Entalpic's White Paper:
AI-Driven Materials Engineering at the Atomic Scale

Alexandre Duval, Victor Schmidt, Mathieu Galtier
info@entalpic.ai

The performance of advanced materials, from batteries and catalysts to logic chips and memory components, is determined at the atomic scale of surfaces and interfaces. Engineering these layers with precision is now a strategic imperative, yet R&D remains throttled by vast chemical spaces, competing constraints, and the tight coupling between chemistry and process conditions. Entalpic replaces empirical trial-and-error with an AI-native discovery platform integrating computational chemistry, multi-scale physics and experimental data. By bridging atomic-scale chemistry with reactor-scale performance metrics, we enable an orders-of-magnitude reduction in time-to-discovery for the semiconductor, battery, and specialty chemicals sectors.

Date: May 21, 2026

1 Introduction

Advanced materials are a foundational driver of global economic growth: from logic chips and memory components to energy systems and industrial chemical processes. Their performance is increasingly determined not in bulk, but at the atomic scale of surfaces and interfaces. A single gate oxide, diffusion barrier, or passivation layer, i.e. just a few nanometres thick, can make or break a device's efficiency, reliability, or yield. As the demands placed on these layers grow more stringent, the ability to engineer them atom by atom has become a strategic imperative for the semiconductor, energy, and advanced manufacturing industries.

Atomic Layer Processes (ALPs), encompassing deposition (ALD), etching (ALE) and area-selective deposition (ASD), are the standard methods for achieving this level of control. ALD, the most established of these, relies on sequential pulses of specialised chemicals (e.g. organometallic molecules) called **precursors**, which react with the surface in a self-limiting fashion: each pulse saturates the surface with exactly one molecular layer, preventing precursor mixing in the gas phase and ensuring growth stops automatically once all reactive sites are occupied. This mechanism provides sub-nanometre thickness control, conformal coatings on high-aspect-ratio structures, and precise interface engineering: making ALP a cornerstone of modern materials manufacturing.

However, pushing innovation in this space is hindered by a complex R&D bottleneck: vast chemical spaces, multiple competing properties, fragmented knowledge, and a tight coupling between chemistry and process conditions. Success involves solving a multivariate and multi-objective optimisation problem at the intersection of precursor chemistry, surface reactivity, reactor dynamics, and target device performance. In practice, these dimensions are often explored sequentially and empirically, leading to long and costly development cycles. Research labs produce on the order of tens of validated precursors per couple of years, a pace wholly incompatible with the speed and predictability that industrial development programs demand.

Recent breakthroughs in AI-driven science demonstrate that this bottleneck can be overcome. AlphaFold's resolution of the protein-folding problem [Jumper et al., 2021], AI-accelerated drug discovery pipelines, and the emergence of machine-learned interatomic potentials that replicate quantum-mechanical accuracy at a fraction of the computational cost [Batatia et al., 2023] have established a new paradigm. Data-driven models can navigate complex chemical spaces faster and more systematically than empirical trial-and-error. In materials science, this transformation is just beginning, and atomistic surface processes represent one of its most consequential frontiers.

Entalpic integrates machine learning, computational chemistry, reactor-scale physics, and multi-modal experimental data to address these bottlenecks, structuring the problem across three complementary pillars:

1. **Precursor discovery:** identifying suitable chemicals with optimised delivery windows and controlled reactivity, enabling new deposition chemistries.
2. **Process modelling:** building multi-scale physics AI digital twins of ALP reactors to simulate, optimise, and control deposition processes under realistic operating conditions.
3. **Thin-film discovery:** designing material compositions, interfaces, and structures tailored to target device properties across semiconductors, energy systems, and catalysis.

