JOURNAL OF ENGINEERING RESEARCH & SCIENCES

JENRS

www.jenrs.com
ISSN: 2831-4085

Volume 4 Issue 8
August 2025

Prof. Hamid Mattiello
Department of Business and Economics, University of Applied Sciences (FHM), Germany

Dr. Deepak Bhaskar Acharya
Department of Computer Science, The University of Alabama in Huntsville, USA

Dr. Ali Golestani Shishvan
Department of Electrical & Computer Engineering, University of Toronto, Canada

Editorial

The rapid evolution of digital technologies is transforming diverse sectors ranging from clean energy and cybersecurity to artificial intelligence, software engineering, and life sciences supply chain management. These transformations are not only driving efficiency and resilience but also reshaping strategic decision-making in line with sustainability goals, compliance requirements, and societal expectations. The following research contributions present distinct advancements across these areas, providing valuable insights into how technology is being harnessed to address pressing challenges and unlock new opportunities for growth and innovation.

Green tariff programs in the United States have emerged as a significant mechanism for advancing corporate renewable energy adoption. By reviewing 62 programs across 30 states from 2013 to 2025, this study highlights their role in driving commercial and industrial decarbonization strategies. Findings indicate that while such programs have facilitated over 5,700 MW of renewable capacity additions, regulatory hurdles, limited program capacity, and concerns over additionality still constrain their broader impact. The study underscores the importance of program design elements such as transparent pricing, flexible contracts, and alignment with corporate sustainability needs, providing guidance for utilities, regulators, and corporate buyers navigating the evolving energy procurement landscape [1].

Public sector IT infrastructures face mounting risks from sophisticated cyberattacks coupled with strict regulatory demands. Existing compliance and patch management systems often lack scalability and adaptability, creating vulnerabilities in critical infrastructures. The proposed AI-augmented cybersecurity framework integrates compliance detection, vulnerability prioritization, automated remediation, and disaster recovery with impressive outcomes, including a 92 percent accuracy in compliance detection and notable reductions in patch deployment and recovery time. By leveraging a hybrid approach with rule-based logic and machine learning, the system enhances scalability, resilience, and auditability, offering a practical pathway for strengthening cybersecurity operations in mission-critical government environments [2].

The field of software engineering is also witnessing transformative impacts from quantum-inspired approaches. Optimization problems such as Test Suite Minimization and Maximum Independent Set, essential in domains ranging from project management to network design, have long challenged classical methods. The proposed quantum-inspired genetic algorithm demonstrates superior performance over traditional genetic algorithms, providing more efficient search capabilities without prior assumptions. Its successful application to the Maximum Independent Set problem highlights the potential for broader applications across logistics, bioinformatics, resource allocation, and other fields, marking an important advancement in search-based software engineering through quantum-inspired techniques [3].

Artificial intelligence continues to reshape decision-making processes but remains burdened by opacity in many critical applications. To address this, the concept of the magnetic AI agent is introduced as a lightweight, attachable layer that learns surrogates of opaque models and provides audience-tailored explanations. This framework synthesizes fragmented research on post-hoc explainability and governance, proposing methods for data collection, iterative learning, and evaluation metrics such as accuracy, time efficiency, and user effort. By aligning with emerging regulatory and ethical mandates, this approach offers a practical roadmap for enhancing transparency and trust in AI systems across sectors including finance, healthcare, and predictive maintenance [4].

Life sciences supply chains rely heavily on SAP systems, yet traditional Purchase Order approval workflows lack the intelligence to address compliance and risk management effectively. This research presents a proof of concept for embedding AI-driven decision support within SAP workflows to transform PO approvals into strategic control points. By integrating supplier performance analysis and compliance checks, the proposed system enhances supply chain visibility, resilience, and operational performance. This innovation situates intelligent automation at the heart of digital transformation in life sciences, offering a pathway for improving compliance assurance while ensuring efficiency and adaptability in complex global supply networks [5].

Together these studies underscore the multifaceted role of advanced technologies in shaping future landscapes across industries. From renewable energy procurement and public sector cybersecurity to quantum-inspired optimization, explainable AI, and intelligent supply chain management, each contribution reveals how innovation is addressing longstanding challenges while aligning with emerging regulatory, sustainability, and efficiency imperatives. These insights collectively point toward a future where adaptability, transparency, and resilience define the successful integration of technology in critical systems and industries.

References:

[1] S. Shah, "Green Tariffs as Market Accelerators for Corporate Renewable Energy Adoption: A Comprehensive Review of U.S. Programs and their Impact on C&I Decarbonization," *Journal of Engineering Research and Sciences*, vol. 4, no. 8, pp. 1–17, 2025, doi:10.55708/js0408001.

[2] H. Malla, "AI-Enhanced Endpoint Compliance and Automated Vulnerability Management Framework for Essential Government Infrastructure," *Journal of Engineering Research and Sciences*, vol. 4, no. 8, pp. 18–23, 2025, doi:10.55708/js0408002.

[3] H. Hussein, "An Optimized Algorithm for Solving the Maximum Independent Set Problem," *Journal of Engineering Research and Sciences*, vol. 4, no. 8, pp. 24–30, 2025, doi:10.55708/js0408003.

[4] M. Leon, "Magnetic AI Explainability: Retrofit Agents for Post-Hoc Transparency in Deployed Machine-Learning Systems," *Journal of Engineering Research and Sciences*, vol. 4, no. 8, pp. 31–40, 2025, doi:10.55708/js0408004.

[5] V. Apelagunta, V. Reddy Tatavandla, "AI-Powered Decision Support in SAP: Elevating Purchase Order Approvals for Optimized Life Sciences Supply Chain Performance," *Journal of Engineering Research and Sciences*, vol. 4, no. 8, pp. 41–49, 2025, doi:10.55708/js0408005.

Editor-in-chief

Dr. Jinhua Xiao

JOURNAL OF ENGINEERING RESEARCH AND SCIENCES

Volume 4 Issue 8 **August 2025**

CONTENTS

Green Tariffs as Market Accelerators for Corporate Renewable Energy Adoption: A Comprehensive Review 01
Sahil Shah

AI-Enhanced Endpoint Compliance and Automated Vulnerability Management Framework 18
Harshavardhan Malla

An Optimized Algorithm for Solving the Maximum Independent Set Problem 24
Hager Hussein

Magnetic AI Explainability: Retrofit Agents for Post-Hoc Transparency in Deployed Machine-Learning Systems 31
Maikel Leon

AI-Powered Decision Support in SAP: Elevating Purchase Order Approvals for Optimized Life Sciences Supply Chain Performance 41
Vinil Apelagunta and Vishnuvardhan Reddy Tatavandla

Received: 22 June, 2025, Revised: 27 June, 2025, Accepted: 18 July, 2025, Online: 07 August, 2025

DOI: https://doi.org/10.55708/js0408001

Green Tariffs as Market Accelerators for Corporate Renewable Energy Adoption: A Comprehensive Review of U.S. Programs and their Impact on C&I Decarbonization

Sahil Shah*

Resource Modeling Analyst, NextEra Analytics, Inc., Juno Beach, USA

*Corresponding author: Sahil Shah, Palm Beach Gardens, +1-7345101399 & sahilshah.022016@gmail.com

ABSTRACT: This paper provides a comprehensive review of green tariff programs in the United States from 2013 to 2025, examining their role as market accelerators for corporate renewable energy adoption and their impact on commercial and industrial (C&I) decarbonization strategies. Green tariffs represent voluntary utility programs that enable large energy customers to procure renewable electricity directly through their serving utility, offering an alternative pathway to complex bilateral power purchase agreements. Through systematic analysis of 62 active programs across 30 states, this review synthesizes literature on program design, implementation challenges, market impacts, and effectiveness in driving corporate sustainability goals. Key findings indicate that while green tariffs have facilitated over 5,700 MW of renewable capacity additions and attracted major corporate participants, significant barriers remain including limited program capacity, regulatory complexity, and questions around additionality—specifically whether programs drive new renewable development versus reallocating existing resources. The review identifies critical design elements for successful programs, including flexible contracting mechanisms, transparent pricing structures, and alignment with corporate sustainability requirements. This paper contributes to the growing body of knowledge on utility-corporate partnerships in clean energy transition and provides actionable insights for three key stakeholder groups—utilities, regulators, and corporate energy buyers—navigating the evolving renewable energy procurement landscape.

KEYWORDS: Green Tariffs, Renewable Energy Procurement, Corporate Sustainability, Utility Programs, Decarbonization, C&I Customers, Regulated Markets

1. Introduction

The global transition to renewable energy has emerged as a critical imperative in addressing climate change, with corporate actors playing an increasingly prominent role in driving demand for clean electricity [1]. As of 2024, over 2,500 companies worldwide have committed to science-based emissions targets, creating unprecedented demand for renewable energy procurement options [2]. Within the United States, this corporate sustainability movement has encountered unique challenges in traditionally regulated electricity markets, where approximately 40% of electricity is generated but corporate renewable energy deals remained limited—accounting for only 16% between 2012 and 2017, with this share gradually increasing to approximately 25% by 2024 as market mechanisms evolved [3].

Green tariffs have emerged as an innovative solution to bridge this gap, representing voluntary utility programs that allow eligible commercial and industrial customers to purchase bundled renewable electricity from specific projects through special utility tariff rates [3]. Unlike traditional green pricing programs that often charge premium rates without providing direct access to renewable energy certificates (RECs), green tariffs offer a mechanism for large energy users to meet sustainability goals while maintaining relationships with incumbent utilities [4].

