Adaptation of LLMs

https://adapt-llm.github.io/



Zixuan Ke **Yifei Ming** **Shafiq Joty**

Acknowledgement: Austin Xu

Minimal LLM Basics

Prerequisites

Training ML Models

- **Learning algorithms related:**
  - SGD, Learning rate, AdamW, Batch size
- **Model architecture related:**
  - Cross and Self Attentions
  - Encoder-Decoder
  - Transformers

Basic LLM concepts

- Transformer decoder
- Next token prediction
- Tokenization, sequence/context length
- In-context learning:
  - Zero- and few-shot learning

This Tutorial

Goals

Build Foundational understanding for LLM Adaptation

- Evaluation methods
- Key concepts of LLM adaptation
- Key techniques for LLM adaptation
  - Data perspective
  - Model perspective
- Key trends

Table of contents

LLM vs. human performance

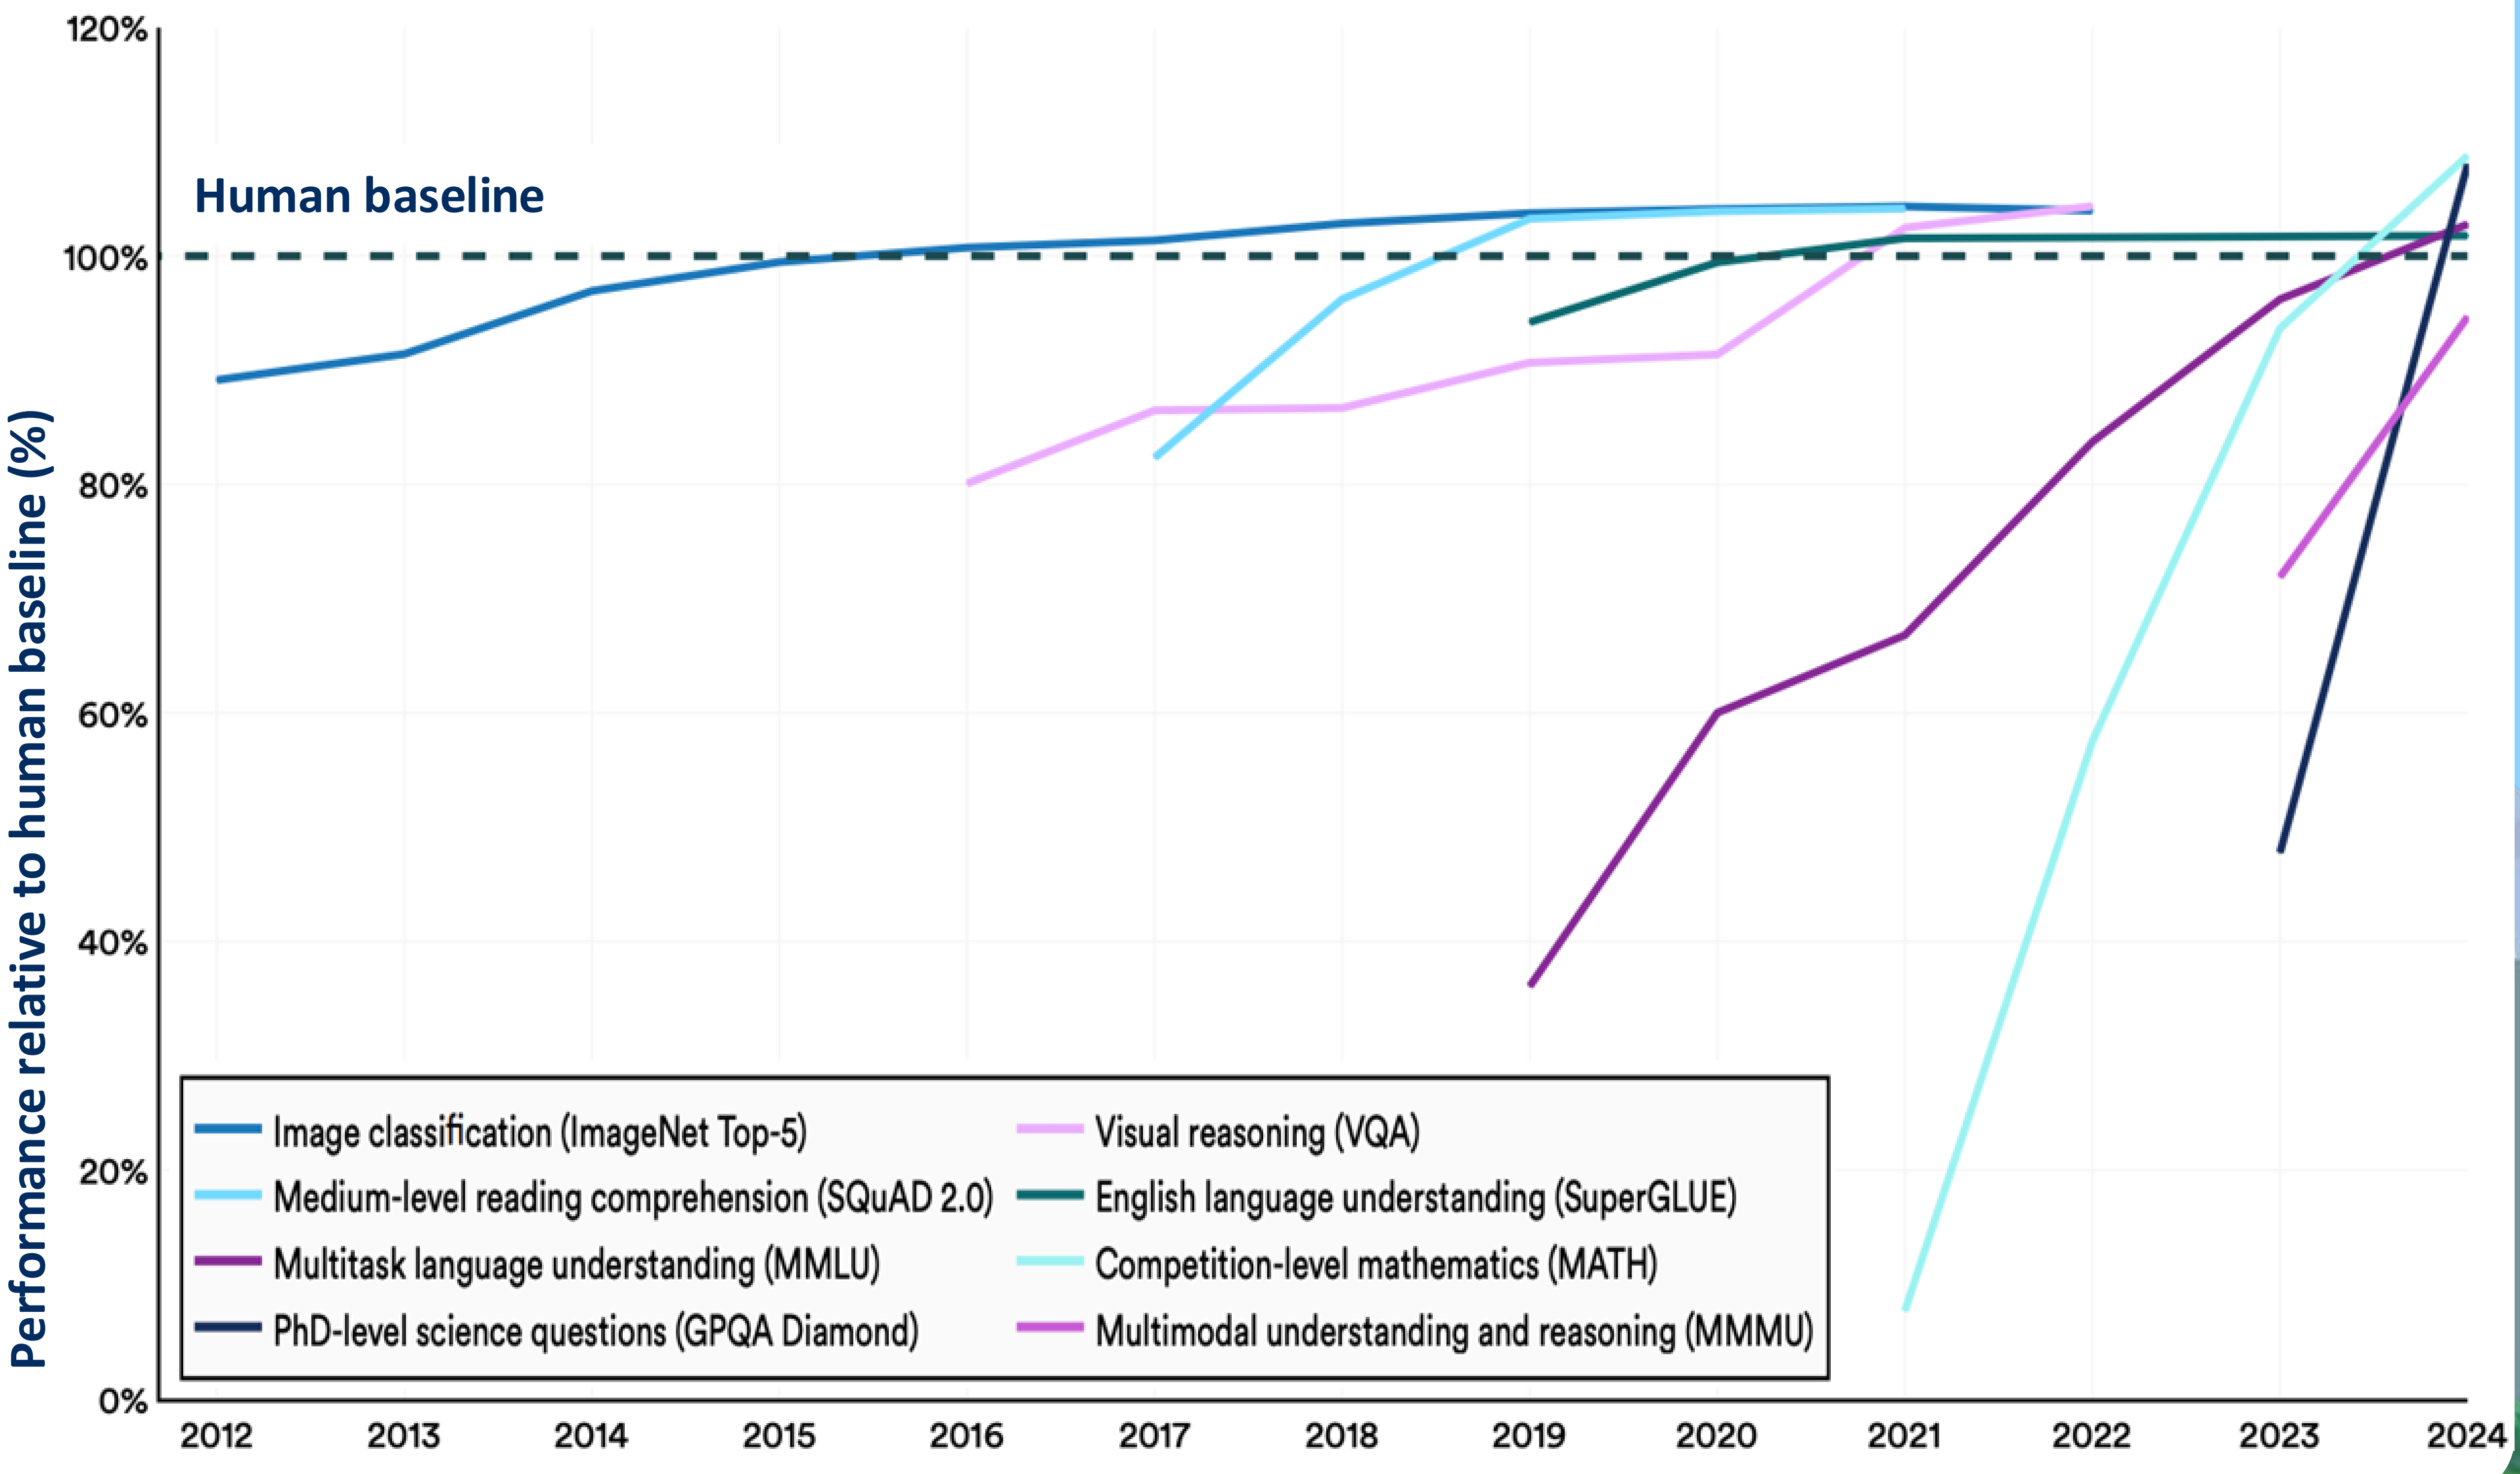

The 2025 AI Index Report, Standford HAI, 2025

Adapted LLM

Base LLM

$$M_{adapt} = F_{adapt}(M_{base})$$

Adaptation Method

Why We *Still* Need *Adaptation*

Adaptation → Performance↑

Domain

SaulLM-54B & SaulLM-141B: Scaling Up Domain Adaptation for the Legal Domain

Pierre Colombo equall Equall | Telmo Pires equall Equall | Malik Boudiaf equall Equall | Rui Melo equall Equall

