Review

Integrating the Theory of Inventive Problem Solving with Large Language Models: Enhancing Reasoning for Innovation in Materials Science at the Molecular Scale

Sergey Gusarov *, Svetlana Sapelnikova, Julio J. Valdes, Anguang Hu, Stanislav R. Stoyanov *

Posted Date: 23 February 2026

doi: 10.20944/preprints202602.1398.v1

Keywords: material design; TRIZ; AI/ML; algorithm of inventive problem solving; reasoning; computational chemistry; chemical engineering

 doi:10.20944/preprints202602.1398.v1

Review

Integrating the Theory of Inventive Problem Solving with Large Language Models: Enhancing Reasoning for Innovation in Materials Science at the Molecular Scale

Sergey Gusarov [1,*], Svetlana Sapelnikova [2], Julio J. Valdes [3], Anguang Hu [4] and Stanislav R. Stoyanov [5,*]

[1] National Research Council Canada, Digital Technologies, 11421 Saskatchewan Drive Edmonton, Alberta T6G 2M9 Canada

[2] SVConsulting , Edmonton, Alberta T6R 0J1

[3] National Research Council Canada, Digital Technologies Research Centre, M-50, 1200 Montreal Road, Ottawa, Ontario K1A 0R6 Canada

[4] Suffield Research Centre, Defence Research Development Canada, Medicine Hat, Alberta T1A 8K6 Canada

[5] Natural Resources Canada, CanmetENERGY Devon, 1 Oil Patch Drive, Devon, Alberta T9G 1A8 Canada

* Correspondence: sergey.gusarov@nrc-cnrc.gc.ca, stanislav.stoyanov@NRCan-RNCan.gc.ca

Abstract

The Theory of Inventive Problem Solving (TRIZ) has long been a cornerstone for systematic innovation in engineering domains, including chemical and materials science. This paper proposes a novel framework that integrates TRIZ principles with large language models (LLMs) to emulate researcher-like reasoning in atomistic materials science. By structuring LLM prompts around TRIZ tools—such as patterns of evolution, contradiction matrices, and inventive principles—we enable models to identify problems, frame contradictions, and generate inventive solutions for challenges like data scarcity, poor interpretability, and unphysical predictions in quantum-chemical simulations and machine learning (ML) models. Drawing on recent artificial intelligence-TRIZ hybrids, like AutoTRIZ and TRIZ-GPT (generative pre-trained transformer), we demonstrate applications in molecular design, such as resolving contradictions in shape-memory polymers. This approach not only amplifies current trends in physics-informed ML and generative design but also democratizes advanced problem-solving, accelerating discoveries toward ideality.

Keywords: material design; TRIZ; AI/ML; algorithm of inventive problem solving; reasoning; computational chemistry; chemical engineering

1. Introduction

In the era of artificial intelligence (AI), large language models (LLMs) like GPT-4(5) (generative pre-trained transformer), DeepSeek and Grok have revolutionized reasoning tasks, enabling chain-of-thought (CoT) prompting for step-by-step problem-solving. However, their outputs often suffer from hallucinations, lack of domain-specific rigor, and inefficiency in handling complex scientific contradictions. This is particularly evident in atomistic materials science, where quantum-mechanical simulations intersect with data-driven machine learning (ML) to predict molecular properties, yet face persistent challenges such as limited datasets and unphysical extrapolations. At the moment, AI models are very bad at highlighting the problem and need strong algorithms for their solution. One of the most pressing issues is data scarcity at the molecular quantum level, where high-fidelity data from methods like density functional theory (DFT), coupled-cluster singles doubles with perturbative

 doi:10.20944/preprints202602.1398.v1

triples (CCSD(T)), molecular dynamics (MD) trajectories, or experimental thermochemistry is prohibitively expensive and time-consuming to generate, often resulting in datasets too small to adequately cover the vast chemical space and thereby severely limiting model generalization [1]. Even datasets comprising 500–1000 molecules are deemed "low-data" regimes, pushing ML models into unstable territories where predictions falter on unseen compounds. Compounding this, molecular representations remain limited or imperfect, as standard descriptors like fingerprints, simple graphs, or 2D features frequently overlook essential physical phenomena such as π-delocalization, long-range polarization, donor-acceptor orbital interactions, hydrogen-bond charge transfer, or aromatic resonance. Graph neural networks (GNNs) excel at capturing bond structures but fall short in encoding electronic intricacies without explicit enhancements, making it challenging for AI models to discern subtle reactivity patterns from raw inputs. This inadequacy in descriptors is a critical bottleneck; they require significant improvements to incorporate richer chemical information, enabling more accurate predictions of dynamic processes like bond formation, energy barriers, or phase transitions in complex molecular systems. For example, Conductor-like Screening Model for Real Solvents (COSMO-RS)-based σ-profile descriptors augmented with dipole moment and induced phosphorus charge have been shown to more effectively capture chemical reactivity and solvation thermodynamics than standard structure-based fingerprints, yielding superior clustering and screening performance for nucleotide analogues targeting the RNA-dependent RNA polymerase (RdRp) of the Severe Acute Respiratory Syndrome Coronavirus (SARS-CoV-2) compared to conventional molecular descriptors and docking scores [2]. Weak extrapolation between chemical scaffolds further exacerbates these issues, with ML models performing admirably on trained scaffolds but failing catastrophically when faced with variations in substituent positions, ring types, or heteroatom environments, as the underlying electronic structures shift in unanticipated ways. This behavior is consistent with recent results on Lewis acidity prediction, where interpretable linear models built on chemically meaningful descriptors (RDKit and Hammett-extended) for a restricted ONO-boron scaffold (ONO = *N*,*N*-bis(3,5-di-*tert*-butyl-2-phenoxy)amide) achieved a mean absolute error of ~6 kJ·mol^{-1} ($R^2 \approx 0.98$), whereas an advanced graph neural network (GNN) trained on a broad set of ~49k Lewis acids performed significantly worse when transferred to the same scaffold (mean absolute error ≈ 23 kJ·mol^{-1}, $R^2 \approx 0.51$), and extrapolation between different scaffolds remained highly structure-dependent despite careful feature selection [3]. The interpretability of deep or black-box models poses another hurdle, where atomistic ML architectures like GNNs provide predictions without chemical rationale, complicating validation for chemists who must discern whether models rely on meaningful electronic trends or spurious correlations, ultimately eroding trust and hindering contributions to mechanistic insights.

The risk of unphysical or chemically incorrect predictions looms large, as pure ML models may violate fundamental rules like orbital ordering, resonance constraints, charge conservation, or thermodynamic trends, leading to erroneous forecasts of properties, such as Lewis acidity or hydrogen bond strength, especially in edge cases where small electronic-level errors cascade into qualitatively flawed conclusions. One way to mitigate such failures is to base descriptors directly on well-defined reactivity indices: for instance, projecting the Fukui function onto the Connolly surface and augmenting it with the work function has been shown to correlate strongly with mapped CO adsorption energies across multiple Ni/NiGa facets ($|R| \approx 0.97$), providing a fast, physically grounded surrogate for expensive DFT adsorption maps in catalyst screening [4], [5]. Sensitivity to computational methods adds variability, with labels depending on basis-set choices, solvation models, and dispersion corrections, introducing noise that undermines training data quality and model reliability. Without uncertainty quantification (UQ), most models deliver point predictions devoid of confidence estimates, leaving researchers unable to gauge reliability in extrapolative scenarios and risking misguided inferences in data-sparse domains. Nonlinear, multi-factor reactivity proves elusive for ML to capture, stemming from intertwined effects like orbital energies, substituent resonance, sterics, charge transfer, and conformational dynamics that exhibit high nonlinearity and non-additivity, demanding datasets far larger than those typically available in chemistry. At the same time, properly chosen electronic descriptors can dramatically reduce complexity: a recent study showed that using only eight natural bond orbital stabilization energies

