GBC: Gradient-Based Connections for Optimizing Multi-Agent Systems

Xiaocheng Yang, Abdulrahman Alrabah,
Dilek Hakkani-Tür, Gokhan Tur

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

CONVAI

Motivation

Multi-agent systems are hard to debug

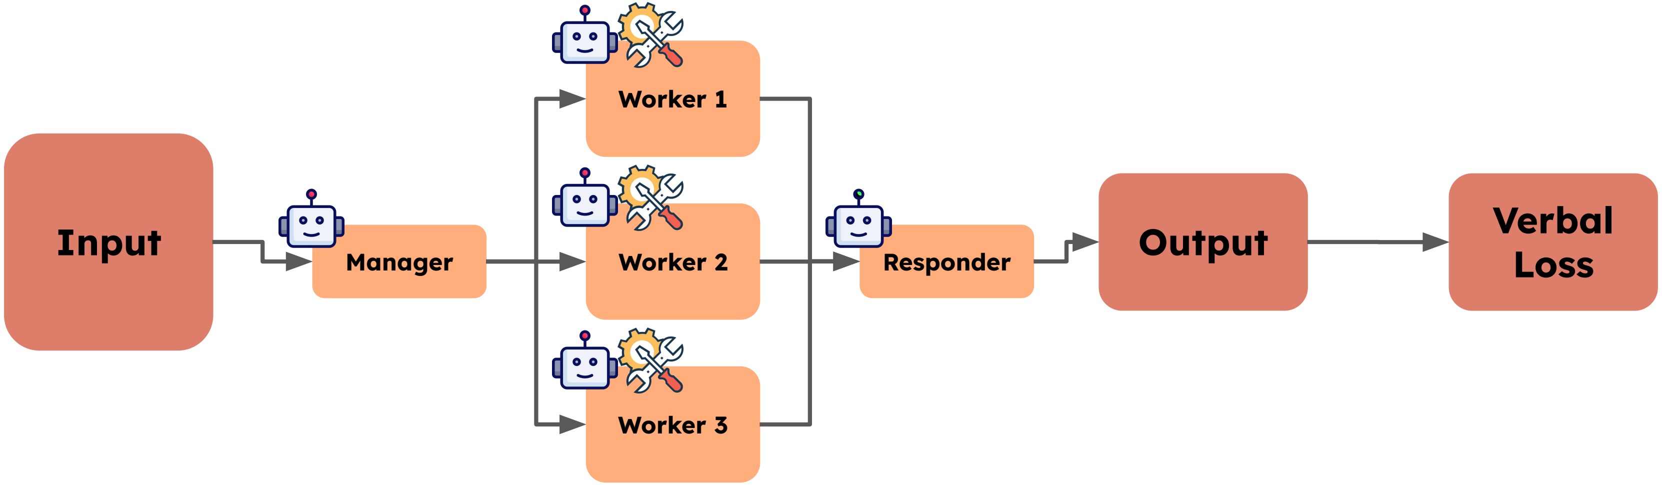

A final failure often hides where the failure actually started.

But optimization needs: which agent, step, or token caused the error?

Without attribution, prompt updates become broad and noisy.

