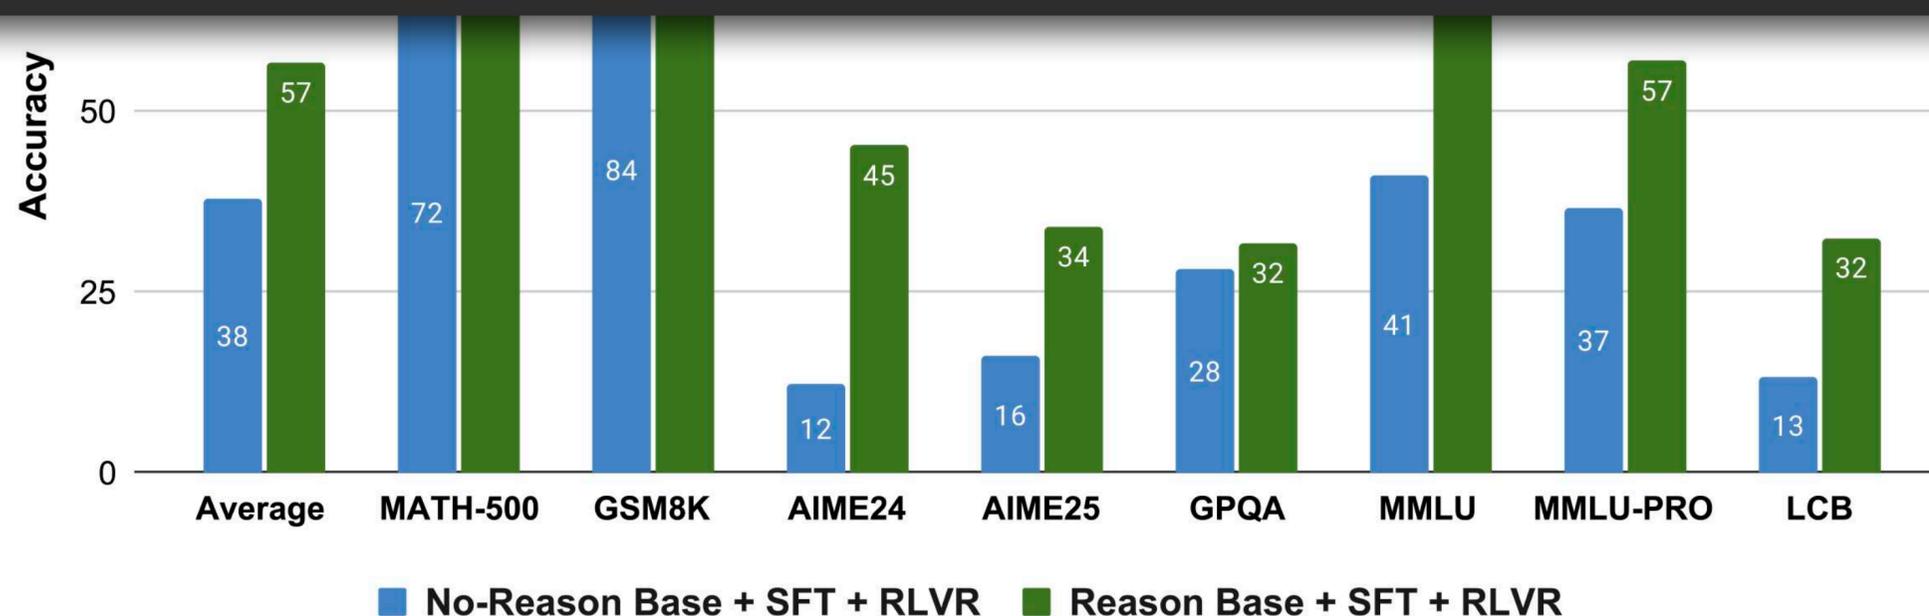
FLR systematically injects reasoning-style data at different phases of training—pretraining, SFT, RL—while varying its diversity, quantity, and quality.



RQ5: Does the gain matter after heavy RLVR?

Lesson VI: **Yes!** "Front-loading" reasoning data creates a **durable, compounding gain**

- 💡 Front-loading reasoning data into pretraining creates a durable, compounding advantage.
- 💡 High-quality pretraining data can have a latent effect unlocked by SFT
- 💡 The optimal data strategy is asymmetric: prioritize diversity in pretraining and quality in SFT.



Reason-Base model finished with a stunning **+19%** average lead. On the most difficult competition math problems (AIME), this advantage ballooned to a **+39.3%** gain.



# Concluding Remarks



# *David vs. Goliath: the Art of Scaling Intelligence in the Era of Extreme-Scale Neural Models*

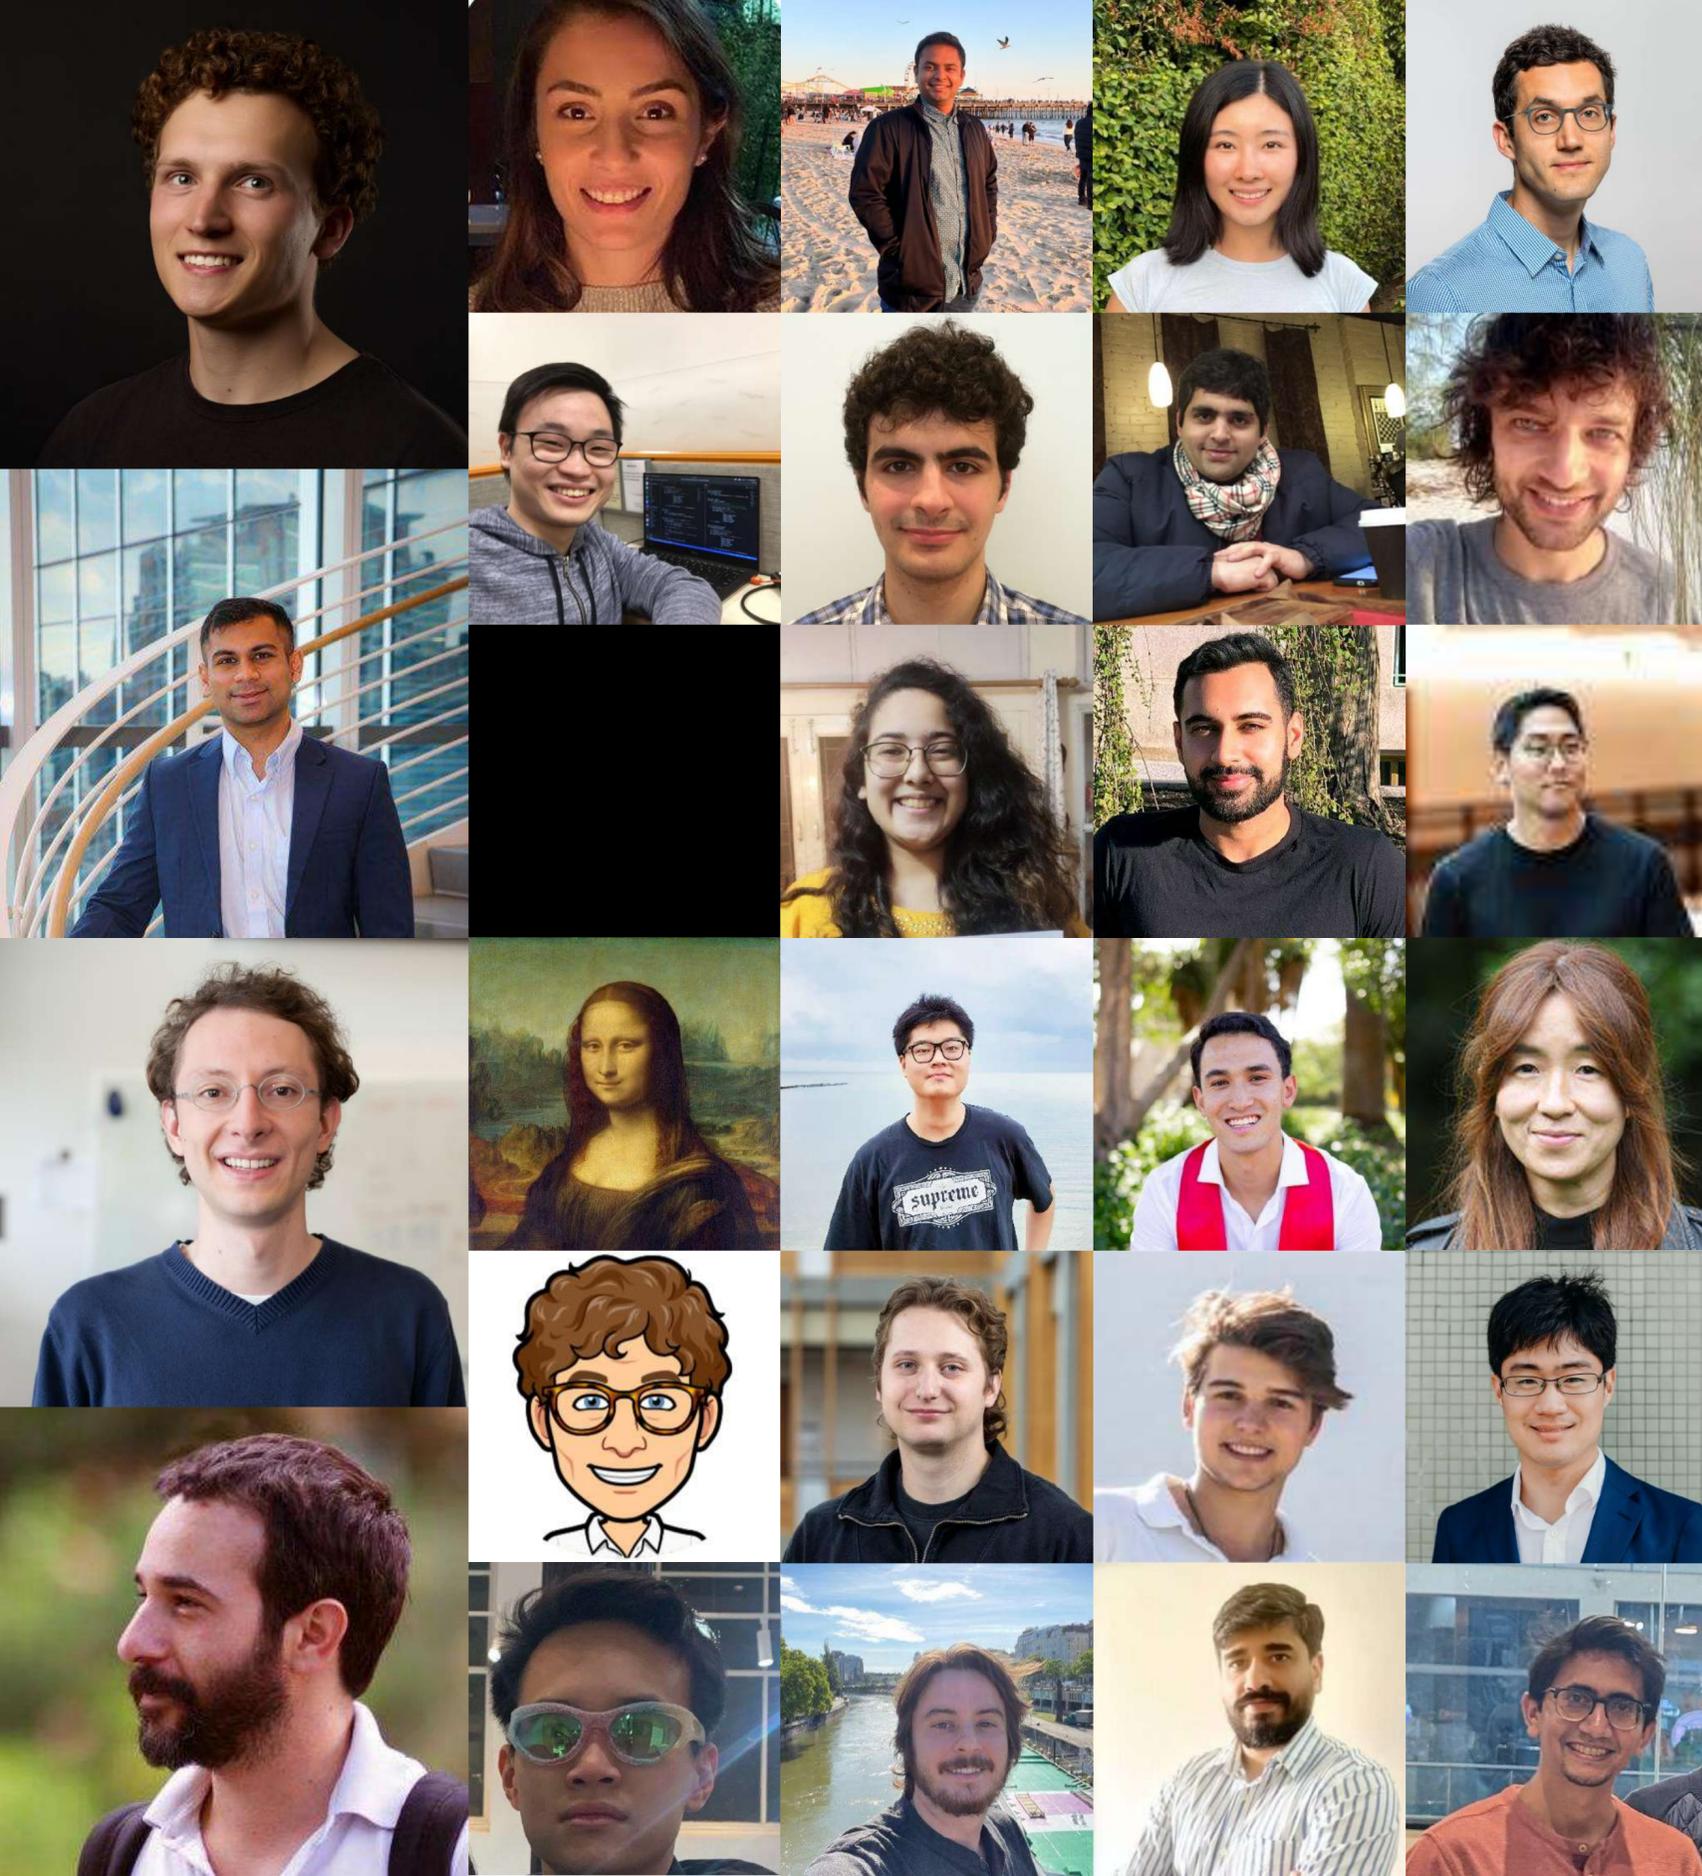
Three components to innovate:

