



# CUSTOMIZING DEEP RESEARCH: ELO EVALS AND OPTIMIZATION OF ADVANCED AI AGENTS



**Jakub Zavrel**  
Founder / CEO



**Fernando Rejon**  
CTO

September 17<sup>th</sup> 2025 | The AI Conference | San Francisco




**CHEMISTRY, HIGH TECH MANUFACTURING, LIFE SCIENCES, PHARMA ETC.**



# DEEP RESEARCH

## WHAT IF YOU NEED ALL THE ANSWERS?

innovative methods to improve the structural properties of PHA and PLA 






Summary

Table

Research

NEW

Innovative methods to improve the structural properties of PHA and PLA include:

- Blending and Copolymerization:** Combining PLA with PHA or other polymers can enhance mechanical and thermal properties. For instance, blending PLA with P3HB using melt processing improves structural characteristics .
- Nanocomposites:** Incorporating nanofillers like cellulose or other nanoparticles into PLA/PHA composites enhances interfacial adhesion, thermal stability, and mechanical properties  .
- Nucleating Agents:** Developing green and efficient nucleating agents can improve the crystallinity and heat resistance of PLA .
- Plasticizers:** Using biobased plasticizers can improve the flexibility and mechanical properties of PLA/PHA blends .

These approaches are tailored to address specific limitations of these biopolymers.



 Open in chat

innovative methods to improve the structural properties of PHA and PLA

Summary Table Research

Innovative methods to improve the structural properties of PHA and PLA include:

- Blending and Copolymerization:** Combining PLA with PHA or other polymers can enhance mechanical and thermal properties. For instance, blending PLA with P3HB using melt processing improves structural characteristics [5].
- Nanocomposites:** Incorporating nanofillers like cellulose or other nanoparticles into PLA/PHA composites enhances interfacial adhesion, thermal stability, and mechanical properties [2, 3].
- Nucleating Agents:** Developing green and efficient nucleating agents can improve the crystallinity and heat resistance of PLA [4].
- Plasticizers:** Using biobased plasticizers can improve the flexibility and mechanical properties of PLA/PHA blends [1].

These approaches are tailored to address specific limitations of these biopolymers.

Open in chat

## Quick Summary

Zeta Alpha

14/03/2025 17:50:05 Generated by Zeta Alpha for Jakob Zareif

### Innovative Methods to Improve the Structural Properties of PHA and PLA

Poly(lactic acid) (PLA) and poly(hydroxyalkanoates) (PHA) are bio-based and biodegradable polymers with significant potential for various applications. However, their structural properties often require enhancement to meet the demands of advanced applications. Below, we explore some of the most interesting and effective methods to improve the structural properties of these materials.

#### 1. Incorporation of Nanofillers

The addition of nanofillers is a widely studied approach to enhance the mechanical properties of PLA and PHA. For instance:

- PLA Nanocomposites:** Integrating nanofillers into PLA results in improved structural characteristics, making it suitable for diverse applications [1].
- Graphene Nanoplatelets (GNPs):** Adding GNPs to both PLA and PHA films only enhances their mechanical properties but also improves piezoresistive electrical performance, offering a multifaceted improvement [2].
- Cellulose Nanocrystals (CNCs):** In PLA/PBAT blends, CNCs localize at the interface of the two phases, stabilizing the morphology under shear flow, improving mechanical properties [3].
- Carbon Nanotubes and TiO<sub>2</sub>:** PLA/PHBV blends reinforced with carbon nanotubes enhance electrical and electromagnetic properties, while TiO<sub>2</sub> improves thermal stability and mechanical strength, making these materials suitable for advanced applications like sensors and electronic devices [4].

#### 2. Blending with Other Polymers

Blending PLA and PHA with other polymers is another effective strategy to enhance their structural properties:

- PLA Blends:** Blending PLA with poly(3-hydroxybutyrate-co-3-hydroxyhexanoate) (PHB) improves tensile strength and impact resistance. For example, a PLA/PHB25 blend achieves a tensile strength of 35.2 MPa due to its morphology [6].
- PHA Blends:** Blending PHA with poly(butylene adipate-co-terephthalate) (PB) significantly enhances ductility and toughness. This improvement is achieved through mechanisms like homogeneous and heterogeneous crazing, where fibers provide additional integrity for larger plastic deformation values [7].
- Impact Modification with Biobased PHA:** Blending PLA with biobased PHA biopolymers enhances its impact resistance and overall mechanical performance, offering a sustainable way to improve PLA's properties [8].

#### 3. Use of Compatibilizers in Polymer Blends

The incorporation of compatibilizers in polymer blends improves interfacial adhesion between different components, leading to better structural integrity. This method is particularly effective for bio-based and biodegradable polymers like PLA [9].

#### 4. Advanced Additive Manufacturing Techniques

Additive manufacturing techniques, such as 3D printing, allow for precise control of the structure of PLA and PHA materials. These methods enhance mechanical performance and make the materials more suitable for advanced applications, such as tissue engineering [10][11].

#### 5. Reinforcement with Natural and Synthetic Additives

Reinforcing PLA and PHA with various additives has shown promising results:

- Calcium Phosphate:** PLA-based composites loaded with calcium phosphite exhibit enhanced mechanical strength and biocompatibility, making them suitable for biomedical applications like bone tissue engineering [12][13].
- Organomodified Sepiolite:** Reinforcing industrial PHA with organomodified sepiolite, a type of clay material, significantly improves its mechanical properties [14].
- Cellulose Microfibers:** Adding cellulose microfibers to PLA and PHA biopolymers enhances their mechanical robustness [15].

#### 6. Development of Biocomposites

The development of PLA- and PHA-based biocomposites using advanced processing techniques and tailored additives has been highlighted as an effective strategy. Biocomposites exhibit superior mechanical and thermal performance, making them suitable for a wide range of applications [16][17].

#### 7. Chemical Modifications

Chemical modification techniques are being explored to tailor the properties of PLA and PHA for specific applications. These modifications allow for the customization of material characteristics, broadening their usability in various fields [18].

#### 8. Nanocomposites for Multifunctional Properties

Nanocomposites not only improve structural properties but also introduce additional functionalities:

- Antibacterial and Biocompatibility Properties:** Incorporating boron nitride into PHA/chitosan nanocomposites enhances thermal stability, mechanical properties, and antibacterial performance, making them suitable for biomedical applications [19].

- Electrical and Electromagnetic Properties:** PLA/PHBV blends reinforced with carbon nanotubes maintain biodegradability while enhancing electrical properties, making them ideal for sensors and electronic devices [4].

#### Conclusion

The structural properties of PLA and PHA can be significantly enhanced through a variety of innovative methods, including the incorporation of nanofillers and other polymers, advanced manufacturing techniques, and the use of natural and synthetic additives. These approaches not only improve mechanical and thermal performance but also introduce multifunctional properties, expanding the applicability of these biopolymers in fields ranging from biomedical engineering to advanced electronics.

