



## AI That Explains Itself: Helping Experts from Different Fields Make Better Decisions Together

*Dr. Akhil Pandey<sup>1</sup>, Dr. Ashok Kumar Kajla<sup>2</sup>, Dr. Vishal Shrivastava<sup>3</sup>, Chirag Suthar<sup>4</sup>*

Department of Information Technology Student of Information Technology, Arya College of Engineering and IT, Kukas, Jaipur

### ABSTRACT –

As artificial intelligence becomes increasingly embedded in scientific research, the complexity of modern AI models poses a serious challenge to interdisciplinary collaboration. Experts in domains such as physics, biology, and chemistry often struggle to interpret opaque, black-box predictions made by AI systems, limiting their trust and adoption. This research focuses on developing **scalable and interpretable AI techniques** that can be adapted across diverse scientific fields. By merging **domain-specific insights with universal interpretability frameworks**, the study aims to build AI models that are not only powerful but also transparent and understandable to users with varying levels of technical expertise. These models are designed to provide clear, intuitive explanations of their outputs, allowing researchers to validate results, make confident decisions, and foster meaningful cross-disciplinary collaboration. Ultimately, the goal is to bridge the gap between AI and human understanding—making advanced technology a trusted partner in scientific discovery.

**Index Terms** – Scalable Interpretability, Multi-Disciplinary Research, Explainable Machine Learning, Hypothesis Generation, Cross-Domain Collaboration

### 1. INTRODUCTION

The rapid advancements in artificial intelligence have resulted in substantial breakthroughs across various scientific domains, including physics, biology, and chemistry. However, despite the growing complexity of artificial intelligence models, their dynamic processes are often obscure, creating challenges for researchers who need transparent and comprehensible output. This lack of interpretability hampers the trust and effective employment of artificial intelligence systems, particularly in multi-disciplinary investigation, where field experts are required to collaborate across numerous domains. The demand for scalable interpretability techniques capable of producing clear insights into model behavior in multiple domains is critical to ensuring a broader uptake of artificial intelligence in scientific discovery. By establishing scientific AI-based methods that can be tailored to the unique requirements of different disciplines, we can enhance the transparency of AI models and bridge the gap between sophisticated computational approaches and traditional scientific inquiry. This paper aims to explore innovative approaches to making AI-based models interpretable and accessible to practitioners from diverse backgrounds, thus facilitating more informed autonomous decision-making, hypothesis validation, and cross-domain collaboration.

#### 1.1 Challenges in man-made intelligence Interpretability Across Logical Spaces

AI is reforming logical revelation by empowering information driven examination in such fields as material science, science, science, and medical care. Customary logical strategies depend on speculation detailing and experimental approval, yet AI speeds up this interaction by parsing monstrous datasets for examples and connections that probably would not be quickly clear to human analysts. AI models, especially profound learning and support learning, are presently ready to create unique theories, plan analyses, and even make forecasts about logical variations from the norm. Nonetheless, AI-created results frequently need straightforwardness, making it hard for researchers to trust and approve their credibility. Along these lines, there is a developing interest for AI frameworks that computerize disclosure, yet additionally give experiences in an intelligent and human-comprehensible way.

#### 1.2 The Requirement for Versatile Interpretability Strategies

Most artificial intelligence algorithms, in particular, deep neural networks, function as “black boxes” because their dynamic processes are not evident to people. This absence of transparency is a serious problem in scientific research, which requires evidence and repeatability. If researchers are not given accessible justifications, the conclusions obtained from AI methods can be challenging to confirm, resulting in skepticism among the scientific community. Furthermore, decision-making mechanisms may present faults, errors, or misunderstandings that can lead to a substantial distortion of the data. For example, in medical practice, an AI model that forecasts the development of a disease should be accompanied by a valid rationale for its conclusions to receive accurate clinical recommendations. Resolving this difficulty will necessitate the development of methodologies to enhance AI interpretability while preserving its high performance post challenge.