Intuition

Core idea: build credit paths backward from loss

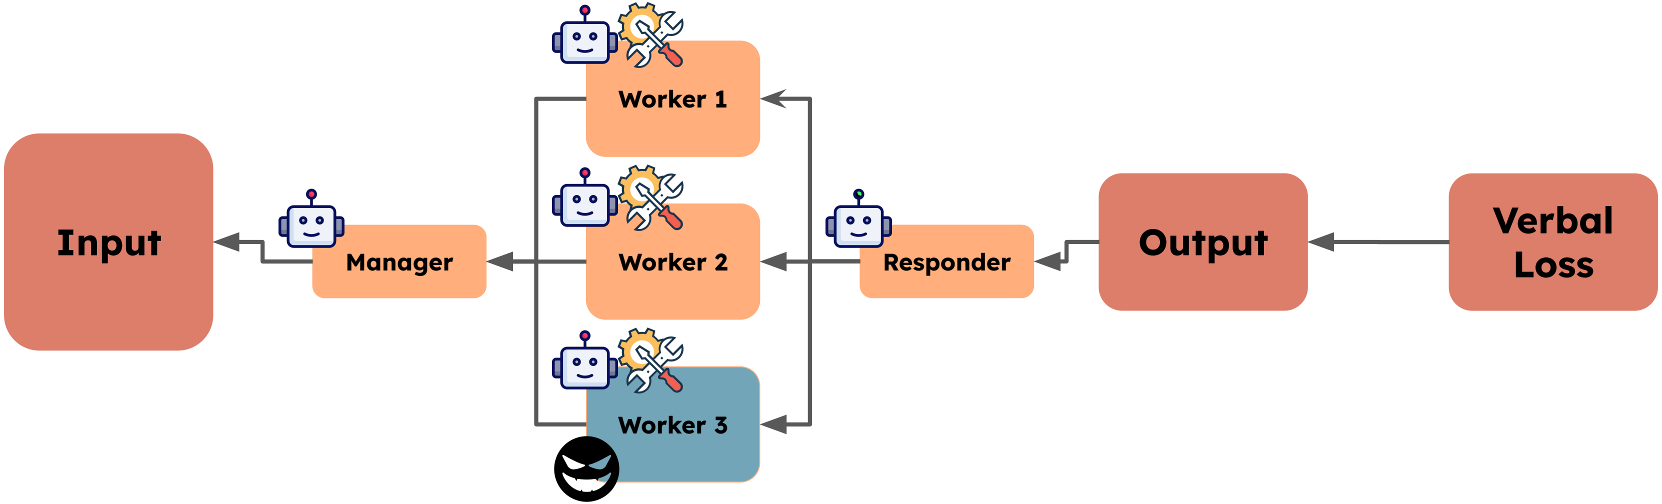

- Can we treat a multi-agent workflow as a computational graph, then estimate how each predecessor output influences downstream generation?

