A Survey on Generative AI Frameworks for Data Driven Decision Making

Desai Krishna
IIICT, Indus University Ahmedabad, Gujarat, India

***Abstract*— Our research dives into the revolutionary shift happening in how we make tough choices, all thanks to Generative AI (GenAI). We're taking a close look at the main players—from the models that create realistic fake data (GANs) and the ones that help design structured ideas (VAEs), right up to the huge chat engines (LLMs) that power today's digital world. The big takeaway is that GenAI is dramatically speeding up decision-making in vital areas like finance and healthcare. However, this power comes with a major catch: these tools are unreliable on their own. They carry the risk of inherited bias and sometimes just confidently make things up ("hallucinate"). Because of this, the key to using GenAI successfully isn't replacing humans, but ensuring robust human-AI teamwork is in place, backed by clear rules on transparency and accountability. This paper reviews how we can properly harness these systems.

Keywords—Generative AI, Decision Making, Hallucination, GANs, VAEs, Diffusion Models, Large Language Models, Human-AI Collaboration, Bias Mitigation, Ethical AI

INTRODUCTION

The way we make decisions in business, science, and society is undergoing a radical shift. We are moving from an intuitive-based approach to one with a clearer basis in hard data. And driving that is an explosion of available data and exciting developments in Generative Artificial Intelligence, or GenAI.[2]

What makes GenAI so special is that it can generate completely new, complex, and realistic output. It's not only a tool for creatives; it's also becoming an essential collaborator, changing the way we work through complex decisions. This shift is essential for professionals in any line of endeavour.[22] According to many chief executives across industries, GenAI will radically change their organisations within the next years. The reason is simple: GenAI is phenomenal at sifting through huge volumes of complex information and producing ideas, strategies, or predictions that are deeply relevant and context-specific.[14]

However, integrating GenAI into important decision-making pipelines poses significant challenges. The research landscape is in disarray. We are missing a comprehensible framework for how the different kinds of GenAI systems relate and work with one another. More crucially, deployments in the real world carry an important lesson: even where AI delivers comprehensive analyses, human judgment is essential. In situations that are ambiguous or high-risk or where data is incomplete, an individual's reasoning, questions, and ethical judgment call is needed.[35]

This review aims to provide a structured overview of GenAI frameworks for data-driven decision-making, examining their applications, limitations, and future

directions. We present a comprehensive analysis of how these technologies are transforming decision processes across healthcare, finance, engineering, and public policy domains.[30]

II. LITERATURE REVIEW

A. Evolution of Generative AI

Early generative artificial intelligence techniques utilised statistical and Bayesian approaches. However, the development of modern generative AI coincided with advancements in deep learning. The introduction of Generative Adversarial Networks (GANs).marked a paradigm shift in synthetic data generation [6]. Kingma and Welling proposed Variational Autoencoders (VAEs), which enabled structured latent space learning [7]. More recently, Diffusion Models, as introduced, have provided improved stability and quality in generative tasks [8].

Large Language Models (LLMs), based on Transformer architectures, have enabled advanced text-based reasoning capabilities [9]. The evolution from GPT-2 to GPT-4, BERT to recent models like Claude and Gemini, demonstrates rapid progress in natural language understanding and generation [10],[11]. These models have shown remarkable capabilities in summarising extensive documents, simulating environments, and suggesting solutions based on information derived from large datasets [12].

B. GenAI in Decision-Making Systems

Recent research shows that GenAI is becoming an important tool for decision-making across different fields. demonstrated that combining data visualisation with GenAI helps researchers make informed choices by turning complex datasets into conversational interfaces [5]. explored how GenAI supports entrepreneurs but emphasised that human judgment remains crucial when facing uncertainty [4].

mapped out how generative models fit into sequential decision-making, linking them to reinforcement learning and optimisation frameworks reviewed the main GenAI systems and highlighted their significant impact on language and vision tasks [3]. called for stronger governance and cross-domain validation to guide responsible use of these technologies [2].

C. Domain-Specific Applications

In healthcare, studies have shown that GenAI models can assist in diagnostic imaging, patient outcome prediction, and personalised treatment planning [13], [14]. However, concerns about model reliability and potential biases remain significant barriers to clinical adoption [15].