Unconventional data 

Unconventional algorithms 

Unconventional collaboration 

- ~ implicit: open-science & open-source community
- ~ explicit: cross-institutional, cross-boarder collaborations

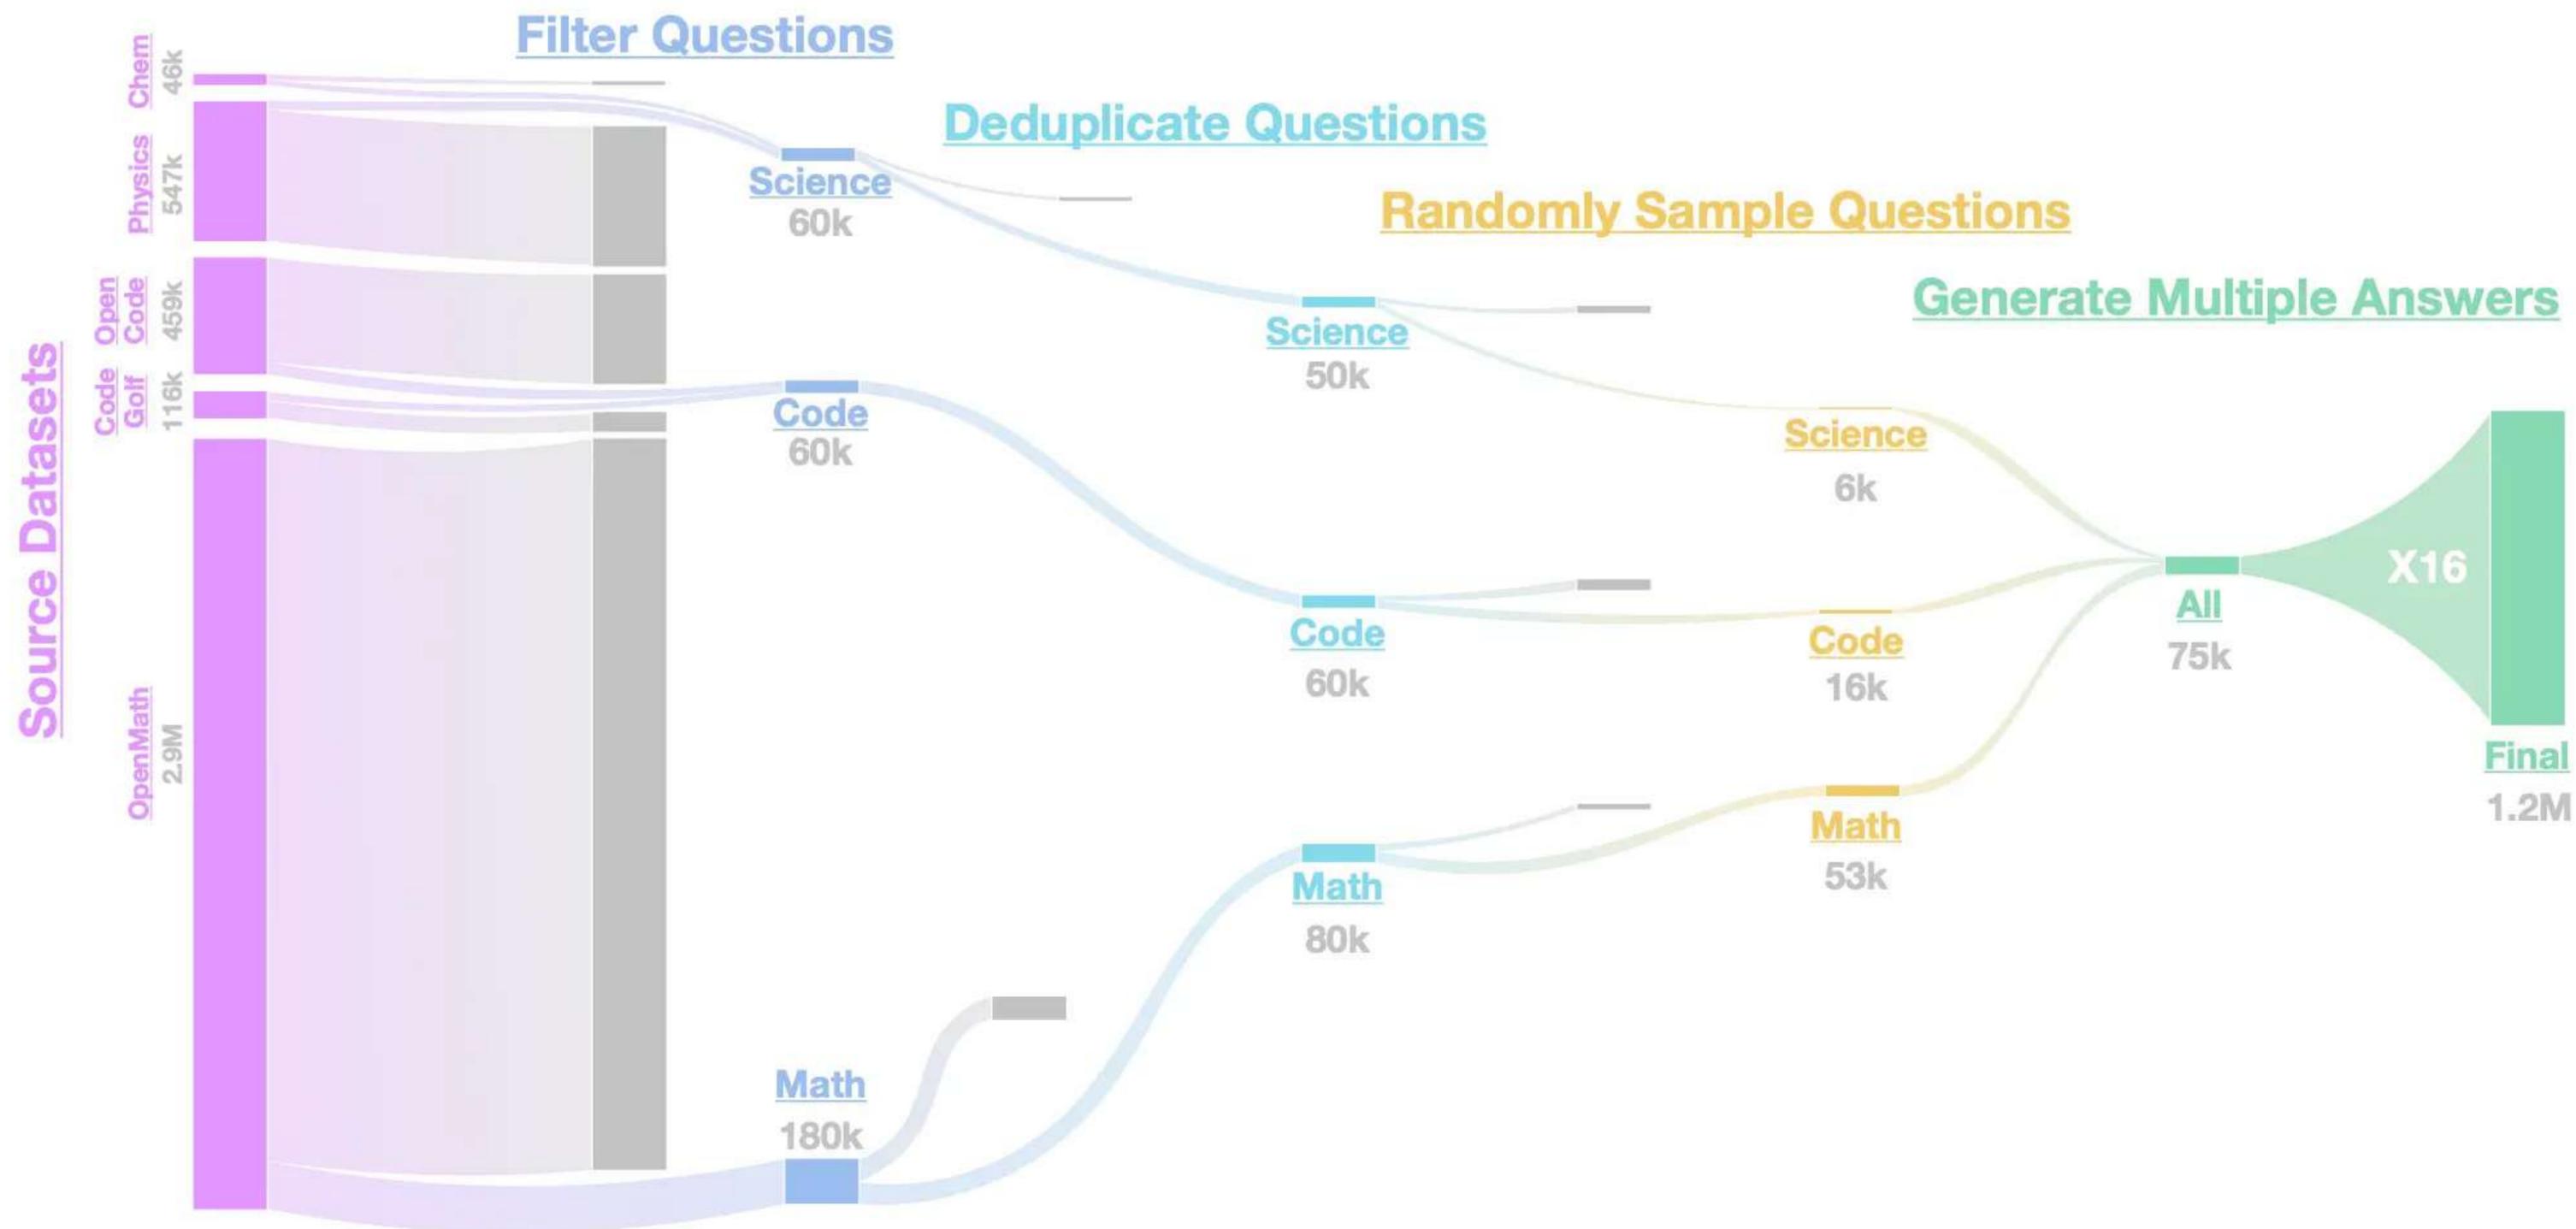


# The OpenThoughts Team

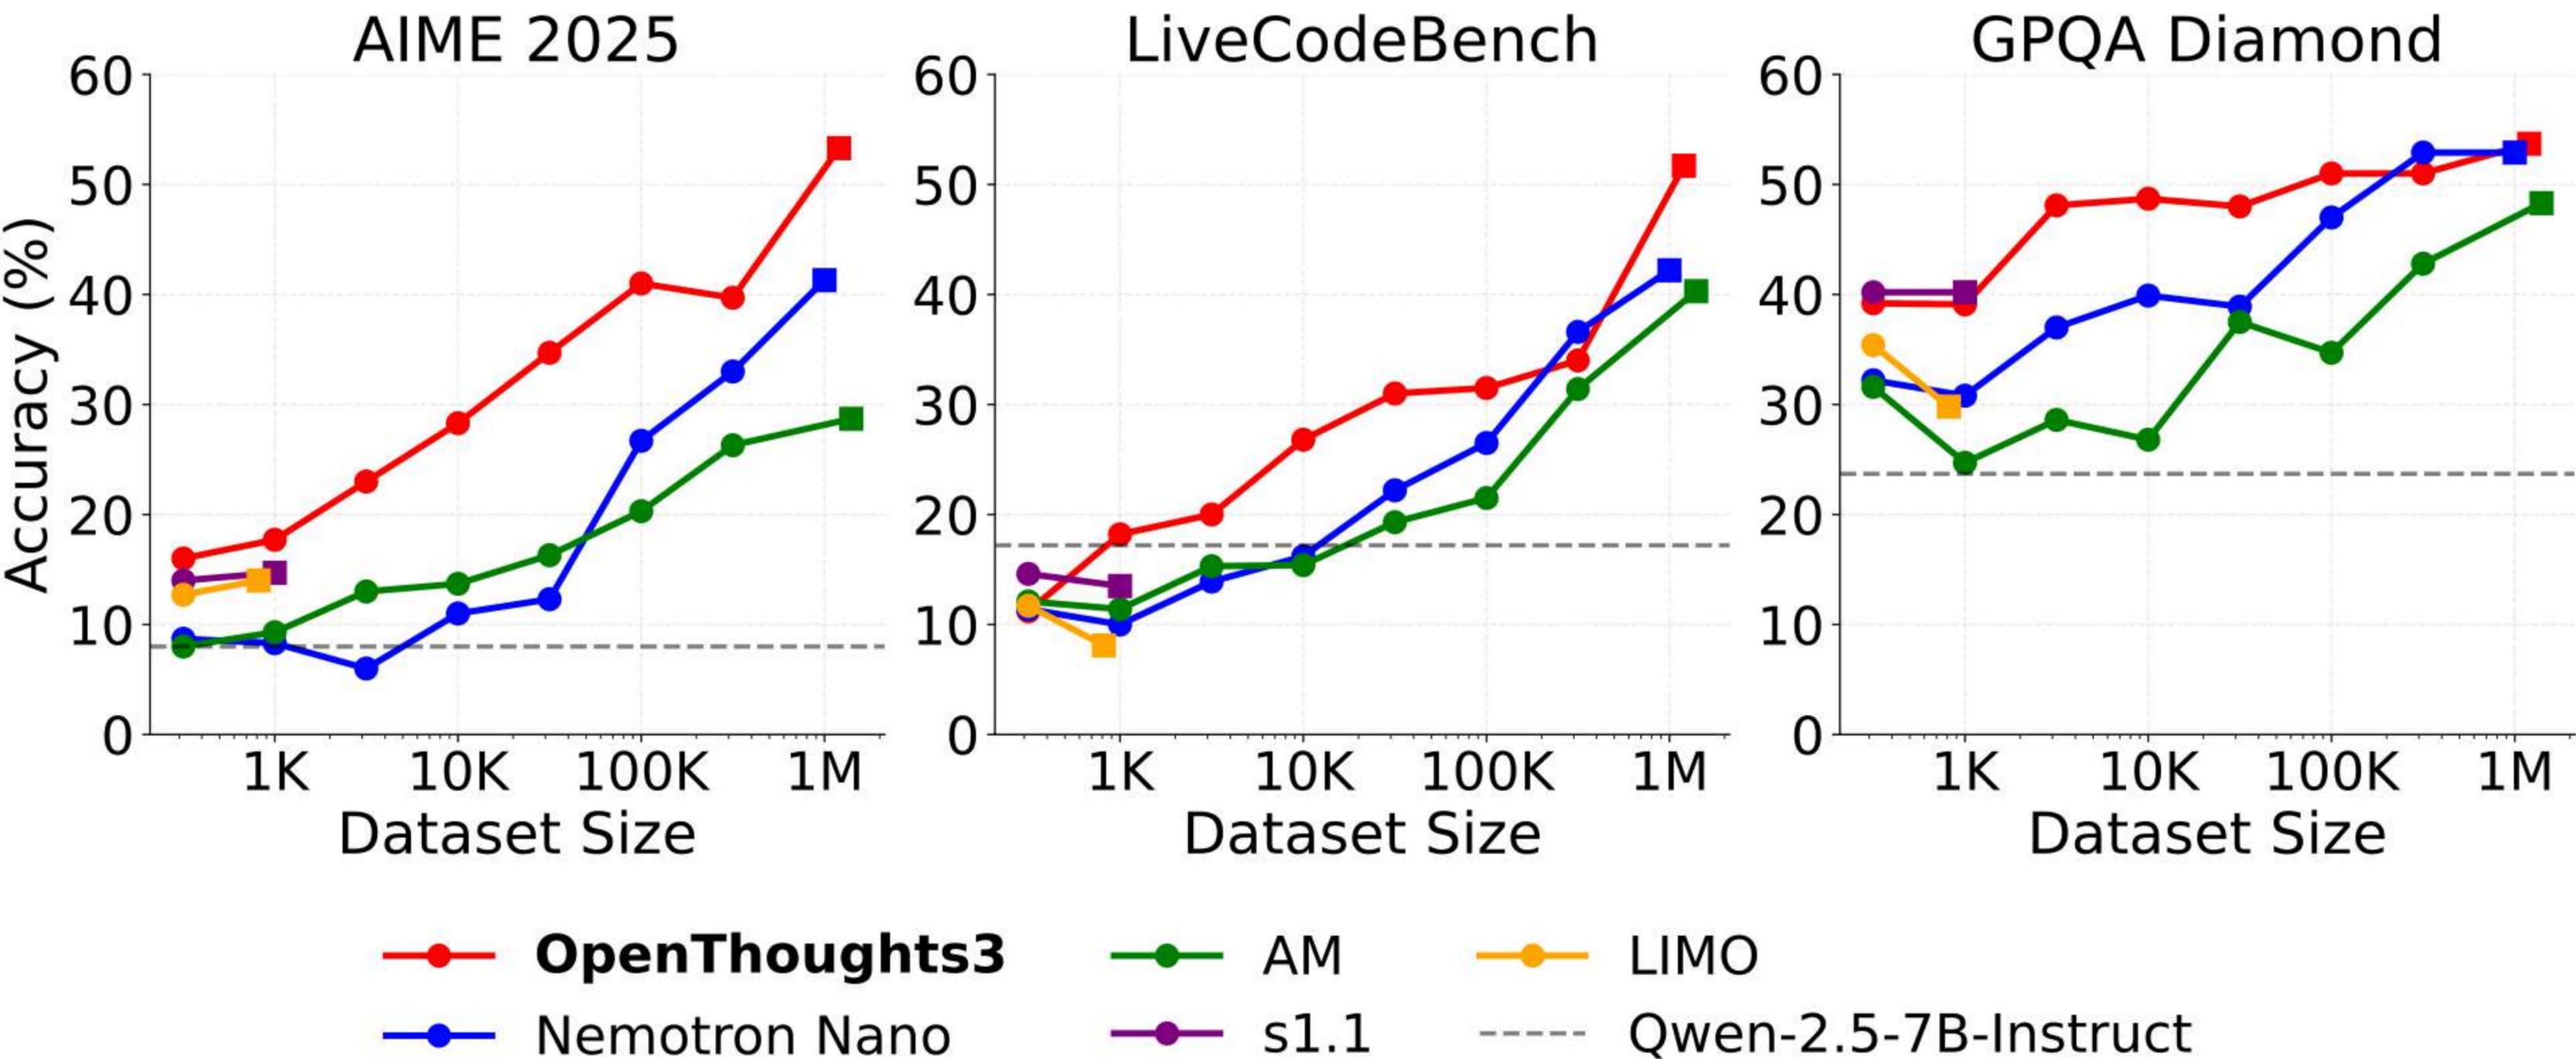
Etash Guha\*, Ryan Marten\*, Sedrick Keh\*, Negin Raoof\*, Georgios Smyrnis\*, Hritik Bansal, Marianna Nezhurina, Jean Mercat, Trung Vu, Zayne Sprague, Ashima Suvarna, Benjamin Feuer, Liangyu Chen, Zaid Khan, Eric Frankel, Sachin Grover, Caroline Choi, Niklas Muennighoff, Shiye Su, Wanjia Zhao, John Yang, Shreyas Pimpalgaonkar, Kartik Sharma, Charlie Cheng-Jie Ji, Yichuan Deng, Sarah Pratt, Vivek Ramanujan, Jon Saad-Falcon, Stutee Acharya, Jeffrey Li, Achal Dave, Alon Albalak, Kushal Arora, Blake Wulfe, Chinmay Hegde, Greg Durrett, Sewoong Oh, Mohit Bansal, Saadia Gabriel, Aditya Grover, Kai-Wei Chang, Vaishaal Shankar, Aaron Gokaslan, Mike A. Merrill, Tatsunori Hashimoto, Yejin Choi, Jenia Jitsev, Reinhard Heckel, Maheswaran Sathiamoorthy, Alexandros G. Dimakis, Ludwig Schmidt\*