 doi:10.20944/preprints202602.1398.v1

E(2), capturing key donor–acceptor and charge-transfer interactions in hydrogen-bonded complexes, is sufficient to train ML models that predict hydrogen-bond basicity on the pK_{BHX} scale with $R^2 \approx 0.92$ and errors below 0.4 kcal·mol^{-1}, outperforming previous approaches that relied on large, heterogeneous descriptor sets [6]. Heterogeneous, inconsistent, or noisy datasets compound these woes, blending disparate computational methods, solvation environments, levels of theory, reaction conditions, or Simplified Molecular Input Line Entry System (SMILES) conventions, necessitating removal of problematic entries and eroding reproducibility. Finally, limited scalability of atomistic deep learning models, such as GNNs, Transformers, and diffusion architectures, demands vast resources unavailable in molecular chemistry, where scaling to larger molecules inflates computational demands disproportionately, creating a chasm between theoretical potential and practical deployment [7], [8]. Decades of work in materials informatics and machine learning—from early applications of supervised/unsupervised learning, active learning, and kernel methods to neural-network interatomic potentials—have demonstrated substantial acceleration of prediction and screening, but have largely focused on pattern extraction rather than explicit scientific reasoning [9]. These challenges underscore a strategic problem: the absence of a truly "reasoning" approach in current AI systems, which often prioritize pattern matching over inventive problem-solving, leading to superficial or erroneous outputs in intricate scientific contexts [10], [11], [12].

The Theory of Inventive Problem Solving (TRIZ), developed by Genrich Altshuller in the 1940s [13], offers a structured antidote. TRIZ reframes problems as contradictions and resolves them using 40 inventive principles, patterns of evolution, and tools like the contradiction matrix. Originally for mechanical systems, TRIZ has been adapted for chemical engineering, as seen in [14] for sustainable processes and [15] for unit operations like mixing, introducing domain-specific characteristics (e.g., environmental impact) and principles (e.g., change process conditions). With its long history of evolution through practical examples and real-world tasks, TRIZ provides well-tuned algorithms and patterns for identifying contradictions—such as improving accuracy while worsening data availability—and resolving them systematically, making it exceptionally suited to address the "thinking" deficit in AI for materials science. By fostering inventive pathways over rote computation, TRIZ can guide AI toward more robust, creative solutions, drawing from decades of validated applications in diverse fields like mineral processing (where contradictions like high yield versus emissions are resolved through reductive pathways) and process optimization (e.g., adapting matrices for mixing to balance efficiency and energy use). The recent systematic review “Tools of Theory of Inventive Problem Solving Used for Process Improvement” [16] highlights a marked rise in scholarly interest in TRIZ, particularly after 2020. The authors report a substantial growth in peer-reviewed publications: 29–48 TRIZ-focused process-improvement papers identified in Web of Science and 193 in Scopus for 2020–2023, followed by rigorous filtering to extract high-quality case-based studies. The review further shows that TRIZ is increasingly applied not as a standalone method but in combination with Lean Six Sigma, function analysis, scientific effects, trends of engineering system evolution, and Algorithm of Inventive Problem Solving (ARIZ), reflecting its expanding role in structured innovation. Importantly, the most frequently used tools are the technical contradiction matrix with inventive principles, indicating continued reliance on TRIZ’s core mechanism for resolving system-level trade-offs. This documented growth underscores a broader trend: the need for systematic inventive reasoning frameworks in complex technical domains. In this context, integrating TRIZ principles with large language models provides a natural next step—embedding contradiction-driven, evolution-based reasoning within LLM workflows to enhance scientific creativity and reliability in molecular-scale materials research.

Current trends in AI for atomistic materials science are evolving to mitigate these issues, with physics-informed ML integrating quantum-aware frameworks to ground data-driven predictions in physical laws, enhancing accuracy in sparse-data scenarios. Generative models like variational autoencoders (VAEs) and diffusion architectures are advancing inverse design, proposing novel molecules with targeted properties while navigating chemical validity [17], [18] . Explainable AI is gaining momentum through techniques like SHapley Additive exPlanations (SHAP) for feature attribution, transforming black-box models into sources of mechanistic insight [19]. Active learning and autonomous discovery loops are optimizing data selection, iteratively refining models with

 doi:10.20944/preprints202602.1398.v1

quantum simulations to bridge low-data gaps and enabling increasingly closed-loop, self-driving experimental workflows in materials science [20]. At larger scales, graph neural networks trained in active-learning regimes, such as the GNU Network Object Model Environment (GNoME) framework, have demonstrated order-of-magnitude expansions of the known inorganic crystal space and emergent out-of-distribution generalization for stability prediction and interatomic potentials [21]. Transfer learning strategies are proving essential, leveraging pre-trained models on large corpora to fine-tune for specialized tasks, amplifying efficiency across molecular systems [22]. These trends align seamlessly with TRIZ's emphasis on ideality and contradiction resolution, offering opportunities for AI augmentation and particularly for generative-AI–driven inverse design pipelines in materials science [23].

This paper proposes integrating TRIZ with LLMs to guide reasoning models toward researcher-like innovation in molecular/atomistic materials science. By prompting LLMs to follow TRIZ workflows—first detecting patterns, then resolving contradictions—we enhance their ability to tackle domain challenges and align with trends like generative molecular design. We build on recent advancements, such as AutoTRIZ [24], which automates TRIZ ideation, and extend it to atomistic contexts, including case studies in smart materials.

2. Background on TRIZ and Its Applications

Pioneered in the 1940s by Genrich Altshuller, TRIZ emerged from an exhaustive analysis of hundreds of thousands of patents, revealing that innovative solutions follow predictable patterns across diverse technical fields [25]. TRIZ asserts that technical evolution adheres to repeatable laws, problems and their resolutions recur across industries, and creativity can be systematically taught and applied, democratizing invention beyond innate talent or chance. Central to TRIZ are concepts like ideality—striving for systems that maximize benefits while minimizing harms and costs, ideally approaching an "ideal final result" (IFR) where functions are performed without additional resources; functionality, ensuring every system component contributes meaningfully; and contradictions, where conflicts arise from improving one parameter at the expense of another, categorized as technical (e.g., enhancing strength reduces flexibility) or physical (e.g., a component must be both hot and cold). Its toolkit, including 40 inventive principles (e.g., segmentation or parameter changes), eight patterns of evolution (e.g., increasing dynamism or transition to micro-levels), and the contradiction matrix—a 39×39 grid mapping conflicting parameters to recommended principles—enables users to abstract specific problems into generic forms, draw from cross-domain solutions, and concretize them, bypassing psychological inertia and trial and- error.

In chemical engineering, TRIZ has been effectively adapted to address the unique complexities of process industries, where mechanical-focused origins are extended to encompass chemical and physicochemical phenomena. For instance, [15] tailored the contradiction matrix for chemical process industries (CPI) by analyzing over a decade of literature on unit operations like mixing, introducing 14 new characteristics (e.g., complexity, environmental impact, duration of operation) and eight inventive principles (e.g., change process conditions, convert harmful effects into benefits) to resolve conflicts in design and operation, such as balancing homogeneity with energy consumption in mixing or improving product quality in crystallization. Similarly, [14] demonstrated TRIZ's utility in sustainable chemical engineering through industrial case studies [26]: applying the matrix to mineral metal carbonate processing resolved the contradiction of high yield versus CO_2 emissions by proposing a reductive pathway, while for biobased effluents, it suggested reactive separations to valorize carboxylic acids, emphasizing TRIZ's role in innovation roadmaps that prioritize system analysis and practical ideation tools like ARIZ for rupture innovations. These adaptations highlight TRIZ's versatility, evolving from patent-based mechanical principles to domain-specific applications that enhance efficiency, safety, and sustainability in CPI.

