Generative Artificial Intelligence for Inverse Materials Design

Junwu Chen[1], Xu Huang[2], Yuanqi Du[3], Gerbrand Ceder[2], Philippe Schwaller[1]

[1]EPFL, Lausanne, Switzerland
[2]University of California, Berkeley, California, USA
[3]Cornell University, Ithaca, NY, USA

junwu.chen@epfl.ch; philippe.schwaller@epfl.ch; xu_huang@berkeley.edu; gceder@berkeley.edu;

Generative artificial intelligence (AI), exemplified by diffusion models and large language models (LLMs), is fundamentally reshaping the landscape of science and technology. Here, we explore how these powerful models can be harnessed for materials design and autonomous research. First, we propose MatInvent, a general and efficient reinforcement learning (RL) workflow that optimizes diffusion models for goal-directed crystal generation. MatInvent achieves robust optimization across 15 single-objective tasks and 2 multi-objective tasks with competing properties, encompassing electronic, magnetic, mechanical, thermal, and physicochemical properties. Compared to state-of-the-art methods, MatInvent exhibits superior generation performance under specified property constraints while dramatically reducing the demand for property computation by up to 378-fold. Second, recognizing that human-specified objectives are imperfect proxies for true scientific goals, we introduce SAGA (Scientific Autonomous Goal-evolving Agent), in which LLM agents autonomously perform solution optimization, analyze outcomes, propose new objectives, and convert them into computable scoring functions, enabling systematic exploration of objective space across inorganic materials design and beyond. Third, we develop CASCADE, a self-evolving agentic framework that acquires skills and codifies knowledge through continuous learning and self-reflection. CASCADE can independently learn how to control various complex tools, and ultimately put them together in an approach to solve research questions. Evaluated on our SciSkillBench benchmark of 116 materials and chemistry research tasks, CASCADE achieves a 93.3% success rate, outperforming all baselines including Claude Code. We further demonstrate CASCADE's real-world impact across four research scenarios: autonomously determine materials properties such as piezoelectricity; formulating and validating hypotheses on systematic differences in various machine learning interatomic potentials; integrating into an autonomous laboratory to drive end-to-end halide materials research, from synthesis and characterization to ionic conductivity analysis; and reproducing simulation results from published papers.

References

[1] Zeni, C., Pinsler, R., Zügner, D. et al., A generative model for inorganic materials design, Nature, 2025, 639, 624–632.
[2] M. Bran, A., Cox, S., et al., Augmenting large language models with chemistry tools., Nat. Mach. Intell., 2024, 6, 525–535.
[3] Szymanski, N.J., et al., An autonomous laboratory for the accelerated synthesis of inorganic materials., Nature, 2023, 624, 86–91.
[4] Chen, J., Guo, J., et al., Accelerating inverse materials design using generative diffusion models with reinforcement learning., 2025, arXiv preprint (arXiv:2511.03112).
[5] Du, Y., et al., Accelerating Scientific Discovery with Autonomous Goal-evolving Agents., 2025 arXiv preprint (arXiv:2512.21782).
[6] Huang, X., Chen, J., et al., G., CASCADE: Cumulative agentic skill creation through autonomous development and evolution., 2025, arXiv preprint (arXiv:2512.23880).

Figures

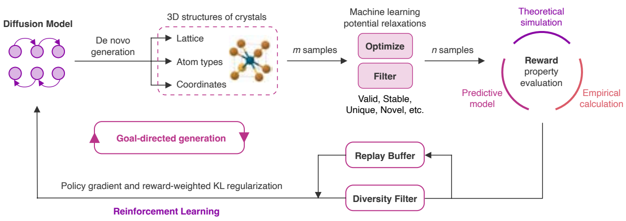

Figure 1. The schematic overview of MatInvent workflow for goal-directed material generation.