And several others with ongoing contributions

# OpenThoughts3-1M



# OpenThoughts3 is the SOTA reasoning dataset recipe



Benchmark	OpenThoughts3-7B	DS-R1-Qwen-7B	NemoNano-1M	AM-1.4M	OpenR1-Distill-7B	SFT vs. RL Nemotron-Nano-8B	AceReason-7B	Skywork-7B	Base Model Qwen2.5-7B-Instruct	
Base Model		 <sup>M</sup>			 <sup>M</sup>				N/A	
Train Size	1.2M	800K	1M	1.4M	350K	3.9M	57K	119K	N/A	
Method	SFT	SFT	SFT	SFT	SFT	SFT/RL	RL	RL	N/A	
Trained by us	Yes	No	Yes	Yes	No	No	No	No	N/A	
Open Data									N/A	
Average	<b>55.3</b>	42.9	47.3	42.1	47.2	53.2	52.9	51.6	24.0	
<i>Math</i>	AIME24	<b>69.0</b>	51.3	55.0	28.3	57.7	62.0	<b>71.0</b>	<b>68.3</b>	15.0
	AMC23	<b>93.5</b>	92.0	87.0	82.2	87.0	<b>94.0</b>	<b>93.8</b>	91.0	53.0
	MATH500	90.0	88.0	86.8	87.4	88.0	89.4	89.8	<b>90.2</b>	70.8
<i>Code</i>	CodeElo	31.0	19.9	28.6	21.0	30.1	30.9	32.9	<b>37.0</b>	5.5
	LCB 05/23-05/24	64.5	48.7	58.0	54.5	37.9	<b>68.0</b>	60.5	60.4	36.2
	CodeForces	<b>32.2</b>	21.1	28.3	24.8	29.3	<b>32.9</b>	<b>30.9</b>	<b>32.5</b>	10.2
<i>Sci</i>	GPQA-D	53.7	33.2	52.9	48.3	<b>58.9</b>	52.9	52.9	50.2	24.6
	JEEBench	<b>72.4</b>	50.4	61.0	61.1	68.7	70.7	64.3	55.3	33.9
<i>Held Out</i>	HMMT 02/25	<b>42.7</b>	25.0	24.7	19.0	25.7	26.7	33.3	32.7	2.0
	HLE MCQ	10.2	12.4	2.1	9.5	12.4	12.0	10.9	10.7	<b>12.7</b>
	AIME25	<b>53.3</b>	38.0	41.3	28.7	39.7	48.0	50.7	47.3	8.0
	LCB 06/24-01/25	<b>51.7</b>	34.5	42.2	40.3	30.7	<b>50.9</b>	44.3	43.8	16.3

Super effortful SFT can win over effortful RL

# Open Thoughts

## DATA RECIPES FOR REASONING MODELS

**Etash Guha<sup>1,2</sup>, Ryan Marten<sup>3</sup>, Sedrick Keh<sup>4</sup>, Negin Raouf<sup>5</sup>, Georgios Smyrnis<sup>6</sup>,  
 Hritik Bansal<sup>7</sup>, Marianna Nezhurina<sup>8,9,16</sup>, Jean Mercat<sup>4</sup>, Trung Vu<sup>3</sup>, Zayne Sprague<sup>6</sup>,  
 Ashima Suvarna<sup>7</sup>, Benjamin Feuer<sup>10</sup>, Liangyu Chen<sup>1</sup>, Zaid Khan<sup>11</sup>, Eric Frankel<sup>2</sup>,  
 Sachin Grover<sup>12</sup>, Caroline Choi<sup>1</sup>, Niklas Muennighoff<sup>1</sup>, Shiye Su<sup>1</sup>, Wanjia Zhao<sup>1</sup>, John Yang<sup>1</sup>,  
 Shreyas Pimpalgaonkar<sup>3</sup>, Kartik Sharma<sup>3</sup>, Charlie Cheng-Jie Ji<sup>3</sup>, Yichuan Deng<sup>2</sup>,  
 Sarah Pratt<sup>2</sup>, Vivek Ramanujan<sup>2</sup>, Jon Saad-Falcon<sup>1</sup>, Jeffrey Li<sup>2</sup>, Achal Dave, Alon Albalak<sup>13</sup>,  
 Kushal Arora<sup>4</sup>, Blake Wulfe<sup>4</sup>, Chinmay Hegde<sup>10</sup>, Greg Durrett<sup>6</sup>, Sewoong Oh<sup>2</sup>,  
 Mohit Bansal<sup>11</sup>, Saadia Gabriel<sup>7</sup>, Aditya Grover<sup>7</sup>, Kai-Wei Chang<sup>7</sup>, Vaishaal Shankar,  
 Aaron Gokaslan<sup>14</sup>, Mike A. Merrill<sup>1</sup>, Tatsunori Hashimoto<sup>1</sup>, Yejin Choi<sup>1</sup>,  
 Jenia Jitsev<sup>8,9,16</sup>, Reinhard Heckel<sup>15</sup>, Maheswaran Sathiamoorthy<sup>3</sup>,  
 Alexandros G. Dimakis<sup>3,5</sup>, Ludwig Schmidt<sup>1</sup>**

<sup>1</sup>Stanford University, <sup>2</sup>University of Washington, <sup>3</sup>BespokeLabs.ai, <sup>4</sup>Toyota Research Institute  
<sup>5</sup>UC Berkeley, <sup>6</sup>UT Austin, <sup>7</sup>UCLA, <sup>8</sup>JSC, <sup>9</sup>LAION, <sup>10</sup>NYU, <sup>11</sup>UNC Chapel Hill  
<sup>12</sup>ASU, <sup>13</sup>Lila Sciences, <sup>14</sup>Cornell Tech <sup>15</sup>TUM <sup>16</sup>Open-Ψ (Open-Sci) Collective

	on-Nano-8B	AceReason-7B	Skywork-7B	Base Model	Qwen2.5-7B-Instruct
	57K	119K			
L	RL	RL			
	No	No			
	52.9	51.6			24.0
	<b>71.0</b>	<b>68.3</b>			15.0
	<b>93.8</b>	91.0			53.0
	89.8	<b>90.2</b>			70.8
	32.9	<b>37.0</b>			5.5
	60.5	60.4			36.2
	22.9	22.5			19.9

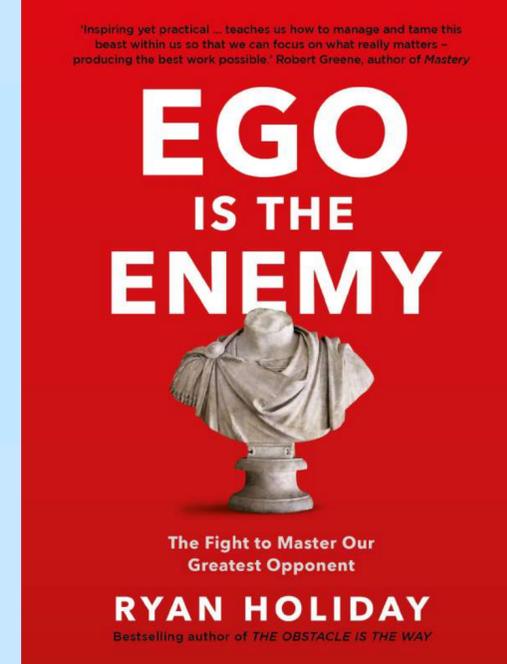
Unconventional collaboration wins!

	HMMT 02/25	42.7	25.0	24.7	19.0	25.7	48.0	50.7	47.5	8.0
<i>Held Out</i>	HLE MCQ	10.2	12.4	2.1	9.5	12.4				
	AIME25	<b>53.3</b>	38.0	41.3	28.7	39.7				
	LCB 06/24-01/25	<b>51.7</b>	34.5	42.2	40.3	30.7	<b>50.9</b>	44.3	43.8	16.3

Super effortful SFT can win over effortful RL

# Concluding Thoughts

- Nothing is easy in life
- No pain no gain
- But everything is doable with efforts
- The power of collaboration
- Ego is the enemy



Three components to innovate:

- Unconventional data 🔥
- Unconventional algorithms 🚀
- Unconventional collaboration 🌍

Chemistry between the base LLM and RL matters

Conclusions from **effortless** RL  $\neq$  **effortful** RL

Conclusions from **effortless** SFT  $\neq$  **effortful** SFT



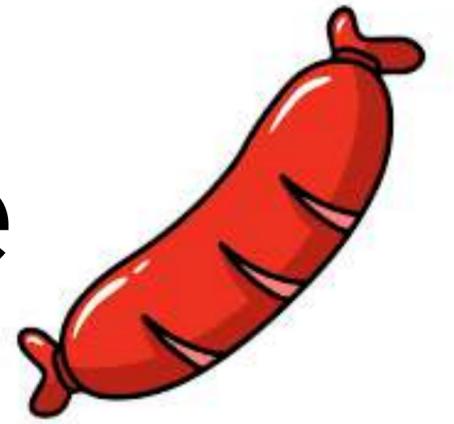


## NVIDIA LEADS OPEN SOURCE AI MOMENTUM AS CHINESE LABS CLOSE IN



1. Small models are in high demand, with daily downloads hitting tens of millions (!!! 🔥) for top repositories
2. China's contribution to the open source ecosystem has skyrocketed, with Alibaba now surpassing Hugging Face 😲
3. NVIDIA has emerged as a rising star ★ (or a risen star? 😎), for open-source contributions ❤️

# LLM 101: How sausages are made



"Fine-tuning"

SFT

RL

Pre-

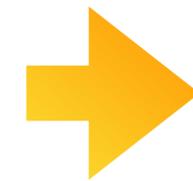
training

(On the internet data)



sequential  
fine-tuning

(On curated exam-style data)