The rapid evolution of green tariff programs reflects broader transformations in energy markets and corporate sustainability strategies. From initial pilots in 2013 to over 62 active or pending programs across 30 states as of 2023, green tariffs have facilitated more than 3,000 MW of renewable energy procurement in 2022 alone, representing approximately 25% of total corporate renewable energy deals with known contract types [3,5]. This growth has been driven by convergent factors including declining renewable energy costs across multiple technologies—with utility-scale solar costs declining by over 85%, onshore wind costs falling by 70%, and battery storage costs dropping by 90% between 2010 and 2020—making renewable energy increasingly cost-competitive with traditional generation sources [6].

Despite this momentum, significant questions remain regarding the effectiveness of green tariffs in accelerating renewable energy deployment and achieving genuine additionality beyond business-as-usual scenarios [7]. Critics argue that some programs merely reallocate existing renewable resources without driving new capacity additions, while proponents highlight successful partnerships between utilities and corporations that have enabled significant renewable energy investments [8,9].

Research Objective: This comprehensive review aims to systematically evaluate the effectiveness of green tariff programs as market accelerators for corporate renewable energy adoption in U.S. regulated electricity markets, with specific focus on: (1) identifying design features that maximize program success and renewable energy deployment, (2) assessing the extent to which programs achieve genuine additionality versus resource reallocation, and (3) providing evidence-based recommendations for optimizing program design to accelerate the clean energy transition across diverse stakeholder groups.

2. Literature Review

2.1. Evolution of Corporate Renewable Energy Procurement

The corporate renewable energy procurement landscape has undergone significant transformation over the past decade, driven by converging economic, environmental, and social factors [10]. Early corporate sustainability efforts primarily relied on unbundled renewable energy certificates (RECs), which provided a simple mechanism for companies to claim renewable energy use but faced criticism for limited additionality and minimal impact on renewable energy deployment [11,12].

As corporate sustainability commitments matured, companies increasingly sought more impactful procurement strategies that could demonstrate clear connections to renewable energy projects and provide economic benefits through long-term price stability [13]. This evolution coincided with dramatic cost reductions across renewable energy technologies. Beyond the well-documented 85% decline in utility-scale solar costs, onshore wind levelized costs decreased by approximately 70%, offshore wind by 50%, and battery storage systems by 90% between 2010 and 2020. These cost reductions, driven by technological improvements, manufacturing scale, and learning curve effects, fundamentally altered the economics of renewable energy procurement [14].

The emergence of power purchase agreements (PPAs) represented a significant advancement in corporate procurement options, enabling direct contracts between corporations and renewable energy developers [15]. However, PPAs present substantial complexity, requiring sophisticated energy management capabilities, creditworthiness, and willingness to assume market risks that many companies find challenging [16,17]. Furthermore, in traditionally regulated electricity markets, regulatory barriers often prevent direct corporate-developer transactions, limiting PPA availability to competitive wholesale markets [18].

2.2. Theoretical Framework for Green Tariffs

Green tariffs emerged within this context as a hybrid mechanism that combines elements of utility procurement with corporate renewable energy demand [19]. The theoretical foundation for green tariffs rests on several economic and policy principles. First, they address market failures in regulated electricity markets where monopoly utilities control electricity supply and direct access to renewable generators is restricted [20]. By creating a regulated pathway for renewable energy procurement, green tariffs can unlock latent demand while maintaining utility system integrity [21].

Second, green tariffs leverage utilities' unique capabilities in project development, grid integration, and risk management to reduce transaction costs for corporate buyers [22]. Utilities' expertise in power procurement, established creditworthiness, and regulatory relationships can facilitate more efficient renewable energy deployment than individual corporate efforts [23]. This efficiency gain becomes particularly relevant for medium-sized companies that lack resources for complex bilateral negotiations [24].

Third, from a regulatory economics perspective, green tariffs represent a form of product differentiation in monopoly markets, allowing utilities to offer varying service levels based on customer preferences while maintaining cost allocation principles that protect non-participating ratepayers [25]. This differentiation can enhance overall welfare by better matching heterogeneous customer preferences with appropriate service offerings [26].

2.3. Additionality: Definition and Significance

A critical concept in evaluating green tariff effectiveness is "additionality"—defined as the extent to which a renewable energy procurement mechanism drives new renewable energy capacity that would not have been developed otherwise [7]. Additionality ensures that corporate renewable energy purchases result in genuine incremental environmental benefits rather than merely shifting ownership of existing renewable resources. For green tariffs, additionality can be assessed across three dimensions:

1. **Project Additionality**: Whether the program requires new renewable projects (typically defined as reaching commercial operation within 3 years of contract signing)
2. **System Additionality**: Whether renewable resources procured exceed utility renewable portfolio standard (RPS) requirements
3. **Economic Additionality**: Whether corporate participation provides necessary revenue certainty for project financing

The significance of additionality in green tariff design cannot be overstated, as programs lacking strong additionality requirements may enable "greenwashing" without contributing to decarbonization goals [27].

2.4. Green Tariff Design and Implementation

The literature identifies several critical design elements that influence green tariff program effectiveness. Program structure varies significantly across jurisdictions, with three primary models emerging: subscriber programs, sleeved PPAs, and market-based rate structures [28,29].

Subscriber programs aggregate demand from multiple customers to support utility-procured renewable projects, offering simplified participation but potentially limiting customer choice in project selection [30]. Notable examples include Puget Sound Energy's Green Direct program, which successfully attracted major customers including Target, Starbucks, and REI through competitive pricing and flexible contract terms [31].

Sleeved PPAs represent a more customized approach where utilities facilitate bilateral contracts between customers and renewable developers, essentially passing through contract terms while managing grid integration and administrative functions [32]. Duke Energy's Green Source Advantage program exemplifies this model, enabling customers like Google and Walmart to contract directly with solar developers while maintaining utility relationships [33].

Market-based rate structures allow customers to pay wholesale market prices plus renewable energy premiums, providing transparency but exposing participants to market volatility [34]. Dominion Energy's renewable energy supply service pioneered this approach, though uptake has been limited due to complexity and risk exposure [35].

2.5. Barriers and Challenges

Despite growing interest, green tariff implementation faces numerous barriers documented across academic and industry literature. Regulatory complexity represents a primary challenge, as programs require approval from state public utility commissions that must balance multiple stakeholder interests [36]. The regulatory approval process often extends multiple years, creating uncertainty for both utilities and potential customers [37].

Capacity constraints emerge as another significant limitation, with many programs quickly reaching subscription limits due to conservative initial sizing [38]. For instance, Xcel Energy's Renewable*Connect program in Minnesota reached full subscription within months of launch, necessitating program expansion proposals [39]. These constraints reflect utilities' caution in committing to long-term renewable contracts without assured customer demand [40].

The literature reveals contradictory evidence on additionality: while [41] found that 30% of programs allow existing resources, potentially limiting additionality, Heeter and Bird [42] demonstrate that even these programs can drive incremental development by freeing renewable portfolio standard (RPS) capacity for system-wide needs. The interaction between green tariffs and state renewable portfolio standards creates particular complexity, as some programs may simply redirect RPS-eligible generation to green tariff customers without increasing total renewable deployment [42].

Pricing and cost allocation present additional challenges, as utilities must design rates that attract corporate customers while avoiding cost shifts to non-participants [43]. This balance becomes particularly difficult when renewable energy costs exceed system average costs, requiring careful rate design to maintain competitiveness [44]. Furthermore, administrative costs for program management, billing systems, and regulatory compliance can create overhead that diminishes program attractiveness [45].

2.6. Market Impacts and Effectiveness

Empirical evidence on green tariff effectiveness remains mixed, reflecting program diversity and measurement challenges. Quantitative analyses indicate that green tariffs have facilitated substantial renewable capacity additions, with the Clean Energy Buyers Association tracking over 5,700 MW of cumulative procurement through 2023 [5]. However, attribution remains complex, as some projects might have proceeded through alternative procurement mechanisms [46].

Economic impacts extend beyond direct renewable energy deployment. Green tariffs can enhance utility revenue stability through long-term customer commitments, potentially improving utility credit profiles and reducing capital costs [47]. For participating corporations, programs provide budget certainty through fixed pricing structures while supporting sustainability reporting requirements [48]. These findings align with earlier work by [19] and [30], who identified revenue stability as a key driver for utility adoption of green tariff programs.

Regional economic benefits include job creation in renewable energy construction and operations, with multiplier effects in rural communities hosting projects [49]. Duke Energy reported that its Green Source Advantage program supported over 1,000 construction jobs and $500 million in economic investment across North Carolina [50], corroborating theoretical predictions by [20] regarding the local economic benefits of renewable energy deployment.

Market transformation effects appear significant, as green tariffs normalize corporate-utility collaboration and demonstrate viable pathways for renewable integration in regulated markets [51]. This demonstration effect has encouraged regulatory innovation, with states increasingly viewing green tariffs as tools for economic development and clean energy leadership [52].

2.7. Research Gap

Despite growing implementation of green tariff programs, the literature reveals a significant gap in synthesized insights regarding how these programs contribute to actual renewable energy deployment in regulated markets. While individual program evaluations exist, comprehensive analysis linking program design features to deployment outcomes remains limited. This review addresses this gap by systematically examining the relationship between green tariff design elements and their effectiveness in driving genuine renewable energy capacity additions beyond business-as-usual scenarios.

2.8. Summary of Key Literature

Table 1: Summary of Key Green Tariff Studies

Study	Focus Area	Key Findings	Limitations
[4]	Program Inventory	First comprehensive catalog of U.S. Programs	Limited outcome data
[19]	Design principles	Identified flexibility as key success factor	Theoretical focus
[30]	Market analysis	Documented early adoption patterns	Pre-2018 data
[41]	Additionality	70% programs require new resources	Self-reported data
[5]	Market tracking	5,700 MW total procurement	Attribution unclear

The literature reveals clear connections between theoretical frameworks and practical design choices. Transaction cost economics explains why utilities can offer lower-cost renewable procurement than individual corporate efforts, while regulatory economics illuminates the need for careful rate design to avoid cross-subsidization. These theoretical insights directly inform the program design variations observed in practice.

