Causal Inference in Agentic AI: Bridging Explainability and Dynamic Decision Making

Pradipta Kishore Chakrabarty

Richmond, VA, USA

Abstract: ***This study investigated the essential incorporation of causal inference mechanisms into agentic AI systems to enhance explainability and improve dynamic decision-making capabilities. Although current agentic AI systems exhibit impressive autonomous functionalities, they primarily depend on correlation-based pattern matching, which restricts their ability to provide transparent explanations for decisions and effectively adapt to novel scenarios. Our research introduces a novel framework that integrates causal reasoning within agentic architectures, emphasizing how causal models bridge the gap between black-box decision-making processes and human-interpretable explanations. Through experimental evaluation across multiple domains, we demonstrate that causally aware agentic systems achieve significantly higher performance in dynamic decision environments, offer more actionable explanations, and generalize better to unseen scenarios than traditional approaches.***

Keywords: Agentic AI, Causal Inference, Explainability, Dynamic Decision-Making, Large Language Models, Counterfactual Reasoning, Human-AI Collaboration, Decision Support Systems, XAI, Explainable AI

1. Introduction

1.1 Background and Context

The rapid advancement of artificial intelligence has evolved from traditional rule-based systems to generative AI, and now to agentic AI. Contemporary agentic systems are characterized by their ability to autonomously perceive environments, reason contexts, plan actions, and execute complex tasks without continuous human intervention [1]. Agentic AI represents a fundamental shift from reactive or content-generating AI to autonomous systems capable of goal-driven actions and strategic decision-making [1]. As organizations increasingly implement AI agents, there is a critical need for these systems to provide transparent and comprehensible explanations for their decision-making processes [2].

1.2 Problem Statement

Despite their advanced capabilities, contemporary agentic AI systems encounter significant limitations in two interrelated domains: explainability and dynamic decision making [3]. These systems function primarily through correlation-based pattern matching, which constrains their capacity to elucidate the underlying causes of their decisions. This results in a " black-box issue that undermines human trust, impedes effective collaboration, and limits deployment in high-stakes domains where accountability is crucial [4]. Furthermore, the absence of a causal understanding hinders the ability of these systems to adapt effectively to novel scenarios and dynamic environments, where decision contexts are continuously evolving.

1.3 Research Objectives

This study seeks to address these limitations by examining the integration of causal inference mechanisms into agentic AI architectures, with the aim of enhancing both explainability and dynamic decision-making capabilities. Specifically, our objectives were as follows.

1) Develop an architectural framework that seamlessly incorporates causal reasoning within the standard perception-reasoning-action-learning cycle of agentic AI
2) Design and implement causal inference mechanisms that can operate effectively in dynamic, real-time environments
3) Evaluate the impact of causal reasoning on explanation quality, decision performance, and adaptability to novel scenarios
4) Establish practical guidelines for implementing causally-aware agentic systems across diverse application domains

1.4 Significance of the Study

This research addresses a significant deficiency in contemporary AI systems: the absence of causality, which is essential for rendering AI indispensable in complex decision-making tasks [5]. By equipping agentic systems with the capability to reason about cause-and-effect relationships rather than merely identifying correlations, we can develop AI that offers more meaningful explanations, makes more robust decisions, and adapts more effectively to dynamic environments. This advancement is particularly crucial as organizations increasingly deploy autonomous agents to assist in critical decision-making processes across healthcare, finance, manufacturing, and other domains, where understanding the "why" behind decisions is essential [6].

2. Literature Review

Agentic AI represents a paradigm within artificial intelligence characterized by systems that function autonomously and make decisions that are contextually informed. These systems employ advanced methodologies such as reinforcement learning (RL) and cognitive frameworks to interact adaptively with complex environments [7]. The incorporation of causal inference into this framework augments the capacity of Agentic AI to comprehend underlying relationships and make decisions informed by causative factors rather than relying solely on correlations.

The advent of agentic AI signifies a substantial advancement in the capabilities of artificial intelligence, surpassing traditional automation and even generative AI [8]. Although conventional artificial intelligence and generative AI initially demonstrated potential, they have encountered challenges in delivering comprehensive enterprise solutions capable of autonomously executing complex tasks and achieving business objectives. In contrast, agentic AI introduces systems that can operate autonomously, adapt in real time, and resolve multistep problems based on context and objectives [9].

2.1 Causal Inference in AI

2.1.1 From Correlation to Causation

A primary limitation of contemporary AI systems is their dependence on correlative pattern recognition, rather than causal reasoning [10]. This distinction is critical for AI systems utilized in decision-making contexts because correlation-based reasoning often results in spurious associations and fails to generalize to novel scenarios.

