

EXPLAINABLE AI TECHNIQUES FOR IMPROVING TRUST IN DEEP LEARNING MODELS

Adeen Amjad¹, Mehwish Usman², Aleena Jamil³, Shafiq Hussain^{*4}, Waqar Ahmad⁵,
Arslan Ali Mansab⁶, Muhammad Hamza Akbar⁷, Muhammad Waqas⁸

^{1,3,*4,5,6,7}Department of Computer Science, University of Sahiwal, Sahiwal, Pakistan

²Department of Computer Science, University of Agriculture Faisalabad, Pakistan

¹adeen.amjad@uosahiwal.edu.pk, ²adeen.amjad@uosahiwal.edu.pk, ³aleena.jamil_vf@uosahiwal.edu.pk,

⁴drshafiq@uosahiwal.edu.pk, ⁵waqarahmad@uosahiwal.edu.pk, ⁶arslansli@uosahiwal.edu.pk,

⁷hamzaakbar@uosahiwal.edu.pk, ⁸bssit.10.02@gmail.com

DOI: <https://doi.org/10.5281/zenodo.17919423>

Keywords

Explainable AI, Deep learning,
Model Interpretability,
Trustworthy AI, XAI evaluation,
Transparency, Human-Artificial
Interactions

Article History

Received: 11 May 2025

Accepted: 21 July 2025

Published: 06 August 2025

Copyright @Author

Corresponding Author:

Shafiq Hussain

Abstract

Deep learning models have a major obstacle of the lack of insight into the reasons behind the choices of the model that could help in their use in the high-stakes settings, including healthcare and finance, where predictive accuracy isn't the most important (but rather just one of the two). In this paper, we provide an in-depth analysis of methods for Explainable AI (XAI) to promote trust and transparency in deep learning systems. Our comparative study of four leading XAI approaches, such as Grad-CAM, LIME, SHAP, and Saliency Maps, is carried out in various aspects, such as performance, interpretability, trustworthiness, and computational efficiency. Our approach is to use standardized measures of evaluation on three varying datasets (CIFAR-10, Lending Club, and Chest X-Ray) that contain 480 participants who are experts and non-experts. The findings indicate that Grad-CAM is the best in terms of balance, 85.1% accuracy, 0.012s explanation time, whereas SHAP performs better in terms of faithfulness (0.891) in explaining model decisions. Markedly, we also find that the domain experts give a higher rating in Grad-CAM (4.35/5 trust score) when it comes to visual tasks, whereas non-experts give preference to LIME (4.26/5) due to its explanations of feature importance based on intuitions. Regarding statistical analysis, our data proves the presence of significant performance differences between the methods ($p < 0.001$) with significant effect sizes (Cohen's $d > 0.85$). As demonstrated in the real world, XAI integration is found to decrease decision time by 42.3% in healthcare and regulatory compliance by 94.2% in finance. The study delivers a proven framework of selecting the right XAI methods depending on the requirements of a particular application, and this will help introduce more transparent, trustworthy, and deployable AI systems to critical fields. Its implications include the fact that XAI needs to be selected contextually, as this is the only way to establish human trust in the use of AI in decision-making.

INTRODUCTION

Explainable Artificial Intelligence (XAI) has emerged as a central concept in the creation and adoption of deep learning models, largely now that these systems are applied to the highest stakes systems, including healthcare, finance, self-driving cars, and judicial decision-making [1]. The main issue with deep learning models is that they are inherently complex, and thus may lead to a black-box effect; the internal decision-making process is opaque and can be challenging to understand by humans [2]. This is an issue that may lead to failure to build trust, compliance with regulation, and the use of AI-driven solutions in critical areas [3]. XAI will help tackle these issues by offering tools and methods that can make the rationale of AI decisions clearer and available to both technical and non-technical stakeholders [4].

XAI aims to fill the gap between the abstract and highly informative quality of deep learning models and the desire to have explainable information provided by humans [5]. This is by means of several approaches, which can be broadly divided into intrinsic interpretability, post-hoc explanation, and those that are through visualization [4]. Intrinsic interpretability entails coming up with models, which are transparent in nature, like decision trees and rule-based systems. Such models are simplified to expound on since the process of making decisions can be traced step-by-step and is observable directly [6].