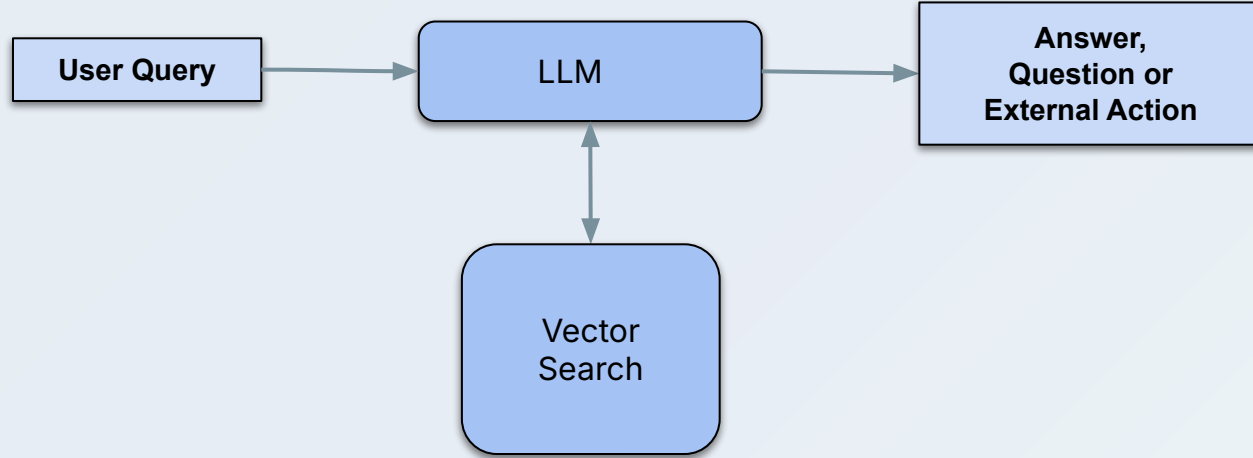
#### References

- [1] ND Bikiaris, I Koumentakou, C Samiotaki, Recent advances in the investigation of poly(lactic acid)(PLA) nanocomposites: incorporation of various nanofillers and their properties and applications, 2023
- [2] G Marmol, UK Sanivada, R Fanguiero, Effect of GNPs on the piezoresistive, electrical and mechanical properties of PHA and PLA films, 2021
- [3] M Mohammad, MC Heuzey, RJ Carreau, A Taguet, Morphological and rheological properties of PLA, PBAT, and PLA/PBAT blend nanocomposites containing CNCs, 2021
- [4] APB Silva, LS Montagna, FR Passador, Biodegradable nanocomposites based on PLA/PHBV blend reinforced with carbon nanotubes with potential for electrical and electromagnetic applications, 2021
- [5] JP Mokkeng, AS Luyt, Morphology and thermal degradation studies of melt-mixed PLA/PHBV biodegradable polymer blend nanocomposites with TiO<sub>2</sub> as filler, 2015
- [6] M Noller, R Salehyan, M Barletta, Tuning the structure-property relationships in binary and ternary blends of PLA/PBAT/PBH, 2024
- [7] R Hosseini-Zadeh, D Elumalai, I Vozniak, Approaches to control crazing deformation of PHA-based biopolymeric blends, 2023
- [8] I Burzic, C Protschku, D Kaineder, G Eder, Impact modification of PLA using biobased biodegradable PHA biopolymers, 2019
- [9] B Imre, B Pukánszky, Compatibilization in bio-based and biodegradable polymer blends, 2013
- [10] RA Ilyas, SM Sapuan, MM Harussani, M Hakim, Poly(lactic acid) (PLA) biocomposites: Processing, additive manufacturing and advanced applications, 2021
- [11] M Mehrooz, H Vahabi, M Barletta, P Lahurte, Additive manufacturing of poly(hydroxyalkanoates) (PHAs) biopolymers: Materials, printing techniques, and applications, 2021
- [12] M Furko, K Balázs, C Balázs, Calcium phosphate loaded biopolymer composites—a comprehensive review on the most recent progress and promising trends, 2023
- [13] M Najah, R Aisyah, S Adzila, Mechanical properties of calcium phosphate reinforced poly(hydroxyalkanoate) (PHA) biocomposite, 2023
- [14] L Garcia-Quiles, A Fernandez-Cuello, P Castel, Sustainable materials with enhanced mechanical properties based on industrial poly(hydroxyalkanoates) reinforced with organomodified sepiolite and ..., 2019
- [15] G Marmol, C Gauss, R Fanguiero, Potential of cellulose microfibers for PHA and PLA biopolymers reinforcement, 2020
- [16] A Farahani, A Zarei-Hanzaki, HR Abedi, Poly(lactic acid) piezo-biopolymers:

## Deep Research

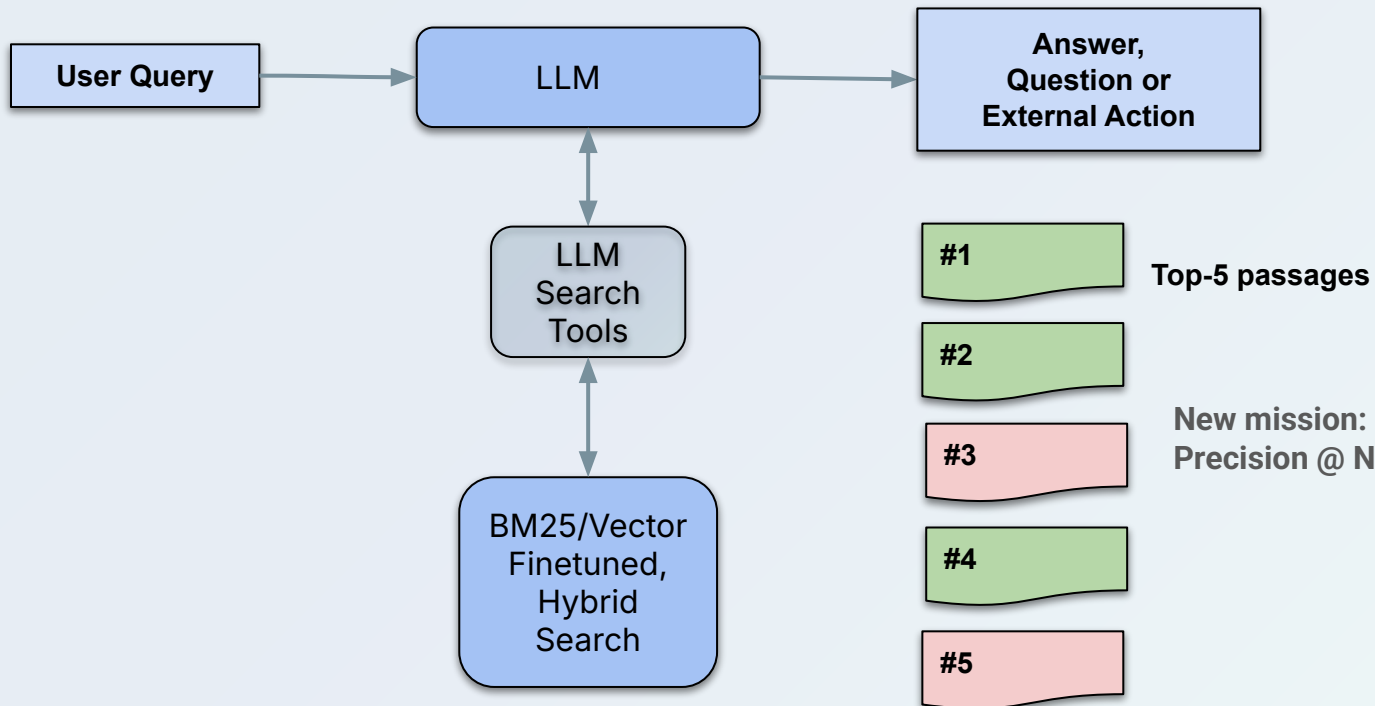


# FROM RAG TO AI AGENTS





# FROM RAG TO AI AGENTS



# COMPLEX AGENTIC SEARCH PIPELINES

Provide all information to report double materiality for banks



Query Rewriting

Intent Recognition

LLM Answer Writing

Chat API

**Keyword Search**  
(BM25, top-150)

**Neural Search**  
(ft E5-base,  
top-150)

**Rank Fusion**  
(RRF, Top-30)

**Re-ranking**  
(ft T5-large, top-10)

Zeta Alpha Search API

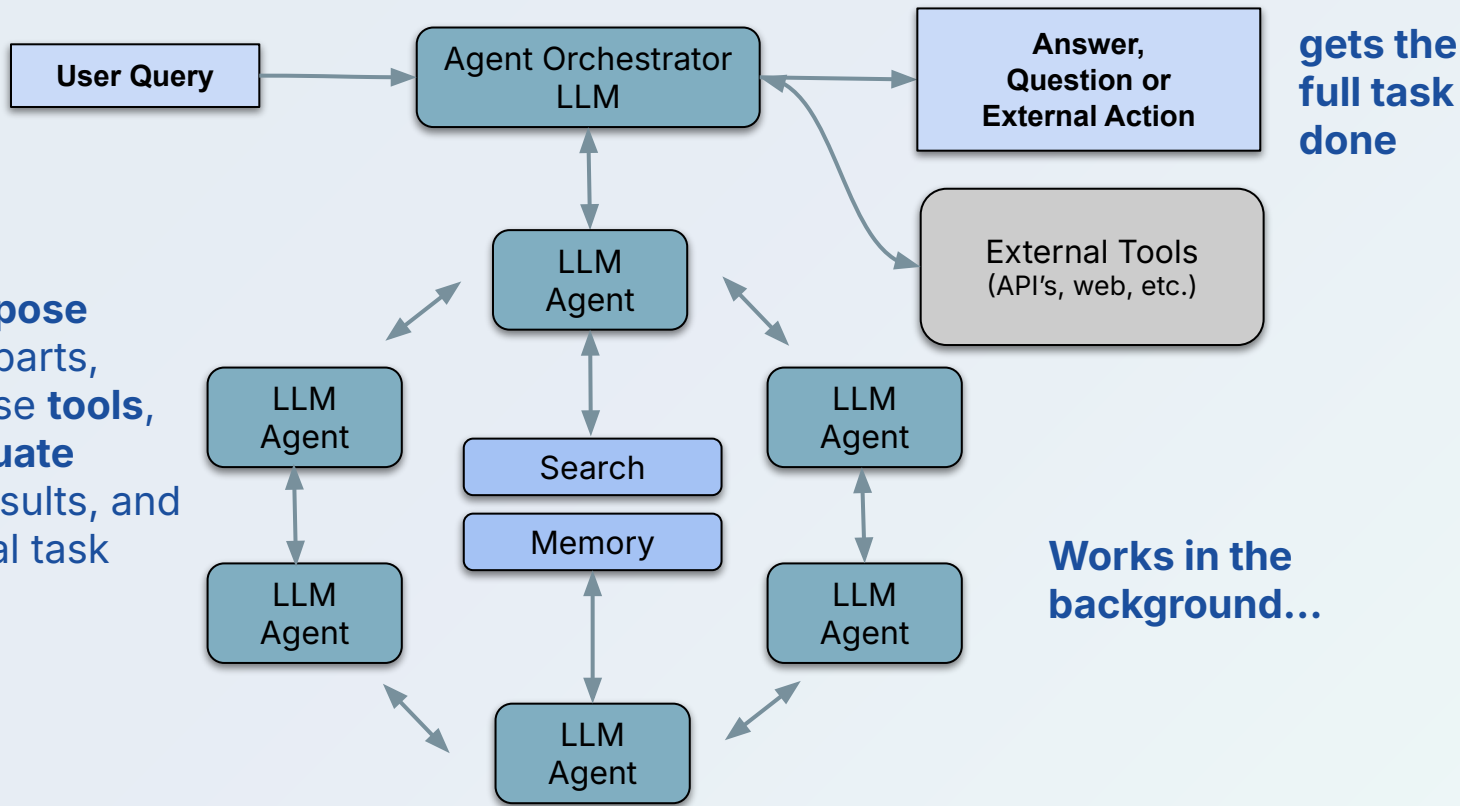




# FROM RAG TO AI AGENTS



Agents **decompose** large task into parts, make a **plan**, use **tools**, **store** and **evaluate** intermediate results, and **synthesize** final task outcome.





