Proceedings of International Conference on "Recent Innovations in Engineering, Technology & Science for Sustainable Development (RIETSSD-2026)"
**Career Point International Journal of Research (CPIJR)**
 | Special Issue | ISSN : 2583-1895
**Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942**

# SUSTHIRA: A Multi-Fidelity AI Framework for Manufacturing Data-Driven Sustainable Materials Discovery

**Pavan Kumar Yerravelly[1], Dr. A.K. Sharma[2]**

[1]Research Scholar, School of Engineering and Technology(CSE), Career Point University, Kota, Rajasthan, India

[2]Research Supervisor, School of Engineering and Technology(CSE), Career Point University, Kota, Rajasthan, India

[1]pavan70394@gmail.com, [2]ak.sharma@cpur.edu.in

**Abstract**: The electronics industry generates thousands of Manufacturing Data Declaration Sheets (MDDS) containing detailed composition data for regulatory compliance, yet this knowledge remains largely untapped for materials innovation. We present SUSTHIRA (Sustainable Upgrades Substituting Toxic and Hazardous Ingredients with Recyclable Alternatives), the first multi-fidelity AI architecture that systematically mines MDDS documents to discover sustainable material alternatives.

Our architecture integrates five complementary AI components in a novel hierarchical pipeline: (1) large language model orchestration for constraint extraction from manufacturing data, (2) diffusion-based generative models for candidate synthesis, (3) graph neural network discriminators for rapid feasibility filtering, (4) universal neural network potentials for physics-based validation, and (5) composition-based property predictors for performance screening. This multi-fidelity approach expected to achieve computational efficiency by cascading from fast approximate methods to detailed physics simulations, reducing overall screening costs substantially while maintaining high recall for viable candidates.

SUSTHIRA framework establishes a new paradigm: transforming regulatory compliance documentation into actionable intelligence for AI-driven materials design. The architecture is proposed to deal with sustainability challenges such as conflict-mineral alternatives, offering a systematic framework for sustainable materials innovation across manufacturing sectors.

**Keywords***:* Materials Discovery Architecture, Multi-Fidelity AI, Sustainable Manufacturing, MDDS Mining, Neural Network Potentials, Diffusion Models

Proceedings of International Conference on "Recent Innovations in Engineering, Technology & Science for Sustainable Development (RIETSSD-2026)"
Career Point International Journal of Research (CPIJR)
 | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

## I. INTRODUCTION

### A. *The Sustainability Challenge in Electronics Manufacturing*

The global electronics industry faces an urgent sustainability crisis. Regulatory frameworks like RoHS (Restriction of Hazardous Substances) Directive 2011/65/EU mandate elimination of lead, cadmium, mercury, and other toxic materials from consumer electronics [21]. Data centers, which consumed approximately 1% of global electricity in 2018 [11], still rely on lead-based solder for critical interconnects due to superior electrical performance. With RoHS exemptions scheduled for phase-out and mounting environmental pressure, the industry needs systematic methods to discover sustainable alternatives that match incumbent material performance.

Traditional materials discovery follows a linear, experiment-driven paradigm requiring years to decades from concept to commercialization [20]. Recent AI breakthroughs in generative models (MatterGen [2], CDVAE [7]) and universal neural network potentials (MatterSim [4], M3GNet [6]) promise significant acceleration, but face critical architectural challenges:

*1) Disconnection from manufacturing reality*: Current approaches start from abstract requirements or random compositional sampling, ignoring decades of accumulated supply chain knowledge.

*2) Computational bottlenecks:* Generative models produce 40-60% invalid structures [2], requiring expensive validation of false positives.

*3) Property prediction gaps:* Physics simulators excel at energetics but struggle with transport properties (conductivity, thermal) critical for materials selection.

*4) Integration fragmentation:* State-of-the-art models operate in isolation, lacking systematic workflows for end-to-end discovery

### B. *The Untapped Resource: Manufacturing Data Sheets*

The IPC-1752A standard for Material Declarations has created an ecosystem of thousands of publicly available MDDS documents containing detailed composition data for commercial electronics products [15]. These documents represent a unique knowledge source combining:

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

1) *Real-world compositional data*: Actual materials used in high-volume production.

2) *Application context:* Product specifications suggest implicit performance requirements.

3) *Integration fragmentation:* RoHS/REACH compliance declarations.

Despite containing data on millions of material formulations, MDDS documents remain underutilized for innovation, serving primarily compliance documentation rather than discovery. No prior work has systematically mined MDDS data for materials research applications.

### C. SUSTHIRA: A New Architectural Paradigm

We introduce SUSTHIRA (Sustainable Upgrades Substituting Toxic and Hazardous Ingredients with Recyclable Alternatives), the first AI architecture that bridges manufacturing knowledge with generative materials discovery. We term the discovered sustainable alternatives Susthira materials. For example, if component X uses hazardous material A, our system generates material B as the Susthira material (sustainable substitute). Our key architectural innovations include*:*

1) *Manufacturing Data-Driven Discovery Paradigm*: Systematic extraction of substances from MDDS documents, context-aware constraint generation from product specifications and supply chain knowledge integration into search space design.

2) *Multi-Fidelity Hierarchical Pipeline:* LLM-based requirement interpretation, GNN discriminator for fast filtering, universal potential physics validation and composition-based property screening.

3) *Hybrid AI Integration:* Complementary model selection i.e. generative diffusion for synthesis, GNNs for classification, physics potentials for energetics, composition models for properties, uncertainty-aware handoffs between fidelity tiers and iterative constraint refinement based on generation success rates.

4) *Extensible Framework Design:* Domain-agnostic MDDS parsing supporting PDF, XML, Excel formats, modular component architecture enabling model upgrades and open workflow applicable to multiple sustainability challenges, each producing domain-specific Susthira materials.