Method

GBC pipeline: Forward→Attribution→Backward→Update

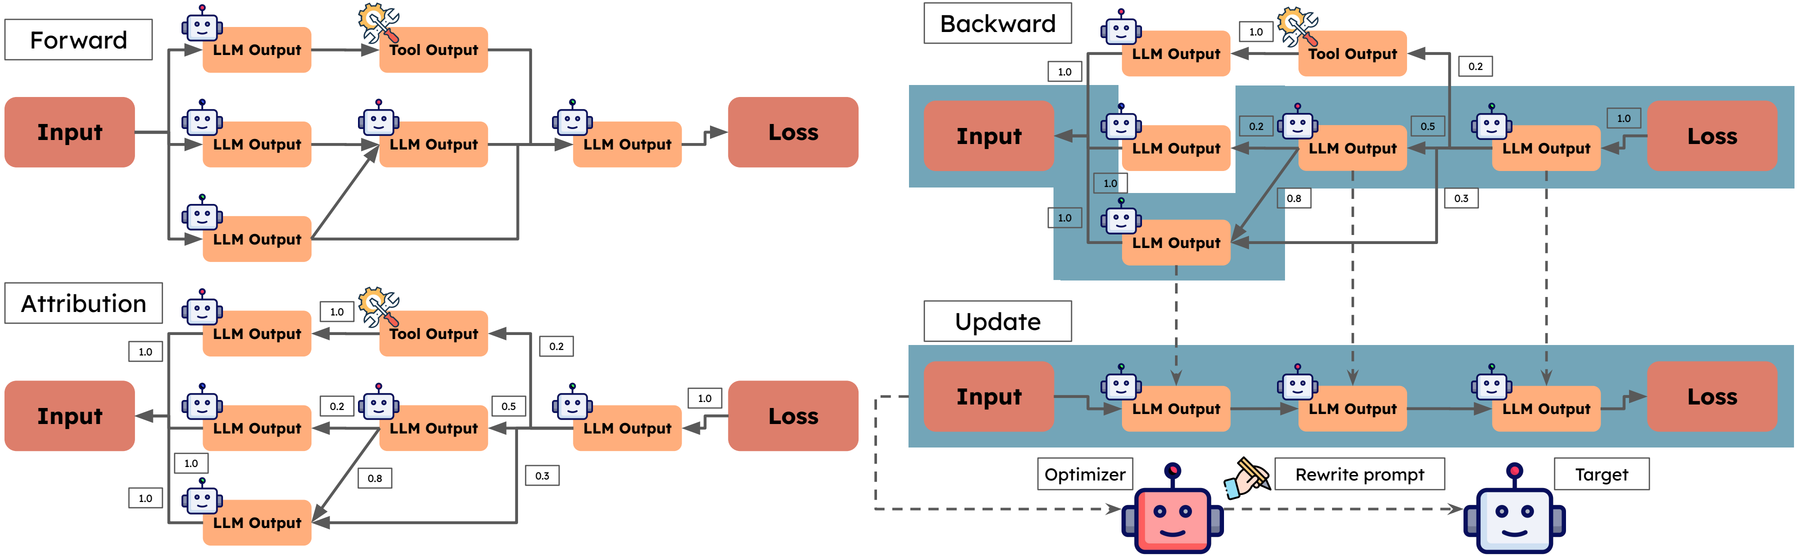

1. Forward

Run the agent graph and produce the output.

2. Attribution

Compute connection weights.

3. Backward

Propagate the loss through the attribution graph.

4. Update

Rewrite the target prompts with an optimizer LLM.

Method

Gradient-based connections quantify influence

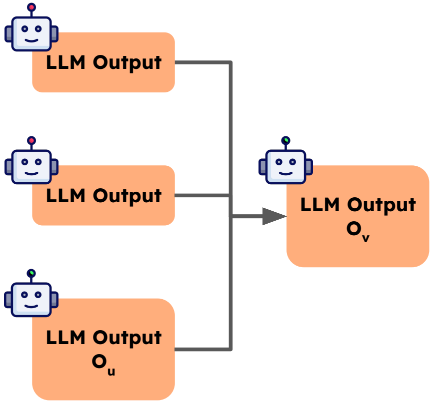

O_u Example

The hotel named Bedouin has been booked. Reference number is FRQZWQL2. Please inform the user of the information.

Predecessor

O_u

(Input sequence length × 1)

Token-level gradient

$\nabla \prod_{w \in O_v} P(w \mid \mathrm{Embed}(O_u))$

(Input sequence length × hidden size)

Token-level weight

Aggregation over dimension of hidden size

(Input sequence length × 1)

- L1 norm

$\| \nabla \prod_{w \in O_v} P(w \mid \mathrm{Embed}(O_u)) \|_{L1}$

- Product with Input

$\nabla \prod_{w \in O_v} P(w \mid \mathrm{Embed}(O_u)) \cdot \mathrm{Embed}(O_u)$

Connection weight

Aggregation over dimension of sequence length

(Scalar)

- Average effect of each token

Avg(·)

- Effect of the most important token

Max(·)

Implementation

Prefix trick makes optimization loop practical

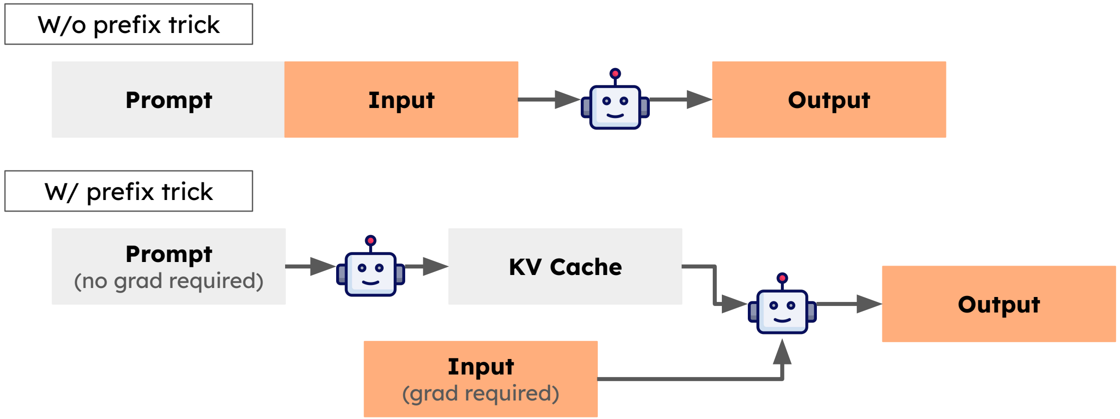

- Prompt tokens are treated as a fixed prefix.
- Gradients are only stored for input tokens.
- Autograd activation memory: $O(n \cdot d \cdot L) \rightarrow O((n - k) \cdot d \cdot L)$
- AgentChord, a practical, lightweight framework is built around this prefix trick.

Gradient-Based Connections for automated multi-agent system optimization.