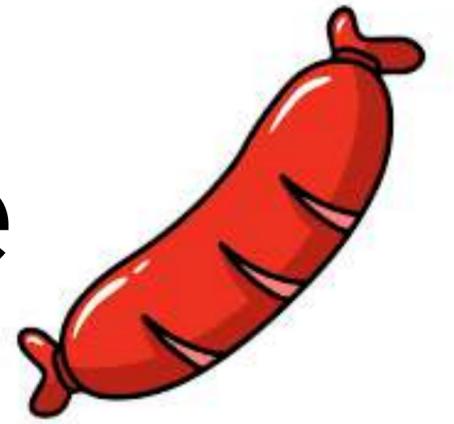
Reinforcement  
learning

(On curated exam-style data)

Imitation learning

Exploration learning

# LLM 101: How sausages are made



"Fine-tuning"

It's all about data

Use entirely of the internet data

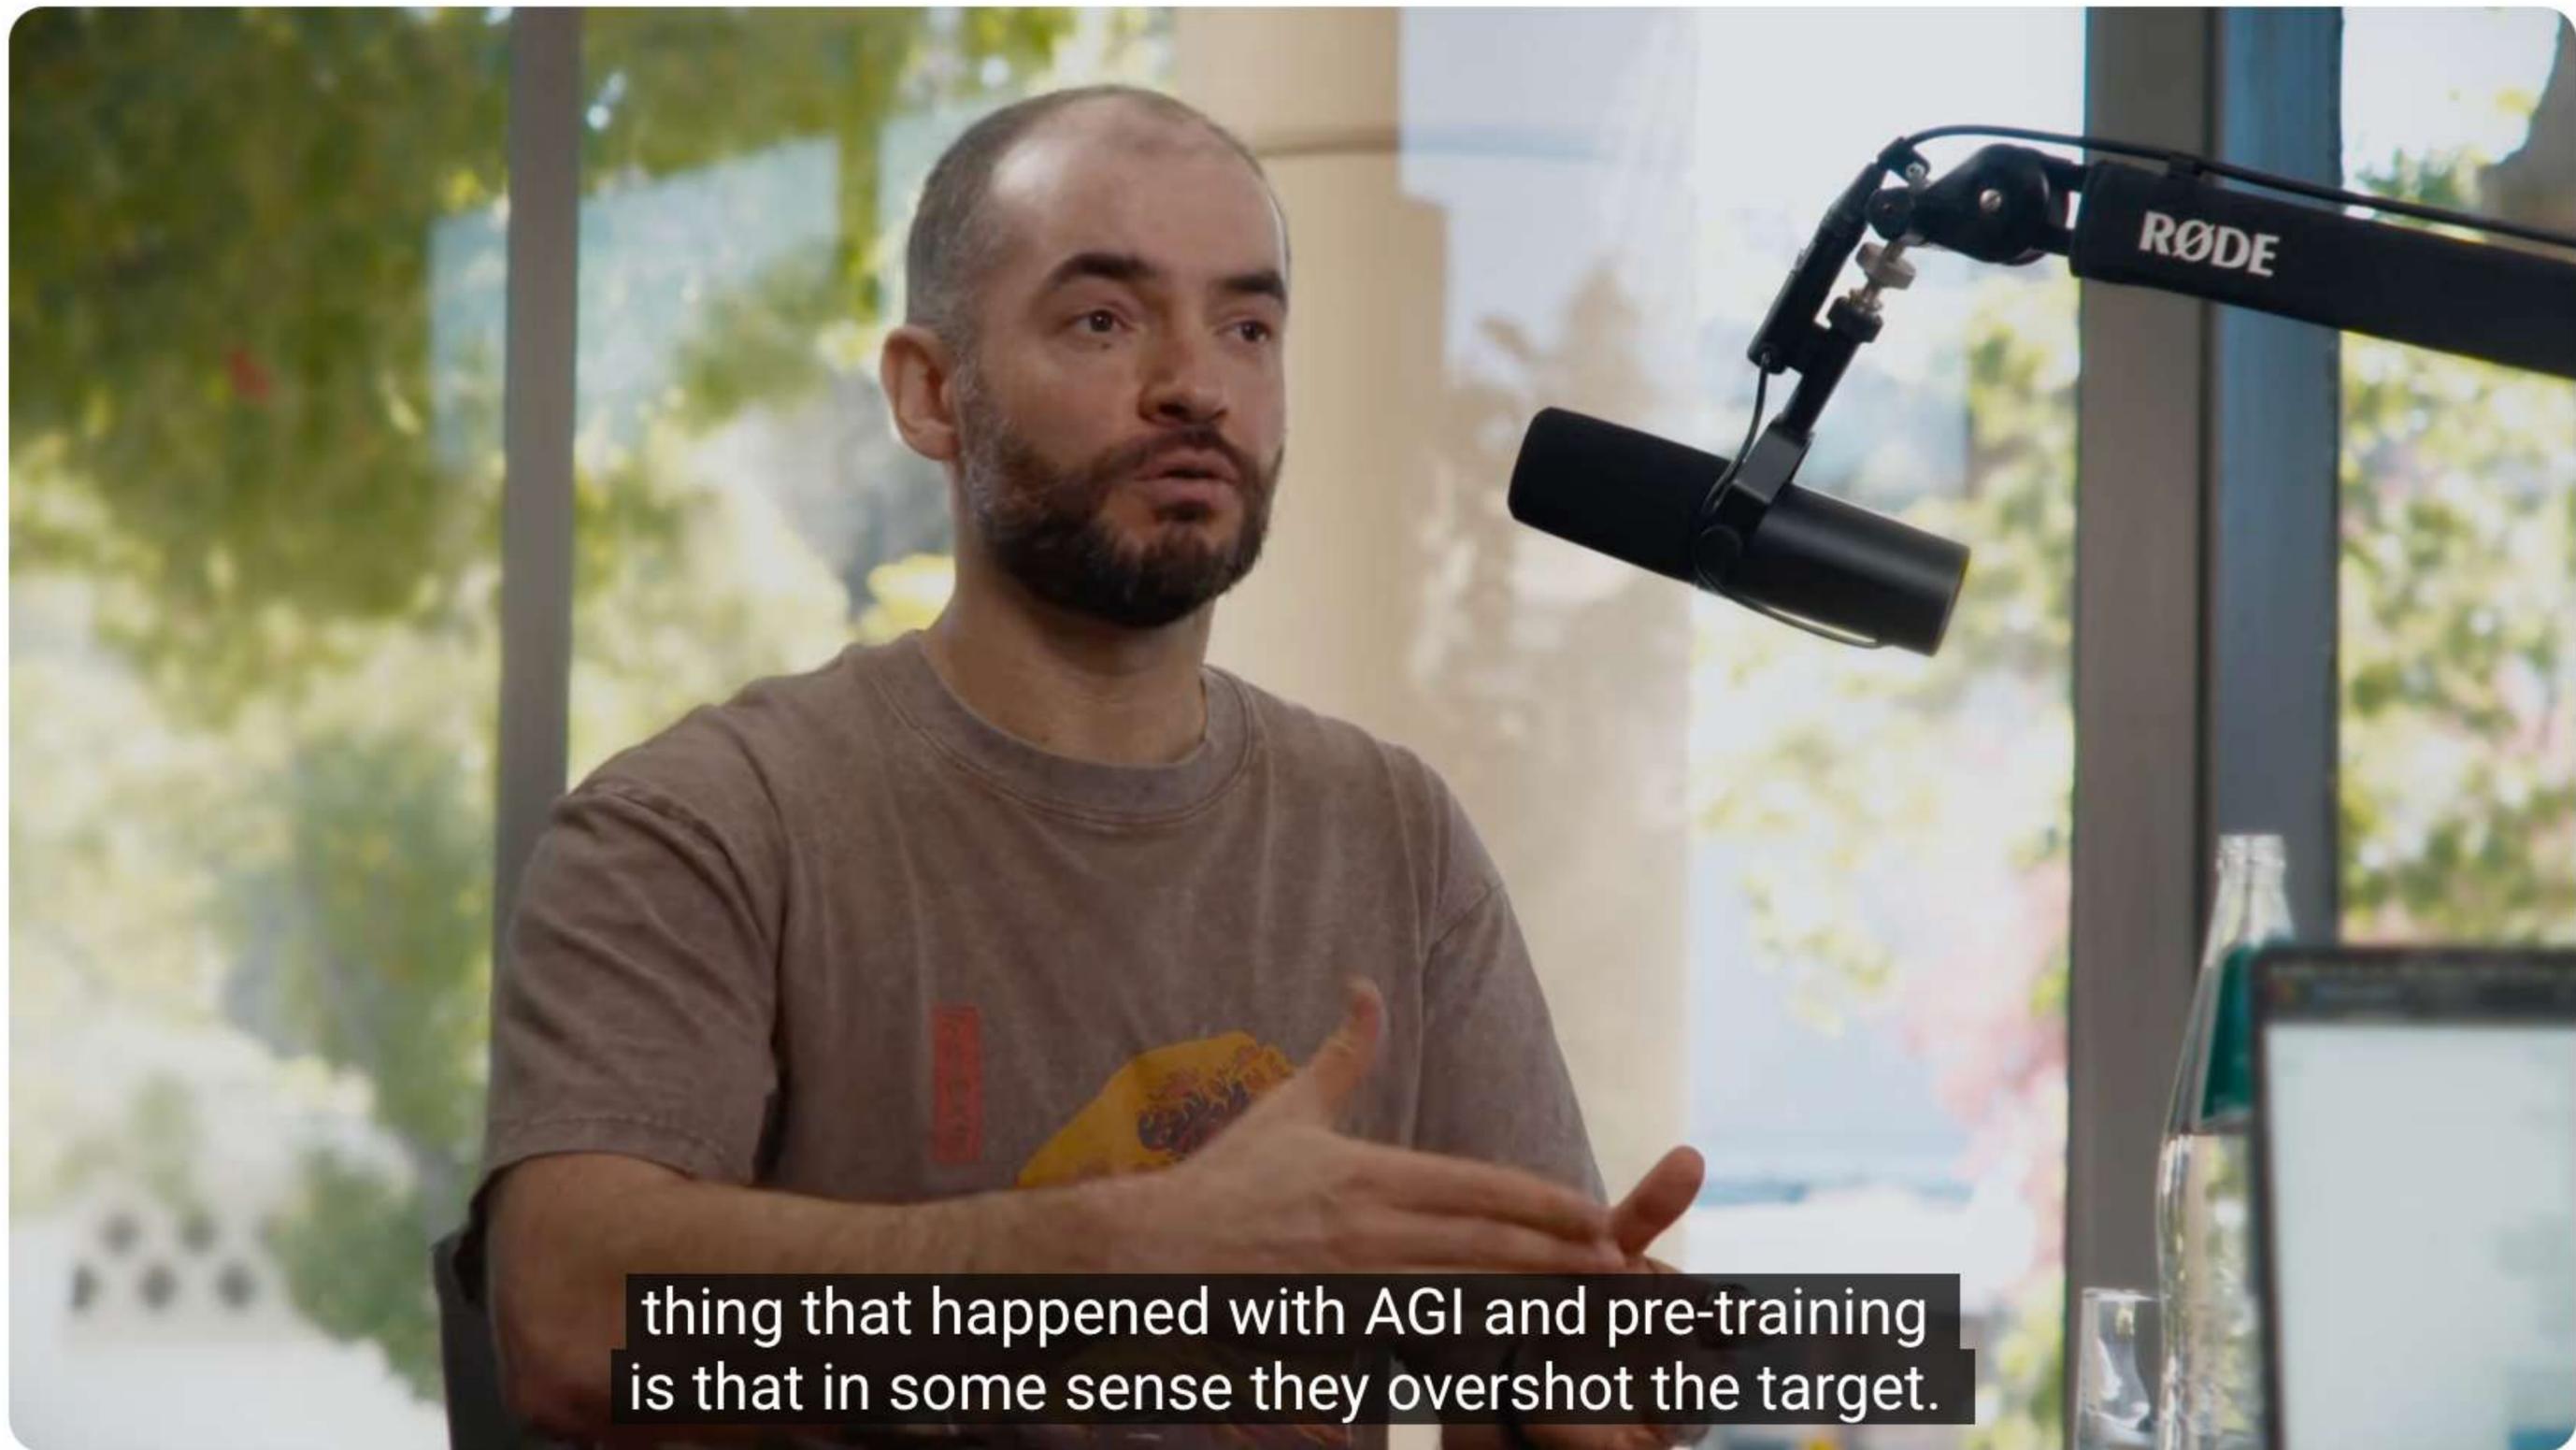
Internet data not enough, so ask humans to write exam data

When even that's not enough, ask AI to synthesize data

=> Make OOD data as in-dist data via brute-force data synth!