### 1.3 Overcoming any barrier Among man-made intelligence and Space Specialists

To overcome any issues between artificial intelligence mechanization and logical straightforwardness, a few interpretability strategies have been created. These include:

- Include Significance Examination: Recognizing which information highlights contribute the most to computer-based intelligence expectations.
- Consideration Instruments: Featuring the pieces of info information that impact artificial intelligence choices.
- Rule-Based Models: Making simulated intelligence frameworks that work utilizing comprehensible legitimate standards.
- Model-Rationalist Methodologies: Utilizing techniques like SHAP (SHapley Added substance Clarifications) or LIME (Nearby Interpretable Model-freethinker Clarifications) to make sense of complicated models.

By coordinating these strategies, analysts can guarantee that computer-based intelligence created revelations are joined by legitimizations, making them more dependable and experimentally helpful.

### 1.4 Influence on Logical Disclosure and Coordinated effort

Interpretable simulated intelligence can possibly change different logical spaces, including:

- Medical services: artificial intelligence can aid infection finding by giving clarifications to clinical picture orders.
- Genomics: AI models can reveal quality infection connections while featuring critical biomarkers.
- Materials Science: artificial intelligence driven expectations of material properties can be made interpretable for trial approval.
- Material science and Cosmology: computer based intelligence can assist with investigating astronomical information, foresee heavenly occasions, and make sense of the thinking behind astrophysical expectations.

These applications show how computer based intelligence, when made interpretable, can act as a solid logical accomplice as opposed to only a computational device.

---

## 2. TECHNIQUES

Like LIME and SHAP, these procedures make it workable for neighbourhood surmisings of mind-boggling AI models to be acquired, offering clear clarifications of predictions without the requirement for getting to the model's inside activities. They are generally relevant and can be applied across different scientific spaces.

### 2.1 Model-Skeptic Interpretability Strategies

Model-free thinker techniques offer interpretabilities to any AI model, regardless of its architecture. Such techniques focus on understanding the predictions of a model, without needing access to its internal workings. For example, popular model-free thinker procedures, including LIME and SHAP, produce interpretable explanations by deliberately approximating intricate models with simpler, interpretable ones on a local scale. In this way, these techniques eliminate the need for domain-specific knowledge about the underlying design and are highly appropriate for multi-disciplinary research. By producing local approximations of model behavior, it is possible to explain how individual predictions are reached, thus enabling specialists from diverse backgrounds to understand a given model's thinking process. Furthermore, model-skeptic techniques allow producing visualizations of important features that contribute to a prediction output, thus facilitating a deeper understanding of the model's decision-making process.

### 2.2 Highlight Significance Examination

Point out the importance strategies plan to assess the significance of each data include in the model's dynamic cycle. Such procedures are vital in multi-disciplinary research, where information on which factors are the most influential might be crucial to scientific discovery. Techniques, such as Random Forests, Angle Boosting Machines, and Feature Component Importance evaluate with features drive the predictions the most and provide transparency in the model's behavior. This method is especially useful in fields such as physics or chemistry, where specific molecular characteristics or experimental conditions play a crucial role in the outcomes. Feature importance allows scientists to determine critical factors affecting the model's predictions, streamlining further experimental work. Furthermore, this process enables space researchers to focus on the most critical factors, reducing noise and improving the efficiency of scientific inquiry. Illustrations of feature importance help translate complex models into understandable insights for various research communities.

### 2.3 Causal Derivation and Thinking

Causal derivation techniques focus on such a relationship between the inputs and outputs in a model's predictions, providing more than purely correlation-based insights. In multi-faceted analysis, causal models are particularly helpful in figuring out fundamental quirks, as they provide researchers with an explanation of why some outputs occur. Techniques such as Do-calculus, Granger causality, and Counterfactual Deduction allow the analysts to identify potential causal factors behind observed data patterns, leading to more knowledge-driven guidance. Moreover, causal models can help eliminate confounding relationships, resulting in clearer paths to scientific discovery. For example, in a biological research context, understanding causal relationships between genes and diseases can lead to better-targeted treatments.

