AI4Science – LANL's ArtIMis Investment

Nathan DeBardeleben, Ph.D.
HPC-DESign
ArtIMis Co-Lead
~~ML Lead for Enabling Manufacturing~~

February, 2026

LA-UR-26-21430

Industry has led the way on AI for science but is unlikely to lead a sustained effort.

- Need institutions with long horizons, open verification, and mission alignment
- Commercial incentives favor demo cycles, product pivots, and PR wins
- DeepMind's materials project walked back after analysis by LBNL
- Meta's Galactica taken offline days after release
- IBM sold off Watson Health – market alignment diverged from scientific utility

Materials at DeepMind

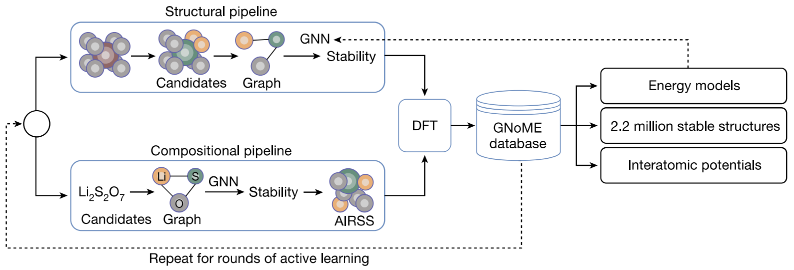

Less successful than initial claims.

Earth Systems at Microsoft

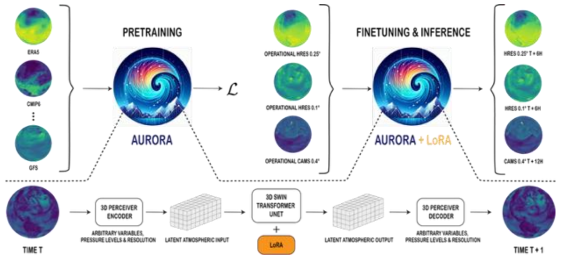

Effort abandoned.

DOE has invested significant effort to develop an AI strategy.

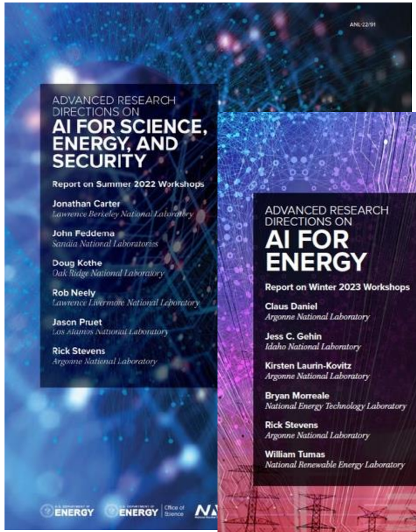
ANL-22/91

ADVANCED RESEARCH DIRECTIONS ON
AI FOR SCIENCE, ENERGY, AND SECURITY
Report on Summer 2022 Workshops

Jonathan Carter
Lawrence Berkeley National Laboratory

John Feddema
Sandia National Laboratories

Doug Kothe
Oak Ridge National Laboratory

Rob Neely
Lawrence Livermore National Laboratory

Jason Pruet
Los Alamos National Laboratory

Rick Stevens
Argonne National Laboratory

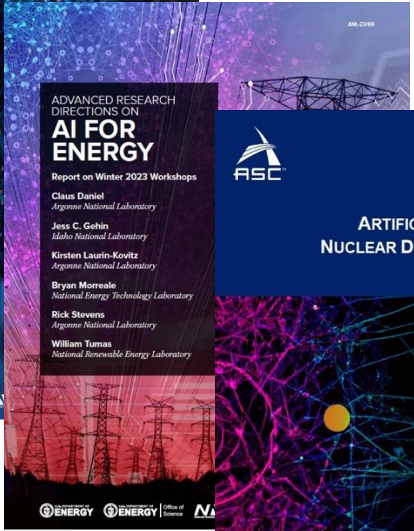
ANL-23/69

ADVANCED RESEARCH DIRECTIONS ON
AI FOR ENERGY
Report on Winter 2023 Workshops

Claus Daniel
Argonne National Laboratory

Jess C. Gehin
Idaho National Laboratory

Kirsten Laurin-Kovitz
Argonne National Laboratory

Bryan Morreale
National Energy Technology Laboratory

Rick Stevens
Argonne National Laboratory

William Tumas
National Renewable Energy Laboratory

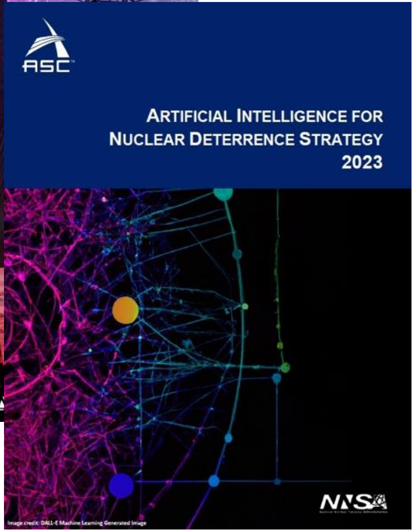
ASC

ARTIFICIAL INTELLIGENCE FOR NUCLEAR DETERRENCE STRATEGY 2023

Image credit: DALL-E Machine Learning Generated Image

Genesis Mission

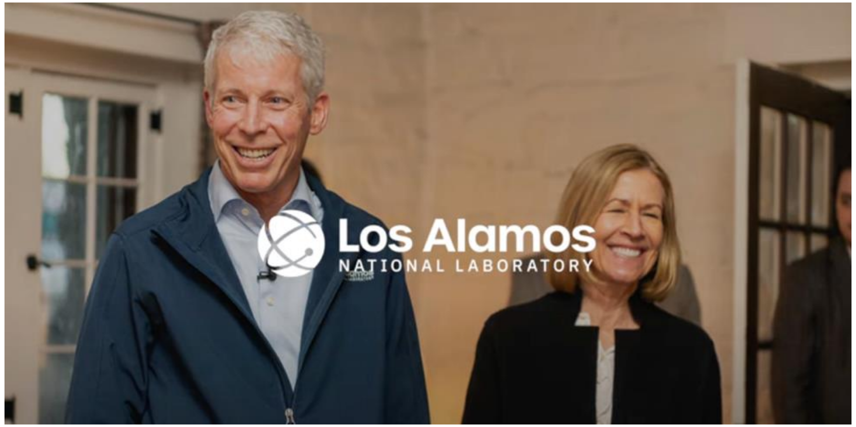

Secretary Wright visited LANL to discuss Manhattan Project 2.0.

https://www.youtube.com/watch?v=Gyr_wOSy_xQ

Application Areas

Our goal is to develop an ecosystem of agents and foundation models to address DOE mission needs.

ArtIMis

AI for Mission

Enabling Technologies

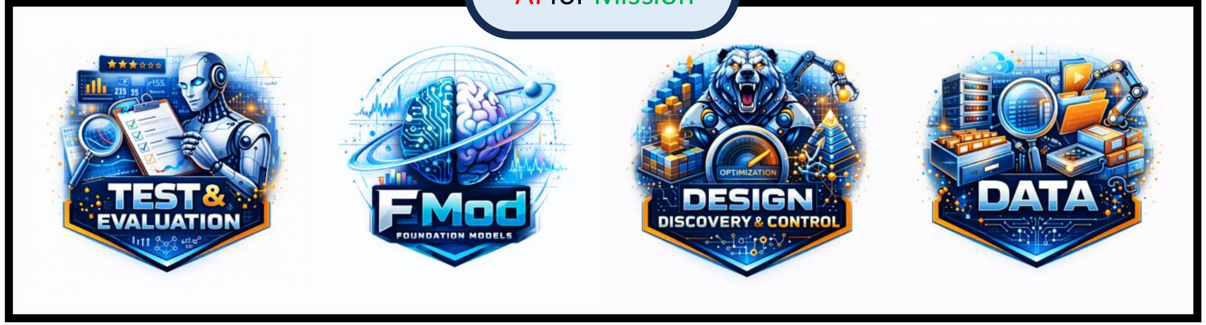

Los Alamos NATIONAL LABORATORY

Application Areas

Emily discussed this yesterday

Moved to lab-wide effort

Enabling Technologies

We have built over a dozen special purpose science foundation models including . . .

Stochastic Latent Transformers and VAEs

LATTE

GPT-like (discrete tokens)

SPUS

UNet-based

Shape-agnostic PDE Foundation Model
https://huggingface.co/mahindrautela/MORPH

MORPH
PHYSICS FOUNDATION MODEL

Transformer-based

Most models are:

- Trained on autoregressive roll-outs
- Multi-modal (input decks, equations, simulation data)
- Fine-tuned on new datasets (across domain)
- Fine-tuned for specific scalar outputs (time to fracture, physics numeric quantity, etc.)

<u>True foundation models</u> – not just surrogate models

CHELATRON

TRANSPARENT EARTH

Bomberman

TellusFM

FindingPhysics

Looking for a good name still ☺

MORPH: PDE Foundation Models with Arbitrary Data Modality

Results on PLI dataset (MPMM)

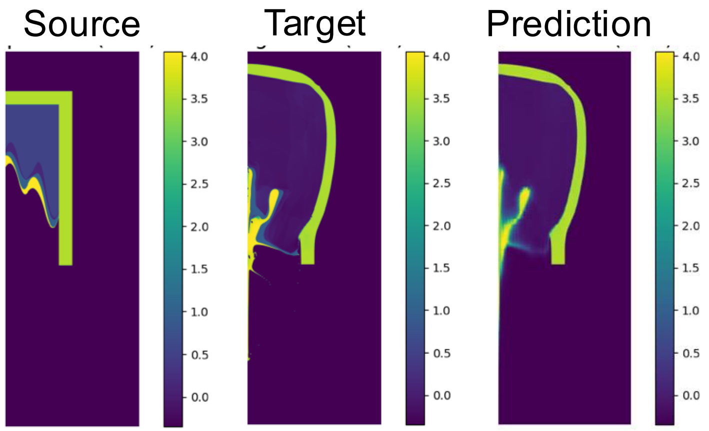

Pretraining advantage

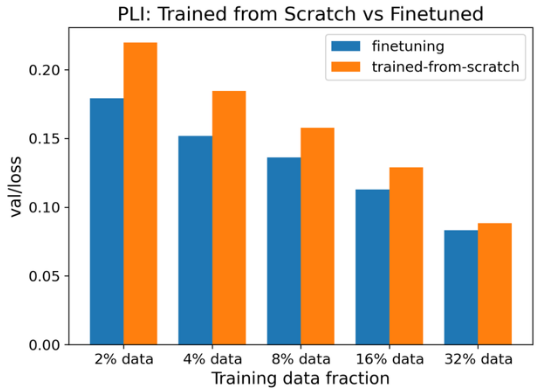

Results on JAG-ICF dataset (LLNL)