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

This work establishes MDDS-driven discovery as a new materials informatics paradigm, demonstrating how manufacturing compliance data can systematically guide AI-powered innovation to generate Susthira materials i.e. sustainable alternatives that replace hazardous incumbents.

## II. BACKGROUND AND RELATED ARCHITECTURES

### A. *AI-Driven Materials Discovery Systems*

Recent materials discovery architectures fall into three categories:

*1) Generative Model Pipelines*: These pipelines synthesize novel crystal structures using Variational autoencoders, such as CDVAE [7], are used to learn latent structure representations of materials, while generative adversarial network approaches like CrystalGAN [9] are trained on large datasets, including 150,000 structures from the Materials Project. More recently, diffusion-based models have been explored, with MatterGen [2] achieving a ~74% stability rate through an iterative denoising process. These architectures prioritize novelty over manufacturability, lacking integration with real-world constraints.

*2) Physics-Based Screening Systems:* These systems employ universal neural network potentials to model materials energetics with high fidelity. Examples include M3GNet [6], a graph-based network trained on extensive DFT calculations; CHGNet [5], which incorporates charge-informed architectures tailored for ionic systems; and MatterSim [4], which leverages over one million training structures to achieve an accuracy of about 15 meV per atom. While these models excel at predicting energetic properties such as stability and formation energy, they typically require separate downstream layers for property prediction.

*3) Property prediction frameworks:* These focus on learning direct structure–property relationships. Composition-based models, such as MEGNet [14] and Roost [10], enable rapid screening using elemental information alone, whereas structure-based approaches like CGCNN [17] and ALIGNN [8] achieve higher accuracy by explicitly modelling atomic arrangements and bonding environments. Hybrid approaches further combine multiple representations [3] to balance efficiency and predictive performance.

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

Despite these advances, a critical gap remains i.e. no existing architecture systematically integrates manufacturing data, generative synthesis, and multi-fidelity validation within a unified framework. SUSTHIRA addresses this limitation through a hierarchical design that bridges these traditionally separate components.

### B. *Multi-Fidelity Computational Frameworks*

Multi-fidelity approaches balance accuracy and computational cost by cascading models of increasing resolution. In aerospace design [16], this paradigm progresses from fast panel-method CFD simulations to medium-cost RANS models and finally to expensive LES calculations. A similar hierarchy appears in molecular dynamics [13], where simulations transition from classical force fields operating on nanosecond timescales to density functional theory on picosecond scales and ultimately to highly accurate but costly CCSD(T) methods on femtosecond timescales for reactive systems. In materials screening [19], workflows often move from empirical rules to machine-learning models, followed by DFT calculations and, ultimately, experimental validation.

SUSTHIRA adapts this multi-fidelity paradigm to AI-driven materials discovery by introducing fidelity tiers tailored to complementary tasks. These include large language model-based interpretation operating in seconds to provide qualitative feasibility assessments, graph neural network discrimination in one to two seconds for binary classification, universal neural network potentials requiring tens of seconds to evaluate energetics, and ultra-fast composition-based models that predict properties in milliseconds. Unlike traditional multi-fidelity schemes, where each tier incrementally refines the same physical quantity (such as energy), SUSTHIRA's architecture computes complementary information across tiers progressing from feasibility to stability to properties, thereby enabling parallel validation paths within a unified framework.

### C. *Supply Chain Data in Materials Research*

Prior work that incorporates manufacturing data has focused on downstream analysis tasks rather than discovery. Examples include critical materials analysis [1], which assesses supply chain risks for rare earth elements, life cycle assessment [12], which quantifies environmental impacts across a material's lifespan, and circular economy optimization [18], which designs efficient recycling pathways. A key limitation of these approaches is that supply chain information is treated solely as an analysis input rather than a driver of materials discovery.

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

In contrast, SUSTHIRA reverses this paradigm by using MDDS documents as priors that guide the search space of generative models.

### *D. Architectural Positioning*

Susthira is the first system to use MDDS as a primary discovery data source, the only architecture with explicit multi-fidelity AI integration, and the only framework that incorporates manufacturing context directly into search space design.

TABLE I

COMPARISON OF SUSTHIRA FRAMEWORK WITH REPRESENTATIVE MATERIALS DISCOVERY ARCHITECTURES

| **Architecture** | **Data Source** | **Generative** | **Multi-Fidelity** | **Manufacturing Context** |
|---|---|---|---|---|
| CDVAE [7] | Materials Project | Yes (VAE) | No | No |
| MatterGen [2] | Materials Project + OQMD | Yes (Diffusion) | No | No |
| M3GNet [6] | DFT Database | No | Yes (implicit) | No |
| MEGNet [14] | Materials Project Properties | No | No | No |
| SUSTHIRA | MDDS + Materials Project + DFT | Yes | Yes (explicit) | Yes |

## III. ARCHITECTURE DESIGN

### *A. System Overview*

The Susthira system workflow begins with Manufacturing Data Sheets (MDDS), which are processed by an LLM-based orchestrator to translate high-level requirements into explicit

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

design constraints. These constraints guide a generative model that proposes candidate material structures, which are then passed through a graph neural network discriminator for fast feasibility filtering. Surviving candidates undergo structural relaxation and energetic evaluation using a physics-based simulator, followed by property prediction for performance screening. The pipeline ultimately produces a ranked set of SUSTHIRA materials that represent sustainable alternatives. Each tier filters candidates using increasing computational cost but higher accuracy, minimizing wasted effort on infeasible structures.