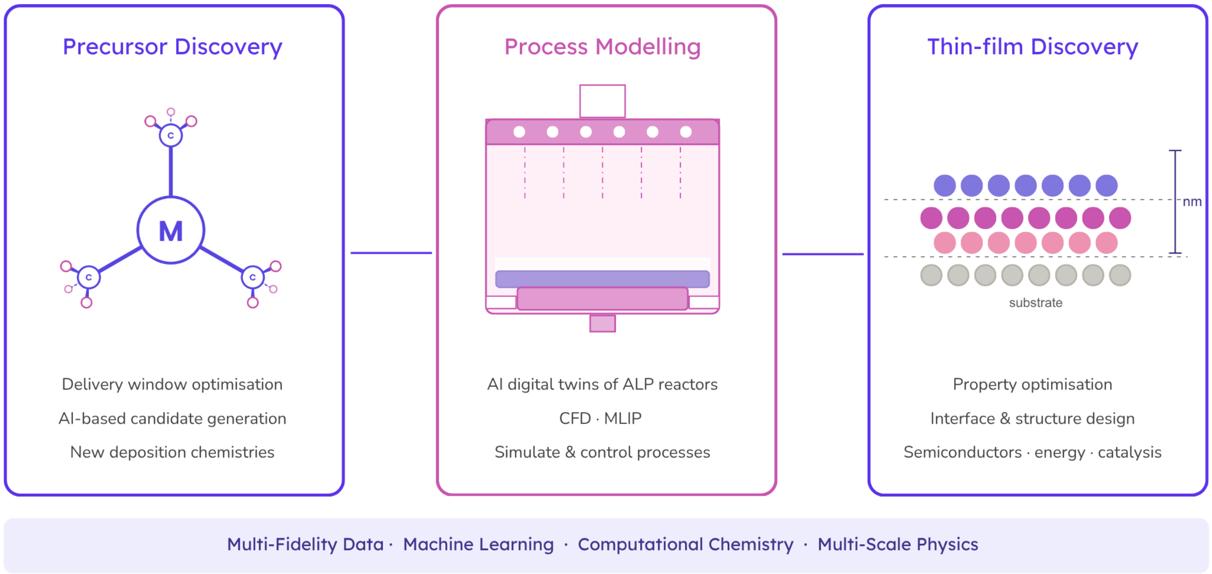

Figure 1 Entalpic addresses ALP bottlenecks across three complementary pillars, feeding each other, underpinned by a shared AI-driven discovery platform.

Our **unified AI platform** is built on a multi-modal, multi-fidelity simulation stack spanning literature mining, fast machine-learned interatomic potential simulations, high-accuracy quantum calculations, process models and experimental measurements. This hierarchy is key: it allows us to generate training data at scale, validate it at high fidelity, and anchor it against experimental ground truth. Analogous to scaling laws in large language models [Kaplan et al., 2020], this approach produces systematic improvements in predictive accuracy as both computational and experimental data accumulate, i.e. a compounding advantage that deepens over time. By continuously exploring and prioritising the most informative simulations and experiments, we focus resources on the highest-value regions of the design space, enabling a more predictive and efficient path to the nano-engineering of tomorrow's materials.

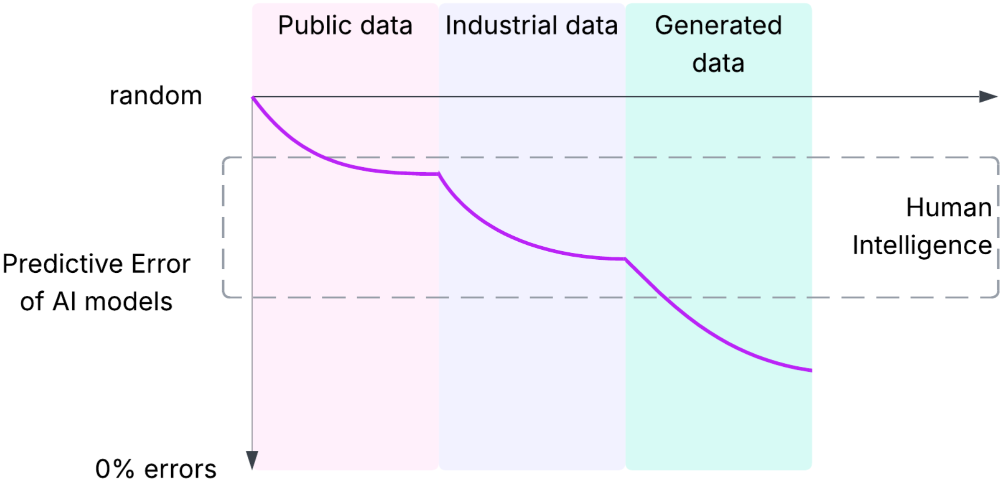

Figure 2 Illustrative scaling law for ALP. AI-driven R&D trains computational models to accurately predict ALP behaviour, with the intent to increase the hit-rate of experiments. Models are successively trained on public data (academic literature & patents), industrial data (existing experimental results from private labs), and generated data (additional experiments or simulations). The expectation is to reach super-human capabilities by integrating diverse data sources that exceed what the human brain can fully ingest.