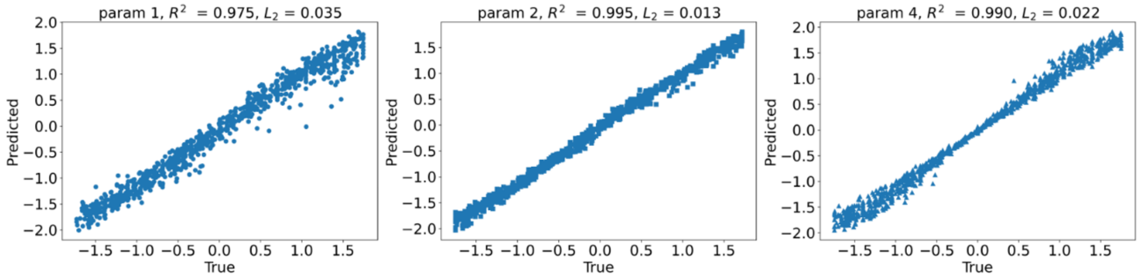

Parameter estimation

- R^2 = 0.995
- MSE = 0.0235

Recons MSE:

- 0.0012

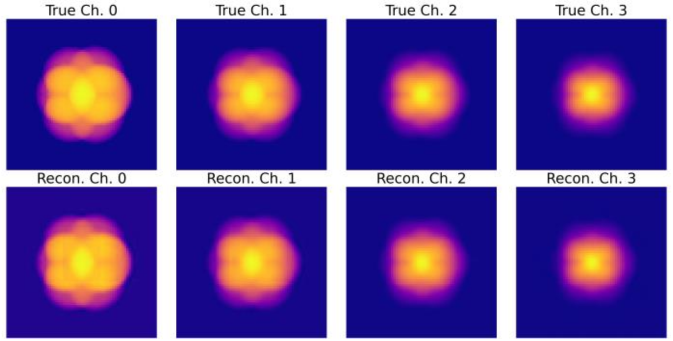

Pretraining advantage

- 60% for 10% data.
- Reduces later

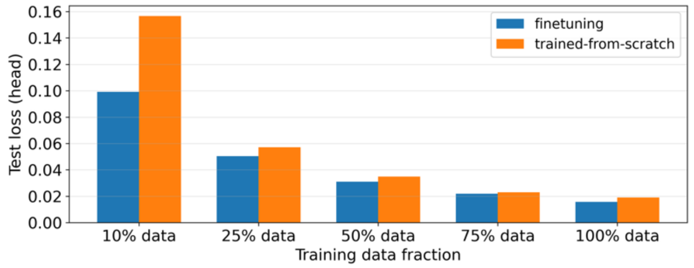

LANL Staff Can Use Surrogate Models to Explore Designs

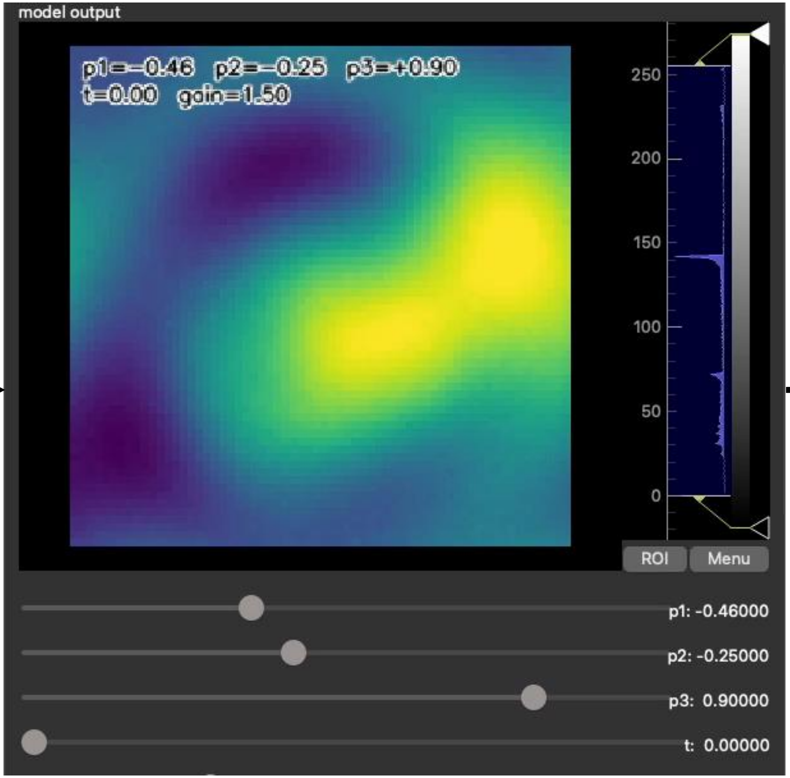

p1: -0.46000
p2: -0.25000
p3: 0.90000
t: 0.00000

But we can also drive those FMs with reasoning agents

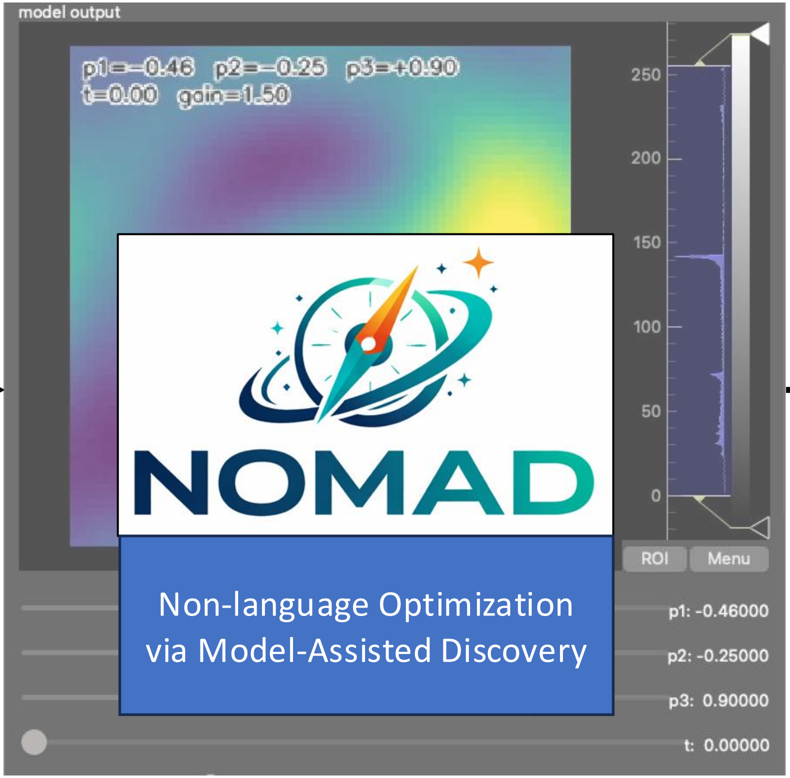

Non-language Optimization via Model-Assisted Discovery

We are building an agentic design and discovery system driven by industry reasoning models.

