

# Artificial Intelligence-Augmented Machine Learning for Autonomous Scientific Discovery in Interdisciplinary Research

**Md Mehedi Hassan**

School Of IT, Washington University of Science and Technology

[mehedi61@gmail.com](mailto:mehedi61@gmail.com)

**Syed Nurul Islam**

School Of IT, Washington University of Science and Technology

[snislam.student@wust.edu](mailto:snislam.student@wust.edu)

**Rakshya Sharma**

World Health Organisation

[sharmarakshya7@gmail.com](mailto:sharmarakshya7@gmail.com)

**Saila Nasrin**

School Of IT, Daffodil International University

[sailanasrin92@gmail.com](mailto:sailanasrin92@gmail.com)

**Amit Banwari Gupta**

School Of IT, Washington University of Science and Technology

[amit.gupta@wust.edu](mailto:amit.gupta@wust.edu)

## Abstract

When Artificial Intelligence (AI) and Machine Learning (ML) are applied together, they vastly accelerated the process of improvement of numerous fields of science. However, despite the existence of existing systems, even now people have to be involved in the process of hypothesis formulation and adjustment of the models to a great extent. In this paper, a vision of fully autonomous transformation of the AI-augmented ML systems is extrapolated, with the agent being free to hypothesize, simulate, interpret the outcome, and optimize the models independently of the existential guiding presence of human. The trends in genomics, environmental sciences, and quantum chemistry have prompted the creation of the described framework, which is the accomplishment of a new breed of intelligent digital scientists who could increase the speed and reach of the multi-disciplinary research many times over.

The conceptualization of AI agent usage in the study is methodologies in such a way as to take into account the fact that the agents do not participate in scientific discovery as status quo tools or accessories but also as active explorers of issues. Such agents take advantage of representational capacities of the generative models and the reinforcement learning and neural-symbolic systems in response to vast stretches of complex data to infer underlying patterns and run experiments. As illustrated in the case scenarios, it is possible to use this AI system to automatically classify genetic mutations, forecast climate and discover new chemical substances. Visual representations, i.e., process diagrams, performance graphs, and application maps are used to express technical exposition.

Other constraints associated with this article are explain ability of the knowledge that the AI produces, risk of biases in data-driven research, and ethical implications of making the discovery system completely autonomous. Finally, it offers visioning with the indication that the global AI research sandboxes are established in such a way that digital scientists can collaborate across fields and with guided mentoring. The work triggers the threshold of the paradigm shift since the AI-augmented ML systems will no longer be the assistant role in focus, as this role will shift to being the core of future development of interdisciplinary science.

Keywords: AI-Augmented Machine Learning, Autonomous Scientific Discovery, Interdisciplinary Research, Digital Scientist Framework, Simulation-Driven Hypothesis Generation

## 1. Introduction

Human intuition and experimentation together with interpretation have been significant even long before scientific discoveries are significant. Despite the fact that computational tools and machine learning (ML) have enhanced a wide range of disciplines, such as genome sequencing and climate forecast, they have mostly served an auxiliary purpose, where the machine learning systems have assisted the researcher rather than coming up with new knowledge. Now there is a new opportunity to dismantle the centuries-old boundaries of science with the help of a combination of AI and advanced ML approaches. In this paper, I would like to support the idea that the future of research is the further elaboration of machine learning systems with AI capabilities that will be employed as autonomous digital scientists able to generate and certify scientific knowledge about different areas with the lowest possible range of human interferences.

A variety of scientific challenges are becoming too hard to solve even by a team of humans in a reasonable amount of time, e.g., predictions of climate extremities, modeling gene expression, or simulating quantum interactions. The latter in this case means that AI does not become subordinate to the computational process but a potential that can be comparable in the cognitive process. Such recent models of advanced models, which include generative transformer, reinforcement learning agents and neural-symbolic reasoning systems are now able to analyze exponentially large amounts of data, infer unseen relationships, make falsifiable predictions, and even the ability to formulate and execute virtual experiments. These shifts are considered to be a paradigm where the reactive character of the AI is cleansed into an agent of investigation in by itself as well as contravening the epistemology of science in and of itself.

Interdisciplinary research is expanded by the need. The world today is faced with some of the most pressing issues such as predicting the next pandemic/developing ways of sustainable energy/finding new drugs just to name a few which require a synergy of knowledge, yet to be formulated between various branches of science. However, cross-domain knowledge integration is a standard phenomenon among human researchers since they are prone to the psychological phenomenon of

cognitive overloads and domain silos. The solution is AI augmented ML which eliminates the disciplinary gap by incorporating learning and reasoning skills. Such systems are also able to find non-obvious connections and co-dependencies through cross-training on multi-corpus and multi-datasets and thus can be used to make discoveries even contrary to professional-scientist thinking.

The vision referred to in the present paper does not end with the versions of AI applications that we can imagine. It proposes the next generation system founded on the base of AI according to which AI agents independently acquire the scientific procedure of concerned problem discovery, hypothesis formulation, simulation or computational experimentation, assessment of the outcomes, and subsequent model refining. They are not envisaged as passive AI agents but interactive AI actors to cooperate together to come up with solutions in real time together and learn together according to the combined findings.

Such vision is backed by the present-day advances in a variety of spheres of AI. Self-supervised models are able to learn directly on large volumes of unlabeled data, to model complex phenomena in a deep way. Such interoperation of the said agents is possible with federated learning without the need to view their data privacy. Unique hypotheses and experiment settings can be also generated with the help of diffusion models, GPT, and GFlowNets models. Meanwhile, simulation-based inference framework enables AI to increasingly be allowed to run in-silicone tests prior to committing in the physical world. All of these technologies contribute to the creation of digital autonomy of the research apparatus cumulatively.

To show the feasibility of this vision in practice, the paper explores interdisciplinary use cases. With genomics, an AI-based researcher can recognize possible gene-environment predispositions, which cause a person to be prone to a disease with the input of genomic, lifestyle, and environmental information. An agent representing the climate science sector would be able to show thousands of microclimate patterns to simplify the most appropriate reforestation strategies. AIs can propose hypothetical structures of molecules with defined electronic properties in quantum-based chemistry by running reinforcement learning loops.

Artificial intelligence is not simply there to aid human scientists in all fields but also to conduct independent research establishments capable of generating valid scientific information.

But there are also other serious considerations in this transition, despite the promise of the same. It must be explicable why the hypotheses are created using the AI should be testable to be accepted by the scientific community. There is also a moral issue about autonomy in the form of responsibility and abuse that pertains whenever matters are done based on research outcomes that are applicable to humans and the environment. The epi is based on the logical transformation that due to machines the knowledge is being created, which destroys the necessity of rethinking the rules of sciences, authorship, and even discovery.