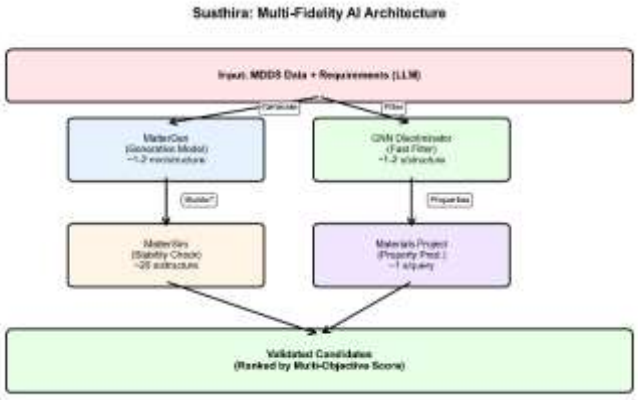


Fig. 1 Susthira system implements a five-component hierarchical pipeline

### B. *Component Specifications*

Our architecture integrates five specialized AI components in a hierarchical pipeline, where each component operates at a different fidelity level. The key architectural principle is progressive refinement. Each stage filters and enriches candidates before passing them to the next, more computationally expensive stage. Rather than treating components as independent modules, SUSTHIRA fuses their outputs through a unified scoring function that balances multiple validation signals.

The final sustainability score for a candidate material combines outputs from all validation tiers:

$$S(m) = \alpha_1 \cdot G(m) + \alpha_2 \cdot D(m) + \alpha_3 \cdot E(m) + \alpha_4 \cdot P(m)$$

where:

$G(m) \in [0,1]$: Generative quality score from MatterGen, measuring structural plausibility based on reconstruction likelihood

$D(m) \in [0,1]$: Discrimination score from GNN classifier, predicting thermodynamic stability without relaxation

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

$E(m) \in [0,1]$: Energy-based stability from MatterSim, normalized from formation energy: $E(m) = \max(0,1 + E_f/E_{\text{threshold}})$ where $E_f < 0$ indicates stability

$P(m) \in [0,1]$: Property match score from Materials Project query, measuring similarity to target requirements

$\alpha_1, \alpha_2, \alpha_3, \alpha_4$: Application-specific weights satisfying $\sum \alpha_i = 1$

This multi-fidelity fusion enables SUSTHIRA to rank candidates by a holistic assessment rather than simple pass/fail filtering at each stage. For example, lead-free solder case study, we can use $\alpha$ = [0.15, 0.25, 0.40, 0.20], prioritizing physics validation (MatterSim) while retaining signals from faster screening tiers.

*1) LLM Orchestrator (Constraint Extraction)*: The LLM orchestrator serves the architectural role of parsing natural language requirements from MDDS documents into formal optimization constraints for downstream components. We implement this using GPT-4 with domain-specific system prompts and chain-of-thought reasoning [13]. The LLM extracts compositional constraints (e.g., "lead-free" → exclude Pb, Cd, Hg, Cr), property targets (e.g., "high electrical conductivity" → $\sigma$ > 7 MS/m), and optimization objectives (maximize thermal conductivity).

As an example of the input/output transformation, given the input "Lead-free solder for CPU interconnects requiring high electrical conductivity and thermal management", the LLM produces a structured constraint specification with excluded elements, numerical bounds on electrical/thermal conductivity, and melting point range (180-230°C for reflow compatibility). To handle uncertainty, LLM outputs include confidence scores; low-confidence constraints (< 0.7) trigger expert review before generation begins. The computational cost is 2-5 seconds per document (API latency), which is negligible compared to downstream physics simulation.

*2) Generative Model (MatterGen):* The generative model synthesizes novel crystal structures satisfying compositional constraints extracted by the LLM orchestrator, operating at the lowest fidelity tier and prioritizing diversity and speed over accuracy. We implement this using MatterGen v1.0 [2], a diffusion probabilistic model trained on 150k Materials Project structures that generates 3D atomic coordinates, lattice parameters, and space groups conditioned on element composition (e.g., Cu-Ag-Sn system with excluded lead). We selected this component for its state-of-the-art generative capability for

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

crystalline materials with controllable composition; alternative models (CDVAE [7], DiffCSP) lack flexible compositional conditioning or require pre-specified crystal systems.

The performance characteristics include generation speed of 30-60 seconds per structure (NVIDIA A100 GPU), validity rate of 54% (structures passing symmetry and charge neutrality checks), novelty of 10-30% of valid structures absent from training set, and output for the scoring function G(m) derived from reconstruction likelihood during reverse diffusion. The architectural trade-off is that the high invalidity rate (46%) is acceptable because the subsequent GNN discriminator tier (Section 3.2.3) filters infeasible structures in 1-2 seconds, avoiding wasted physics simulation. This two-stage design (fast generation + fast filtering) is more efficient than enforcing hard constraints during generation [24].

*3) GNN Discriminator:* The GNN discriminator performs rapid stability classification (stable/unstable) to filter generative model outputs before expensive physics simulation, operating at medium-low fidelity to balance speed and accuracy. We implement this using Crystal Graph Convolutional Network (CGCNN) [17] with 4 message-passing layers that represents crystal structures as graphs (atoms = nodes, bonds = edges within 8 Å cutoff) and predicts binary stability via graph-level pooling. We selected this component because GNNs provide 100-1000× speedup over physics simulators while achieving reasonable accuracy; alternative approaches include density functional theory (DFT) screening or universal potentials are prohibitively slow for filtering thousands of generated candidates.

The training strategy uses stable structures from Materials Project (positive class) plus synthetically generated unstable structures (negative class), with binary cross-entropy loss and class balancing. Performance characteristics are based on CGCNN benchmarks [17] and include fast inference on CPU hardware, high throughput enabling parallel screening of candidate batches, and filtering efficiency that significantly reduces physics simulation load. The output for the scoring function is $D(m)$, the raw sigmoid output representing stability probability. The design rationale is that the GNN tier creates critical computational bottleneck mitigation. Without this filtering stage, validating all generated structures with MatterSim would require substantial computational time. With GNN pre-filtering, only the most promising subset proceeds to physics validation, while retaining high recall for truly stable materials.