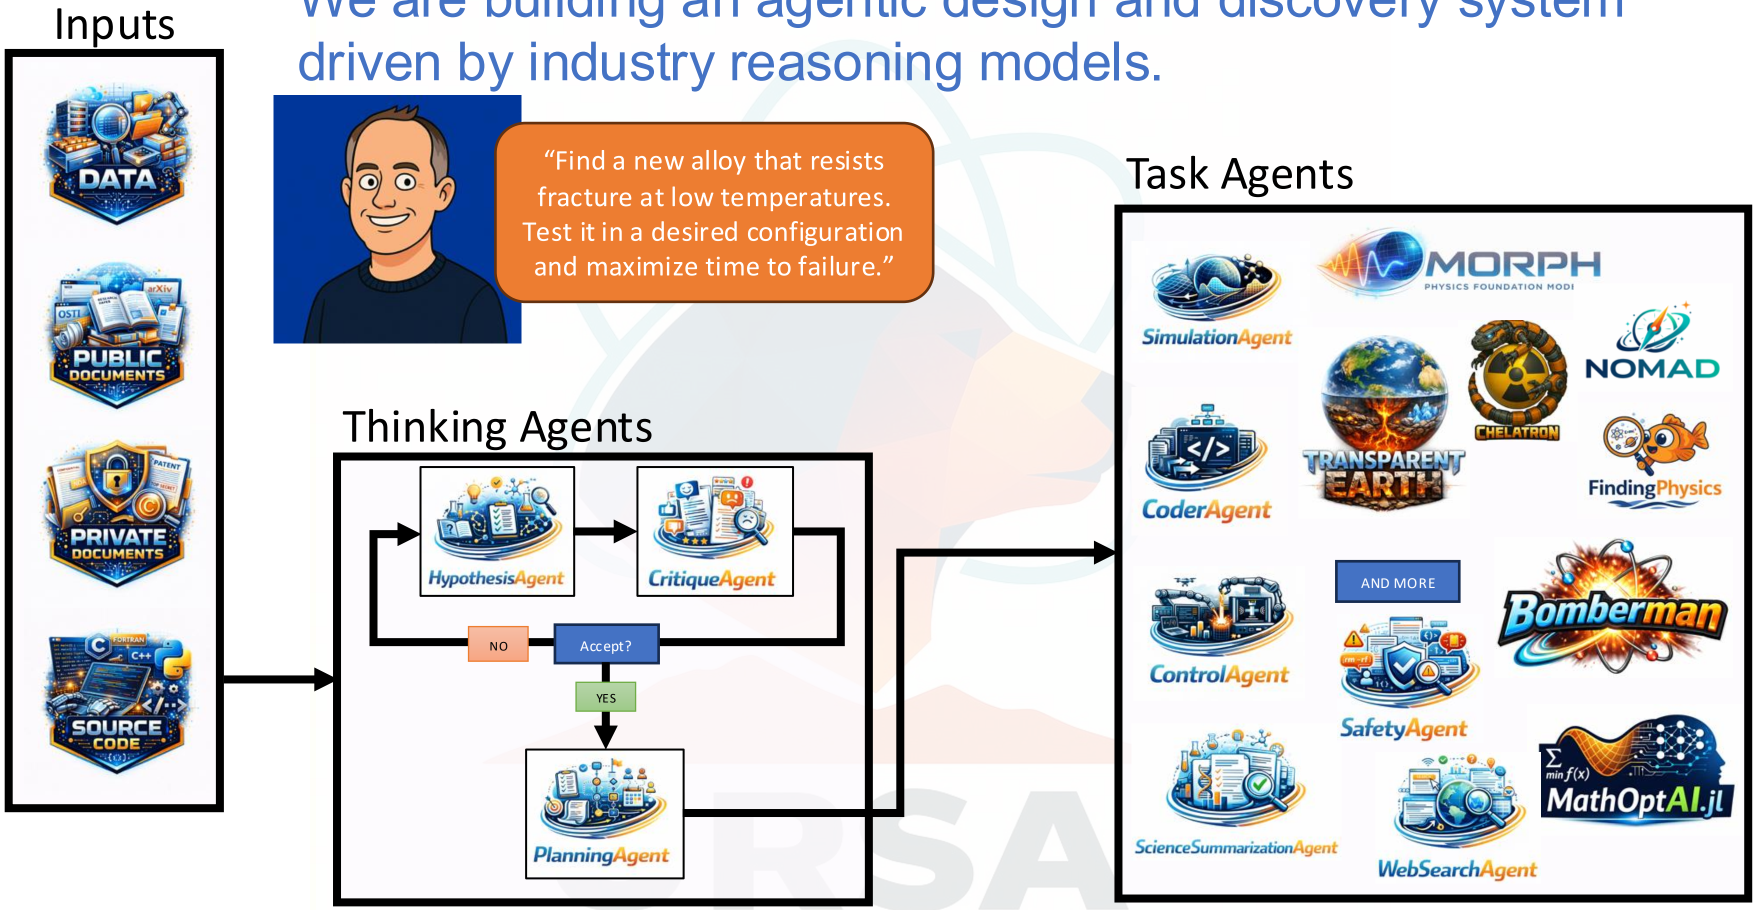

"Find a new alloy that resists fracture at low temperatures. Test it in a desired configuration and maximize time to failure."

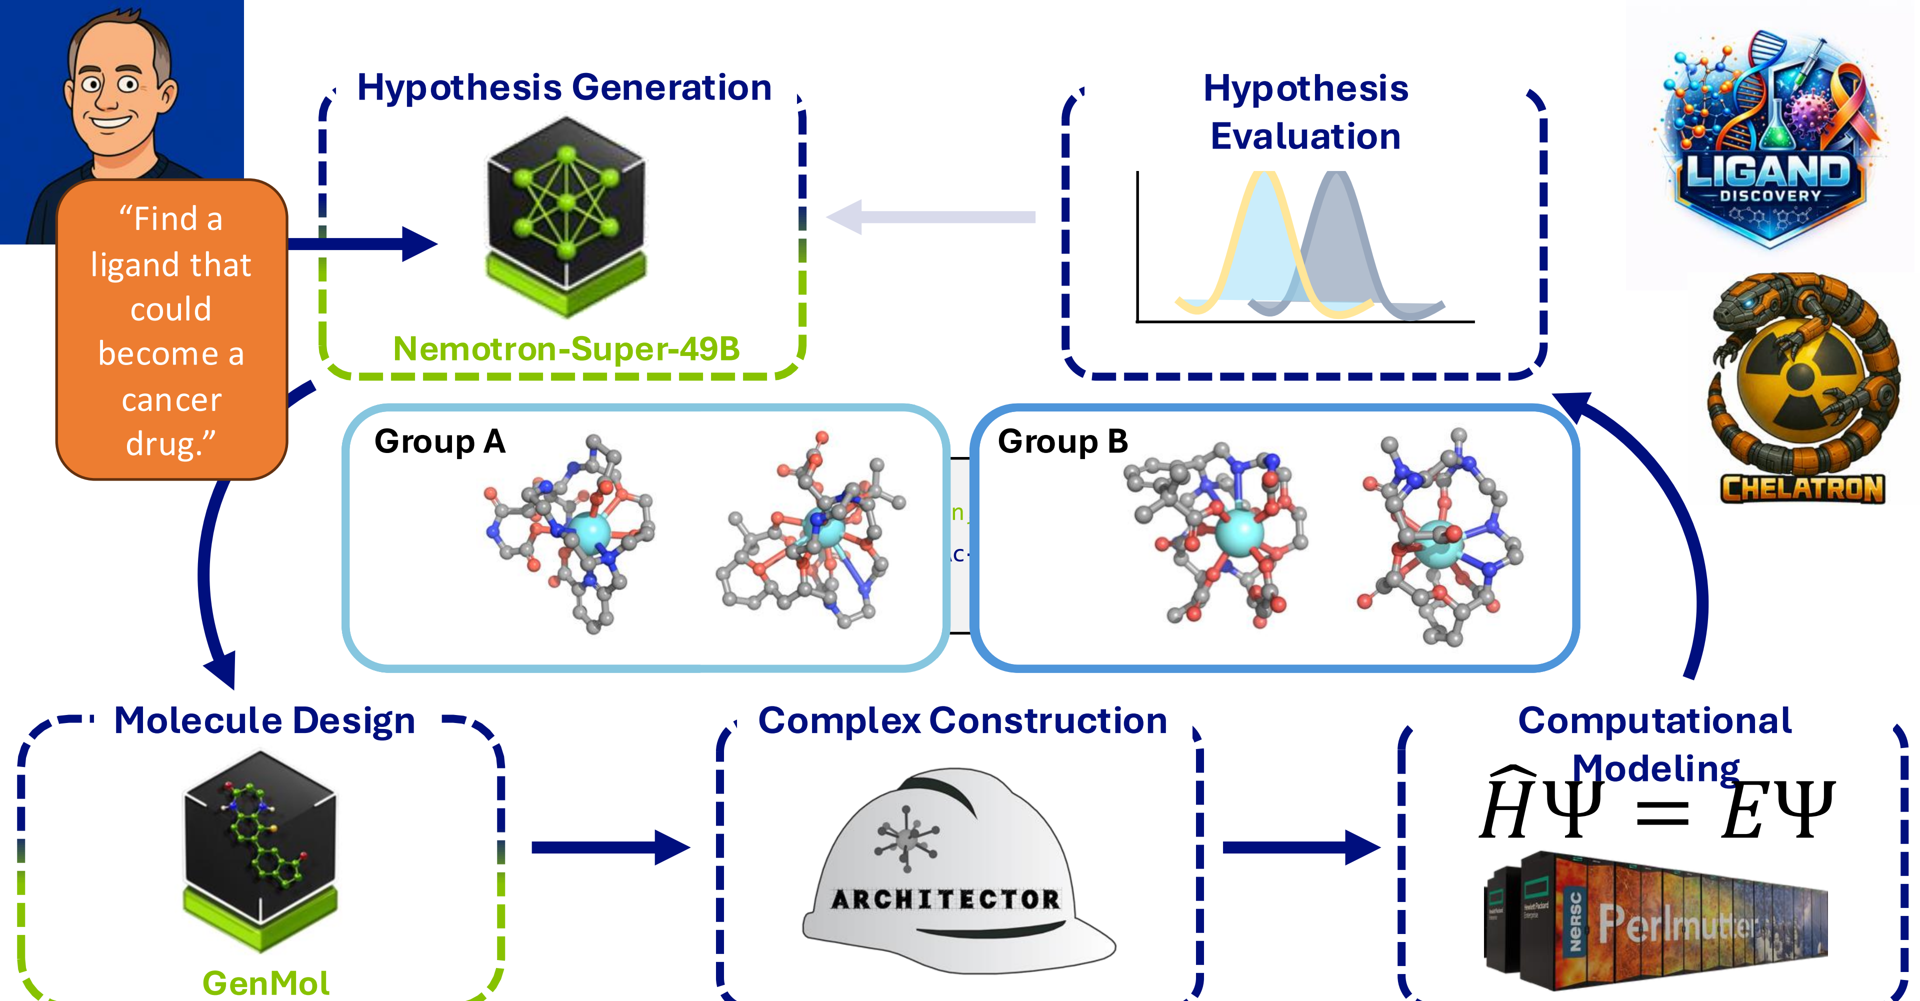
Hypothesis Generation
Nemotron-Super-49B
Hypothesis Evaluation
LIGAND
DISCOVERY
CHELATRON
"Find a ligand that could become a cancer drug."
Group A
Group B
Molecule Design
GenMol
Complex Construction
ARCHITECTOR
Computational Modeling
$\hat{H}\Psi = E\Psi$
NERSC
Perlmutter

We are automating a materials design lab for alloy discovery.

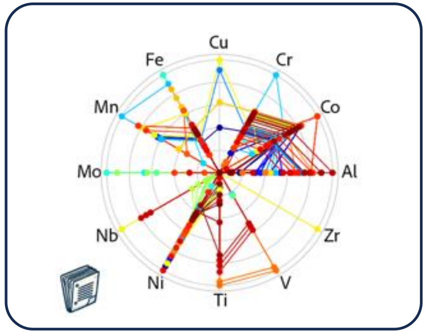

Data collection from literature. Automated.

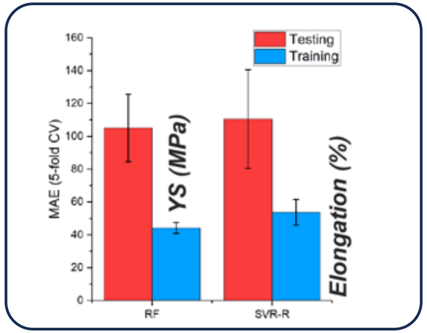

Property prediction. Partially automated

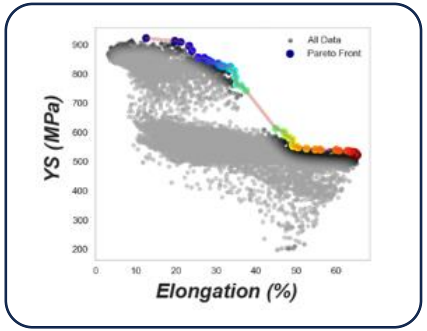

Model Inversion. Partially automated.

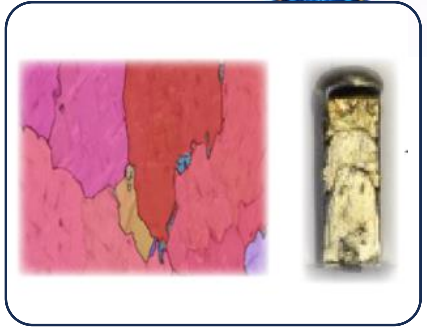

Manufacturing and testing. Partially automated.