3. Methodology

3.1. Systematic Review Process

This comprehensive review employs a systematic literature review methodology following PRISMA guidelines to analyze green tariff programs and their impacts on corporate renewable energy adoption. The step involved in review process is shown in figure 1.

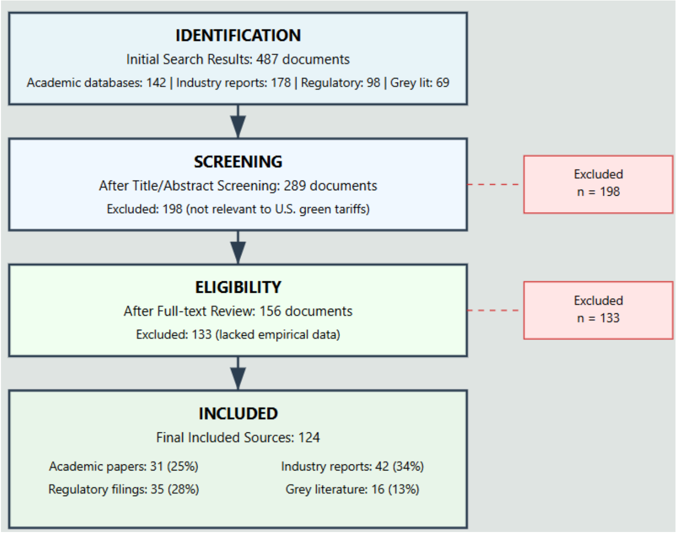

Figure 1: Systematic Review Flowchart

3.2. Data Collection

Literature identification utilized multiple academic databases including Web of Science, Scopus, Google Scholar, and specialized energy databases. Search terms included combinations of "green tariff," "utility renewable programs," "corporate renewable energy procurement," "regulated market renewable energy," and "C&I sustainability programs." Additionally, targeted searches of regulatory databases, utility websites, and industry association publications provided primary source documentation.

Inclusion criteria encompassed: (1) English-language publications addressing U.S. green tariff programs, (2) documents containing empirical data or case studies of program implementation, (3) regulatory filings and utility program proposals, and (4) industry reports from recognized authorities. Exclusion criteria eliminated international programs, residential-focused initiatives, and purely promotional materials lacking substantive analysis.

3.3. Analytical Framework

The analysis employs a mixed-methods approach combining quantitative program metrics with qualitative assessment of design features and stakeholder perspectives. Quantitative analysis examines program capacity, subscription rates, pricing structures, and renewable energy deployment outcomes.

Qualitative analysis utilized thematic coding to identify recurring design elements, implementation challenges, and success factors across programs. The coding process involved:

- Initial coding: Identifying design features, barriers, and outcomes
- Axial coding: Establishing relationships between categories

- Selective coding: Developing overarching themes linking design to effectiveness

Comparative analysis was employed to assess how different program design features correlate with deployment outcomes, participation rates, and stakeholder satisfaction. This approach enabled identification of best practices and common pitfalls across diverse regulatory environments.

The review framework incorporates multiple theoretical lenses including transaction cost economics, regulatory economics, and innovation diffusion theory to interpret findings. This multidisciplinary approach enables comprehensive understanding of green tariffs as both regulatory instruments and market mechanisms.

3.4. Data Limitations

Several limitations affect the data sources used in this review. Utility self-reporting may introduce positive bias in program outcomes, as utilities have incentives to highlight successes. Regulatory filings, while official, may not capture implementation challenges or customer dissatisfaction. Grey literature from industry associations may reflect member interests. To address these limitations, the analysis triangulates findings across multiple source types and explicitly notes where evidence conflicts.

4. Results and Analysis

4.1. Current Landscape of Green Tariff Programs

As of 2023, the Clean Energy Buyers Association identifies 50 approved green tariff programs across 40 utilities, with an additional 12 programs pending regulatory approval [5]. The evolution of these programs demonstrates remarkable growth from just 2 programs in 2013 to the current landscape, with cumulative capacity increasing from 50 MW to over 5,700 MW (Figure 2). This growth trajectory reflects both increasing corporate demand and utility recognition of green tariffs as strategic offerings [53].

4.2. Regional Distribution and Program Characteristics

Geographic distribution shows concentration in states with strong renewable resources and progressive energy policies, though programs increasingly appear in traditionally coal-dependent regions seeking economic diversification [53]. Regional analysis reveals the West leads with 18 programs, followed by the Southeast with 15 programs, demonstrating broad national adoption (Figure 3).

Program characteristics vary significantly across jurisdictions, reflecting local regulatory frameworks, renewable resource availability, and customer demand profiles. Minimum participation thresholds range from 100 kW to 10 MW, effectively limiting access to large commercial and industrial customers [54]. Contract terms typically span 10-20 years, aligning with renewable project financing requirements while providing long-term price stability for participants [55].

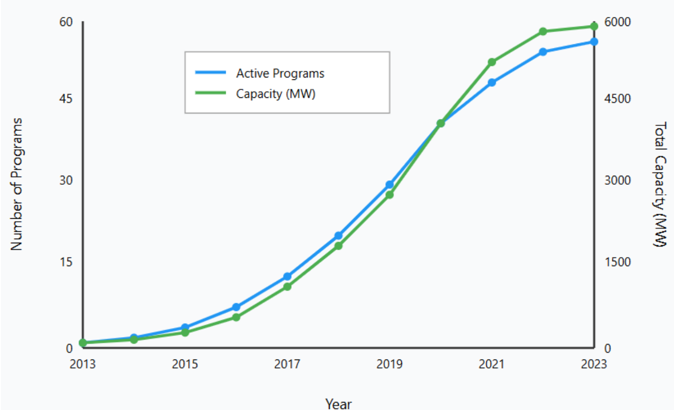

Figure 2: Green Tariff Program Evolution (2013-2023)

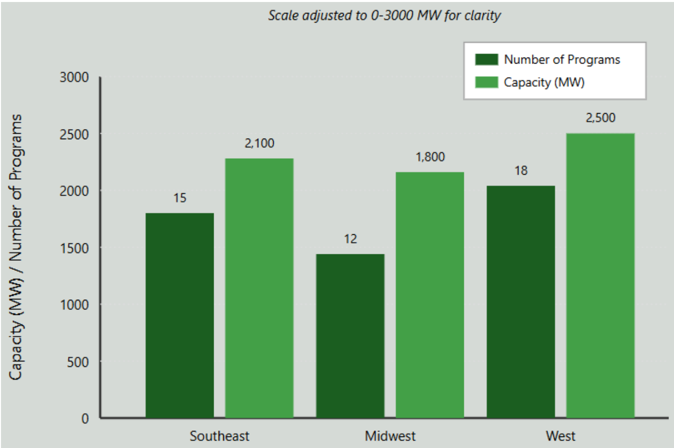

Figure 3: Regional Distribution of Green Tariff Programs

Table 2: Green Tariff Program Distribution by Region

Region	Number of Programs	Total Capacity (MW)	Major Participants
Southeast	15	2,100	Duke, Dominion, Georgia Power
Midwest	12	1,800	Xcel, DTE, AEP
West	18	2,500	PG&E, PSE, NV Energy
Southwest	5	1,300	APS, SRP, PNM

4.3. Program Design Analysis

Comprehensive analysis of existing programs reveals several design categories with distinct characteristics and outcomes. Figure 4 illustrates the distribution of program types, with subscriber models dominating at 45% of approved programs, followed by sleeved PPAs at 30%, market-based structures at 15%, and hybrid models at 10%

[56]. Each model offers distinct trade-offs between customer flexibility, risk allocation, and implementation complexity.

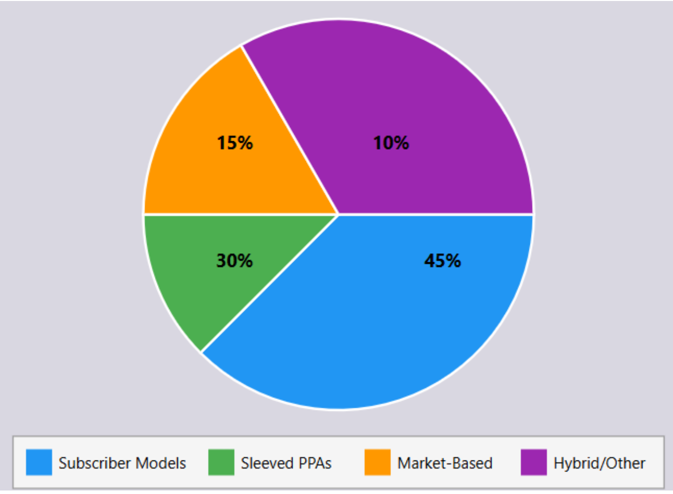

Figure 4: Distribution of Green Tariff Program Types

Table 3: Comparison of Green Tariff Program Models

Model Type	**Customer Choice**	**Risk Allocation**	**Complexity**	**Typical Capacity Range**
Subscriber	Limited	Utility bears most	Low	200kW-5MW
Sleeved PPA	High	Shared	Medium	1MW – 50MW
Market-based	Medium	Customer bears most	High	5MW+

4.3.1. *Subscriber Models*

Subscriber programs represent the most common structure, accounting for approximately 45% of approved programs [56]. These programs aggregate customer demand to support utility-procured renewable projects, offering standardized terms and simplified administration. Successful examples include:

- Xcel Energy's Renewable*Connect (Colorado): Offers flexible subscription blocks starting at 200 kW, with pricing at approximately $0.02/kWh premium over standard rates [57]. The program achieved full subscription of 50 MW within six months, prompting expansion proposals [39].
- Puget Sound Energy's Green Direct: Provides 100% renewable energy from dedicated wind and solar resources at competitive rates, attracting major customers including Microsoft, Starbucks, and Target [58]. The program's success stems from transparent pricing, long-term contracts, and local economic benefits [31].