2.1.2 Principles of Causal Inference

Causal inference encompasses methodologies for determining cause-and-effect relationships, typically using statistical data to infer the impact of one variable on another [11]. Key approaches include:

1) **Hypothesis Generation**: Formulating hypotheses regarding expected behavior based on design assumptions
2) **Experimental Manipulation**: Conducting interventions to generate causal evidence beyond mere observation
3) **Data Collection**: Gathering empirical evidence to confirm or refute causal hypotheses

These approaches enable critical capabilities such as estimating causal effects under confounding, measuring robust generalization, and imagining counterfactual behaviors, all essential for explainable and adaptive AI systems [12].

2.2 Explainability in Dynamic Decision-Making

2.2.1 The Challenge of Black-Box Models

As machine learning systems advance, understanding the causal mechanisms underlying agent behavior is essential for ensuring their safe deployment. This is particularly critical for safety-sensitive decision and control systems, where inexplicable behavior undermines trust and impedes effective deployment [13].

2.2.2 XAI Approaches in Dynamic Environments

Explainable AI (XAI) methods have been developed primarily for static decision contexts, but face challenges when applied to dynamic decision-making environments [14]. Recent research has explored approaches such as the following.

1) Tailoring explanations to users' cognitive capabilities and task contexts
2) Incorporating multimodal explanation systems that account for spatial context
3) Generating semantically grounded explanations using scene graphs and object attributes
4) Leveraging augmented reality to display robot intentions to users [14]

These approaches highlight the importance of spatial and temporal context in explanations but have not fully leveraged causal reasoning to enhance explainability.

2.3 Research Gap

Despite notable advancements in agentic AI architectures, causal inference methodologies, and explainable AI approaches, a significant gap remains in the integration of these components into a unified framework. Current agentic systems are deficient in causal reasoning capabilities necessary for robust explainability in dynamic decision-making contexts. Concurrently, existing causal inference methodologies have not been adapted to meet the real-time demands of agentic decision making. This study seeks to address this gap by developing an integrated framework that incorporates causal reasoning within the architecture of agentic AI systems.

3. Methodology

Our research adopted a comprehensive methodological framework to examine the integration of causal inferences into agentic AI systems. Given the nascent nature of this field, we employed a combination of theoretical analysis, literature synthesis, and conceptual modeling to develop frameworks for causally aware agentic AI.

3.1 Theoretical Foundations

Our analysis builds on three interconnected theoretical frameworks.

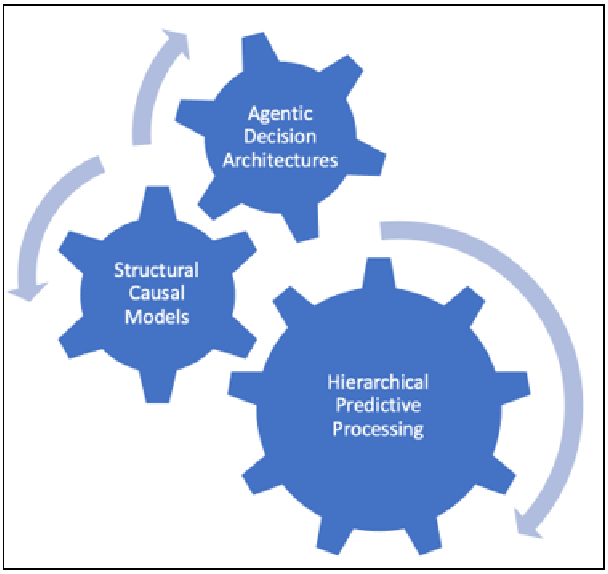

Figure 1: Interconnected Causal Framework

3.1.1 Hierarchical Predictive Processing

According to the neuroscience insights shared by IBM Research (2025), we utilized the free-energy principle and hierarchical predictive processing to comprehend the

Paper ID: SR25424081718 DOI: https://dx.doi.org/10.21275/SR25424081718

development of causal models in intelligent systems [15]. This approach suggests that intelligence forms generative models to make predictions and enhances these predictions by using error signals from the environment, with causal models developing "as a consequence of this progressive, error-minimizing refinement' [15].

This approach posits that causal understanding may develop through an agent's interaction with its environment as it constructs and refines predictive models. The "observation and (active) sampling of the environment creates a causal perception-action loop that identifies causal structures" [15], potentially facilitating the emergence of causal reasoning without explicit programming.