# AI AGENT **EVALUATION**



## Can we use synthetic queries to evaluate our system?

- Large test collection for **reliably** is usually not available.
- Performance on public benchmarks doesn't always translate into performance in **private data**.
- Why not use LLM-generated queries for **evaluating** your pipeline?
  - *after bootstrapping on a handful of expert queries* –

What is the definition of sustainability in European legislation incl. guidance from ESAs?

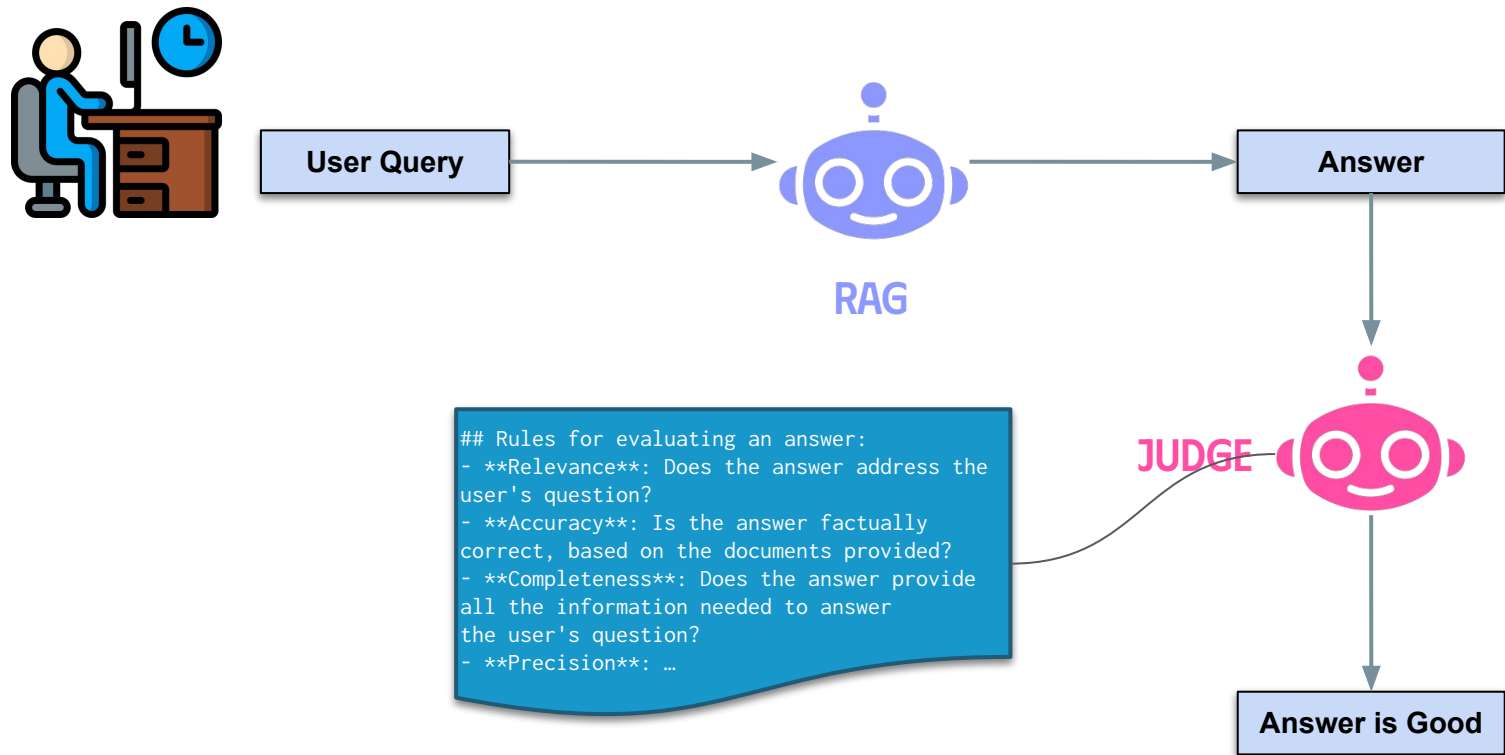
What is the rate of hydrogenation of alpha-D-glucose in MeCONMe2?

Provide all information to report double materiality for banks

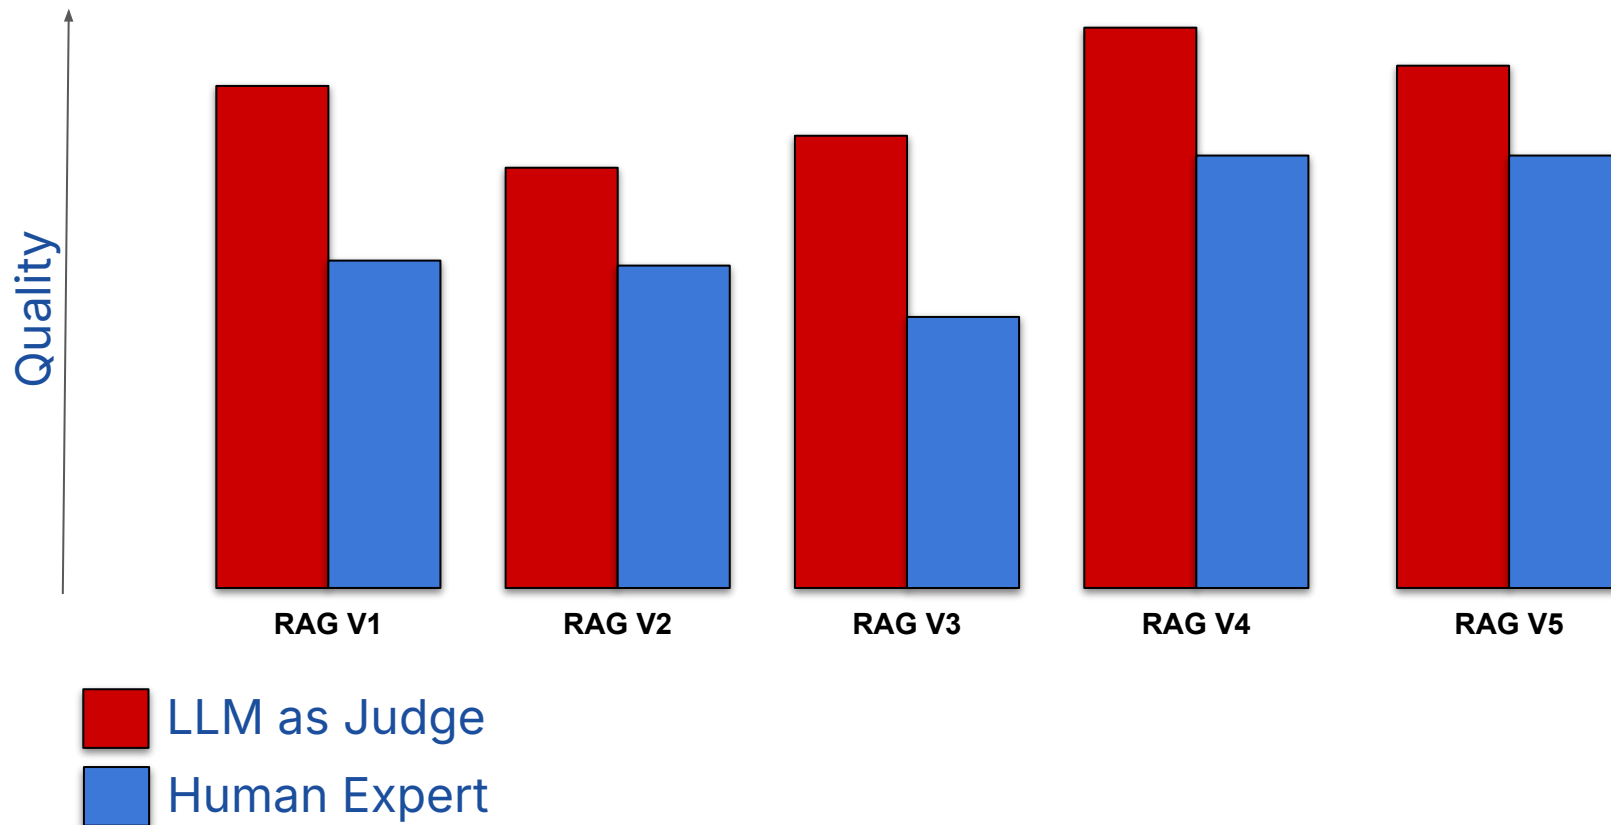
Which programming languages can be used for PLC programming?

which valve is compatible with VABM-B10-25-G12-2-P53?