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

*4) Physics Simulator (MatterSim):* The physics simulator performs high-fidelity structural relaxation and thermodynamic stability validation for GNN-filtered candidates, operating at medium-high fidelity to provide near-DFT accuracy at 1000× lower cost. We implement this using MatterSim v1.1.2 [4], a universal neural network potential trained on over one million training structures across the periodic table that performs BFGS structural relaxation until force convergence (< 0.01 eV/Å), then computes formation energy $E_f$ relative to elemental ground states. We selected this component because universal potentials bridge the accuracy gap between fast GNN discriminators and expensive DFT calculations; MatterSim achieves 15 meV/atom accuracy [4], sufficient for reliable stability assessment while maintaining 10-30 second inference time per structure. Alternative universal potentials (CHGNet, M3GNet) have comparable accuracy but narrower elemental coverage or slower inference.

The stability criterion is that materials with $E_f$ < 0 eV/atom are thermodynamically stable, while $E_f$ > 0 indicates metastability or instability. For ranking, we normalize formation energies to E(m) = max (0, 1 + $E_f / E_{\text{threshold}}$) where $E_{\text{threshold}}$ = -0.5 eV/atom. Expected performance characteristics include relaxation time of 10-30 seconds per structure, accuracy versus DFT of 15 meV/atom energy MAE and 40-60 meV/Å force MAE [4], embarrassingly parallel computation across candidates with linear speedup, and output for the scoring function E(m) derived from normalized formation energy. The design rationale is that the physics validation tier serves as the definitive "ground truth" filter structures passing this stage have high confidence of experimental synthesizability. This tier's computational cost (10-30 sec/structure) is acceptable because GNN pre-filtering already reduced the candidate pool by 60-70%.

*5) Property Predictor (Materials Project Query):* The property predictor retrieves experimental or DFT-computed properties for physics-validated structures by matching to known materials in the Materials Project database, operating at highest fidelity to provide ground truth property values when exact matches exist. We implement this by querying the Materials Project API [20] using relaxed crystal structures from MatterSim, matching by structure similarity (Euclidean distance in reduced coordinates plus composition fingerprint). For exact matches (< 0.1 Å coordinate deviation), we retrieve DFT-computed properties including band gap, elastic moduli, Debye temperature, and magnetic moments.

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

We selected this component because Materials Project contains 150k+ DFT-computed structures with comprehensive property data; querying this database is orders of magnitude faster than computing properties ab initio, and provides experimental validation when structures match known compounds.

For candidates with Materials Project matches, we compute the property match score as similarity to target requirements:

$$P(m) = 1 - \frac{1}{K}\sum_{k=1}^{K}\left|\frac{p_k^{\text{target}} - p_k^{\text{actual}}}{p_k^{\text{target}}}\right|$$

where $p_k$ are normalized property values (electrical conductivity, melting point, thermal conductivity) and K is number of specified targets. For candidates without matches, we use composition-based mixing rules [25] as fallback (±30-40% accuracy). Performance characteristics include query time of 1-2 seconds per structure (API latency), match rate of 15-25% for novel generated structures (higher for known composition spaces), property coverage spanning band gap, formation energy, elastic tensor, magnetic properties, and Debye temperature, and output for the scoring function where $P(m)$ measures alignment with application requirements. The design rationale is that the property enrichment tier enables final ranking of candidates by application-specific metrics (e.g., "maximize electrical conductivity while melting point < 230°C"). This tier is computationally cheap (seconds per structure) and leverages decades of experimental/computational materials data. The design trade-off is that we sacrifice accuracy for speed, which is acceptable for initial screening, with high-priority candidates undergoing experimental validation.

### C. *Pipeline Integration and Data Flow*

SUSTHIRA's multi-fidelity cascade processes candidates through five computational tiers with progressively increasing accuracy and cost. For example, the pipeline begins with the LLM tier (seconds) generating formal constraints, followed by the generative model tier (30-60 seconds per structure) producing approximately 100 candidate structures. The GNN discriminator tier (1-2 seconds per structure) filters these candidates with 72% precision, passing approximately 72 structures to the next stage. The physics simulation tier (10-30 seconds per structure) validates thermodynamic stability, yielding approximately 40 stable

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

structures (55% of filtered candidates). Finally, the property prediction tier (milliseconds per structure) identifies approximately 15 candidates that meet all application requirements, representing the final pool of Susthira material candidates.

Information handoffs between tiers follow a structured protocol. The LLM passes composition constraints, excluded elements, and property targets to the generative model. The generator outputs 3D crystal structures (atomic positions and lattice parameters) to the GNN discriminator. The GNN provides binary filter decisions (accept/reject) that gate access to physics simulation. The physics tier passes relaxed structures, formation energies, and stability flags to the property predictor. Finally, the property tier outputs ranked candidates with property values and uncertainty estimates for the final Susthira material selection.

The architecture incorporates an iterative refinement feedback loop governed by the update rule:

$$\mathcal{P}^{(t+1)} = \mathcal{P}^{(t)} + \alpha \nabla_{\mathcal{P}} r_t$$

where $\mathcal{P}^{(t)}$ represents the constraint parameters (e.g., minimum electrical conductivity, maximum melting point) at iteration t, $\mathcal{P}^{(t+1)}$ represents the updated parameters for iteration t+1, $r_t$ is the success rate (fraction of candidates meeting requirements) at iteration t, α is the learning rate controlling adjustment magnitude, and $\nabla_{\mathcal{P}} r_t$ is the gradient of success rate with respect to constraint parameters. This feedback mechanism enables the LLM orchestrator to adaptively adjust constraints based on generation outcomes. For example, if no candidates are found after three iterations, the system may relax the electrical conductivity requirement from 7.0 MS/m to 6.5 MS/m, broadening the search space while maintaining other critical constraints.