BioMedLM: A 2.7B Parameter Language Model Trained On Biomedical Text

Elliot Bolton[1†], Abhinav Venigalla[2], Michihiro Yasunaga[1], David Hall[1], Betty Xiong[1], Tony Lee[1], Roxana Daneshjou[1], Jonathan Frankle[2],

Demystifying Domain-adaptive Post-training for Financial LLMs

Zixuan Ke, Yifei Ming, Xuan-Phi Nguyen, Caiming Xiong and Shafiq Joty
Salesforce AI Research
{zixuan.ke,yifei.ming,xnguyen,cxiong,sjoty}@salesforce.com
🧠 Project Page: https://github.com/SalesforceAIResearch/FinDAP
🤗 Datasets: https://huggingface.co/datasets/Salesforce/FinEval

Task

SFR-RAG: Towards Contextually Faithful LLMs

Foundational Autoraters: Taming Large Language Models for Better Automatic Evaluation

🔥 PROMETHEUS: INDUCING FINE-GRAINED EVALUATION CAPABILITY IN LANGUAGE MODELS

Seungone Kim[1,2*†] **Jamin Shin**[2,3*†] **Yejin Cho**[1*†] **Joel Jang**[4] **Shayne Longpre**[5] **Hwaran Lee**[2,3] **Sangdoo Yun**[2,3] **Seongjin Shin**[3] **Sungdong Kim**[1,2,3] **James Thorne**[1] **Minjoon Seo**[1†]

[1]KAIST AI [2]NAVER AI Lab [3]NAVER Cloud [4]University of Washington [5]MIT

Adaptation → Performance↑

salesforce

Domain/Language

Code Llama: Open Foundation Models for Code

Baptiste Rozière[†], Jonas Gehring[†], Fabian Gloeckle[†,*], Sten Sootla[†], Itai Gat, Xiaoqing Ellen Tan, Yossi Adi[◇], Jingyu Liu, Romain Sauvestre, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, Gabriel Synnaeve[†]

Meta AI

CHIMED-GPT: A Chinese Medical Large Language Model with Full Training Regime and Better Alignment to Human Preferences

Yuanhe Tian[♠♥*], Ruyi Gan[♣♠*], Yan Song[♠†], Jiaxing Zhang[♣], Yongdong Zhang[♠]

ALLaM: Large Language Models for Arabic and English

NCAI
المركز الوطني
للذكاء الاصطناعي
National Center for AI

SDAIA
الهيئة السعودية للبيانات
والذكاء الاصطناعي
Saudi Data & AI Authority

Task

How to Train Long-Context Language Models (Effectively)

Tianyu Gao[*] Alexander Wettig[*] Howard Yen Danqi Chen
Princeton Language and Intelligence, Princeton University
{tianyug,awettig,hyen,danqic}@cs.princeton.edu

DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning

DeepSeek-AI

research@deepseek.com

Toolformer: Language Models Can Teach Themselves to Use Tools

Timo Schick Jane Dwivedi-Yu Roberto Dessì[†] Roberta Raileanu
Maria Lomeli Luke Zettlemoyer Nicola Cancedda Thomas Scialom
Meta AI Research [†]Universitat Pompeu Fabra

Adaptation →
Performance↑
Cost↓

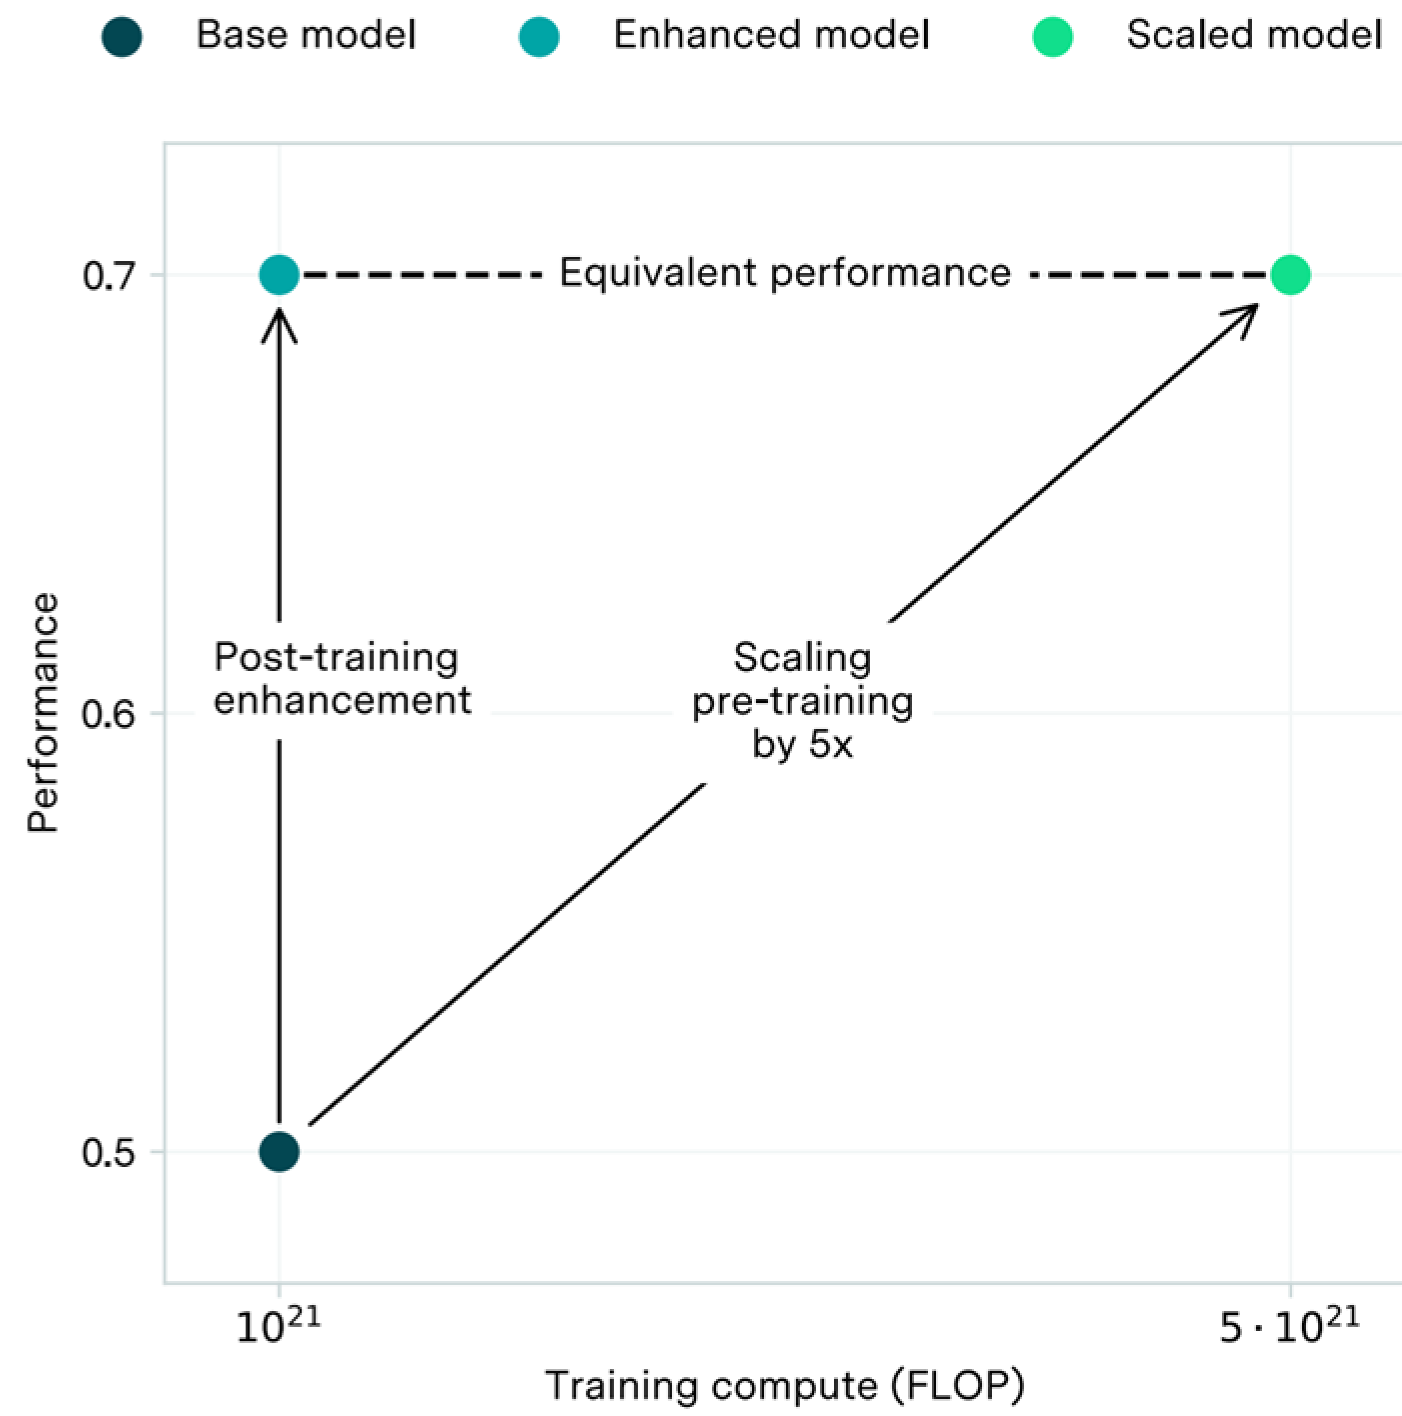

AI capabilities can be significantly improved without expensive retraining, Davidson et al., 2023

