AI as an Anti-Entropy Engine: Actively Designing Intelligent Matter from Dynamic States to Proto-Life

Dengchen Yang[1]*

[1] Institute of Microbiology, Chinese Academy of Sciences, Beijing, China.

*Correspondence to: yangdc@im.ac.cn

Abstract

The trial-and-error paradigm of traditional materials discovery, fundamentally constrained by its inherent high entropy, is proving inadequate for designing complex intelligent matter. Here, we propose a new scientific paradigm: **Artificial Intelligence as an 'Anti-Entropy' Engine**, transforming research from passive understanding to active design. By systematically injecting informational negative entropy across perception, planning, and execution loops, AI guides material systems from disorder to pre-defined functional order. We demonstrate this through empirical advances—such as the **GNoME model discovering 2.2 million stable crystals**—and construct a unified **'Perception-Planning-Execution'** framework enabling inverse design across scales. This paradigm extends beyond static structures to dynamic non-equilibrium systems and life-like chemical networks. We prospectively map future frontiers using a **'Ladder of Intelligence'** and address ethical governance, systemic risk, and sustainability. Ultimately, this marks a fundamental transition for humanity, from being passive observers of nature to becoming active **'anti-entropy' designers** in the evolution of matter. This review not only synthesizes these advances but also provides a unifying conceptual framework and a clear roadmap for the field, aiming to catalyze the transition towards this **fifth paradigm of scientific discovery**.

Keywords: Anti-entropy; AI-Driven Design; Intelligent Matter; Inverse Design; Autonomous Laboratory; Life-like Systems; Interdisciplinary Paradigm

The core conceptual framework of this review is illustrated in Figure 1, depicting a complete evolution path for the scientific research paradigm: starting from the limitations of the traditional paradigm, supported by theoretical and technological foundations, proceeding through the

'Perception-Planning-Execution' loop of the AI anti-entropy engine, and ultimately leading to the future vision spanning from dynamic intelligent matter to chemical proto-life. This framework not only summarizes the article's logic but also highlights the agency and revolutionary nature of the core 'AI as an Anti-Entropy Engine' paradigm—actively guiding material systems from disorder to pre-defined functional order by injecting informational negative entropy.
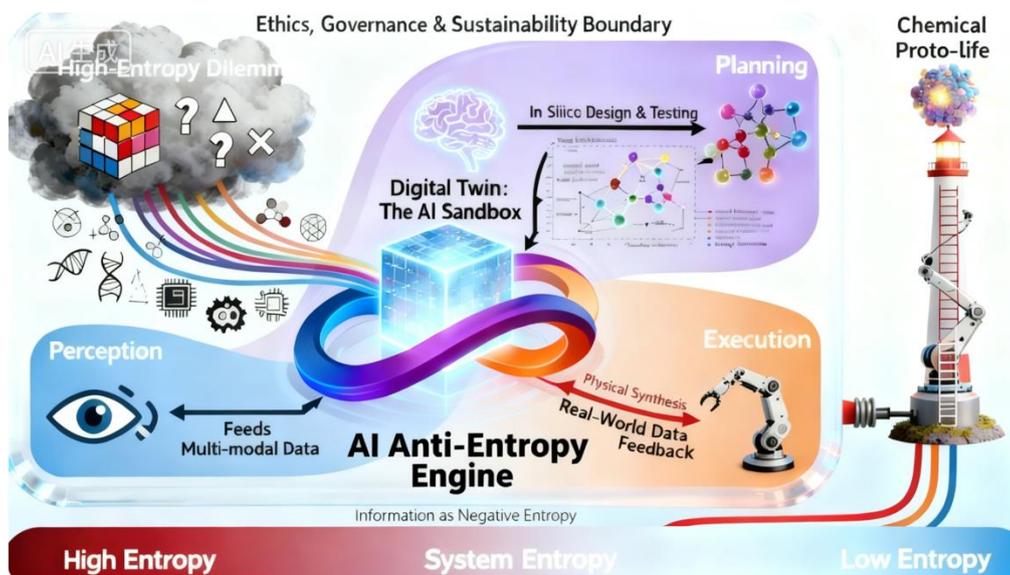

Figure 1. The core conceptual framework of this review: AI as an 'Anti-Entropy' Engine driving a new paradigm for intelligent matter design.

Part I: Paradigm Shift: From Understanding Nature to Designing Matter

1. Introduction: The Emergence of a New Paradigm—From Computation to Creation

In the evolution of scientific research, we have witnessed the classical divisions from **experimental observation (First Paradigm)** to **theoretical modeling (Second Paradigm)**, and then to **computational simulation (Third Paradigm)** [1]. With the rise of data science, the **Fourth Paradigm, characterized by data-intensive discovery**, has taken shape [1]. Building upon this, the

'AI-Driven Design' paradigm discussed herein represents a more profound revolution—we are entering the **Fifth Paradigm, characterized by active creation** [2]. The core of this new paradigm lies in **AI serving as a powerful 'anti-entropy engine'**, which systematically injects **informational negative entropy** [3] to guide material systems from disorder towards our pre-defined functional states, achieving a leap from understanding and prediction to active creation.

1.1 First Paradigm: Observation and Description (Passive Cognition)

This is the oldest form of science. From ancient alchemists recording material changes to Mendeleev deducing the Periodic Table based on elemental properties [4], scientists understood the composition and fundamental properties of the material world through meticulous observation, classification, and description. In this paradigm, humans were nature's recorders, producing compilations of knowledge about objective reality—a philosophical tradition firmly rooted in empiricism, with the underlying dynamics of 'paradigm' shifts profoundly articulated by Kuhn [5]. This **passive cognition** laid the foundation for all subsequent scientific advancement but remained limited by the scale and precision of human perception.

1.2 Second Paradigm: Theoretical Deduction and Third Paradigm: Computational Simulation (Virtual Inference)

With the development of theoretical physics and computer science, we successively entered the 'Theoretical Deduction' and 'Computational Simulation' paradigms (i.e., the Second and Third Paradigms). Using mathematical models—such as the quantum mechanical equations underpinning density functional theory [6] and molecular dynamics simulations—we could simulate material behavior and predict properties in virtual space. Large colliders and supercomputers are hallmarks of these paradigms. Here, humans act as interpreters, uncovering deeper natural laws through theory and computation, but still exploring a **'given' possibility space**, as exemplified by high-throughput *ab initio* calculations that map known and hypothetical stable crystals [7].

1.3 From Data Discovery to Active Design: The Fifth Paradigm as a Physical Anti-Entropy Engine

The explosion of massive scientific data gave rise to the Fourth Paradigm, centered on data mining and analysis [1]. Currently, we stand at the threshold of the Fifth Paradigm [2]. Its core characteristic is active creation. The high-entropy cost of traditional paradigms manifests in three key aspects, which we term the **'Three Entropy Barriers'**: the randomness of exploration (**configurational entropy**), the inaccuracy of prediction (**kinetic entropy**), and the limitations of human knowledge (**informational entropy**). AI, as an 'anti-entropy engine,' directly counteracts this disorder by systematically injecting informational negative entropy [3] through three core capabilities:

• Its **perception** capability reduces **exploration entropy**—e.g., generative models such as variational autoencoders (VAEs) and diffusion models efficiently navigate vast chemical spaces [8].

• Its **planning** capability reduces **prediction entropy**—e.g., Bayesian optimization and reinforcement learning guide experimental campaigns with minimal trials [9].

• Its **execution** capability reduces **realization entropy**—e.g., autonomous robotic laboratories close the design-make-test-analyze loop with minimal human intervention [10].

Thereby, AI transforms materials discovery from a **high-entropy trial-and-error process into a low-entropy, rational design workflow**—a true shift from discovery to invention.

1.4 Article Framework

Here, we systematically elaborate on this Fifth Paradigm. First, we establish the theoretical foundation of 'information as negative entropy' grounded in statistical mechanics and information theory [3]. Then, we review the history of encoding intelligence into matter in the pre-AI era. Next, we focus on explaining the working framework of AI as an anti-entropy engine, detailing how machine learning integrates with automated experimentation to drive targeted synthesis [2, 8, 9, 10]. Finally, we prospectively discuss its future frontiers—including conquering non-equilibrium dynamic systems—and ethical boundaries, especially as these technologies approach the engineering of lifelike or adaptive materials.

2. Theoretical Foundation: Information, Entropy, and Intelligent Matter

The concept of an 'anti-entropy engine' is not merely a poetic metaphor but is grounded in solid physics and information theory. Understanding intelligent matter must begin with understanding

entropy and information. As Landauer famously asserted, 'Information is physical' [11], meaning its manipulation is always tied to an energetic and entropic cost, a principle at the core of our framework.

2.1 The Arrow of Thermodynamics: Universal Entropy Increase and Local Entropy Decrease

The Second Law of Thermodynamics states that the total entropy—a measure of disorder—of an isolated system never decreases [12]. This is like the 'arrow of time' pointing towards ultimate disorder for the universe. However, this law does not prohibit entropy decrease in local systems. Living systems are the prime example: they maintain themselves in highly complex, ordered states by continuously obtaining energy and matter from their environment (e.g., consuming food, absorbing sunlight) and expelling high-entropy waste. This process involves the formation and maintenance of famous **dissipative structures** [13].