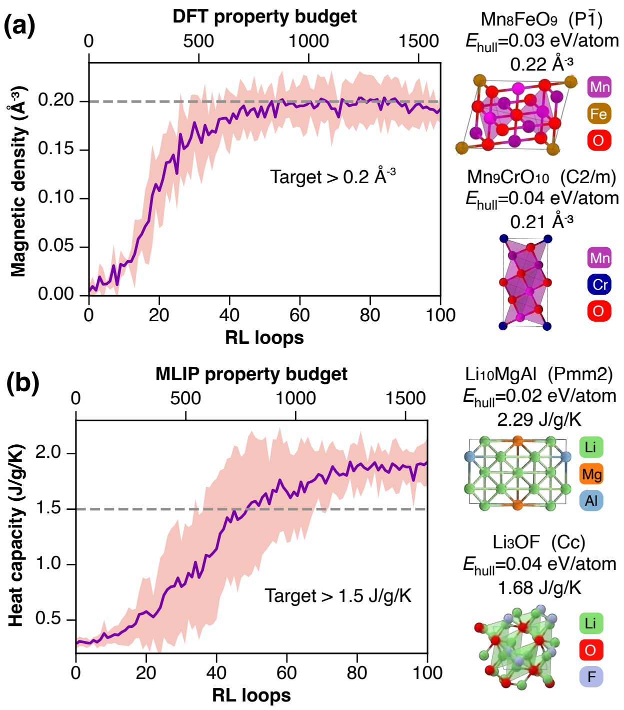

Figure 2. MatInvent performance on single property optimization. The RL optimization curves (left) and visualizations of some generated crystal structures (right) on different inverse design tasks.

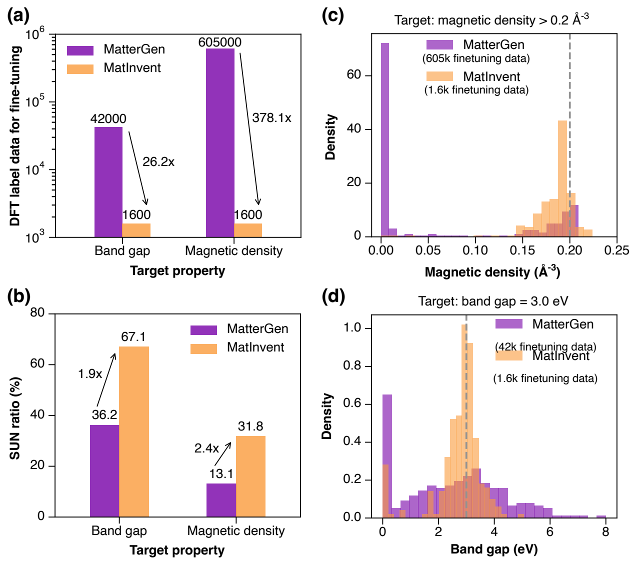

Figure 3. (a) Number of DFT-labeled data used for model fine-tuning in the MatInvent workflow and conditional generation of MatterGen on two tasks. (b) SUN ratios of generated structures from MatterGen conditional generation and RL-finetuned diffusion model. Probability density distributions of property values of SUN structures generated by RL-finetuned diffusion models and MatterGen's conditional generation, respectively.