Training is Becoming Increasingly Affordable

Size↓

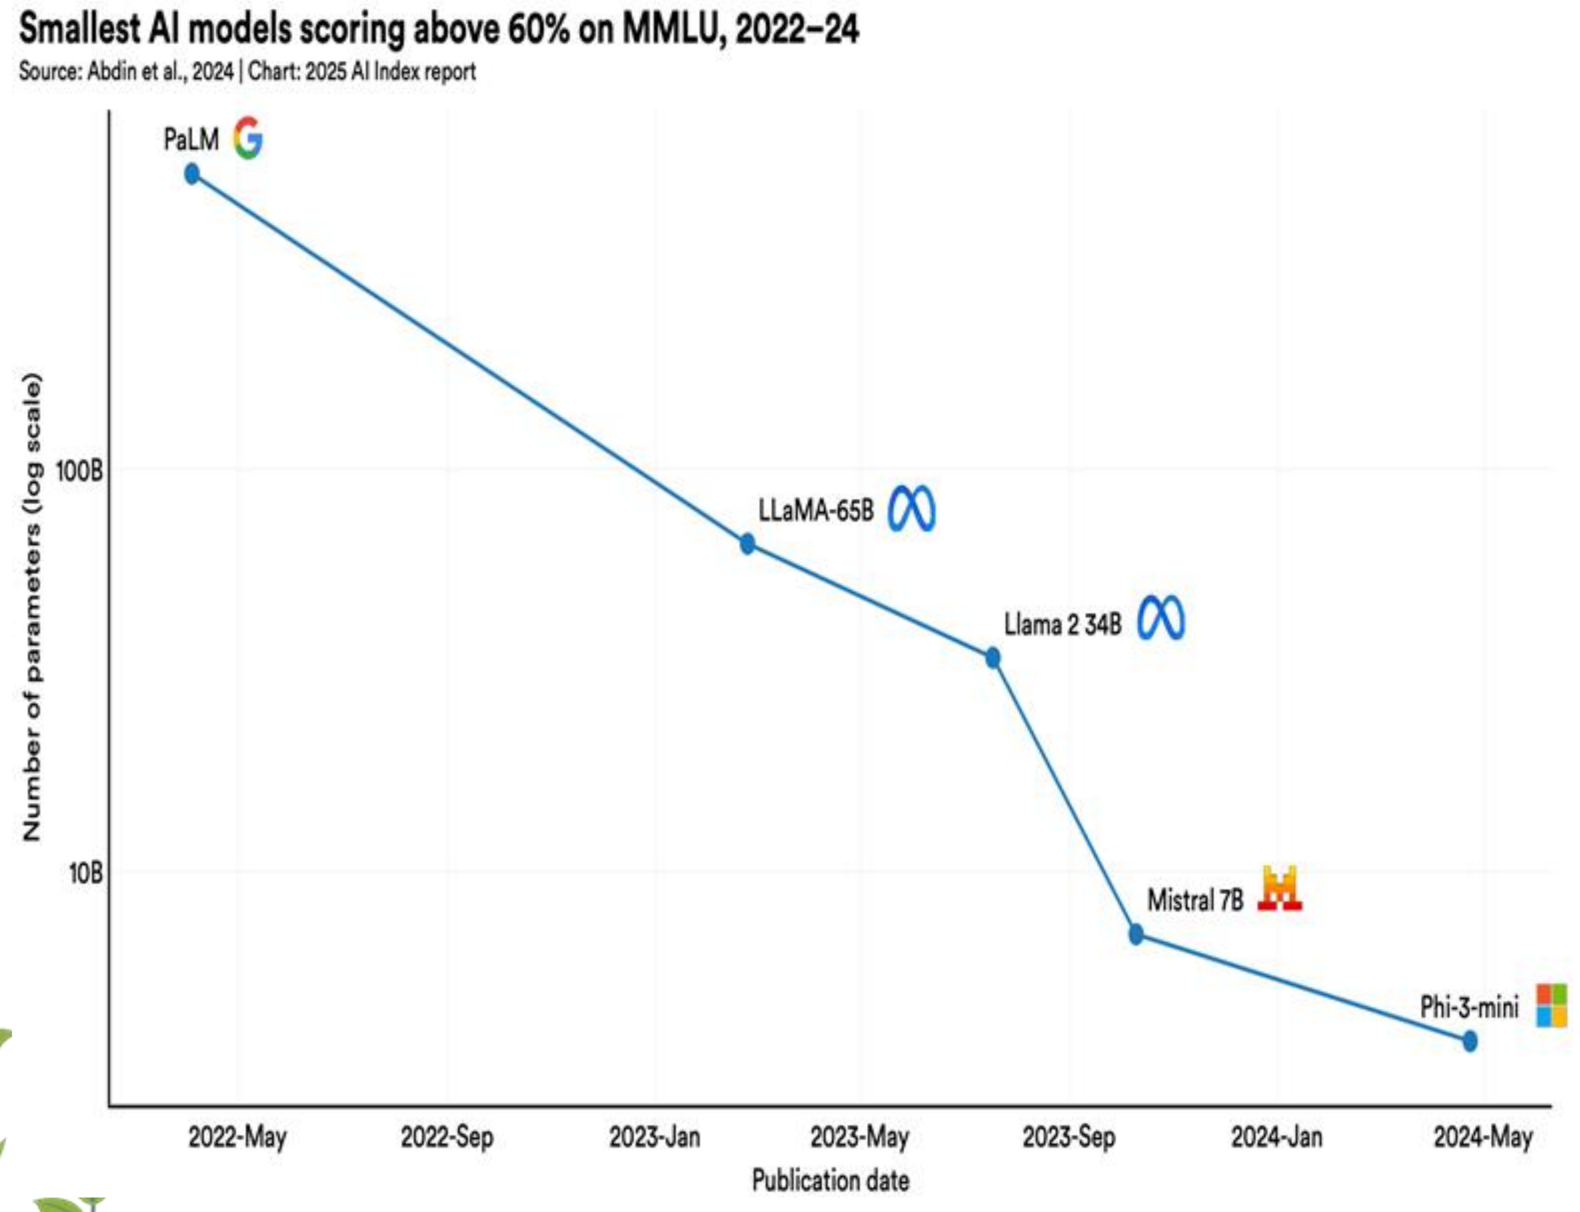

Cost↓

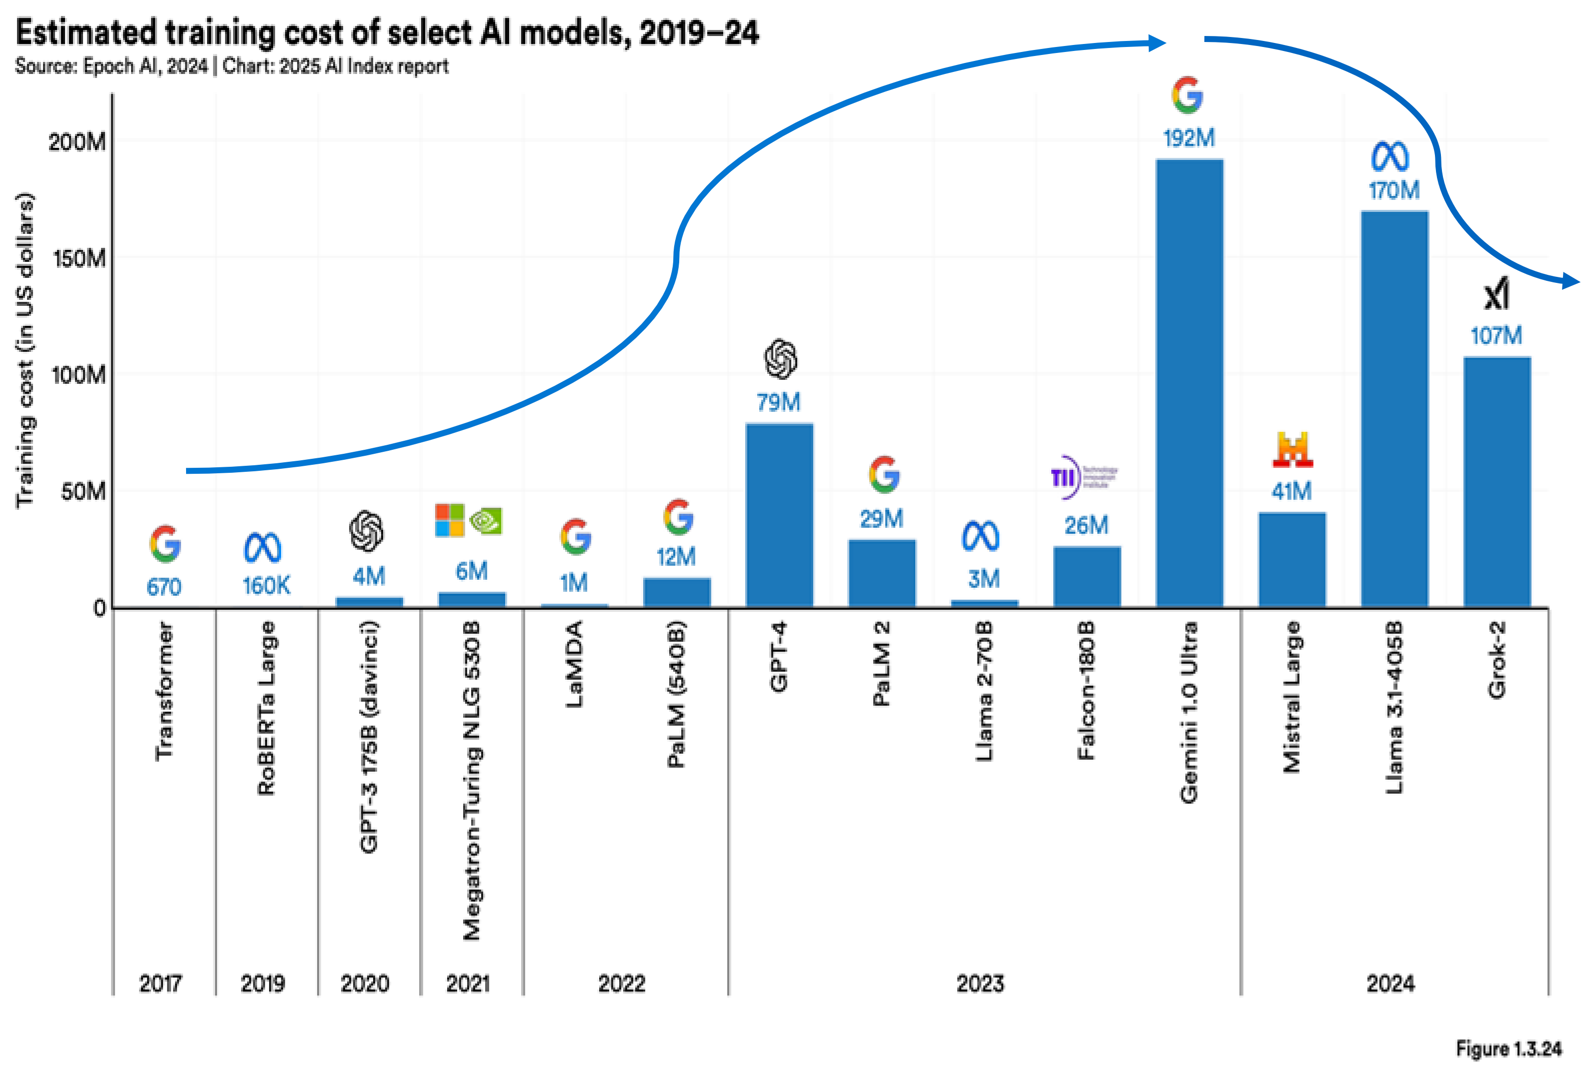

Lower cost- to-serve for small domain or task specific models

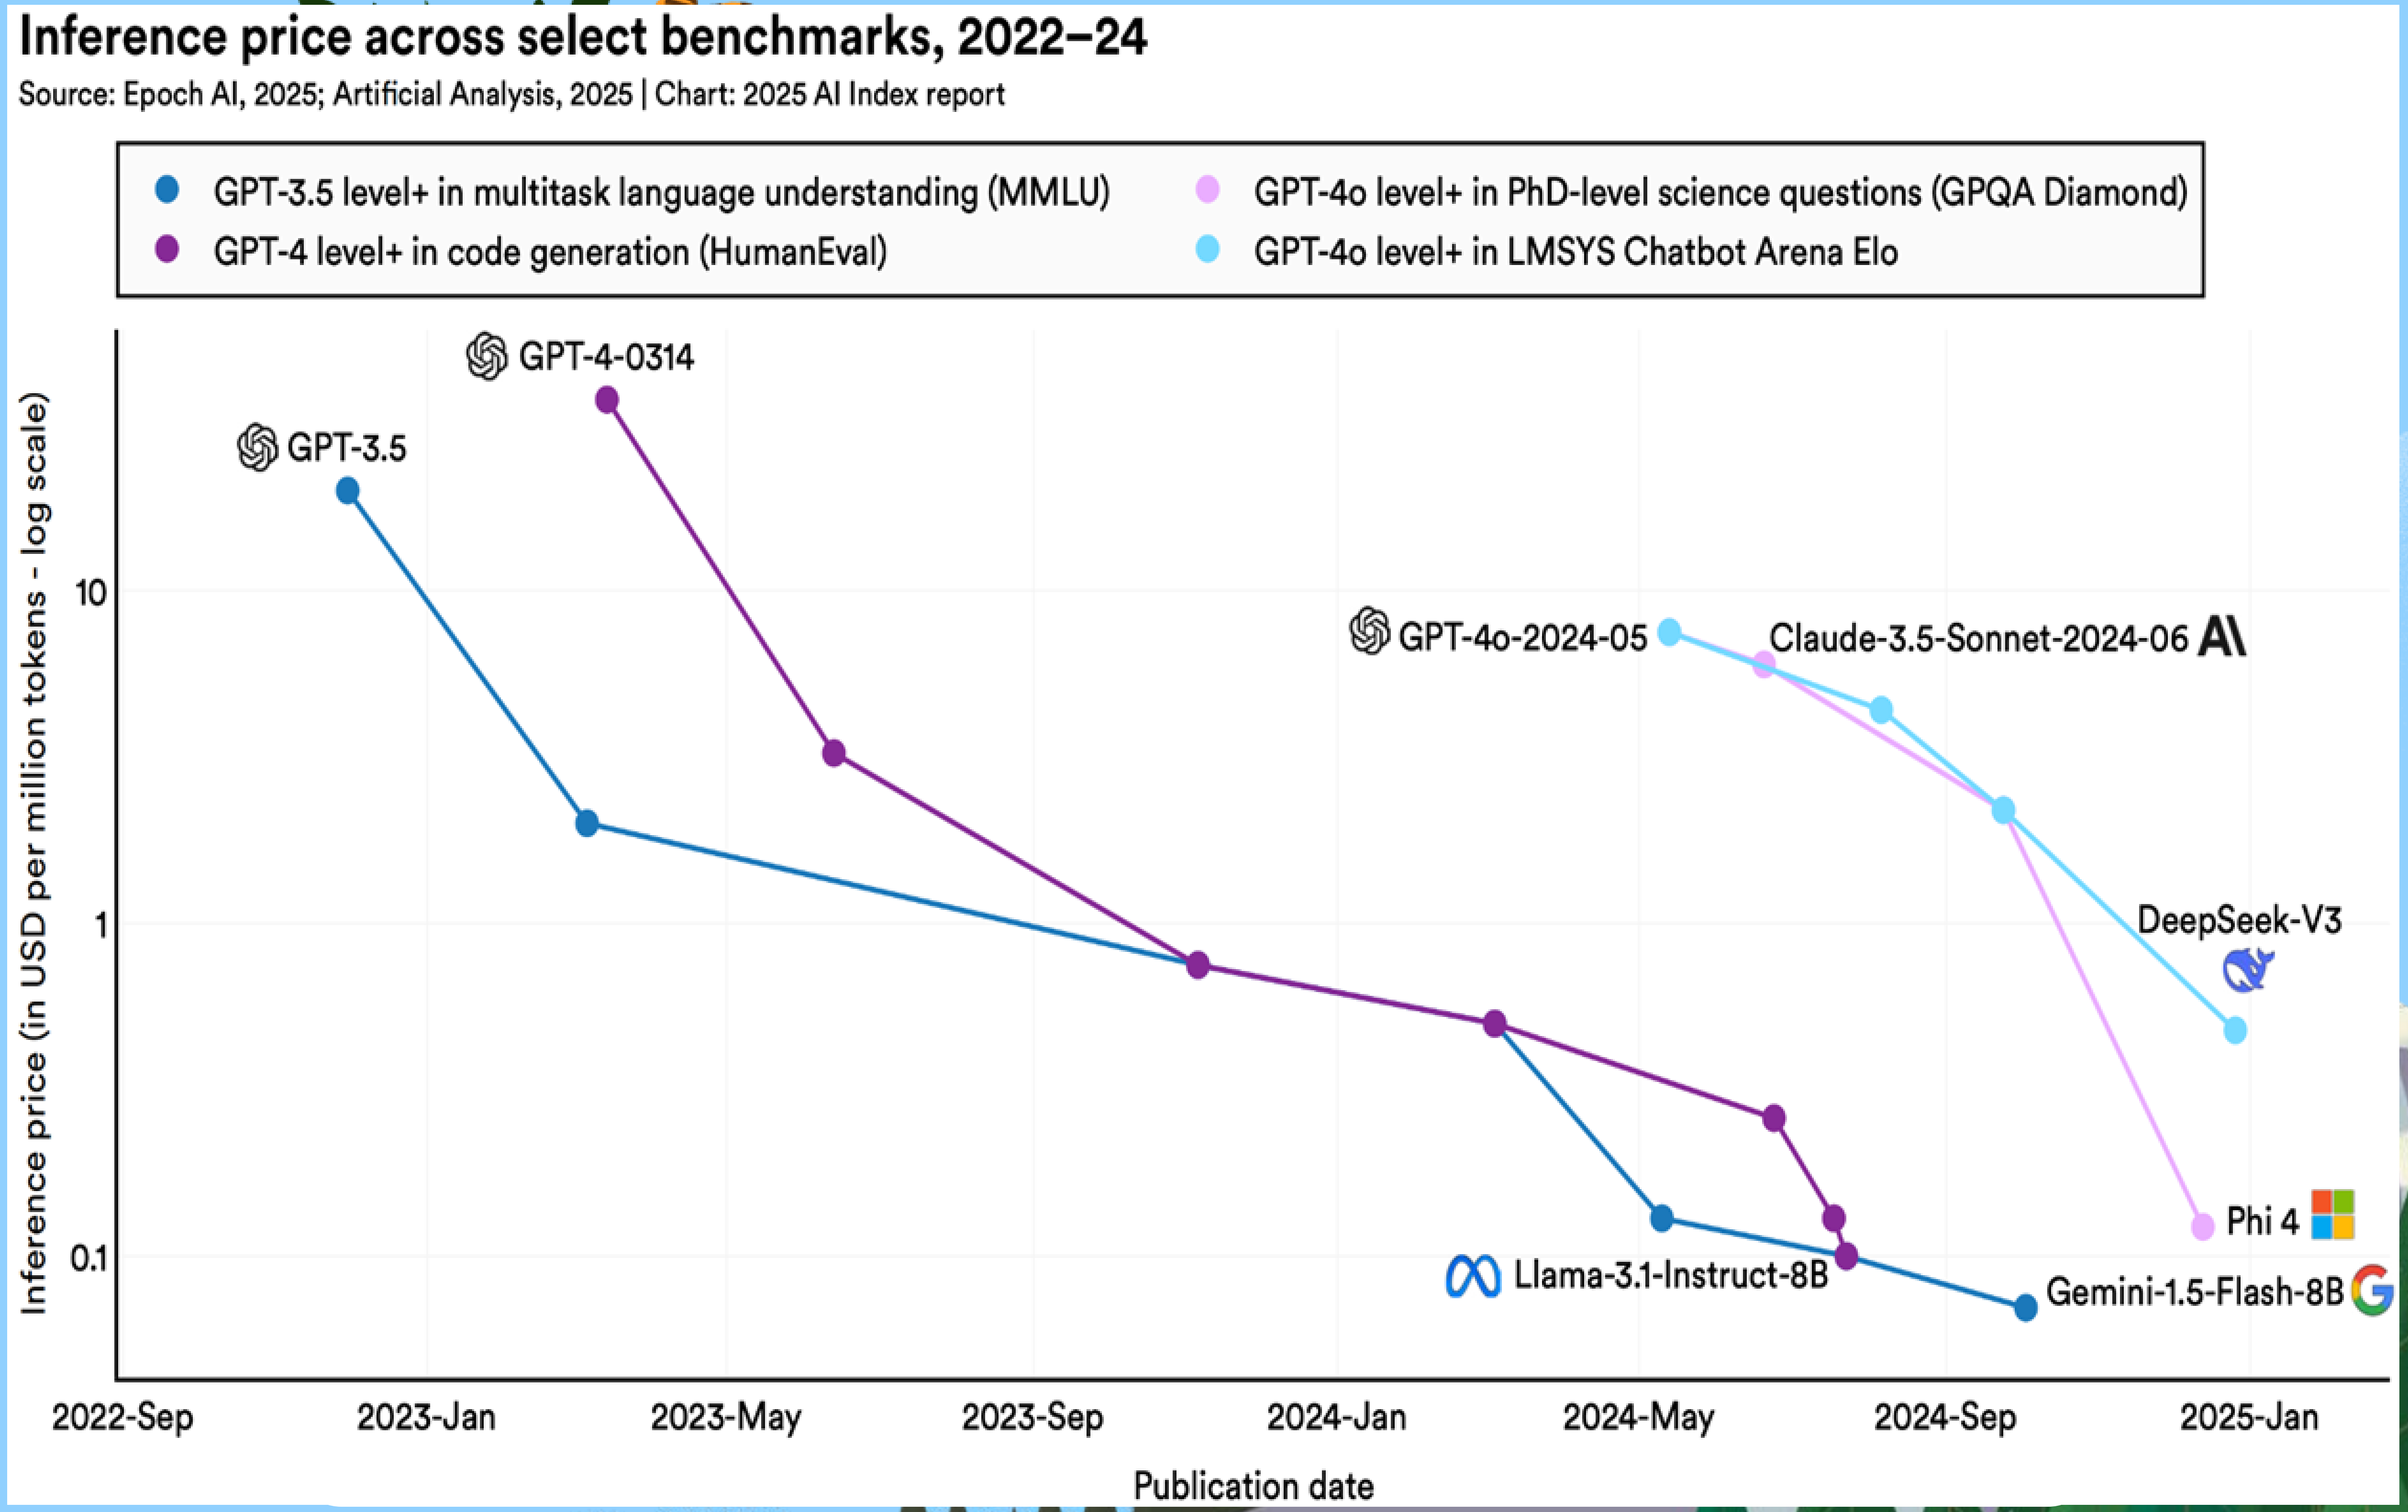

Adaptation in the Era of Experience

Our World is changing — LLMs must adapt accordingly

- Long-tail domains/tasks
- Emerging domains/tasks

To go beyond human data, LLMs need to adapt through their own experience

- Self-discover own knowledge + adaptation

My personal bet is we’re going to see a mixture of general models and specialist models that are much more focused

Dan Klein, professor at UC Berkeley (Mar, 2025)

https://alumni.berkeley.edu/california-magazine/online/data-is-fueling-the-ai-revolution-what-happens-when-it-runs-out/#

