Accelerating Ammonia Decomposition Catalyst Discovery with AI

Samuel P. Gleason [1] **Martin Siron** [1] **Alan Barrios** [2] **Richard Tran** [1] **Joseph Musielewicz** [1]
Piyusa Priyadarsan Pattanaik [2] **Anvesh Nathani** [1 3] **Tristan Deleu** [1] **Etienne Du Fayet** [1]
Jean-Gabriel Argaud [1] **Sagar Pal** [1] **Hannah Bull** [1] **Cory Hargus** [1] **Svetlana Heyte** [2] **Victor Schmidt** [1]
Egon Heuson [2] **Sébastien Paul** [2] **Alexandre Duval** [1]

[1]*Entalpic, France* [2]*Univ. Lille, CNRS, Centrale Lille, Univ. Artois, UMR 8181-UCCS-Unité de Catalyse et Chimie du Solide, F 59000 Lille, France* [3]*School of Aerospace and Mechanical Engineering, University of Oklahoma, Norman, OK 73019-1052, USA.* Correspondence to: Samuel P. Gleason samuel.gleason@entalpic.ai, Sébastien Paul sebastien.paul@centralelille.fr , Alexandre Duval alexandre.duval@entalpic.ai .

Fig. 1: Overview of the catalyst screening and experimental validation workflow. Top: Candidate bulk crystal structures are generated by a generative model or sampled from the LeMaterial database [1]. Predictive models are then used to identify stable surface terminations and predict catalytic performance for NH_3 decomposition. Bottom: Experimental validation of the computational pipeline.

1. Introduction

The development of efficient catalysts for ammonia (NH_3) decomposition is a key challenge in enabling transportable, carbon-free hydrogen energy systems [2, 3]. State of the art industrial catalysts for NH_3 decomposition are ruthenium-based, which are highly active but suffer from prohibitive cost and limited elemental availability [4]. Nickel-based catalysts represent the most viable low cost alternative, however, their substantially lower catalytic efficiency is a major drawback [4]. This performance–cost gap highlights the need for new NH_3 decomposition catalysts that outperform nickel while remaining economically competitive.

In this work, we introduce a Machine Learning (ML) accelerated catalyst discovery pipeline that combines large scale computational screening with experimental synthesis, catalytic performance benchmarking, and structural validation. Our high-throughput inference workflow enables rapid identification of low cost and highly active NH_3 decomposition catalyst candidates. Using this approach, we experimentally synthesize catalysts that outperform conventional nickel-based benchmarks, validating the effectiveness of our materials discovery pipeline. Our results highlight how integrated ML and experimental workflows can accelerate materials discovery and provide a scalable, generalizable framework for advancing innovation in sustainable energy and other domains.

2. Results

2.1 Catalyst Generation

In this work, we use a generative ML model to propose candidate bulk crystal structures for catalytic NH_3 decomposition. We focus on bi-metallic systems composed of 3d-, 4d-, and 5d- transition metals, balancing catalytic performance with elemental abundance and cost considerations. The model is fine-tuned on the metastable structures in the LeMaterial dataset [1], comprising over 800,000 bulk crystal structures annotated with stability values calculated using density functional theory (DFT). To further align generation with task specific objectives, we additionally fine-tune the model on approximately 12,000 structures labeled with predicted NH_3 decomposition performance from our screening pipeline, as shown in Fig. 1, and incorporate reward signals that explicitly encode synthesizability and material cost. This multi-objective training strategy steers the generative model toward the discovery of stable, high-performing, and experimentally actionable catalyst candidates.

2.2 Catalyst Screening Pipeline

We employ a ML based screening workflow to predict catalytic performance for NH_3 decomposition. Slabs are generated from the unit cell using the Pymatgen surface module [5] from the unit cell, and surface stabilities are evaluated with machine-learned interatomic potentials (MLIPs) [6]. These surfaces are used to construct size dependent Wulff shapes,

enabling the assessment of nanoparticle catalyst stability at a representative diameter of 5 nm. Catalytic performance is then estimated from nitrogen and hydrogen adsorption energies, leveraging the established correlation between N* adsorption energy and NH_3 conversion efficiency [7]. For each surface, the lowest energy N* adsorption configuration is identified and an overall catalyst performance reward is assigned by comparing the catalyst's N* adsorption profile to an optimal target derived from the literature [7].

We select candidates for experimental validation by first confirming performance and stability values using DFT on the most promising materials determined by our MLIP pipeline. Rather than experimentally targeting individual crystal structures, we prioritize material families to ensure experimental feasibility and compatibility with high-throughput synthesis. As illustrated in Fig. 2, the Ni–Mo system is selected due to the strong correlation between predicted activity and catalyst stability. This correlation is particularly advantageous, as it relaxes the requirement for precise composition and phase pure synthesis: achieving approximate compositional control (e.g. Ni_8Mo, Ni_5Mo, Ni_4Mo) is sufficient to ensure the presence of a high-performing phase. This strategy enables robust closure of the automated discovery loop, with characterization directly linked to catalytic performance. In particular, the identification of the pipeline of Ni_4Mo is independently supported by recent experimental studies reporting high NH_3 decomposition activity for this material [8].