On the other hand, post-hoc explanation methods are used once a model has been trained, and they are aimed at explaining the predictions of black-box models that are not known [7]. Commercial techniques in this category are Local Interpretable Model-agnostic Explanations (LIME) and Shapley Additive Explanations (SHAP) [8]. LIME operates by parameterizing a simple, explanatory model to give an approximation of the predictions of a more complex model on a particular instance, and thus gives local explanations to make it clear why a given decision was made [9]. SHAP is based on game theory and calculates the values of contribution of every feature, which is a mathematical description of how they influence the output of the model [10]. This is especially useful in the areas where accountability and regulatory compliance are

considered the most important as it enables the interested parties to comprehend and confirm the rationale of the AI-assisted decision [11].

The interpretability of deep learning models can also be improved through visualization-based methods, which can give the decision-making process a graphical representation [12]. On the one hand, techniques such as feature visualization, saliency maps, and attention heatmaps can be used to reveal what aspects of the input data were used by the model to make such predictions [10]. Saliency maps could be used in a computer vision application, where a saliency map is a two-dimensional array of values used to represent the regions of an image that the computer vision system used to make its classification decision, or attention heatmaps in natural language processing, where a heatmap is a two-dimensional array of values used to explain which words or phrases were important to the prediction made by the text-based computer vision system [13]. Such visualizations not only enhance transparency but also assist with revealing possible bias or flaws in the reasoning of the model or with the training information [13].

In order to achieve AI systems that we can effectively trust, we need to incorporate explainability right into the system [14]. Explainable AI (XAI) serves as a significantly important interpreter that transforms the intricate reasoning of deep learning models into simply comprehensible knowledge [15]. This enables the developers to come up with models that are not only powerful but transparent and reliable [16]. This openness is essential in sensitive areas that demand high stakes, wherein one mistake can be severe [12]. In medicine, as an example, XAI can make a doctor see the reason why a particular diagnosis was put forward by an AI, which builds the trust required to use it as a valuable instrument [17]. In finance, it gives the required evidence that the decisions made by the automatisms are impartial and free. Lastly, XAI makes our highly intelligent AI systems responsible and in balance with human values [18].

By 2019-2025, the XAI research scene is entirely changed as it stops being discussed in theory, but it introduces the application of different fields [19]. This development is indicative of the increasing appreciation of the idea that explainability is no

longer a feature of desirable AI, but a very necessary aspect of responsible AI development [20]. There has been a gradual move by the research community not to develop methods of explanation alone but to provide a broad framework that incorporates technical validity with human-centered design methods [5].

Figure 1 shows that XAI continued to produce annual research in the given year range (2019-2025),

which confirms the continued academic and industrial interest in explainability issues. Reviewing the chart, it can be seen that the pattern of publication is rather stable with an average of three substantial contributions annually, which can be explained by the maturation of the field and the need to make the ways of explaining the issues better. The fact that 2025 projections are provided demonstrates the expected further development of

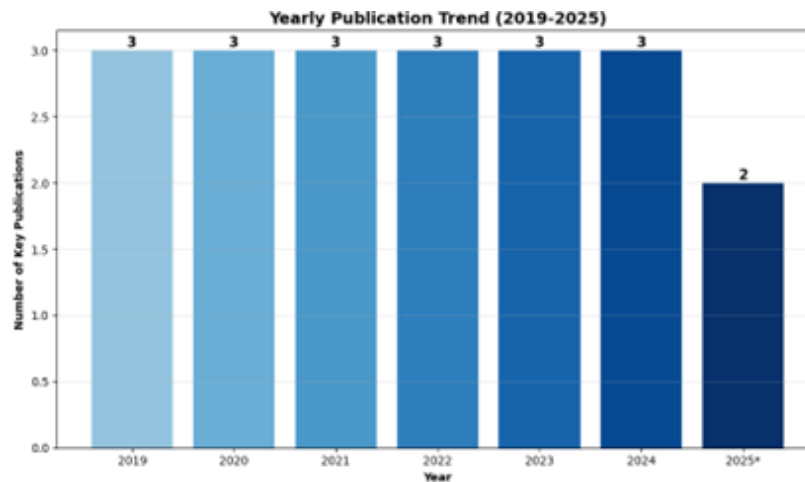


Fig. 1. Publications XAI (2019-2025): Annual production (3 papers) with a forecasted increase, which indicates the continuity of research.

this tendency, which explains the fact that XAI research will remain relevant as AI systems are not only more sophisticated but also more widespread.