Key Concepts in Adaptation

LLM Workflow

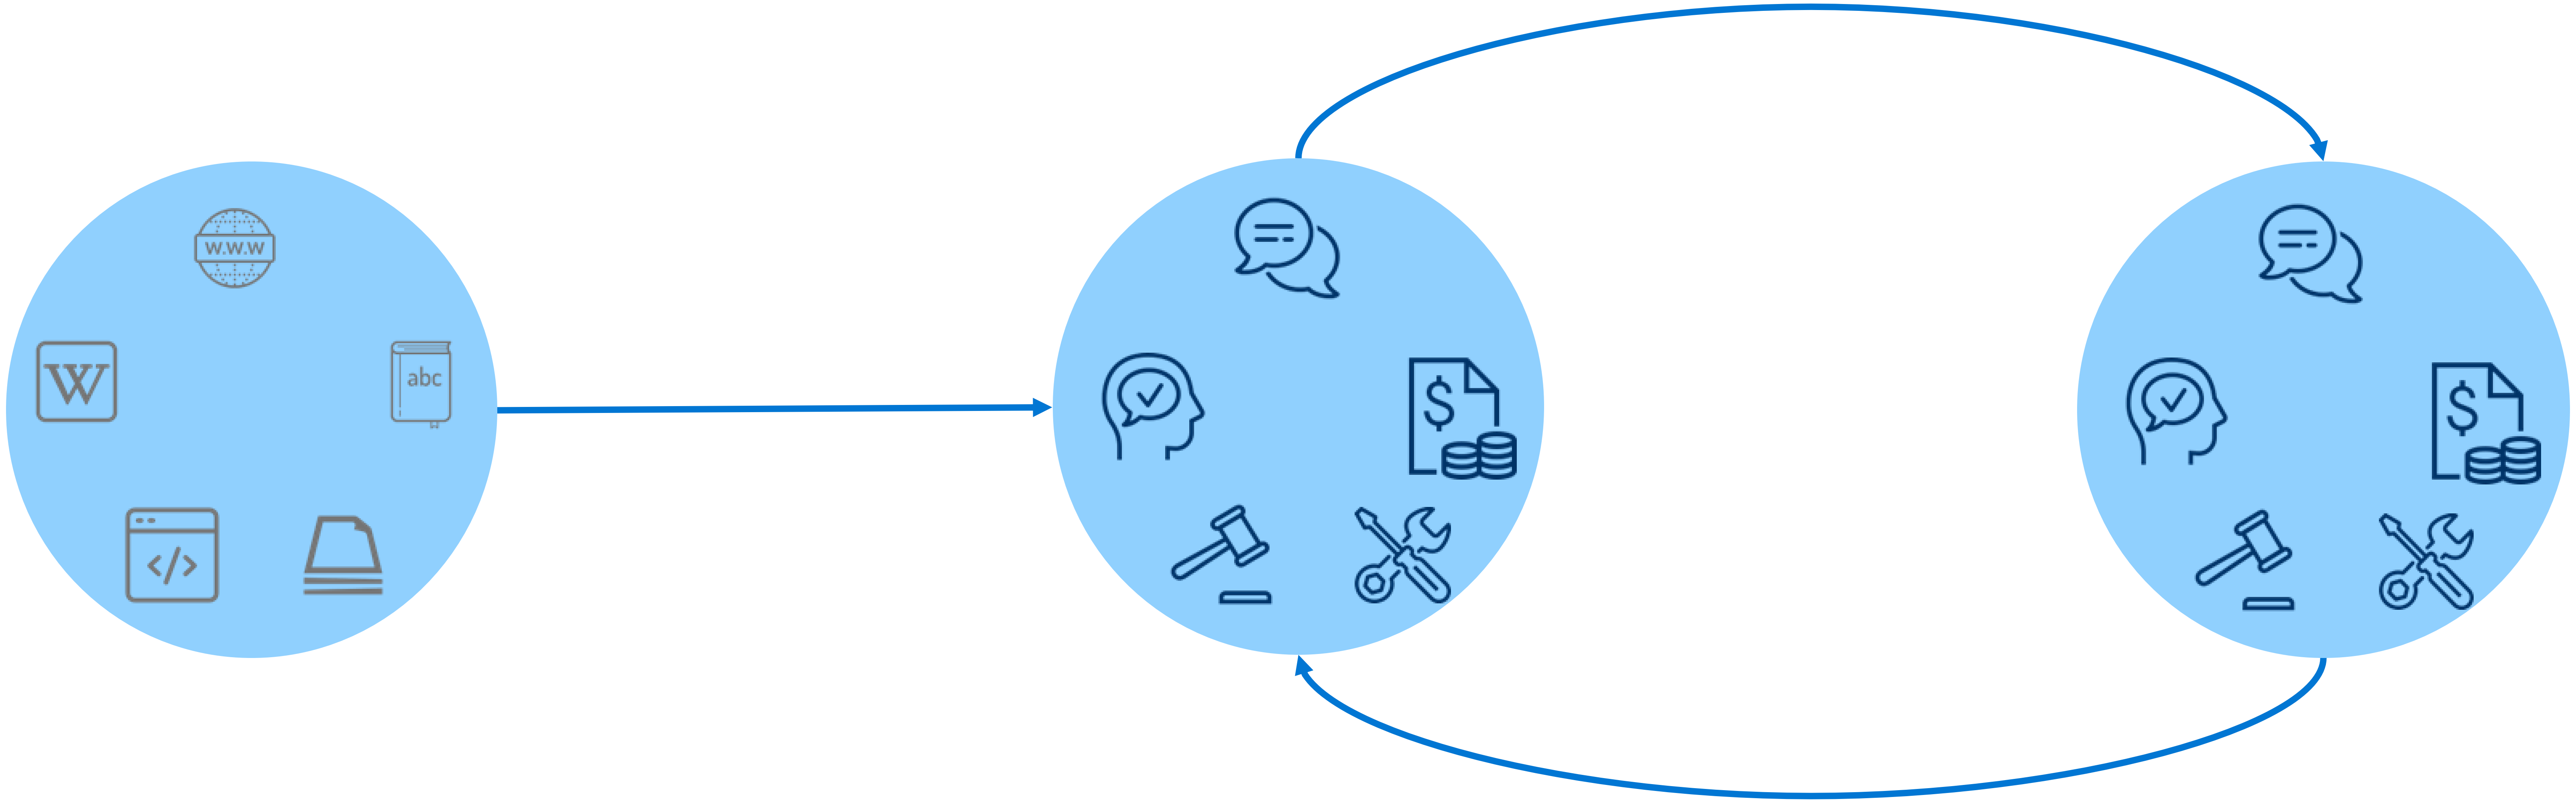

Pre-training

Large-scale data,
Extensive computation

Adaptation

Evaluation

General capabilities
(e.g., chat, reasoning)

Specialized capabilities (e.g., finance, tool-use)

Adaptation – Regimes

In-context Learning

Single LLM, zero-shot, few-shot, **No parameters updated**

Learning to Adapt

Update the LLM parameters to adapt LLM to specific task/domain/environment

Main focus of this tutorial

Inference Scaling

Multiple LLM calls, **No parameters updated**

Adaptation – Paradigms

Parametric Knowledge

Update LLM parameters, without interacting with external environment (e.g., domain- and task-specific LLMs)

Semi-Parametric Knowledge

Update LLM parameters to interact with external environment (e.g., RAG)

This represents the shift from standalone LLMs → **agents**

Adaptation – A Comparison

Pre-training	Post-training	Continual Learning
Learn the foundation knowledge, but the raw pre-trained LLMs are **neither** safe **nor** robust for public use and interactions (thus "alignment/adaptation" is required)	**Convention:** Adaptation **=** Adapt model from source to target distribution **LLM Era:** Adaptation **≈** Post-training	**Convention:** Learning a sequence of disjoint tasks; **Main focus:** prevent forgetting **Side focus:** encourage transfer **LLM era:** Tasks not disjoint; **Main focus:** encourage transfer **+** prevent forgetting

Continual Pre-training of Language Models Ke et al., 2023
Continual Learning of Natural Language Processing Tasks: A Survey, Ke et al., 2023

Adaptation – Four Most Popular Methods

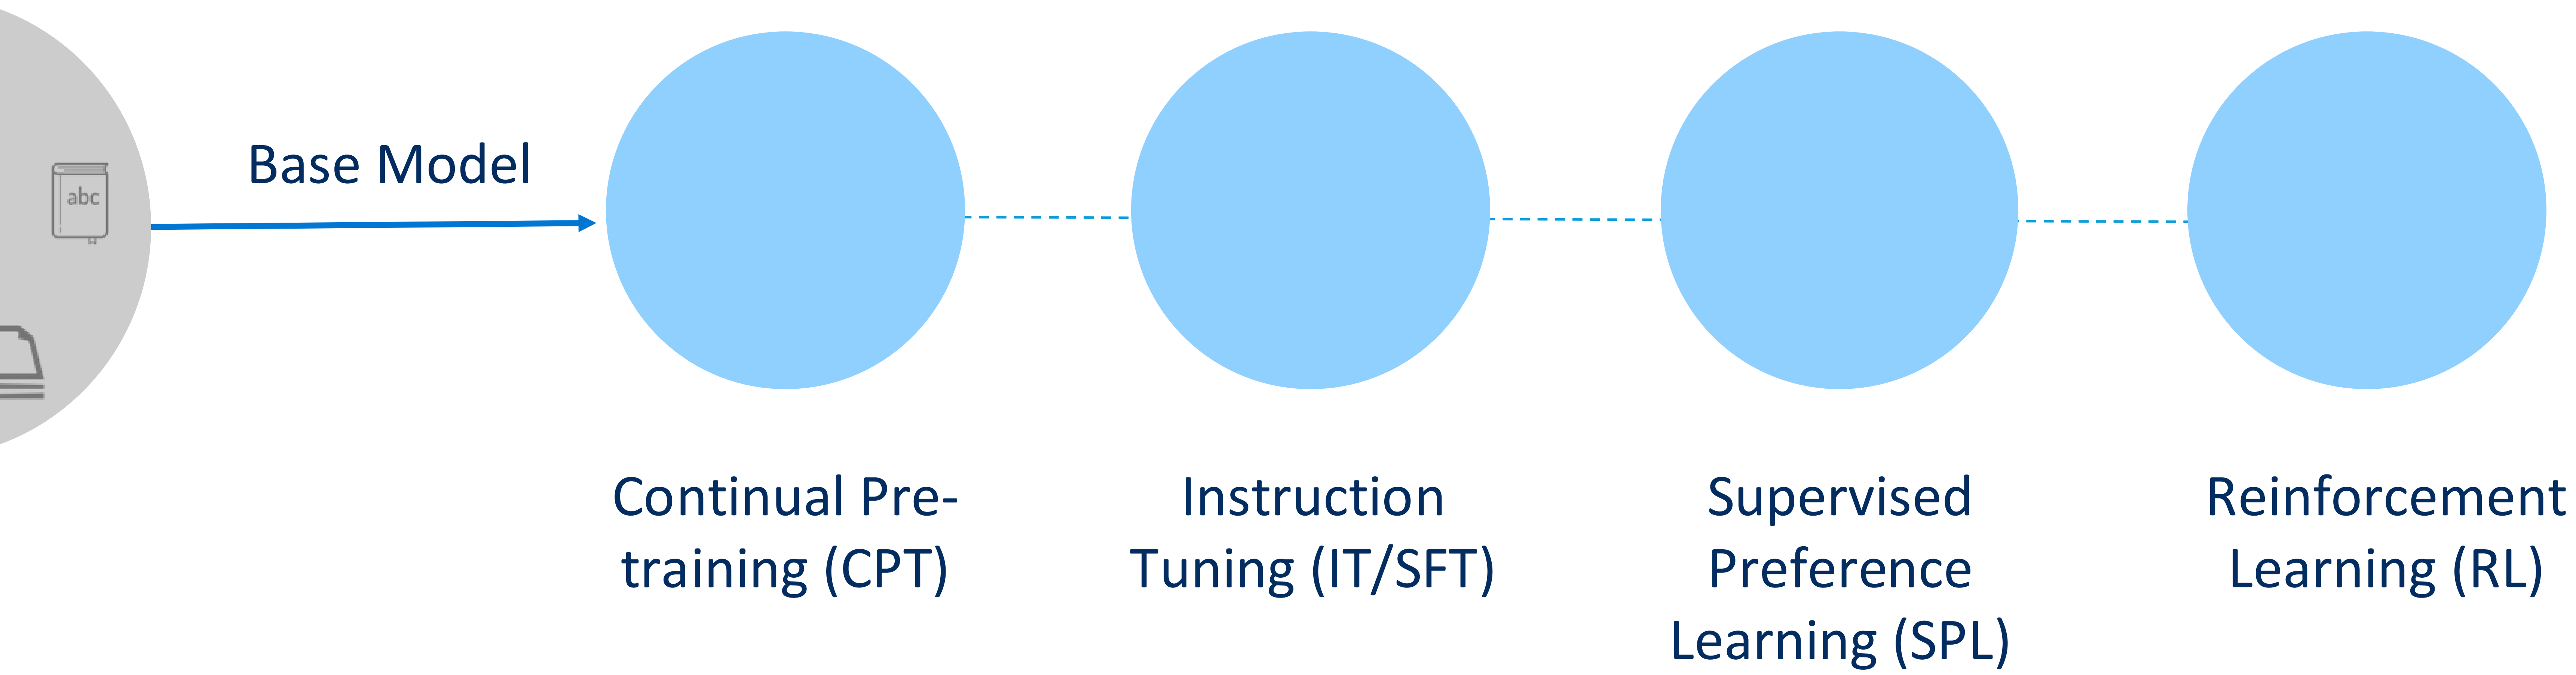

Adaptation – Four Most Popular Methods

```
<|begin_of_text|>
SEC Finalizes ARS Settlement
to Provide $7 Billion in
Liquidity to Wachovia
Investors...
<|end_of_text|>
```

Continual Pre-training

Inject or emphasize target knowledge (e.g., domain knowledge)

```
<|system|>
You are a helpful assitant
<|end|>
<|user|>
How many helicopters can you eat?
<|end|>
<|assistant|>
{Answer goes here}
```

Instruction Tuning

Formatting and instruction following

```
<|prompt|>what are the minimum
lease payments in 2022
<|end|>
<|rejected|>
$17,188 / $34,356 * 100
= 49.98%.
<|end|>
<|chosen|>
$17,188 / $34,356 * 100
= 49.99%.
<|end|>
```