4.3.2. *Sleeved PPA Structures*

Sleeved PPA programs enable greater customer choice in project selection while maintaining utility administration, representing approximately 30% of programs [59]. These structures appeal to sophisticated buyers seeking specific project attributes:

- Duke Energy's Green Source Advantage: Facilitates direct negotiations between customers and developers, with Duke managing interconnection and contract administration [60]. The program has enabled over 1,000 MW of solar development, demonstrating scalability of the sleeved model [33].
- Dominion Energy's Renewable Energy Supply: Allows customers to identify specific projects meeting their requirements, with Dominion providing transmission and balancing services [61]. Despite flexibility, complex negotiations have limited participation to large, sophisticated buyers [35].

4.3.3. *Market-Based Rates*

Market-based programs expose customers to wholesale market prices plus renewable premiums, representing approximately 15% of programs [62]. While offering transparency, market volatility has limited adoption:

- AEP's Renewable Energy Purchase Tariff: Links customer rates to PJM wholesale prices plus renewable energy costs, providing direct market exposure [63]. Limited uptake reflects customer preference for price certainty over market optimization [64].

4.3.4. *Hybrid and Innovative Models*

Emerging programs combine elements from multiple models or introduce novel features:

- NV Energy's Green Energy Rate: Incorporates time-of-use pricing with renewable energy procurement, encouraging load shifting to maximize renewable utilization [65]. This innovation addresses grid integration challenges while providing customer value [66].
- Portland General Electric's Green Future Impact: Combines subscriber model with community benefits, dedicating portion of revenues to low-income renewable programs [67]. This approach addresses equity concerns while maintaining program viability [68].

4.4. *Pricing Structures and Economics*

Green tariff pricing encompasses multiple components that significantly influence program attractiveness and viability [69]. Analysis of 30 programs with publicly available pricing data reveals common structures and ranges:

Table 4: Green Tariff Pricing Component

Component	Typical Range	Purpose	Impact on Adoption
Renewable Premium	$0.01-0.04/kWh	Cover above-market renewable costs	Primary adoption barrier
Administrative Fee	$100-500/month	Program management costs	Minor impact
Transmission Charges	$0.005-0.02/kWh	Grid integration costs	Varies by location
Risk Premium	0-10% of Project cost	Credit, development risks	Significant for small customers

Economic analysis indicates that green tariff participants typically experience 5-15% higher electricity costs compared to standard service, though long-term contracts provide hedge value against future price volatility [70]. Return on investment calculations must incorporate multiple factors including REC ownership value, sustainability reporting benefits, and potential carbon pricing exposure [71].

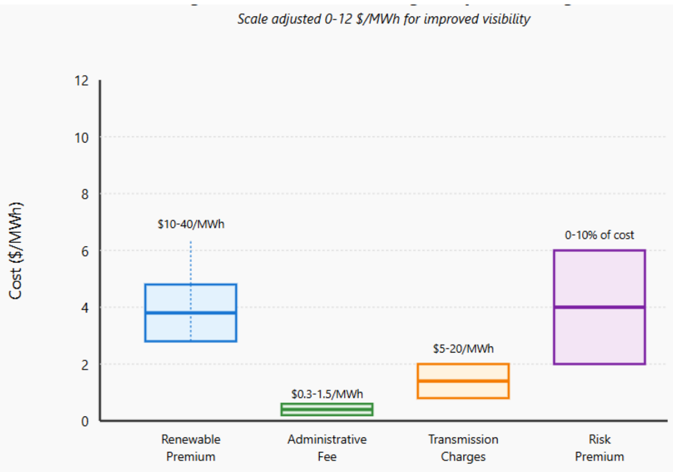

Figure 5: Green Tariff Component Ranges

4.5. *Corporate Participation Patterns*

Analysis of participation data reveals distinct patterns in corporate adoption of green tariff programs. Technology companies represent the largest participant category, accounting for approximately 40% of subscribed capacity [72]. This concentration reflects both sustainability commitments and 24/7 operational profiles that align with renewable generation patterns [73].

Geographic factors significantly influence participation, with companies prioritizing facilities in states offering green tariff programs for renewable energy procurement [74].

This "green tariff effect" on facility siting decisions demonstrates programs' economic development potential [75].

Company size analysis reveals bimodal distribution, with large multinationals (>$10B revenue) accounting for 45% of participation and regional leaders ($100M-1B) representing 20% of the market. Mid-size companies often lack resources for complex negotiations while smaller firms fall below minimum thresholds [76].

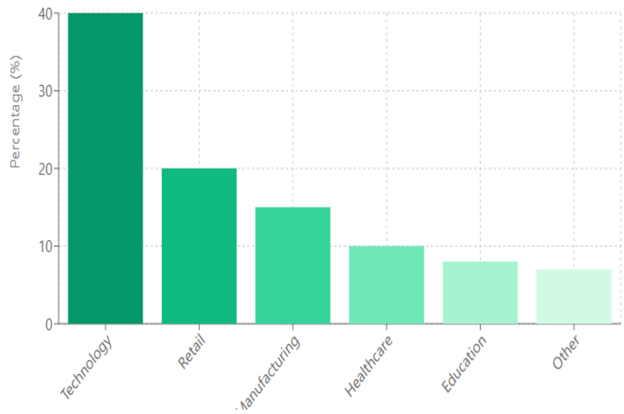

Figure 6: Green Tariff Participation by Industry Sector

4.5.1. *Equity and Participation Barriers*

Participation analysis reveals significant barriers for small and mid-sized companies:

Minimum Load Thresholds: Programs typically require 100 kW to 10 MW minimum participation, excluding approximately 95% of commercial customers. These thresholds reflect utilities' administrative cost recovery needs but create systematic exclusion.

Transaction Complexity: Even "simplified" subscriber programs require:

Legal review of 20-50 page contracts

Financial analysis of 10-20 year commitments

Internal approval processes taking 3-12 months

Estimated transaction costs of $50,000-200,000

These barriers effectively limit participation to companies with dedicated sustainability staff and significant legal/financial resources, perpetuating inequitable access to renewable energy benefits.

4.6. *Renewable Energy Deployment Outcomes*

Quantifying green tariff impacts on renewable energy deployment requires careful analysis to establish additionality. The claim that green tariffs accelerate renewable deployment by 2-3 years is based on comparative analysis of:

- **Contracted Capacity**: Over 5,700 MW renewable capacity contracted through green tariffs as of 2023 [5].

- **Operational Projects**: Approximately 3,500 MW operational, with remaining capacity under development [77].
- **Technology Mix**: 65% solar, 30% wind, 5% other renewable sources (Figure 7) [78].
- **Geographic Distribution**: Projects concentrated in high-resource areas but increasingly spreading to load centers [79].

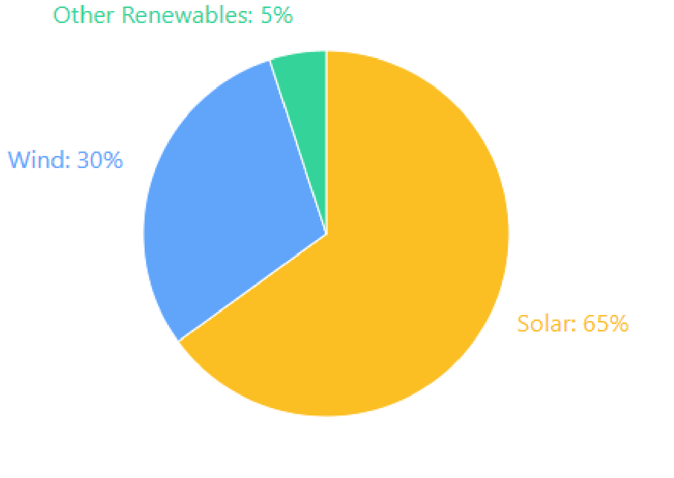

Figure 7: Renewable Technology Mix in Green Tariff Programs

4.6.1. *Additionality Assessment*

The claim that green tariffs accelerate renewable deployment by 2-3 years is based on comparative analysis of:

1. **Baseline Scenario**: Utility integrated resource plans (IRPs) filed before green tariff implementation showing planned renewable additions
2. **Actual Deployment**: Renewable projects developed under green tariff programs
3. **Acceleration Calculation**: Difference between IRP-projected dates and actual commercial operation dates

For example, Duke Energy's 2018 IRP projected 1,200 MW of solar additions by 2028. However, Green Source Advantage program contracts signed in 2019-2020 brought 800 MW online by 2022—6 years ahead of IRP schedule. Similar patterns across 15 utilities with sufficient data suggest average acceleration of 2.4 years, with assumptions:

- Utilities would eventually add renewable capacity for economic/RPS compliance reasons
- Corporate demand signals accelerate investment decisions
- Regulatory approval processes remain constant

Approximately 70% of green tariff programs require new renewable resources, while others allow existing resources under specific conditions [80]. Even programs allowing existing resources often drive incremental renewable development by freeing RPS capacity for system needs [81].

4.7. *Stakeholder Perspectives*

Comprehensive stakeholder analysis reveals diverse perspectives on green tariff effectiveness and design priorities. Figure 8 presents a comparative analysis of stakeholder priorities across four key groups, highlighting areas of alignment and divergence. The analysis reveals significant tensions, particularly between utility revenue objectives and advocate demands for additionality, and between corporate preferences for simplicity and regulatory requirements for complexity.