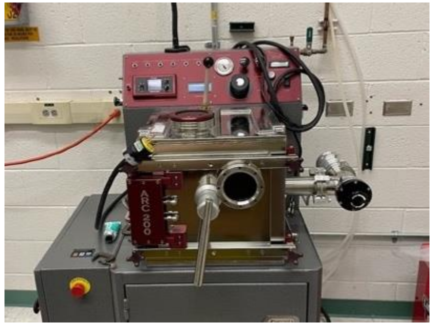

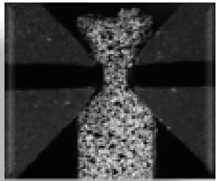

Team Leads: Saryu Fensin and Tim Germann

Testing is automated. Automated manufacturing this summer.

The fracture foundation model predicts PBX fractures that match physics sims, but 10,000x times faster.

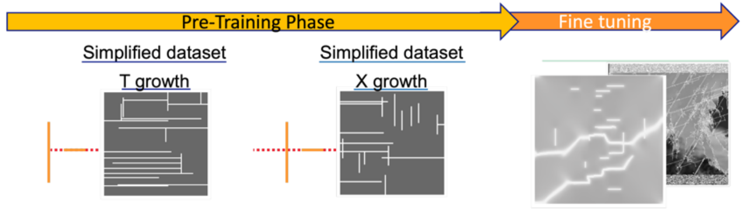

The transformer-based model is pre-trained on simplistic fracture data and then fine-tuned with physics-based simulations.

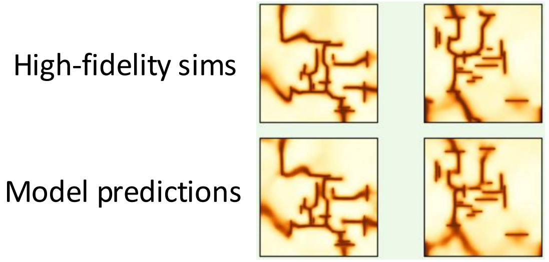

The fine-tuned PBX model matches high-fidelity simulations with ~10% error, but is 10,000x faster.

The foundation model is more accurate fine-tuned with 750 simulations than training from scratch with 15,000 simulations.

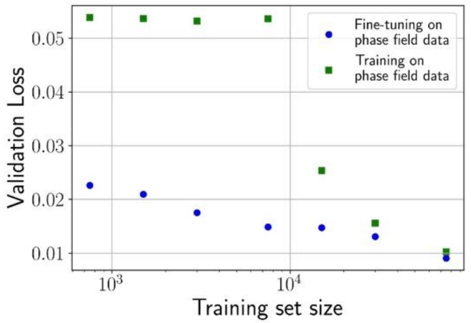

With 15,000 sims, the foundation model gives lower RMSE and more visually correct fracture patterns than training from scratch.

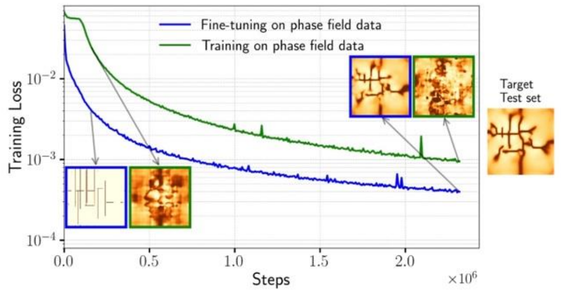

Team Leads: Hari Viswanathan and Dan O'Malley

Models include tuning we have found useful

- Curriculum learning – tasks get progressively harder
- Scheduled sampling – sample from either ground truth or rollout during training
- "Pushforward trick"
- Data augmentation (including noise injection)
- Domain-relevant loss functions
- And more

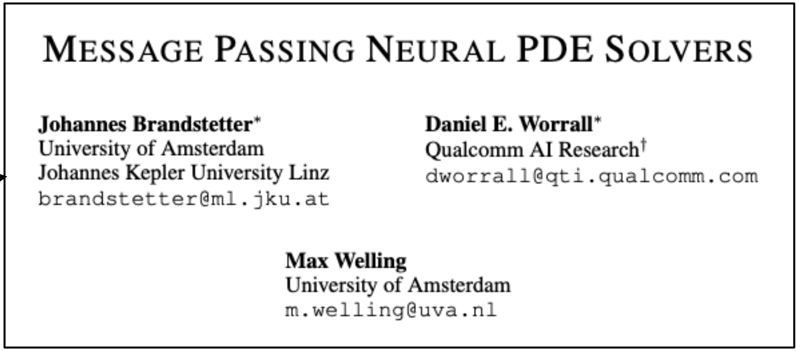

MESSAGE PASSING NEURAL PDE SOLVERS

Johannes Brandstetter*
University of Amsterdam
Johannes Kepler University Linz
brandstetter@ml.jku.at

Daniel E. Worrall*
Qualcomm AI Research†
dworrall@qti.qualcomm.com

Max Welling
University of Amsterdam
m.welling@uva.nl

We have invested heavily in improving auto-regressive rollouts

BEFORE OPTIMIZATION

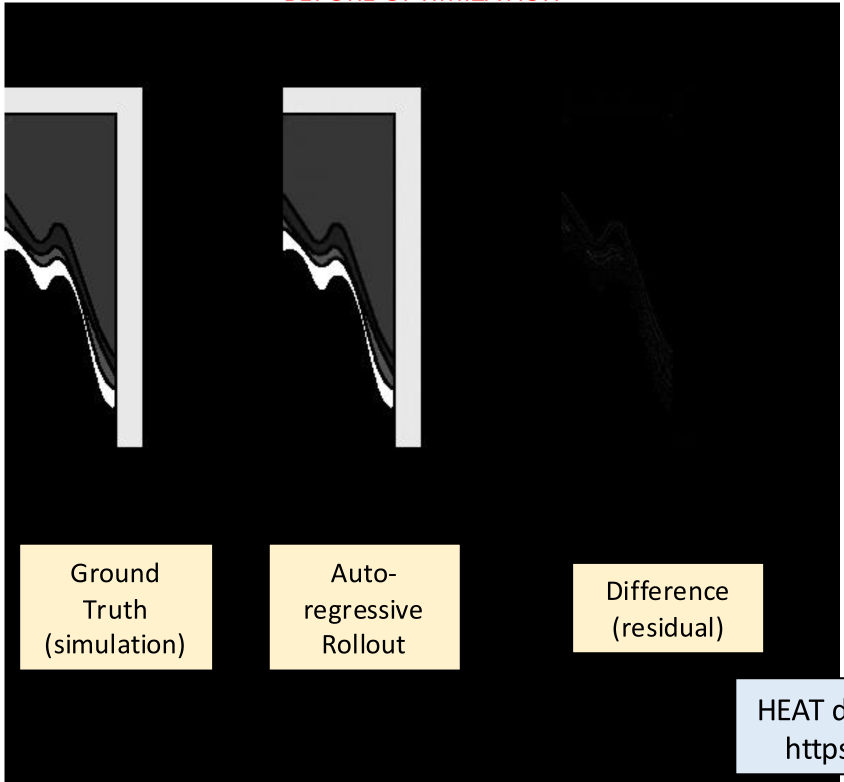

AFTER OPTIMIZATION

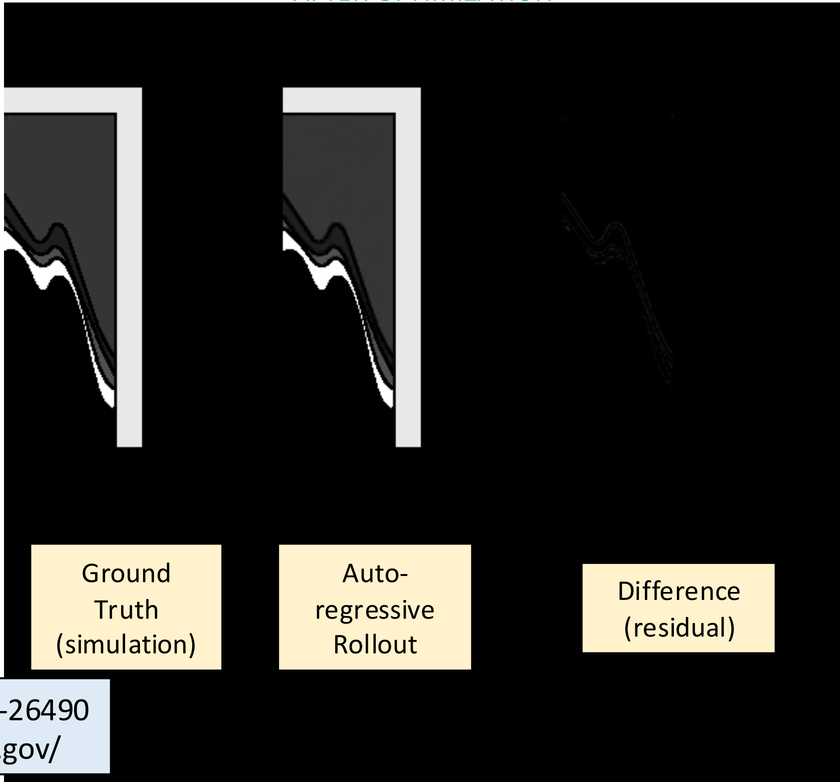

HEAT dataset: LA-UR-25-26490
https://oceans11.lanl.gov/

Michigan’s MIST model has found application in our bio and HE work.

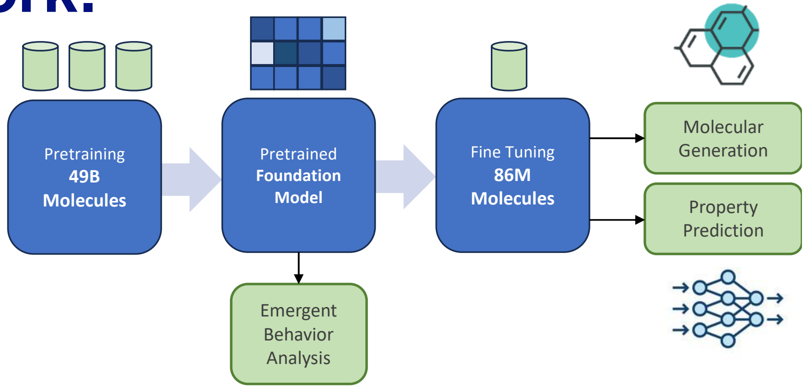

Pretraining 49B Molecules
Pretrained Foundation Model
Fine Tuning 86M Molecules
Molecular Generation
Property Prediction
Emergent Behavior Analysis

Fine-tuned with LANL-curated data, MIST can predict properties for energetic molecules.