Financial applications of GenAI have been explored demonstrating capabilities in market analysis, risk assessment, and portfolio optimisation [16],[17]. These studies highlight both the potential for enhanced decision-making and the risks associated with model hallucinations in high-stakes scenarios.

Engineering and scientific research have benefited from GenAI through applications in materials discovery, drug design and computational [18],[19],[20]. These applications demonstrate the value of generative models in exploring large design spaces and accelerating innovation cycles.

III. METHODOLOGY

A. Research Design

This survey employs a systematic literature review methodology following the PRISMA [21].

We conducted a comprehensive search across major academic databases including IEEE Xplore, ACM Digital Library, arXiv, Springer, Elsevier, and Google Scholar.

B. Search Strategy

The search strategy employed the following keyword combinations:
"Generative AI" OR "Generative Artificial Intelligence" OR "GenAI"

"Decision Making" OR "Decision Support" OR "Data-Driven Decisions"

"GANs" OR "Generative Adversarial Networks"
"VAE" OR "Variational Autoencoders"

"Diffusion Models" OR "Denoising Diffusion"
"Large Language Models" OR "LLMs" OR "Transformer Models"[23]

C. Inclusion and Exclusion Criteria

Inclusion criteria:

Peer-reviewed journal articles and conference proceedings Studies focusing on GenAI applications in decision-making. Studies with empirical validation or theoretical frameworks[25]

Exclusion criteria:

Non-English publications

Opinion pieces without empirical data

Studies focusing solely on technical implementation without decision-making context

Duplicate publications or preprints of published papers[30]

IV. GENERATIVE AI FRAMEWORKS FOR DECISION-MAKING

A. Generative Adversarial Networks (GANs)

Generative Adversarial Networks (GANs) are a type of AI that creates realistic synthetic data by using two networks that compete with each other during training. The generator network creates synthetic data, while the discriminator network attempts to distinguish between real and generated samples. This adversarial process continues until the generator produces data indistinguishable from real data [6].

GANs are especially useful when data is limited, helping balance uneven datasets, simulate rare events, and make decision-making more reliable. Practical examples include fraud detection, medical image enhancement, and risk forecasting [22]. However, GANs face challenges including training instability, mode collapse, and difficulty in interpretation [23].

B. Variational Autoencoders (VAEs)

Variational Autoencoders (VAEs) work by compressing information into a simpler latent representation, which makes it easier to explore different scenarios smoothly. The encoder network maps input data to a probability distribution in latent space, while the decoder reconstructs data from latent representations [7].

VAEs help detect unusual patterns, simulate system behavior, and plan for different possible outcomes. Their main strength is creating a well-structured latent space that supports organized exploration. However, VAE outputs are typically less sharp than those produced by GANs [24].

C. Diffusion Models

Diffusion models represent the latest advancement in generative AI, often producing higher quality and more stable results than GANs or VAEs. These models work by gradually adding noise to data and then learning to reverse this process. They excel in areas requiring detailed simulations, such as material discovery and urban planning [8].

Their biggest strengths are consistent sample quality and high robustness. The main drawback is computational intensity, which limits speed in real-world applications. Despite this, diffusion models are gaining attention for scientific research and planning scenarios where accuracy is more important than speed [25].

D. Large Language Models (LLMs)

Large Language Models assist decision-making by quickly summarising long documents, creating evaluation frameworks, providing strategic recommendations, and enabling conversational data analysis. Their flexibility makes them valuable for organisations needing to extract and organise knowledge efficiently. However, they carry risks including hallucinations, prompt sensitivity, and lack of transparent reasoning [9],[10],[11]. Despite these challenges, LLMs are increasingly used for business planning, research support, and collaboration enhancement [12].