Imitation learning

Exploration learning



## Ilya Sutskever – We're moving from the age of scaling to the age of research



**Dwarkesh Patel**  
1.11M subscribers

**Subscribe**

24K |

Share

Ask

Download





Supervised learning is not something that happens in nature.

## Richard Sutton – Father of RL thinks LLMs are a dead end



**Dwarkesh Patel**  
1.11M subscribers

**Subscribe**

 13K | 

 Share

 Ask

 Download





## Andrej Karpathy — “We’re summoning ghosts, not building animals”



**Dwarkesh Patel**  
1.11M subscribers

**Subscribe**

 23K



 Share

 Ask

 Download



# The Universe of Knowledge

= The Universe of Synthetic Data

Knowledge in the **Internet** Data



Knowledge in the (Conventional) **Distillation**

Knowledge in the **Human/Experts'** Annotation

Knowledge in the **Extreme Scaling** of Reasoning

Knowledge in the **Unconventional** Simulation

Discovery of new knowledge is the act of a scientist, even if the target domain is not for scientific fields per say

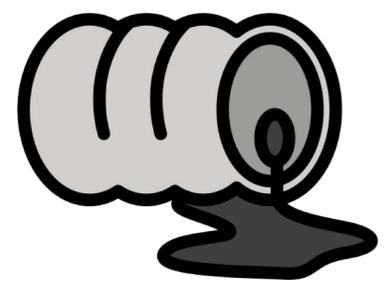
# Universe of Knowledge

Many real-world use cases require the capability to fill in this data gap

= The Universe of Synthetic Data

Those with a stronger synthetic data flywheel will win. It will be increasingly reasoning & compute-heavy, and often not be open.

Knowledge in the **Internet** Data



Knowledge in the (Conventional) **Distillation**

Knowledge in the **Human/Experts'** Annotation

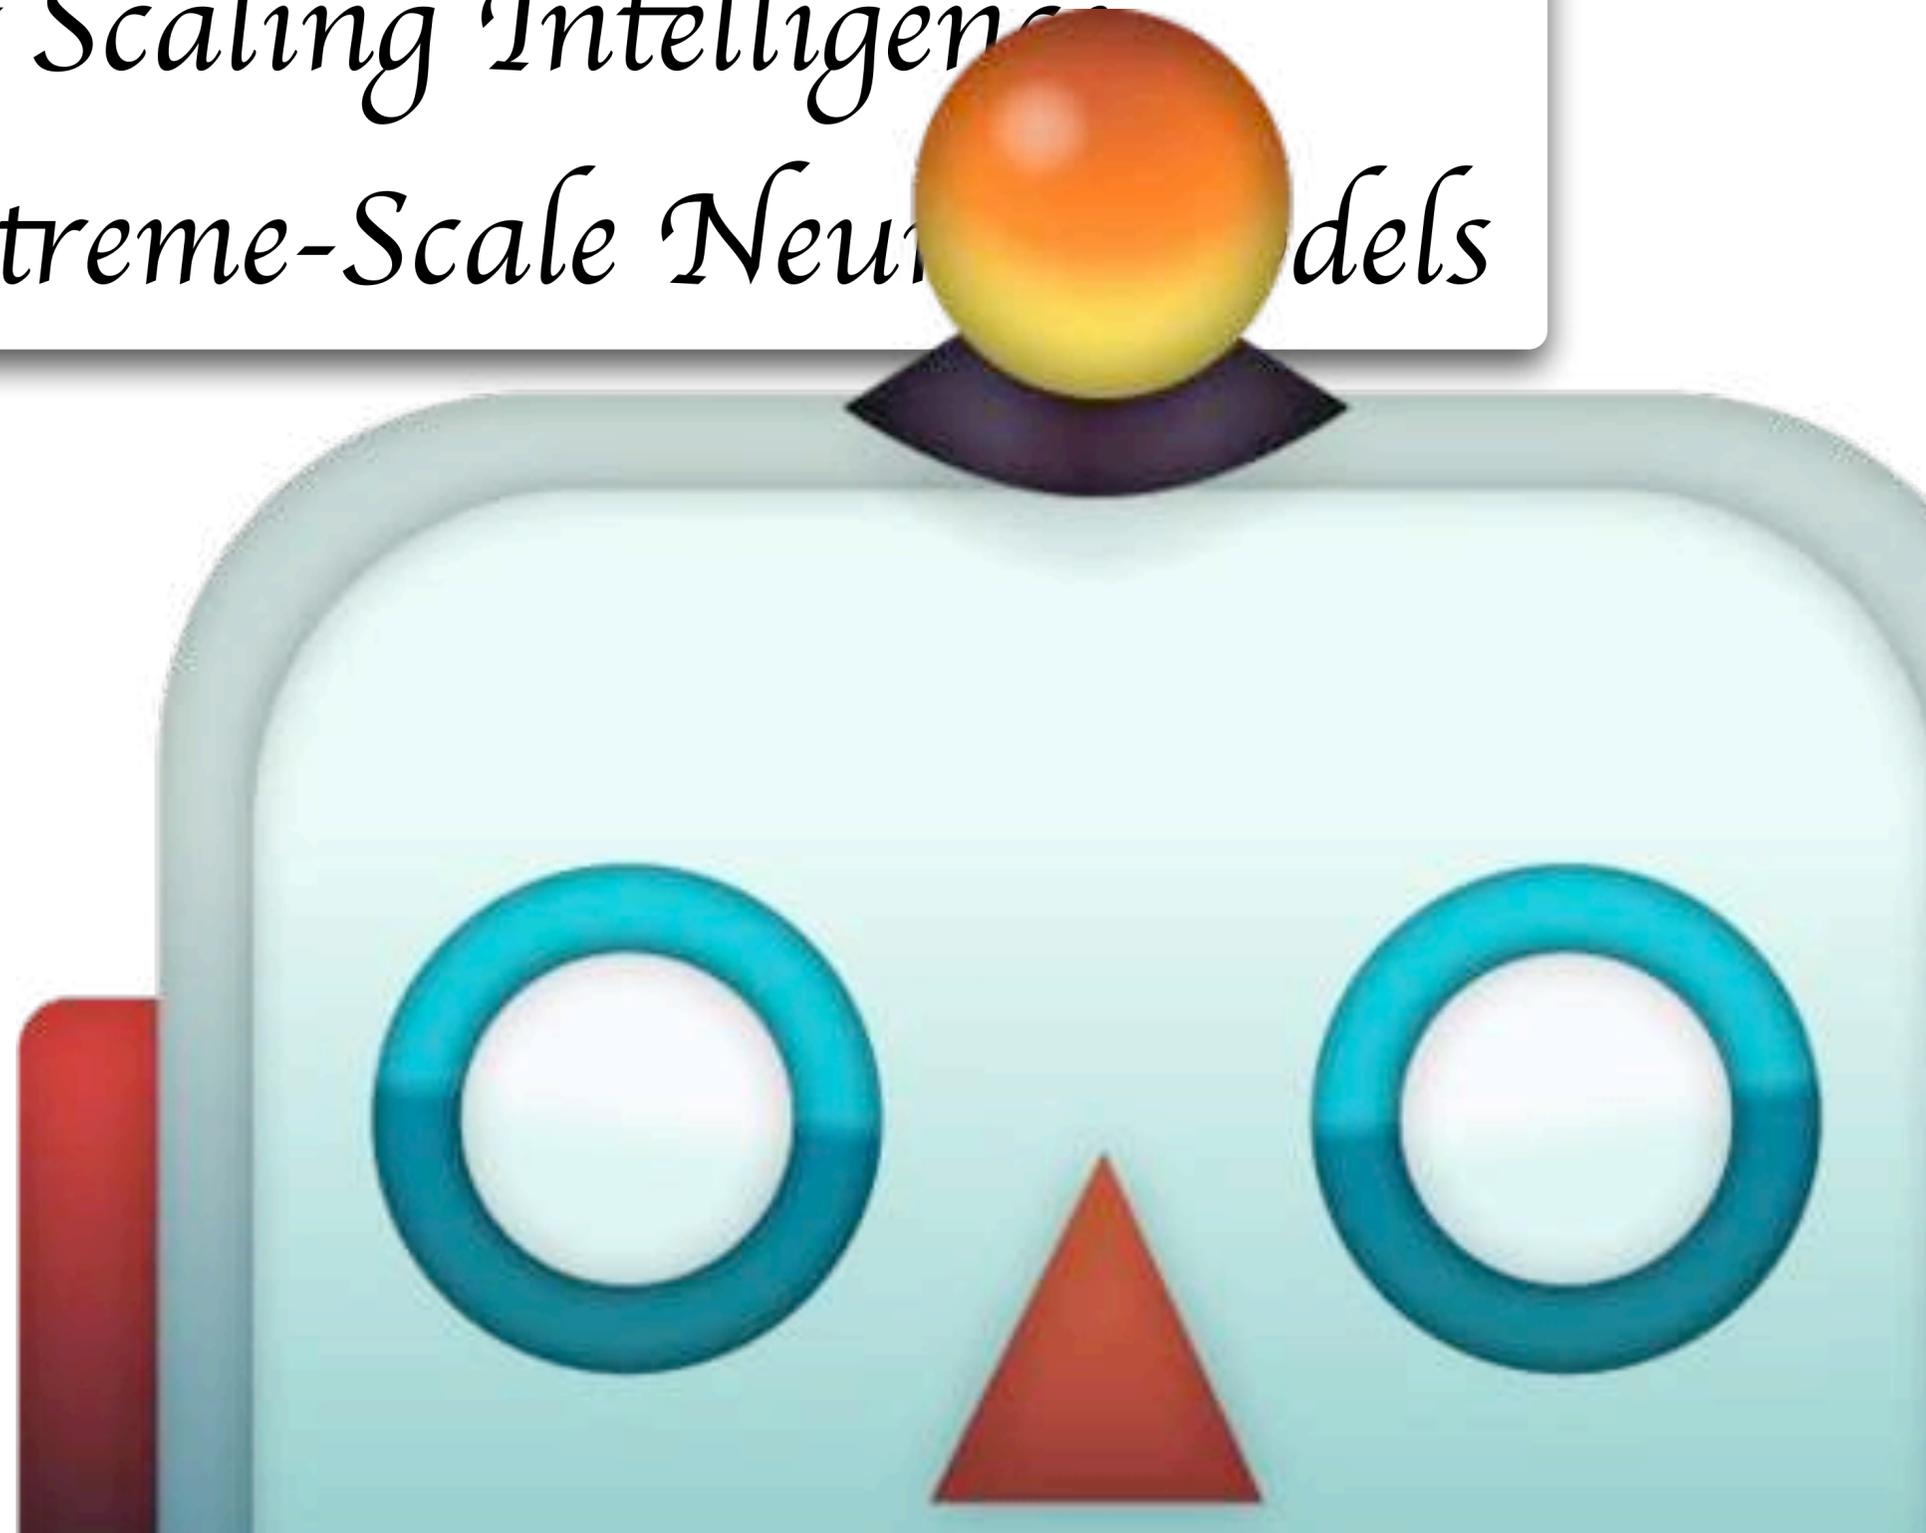
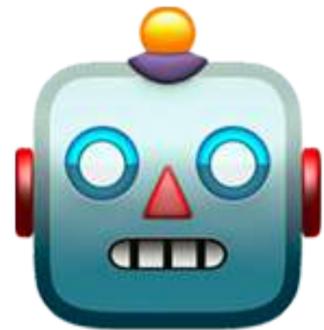