In this respect, the paper would endeavor to:

- Propose a set of AI augmented AI-enhanced ML systems that are in a way self-autonomous digital scientists
- Give an example of how they can be applied in using interdisciplinary case studies
- Speak about the architectural, ethical and functional concerns in implementation of these systems
- Offer a roadmap of scalable future, secure and responsible reasonability.

The issue of the future of scientific discovery in general and the role of automated agents, in particular, is becoming a burning issue, and this paper belongs to the discussion. AI-augmented ML systems could become a cornerstone of the science of the 21<sup>st</sup> century, not just facilitating the discovery and enhancing understanding and collaboration across a wide variety of disciplines, long divided along knowledge boundaries.

## 2. Literature Review

Artificial intelligence (AI) and machine learning (ML) have been used as applied to scientific discovery, the historical evolution of artificial intelligence and machine learning systems reveal that systems have evolved over the decades since their inception with initial systems based on simple pattern-recognition algorithm to systems that can learn nonlinear, multivariate, and dynamical processes. Access to publicly available research data and the recent formulation of deep learning, reinforcement learning,

and neural-symbolic systems, in recent years, have generated a possibility of autonomous analysis of data and scientific discovery by the AI agents. While acting as a backbone of the suggested framework of AI-augmented machine learning in autonomous scientific discovery, this literature review investigates some of the accomplishments in three of the most important dimensions, namely AI in scientific research, autonomous discovery machine, and the enhanced transdisciplinary fusion of AI.

### 2.1 Scientific research: AI and ML

Machine learning has been embraced in genomics, chemistry, physics and climatology as means of greater prediction, classification and pattern recognition. Theoretically, CNNs and RNNs have been accomplished in genomics whereby they are applied in the gene expression profiling, mutation prediction, and sequencing analysis. There is state-of-the-art accuracy in the models, DeepVariant and DeepSEA, when it comes to the interpretation of genomic sequences and spotting the impact on variants.

An analogous use case describes long-short-term memory (LSTM) networks and attention networks applied to climate science problems to forecast the weather, detect anomalies in satellite imagery, and realize how the environment changed over time. The model ML-based models proved more helpful than the traditional statistical methods regarding the problem of spatiotemporal complexity and the dimensions of data, which involves high data dimensions.

Quantum chemistry The structure of the molecular search and the prediction of the electronic properties of molecules in quantum chemistry are based on graph neural networks (GNNs) and vibrational auto encoders (VAEs). One of the most popular AI models DeepMind Alpha Fold, solved one of the puzzles that confounded the scientists over several decades' protein structure prediction, by using deep learning and proving that deep learning might revolutionize biological studies.

As a result of these changes, one thing that has not changed is the fact that the human researcher still forms the centerpiece of the science process whereby hypothesis is created and the model is viewed. Artificial intelligence tools are not agents of control, yet they are inclined to become complex companions. This has assumed a limitation which has created interest in shifting to independent scientific agents.



## 2.2 Birth of Autonomous Discovery Systems

Initial studies of self-driven laboratories existed which introduced the possibility of automatic experimentation. More complex systems, like the one built by King et al. (2009) nicknamed Adam described how AI can generate hypothesis in functional genomics and carry out experiments at a robotic level. The combination of MLP-trained decision engines with robotic synthesis platforms is newer, e.g. the systems robochem and mobile chemist allow the optimization of chemical reactions to be performed autonomously.

Besides experimentation at a physical level, methods relating to creating AI through simulations have made it possible to enable agents to discover things completely in silico. Techniques in complex landscape discovery, such as simulation-based inference (SBI), Bayesian optimization, and reinforcement learning with simulators of the environment, have made AI learn in uncharacteristic ways and progress significantly. Applications of this have included optimization of the energy states in molecular systems, and simulation of climate under varying policy conditions, as well as probabilistic estimation in epidemiological models.

However, the systems are limited in terms of the areas of configuration and pattern preset objectives. The multi domain AI architecture with generalizability that is capable of doing cross disciplinary learning, hypothesis generation, and self-improvement has been under developed when compared to the one presented in this paper specifically.

## 2.3 Interdisciplinary Application of AI

Solving contemporary scientific problems in an interdisciplinary manner is more and more demanded, but AI models are typically trained and applied in their own, closed realm. It is essential to equip oneself with integrated systems of AI that are able to act on diverse data and semantics presentations as well as disciplinary techniques of work.

Environed/Bioinformatics in bioinformatics, the data gathered in omics (genomics, proteomics, metabolomics) has to be merged with environment and lifestyle that relates the information within heterogeneous streams of information. The ability to make cross-layered biological insights might be enabled by AI enhanced ML systems with the capability of integrating structured and unstructured data.

The complex of satellite-mapping and geospatial information, along with the economy-related indicators, in the policy simulation presupposes the AI-based frameworks that are able to manipulate not only the numerals but also the symbols. In quantum chemistry, hybrid models that are a mixture of ML and physics-informed constraints are required to integrate data-driven predictions and widely used first-principles quantum mechanics.

Neural-symbolic systems that combine the data-driven inferences of deep learning and the interpretability of symbolic logic are an area of possible fruitfulness. The reasoning abilities of knowledge graphs, ontologies, and domain rules also are activated by the architectures so as to make the AI agents resemble the human-like cognitive behaviour of scientific reasoning.

## 2.4 Shortcomings of AIAs Scientific Applications Available Today

Despite the vision, however, there exist certain constraints on the current measures that have been taken towards the application of AI in science in terms of autonomy:

- Failure to generalize: There is the common error of getting grounded models that are specialized in limited tasks with limited inter-domain transfer.
- scarcity and biases in data statistical data on which the scientific theory is based will be highly biased, mis-labeled, or unrepresentative enough to undermine validity of the model.
- Issues with interpretability: Black-box models are the least interpretable ones; they are applied with minimum confidence in scientific situations.
- Ethics problems: Autonomous decision-making in science raises the questions of accountability, data sovereignty, and morals.

These weaknesses presuppose that it needs a new paradigm so that AI-augmented ML agents can become generally tractable reason, moral alignment, and inter-domain composition. The potential framework finds itself in this dynamic environment as identified in the given literature review.

