TECHNICAL REPORT

MatterSpace Lattice: Hosted Blind Rediscovery of Single-Atom Alloy Catalysts

Faruk Guney

Vareon Inc., Irvine, California, USA

March 2026

ABSTRACT

MatterSpace Lattice is a delivered materials discovery product — not a research prototype, not a conceptual framework, but a hosted deployment that runs constraint-guided, goal-driven, AI-native discovery campaigns against real computational targets. In blind rediscovery benchmarks on its live production deployment, MatterSpace Lattice recovered Re_1@Ni(111)+CH_4 and Ir_1@Ni(111)+CH_4 single-atom alloy catalysts to sub-angstrom accuracy: 655 milliangstrom and 508 milliangstrom RMSD respectively, both below the 1.00 Å threshold. The product passed all 82 of 82 surface checks spanning API correctness, constraint enforcement, campaign orchestration, and result integrity.

Executive Summary

This document positions MatterSpace Lattice as a delivered materials discovery product whose central validation claim is blind rediscovery of two sealed catalyst targets through a hosted deployment. Two sealed targets — Re_1@Ni(111)+CH_4 and Ir_1@Ni(111)+CH_4 — were recovered to sub-angstrom RMSD through the production API, and the deployment passed all 82 of 82 product surface checks covering constraint enforcement, campaign orchestration, validation integrity, and API correctness.

The significance of blind rediscovery as a product validation strategy is straightforward: if a materials discovery product can autonomously find known catalysts that human researchers took years to identify, without any target information leaking into the generation process, then its novel discoveries on unknown targets deserve serious scientific consideration. Blind rediscovery is the product's north star — the single benchmark that ties research claims to live, deployable capability.

Product Architecture

MatterSpace Lattice is built on three architectural principles that distinguish it from conventional computational materials discovery tools:

Constraint-Guided: Explicit physical and chemical constraints shape the search process from the outset. Rather than generating unconstrained candidates and filtering them post-hoc, MatterSpace embeds constraints — stoichiometry, lattice geometry, bonding rules, surface chemistry — directly into the generative dynamics. This produces valid-by-construction candidates, eliminating the combinatorial waste that characterizes conventional screening workflows.

Goal-Driven: Every campaign begins with declared objectives: target properties, composition ranges, performance thresholds, and validation criteria. The discovery engine uses these objectives to allocate computational resources, select generation strategies, and evaluate candidates against user-defined success criteria. Goal-driven operation means the product works toward specific outcomes rather than exploring aimlessly.

AI-Native Architecture: MatterSpace Lattice is designed to be operated by AI agents as its primary interface modality. The REST API, Python SDK, and Model Context Protocol (MCP) integration provide structured, programmatic access to every product capability. Human users interact through the same surfaces — the product does not maintain separate "human" and "machine" interfaces that drift out of sync.

Campaign Modes

MatterSpace Lattice supports four campaign modes, each addressing a distinct discovery scenario:

Campaign Mode	Objective	Target Knowledge
Open Discovery	Explore novel compositions and structures without a predefined target	None — unconstrained search
Prototype Optimization	Improve a known candidate structure toward better performance	Starting structure provided
Guided Rediscovery	Rediscover a known target with partial structural hints	Partial constraints provided
Blind Rediscovery Benchmark	Recover a sealed target with zero structural information	None — sealed reference

Blind Rediscovery Benchmark is the most demanding mode and serves as the product's primary validation benchmark. In this mode, the generation engine has no access to the target structure, composition, or properties during the discovery process. All comparisons are performed post-hoc against sealed references, ensuring that any match is a genuine independent discovery rather than an optimization toward a known answer.