Knowledge in the **Extreme Scaling of Reasoning**



Knowledge in the **Unconventional Simulation**

*David vs. Goliath:  
the Art of Scaling Intelligence  
in the Era of Extreme-Scale Neural Models*



*David vs. Goliath:  
the Art of Scaling Intelligence  
in the Era of Extreme-Scale Neural Models*





AI mirrors human intelligence:

- Reasoning is often memorized knowledge
- Exploitation vs exploration trade offs
- Diverse examples/experiences enhance learning

AI diverges from human intelligence:

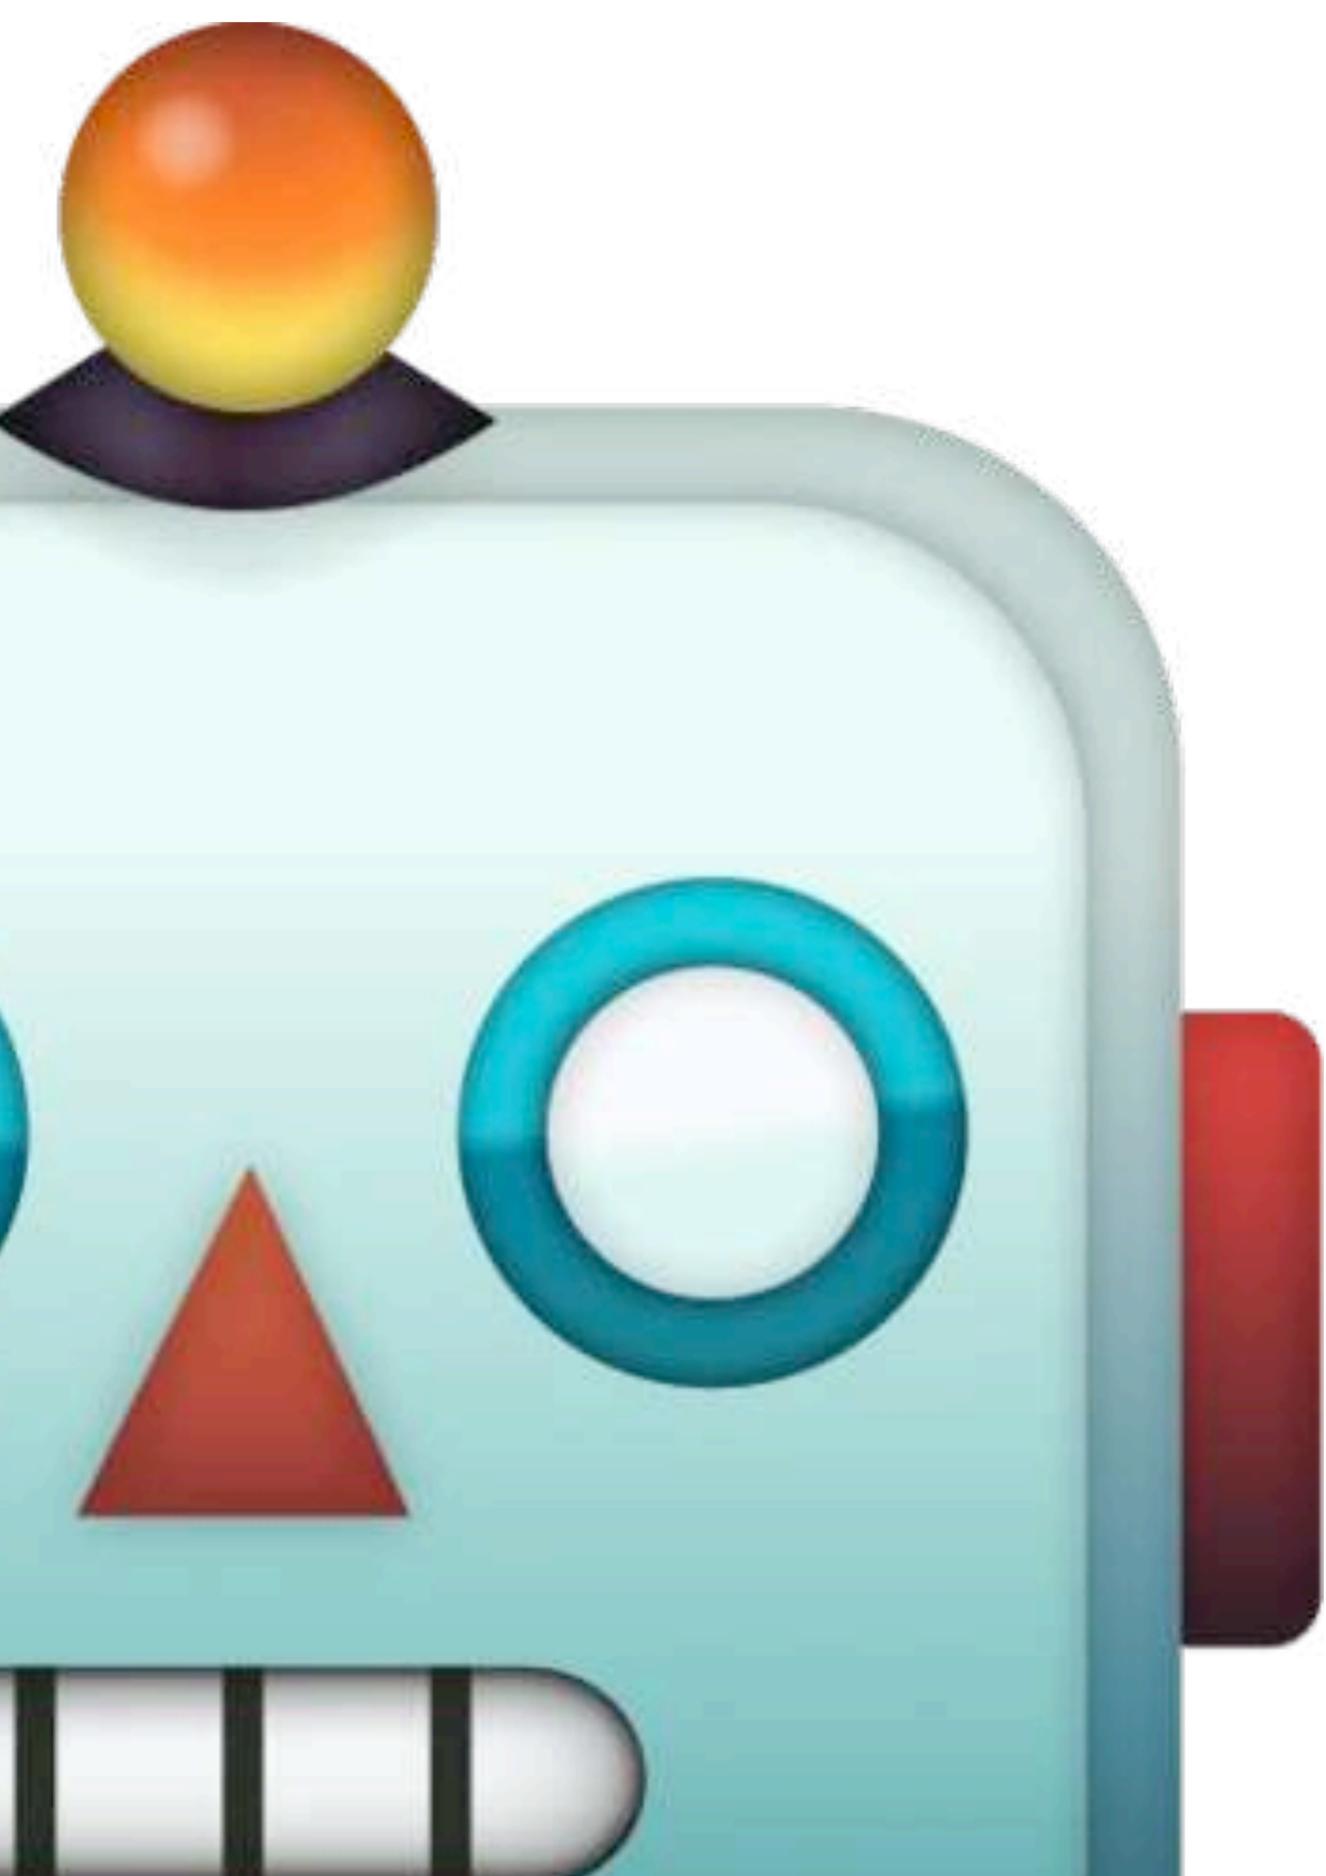
- More data, compute, and memory
- Less abstraction and conceptualization
- Clear separation between training and testing