Sup. Preference Learning

Align to human or AI preferences

```
<|prompt|>
I'm not sure if it's the right
to do and could use some
outside opinions.
TL;DR:
<|end|>
```

Reinforcement Learning

Boost performance on complicated (and verifiable) tasks (e.g., reasoning)

Adaptation – Example Training Workflow

E.g., Tulu 1,2; Instruct GPT

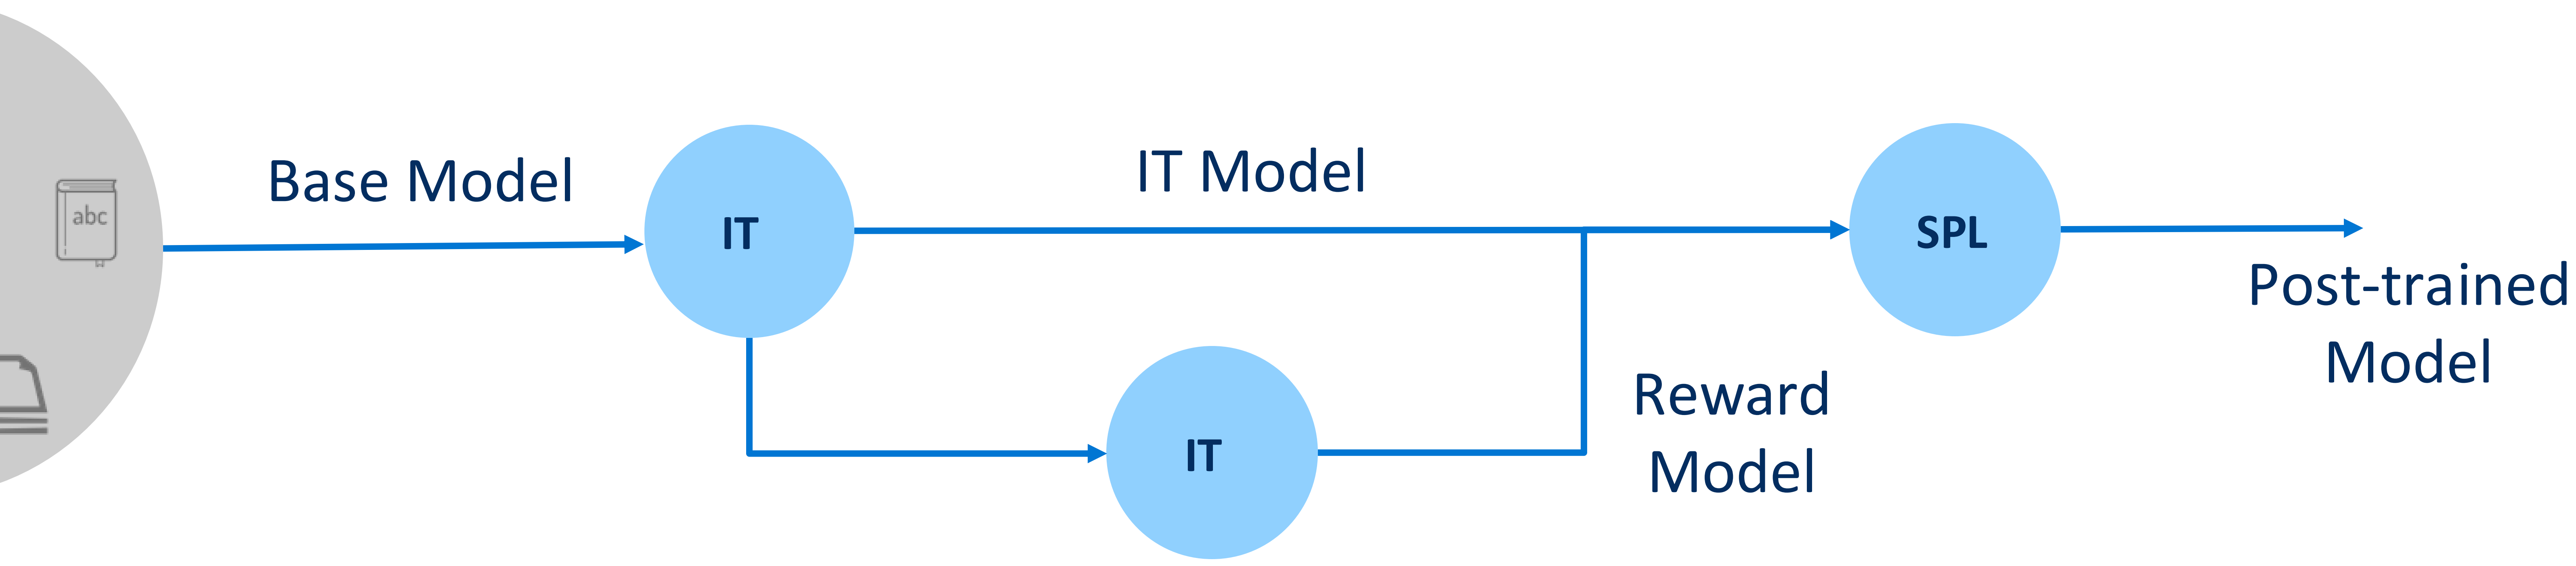

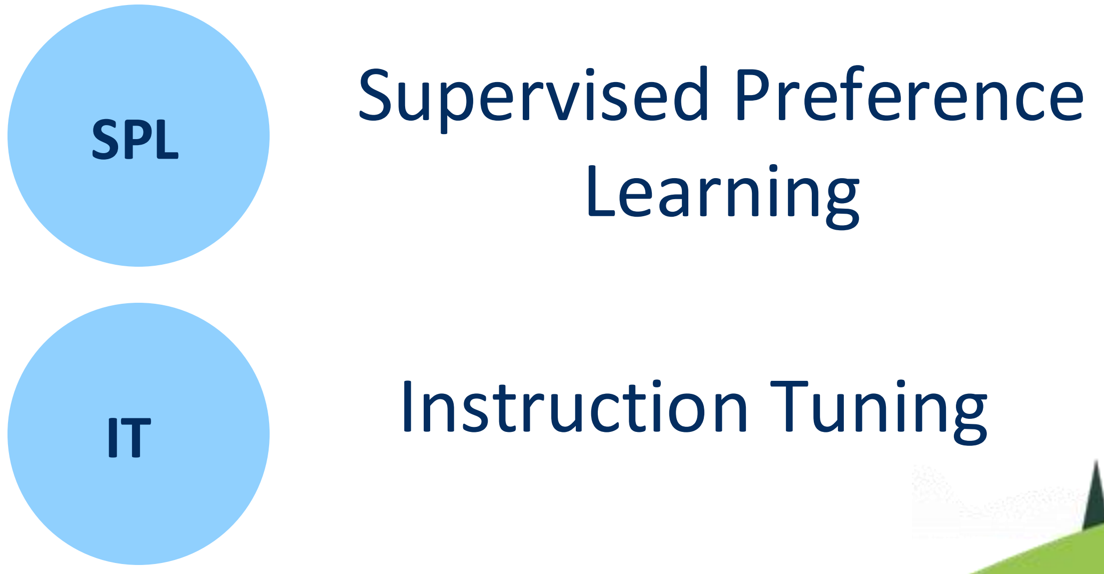

Training language models to follow instructions with human feedback, Ouyang et al., 2022

