



LoRA Soups: Merging LoRAs for Practical Skill Composition Tasks

Akshara Prabhakar¹, Yuanzhi Li², Karthik Narasimhan¹, Sham Kakade³, Eran Malach³, Samy Jelassi³



Q. Given LoRAs trained on specialized domains, is it possible to merge them to effectively solve a new problem that requires a combination of these domains?

Motivation

- Observation #1:** Many target tasks can be decomposed into subtasks requiring diff skills
- Observation #2:** It is difficult to procure expert-curated training data for specialized tasks
- Observation #3:** LoRA is a cost-effective way to make a model learn a skill

RQ: Can we combine skill LoRAs in an efficient way to solve binary skill-composition problems?

Our approach: We introduce the skill composition problem and model merging technique **Learnable Concatenation (CAT)**

Skill Composition in NL

A target task that involves combining multiple skills, each skill coming from a single model

Challenging task

- **OOD** since the skill LoRAs are trained on datasets that are very different from the target task
- model needs to “know” the different skills and also needs to **understand the appropriate context** for applying each skill

Settings

We analyze 4 practical settings

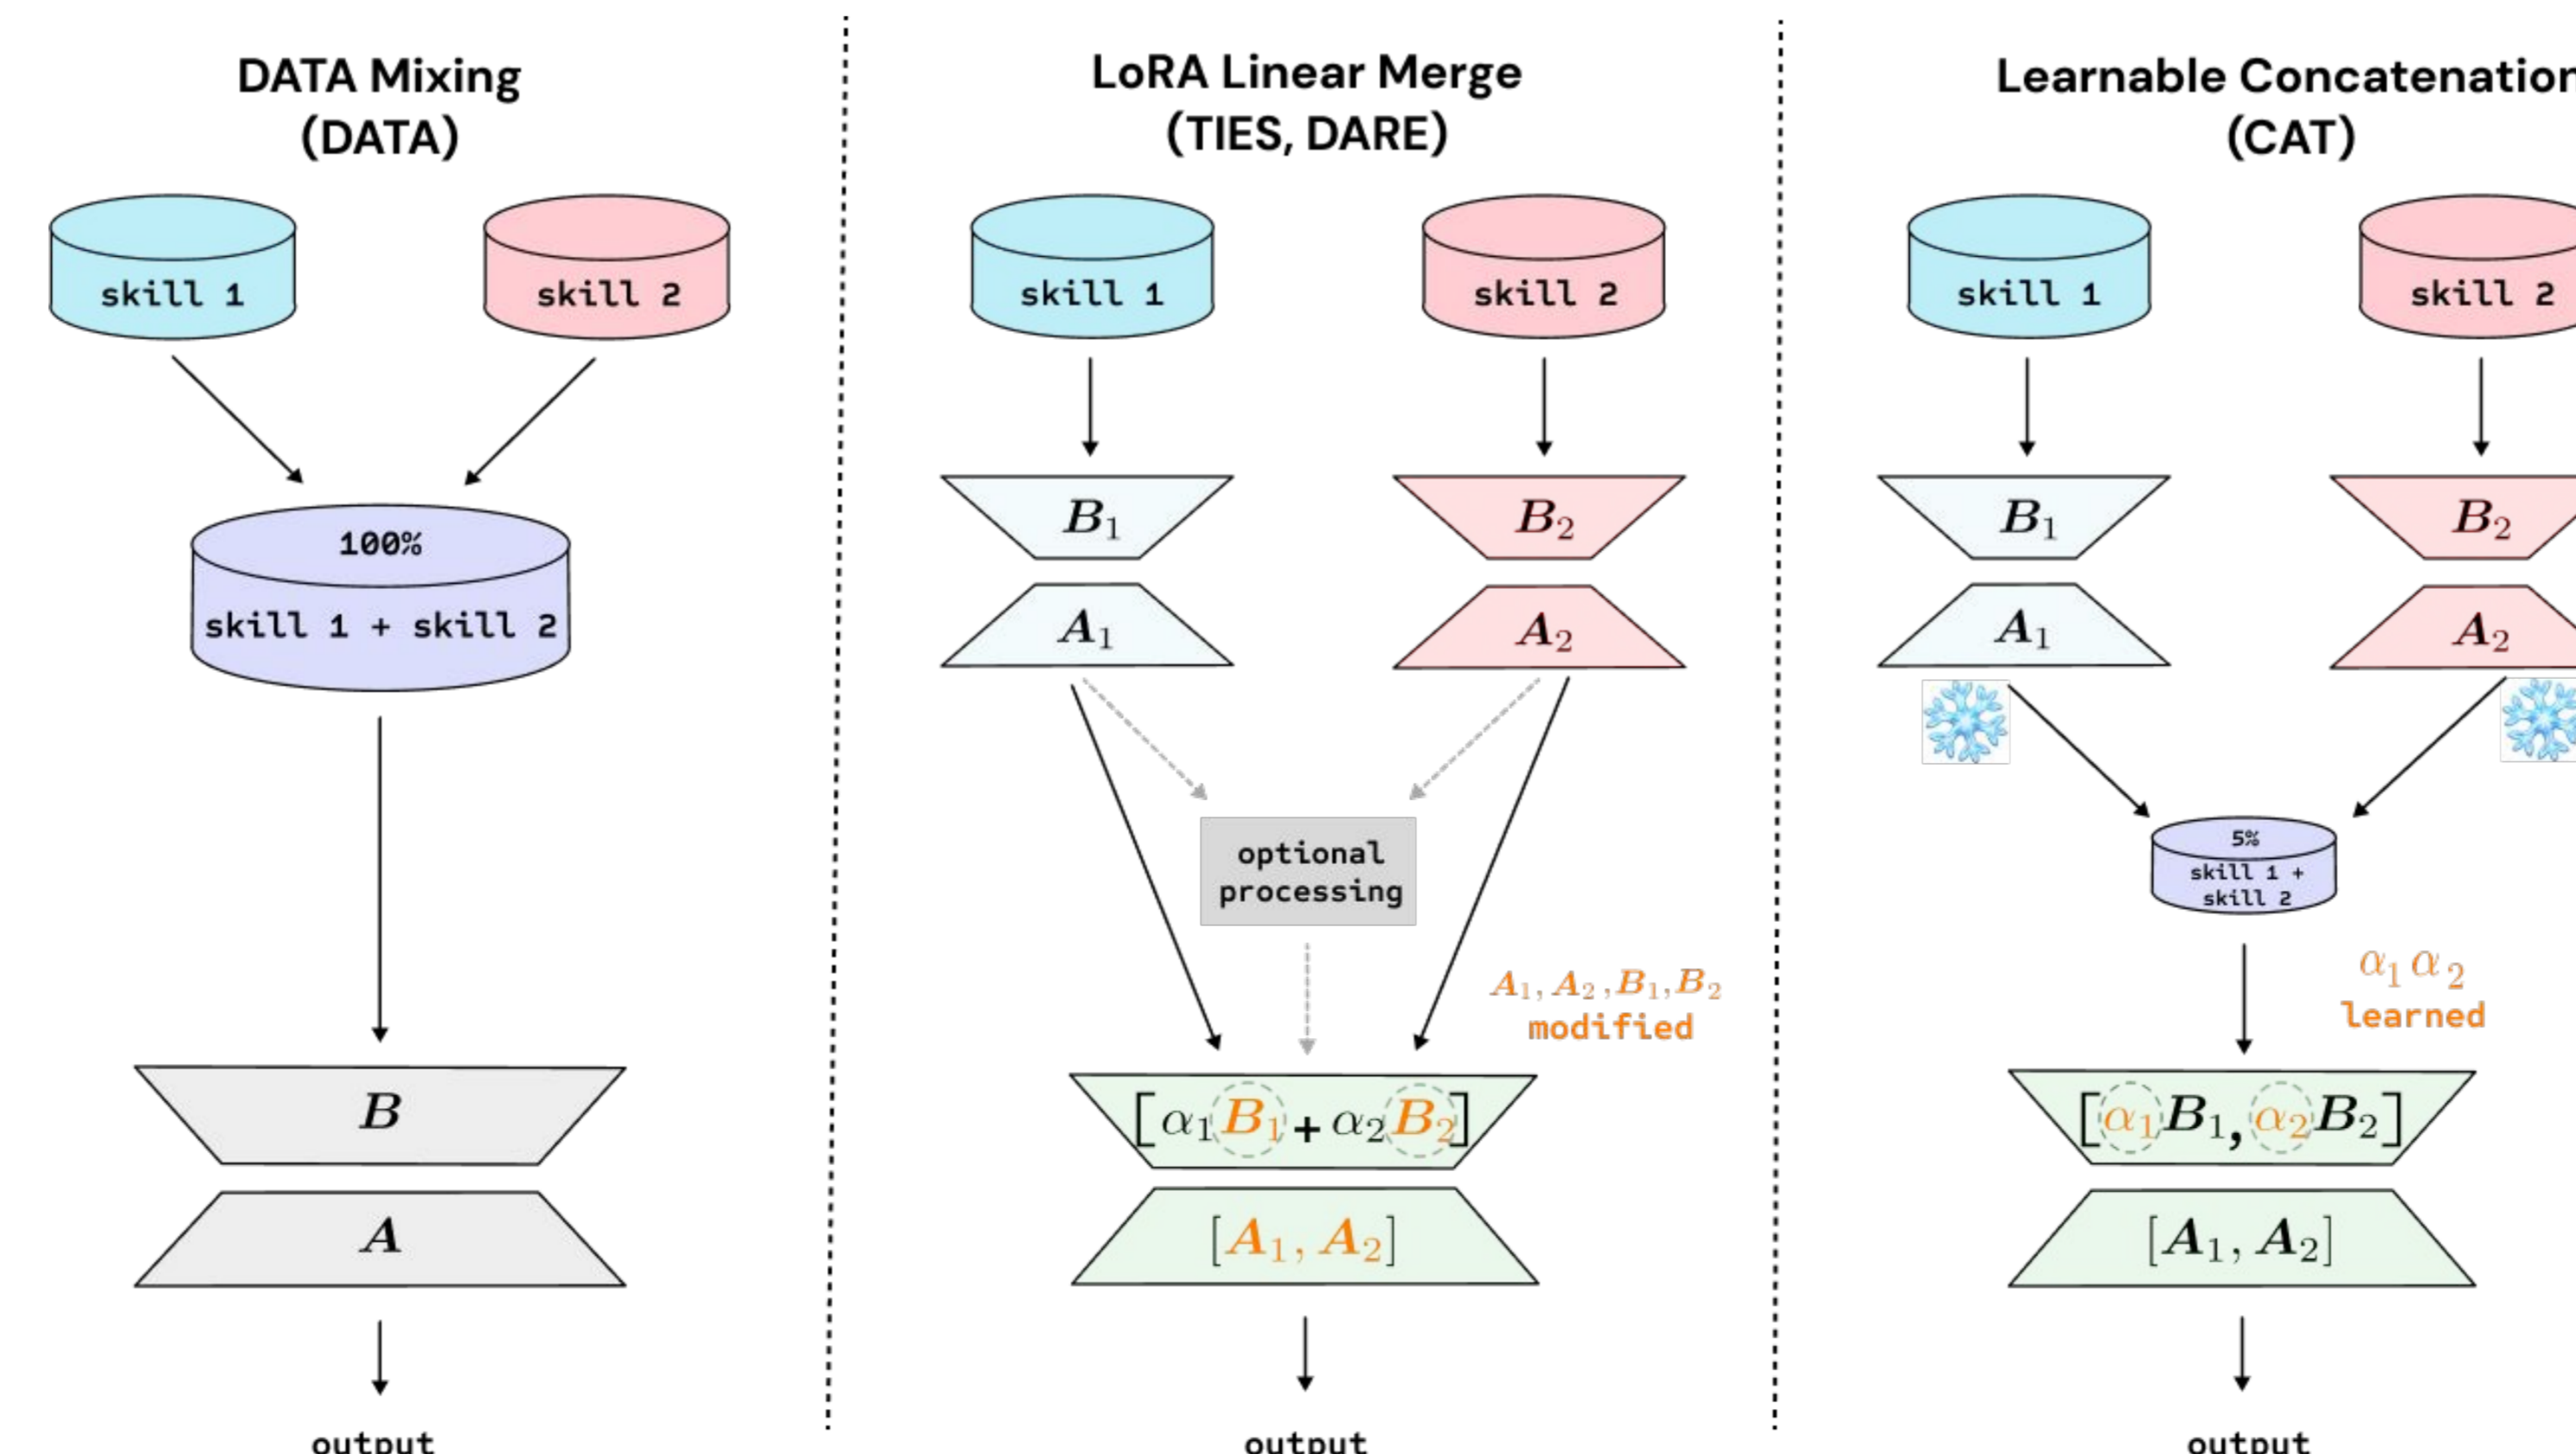
- solving hard math problems with code
- domain specific question-answering bots
- comprehension on technical documents
- robustness to prompt formats