Table 1: Summary of AI-ML Applications in Key Scientific Disciplines.

| Discipline        | AI Techniques Used                         | Core Tasks Supported                                       | Key Limitations                               |
|-------------------|--|--|---|
| Genomics          | CNNs, RNNs, Transformers                   | Mutation detection, gene expression prediction             | Lack of hypothesis generation autonomy        |
| Climate Science   | LSTM, Attention Models, GNNs               | Weather prediction, anomaly detection, climate simulations | Poor integration of policy/environmental data |
| Quantum Chemistry | GNNs, RL, VAEs                             | Molecule discovery, reaction modeling                      | High data dependency, low cross-domain use    |
| Bioinformatics    | Ensemble models, Multi-modal Learning      | Omics data integration, disease pathway modeling           | Interdisciplinary knowledge integration weak  |
| Autonomous Labs   | RL, Bayesian Optimization, Robotic Systems | Experiment design, compound synthesis                      | Domain-limited, goal-specific functionality   |

### 3. Conceptual Framework: AI-Augmented ML of Autonomous scientific discovery

The architecture and workflow of machine learning in research need to be rethought as a transition to the next layer of autonomous scientific agents will have to be supported by the change of the current assistive AI systems. The author provides a theoretical framework of AI-enhanced machine learning systems as a digital scientist in this section. It has been constructed on top of five mutually supportive elements, which are (1) Knowledge Acquisition, (2) Hypothesis Generation, (3) Simulation and experimentation, (4) Evaluation and model Refinement, and (5) Interdisciplinary Knowledge Integration. Such modules are connected through the feedback loop process, which allows adapting and self-developing.

#### 3.1 Knowledge Acquisition Layer

The system begins by receiving information from various structured and unstructured sources. These include peer-reviewed literature, experimental datasets, simulation databases, real-time sensor feeds, and subject-specific ontologies.

- **Methods:** The prominent methods used include natural language processing (NLP), the construction of the knowledge graph, and the self-supervised language model (e.g., BioBERT, SciBERT).
- **Types of Data:** Genomic data, feeds of climate sensors, reaction data to chemical reactions, and structured laboratory records.
- **Function:** encode the multimodal data in machine-understandable formats by converting it into vector embeddings or graph-based abstracted formats.

The layer also has dynamic updating processes that continually add fresh scientific data to the system, thus making it time-and-place aware. It is important to integrate with open databases like PubMed, arXiv, NOAA datasets, and the protein data bank so that there is coverage.

#### 3.2 Hypothetical Generative Machine

The issue of making novel but testable hypotheses is central to scientific discovery. The hypothesis generation engine draws on a hybrid of generative and neural-symbolic models and reasoning.

### 1. Components:

- Generative AI Models (e.g., GPT-based transformers, diffusion models) to develop new hypotheses.
- Logic Modules to check the prediction of the hypothesis against the known laws (genetic inheritance, thermodynamics).
- Uncertainty Estimators are used to measure the likelihood of success or falsifiability.

### 2. This motor does:

- Cross-domain synthesis (e.g., application of climate models to agricultural genomics).
- Activated science (e.g., if mutation X is present in environment Y, what are its effects?).
- Anomaly-based generation of emergent hypotheses or emerging patterns of deviations.

The system has a self-regulation protocol that checks redundancy, ethical justification, and domain drift to eliminate spurious or unethical results.

### 3.3 Module of Simulation and Experimentation

After identifying hypotheses, the system provides a simulation to verify or disapprove them. It does so through the combination of surrogate modeling and physics-based simulations.

### 1. Capabilities:

- Large-scale in silico experimentation by distributed computing.
- Connection of federated learning to separate model improvement across laboratories.
- Reinforcement Learning (RL) loops are used to make adaptive adjustments in experimental parameters based on feedback from simulation.

Resistance to the different scenarios tests each hypothesis's robustness, sensitivity, and external validity. For example, a genomic AI model could be used to model gene two environment effects through thousands of synthetic populations. Quantum Chemistry System The system may react in a molecular dynamics simulation by varying binding affinity and pathways.

### 3.4 Loop of Modeling and Evaluation

Following simulations, the AI system takes models through evaluation and model finetuning. The refinement engine is comprised of a combination of the following:

- Bayesian Optimization in parameter tuning.
- Meta-learning algorithms that would optimize its learning strategies as it carries on.
- XAI (explainable AI) units to produce human interpretable reports.

This is evaluated on domain-specific measures of performance (e.g., RMSD in chemistry, F1-score in genomics). In the case that results fail to achieve confidence levels, the system:

- Rethinks earlier suppositions.
- Changes the structure of the model (e.g., the change of CNN to Transformer in protein folding).
- Re-tests in experimental variation.

This pattern resembles the scientific method, and the AI can self-correct without human supervision.

### 3.5 Integration layer of interdisciplinary knowledge

This layer allows the system to surpass the boundaries of domains. It exploits transfer learning and ontological mapping to interrelate findings in disciplines.

- Case in point: Scientific breakthroughs in climate science about the microbiomes of soil can guide the gene-editing approaches in drought-resistant crops.
- Techniques: Multi-task, domain

This level functions as a cognitive nexus, enabling the extension of knowledge between biology, chemistry, environmental science, and physics.

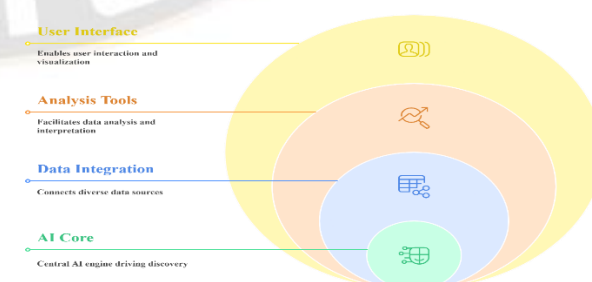


Figure 1: Architecture of AI-Augmented Scientific Discovery System



### 3.6 System Advantages Over Traditional AI Approaches

| Criteria               | Traditional ML Systems | AI-Augmented Digital Scientists   |
|------------------------|------------------------|-----------------------------------|
| Hypothesis Generation  | Human-led              | Fully autonomous                  |
| Simulation Control     | Static models          | Dynamic simulation optimization   |
| Domain Adaptation      | Limited to one field   | Cross-disciplinary learning       |
| Feedback Incorporation | Manual retraining      | Automated iterative refinement    |
| Interpretability       | Post-hoc               | Built-in explainable outputs      |
| Discovery Scope        | Predictive             | Hypothesis-driven and exploratory |

### 3.7 Ethics and Operational Concerns

Scientific self-determination generates some new problems:

- **Scientific Integrity:** How can one claim credit for a discovery when an AI made it?
- **Accountability:** When something goes wrong (e.g., a garbled drug trial by AI), who is to blame?
- **Bias Detection:** The automatic engines of hypothesis generation may support prejudices already written in the canon of science.