Adaptation – Example Training Workflow

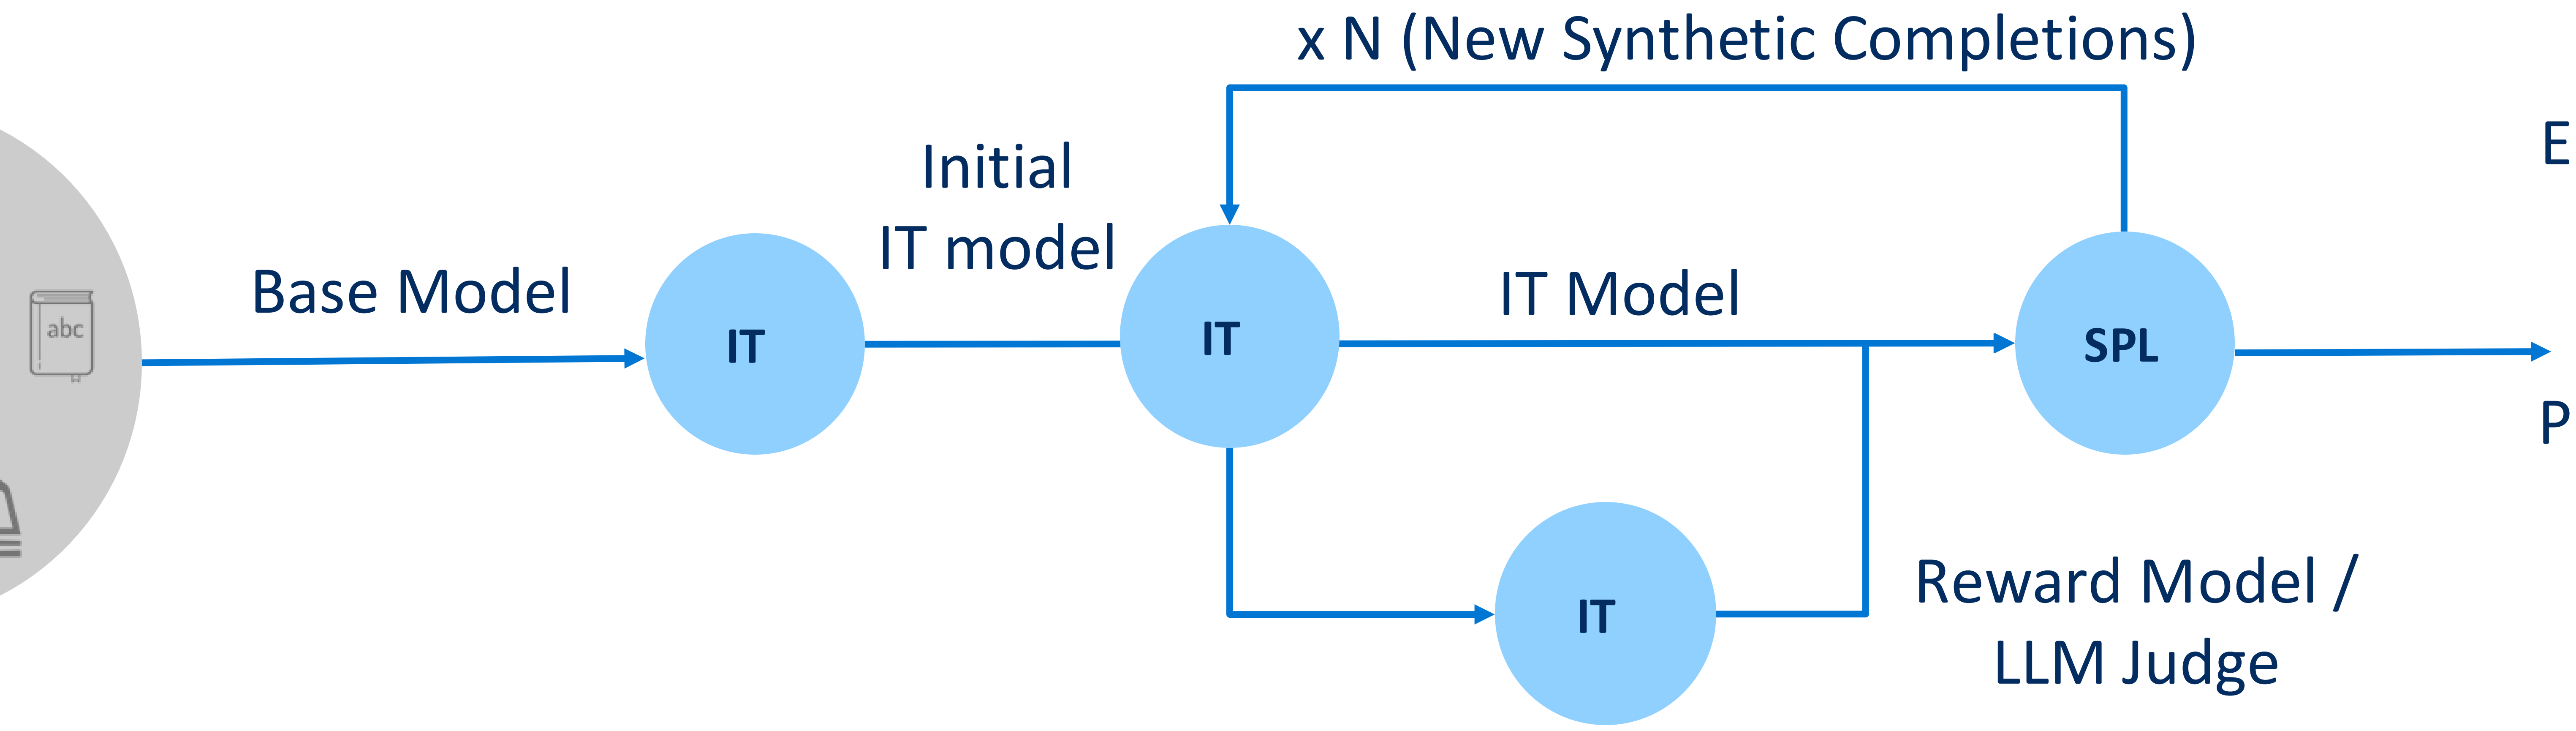

E.g., Tulu 3; Llama3.1

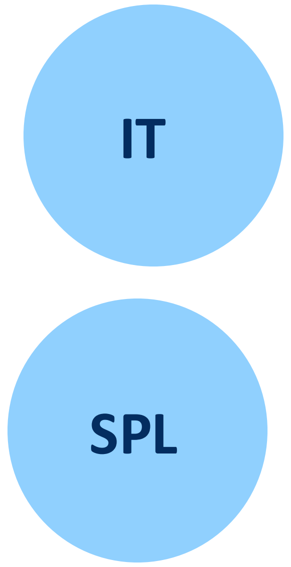

Instruction Tuning

Supervised Preference Learning

Tülu 3: Pushing Frontiers in Open Language Model Post-Training, Lambert et al., 2025