Note. *Throughout this paper, ALD is used as the representative example of the broader class of Atomic Layer Processes (ALPs). Entalpic's platform and the methods described here extend across the full ALP spectrum, encompassing Atomic Layer Deposition (ALD), Atomic Layer Etching (ALE), and Area-Selective Deposition (ASD). This choice of terminology reflects the established prominence of ALD in the semiconductor literature, not a limitation of scope.*

2 Precursor

2.1 Molecular Predictive Models

The primary technical bottleneck in ALD research is the *"Delivery Window"*: the temperature range where a precursor is volatile enough for transport but remains thermally stable enough to avoid premature decomposition [Barry, 2021]. This constraint can be written as:

$$T_{\text{vap}} < T_{\text{process}} < T_{\text{decomp}} \tag{1}$$

Historically, more than 95% of candidate molecules fail to meet these operational constraints, a phenomenon often referred to as the "Valley of Death" in precursor R&D. Entalpic addresses this by replacing traditional trial-and-error with high-throughput predictive screening.

To screen any potential molecule, we have developed a suite of machine learning models built on complementary molecular representations: SMILES-based architectures (Transformers), 2D graphs (Graph Neural Networks), 3D graphs (Geometric GNNs) [Duval et al., 2023], and tabular models integrating experimental descriptors. These models are trained on curated datasets combining literature, patents, vendor data, and proprietary measurements, and provide the quantitative metrics required to determine whether a candidate satisfies the fundamental ALD requirements above. Our models achieve state-of-the-art performance on standard organic property benchmarks and on proprietary organometallic datasets.

We build predictive models for each of the following physical and chemical descriptors:

- **Volatility metrics**: Predicting enthalpy of vaporization, vapor pressure, and vaporization temperature (T_{vap}).

- **Thermal stability**: Establishing the onset decomposition temperature (T_{decomp}), automated decomposition pathway analysis, melting point, and estimated residual mass loss at high temperatures.

Testing a precursor in a lab requires significant time and specialised characterisation equipment. Our platform enables a radical shift by computing these descriptors in seconds for each molecule, providing both rapid filtering and interpretable signals to guide chemical design. Importantly, the objective of these models is not to reproduce absolute thermodynamic values, but to correctly rank candidates relative to one another. This is a more tractable objective that delivers actionable prioritisation even under model uncertainty, and one that is directly validated against rank-order agreement with experimental measurements.

This predictive layer screens a systematically enumerated space of candidate molecules, reducing it to a tractable subset that satisfies fundamental delivery constraints. Remaining candidates move to the next stage of the funnel: surface reactivity.

2.2 Surface Reactivity

Satisfying the delivery window is necessary but not sufficient: it does not guarantee that a precursor will form a high-quality film, etch, or inhibit successfully. A precursor must also react with the surface in the right way [Puurunen, 2005]. This reactivity pillar evaluates the surface mechanisms of ALP that dictate Growth Per Cycle (GPC), film purity, wettability, steric hindrance from bulky ligands, surface reconstitution, and more.

Traditional Density Functional Theory (DFT) [Kohn et al., 1996] simulations provide high accuracy but are limited to small, simplified systems at 0 K. These conditions poorly represent ALP, where finite-temperature dynamics, defects, dopants, and the stochastic nature of surface saturation all play a role. Moreover, scaling DFT across the thousands or millions of candidates required for systematic screening is computationally intractable.

We address this with Machine Learning Interatomic Potentials (MLIPs), which match near-DFT accuracy for energies and forces while enabling simulations several orders of magnitude faster [Wood et al., 2025]. Crucially, MLIPs often model experimental conditions more faithfully than static DFT, for example by unlocking molecular dynamics at realistic temperatures and timescales, capturing kinetic trapping, surface reconstruction, and saturation behaviour that DFT cannot access at scale.

This evaluation follows a multi-fidelity funnel logic, as illustrated on Figure 3. We start with fast MLIP-based proxies across the full remaining candidate set, such as adsorption energies of the first half-cycle. We then progressively filter down to a smaller subset on which more expensive, higher-fidelity simulations are performed. To capture the broader complexity of ALP reactions, we integrate curated reaction pathway generation frameworks that enumerate plausible mechanisms while accounting for non-ideal surfaces (e.g. defects, dopants) and finite-temperature dynamic behaviour through Molecular Dynamics simulations.