Fig. 2: The selection process for experimental investigation. For a given material family (Ni–Mo shown), all sampled compositions and crystal phases are evaluated according to predicted phase stability (color scale from blue to red, with blue indicating higher stability) and NH_3 decomposition performance (vertical axis, with higher positions corresponding to higher predicted activity).

2.3 Experimental Validation

Experimental validation is carried out using a Chemspeed Catimpreg high-throughput synthesis platform, an Avantium Flowrence XL high-throughput catalyst testing system, and a JEOL ARM 200F aberration-corrected electron microscope [9,

10]. Catalysts are synthesized via co-precipitation and impregnation methods. This workflow has produced five catalyst formulations, identified by the screening pipeline, that experimentally outperform a nickel-based benchmark while remaining significantly lower in cost than ruthenium based catalysts, as shown in Fig 3. High-resolution scanning transmission electron microscopy (HR-STEM) is used to confirm catalyst composition and surface structure. Discrepancies between predicted stable surface terminations and observed catalytic performance are systematically fed back into the computational pipeline, enabling iterative model refinement and improved candidate selection in subsequent screening cycles.

Fig. 3: An example of our experimental catalysts which outperform our Ni baseline.

3. Conclusions and Outlook

In this work, we demonstrate a machine-learning-driven catalyst discovery pipeline which accelerates the identification of low cost, high performance catalysts for NH_3 decomposition. By combining generative models, physics-informed screening, high-throughput experimentation, and iterative feedback, our approach bridges the gap between computational prediction and experimental validation. The successful identification and experimental validation of multiple catalysts that outperform nickel-based benchmarks, while avoiding scarce noble metals, highlights the practical impact of this workflow for sustainable hydrogen technology.

Beyond the application to NH_3 decomposition, this work illustrates a broader framework for AI-enabled materials discovery. The emphasis on material families rather than individual crystal structures further enhances experimental robustness and scalability, enabling reliable closure of the discovery loop under realistic synthesis constraints. Our results demonstrate the potential of integrating machine learning and experimentation to address pressing challenges in energy, sustainability, and materials science.

References

[1] Martin Siron, Inel Djafar, Ali Ramlaoui, Etienne du Fayette, Amandine Rossello, Edvin Fako, Matthew McDermott, Felix Therrien, Luis Barroso-Luque, Flaviu Cipcigan, et al. Lematbulk: aggregating, and de-duplicating quantum chemistry materials databases. *arXiv preprint arXiv:2511.05178*, 2025.

[2] T. Han, L. Wei, S. Xie, Y. Liu, H. Dai, and J. Deng. Catalyst design for ammonia decomposition: An overview. *Industrial Chemistry & Materials*, 3:311–331, 2025.

[3] W. U. Khan, H. S. Alasiri, S. A. Ali, and M. M. Hossain. Recent advances in bimetallic catalysts for hydrogen production from ammonia. *The Chemical Record*, 22(7), 2022.

[4] Tianxu Su, Bin Guan, Jiefei Zhou, Chunzheng Zheng, Jiangfeng Guo, Junyan Chen, Yaoyao Zhang, Yuheng Yuan, Wenkai Xie, Nanxin Zhou, Hongtao Dang, Bingyu Xu, and Zhen Huang. Review on ru-based and ni-based catalysts for ammonia decomposition: Research status, reaction mechanism, and perspectives. *Energy & Fuels*, 37:311–331, 2023.

[5] Richard Tran, Zhantao Xu, Bala Radhakrishnan, Dylan Winston, Wei Sun, Kristin A. Persson, and Shyue Ping Ong. Surface energies of elemental crystals. *Scientific Data*, 3:160080, 2016.

[6] B. M. Wood et al. Uma: A family of universal models for atoms. *arXiv preprint*, 2025.

[7] Tong Han, Lu Wei, Shaohua Xie, Yuxi Liu, Hongxing Dai, and Jiguang Deng. Catalyst design for ammonia decomposition: An overview. *Industrial Chemistry & Materials*, 3:311–331, 2025.

[8] Jiaqi Yang, Kailong Ye, Shaohua Xie, Qiang Li, Charles Milhans, Fudong Liu, and Fanglin Che. Diffusion model-guided inverse design of bimetallic catalysts for ammonia decomposition. *Journal of the American Chemical Society*, 148:537–546, 2025.

[9] REALCAT. Realcat — catalyseur de progrès. `http://realcat.fr/`, 2026.

[10] The REALCAT platform is benefiting from a governmental subvention administrated by the French National Research Agency (ANR) within the frame of the 'Future Investments program' (PIA), with the contractual reference 'ANR-11-EQPX-0037'. The Hauts-de-France region, the FEDER, the CPER CHEMACT and BIOHAUTE-CODEFRANCE, Centrale Lille Institut and Centrale Initiatives Foundation are warmly acknowledged for their financial contributions to the acquisition of the equipment of the REALCAT platform.