Handy visualization of Gradient-Based Connections.

Connected to W&B to easily monitor the optimization process.

Experiment

Evaluation: dialogue + tool-use multi-agent systems

CONVAI

MultiWOZ 2.4

Task-oriented dialogue.

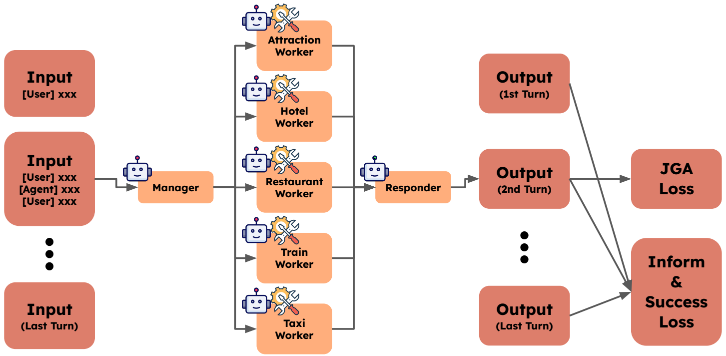

6

Metrics

Inform, Success, JGA, and Slot metrics

10

Steps

Update every 3 samples

τ-bench

Interactive multi-step tool-use with a user simulator.

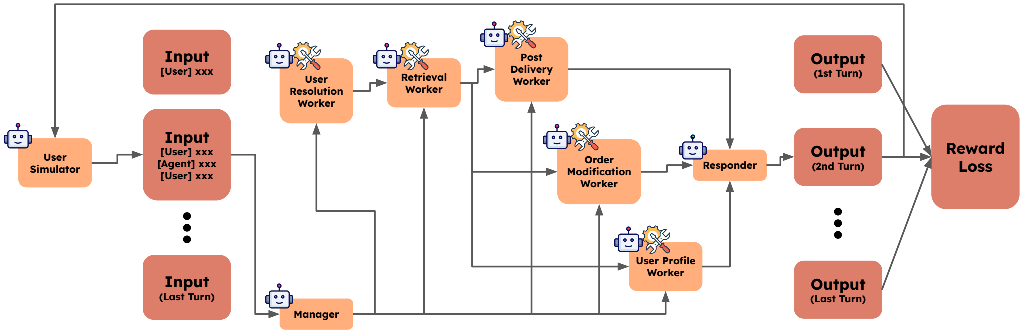

3

Metrics

Action, Output, and Overall Reward

10

Steps

Update after each complete conversation

Result

MultiWOZ: GBC builds a strong TOD system

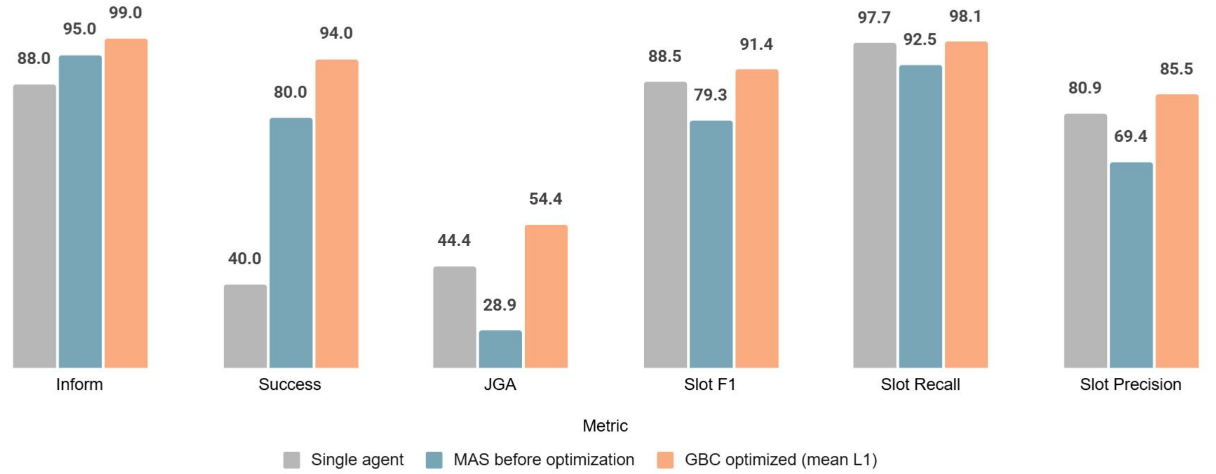

+25.5 JGA

28.9→54.4

+12.1 Slot F1

79.3→91.4

99.0 Inform

Best score

94.0 Success

2nd best score

- Qwen-3-32B multi-agent system before optimization had high Inform/Success but weaker state tracking; GBC fixes the bottleneck.
- The optimized MAS surpasses the Qwen-3-32B single-agent baseline on Inform, Success, JGA, and Slot metrics.