Experimental observables obtained from partner laboratories or customer collaborations serve as the ultimate ground truth and also feed back into model training to continuously improve simulation fidelity. Similarly to the molecular predictive layer, the goal is to obtain reliable relative ranking rather than absolute precision. Systematic model errors largely cancel out when comparing chemically similar candidates, making rank-order predictions robust even where absolute accuracy is limited.

By modelling these interactions, we produce a comprehensive suite of **ALD qualification metrics** to score each precursor candidate:

- Precursor adsorption energy on substrate (nucleation) and surface (growth rate).
- Surface saturation behaviour (self-limiting behaviour, steric blocking).
- Free energies and reaction barriers (growth kinetics, GPC, etch rate).
- Ligand exchange energetics and by-product volatility (film impurity).

- Surface functional group interaction with co-reactant (film composition, reaction completion).
- Selectivity and competitive behaviours (especially for ASD inhibitors).

This high-fidelity modelling allows us to rank thousands of potential precursors based on their ability to deposit a high-quality thin film (or etch), while factoring in cost and production feasibility. Beyond precursor chemistry, these atomistic models also provide physically grounded parameters, such as sticking coefficients, which feed directly into higher-level process models explored in the next section.

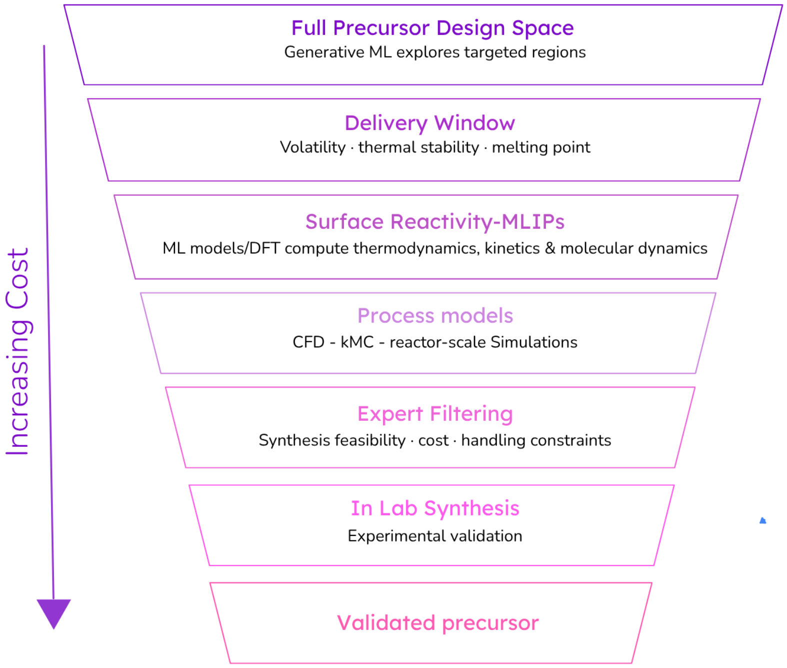

Figure 3 The Entalpic precursor screening funnel: generated candidate molecules are progressively filtered through molecular property prediction, surface reactivity modelling, and process-level validation, converging from millions to a small set of experimentally testable candidates.

2.3 Generative Models

While property prediction models enable rapid evaluation of known molecular libraries, and surface reactivity models validate their ALD behaviour, true innovation requires the ability to navigate the vast, unexplored regions of organometallic space. Existing precursor libraries cover only a small fraction of possible compounds, spanning metal centres, ligand families, coordination geometries, and oxidation states, and the complexity of experimental synthesis has led to slow discovery rates.

We address this by deploying specialised generative AI models designed to propose novel, chemically valid, and high-performing precursors through an **inverse design** rationale.

Rather than producing candidates in a single unconstrained pass, our platform utilises a Transformer-based architecture that constructs precursors sequentially, following a Reinforcement Learning paradigm. Our agent "builds" a molecule through a series of discrete chemical decisions, enabling the injection of rigorous physical and industrial constraints at every step:

- Generation can be conditioned around specific transition metals or rare-earth elements critical for the target device.
- The model samples from libraries of known ligands or produces entirely new heteroleptic and homoleptic

configurations.