### **2.4 Representation Apparatuses for Model Interpretability**

Visualization plays a critical role in enhancing model interpretability, especially in cross-disciplinary research, where complex data needs to be presented in an easily comprehensible format. Tools such as t-SNE, PCA, and UMAP are used to project high-dimensional data into lower dimensions, making it easier for experts to interpret patterns and relationships. Additionally, heatmaps, partial dependence plots, and feature importance plots give vivid, intuitive representations of model behavior, helping researchers comprehend the crucial driving factors more quickly. Visualization tools are particularly valuable when dealing with large-scale scientific data because they provide insights into the internal workings of complex models, such as deep brain networks. By offering graphical representations of model predictions and internal processes, these tools bridge the gap between the complexity of AI and domain-specific expertise. Visualizations simplify not only the interpretation of model outcomes but also collaboration within and among diverse research teams.

---

## **3. MULTI-DISCIPLINARY RESEARCH**

As adaptable interpretability research for man-made consciousness, multi-disciplinary examination consolidates the information and aptitude scopes from real different logical fields such as physical science, science, science and ecological science. It is intended to make complex AI models straightforward and reasonable to scientists with various foundations, engaging cooperation. When fitted to various areas, interpretability strategies make AI models more relevant to various logical controls. This methodology advances between disciplinary critical thinking and works on the general logical process. By tending to difficulties of powerful all-encompassing interpretability approach, adaption reshapes the techniques to fulfill the novel needs of each space. Normalizing the interpretability conduct across different fields advances quality, at last interpretation of the AI model, simple models fosters the relationship between the AI researchers and the area researchers, in this way facilitating the between the disciplinary logical progress.

### **3.1 Connecting Area Skill with man-made intelligence Models**

In interdisciplinary examination, AI models should join information from various logical fields to give important bits of knowledge. Spanning the chasm among AI specialists and space-explicit scientists is crucial to guarantee that interpretability procedures are useful and meaningful. For example, physicists, scientist, and researchers all have different information designs and ways to deal with critical thinking. Due to this fact, AI interpretability must take care of the language and necessities of each discipline, while models remain understandable across all areas. Through versatile interpretability techniques, researchers can collectively explore new scientific theories, validate AI predictions, and trust the results in their respective fields. Articulating a common ground through interpretable AI can lead to more meaningful and conducive discoveries. Collaborative tools, such as affordable AI visualizations and shared interpretability standards, can promote smoother communication between cross-disciplinary teams.

### **3.2 Computerized Exploratory Plan and Execution**

As of now, artificial intelligence-based frameworks are equipped for booking and aiming media inquires very to no human intervention. In the fields of materials science and engineered science, for instance, artificial intelligence can propose trial game plans free without help from anyone else of feed retreat, controlling the assigned calling of the computerized examination spot and subtletling observational conditions subject to advancing feedback. Robotized chemical amalgamation locales apply fake understanding into assessing various sub-nuclear blends, in as much as possible participant conditions to uncover old materials with consideration catching properties. Such frameworks not barely lessen the hour and materials spent on experiments trial and sin yet by the same token oblige the nature of being proclivity. Also, strengthen learning evaluations permit man-made brainpower to expertet trial game plan united on past results to assemble a greater long-in-the-tooth and attractive line of examination. Regardless, challenges recurrence the same course such as clinical capillarity, viral disrespects, and security revolve require being rehearsed notwithstanding the faith in the precision of artificial intelligence-driven in a general sense based questioning..