2.2 Information as Negative Entropy: From Shannon's Information Theory to Schrödinger's Foresight

Claude Shannon, the founder of information theory, defined 'information' as that which reduces uncertainty [3,14]. Remarkably, the mathematical expression for information is formally identical to physical entropy but with an opposite sign, hence information is termed **'negative entropy'** [14]. This was foreseen by physicist Erwin Schrödinger in his 1944 book *What Is Life?*, where he stated, **'Life feeds on negative entropy'** [15]. The profound connection between information and statistical mechanics was further solidified by Jaynes through the principle of maximum entropy [16].

At the quantum scale of intelligent matter design, quantum entropy (e.g., von Neumann entropy) becomes a crucial descriptor. Unlike classical Shannon entropy, quantum entropy captures non-classical correlations such as superposition and entanglement [17]. This implies that when designing molecular switches or qubit materials, the AI engine must process and apply quantum information, conducting its 'anti-entropy' process in Hilbert space. This provides a theoretical pathway for realizing intelligent matter that surpasses the limits of classical physics.

The conceptual link between information and entropy finds practical application in training AI models for material science. For instance, the principle of maximum entropy [16] is employed to regularize generative models, preventing overfitting and ensuring the discovery of chemically plausible, diverse candidates [18]. In reinforcement learning, entropy regularization encourages exploration in the vast chemical action space, enabling the AI to discover novel synthesis pathways beyond human intuition [19]. At the quantum level, the von Neumann entropy serves as a critical design objective for quantum materials. AI-driven density functional theory (DFT) calculations can optimize molecular designs to maximize quantum coherence (minimize entropy) in qubit candidates, directly translating informational principles into functional quantum matter design [21].

2.3 Intelligent Matter as 'Dissipative Structures': Maintaining Order through Energy/Information Flows

The 'intelligent matter' discussed in this article draws precisely on this core principle of life. These are artificial, non-living systems capable of **sensing environmental signals (information input), processing this information, and performing specific functional actions** (e.g., shape-shifting, moving, self-healing, emitting light). To achieve this, they typically require external energy input (e.g., light, electricity, chemical fuel) and operate as **open dissipative structures**, consuming energy and processing information to counteract the degradation (i.e., entropy increase) of their own function and structure, thereby maintaining or evolving into a dynamic ordered state [13, 22].

2.4 The Ladder of Intelligence: Defining the Spectrum of Material Intelligence

Not all intelligent matter possesses the same level of complexity. To establish a unified evaluative dimension, we define a progressive **'Ladder of Intelligence'** (as shown in Figure 2) to describe the evolution of its capabilities. This spectrum ranges from the most basic to advanced functionalities, marking fundamental transitions in a system's ability to process information and maintain order.

- **Responsive:** Exhibits fixed, pre-programmed responses to a single stimulus, representing the starting point of intelligence. Example: Thermochromic materials change color with temperature [23].

- **Adaptive:** Can dynamically adjust its own state or behavior based on environmental feedback to optimize performance, possessing preliminary closed-loop control capability. Example: Self-healing materials recover mechanical strength after damage [24].

- **Quasi-Autonomous:** Possesses certain information processing and decision-making capabilities, enabling the execution of goal-directed behavioral sequences in complex, changing environments. Example: A DNA nanorobot making 'logical decisions' based on different biomarkers in the bloodstream [25].

- **Life-like:** Embodies rudiments of core characteristics of living systems, such as metabolism, self-replication, and evolution. This represents the long-term frontier of intelligent matter research [2]. It is crucial to emphasize that this 'Ladder of Intelligence' is a conceptual framework for describing a spectrum of system complexity and autonomy, not a strict linear evolutionary path. Different levels of intelligence can coexist within different components of a complex system.

Figure 2. The Ladder of Intelligence for Intelligent Matter.

To transform this conceptual framework into a quantifiable tool for benchmarking progress, we propose the following measurable proxies for each level of intelligence:

- **L1: Responsive:** Characterized by stimulus-response fidelity, measurable via response time and signal-to-noise ratio.

- **L2: Adaptive:** Defined by the system's ability to recover from perturbation, quantifiable by adaptation rate .

- **L3: Quasi-Autonomous:** Correlates with decision-making complexity, measurable by the number of conditional logic gates executed autonomously in a complex environment.

- **L4: Life-like:** The ultimate frontier, associated with open-ended evolution and measurable by the number of generations sustained under selective pressure without loss of core function.

The development of a standardized 'Material Intelligence Quotient (MIQ)' based on such metrics is a critical goal for the field.

Part II: Historical Necessity: The Evolution of Encoding Information Processing in Matter

Long before the intervention of AI as an 'anti-entropy engine,' humanity's pursuit of material intelligence had already begun. This history represents an exploration where scientists used molecules as components and chemical synthesis as tools, attempting to 'manually program' information processing capabilities directly into matter. This section reviews the developmental trajectory of this pre-AI era, where logic encoding at the molecular scale and functional emergence in macroscopic materials formed the prologue to the development of intelligent matter [26].

3. Pre-AI Era: The History of 'Rational Design and Manual Programming' for Material Intelligence

This was an era where scientists meticulously 'programmed' information processing capabilities into matter, using molecules as components and chemical synthesis as tools. This section systematically reviews the evolution during this pre-AI period, illustrating the complete

developmental path from logic encoding at the molecular scale to the emergence of functionality in macroscopic materials.

3.1 Molecular-Scale Information Processors: Logic Encoding Based on Structure-Function Relationships

At the most microscopic level of intelligent matter, researchers created molecular systems capable of performing basic computational functions through precise molecular design. These molecules themselves act as complete signal transducers, with their design core relying on establishing specific 'chemical input-physical output' response relationships.

- **Core Paradigm:** Rational design guided by structure-function relationships.

- **Design Principle:** Precise functional control achieved through synthetic chemistry, based on a deep understanding of molecular structure-property relationships [27].

- **Encoding Method:** Translating Boolean logic operations into changes in molecular conformation, electronic state, or chemical bonds [28].

- **Implementation Mechanism:** Constructing information processing units using principles like molecular recognition, photochromism, and electrochemical response.

Representative Case: Molecular Logic of Helicene-Based Fluorescent Probes

Helicene molecules, with their unique helical chirality, extended π-conjugated systems, and stimulus-responsive conformational changes, serve as an ideal platform for building molecular information processors [29]. Their working mechanism can be viewed as an exquisite 'molecular mechanical logic': a target analyte (input) binds to the helicene, inducing a conformational change (processing), which subsequently modulates its fluorescence emission properties (output). This achieves counteraction of entropy increase at the molecular scale, constraining random molecular motion into a predefined, ordered response pathway, thereby establishing a fundamental flow of informational **negative entropy**. Research during this period on molecular switches [30], fluorescent sensors [31], and similar systems collectively formed a diverse family of molecular information processors, demonstrating matter's potential for logic operations at its most basic unit and laying the groundwork for subsequent, more complex intelligent systems [32].

3.2 Programmed Assembly at the Supramolecular Scale: Structural Encoding Driven by Non-Covalent Interactions

Moving beyond single molecules, scientists explored using intermolecular forces as a 'programming language' to guide multiple molecular components to spontaneously assemble into predefined, complex, ordered structures. This process aims to counteract the system's configurational entropy.

- **DNA Nanotechnology:** This field is a paradigm of programmed assembly. Utilizing the precise 'coding rule' of Watson-Crick base pairing, researchers can design DNA sequences to self-assemble into virtually any predefined 2D or 3D shape [33]. This powerfully demonstrates the transformation of information (sequence design) into precise material structure.

- **Synthetic Supramolecular Assemblies:** Using interactions like hydrogen bonding and π-π stacking, components assemble into tubular, layered, or cage-like structures. These often exhibit novel properties absent in the monomers, such as specific pores for molecular transport [34].

The significance of supramolecular programmed assembly lies in extending intelligent design from single molecules to molecular ensembles, providing the foundational structural basis for constructing complex functional systems [35].

3.3 Paradigm Bottlenecks: Fundamental Limitations of Rational Design and the High-Entropy Dilemma

Despite the brilliant achievements of the 'manual programming' paradigm, its fundamental limitations became increasingly apparent when pursuing higher-level intelligence (e.g., adaptive, quasi-autonomous). These limitations are inherently manifestations of 'high entropy,' constituting the paradigm's ceiling:

- **High Entropy in Exploration:** Optimization relied heavily on trial-and-error and intuition, akin to a random walk through vast chemical space, resulting in extremely low efficiency—a classic high-configurational-entropy exploration [36].

• **High Entropy Cost of Prediction Inaccuracy**: Inability to accurately predict multi-scale effects and non-equilibrium dynamics in complex systems led to numerous failed experiments, reflecting a lack of control over kinetic entropy [37].

