TATA CONSUMER PRODUCTS

AAAI 2026

WMAC - AAAI 2026 Bridge Program on Advancing LLM-Based Multi-Agent Collaboration

Hierarchical Multi-Agent System for Data-Efficient Alloy Discovery

With Closed-Loop Experimental Feedback

Hierarchical Architecture | Closed-Loop Feedback | 21 Novel Alloys

Mahule Roy : University of Oxford
Subhas Roy : TATA Consumer Products Limited

Executive Summary

The Challenge: Traditional AI-driven materials discovery employs monolithic architectures that create a brittle interface with physical experimentation, leading to high costs and slow iteration.

Our Contribution: ODL-DSP v4.0

We present a **Hierarchical Multi-Agent System (MAS)** that redesigns this architecture through three innovations:

1. Furnace-to-agent feedback loops (Continuous Online Learning).
2. Curiosity-annealing scheduler (Adaptive Exploration).
3 Memory-injected composition generators.

.

7x	21	97%
Reduction in Lab Iterations	Novel Pareto Alloys	Metallurgical Feasibility

Context: The Materials Discovery Bottleneck

Exponential Search Spaces

Materials discovery is constrained by the combinatorial explosion of possible elements. For Multi-Principal Element Alloys (MPEAs), the design space is vast and non-linear.

Costly Experimentation

Physical synthesis ("The Furnace") is the bottleneck. Traditional iterative cycles are slow and expensive ($50k+ per campaign).

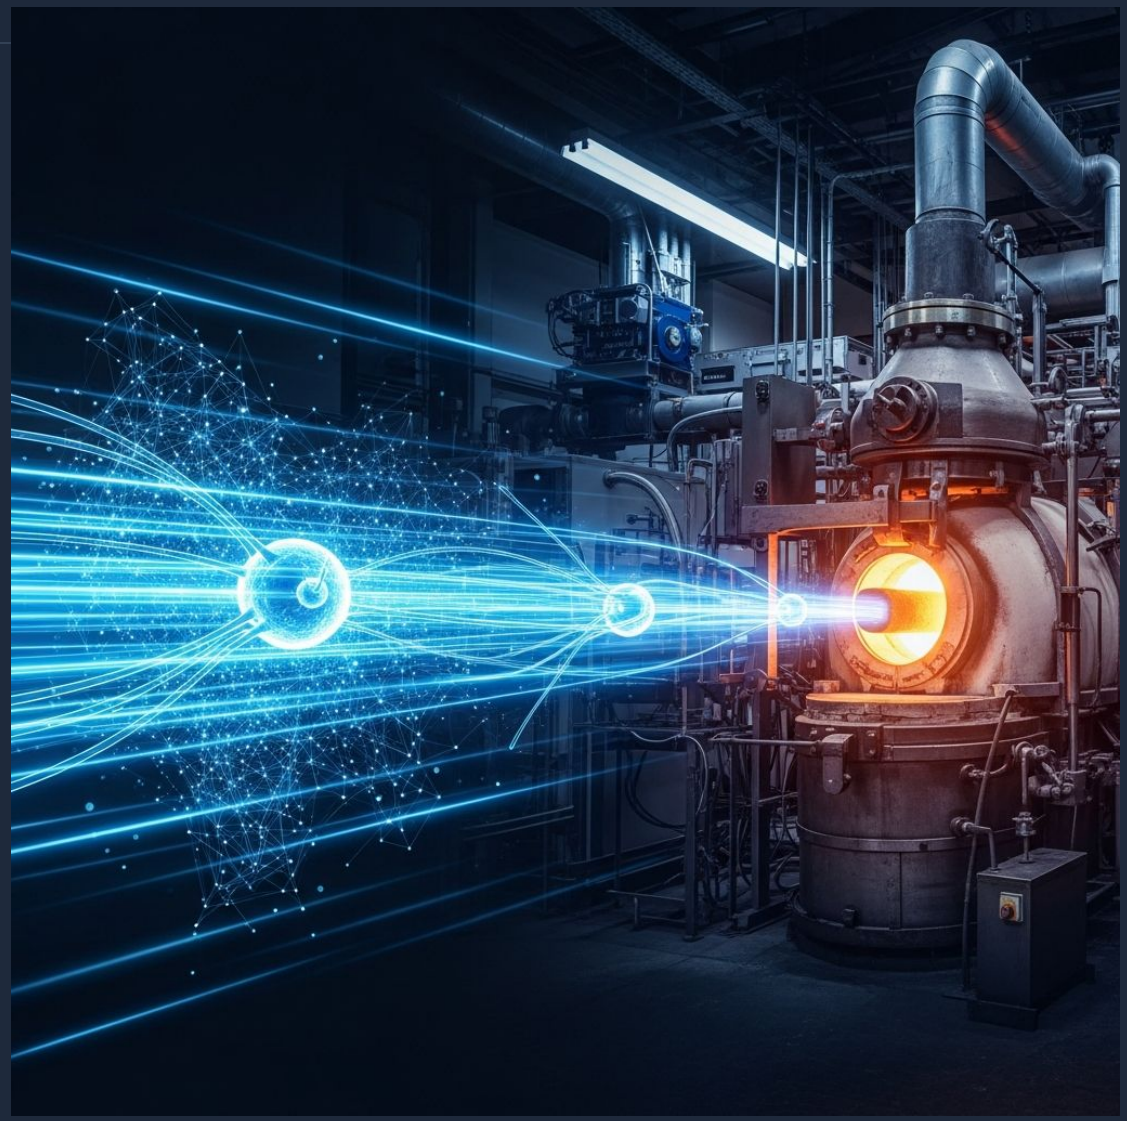

The Gap: Monolithic AI Architectures

Monolithic

Single surrogate models are trained once and deployed statically.

Problem: They do not learn from ongoing experimental failures.

Flattened

Multi-objective problems (Strength vs Ductility) are often scalarized into a single number.

Problem: Loss of nuance in trade-offs.

Brittle

The interface between the AI and the Physical Lab is rigid.

Problem: Requires expensive retraining campaigns.

Solution: Hierarchical MAS

We introduce a system where specialized agents continuously adapt through experimental feedback.

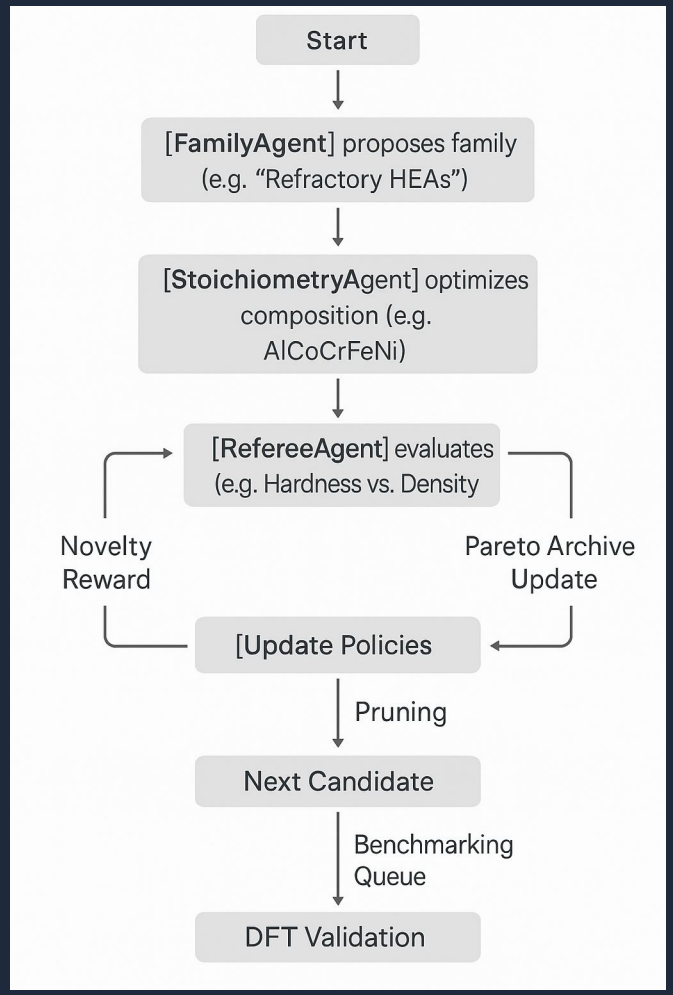

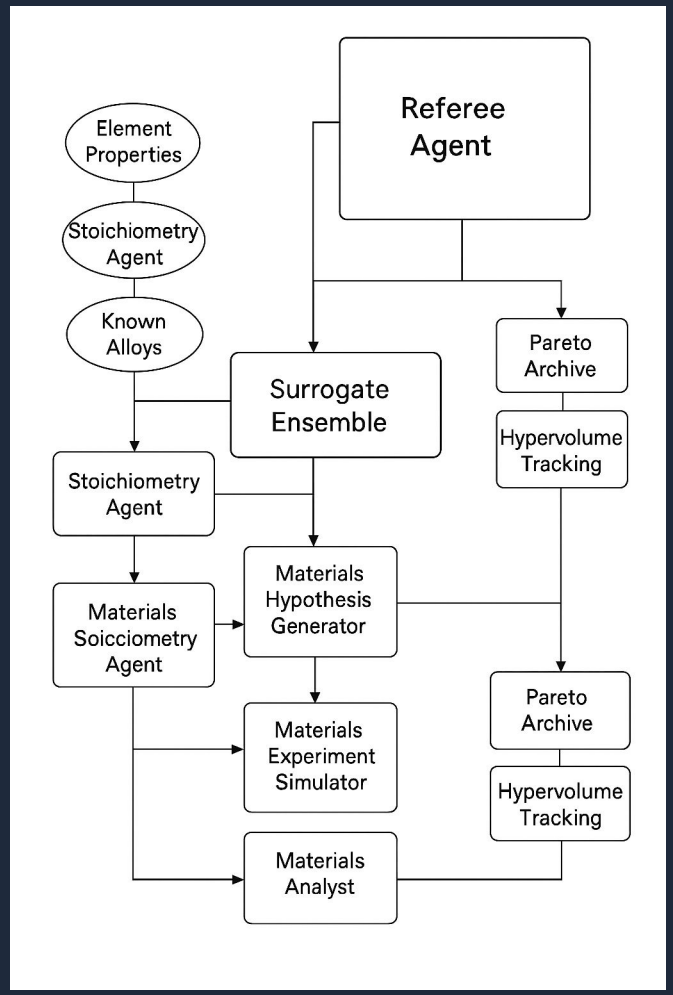

"Furnace-aware learning enables seven-fold reductions in experimental costs."