Domain Packs

MatterSpace Lattice ships with nine domain packs, each encoding the physical constraints, validation criteria, property targets, and computational parameters appropriate to a specific class of materials:

#	Domain Pack	Target Applications
1	magnets	Permanent magnets, magnetic storage, spintronic devices
2	battery_cathodes	Li-ion and post-Li cathode materials
3	solid_electrolytes	Solid-state battery electrolytes
4	catalysis	Heterogeneous catalysts, SAA catalysts, electrocatalysts
5	thermoelectrics	Waste-heat recovery, Peltier devices
6	photovoltaics	Solar cell absorbers, transparent conductors
7	high_entropy_alloys	Multi-principal-element alloys, extreme-environment materials
8	superconductors	High-T_c superconductors, quantum computing materials
9	thermal_barrier_coatings	Turbine blade coatings, aerospace thermal protection

Each domain pack defines the physical constraints, target property ranges, scoring functions, and validation criteria specific to its material class. The catalysis domain pack was used for the blind rediscovery benchmark documented in this report.

Blind Rediscovery Results

Blind rediscovery was executed on the hosted MatterSpace Lattice deployment using NVIDIA A10G hardware — demonstrating that the product achieves research-grade results on production-class (non-research) GPUs.

Target	Level A (50 candidates)	Level B (50 candidates)	Best RMSD	Status
Re_1@Ni(111)+CH_4	50 / 50 pass	26 / 50 pass	655 mÅ	✓ PASS
Ir_1@Ni(111)+CH_4	50 / 50 pass	30 / 50 pass	508 mÅ	✓ PASS

Both targets achieved 100% Level A pass rates (all 50 candidates below the adsorption energy threshold) and majority Level B pass rates (26/50 and 30/50 respectively above the fingerprint similarity threshold). The best RMSD values — 655 mÅ for Re_1@Ni and 508 mÅ for Ir_1@Ni — both fall below the 1.00 Å (1000 mÅ) threshold defined for the hosted deployment protocol.

These results were achieved on NVIDIA A10G GPUs — production-class hardware significantly less powerful than the A100 used in the research paper. The fact that sub-angstrom rediscovery succeeds on A10G hardware confirms that MatterSpace Lattice's computational requirements are compatible with standard cloud deployment infrastructure.

Product Interfaces

MatterSpace Lattice exposes three interface surfaces, all providing equivalent access to the full product capability:

REST API: The primary programmatic interface. RESTful endpoints organized into resource groups: campaigns, candidates, constraints, validation, domain packs, and results. All operations are stateless and authenticated. Request and response payloads use JSON with typed schemas.

Python SDK: A thin client library wrapping the REST API with Pythonic conventions. Provides typed data classes for campaigns, candidates, and results. Supports both synchronous and asynchronous usage patterns. Available via pip.

Model Context Protocol (MCP): Native MCP integration enabling AI agents to discover and invoke MatterSpace capabilities through the standard MCP tool protocol. MCP-compatible agents can list available tools, inspect schemas, and execute discovery campaigns without custom integration code.

All three interfaces share the same underlying API surface. A campaign launched via the Python SDK produces identical results to one launched via REST or MCP. This interface equivalence is verified as part of the 82/82 product surface checks.

Limitations and Roadmap

Current Limitations:

- No DFT confirmation has been performed on generated candidates. All energetic evaluations use machine-learned interatomic potentials. DFT validation is planned as a near-term milestone.
- No experimental validation (synthesis and characterization) has been conducted. The blind rediscovery results are computational predictions awaiting laboratory confirmation.
- The current deployment supports single-adsorbate configurations. Multi-adsorbate and multi-site catalysis scenarios are not yet supported.

Planned Extensions:

- **MatterSpace Pharma:** Drug-like molecule generation with constraint-guided conformer search and ADMET property targeting.
- **MatterSpace Chip:** Semiconductor material discovery for next-generation logic and memory devices.
- **MatterSpace Compute:** Materials optimization for quantum computing substrates and interconnects.

Each planned extension will use the same core constraint-guided generative engine with domain-specific packs, validation protocols, and property targets. The product architecture is designed for domain extensibility — new material classes require new domain packs, not new engines.

AI-native research and engineering, built from the ground up on first principles.