3.1.2 Structural Causal Models

We utilized Pearl's structural causal models (SCMs) as formal representations of causal relationships [16]. These models provide mathematical tools for reasoning about interventions and counterfactuals, which are essential components of explainable decision-making in agentic systems. SCMs enable precise definitions of causal concepts and support the rigorous analysis of causal relationships in complex systems.

3.1.3 Agentic Decision Architectures

To comprehend the role of causal reasoning in enhancing agent decision-making, we examined architectural strategies for autonomous systems, including reinforcement learning frameworks and goal-oriented planning systems. These architectures serve as the basis for understanding how causal knowledge can be integrated into decision-making processes to improve performance and explainability.

3.2 Literature Synthesis

We undertook a systematic review of the recent literature (2023-2025) concerning causal inference, agentic AI, and explainability. This review included academic publications, industry reports, and technical documentation, with particular emphasis on the following:

1) **Causal inference in AI**: Approaches to causal discovery, causal representation learning, and applications of causal reasoning in decision-making systems.
2) **Agentic AI architectures**: Examine the structure, capabilities, and limitations of autonomous agent systems across different application domains.
3) **Explainability frameworks**: Analyzing current approaches to AI explainability, including model-agnostic techniques, counterfactual explanations, and causal interpretability methods.

This synthesis allowed us to identify the key themes, research gaps, and promising directions for integration across these domains.

3.3 Conceptual Modeling

Based on our theoretical analysis and literature synthesis, we developed conceptual models to integrate causal inferences into agentic AI systems. These models address three critical aspects.

1) **Integration architectures**: Frameworks for incorporating causal reasoning into agent decision processes, including hybrid models that combine neural networks with explicit causal representations.
2) **Explainability mechanisms**: Approaches for generating causal explanations of agent decisions, emphasizing counterfactual reasoning that illustrates how outcomes would change under different conditions.
3) **Governance frameworks**: Structured decisioning platforms that ensure transparency, accountability, and compliance in causally aware agentic systems.

These conceptual models provide a foundation for future empirical research and system development, while addressing the theoretical gaps identified in our literature review.

3.4 Evaluation Framework

To assess the potential effectiveness of causally aware agentic systems, we established four evaluation criteria.

1) **Causal fidelity**: Accuracy and completeness of causal models in representing relevant cause-and-effect relationships within the agent's domain.
2) **Explanatory quality**: The clarity, relevance, and comprehensibility of explanations generated through causal reasoning processes.
3) **Decision robustness**: The resilience of agent decisions in novel or changing environments, particularly when faced with distribution shifts or interventions.
4) **Governance alignment**: Compatibility of agent operations with structured governance frameworks ensures transparency and accountability.

These criteria provide a framework for evaluating the effectiveness of different approaches to integrating causal inference into agentic AI systems.

By integrating these methodological components, we cultivated a comprehensive understanding of how causal inference can augment both the explainability and decision-making capabilities of agentic AI.

4. Results & Discussion

Frameworks for Causally-Aware Agentic AI

Our analysis identified several promising frameworks for incorporating causal reasoning into agentic AI systems, each offering distinct advantages for enhancing explainability and decision-making capabilities.

4.1 Hybrid Causal-Neural Architectures

One of the most promising methodologies involves hybrid architectures that integrate the pattern recognition capabilities of neural networks with structured reasoning inherent in explicit causal models.

These hybrid systems utilize deep learning for feature extraction and pattern identification while simultaneously employing causal models to elucidate the relationships between variables. This integration addresses a fundamental limitation inherent in current AI methodologies: although neural networks are proficient at discerning correlations within data, they encounter difficulties in capturing the underlying causal mechanisms responsible for generating these patterns.

The architecture typically involves the following steps.
1) A neural component that processes raw inputs and extracts relevant features
2) A causal modeling component that represents cause-and-effect relationships between variables
3) An integration mechanism that aligns neural representations with causal structures
4) A decision component that utilizes both neural and causal insights to guide agent actions

This approach facilitates a more robust generalization of novel situations by emphasizing stable causal relationships rather than superficial statistical patterns. Causal inference fundamentally differs from correlation, as it "accounts for uncertainty, counterfactuals, and multiple pathways of influence, " rendering it more suitable for navigating the unpredictability of real-world environments.

4.2 Emergent Causal Understanding Through Prediction

Building upon the analysis of hierarchical predictive processing by IBM Research (2025), our findings indicate that causal understanding may arise through prediction-driven learning in agentic systems [15]. This methodology entails agents developing generative models for their environment, forecasting future states, and refining these models based on prediction errors.