This structured, algorithmic nature of TRIZ makes it particularly well-suited for integration with LLMs to enhance reasoning, addressing LLMs' tendencies toward shallow or biased outputs by imposing rigorous, inventive workflows. By prompting LLMs to detect evolution patterns (e.g., uneven system development in data-scarce AI models) and resolve contradictions via the matrix (e.g., accuracy versus computational cost in quantum simulations), TRIZ guides models toward

 doi:10.20944/preprints202602.1398.v1

researcher-like thinking, systematically exploring creative paths while ensuring solutions align with ideality. Recent AI-TRIZ hybrids, such as AutoTRIZ (which automates contradiction identification and principle application), PaTRIZ (for patent analysis using TRIZ) [27], and TRIZ-GPT (augmenting LLMs for problem abstraction), exemplify this synergy, outperforming manual methods in ideation speed and novelty, and paving the way for LLMs to tackle atomistic materials science challenges with the precision of TRIZ's proven algorithms.

ARIZ: Algorithm for Inventive Problem Solving in TRIZ

The ARIZ is a core analytical tool within TRIZ, developed by Genrich Altshuller and his collaborators starting in the 1940s. It represents a structured, multi-step methodology designed to tackle complex, non-routine inventive problems by systematically intensifying conflicts, mobilizing resources, and applying TRIZ principles to achieve breakthrough solutions without relying on trial and error (Figure 1) or psychological inertia. ARIZ evolved through several versions (e.g., ARIZ-56, ARIZ-85c), with ARIZ-85c being a prominent iteration that organizes the process into 9 parts, encompassing around 40 steps. It integrates seamlessly with other TRIZ tools, such as the contradiction matrix (for mapping technical contradictions to inventive principles), the 40 inventive principles (e.g., segmentation or parameter changes), and patterns of evolution (e.g., transition to micro-levels or increasing dynamism). The goal is to guide users toward the IFR—a solution where desired functions are achieved with minimal resources, no trade-offs, and elimination of harms—by reframing problems as contradictions and resolving them through physical, technical, or systemic transformations. ARIZ is particularly effective in engineering and technical contexts, as seen in adaptations for chemical processes (e.g., resolving yield vs. emissions in mineral processing), but its logical, algorithmic flow makes it adaptable to modern applications like AI reasoning [28].

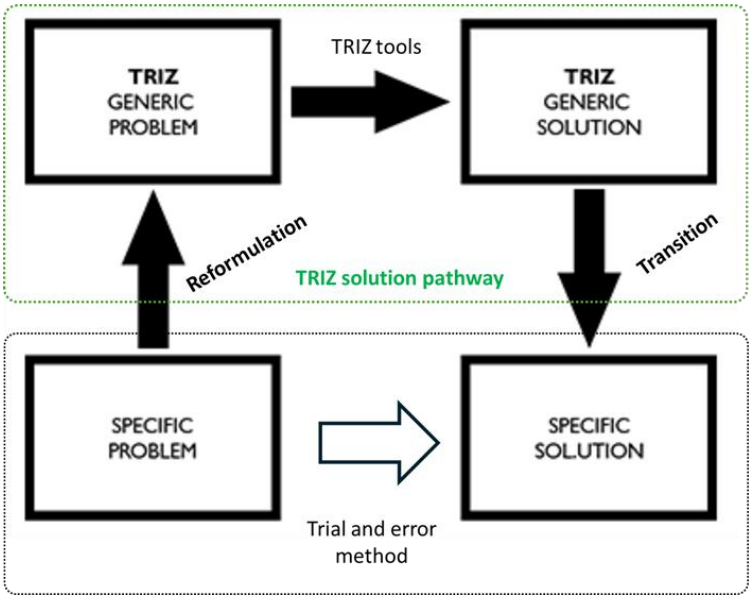

Figure 1. TRIZ solution pathway contrasted with the trial-and-error method (adapted from [40]).

Structure and Key Concepts of ARIZ

ARIZ operates as a sequential algorithm that progresses from vague problem statements to verified solutions, emphasizing the identification and resolution of contradictions: technical contradictions (TC, where improving one system parameter worsens another) and physical contradictions (PhC, where a single element must satisfy opposing states, like being present and absent). Resources—substance-field resources (SFR), such as materials, fields (e.g., mechanical, thermal), space, time, or environmental elements—are mobilized minimally to resolve these without introducing new complexities. The process draws on TRIZ's knowledge base, including standards for problem-analogues, physical effects, and evolution patterns, to ensure solutions are inventive rather than compromises. For example, in engineering, ARIZ might resolve a contradiction in a telescope antenna (protection from lightning vs. signal interference) by using air ionization as a resource, applying principles like phase transitions.

 doi:10.20944/preprints202602.1398.v1

The 9 parts of ARIZ-85c are as follows, with summarized steps and purposes:

- Part 1: Analyzing the Problem – Purpose: Formulate a mini-problem model by intensifying the conflict and eliminating jargon. Steps include defining the conflicting elements (product and tool), building graphic models for TCs using a table of contradiction types (e.g., counteraction), selecting and intensifying a model, and checking inventive standards. This integrates the contradiction matrix early to abstract the issue.

- Part 2: Analyzing the Problem Model – Purpose: Identify operational zones (OZ, conflict space), times (OT, conflict duration), and SFRs (internal/external resources). Steps focus on listing and prioritizing resources, setting the stage for resolution without external additions.

- Part 3: Defining the IFR and PhC – Purpose: Shift to an ideal solution and pinpoint PhC. Steps involve formulating IFR-1 (X-element resolves harm while preserving benefits), intensifying it (using only SFRs), identifying macro/micro PhC (opposing states), and applying standards. This part bridges TC to PhC, often using evolution patterns like separation in time/space.

- Part 4: Mobilizing and Using SFR – Purpose: Derive new resources creatively. Steps include "little creatures" simulation (modeling with particles), stepping back from IFR, combining substances, using voids, derived resources, fields, or field-sensitive substances. Examples: Rarefied air for ionization resolves conductivity contradictions.

- Part 5: Applying the Knowledge Base – Purpose: Use TRIZ's full arsenal for stubborn problems. Steps apply standard solutions, problem-analogues, principles for PhC elimination (via a table of separation methods), and pointers to physical effects.

- Part 6: Changing or Substituting the Problem – Purpose: Reformulate if stuck, removing inertia. Steps transition to technical solutions, combine/decompose problems, change the mini-problem, or shift to super-systems, leveraging evolution patterns.

- Part 7: Analyzing the Method of Resolving the PhC– Purpose: Verify solution quality. Steps check for IFR alignment, patent novelty, and sub-problems.

- Part 8: Applying the Obtained Solution – Purpose: Extend beyond the original issue. Steps estimate super-system changes, find new applications, and generalize using morphology or opposites.

- Part 9: Analyzing the Problem-Solving Process – Purpose: Reflect for learning. Steps compare the process to standards and document new insights.

This structure ensures ARIZ is iterative and adaptive, with checkpoints to apply TRIZ tools like the matrix for TCs or standards for PhCs.

3. TRIZ Perspective on Modern Reasoning AI Models

Reasoning models are AI systems explicitly optimized to perform structured, multi-step inference rather than just pattern-based prediction. In contrast to “next-token” language models that primarily learn statistical correlations, reasoning models are designed to decompose problems into sub-tasks, maintain intermediate representations, and follow explicit procedures (e.g., plans, graphs, or symbolic constraints) to reach conclusions. This architectural shift is visible in recent work on tool-augmented LLMs, program-synthesis–based reasoning, and neuro-symbolic systems, where a model not only generates text but also calls external tools, executes code, or manipulates symbolic structures as part of its internal reasoning loop [29], [30], [31].

A core feature of modern reasoning models is the integration of deliberate inference processes—such as chain-of-thought, self-consistency, and planning—with external environments. For example, models can be trained or prompted to generate explicit reasoning traces, explore multiple candidate reasoning paths, and select among them based on verification signals [32], [33]. When combined with tool use (code interpreters, search engines, simulators), this enables them to solve significantly harder tasks in mathematics, scientific question answering, and decision making than standard LLMs of similar size [34], [35]. In scientific domains, this means a reasoning model can iteratively propose a hypothesis, call a computational chemistry package, inspect the result, and refine the hypothesis in a closed loop—essentially emulating parts of the human scientific method.