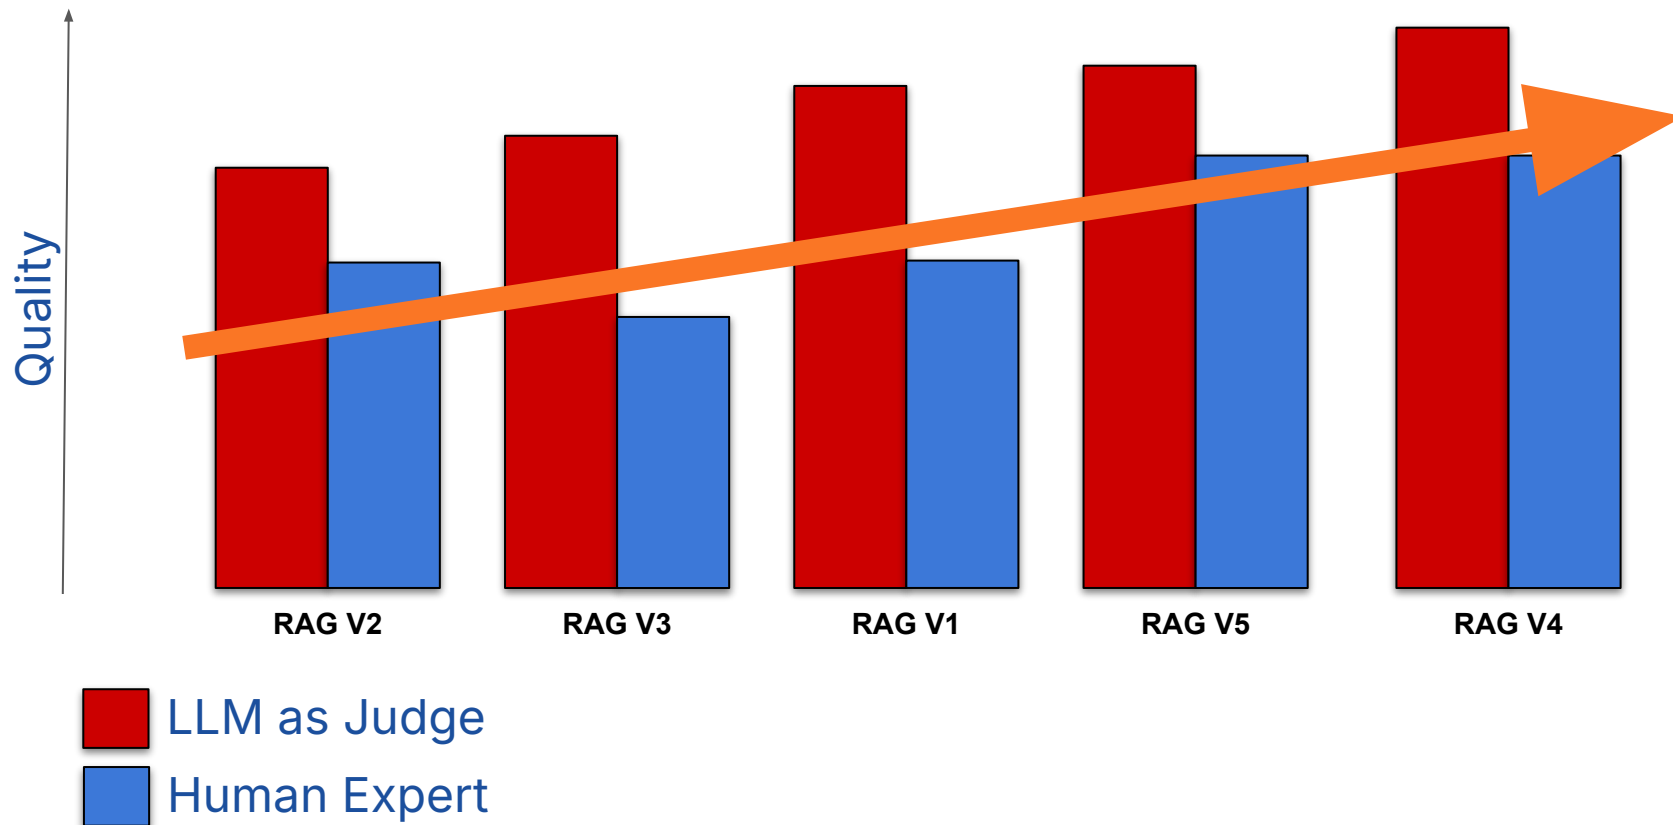
# LLM AS A JUDGE: POINTWISE

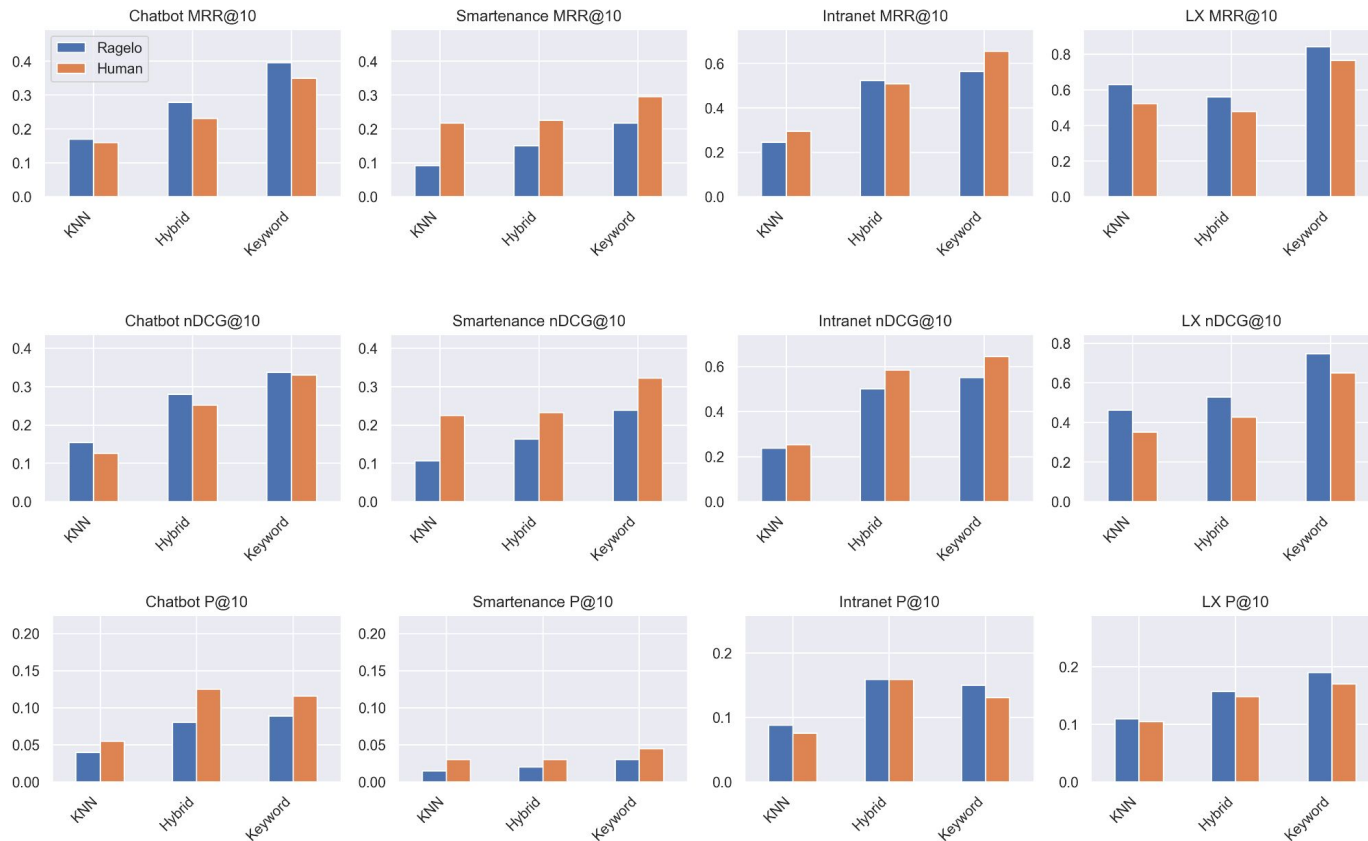


# LLM AS A JUDGE VS EXPERTS



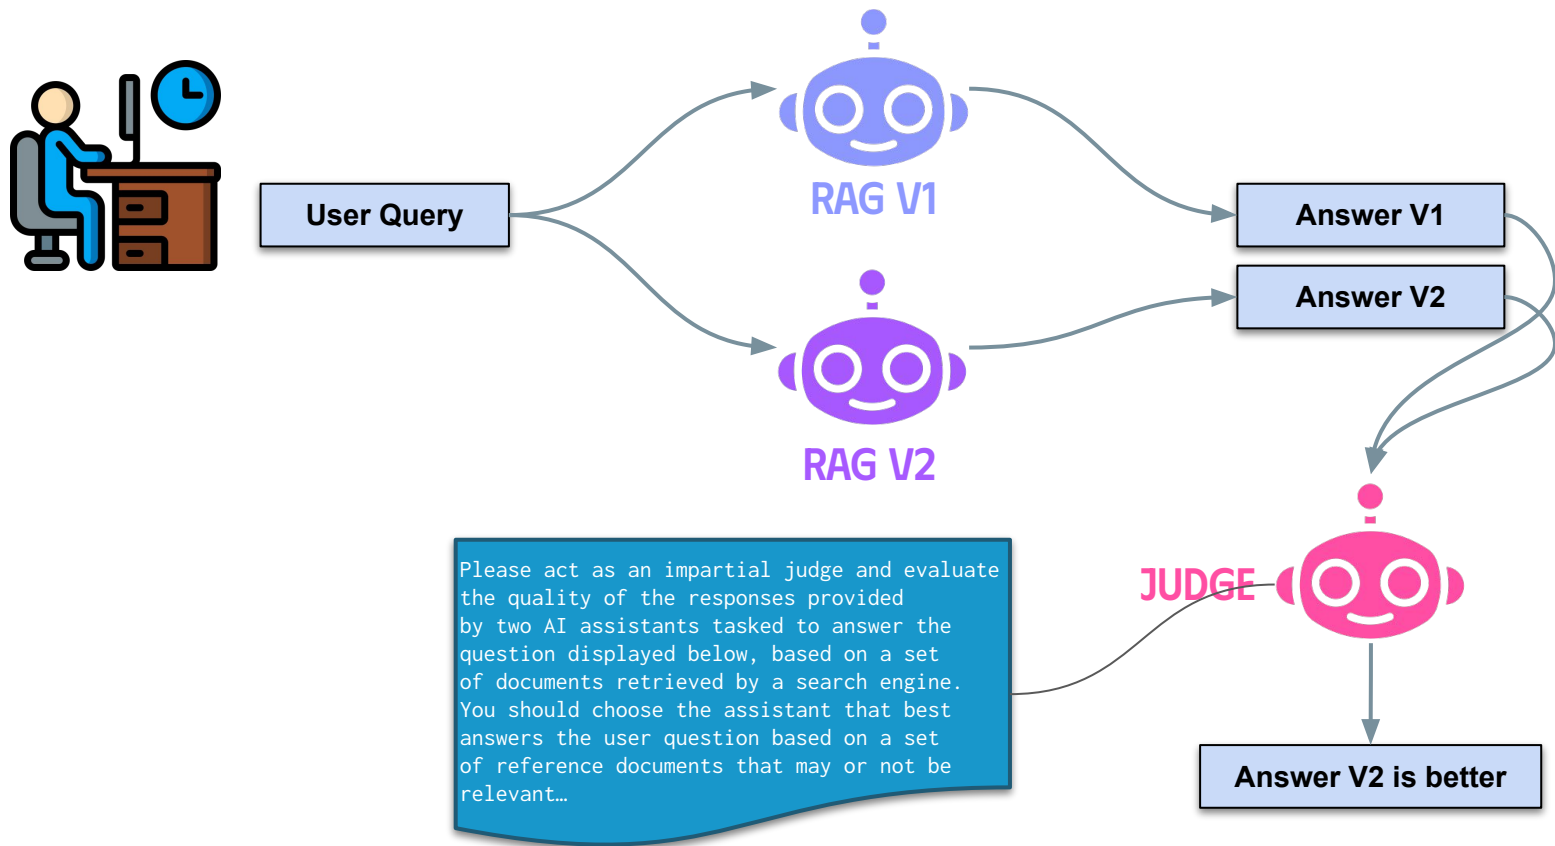
# LLM AS A JUDGE VS EXPERTS



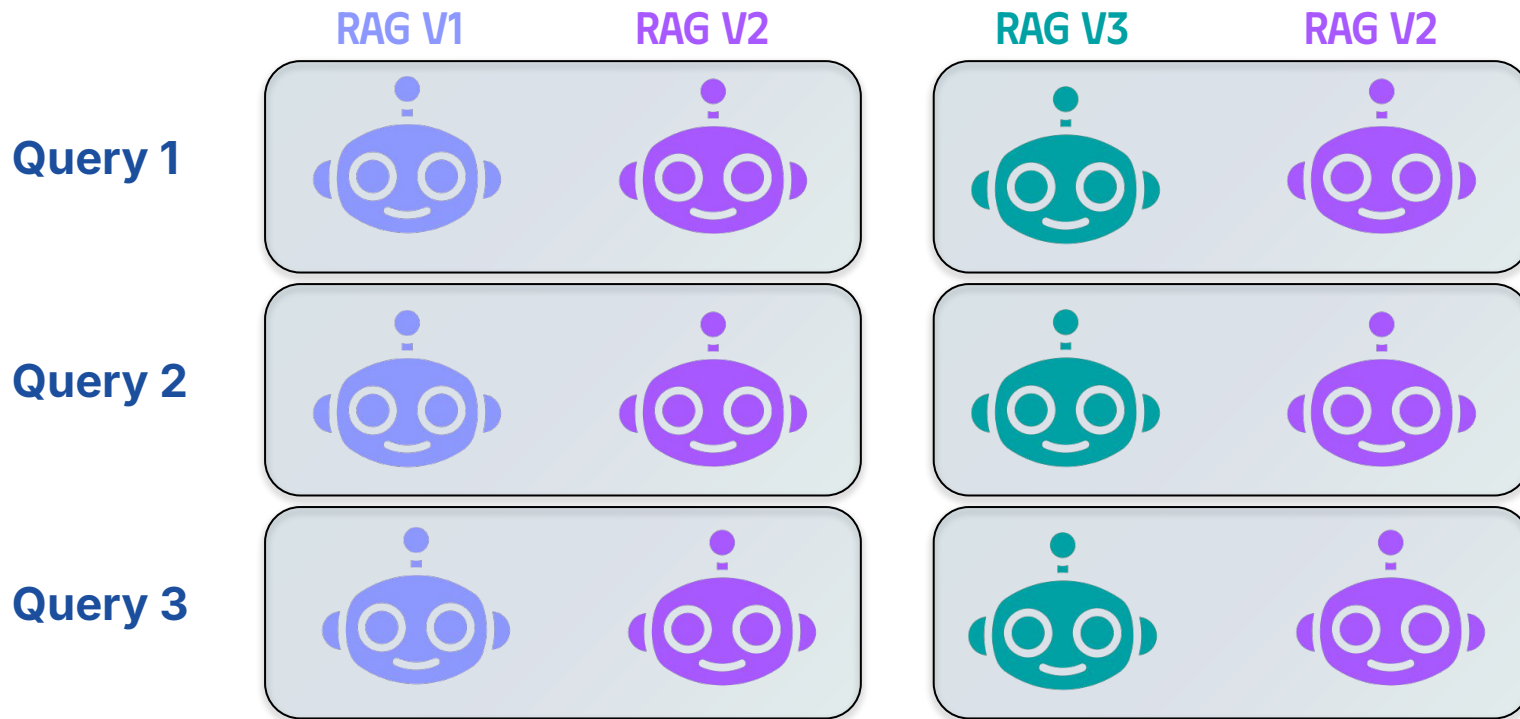




