# MAGDi: Structured Distillation of Multi-Agent Interaction Graphs Improves Reasoning in Smaller Language Models

Justin Chih-Yao Chen*, Swarnadeep Saha*, Elias Stengel-Eskin, Mohit Bansal
UNC Chapel Hill

## Motivation & Contributions

- Multi-agent interactions with LLMs **improve reasoning but are costly** and lack a unified model for efficient inference.
- We propose MAGDi, a structured distillation method which distills the interactions between multiple LLMs into smaller student models.
- The teacher LLMs interactions are represented as graphs.
- Tested on seven benchmarks, MAGDi boosts smaller models' reasoning abilities and achieving significant efficiency gains.
- MAGDi also shows improved generalizability, scalability and diversity.

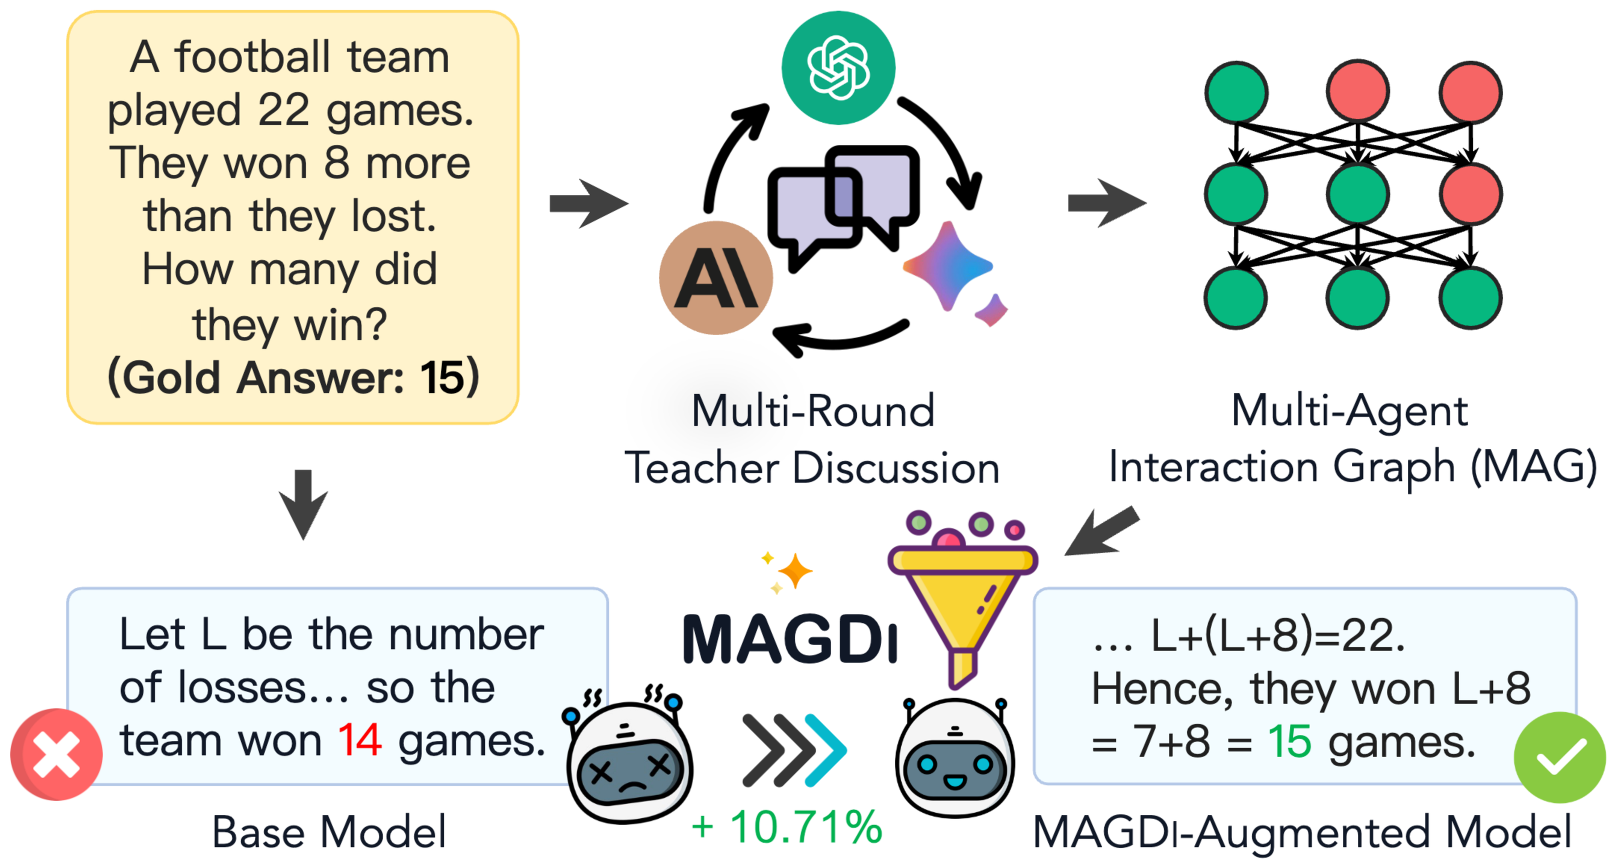


## Problem Setup

We present Multi-Agent Interaction as a graph (MAG). We capture rich knowledge from MAGs via the following four levels of MAG components.