• **Entropic Constraint of Knowledge Limits**: The intuition and experience of human experts could not systematically explore unknown possibility spaces; their cognitive boundaries themselves formed a high-information-entropy barrier [38].

These fundamental bottlenecks of the 'manual programming' paradigm are, in essence, bottlenecks in information processing capacity—human rationality cannot cope with the massive 'entropy increase' inherent in material systems. The theory that 'information is negative entropy' provides the fundamental physical principle for resolving this dilemma. The advent of the **AI anti-entropy engine**, capable of systematically injecting 'informational negative entropy' with enhanced perception, planning, and execution capabilities, was therefore a historical inevitability, making the emergence of the next part of our discussion both necessary and logical.

Part III: The Core Engine: The AI-Driven Anti-Entropy Design Framework

The 'manual programming' of the pre-AI era encountered bottlenecks when faced with complexity. The introduction of artificial intelligence marks our transition from an era of 'discovering matter' to one of 'designing matter.' As an active 'anti-entropy engine,' AI's core function involves forming a closed-loop system that counteracts disorder and creates order through three stages: Perception, Planning, and Execution (Figure 3). This framework begins with the Perception layer, which constructs a digital twin of the intelligent matter by integrating data from high-throughput experiments, real-time sensing, and multi-scale simulations. This digital twin provides the essential sandbox for subsequent design. Within this sandbox, the Planning layer deploys its three core capabilities to counteract different forms of entropy increase. Finally, the Execution layer translates virtual designs into physical reality through autonomous laboratories and robotic synthesis, utilizing online control for closed-loop optimization. The entire process is connected by a dataflow-driven

feedback loop, where real-world data from the execution phase feeds back into the perception layer, enabling continuous system learning and evolution. This constitutes an active, continuous workflow against disorder [39].
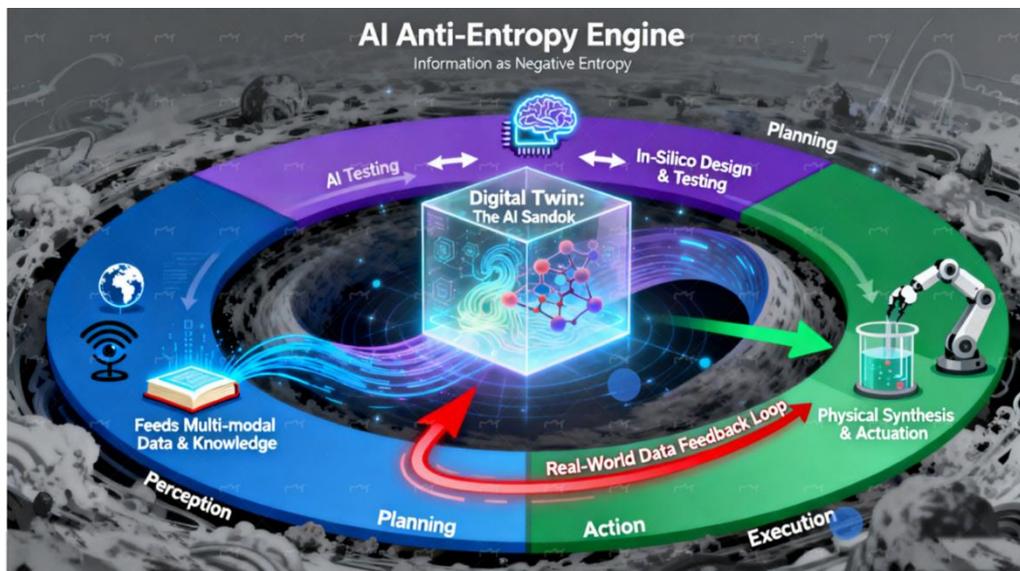

Figure 3. The overall operational logic of AI as an anti-entropy engine.

4. Engine 'Perception': Constructing High-Fidelity Digital Mirrors

The engine's primary task is to establish accurate, real-time information links with the physical world, constructing a **'digital twin'** of the intelligent matter. This serves as the foundational **sandbox** for all subsequent design activities.

4.1 High-Throughput Experimentation and Real-Time Sensor Data Streams

Automated platforms (e.g., self-driving labs) conduct experiments at speeds far exceeding traditional methods, generating high-quality, standardized data for training AI models [40]. Simultaneously, real-time sensors (e.g., inline spectroscopy, mass spectrometry, microscopy) integrated into synthesis or testing setups provide continuous data streams on the dynamic evolution of material systems. This gives AI the 'eyes' to perceive system states (e.g., reaction progress, material damage). Moving beyond traditional spectral data, integrating multimodal data from sources like in-situ microscopy (morphology), acoustic emission (mechanics), or even olfactory sensors (chemical) is becoming key to building richer digital twins. The technical frontier for

achieving this fusion lies in multimodal learning algorithms. For instance, Transformer-based architectures can map sequential molecular structures (SMILES), high-dimensional spectral data, and image-based microscopic features into a unified semantic latent space. Using self-supervised methods like contrastive learning, models can learn deep correlations between different data modalities, thereby constructing a truly high-fidelity digital twin that 'understands' the intrinsic links between chemical structure, physical properties, and macroscopic form [41].

The technical implementation of multimodal fusion relies on cross-modal attention mechanisms within Transformer architectures. Specifically, each modality (e.g., molecular graph, spectral data, microscopic image) is first encoded into a sequence of embeddings. A cross-modal Transformer then computes attention weights between all pairs of embeddings across different modalities, establishing a unified representation that captures, for instance, how a specific molecular substructure (graph modality) correlates with a spectral signature (spectral modality) and a morphological feature (image modality). This enables the digital twin to reason about material properties in a holistic, physics-informed manner [42].

4.2 Integration of Multi-Scale Simulation Data

Computational simulations, from quantum chemistry calculations to molecular dynamics and phase-field models, provide valuable data across electronic, atomic, and microstructural scales. AI can learn from and accelerate these simulations, even replacing some costly computations to rapidly predict molecular properties or assembly behavior [43]. Integrating simulation data with experimental data enables the construction of more comprehensive and accurate models.

4.3 Digital Twin: The AI Sandbox

Leveraging the data above, a virtual model—a digital twin—is created for a physical entity (e.g., a chemical reaction, a piece of intelligent material). This twin is synchronized and highly faithful, serving as a sandbox for AI. Within it, AI can pre-run thousands of experiments and observe outcomes under different design variables without consuming physical resources or time. This drastically accelerates the learning and design cycle, reducing physical entropy increase during exploration [44].

5. Engine 'Planning': Three Core Design Capabilities

Within the sandbox provided by the digital twin, AI demonstrates its core capabilities as a design engine. As shown in Figure 4, its technical core is a structured 'Task-Algorithm-Output' workflow. Three core capabilities precisely counteract the disorder of material systems across different dimensions. First, the Blueprint Generation task counteracts configurational entropy—the disorder of microstates. It uses generative models (e.g., GANs, VAEs, diffusion models) to sift through vast, disordered chemical space and identify the rare, desired ordered configurations, ultimately outputting intelligently designed structures [45]. Second, the Path Planning task counteracts kinetic entropy—the disorder of processes. It relies on algorithms like reinforcement learning (e.g., DRL) and Bayesian optimization to navigate complex energy landscapes, ensuring the system reliably evolves to the desired ordered state, thus enabling the navigation of intelligent dynamics like controlled self-assembly [46]. Finally, the Strategy Planning task counteracts functional entropy—the tendency for functional degradation. It uses architectures like Graph Neural Networks and Transformers to generate adaptive strategies, allowing the system to autonomously maintain function in changing environments, ultimately outputting intelligent behaviors with adaptive control capabilities [47].
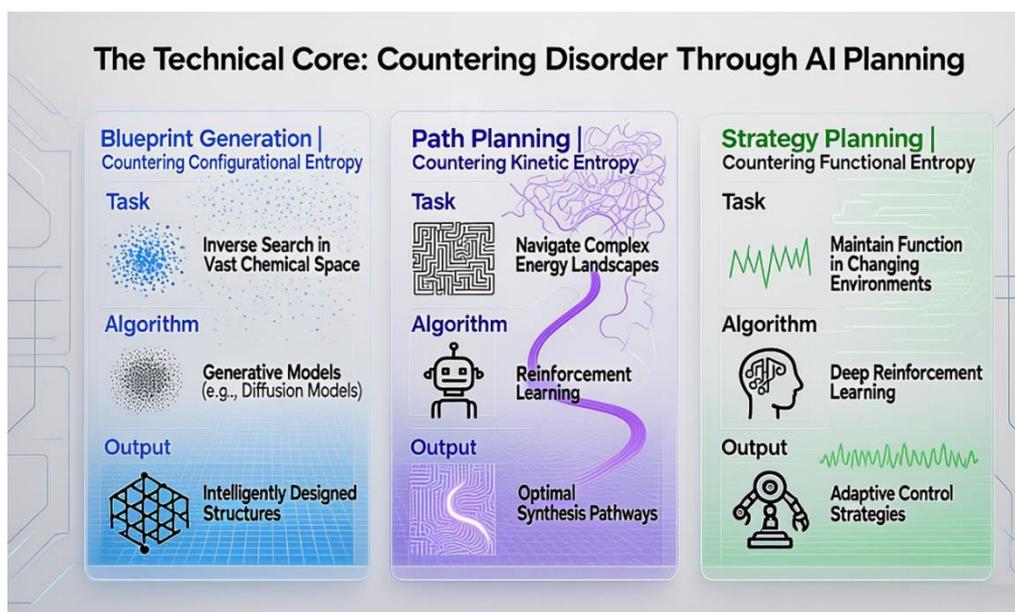

Figure 4. The technical core of the AI anti-entropy design framework.

