

Adventures with Agentic AI (in particle physics)

New Opportunities in Particle Physics
IHEP Beijing

David Shih
August 15, 2026



RUTGERS
THE STATE UNIVERSITY
OF NEW JERSEY

Agentic AI — a watershed moment

Claude Code, Codex — early 2025

Giving *agency* to LLMs — from question answering to reading/writing/executing code



Home News Sport Business Technology Health Culture Arts Travel Earth Audio Video Live

'Vibe coding' named word of the year by Collins Dictionary

tom's**HARDWARE**

Follow

Claude-powered AI coding agent deletes entire company database in 9 seconds — backups zapped, after Cursor tool powered by Anthropic's Claude goes rogue



The New York Times



Artificial Intelligence >

Apple Settlement

Job Losses

Meta Sued

Vetting A.I. Models

A.I. Spending Rec

SHOP TALK


With 'Vibecoding,' A.I. Can Help Anyone Build an App

Bringing on artificial intelligence as a collaborator can make coding feel more accessible to those with little training in it, but there are trade-offs.

TECH

20,000 job cuts at Meta, Microsoft raise concern that AI-driven labor crisis is here

Agentic AI explained



ZBrain

- Agentic AI consists of
 - LLM backend (Claude, GPT, Gemini ...)
 - a set of ***tools*** (bash, python, grep, text editor, ...)
 - a ***harness/scaffold*** (traditional code infrastructure managing model responses and tool use)
- Examples include Claude Code, Codex, ...

How are people using agentic AI?

- Particle physicists are building custom harnesses for automating

- anomaly detection

Agents of Discovery

Sascha Diefenbacher¹, Anna Hallin²,
Gregor Kasieczka², Michael Krämer³, Anne Lauscher⁴, Tim Lukas²,

- pheno tasks [ArgoLOOM, HEPTAPOD, MadAgents, CoLLM, FERMIACC, ColliderAgent/Magnus...]

The FERMIACC: Agents for Particle Theory

- experimental data analysis

AI Agents Can Already Autonomously Perform Experimental High Energy Physics

Eric A. Moreno^{*†1,2}, Samuel Bright-Thonney^{*†1,2}, Andrzej Novak^{*§1,2}, Dolores Garcia^{#3},
and Philip Harris^{b1,2}

Prateek Agrawal,^a Nathaniel Craig,^{a,b} Amalia Madden,^b and Iñigo Valenzuela Lomber^a

^aDepartment of Physics, University of California, Santa Barbara, CA 93106, USA

^bKavli Institute for Theoretical Physics, University of California, Santa Barbara, CA 93106, USA

An End-to-end Architecture for Collider Physics and Beyond

Shi Qiu,¹ Zeyu Cai,¹ Jiashen Wei,¹ Zeyu Li,^{1,2} Yixuan Yin,¹ Qing-Hong
Cao,^{1,3,2} Chang Liu,^{1,4} Ming-xing Luo,⁵ Xing-Bo Yuan,⁶ and Hua Xing Zhu^{1,2}

How are people using agentic AI?

- Particle physicists are building custom harnesses for automating

- anomaly detection

Agents of Discovery

Sascha Diefenbacher¹, Anna Hallin²,
Gregor Kasieczka², Michael Krämer³, Anne Lauscher⁴, Tim Lukas²,

- pheno tasks [ArgoLOOM, HEPTAPOD, MadAgents, CoLLM, FERMIACC, ColliderAgent/Magnus...]

The FERMIACC: Agents for Particle Theory

- experimental data analysis

AI Agents Can Already Autonomously Perform Experimental High Energy Physics

Eric A. Moreno^{*†1,2}, Samuel Bright-Thonney^{*†1,2}, Andrzej Novak^{*§1,2}, Dolores Garcia^{#3},
and Philip Harris^{b1,2}

Prateek Agrawal,^a Nathaniel Craig,^{a,b} Amalia Madden,^b and Iñigo Valenzuela Lomber^a

^aDepartment of Physics, University of California, Santa Barbara, CA 93106, USA

^bKavli Institute for Theoretical Physics, University of California, Santa Barbara, CA 93106, USA

- Questions:

An End-to-end Architecture for Collider Physics and Beyond

Shi Qiu,¹ Zeyu Cai,¹ Jiashen Wei,¹ Zeyu Li,^{1,2} Yixuan Yin,¹ Qing-Hong
Cao,^{1,3,2} Chang Liu,^{1,4} Ming-xing Luo,⁵ Xing-Bo Yuan,⁶ and Hua Xing Zhu^{1,2}

How are people using agentic AI?