This makes reasoning models a natural host for TRIZ algorithms in materials science at the molecular scale. TRIZ already provides a highly structured reasoning pipeline: identify and formalize

 doi:10.20944/preprints202602.1398.v1

the core contradiction, abstract it into TRIZ parameters, consult contradiction matrices or patterns of evolution, generate candidate solution concepts, and then refine them under real-world constraints. Embedding this pipeline inside a reasoning model allows the system to: (i) treat contradiction identification as an explicit sub-task; (ii) use TRIZ tools as "internal oracles" for proposing inventive directions; and (iii) couple those directions to domain-specific tools such as quantum-chemical solvers or molecular dynamics engines. For example, given a conflict like "increase ionic conductivity without sacrificing mechanical integrity in a polymer electrolyte", a TRIZ-guided reasoning model can decompose the problem, query a contradiction matrix to obtain abstract principles (e.g., segregation, composite materials, phase transitions), instantiate these as concrete molecular design moves, simulate their properties, and iteratively converge toward viable chemistries—thereby combining human-crafted inventive heuristics with machine-scale search and verification.

Adapting ARIZ for AI Reasoning

The algorithmic rigor of ARIZ makes it ideal for enhancing AI reasoning, particularly in large language models (LLMs) like GPT-4 or Grok, where CoT prompting can emulate its steps to foster inventive, structured thinking over hallucinatory or biased outputs. In AI contexts, ARIZ can be implemented as a prompted workflow (Figure 2): Start with Part 1 to model vague queries (e.g., "improve ML accuracy in data-scarce materials science") as mini-problems with TCs (accuracy vs. data availability); use Parts 2-3 to define IFR (e.g., self-generating data without cost) and PhC (model must be data-hungry yet data-independent); apply Parts 4-5 to mobilize "resources" like existing datasets or physics priors, integrating TRIZ tools via embedded knowledge (e.g., prompt LLM to simulate contradiction matrix for principles like "parameter changes" to suggest active learning). Parts 6-9 ensure verification and generalization, e.g., reformulating for edge cases or analyzing the reasoning process for self-improvement.

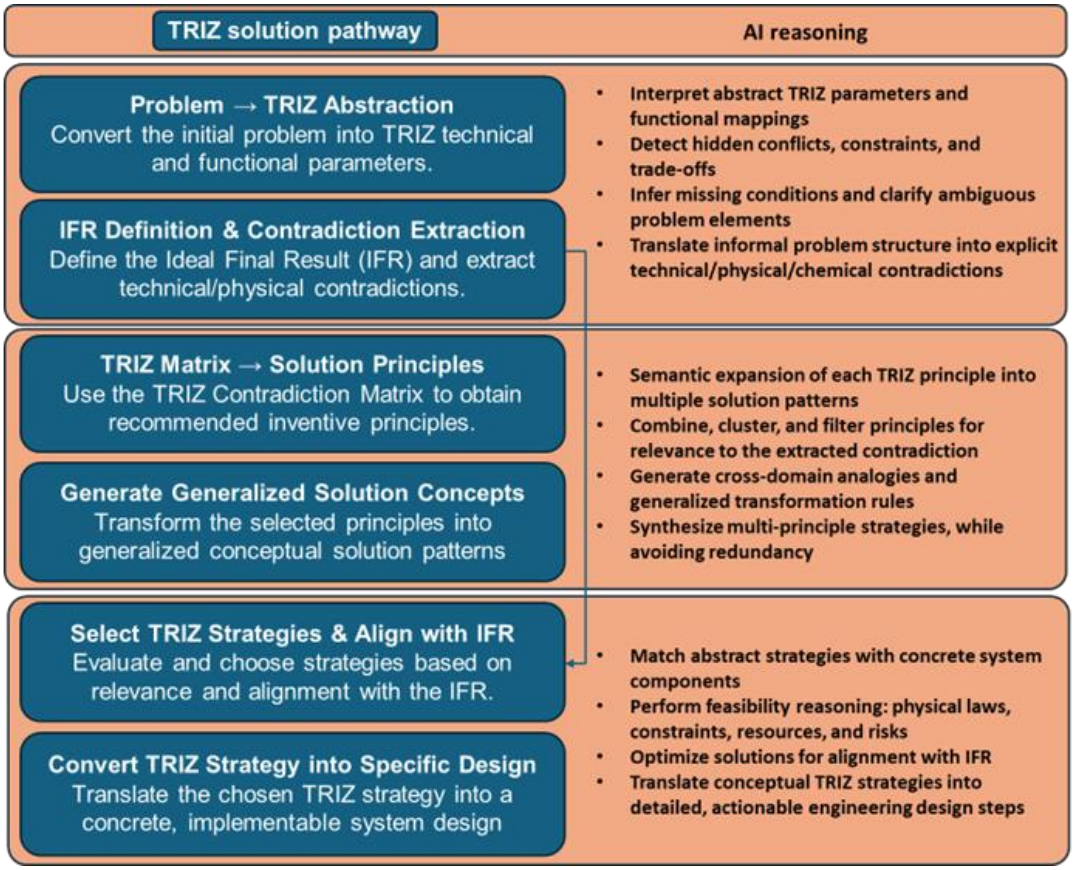

Figure 2. AI reasoning + TRIZ for material science applications.

For example, in AI-driven materials science, ARIZ could guide reasoning on unphysical predictions: Formulate TC (flexibility vs. correctness), intensify to PhC (model must violate/ obey physics), resolve via resources (embed priors as "little creatures" simulating particles), and verify IFR

 doi:10.20944/preprints202602.1398.v1

(physics-enforced predictions without added complexity). Modern adaptations, like AutoTRIZ or TRIZ-GPT, automate this in LLMs, outperforming humans in ideation speed. This turns AI into "researcher-like" agents, systematically resolving contradictions in tasks like generative design or uncertainty quantification, aligning with TRIZ's evolution patterns for AI systems.

4. Combining TRIZ with Reasoning Models: Core Idea

The core idea of combining TRIZ with reasoning models, such as LLMs equipped with CoT prompting, revolves around leveraging TRIZ's systematic, algorithmic framework to elevate AI from pattern-matching tools to inventive, researcher-like agents. Recent studies in chemical text mining demonstrate that fine-tuned LLMs can reliably extract structured, domain-specific information from complex scientific prose, substantially reducing prompt-engineering overhead and serving as flexible front-ends for automated data curation [36], which makes them natural hosts for TRIZ-guided reasoning workflows. TRIZ, with its emphasis on identifying contradictions, mobilizing resources, and applying principles to achieve ideality, provides a scaffold that structures LLM reasoning, prompting models to first detect systemic patterns (e.g., uneven evolution in data-scarce AI systems) and then resolve them through tools like the contradiction matrix. This integration addresses LLMs' common pitfalls—such as hallucinations, bias, and repetitive outputs—by enforcing logical workflows that mimic human inventive rigor, as seen in adaptations for chemical engineering where TRIZ resolves domain-specific conflicts like high yield versus environmental harm [14] or efficiency versus energy in mixing operations [15]. By embedding TRIZ steps into prompts, reasoning models can emulate a researcher's process: abstracting problems, framing technical contradictions (e.g., improving model accuracy worsens data availability), and generating solutions via principles like parameter changes or segmentation, ultimately accelerating innovation in atomistic materials science.