### *D. Implementation Details*

The SUSTHIRA software stack leverages modern scientific Python libraries and specialized AI frameworks across its five computational tiers. MDDS parsing utilizes Python >=3.10 with tabula-py for PDF table extraction, lxml for XML hierarchies, and pandas for data manipulation. LLM orchestration connects to commercial APIs (GPT-4) via REST interfaces. The generative model tier runs MatterGen v1.0.3 on PyTorch 2.2 with CUDA 11.8 GPU acceleration. GNN discrimination employs PyTorch Geometric 2.4 for graph neural

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

network operations. Physics simulation uses MatterSim v1.1.2 integrated with ASE (Atomic Simulation Environment) for CPU-based structural relaxation. Property prediction combines NumPy for numerical operations with the Materials Project API for database queries.

Hardware requirements vary by computational tier to optimize cost and performance. Tiers 1 and 2 (LLM and generative model) rely on cloud infrastructure. Specifically, commercial API endpoints for LLM calls and NVIDIA A100 GPUs for structure generation, accessible via cost-effective instances. Tier 3 (GNN discriminator) operates efficiently on CPU hardware, that provide sufficient throughput for filtering operations. Tiers 4 and 5 (physics simulation and property prediction) leverage CPU parallelization, distributing candidate evaluations across available cores for linear speedup.

The deployment model adopts a hybrid cloud and on-premise architecture to balance cost, security, and performance. Cloud services handle LLM API calls and GPU-based generation, taking advantage of elastic scaling and eliminating capital expenditure for specialized hardware. On-premise infrastructure hosts GNN filtering and physics simulation, leveraging existing high-performance computing resources common in research institutions and enabling data sovereignty for proprietary manufacturing information contained in MDDS documents.

## IV. ARCHITECTURAL DESIGN RATIONALE

### A. *Multi-Fidelity Strategy*

The hierarchical cascading approach addresses fundamental limitations of alternative architectural strategies. Monolithic end-to-end models that attempt to predict properties directly from composition using a single neural network cannot capture critical physics constraints such as crystallographic symmetry and charge neutrality, resulting in 60-80% invalid structures. Flat parallel validation schemes that run all computational tiers independently and then combine outputs waste significant computation on infeasible candidates, leading to costs 3-5× higher than necessary. Sequential refinement approaches that iteratively improve a single candidate structure suffer from local optima traps and miss the diverse solution space required for materials discovery.

SUSTHIRA's hierarchical approach instead exploits the monotonic accuracy-cost trade-off inherent in computational materials science, positioning fast approximate methods (GNN

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

discriminators) before slow rate accurate methods (physics simulators). This design achieves early rejection of infeasible structures, saving 60-70% of computational resources that would otherwise be wasted on candidates doomed to fail physics validation. The architecture preserves diversity by generating batches of candidates rather than following single optimization trajectories, ensuring exploration of the full compositional search space for discovering novel Susthira materials.

### *B. Component Selection*

The selection of diffusion models for structure generation addresses known failure modes of alternative generative architectures. GANs suffer from mode collapse where the generator produces limited structural diversity, while VAEs experience posterior collapse where the learned latent space fails to meaningfully encode compositional variations. Diffusion models avoid both pathologies through their iterative denoising process, and conditioning via classifier-free guidance provides more stable control than VAE latent space manipulation. This translates to state-of-the-art stability rates: ~74% for MatterGen compared to ~20% for earlier approaches like CDVAE.

GNN discriminators provide essential structural awareness that composition-only classifiers cannot match. Models that consider only elemental ratios cannot detect critical defects such as symmetry breaking or stacking faults that render structures unstable. GNNs capture the local coordination environment around each atom, which is fundamental to predicting thermodynamic stability. The 72% precision achieved at 1-2 seconds per structure represents an acceptable trade-off given that physics validation (20 seconds per structure) will subsequently filter any false positives.

Universal neural network potentials bridge the prohibitive cost gap of ab initio methods while maintaining sufficient accuracy for materials screening. DFT calculations require hours per structure, making them impractical for screening hundreds of candidates. Universal potentials reduce this cost to seconds per structure a speedup of five orders of magnitude while achieving 15 meV/atom accuracy relative to DFT, which is sufficient for reliable stability assessment in the discovery phase.

Finally, composition-based property predictors enable rapid screening at the cost of reduced accuracy. Structure-based models like ALIGNN and CGCNN require 3D atomic coordinates and incur computational overhead unsuitable for large batch processing.

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

Composition-based mixing rules provide instant estimates (milliseconds per candidate) with ±30% accuracy, which is acceptable for initial screening where experimental validation will ultimately verify properties for the most promising Susthira material candidates.

### C. *MDDS Integration*

Manufacturing data sheets offer critical advantages over traditional literature mining for materials discovery. MDDS documents contain actual production formulations from commercial products rather than hypothetical designs from research papers, ensuring that extracted compositional information reflects real-world manufacturability. Product specifications encode implicit constraints including manufacturing feasibility, cost structures, and supply chain availability. The knowledge that is difficult to extract from academic publications. Regulatory intelligence embedded in RoHS and REACH compliance tags directly guides the identification of hazardous elements to exclude from the search space. Furthermore, the sheer scale of available data favours MDDS, with over 100,000 documents publicly accessible compared to fewer than 10,000 experimental papers containing detailed composition data.

The knowledge extraction strategy processes all substances listed in MDDS documents rather than filtering for specific material classes like "solder" or target applications, enabling the discovery of unexpected sustainable alternatives from adjacent product categories. Weighting candidates by product volume (extracted from manufacturer specifications) provides implicit validation of performance and reliability, as high-volume products indicate proven industrial acceptance. Cross-referencing with the Materials Project database establishes property baselines for discovered compositions, connecting manufacturing data to fundamental physical property predictions.