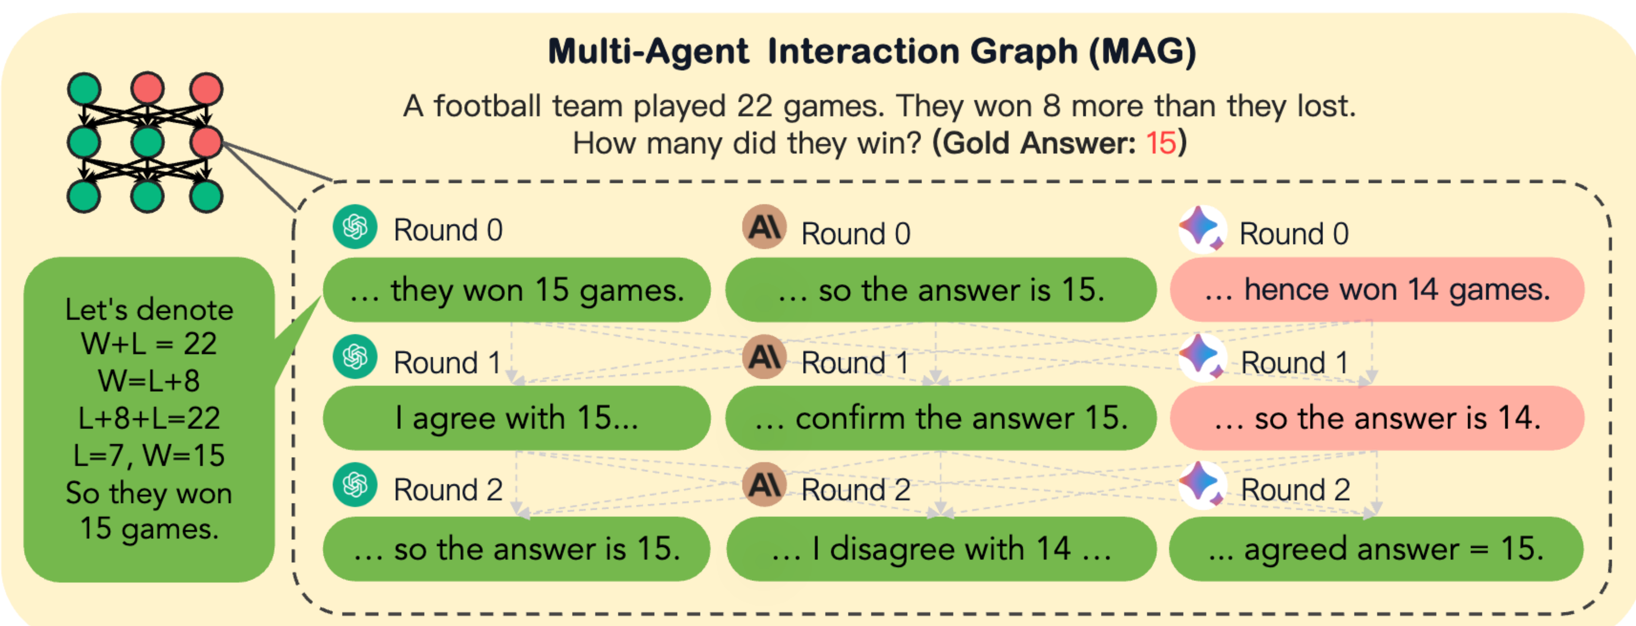


- **Level 1:** Learning from multiple teachers. The student learns from the correct reasoning of multiple teachers, rather than one.
- **Level 2:** Learning from teacher interactions. The student learns from both pre- and post-interaction data between multiple teachers.
- **Level 3:** Learning from negative reasoning. The student additionally distills from negative or incorrect reasoning from the teacher models.
- **Level 4:** Learning from structure. The student learns from the output and graph-structure of teacher LLM interactions.