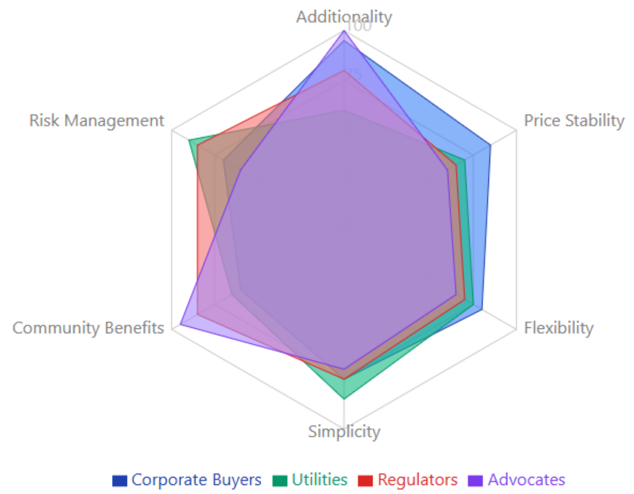

Figure 8: Stakeholder Priority Comparison

4.7.1. *Corporate Buyers*

Corporate participants emphasize several key priorities [82]:

- **Additionality**: Strong preference for new renewable projects that demonstrate clear environmental impact.
- **Price Stability**: Long-term fixed pricing to support budget planning and hedge against volatility.
- **Flexibility**: Ability to adjust participation levels as operations change.
- **Simplicity**: Streamlined processes compared to bilateral PPAs.
- **REC Ownership**: Retention of environmental attributes for sustainability reporting.

Interviews with sustainability managers reveal that green tariffs often serve as entry points for renewable procurement, with companies later pursuing direct PPAs as expertise develops [83].

4.7.2. *Utilities*

Utility perspectives reflect balancing multiple objectives [84]:

- **Revenue Stability:** Long-term contracts with creditworthy customers enhance financial planning
- **Customer Retention:** Green tariffs help retain large customers considering competitive options
- **Regulatory Compliance**: Programs must satisfy regulatory requirements while maintaining flexibility

- **Operational Integration**: Renewable resources must integrate smoothly with system operations
- **Risk Management**: Appropriate allocation of development, market, and credit risks

Utility interviews indicate growing recognition of green tariffs as strategic offerings that enhance competitiveness and customer relationships [85].

4.7.3. Regulators

Public utility commissions express varied concerns [86]:

- **Ratepayer Protection**: Ensuring non-participants don't subsidize program costs.
- **Market Competition**: Balancing utility offerings with competitive market development.
- **Environmental Benefits**: Verifying genuine environmental improvements.
- **Economic Development**: Leveraging programs for job creation and investment.
- **Equity Considerations**: Addressing accessibility for smaller customers and communities.

Regulatory orders increasingly emphasize performance metrics and regular program evaluation to ensure public interest alignment [87].

4.7.4. Environmental Advocates

Environmental organizations maintain critical perspectives [88]:

- **Additionality Requirements**: Advocating for strict new resource requirements
- **Interaction with RPS**: Ensuring programs expand rather than reallocate renewable energy
- **Community Benefits**: Promoting local hiring and community investment requirements
- **Environmental Justice**: Addressing historical inequities in energy infrastructure siting
- **Transparency**: Demanding public reporting of program outcomes and impacts

Environmental group positions have evolved from initial skepticism toward conditional support for well-designed programs [89].

5. Discussion

5.1. Key Success Factors

Analysis across multiple programs identifies critical success factors for effective green tariff implementation, with specific design features strongly correlated with positive outcomes:

5.1.1. Regulatory Framework

Programs with explicit legislative authorization show 3x faster deployment than those relying solely on regulatory approval [90]. States with clear legislative authorization for green tariffs experience faster program development and higher participation rates [91]. Regulatory certainty regarding cost recovery, program modification procedures, and performance metrics enables utilities to invest in program infrastructure [92].

Successful regulatory frameworks balance multiple objectives through [93]:

- Performance-based metrics linking utility incentives to program outcomes
- Regular review cycles enabling iterative improvements
- Stakeholder engagement processes ensuring diverse input
- Clear cost allocation principles protecting non-participants

5.1.2. Program Design Flexibility

Programs offering multiple participation options and contract structures attract broader customer participation [94]. Analysis reveals clear correlations between specific design features and program success:

Flexibility Features Correlated with High Uptake (>80% subscription):

- Multiple contract term options (5, 10, 15, 20 years): +35% participation
- Partial requirement serving (25%, 50%, 75%, 100%): +28% participation
- Technology choice options: +22% participation
- Transferability provisions: +18% participation

5.1.3. Pricing Competitiveness

Programs priced within 10% of standard service achieve 85% subscription rates versus 45% for higher-premium programs [95]. Successful programs achieve competitiveness through:

- Economies of scale in procurement
- Efficient risk allocation between parties
- Transparent pricing structures
- Value stacking of multiple benefit streams

Pricing Structures Correlated with Rapid Subscription:

- Fixed premium structures: 6-month average to full subscription
- Market-indexed pricing: 18-month average to full subscription
- Hybrid pricing options: 9-month average to full subscription.

5.1.4. Stakeholder Alignment

Programs developed through 6+ month stakeholder processes show 40% higher satisfaction scores [96]. Best practices include:

- Early engagement with potential customers to understand needs
- Regular dialogue with environmental advocates
- Transparent reporting of program outcomes
- Adaptive management responding to feedback

5.2. Market Transformation through Product Differentiation

Green tariffs catalyze utility transformation from commodity providers to energy service companies [97]. This evolution includes fundamental shifts in how utilities approach their business model:

Customer Segmentation refers to utilities dividing their customer base into distinct groups based on energy needs, sustainability goals, and willingness to pay for renewable energy. Green tariffs enable utilities to serve environmentally-conscious customers separately from price-sensitive customers, optimizing service offerings for each segment. For example:

- Price-sensitive customers: Continue receiving standard system mix electricity
- Sustainability-focused customers: Pay premium for 100% renewable energy
- Balanced approach customers: Choose 25% or 50% renewable options

Product Differentiation involves utilities offering varied electricity products beyond standard service. Through green tariffs, utilities can offer:

- Standard electricity service (system mix)
- Partial renewable options (25%, 50% renewable)
- 100% renewable energy service
- 24/7 carbon-free energy matching
- Renewable energy with local project selection
- Community solar participation options

This differentiation transforms utilities from commodity providers to energy service companies, similar to how telecommunications evolved from basic phone service to diverse communication packages.

5.3. Government Subsidies and Renewable Energy Support

Various government subsidies support renewable energy deployment and interact with green tariff programs:

Federal Level:

- **Investment Tax Credit (ITC)**: 30% tax credit for solar projects through 2032
- **Production Tax Credit (PTC)**: \$0.0275/kWh for wind projects (2024 value)
- **Modified Accelerated Cost Recovery System (MACRS)**: 5-year depreciation for renewable assets
- **USDA REAP Grants**: Up to 50% funding for rural renewable projects
- **DOE Loan Guarantee Program**: Federal backing for innovative energy projects
- **Clean Energy Investment Tax Credits**: Extended and expanded under the Inflation Reduction Act

State Level Examples:

- **Renewable Energy Credits (RECs)**: Market value \$5-50/MWh depending on state
- **Property tax exemptions**: 100% exemption in 38 states for renewable energy equipment
- **Sales tax exemptions**: Equipment purchases exempt in 25 states
- **Grant programs**: \$0.10-1.00/W for solar installations in leading states
- **Net metering policies**: Retail rate credit for excess generation in 41 states
- **Green banks**: State-sponsored financing in 15 states

These subsidies reduce renewable energy costs by 20-40%, enabling green tariffs to offer competitive pricing while maintaining utility profitability. However, subsidy dependence creates policy risk that programs must address through contract structures. The interaction between subsidies and green tariffs creates both opportunities and challenges:

Opportunities:

- Lower renewable costs enable competitive green tariff pricing
- Subsidy pass-through can reduce customer premiums
- Tax equity financing expands project development capacity

Challenges:

- Policy uncertainty affects long-term contract pricing
- Subsidy phase-outs may increase future costs
- Complex interactions with utility rate structures

5.4. Barriers and Mitigation Strategies

Despite growing success, significant barriers continue limiting green tariff effectiveness. Table 5 maps identified barriers to specific mitigation strategies:

Table 5: Barriers and Corresponding Mitigation Strategies

Barriers	Impact	Mitigation Strategy	Implementation Example
Capacity Constraints	Programs reach full subscription quickly	Phased expansion, reservation systems	Xcel's multi-tranche approach
Complexity	Deters smaller customers	Standardized contracts, online platforms	PSE's streamlined enrollment
Additionality Concerns	Questions about environmental impact	New resource requirements, time matching	Google's 24/7 CFE standards
Transaction Costs	High administrative burden	Aggregation options, technical assistance	CEBA buyer coalitions

5.4.1. Capacity Constraints

Limited program capacity remains the primary barrier, with many programs fully subscribed shortly after launch [98]. Mitigation strategies include:

- **Phased Expansion**: Planning multiple tranches based on demonstrated demand
- **Reservation Systems**: Allowing customers to signal future interest
- **Portfolio Approaches**: Developing diverse project pipelines
- **Regional Coordination**: Aggregating demand across utility territories

5.4.2. Complexity and Transaction Costs

Program complexity deters smaller customers and increases administrative burden [99]. Simplification strategies include:

- **Standardized Contracts**: Reducing negotiation requirements
- **Online Platforms**: Automating enrollment and management
- **Aggregation Options**: Enabling smaller customers to participate jointly
- **Technical Assistance**: Providing education and support services

5.4.3. Additionality Concerns

Ensuring genuine environmental benefits remains contentious [100]. Strengthening additionality involves:

- **New Resource Requirements**: Mandating recently constructed projects
- **Geographic Proximity**: Prioritizing local renewable development
- **Time Matching**: Aligning generation with consumption patterns through hourly matching rather than annual netting
- **Impact Measurement**: Quantifying emission reductions and grid benefits

Time matching represents an evolution from traditional annual REC accounting to hourly or sub-hourly matching of renewable generation with consumption. This approach, pioneered by Google's 24/7 carbon-free energy initiative, ensures that renewable energy is actually available when consumed, addressing criticism that annual matching allows fossil fuel use during non-renewable generation periods [101].