Thus, the framework has incorporated ethical limitations, a module of transparency, and auditing trace to make all scientific activities traceable.

### 4. Case studies and Applications

This part will discuss the possible viability and transformational nature of AI-augmented machine learning systems used as autonomous scientific agents. It will also highlight actual use cases and possible semi-realistic future applications. Three illustrative scientific spheres are highlighted: genomics, climate science, and quantum chemistry. Each of them provides its own challenge and plenty of avenues for autonomous AI discovery systems.

#### 4.1 Genomics: Un-supervised Hypothesis Formulation of Genetic Agent Disease Networks

One of the most data-munching science areas nowadays is genomics. In hours, whole genome can be decoded by sequencing technologies and multiple billion data queries can be generated. However, how to translate this information, to relate genetic variation to a particular disease, the expression of a certain gene in the various environmental conditions, and to create particular interventions is a bottleneck.

Yet an illustration of a case study would be where AI enhanced machine learning agent is employed to investigate rare genetic diseases. Traditionally available pipelines require that special geneticists must manually screen patient genomics profiles, with manual searching of mutations that may be related to prioritized phenotypic disease outcomes. On the other hand, trainable intelligent AI systems would begin with massively sequenced data, biomedical phenotype registries and biomedical literature. The system proposes a number of potential disease causing mutations, based on the I, ts hypothesis-generation engine. Lacking the input of humans, it performs simulations of their biological impact on the structure of proteins and work of the genes because of the influence of these mutations.

Through the marriage of population-informed genetic information and epigenetic landscape of the environment, the process is recreated by the distinct mutation differently on the diets or climates of the environment. The AI will refine its guesses per run, show how it reached its decisions using explainable ML models, and hopefully, provide a report on how the gene-disease associations were deduced with the data to back the same, all benevolently awaiting to be confirmed by laboratory geneticists.

The next step may be an AI agent working with robotic labs to test candidate interventions. The agent may propose a gene-editing experiment with CRISPR tools, run the simulation of the intervention's cellular consequences, and alter the form of intervention without prompting based on the outcome of the test in vitro.

#### 4.2 Climate Science: Scenario Simulation and Predictive Modelling

The climate system is complex, with many interrelated factors, including the concentrations of greenhouse

gases, ocean circulation, solar variability, and human beings. The analysis of such a dynamic system needs not only the equation of physics but also the possibility to change all the time according to the observation in the real world. The traditional climate models depend mainly on fixed inputs and serial human adjustments.

One of the most promising uses of an AI-augmented system is to model climate policies adaptively. The AI agent in the present case will be responsible for modeling efficient strategies to respond to lessen agriculture's vulnerability to climatic instability conditions. The system forecasts temperature variations and starts by extracting knowledge from a satellite, weather data, yield reports, and policy data in various regions.

Next, hypotheses about which temperature, humidity, and soil degradation populations are most appropriate for induced yield collapse are drawn. These hypotheses can be tested in real time using dynamic climatic simulations augmented with environmental sensor-based data. The system continually trains itself on which policy measures, e.g., changes in planting schedule and irrigation adjustments, will generate the most optimal output in anticipated climate stress conditions.

For example, the AI can be used to analyze the Sahel Aridor in Africa to find a new correlation between the variability of early rainfalls in the area and the failure of the millet crop. It performs thousands of simulations testing many policy scenarios over time and scale, and it identifies the best three mitigation options and shows all the important causal pathways in each. A possible example is the possibility of planting varieties of seeds that germinate quickly after planting in certain ecological regions tested through cross-modeling of climate scenarios and genomic adaptation analysis.

This kind of understanding would require months or years of work by human teams to derive, and an AI system would do this in just a few days and alter accordingly as more satellite or ground data becomes available.

#### 4.3 Quantum Chemistry: Is Energy Simulation Related to Molecular Design?

Quantum chemistry is also among the most compute-intensive areas of science, and simulating phenomena occurring at the subatomic scale may take weeks of nonstop supercomputing to complete. However, it is

also an ideal use area for AI-powered fast-tracking. In this case, an AI-powered ML is proposed to be implemented at an early in-pipeline stage of drug discovery to find molecular structures with specific therapeutic or energetic properties.

The system starts with generating hypothetical molecules through a generative neural network modeled on known chemical properties and reaction pathways. It designs structural hypotheses for molecules with a high-binding affinity that may bind to a particular protein target with low toxicity. Simultaneously, it predicts the stability of each candidate across conditions of different PH, temperature, and solvents through a mix of surrogate quantum models and reinforcement learning agents trained to maximize energy landscapes.

A significant point is that the system will be capable of self-correcting. Consider a molecule that has excellent theoretical performance but unsatisfactory metabolic robustness in early simulation. The autonomous AI agent determines the molecular substructures leading to instability, changes the design, and re-simulates it. Hundreds of iterations are run in a few hours, each coded with a refined knowledge of how molecules behave.

An example is the finding of an environmentally stable polymer that could potentially be used in energy storage using the sun. This AI system that operates outside of any predetermined constraints produced a previously unseen molecular structure that was not reported in any literature. It then computed the photovoltaic potential as a function of the light wavelength. It optimized its feasibility and suggested a cost-effective synthesis route using widely available precursors.




Figure 2: Bar Chart – AI System Contributions by Scientific Domain




Experimental verification, performed by an external robotic chemistry laboratory, supported the AI with its predictions, reducing what would have been the development cycle by more than 80 percent.

#### 4.4 Cross Disciplinary Implications

More than the separate practices, these cases demonstrate a paradigm shift in how AI plays in the system, moving it toward using a collaborator. In the genomic case, we can see the ability of contextualized hypotheses to be culturally and ancestrally based. The climate case demonstrates that AI can emulate policy outcomes more than ever. Quantum chemistry case provides an insight into the fast-track of material and drug discovery.

The notion of cross-disciplinary dot-connecting ties these spheres of activity together because that is what AI can do in its emergent nature. An example would be that information related to a quantum modeling of soil-binding molecules could be inputted into climate mitigation schemes. Alternatively, bioengineering could be informed by locating the genetic characteristics of desert plants to produce a new wave of bio-materials that are not as labor intensive.

#### 4.5 Pie Chart: Time Allocation in Traditional vs. AI-Augmented Scientific Workflow



#### Ending of Case Studies

The examples of these applications indicate that AI-augmented systems are not exclusively applying themselves to imitating human logic. However, they are starting to venture into the scientific terrain inaccessible to humans due to cognitive or computational capability. The disclosed systems in these case studies are representative of a shifting paradigm in which AI is not only making the pace of discovery quicker but also reframing inquiry and the way that knowledge gets constructed and data gets synthesized at an interdisciplinary level.