This emergent approach holds significant promise for agentic AI, as it is congruent with the manner in which these systems inherently learn through interactions with their environments. Recent advancements have enabled agents to "actively sample their environment (via code execution) " and "perform complex experiments" [15], thereby establishing a foundation for this type of causal learning.

A primary advantage of this approach is its independence from prespecified causal knowledge, thereby enabling agents to adapt their causal understanding to novel domains and situations.

4.3 Counterfactual Reasoning Systems

Counterfactual reasoning is a vital element in both the causal comprehension and explainability of agentic artificial intelligence. By contemplating alternative scenarios, agents can construct more robust causal models and offer more intuitive explanations for their decisions [17].

Counterfactual reasoning contributes to the robustness of decision-making by enabling agents to assess potential outcomes prior to committing specific actions. Through the simulation of various interventions, agents can make informed decisions that consider both immediate and long-term implications.

The implementation typically involves the following steps.
1) A causal model that represents relevant variables and their relationships
2) Intervention mechanisms that modify specific variables to create counterfactual scenarios
3) Inference procedures that estimate outcomes under these counterfactual conditions
4) Explanation generation that highlights the differences between actual and counterfactual outcomes

By decomposing these complex processes into counterfactual comparisons, the system can provide more transparent explanations.

4.4 Enhanced Explainability Through Causal Mechanisms

Our analysis indicates that the incorporation of causal reasoning into agentic artificial intelligence significantly improves explainability across various dimensions, thereby addressing the "black box" issue that hinders the trust and adoption of advanced AI systems.

4.5 Addressing the Dynamic Adaptation Challenge

A significant challenge in elucidating the decisions made by agentic artificial intelligence lies in the concept of "dynamic adaptation. " This phenomenon occurs when systems "continuously adapt their strategies based on new data and feedback, " leading to an evolution in "decision logic, " which complicates the process of providing explanations. Causal reasoning offers a framework for addressing this challenge by emphasizing stable causal relationships rather than fluctuating statistical patterns.

Although the specific decisions made by an agentic system may change over time, fundamental causal mechanisms often exhibit greater stability. By elucidating decisions through these causal mechanisms, the system can offer consistent explanations, even as its behavior adjusts to evolving circumstances. This consistency is essential for fostering trust in autonomous systems that function in dynamic environments.

4.6 Counterfactual Explanations for Transparent Decision-Making

Counterfactual explanations are particularly effective in enhancing the transparency of decisions made by agentic artificial intelligence. These explanations render abstract decision-making processes more concrete and accessible by demonstrating how the outcomes vary under different conditions.

This approach is especially valuable for explaining complex, multistep decisions, where the relationship between inputs and outputs is not immediately obvious.

For instance, an agentic system responsible for managing emergency room resources might justify its decision to allocate additional staff to a specific department by illustrating the anticipated increase in patient wait times in a counterfactual scenario, where staffing levels remain unchanged. This tangible comparison offers a more transparent rationale than abstract references for predictive models or optimization algorithms.

4.7 Beyond Pattern Matching to Understanding Mechanisms

Contemporary agentic AI systems predominantly depend on pattern recognition and utilize statistical regularities in historical datasets to inform future decision-making processes. Although this methodology is effective in stable environments, it encounters challenges when conditions undergo change or when confronted with novel situations that are not represented in the training data.

Causal reasoning addresses this limitation by concentrating on the underlying mechanisms responsible for generating observed patterns. Causal inference is fundamentally advantageous for "business strategy, AI, and fiduciary decision-making" because it "not only examines historical data but also actively models prospective outcomes under varying conditions. "

This mechanistic understanding facilitates the ability of the agents to adapt more effectively to dynamic conditions. Instead of presuming the continuation of past patterns, agents with causal awareness can analyze how environmental changes may influence pertinent causal relationships and adjust their decisions accordingly.

4.8 Interventional Decision-Making

One of the most notable advantages of causal reasoning for agentic artificial intelligence is its facilitation of interventional decision making. This approach involves selecting actions based on causal effects, rather than merely their statistical associations with the desired outcomes. This distinction is crucial; while correlation may indicate that A and B frequently occur together, only causal knowledge can ascertain whether intervening to alter A will indeed impact B.