Result

Optimization dynamics: most metrics improve steadily

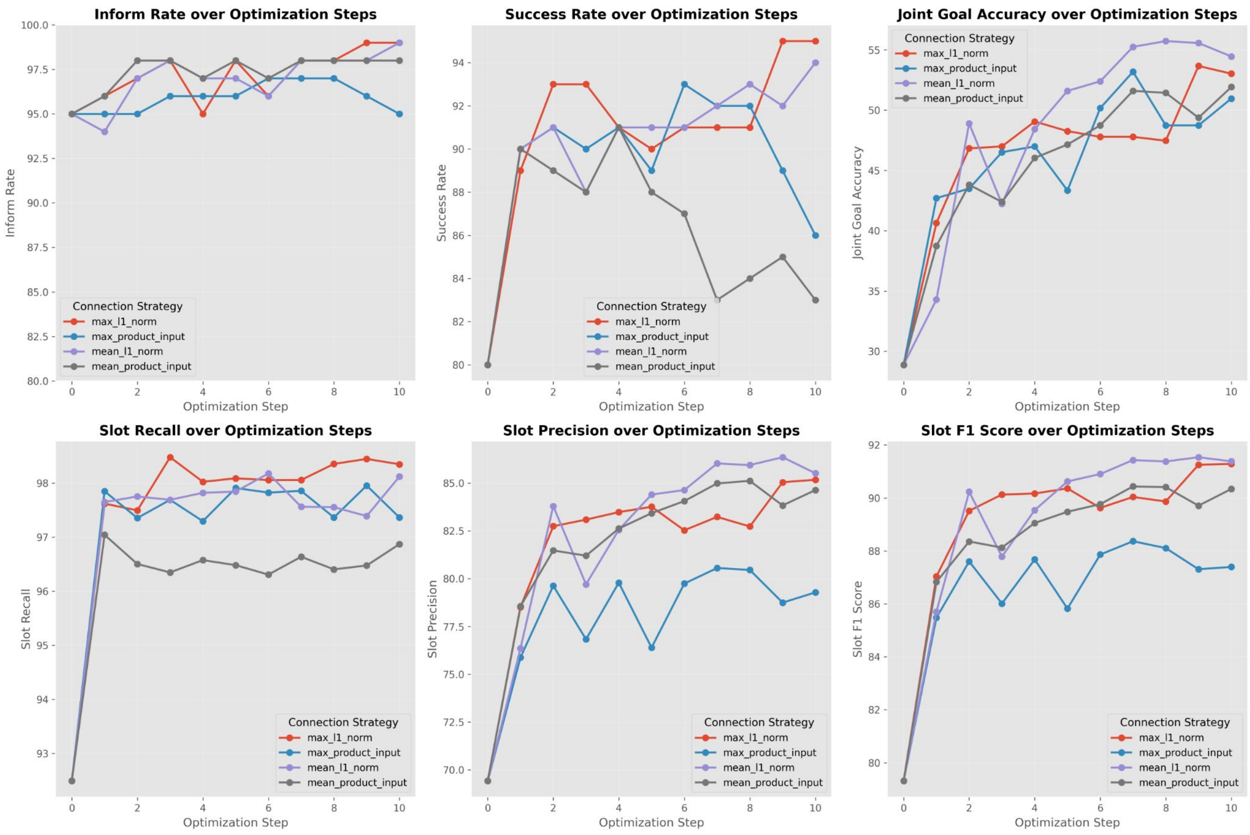

JGA rises and continues trending upward.

Slot precision and slot F1 improve, indicating fewer over-predicted slots.

Success is more variable, reflecting higher requirement of the metric.

L1-norm-based connection weights yield the strongest performance.