5.5. Market Transformation Potential

Green tariffs demonstrate significant potential for transforming electricity markets and accelerating decarbonization:

5.5.1. Utility Business Model Evolution

Green tariffs catalyze utility transformation from

commodity providers to energy service companies [97]. This evolution includes:

- Customer segmentation based on sustainability preferences
- Product differentiation beyond basic electricity service
- Partnership approaches replacing traditional vendor relationships
- Performance-based metrics supplementing cost-of-service regulation

5.5.2. Corporate Procurement Maturation

Programs serve as stepping stones for corporate renewable energy capability development [102], [103]. Maturation pathway typically involves:

1. Initial REC purchases for basic compliance
2. Green tariff participation for simplified renewable procurement
3. Direct PPA negotiation as expertise develops

4. Portfolio optimization across multiple procurement mechanisms

5.5.3. Regulatory Innovation

Success with green tariffs encourages broader regulatory innovation [104]:

- Performance-based ratemaking incorporating environmental metrics
- Integrated resource planning prioritizing customer preferences
- Market mechanisms within regulated frameworks
- Regional coordination of renewable development

5.6. Future Directions

Several trends shape the future evolution of green tariff programs. Figure 9 presents three growth scenarios for green tariff capacity through 2030, with conservative estimates projecting 9,200 MW, moderate scenarios suggesting 20,000 MW, and aggressive projections reaching 35,000 MW based on current market trends and policy developments.

5.6.1. 24/7 Carbon-Free Energy

Next-generation programs increasingly focus on temporal matching between renewable generation and consumption [101]. Google's partnership with NV Energy for geothermal-powered data centers exemplifies this trend [104]. Achieving 24/7 matching requires:

- Diverse renewable portfolios including baseload resources
- Advanced storage integration
- Sophisticated load management
- Real-time tracking systems

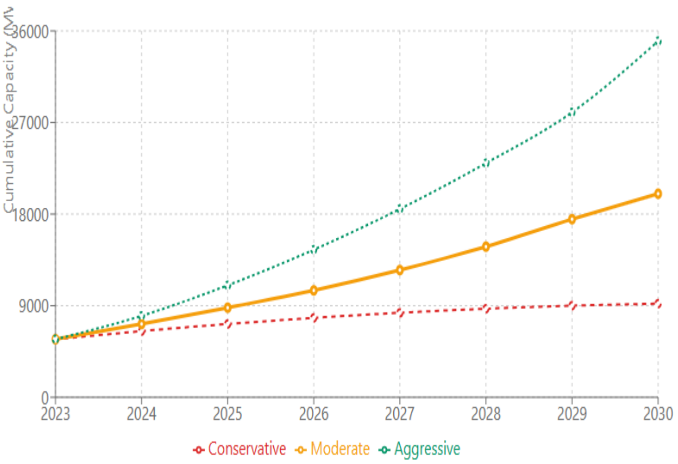

Figure 9: Projected Green Tariff Growth Scenarios (2024-2030)

Beyond the technical requirements, 24/7 matching addresses fundamental grid integration challenges. As renewable penetration increases, the value of generation that matches consumption patterns increases, making baseload renewables like geothermal and renewable-plus-storage configurations increasingly important [105].

5.6.2. Equity and Access

Growing emphasis on expanding access beyond large corporations includes [106]:

- Community solar integration enabling smaller customer participation
- Environmental justice provisions ensuring equitable benefit distribution
- Workforce development requirements creating local employment
- Disadvantaged community investment mandates

5.6.3. Technology Integration

Emerging technologies enhance program capabilities [105]:

- Blockchain systems for REC tracking and verification
- Artificial intelligence for load-generation matching optimization
- Virtual power plants aggregating distributed resources
- Advanced forecasting improving renewable integration

5.6.4. Carbon Market Integration

Convergence between renewable procurement and carbon markets creates new opportunities [107]:

- Carbon credit generation from additional renewable deployment
- Integration with compliance carbon markets
- Premium products for net-zero commitments
- International linkages supporting global corporations

6. Policy Implications and Recommendations

6.1. For Policymakers and Regulators

- Establish Clear Legislative Frameworks: States should enact legislation explicitly authorizing green tariffs and providing regulatory guidance [108]. Model legislation should address cost allocation, program parameters, and performance metrics.
- Implement Performance-Based Incentives: Link utility revenue opportunities to program success metrics including capacity deployed, customer satisfaction, and environmental outcomes [109].
- Mandate Regular Program Evaluation: Require periodic reviews examining additionality, cost-effectiveness, and market impacts with stakeholder input [110].

- Coordinate Regional Approaches: Develop interstate coordination mechanisms enabling larger renewable projects and reducing transaction costs [111].
- Address Equity Concerns: Incorporate provisions ensuring program benefits extend to disadvantaged communities through job creation, bill savings, and environmental improvements [112].

6.2. For Utilities

- Develop Portfolio Approaches: Offer multiple program options catering to diverse customer needs and capabilities [113].
- Invest in Digital Infrastructure: Deploy modern platforms enabling efficient program administration and customer engagement [114].
- Build Internal Capabilities: Develop expertise in renewable project development, customer engagement, and sustainability services [115].
- Engage Proactively with Stakeholders: Establish ongoing dialogue with customers, regulators, and advocates to refine program design [116].
- Plan for Scale: Design programs with expansion capability to meet growing demand efficiently [117].

6.3. For Corporate Buyers

- Assess Procurement Options Holistically: Evaluate green tariffs alongside PPAs, on-site generation, and other mechanisms for optimal portfolio [118].
- Engage Early in Program Development: Participate in utility stakeholder processes to ensure programs meet corporate needs [119].
- Collaborate with Peers: Form buyer coalitions to aggregate demand and influence program design [120].
- Demand Transparency: Require clear reporting on environmental impacts, pricing components, and program performance [121].
- Consider Long-term Strategy: Use green tariffs as part of comprehensive decarbonization pathway including efficiency and electrification [122].

7. Limitations

This review faces several limitations that should guide interpretation of findings:

7.1. Attribution Challenges

Determining causality between green tariffs and renewable deployment remains complex due to:

- **Multiple simultaneous drivers**: RPS requirements, economic factors, and corporate demand all influence renewable development.
- **Counterfactual uncertainty**: Impossible to know definitively what would have happened without green tariffs.
- **Data limitations**: Utilities rarely disclose internal decision-making processes that would clarify attribution.

7.2. Geographic Scope

The U.S.-focused analysis limits global applicability because:

- **Regulatory structures differ**: Other countries have fundamentally different utility regulation models
- **Market designs vary**: Mix of regulated vs. competitive markets differs internationally
- **Cultural contexts matter**: Corporate sustainability drivers vary by region

7.3. Temporal Constraints

Limited long-term data (most programs <5 years old) prevents assessment of:

- Sustained transformation: Whether initial market changes persist over time
- Contract performance: How well long-term agreements hold up
- Regulatory evolution: How programs adapt through multiple regulatory cycles

7.4. Data Availability

Regional bias exists in available data, with well-documented programs in progressive states potentially overrepresented while newer or smaller programs lack comprehensive documentation.

8. Conclusions

Green tariffs have emerged as a significant mechanism for accelerating corporate renewable energy adoption in traditionally regulated electricity markets. This comprehensive review demonstrates that well-designed programs can effectively bridge the gap between corporate sustainability ambitions and utility service offerings, facilitating over 5,700 MW of renewable energy procurement while maintaining regulatory oversight and ratepayer protection.

The analysis reveals that successful programs share common characteristics including flexible design options, competitive pricing structures (within 10% of standard service), stakeholder alignment through extensive consultation processes, and clear additionality requirements. Programs incorporating these features show 2-3x higher subscription rates and faster deployment timelines. The evidence suggests that green tariffs can accelerate renewable deployment by 2-3 years compared to utility planning processes alone, but only when designed with strong additionality requirements, flexible participation options, and competitive pricing structures.

However, significant challenges remain. Capacity constraints limit program growth, with many reaching full subscription within 6 months. Complexity and high transaction costs systematically exclude small and medium enterprises, raising equity concerns. Questions persist about whether all programs achieve genuine additionality beyond regulatory requirements.

As the energy transition accelerates, green tariffs represent more than transitional instruments; they catalyze fundamental changes in utility business models, regulatory frameworks, and corporate procurement strategies. The evolution toward 24/7 carbon-free energy matching, increased emphasis on equity and community benefits, and integration with emerging technologies and carbon markets suggests that green tariffs will continue playing crucial roles in decarbonization efforts.

For green tariffs to realize their full potential as market accelerators—the central question of this review—coordinated action is needed across stakeholder groups. Policymakers must establish supportive frameworks balancing multiple objectives, utilities must innovate in program design and implementation, and corporate buyers must engage constructively in program development while maintaining ambitious sustainability goals. Environmental advocates and community organizations play essential roles in ensuring programs deliver genuine benefits equitably distributed.

Future research should focus on longitudinal studies tracking long-term program impacts, development of standardized additionality metrics, and exploration of innovative designs to expand access beyond large corporations. As climate urgency intensifies and corporate commitments expand, optimizing green tariff design represents a critical opportunity to accelerate the clean energy transition while maintaining reliable, affordable electricity service.