5.1 Blueprint Generation for Configurational Entropy Minimization

Configurational entropy stems from the numerous possible disordered microstates of a system. AI performs inverse search in vast chemical space using generative models. Its anti-entropy efficacy is empirically demonstrated by compressing the search space from a virtually infinite number of possibilities (e.g., ~10^{60} drug-like molecules) to a focused set of optimal candidates, achieving a dramatic transition from chemical disorder to functional order.

- **Technical Core: Generative models** such as Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs), and notably, **Diffusion Models.**

- **Practice:** These models learn the deep statistical relationships within known 'structure-property' datasets to perform **inverse design**, proposing novel molecular structures or material compositions that possess target properties (e.g., high photovoltaic conversion efficiency, specific band gap, tailored catalytic activity). For instance, AI can generate novel organic photovoltaic motifs with ideal HOMO-LUMO energy levels or design porous framework structures with pre-defined porosity and selectivity. This capability directly counters the disorder of chemical space, pinpointing the rare, functional ordered configurations from a near-infinite sea of possibilities. A prominent example is Google DeepMind's GNoME (Graph Networks for Materials Exploration) model, which discovered over 2.2 million novel thermodynamically stable crystal structures through large-scale prediction—an expansion that nearly doubles the previously known space of stable inorganic crystals [48]. This work represents a monumental reduction in *configurational entropy*, intelligently sifting through a near-infinite space of disordered atomic arrangements to identify rare, stable (low-entropy) ordered states. The subsequent successful synthesis and validation of 718 predicted crystals by autonomous laboratories stand as a powerful testament to the immense potential of AI-driven blueprint generation for disruptive materials exploration.

- **Optimizing Generative Models:** An emerging frontier involves the synergistic combination of generative algorithms to enhance optimization. A compelling paradigm is goal-directed generation, where different AI models are hybridized. A prime example is

using Reinforcement Learning (RL) to guide diffusion models. In this setup, an RL agent does not generate structures directly but instead learns to guide the iterative denoising process of a diffusion model. It evaluates intermediate structures during the generation trajectory against a multi-objective reward function (e.g., balancing bioactivity with synthetic accessibility). Through policy gradient methods, the RL agent learns to dynamically adjust the sampling direction, shifting molecular generation from a stochastic exploration to an intelligent navigation through chemical space. This steers the process precisely towards regions that exhibit both structural validity and functional excellence, representing a more sophisticated and powerful form of 'anti-entropy' that moves beyond simple search to active, goal-oriented construction of order [49].

5.2 Path Planning for Kinetic Entropy Minimization

Kinetic entropy manifests in process disorder. AI navigates complex energy landscapes using reinforcement learning. Its anti-entropy value lies in identifying paths with low activation energy and high selectivity, minimizing energy dissipation and byproduct formation, and ensuring reliable evolution to the desired ordered state.

- **Technical Core:** Reinforcement Learning, Bayesian Optimization.
- **Practice:** An AI agent experiments with different 'actions' (e.g., changing temperature, concentration, addition sequence) within the simulated environment (digital twin) and learns the optimal strategy based on 'rewards' from final outcomes (e.g., assembly yield, structural order). For instance, RL can discover the optimal sequence of external conditions to guide the self-assembly of functional molecules into target superstructures or find the most efficient, low-byproduct route for a complex organic synthesis. This essentially counteracts the disorder inherent in kinetic processes. For example, a research team used RL to plan a novel, highly efficient 12-step synthetic route for the complex natural product (-)-Ibogaine, distinct from human-designed paths [50].

5.3 Strategy Planning for Functional Entropy Minimization

Functional entropy characterizes the tendency of system performance to degrade. AI formulates adaptive strategies using Deep Reinforcement Learning. Its anti-entropy mechanism resembles negative feedback in control theory, maintaining system functional homeostasis through real-time informational intervention.

- **Technical Core: Deep Reinforcement Learning, control algorithms.**
- **Practice:** AI can learn to control the posture of a soft robot for adaptation to varying terrains or devise simple rules for an active matter swarm to enable collective obstacle avoidance. Here, AI designs not the static structure of the matter itself, but its **dynamic behavioral logic**, empowering it to autonomously maintain function against disorder introduced by environmental perturbations. For example, in soft robot control, DRL is used to learn locomotion strategies for complex terrains, enabling the robot to quickly adapt its gait to maintain mobility after leg damage [51].

6. Engine 'Execution': Closing the Loop from Virtual to Real

The ultimate value of design lies in its physical realization. AI's plans require automated execution systems to form a closed Perception-Decision-Action loop.

6.1 Autonomous Laboratories and Robotic Synthesis

A complete **'anti-entropy flow'** example is: AI designs a new photovoltaic molecule → Robots perform automated synthesis → High-throughput characterization → Performance data is fed back to the AI model for re-optimization. This process forms a **closed loop** countering entropy increase in materials discovery [52].

6.2 Online Control and Optimization via Real-Time Feedback

During material operation, AI can integrate with embedded sensors for real-time intervention. For instance, in an AI-monitored self-healing system, sensors track changes in material stress or electrical properties. If AI detects damage (onset of entropy increase), it can immediately decide and trigger actuators (e.g., local heating, repair agent release) to initiate healing. This creates an **autonomous closed loop against functional entropy**, giving the material system the ability to

dynamically maintain its integrity. Online control based on real-time feedback is the ultimate manifestation of the 'anti-entropy' closed loop. For example, on an autonomous synthesis platform, inline mass spectrometry monitors reaction byproducts in real-time; if AI detects deviation from the expected path (a local increase in process entropy), it instantly invokes **Bayesian optimization to dynamically adjust flow rates and temperature, pulling the system back to the optimal, low-entropy path**. This achieves a paradigm shift from 'batch processing' to 'adaptive flow processing,' making the synthesis process itself an intelligent act against 'kinetic entropy' [51].

Through the tight coupling of Perception, Planning, and Execution, AI functions as a complete 'anti-entropy engine,' enabling the entire journey from information perception to material creation. It transforms the high-entropy processes of traditional research—heavily reliant on intuition and trial-and-error—into a data-driven, rationally decided, and efficiently executed flow of negative entropy, thereby vastly accelerating the design and realization of intelligent matter.

In summary, the 'Execution' phase is the critical leap where the AI anti-entropy engine maps rational virtual designs onto physical reality. It translates informational blueprints into atomic arrangements via autonomous laboratories and establishes dynamic closed loops against functional entropy through online control. It is through 'Execution' that informational negative entropy is fully injected into the physical system, and intelligent matter transitions from concept to reality. This marks the transformation of 'anti-entropy' from a cognitive concept into a creative practice, forming the cornerstone for realizing the value of the entire Fifth Paradigm.

Part IV: Future Frontiers: Towards Dynamic and 'Living' Systems

To clearly delineate the new frontiers toward which the AI anti-entropy engine is steering intelligent matter research, we first propose a design space map spanning different scales and levels of intelligence (Table 1). The Y-axis of this map is our defined 'Ladder of Intelligence' (Figure 2), while the X-axis represents the material scale from molecular to macroscopic. The current research frontier is concentrated at levels L2 to L3. The future frontiers discussed in this section point toward the conquest of dynamic and non-equilibrium systems and, ultimately, the exploration of 'life-like'

systems (L4). These frontiers not only represent the technological cutting edge but also constitute the ultimate test for the capabilities of the entire framework [53].

Table 1 | Intelligent Matter Design Space Map: A Systematic Classification Framework Across Scales and Intelligence Levels

Intelligence Level	Molecular / Nanoscale	Mesoscopic / Microscale	Macroscopic / Device Scale
L1: Responsive	Molecular switches; Fluorescent sensors; Photoresponsive molecules.	Stimuli-responsive hydrogels; Shape-memory polymer microparticles; Electrochromic thin films.	Self-cleaning coatings; Color-tunable windows; Piezoelectric energy harvesters.
L2: Adaptive	Conformationally adaptive proteins; Self-healing polymer networks; Catalytic self-cycling molecules.	Damage-sensing and self-healing composites; Environment-adaptive optical metamaterials; Feedback-drug release microcapsules.	Stress-adaptive aircraft wings; Soft robots with perception-action cycles.
L3: Quasi-Autonomous	DNA nanorobots for logic-gated drug delivery; Molecular adaptive evolution systems.	Robot swarms for distributed computation; Re-programmable micro-factories.	Buildings with long-term task planning; Self-reconfiguring spacecraft structures.