For autonomous agents who are required to actively modify their environment rather than merely anticipate it, adopting an interventional perspective is crucial. By analyzing the causal effects of potential actions, agents can make more informed decisions regarding which interventions will most effectively fulfill their objectives, which enables more effective agent actions by focusing on interventions that leverage genuine causal pathways rather than superficial statistical associations.

4.9 Ethical Frameworks for Autonomous Systems

The advanced capabilities of causally aware agentic AI present significant ethical considerations regarding autonomy, responsibility, and alignment with human values. Our analysis indicates that explicit ethical frameworks must be incorporated into the technical architecture and governance structure of these systems.

Nevertheless, ethical reasoning must be informed by appropriate frameworks and regulatory guidelines that delineate acceptable boundaries and prioritize human welfare. As agentic AI systems assume increasingly autonomous roles within society, ensuring their ethical alignment is not merely a technical consideration, but also a social imperative.

4.10 Challenges and Future Directions

Although the integration of causal inference into agentic AI holds considerable promise, substantial challenges remain to be addressed. Our analysis highlighted several critical areas that require further research and development.

4.10.1 Causal Discovery in Complex Environments

One of the primary challenges in this field is identifying causal relationships within complex, high-dimensional environments characterized by numerous interacting variables. Although advancements have been made in the development of causal discovery algorithms, these algorithms frequently depend on assumptions that may not be applicable in real-world scenarios, particularly those involving unobserved confounders or feedback loops.

This issue is particularly pronounced for agentic AI systems, which are required to function in unstructured environments where causal relationships are not explicitly delineated. Future research should prioritize the development of more robust causal discovery methods that can address the complexity and uncertainty inherent in real-world environments.

4.10.2 Computational Efficiency for Real-Time Decisions

Causal reasoning, particularly counterfactual reasoning, is often more computationally demanding than simple correlative methods. This complexity poses challenges for real-time decision making in environments with limited resources, where agents are required to respond swiftly to dynamic conditions.

Further research is required to develop more efficient algorithms for causal inference and identify suitable trade-offs between causal precision and computational tractability. Achieving this balance is crucial for practical deployment of causally aware agentic systems in time-sensitive applications.

4.10.2 Integration with Deep Learning Architectures

Our analysis identified promising hybrid architectures; however, the complete integration of causal reasoning with deep learning remains a complex challenge. Current neural network architectures are not inherently designed to facilitate causal reasoning, and substantial modifications may be required to enable effective causal inferences within these systems.

Future research should investigate innovative architectures that effectively integrate the pattern-recognition capabilities of neural networks with the structured reasoning inherent in causal models. This integration is crucial for the development of agentic systems capable of identifying patterns in complex data and reasoning about the causal mechanisms that generate these patterns.

5. Conclusion

This study investigates the integration of causal inference into agentic AI systems and elucidates how this synthesis enhances explainability and decision-making capabilities. Our analysis indicates that although current agentic AI exhibits significant potential for autonomous operation, its dependence on correlative approaches constrains both transparency and robustness in dynamic environments.

Causal inference presents a promising approach to address these limitations by enabling artificial intelligence systems to reason cause-and-effect relationships rather than merely identifying statistical patterns. This causal understanding facilitates more transparent explanations that align with human cognitive frameworks, thereby fostering trust among stakeholders and enhancing the efficacy of human-AI collaboration. As AI continues to permeate high-stakes domains such as healthcare, finance, and autonomous systems, this transparency is not merely desirable, but essential for responsible deployment.

Our research highlights several promising frameworks for implementing causally aware agentic systems, such as hybrid causal-neural architectures, emergent causal understanding through prediction-driven learning, and counterfactual reasoning systems. Although these frameworks lay a solid foundation for future development, significant challenges persist in areas such as causal discovery in complex environments, computational efficiency for real-time decisions, and integration with deep learning architectures.

Future research should prioritize the empirical validation of the conceptual frameworks introduced in this study, development of more efficient causal discovery algorithms for complex environments, and exploration of innovative architectural approaches that effectively integrate causal reasoning with the pattern recognition capabilities of deep learning. Furthermore, interdisciplinary collaboration among AI researchers, cognitive scientists, ethicists, and domain experts is crucial for the development of governance frameworks that ensure the responsible deployment of increasingly autonomous AI systems.

In conclusion, merging causal inference with agentic AI presents a promising avenue for developing artificial intelligence systems capable of autonomous action while ensuring transparency, explainability, and alignment with human values. By bridging the gap between correlation and causation, these systems have the potential to overcome the limitations of current approaches and fully realize AI's potential of AI as a partner in tackling complex challenges across various domains.