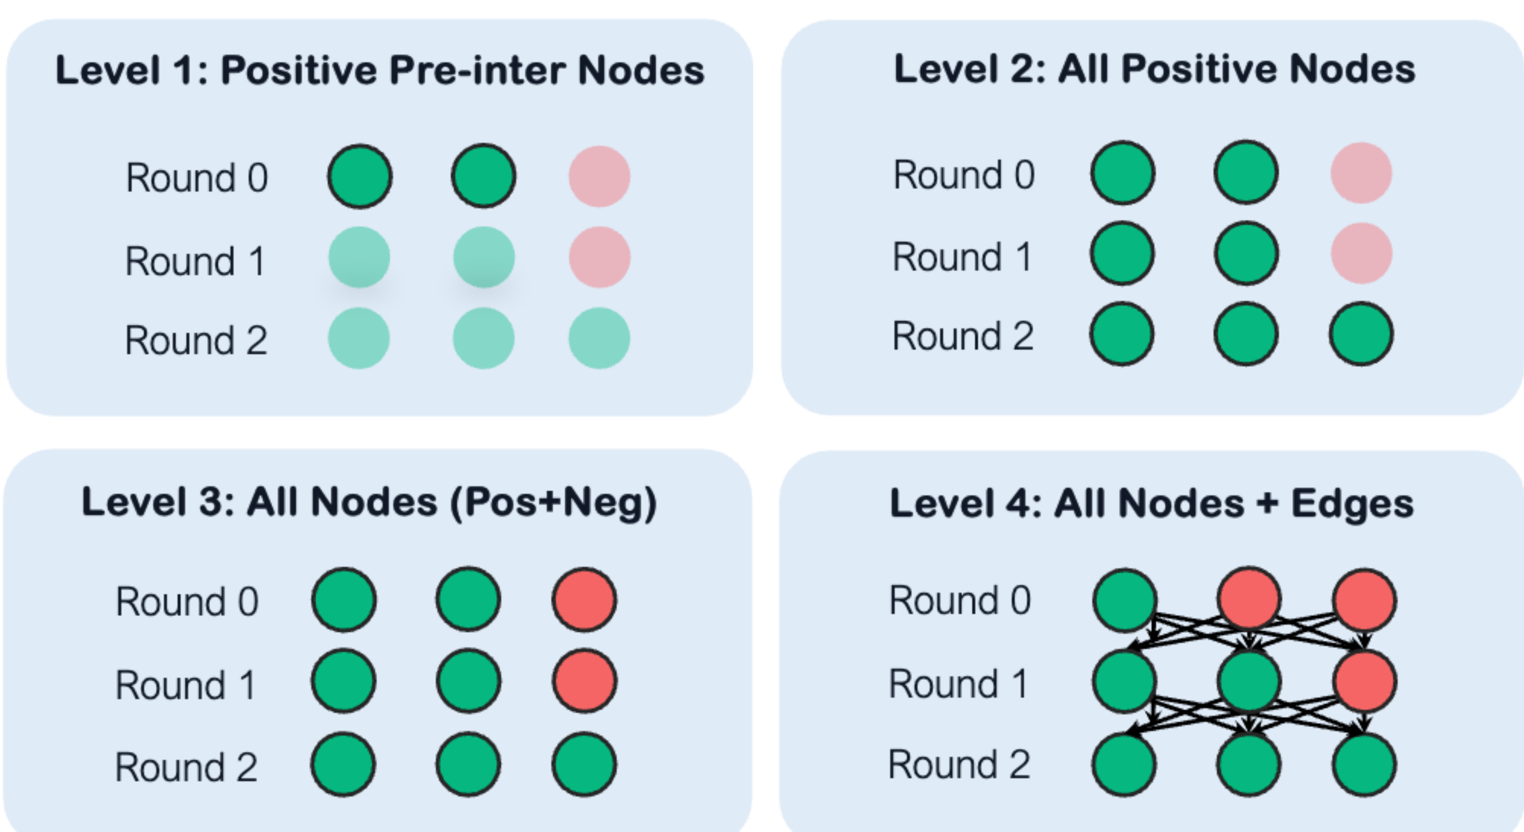


## Methodology

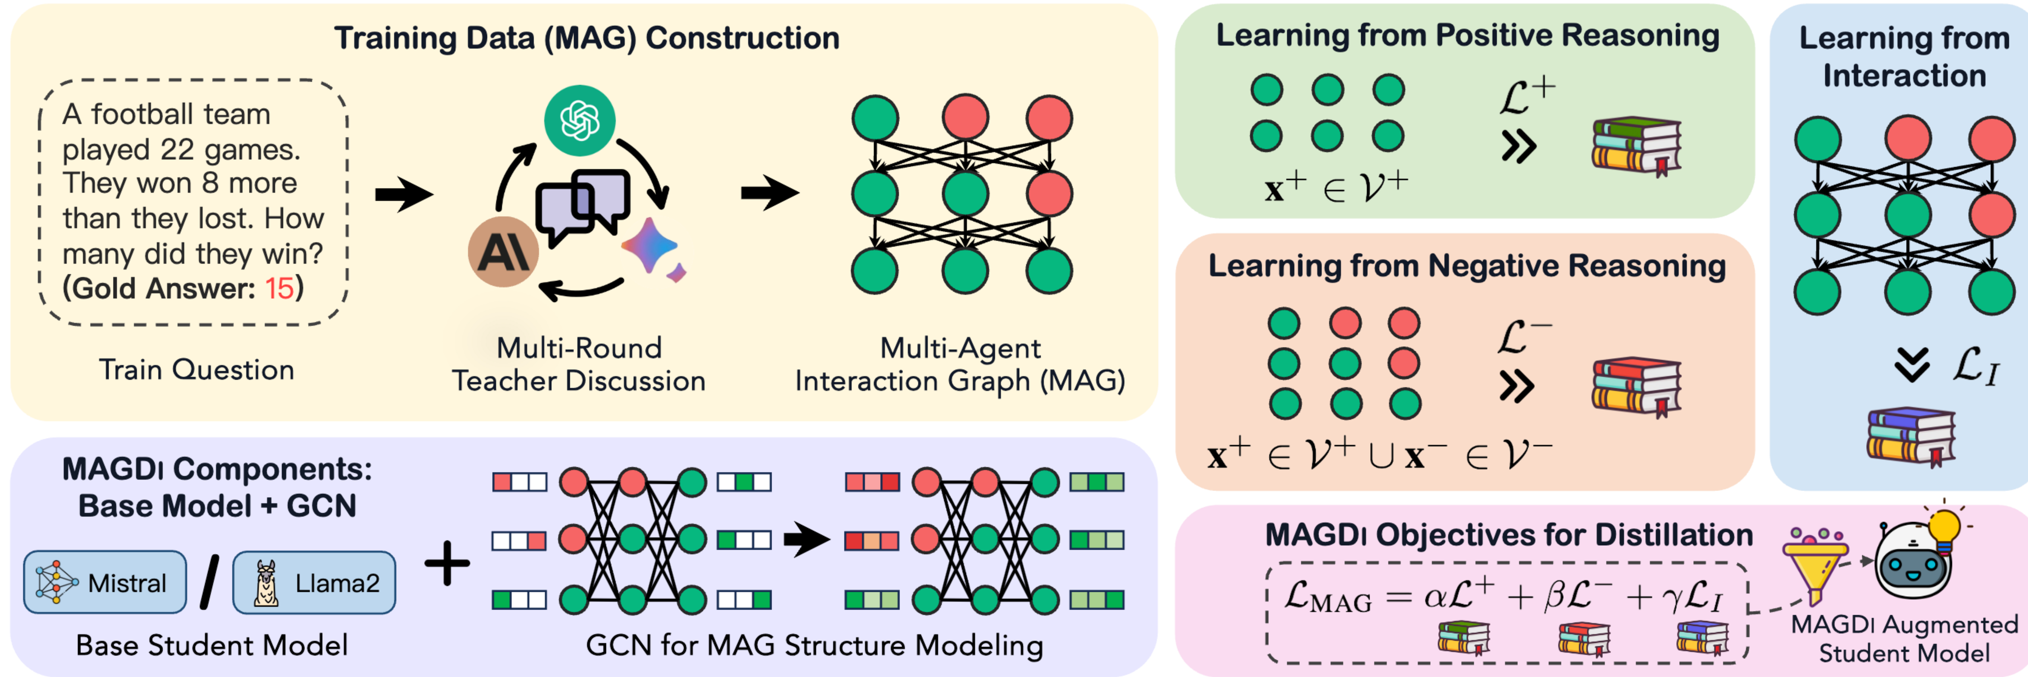


We employ three objectives to learn increasing levels of interaction structure in a MAG.

- Learning from Positive Reasoning: Next-token prediction
- Learning from Negative Reasoning: Margin-based ranking loss
- Learning from Interaction: Node classification using a GCN

## Main Result

| Distillation Data | Distilled Model | Datasets: StrategyQA | CSQA | ARC-c | GSM8K | MATH | Average Acc |