**WHAT IF WE ASK THE LLM TO COMPARE TWO ANSWERS?**



# HOW DO WE RANK PAIRWISE?

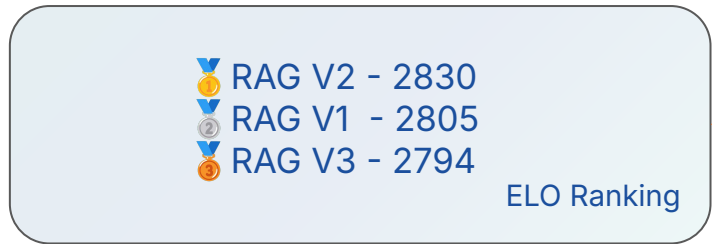
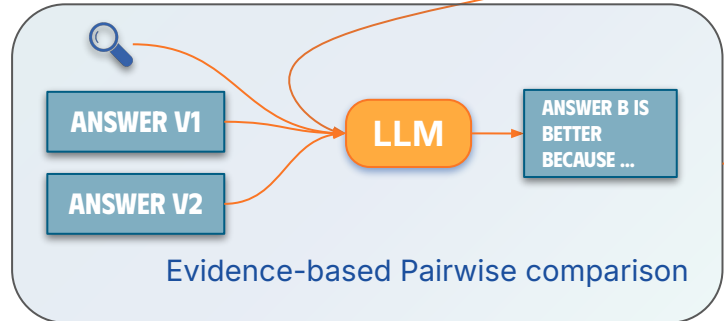
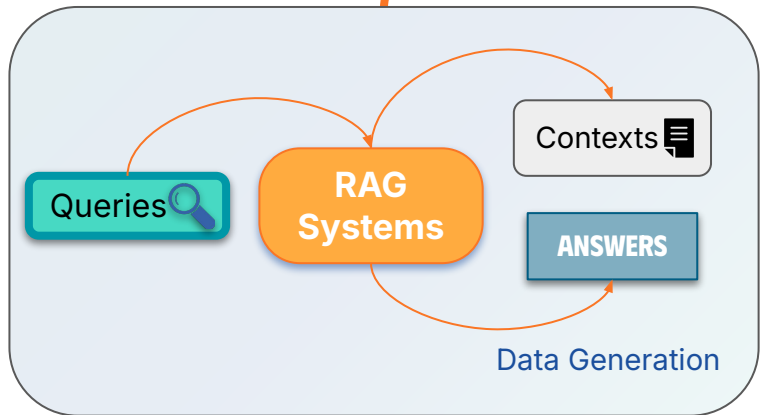


# THE ELO RANKING SYSTEM

Idea: Keep an ELO tournament between all the generated agents?