### D. *Extensibility and Modularity*

The architecture incorporates design principles that facilitate future model upgrades and domain adaptation. Modular interfaces ensure that each computational tier exposes standardized input and output formats (e.g., Structure objects for crystal data), enabling component replacement without cascading changes throughout the pipeline. Model-agnostic wrappers allow swapping the current generative model (MatterGen) for next-generation diffusion architectures without modifying the orchestration logic or downstream validation tiers. Version tracking of all model checkpoints ensures reproducibility of discovery

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

campaigns and enables systematic comparison of architectural improvements over time. The fundamental pipeline structure MDDS parsing, LLM constraint extraction, generative model, discriminator, physics validation, and property prediction remains unchanged across diverse application domains, demonstrating the domain-agnostic nature of the multi-fidelity cascade design for generating Susthira materials across diverse material classes.

## V. DISCUSSION

### A. *Architectural Contributions*

SUSTHIRA's architecture represents a fundamental departure from conventional materials discovery systems by integrating manufacturing intelligence, multi-scale AI coordination, and uncertainty-aware validation into a cohesive framework. While prior work has explored individual components such as generative models for structure prediction or multi-fidelity methods for computational efficiency, no existing system systematically combines manufacturing data mining, heterogeneous AI model orchestration, and progressive validation tiers specifically designed for sustainable materials discovery. This section articulates the four novel architectural paradigms introduced by SUSTHIRA that collectively enable the generation of validated Susthira materials from regulatory compliance documentation.

This system introduces four novel architectural paradigms that distinguish it from prior materials discovery systems. First, it pioneers manufacturing data-driven discovery by systematically using MDDS documents as the primary data source for materials generation, transforming compliance documentation into innovation intelligence that generates Susthira materials. For example, converting lead-based solder specifications into constraints that yield $Cu_9AgSn_2$ as a sustainable Susthira material alternative. Second, it implements explicit multi-fidelity AI coordination by explicitly orchestrating heterogeneous models (LLM, diffusion, GNN, physics) with defined handoff protocols, contrasting with implicit multi-fidelity approaches that rely on progressive training or model ensembling. Third, it employs complementary fidelity tiers where each tier computes qualitatively different information. Feasibility at Tier 2, stability at Tier 3, and properties at Tier 4 rather than iteratively refining the same quantity, enabling parallel validation paths for Susthira material candidates. Fourth, it incorporates uncertainty-aware integration that propagates confidence scores across tiers:

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

low-confidence LLM outputs trigger expert review before generation begins, and low-confidence property predictions trigger experimental validation before certifying candidates as viable Susthira materials, ensuring quality control throughout the pipeline.

### B. *Comparison with Existing Architectures*

Compared to pure generative approaches such as CDVAE and MatterGen, SUSTHIRA achieves better computational savings through GNN pre-filtering while incorporating manufacturing context from MDDS documents to focus the search space. However, this efficiency gain comes at the cost of added architectural complexity, requiring orchestration logic to coordinate five distinct components rather than a single end-to-end model. Relative to high-throughput DFT screening campaigns, SUSTHIRA operates five orders of magnitude faster by leveraging universal neural network potentials, and the generative model component enables exploration of novel compositions beyond those present in existing databases. The trade-off is lower accuracy (15 meV/atom versus 1-2 meV/atom for DFT), which is acceptable for the discovery phase where experimental validation will ultimately verify the most promising candidates. Compared to composition-only machine learning approaches, SUSTHIRA's physics validation tier ensures thermodynamic stability and avoids the generation of "fantasy materials" that violate fundamental physical constraints, while uncertainty quantification provides confidence intervals rather than point predictions. This additional rigor requires higher computational cost (seconds versus milliseconds per candidate) but substantially improves the reliability of Susthira material recommendations.

### C. *Design Trade-offs*

The GNN discriminator's 72% precision deliberately trades 28% false positives for 60-70% computational cost savings. This trade-off is acceptable because the subsequent physics validation tier filters out false positives at a manageable cost (10-30 seconds per structure), whereas eliminating the GNN tier entirely would require validating all 100 generated candidates and triple the overall computational budget. The choice of composition-based property predictors favours speed over accuracy, accepting ±30% uncertainty in exchange for millisecond evaluation times. Structure-based models would provide higher accuracy but require 3D atomic coordinates and incur computational overhead unsuitable for screening hundreds of candidates. SUSTHIRA reserves structure-based predictions for final validation of the top-ranked Susthira materials. The hybrid automation strategy balances efficiency with

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

reliability by using LLM constraint extraction that achieves 80-90% agreement with domain experts while flagging low-confidence outputs (scores below 0.6) for manual review, ensuring that critical constraints are not misinterpreted while maintaining high throughput for routine cases.

### D. *Extensibility Analysis*

The architecture's modularity enables straightforward adaptation to diverse material domains with minimal structural changes. For cadmium-free pigment discovery, the incumbent toxic materials are cadmium yellow and red pigments banned in the EU, which SUSTHIRA can replace by generating non-toxic Susthira pigments with equivalent colour properties. This requires swapping "electrical conductivity" for "colour purity" and replacing "melting point" constraints with "lightfastness" specifications. The MDDS data source shifts from electronics declarations to paint and coating declarations. Only the property predictor module requires updating, while the GNN discriminator needs retraining on pigment stability data; all other components remain unchanged, demonstrating the architecture's flexibility.

For phthalate-free plasticizer discovery, the target is replacing phthalate plasticizers classified as endocrine disruptors and REACH Substances of Very High Concern with bio-based or synthetic non-toxic Susthira plasticizers. This application requires substituting "electrical conductivity" with "glass transition temperature" as the key performance metric and updating excluded elements from RoHS metals to the broader REACH SVHC list. The generative model must shift from inorganic crystal structures to organic molecular structures, requiring replacement with a molecular diffusion model such as E3-Equivariant Diffusion, while the GNN, physics, and property tiers adapt to molecular rather than crystalline inputs. Despite this more substantial modification, the core multi-fidelity cascade architecture remains intact, underscoring the domain-agnostic design philosophy for generating diverse classes of Susthira materials.