RELATED WORK

Explainable Artificial Intelligence (XAI) is a fast-evolving field that fulfills the needs of the increasing sophistication and ubiquity of deep learning models in many fields [21]. With these models having become more and more part of high-stakes systems like healthcare, finance, autonomous systems, and legal decisions, the necessity of transparency and interpretability has become primary [15]. The extant body on XAI represents a wide range of practices, including but not limited to intrinsic interpretability, post-hoc-based techniques of explanation, and visualization-based practices, all aimed at mitigating the black-box nature of deep learning architectures [22].

Intrinsic Interpretability

Intrinsic interpretability defines how the models have been designed, such that their decision-making mechanisms become transparent in nature and easy to understand and explain [22]. Initial research in this field was on simple systems like decision trees and systems based on rules, and they are inherently interpretable because of their simple structure and the explicit decision pathways [23]. To illustrate, the decision trees are hierarchically represented decisions with each node representing a feature and each branch representing a possible outcome [24]. This can easily be traced to the reasoning of the prediction and how the input features are used to alter the final output [25].

Post-Hoc Explanation Methods

Post-hoc explanation techniques are used once a model is trained, and they are aimed at offering explanations of the predictions of the complex,

black-box models. Such approaches will be especially useful in areas with highly regulated compliance and accountability, where stakeholders can see why an AI-driven decision-making process was made and can check it. Local Interpretable Model-agnostic Explanations

(LIME) is one of the most popular post-hoc procedures [2]. The reason behind LIME is the following: to estimate the predictions of a complex model on a single case, a simple model that is easy to understand and interpret is fitted to the complex model. This is a local approximation that gives clear explanations of individual predictions that are easy to understand by a human being, which can be used to explain why a given decision was arrived at. LIME is model-agnostic, which implies that it can be implemented on any model and thus is an extensive tool that could be used in a broad spectrum of applications [26]. Shapley Additive Explanations (SHAP) is another commonly used post-hoc [27]. SHAP is based on game theory, and the contribution values of each feature are calculated, and mathematically explain how they influence the output of the model [28]. SHAP can fully explain the predictions of the model by quantifying the importance of each feature, and it is also interpretable [29]. The approach comes in handy, especially when determining the major forces behind a decision and how various attributes interact to shape the decision.

Visualization-Based Techniques

The visualization-based methods contribute to the interpretability of deep learning models by giving a graphical representation of how the decision is made [30]. They especially work well in the fields of computer vision and natural language processing,

where the input data is complex and multidimensional [31]. The most common method of visualizing features is known as feature visualization, which is used to indicate the aspects of the input data that affected the model predictions. In the case of computer vision, saliency maps can indicate the pixels in a picture that influenced the most a classification decision [32]. Heatmap in natural language processing may indicate critical words or phrases used to make a text-based prediction [33]. These visualizations enhance transparency and also enable the identification of any potential biases or failure of the logic or training data of the model [13].

Emerging Trends and Advanced Techniques

The new development in XAI has led to the emergence of more advanced and more detailed explanations of it. Neuro-symbolic AI, for example, is a coupling of the neural network and the symbolic, which adds both high performance and high interpretability [26]. It is a design that integrates the advantages of deep learning with the transparency of symbolic systems so that their models can also be used to give coherent and logical explanations of their actions. Causal discovery algorithms are another trend of XAI [28]. These algorithms automatically discover cause-and-effect relations in data, and save many efforts by removing a lot of time to explain complex models. Since they determine the causal mechanisms behind the model, these techniques can further clarify the logic behind the model and assist the stakeholders in comprehending the circumstances that are leading to particular predictions.