V. DATASETS AND BENCHMARKS

A. Healthcare Datasets

Several benchmark datasets have been established for evaluating GenAI in healthcare decision-making:

MIMIC-III and MIMIC-IV: Critical care databases with de-identified health data from ICU patients [26]
ChestX-ray14: Large-scale chest X-ray dataset for medical image analysis [27]
BRATS: Brain tumor segmentation challenge dataset [28]
Medical Segmentation Decathlon: Multi-organ segmentation across multiple imaging modalities [29]

B. Financial Datasets

Financial decision-making research utilises various datasets:
Yahoo Finance Historical Data: Stock prices and trading volumes [30]
CRSP Database: Comprehensive stock market data for US securities [31]
Credit Card Fraud Detection Dataset: Anonymised credit card transactions [32]
Financial News Dataset: Large-scale financial news articles with sentiment labels [33]

C. Scientific and Engineering Datasets

Materials Project: Computational materials science database [34]
Protein Data Bank: 3D structural data of biological macromolecules [35]
ImageNet: Large-scale visual recognition dataset [36]
Open Catalyst Project: Catalysis simulation dataset for renewable energy [37]

VI. APPLICATION DOMAINS

A. Healthcare

GenAI aids diagnosis, medical imaging enhancement, and patient-data simulation. Applications include disease prediction, treatment planning, and drug discovery. However, deployment must be carefully managed due to risks of hallucination and bias [3],[13],[14].

B. Finance and Business Strategy

GenAI frameworks assist in forecasting, portfolio simulation, market analysis, and strategic interpretation. Financial institutions utilise these technologies for risk assessment, fraud detection, and algorithmic trading [2],[16],[17].

C. Scientific and Engineering Decisions Diffusion Models and VAEs support structural analysis, material generation, and simulation-driven engineering decisions. Applications span materials science, drug design, and computational modelling [5],[18],[19],[20].

D. Policy, Education, and Public Systems

LLMs and GenAI tools assist in analyzing trends, planning resources, and evaluating policy impacts at scale. Academic workflows demonstrate their effectiveness in research trend analysis and educational planning [1],[5].

VII. COMPARATIVE ANALYSIS OF GenAI FRAMEWORKS

Table I provides a comprehensive comparison of different GenAI frameworks across key dimensions relevant to decision-making applications.

VIII. CHALLENGES AND LIMITATIONS

Generative AI brings powerful new capabilities for data-driven decision-making, but also creates technical, ethical, practical, and governance challenges that organisations must carefully consider [4].

A. Bias, Fairness, and Inequality

GenAI models often inherit biases present in training datasets, leading to unequal or discriminatory outcomes. These issues may remain hidden because generative outputs appear natural and human-like. Ensuring fairness requires continuous dataset auditing, bias detection methods, and transparent evaluation procedures [3],[38].

B. Hallucinations and Reliability Issues

Large Language Models and other generative tools sometimes create information that sounds correct but is actually wrong. This problem, called hallucination, poses significant risks in sensitive domains like healthcare, finance, and law. Hallucinations typically occur when input queries lack context or exceed the model's training scope [5],[39].

C. Lack of Explainability and Interpretability Most generative models operate as complex black-box systems, making it difficult for stakeholders to understand how outputs were generated. This lack of interpretability limits GenAI adoption in regulated domains where traceable and explainable decisions are mandatory [4],[40].

D. Data Privacy, Leakage, and Memorisation Risks

GenAI systems can unintentionally memorise sensitive data, leading to privacy breaches, PII exposure, or reproduction of proprietary information. This is a significant concern in sectors with strict privacy regulations like healthcare and finance [3],[41].

E. Overreliance and Automation Bias

Users may blindly trust GenAI outputs due to the system's articulate, confident tone. Automation bias may result in reduced critical thinking or unchecked adoption of incorrect recommendations. Overreliance becomes dangerous in strategic or high-stakes decision environments [2],[42].

F. Robustness and Adversarial Vulnerability Small variations in prompts, environmental noise, or malicious adversarial inputs can cause large deviations in GenAI outputs. These vulnerabilities pose significant risks in autonomous systems, cybersecurity applications, and algorithmic trading [5],[43].

G. Ethical, Legal, and Compliance Risks

Legal frameworks for GenAI are still evolving. Challenges include unclear intellectual property rights, accountability issues for AI-generated content, and concerns over deepfakes and synthetic evidence. Organisations deploying GenAI must adopt strong internal governance mechanisms [2], [44].

