Bridging the AI4Materials Innovation Gap — A Startup's Blueprint for Industrial Impact

Yu Yang Fredrik Liu[a]

[a] *DeepVerse PTE. LTD., Singapore* fred@deepverse.tech

* Presenting author

1. Introduction

Artificial-intelligence methods have expanded the frontiers of materials discovery at an unprecedented pace: graph-network surrogates such as DeepMind's *GNoME* enumerated ∼380 000 previously unknown stable crystals in silico [1], while self-driving laboratories are compressing years of iterative synthesis into weeks [2]. Argonne's *Polybot*, for instance, tripled the discovery rate of conductive polymer films compared with manual workflows [3].

Despite this momentum, industrial programmes rarely progress beyond *TRL 4*, where components are proven in the lab but not integrated into production [4]. Surveys highlight fragmented software stacks, inconsistent data governance and volatile economics as primary blockers; more than 40 % of manufacturers cite data-integration costs as the reason AI pilots stall [5]. These obstacles persist even a decade after the Materials Genome Initiative (MGI) called for a unified digital infrastructure to halve development timelines [4].

DeepVerse, a startup specializing in AI-driven materials informatics, identifies critical yet overlooked challenges at the AI-materials-industry nexus, leveraging insights from over 120 industry pilots, and proposes an actionable framework *Industrial Materials Intelligence* (IMI) to align academic innovation with manufacturing realities.

2. Motivation

The MGI envisioned a transition from trial–and–error to simulation-guided design, yet industrial uptake lags because of the structural bottlenecks *(i) Data chaos in addition to data sparsity*: industrial materials data are scattered across fragmented schemas (industrial LIMS, ad-hoc spreadsheets, hand-written lab notebooks and IoT sensor logs), while metadata are often undocumented, crippling AI/ML reproducibility. Industry partners have to spend months on data harmonisation before modelling can begin. *(ii) "Sim-to-real" safety gaps*: in self-driving labs, optimisation policies generated *in silico* can breach equipment safety thresholds, causing rig faults and costly stoppages. A pre-validation protocol is therefore needed to filter unsafe recipes. *(iii) Talent bifurcation*: materials scientists rarely master MLOps, whereas data engineers often lack process-chemistry intuition, creating a "missing-middle" skills gap[6, 7].

Beyond these hurdles lie industrial KPIs tied to factory economics that academia rarely optimises. A multi-objective active-learning agent embedding commodity price forecasting model based on supply network and structured market sentiment captures hidden supply-chain risks can outperform pure-performance models. Likewise, scale-up metrics—such as *time-to-first-kilogram*, *performance consistency* and *yield stability*—are absent from most academic leaderboards, yet they are critical to pilot-plant funding decisions.

IMI addresses these barriers through auditable data metrics, ontology-based semantic APIs and low-code interfaces that inject cost and supply-risk considerations directly into every learning loop, thereby converting MGI's simulation-guided vision into factory-floor reality.

Domain	Exemplar
AI discovery models	GNoME large-scale graph networks [1]
Autonomous labs	Polybot [3]; A-Lab [2]
Alloy design	MagNex rare-earth-free magnet [8]
Semantic interoperability	PMD-Core ontology [6]
Production MLOps	Drift-observability studies [9, 7]

Table 1: Summary of related work.

3. Conceptual Framework

We propose the *Industrial Materials Intelligence* (IMI) framework, that is a modular software-and-workflow combines: (i) Data-Readiness Index, (ii) Materials Ontology-driven Middleware, (iii) Multi-Objective Active Learning, (iv) Shadow-Execution Queue, (v) Human-in-the-Loop Dashboards, and (vi) Self-Driving Laboratory Connectors.

3.1 Data-Readiness Index

A data evaluation tool assessing raw data using various metrics—volume, completeness, consistency, correlation, constraints, domain knowledge, label difficulty, signal-to-noise ratio, and hypervolume fill—to quantify dataset fitness, guiding data cleaning process and model trust.

3.2 Semantic Middleware

A graph adaptor aligns heterogeneous LIMS/ELN schemas to PMD-Core and MGI ontologies, exposing lineage-aware REST and OPC-UA endpoints[10] and cutting data-wrangling time from months to weeks in pilot projects.

3.3 Multi-Objective Active Learning

A Bayesian optimiser embeds live commodity prices and supply-risk indices into optimization loops. For our battery electrolyte design pilot with Envision AESC, this reduces reliance on expensive compounds by 17% while maintaining the performance.

3.4 Shadow-Execution Queue

Each AI-proposed recipe is first validated in simulations before robotic execution, cutting laboratory cost by ≈20 % [11]. This "shadow-execution" step filters hazardous or low-yield pathways, mirroring the validation philosophy adopted by Coley *et al.*[12], where they integrating AI planning with a reconfigurable flow-chemistry robot with minimal human intervention—establishing the feasibility of simulation-to-robot hand-offs in chemical synthesis workflows.

3.5 Human-in-the-Loop Dashboards

IMI deploys pipelines for CI/CD and automatic retraining triggered by DRI changes or drift alarms. Engineers or Scientists receive a ranked queue of AI-proposed recipes, but they can also override decisions, and every action feeds back into the registry—closing the learning loop.