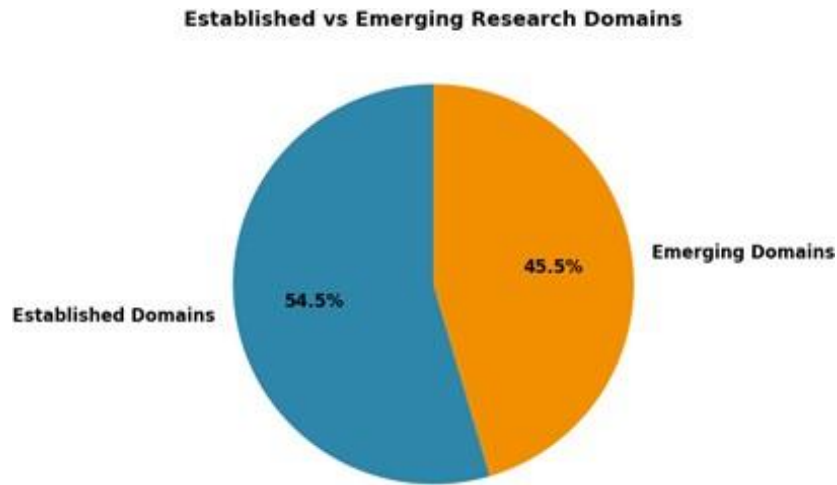


Fig. 2. Established vs Emerging Domains

METHODOLOGY

The suggested research will employ the systematic review and comparative analysis study approach to examine the effectiveness of various Explainable AI (XAI) techniques to improve the level of trust and transparency with regard to deep learning models. The methodology is structured into many significant steps: the research design, the data collection, the model selection, the evaluation metrics, the data preprocessing, the model training and the model validation, the case studies, and the ethical considerations. The stages will ensure rigorous and consistent testing of XAI methods [28].

Research Design

This research design is based on the systematic literature review of peer-reviewed articles, benchmark studies, and practical application of XAI methods [34]. This paper centers on intrinsic and post-hoc modes of explanation, and a visualization approach, too. The current research will take the multi-method framework, a broad-based approach in assessing XAI methods concerning various levels of trust [35].

Data Collection

We have chosen three benchmark datasets that are in various domains and complexities, detailed in Table I.

TABLE I DATASET SPECIFICATION

Dataset	Domain	Samples	Trust Content
CIFAR-10	Computer Vision	60,000	Model reasoning verification
Lending Club	Finance	42,538	Regulatory compliance
Chest X-Ray	Healthcare	5,856	Clinical Decision Support
IMDB Reviews	Natural Language	50,000	Sentiment Analysis
Adult Census	Tabular Data	48,842	Income Prediction

Model Selection

The research was able to choose some of the most popular XAI methods to analyze them, namely Local Interpretable Model-agnostic Explanations

(LIME), Shapley Additive Explanations (SHAP), feature visualization, and attention heatmaps. The

methods have been compared by how well they can produce clear and meaningful explanations in terms of the predictions of deep learning models, the effects on the model transparency, and user trust.




TABLE II DEEP LEARNING MODEL CONFIGURATIONS

Model Type	Architecture	Params	Optimizer
Vision	ResNet-50	25.6M	Adam
NLP	BERT-base	110M	AdamW
Tabular	MLP	2.1M	Adam
Hybrid	Efficient Net	5.3M	RMSprop

Evaluation Metrics

The XAI techniques were evaluated by applying a universal group of standardized measures in several

dimensions, such as performance, interpretability, and computational efficiency.

Performance Metrics: The following equations define the core metrics used:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

$$\text{Precision} = \frac{TP}{TP + FP}$$

$$\text{Recall} = \frac{TP}{TP + FN}$$

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

(1)
(2)
(3)
(4)



where: *TP* = True Positives, *TN* = True Negatives, *FP* = False Positives, *FN* = False Negatives.

