LANGUAGE-GUIDED EXPERTISE EVOLUTION FOR PROTEIN OPTIMIZATION

Xingyue Liu, Zijie Xing, Runze Wang, Luoming Hu, Yanming Shen
Dalian University of Technology

Introduction

- LLMs show strong general reasoning but are limited in specific hard scientific problems like protein mutation tasks.
- RL-based fine-tuning is common but suffers from sparse rewards and high computational costs.
- This paper advocates that, for many scientific tasks, the primary bottleneck is not insufficient model capacity, but insufficient explicit expertise.
- We propose adapting LLMs by optimizing domain expertise rather than model parameters.

Experiments & Results

- **Task:** Optimizing protein stability using REF 2015 as the main reward, and its decomposition terms as the dense reward.
- **Dataset**: Processed version of the Brenda dataset.
- **Evaluation:** Pass@1 metric, defined as the success rate of improving Rosetta Score over a held-out validation set of 200 proteins.

Method	Model	Pass@1(%)				Rollout
		Seed 42	Seed 1145	Seed 1919	Average	
Vanilla	LLaMA 8B	36.0	38.5	34.5	36.3	-
	Deepseek	41.5	44.5	43.0	43.0	-
	Qwen-plus	40.5	43.5	45.5	43.2	-
	Qwen-max	43.0	43.5	46.5	44.3	-
RL	Pro-1 8B	43.5	44.0	44.0	43.8	> 5000
Our method	Qwen-max	**51.5**	**50.0**	**56.0**	**52.3**	**2037**

- **Results:** Average Pass@1 of 52.3%, 8% improvement over the backbone model.
- Superior performance over RL-based Pro-1 Hla (2025) with lower resource consumption.
- Achieves better results with 2037 rollouts, compared to Pro-1's 5000+ rollouts.