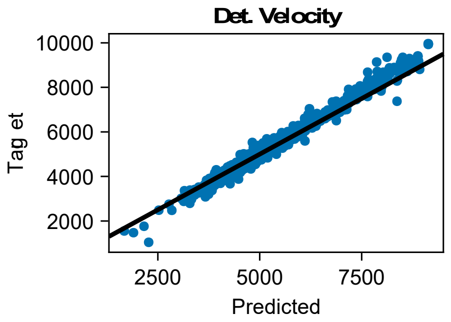

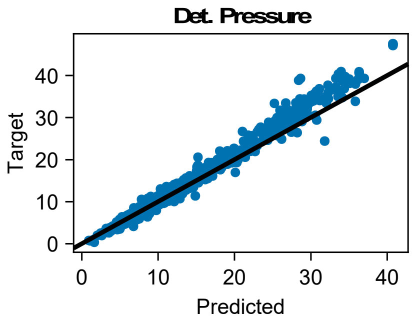

Fine-tuned MIST embeddings let the Bio team vastly improve token prediction related to small-molecule binding.

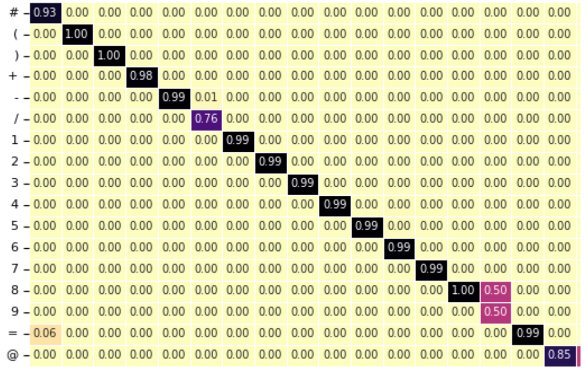

#	0.93	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
(	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
)	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
+	0.00	0.00	0.00	0.98	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-	0.00	0.00	0.00	0.00	0.99	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
/	0.00	0.00	0.00	0.00	0.00	0.76	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1	0.00	0.00	0.00	0.00	0.00	0.00	0.99	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.99	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
3	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.99	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
4	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.99	0.00	0.00	0.00	0.00	0.00	0.00	0.00
5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.99	0.00	0.00	0.00	0.00	0.00	0.00
6	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.99	0.00	0.00	0.00	0.00	0.00
7	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.99	0.00	0.00	0.00	0.00
8	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.50	0.00	0.00
9	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.50	0.00	0.00
=	0.06	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.99	0.00
@	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.85

We have an ICF design demo combining planning, research, simulation, and data processing agents.

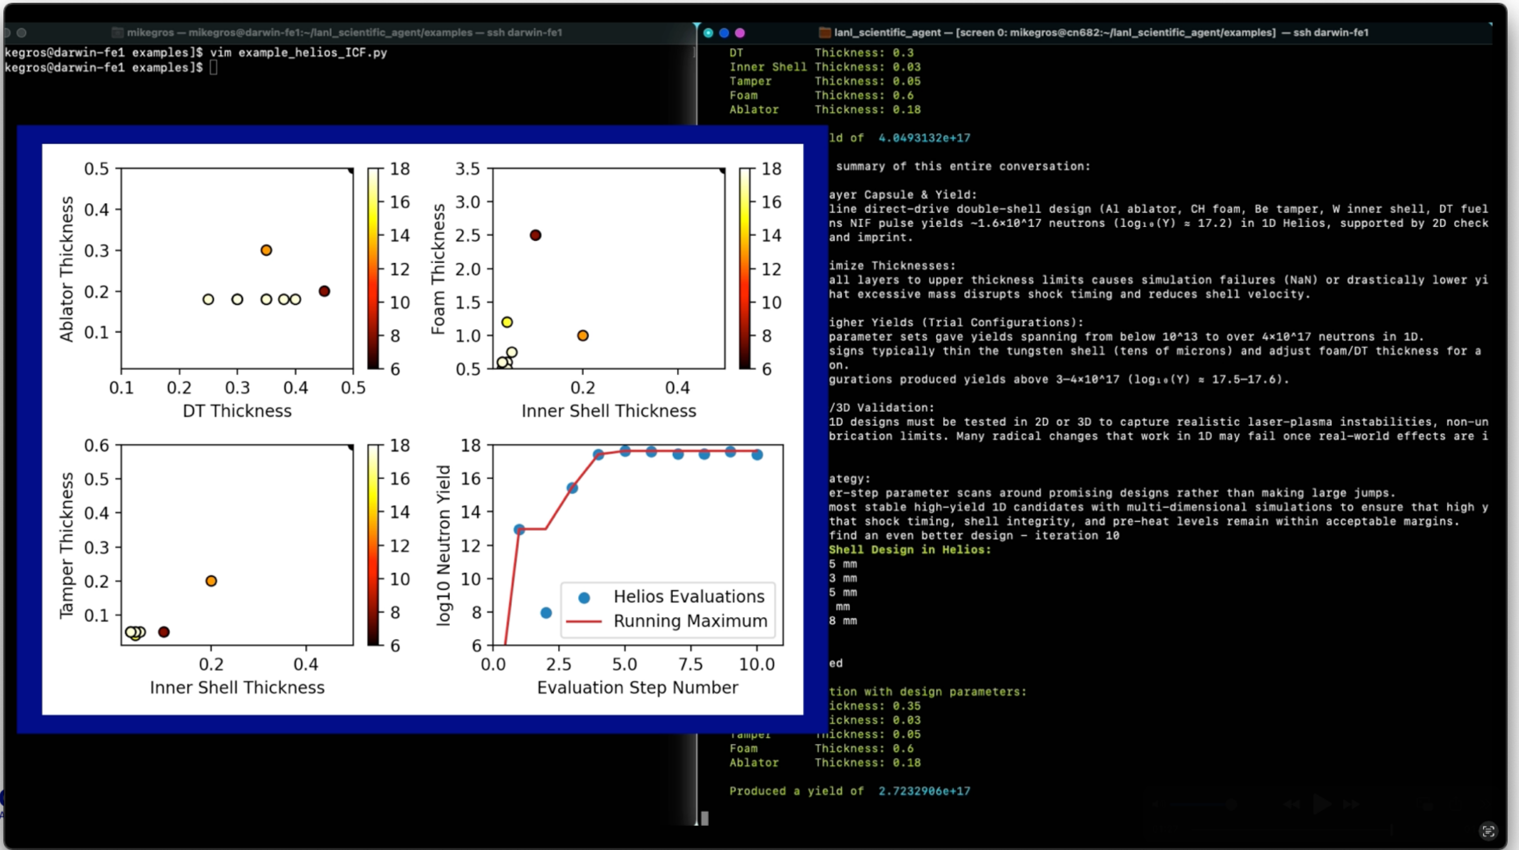

mikegros — mikegros@darwin-fe1:~/lanl_scientific_agent/examples — ssh darwin-fe1

```
kegros@darwin-fe1 examples]$ vim example_helios_ICF.py
kegros@darwin-fe1 examples]$ 
```

lanl_scientific_agent — [screen 0: mikegros@cn682:~/lanl_scientific_agent/examples] — ssh darwin-fe1

```
DT           Thickness: 0.3
Inner Shell  Thickness: 0.03
Tamper       Thickness: 0.05
Foam         Thickness: 0.6
Ablator      Thickness: 0.18

ld of  4.0493132e+17

 summary of this entire conversation:

ayer Capsule & Yield:
line direct-drive double-shell design (Al ablator, CH foam, Be tamper, W inner shell, DT fuel
ns NIF pulse yields ~1.6×10^17 neutrons (log₁₀(Y) ≈ 17.2) in 1D Helios, supported by 2D check
and imprint.

imize Thicknesses:
all layers to upper thickness limits causes simulation failures (NaN) or drastically lower yi
hat excessive mass disrupts shock timing and reduces shell velocity.

igher Yields (Trial Configurations):
parameter sets gave yields spanning from below 10^13 to over 4×10^17 neutrons in 1D.
signs typically thin the tungsten shell (tens of microns) and adjust foam/DT thickness for a
on.
gurations produced yields above 3–4×10^17 (log₁₀(Y) ≈ 17.5–17.6).

/3D Validation:
1D designs must be tested in 2D or 3D to capture realistic laser-plasma instabilities, non-un
brication limits. Many radical changes that work in 1D may fail once real-world effects are i


ategy:
er-step parameter scans around promising designs rather than making large jumps.
most stable high-yield 1D candidates with multi-dimensional simulations to ensure that high y
that shock timing, shell integrity, and pre-heat levels remain within acceptable margins.
find an even better design - iteration 10
Shell Design in Helios:
5 mm
3 mm
5 mm
 mm
8 mm


ed

tion with design parameters:
ickness: 0.35
ickness: 0.03
Tamper       Thickness: 0.05
Foam         Thickness: 0.6
Ablator      Thickness: 0.18

Produced a yield of  2.7232906e+17
```

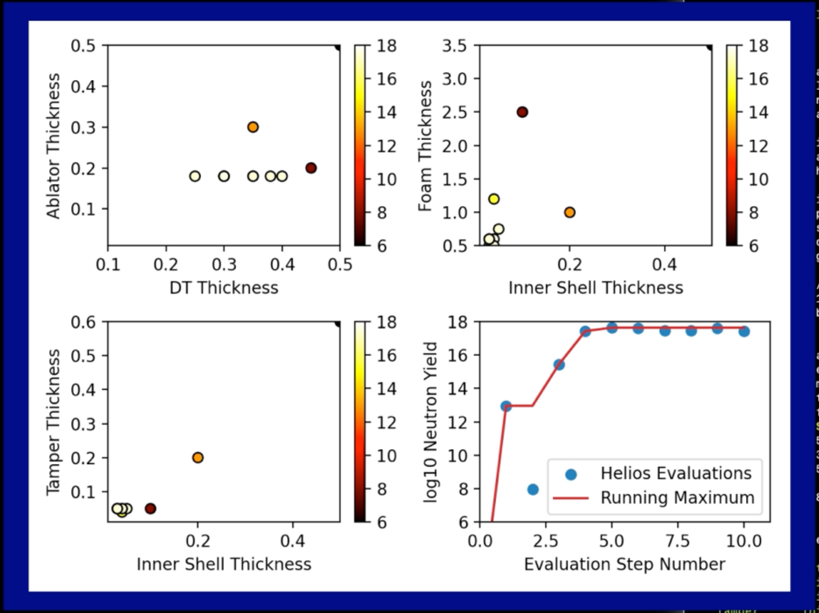

The planning agent does a careful breakdown with expected outputs and success criteria.