Three Core Architectural Innovations

1. Continuous Online Learning

Updates agents via furnace feedback after **each** experiment. The AI model is not static; it evolves with the campaign.

2. Adaptive Exploration

Uses **Bayesian Optimization** to balance exploration (curiosity) and exploitation (known high-yields) dynamically as the campaign progresses.

3. Memory of Success

Memory-injected composition generators bias new proposals toward historically successful regions while avoiding repeated failures.

MAS Framework: The Orchestrator

The Orchestrator (**Orchestrator (O)**) coordinates the closed discovery loop consisting of Hypothesis Generators (H), Simulators (E}, and Analysts (A).

$$O : h_t \xrightarrow{E} \hat{y}_t \xrightarrow{A} s_t \quad \text{with} \quad h_{t+1} \sim \pi(h \mid s_{1:t}) \tag{Equation 1}$$

Where pi denotes the adaptive proposal policy updated via historical scores. This role specialization preserves multi-objective trade-offs.

Strategic Layer: FamilyAgent

Responsibility

Selects broad metallurgical families to explore (e.g., Refractory, HEA, Ni-superalloy).

Mechanism

Uses **Curiosity-Weighted Sampling**. It does not look at specific percentages but rather the combination of elements.

Adapts its exploration coefficient β online via Bayesian Optimization.

Tactical Layer: StoichiometryAgent

Responsibility

Generates specific compositions within the families selected by the Strategic layer.

Mechanism

Guided by **Success Memory**. It fine-tunes element fractions.

$$\pi(x \mid H_t) \propto \exp(\kappa R_t(x) + \gamma \mathrm{Nov}(x) + \varepsilon_t)$$

Generative policy balancing reward (R) and novelty (Nov).

Evaluative Layer: RefereeAgent

Responsibility

Performs multi-objective evaluation and maintains the Pareto Archive.

Mechanism

Computes composite rewards and decides which candidates proceed to physical synthesis (The Furnace).

Updates the archive $\mathcal{A}$ after every cycle with empirical ground truth.

Feature Engineering

Alloys are represented by a 48-dimensional feature vector balancing metallurgical completeness with computational efficiency.

Elemental Properties

- Electronegativity (Pauling)
- Atomic Radius (pm)
- Valence Electron Conc. (VEC)
- Melting Temp (K)
- Density (g/cm^3)

Thermodynamic

- Mixing Enthalpy (ΔH_{mix})
- Mixing Entropy (ΔS_{mix})
- Atomic Size Mismatch (δ)
- Electronegativity Diff ($\Delta \chi$)

Compositional

- Element Fractions
- Hume-Rothery Parameters
- Phase Stability (PHACOMP)
- Solid Solution Strengthening

Surrogate Model: Gradient Boosting

The property predictor $\mathcal{E}$ uses XGBoost ensembles.

- **Targets:** Vickers Hardness, Corrosion Rate, Electrical Conductivity.
- **Initial Training:** 500 characterized alloys.
- **Dynamics:** Significant online learning improvement through furnace feedback.

Appendix A3: Model Config

Parameter	Value
Model Type	XGBoost
Estimators	1000
Max Depth	6
Learning Rate	0.01
Objective	reg:squarederror

Innovation 1: Furnace-Aware Learning

After each furnace run, agents update using empirical performance discrepancies. This grounds the computation in physical reality.

$$\Delta\theta_{\text{agent}} = -\eta \nabla_{\theta} E_{x \sim \pi_{\theta}} [(r_{\text{empirical}}(x) - r_{\text{predicted}}(x))^2]$$

(Equation 6: Feedback Loop)

Hardness prediction RMSE decreased from **28.5 HV to 18.2 HV** over 50 discovery cycles.

Innovation 2: Success Memory

Agents maintain rolling success memories updated via exponential smoothing. This allows the system to "remember" what worked well in previous cycles.

$$R_{t+1}(x) = (1-\alpha) R_t(x) + \alpha s_t(x)$$

(Equation 2: Memory Update)

Where $\alpha = 0.05$. This prevents the catastrophic forgetting often seen in standard RL approaches.

Innovation 3: Adaptive Policy

The FamilyAgent adapts its exploration coefficient online via Bayesian Optimization over campaign performance.

$$\beta_t = \mathrm{BO}_{\mathrm{EI}}\left(f(\text{performance}, \text{novelty})\right)$$

This balances the need to find new alloy families (Exploration) with the need to refine existing high-performers (Exploitation).

Optimization Logic

We optimize strength, toughness, and corrosion resistance simultaneously using **Pareto dominance**.

$$x \prec y \leftrightarrow \forall j , f_j (x) \geq f_j (y) \wedge \&exists; j , f_j (x) > f_j (y)$$

The composite reward combines normalized objectives with novelty:

$$r (x) = \lambda_S \frac{S (x)}{S_{max}} + \lambda_T \frac{T (x)}{T_{max}} - \lambda_C \frac{C (x)}{C_{max}} + \beta Nov (x)$$

Experimental Setup

Conditions

- **Initial Data:** 500 alloys.
- **Budget:** 50 furnace iterations (cycles).
- **Comparison:** Matched computational budgets across all methods.

Evaluation Metrics

- **Pareto-optimal Yield (#P):** Count of validated top-tier alloys.
- **Feasibility (F%):** Percentage of proposals metallurgically viable.
- **Convergence Speed:** Iterations to reach target performance.

Hyperparameters (Appendix A3)

Parameter	Value	Description
Cycles (M)	50	Furnace melts per campaign
Proposals (N_{prop})	100	In-silico evaluations per cycle
Learn rate (η)	0.05	Agent parameter updates
Memory (α)	0.05	Success memory weight
Explore (β_0)	0.8	Initial novelty weight
BO window	50	Past cycles for β update

The system demonstrates strong robustness, with <2% performance degradation across ±20% parameter changes.

Model Config (Appendix A3)

Configuration for the XGBoost Surrogate Model.

Parameter	Value	Description
Model Type	XGBoost	Gradient boosted trees
Estimators	1000	Boosting rounds
Max Depth	6	Tree depth limit
Learning Rate	0.01	Boosting rate
Objective	reg:squarederror	Continuous prediction
Feature Dim	48	Metallurgical descriptors

Main Results: Performance Comparison

Our hierarchical MAS outperforms all benchmarks.

Method	RMSE	#P (Pareto)	F% (Feasible)	Hits/100
ODL-DSP v4.0	**0.043**	**21**	**97.3**	**34**
MatGPT	0.055	15	72	12
AtomAgent	0.049	18	68	9
AlloyDB RF	0.061	11	55	7
Random	0.089	0	43	0

Result: 34 "Hits" per 100 proposals vs 9 for the best static MAS baseline.

Efficiency Across Budgets

Performance at different iteration checkpoints (Table 3).

Method	10 it.	25 it.	50 it.	Efficiency
Our MAS	**8.2**	**15.7**	**21.3**	**0.89**
MatGPT	5.1	9.3	15.2	0.54
AtomAgent	4.8	8.7	14.9	0.51
Random	2.3	4.1	6.8	0.18

Our method is nearly twice as efficient (0.89 vs 0.51) as the SOTA AtomAgent.

Agent Contribution & Ablation

Agent Contributions

Agent	Gain
Family	2.1x
Stoich.	2.8x
Referee	1.9x
Full MAS	**7.2x**

Extended Ablation (App A4)

- **No Feedback:** #P drops to 12.6.
- **No Memory:** Feasibility drops to 70.2%.
- **Flat MAS:** #P drops to 16.1.

Removing furnace feedback caused a 40% drop in Pareto-optimal yield.

Surrogate Model Evolution (Appendix A10)

RMSE Evolution across the 50-cycle campaign.

Cycle	Hardness (HV)	Corrosion	Conductivity
Initial	28.5	0.0087	1.45
Cycle 10	24.2	0.0072	1.28
Cycle 30	19.5	0.0059	1.07
Final	**18.2**	**0.0051**	**0.94**

Corrosion rate prediction showed the most significant relative improvement (41.4%).

Experimental Validation (Part 1)

Validation of Benchmarks vs Novel Discoveries (Hardness & Corrosion).

Composition	HV (Pred/Exp)	Status
Ti-6Al-4V (Benchmark)	255 / 255	B
Fe 10.4, Co 73.8, Mo 15.8	318 / 305	P (Pareto)
Ti 51.0, Cu 16.5, Ni 32.5	245 / 262	P (Pareto)
Co 40.0, Cr 30.0, Ni 20.0, W 10.0	366 / 351	P (Pareto)
Ti 55.0, Al 25.0, V 15.0, Sn 5.0	275 / 264	P (Pareto)

Experimental Validation (Part 2)

Further validation of novel compositions.

Composition	HV (Pred/Exp)	Status
Co 35.0, Cr 25.0, Ni 20.0, Mo 20.0	346 / 332	P (Pareto)
Ti 60.0, V 20.0, Al 15.0, Cr 5.0	265 / 254	P (Pareto)
Ni 55.0, Cr 20.0, Co 15.0, Mo 10.0	326 / 313	P (Pareto)
Co 25.0, Cr 30.0, Ni 25.0, W 20.0	365 / 352	P (Pareto)

Status P = Pareto-optimal. Status F = Feasible but not Pareto.