Task	Skills Composed	Test Dataset	Evaluation Metric
Hard math-word	{math, code}	GSM-Hard [15] (1319)	Exec. Accuracy
Question-answering	{biology/game, closed-book QA}	Open-source book/manual SB.2 (95)	Accuracy (GPT-4)
Reading comprehension	{biomedical, open-book QA}	BioASQ-QA [31] (200)	Elo Rating (GPT-4)
Prompt robustness	{prompt _i , prompt _j }	SuperNaturalInstructions [58] (417)	Accuracy

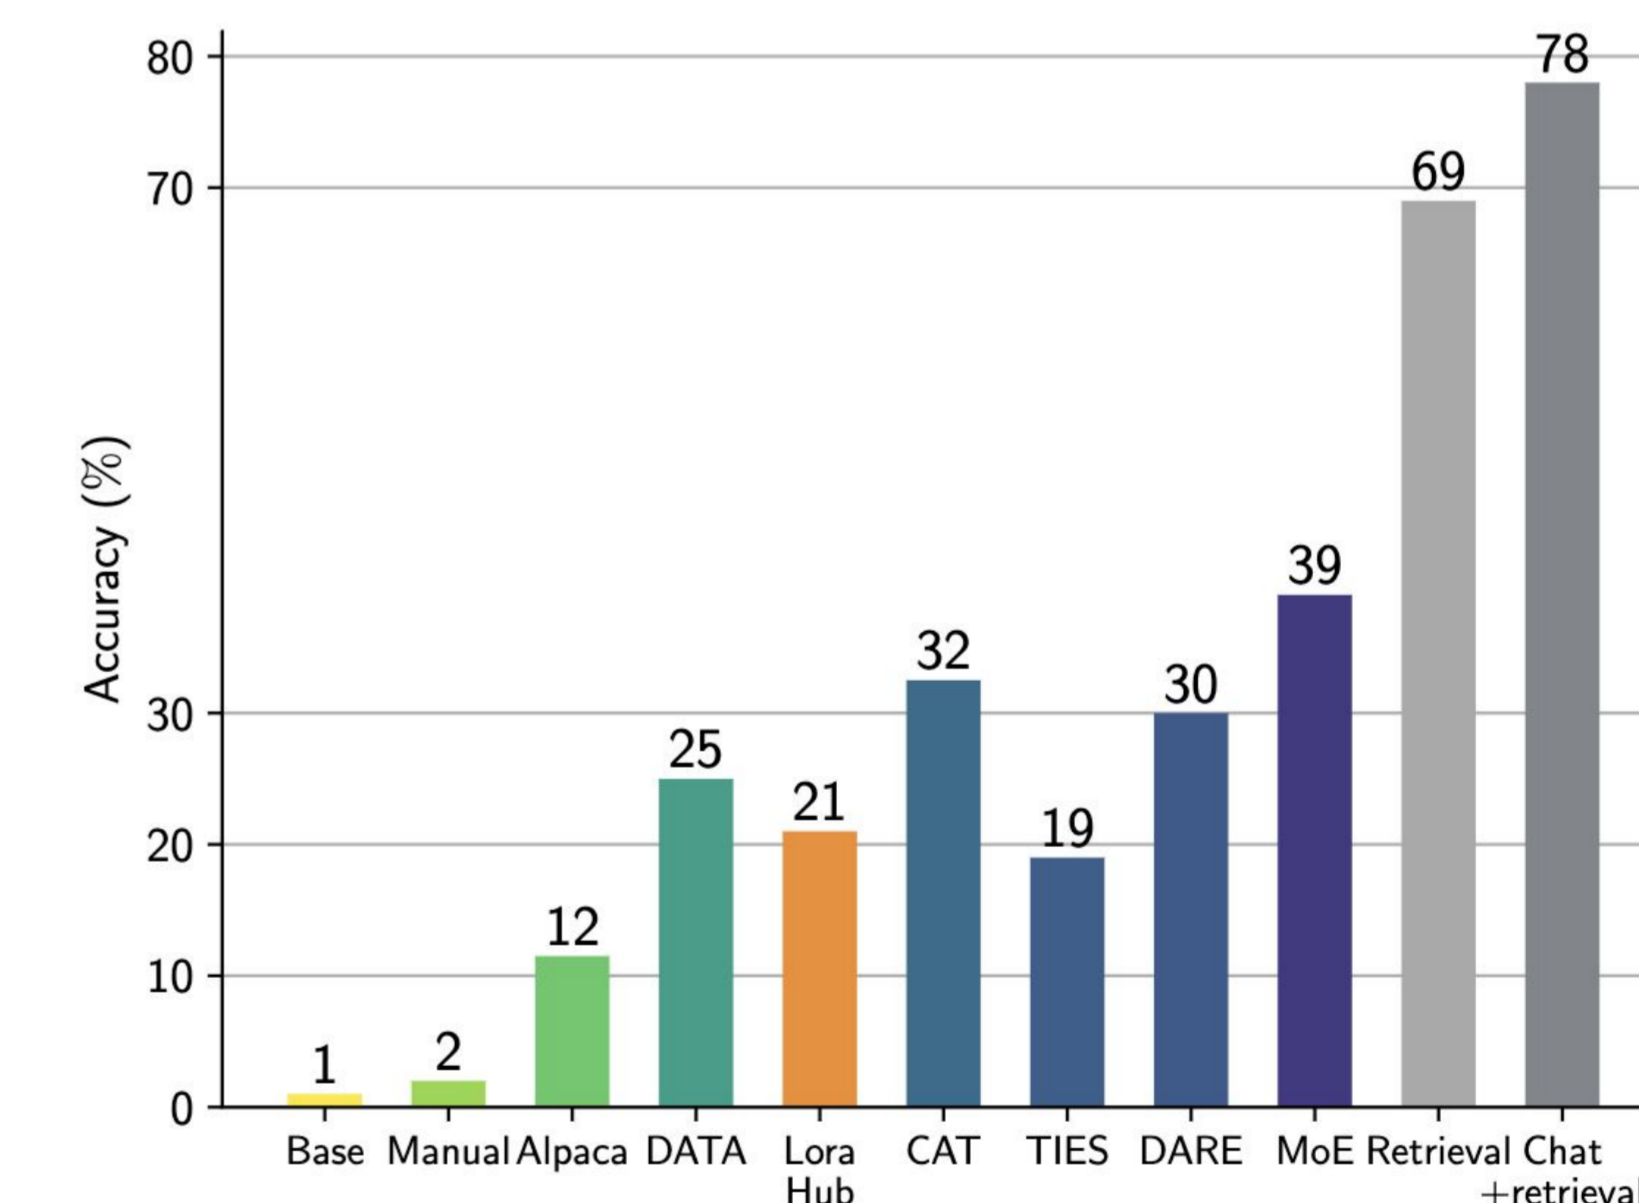
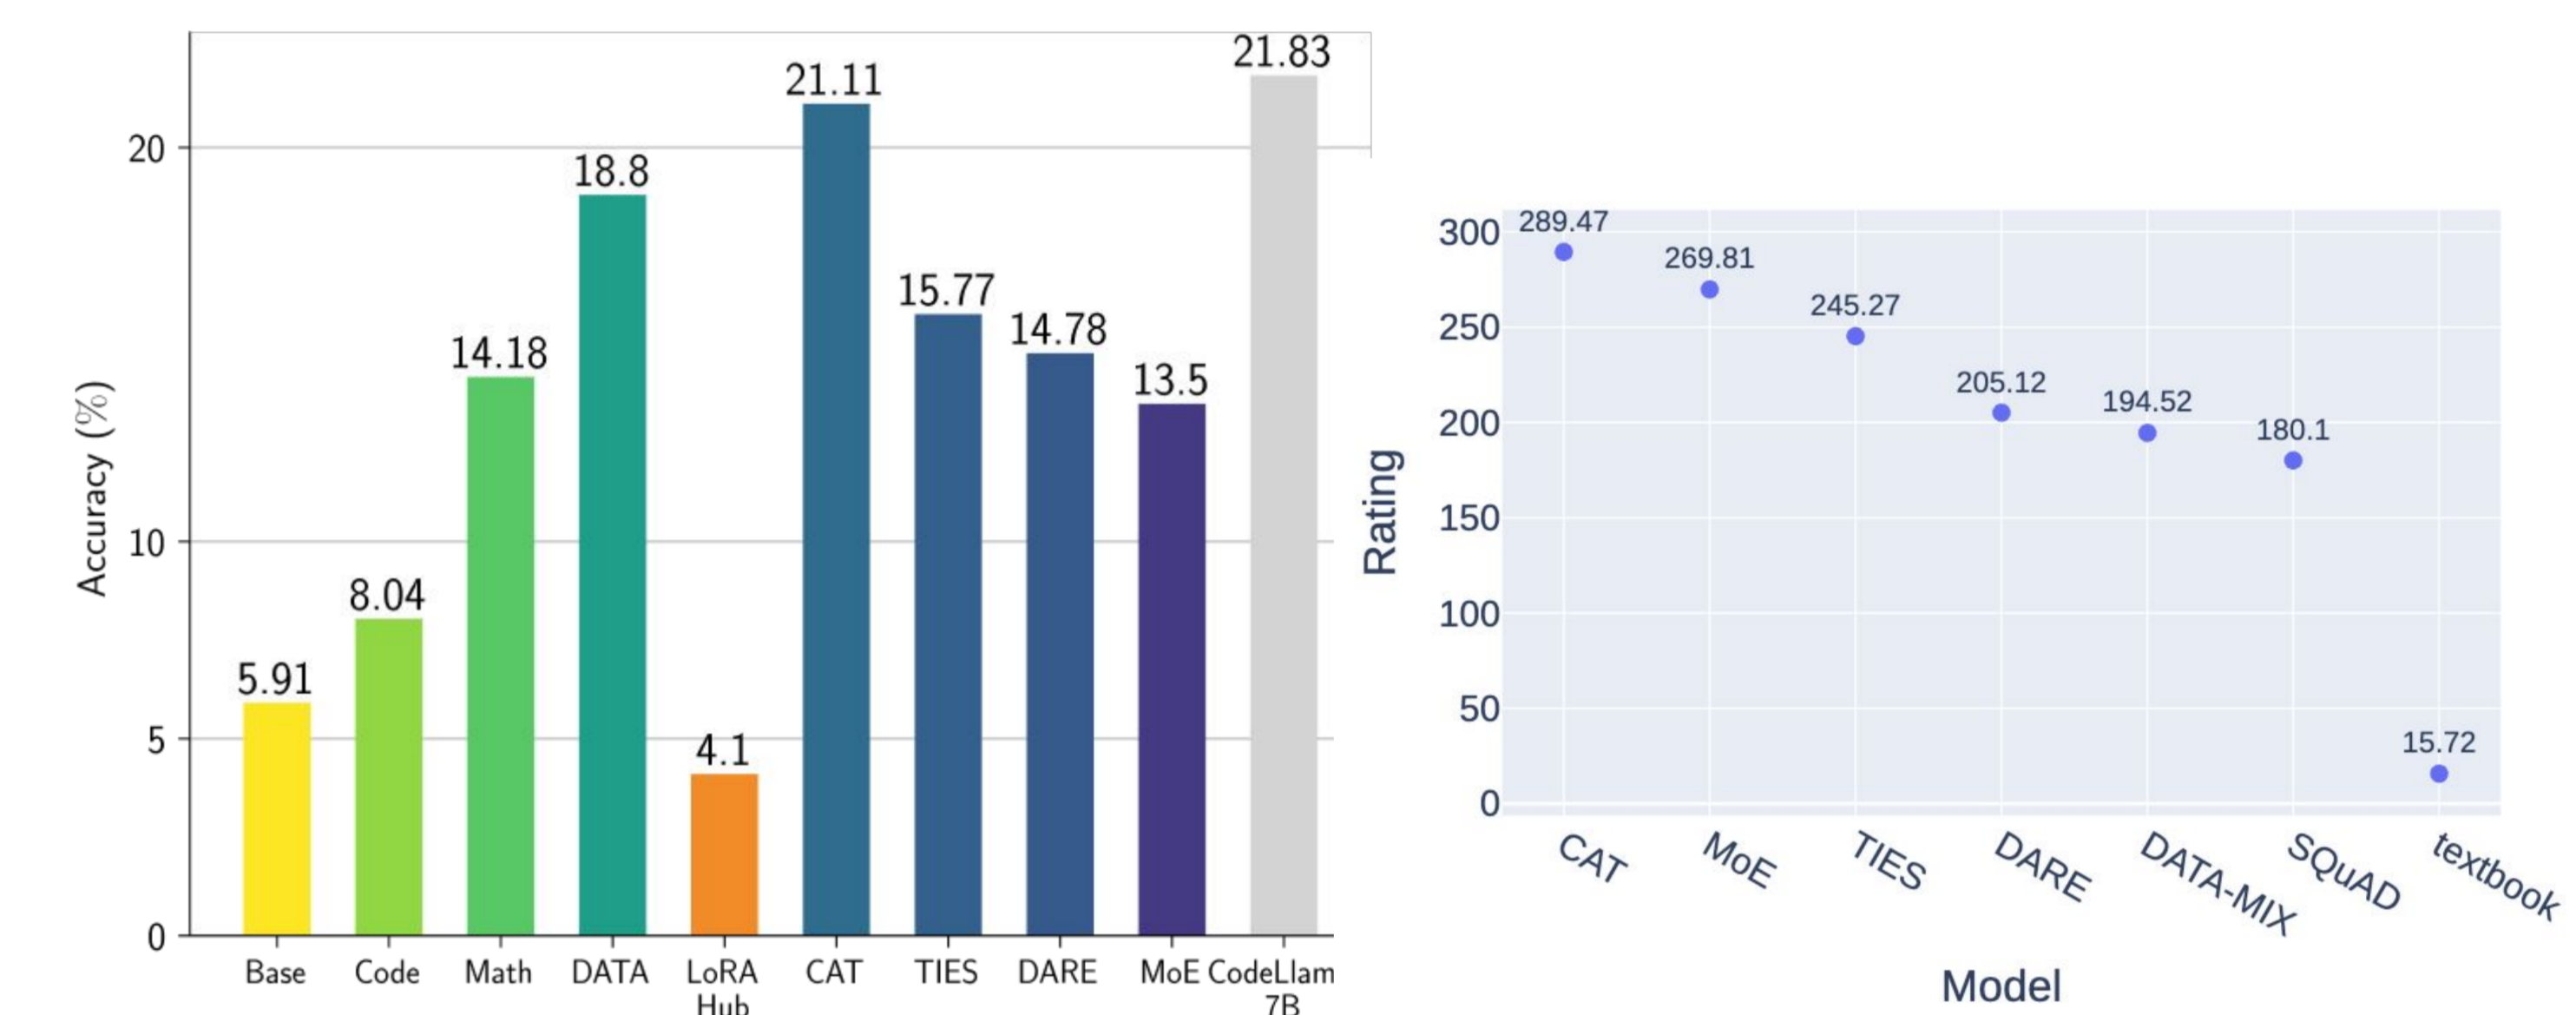
Learnable Concatenation (CAT)

Traditional approach: Training over a mixture of the two datasets (DATA-MIX)

CAT: a simple weighted average of the skill LoRAs by cheaply learning the weights layer-wise



Results



CAT outperforms DATA-MIX and other merging methods (TIES, DARE) on most settings.

Key Takeaways

Model mixing via a small, learnable router on top of independently trained skills generalizes better than dataset mixing for binary tasks.

Future: identifying practical tasks decomposable into and more detailed analysis on more than two skills