

Tuesday 13 October 2026
<https://indico.global/event/19255/>

CLARIPHY

AI Collaboration Meeting

AGENTIC ANALYSIS






The Conveners



Petra Maksimovic



Artur Choi



Peter Fackeldey



Andrzej Novak




Giordon Stark



Haichen Wang



Code of Conduct



CLARIPHY

*Community Laboratory for AI-native
Research In PHYsics*

As a reminder: this is a public meeting open to **ALL** members of the HEP community. Help us create a respectful, inclusive, and collaborative environment.

Be respectful and inclusive: Treat everyone with respect. We welcome diverse perspectives, backgrounds, career stages, and areas of expertise.

Make space for everyone: Share airtime, avoid interruptions, listen actively, and give others room to contribute.

Respect accessibility needs: Follow accessibility requests, use Zoom hands/reactions, and the chat box.

Keep the space safe: No harassment, discrimination, intimidation, or unwelcome conduct. Respect personal boundaries and confidentiality where appropriate.

Help us improve: Be constructive, assume good faith while recognizing impact, and speak up or seek help if something goes wrong. Need help or want to raise a concern? Contact any of the conveners privately.



Goal of this forum

From the CLARIPHY meeting, we resolved to start biweekly “AI Agentic Analysis Systems” (AAAS) meetings.

- A standing cross-experiment informal forum to improve communication and coordination
- Do you build skills or agents or MCP tooling, or use AI on analysis notes?
Reach out to us!
 - Contribute to the awesome-hep-agentic-analysis collection:
<https://github.com/awesome-hep-agentic-analysis>
- Small & mid-scale experiments are in scope: please do not leave them out!
- Join the Slack channel:
<https://clariphyworkspace.slack.com/archives/C0BB1JV7FME>

The first thing people would hand to an agent

Repetitive work nobody enjoys

- Find the datasets and current recommendations, then use with them
- Triage thousands of failed batch jobs; then fix the trivial ones
- Resurrect last publication's code against the current release
- Check that the analysis note is self-consistent

Things we can't mentally keep track of

- Enumerate every background process for a final state
- Look up trigger and detector metadata on demand
- Literature review: prior results, limits, reinterpretations
- Overlap detection: where do existing analyses already cover this phase space?
- Critique the plan: cuts, statistics, alternatives from other analyses

Almost none of these are physics problems. They are access, metadata, provenance and retrieval problems.

Eight groups, eight demonstrators

AREA	DEMONSTRATOR	MAIN BLOCKER	HORIZON
Revisiting how analysis is done	Tech notes that are easy to review — AI edits, student still writes the first draft; test with a real review committee	Social — collaboration & pub-board buy-in	~6 mo if buy-in
Noodling (idea → plot)	Authenticated web interface for fast questions on a small experiment's data; returns the plot and the source code	Access to files, metadata, CLI; auth	~6 mo, 1 FTE
Discoverability	Public skill marketplace for the field + internal versions per experiment; MCP servers to reach existing infrastructure	Sociological; no standards yet	0–3 mo tech, 6+ social
Benchmarking	A standard set of measurements across experiments, scored against published values with a scalar	No clean metric; open data scarce off-collider	~6 mo collider
Review	Split an open-data analysis into analysis & reviewer teams; write the paper about how the review was done	Objective benchmark; expertise; acceptance	Start now
Debugging / reproducibility	Checkpoints that validate individual elements along a workflow providing debugging and reproducibility at once	No common thrust across experiments	~6 mo template
Data standardization	Agentified fit exploring treatments of correlated systematics; recommendations incl. for tokenization	Scale; person-power, coordination, tokens	~6 mo demo
AI use / policy	Supervisor LLM over task-specific small language models, looping until success or a capped return to a human	Sub-agents that are themselves LLMs	now – 6 mo



All the news that's fit to ... print(?)

- <https://reflection.ai/blog/introducing-beam>: 501b parameter open weight model for agentic analysis, coding, and science
- NSF 26-512 solicitation “AI-ready Datasets”:
<https://www.nsf.gov/funding/opportunities/ai-datasets-unlocking-dataset-value-ai-enabled-scientific-discovery/nsf26-512/solicitation>
- NSF 26-513 solicitation “AI Regional Hubs”:
<https://www.nsf.gov/funding/opportunities/us-national-science-foundation-state-regional-artificial/nsf26-513/solicitation>
- Genesis Graduate Fellowship Pilot FOA (“4-year PhDs”):
<https://science.osti.gov/-/media/grants/pdf/foas/2027/DE-FOA-0003673.pdf>
- More Genesis Phase 1 / 2 proposals announced:
https://science.osti.gov/-/media/funding/pdf/Awards-Lists/2026/GM-RFA-Awards-List_09-30-26.pdf