Visualizing Discovery

Pareto Front Expansion

Our system systematically pushes the boundary of the known Pareto front.

- **Upper Right Quadrant:** Discovery of alloys that balance high performance with high novelty.
- **Diversity:** Parallel investigation of multiple alloy families (Refractory, HEA) simultaneously.

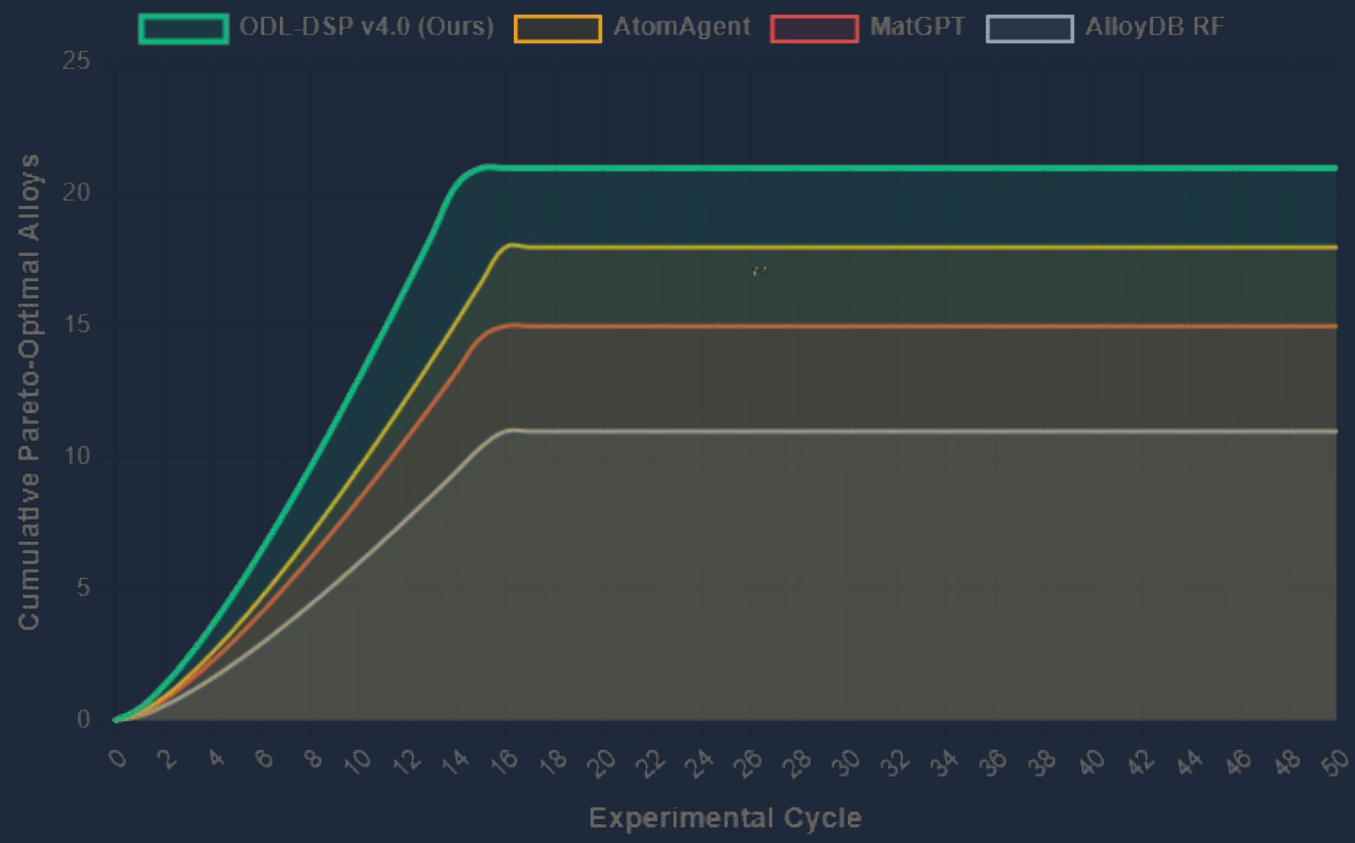

Cross-Domain Generalizability

The architecture is domain-agnostic. We tested it across different material classes.

Table 5: Cross-Domain Transfer Performance

Mat. Class	Our MAS	MatGPT	Fine-Tuned
Refractory	21.3 ± 0.6	15.2	17.8
HEAs	19.8 ± 0.8	13.7	16.2
Ni-super	22.1 ± 0.5	16.3	18.9
Intermet.	18.7 ± 0.9	12.4	15.1

The hierarchical decomposition applies to structured search spaces beyond just simple alloys.