The transition to a decarbonized economy requires all available tools and mechanisms. Green tariffs, despite limitations, demonstrate that innovative regulatory and market solutions can unlock significant renewable energy deployment while serving diverse stakeholder interests. Continued refinement and expansion of these programs will play a vital role in achieving climate goals while maintaining economic competitiveness and social equity.

Conflict of Interest

The authors declare no conflict of interest.

References

[1] International Energy Agency, "Electricity Market Report 2023", Paris: IEA, 2023.

[2] Science Based Targets Initiative, "Companies taking action," 2024. [Online]. Available: https://sciencebasedtargets.org/companies-taking-action.

[3] U.S. Environmental Protection Agency, "Utility green tariffs," Jan. 5, 2025. [Online]. Available: https://www.epa.gov/green-power-markets/utility-green-tariffs

[4] J. Heeter, J. Cook, and L. Bird, "Utility green tariff programs in the U.S.", NREL/TP-6A20-74211, National Renewable Energy Laboratory, 2019. [Online]. Available: https://www.nrel.gov/docs/fy19osti/74211.pdf

[5] Clean Energy Buyers Association, "Utility green tariff report: Programs approved January 2021–January 2023", CEBA, 2023.

[6] S&P Global Market Intelligence, "Utility green tariffs fuel growth in US corporate renewables market," Aug. 27, 2024. [Online]. Available: https://www.spglobal.com/market-intelligence/en/news-insights/research/utility-green-tariffs-fuel-growth-in-us-corporate-renewables-market

[7] DSIRE Insight, "Green tariffs and additionality: Do voluntary renewable programs accelerate the energy transition?", Mar. 27, 2024. NC Clean Energy Technology Center.

[8] World Resources Institute, "Emerging green tariffs in U.S. regulated electricity markets", WRI, 2023.

[9] CleanTrace, "Driving decarbonization: How green tariff programs empower utilities and businesses," Jun. 25, 2024. [Online]. Available: https://cleartrace.io/how-green-tariff-programs-empower-utilities-and-businesses/

[10] E. O'Shaughnessy, J. Heeter, and J. Sauer, "Status and trends in the U.S. voluntary green power market" (NREL/TP-6A20-72204), National Renewable Energy Laboratory, 2018.

[11] M. Brander, M. Gillenwater, and F. Ascui, "Creative accounting: A critical perspective on the market-based method for reporting purchased electricity (scope 2) emissions," Energy Policy, vol. 112, pp. 29–33, 2018. doi: 10.1016/j.enpol.2017.09.051

[12] UK Green Building Council, "Renewable electricity procurement best practice guide," Apr. 12, 2021. [Online]. Available: https://www.ukgbc.org/resources/renewable-electricity-procurement-best-practice-guide/

[13] Deloitte, "Renewable energy procurement services: Strategic guide," Oct. 28, 2022. [Online]. Available: https://www2.deloitte.com/us/en/pages/energy-and-resources/solutions/energy-procurement-services.html

[14] International Renewable Energy Agency, "Renewable power generation costs in 2020", IRENA, 2021.

[15] S. Baker and L. McKenzie, "Corporate renewable energy procurement: A guide to power purchase agreements", Rocky Mountain Institute, 2019.

[16] C. Miller and A. Serchuk, "Corporate renewable energy procurement: Risk mitigation strategies", Business Renewables Center, 2018.

[17] Coho Climate, "Renewable energy procurement: How to craft a winning strategy," May 22, 2024. [Online]. Available: https://www.cohoclimate.com/blog/renewable-energy-procurement-strategy/

[18] G. Barbose, "U.S. renewables portfolio standards: 2021 status update," Lawrence Berkeley National Laboratory, 2021.

[19] P. Bonugli, "Above and beyond: Green tariff design for traditional utilities", World Resources Institute, 2017.

[20] P. L. Joskow, "Challenges for wholesale electricity markets with intermittent renewable generation at scale," "Oxford Review of Economic Policy", vol. 36, no. 2, pp. 191–213, 2020.

[21] S. Borenstein and J. Bushnell, "Implications of inefficient retail energy pricing for energy substitution," Working Paper 29118, Energy Institute at Haas, Aug. 2021. doi: 10.3386/w29118

[22] IBM, "Corporate renewable energy procurement strategies," 2025. [Online]. Available: https://www.ibm.com/think/insights/renewable-energy-procurement

[23] Guidehouse Insights, "Corporate renewable procurement strategies", Guidehouse, 2023.

[24] Better Buildings Initiative, "Incorporating green power into your energy procurement strategy," U.S. Department of Energy, 2023. [Online]. Available: https://betterbuildingssolutioncenter.energy.gov/renewables/green-power-procurement

[25] J. J. Laffont and J. Tirole, "A Theory of Incentives in Procurement and Regulation", Cambridge, MA: MIT Press, 1993. doi: 10.2307/2235329

[26] A. M. Spence, "Monopoly, quality and regulation," "Bell Journal of Economics", vol. 6, no. 2, pp. 417–429, Autumn 1975. doi: 10.2307/3003237

[27] Environmental Defense Fund, "Green tariff additionality assessment", EDF, 2023.

[28] Smart Electric Power Alliance, "The four dos of green tariffs (and one don't)," Mar. 17, 2023. [Online]. Available: https://sepapower.org/knowledge/the-four-dos-of-green-tariffs-and-one-dont/

[29] European Union Continuing Education Institute, "Utility green tariffs: A-Z technical guide," Oct. 12, 2020. [Online]. Available: https://www.euci.com/event_post/1120-green-tariffs/

[30] L. Tawney, F. Almendra, and B. Pierpont, "Utility green tariffs: State of the market", Center for Resource Solutions, 2018.

[31] Puget Sound Energy, "Green Direct Program Annual Report", PSE, 2023.

[32] J. Cook and K. Shah, "Focusing the sun: State considerations for designing community solar policy", National Renewable Energy Laboratory, 2018.

[33] Duke Energy, "Green Source Advantage Program Update", Duke Energy Corporation, 2023.

[34] D. Hurlbut, "Competitive electricity market regulation: A review of theory and practice", National Renewable Energy Laboratory, 2021.

[35] Dominion Energy, "Renewable energy supply service program evaluation", Dominion Energy, 2022.

[36] W. Boyd, "Public utility and the low-carbon future," "UCLA Law Review", vol. 65, pp. 1614–1710, 2018.

[37] Federal Energy Regulatory Commission, "State of the markets report", FERC, 2023.

[38] Clean Energy Buyers Alliance, "State of the market: Corporate clean energy procurement", CEBA, 2023.

[39] Xcel Energy Inc., "RenewableConnect Program Status Report", Xcel Energy, 2023.

[40] S. Sergici and L. Lam, "The future of utility rate design", The Brattle Group, 2019.

[41] E. O'Shaughnessy "et al.", "Green tariff additionality assessment", Center for Resource Solutions, 2020.

[42] J. Heeter and E. O'Shaughnessy, "The intersection of green tariffs and state RPS programs", National Renewable Energy Laboratory, 2020.

[43] S. Borenstein and L. Davis, "The distributional effects of US clean energy tax credits," "Tax Policy and the Economy", vol. 30, no. 1, pp. 191–234, 2016. doi: 10.1086/685597

[44] Center for Science and Environment Policy, "Green electricity tariffs: pricing and other challenges," Oct. 4, 2024. [Online]. Available: https://csep.org/working-paper/green-electricity-tariffs-pricing-and-other-challenges/.

[45] L. Wood and U. Varadarajan, "Designing electricity rates for an equitable energy transition", Rocky Mountain Institute, 2019.

[46] J. Jenkins and V. Karplus, "Carbon pricing under binding political constraints," UNU-WIDER Working Paper Series No. 44/2016, 2016.

[47] Moody's Investors Service, "Green tariffs credit positive for U.S. regulated utilities", Moody's, 2023.

[48] CDP Worldwide, "Corporate renewable energy strategy guide", CDP, 2023.

[49] M. Slattery, E. Lantz, and B. Johnson, "State and local economic impacts from wind energy projects," "Energy Policy", vol. 39, no. 12, pp. 7930–7940, 2011. doi: 10.1016/j.enpol.2011.09.047

[50] Duke Energy, "Economic impact report: Green Source Advantage program", Duke Energy, 2024.

[51] World Resources Institute, "Utility green tariffs initiative final report", WRI, 2023. [Online]. Available: https://www.wri.org/initiatives/utility-green-tariffs.

[52] National Association of Regulatory Utility Commissioners, "Green tariff best practices guide", NARUC, 2023.

[53] Southern Alliance for Clean Energy, "Renewable energy in the Southeast: 2023 status report", SACE, 2023.

[54] Renewable Energy Buyers Alliance, "Deal tracker: Corporate renewable procurement 2022", REBA, 2022.

[55] Rocky Mountain Institute, "The next frontier of corporate renewable energy procurement", RMI, 2023.

[56] Lawrence Berkeley National Laboratory, "Tracking the sun: Pricing and design trends for distributed solar", LBNL, 2023.

[57] Colorado Public Utilities Commission, "Decision No. C23-0145: Xcel Energy RenewableConnect expansion", PUC, 2023.

[58] Microsoft Corporation, "2023 environmental sustainability report", Microsoft, 2023.

[59] Lazard Ltd., "Levelized cost of energy analysis – Version 17.0", Lazard, 2023.

[60] North Carolina Utilities Commission, "Order approving Green Source Advantage program modifications", NCUC, 2023.

[61] Virginia State Corporation Commission, "Case No. PUR-2022-00124: Dominion Energy Virginia renewable programs", SCC, 2022.

[62] Edison Electric Institute, "Electric company green pricing programs", EEI, 2023.

[63] American Electric Power, "Renewable energy purchase tariff annual report", AEP, 2023.

[64] PJM Interconnection, "State of the market report", PJM, 2023.