Overall framework

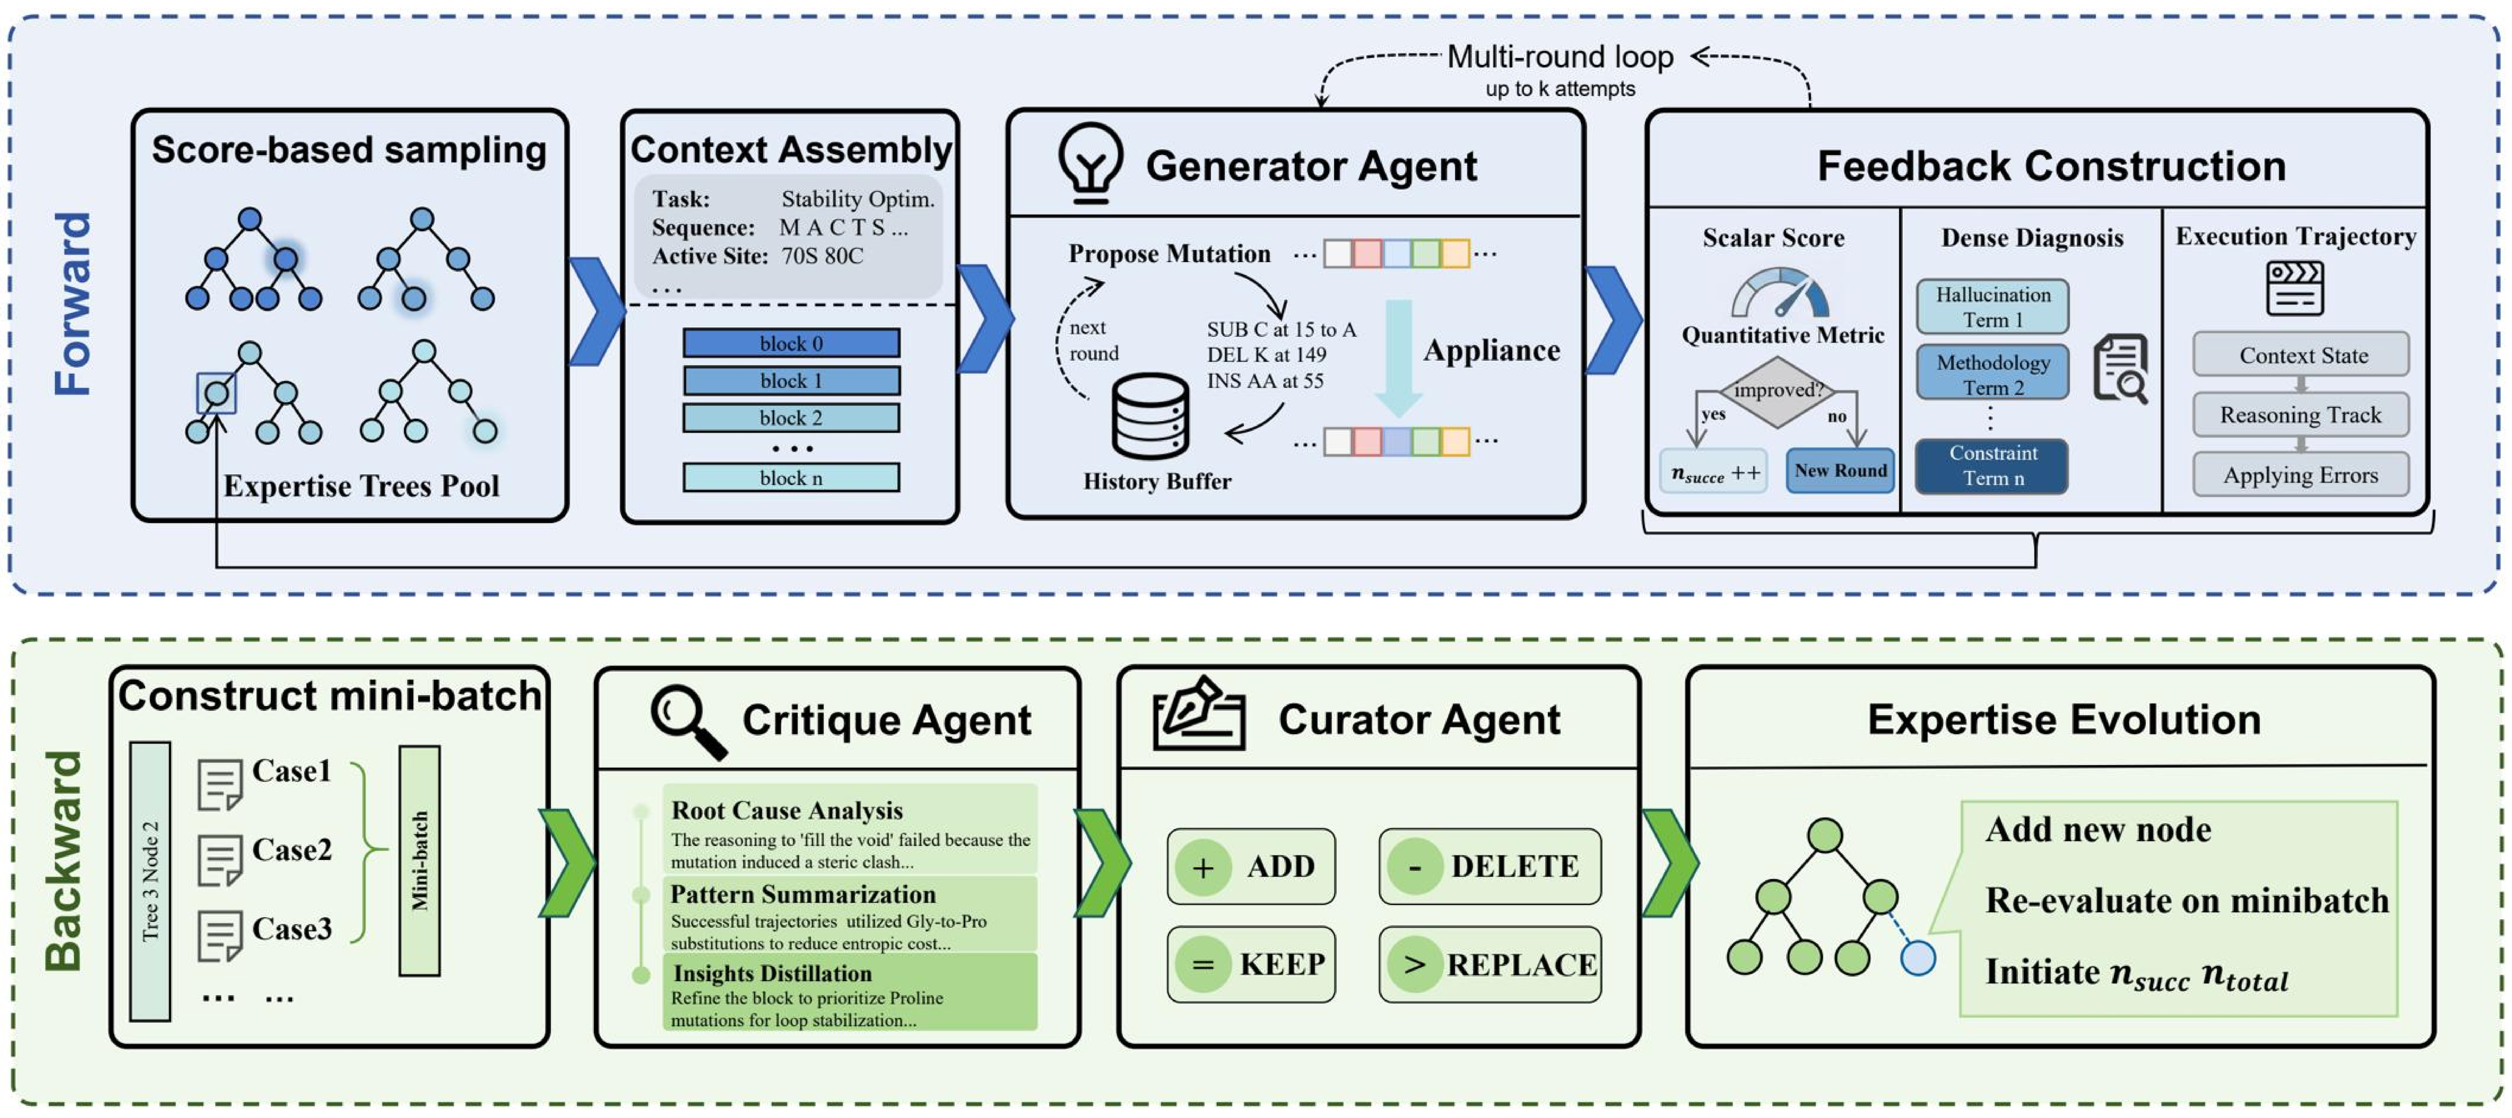

Mimics the scientific cycle of hypothesis generation, verification, and knowledge accumulation.
Generator: Proposes mutations to explore the sequence space based on sampled domain expertise.
Critique: Analyzes experimental feedback to distill root-cause insights from successes and failures.
Curator: Refines knowledge blocks via incremental updates to codify winning patterns.

Expertise Pool Management

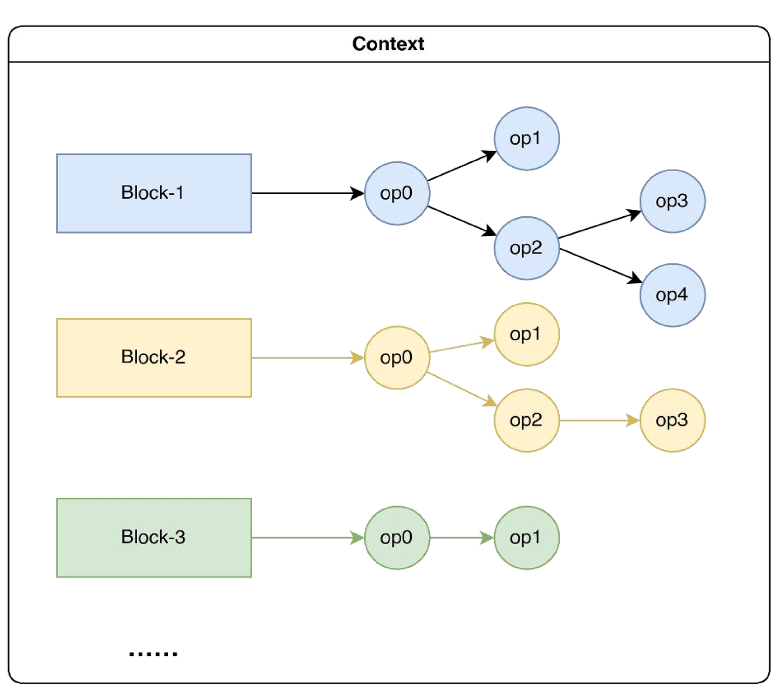

- **Context Blocking**: The learnable context is divided into multiple evolving expertise blocks.
- **Evolutionary Expertise Tree**: Each modification by the Curator creates a child node that inherits expertise from its parent.

1. Evolution via 4 operations: ADD, REPLACE, DELETE, and KEEP .
2.Node Evaluation: Success rates are tracked via a Beta Distribution based on trial history.

Algorithmic Innovations:

1. Uses Thompson Sampling with Depth Bonus to explicitly bias the tree search, favoring deeper nodes that have undergone more refinement iterations.
2. Employs the Lower Confidence Bound of the posterior distribution to deterministically select nodes with high success rates and low variance.

Conclusion & Discussion

We present a framework that:

- Reframes protein optimization as an explicit optimization over external expertis rather than model parameters.
- Evolving expertise pools achieve higher success rates and faster convergence than reinforcement learning baselines, while substantially reducing computational cost.
- Explicit expertise evolution offers a scalable and interpretable alternative to parameter-heavy fine-tuning for scientific optimization tasks.

Contact us:

Email: {aroma, kokojjj}@mail.dlut.edu.cn