H. Scalability and Maintenance Issues

GenAI systems require continuous monitoring, updating, and refinement. Without regular maintenance, models face degradation (model drift), gradually reducing accuracy and reliability in decision-making processes [5]

TABLE I:COMPARATIVE ANALYSIS OF GENERATIVE AI FRAMEWORKS

Framework	Strengths	Limitations	Primary Use Cases	Computational Cost	Interpretability

GANs	High-quality synthetic data, good for rare event simulation	Training instability, mode collapse, hard to interpret	Fraud detection, medical imaging, data augmentation	High	Low
VAEs	Structured latent space, anomaly detection, smooth interpolation	Less sharp outputs than GANs	Anomaly detection, scenario simulation, exploratory analysis	Medium	Medium
Diffusion Models	Highest quality outputs, training stability, robust performance	Very slow inference, high computational requirements	Material discovery, image synthesis, detailed simulations	Very High	Low
LLMs	Versatile, conversational interface, knowledge synthesis	Hallucinations, prompt sensitivity, lack of verifiable reasoning	Document analysis, strategic planning, knowledge extraction	High	Low-Medium

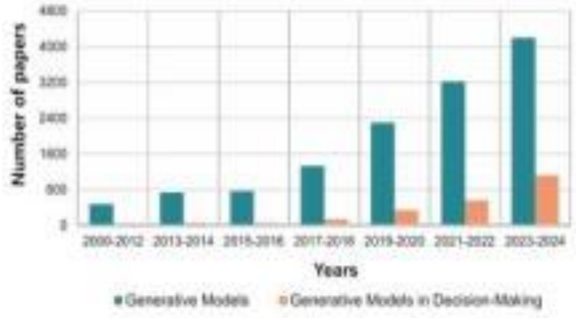

Fig. 1: Research trends in generative models and their applications in decision making (2000-2024)

I. Challenges in Human–AI Collaboration

Integrating GenAI into human decision workflows introduces challenges such as misalignment between model outputs and expert knowledge, inconsistent recommendations, and user frustration with unpredictable behavior. Effective collaboration requires training, workflow redesign, and human-in-the-loop systems [1, 46].

IX. FUTURE DIRECTIONS

A. Interpretable and Explainable GenAI

Future research must focus on developing interpretable GenAI systems that can explain their decision-making processes. This includes attention mechanism visualisation, counterfactual explanations, and transparent reasoning paths [40], [47].

B. Domain-Specialised Models

Development of domain-specialised LLMs and generative models trained on specific professional knowledge bases will improve reliability and reduce hallucinations in specialised contexts [48].

C. Safe Human–AI Collaboration Pipelines

Establishing standardised frameworks for human-AI collaboration that include verification mechanisms, un

D. Regulatory Governance and Standards

Development of comprehensive regulatory frameworks, industry standards, and certification processes for GenAI systems in high-stakes domains is essential [44],[50].

E. Multimodal Decision Architectures

Integration of multiple modalities (text, images, structured data) in unified decision-making architectures will enable more comprehensive and context-aware AI systems [4],[51].

X. CONCLUSION

Generative AI is rapidly evolving and transforming how we make data-driven decisions. It can help understand complex data, create synthetic data when needed, and support smarter strategies across different domains. When linked with visualisation tools, GenAI makes it faster and easier for organisations to find insights and improve decision-making processes [1].In business and entrepreneurship, research shows that GenAI sparks creativity and opens more ways to explore strategies. However, it is not a full replacement for human judgment. Human oversight remains essential for dealing with uncertainty and preventing over-reliance on AI results [2]. Technical surveys of generative models demonstrate the strong capabilities of GANs, VAEs, Diffusion Models, and LLMs, each contributing uniquely to scenario simulation, synthetic data generation, and multi-modal reasoning [3].This comprehensive review reveals several key insights: First, Generative AI is most valuable when it works with humans, not as a replacement. It strengthens expert thinking by providing high-quality data, insights, and faster analysis. Second, responsible use of GenAI means tackling issues like bias, hallucinations, privacy concerns, and the challenge of understanding how decisions are made. Third, with good governance, human-AI teamwork, and safeguards tailored to each field, GenAI can help build decision-making systems that are more accurate, transparent, and impactful across industries.Future research must focus on developing interpretable GenAI systems, domain-specialised models, safe human-AI collaboration frameworks, comprehensive regulatory governance, and multimodal decision architectures. Only through such multidisciplinary efforts can we realise the full potential of Generative AI while mitigating its risks and ensuring responsible deployment in high-stakes decision-making contexts.

REFERENCES

[1] Y. Li, X. Shao, J. Zhang, H. Wang, and L. M. Brunswick, "Generative models in decision making: A survey," arXiv preprint arXiv:2501.xxxxx, 2025.