- Particle physicists are building custom harnesses for automating

- anomaly detection

Agents of Discovery

Sascha Diefenbacher¹, Anna Hallin²,
Gregor Kasieczka², Michael Krämer³, Anne Lauscher⁴, Tim Lukas²,

- pheno tasks [ArgoLOOM, HEPTAPOD, MadAgents, CoLLM, FERMIACC, ColliderAgent/Magnus...]

- experimental data analysis

AI Agents Can Already Autonomously Perform Experimental High Energy Physics

Eric A. Moreno^{*†1,2}, Samuel Bright-Thonney^{*†1,2}, Andrzej Novak^{*§1,2}, Dolores Garcia^{#3},
and Philip Harris^{b1,2}

The FERMIACC: Agents for Particle Theory

Prateek Agrawal,^a Nathaniel Craig,^{a,b} Amalia Madden,^b and Iñigo Valenzuela Lomber^a

^aDepartment of Physics, University of California, Santa Barbara, CA 93106, USA

^bKavli Institute for Theoretical Physics, University of California, Santa Barbara, CA 93106, USA

- Questions:

- Frontier LLM models and harnesses improving rapidly
— how much value does a custom harness add?

An End-to-end Architecture for Collider Physics and Beyond

Shi Qiu,¹ Zeyu Cai,¹ Jiashen Wei,¹ Zeyu Li,^{1,2} Yixuan Yin,¹ Qing-Hong
Cao,^{1,3,2} Chang Liu,^{1,4} Ming-xing Luo,⁵ Xing-Bo Yuan,⁶ and Hua Xing Zhu^{1,2}

How are people using agentic AI?

- Particle physicists are building custom harnesses for automating

- anomaly detection

Agents of Discovery

Sascha Diefenbacher¹, Anna Hallin²,
Gregor Kasieczka², Michael Krämer³, Anne Lauscher⁴, Tim Lukas²,

- pheno tasks [ArgoLOOM, HEPTAPOD, MadAgents, CoLLM, FERMIACC, ColliderAgent/Magnus...]

- experimental data analysis

AI Agents Can Already Autonomously Perform Experimental High Energy Physics

Eric A. Moreno^{*†1,2}, Samuel Bright-Thonney^{*†1,2}, Andrzej Novak^{*§1,2}, Dolores Garcia^{#3},
and Philip Harris^{b1,2}

The FERMIACC: Agents for Particle Theory

Prateek Agrawal,^a Nathaniel Craig,^{a,b} Amalia Madden,^b and Iñigo Valenzuela Lomber^a

^aDepartment of Physics, University of California, Santa Barbara, CA 93106, USA

^bKavli Institute for Theoretical Physics, University of California, Santa Barbara, CA 93106, USA

- Questions:

- Frontier LLM models and harnesses improving rapidly
— how much value does a custom harness add?
 - Which frontier models/harnesses perform better than others? How much human-in-the-loop is needed?

An End-to-end Architecture for Collider Physics and Beyond

Shi Qiu,¹ Zeyu Cai,¹ Jiashen Wei,¹ Zeyu Li,^{1,2} Yixuan Yin,¹ Qing-Hong
Cao,^{1,3,2} Chang Liu,^{1,4} Ming-xing Luo,⁵ Xing-Bo Yuan,⁶ and Hua Xing Zhu^{1,2}

How are people using agentic AI?

- Particle physicists are building custom harnesses for automating

- anomaly detection

Agents of Discovery

Sascha Diefenbacher¹, Anna Hallin²,
Gregor Kasieczka², Michael Krämer³, Anne Lauscher⁴, Tim Lukas²,

- pheno tasks [ArgoLOOM, HEPTAPOD, MadAgents, CoLLM, FERMIACC, ColliderAgent/Magnus...]