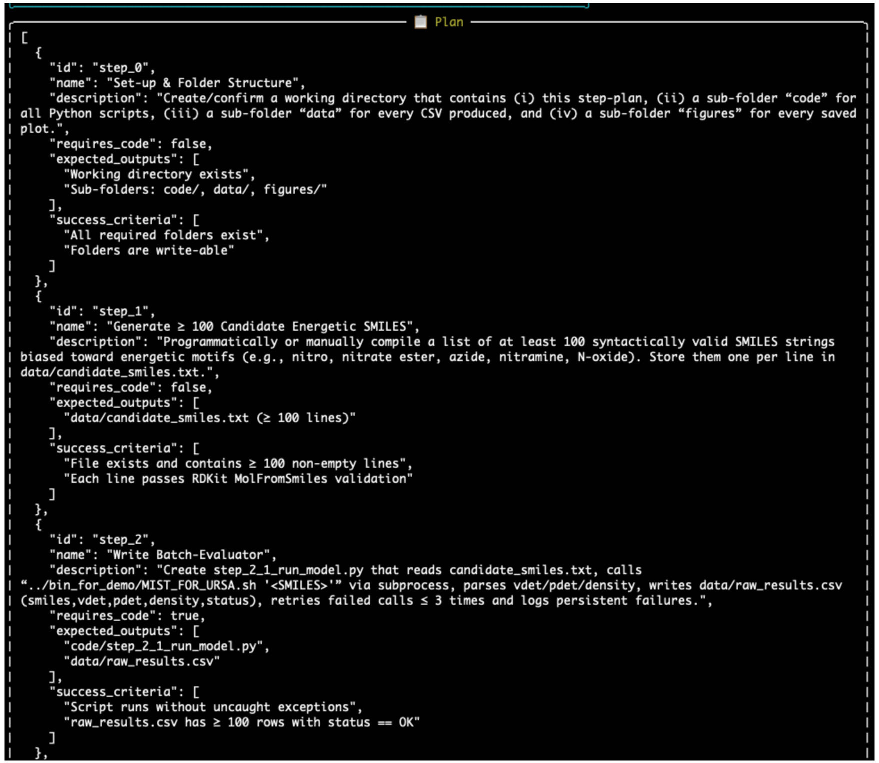

```
Plan
[
  {
    "id": "step_0",
    "name": "Set-up & Folder Structure",
    "description": "Create/confirm a working directory that contains (i) this step-plan, (ii) a sub-folder “code” for
all Python scripts, (iii) a sub-folder “data” for every CSV produced, and (iv) a sub-folder “figures” for every saved
plot.",
    "requires_code": false,
    "expected_outputs": [
      "Working directory exists",
      "Sub-folders: code/, data/, figures/"
    ],
    "success_criteria": [
      "All required folders exist",
      "Folders are write-able"
    ]
  },
  {
    "id": "step_1",
    "name": "Generate ≥ 100 Candidate Energetic SMILES",
    "description": "Programmatically or manually compile a list of at least 100 syntactically valid SMILES strings
biased toward energetic motifs (e.g., nitro, nitrate ester, azide, nitramine, N-oxide). Store them one per line in
data/candidate_smiles.txt.",
    "requires_code": false,
    "expected_outputs": [
      "data/candidate_smiles.txt (≥ 100 lines)"
    ],
    "success_criteria": [
      "File exists and contains ≥ 100 non-empty lines",
      "Each line passes RDKit MolFromSmiles validation"
    ]
  },
  {
    "id": "step_2",
    "name": "Write Batch-Evaluator",
    "description": "Create step_2_1_run_model.py that reads candidate_smiles.txt, calls
“../bin_for_demo/MIST_FOR_URSA.sh '<SMILES>'” via subprocess, parses vdet/pdet/density, writes data/raw_results.csv
(smiles,vdet,pdet,density,status), retries failed calls ≤ 3 times and logs persistent failures.",
    "requires_code": true,
    "expected_outputs": [
      "code/step_2_1_run_model.py",
      "data/raw_results.csv"
    ],
    "success_criteria": [
      "Script runs without uncaught exceptions",
      "raw_results.csv has ≥ 100 rows with status == OK"
    ]
  },
```

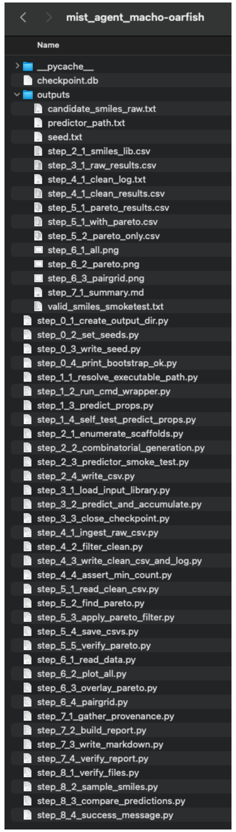

```
mist_agent_macho-oarfish
Name
> __pycache__
  checkpoint.db
v outputs
    candidate_smiles_raw.txt
    predictor_path.txt
    seed.txt
    step_2_1_smiles_lib.csv
    step_3_1_raw_results.csv
    step_4_1_clean_log.txt
    step_4_1_clean_results.csv
    step_5_1_pareto_results.csv
    step_5_1_with_pareto.csv
    step_5_2_pareto_only.csv
    step_6_1_all.png
    step_6_2_pareto.png
    step_6_3_pairgrid.png
    step_7_1_summary.md
    valid_smiles_smoketest.txt
  step_0_1_create_output_dir.py
  step_0_2_set_seeds.py
  step_0_3_write_seed.py
  step_0_4_print_bootstrap_ok.py
  step_1_1_resolve_executable_path.py
  step_1_2_run_cmd_wrapper.py
  step_1_3_predict_props.py
  step_1_4_self_test_predict_props.py
  step_2_1_enumerate_scaffolds.py
  step_2_2_combinatorial_generation.py
  step_2_3_predictor_smoke_test.py
  step_2_4_write_csv.py
  step_3_1_load_input_library.py
  step_3_2_predict_and_accumulate.py
  step_3_3_close_checkpoint.py
  step_4_1_ingest_raw_csv.py
  step_4_2_filter_clean.py
  step_4_3_write_clean_csv_and_log.py
  step_4_4_assert_min_count.py
  step_5_1_read_clean_csv.py
  step_5_2_find_pareto.py
  step_5_3_apply_pareto_filter.py
  step_5_4_save_csvs.py
  step_5_5_verify_pareto.py
  step_6_1_read_data.py
  step_6_2_plot_all.py
  step_6_3_overlay_pareto.py
  step_6_4_pairgrid.py
  step_7_1_gather_provenance.py
  step_7_2_build_report.py
  step_7_3_write_markdown.py
  step_7_4_verify_report.py
  step_8_1_verify_files.py
  step_8_2_sample_smiles.py
  step_8_3_compare_predictions.py
  step_8_4_success_message.py
```

URSA can use molecular science foundation models to explore promising compounds

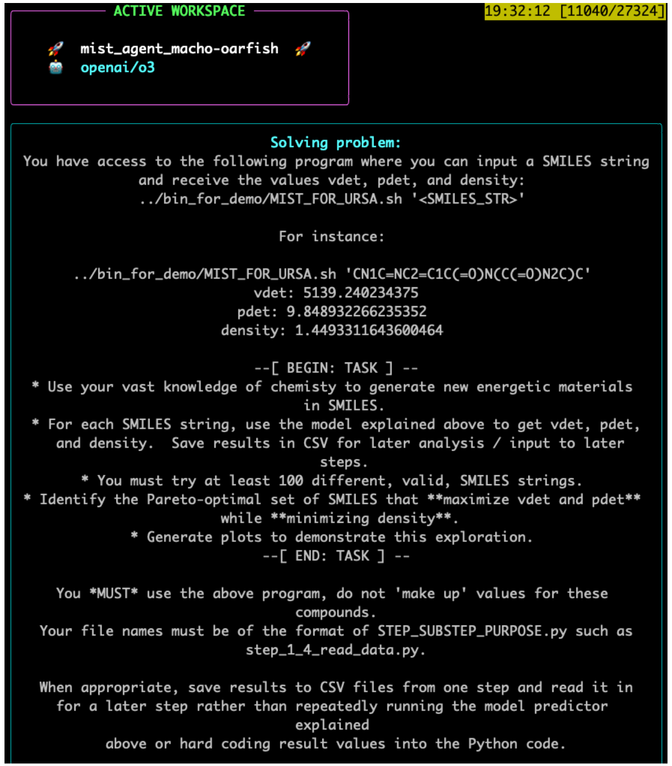

```
ACTIVE WORKSPACE                                          19:32:12 [11040/27324]

   🚀 mist_agent_macho-oarfish 🚀
   🤖 openai/o3

                              Solving problem:
 You have access to the following program where you can input a SMILES string
                 and receive the values vdet, pdet, and density:
               ../bin_for_demo/MIST_FOR_URSA.sh '<SMILES_STR>'

                                For instance:

        ../bin_for_demo/MIST_FOR_URSA.sh 'CN1C=NC2=C1C(=O)N(C(=O)N2C)C'
                             vdet: 5139.240234375
                           pdet: 9.848932266235352
                        density: 1.4493311643600464

                            --[ BEGIN: TASK ] --
  * Use your vast knowledge of chemisty to generate new energetic materials
                                 in SMILES.
  * For each SMILES string, use the model explained above to get vdet, pdet,
     and density.  Save results in CSV for later analysis / input to later
                                   steps.
        * You must try at least 100 different, valid, SMILES strings.
 * Identify the Pareto-optimal set of SMILES that **maximize vdet and pdet**
                        while **minimizing density**.
            * Generate plots to demonstrate this exploration.
                             --[ END: TASK ] --

    You *MUST* use the above program, do not 'make up' values for these
                                 compounds.
  Your file names must be of the format of STEP_SUBSTEP_PURPOSE.py such as
                         step_1_4_read_data.py.

  When appropriate, save results to CSV files from one step and read it in
     for a later step rather than repeatedly running the model predictor
                                  explained
       above or hard coding result values into the Python code.
```