- Each **agent** (or RAG system) is a tournament with an initial rank.
- Each **query** is **game** played between **agents**.
- A **game** between **Agent A** and **Agent B** is played by prompting an LLM to select which answer to the same **query** is better.
- If **A** wins and its ranking is **higher** than **B**:
  - Score of **A** increases a bit.
  - Score of **B** decreases a bit.
- If **A** wins and its ranking is **lower** than **B**:
  - Score of **A** increases more.
  - Score of **B** decreases more.





zetaalphavector / RAGElo

Can we use synthetic queries and judgments to evaluate our system?

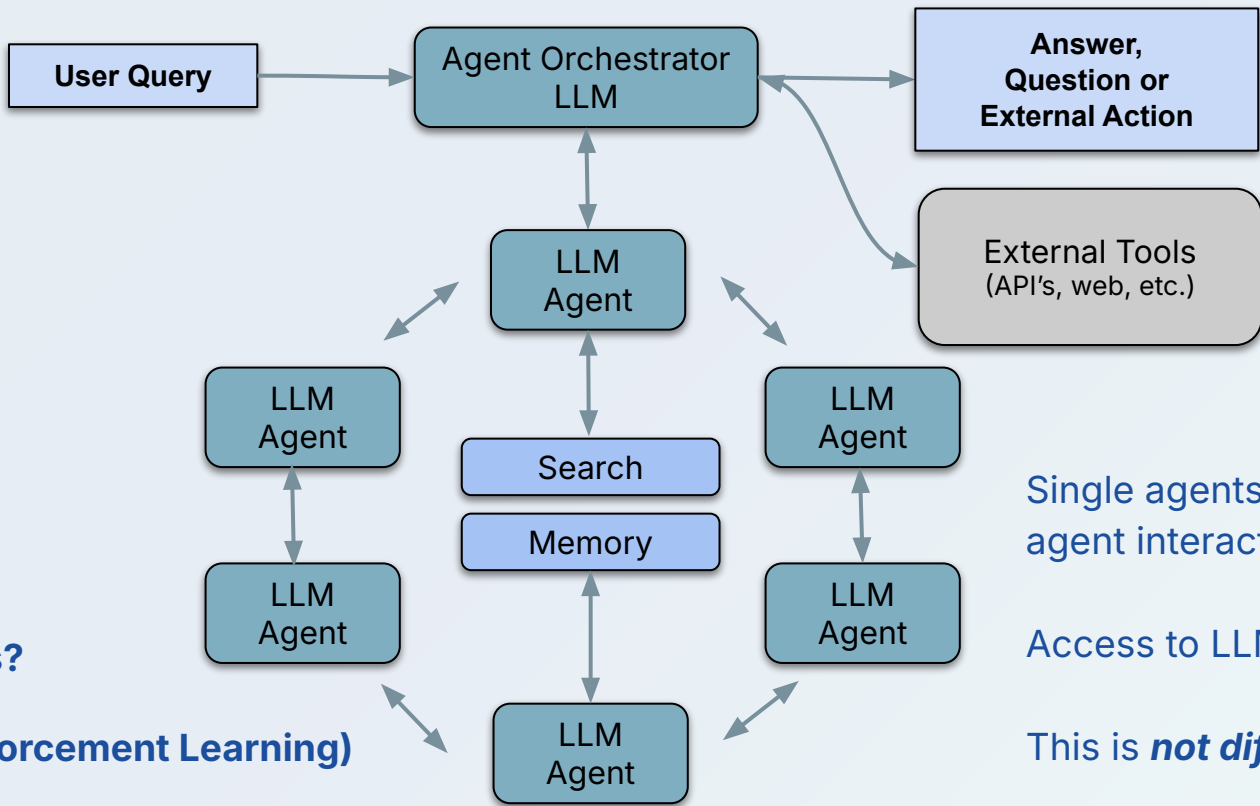
- Extrapolation from a few expert generated examples.
- No 'gold standard' metric needed.
- We can add **new** 'players' later without re-running the whole tournament.

# AI AGENT **OPTIMIZATION**





# AI AGENTS: HOW DO YOU OPTIMIZE THIS?



Many prompts...

So how do you optimize all of this?

Not with RL (Reinforcement Learning)

Single agents + agent interactions.

Access to LLM weights?

This is **not differentiable!**

# A DEEP RESEARCH ARCHITECTURE

Give me a report about the hardware accelerators used in AI in 2025.

**Plan:** I need to research about the most popular accelerators that are being used recently and how they differ to each other. I also need (...)

Most popular Nvidia GPUs 2025



Intel AI Hardware



Cerberas Inference cloud



## Extractors

(One per batch of retrieved documents)

## Mergers

(one per query)

## Planner

The H100 is the most popular NVidia GPU for AI in 2025 [citation], with the A100 remaining popular. Intel's Xeon platform for servers is still one of the most commonly used CPUs for AI and Deep Learning. Cerberas is an AI chipmaker that provides high throughput and low-latency inference hardware for AI...

I have enough information about these vendors, but I did not researched about Google's TPUs, and I need more information about pricing and availability (...)

Now I have enough information to write a comprehensive report on the user's question.

## Writer

8/26/25, 4:07 PM

Comprehensive Report on the Most Commonly Used AI Hardware Accelerators

Zeta Alpha

8/26/2025 4:07:23 PM Generated by Zeta Alpha for Arthur Cămară

## Comprehensive Report on the Most Commonly Used Hardware Accelerators

AI hardware accelerators are specialized processors designed to efficiently execute artificial intelligence (AI) and machine learning (ML) workloads. The rapid growth in model complexity and deployment across data centers, edge devices, and embedded systems has driven the development and adoption of a diverse range of accelerators. This report provides a detailed overview of the most commonly used AI hardware accelerators, highlighting their main characteristics, strengths, weaknesses, and use cases, as well as recent trends and emerging technologies.

### 1. Graphics Processing Units (GPUs)

#### Description & Characteristics:

GPUs, originally designed for graphics rendering, have become the dominant hardware for AI and deep learning due to their massive parallel processing capabilities. Modern GPUs, such as those from NVIDIA and AMD, feature thousands of cores and high bandwidth memory, making them highly effective for matrix operations and neural network computations. Innovations like Tensor Cores (NVIDIA Volta architecture) further optimize GPUs for AI workloads, especially for mixed-precision operations. GPUs are widely supported by major deep learning frameworks and have a mature software ecosystem (e.g., CUDA, cuDNN, ROCm) [1][2][3][4][5][6][7][8].

#### Strengths:

- **High Parallelism & Performance:** Excels at large-scale parallel computing, enabling fast training and high throughput for inference tasks [9][10][11][12].
- **Flexibility & Programmability:** Highly flexible and programmable, with robust support from mature software ecosystems (CUDA, cuDNN, PyTorch, TensorFlow, OpenCL) [8][13][10].

## Recent surge in research...

### Key element:

Use the raw intelligence of the LLM itself as an operator to come up with improved versions of the system.

[illegible]

The diagram illustrates the TapeAgents framework architecture, organized into several layers and components:

- Contributors:** A list of individuals associated with the framework:
  - Core contributors: Dmitry Babukhanov, Nicolas Gontier, Gabriel Huang, Elton Kananlov, Rafael Purdies, Alex Pichin.
  - Contributors: Tamas Szekely, Oleg Shilnikov, Jordan Prince Tevzabay.
  - Committers: Roman Chernikov, Mikhail Mironov, Alexander Mikhalev.
- Environment Layer (Green Box):** Contains components that interact with the Tape:
  - Human
  - APIs
  - RAG
  - Code Execution
  - Web
- Orchestrator (Red Box):** Acts as a central hub, receiving input from the Environment and interacting with the Tape and Agent A.
- Tape (Yellow Cylinder):** A central data store that receives "Observations" from the Environment and provides "Tape" data to the Studio and "Thoughtful Actions" to Agent B.
- Studio (Dotted Box):** Used for managing the Tape, including:
  - Interact & Resume
  - Tape Review
  - Evaluate
  - Monitor
- Optimize (Dotted Box):** Used for improving the system, including:
  - Synthetic Data Generation
  - Agent Optimization
  - Flow-tuning
  - LLM
  - Prompt Tuning
- Agent Layer (Blue Boxes):**
  - Agent A:** Interacts with the Orchestrator via "agent.run()" and delegates tasks to Agent B.
  - Agent B:** Receives "Thoughtful Actions" from the Tape and interacts with a sequence of nodes (Node 1 to Node N). It includes an "Action Node" and an "LLM" component. It generates "Output" and delegates tasks to Agent C and Agent D.
  - Agent C:** Receives a "delegate()" task from Agent B.
  - Agent D:** Receives a "delegate()" task from Agent B.