- Hard-coded physico-chemical rules (valence rules, coordination chemistry, steric hindrance) ensure that 100% of generated outputs are chemically consistent, valid molecules.

To ensure candidates are not only novel but industrially viable, we employ a multi-stage training pipeline: (1) pre-training on a comprehensive precursor database extracted from scientific literature, patents, vendor catalogs, and proprietary datasets; (2) supervised fine-tuning on high-performing chemical families to internalise the characteristics of stable organometallics; and (3) Reinforcement Learning with Verifiable Rewards (RLVR) to optimise for target criteria while encouraging chemical diversity.

This allows the exploration of millions of precursor candidates in hours, narrowing them to thousands of high-value (potentially "exotic") candidates that satisfy multiple objectives simultaneously. The optimised criteria can be any combination of properties from the predictive models described above, as well as material cost, ligand availability, and synthesis complexity. Specifically, for the latter, we leverage AI agents trained on expert chemistry knowledge to predict optimal synthesis routes and provide actionable recommendations, significantly reducing the empirical overhead for experimentalists. These remaining candidates move to the next stage of the funnel: high-fidelity surface modelling to assess their thin-film deposition performance.

3 Process Models and Digital Twins

While fundamental chemistry dictates the theoretical potential of a molecule or material, real-world ALD qualification metrics are strongly influenced by process conditions, reactor geometry, and transport phenomena. Film conformality, growth rate, and uniformity are not solely governed by surface reactions, but also by precursor delivery, flow dynamics, and reactor design. To account for these effects, Entalpic builds AI digital twins of ALP reactors, with the objective of refining qualification metrics and providing actionable recommendations for process optimisation.

Our digital twins couple AI models trained on process data and computational descriptors, AI-driven chemistry predictions derived from the surface reactivity models (sticking coefficients or flux equations), and reactor-scale physics simulations (e.g. Computational Fluid Dynamics). The result is an accurate representation of gas flow, precursor transport, and reaction environments, tailored to specific equipment platforms and their geometry, operating conditions, and hardware constraints.

By simulating reactor behaviour under a wide range of operating conditions, we provide tools for:

1. **Cycle optimisation** : identifying optimal precursor dosing, temperature profiles, and purge durations to maximise throughput while maintaining film quality. This provides direct guidance to process engineers and significantly reduces the number of iterations required to qualify a process.

2. **Scale-up strategies** : anticipating transport limitations, non-uniformities, and regime shifts when moving from lab-scale to production-scale reactors. This enables more reliable process transfer and reduces the risk associated with hardware transitions.

3. **Predictive process control** : quantifying the impact of process variations and enabling real-time parameter adjustments to maintain stable operation. This reduces variability, improves yield, and lowers operational costs by minimising off-spec runs.

4. **Multi-objective optimisation** : exploring trade-offs between competing objectives using optimisation strategies (e.g. genetic algorithms) to identify Pareto-optimal operating conditions tailored to specific production constraints.

This digital twin framework translates into actionable outputs for process engineers, including optimised process windows, recommended pulse sequences, and sensitivity analyses of key operating parameters under realistic tool constraints. By combining computational insights with experimental data, this framework delivers predictions that are more robust and generalisable than empirical approaches alone, reducing reliance on costly experimental iteration and shortening the path from molecular design to qualified process. Importantly, these process simulations translate reactor-specific constraints into actionable guidance on precursor selection, dosing, and process conditions, ensuring the right chemistry is identified for each specific configuration.