Extended Visualizations

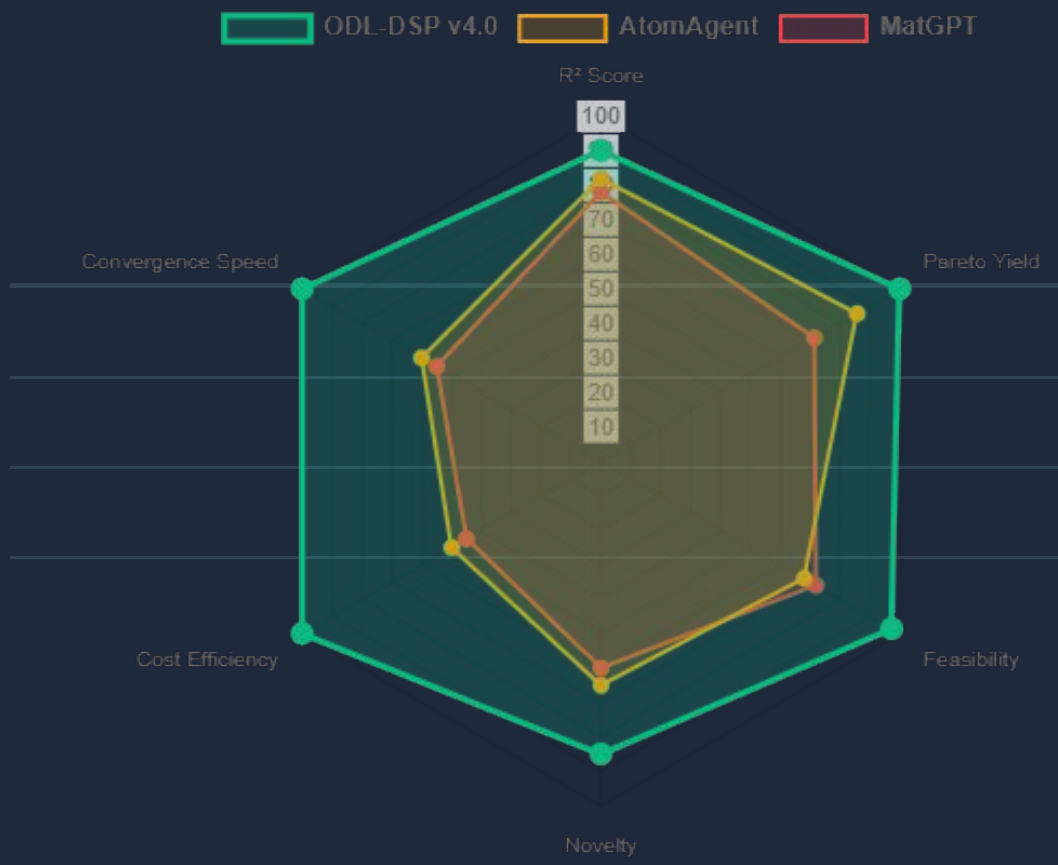

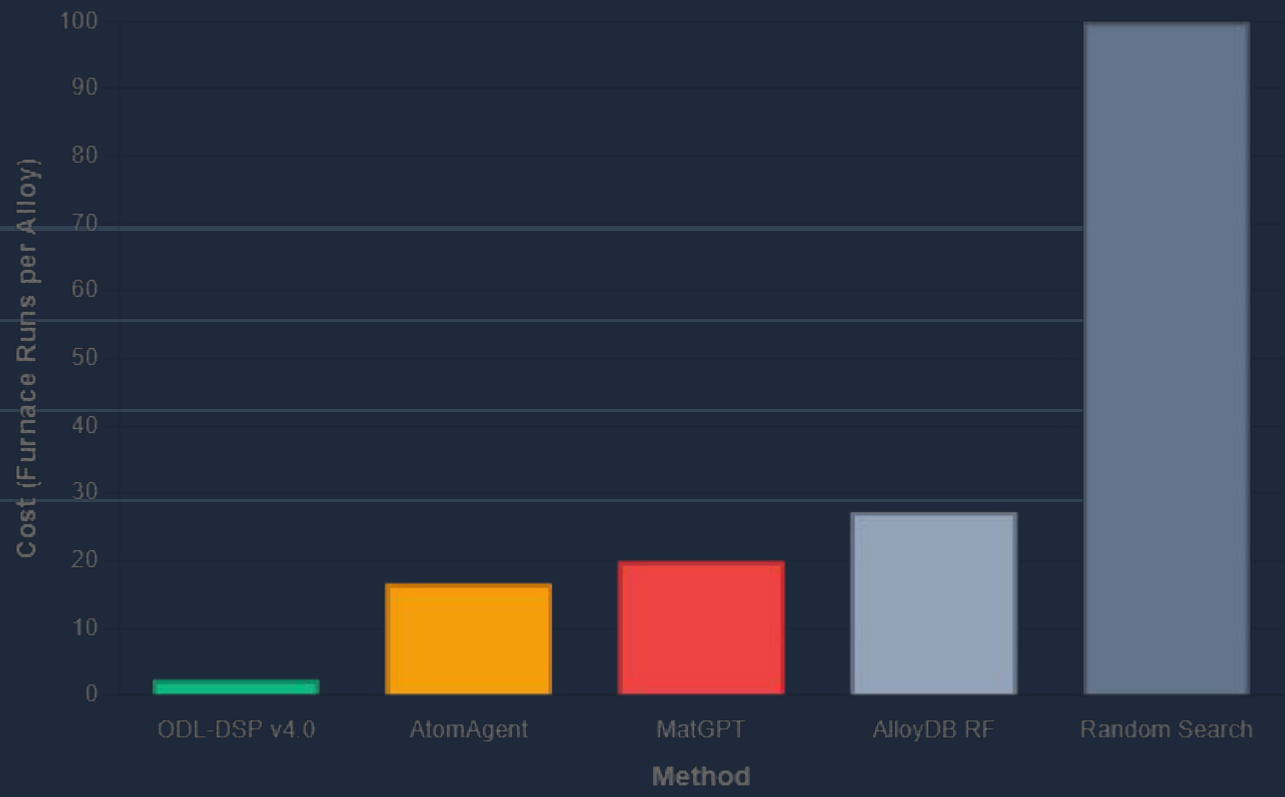

Extended Visualizations

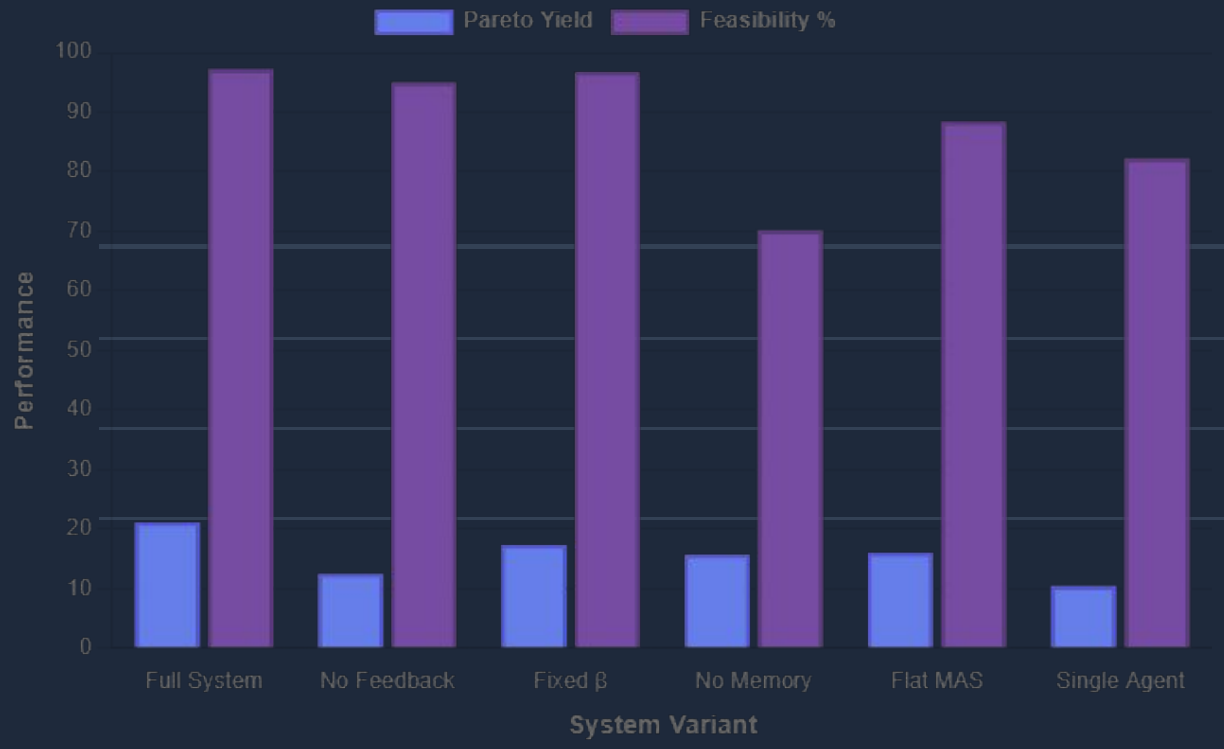

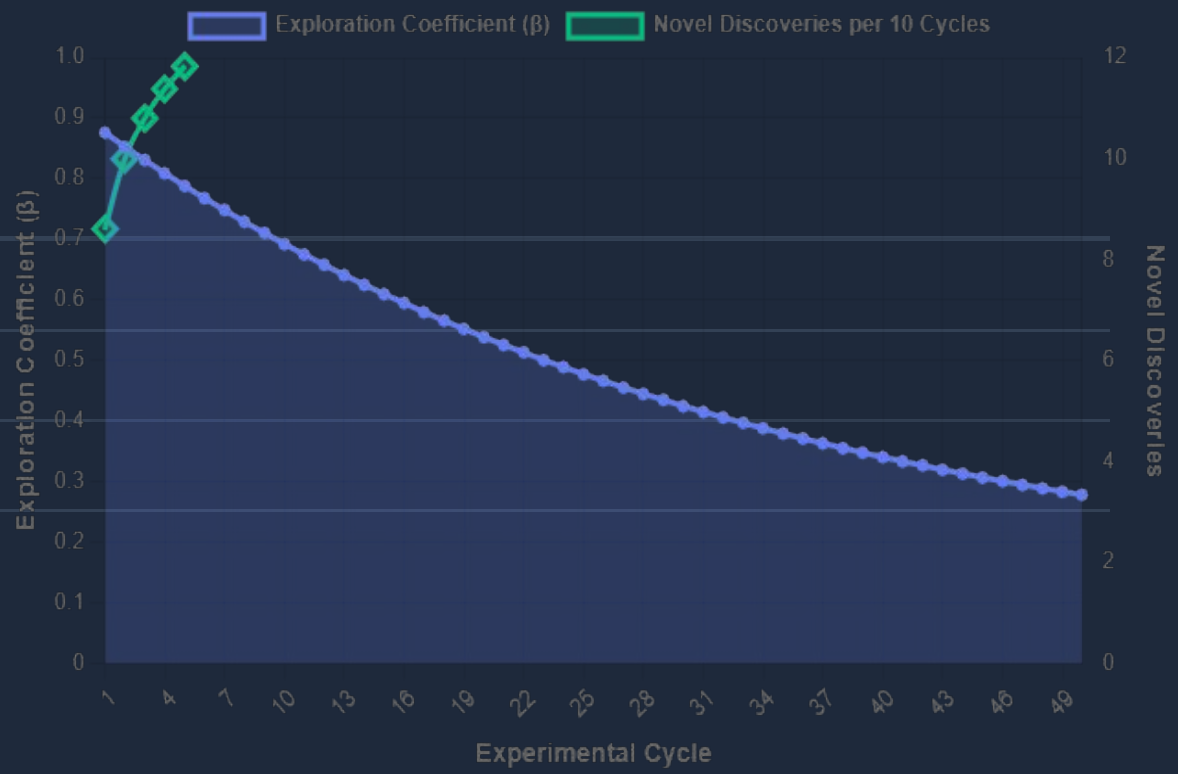

Extended Visualizations

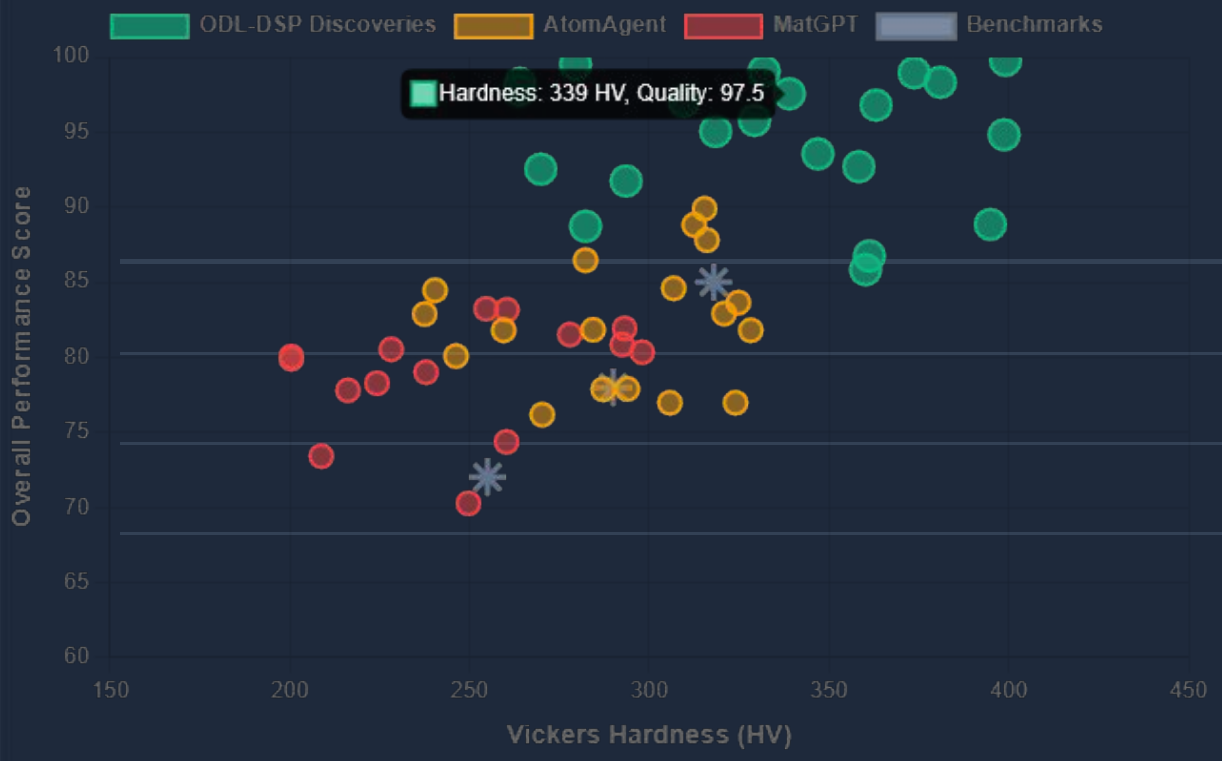

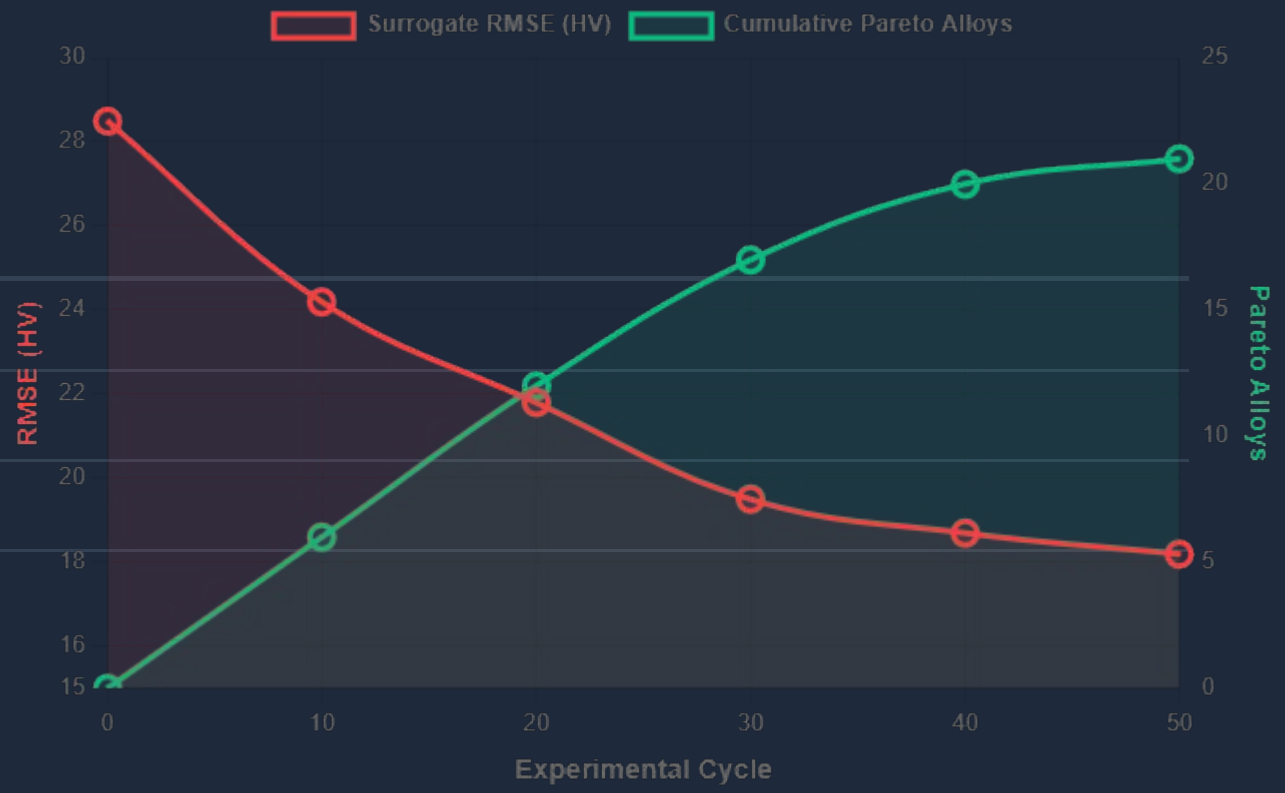

Extended Visualizations

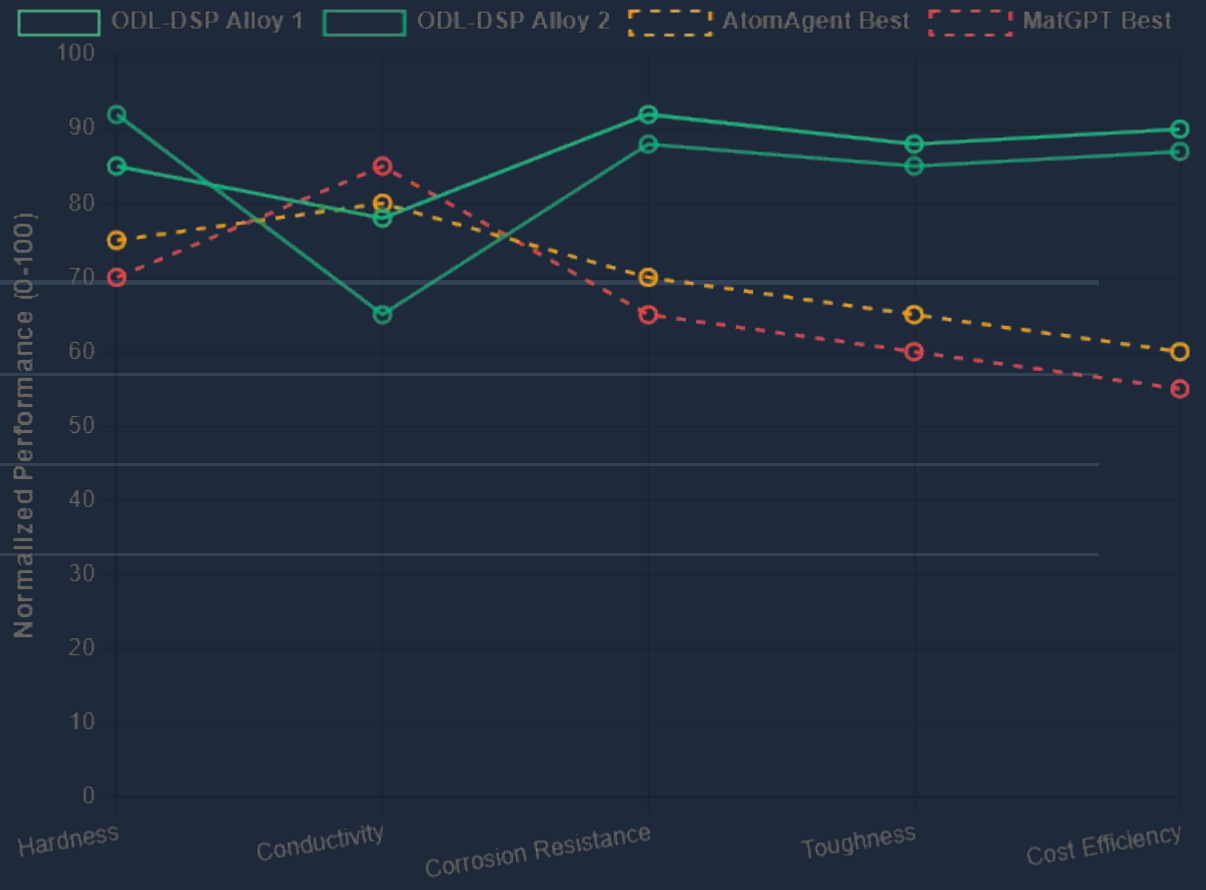

Extended Visualizations

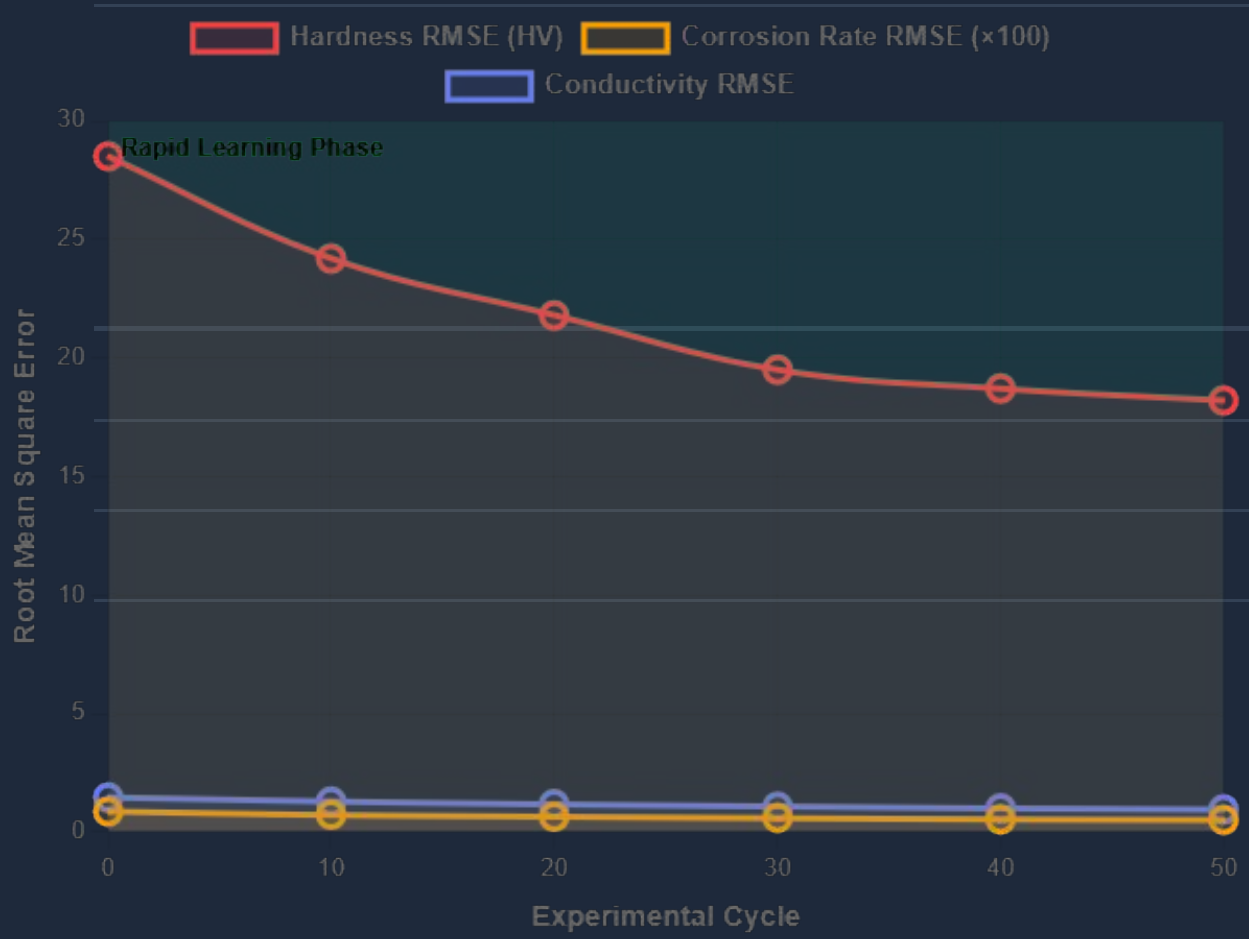

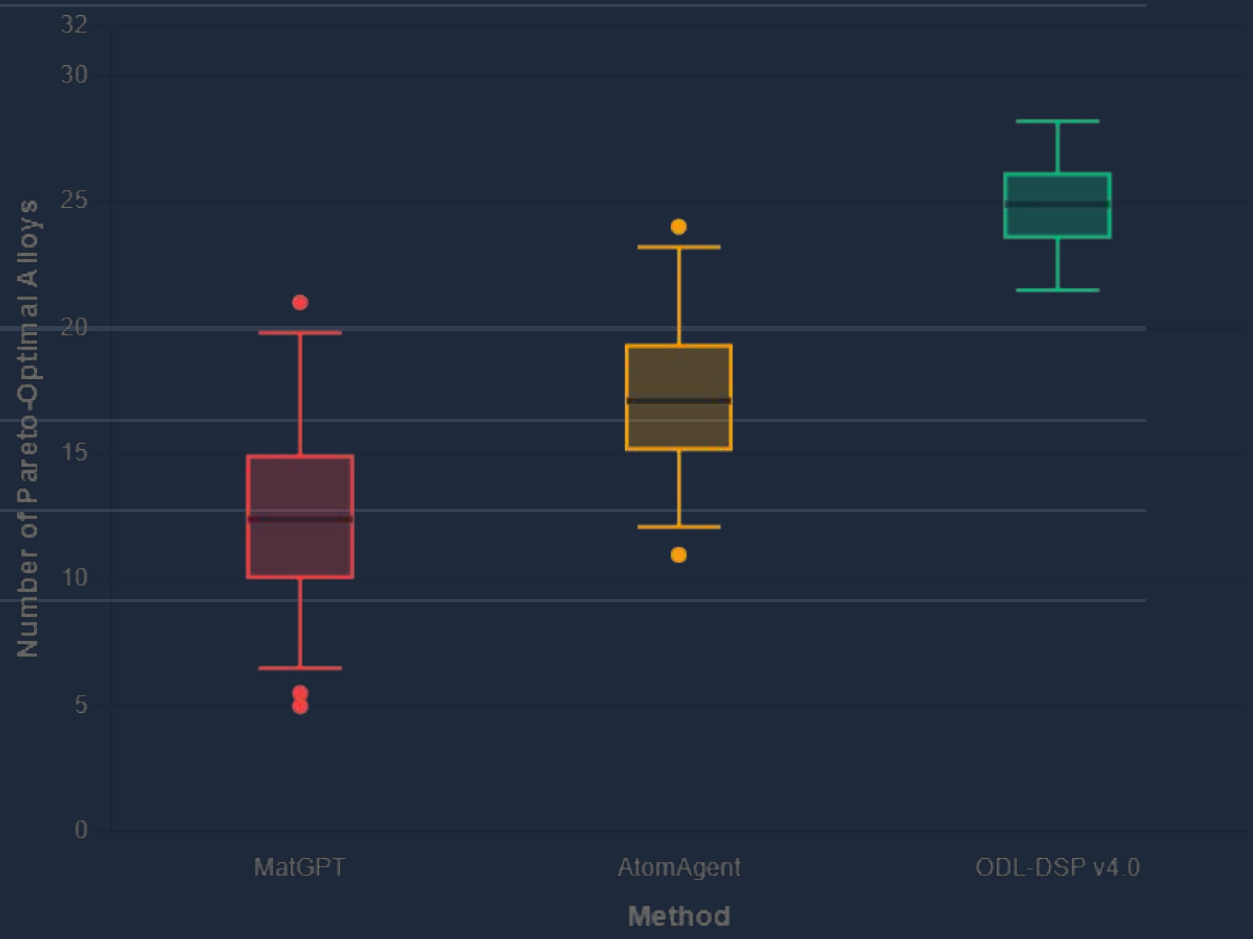

Extended Visualizations

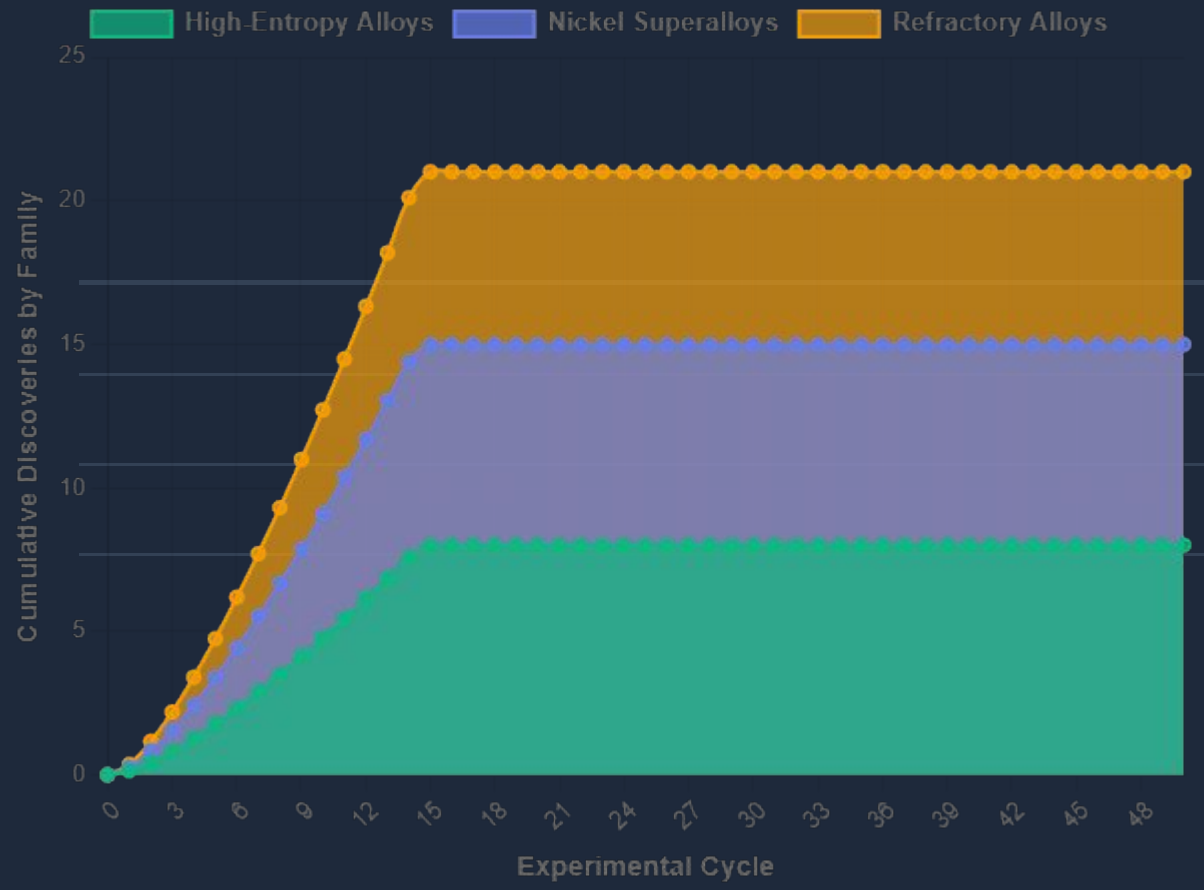