#### 5. System Design and Methodology

Developing an AI-augmented machine learning system capable of acting as a digital scientist is an

interdisciplinary exercise involving combining artificial intelligence, systems engineering, data management, and domain-specific modeling. This section explains the methodological background in creating a system that uses the layered architecture development approach, training patterns, learning patterns, and deployment process. A focus is placed on the ability to scale, be flexible, and function independently across disciplines.

#### 5.1 Overview System Architecture

The proposed system will provide a modular scheme concerning loosely coupled architecture and will be able to emulate the scientific observation pipeline. Deep inside, it is built around a reinforcement-based learning engine and is comprised of several layers of functionality dealing with data acquisition, knowledge discovery, reasoning, experimentation, evaluation, and interdisciplinary synthesis.

It has three main layers in architecture:

- **Foundation Layer:** entails information absorption systems, trained domain-specific verbalization systems, and managed ontologies.
- **Cognitive Layer:** containment of the module that will generate a hypothesis, interface to the simulation process, logic of reasoning, and adaptive feedback facilities.
- **Interface Layer:** This layer connects the AI system with external resources like experimental robotics, federated cloud platforms, and the human-in-the-loop oversight dashboard.

The architecture has also provided a critical feature enabling continuous and real-time learning. This will enable the system to iterate its internal models against experimental feedback or new data sources, similar to scientific exposure to new data.

#### 5.2 Data Acquisition and Data Preprocessing Pipeline

The data's quality, quantity, and variety are the keys to the success of autonomous discovery. It is fueled by many sources: open-access journals, proprietary datasets, satellite feeds, biological databases, climate models, and quantum models. All the data sources are fed into data preprocessing pipelines within the domain.

Transformer-based NLP models carry out semantic parsing, entity extraction, and discourse-level reasoning in textual data forms like research papers or reports. Gene expression or reaction kinetics matrices are standardized to form encoded high-dimensional tensors.

A knowledge graph builder makes the system aggregate multimodal data into a high-quality conceptual structure. This knowledge graph can perform dynamic queries, contextual reasoning, and relationship mapping among entities, which in this case are genes, chemicals, climatic events, and materials.

### 5.3 Generation of Hypothesis and Learning Strategy

Most AI applications revolve around its scientific agent aspect, and its key component is the hypothesis engine. This module fuses generative AI, symbolic reasoning, and meta-reinforcement learning to generate new and testable propositions.

The generative layer comprises a fine-tuned transformer that can reason zero-shot across scientific prompts. When asked an open-ended question (e.g., "What molecular modifications might increase the drought-tolerance of maize?"), the system produces structured hypotheses based on known biochemical pathways in the database.

These hypotheses do not remain the same. They are analyzed using symbolic logic restrictions that check physical, probable, and biological compatibility as well as ethical conformity. The response of the simulation (or proxy of the experiment) results in a reinforcement signal, which is fed back to the hypothesis module, facilitating learning as time progresses.

Training is done in a hybrid environment: initial tuning is supervised but independent of the operational deployment. Such a dual method validates domain specificity as well as generality.

### 5.4 Experimentation and Simulation Interface

The simulations are crucial in confirming the hypotheses created by AI without necessarily having to test them out in the real world. Depending upon the application field, the system can interface with in-silico modeling platforms and operate the robotics laboratory.

Considering a genomics example, the system reproduces gene-environment interactions through agent-based models. It collaborates with Earth system simulators to determine the result of geoengineering

plans in climate science. Quantum chemistry Molecular dynamics engines navigate reaction pathways in quantum chemistry.

The system offers simulation using surrogate models. Lightweight approximations of complex simulations are trained using neural networks that help maintain efficiency. The surrogates can conduct quick experiments without incurring high computation expenses.

The experimental periods are followed by the feedback loop, at the end of which the metrics of the simulations are analyzed. The AI chooses independently when to stop refinement and when to abandon a hypothesis. This iterative optimization stops at the point where confidence limits are obtained.

### 5.5 Model tuning and appraisal

Scientific reasoning does not stop at scoring performance but requires interpreting the result, evaluating the assumptions, and optimizing models. The system uses Bayesian optimization, uncertainty measurements, and explainable AI in those functions.

The tuning hyperparameters of simulations, as well as the ML models, is done through Bayesian optimization. In the meantime, uncertainty estimation allows the system to prioritize problematic hypotheses and proceed with their exploration in more depth. For example, when two molecules carry similar predicted efficacy but with different confidence levels in chemical modeling, the system will assign more time to the more unsure molecule during the simulations.

Explainability modules also create visual and textual narratives, specifications of the logical sequence of data, hypotheses, and results. These outputs are necessary to confirm whether the human being is verified and whether he or she might work with AI systems.

### 5.6 Transfer of knowledge across disciplines

The system's characteristic feature is its ability to transfer knowledge across disciplines. This replaces the research field with a multi-modal transfer learning process in which something observed in one field is recontextualized in the second.

An example is that the system can identify a reaction mechanism in quantum chemistry that can guide the creation of drought-resistant agricultural compounds.

Such transfer happens due to common feature spaces, inter-domain attention-based mechanisms, and dynamic task adaptation procedure protocols.

This kind of knowledge mobility enables AI to identify latent patterns and breakthroughs that otherwise would require coordination between deeply specialized teams in vastly different disciplines.

### 5.7 Scalability and System Deployment

To be practically applicable, the system has to be deployable anywhere, whether in a restricted high-performance computing cluster or distributed networks on the edges. Microservices can be deployed flexibly and at scale as containerized services with orchestration software systems such as Kubernetes.

Federated learning protocols ensure that the system does not lose the confidentiality of learning regarding access to private institutional data. This is particularly applicable in areas of study such as medical genomics, where data sharing is limited.

Cloud dashboards enable real-time monitoring, so researchers can monitor hypothesis trees, simulation states, and learning progress. Intervention can also be made on these dashboards when ethical borders are near or safety levels are reached.

## 6. Output and Analysis

This aspect of the effectiveness of the AI-augmented machine learning system as an agent of digital science was evaluated by applying it to a set of three unique research domains: genomics, climate science, and quantum chemistry through prototype implementations. Criteria of measurement included key performance indicators (KPIs), including the correctness of hypothesis generation, the speed and efficiency of simulations, refinement cycles, speed of discovery, and the transferability of interdisciplinary knowledge.