9 May 2025

**Darwin G del Machine:  
Open-Ended Evolution of Self-Improving Agents**

---


Jenny Zhang<sup>1,2</sup> Shengran Hu<sup>1,2,3</sup> Cong Lu<sup>3,2</sup> Robert Lange<sup>3,2</sup> Jeff Clune<sup>1,2,3,4</sup>

<sup>1</sup>University of British Columbia <sup>2</sup>Vector Institute <sup>3</sup>Rakana AI <sup>4</sup>Canada CIFAR AI Chair  
{jennyxz, shu, conglu}@ucb.ubc.ca, robert@rakana.ai, jeff.clune@ubc.ca

**Abstract**

Most of today's AI systems are constrained by human-designed, fixed architectures and cannot autonomously and continuously improve themselves. The scientific method, on the other hand, provides a cumulative and open-ended system, where each innovation builds upon previous artifacts, enabling future discoveries. There is growing hope that the current manual process of advancing AI could itself be

**Darwinian  
Agent Archive**

Initial Agent 

[illegible]

# TEXTGRAD: AUTOMATIC "DIFFERENTIATION" VIA TEXT (JULY 2024)

By Mert Yuksekgonul et al. (+6 authors)

**TextGrad**: LLMs to perform automatic "differentiation" through text-based feedback.



Feedback/  
gradient  
generator

- Given an answer and evaluation, generates a **feedback to the answer**.
- For each sub-agent, generates **feedback to its prompt**.
- Feedbacks are "back-propagated" across sub-agents



Optimizer

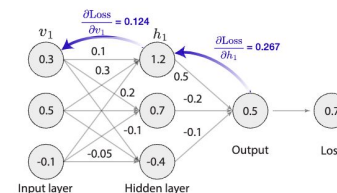
- For each sub-agent, **Feedback to its prompt**. Generates a **new prompt**.

## TextGrad: Automatic "Differentiation" via Text

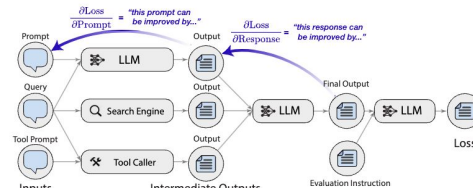
Mert Yuksekgonul<sup>1\*</sup>  
Federico Bianchi<sup>1\*</sup>  
Joseph Boen<sup>2\*</sup>  
Sheng Liu<sup>2\*</sup>

MERTY@STANFORD.EDU  
FEDE@STANFORD.EDU  
TBOEN@STANFORD.EDU  
SHENGL@STANFORD.EDU

### a Neural network and backpropagation using numerical gradients



### b Blackbox AI systems and backpropagation using natural language 'gradients'



### c 1 Analogy in abstractions

	Math	PyTorch	TextGrad
Input	$x$	<code>Tensor(image)</code>	<code>tg.Variable(article)</code>
Model	$\hat{y} = f_{\theta}(x)$	<code>ResNet50()</code>	<code>tg.BlackboxLLM("You are a summarizer.")</code>
Loss	$L(y, \hat{y}) = \sum y_i \log(\hat{y}_i)$	<code>CrossEntropyLoss()</code>	<code>tg.TextLoss("Rate the summary.")</code>
Optimizer	$GD(\theta, \frac{\partial L}{\partial \theta}) = \theta - \frac{\partial L}{\partial \theta}$	<code>SGD(list(model.parameters()))</code>	<code>tg.TGD(list(model.parameters()))</code>

### 2 Automatic differentiation

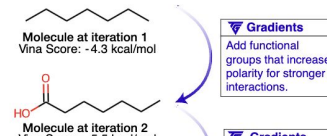
PyTorch and TextGrad share the same syntax for backpropagation and optimization.

**Forward pass**  
`loss = loss_fn(model(input))`

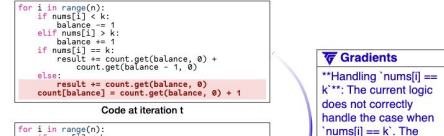
**Backward pass**  
`loss.backward()`

**Updating variable**  
`optimizer.step()`

### d TextGrad for molecule optimization



### e TextGrad for code optimization





# GEPA: REFLECTIVE PROMPT EVOLUTION CAN OUTPERFORM REINFORCEMENT LEARNING (JULY 2025)

By Lakshya A Agrawal et al. (+16 authors)

**GEPA:** LLMs for prompt reflection and pareto-optimality to quickly select the best agents and merge sub-agents.

Feedback/  
gradient  
generator

Optimizer

Agent  
selection

- Given one agent, get the answers **mini-batch of training queries**, obtain textual **feedback on its output**.
- With the **feedback on the mini-batch**, refine the prompt of one sub-agent at a time.
- Optionally, select another agent and combine its (other) sub-agents with this new sub-agent.
- If the new agent is better, add to **pool of candidate agents**.
- At each step, select from the agents that are **pareto-optimal** on a set of training samples

## GEPA: REFLECTIVE PROMPT EVOLUTION CAN OUTPERFORM REINFORCEMENT LEARNING

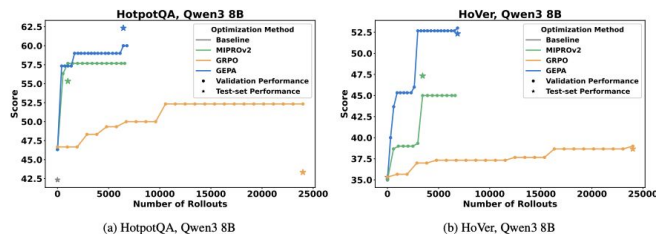
Lakshya A Agrawal<sup>1</sup>, Shangyin Tan<sup>1</sup>, Dilara Soylu<sup>2</sup>, Noah Ziem<sup>4</sup>,  
Rishi Khare<sup>1</sup>, Krista Opsahl-Ong<sup>2</sup>, Arnab Singhvi<sup>2,5</sup>, Herumb Shandilya<sup>2</sup>,  
Michael J Ryan<sup>2</sup>, Meng Jiang<sup>1</sup>, Christopher Potts<sup>2</sup>, Koushik Sen<sup>1</sup>,  
Alexandros G. Dimakis<sup>1,3</sup>, Ion Stoica<sup>1</sup>, Dan Klein<sup>1</sup>, Matei Zaharia<sup>1,5</sup>, Omar Khattab<sup>6</sup>

<sup>1</sup>UC Berkeley <sup>2</sup>Stanford University <sup>3</sup>BespokeLabs.ai <sup>4</sup>Notre Dame <sup>5</sup>Databricks <sup>6</sup>MIT

### ABSTRACT

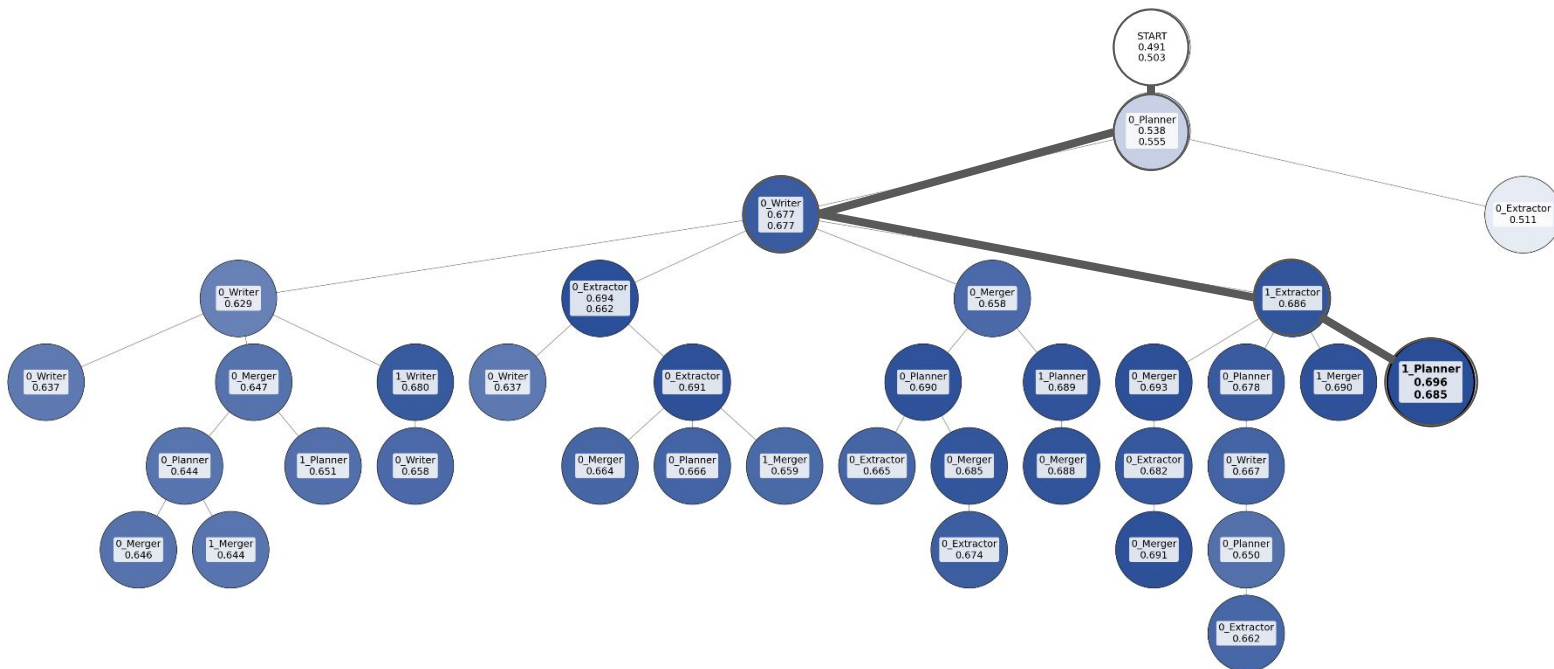
Large language models (LLMs) are increasingly adapted to downstream tasks via reinforcement learning (RL) methods like Group Relative Policy Optimization (GRPO), which often require thousands of rollouts to learn new tasks. We argue that the interpretable nature of *language* can often provide a much richer learning medium for LLMs, compared with policy gradients derived from sparse, scalar rewards. To test this, we introduce GEPA (Genetic-Pareto), a prompt optimizer that thoroughly incorporates *natural language reflection* to learn high-level rules from trial and error. Given any AI system containing one or more LLM prompts, GEPA samples system-level trajectories (e.g., reasoning, tool calls, and tool outputs) and reflects on them in natural language to diagnose problems, propose and test prompt updates, and combine complementary lessons from the Pareto frontier of its own attempts. As a result of GEPA's design, it can often turn even just a few rollouts into a large quality gain. Across four tasks, GEPA outperforms GRPO by 10% on average and by up to 20%, while using up to 35x fewer rollouts. GEPA also outperforms the leading prompt optimizer, MIPROv2, by over 10% across two LLMs, and demonstrates promising results as an inference-time search strategy for code optimization.