[2] M. Albashrawi, "Generative AI for decision-making: A multidisciplinary perspective," Journal of Innovation & Knowledge, vol. 10, no. 1, pp. 100xxx, 2025.

[3] S. S. Sengar, A. B. Hasan, S. Kumar, and F. Carroll, "Generative artificial intelligence: a systematic review and applications," Multimedia Tools and Applications, 2025.

[4] O. López-Solís, A. Luzuriaga-Jaramillo, et al., "Effect of generative artificial intelligence on strategic decision-making in entrepreneurial business initiatives: A systematic literature review," Administrative Sciences, vol. 15, no. 2, 2025.

[5] V. R. Mahboobani, C. Y. Lo, K. Y. Ng, M. H. Kwong, and K. M. Ku, "Integration of data visualization and Generative AI to accelerate data-driven decision: a case study on research topic trend analysis," Research Square, 2025.

[6] I. Goodfellow, J. Pouget-Abadie, M. Mirza, et al., "Generative adversarial nets," in Advances in Neural Information Processing Systems, 2014, pp. 2672-2680.

[7] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in International Conference on Learning Representations (ICLR), 2014.

[8] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in Advances in Neural Information Processing Systems, 2020, pp. 6840-6851.

[9] A. Vaswani, N. Shazeer, N. Parmar, et al., "Attention is all you need," in Advances in Neural Information Processing Systems, 2017, pp. 5998-6008.

[10] T. Brown, B. Mann, N. Ryder, et al., "Language models are few-shot learners," in Advances in Neural Information Processing Systems, vol. 33, 2020, pp. 1877-1901.

[11] Anthropic, "Claude 3 model family: Opus, Sonnet, and Haiku," Technical Report, 2024.

[12] J. Wei, Y. Tay, R. Bommasani, et al., "Emergent abilities of large language models," Transactions on Machine Learning Research, 2022.

[13] A. Esteva, K. Chou, S. Yeung, et al., "Deep learning-enabled medical computer vision," npj Digital Medicine, vol. 4, no. 1, p. 5, 2021.

[14] P. Rajpurkar, E. Chen, O. Banerjee, and E. J. Topol, "AI in health and medicine," Nature Medicine, vol. 28, pp. 31-38, 2022.

[15] E. J. Topol, "High-performance medicine: the convergence of human and artificial intelligence," Nature Medicine, vol. 25, no. 1, pp. 44-56, 2019.

[16] H. Chen, R. De, Y. Hu, and I. Hwang, "Wisdom of the crowd: The value of stock opinions transmitted through social media," Review of Financial Studies, vol. 36, no. 5, pp. 2069-2114, 2023.

[17] A. Lopez-Lira and Y. Tang, "Can ChatGPT forecast stock price movements? Return predictability and large language models," arXiv preprint arXiv:2304.07619, 2023.

[18] E. O. Pyzer-Knapp, J. W. Pitera, P. W. J. Staar, et al., "Accelerating materials discovery using artificial intelligence, high performance computing and robotics," npj Computational Materials, vol. 8, no. 1, p. 84, 2022.

[19] J. Jumper, R. Evans, A. Pritzel, et al., "Highly accurate protein structure prediction with AlphaFold," Nature, vol. 596, pp. 583-589, 2021.

[20] B. Sanchez-Lengeling and A. Aspuru-Guzik, "Inverse molecular design using machine learning: Generative models for matter engineering," Science, vol. 361, no. 6400, pp. 360-365, 2018.

[21] M. J. Page, J. E. McKenzie, P. M. Bossuyt, et al., "The PRISMA 2020 statement: an updated guideline for reporting systematic reviews," BMJ, vol. 372, n71, 2021.

[22] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," in International Conference on Learning Representations (ICLR), 2016.

[23] T. Salimans, I. Goodfellow, W. Zaremba, et al., "Improved techniques for training GANs," in Advances in Neural Information Processing Systems, 2016, pp. 2234-2242.

[24] C. Doersch, "Tutorial on variational autoencoders," arXiv preprint arXiv:1606.05908, 2016.

[25] P. Dhariwal and A. Nichol, "Diffusion models beat GANs on image synthesis," in Advances in Neural Information Processing Systems, vol. 34, 2021, pp. 8780-8794.