- experimental data analysis

AI Agents Can Already Autonomously Perform Experimental High Energy Physics

Eric A. Moreno^{*†1,2}, Samuel Bright-Thonney^{*†1,2}, Andrzej Novak^{*§1,2}, Dolores Garcia^{#3},
and Philip Harris^{b1,2}

The FERMIACC: Agents for Particle Theory

Prateek Agrawal,^a Nathaniel Craig,^{a,b} Amalia Madden,^b and Iñigo Valenzuela Lomber^a

^aDepartment of Physics, University of California, Santa Barbara, CA 93106, USA

^bKavli Institute for Theoretical Physics, University of California, Santa Barbara, CA 93106, USA

- Questions:

- Frontier LLM models and harnesses improving rapidly — how much value does a custom harness add?
 - Which frontier models/harnesses perform better than others? How much human-in-the-loop is needed?
 - Automation vs quality control — how to avoid “AI slop”?

An End-to-end Architecture for Collider Physics and Beyond

Shi Qiu,¹ Zeyu Cai,¹ Jiashen Wei,¹ Zeyu Li,^{1,2} Yixuan Yin,¹ Qing-Hong
Cao,^{1,3,2} Chang Liu,^{1,4} Ming-xing Luo,⁵ Xing-Bo Yuan,⁶ and Hua Xing Zhu^{1,2}

How are people using agentic AI?

- Particle physicists are building custom harnesses for automating

Increasing interest in rigorous benchmarking

- anomaly detection
- pheno tasks [ArgoLOOM, HEPTAPOD, MadAgents, CoLLM, FERMIACC, ColliderAgent/Magnus...]

Sascha Diefenbacher¹, Anna Hallin²,
Gregor Kasieczka², Michael Krämer³, Anne Lauscher⁴, Tim Lukas²,

- experimental data analysis

AI Agents Can Already Autonomously Perform
Experimental High Energy Physics

Eric A. Moreno^{*†1,2}, Samuel Bright-Thonney^{*†1,2}, Andrzej Novak^{*§1,2}, Dolores Garcia^{#3},
and Philip Harris^{b1,2}

The FERMIACC: Agents for Particle Theory

- Questions:

- Frontier LLM models and harnesses improving rapidly — how much value does a custom harness add?
- Which frontier models/harnesses perform better than others? How much human-in-the-loop is needed?
- Automation vs quality control — how to avoid “AI slop”?

An End-to-end Architecture for Collider Physics and Beyond

Shi Qiu,¹ Zeyu Cai,¹ Jiashen Wei,¹ Zeyu Li,^{1,2} Yixuan Yin,¹ Qing-Hong
Cao,^{1,3,2} Chang Liu,^{1,4} Ming-xing Luo,⁵ Xing-Bo Yuan,⁶ and Hua Xing Zhu^{1,2}

Prateek Agrawal,^a Nathaniel Craig,^{a,b} Amalia Madden,^b and Iñigo Valenzuela Lomber^a

^aDepartment of Physics, University of California, Santa Barbara, CA 93106, USA

^bKavli Institute for Theoretical Physics, University of California, Santa Barbara, CA 93106, USA

How are people using agentic AI?

- Particle physicists are building custom harnesses for automating

Increasing interest in rigorous benchmarking

- anomaly detection
- pheno tasks [Argo, CoLLM, FERMIACC]
- experimental data

Sascha Diefenbacher¹, Anna Hallin²,
Goran Kaniadze², Michael Kerner³, Anna Lammertner⁴, Tim Lukas²,

See: **ColliderBench** [2605.13950]

Faroughy, Palacios Schweitzer, Pang, Mishra-Sharma & DS

the FERMIACC: Agents for Particle Theory

A new benchmark for agentic AI systems based on reinterpreting LHC searches (“recasting”)

Teek Agrawal,^a Nathaniel Craig,^{a,b} Amalia Madden,^b and Iñigo Valenzuela Lomber^a
^aDepartment of Physics, University of California, Santa Barbara, CA 93106, USA
^b Kavli Institute for Theoretical Physics, University of California, Santa Barbara, CA 93106, USA

- Questions:

An End-to-end Architecture for Collider Physics and Beyond

Shi Qiu,¹ Zeyu Cai,¹ Jiashen Wei,¹ Zeyu Li,^{1,2} Yixuan Yin,¹ Qing-Hong Cao,^{1,3,2} Chang Liu,^{1,4} Ming-xing Luo,⁵ Xing-Bo Yuan,⁶ and Hua Xing Zhu^{1,2}

- Frontier LLM models and harnesses improving rapidly — how much value does a custom harness add?
- Which frontier models/harnesses perform better than others? How much human-in-the-loop is needed?
- Automation vs quality control — how to avoid “AI slop”?