At its heart, this combination harnesses TRIZ's flexibility, demonstrated in chemical process industries through extensions like new characteristics (e.g., "duration of operation") and principles (e.g., "change process conditions") tailored for unit operations [15] , to make AI reasoning more adaptive and domain-aware [37]. For instance, in materials science challenges like weak extrapolation between scaffolds or poor interpretability of GNNs, TRIZ-guided prompts instruct models to scan for evolution patterns (e.g., "transition to micro-level" for molecular representations) and simulate matrix resolutions, producing non-hallucinatory outputs aligned with ideality—maximizing predictive power while minimizing computational harms. This mirrors how [14] used the matrix for sustainable innovations, such as reductive pathways in mineral processing, and extends it to AI: prompting LLMs to output structured responses like "Detected pattern: Uneven development; Contradiction: Fidelity vs. availability; Principles: 35 (physics-informed priors), 40 (hybrid quantum mechanics / machine learning (QM/ML))." The result is AI that "thinks like a researcher," overcoming psychological inertia (the tendency toward familiar solutions) and fostering breakthroughs in trends like physics-informed ML or generative design.

To make this integration more concrete, the reasoning model is first instructed to define the contradiction by using TRIZ's eight patterns of evolution (e.g., systems becoming more dynamic or segmented), which helps analyze a query for underlying issues—such as recognizing "weak extrapolation between scaffolds" as a "transition to micro-level" evolution in molecular materials science, mirroring how researchers spot systemic flaws. On the next step, the model solves it with the use of the TRIZ matrix, mapping the contradiction to inventive principles (e.g., local quality or intermediary), while involving the TRIZ concept of the ideal solution (IFR)—a state where the problem is resolved without trade-offs, using only existing resources to maximize benefits and eliminate harms. This two-step process fits seamlessly into the idea of reasoning models, as CoT prompting naturally supports sequential logic: the initial pattern detection builds contextual understanding to avoid shallow outputs, while matrix resolution and IFR evaluation enforce inventive, verifiable steps, reducing hallucinations and enhancing domain-specific rigor, much like AutoTRIZ automates TRIZ in LLMs for faster ideation (Jiang et al., 2024).

Step 1: Identify the Problem via Patterns. TRIZ's eight patterns of evolution (e.g., systems become more dynamic or segmented) can prompt the model to analyze a query for underlying issues. For instance, in materials science at the molecular level (as in our previous discussions), a model could

 doi:10.20944/preprints202602.1398.v1

scan for patterns like "weak extrapolation between scaffolds" (a challenge from earlier) by recognizing it as a "transition to micro-level" evolution. This mirrors how researchers spot systemic flaws before diving in—LLMs can be prompted to output: "Detected pattern: Increasing ideality requires resolving data scarcity vs. model accuracy."

Step 2: Resolve via Contradiction Matrix. Once identified, frame the problem as a technical contradiction (improving one parameter worsens another), then apply the matrix to suggest principles. The model could simulate this: Map features to matrix axes (e.g., "data fidelity" improving but "computational cost" worsening), retrieve principles (e.g., "segmentation" for modular data augmentation), and generate solutions. This enforces logical, non-hallucinatory reasoning, much like a researcher's hypothesis-testing.

Evidence from emerging AI-TRIZ hybrids underscores the feasibility and promise of this idea. Tools like AutoTRIZ (Jiang et al., 2024) automate TRIZ within LLMs, identifying contradictions and applying principles faster than human experts, while multi-agent systems [38] distribute tasks like pattern detection across agents for collaborative reasoning. In atomistic AI, this could resolve unphysical predictions by framing them as contradictions (flexibility vs. correctness) and suggesting intermediaries like SHAP layers, enhancing explainability without sacrificing depth. Domain adaptations, as in Pokhrel et al.'s chemical extensions, suggest creating AI-specific matrices (e.g., "prompt complexity" vs. "response coherence") to further refine reasoning, turning models into self-improving entities that evolve toward "ideal" AI researchers through TRIZ loops.

However, the integration's strength lies in its strategic fit for materials science's challenges, where TRIZ's long-validated algorithms—honed on practical tasks like process optimization—provide the "thinking" approach AI lacks. Limitations, such as potential misapplication of abstract concepts without fine-tuning, can be mitigated by iterative prompting, ensuring outputs remain logical and inventive. Overall, this combination democratizes research-like thinking, accelerating discoveries in fields like molecular design by augmenting AI with TRIZ's inventive core, much like its successful applications in sustainable engineering.

A critical practical consideration when deploying TRIZ-guided reasoning models in atomistic materials science concerns the trade-off between local and cloud-based LLMs—a decision that itself embodies a TRIZ-style technical contradiction. Cloud-based platforms like GPT-4, DeepSeek, or Grok offer substantially larger parameter counts and more extensive pre-training, potentially delivering superior reasoning depth for complex contradiction resolution and principle application. However, materials science research frequently operates under stringent privacy, security, and confidentiality constraints, particularly in industrial settings involving proprietary molecular designs, unpublished quantum-chemical data, or patent-sensitive synthetic routes where data leakage could compromise competitive advantage or violate intellectual property agreements. Cloud-based LLMs, despite their power, introduce vulnerabilities to security breaches and lack guarantees of data sovereignty, making them unsuitable for defense-related catalysis research, pharmaceutical discovery, or advanced materials development in closed corporate environments.

Moreover, cloud systems operating under high load with millions of concurrent requests often optimize computational costs by limiting inference time—employing aggressive early-stopping heuristics, truncated reasoning chains, or reduced beam search strategies to minimize response latency and service costs, which can degrade the quality of solving complex scientific tasks that require deep multi-step contradiction analysis and careful evaluation of inventive principles. Local LLMs—deployable on institutional servers or even air-gapped workstations—offer complete control over data flows and can operate without internet connectivity, aligning with security protocols in sensitive domains. Critically, local deployments allow researchers to allocate substantially more compute per query—employing extended chain-of-thought reasoning, exhaustive sampling of solution candidates, or iterative refinement loops—without cost penalties or time constraints imposed by commercial API rate limits. While this increases wall-clock time for individual queries, such delays are often negligible in research contexts where a single well-reasoned molecular design proposal, generated over minutes rather than seconds, can guide weeks of experimental validation, making thoroughness far more valuable than speed.

 doi:10.20944/preprints202602.1398.v1

While historically constrained by smaller model sizes, recent advances in model compression techniques (quantization, pruning), domain-specific fine-tuning on chemistry corpora, and retrieval-augmented generation suggest that local models can achieve performance levels sufficient for TRIZ workflows in molecular design, where structured reasoning, domain knowledge retrieval, and adherence to inventive principles often matter more than raw parameter scale. This mirrors TRIZ's own emphasis on resource optimization and ideality: achieving inventive solutions—here, secure, privacy-preserving AI reasoning—without excessive infrastructure or compromising core functional requirements, effectively resolving the "accuracy versus security" contradiction through segmentation (Principle 1) and local quality (Principle 3) by tailoring models to specific, controlled environments, while also applying "turn harm into benefit" (Principle 22), transforming the apparent disadvantage of slower computation into the advantage of more thorough analysis.

5. Integrating TRIZ with Reasoning Models

We propose a TRIZ-guided framework for LLMs:

- **Step 1: Pattern Detection**: Prompt LLMs to identify evolution patterns (e.g., "transition to micro-level" for scaffold extrapolation).
- **Step 2: Contradiction Framing**: Map challenges to matrix axes (e.g., "physical fidelity" improving worsens "computational simplicity").
- **Step 3: Principle Application**: Generate solutions using principles (e.g., Principle 40: Composites for hybrid embeddings).
- **Step 4: Ideality Evaluation**: Refine toward IFR, minimizing harms.

The TRIZ contradiction matrix, which systematically maps technical conflicts to inventive principles for generating potential solutions, can be readily adapted for solving chemical problems by incorporating domain-specific characteristics and principles tailored to physicochemical phenomena. As demonstrated by [15], the original matrix—designed primarily for mechanical systems—has been extended for chemical process industries through the addition of 14 new parameters (e.g., complexity, environmental impact, duration of operation) and eight inventive principles (e.g., change process conditions, convert harmful effects into benefits), enabling it to address challenges in unit operations like mixing, where contradictions such as homogeneity versus energy consumption are resolved more effectively. Similarly, [14] applies an adapted matrix to sustainable chemical engineering, resolving conflicts like high product yield versus CO_2 emissions in mineral processing by suggesting innovative pathways such as reductive calcination, illustrating how the matrix's flexibility allows it to evolve beyond its origins to provide practical, inventive solutions in chemistry.