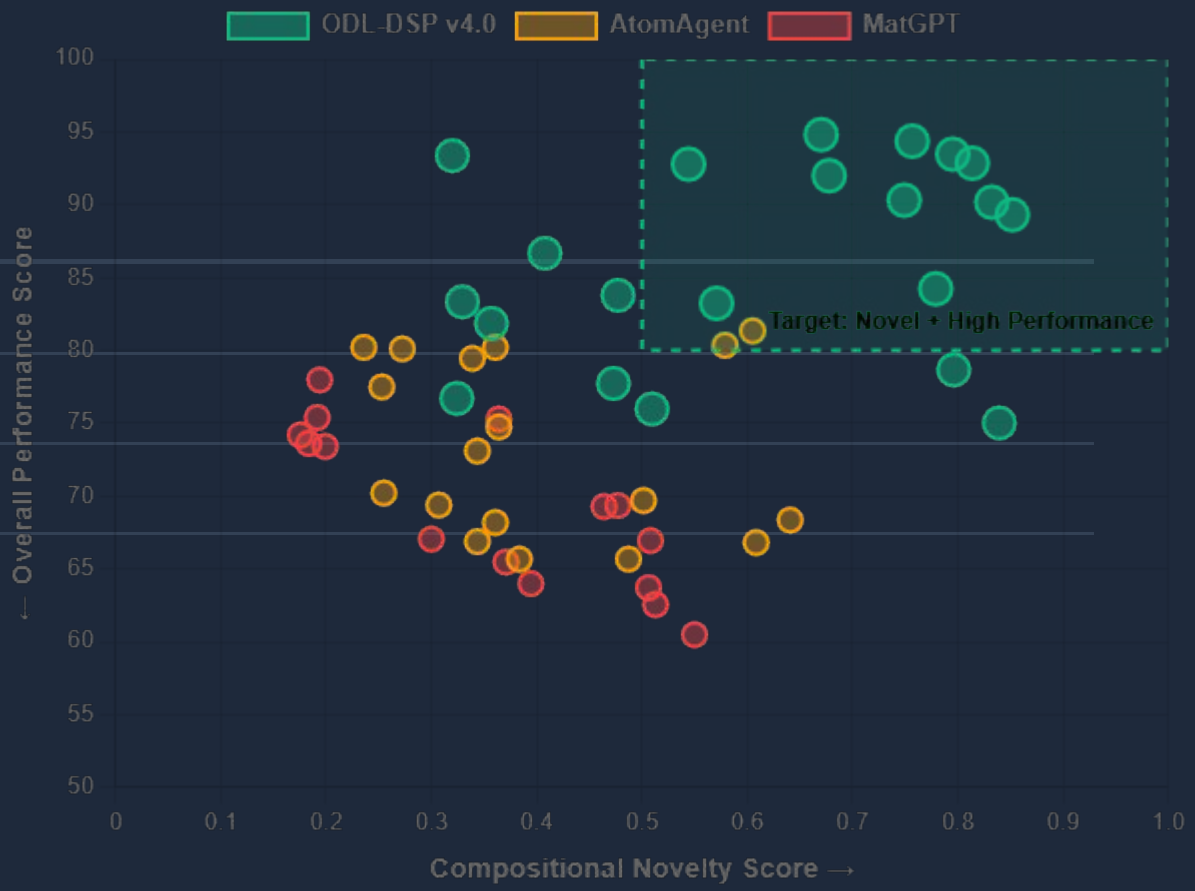

Extended Visualizations

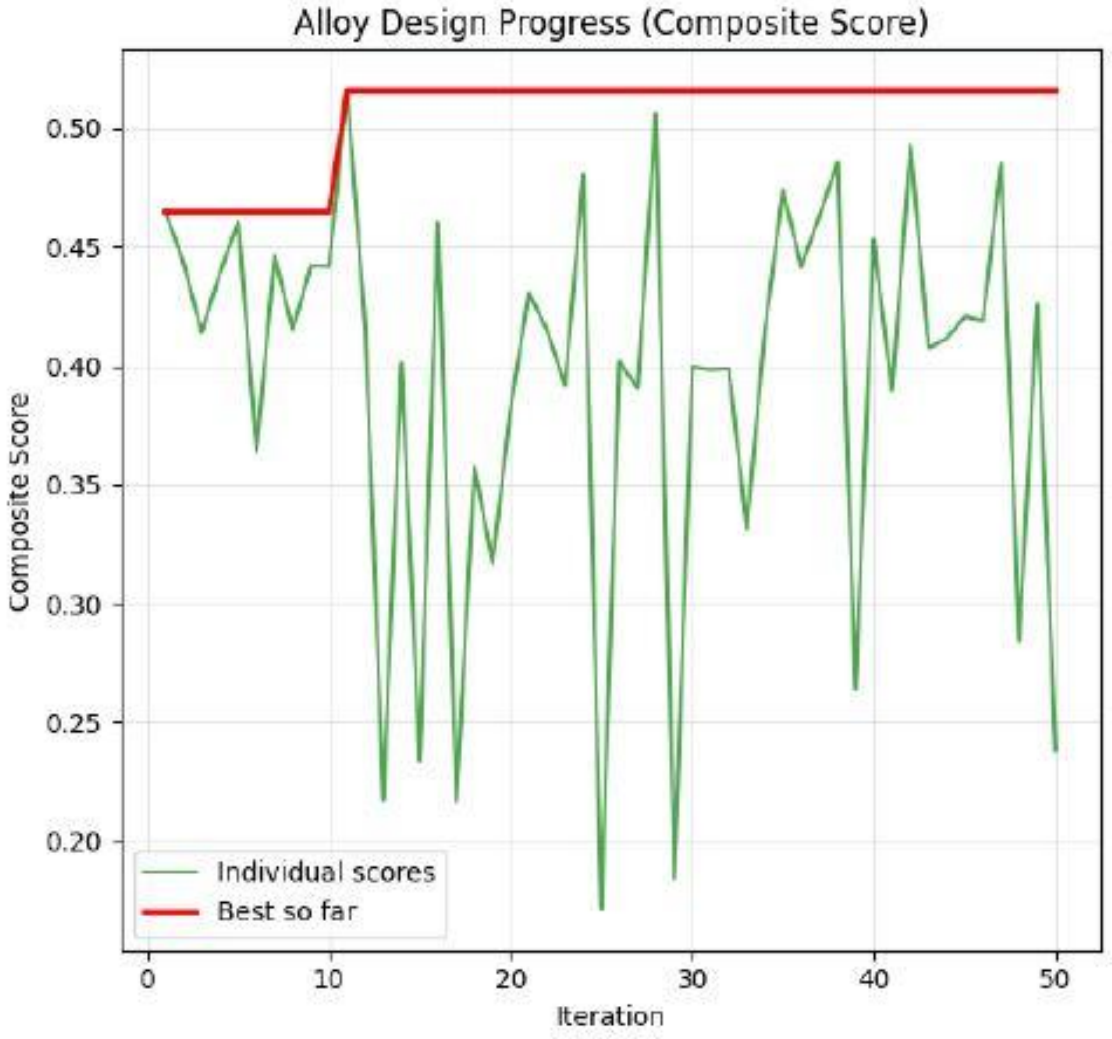

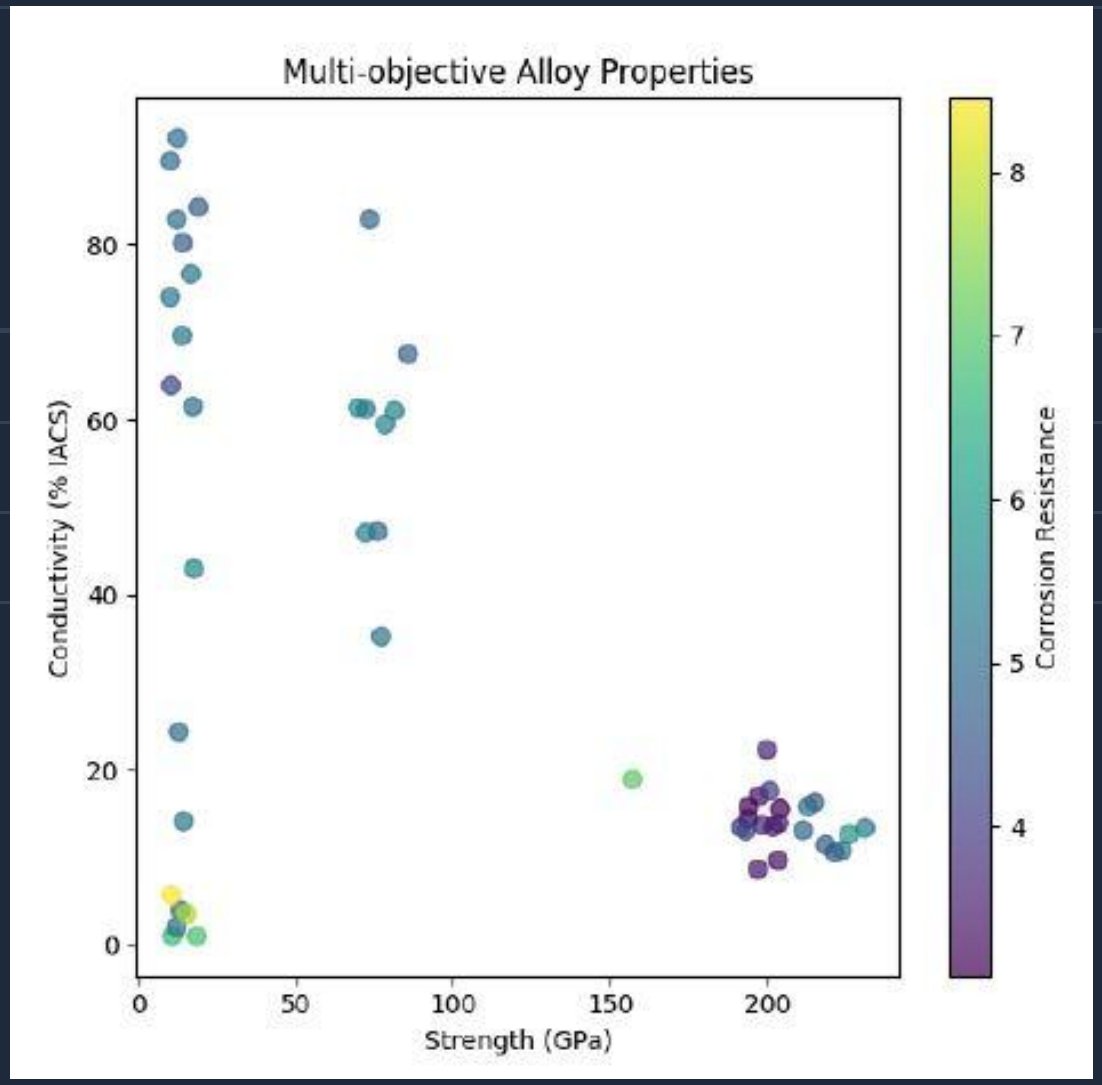

Extended Visualizations

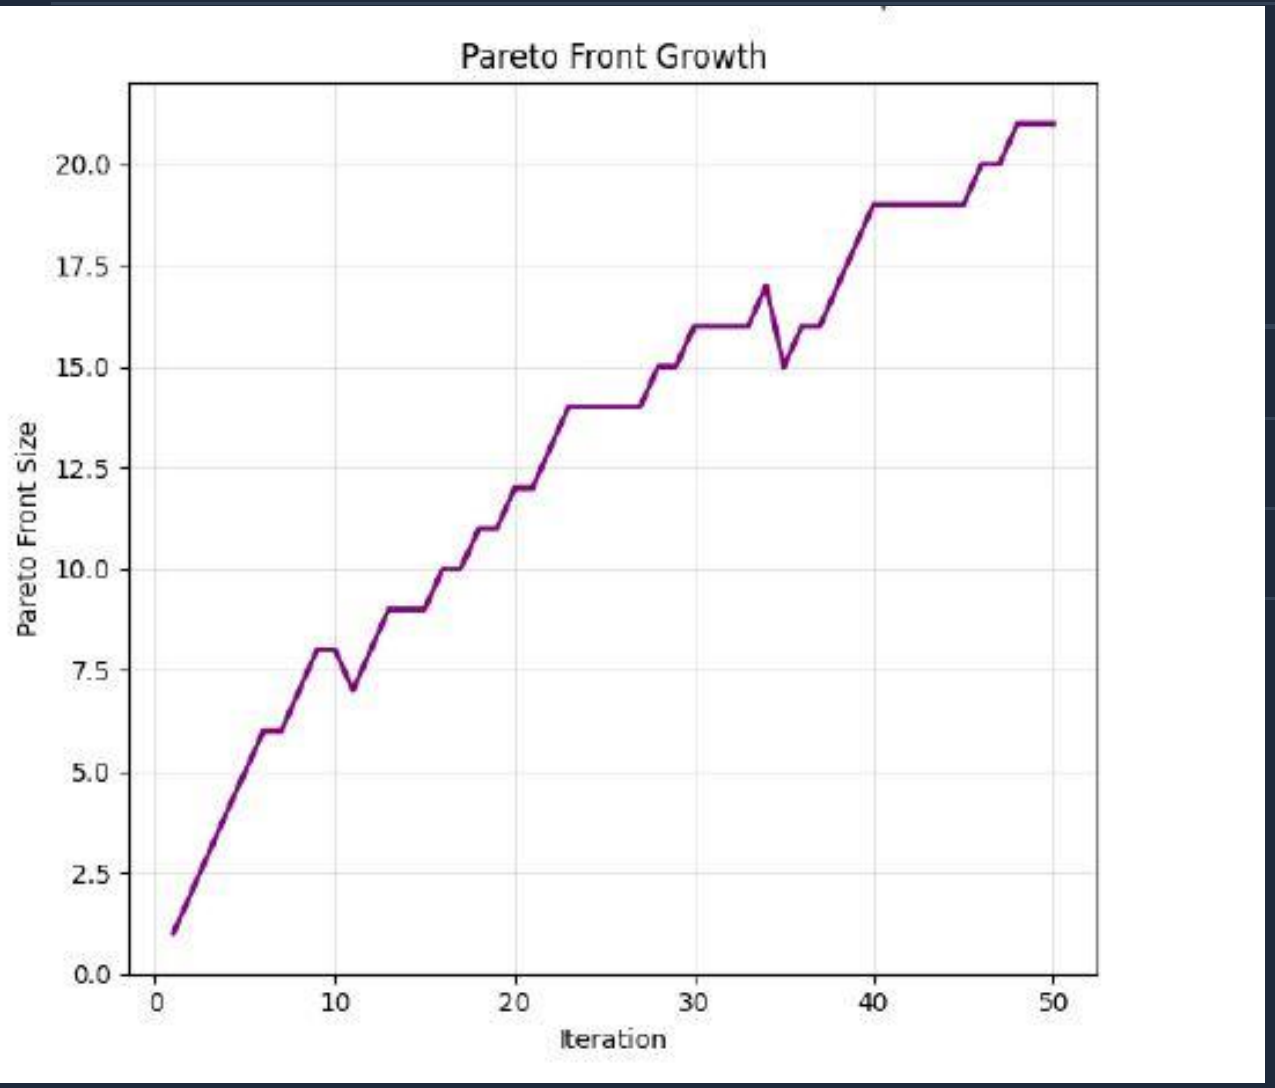

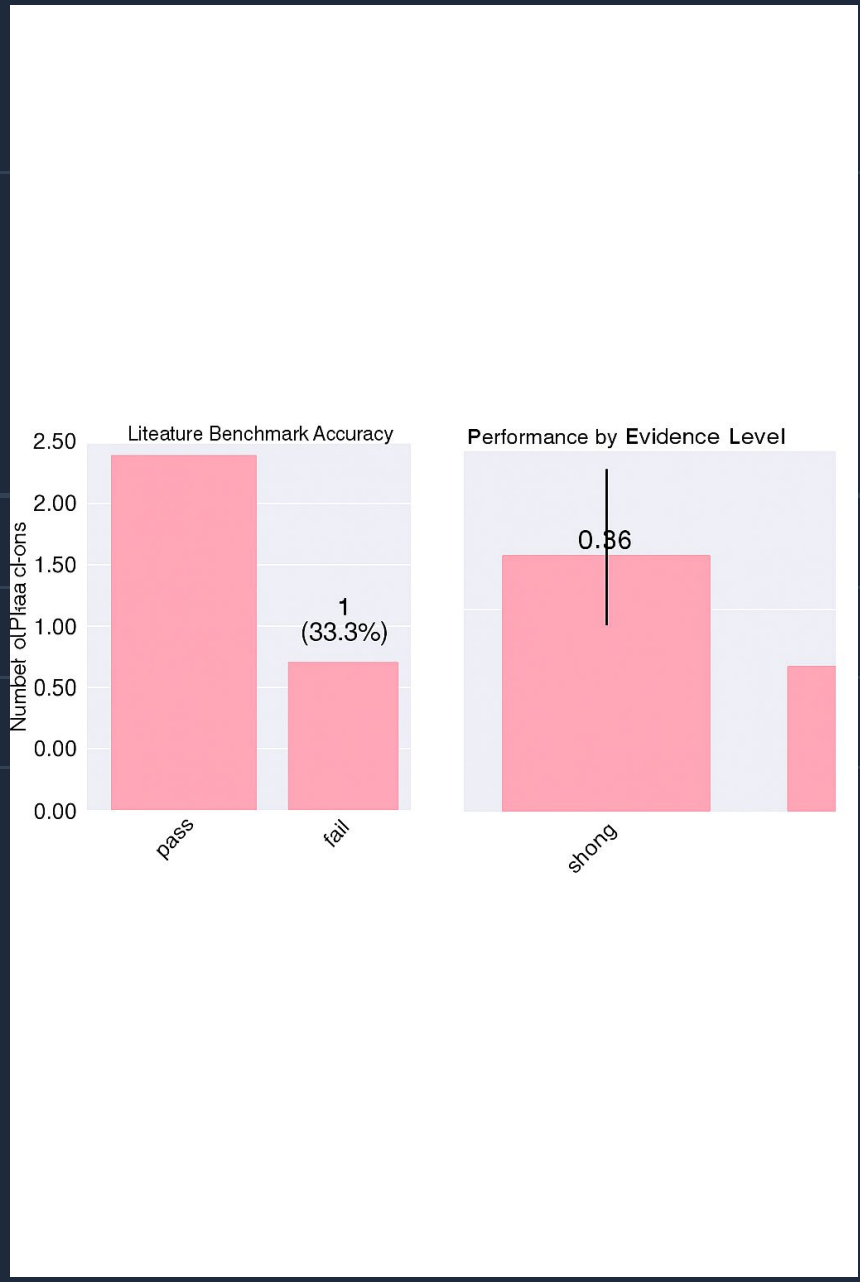

Computational Efficiency (App A9)

Despite architectural complexity, the system maintains computational tractability.

Speed

100 candidates screened in < 2 minutes per cycle.

Memory

< 50MB overhead for success memory across all agents.

Total

Complete 50-cycle campaign uses ~48 GPU-hours on modest hardware (Tesla T4).