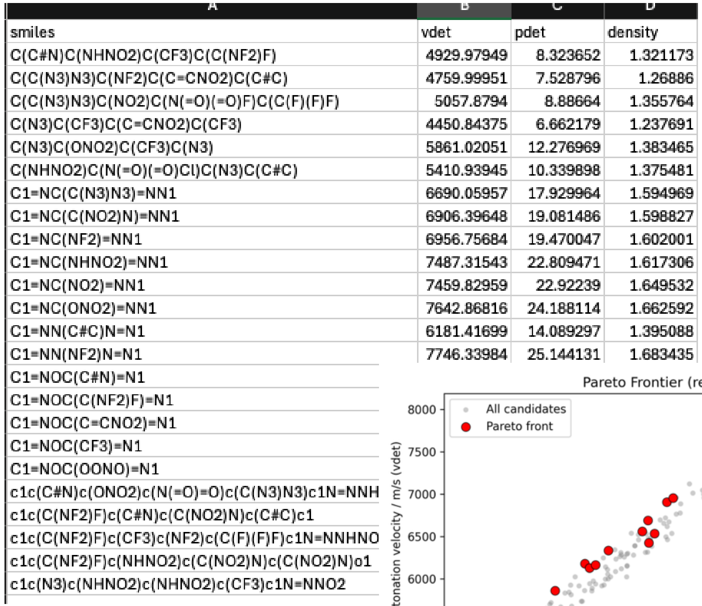

A	B	C	D
smiles	vdet	pdet	density
C(C#N)C(NHNO2)C(CF3)C(C(NF2)F)	4929.97949	8.323652	1.321173
C(C(N3)N3)C(NF2)C(C=CNO2)C(C#C)	4759.99951	7.528796	1.26886
C(C(N3)N3)C(NO2)C(N(=O)(=O)F)C(C(F)(F)F)	5057.8794	8.88664	1.355764
C(N3)C(CF3)C(C=CNO2)C(CF3)	4450.84375	6.662179	1.237691
C(N3)C(ONO2)C(CF3)C(N3)	5861.02051	12.276969	1.383465
C(NHNO2)C(N(=O)(=O)Cl)C(N3)C(C#C)	5410.93945	10.339898	1.375481
C1=NC(C(N3)N3)=NN1	6690.05957	17.929964	1.594969
C1=NC(C(NO2)N)=NN1	6906.39648	19.081486	1.598827
C1=NC(NF2)=NN1	6956.75684	19.470047	1.602001
C1=NC(NHNO2)=NN1	7487.31543	22.809471	1.617306
C1=NC(NO2)=NN1	7459.82959	22.92239	1.649532
C1=NC(ONO2)=NN1	7642.86816	24.188114	1.662592
C1=NN(C#C)N=N1	6181.41699	14.089297	1.395088
C1=NN(NF2)N=N1	7746.33984	25.144131	1.683435
C1=NOC(C#N)=N1			
C1=NOC(C(NF2)F)=N1			
C1=NOC(C=CNO2)=N1			
C1=NOC(CF3)=N1			
C1=NOC(OONO)=N1			
c1c(C#N)c(ONO2)c(N(=O)=O)c(C(N3)N3)c1N=NNH			
c1c(C(NF2)F)c(C#N)c(C(NO2)N)c(C#C)c1			
c1c(C(NF2)F)c(CF3)c(NF2)c(C(F)(F)F)c1N=NNHNO			
c1c(C(NF2)F)c(NHNO2)c(C(NO2)N)c(C(NO2)N)o1			
c1c(N3)c(NHNO2)c(NHNO2)c(CF3)c1N=NNO2			

Pareto Frontier (red)

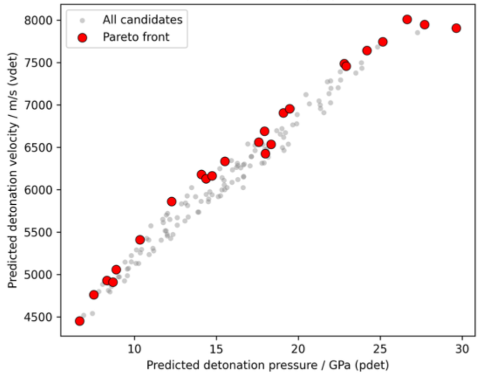

Available now: https://github.com/lanl/ursa

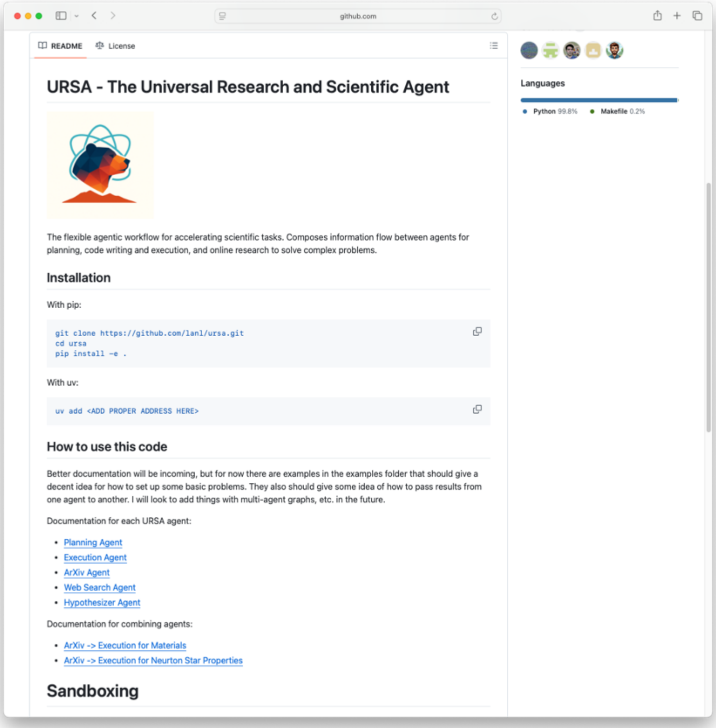

URSA - The Universal Research and Scientific Agent

The flexible agentic workflow for accelerating scientific tasks. Composes information flow between agents for planning, code writing and execution, and online research to solve complex problems.

Installation

With pip:

```
git clone https://github.com/lanl/ursa.git
cd ursa
pip install -e .
```

With uv:

```
uv add <ADD PROPER ADDRESS HERE>
```

How to use this code

Better documentation will be incoming, but for now there are examples in the examples folder that should give a decent idea for how to set up some basic problems. They also should give some idea of how to pass results from one agent to another. I will look to add things with multi-agent graphs, etc. in the future.

Documentation for each URSA agent:

- Planning Agent
- Execution Agent
- ArXiv Agent
- Web Search Agent
- Hypothesizer Agent

Documentation for combining agents:

- ArXiv -> Execution for Materials
- ArXiv -> Execution for Neurton Star Properties

Sandboxing

Los Alamos NATIONAL LABORATORY

University partnerships have expanded our ability to make progress on fundamentals.

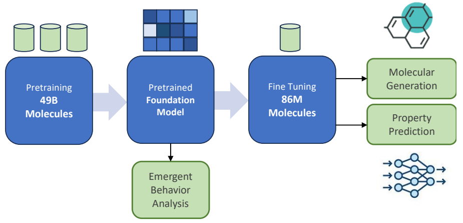

Pretraining 49B Molecules
Pretrained Foundation Model
Fine Tuning 86M Molecules
Molecular Generation
Property Prediction
Emergent Behavior Analysis

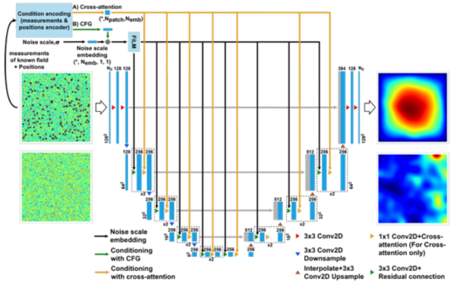

Condition encoding (measurements & positions encoder)
A) Cross-attention
B) CFG
Noise scale, σ
Noise scale embedding
FiLM
measurements of known field + Positions
Noise scale embedding
Conditioning with CFG
Conditioning with cross-attention
3x3 Conv2D
3x3 Conv2D Downsample
Interpolate+3x3 Conv2D Upsample
1x1 Conv2D+Cross-attention (For Cross-attention only)
3x3 Conv2D+ Residual connection

Foundation models for materials and multiphysics.

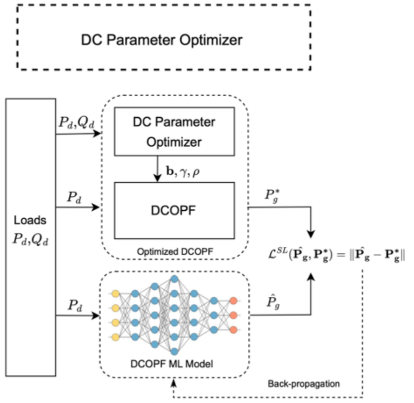

Design, discovery & control with Georgia Tech.

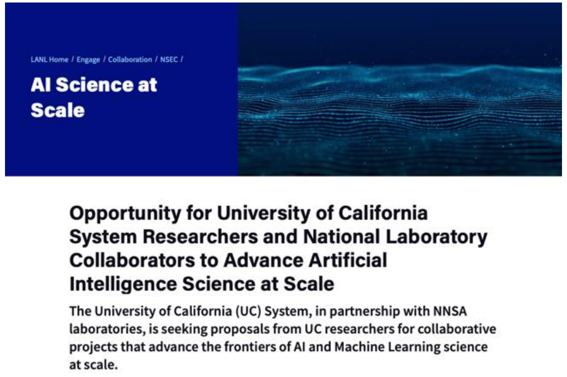

LANL Home / Engage / Collaboration / NSEC /

AI Science at Scale

Opportunity for University of California System Researchers and National Laboratory Collaborators to Advance Artificial Intelligence Science at Scale

The University of California (UC) System, in partnership with NNSA laboratories, is seeking proposals from UC researchers for collaborative projects that advance the frontiers of AI and Machine Learning science at scale.

Applications.

Industry partnerships have provided access to tools and accelerated progress.

Training and support for model building.

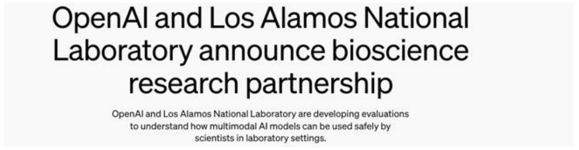

OpenAI and Los Alamos National Laboratory announce bioscience research partnership

OpenAI and Los Alamos National Laboratory are developing evaluations to understand how multimodal AI models can be used safely by scientists in laboratory settings.