References

[1] L. A. Dennis and M. Fisher, “Verifiable Self-Aware Agent-Based Autonomous Systems,” *Proceedings of the IEEE*, vol.108, no.7, pp.1011–1026, Jul.2020, doi: 10.1109/jproc.2020.2991262.

[2] B. C. Cheong, “Transparency and accountability in AI systems: safeguarding wellbeing in the age of algorithmic decision-making,” *Frontiers in Human Dynamics*, vol.6, Jul.2024, doi: 10.3389/fhumd.2024.1421273.

[3] J. Maclure, “AI, Explainability and Public Reason: The Argument from the Limitations of the Human Mind,” *Minds and Machines*, vol.31, no.3, pp.421–438, Aug.2021, doi: 10.1007/s11023-021-09570-x.

[4] M. Balbaa and M. Abdurashidova, “THE IMPACT OF ARTIFICIAL INTELLIGENCE IN DECISION MAKING: A COMPREHENSIVE REVIEW, ” *EPRA International Journal of Economics, Business and Management Studies*, pp.27–38, Feb.2024, doi: 10.36713/epra15747.

[5] C. Gomez, M. Unberath, and C.-M. Huang, “Mitigating knowledge imbalance in AI-advised decision-making through collaborative user involvement,” *International Journal of Human-Computer Studies*, vol.172, p.102977, Dec.2022, doi: 10.1016/j. ijhcs.2022.102977.

[6] M. Cummings, “Automation Bias in Intelligent Time Critical Decision Support Systems,” Jun.2004. doi: 10.2514/6.2004-6313.

[7] M. K. K, S. Mehta. S, A. H. L, S. N, G. Jadhav, and A. Mitra, “Neuromorphic-Driven Agentic AI for Autonomous Decision-Making Systems,” Dec.2024. doi: 10.1109/icmnwc63764.2024.10872131.

[8] F. Sado, M. Kerzel, S. Wermter, W. S. Liew, and C. K. Loo, “Explainable Goal-driven Agents and Robots-A Comprehensive Review,” *ACM Computing Surveys*, vol.55, no.10, pp.1–41, Feb.2023, doi: 10.1145/3564240.

[9] K. Suzanne Barber, A. Goel, and C. E. Martin, “Dynamic adaptive autonomy in multi-agent systems, ” *Journal of Experimental & Theoretical Artificial Intelligence*, vol.12, no.2, pp.129–147, Apr.2000, doi: 10.1080/095281300409793.

[10] M. K. Mannava, “Causal Inference in AI Based Decision Support: Beyond Correlation to Causation, ” Dec.2024, pp.176–181. doi: 10.1109/icuis64676.2024.10866582.

[11] L. Cox, “Information Structures for Causally Explainable Decisions., ” *Entropy*, vol.23, no.5, p.601, May 2021, doi: 10.3390/e23050601.

[12] M. Antoniadi *et al.*, “Current Challenges and Future Opportunities for XAI in Machine Learning-Based Clinical Decision Support Systems: A Systematic Review, ” *Applied Sciences*, vol.11, no.11, p.5088, May 2021, doi: 10.3390/app11115088.

[13] J. Schoeffer, Y. Machowski, and N. Kuehl, “‘There Is Not Enough Information’: On the Effects of Explanations on Perceptions of Informational Fairness and Trustworthiness in Automated Decision-Making, ” Jun.2022, pp.1616–1628. doi: 10.1145/3531146.3533218.

[14] J. Heuvel, S. Müller, M. Wessels, A. Akhtar, C. Bauckhage, and M. Bennewitz, “Immersive Explainability: Visualizing Robot Navigation Decisions through XAI Semantic Scene Projections in Virtual Reality.” Apr.01, 2025. doi: 10.48550/arxiv.2504.00682.

[15] E. Miehling *et al.*, “Agentic AI Needs a Systems Theory.” Feb.28, 2025. doi: 10.48550/arxiv.2503.00237.

[16] L. Hayduk *et al.*, “Pearl’s D-Separation: One More Step Into Causal Thinking,” *Structural Equation Modeling: A Multidisciplinary Journal*, vol.10, no.2, pp.289–311, Apr.2003, doi: 10.1207/s15328007sem1002_8.

[17] E. Rafetseder, J. Perner, and R. Cristi-Vargas, “Counterfactual Reasoning: Developing a Sense of ‘Nearest Possible World,” *Child Development*, vol.81, no.1, pp.376–389, Jan.2010, doi: 10.1111/j.1467-8624.2009.01401. x.