Adapting TRIZ Tools for Reasoning Models in Molecular Materials Design

To make TRIZ really useful for reasoning-oriented LLMs in molecular and materials design, the classical tools need to be adapting further, made chemistry-aware, and exposed in a form that a model can call during its step-by-step reasoning. In our framework “Integrating TRIZ with LLMs: Enhancing Reasoning for Innovation in Materials Science at the Molecular Scale,” we treat the extended contradiction matrix and domain-specific principles not as fixed structures, but as flexible building blocks that can be scaled up or down depending on the task and the available data.

A practical starting point is a compact, chemistry-aware contradiction matrix that covers the trade-offs most relevant for atomistic and ML-based design. The 57×57 scheme described above is one possible, “maximal” version, but in many LLM applications a much smaller subset of parameters is sufficient. For example, a matrix fragment like the one in Table 1 can already guide the model toward reasonable solution patterns by linking typical contradictions (reactivity vs. stability, accuracy vs. cost, etc.) to a small set of inventive principles. This fragment should be viewed as an illustration of how to encode such knowledge in a machine-readable way, not as a rigid prescription for matrix size.

Table 1. Illustrative contradiction fragment for molecular/ML design.

Improve ↓ / Worsen →	Stability ↑	Atom economy ↑	Computational cost ↓

 doi:10.20944/preprints202602.1398.v1

Reactivity ↑	27, 41	35, 6	40, 48
Selectivity ↑	15, 44	27, 43	24, 46
Model accuracy ↑	40, 35	1, 48	15, 37

Numbers refer to classical TRIZ principles (1–40) and chemistry-specific ones (41–48). For instance, the entry for "increase model accuracy without increasing computational cost" suggests principles 15 (Dynamism) and 37 (Phase transition–like change), which an LLM can interpret as, e.g., adaptive-resolution simulations, coarse-to-fine model cascades, or active learning focused on the most informative regions of chemical space.

In parallel, a small set of chemistry-adapted inventive principles can be turned into reusable "reasoning operators" that the model applies once a contradiction is identified. Here again, the extended list of 48 principles is a rich starting point, but for everyday LLM use it is often enough to work with a curated subset that maps cleanly onto domain actions. Table 2 shows an example of such a mapping, where each principle is associated with simple, repeatable transformations that the model can concretize into molecular edits, architectural changes in ML models, or adjustments to simulation protocols.

Table 2. Example of chemistry-aware TRIZ principles as reasoning operators.

Principle	Chemistry-adapted meaning	Typical LLM-level action in molecular design
1. Segmentation	Molecular segmentation	Propose splitting a bulky ligand into modular fragments or separating functional blocks in a copolymer.
15. Dynamism	Conformational / environmental switching	Suggest stimuli-responsive groups, flexible linkers, or solvent/temperature-dependent conformations
24. Feedback	Autocatalysis / self-regulation	Introduce product-driven activation, catalyst resting states, or self-correcting reaction cycles
40. Composite	Hybrid QM/ML or multi-phase systems	Combine high-level QM on a reactive center with ML surrogates for the environment; design composite electrolytes.
41. Electronic separation	Spatial/temporal separation of electronic effects	Move reactive orbitals away from fragile motifs; decouple activity and stability in different domains
48. Scale-bridging	Linking molecular and mesoscale structure	Propose motifs that self-assemble into desired morphologies; connect atomistic features to bulk properties.

Within a TRIZ-guided prompt, a reasoning model can then follow a lightweight workflow: (1) restate the problem as a contradiction in terms of a chosen parameter set (e.g., rows and columns from Table 1); (2) retrieve 2–4 candidate principles from the corresponding matrix entries; (3) expand these principles via operators like those in Table 2 into concrete design moves (molecular changes, new descriptors, hybrid QM/ML schemes); and (4) plan how to test them using quantum-chemical calculations, molecular dynamics, or active-learning loops. The exact granularity of the matrix and the number of principles can be adjusted to the application: small, task-specific matrices for focused molecular design problems, or richer, multi-level matrices when a broader range of contradictions must be handled. This flexibility is key for integrating TRIZ into LLM-based reasoning without over-constraining either the human user or the model.

 doi:10.20944/preprints202602.1398.v1

6. Validation

To evaluate the effectiveness of the proposed method, we compared the computational cost for the selected set of chemical problems under identical conditions, using a thinking AI model with and without the incorporation of TRIZ-based algorithm. All calculations were performed using the free local model, qwen3:8b, obtained from the Ollama library. This relatively small model fits on a standard desktop GPU with 8 GB of video memory.

The TRIZ-guided reasoning process was implemented as a three-stage pipeline (a simplified version of the algorithm shown in Fig. 1) that combines language-model inference with deterministic TRIZ knowledge lookup (the corresponding Python script implementing this pipeline is provided in the Supporting Information.):

Mode A: Contradiction Identification (LLM-assisted). In the first stage, the language model was used exclusively to interpret the natural-language problem description and map it onto a formal TRIZ contradiction. Specifically, the model selected one improving parameter and one worsening parameter from a predefined list of TRIZ parameters. The output was constrained to a strict, structured format and did not involve any solution generation. This mode served solely as a semantic translation from the problem statement to a formal contradiction representation.

TRIZ Matrix Resolution (Deterministic, non-LLM). In the second stage, the selected improving and worsening parameters were used as inputs to the TRIZ contradiction matrix. The matrix lookup was performed deterministically, without any involvement of the language model. This step produced a finite set of inventive principles associated with the identified contradiction. Because the matrix is fixed, this stage is fully reproducible and independent of model behavior.

Mode B: Principle-Based Solution Synthesis (LLM-assisted). In the final stage, the language model was invoked again to synthesize solution concepts guided by the retrieved TRIZ principles. The model was instructed to explicitly address each principle, analyze the underlying physical or chemical mechanisms, and propose a concrete solution or indicate its limited applicability. No additional principles beyond those returned by the TRIZ matrix were permitted. This constraint ensured that solution generation was guided by formal TRIZ reasoning rather than unconstrained ideation.

The set of TRIZ parameters and principles was extended beyond the classical values with chemistry-specific additions, following the approaches proposed in the works [14], [15],[26], [39]. The additional chemical parameters included:

- Barrier height (ΔG‡)
- Driving force (ΔG, equilibrium)
- Mechanistic uncertainty
- Competing pathways/side reactions propensity
- Selectivity (chemo/regio/stereo)
- Binding/adsorption specificity
- Catalyst activity (TOF/TON proxy)
- Catalyst stability (deactivation/poisoning/leaching)
- Active site accessibility/confinement
- Solvent/matrix effect strength
- Transport limitation (mass/heat transfer limitation)
- Phase behavior complexity (multiphase/interfaces)
- Compute cost
- Model accuracy
- Model transferability (domain of applicability)
- Interpretability/causality
- Data quality and bias
- Reproducibility (protocol/software/lab)

The additional chemical TRIZ principles included:

- Microenvironment engineering
- Mechanism/pathway engineering
- Selectivity by recognition

 doi:10.20944/preprints202602.1398.v1

- Modular chemical representation
- Calibration & transfer (multi-fidelity, Δ learning)
- Uncertainty as a design variable (UQ/AL/BO)
- Controlled reactivity/gating
- Design for observability

Complete lists of the extended parameters, principles and TRIZ resolution matrix are provided in the Supporting Information as JSON files.

For comparison with the TRIZ-guided workflow, each chemical task was also processed using a baseline mode without any TRIZ constraints or structured reasoning. In this mode, the language model generated several chemically reasonable solutions directly from the task description/prompt. No explicit problem decomposition, parameter identification, or principle-based guidance was applied. This baseline reflects standard thinking LLM-based ideation and serves as a reference for evaluating the effect of TRIZ guidance.