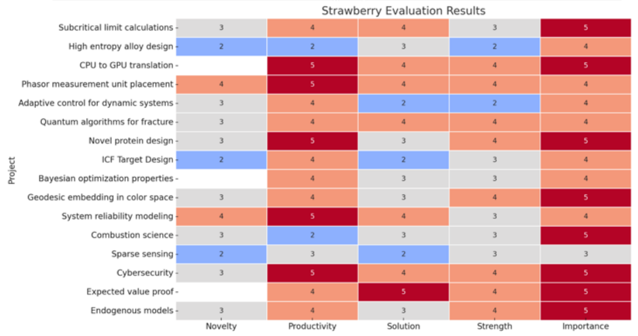

Strawberry Evaluation Results

Project	Novelty	Productivity	Solution	Strength	Importance
Subcritical limit calculations	3	4	4	3	5
High entropy alloy design	2	2	3	2	4
CPU to GPU translation		5	4	4	5
Phasor measurement unit placement	4	5	4	4	4
Adaptive control for dynamic systems	3	4	2	2	4
Quantum algorithms for fracture	3	4	4	4	4
Novel protein design	3	5	3	4	5
ICF Target Design	2	4	2	3	4
Bayesian optimization properties		4	3	3	4
Geodesic embedding in color space	3	4	3	4	5
System reliability modeling	4	5	4	3	4
Combustion science	3	2	3	3	5
Sparse sensing	2	3	2	3	3
Cybersecurity	3	5	4	4	5
Expected value proof		4	5	4	5
Endogenous models	3	4	3	4	5

Early access for test and evaluation.

We have the only on-prem installation of an OpenAI model anywhere in the world.

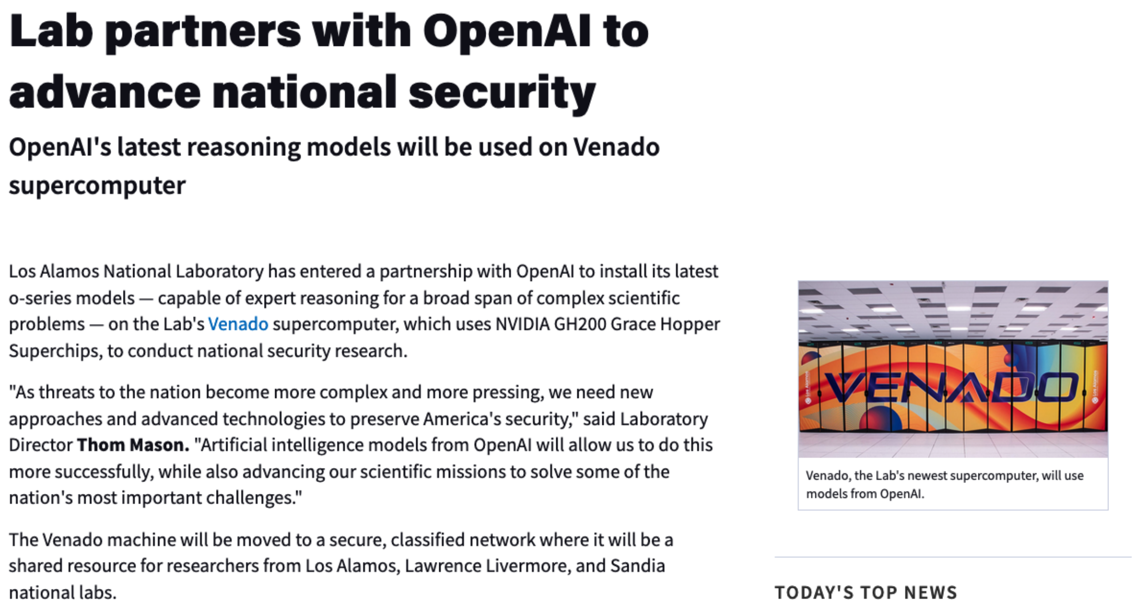

Lab partners with OpenAI to advance national security

OpenAI's latest reasoning models will be used on Venado supercomputer

Los Alamos National Laboratory has entered a partnership with OpenAI to install its latest o-series models — capable of expert reasoning for a broad span of complex scientific problems — on the Lab's Venado supercomputer, which uses NVIDIA GH200 Grace Hopper Superchips, to conduct national security research.

"As threats to the nation become more complex and more pressing, we need new approaches and advanced technologies to preserve America's security," said Laboratory Director **Thom Mason.** "Artificial intelligence models from OpenAI will allow us to do this more successfully, while also advancing our scientific missions to solve some of the nation's most important challenges."

The Venado machine will be moved to a secure, classified network where it will be a shared resource for researchers from Los Alamos, Lawrence Livermore, and Sandia national labs.

Venado, the Lab's newest supercomputer, will use models from OpenAI.

TODAY'S TOP NEWS

Utility vs. Spectacle

We don't need **THIS**,
for **THIS** to be useful

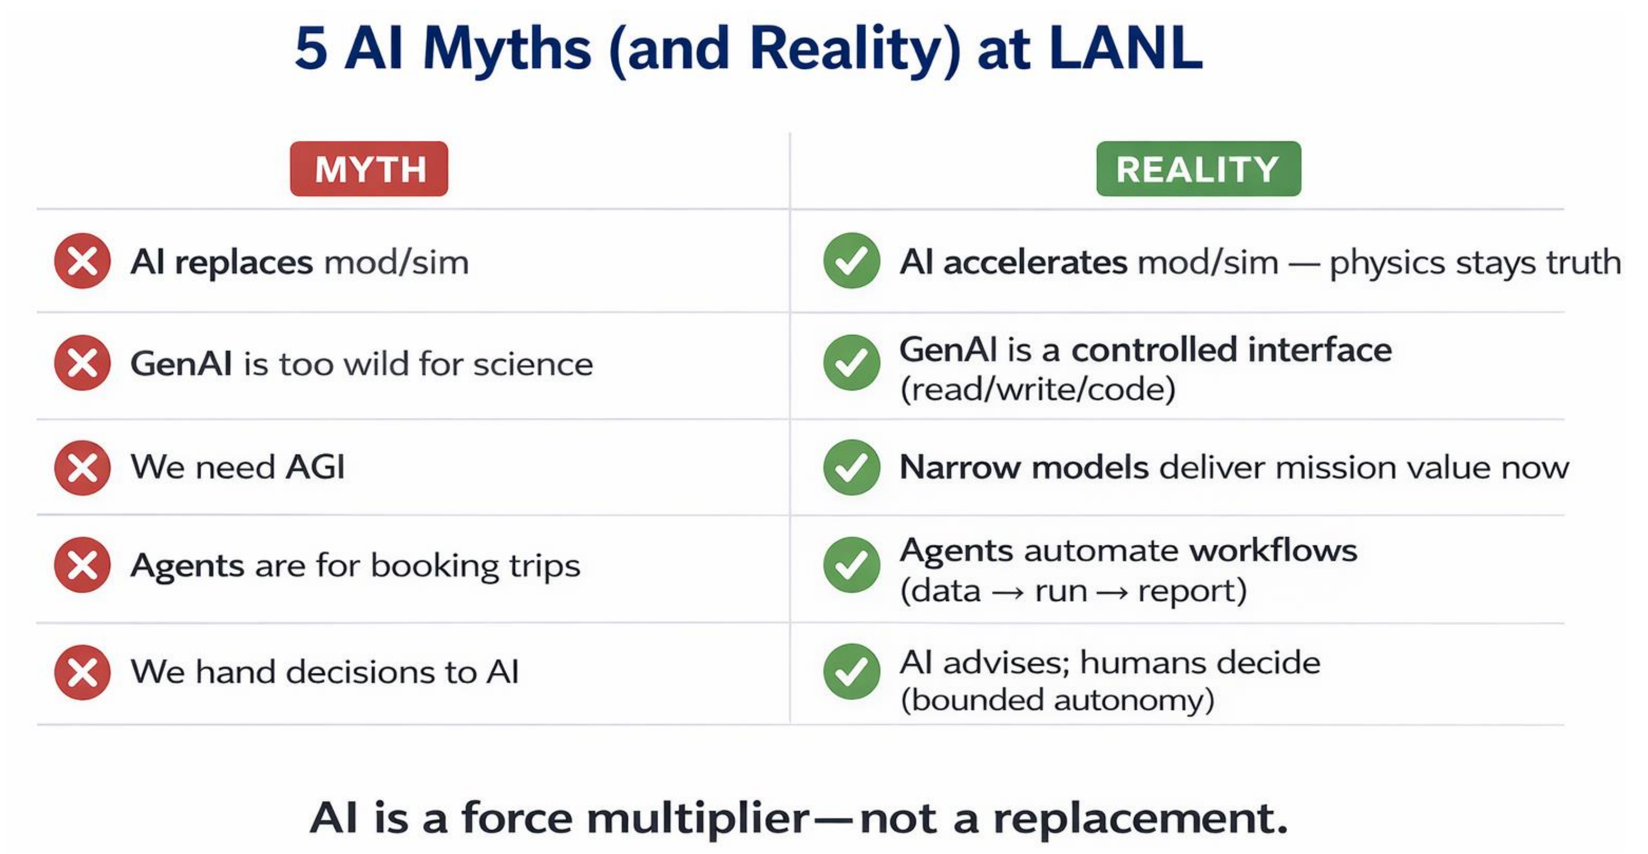

5 AI Myths (and Reality) at LANL

MYTH	REALITY
✗ **AI replaces** mod/sim	✓ **AI accelerates** mod/sim — physics stays truth
✗ **GenAI** is too wild for science	✓ **GenAI** is a **controlled interface** (read/write/code)
✗ We need **AGI**	✓ **Narrow models** deliver mission value now
✗ **Agents** are for booking trips	✓ **Agents** automate **workflows** (data → run → report)
✗ We hand decisions to AI	✓ AI advises; humans decide (bounded autonomy)

AI is a force multiplier—not a replacement.

LANL’s ArtIMis in Summary . . .

- Leverage frontier industry reasoning models
- AI ecosystem of reasoning agents powering non-language science foundation models
- Partnerships with industry and academia
- Dozens of papers – CS as well as in applied domains
- Deep application to LANL mission
- Feeding directly into Genesis Mission project

URSA

https://github.com/lanl/ursa

Thanks for the Invite!

Feel free to reach out

Nathan DeBardeleben, Ph.D.

High Performance Computing Design (HPC-DES)

ndebard@lanl.gov