Reproducibility Checklist (App A12)

Data Availability

- **Dataset S1:** Initial training data (500 alloys).
- **Dataset S2:** 48-dim feature vectors.
- **Dataset S3:** Raw measurements.

Software Stack

- Python 3.10.12
- PyTorch 1.13.1
- XGBoost 1.7.6
- Scikit-learn 1.2.2

Random Seed: **42** fixed for all stochastic processes.

Protocol (Appendix A6)

- **Synthesis:** Arc melting under argon (purity >99.9%). Flipped/re-melted 5x.
- **Treatment:** Annealed at 1000°C for 48h, water quenched.
- **Hardness:** Vickers (500gf load, 15s dwell).
- **Corrosion:** Potentiodynamic polarization in 3.5 wt% NaCl.
- **Structure:** SEM/EDS analysis for phase distribution.

Limitations

Scaling

Scaling to High-Entropy Alloys (7+ elements) requires addressing the curse of dimensionality via advanced feature selection.

Phases

Phase prediction is limited to solid solutions. We lack integration with Crystal Graph Neural Networks (CGNN) for intermetallics.

Process

Focused solely on composition; lacks integration with post-processing parameters (heat treatment, cooling rates).

Future Work

- **Generative Agents:** Replace Hypothesis Generator with Foundation Models (LLMs) for novel element combinations.
- **Multi-Scale Validation:** Incorporate microstructural prediction models to bridge atomic features with macroscopic performance.
- **Transfer Learning:** Cross-domain knowledge transfer to accelerate initial cycles.
- **Multi-Furnace Scalability:** Adapt RefereeAgent for parallel, asynchronous global feedback.

Conclusion

We demonstrated that **Architectural Innovation**, not model scale, drives the next frontier of AI-accelerated scientific discovery.

- 7-Fold Reduction in Cost
- 21 Novel Alloys
- Closed-Loop Resilience

Questions?