Result

τ-bench: GBC improves interactive tool use

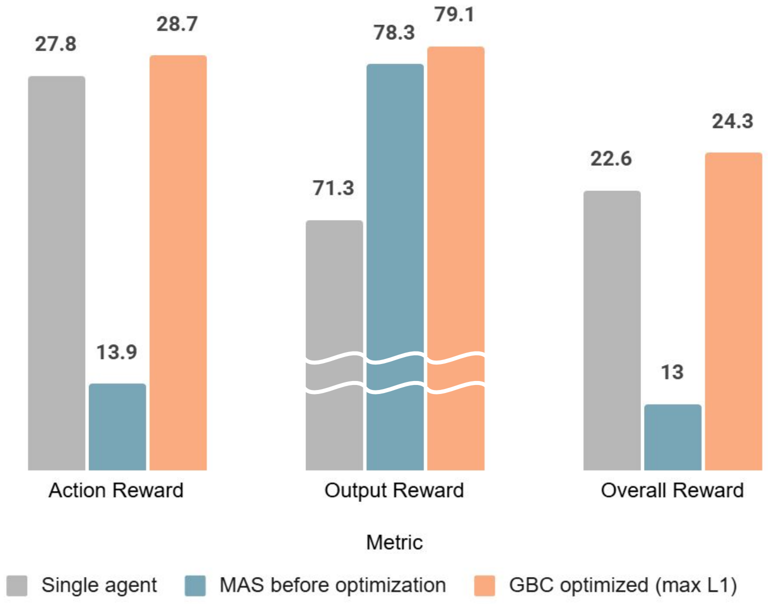

*Output Reward shown with axis break.

Retail

Task Domain

Availability of training data

13.9 → 28.7

Action Reward

Best score

13.0 → 24.3

Overall Reward

Best score

- GBC enables the Qwen multi-agent system to surpass the strong single-agent baseline.
- The gain mainly comes from improved interactive tool-use capability.

Analysis

Diagnostics: GBC shows what still fails and why

MultiWOZ 2.4

Cross-domain errors, information omission, and over-prediction dominate.

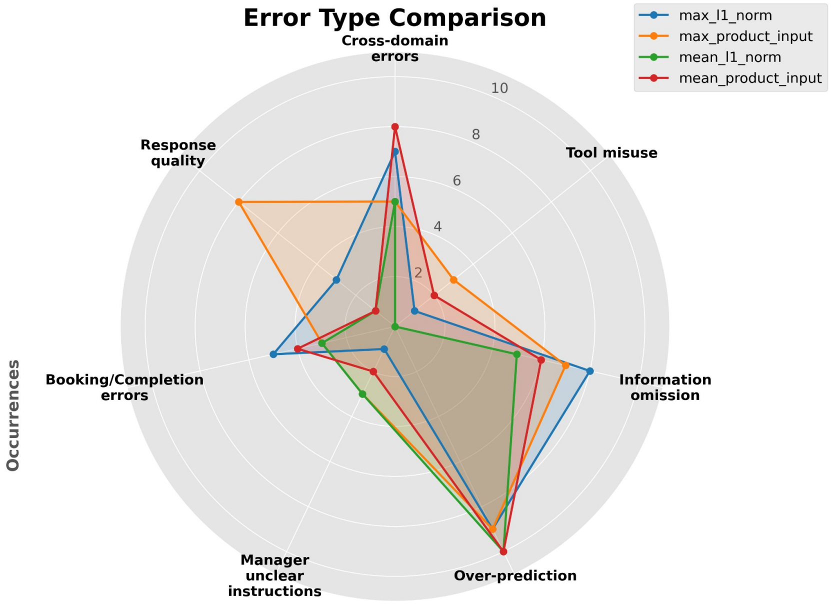

τ-bench

Retrieval and identification failure, tool misuse, and premature escalation or stalling dominate.

Conclusion

- **GBC makes multi-agent optimization more targeted by assigning credit across agents using gradient signals at token level.**

Model the MAS as a computational graph.

Use gradient-based connections to construct attribution trajectories.

Update only the prompts implicated by those trajectories.

Empirical message

GBC improves MultiWOZ and τ-bench performance, and can enable optimized multi-agent systems to surpass strong single-agent baselines.

Open directions

More efficient gradient approximations, improved verbal loss design, reinforcement learning integration, and dynamic agent topologies.

GBC:

Gradient-Based Connections for Optimizing Multi-Agent Systems

Code

Paper

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

CONVAI