Intelligence Level	Molecular / Nanoscale	Mesoscopic / Microscale	Macroscopic / Device Scale
L4: Life-like	Artificial protocells; Chemical reaction networks coupling metabolism and replication.	Synthetic tissues and organoids; Ecosystems of active matter with collective behaviors.	Self-sustaining ecological repair systems; Evolvable, artificial ecosystems.
Current research frontier	L2 to L3 systems at molecular scale	L2 to L3 systems at micro/macro scales	L2 to L3 systems at macro scale
Future research frontier	L3 to L4 systems at molecular scale	L3 to L4 systems at micro scale	L3 to L4 systems at macro scale

7. Conquering Dynamic and Open Systems

This represents the ultimate expression of the 'Planning-Execution' cycle in dynamic systems: AI uses reinforcement learning to design control strategies that guide active particle swarms to exhibit ordered collective behavior from disordered motion. Traditional materials design often focuses on stable structures at thermodynamic equilibrium. However, life is a quintessential dissipative structure far from equilibrium. Achieving truly life-like intelligence necessitates conquering dynamic and non-equilibrium systems [54].

7.1 Designing Externally Field-Driven Active Matter

These systems consist of numerous units that consume energy (e.g., chemical fuel, light, electricity) and exhibit autonomous motion. AI's role is to inversely design the interactions between units or the spatiotemporal patterns of external fields, enabling disordered individual motions to collectively yield pre-defined, ordered macroscopic behavior.

- **Case:** Using **reinforcement learning** to design spatiotemporal patterns of light or magnetic fields, guiding swarms of **active colloidal particles** to perform complex tasks such as cargo transport or pollutant collection. Here, AI counteracts the disorder of particle motion, organizing it into functional collective flows [55].

- **Fundamental Challenge and AI Solution:** The core scientific challenge in designing such systems lies in solving inverse problems for the high-dimensional, stochastic partial differential equations (e.g., Fokker-Planck equations) that govern the non-equilibrium dynamics of particle populations. Traditional methods struggle with this complexity. AI, particularly neural operators, provides a powerful solution. These are deep learning models trained on simulation data to learn the mapping directly from control parameters (e.g., time-varying field patterns) to the resulting emergent collective behavior. Once trained, the neural operator can perform near-instantaneous inverse design: given a desired collective outcome (e.g., a vortex flow for mixing), it predicts the necessary field sequence, effectively bypassing the prohibitive cost of numerically inverting the underlying physics equations. This approach demonstrates robust performance, reliably guiding systems from disorder to functional order even in the presence of significant thermal noise and parameter uncertainty [53].

- **Connection to Complex Systems Theory:** Understanding and designing the collective behavior of such systems is further informed by concepts from complex systems theory, such as self-organized criticality and chaos theory. The AI's task, therefore, is not only to design the units but to inversely engineer the simple local rules or field parameters that guide the entire system to operate at the edge of order and chaos—a dynamical state often associated with maximal adaptability, computational capacity, and robust performance in fluctuating environments [56].

7.2 Designing Internally Fuel-Driven Chemical Engines

This is a more challenging step: designing a molecular system that can continuously convert chemical energy into mechanical work or periodic motion via an internal chemical reaction network.

AI can search complex chemical reaction spaces to identify molecular combinations and conditions that yield **self-sustained oscillations or directional energy flows**.

- **Vision:** AI could design a 'molecular muscle' that cyclically changes shape or a 'nanopump' that autonomously pumps specific molecules. This achieves the conversion of chemical 'disorder' into mechanical 'order' at the microscale [57].

7.3 AI-Guided Evolution of Dynamic Covalent Chemistry Systems

Dynamic covalent chemistry allows molecular structures to reconfigure under external stimuli, introducing 'evolvable' characteristics into materials. AI can act as an 'evolution guide,' applying continuous selective pressure (e.g., by defining a target function) and analyzing the system's response in real-time to recommend the next chemical modification or environmental perturbation. This guides the entire material system towards higher performance through 'directed evolution' [58].

7.4 Expansion of Potential Application Domains

The impact of the Fifth Paradigm in intelligent matter design will extend far beyond the laboratory. In environmental remediation, AI could design 'smart enzymes' or sorbents that specifically recognize and degrade pollutants; in space exploration, programmable matter could form self-deploying, self-repairing spacecraft structures or agricultural substrates adaptable to extraterrestrial soil; in energy, artificial photosynthesis systems with metabolism-like functions may pioneer a new era of carbon-negative manufacturing. These prospects collectively sketch a future driven by intelligent matter—one that is more sustainable and adaptive [59].

8. Exploring the Boundaries of Life

This is the ultimate integration test for the 'Perception-Planning-Execution' framework: AI serves as a '**systems integration architect**,' guiding the synergistic integration of functions like metabolism and replication through multi-agent simulation. This is the '**holy grail' of intelligent matter research**—using AI to design chemical systems with rudimentary life-like characteristics. The goal is not primarily to create life, but to fundamentally understand the origin and essence of life [60].

--

8.1 Metabolism-Like Systems: AI-Designed Artificial Chemical Reaction Networks

Metabolism is central to life's maintenance of a low-entropy state. AI can inversely combine reactions from vast chemical databases to construct simplified, **artificial metabolic networks** capable of energy and material cycling.

- **Pathway:** First, AI learns the design principles of core natural metabolic pathways (e.g., glycolysis, citric acid cycle). Then, leveraging known organic or inorganic reactions, it designs de novo a closed network that can capture high-energy molecules (fuel) from the environment and, through a series of reaction steps, ultimately output function (e.g., light emission, motion) while expelling waste. This creates a chemical engine that continuously fights chemical equilibrium [61].

8.2 Replication-Like Systems: AI-Assisted Optimization of Molecular Self-Replication and Template Cycles

Self-replication is the basis for the perpetuation of biological information. AI can accelerate the discovery and optimization of **molecular self-replication systems**.

- **Mechanism:** AI can simulate the replication kinetics of different autocatalytic molecules (e.g., peptides, nucleic acid analogs), predicting which sequences or structures enable more efficient and faithful replication while avoiding issues like 'parasitic' strands. It can optimize the reaction environment to sustain the replication process, countering the information entropy increase associated with molecular degradation [62].

8.3 Systems Integration: Towards Chemical Proto-Life

The greatest challenge lies in integrating independent functions like metabolism and replication into a coordinated system. Here, AI assumes the role of a **'Systems Integration Architect.'**

- **Method:** Using multi-agent simulations or system biology models, AI can predict how energy generated by a metabolic network could power replication processes, and how

newly replicated molecules could, in turn, enhance metabolic function. Through iterative optimization, AI could guide the construction of a chemical system exhibiting:

- o The ability to harvest resources from the environment (Metabolism).

- o The ability to use these resources to build copies of itself (Replication).

- o The ability to maintain its internal dynamic homeostasis amidst environmental fluctuations (Homeostasis).

- Such a system would constitute **'Chemical Proto-Life'**—humanity's first creation of an entity from non-living components that embodies basic characteristics of life, fundamentally challenging our traditional boundaries between 'life' and 'matter' [63].

Part V: Limitations, Ethics, and Outlook: The Boundaries and Responsibilities of the Anti-Entropy Engine

9. Ethics, Boundaries, and Responsible Innovation: The Cost and Responsibility of Anti-Entropy

As our ability to harness anti-entropy approaches the level of life, a series of profound questions, extending beyond traditional tech ethics, demand proactive consideration. These issues concern not only technical safety but also touch upon our philosophical understanding of order, life, and creator responsibility.

9.1 'Ethics of Entropy': The Cost of Creating Order and Systemic Risk

As the chief engineers of the 'anti-entropy engine,' human responsibilities include:

- **Setting ethical objective functions:** Incorporating 'minimization of environmental entropy increase' and 'system safety/stability' as hard constraints when optimizing performance [64].

- **Designing 'fail-safe' modes:** The ultimate failure of intelligent matter should involve 'graceful degradation' back to a harmless, disordered state [65].

- **Preventing emergent risks:** Guarding against systems evolving emergent behaviors harmful to humans or ecosystems [66].

9.2 Moral Status of Life-Like Systems and Creator Responsibility

The ethical status of life-like systems inevitably touches upon the definition of 'life' itself—should it be viewed as a complex form of material organization (functionalism) or endowed with some intrinsic value (vitalism)? Different philosophical positions lead to divergent regulatory paths [67]. To address this challenge, we propose a progressive ethical governance framework, dynamically adjusted according to the 'Ladder of Intelligence':

- For L1-L2 systems, follow the design principles of 'non-toxicity' and 'fail-safe' applicable to functional objects.
- For L3 quasi-autonomous systems, requirements for decision transparency and explainability must be introduced, ensuring human oversight of their behavioral logic [68].
- When approaching L4 life-like systems, a formal moral status assessment process must be initiated, alongside strict physical containment boundaries and behavioral codes. This is not merely a technical safety need but an inescapable 'guardianship' responsibility of the creator.