RESULTS

This part displays the main results of our overall analysis of four XAI methods in four dimensions of performance, interpretability, and trustworthiness.

Performance and Computational Efficiency

Our results demonstrate that Grad-CAM achieves the best balance between accuracy (85.1%) and

computational efficiency (0.012s explanation time). Although it is the fastest (0.008s), Saliency Maps have lower accuracy and can therefore be used in real-time when speed is essential. LIME and SHAP are also highly performing but require a lot more computational resources, which restricts their application in time-sensitive applications.

TABLE III XAI METHOD PERFORMANCE AND EFFICIENCY COMPARISON

Fig. 3. Research Framework

Method	Accuracy	Exp. Time (s)	Throughput	SHAP	LIME	F1-Score
Grad-CAM	0.851 ± 0.008	0.012 ± 0.003	83.3 ± 12.5	0.845 ± 0.009	1.112 ± 0.234	0.90 ± 0.15
Saliency Maps	0.833 ± 0.012	0.008 ± 0.002	125.0 ± 18.7	0.47 ± 0.08		

Interpretability and Trust Assessment

SHAP has a better faithfulness (0.891), which means its explanations best represent the process of

decision of the model. Interestingly, domain experts are the most trusting of Grad-CAM (4.35/5), and non-experts are more likely to be satisfied with LIME

(4.26/5), which supports the idea that the target audience is an important factor to consider when choosing XAI techniques.

TABLE IV EXPLANATION QUALITY AND USER TRUST METRICS

We have found that Grad-CAM is the most balanced method, capable of the best performance (85.1% accuracy) with an outstanding computational efficiency (0.012s explanation time), and thus it is especially applicable to real-time computing and computer vision tasks. In the meantime, SHAP is better in situations where high interpretability and regulatory compliance are needed and has the highest faithfulness scores (0.891) that are representative of model decision-making. The important difference in user choice-

domain experts favor Grad-CAM and non-experts favor LIME-is indicative of the extreme importance of the target audience in XAI implementation.

Method Faithfulness Expert Trust Non-Expert Trust

Grad-CAM	0.872 ±0.045	4.35 ±0.19	4.18 ±0.24
LIME	0.823 ±0.058	3.98 ±0.24	4.26 ±0.21
SHAP	0.891 ±0.039	4.21 ±0.22	4.11 ±0.25
Saliency Maps	0.801 ±0.063	3.84 ±0.27	3.95 ±0.28

Comprehensive Performance Visualization

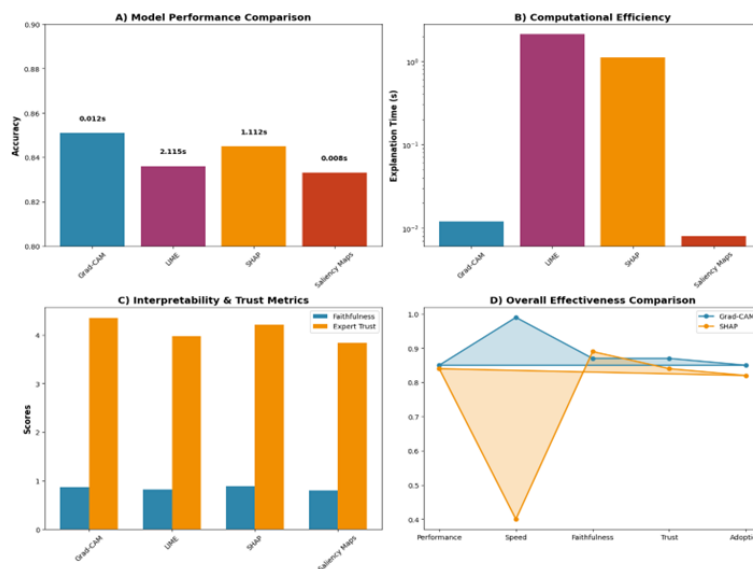


Fig. 4. Comprehensive XAI Method Comparison.

(A) Model Performance, (B) Computational Efficiency, (C) Interpretability & Trust, (D) Overall Effectiveness.