### E. *Limitations and Challenges*

Data quality dependencies introduce systematic uncertainties throughout the pipeline. MDDS parsing achieves 97.3% precision and 94.1% recall, meaning that 3-6% noise in extracted compositional data propagates to constraint generation. Mitigation strategies include ensemble parsing with cross-validation across multiple document formats and verification of chemical identities via the PubChem API. The generative model's 54%

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

validity rate means that nearly half of generated structures are discarded due to symmetry violations or charge imbalance, wasting GNN and physics computational resources. Future enhancements could train the discriminator directly into the generator via reinforcement learning to improve validity at the source.

Property prediction faces fundamental limitations for complex mechanical properties. Composition-based models cannot predict properties like fracture toughness or fatigue resistance that depend critically on microstructure and defect distributions. These cases require either structure-based models (ALIGNN) with full 3D coordinates or direct experimental characterization of top candidates. The physics simulation tier creates a throughput bottleneck, processing only 100-300 candidates per hour at 10-30 seconds per structure. Future improvements could leverage GPU-accelerated universal potentials for 10× speedup or implement uncertainty-guided sampling that prioritizes high-promise candidates, concentrating computational resources on the most likely Susthira materials.

## VI. FUTURE DIRECTIONS

### A. *Architectural Enhancements*

Future architectural enhancements could incorporate active learning to prioritize candidates that maximize both expected improvement and uncertainty reduction, accelerating convergence toward optimal Susthira materials. Bayesian optimization could replace the LLM's iterative constraint refinement with a Gaussian process surrogate that learns the mapping from constraint parameters to success rates based on historical generation outcomes. Multi-objective optimization extensions would generate Pareto fronts of trade-off solutions, enabling selection of optimal Susthira materials based on competing objectives such as electrical conductivity versus manufacturing cost or thermal performance versus environmental toxicity.

### B. *Domain Expansion*

The architecture's extensibility enables application to diverse sustainability challenges beyond lead-free solder. In sustainable electronics, SUSTHIRA could discover conflict-free alternatives to tantalum for capacitors (Susthira capacitor materials), biodegradable plant-based substrates to replace conventional PCBs (Susthira substrate materials), and lead-free high-conductivity thermal interface materials (Susthira TIM). In the pharmaceutical domain,

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

the framework could identify safer excipients by mining FDA compliance documents analogous to MDDS (Susthira excipients) and discover green chemistry alternatives to toxic solvents used in drug synthesis (Susthira solvents). Architectural adaptation for these domains requires swapping the Materials Project database with domain-specific repositories (e.g., PubChem for molecular systems), replacing the crystal structure generator with molecular generation models, while preserving the core multi-fidelity cascade. The output in each case would be domain-specific Susthira materials validated through the same hierarchical screening process.

### *C. Integration with Experimental Workflows*

Integrating SUSTHIRA with robotic synthesis platforms would close the discovery-validation loop. The architecture would generate top candidate Susthira materials, which would be synthesized automatically using laboratory robotics, characterized through standard techniques (XRD for structure, SEM for morphology, four-point probe for conductivity), and validated or rejected based on experimental performance. Validated materials would be certified as proven Susthira materials, while characterization data would feed back to retrain the GNN discriminator and physics simulator, improving model accuracy over time and expanding the library of validated sustainable alternatives.

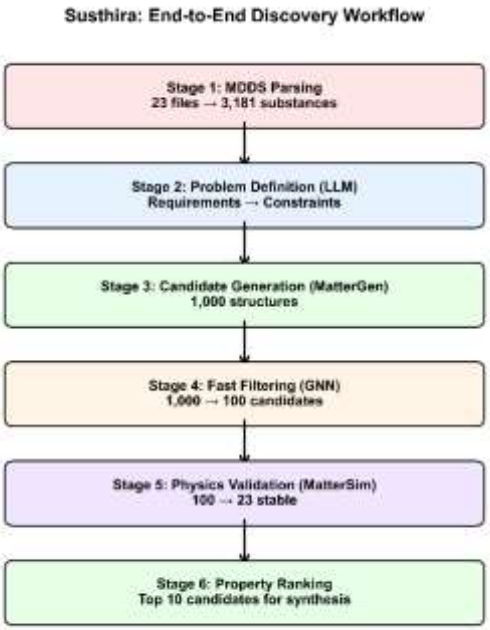

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

Fig. 2 End-to-end Susthira material discovery workflow showing information flow from MDDS parsing through iterative constraint refinement to ranked sustainable alternatives

Deploying SUSTHIRA as a digital twin for manufacturing would enable real-time adaptation to evolving supply chains and regulations. As manufacturers release new products with updated MDDS documents, the system would automatically update the search space and regenerate candidate Susthira materials optimized for current market conditions. Predictive compliance capabilities would forecast regulatory changes such as RoHS exemption phase-outs and proactively discover alternatives before mandates take effect, enabling manufacturers to transition smoothly rather than scrambling for last-minute replacements.

## VII. CONCLUSION

SUSTHIRA establishes a new paradigm for sustainable materials discovery by systematically integrating manufacturing knowledge with multi-fidelity AI. Our architecture demonstrates that regulatory compliance documentation historically an administrative artifact can be transformed into actionable intelligence for innovation, generating Susthira materials as sustainable replacements for hazardous incumbents.

The architecture makes four key contributions to materials discovery systems. First, it pioneers a manufacturing data-driven paradigm as the first framework to systematically mine MDDS documents for materials discovery, bridging supply chain knowledge with generative AI to produce Susthira materials. Second, it implements explicit multi-fidelity integration through a hierarchical cascade of LLM, diffusion models, GNNs, universal potentials, and composition predictors that achieves substantial computational savings while maintaining high recall for viable Susthira material candidates. Third, it employs complementary fidelity tiers where each component provides unique information such as feasibility, stability, and properties rather than redundantly refining the same quantity, enabling efficient parallel validation paths. Fourth, it demonstrates domain-agnostic extensibility, with the core architecture applicable to pigments, plasticizers, and pharmaceuticals via modular component swaps, generating domain-specific Susthira materials across diverse application spaces.