3.6 Self-Driving Laboratory Connectors

Accepted recipes are serialised in domain-specific languages such as XDL [13, 14] or ChemOS[15] formats and dispatched to autonomous assets such as ChemOS-enabled benches or ESCALATE workflows [16]. To mainstream this hand-off, the OPC Foundation has launched the "OPC-UA for AI" working group, standardising metadata and control profiles so that SDLs can plug seamlessly into existing plant automation networks[10, 17].

Together, these six elements close the loop from data curation to physical execution and back, ensuring every new data point is captured, validated and re-fed into the IMI framework—without sacrificing safety, traceability or human insight.

4. Results

We benchmark IMI on four industrial pilots. Key performance indicators (KPIs)—*search-space, number of physical experiments, and time-to-spec*—are summarised in Table 2.

Across all pilots, the DRI and material ontology middleware layer reduced data-wrangling lead-time from 6 months to less than 7 weeks (-73%). The shadow-execution queue protocol rejected 20 % of unsafe candidate recipes. Taken together, these results confirm that IMI's disciplined data governance, cost-aware optimisation and closed-loop MLOps has the potential to translate academic AI4M advances into measurable value on the factory floor.

5. Future Directions

IMI demonstrates that rigorous data governance and modular software are critical for scaling AI4Materials. Next steps include open benchmarks for *time-to-first-kilogram*, pilot-as-a-service sandboxes, cross-training programmes and formal safety verification for autonomous synthesis. DeepVerse invites AI4X participants to establish a *Open Industrial Benchmark Consortium* and co-define the metrics that matter for industrial deployment.

References

[1] A. Merchant, S. Batzner, S. Schoenholz, et al. Scaling deep learning for materials discovery. *Nature*, 624:80–85, 2023.

[2] N. J. Szymanski, C. Bartel, K. Persson, et al. An autonomous laboratory for the accelerated synthesis of novel materials. *Nature*, 624:86–91, 2023.

[3] J. Xu, H. Chan, Y. Liu, et al. Autonomous platform for solution processing of electronic polymers. *Nature Communications*, 16:55655, 2025.

[4] National Science and Technology Council. Materials genome initiative strategic plan, 2021. White House Office of Science and Technology Policy.

[5] McKinsey & Company. The state of materials informatics 2024, 2024. Industry report.

[6] G. Schmitz, P. Boehnke, and M. Greiner. Pmd core ontology: Achieving semantic interoperability in materials science. *Materials & Design*, 230:112438, 2023.

[7] K. Kreuzberger, F. Schmidt, and B. Wilder. Machine learning operations landscape: platforms and tools. *Artificial Intelligence Review*, 2025.

[8] C. C. Weiss. Ai discovers new rare-earth-free magnet at 200 times the speed of man, 2024. New Atlas, 12 June 2024.

[9] W. Zhang, S. Kim, and V. Gadepally. Augur: A step towards realistic drift detection in production ml systems. *Proceedings of the ACM SIGMOD Conference*, pages 2501–2512, 2022.

[10] OPC Foundation. Opc foundation launches new working group "opc ua for ai"—revolutionising manufacturing solutions, 2024. Press release.

[11] Anirudh M. K. Nambiar, Christopher P. Breen, Travis Hart, Timothy Kulesza, Timothy F. Jamison, and Klavs F. Jensen. Bayesian optimization of computer-proposed multistep synthetic routes on an automated robotic flow platform. *ACS Central Science*, 8(6):825–836, 2022.

[12] J. M. Coley, D. A. Thomas, J. A. M. Lummiss, et al. A robotic platform for flow synthesis of organic compounds informed by ai planning. *Science*, 365:eaax1566, 2019.

Pilot / Use-case	**Search-space**	**Experiments ↓**	**Time-to-spec ↓**
Thermo catalysts	$\approx 10^8$	≈400 tests	10 yr → 2 yr
Conductive-polymer films	$\approx 10^{12}$	60 → 18	6 mo → 7 wk
Battery electrolytes	$\approx 10^{50}$	120 → 25	9-12 mo → 3 mo
Bio-based cosmetics	$\approx 10^{11}$	50 → 16	6 mo → 5 wk

Table 2: Performance of IMI across different industrial pilot cases. "Experiments ↓" and "Time-to-spec ↓" denote reductions versus historical baselines; *time-to-spec* is the calendar time from project kickoff to the first sample that meets all pre-agreed performance specifications

.

[13] Sami Mehr, Leroy Cronin, et al. A universal system for digitization and automatic execution of chemical synthesis. *Science*, 370:101–108, 2020.

[14] Marc Seifrid, Loïck Roch, and Alán Aspuru-Guzik. Large language models for chemistry robotics. *Autonomous Robots*, 47(3):405–422, 2023.

[15] Florian Häse, Loïck M. Roch, and Alán Aspuru-Guzik. Chemos: Orchestrating autonomous experimentation. *Science Robotics*, 4(26):5559, 2019.

[16] Paul Dietrich, John McCloy, James Martin, et al. Experiment specification, capture and laboratory automation technology (escalate): A software pipeline for automated chemical experimentation and data management. *MRS Communications*, 9(3):795–801, 2019.

[17] OPC Foundation Working Group. Working group artificial intelligence—opc ua profiles, 2024. Accessed 26 Apr 2025.