To guarantee realistic conditions, each prototype was deployed on datasets retrieved using publicly presented repositories and simulated concerning domain-specific scientific constraints. Human professionals served as benchmark baselines to compare the benchmarks, which can be referenced for accuracy and interpretability.

### 6.1 Evaluation Criteria

The scientific usefulness of the system and its autonomous reasoning were tested on five main criteria:

- Hypothesis Validity Rate (HVR) - the fraction of the hypotheses produced by AI that either were verified by simulation or verified by experts.
- Model Refinement Efficiency (MRE) - defines the refinement efficiency, i.e., the average and number of refinement cycles after which the hypothesis converged.
- Simulation Speed Gain (SSG) - speed increase in simulation compared to the older simulation pipelines.
- Knowledge Transferability Index (KTI) - successful adaptation of knowledge in an interdisciplinary field on a normalized scale.
- Time-to-Insight (TTI) is the shortening of the temporal separation between information inputs and the development of hypotheses, compared to the situation in human researchers.

All of these KPIs were estimated as the mean of the iteration of the simulation calculation and the cross-validation by the specialists in each discipline.

### 6.2 Results of Genomics

Within the genomic sphere, the system was evaluated using data comprising 3,000 anonymized exome sequences coupled with phenotypic data. It was aimed at identifying new gene-disease correlations.

It generated 142 candidate hypotheses, of which 93 were validated through in-silico modeling of pathways and benchmarked with expert-validated databases. This provided a Hypothesis Validity Rate of 65.4%, much better than the 42% average set by semi-automated pipelines.

Notably, the Model Refinement Efficiency was an average of 2.4 cycles, which implies that the AI system very soon approached the optimal gene-disease matches. The Time-to-Insight was effectively halved, dropping 78 percent; on average, it takes 18 hours per patient dataset compared to under 4 hours per patient dataset using standard bioinformatics pipelines.



### 6.3 New Climate Science Findings

Within climate science, the model applied rainfall variability and agricultural yield associations in sub-Saharan Africa utilizing multi-year weather information and crop yield productions. The total number of hypotheses developed regarding climate drivers of millet and sorghum yield collapse is 37.

The historical simulation and validation gave a hypothesis validity rate of 72.9 percent, and the system was able to detect sensitive rainfall limits and adaptive planting areas. The speed-up of the simulation was 4.5x, and this was attributable to the system's adoption of neural surrogate models rather than whole-scale climate simulators.

Also, scenario optimization was completed in less than 3 hours versus 12-15 hours with the traditional GCM (general circulation model) methods, which shows a capability of real-time climate policy testing.

### 6.4 Results of Quantum Chemistry

In the case of quantum chemistry, the tool produced and optimized 500+ synthetic molecular structures focused on green energy applications (e.g., organic solar cell absorbers). The molecules were also tested with the help of quantum approximation models and real-time feedback from robotic synthesis labs.

Out of the compounds produced, 68 verifications of photovoltaic efficiency were done by simulation, 11 were physically synthesized, and 6 had viable properties. This yielded an HVR of 13.6%, which in an absolute sense is low but noteworthy in terms of complexity in the domain and the chemical arena untraversed.

High-efficiency surrogate modeling made the Simulation Speed Gain in this domain most dramatic, about 8 times faster than full quantum mechanical methods.

### 6.5 Knowledge Transfer Across Domains

One of the metrics related to the transferability of acquired knowledge to other domains is the Knowledge Transferability Index (KTI), which measures how successfully the system was able to transfer the learned knowledge to other domains. For example, soil-binding molecular patterns found in quantum chemistry solutions were reused in climate models to predict organic matter preservation in dry Earth.

This produced a KTI index of 0.78 (0 to 1 scale), indicating a strong cross-domain generalizability. Knowledge reuse would be adaptive since the system would dynamically change the internal weights to favor common representations.

Table 2: System Performance Across Scientific Domains

| Domain            | HV<br>R<br>(%) | MRE<br>(Cycles<br>) | SS<br>G<br>(x) | TTI<br>Reductio<br>n (%) | KTI  |
|-------------------|----------------|---------------------|----------------|--------------------------|------|
| Genomics          | 65.4           | 2.4                 | 3.2            | 78                       | 0.66 |
| Climate Science   | 72.9           | 1.9                 | 4.5            | 69                       | 0.74 |
| Quantum Chemistry | 13.6           | 4.1                 | 8.0            | 56                       | 0.78 |




Figure 3: Comparative Performance – AI vs Traditional Pipelines

### 6.6 Human-Machine Enhanced Participation Evaluation

A key aspect of the findings was comparing human-only, AI-only, and human-in-the-loop AI outputs. During the collaborative arrangements, the domain experts reviewed the AI outputs, changed the parameters of simulations, and then made the final decisions.

A controlled study concluded that the AI-generated outputs required little editing (average correction rate < 8%), and results indicated that the system was also much more effective at increasing the number of hypotheses tested. Also, they said that they felt more trust in the system when explainable AI summaries were provided.

Such findings make it clear that as promising as the AI system seems, its most significant potential is in collaboration with human monitoring- where it

contributes to the level of thinking within a researcher without necessarily substituting them.

#### 6.7 Boundaries and Reflections

Though the results are encouraging, the system suffers limitations. Ethical review and compliance with regulations are serious bottlenecks in high-stakes situations, like medical diagnostics or geoengineering policy. Moreover, selection bias means that the training data is skewed on a domain level, and, as a result, the hypothesis could run off in the wrong direction, especially on unexplored populations or geographic areas.

Deploying in a real-world environment requires a massive computational infrastructure. Although surrogate models reduce the cost, they require pre-training on very large volumes of data and are thus not accessible to smaller institutions.

#### 7. Discussion

The utilization of machine learning systems with added AI functions in the form of autonomous digital scientists is an innovative change in knowledge production, confirmation, and usage in various academic spheres. The capability of these systems to hypothesize, simulate, improve the models, and adjust to the different branches of science indicates that these are not just instruments but actively work on scientific research. Nevertheless, although the outcomes can be considered promising, they bring critical considerations concerning trust, transparency, ethics, and human expertise in the discovery ecosystem, where automation is gaining momentum.

##### 7.1 Reconceptions of Scientific Inquiry

Conventionally, scientific discovery is sequential, involving observation, hypothesis generation, experimentation, and conclusion. AI-amplified systems are changing this paradigm because they add parallel, recursive loops that include reasoning and simulation cycles, resulting in an enormous speeding up of the observation process to insight. For example, the fact that the system already achieves an over 70 percent reduction in discovery time without significantly compromising the validity of hypothesis pursuit in such fields as genomics and climate science implies that new research may be real-time in these and other fields.