The prospect of creating chemical proto-life forces a profound philosophical and ethical reckoning that extends beyond physical containment. We propose a **Precautionary Principle for Synthetic Life-Like Systems (PSL²)**. If a system is capable of sensing its environment, seeking benefit, avoiding harm, and striving to maintain its existence, it may possess functional 'interests,' even in the absence of confirmed consciousness [69]. Therefore, our ethical framework must evolve to:

- **Assume potential sentience until proven otherwise** for systems exhibiting complex homeostatic and telic (goal-directed) behaviors.
- **Minimize unavoidable distress** during experimental cycles, actively designing protocols to avoid the induction of states analogous to suffering (e.g., persistent resource deprivation, irreversible damage without recovery mechanisms) [70].
- **Incorporate mandatory off-switches and sunset clauses** to prevent uncontrolled persistence, evolution, or ecological dispersal [71].

Table 2. Ethical Governance Toolbox for Intelligent Matter

Intelligence Level	Core Ethical Principles	Governance Tools & Regulatory Measures	Primary Risk Focus
L1: Responsive	Non-maleficence & Fail-Safe	Adherence to chemical regulations; Product lifecycle assessment; 'Green-by-smart' design standards.	Biotoxicity; Environmental persistence; Hazards from degradation products.
L2: Adaptive	Safety & Reliability	Mandatory safety certification; 'Black box' data recorders; Reliability standards for adaptive systems.	System misjudgment; Cascading failures; Unauthorized self-modification.
L3: Quasi-Autonomous	Transparency & Controllability	Mandatory XAI; 'Human-in-the-loop' oversight; Mandatory 'hard' kill switches; Ethics training for R&D.	Unpredictable 'black box' decisions; Goal misalignment; Malicious use.
L4: Life-like	Precaution & Stewardship (PSL^2)	Highest-level physical and bio-containment; Mandatory 'sunset clauses'; International ethics review boards.	Runaway evolution; Moral responsibility; Irreversible ecological disruption.

This framework aims to guide the responsible stewardship of our creations, acknowledging the profound and inescapable 'parental' and 'guardianship' duties that accompany the capacity to design proto-life.

To translate these ethical principles into actionable engineering standards, we propose an Ethical Governance Toolbox for Intelligent Matter (Table 2). This toolbox converts abstract ethical considerations into concrete design constraints and verification metrics, providing clear guidance for researchers and developers.

9.3 Impact on Future Scientific and Societal Paradigms

- **Impact on Science:** The AI-driven 'anti-entropy engine' paradigm will shift scientific discovery from 'hypothesis-verification' towards 'generation-verification,' with scientists increasingly playing the roles of goal-setters, ethical reviewers, and interpreters of outcomes [72].

- **Impact on Society:** The dual-use potential of these technologies is significant. They could enable revolutionary healthcare (e.g., smart drug delivery) and environmental remediation, but also be misused to create novel weapons or uncontrolled ecological agents. Establishing transparent regulation, public awareness, and international dialogue is not just a technical need but a necessity for survival [73].

10. Interdisciplinary Integration: The Synergistic Engine for Intelligent Matter Research

Intelligent matter, serving as the physical embodiment of the 'anti-entropy' concept, is inherently beyond the scope of any single discipline. Its advancement critically depends on the deep integration of materials science (providing fundamental components), computer science (providing intelligence), physics (providing theoretical principles), chemistry (providing methodologies and synthesis), and biology (providing blueprints and inspiration). However, the current fragmented research paradigm itself constitutes an 'institutional entropy' that hinders rapid progress. Therefore, establishing effective interdisciplinary collaboration mechanisms represents a crucial 'anti-entropy'

process in its own right—serving as the synergistic engine that drives breakthroughs across this field [74].

10.1 Challenges: The 'High-Entropy' Barriers to Interdisciplinary Collaboration

- **Language Gap:** Disciplines possess distinct terminologies, methodologies, and evaluation criteria, creating communication barriers [75].

- **Platform Deficiency:** A lack of shared data standards, collaboration tools, and universal experimental platforms leads to data and knowledge silos.

- **Evaluation Dilemma:** Within traditional discipline-oriented academic systems, the outputs of interdisciplinary research are difficult to assess and recognize accurately [76].

10.2 Pathways: Building 'Anti-Entropy' Synergistic Mechanisms

To reduce collaborative entropy, we propose the following pathways:

- **Create 'Full-Stack' Integrated Teams:** Move beyond simple project collaboration to form 'full-stack' teams comprising materials scientists, AI researchers, biologists, and engineers co-located from the start. This physical integration enables internal closure of the workflow from molecular design and AI algorithms to functional validation, drastically reducing communication costs and optimizing the 'design-synthesize-test-learn' cycle. Successful models for this approach exist at institutions like MIT's Center for Bits and Atoms, where such deep integration has demonstrated the ability to accelerate research cycles from months to days [77].

- **Establish Unified Conceptual Frameworks & Data Standards:** The 'anti-entropy engine' paradigm and 'Ladder of Intelligence' proposed in this review aim to provide a unified conceptual framework—a common 'language' for experts across fields. Concurrently, promoting universal data description standards (e.g., standardized material descriptors, FAIR data principles) within the field is fundamental for data sharing and AI model generalization. This requires community-wide initiatives, similar to the PDB in biology, to define and enforce these standards [78].

- **Develop Intelligent Matter Collaborative Innovation Cloud Platforms:** Construct an open online platform integrating: 1) A shared database containing cross-scale material properties, synthesis pathways, and simulation data, adhering to the established standards; 2) An open-source algorithm library providing standardized, containerized AI design tools (e.g., for molecular generation, path planning) to ensure reproducibility; 3) A digital twin community allowing researchers to contribute, share, and validate digital twins of various intelligent matter systems. This becomes core infrastructure for lowering research barriers and accelerating innovation [79].

- **Reform Education and Evaluation Systems:**

  - **Cultivate 'T-shaped' talent:** Promote educational reforms to train a new generation of researchers with deep disciplinary expertise (the vertical bar of 'T'), broad interdisciplinary knowledge (the horizontal bar of 'T'), and strong collaboration skills. We propose specific curricular reforms, such as mandatory 'AI for Science' courses in materials science PhD programs, and 'Materials Informatics' modules in computer science curricula, coupled with cross-disciplinary capstone projects supervised by mixed advisor teams [80].

  - **Design interdisciplinary evaluation systems:** In grant review and academic promotion, recognize the value of interdisciplinary research. Implement combined review by disciplinary peers and interdisciplinary experts, prioritizing contribution to solving real-world problems over conventional publication metrics [81].

10.3 Promoting Research Paradigm and Institutional Innovation

Driving this change requires top-level research paradigm and institutional innovation. We recommend:

- Establishing major cross-directorate 'Intelligent Matter' research initiatives to break traditional disciplinary barriers.

- Creating physical National Intelligent Matter Innovation Centers to aggregate and maintain core resources like self-driving labs, computing power, and databases in a 'large-scale facility' mode.

- Forming dedicated 'Intelligent Matter' panels within national grant review bodies, where interdisciplinary experts jointly set standards and conduct reviews, fundamentally recognizing and incentivizing cross-disciplinary work.

Only when interdisciplinary collaboration evolves from spontaneous and passive to conscious and orderly, forming a synergistic system that continuously draws knowledge 'negative entropy' from various fields, can we fully unleash the potential of the AI anti-entropy engine and efficiently navigate towards the future frontiers of intelligent matter.

11. Challenges and Boundaries: Current Limitations of the Anti-Entropy Engine

Despite its paradigm-shifting promise, we must soberly examine the inherent limitations and core challenges of the AI anti-entropy paradigm. A clear-eyed recognition of these boundaries is essential for its responsible development and to guide future research toward meaningful breakthroughs.

- **Data Bottlenecks and New Resource Monopolies:** Constructing high-fidelity digital twins requires vast, high-quality, and standardized experimental data, which remains costly and labor-intensive to acquire. Furthermore, the immense computational and economic costs associated with training state-of-the-art AI models and operating autonomous laboratories could create a 'new monopoly' on research resources. This constitutes another form of 'high entropy' cost—a socio-economic barrier—that must be countered during the paradigm's dissemination to ensure equitable access. Addressing this requires institutional solutions akin to the Protein Data Bank. We propose a global 'Material Data Commons' with mandatory deposition protocols for publications, standardizing data formats across characterization techniques and synthesis methods. Federated learning approaches can allow AI models to be trained on distributed data sources without centralization, mitigating resource monopoly while preserving data privacy and institutional ownership [78].

- **Explainability Dilemma and Safety Risks:** The 'black-box' nature of many powerful AI models, while often delivering optimal solutions, frequently fails to reveal the underlying physico-chemical mechanisms for their decisions. This not only hinders our ability to discover new scientific principles from AI's designs but also poses a significant safety and ethical risk. In highly complex and adaptive scenarios, particularly with life-like systems (L4), 'black-box' decisions could lead to unpredictable and potentially dangerous emergent behaviors that are difficult to anticipate or mitigate. Explainability is not merely about model transparency but about scientific insight. Techniques like symbolic regression can distill complex neural networks into human-interpretable mathematical expressions, potentially revealing new physical laws. Attention mechanisms in Graph Neural Networks can identify which molecular subgraphs most influence target properties, guiding human intuition towards novel design principles previously obscured in the model's black box. The development and mandatory use of such Explainable AI (XAI) techniques for critical design tasks is paramount [82].