|---|---|---|---|---|---|---|---|
| - | Mistral-7B-Instruct | 61.57 | 57.89 | 60.32 | 44.05 | 7.02 | 46.17 |
| Claude2 | SIT-Claude2 | 64.39 | 64.18 | 68.24 | 45.34 | 7.24 | 49.89 |
| Bard | SIT-Bard | 68.56 | 65.06 | 66.87 | 45.61 | 7.06 | 50.63 |
| GPT-4 | SIT-GPT4 | 69.96 | 66.87 | 68.91 | 47.38 | 8.24 | 52.27 |
| Round-0 Nodes | MAGDI-R0 [Level 1] | 71.18 | 67.36 | 72.06 | 48.52 | 9.72 | 53.77 [+ 1.50%] |
| Correct Nodes | MAGDI-CN [Level 2] | 71.62 | 69.31 | 72.34 | 50.11 | 10.66 | 54.81 [+ 2.54%] |
| All Nodes | MAGDI-AN [Level 3] | 72.10 | 70.65 | 71.92 | 50.69 | 11.98 | 55.47 [+ 3.20%] |
| MAG | MAGDI [Level 4] | **74.24** | **72.56** | **72.61** | **52.27** | **12.76** | **56.88 [+ 4.61%]** |

- Level 1: Distillation from multiple teachers > distillation from strongest teacher only.
- Level 2: Distillation from pre- and post-interaction reasoning > only pre. reasoning.
- Level 3: Negative reasoning chains further improve distillation.
- Level 4: Structured distillation from interactions > all multi-teacher baselines.