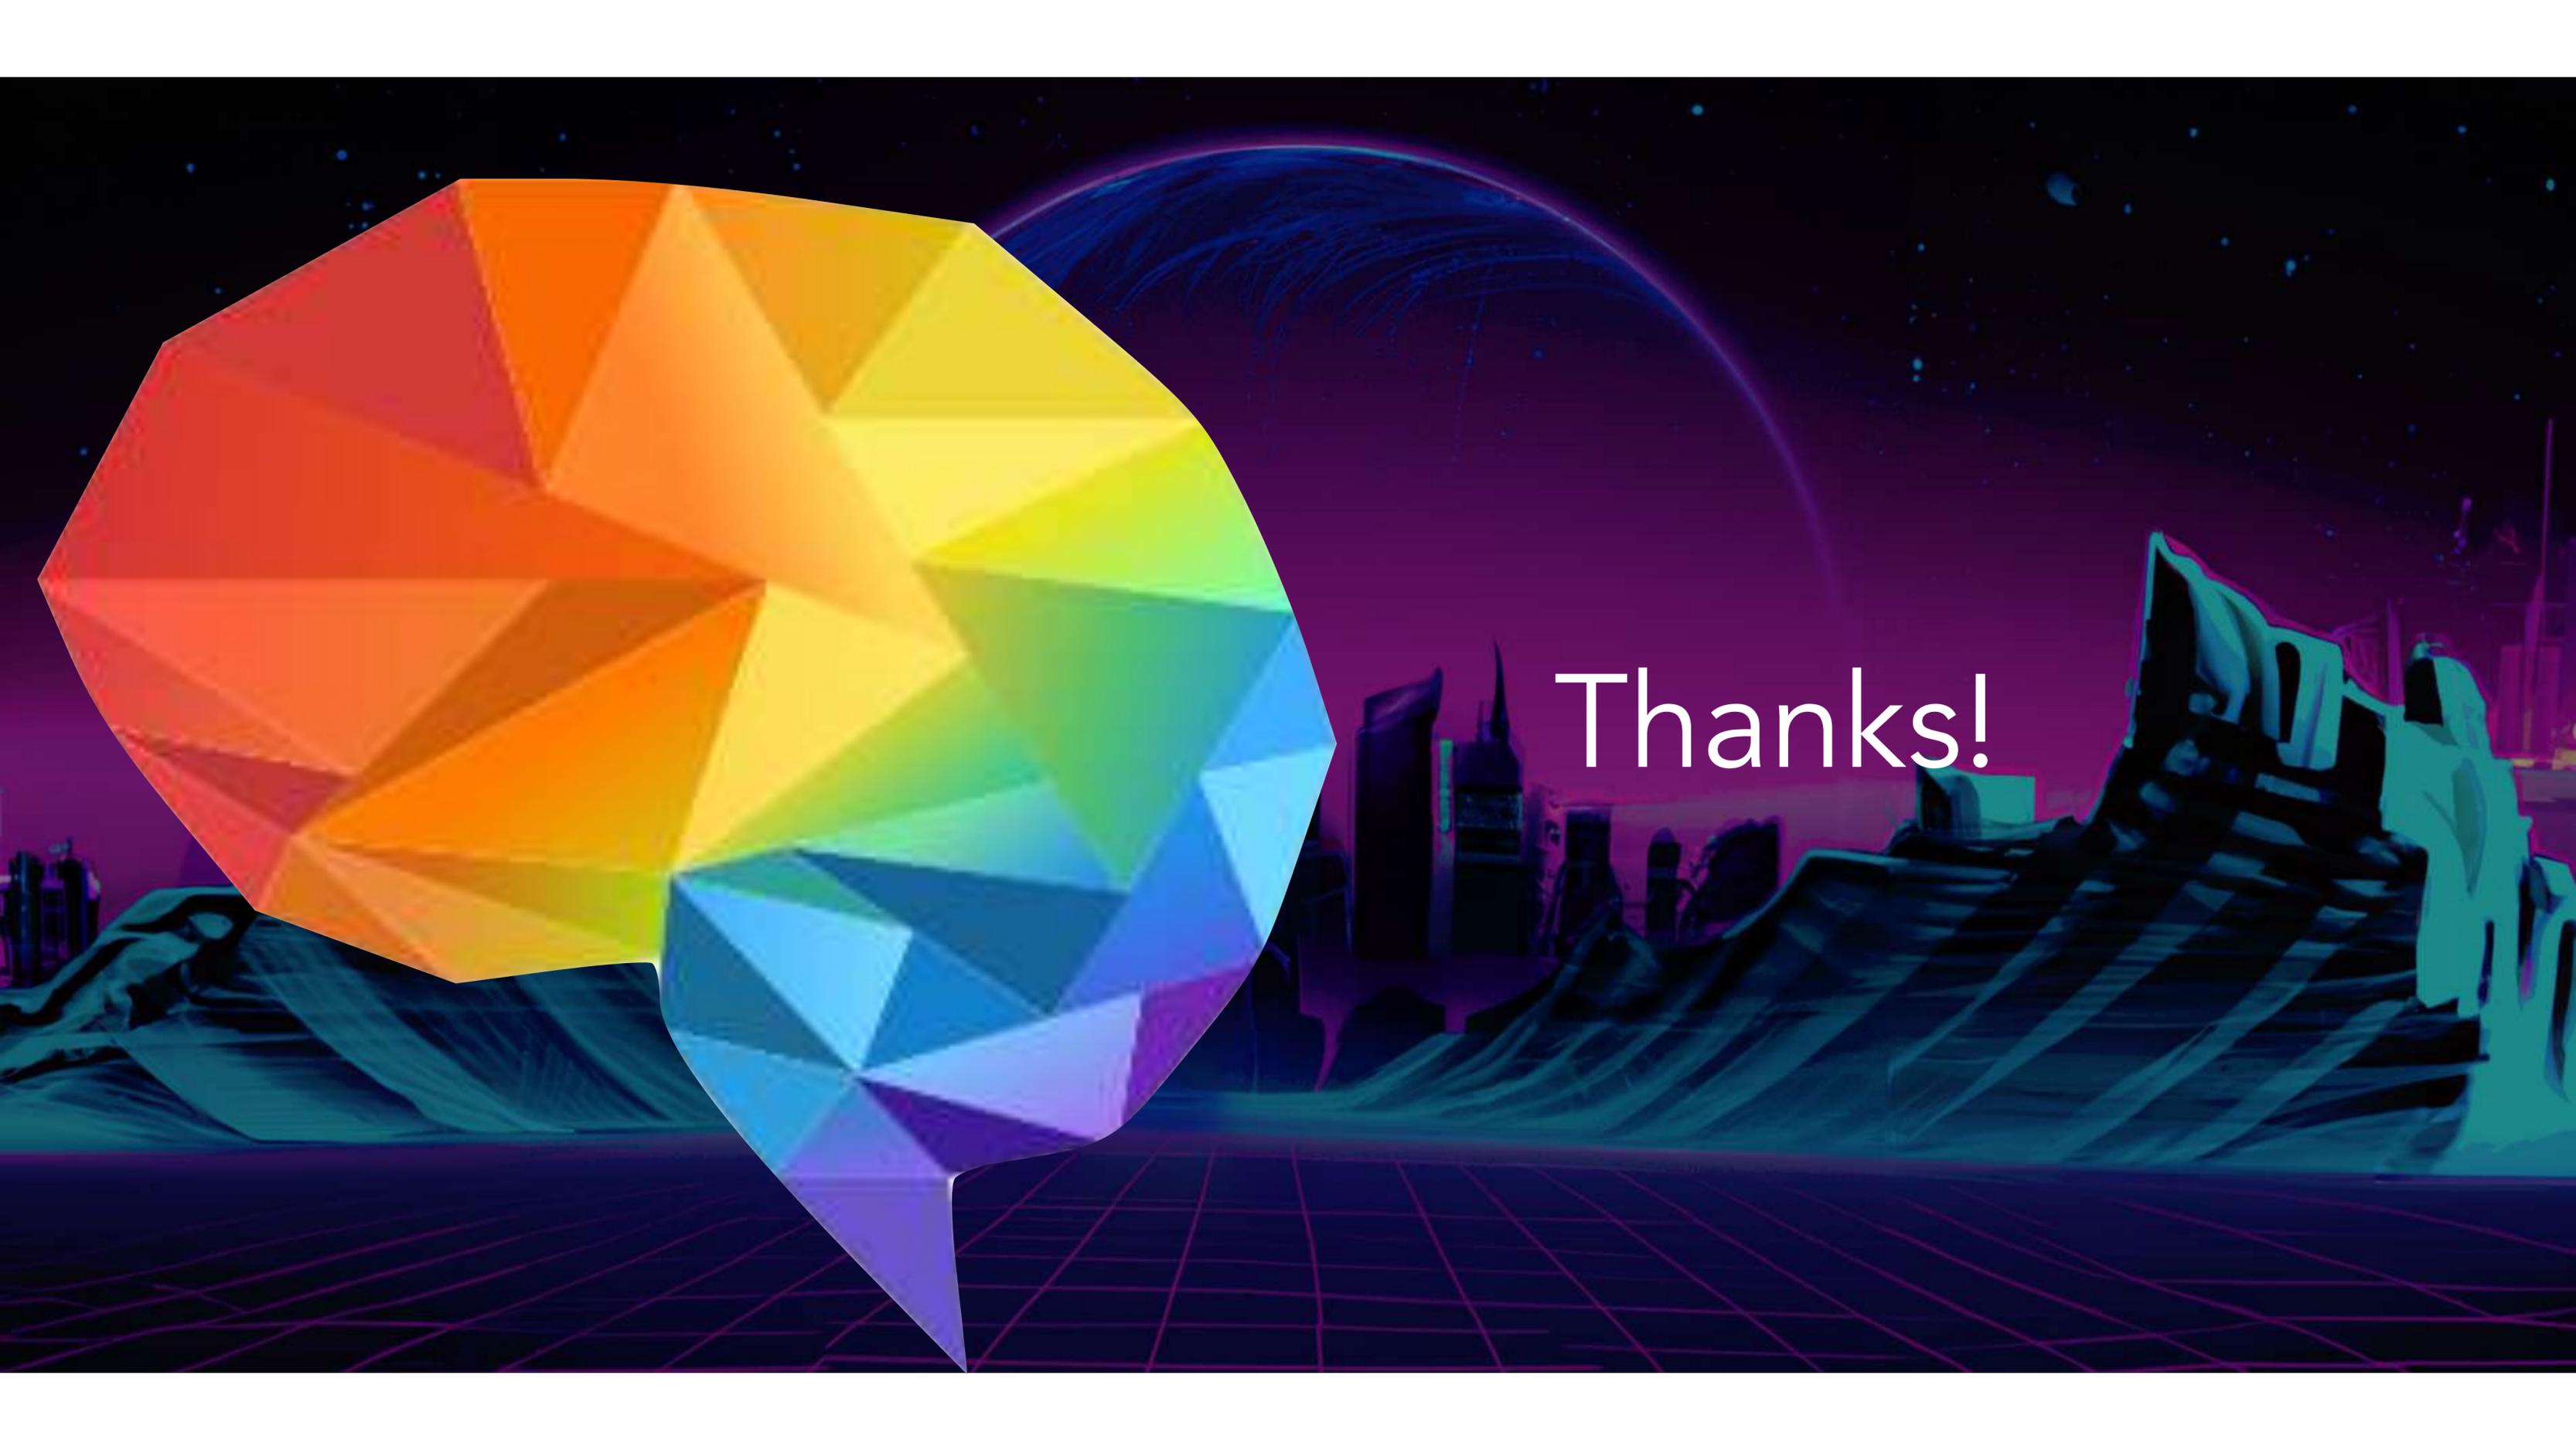
How sausages are made:

-  Human's internet data, the artifact of human intelligence
-  Human's audacity to throw in \$\$\$\$\$\$\$
-  Human's annotations at scale for alignment
-  Human's intuitions and insights on a lot of engineering details



## Open Research Questions

- New theories of intelligences?
- New theories of knowledge and reasoning?
- That humans do not have a context window size of 1M tokens — is this a limitation or blessing?



Thanks!