Solution quality was assessed using a separate scoring procedure, with two quantitative scores: **novelty** and **feasibility** (relative from 0.0 to 1.0). Novelty measures how different a solution is from common or textbook approaches, while feasibility reflects the technical plausibility of implementing the solution under current conditions. Scoring was performed using a language model with low stochasticity to improve reproducibility.

Both novelty and feasibility scores were treated as heuristic, expert-like evaluations rather than precise or objective metrics. They enabled an initial, systematic comparison of solution sets and supported consistent evaluation across different tasks and generation modes. In this context, the scores are interpreted as relative indicators rather than absolute measurements.

Table 3. Computational Cost, Number of Solutions, and Relevance Assessment for a Set of Chemical Queries.

Task 1	Hydrogenation of alpha,beta-unsaturated aldehydes (e.g. crotonaldehyde). Lowering the barrier for 1,2-addition to C=O also promotes 1,4-addition to C=C.									
	tokens	solutions	novelty	feasibility	novelty	feasibility	novelty	feasibility	novelty	feasibility
TRIZ	283	4	0.8	0.9	0.7	0.7	0.9	0.7	0.6	0.8
baseline	401	4	0.8	0.9	0.7	0.9	0.6	0.8	0.5	0.9
Task 2	Methanol oxidation on Cu-based catalysts. High activity requires high temperature and oxidizing conditions, but Cu sinters and oxidizes to inactive CuO.									
	tokens	solutions	novelty	feasibility	novelty	feasibility	novelty	feasibility	novelty	feasibility
TRIZ	273	4	0.6	0.7	0.4	0.9	0.8	0.7	0.9	0.5
baseline	362	4	0.7	0.8	0.8	0.7	0.7	0.8	0.7	0.8
Task 3	Selective oxidation of toluene to benzaldehyde on zeolites. Small pores improve selectivity but limit diffusion and reduce conversion.									
	tokens	solutions	novelty	feasibility	novelty	feasibility				
TRIZ	225	2	0.8	0.5	0.5	0.7				
baseline	292	2	0.7	0.8	0.6	0.8				
Task 4	Prediction of activation energies for catalyst screening. DFT (PBE+D3) is accurate but too expensive to screen 10,000 candidates.									
	tokens	solutions	novelty	feasibility	novelty	feasibility	novelty	feasibility		
TRIZ	246	3	0.8	0.9	0.6	0.9	0.7	0.8		
baseline	322	3	0.6	0.9	0.7	0.9	0.5	0.8		

 doi:10.20944/preprints202602.1398.v1

Task 5	Propylene hydroformylation on Rh-phosphine catalysts. Open Rh sites increase TOF but reduce n-butanal versus iso-butanal selectivity.									
	tokens	solutions	novelty	feasibility	novelty	feasibility	novelty	feasibility		
TRIZ	251	3	0.7	0.9	0.6	0.7	0.9	0.6		
baseline	313	3	0.7	0.8	0.8	0.9	0.9	0.9		

The validation results are presented both in Table 3 and as detailed log files in the Supporting Materials. They show that, at comparable solution quality—evaluated in terms of novelty and feasibility—the TRIZ-guided workflow consistently outperforms the baseline approach, achieving an overall speedup of approximately 20–25%. This improvement is notable, as it is obtained without any observable loss in solution quality under the adopted evaluation criteria. These findings indicate that incorporating TRIZ-based reasoning into the generation pipeline provides a clear efficiency advantage.

At the same time, a detailed and fully objective comparison of qualitative differences between the generated solutions is challenging and would require a more extensive analysis. Such an analysis is beyond the scope of the present study and will be addressed in future work. Importantly, the current implementation represents only a partial realization of the proposed framework, and further improvements in both performance and solution quality are expected with implementation of the full workflow described above.

A key practical advantage of the proposed approach is that it avoids the need for costly and time-consuming model retraining, while enabling the use of relatively small local models. This feature is particularly attractive for industrial applications, where confidentiality, data control, and reproducibility are critical considerations. In addition, the framework naturally supports the use of different models for Mode A and Mode B, providing flexibility to optimize computational resources and offering a clear pathway for further performance improvements.

7. Conclusions

This paper treats TRIZ not as a separate add-on to AI, but as a way to organize how large language models think about problems in molecular and atomistic materials science. We argue that many of the field's key difficulties—too little high-quality data, weak transfer between chemical scaffolds, unphysical predictions, and "black-box" models—can be seen as contradictions in the TRIZ sense. Once they are framed this way, familiar tools such as better descriptors, physics-based constraints, active learning, generative design, and uncertainty estimates stop being isolated tricks and instead become resources that are used deliberately within a structured problem-solving process.

In concluding, it is important to distinguish between better asking and better problem solving. A well-crafted prompt can markedly improve an LLM's output by clarifying intent, constraints, and the desired format, but prompt engineering primarily optimizes the communication channel between the user and the model. By contrast, integrating TRIZ shifts the interaction from ad hoc "ask-and-answer" to a structured design workflow: the problem is first reformulated in terms of an Ideal Final Result (IFR), explicit technical or physical contradictions ("improving X worsens Y"), and an inventory of available resources (substance/field, space, time, information). This scaffold is not merely rhetorical; it provides reusable operators—TRIZ principles, standard solutions, and separation strategies—that systematically expand the search space while keeping it anchored to the core trade-off.

From this perspective, the central advantage of a TRIZ-integrated LLM is not that it "sounds smarter," but that it becomes more traceable and actionable. Ideas are generated with an explicit rationale (which contradiction they resolve and by what principle), enabling more reliable critique, comparison, and iteration. The same mechanism also promotes structured diversity: instead of sampling many variations of the same intuition, the model is guided to explore qualitatively different transformation routes suggested by different principles/standards. Finally, TRIZ naturally encourages verification-minded outputs—because contradictions translate into measurable trade-

 doi:10.20944/preprints202602.1398.v1

offs, proposed concepts can be paired with concrete checks (e.g., simulation, rule-based validators, or targeted experiments) and refined through feedback loops. In domains such as chemistry, biology, and materials science—where feasibility constraints, multi-objective optimization, and costly validation dominate—this combination of systematic ideation and grounded evaluation can make the difference between plausible text and a practical pathway to testable hypotheses.

Building on earlier work like AutoTRIZ, PaTRIZ, and multi-agent TRIZ systems, we show how TRIZ concepts can be built directly into prompts and workflows for large language models. In practical terms, this means asking the model first to identify the core contradiction, then to express it in TRIZ terms, choose suitable inventive principles, and only after that suggest concrete molecular or materials changes. This extra structure helps to cut down on unsupported or physically unreasonable answers and, at the same time, encourages more creative but still testable ideas—for example, hybrid quantum-mechanical and machine-learning schemes or new ways of representing molecules at the micro level.

Looking ahead, several directions seem promising. One is to develop contradiction matrices tailored specifically to AI and materials tasks, such as balancing model complexity and interpretability, or rich prompting and robustness. Another is to turn ARIZ-like procedures into reusable prompt patterns that can be combined with quantum-chemical and molecular-dynamics software in closed discovery loops. Carefully chosen case studies—such as shape-memory polymers with conflicting mechanical requirements or catalytic systems that must combine activity, selectivity, and stability—will be important to test this approach. If it works in practice, integrating TRIZ with large language models could make advanced inventive thinking more widely available in materials science and speed up progress toward more "ideal" materials and more reliable AI tools.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/doi/SI.zip.

Author Contributions: Conceptualization, S.G., S.S.; Methodology, S.G.; Validation and Applications, A.H. and J.J.V.; Writing—original draft, Review, Results Analysis, and Editing, all authors equally. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the National Research Council of Canada (NRC) through the CBMI project "Critical Battery Materials Initiative.".

Data Availability Statement: No new datasets were generated or analyzed during the current study.