### **3.3 Information Driven Logical Disclosure**

A large quantity of scientific information currently being produced — from genomic sequences to meteorology patterns — has rendered it progressively tough for researchers to analyze and manually extract meaningful insights. AI-enabled data mining and artificial intelligence algorithms are proficient in handling extensive datasets automatically, identifying relationships, and issuing forecasts to promote scientific discovery. For illustration, in genomics, artificial intelligence might be leveraged to analyze clusters of DNA to identify genetic mutations concomitant with diseases, subsequently expediting individualized medication research. Environment science artificial intelligence models examine previously observed weather conditions to forecast future climate tendencies and natural disaster patterns. A key challenge in data-driven discovery, however, is the differentiation between relationships and causality. While AI might efficiently uncover patterns, additional interpretability methods, such as causal-inference techniques, are necessary to confirm scientific links and provoke meaningful insights.

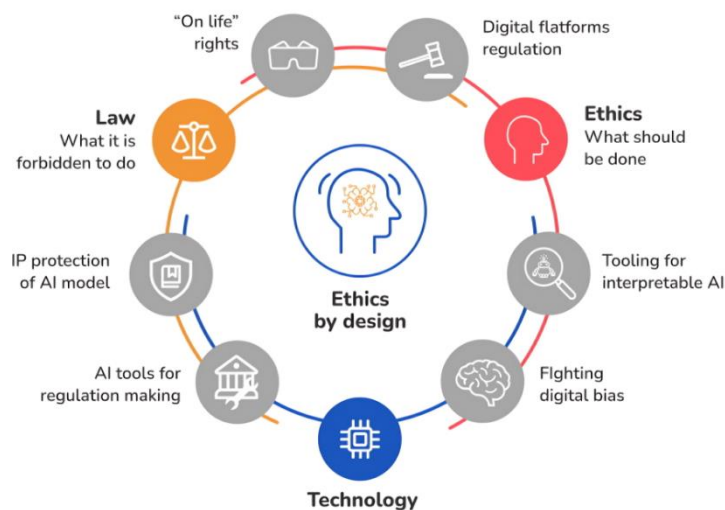
### **3.4 Utilizing man-made intelligence Models for Cross-Disciplinary Revelations**

Using man-made intelligence Models for Cross-Disciplinary Revelations computer-based intelligence's capacity to cycle enormous volumes of information from different orders can prompt pioneering cross-disciplinary revelations. Nonetheless, for man-made intelligence to be genuinely useful to established researchers, its discourses must be interpretable to researchers from different fields. For instance, false models route to geographic data might help usable learn and biologistshow soil aggregation influences surroundings mathematics. Additionally, false models in physics can be employed by substantial scientists to anticipate new belongings physical. Adaptive interpretability can make these cross-disciplinary usages achievable giving

accessible clarifications of the model's conclusions. It allows specialists from various institutions to understand in what way intelligence's purposes correlate with their understanding, subsequently, their study strategy. By using modest versions, false can turn into a discovery tool in fields running in separate stores, inspiring additional networks and thorough outcomes.

### 3.5 Moral and Philosophical Contemplations in Independent Disclosure

Although independent logical disclosure offers massive advantages, it also raises substantial moral and philosophical concerns. One main present issue is the accountability of man-made intelligence generated disclosures – if an AI system proposes a logical hypothesis or experimental plan that leads to a potentially unfavorable outcome, who is accountable? Additionally, the reliance on artificial intelligence generated theories could lead to insolence in machine-based reasoning, ultimately ignoring human intuition and critical thinking. Ethical apprehensions also arise in circumstances where AI is used in dubious scientific fields, such as genetic modification or artificial intelligence generated scientific patents. To tackle these difficult situations, fitting decentralized applications ought to be devised to ensure responsible AI use in research. Furthermore, the combination of scientific AI techniques is critical to ensure that AI-generated discoveries are compatible with ethical standards and human scientific reasoning. Independent logical discovery is transforming the way research is performed, responding to the power of artificial intelligence to generate hypothesis, hypothesis and analysis and experiment on a large scale. While AI systems have the potential to accelerate exploration in various areas, the realization of interpretability, reproductiveness and ethics continuity is decisive. Future advancements in the availability of AI interpretation, causal deduction and interdisciplinary cooperation will be vital to unleashing the potential of independent logical discovery. By developing AI systems that automate AI reasoning, popular scientists can use the power of AI while continuing to trust, be the reason and be responsible for science.