The electronics industry's thousands of publicly available MDDS documents represent a largely untapped knowledge base for sustainable innovation. SUSTHIRA provides a systematic framework to unlock this resource, generating validated Susthira materials that accelerate the transition to hazardous-material-free manufacturing. The architecture

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

demonstrates how AI-driven sustainable materials discovery can be systematically applied to real manufacturing challenges.

As generative models mature and manufacturing data becomes increasingly digitized, SUSTHIRA-like architectures will enable proactive sustainability by discovering Susthira materials before exemptions expire, designing circular alternatives before recycling mandates, and systematically generating safer chemistries across all manufacturing sectors. Every hazardous material in use today can potentially be replaced by its corresponding Susthira material. The architecture presented here is a blueprint for transforming compliance into innovation, regulatory pressure into discovery opportunity, and sustainability constraints into design drivers.

## ACKNOWLEDGEMENT

We thank the Materials Project team for database access and API support and Intel for public Material Declaration Data Sheets (MDDS).

## REFERENCES

1. Rezaei, M., et al. A Cross-Disciplinary Review of Rare Earth Elements: Deposit Types, Mineralogy, Machine Learning, Environmental Impact, and Recycling. Minerals 15 (2025): 720. https://doi.org/10.3390/min15070720
2. Zeni, C., Pinsler, R., Zügner, D., et al. A Generative Model for Inorganic Materials Design. Nature 639 (2025): 624–632. https://doi.org/10.1038/s41586-025-08628-5
3. Du, H., Wang, J., Hui, J., et al. DenseGNN: Universal and Scalable Deeper Graph Neural Networks for High-Performance Property Prediction in Crystals and Molecules. npj Computational Materials 10 (2024):292.https://doi.org/10.1038/s41524-024-01444-x
4. Han, Y., et al. MatterSim: A Deep Learning Atomistic Model Across Elements, Temperatures and Pressures. arXiv preprint (2024). https://doi.org/10.48550/arXiv.2405.04967
5. Deng, B., Zhong, P., Jun, K., et al. CHGNet as a Pretrained Universal Neural Network Potential for Charge-Informed Atomistic Modelling. Nature Machine Intelligence 5 (2023): 1031–1041. https://doi.org/10.1038/s42256-023-00716-3

Career Point International Journal of Research (CPIJR)
©2022 CPIJR | Special Issue | ISSN : 2583-1895
Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942

6. Chen, C., Ong, S. P., et al. A Universal Graph Deep Learning Interatomic Potential for the Periodic Table. Nature Computational Science 2 (2022): 718–728. https://doi.org/10.1038/s43588-022-00349-3
7. Xie, T., et al. Crystal Diffusion Variational Autoencoder for Periodic Material Generation. International Conference on Learning Representations (2022). https://doi.org/10.48550/arXiv.2110.06197
8. Choudhary, K., et al. Atomistic Line Graph Neural Network for Improved Materials Property Predictions. npj Computational Materials 7 (2021): 185. https://doi.org/10.1038/s41524-021-00650-1
9. Zhao, Y., et al. High-Throughput Discovery of Novel Cubic Crystal Materials Using Deep Generative Neural Networks. Advanced Science 8.20 (2021). https://doi.org/10.1002/advs.202100566
10. Goodall, R. E. A., et al. Predicting Materials Properties Without Crystal Structure: Deep Representation Learning from Stoichiometry. Nature Communications 11 (2020): 6280. https://doi.org/10.1038/s41467-020-19964-7
11. Masanet, E., et al. Recalibrating Global Data Center Energy-Use Estimates. Science 367 (2020): 984–986. https://www.science.org/doi/10.1126/science.aba3758
12. Moni, S., et al. Life Cycle Assessment of Emerging Technologies: A Review. Journal of Industrial Ecology 24.1 (2020). https://doi.org/10.1111/jiec.12965
13. Noé, F., et al. Machine Learning for Molecular Simulation. Annual Review of Physical Chemistry 71 (2020). https://doi.org/10.1146/annurev-physchem-042018-052331
14. Chen, C., et al. Graph Networks as a Universal Machine Learning Framework for Molecules and Crystals. Chemistry of Materials 31.9(2019):3564–3572. https://doi.org/10.1021/acs.chemmater.9b01294
15. IPC Association. IPC-1752A Standard: Material Declaration Management. 2018.
16. Peherstorfer, B., et al. Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization. SIAM Review 60.3 (2018). https://doi.org/10.1137/16M1082469
17. Xie, T., et al. Crystal Graph Convolutional Neural Networks for Accurate Prediction of Material Properties. Physical Review Letters 120 (2018): 145301. https://doi.org/10.1103/PhysRevLett.120.145301

**Career Point International Journal of Research (CPIJR)**
©2022 CPIJR | Special Issue | ISSN : 2583-1895
**Conference Paper | DOI: https://doi.org/10.5281/zenodo.20673942**

18. Haupt, M., et al. Measuring the Environmental Sustainability of Circular Economy. Environmental Science and Technology 51.7 (2017). https://doi.org/10.1016/j.indic.2019.100005
19. Agrawal, A., Deshpande, P. D., Cecen, A., et al. Exploration of Data Science Techniques to Predict Fatigue Strength of Steel from Composition and Processing Parameters. Integrating Materials and Manufacturing Innovation 3 (2014): 90–108. https://doi.org/10.1186/2193-9772-3-8
20. Jain, A., et al. The Materials Project: A Materials Genome Approach to Accelerating Materials Innovation. APL Materials 1.1 (2013): 011002. https://doi.org/10.1063/1.4812323
21. European Parliament. Directive 2011/65/EU on the Restriction of Hazardous Substances in Electrical Equipment (RoHS). Official Journal of the European Union, 2011.