Acknowledgments: Sergey Gusarov gratefully acknowledges Professor Randy Goebel (University of Alberta) and Dr. Alain Tchagang (NRC) for valuable discussions and support. Stanislav. R. Stoyanov acknowledges the support of the Government of Canada's Program of Energy Research and Development (PERD) and Defense Research and Development Canada.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Madika, B., et al., Artificial Intelligence for Materials Discovery, Development, and Optimization. ACS Nano, 2025. **19**(30): p. 27116-27158.
2. Gusarov, S. and S.R. Stoyanov, COSMO-RS-Based Descriptors for the Machine Learning-Enabled Screening of Nucleotide Analogue Drugs against SARS-CoV-2. J Phys Chem Lett, 2020. **11**(21): p. 9408-9414.
3. Fenogli, J., L. Grimaud, and R. Vuilleumier Constructing and explaining machine learning models for chemistry: example of the exploration and design of boron-based Lewis acids. 2025. arXiv:2501.01576 DOI: 10.48550/arXiv.2501.01576.
4. Gusarov, S., S.R. Stoyanov, and S. Siahrostami, *Development of Fukui Function Based Descriptors for a Machine Learning Study of CO2Reduction.* The Journal of Physical Chemistry C, 2020. **124**(18): p. 10079-10084.
5. Siahrostami, S., et al., 2 Artificial intelligence for catalysis, in Accelerated Materials Discovery. 2022. p. 27-64.
6. Melo, D.U., et al., Developing a Machine Learning Model for Hydrogen Bond Acceptance Based on Natural Bond Orbital Descriptors. J Org Chem, 2025. **90**(28): p. 9776-9788.

 doi:10.20944/preprints202602.1398.v1

7. Peivaste, I., et al., Artificial intelligence in materials science and engineering: Current landscape, key challenges, and future trajectories. Composite Structures, 2025. **372**.
8. Batool, M., O. Sanumi, and J. Jankovic, Application of artificial intelligence in the materials science, with a special focus on fuel cells and electrolyzers. Energy and AI, 2024. **18**.
9. Mueller, T., A.G. Kusne, and R. Ramprasad, Machine Learning in Materials Science, in Reviews in Computational Chemistry. 2016. p. 186-273.
10. Gomes, C.P., B. Selman, and J.M. Gregoire, *Artificial intelligence for materials discovery*. MRS Bulletin, 2019. **44**(7): p. 538-544.
11. Chávez-Angel, E., et al., Applied Artificial Intelligence in Materials Science and Material Design. Advanced Intelligent Systems, 2025. **7**(8).
12. Belski, I. and I. Belski, *Application of TRIZ in Improving the Creativity of Engineering Experts.* Procedia Engineering, 2015. **131**: p. 792-797.
13. Al'tshuller, G.S., *The innovation algorithm : TRIZ, systematic innovation and technical creativity*. 1999: Worcester, Mass. : Technical Innovation Center, Inc.
14. Lux, S., Application of the TRIZ Contradictory Matrix to Foster Innovation for Sustainable Chemical Engineering. Chemie Ingenieur Technik, 2022. **94**(8): p. 1071-1079.
15. Pokhrel, C., et al., Adaptation of TRIZ contradiction matrix for solving problems in process engineering. Chemical Engineering Research and Design, 2015. **103**: p. 3-10.
16. Sojka, V. and P. Lepsik, Tools of Theory of Inventive Problem Solving Used for Process Improvement—A Systematic Literature Review. Processes, 2025. **13**(1).
17. Denny Handoko, A. and R.I. Made *Artificial Intelligence and Generative Models for Materials Discovery -- A Review*. 2025. arXiv:2508.03278 DOI: 10.48550/arXiv.2508.03278.
18. Cai, J., et al., *Machine learning-driven new material discovery.* Nanoscale Adv, 2020. **2**(8): p. 3115-3130.
19. Oviedo, F., et al., Interpretable and Explainable Machine Learning for Materials Science and Chemistry. Accounts of Materials Research, 2022. **3**(6): p. 597-607.
20. Stein, H.S. and J.M. Gregoire, Progress and prospects for accelerating materials science with automated and autonomous workflows. Chem Sci, 2019. **10**(42): p. 9640-9649.
21. Merchant, A., et al., *Scaling deep learning for materials discovery.* Nature, 2023. **624**(7990): p. 80-85.
22. Goswami, L., M. Deka, and M. Roy Artificial Intelligence in Material Engineering: A review on applications of AI in Material Engineering. 2022. arXiv:2209.11234 DOI: 10.48550/arXiv.2209.11234.
23. Pradeep, *Generative AI in Materials Science: Accelerating Discovery through Inverse Design.* International Journal of Scientific Research and Engineering Development-, 2025. **8**(3): p. 545-550.
24. Jiang, S., et al. AutoTRIZ: Automating Engineering Innovation with TRIZ and Large Language Models. 2024. arXiv:2403.13002 DOI: 10.48550/arXiv.2403.13002.
25. Vaneker, T.H.J., R.G.J. Damgrave, and J.G. Kuster, *TRIZ as an Enabler for Intellectual Property Protection During Product Development*. Procedia Engineering, 2015. **131**: p. 731-736.
26. Abramov, O., et al., *TRIZ-based approach for accelerating innovation in chemical engineering.* Chemical Engineering Research and Design, 2015. **103**: p. 25-31.
27. Guarino, G., A. Samet, and D. Cavallucci, *PaTRIZ: A framework for mining TRIZ contradictions in patents.* Expert Systems with Applications, 2022. **207**.
28. Harlim, J. and I. Belski, *On the Effectiveness of TRIZ Tools for Problem Finding*. Procedia Engineering, 2015. **131**: p. 892-898.
29. Wei, J., et al. *Chain-of-Thought Prompting Elicits Reasoning in Large Language Models*. 2022. arXiv:2201.11903 DOI: 10.48550/arXiv.2201.11903.
30. Gao, L., et al. *PAL: Program-aided Language Models*. 2022. arXiv:2211.10435 DOI: 10.48550/arXiv.2211.10435.
31. Creswell, A., M. Shanahan, and I. Higgins *Selection-Inference: Exploiting Large Language Models for Interpretable Logical Reasoning*. 2022. arXiv:2205.09712 DOI: 10.48550/arXiv.2205.09712.
32. Wang, L., et al. Plan-and-Solve Prompting: Improving Zero-Shot Chain-of-Thought Reasoning by Large Language Models. 2023. arXiv:2305.04091 DOI: 10.48550/arXiv.2305.04091.
33. Kojima, T., et al. *Large Language Models are Zero-Shot Reasoners*. 2022. arXiv:2205.11916 DOI: 10.48550/arXiv.2205.11916.

 doi:10.20944/preprints202602.1398.v1

34. Srivastava, A., et al. Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models. 2022. arXiv:2206.04615 DOI: 10.48550/arXiv.2206.04615.
35. Schick, T., et al. *Toolformer: Language Models Can Teach Themselves to Use Tools*. 2023. arXiv:2302.04761 DOI: 10.48550/arXiv.2302.04761.
36. Zhang, W., et al., *Fine-tuning large language models for chemical text mining.* Chem Sci, 2024. **15**(27): p. 10600-10611.
37. Wu, Y., F. Zhou, and J. Kong, Innovative design approach for product design based on TRIZ, AD, fuzzy and Grey relational analysis. Computers & Industrial Engineering, 2020. **140**.
38. Szczepanik, K. and J.A. Chudziak *TRIZ Agents: A Multi-Agent LLM Approach for TRIZ-Based Innovation*. 2025. arXiv:2506.18783 DOI: 10.48550/arXiv.2506.18783.
39. Kim, J., et al., *Application of TRIZ creativity intensification approach to chemical process safety.* Journal of Loss Prevention in the Process Industries, 2009. **22**(6): p. 1039-1043.
40. Cheng, F.M., et al., Product concept design method for improving user satisfaction based on case-based reasoning. Sci Rep, 2025.