## Efficiency Analysis

| | RECONCILE | MAGDI | Reduction |
|---|---|---|---|
| StrategyQA | 924.5 | 107.5 | 8.6x |
| CSQA | 936.9 | 104.2 | 9.0x |
| ARC-c | 448.3 | 86.4 | 5.2x |
| GSM8K | 642.3 | 141.6 | 4.5x |
| MATH | 1900.1 | 645.0 | 2.9x |
| Average | 970.4 | 216.9 | 4.5x |

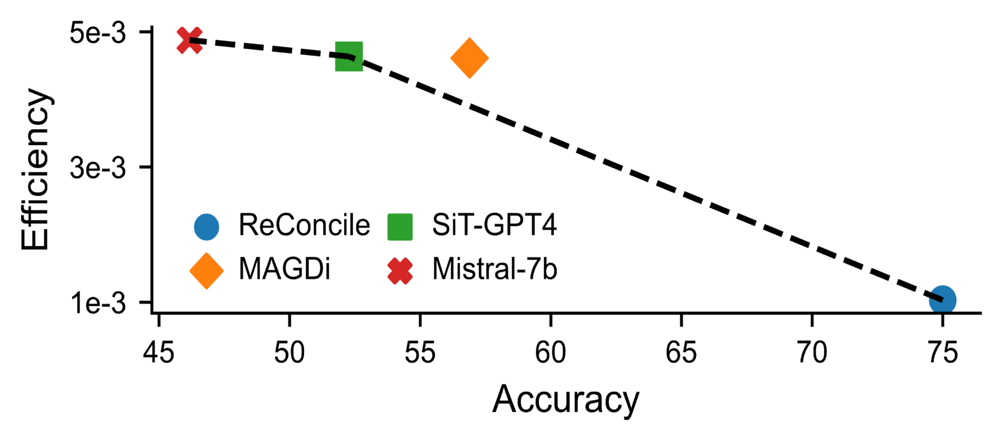


- MAGDi achieves up to a 9x reduction in token count.
- MAGDi achieves a better balance of efficiency and performance.
- More efficient than its teacher system ReConcile
- Performs better than zero-shot and prior single-teacher distillation methods.

## Generalizability, Scalability and Diversity

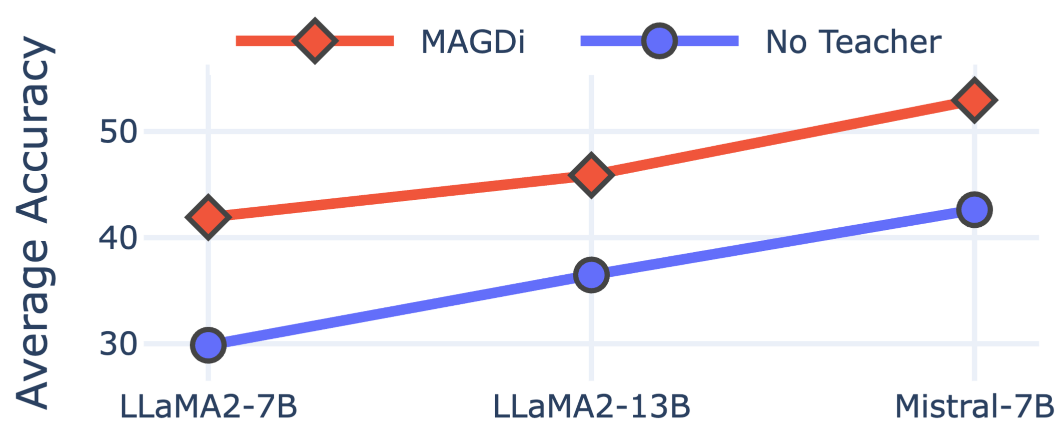


| | BoolQ | SVAMP |
|---|---|---|
| SiT-GPT4-MT | 60.70 | 57.52 |
| MAGDi-MT | **63.98** | **64.30** |

MAGDi scales positively with the student size and performs better on OOD datasets.

| | Mistral-7B | SIT-GPT4 | MAGDI |
|---|---|---|---|
| w/o SC | 44.05 | 47.38 | 52.27 |
| w/ SC | 48.44 [+ 4.39%] | 58.62 [+ 11.24%] | **67.42 [+ 15.15%]** |

MAGDi obtains larger improvements w/ self-consistency, which relies on model diversity.

https://murgelab.cs.unc.edu/
cychen@cs.unc.edu