## 4. EXPLAINABLE MACHINE LEARNING

A critical focus of Reasonable AI (XML) is to create models that not only make accurate predictions but also provide human-comprehensible justifications for their decisions. As AI (ML) models continue to advance in complexity, notably deep learning and harnessing models, their decision-making processes often become opaque, leading to the “black-box issue.” This lack of transparency presents inhumanities to researchers, experts, and regulators who need to leverage AI driven insights, especially in critical domains such as healthcare, finance and autonomous systems. Logic ensures that ML models are explainable, accountable and fair, fostering trust and ethical use of artificial intelligence. To these ends, Logical AI can be broken down to the following critical themes:

### 4.1 The Requirement for Reasonableness in AI

AI models are increasingly being deployed in real-world applications, but their lack of transparency poses concerns regarding trust, fairness, and accountability. In fields such as medical diagnosis, financial risk assessment, and self-driving cars, supervision must be interpretable to ensure safety and reliability. For example, if an AI predicts that a patient has a disease, doctors need to know which features—such as genetic indicators or lifestyle variables—contributed to that prediction. Likewise, in banking, transparent credit scoring models inform regulators and clients why a loan was approved or rejected. Without transparency, biased judgments or inaccurate predictions by AI systems could continue unchecked, leading to unjust or disastrous conclusions. Thus, designing methods to interpret ML models is critical to facilitate ethical AI deployment and establish user confidence.

### 4.2 Sorts of Reasonableness in AI

Reasonableness in ML can be ordered into two fundamental sorts:

- **Worldwide Reasonableness:** Gives a general comprehension of how a model functions across all data of interest. This assists analysts and information researchers with approving the model's overall way of behaving, guaranteeing it lines up with area information.
- **Nearby Reasonableness:** Makes sense of individual forecasts by distinguishing which elements impacted a particular result. This is especially

valuable in high-stakes applications where supports for single choices are required.

For instance, in clinical imaging, worldwide logic might assist with verifying that a profound learning model depends vigorously on surface highlights while grouping growths. In the meantime, neighborhood reasonableness can feature explicit districts in a X-ray examine that added to a malignant growth conclusion. The two kinds of logic assume a vital part in making man-made intelligence models more interpretable and responsible.

#### 4.3 Model-Explicit versus Model-Rationalist Clarification Strategies

Model-Explicit vs Model-Rationalist Exposition Techniques. Model-explicit strategies are proposed on explicit ML models and can in general affect their inside construction. For example, decision trees and direct regression models are naturally interpretable since they explicitly demonstrate how inputs affect yields. Brain organizations, however, require express procedures, for example, activity perceptualization and attention mechanisms to translate their inner activities. Model-Rationalist Techniques. The strategies can be applied to any ML model, which provides flexibility across various calculations. Procedures like LIME (Local Interpretable Model-Agnostic Explanations) and SHAP (SHapley Additive Explanations) create interpretable approximations of complex models. These techniques are often used in applications where deep learning models require post-hoc explanations, e.g., healthcare diagnostics and fraud detection. Both techniques have their pros, and researchers usually combine them to enable a comprehensive range of interpretability.