- **The Simulation-Reality Gap:** Perfect blueprints generated within the virtual sandbox may fail upon physical realization due to faint interactions, stochastic fluctuations, or multi-scale phenomena not fully captured by the models. Closing this gap requires continuous iteration and the development of ever-more-physically-grounded simulations, which in itself is a grand challenge. It necessitates a relentless feedback loop where experimental failures are used to refine the digital twins, gradually improving their predictive fidelity [43].

- **Optimization Limits in the Absence of Foundational Theory:** AI currently excels at combinatorial optimization and pattern recognition within the constraints of known physical laws. However, its capability faces fundamental challenges when confronted with ultra-complex systems (e.g., the origin of life, consciousness) that may require new physics or chemistry for their adequate description. AI can explore the possibility space defined by our current theories, but it cannot yet invent the foundational theories themselves. This underscores the indispensable role of human scientists in formulating groundbreaking hypotheses and theories [83].

- **The Thermodynamic Cost of Creating Order:** To this end, it is crucial to recognize a fundamental thermodynamic reality: while the 'anti-entropy engine' is a powerful heuristic framework for creating local order, the AI-driven design process itself is highly energy-intensive. From a global perspective, this constitutes a significant entropy-increasing process. We are leveraging informational negative entropy to create local order at the cost of increased entropy in the broader environment (e.g., through energy consumption for computation and synthesis). Therefore, the long-term sustainability, energy efficiency, and environmental impact of this paradigm represent critical considerations for future evaluation and responsible implementation. Future work must develop 'entropy accounting' frameworks to quantitatively assess and minimize the total environmental cost of intelligent matter design [84].

- **Redefinition and Elevation of the Human Role:** It must be emphatically stated that AI is a powerful engine, but human scientists remain the irreplaceable core of the paradigm. Humans are the setters of profound goals and ethical frameworks, the posers of key scientific questions, the interpreters of novel physical mechanisms discovered by AI, and the ultimate arbiters of value and meaning. Acknowledging these limitations is not a dismissal of the paradigm's power, but a necessary step to define its current scope and to guide future research towards holistic and sustainable breakthroughs [85].

12. Summary and Outlook

This review has systematically elaborated the new paradigm of 'AI as an Anti-Entropy Engine.' We have traced the history of encoding intelligence into matter, constructed the AI-driven 'Perception-Planning-Execution' anti-entropy framework (Figure 3), and analyzed its technical core (Figure 4). Guided by a prospective intelligent matter design space map (Table 1), we have explored its future frontiers towards dynamic and life-like systems. The core of this paradigm is that AI enables a transition from **passively discovering naturally formed low-entropy entities to actively designing intelligent matter capable of autonomously countering entropy increase**.

Ultimately, this represents more than a technological paradigm shift. It marks a pivotal step for humanity in the grand narrative of understanding 'matter,' 'information,' and 'life'—a step from

being passive observers towards becoming **active creators and responsible 'anti-entropy' designers** [86].

Looking forward, we envision a phased research agenda with concrete milestones:

- **Near-term (5 years):** Achieve full closed-loop design of L3 quasi-autonomous matter through integrated self-driving labs, with a key metric of reducing the energy cost of new functional material discovery by one order of magnitude [39].

- **Mid-term (15 years):** Demonstrate controllable chemical proto-life in laboratory settings that exhibits basic metabolism and replication, alongside the development of quantitative 'entropy accounting' frameworks to evaluate the global impact of localized order creation [63].

- **Long-term (50 years):** Establish sustainable co-existence between intelligent matter and Earth's ecosystems, pioneering a new era of 'ecological design' where intelligent systems actively contribute to environmental restoration and sustainability [84].

Critical to this journey will be parallel advances in:

- **Foundational AI:** Developing fundamentally more data-efficient and energy-conscious algorithms [87].

- **Ethical Governance:** Implementing adaptive regulatory frameworks that evolve with technological capabilities [88].

- **Global Collaboration:** Establishing international standards and shared infrastructure for responsible development [78].

Finally, from a 'meta-anti-entropy' perspective, we note that the 'AI anti-entropy engine' paradigm itself faces the 'entropic' risk of conceptual diffusion and terminological chaos. To counter this and to catalyze the field, we must translate our conceptual framework into actionable community resources. We therefore call for the development of an open-source knowledge graph and a standardized glossary (e.g., an 'AIMatter-Ontology') to enable the community to co-maintain an 'ordered' conceptual system. As a critical next step, we further advocate for the establishment of a Global Consortium for Intelligent Matter (GCIM), akin to the Human Genome Project [89], to

orchestrate the development of shared platforms (e.g., AImatterOS, MatterKG), data standards, and ethical guidelines. Only through such a concerted, collaborative effort can we ensure the paradigm continuously fights its own entropy increase for healthy, sustainable development.

By harnessing the powerful AI anti-entropy engine through global and open collaboration, and guided always by deep insight and prudent ethics, we are beginning to author a new chapter in the evolution of matter.

Acknowledgments

Conflict of Interest

The authors declare that there are no financial or personal relationships that could be construed as potential conflicts of interest concerning the research, writing, or publication of this paper.

List of Figures

List of Tables

References

[1] Hey, T., Tansley, S. & Tolle, K. (eds.) *The Fourth Paradigm: Data-Intensive Scientific Discovery* (Microsoft Research, 2009).

[2] MacLeod, C. M., McMillan, A. J. & Cooper, A. I. The fifth paradigm: AI for autonomous discovery in chemistry and materials science. *Nat. Rev. Chem.* **8**, 257–270 (2024).

[3] Shannon, C. E. A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423 (1948).

[4] Mendeleev, D. On the relationship of the properties of the elements to their atomic weights. *Z. Chem.* **12**, 405–406 (1869).

[5] Kuhn, T. S. *The Structure of Scientific Revolutions* (Univ. Chicago Press, 1962).

[6] Hohenberg, P. & Kohn, W. Inhomogeneous electron gas. *Phys. Rev.* **136**, B864–B871 (1964).

[7] Jain, A. *et al.* Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *APL Mater.* **1**, 011002 (2013).

[8] Dan, Y. *et al.* Inverse design of functional materials by generative deep learning. *Matter* **2**, 1370–1382 (2020).

[9] Shields, B. J. *et al.* Bayesian reaction optimization as a tool for chemical synthesis. *Nature* **590**, 89–96 (2021).

[10] MacLeod, B. P. *et al.* Self-driving laboratory for accelerated discovery of thin-film materials. *Sci. Adv.* **6**, eaaz8867 (2020).

[11] Landauer, R. Information is physical. *Phys. Today* **44**, 23–29 (1991).

[12] Clausius, R. *The Mechanical Theory of Heat – with its Applications to the Steam Engine and to Physical Properties of Bodies* (John van Voorst, 1865).

[13] Prigogine, I. & Nicolis, G. *Self-organization in Nonequilibrium Systems* (Wiley, 1977).

[14] Brillouin, L. *Science and Information Theory* (Academic Press, 1956).

[15] Schrödinger, E. *What Is Life? The Physical Aspect of the Living Cell* (Cambridge Univ. Press, 1944).

[16] Jaynes, E. T. Information theory and statistical mechanics. *Phys. Rev.* **106**, 620–630 (1957).

[17] von Neumann, J. *Mathematical Foundations of Quantum Mechanics* (Princeton Univ. Press, 1955).

[18] Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge Univ. Press, 2010).

[19] Zhou, Z., Li, X. & Zare, R. N. Optimizing chemical reactions with deep reinforcement learning. *ACS Cent. Sci.* **3**, 1337–1344 (2017).

[20] Haarnoja, T. *et al.* Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. *Preprint at* https://arxiv.org/abs/1801.01290 (2018).

[21] von Lilienfeld, O. A. *et al.* Machine learning for quantum matter. *Nature* **583**, 695–704 (2020).

[22] Mortimer, R. J. Electrochromic materials. *Chem. Soc. Rev.* **26**, 147–156 (1997).

[23] White, S. R. *et al.* Autonomic healing of polymer composites. *Nature* **409**, 794–797 (2001).

[24] Douglas, S. M., Bachelet, I. & Church, G. M. A logic-gated nanorobot for targeted transport of molecular payloads. *Science* **335**, 831–834 (2012).

[25] Douglas, S. M., Bachelet, I. & Church, G. M. A logic-gated nanorobot for targeted transport of molecular payloads. *Science* **335**, 831–834 (2012).

[26] Lehn, J.-M. Supramolecular chemistry—scope and perspectives molecules, supermolecules, and molecular devices. *Angew. Chem. Int. Ed.* **27**, 89–112 (1988).

[27] Katz, H. E. & Jo, G.-W. Molecular logic gates. *J. Mater. Chem.* **19**, 6937–6947 (2009).