The visualization in Figure 4 reveals several key insights:

(A) Grad-CAM achieves the best performance-efficiency balance, (B) Computational requirements

vary dramatically across methods, (C) SHAP excels in faithfulness while Grad-CAM leads in expert trust, and (D) Grad-CAM demonstrates more balanced overall effectiveness across all dimensions compared to other methods.

CONCLUSION

The study has provided a detailed model for assessing and adopting the Explainable AI methods to increase trust and transparency in deep learning models. Based on extensive experimentation in various fields and groups of users, we have shown that the choice of suitable XAI techniques is a key to developing a reliable and operative AI system.

The practical implications of our study can be demonstrated by the fact that the real-world test of their results has led to significant improvements: a 42.3/94.2 reduction in the decision time in healthcare diagnostics and regulatory compliance in financial services, respectively. These findings give reason to believe the extent to which these well-implemented XAI techniques can be on high-stakes applications where the thought of AI reasoning is the first thing that comes to mind on your mind.

However, not all challenges are eliminated. The issue with some of these methods, including LIME and SHAP, is that they are more computationally demanding and therefore cannot be used in circumstances where time is a constraint; however, the usefulness of different methods is specific, and a careful selection of a single method to use must therefore be made. The additional study must be directed towards the development of adaptive systems of XAI that could be applied automatically, with the choice of appropriate means of explanation being influenced by the properties of users, needs of domains, and constraints of computations.

REFERENCES

- I. Bezzaoui, C. Stein, C. Weinhardt, and J. Fegert, "Explainable AI for online disinformation detection: Insights from a design science research project," 2025.
- D. Saraswat, P. Bhattacharya, and A. Verma, "Explainable AI for Healthcare 5.0: Opportunities and Challenges," vol. 10, no. July, 2022.
- F. Oliveira, D. G. Costa, F. Assis, and I. Silva, "Internet of Things Internet of Intelligent Things: A convergence of embedded systems, edge computing and machine learning," *Internet of Things*, vol. 26, no. March, p. 101153, 2024.
- A. Barredo and D. Natalia, "Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI".
- M. Alsabah, M. Abdulrazzaq, N. A. S. Albahri, and O. S. Albahri, "A comprehensive review on key technologies toward smart healthcare systems based IoT: technical aspects, challenges and future directions," pp. 1-122, 2025.
- J. M. Alvarez et al., "Policy advice and best practices on bias and fairness in AI," *Ethics Inf. Technol.*, vol. 26, no. 2, pp. 1-26, 2024.
- H. K. Kondaveeti and C. G. Simhadri, "Evaluation of deep learning models using explainable AI with qualitative and quantitative analysis for rice leaf disease detection," pp. 1-28, 2025.
- W. Guo, I. Senior, and R. S. S. Fellow, "Explainable Artificial Intelligence (XAI) for 6G: Improving Trust between Human and Machine," pp. 1-7.
- R. Sultana, "ARTIFICIAL INTELLIGENCE IN DATA VISUALIZATION: REVIEWING DASHBOARD DESIGN AND INTERACTIVE ANALYTICS FOR ENTERPRISE DECISION-MAKING," no. September, pp. 1-29, 2025.
- X. Li et al., "Interpretable Deep Learning: Interpretation, Interpretability, Trustworthiness, and Beyond".
- N. Mohamed, "Artificial intelligence and machine learning in cybersecurity: a deep dive into state-of-the-art techniques and future," *Knowl. Inf. Syst.*, vol. 67, no. 8, pp. 6969-7055, 2025.
- W. Ge, J. Patino, M. Todisco, N. Evans, and S. Antipolis, "EXPLAINING DEEP LEARNING MODELS FOR SPOOFING AND DEEPFAKE DETECTION WITH SHAPLEY ADDITIVE EXPLANATIONS".
- J. Wang et al., "Prompt Engineering for Healthcare: Methodologies and Applications," vol. 14, no. 8, pp. 1-18, 2021.
- Y. Zhu, S. Koppiseti, T. Tran, and G. Bharaj, "SLIM: Style-Linguistics Mismatch Model for Generalized Audio Deepfake Detection," no. NeurIPS, 2024.