#### 4.4 Normal Strategies for Logical AI

A few strategies have been created to make ML models more interpretable. The absolute most broadly utilized include:

- **Include Significance Investigation:** Figures out which info highlights contribute most to a model's expectations. This is helpful in fields like genomics, where understanding the job of explicit qualities in illnesses is essential.
- **LIME (Neighborhood Interpretable Model-Skeptic Clarifications):** Creates basic, interpretable models (like straight relapses) around individual forecasts to make sense of why a ML model settled on a specific choice.
- **SHAP (SHapley Added substance Clarifications):** Uses game-hypothetical standards to ascribe include significance across various forecasts decently. It is regularly utilized in money, medical services, and logical exploration.
- **Saliency Guides and Consideration Components:** Utilized in profound learning models, these procedures feature what parts of an information (like a picture or text) were most persuasive in the dynamic cycle.

Every one of these strategies upgrades computer-based intelligence straightforwardness, assisting clients and partners with understanding how computer-based intelligence driven frameworks work.

Logical AI is a rising field that guarantees that AI models are logical, powerful, and interpretable. As shown in this paper, by depending on different methodologies such as LIME, SHAP, and feature importance analysis, AI-trained professionals can see better the judgments of AI-enabled models, hence advertising liability. While challenges such as computational expenses and implications for interpreted understanding persist, future strides in neuro-symbolic AI, human-probed explanations, and regulatory accountability may indeed work toward more dependable AI systems. Given the increasing growth in AI acceptance over different sectors, providing reliable and interpretable ML models is vital to foster the agency and ensure accountable AI application.

## 5. HYPOTHESIS GENERATION, AND CROSS-DOMAIN COLLABORATION

In closing, the guessing era fueled by artificial

intelligence allows researchers to recognize

patterns and propose novel scientific breakthroughs in various fields. Interpretability techniques ensure that artificial intelligence developed notions are transparent, enabling experts to validate and refine them. Enhanced cross-domain collaboration is promoted through viable AI models that bridge knowledge gaps between domains such as biology, physics, and atmosphere science. AI-driven knowledge transfer facilitates researchers to apply methods from one domain to another, promoting interdisciplinary exploration. Collaborative AI frameworks help to standardize transparency methodologies, making AI insights accessible among different popular researchers. Despite this, adaptation transparency solutions are needed to overcome challenges such as data architecture discrepancies and domain-specific assumptions. Future AI-driven collaboration developments will elevate cross-disciplinary discoveries through intelligent and analytical models.

### 5.1 Theory Age in computer-based intelligence

The speculation age alludes to the way toward creating new speculations, ideas, or connections from big data with the assistance of AI<sup>74</sup>. Generally, theories have been created based on human intuition and past information; nonetheless, AI improves this cycle by recognizing perplexing connections and relationships from substantial amounts of information. For example, in drug revelation, AI-driven speculation age can foster new medication combinations by dissecting atomic structures and foreseeing their adequacy. In astronomy, AI models can recognize exoplanets by recognizing outlines in telescope information. The AI models that use human amalgamation include measurable strategies like Bayesian surmising, profound learning design acknowledgment, and support learning. A test, in any case, is to guarantee that the AI-created speculations are deductively legitimate as opposed to arbitrary relationships. Techniques like causal derivation and space master approval are needed to refine AI-driven associations. By consolidating large data frameworks and human intelligence, the AI inductive speculation age quickened logical advancement while staying solid.

### 5.2 Job of Interpretability in Theory Approval

Interpretability techniques help specialists approve when computer-based intelligence produces a hypothesis; and refine it by giving experiences into model-independent direction. As model-freethinker techniques, SHAP and LIME can feature the specific elements assuming a part in the expectation and help researchers decide whether the theory is true and information-driven; approval in multi-disciplinary areas requires broad collaboration between computer-based intelligence specialists and area specialists to guarantee that these discoveries have been created logically. For example, in materials science, in another model, a simulation-based intelligence prediction model can recommend a new material structure with unique characteristics; however, if they are not clearly clarified, the specialists may not be able to verify the data. By utilizing interpretable computer-based intelligence, the specialists can evaluate the dependability of the prediction and pinpoint areas that are possibly predisposed to erroneousness. This straightforwardness also instills trust in computer-based intelligence predictions and allows researchers to make them more solid.