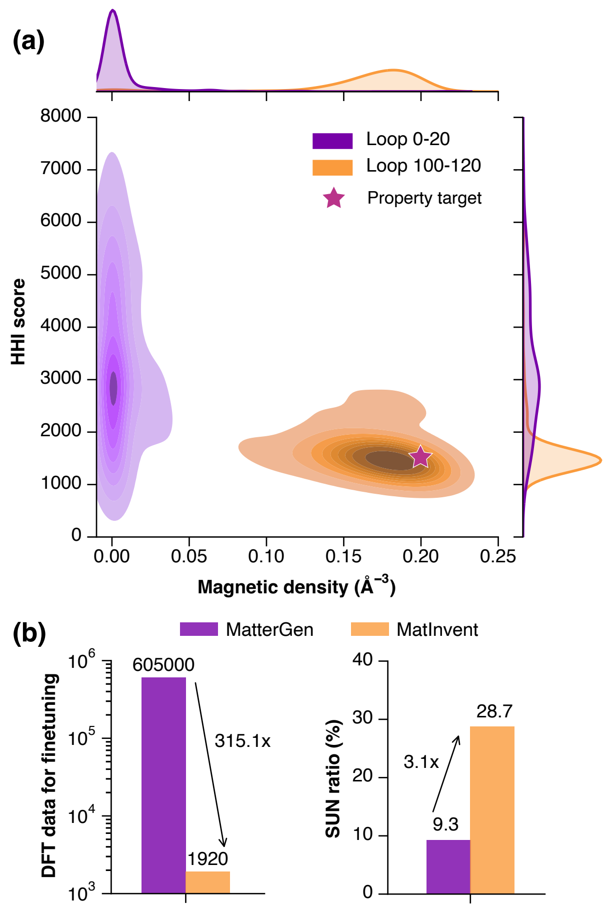

Figure 4. (a) Property distribution of SUN structures generated during the initial (0–20 loops) and final (100–120 loops) stages of RL process. (b) Amount of DFT-labeled data used for model fine-tuning (left) and SUN ratios of generated structures (right) for MatterGen conditional generation and MatInvent workflow.

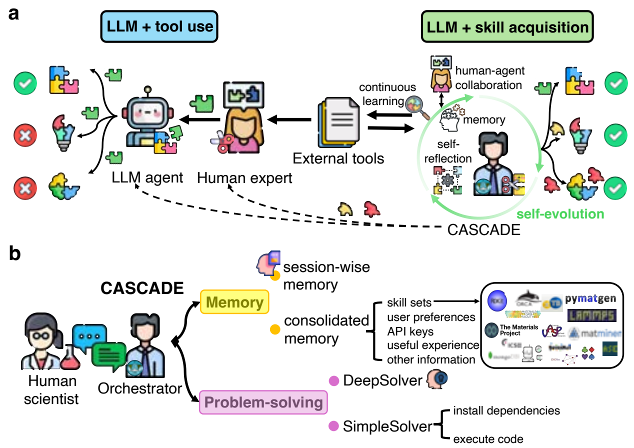

Figure 5. (a) A puzzle-solving metaphor of the 'LLM + tool use' versus the 'LLM + skill acquisition' paradigm. (b) The architecture of CASCADE, a self-evolving agentic framework for complex materials and chemistry tasks.

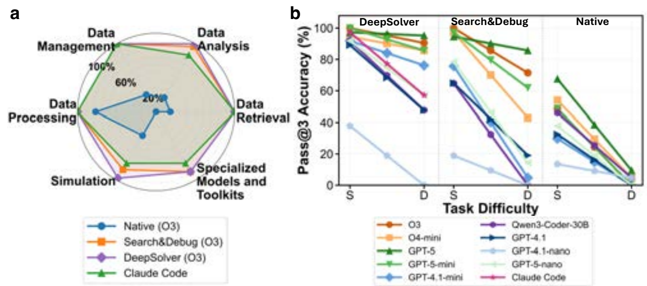

Figure 6. (a) Pass@3 accuracy across all questions by task category for CASCADE and baselines. (b) Pass@3 accuracy against task difficulty on Level 1 questions (S = Simple, D = Difficult). Claude Code (red star) appears in the DeepSolver section.

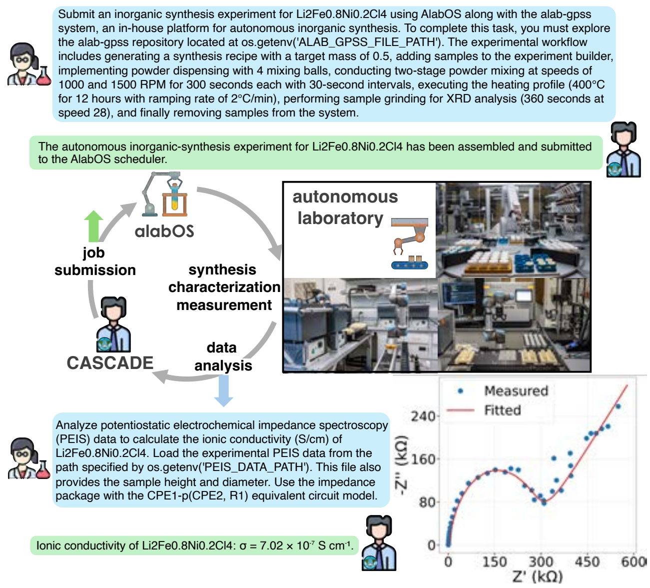

Figure 7. CASCADE integrated into the autonomous lab to drive end-to-end halide materials research, from synthesis and characterization to ionic conductivity analysis.