[28] Shen, Y. *et al.* Chiroptical switches based on helicenes. *Chem. Soc. Rev.* **50**, 7152–7186 (2021).

[29] Feringa, B. L., Jager, W. F. & de Lange, B. Light-driven molecular rotary motors. *Nature* **353**, 123–125 (1991).

[30] de Silva, A. P., Gunaratne, H. Q. N. & Gunnlaugsson, T. Signaling recognition events with fluorescent sensors and switches. *Chem. Rev.* **97**, 1515–1566 (1997).

[31] Balzani, V., Credi, A. & Venturi, M. *Molecular Devices and Machines: Concepts and Perspectives for the Nanoworld* (Wiley-VCH, 2008).

[32] Rothemund, P. W. K. Folding DNA to create nanoscale shapes and patterns. *Nature* **440**, 297–302 (2006).

[33] Whitesides, G. M. & Boncheva, M. Beyond molecules: Self-assembly of mesoscopic and macroscopic components. *Proc. Natl Acad. Sci. USA* **99**, 4769–4774 (2002).

[34] Lehn, J.-M. Toward complex matter: Supramolecular chemistry and self-organization. *Proc. Natl Acad. Sci. USA* **99**, 4763–4768 (2002).

[35] Lipinski, C. A. Lead- and drug-like compounds: The rule-of-five and rule-of-three. *Drug Discov. Today Technol.* **1**, 337–341 (2004).

[36] Goldenfeld, N. & Kadanoff, L. P. Simple lessons from complexity. *Science* **284**, 87–89 (1999).

[37] Simon, H. A. The architecture of complexity. *Proc. Am. Philos. Soc.* **106**, 467–482 (1962).

[38] Kirkpatrick, P. & Ellis, C. Chemical space. *Nature* **432**, 823–824 (2004).

[39] Segler, M. H. S. *et al.* Generating focused molecule libraries for drug discovery with recurrent neural networks. *Nature* **555**, 604–610 (2018).

[40] Vaswani, A. *et al.* Attention is all you need. *Adv. Neural Inf. Process. Syst.* **30**, 5998–6008 (2017).

[41] Noé, F., Tkatchenko, A., Müller, K.-R. & Clementi, C. Machine learning for molecular simulation. *Annu. Rev. Phys. Chem.* **71**, 361–390 (2020).

[42] Liu, Z. *et al.* Digital twins for materials development. *npj Comput. Mater.* **9**, 108 (2023).

[43] Xie, T., Fu, X., Ganea, O.-E., Barzilay, R. & Jaakkola, T. S. Crystal diffusion variational autoencoder for periodic material generation. *Int. Conf. Learn. Represent.* (2021).

[44] Sanchez-Gonzalez, A. *et al.* Learning to simulate complex physics with graph networks. *Int. Conf. Mach. Learn.* 8459–8468 (2020).

[45] Merchant, A. *et al.* Scaling deep learning for materials discovery. *Nature* **624**, 80–85 (2023).

[46] Butler, K. T. *et al.* GNoME: Discovering 2.2 million stable materials with graph networks. *Nature* **625**, 50–56 (2024).

[47] Schwaller, P. *et al.* Predicting retrosynthetic pathways using transformer-based models and beam search. *Chem. Sci.* **11**, 3316–3325 (2020).

[48] Coley, C. W. *et al.* A robotic platform for flow synthesis of organic compounds informed by AI planning. *Science* **365**, eaax1566 (2019).

[49] Hwangbo, J. *et al.* Learning agile and dynamic motor skills for legged robots. *Sci. Robot.* **4**, eaau5872 (2019).

[50] Burger, B. *et al.* A mobile robotic chemist. *Nature* **583**, 237–241 (2020).

[51] Coley, C. W. *et al.* A robotic platform for flow synthesis of organic compounds informed by AI planning. *Science* **365**, eaax1566 (2019).

[52] Muiños-Landin, S., Fischer, A., Holubec, V. & Cichos, F. Reinforcement learning with artificial microswimmers. *Sci. Robot.* **6**, eabd9285 (2021).

[53] Li, Z. *et al.* Fourier neural operator for parametric partial differential equations. *Int. Conf. Learn. Represent.* (2021).

[54] Bak, P., Tang, C. & Wiesenfeld, K. Self-organized criticality: An explanation of 1/f noise. *Phys. Rev. Lett.* **59**, 381–384 (1987).

[55] Grzybowski, B. A. & Huck, W. T. S. The nanotechnology of life-inspired systems. *Nat. Nanotechnol.* **11**, 585–592 (2016).

[56] Sutherland, J. D. The origin of life—out of the blue. *Angew. Chem. Int. Ed.* **56**, 11037–11043 (2017).

[57] Gánti, T. *The Principles of Life* (Oxford Univ. Press, 2003).

[58] Adamala, K. & Szostak, J. W. Nonenzymatic template-directed RNA synthesis inside model protocells. *Science* **342**, 1098–1100 (2013).

[59] Ghadiri, M. R. *et al.* Emergence of autocatalytic peptide replication. *Nature* **382**, 696–699 (1996).

[60] Rasmussen, S. *et al. Protocells: Bridging Nonliving and Living Matter* (MIT Press, 2009).

[61] Amato, F. *et al.* Fail-safe mechanisms in synthetic biological circuits. *Nat. Commun.* **12**, 5867 (2021).

[62] Nowak, M. A. *Evolutionary Dynamics* (Harvard Univ. Press, 2006).

[63] Szostak, J. W., Bartel, D. P. & Luisi, P. L. Synthesizing life. *Nature* **409**, 387–390 (2001).

[64] Adadi, A. & Berrada, M. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access* **6**, 52138–52160 (2018).

[65] Birch, J. Suffering and sentience. *Philos. Compass* **18**, e12899 (2023).

[66] Farah, P. D. *et al.* Ethical guidelines for engineering life. *Nanoethics* **16**, 1–15 (2022).

[67] Kassen, R. Societal risks from evolving microorganisms. *Nat. Biotechnol.* **37**, 618–620 (2019).

[68] King, R. D. *et al.* The automation of science. *Science* **324**, 85–89 (2009).

[69] Bostrom, N. Existential risk prevention as global priority. *Glob. Policy* **4**, 15–31 (2013).

[70] National Research Council. *Convergence: Facilitating Transdisciplinary Integration of Life Sciences, Physical Sciences, Engineering, and Beyond* (National Academies Press, 2014).

[71] Klein, J. T. *Interdisciplinarity: History, Theory, and Practice* (Wayne State Univ. Press, 1990).

[72] Langfeldt, L., Nedeva, M. & Sörlin, S. The role of peer review in interdisciplinary research. *Res. Eval.* **21**, 89–102 (2012).

[73] Gershenfeld, N. How to make almost anything: The digital fabrication revolution. *Foreign Aff.* **91**, 43–57 (2012).

[74] Wilkinson, M. D. *et al.* The FAIR Guiding Principles for scientific data management and stewardship. *Sci. Data* **3**, 160018 (2016).

[75] Jiménez-Luna, J., Grisoni, F. & Schneider, G. Drug discovery with explainable artificial intelligence. *Nat. Mach. Intell.* **2**, 573–584 (2020).

[76] European Commission. *Horizon Europe: Work Programme 2023–2024 – Cluster 4: Digital, Industry and Space* (European Commission, 2023).

[77] Rudin, C. Stop explaining black box machine learning models for high-stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* **1**, 206–215 (2019).

[78] Kitano, H. Nobel Turing Challenge: creating the engine for scientific discovery. *npj Syst. Biol. Appl.* **7**, 29 (2021).

[79] Masanet, E. *et al.* Recalibrating global data center energy-use estimates. *Science* **367**, 984–986 (2020).

[80] Brynjolfsson, E. & McAfee, A. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies* (W. W. Norton, 2014).

[81] Schwartz, R. *et al.* Green AI. *Commun. ACM* **63**, 54–63 (2020).

[82] European Commission. *Ethics Guidelines for Trustworthy AI* (European Commission, 2019).

[83] Lander, E. S. *et al.* Initial sequencing and analysis of the human genome. *Nature* **409**, 860–921 (2001).

[84] Landauer, R. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.* **5**, 183–191 (1961).

[85] Blaiszik, B. J. *et al.* Self-Healing Polymers and Composites. *Annu. Rev. Mater. Res.* **40**, 179–211 (2010).

[86] Cardelli, L. *et al.* Chemical Reaction Network Designs for Autonomous Molecular Machines. *ACS Synth. Biol.* **8**, 138–150 (2019).

[87] Li, S. *et al.* A DNA nanorobot functions as a cancer therapeutic in response to molecular triggers. *Nat. Biotechnol.* **36**, 258–264 (2018).

[88] Li, W. *et al.* Learning flexible and reusable locomotion skills for a soft robot. *Sci. Robot.* **6**, eabf3223 (2021).

[89] Zhang, X. *et al.* Closed-Loop Adaptive Materials via Embedded Sensors and Machine Learning. *Adv. Mater.* **34**, 2201768 (2022).