Outline

- Learning to unscramble: general framework
- **Application 1:** Dilog identities [skip for lack of time]
- **Application 2:** Tree-level YM scattering amplitudes
- **Application 3:** IBP reduction of Feynman loop integrals
- Lessons learned
- Bonus material: more details on **ColliderBench**

Learning to Unscramble: general framework

- Goal: ML methods for mathematical (symbolic) simplification

$$E_{\text{complicated}} = X_1 + X_2 + X_3 + X_4 + \dots$$



mathematical identities

$$E_{\text{simple}} = Y_1$$

- Test beds: dilog sums, tree-level YM scattering amplitudes

(Dersy, Schwartz, Zhang [2206.04115](#); Cheung, Dersy, Schwartz [2408.04720](#))

Training data




- ML methods need training data
- **Key insight:** can think of complicated expressions as *scrambles* of simple ones by mathematical identities.
- Process of simplification is then *unscrambling*. Dersy, Schwartz, Zhang [2206.04115](#)
- *Strong analogy with Rubik's cubes!*
- Can cheaply create training data by applying random identities to simple expressions



Prior approaches: RL

Dersy, Schwartz, Zhang 2206.04115

MDP: Markov Decision Process
action classifier $\pi(a | s)$,
depends only on previous state




RL (PPO, TRPO) could not beat classical algorithm on even very simple simplification tasks (small dlog sums simplifying to zero).

Challenges: very sparse reward signal, non-monotonic reduction paths

Successful RL is very compute intensive and requires a lot of fine tuning.

Prior approaches: end-to-end regression

Cheung, Dersy, Schwartz [2408.04720](#)




End-to-end regression with a seq2seq transformer achieved decent performance (80-90%) for dilog sums and scattering amplitudes

Drawback: mathematical correctness not guaranteed, transformer can hallucinate simpler but incorrect expressions!

Our approach: “Learning to Unscramble”

DS 2603.11164



Instead of unsupervised MDP of RL, train a **self-supervised MDP** on **reversed random scramble sequences**

Learns to predict the best action from unscrambling steps

***Addresses limitations of both previous approaches:
highly dense reward signal and no hallucinations!***

Application 2: scattering amplitudes

[DS 2603.11164]

Laporta algorithm

r: total denominator
(propagator) weight

$$w(I[\mathbf{a}]) = \left(\sum_i \max(a_i, 0), \sum_i |\min(a_i, 0)|, \mathbf{w}_3 \right)$$

s: total numerator
weight

- Basis for general purpose IBP reduction programs like **Kira**

- Make a giant linear system of IBPs
- Use Gaussian elimination to target highest weight IBP
- Continue until fully reduced to master integrals

$$\begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1k} \\ a_{21} & a_{22} & \cdots & a_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ a_{k1} & a_{k2} & \cdots & a_{kk} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_k \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_k \end{bmatrix},$$

$$\left[\begin{array}{cccc|c} a'_{11} & a'_{12} & \cdots & a'_{1k} & b'_1 \\ 0 & a'_{22} & \cdots & a'_{2k} & b'_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & a'_{kk} & b'_k \end{array} \right].$$

- Guaranteed to work: number of IBP identities grows faster than the number of seeds
- But very computationally inefficient — **memory wall**

Self-supervised AI for Loop Integral Reduction (SAILIR)

- Learning to Unscramble applied to IBP reduction
- Some challenges unique to IBP reduction:
 - weight reduction instead of term reduction
 - much bigger action space (number of IBPs combinatorially large)
 - but set of valid actions much smaller (10-100 per expression) but variable (depends on expression)
 - need special classifier architecture to handle variable action space

Thanks!