Adaptation – Example Training Workflow

E.g., DeepSeek-R1

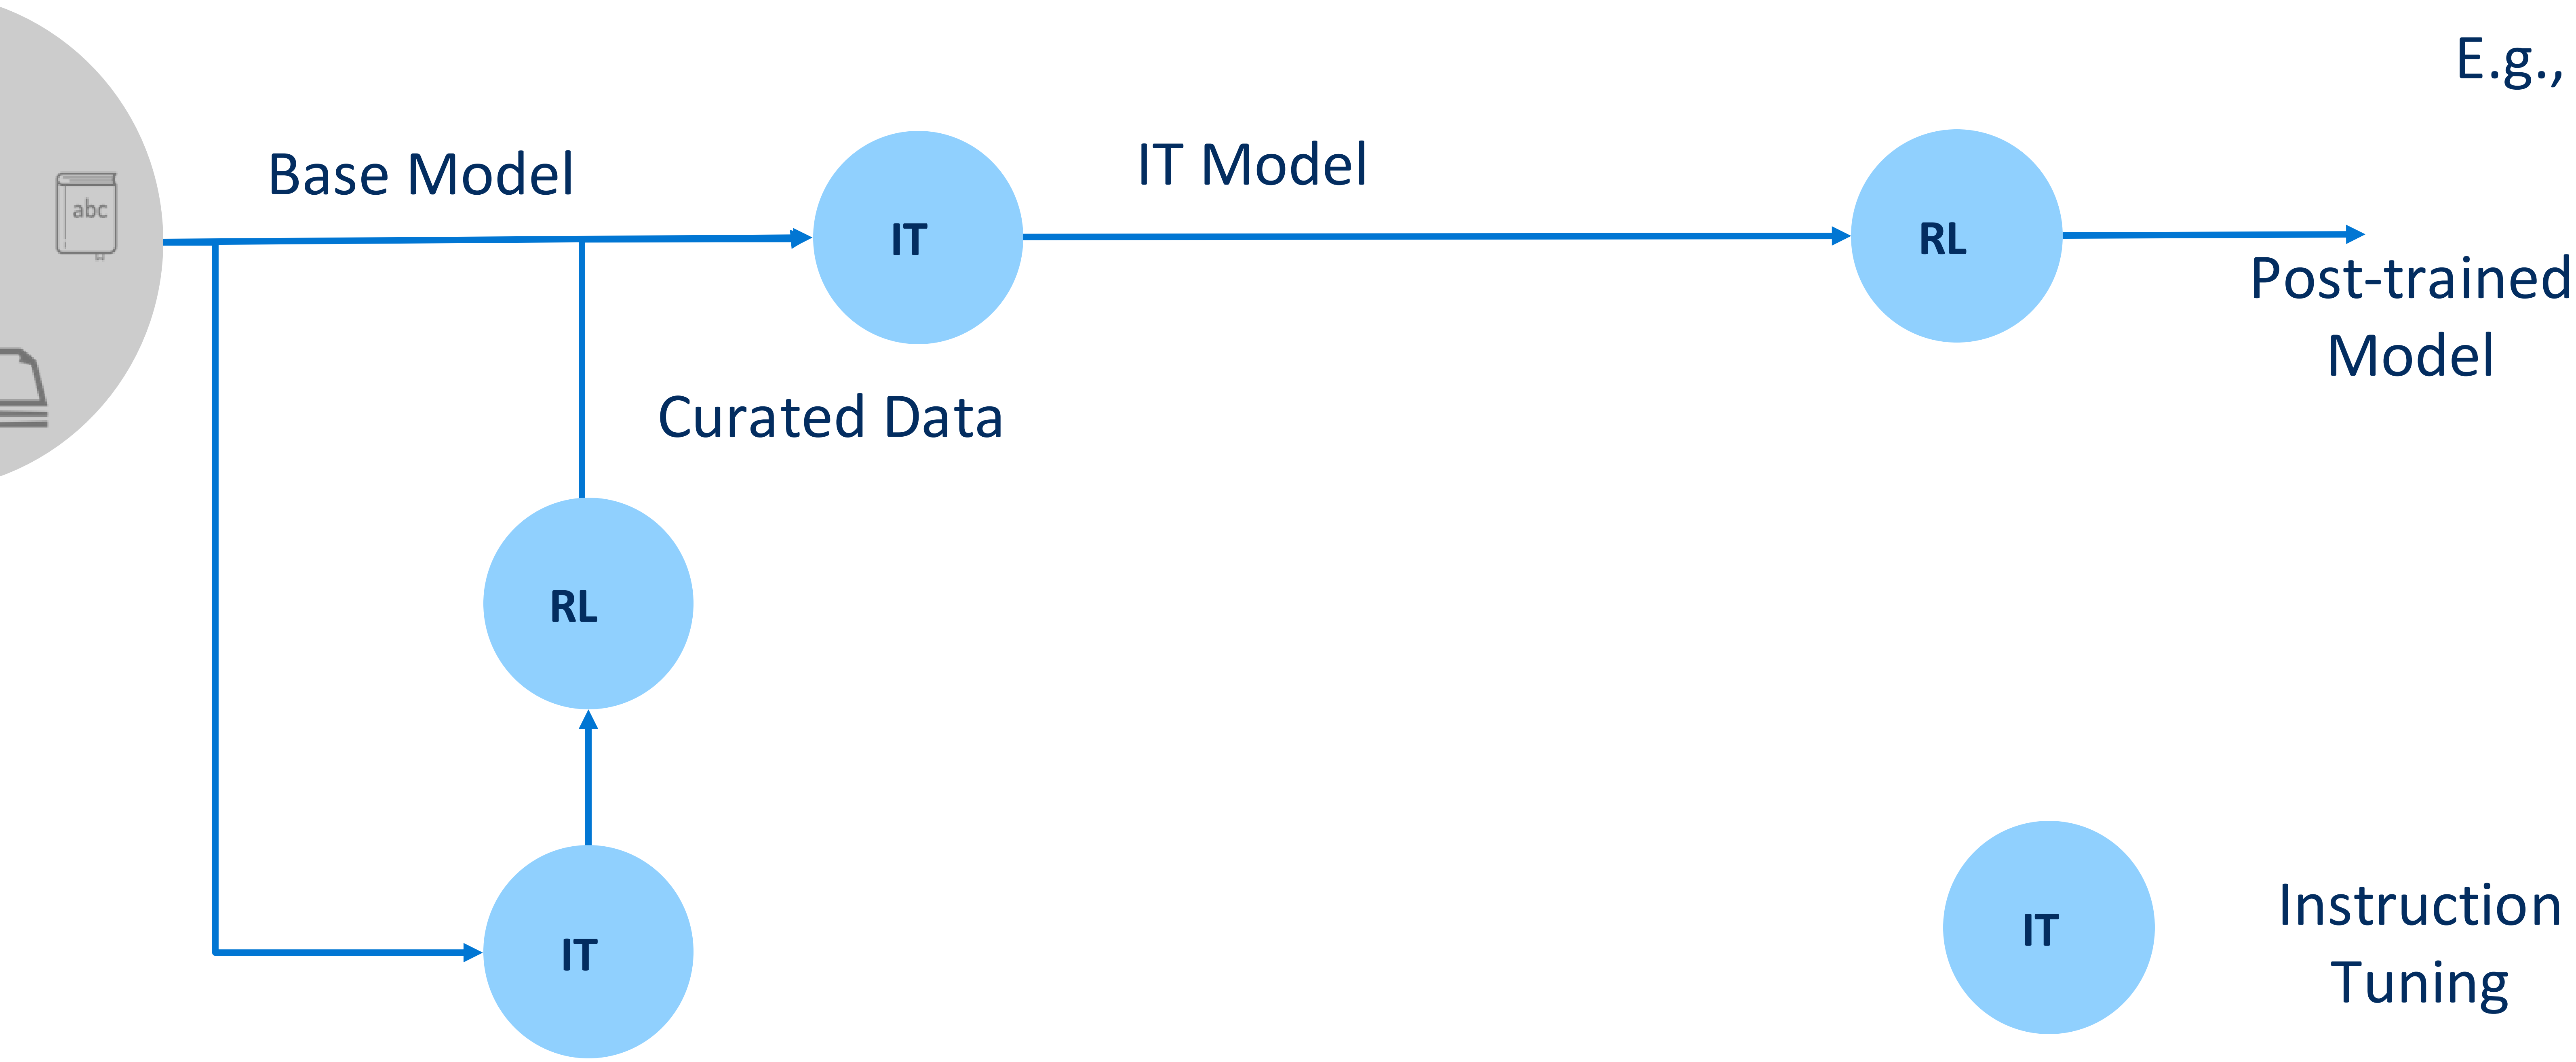

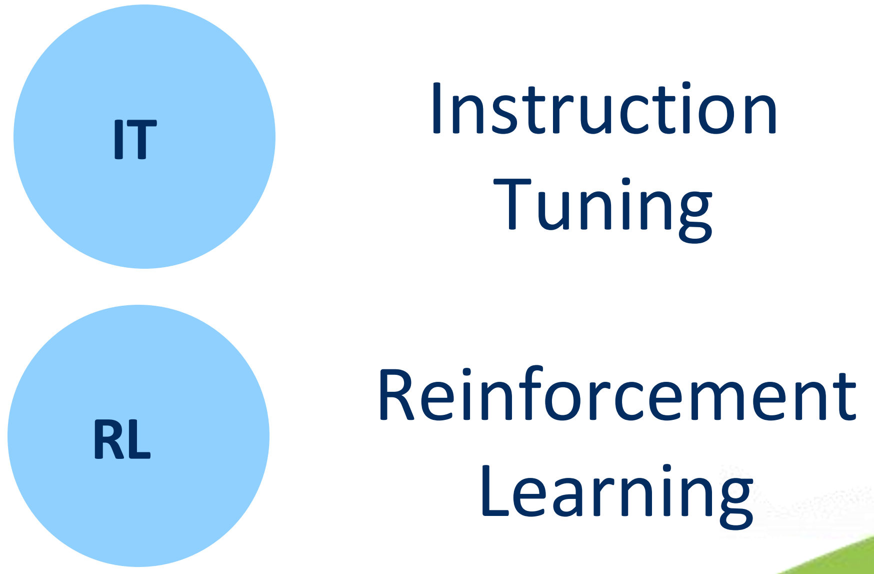

DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning, Deepseek-AI, 2025

Adaptation – Example Training Workflow

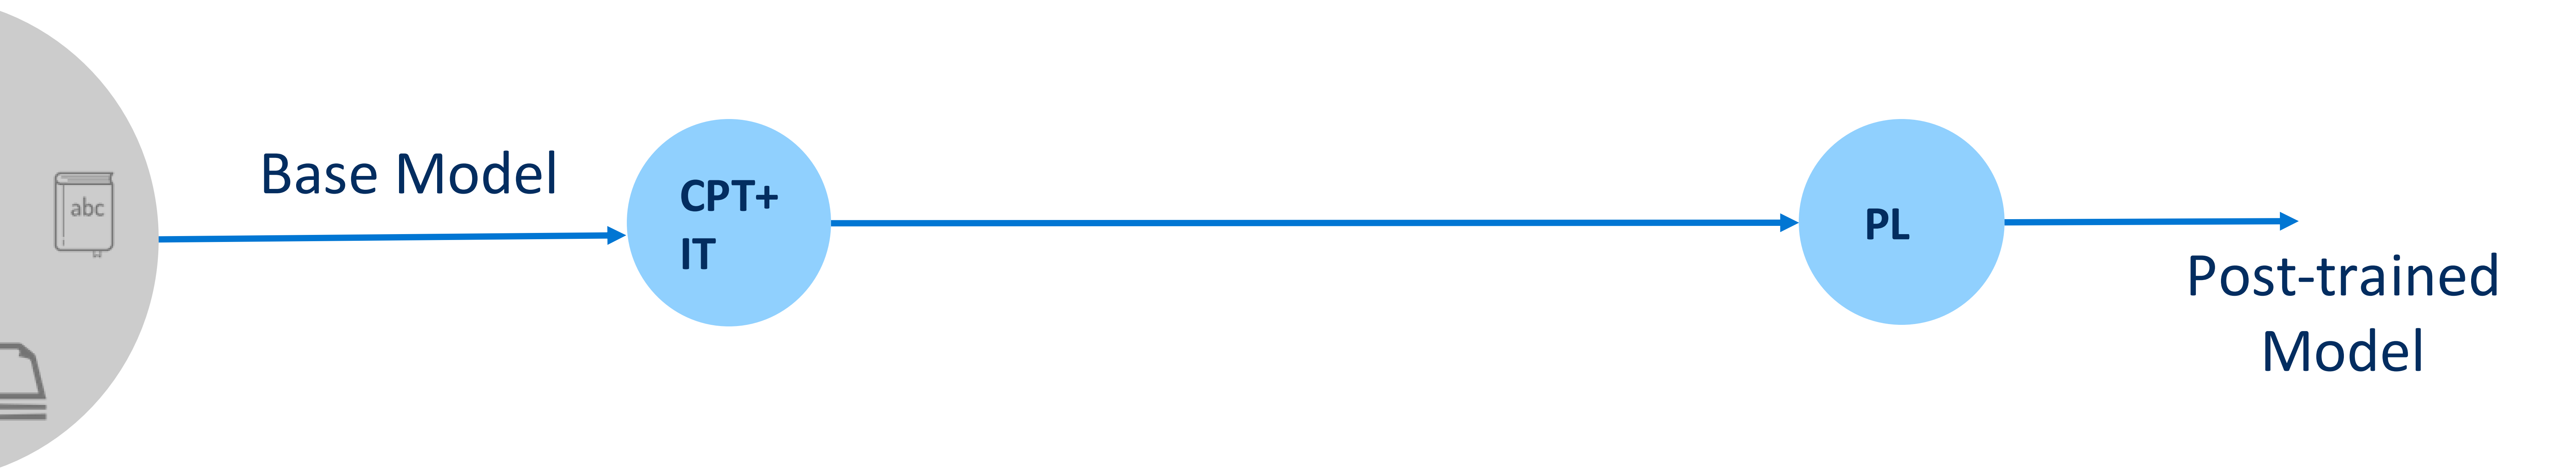

E.g., FinDAP

Demystifying Domain-adaptive Post-training for Financial LLMs, Ke et al., 2025

CPT Continual Pre-training

IT Instruction Tuning

SPL Supervised Preference Learning

Adaptation – Example Training Workflow

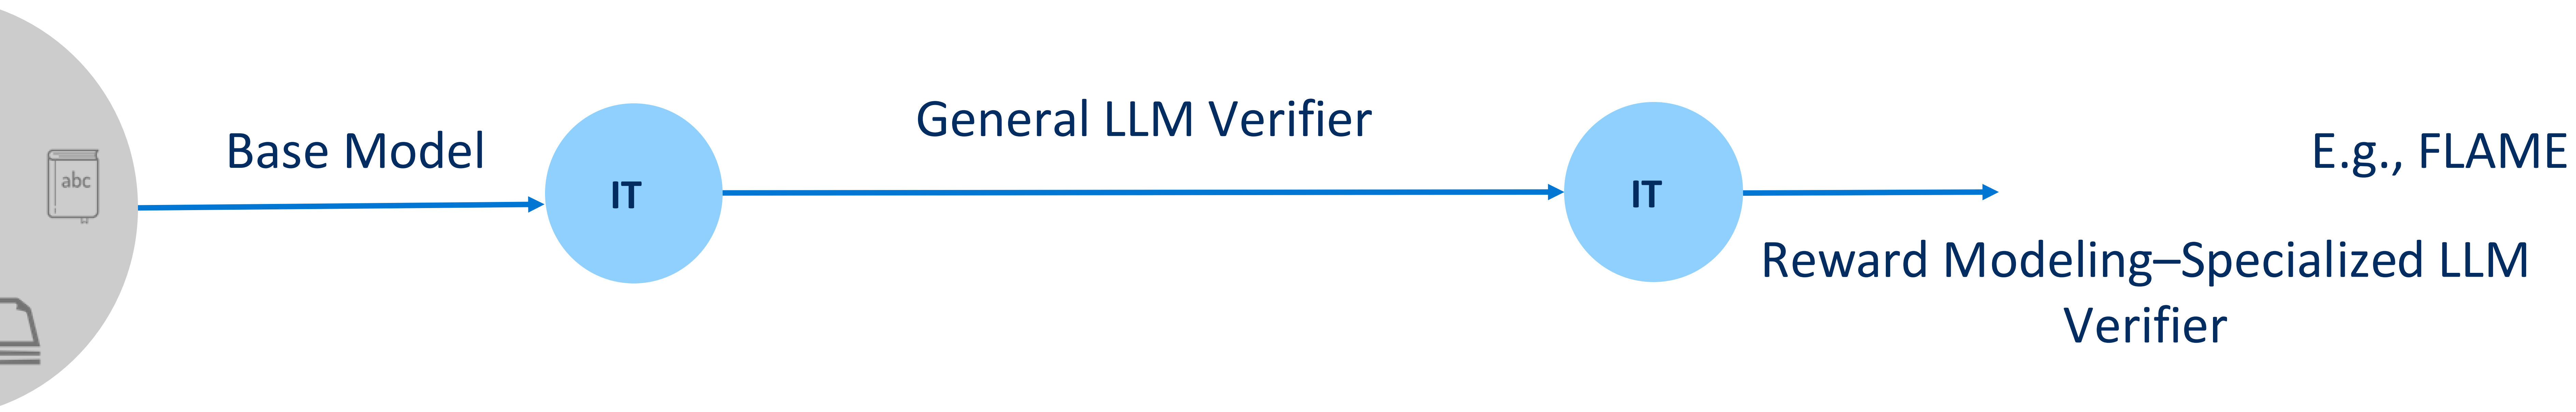

…… We should expect more to come

IT Instruction Tuning

Foundational Autoraters: Taming Large Language Models for Better Automatic Evaluation, Vu et al., 2024

Research Questions in LLM Adaptation

Data Perspective

Seed Data: What gives a good data mixture and how to obtain high-quality data? (often limited in amount)

Data Recipe: Given the limited amount of seed data, how to synthesize or construct high-quality data?

Model Perspective

Methods: What are the basic methods and their variants of LLM adaptation?

Training Workflow: What is the effective workflow to connect those basic methods?

Adaptation – Four Considerations

Core Capabilities

What capabilities do you actually care about?

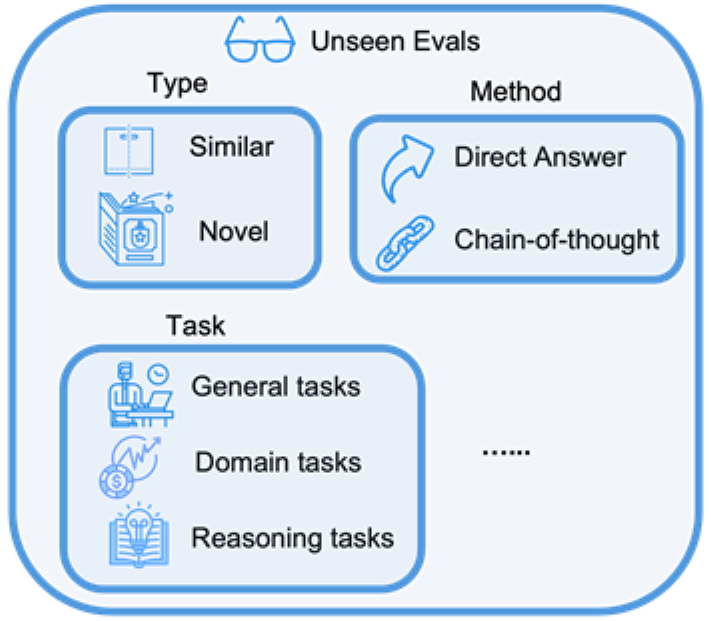

Evaluation

How do you measure the progress toward targeted capabilities?

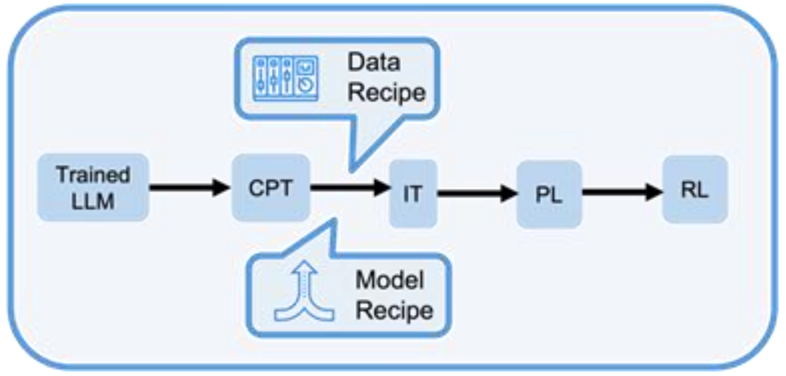

Training Recipe

How do you construct useful data from your seed data and what is your model recipe?

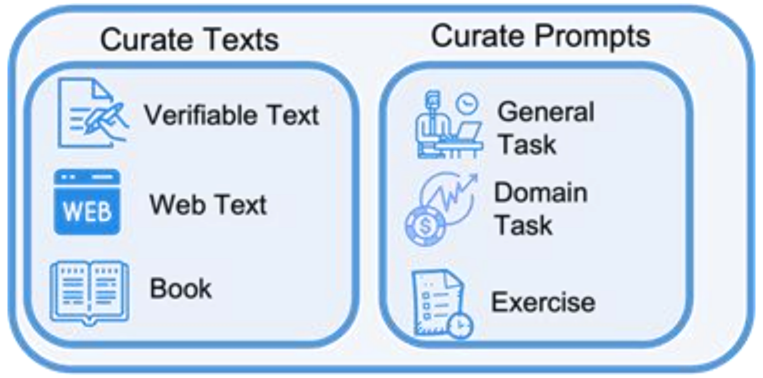

Seed Data

What seed data should be used to implement your training recipe?