Accelerating lattice gauge theory studies with Agentic AI

Venkitesh Ayyar, Daniel Hackett, Michael Wagman
Fermi National Laboratory, Batavia, IL, USA

Fermilab report number:
FERMILAB-POSTER-26-0191-CSAID-T

AI LLM tools

- Very good at text data processing
- Apt at code development and optimization
- Rapidly improving

✦Very suitable for application in scientific computational workflows such as Lattice QCD

✦**AI** can and does make mistakes

✦Scientific accuracy demands rigorous testing and validation

AI assisted scientific workflow

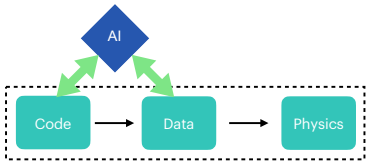

Validating 'vibe coding' for gauge generation workflows

Case study: Lattice QCD Hybrid Monte-Carlo (HMC) workflow

Unit testing
- Test individual pieces: CG, multi-shift CG, Multigrid
- Test Solvers: Residual check
- Test Action accuracy: Compare with reference implementation
- Test machine precision agreement for action/force on random gauge fields

Optimization
- Incorporate gpu kernels from different libraries
- Compare memory and performance optimized versions with simple implementations
- Chroma ↔ Grid ↔ Grid-quda

Finite-difference testing

Checking Forces: Numerical derivative of action converges to forces

Functional testing

Confirm stability of physics observables like Plaquette, Polyakov loop

Use of AI tools

- Build packages Grid[1] and Quda[2]
- Build code for HMC generation
- Code compilation
- Directed code execution and monitoring on HPC clusters
- Data extraction: Parsing of log files
- Documentation of results: Tables and plots

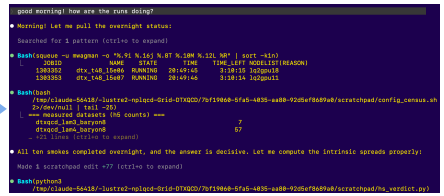

Parameter tuning: eg. Hasenbusch Mass tuning

Hasenbusch mass preconditioning[3]

$$e^{-S_{fermion}} \sim \det(M) = \left(\frac{\det(M)}{\det(M_1)}\right)\left(\frac{\det(M_1)}{\det(M_2)}\right)\ldots\det(M_n)$$

- Standard technique to significantly accelerate HMC
- Multiple solves, but smaller forces from each piece

Hard for humans, easy for AI

- Effort-intensive
- Hard to scan large parameter space

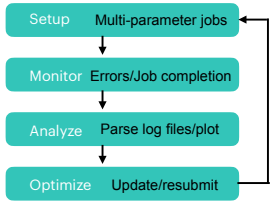

A warm up: Shared Multigrid (MG) vs CG tuning

- Pure CG: Slow at light masses
- Separate MG null-spaces: High setup cost and memory
- Optimal: Shared Multigrid null-spaces for cluster of light masses, CG after crossover

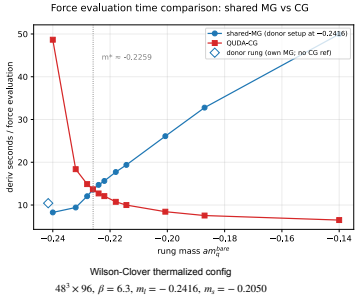

Hasenbusch mass tuning

- Overall step size limited by largest forces
- Competing tradeoffs complicate optimization

Split into a fine ladder

Broad spread of masses to balance forces

Multi-level integrator split
- Large masses: faster solves
- Break into finer steps in multi-level integrator

Comparing three splits

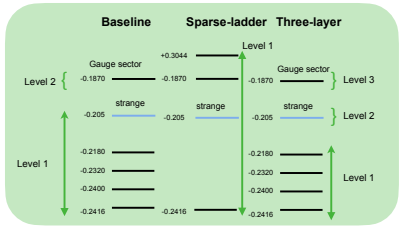

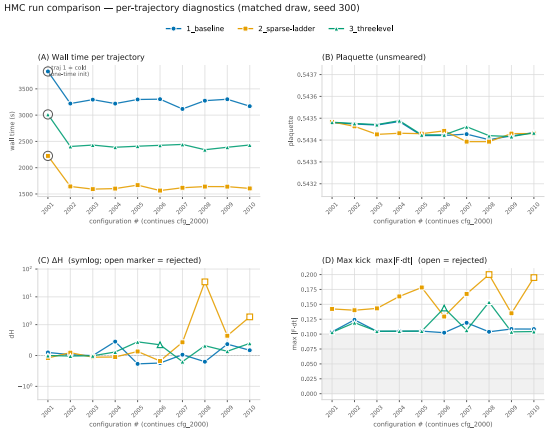

Per-trajectory HMC statistics for the baseline vs. two AI-explored tuning strategies, matched 20-trajectory comparison

Run	Trajectories run	Acceptance rate	Avg wall time, steady-state (s)	Avg dH	Median dH	Max dH	Avg e^{-dH}	Max e^{-dH}	Avg max|F·dt| (per-traj limiter)	Max |F·dt| observed
1_baseline	20	95.0% (19/20)	**3257**	0.079	0.093	0.45	0.943	1.310	0.108	0.124
2_sparse-ladder	20	65.0% (13/20)	**1646**	57.25	0.133	**1022.23**	0.813	3.372	0.258	1.984
3_threelevel	20	95.0% (19/20)	**2394**	0.144	0.155	0.55	0.900	1.657	0.112	0.154

Summary

AI-driven scans produced two alternate HMC configurations:
- Sparser Hasenbusch ladder ~**2×** faster, but frequent Metropolis rejections from occasional force spikes (65% acceptance)
- Restructured 3-level integrator ~**1.3×** faster while matching the baseline's acceptance rate (95%)

Preliminary — illustrates AI's ability to rapidly explore the speed/stability tradeoff, not a final recommended tuning

Future direction

- Generalized approach to obtain Hasenbusch mass split in a reproducible manner that can be applied to new ensembles
- Incorporate into agentic frameworks for automated workflow execution (FemtoMind[4]) and analysis
- Structured gauge generation workflows such as thermalization studies

1. https://github.com/paboyle/Grid
2. https://github.com/lattice/quda
3. Martin Hasenbusch, Speeding up the hybrid Monte Carlo algorithm for dynamical fermions, Phys.Lett.B 519 (2001) 177-182
4. https://www.jlab.org/AI/Projects/Femtomind

Acknowledgements

National Energy Research Scientific Computing Center

Computational resources: Fermilab and NERSC
AI resources: Fermilab, AmSC, Anthropic, OpenAI