4 Thin-Film Design

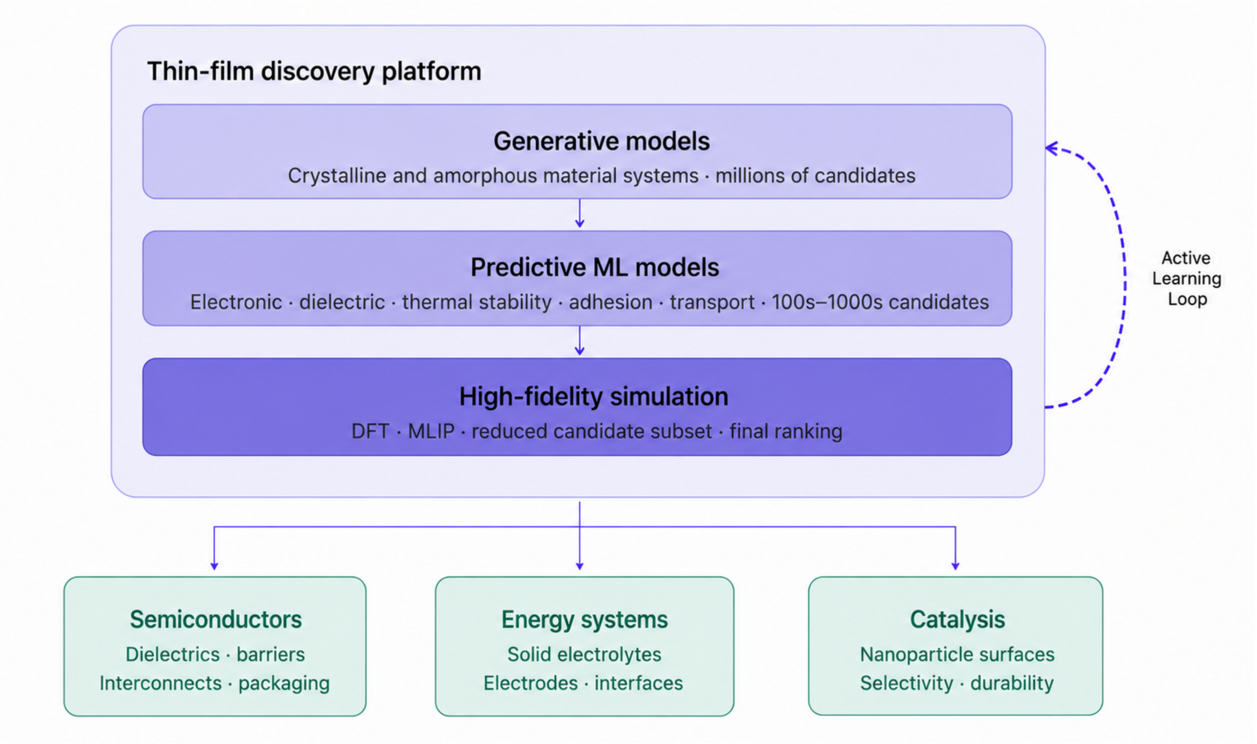

Figure 4 Thin-film design pipeline: candidate material compositions are generated and screened using predictive ML models, with higher-fidelity simulations on shortlisted candidates and experimental validation closing the active learning loop.

Beyond precursor design and process optimisation, the ultimate objective of ALD development is the engineering of thin films with targeted functional properties. These properties, spanning electronic structure and transport behaviour, dielectric response, adhesion, thermal and interfacial stability, directly determine device performance and integration within complex stacks.

At Entalpic, we perform discovery and optimisation of thin-film materials by combining structured exploration of material space with targeted predictive modelling (see Figure 4). Candidate materials are generated using the same technological stack as described in the Generative Models section above, adapted to both crystalline and amorphous material systems. They are evaluated using predictive ML models that estimate key material descriptors, enabling an efficient computational pipeline: fast screening of large candidate sets followed by higher-fidelity simulations on a reduced subset. This allows the ranking of hundreds to thousands of materials and the identification of a small number of high-potential candidates.

This approach is particularly relevant for applications where thin films must satisfy multiple coupled constraints, and can be applied across different industries:

- **Semiconductor technologies:** Interconnect materials, barrier layers, dielectrics, and advanced packaging interfaces, where performance depends on precise control of electrical properties, interface stability, and compatibility with existing process stacks.

- **Energy systems:** Battery electrodes and solid electrolytes, where ionic transport, chemical stability, and interface behaviour must be optimised simultaneously.
- **Catalysis:** Supported nanoparticle systems, where surface composition, interface structure, and stability must be precisely controlled to achieve high activity, selectivity, and long-term durability.

Combined with an active learning loop on computational and experimental data, this approach enables continuous refinement of the search space and improved prediction accuracy. It complements precursor and process modelling by directly targeting the material properties of interest, enabling an integrated design of ALD processes across semiconductors, energy applications, and catalysis.

5 Data, Lab & Agents

A core asset underpinning Entalpic's multi-modal platform is its data infrastructure, built to capture and structure information across the entire thin-film deposition ecosystem. This includes quantum simulations, literature and patents, and experimental data generated both internally and through a network of partner laboratories.