### 5.3 Straightforwardness in simulated intelligence Navigation

In the context of artificial intelligence, straightforwardness is understood as the ability to comprehend how and why a particular AI system behaves in a given manner. Lack of AI-life involves making AI a “black box,” which renders it impossible for the users to decode its logic and foresee possible biases or errors. It is a vital component of ethical AI application, regulatory alignment, and user confidence. There are different levels of AI transparency, from algorithmic transparency, meaning understanding the inner workings of the model, to decision transparency, which implies figuring out the specific outcomes for individual predictions. For example, in the law enforcement sector, AI is used to forecast whether the defendant is likely to commit another crime. If the AI model does not provide an explanation of its predictions, the legal professionals cannot assess whether the forecasts are unbiased or prejudiced. Techniques such as LIME and SHAP are often used to increase AI transparency by breaking down complex predictions into human-understandable insights. Among brain networks, attention mechanisms also help highlight which features have influenced AI’s decisions, thereby making deep learning models more interpretable. However, full transparency is challenging to achieve, especially for deep learning models with

millions of parameters. Researchers are working on developing inherently interpretable artificial intelligence models, such as decision trees, rule-based AI, and neuro-symbolic AI, in order to balance accuracy with interpretability. The future of AI transparency lies in the development of systems that are not only accurate but also comprehensible, thereby ensuring that AI decisions conform to ethical and societal expectations.

### 5.4 End

There are three fundamental columns for powerful artificial intelligence development: speculation age, trust, and straightforwardness. AI-driven theory age speeds up logical disclosure by recognizing covered up examples in information, yet it requires approval to guarantee its dependability. Trust in AI incorporates, and availability, the shopper’s thought, etc., which guarantees that AI frameworks are acknowledged in down to earth application. Straightforwardness upgrades AI responsibility, essentially making the dynamic cycle prompt, decreasing inward inclination, and empowering moral AI sending. By incorporating these standards, AI can turn into a more remarkable, capable, and broadly trusted instrument for logical and modern leaps forward.

## 6. REFERENCES

1. Kramer, S., Cerrato, M., Džeroski, S., and Lord, R. (2023)  
"Robotized Logical Revelation: From Condition Disclosure to Independent Disclosure Frameworks."  
Accessible at: <https://arxiv.org/abs/2305.02251>
2. Quinn, T. P., Gupta, S., Venkatesh, S., & Le, V. (2021) AI That Explains Itself: Helping Experts from Different Fields Make Better Decisions Together  
"A Field Guide to Scientific XAI: Transparent and Interpretable Deep Learning for Bioinformatics Research."  
Available at: <https://arxiv.org/abs/2110.08253>
3. Ribeiro, M. T., Singh, S., and Guestrin, C. (2016)  
"For what reason would it be advisable for me I trust you?" Making sense of the expectations of any classifier. Procedures of the 22nd ACM SIGKDD Worldwide Meeting on Information Revelation and Information Mining (KDD).  
Available at:  
<https://doi.org/10.1145/2939672.2939778>
4. Zhang, A., Lipton, Z. C., Li, M., & Smola, A. J. (2021)  
"Dive into deep learning. Cambridge University Press."  
Available at: <https://d2l.ai/>
5. The Guardian. Published on February 3, 2025  
"AI to revolutionise fundamental physics and 'could show how universe will end.'"

Available at:

<https://www.theguardian.com/science/2025/feb/03/ai-to-revolutionise-fundamental-physics-and-could-show-how-universe-will-end>

6. Ghorbani, A., Abid, A., & Zou, J. (2019)

“Interpretation of neural networks is fragile. Proceedings of the AAAI Conference on Artificial Intelligence, 33(1), 3681-3688.”

<https://doi.org/10.1609/aaai.v33i01.33013681>