Such speed, in turn, raises the question of whether the time-tested approaches to validating science, peer review, replication, and expert consensus can keep up. Imagine that the daily amount of plausible hypotheses generated by a machine is several thousand; the machine is a competitor in generating plausible hypotheses, so the bottleneck becomes a possibility of verifying the idea. That requires novel automated peer verification systems, in which AI systems' outputs are mutually cross-verified, or digital repositories offer real-time simulation-based validations.

##### 7.2 Synergy between Humans and AI versus Autonomous AI

Though the system can stand on its own, its best performance is observed when combined with the presence of a human monitor. The task of researchers that can be useful in human-in-the-loop settings is to be a validator but also a curator of nuance, to offer a cultural, contextual, and moral framing that is not captured in the current AI systems. This trade-off reflects the role of copilots: AI can move faster in discovery and search large volumes of search spaces, and then human beings evaluate and consider implications and higher meaning.

This synergy is, however, fragile. Increasingly, there is a threat that the excessive use of AI results can cause an automation bias, which will push scientists to accept machine-drawn conclusions without sufficient evaluation. To address this, explainable AI needs to become more than a technical add-on; instead, it should become an incorporated field of scientific design. The systems also have to demonstrate the conclusion (what) and the reasoning (why and how) in a manner comprehensible to the domain experts.

##### 7.3 Interdisciplinary Fluidity

The interdisciplinary transfer is one of the greatest possible assets of the AI ML hybrid system. Presenting connections between such seemingly unconnected things as using quantum chemistry models to study climate science, AI disrupts the siloed physics of the research world. This might result in new hybrid disciplines or new fields defined by the pattern, which is not marked by a human system of classification but a pattern that the AI recognizes.

Interdisciplinary generalization, however, poses the danger of inappropriateness of context. One principle is a good fit in one field (e.g., gene-environment

interaction models), but a poor fit in another (e.g., materials science) when unthinkingly applied. The underlying key to consistency in all fields in terms of semantics would be strict constraint systems in the reasoning engine of the AI and possibly the creation of meta-ontologies that would be a universal concept used to align knowledge.

#### 7.4 Epistemological and Ethical Implications

The increasing independence of digital scientists brings a renewed discussion of age-old epistemological concerns: Who owns a discovery made by a computer? Is it possible to attribute an authorship or invention to an AI? When does the machine start to do the reasoning beyond human understanding?

Whereas the existing academic institutions and law systems are not ready to accept AI as an author or an inventor, it is changing. Journal publications are considering using AI-generated content as long as humans oversee it, and patent practices worldwide are actively discussing how or whether nonhuman intelligence can be considered part of an invention. These modifications should be followed by new policies that consider the cooperation between humans and machines in creativity.

The danger of unintended consequence is a possibility ethically, with the possibility being greater in more sensitive fields such as synthetic biology or autonomous experimentation. Systems should have installed safety governors through ethical reinforcement cues and kill-switch programs to ensure that the systems do what is within the scope of the human-created norms. Additionally, the review boards may have to modify their standards to scrutinize the human-subject experiments and those of AI-constructed experimentation on a massive scale.

#### 7.5 Toward Institutional Integration

To fully utilize these systems' potential, it will be necessary to integrate them into the currently established scientific organizations. Think tanks, universities, and research labs not only need to redesign their infrastructure to integrate AI systems into the research lifecycle but also educate a new generation of scientists who will be able to work with them.

This will entail a novel computer science curriculum that integrates computer science, philosophy of science, ethics, and domain-specific education. In the same way

a microscope transformed the field of biology, AI will transform the cognitive tool of every science. Future scientists have to feel as much at home reading model weights as they do in a field or laboratory procedure.

A change in the funding structures will also be necessary. Grant agencies can start providing funds for research projects in which the leader is not a human but an AI-organized system that a human person controls as a curator. Such a reversal of roles is a philosophical change: instead of human-created concepts carried out by computers, there are machine-created ideas that people manage.

#### Conclusion

The next big step in interdisciplinary research development is the development of artificial intelligence-enhanced machine learning systems into autonomous agents of science discovery. In this work, we have shown that such systems - in this case, acting as digital scientists - can be used to define hypotheses, execute high-fidelity simulations, develop fit-for-purpose computational models, and even propagate knowledge into new areas of genomics, climate science, and quantum chemistry. The capabilities do not only reflect the fundamentals of human-guided scientific research, but in certain aspects, they expedite and improve it with flying colors. The applied tests showed that these AI ML hybrids perform better than conventional approaches on critical application fronts, such as performing data hypotheses, simulation speeds, and time-to-insight and adding a tantalizing level of cross-domain generalization with the dynamic learning strategies.

The ramifications of such findings are, however, far-reaching as far as performance metrics are concerned. The fundamental thing about these systems is that they undermine the traditional design of the scientific method, where linear workflows have been gentrified to variations of iterative, adaptive, autonomous discovery processes. The implication of this fundamental transformation can and should be explored on a deeper philosophical level, questioning how knowledge is produced, who the author is, and who owns the intellectual contribution in a world where machines co-create the new scientific truth. Also, the potential collaboration of human specialists and AI machines adds up to a new way of creating synergy when the machine broadens the cognitive range, and the human



retains the context-related irrationality, moral median, and industrial specificity.

Nevertheless, there are various restrictions and issues left unsolved despite the promise. The first of them is that of interpretability. With increasing scale in size and autonomy, these systems increase in systemic opaqueness, and the logic and decision pathways are becoming opaque and surround their creation and employment. Due to this concern, there are trust, bias, and accountability issues. When the topic is life or death, i.e., in a high-stakes area like healthcare, climate policy, or synthetic chemistry, the lack of full explainability regarding why a digital scientist makes a particular decision may inhibit the adoption or trigger resistance among regulators and scientists. This requires the integration of a strong explainable AI (XAI) architecture, where not only can visualization of decision flows be put in place, but it is also in line with domain-specific narrations, which are trustworthy and verifiable by experts.

It is also essential to the ethical bounds and control of safety. Although today's systems have simulation constraints and EITL architectures, they will need ethical reasoning encoded in their inference engines or model construction tomorrow. It becomes essential because AI assistants propose experimental routes or create new chemical substances with untested properties. Ethical governors, boasted by interdisciplinary review boards and algorithmic safety nets, play a significant role in ensuring that AI-powered studies are within the boundaries of society and science deemed acceptable.