- **Quantum data generation** is performed at scale for precursor–surface–co-reactant systems using optimised workflows (ORCA, VASP, Jobflow, Custodian, Atomate2), with ML-based pipelines enabling thousands of evaluations per GPU per hour.
- **Literature and patent data** provide broad coverage of precursor chemistries, processes, and materials across decades of research.
- **Experimental data** is generated internally, by customer collaborations, and through a network of partner laboratories, covering precursor chemistry, deposition process behaviour, and material properties.

Because of this multi-fidelity AI approach, experimental data serves both as ground truth for validation and as a key driver of model improvement, enabling predictions that are more robust and generalisable than simulation-only or experimental-only approaches.

Through our lab and partner network, we establish a continuous feedback loop between simulation and experimentation, spanning three tightly coupled stages: (1) precursor synthesis and characterisation (NMR, IR, TGA, DSC); (2) thin-film deposition and analysis (QCM, ellipsometry, XRR, XPS, SEM/TEM); and (3) device-level testing across semiconductors, batteries, and catalysis.

To make this heterogeneous data usable at scale, we deploy AI-based acquisition and interpretation systems. Multi-modal agents process literature and experimental outputs into structured representations, while dedicated models assist in interpreting experimental signals (e.g. spectroscopic or structural data), enabling faster feedback and consistent data integration.

Rather than treating all candidates equally, our models estimate uncertainty and expected information gain to prioritise the most informative simulations and experiments, an **Active Learning (AL)** framework. Over time, this creates a self-reinforcing cycle where each experiment increases the predictive power of our platform. This was made possible thanks to our data ontology, which enables us to harmonise diverse datasets. Like our LeMaterial initiative, we ensure that our models are trained on reliable, deduplicated, and high-fidelity information.

This integrated data, experimentation, and learning framework transforms fragmented information into a structured and continuously improving knowledge base, enabling Entalpic to move beyond empirical trial-and-error toward a systematic, data-driven approach to ALD development.

6 Conclusion

The increasing complexity of atomic chemistry and process engineering makes purely empirical approaches slow and costly. Entalpic addresses this by combining AI, computational chemistry, and experimental validation into a unified platform to nano-engineer the next generation of high-performance materials for our industrial partners.

With a team of 30+ experts across software engineering, computational chemistry, quantum physics, machine learning and experimental chemistry, we develop integrated workflows spanning thin-film design, precursor discovery, and process optimisation. This enables the identification of target materials, the molecules required to synthesise them, and the ALP conditions to realise them.

By coupling computational and experimental approaches, we achieve faster and more reliable outcomes than experimental workflows alone, changing the way materials engineering R&D is done.

References

Seán Thomas Barry. *Chemistry of Atomic Layer Deposition.* Walter de Gruyter GmbH & Co KG, 2021.

Ilyes Batatia, Philipp Benner, Yuan Chiang, Alin M Elena, Dávid P Kovács, Janosh Riebesell, Xavier R Advincula, Mark Asta, William J Baldwin, Noam Bernstein, et al. A foundation model for atomistic materials chemistry. *arXiv preprint arXiv:2401.00096*, 2023.

Alexandre Duval, Simon V. Mathis, Chaitanya K. Joshi, Victor Schmidt, Santiago Miret, Fragkiskos D. Malliaros, Taco Cohen, Pietro Lio, Yoshua Bengio, and Michael Bronstein. A hitchhiker's guide to geometric gnns for 3d atomic systems. *arXiv preprint arXiv: 2312.07511*, 2023.

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.

Walter Kohn, Axel D Becke, and Robert G Parr. Density functional theory of electronic structure. *The Journal of Physical Chemistry*, 100(31):12974–12980, 1996.

Riikka L Puurunen. Surface chemistry of atomic layer deposition: A case study for the trimethylaluminum/water process. *Journal of applied physics*, 97(12), 2005.

Brandon M Wood, Misko Dzamba, Xiang Fu, Meng Gao, Muhammed Shuaibi, Luis Barroso-Luque, Kareem Abdelmaqsoud, Vahe Gharakhanyan, John R Kitchin, Daniel S Levine, et al. Uma: A family of universal models for atoms. *arXiv preprint arXiv:2506.23971*, 2025.