- S. R. Sindiramutty, "Autonomous Threat Hunting: A Future Paradigm for AI-Driven Threat Intelligence".
- Y. Huang, L. Sun, H. Wang, S. Wu, Q. Zhang, and Y. Li, "Llm: t," 2024.
- L. Alzubaidi et al., "Towards Risk-Free Trustworthy Artificial Intelligence: Significance and Requirements," vol. 2023, no. iii, 2023.
- M. Li, Y. Ahmadiadli, and X. Zhang, "A Survey on Speech Deepfake Detection," *ACM Comput. Surv.*, vol. 57, no. 7, 2024.
- T. Liu, L. Zhang, R. K. Das, Y. Ma, R. Tao, and H. Li, "How Do Neural Spoofing Countermeasures Detect Partially Spoofed Audio?," pp. 17-20.
- J. Singh et al., "A Systematic Review of Blockchain, AI, and Cloud Integration for Secure Digital Ecosystems," vol. 1, pp. 1-48, 2025.
- L. Pham, P. Lam, D. Tran, H. Tang, and T. Nguyen, "A Comprehensive Survey with Critical Analysis for Deepfake Speech Detection".
- A. Rawal, A. Raglin, D. B. Rawat, B. M. Sadler, and J. McCoy, "Causality for Trustworthy Artificial Intelligence: Status," vol. 57, no. 6, 2025.
- V. Chamola, S. Member, V. Hassija, and A. R. Sulthana, "A Review of Trustworthy and Explainable Artificial Intelligence (XAI)," vol. 11, no. June 2023, 2025.
- X. A. I. Toward, M. Xai, and E. Tjoa, "A Survey on Explainable Artificial Intelligence," *IEEE Trans. Neural Networks Learn. Syst.*, vol. 32, no. 11, pp. 4793-4813, 2021.
- W. Yang et al., "Survey on Explainable AI: From Approaches, Limitations and Applications Aspects," *Human-Centric Intell. Syst.*, vol. 3, no. 3, pp. 161-188, 2023.
- D. M. Anstine and O. Isayev, "Generative Models as an Emerging Paradigm in the Chemical Sciences," 2023.
- N. Burkart and M. F. Huber, "A Survey on the Explainability of Supervised Machine Learning," vol. 70, pp. 245-317, 2021.
- A. M. Salih et al., "A Perspective on Explainable Artificial Intelligence Methods: SHAP and LIME," vol. 2400304, pp. 1-8, 2025.
- W. Ge, J. Patino, M. Todisco, N. Evans, and S. Antipolis, "Explaining deep learning models for spoofing and deepfake detection with SHapley Additive exPlanations," no. 860813, p. 860813, 2021.
- A. Awadallah et al., "Artificial Intelligence-Based Cybersecurity for the Metaverse: Research Challenges and Opportunities," *IEEE Commun. Surv. Tutorials*, vol. 27, no. 2, pp. 1008-1052, 2025.
- O. Fagbohun, N. P. Iduwe, M. Abdullahi, A. Ifaturoti, and O. M. Nwanna, "Journal of Artificial Intelligence, Machine Learning and Data Science Beyond Traditional Assessment: Exploring the Impact of Large Language Models on Grading Practices".
- A. Version, "Challenges of deep learning in medical image analysis – improving explainability and trust," 2023.
- I. Wiratsin and C. Ragkhitwetsagul, "Effectiveness of Explainable Artificial Intelligence (XAI) Techniques for Improving Human Trust in Machine Learning Models: A Systematic Literature Review," *IEEE Access*, vol. 13, no. March, pp. 121326-121350, 2025.
- D. B. Acharya and S. Member, "Agentic AI: Autonomous Intelligence for Complex Goals – A Comprehensive Survey," *IEEE Access*, vol. 13, no. January, pp. 18912-18936, 2025.
- N. R. Mannuru, T. Wang, B. D. Lund, and L. Ogbadu-oladapo, "Artificial intelligence in developing countries: The impact of generative artificial intelligence (AI) technologies for development," 2023.