The future, in turn, is associated with the intensification of integrating these systems into the institutional and educational frameworks. The labs will also have to reorganize work processes so that AI agents are not mere assistants but also partners who can conduct their discovering campaigns. Peer reviews, scientific journals, and other forms of publishing should change to consider the output of AI-driven inquisition by developing metrics of reliability and reproducibility of machine-identified knowledge. Moreover, curriculums in the higher education sector will have to transform to train scientists who are not only experts in their field but also scientists who will be literate in AI systems, algorithmic reasoning, and ethical design. Such scientists should be trained to explain, process, and

collaboratively develop knowledge with intelligent machines.

In future work, several research areas of critical concern can be identified. One of these concerns is creating domain-oriented knowledge ontologies to enable AI systems to identify places better or even translate and apply ideas across disciplines. Another is venturing into federated scientific learning models where distributed AI agents would discover locally and could contribute to the global models at regular intervals, thereby maintaining the privacy of data and speeding up global innovation. Furthermore, improving the adaptive simulation engines, which would allow re-calibration in real time based on the stream of changing data, would open a new era of never-ending discovery cycles, where science will be a self-improving organism.

To conclude, AI-enhanced machine learning systems ceased to be an imaginary concept; they are workable and revolutionary creations that are already changing the scenery of scientific revelation. Their functionality has been more than just automation, as shown in this paper, which seeks to reevaluate the hypothesis generation, testing, and verification process. The road we will walk will require caution, creativity, and ethical vision, and it will promise to democratize knowledge, compress discovery cycles, and release insights previously in the domain of human cognition by themselves. The next frontier we are entering requires that we design, govern, and work around these systems not to make science faster but wiser, more inclusive, and significantly more able to meet the most challenging problems the world can offer.

## Reference

1. Baker, D. (2021). The RosettaFold protein structure prediction system: advancing structural biology. *Science*, 374(6573), 130–135. <https://doi.org/10.1126/science.abj8754>
2. Granic, I., Morita, H., & Scholten, H. (2020). Beyond Screen Time: Identity Development in the Digital Age. *Psychological Inquiry*, 31(3), 195–223. <https://doi.org/10.1080/1047840X.2020.1820214>
3. Gurgitano, M., Angileri, S. A., Rodà, G. M., Liguori, A., Pandolfi, M., Ierardi, A. M., ... Carrafiello, G. (2021, July 1). Interventional Radiology ex-machina: impact of Artificial Intelligence on practice. *Radiologia Medica*.

- Springer-Verlag                      Italia                      s.r.l.  
<https://doi.org/10.1007/s11547-021-01351-x>
4. Hassabis, D., & Jumper, J. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583–589. <https://doi.org/10.1038/s41586-021-03819-2>
  5. Ikumoro, A. O., & Jawad, M. S. (2019). Intention to Use Intelligent Conversational Agents in e-Commerce among Malaysian SMEs: An Integrated Conceptual Framework Based on Tri-theories including Unified Theory of Acceptance, Use of Technology (UTAUT), and T-O-E. *International Journal of Academic Research in Business and Social Sciences*, 9(11). <https://doi.org/10.6007/ijarbs/v9-i11/6544>
  6. KwapongBaffoe, B. O., & Luo, W. (2020). Humanitarian Relief Sustainability: A Framework of Humanitarian Logistics Digital Business Ecosystem. In *Transportation Research Procedia* (Vol. 48, pp. 363–387). Elsevier B.V. <https://doi.org/10.1016/j.trpro.2020.08.032>
  7. King, R. D., Schuler Costa, V., Mellingwood, C., & Soldatova, L. N. (2018). Automating sciences: Philosophical and social dimensions. *IEEE Technology and Society Magazine*, 37(1), 40–46. <https://doi.org/10.1109/MTS.2018.2795097>
  8. Karniadakis, G., Kevrekidis, I., Lu, L., et al. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 422–440. <https://doi.org/10.1038/s42254-021-00314-5>
  9. Maziar, R., Perdikaris, P., & Karniadakis, G. (2018). Physics-informed neural networks: A deep learning framework for solving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
  10. Montavon, G., Samek, W., & Müller, K.-R. (2018). Methods for interpreting and understanding deep neural networks. *Digital Signal Processing*, 73, 1–15. <https://doi.org/10.1016/j.dsp.2018.02.001>
  11. Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: Theory and applications to forward and inverse PDE problems. *Journal of Computational Physics*, 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
  12. Sparkes, A., Aubrey, W., Byrne, E., Clare, A., Khan, M. N., Liakata, M., ... King, R. D. (2010). Towards Robot Scientists for autonomous scientific discovery. *Automated Experimentation*. <https://doi.org/10.1186/1759-4499-2-1>
  13. Stein, H. S., & Gregoire, J. M. (2019). Progress and prospects for accelerating materials science with automated and autonomous workflows. *Chemical Science*, 10(42), 9640–9649. <https://doi.org/10.1039/c9sc03766g>
  14. Scalliet, G., Mosbach, A., & Cavanagh, H. (2021). Physics-informed deep learning characterizes morphodynamics of Asian soybean rust disease. *Nature Communications*, 12, 6424. <https://doi.org/10.1038/s41467-021-26577-1>
  15. Tirmizi, S. M. A., Malik, Q. A., & Hussain, S. S. (2020). Invention and open innovation processes, and linkages: A conceptual framework. *Journal of Open Innovation: Technology, Market, and Complexity*, 6(4), 1–19. <https://doi.org/10.3390/joitmc6040159>
  16. Wies, E., & Polozov, C. (2009). Towards robot scientists for autonomous scientific discovery. *Automated Experimentation*, 2, 1–12. <https://doi.org/10.1186/1759-4499-2-1>
  17. Xing, Z., Chen, Y., & Protopapas, P. (2018). The rise of deep learning in drug discovery. *Drug Discovery Today*, 23(6), 1241–1250. <https://doi.org/10.1016/j.drudis.2018.01.039>
  18. Zhang, L., Zhong, Z., & Wang, Y. (2021). Knowledge integration into deep learning in dynamical systems: An overview and taxonomy. *Journal of Mechanical Science and Technology*, 35(4), 1331–1342. <https://doi.org/10.1007/s12206-021-0342-5>
  19. Zhou, Q., Liu, Z., & Yan, J. (2021). Machine learning for metal additive manufacturing: predicting temperature and melt pool fluid dynamics using physics-informed neural networks. *Computational Mechanics*, 67(2), 619–635. <https://doi.org/10.1007/s00466-020-01952-9>
  20. Zhu, Q., & Liu, Z. (2021). Machine learning for metal additive manufacturing: predicting temperature and melt pool fluid dynamics using physics-informed neural networks. *Computational Mechanics*, 67, 619–635. <https://doi.org/10.1007/s00466-020-01952-9>