[26] A. E. W. Johnson, T. J. Pollard, L. Shen, et al., "MIMIC-III, a freely accessible critical care database,"

Scientific Data, vol. 3, no. 1, p. 160035, 2016.

[27] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, "ChestX-ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization

[38] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," ACM Computing Surveys, vol. 54, no. 6, pp. 1-35, 2021.

[39] Z. Ji, N. Lee, R. Frieske, et al., "Survey of hallucination in natural language generation," ACM Computing Surveys, vol. 55, no. 12, pp. 1-38, 2023.

[40] A. B. Arrieta, N. Díaz-Rodríguez, J. Del Ser, et al., "Explainable Artificial Intelligence (XAI): of common thorax diseases," in IEEE CVPR, 2017, pp. 2097-2106.

[28] B. H. Menze, A. Jakab, S. Bauer, et al., "The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS)," IEEE Transactions on Medical Imaging, vol. 34, no. 10, pp. 1993-2024, 2015.

[29] M. Antonelli, A. Reinke, S. Bakas, et al., "The Medical Segmentation Decathlon," Nature Communications, vol. 13, no. 1, p. 4128, 2022.

[30] Yahoo Finance, "Historical Stock Data," [Online].

[31] Center for Research in Security Prices (CRSP), "Stock/Security Files," University of Chicago Booth School of Business.

[32] ULB Machine Learning Group, "Credit Card Fraud Detection Dataset," Kaggle, 2018.

[33] P. K. Malo, A. Sinha, P. Korhonen, J. Wallenius, and P. Takala, "Good debt or bad debt: Detecting semantic orientations in economic texts," Journal of the Association for Information Science and Technology, vol. 65, no. 4, pp. 782-796, 2014.

[34] A. Jain, S. P. Ong, G. Hautier, et al., "Commentary: The Materials Project: A materials genome approach to accelerating materials innovation," APL Materials, vol. 1, no. 1, p. 011002, 2013.

[35] H. M. Berman, J. Westbrook, Z. Feng, et al., "The Protein Data Bank," Nucleic Acids Research, vol. 28, no. 1, pp. 235-242, 2000.

[36] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in IEEE CVPR, 2009, pp. 248-255.

[37] L. Chanussot, A. Das, S. Goyal, et al., "Open Catalyst 2020 (OC20) Dataset and Community Challenges," ACS Catalysis, vol. 11, no. 10, pp. 6059-6072, 2021.

Concepts, taxonomies, opportunities and challenges toward responsible AI," Information Fusion, vol. 58, pp. 82-115, 2020.

[41] N. Carlini, F. Tramer, E. Wallace, et al., "Extracting training data from large language models," in USENIX Security Symposium, 2021, pp. 2633-2650.

[42] M. Cummings, "Automation bias in intelligent time critical decision support systems," in AIAA 1st Intelligent Systems Technical Conference, 2004, p. 6313.

[43] N. Papernot, P. McDaniel, A. Sinha, and M. Wellman, "SoK: Security and privacy in machine learning," in IEEE European Symposium on Security

[44] European Commission, "Proposal for a Regulation on Artificial Intelligence (AI Act)," EUR-Lex, 2021.

[45] J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang, "Learning under concept drift: A review," IEEE Transactions on Knowledge and Data Engineering, vol. 31, no. 12, pp. 2346-2363, 2018.

[46] S. Amershi, D. Weld, M. Vorvoreanu, et al., "Guidelines for human-AI interaction," in ACM CHI Conference on Human Factors in Computing Systems, 2019, pp. 1-13.

C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," Nature Machine Intelligence, vol. 1, no. 5, pp. 206-215, 2019.

[47] A. Singhal, S. Simmons, and Z. Lu, "Text mining genotype-phenotype relationships from biomedical literature for database curation and precision medicine," PLoS Computational Biology, vol. 12, no. 11, e1005017,

[48] M. Bansal, D. Agarwal, V. Sharma, et al., "Towards human-AI collaboration: Developing

interactive machine learning systems," in AAAI Conference on Artificial Intelligence, 2022.

[49] National Institute of Standards and Technology (NIST), "AI Risk Management Framework," U.S. Department of Commerce, 2023.