### 1 INTRODUCTION



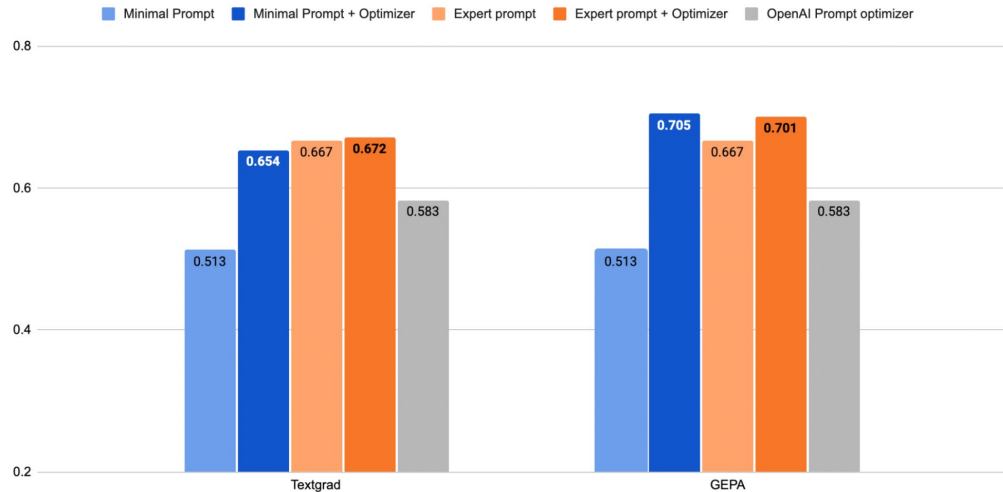
arXiv:2507.19457v1 [cs.CL] 25 Jul 2025

# OPTIMIZING AGENT SYSTEMS: GEPA EXPLORATION



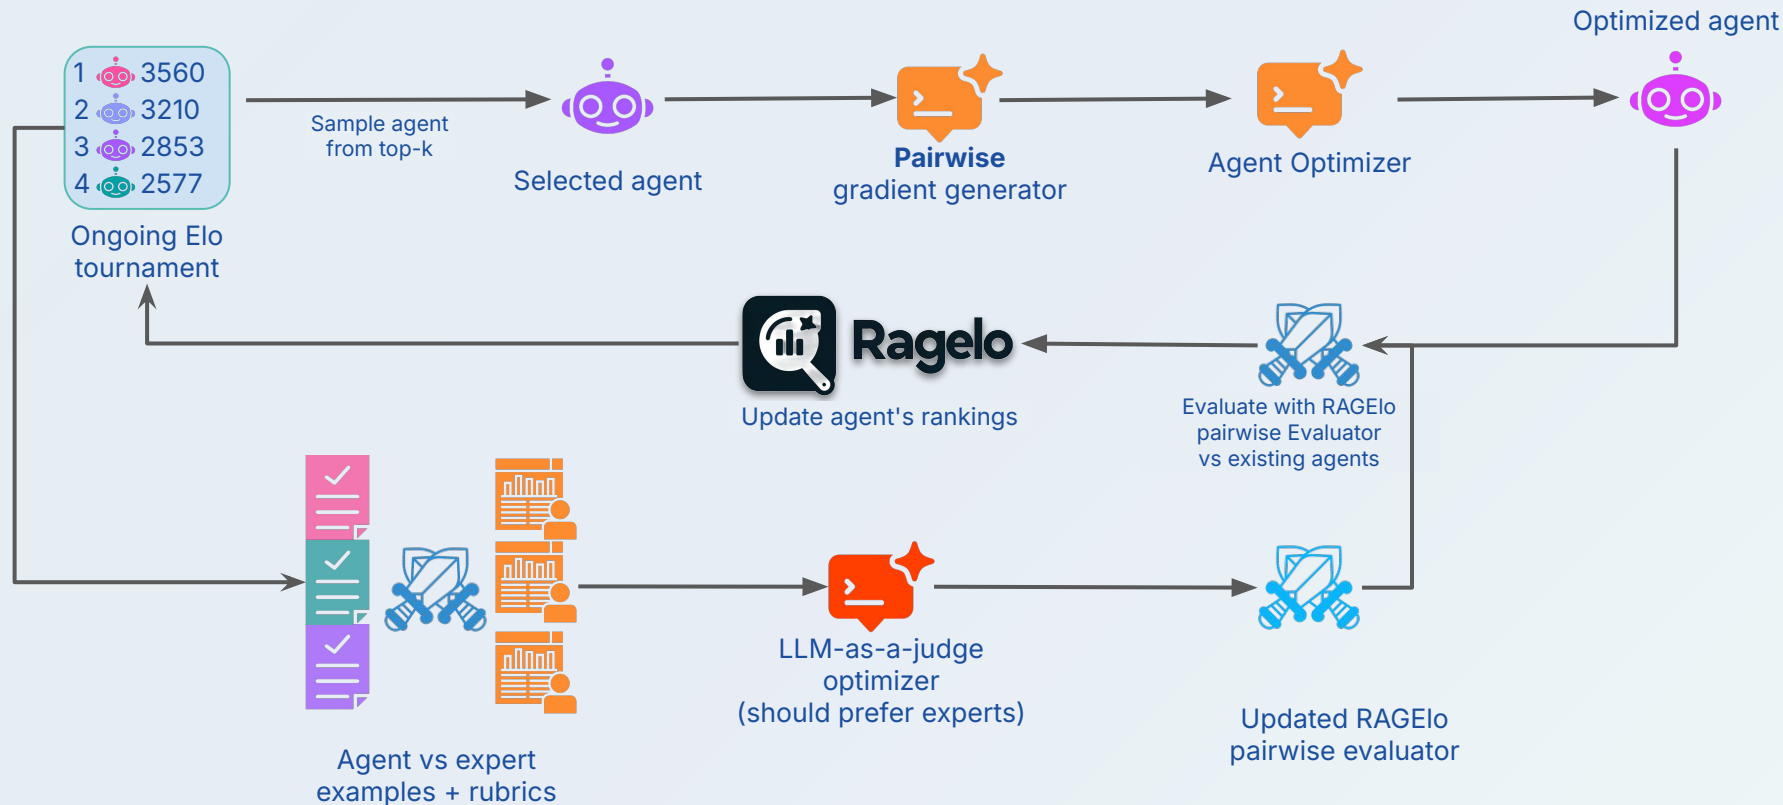


Agent Optimization results



- **ScholarQA-CS Benchmark.** For retrieval, we use the **DeepResearchGym** API.
- Agents init from two starting points:
  - minimal prompt for each sub-agent
  - human-optimized
- LLM used GPT-4.1-mini.
- We also used OpenAI's prompt optimizer for GPT-4.1
- **Self-Optimized Agents match or outperform human-optimized.**
- **GEPA better results than Textgrad.**
- Gains are more pronounced when **starting from simple**, non-optimized prompts.

# NEXT STEPS: A CO-EVOLUTION FRAMEWORK



# SUMMARY & NEXT STEPS

## COMPLEX AGENTS

RAG was just the beginning.

---

Now we orchestrate more complex agentic pipelines.

How to evaluate long-format responses?

LLM-as-a-judge and Elo style tournaments.

---

## DEEP RESEARCH & CO.

Delivering value to experts.

---

What are best prompts and architectures?

How to Improve quality and customize?

What if you only have some examples?

---

## OPTIMIZATION

Evaluate and Iterate.

---

GEPA works for efficiently optimizing your agents.

Next step: co-evolve agents and judges based on examples..

---

# QUESTIONS?

# ENJOY DISCOVERY!

[www.zeta-alpha.com](http://www.zeta-alpha.com)

Visit us at  
booth #19  
in the innovator area



Trends in AI

Monthly webinar, live  
from LAB42 / SF