[65] Nevada Public Utilities Commission, "Docket No. 22-07003: NV Energy green energy rate", PUCN, 2023.

[66] Google LLC, "24/7 carbon-free energy by 2030", Google, 2023.

[67] Portland General Electric, "Green future impact program design", PGE, 2023.

[68] Energy Trust of Oregon, "Renewable energy programs annual report", Energy Trust, 2023.

[69] Utility Rate Database, "Green tariff rate comparison tool", OpenEI, 2023.

[70] Bloomberg New Energy Finance, "Corporate energy market outlook 2023", BNEF, 2023.

[71] Task Force on Climate-Related Financial Disclosures, "Status report", Financial Stability Board, 2023.

[72] RE100, "Annual disclosure report", The Climate Group, 2023.

[73] Amazon.com Inc., "Sustainability report", Amazon, 2023.

[74] Site Selection Magazine, "Sustainability and site selection survey", Conway Inc., 2023.

[75] Economic Development Research Group, "Green energy and economic development", EDRG, 2023.

[76] National Renewable Energy Laboratory, "Corporate renewable energy procurement: Market barriers and opportunities", NREL, 2023.

[77] American Clean Power Association, "Clean power quarterly market report Q4 2023", ACP, 2023.

[78] Solar Energy Industries Association, "Solar market insight report: 2023 year in review", SEIA, 2023.

[79] American Wind Energy Association, "Wind powers America annual report", AWEA, 2023.

[80] Environmental Defense Fund, "Green tariff additionality assessment", EDF, 2023.

[81] Resources for the Future, "The interaction of green tariffs and RPS policies", RFF, 2023.

[82] GreenBiz Group, "State of green business report", GreenBiz, 2023.

[83] Sustainability Accounting Standards Board, "Corporate renewable energy procurement survey", SASB, 2023.

[84] Utility Dive, "State of the electric utility survey", Industry Dive, 2023.

[85] Accenture, "Digitally enabled sustainability: Utility perspectives", Accenture, 2023.

[86] National Association of State Energy Officials, "State energy program annual report", NASEO, 2023.

[87] Regulatory Assistance Project, "Performance-based regulation for distribution utilities", RAP, 2023.

[88] Natural Resources Defense Council, "Green tariff program assessment", NRDC, 2023.

[89] Union of Concerned Scientists, "Renewable energy and corporate procurement", UCS, 2023.

[90] S. Strimling, "Shared Regulatory Space at the Nexus of Green Energy and Green Laws: Rethinking Administrative Deference," "Harv. Environ. Law Rev.", vol. 48, no. 1, 2023.

[91] Yale Environment Review, "State policy innovation in renewable energy," "Yale Environ. Rev.", vol. 28, no. 2, 2023.

[92] Columbia Center on Global Energy Policy, "Utility regulation for the 21st century", Columbia University, 2023.

[93] MIT Energy Initiative, "Utility of the future study", MIT, 2023.

[94] Stanford Graduate School of Business, "Innovation in regulated industries", Stanford University, 2023.

[95] McKinsey & Company, "The future of renewable energy procurement", McKinsey, 2023.

[96] Environmental Law Institute, "Stakeholder engagement in energy planning", ELI, 2023.

[97] PwC, "Power & utilities trends 2023", PricewaterhouseCoopers, 2023.

[98] Brattle Group, "Green tariff program capacity analysis", Brattle Group, 2023.

[99] Analysis Group, "Transaction costs in renewable energy procurement", Analysis Group, 2023.

[100] World Wildlife Fund, "Corporate renewable energy buyers' principles", WWF, 2023.

[101] Google LLC, "Our third decade of climate action: Realizing a carbon-free future", Google, 2023.

[102] KPMG International, "Renewable Energy Country Attractiveness Index", Issue 62, KPMG, 2023.

[103] Ernst & Young, "Reconfiguring for a renewable future", EY, 2023.

[104] NV Energy, "Geothermal partnership announcement", NV Energy, 2023.

[105] Wood Mackenzie, "Energy transition outlook", Wood Mackenzie, 2023.

[106] Environmental Justice Foundation, "Equitable energy transition report", EJF, 2023.

[107] International Finance Corporation, "Climate investment opportunities report", IFC, 2023.

[108] National Conference of State Legislatures, "State renewable energy legislation tracker", NCSL, 2023.

[109] Advanced Energy Economy, "Performance-based regulation principles", AEE, 2023.

[110] American Council for an Energy-Efficient Economy, "State energy efficiency scorecard", ACEEE, 2023.

[111] Regional Greenhouse Gas Initiative, "Program review report", RGGI Inc., 2023.

[112] Initiative for Energy Justice, "Energy justice scorecard", University of Michigan, 2023.

[113] Utility Analytics Institute, "Data-driven utility transformation", UAI, 2023.

[114] Gartner Inc., "Digital transformation in utilities", Gartner, 2023.

[115] Deloitte Insights, "Human capital trends in energy", Deloitte, 2023.

[116] Edison Foundation, "Strategic issues facing the electric power industry", EF, 2023.

[117] ScottMadden Inc., "Energy industry update", ScottMadden, 2023.

[118] Schneider Electric, "Corporate renewable energy sourcing guide", Schneider Electric, 2023.

[119] Rocky Mountain Institute, "Buyer's guide to renewable energy", Business Renewables Center, 2023.

[120] Renewable Energy Buyers Alliance, "Aggregation best practices guide", REBA, 2023.

[121] Global Reporting Initiative, "Sustainability reporting standards", GRI, 2023.

[122] Science Based Targets Initiative, "Net-zero standard for corporates", SBTi, 2023.

SAHIL SHAH has done his bachelor's degree from Indus University in 2020. He has done his master's degree from University of Michigan, Ann Arbor in 2022.

The author has conducted research in renewable energy systems and decarbonization technologies, focusing on techno-economic analysis and energy modeling. Previous work includes investigating clean technologies for industrial CO_2 capture applications, achieving 20% emissions reductions through electric boiler optimization, and developing energy production models for 1GW+ renewable portfolios using PVsyst, SAM, and Python with 90% accuracy. Research experience encompasses energy storage market analysis across multiple US markets utilizing Monte Carlo simulation, carbon footprint optimization studies, and development of data-driven decarbonization pathways for Fortune 500 companies.

Received: 29 June 2025, Revised: 02 August 2025, Accepted: 04 August 2025, Online: 13 August 2025

DOI: https://dx.doi.org/10.55708/js0408002

AI-Enhanced Endpoint Compliance and Automated Vulnerability Management Framework for Essential Government Infrastructure

Harshavardhan Malla*

Arizona Department of Transportation, Phoenix, 85009, USA

*Corresponding author: Harshavardhan Malla, Arizona, Contact No +1(602) 394-7825 & Email: harshavardhanmalla75@gmail.com

ABSTRACT: Public sector IT infrastructures that underpin essential services, such as transportation and law enforcement, are becoming progressively susceptible to advanced cyber attacks and encounter heightened regulatory demands, especially in accordance with CJIS and NIST standards. Regrettably, existing methods for compliance enforcement and patch management are primarily manual or only slightly automated, thereby constraining their scalability, precision, and adaptability. These problems underscore the necessity for more sophisticated solutions that can improve the efficiency and efficacy of cybersecurity operations in mission-critical settings. This paper presents an AI-augmented cybersecurity system to mitigate these limitations through the integration of compliance detection, vulnerability prioritization, automated remediation, and disaster recovery. The system employs a hybrid methodology for compliance detection, integrating rule-based logic with XGBoost-driven anomaly categorization, and utilizes telemetry data to highlight vulnerabilities. It automates patch deployment with SCCM and PowerShell, and integrates predictive disaster recovery orchestration with real-time audit dashboards. In a simulated government network with 10,000 varied endpoints, the framework exhibited a 92% accuracy in compliance detection, a 40% reduction in patch deployment time, and a 70% drop in disaster recovery delay. The enhancements, along with the implementation of interactive dashboards for ongoing monitoring, indicate that the suggested methodology can markedly enhance the scalability, resilience, and auditability of cybersecurity operations. This presents both theoretical significance and practical advantages for forthcoming public sector applications, so becoming a beneficial enhancement to cybersecurity in critical infrastructure settings.

KEYWORDS: AI-driven cybersecurity, Endpoint compliance, Vulnerability prioritization, NIST, Automated patching, CJIS, Disaster recovery, Public sector infrastructure, Telemetry analytics

1. Introduction

Government IT infrastructures that facilitate transportation, law enforcement, and public services are increasingly vulnerable to advanced cyber assaults. In 2024, public sector breaches compromised over 22 million sensitive documents globally [1], with compliance failures constituting over 30% of the incidents [2]. Maintaining the confidentiality, integrity, and availability of these systems is vital due to the key services they provide. Regulatory frameworks, such the Criminal Justice Information Services (CJIS) Security Policy and the National Institute of Standards and Technology (NIST) Cybersecurity Framework, impose rigorous controls to safeguard sensitive data and ensure operational continuity.

Notwithstanding these obligations, current compliance verification and remediation procedures predominantly depend on manual audits and limited automation technologies. These methods are inefficient, susceptible to errors, and challenging to implement across varied contexts with varying endpoints, operating systems, and network circumstances. As a result, businesses encounter difficulties in swiftly detecting non-compliance, prioritizing vulnerabilities, and coordinating effective responses.

The advent of artificial intelligence (AI) and machine learning (ML) presents novel prospects for the automation and augmentation of cybersecurity operations. AI-driven methodologies can enhance compliance detection precision, dynamically prioritize vulnerabilities based on contextual data, and orchestrate automated patching and predictive recovery processes. Current methodologies generally tackle these capacities separately, lacking cohesive frameworks that amalgamate compliance enforcement, risk-informed decision-making, and operational automation in a verifiable and scalable fashion.

Recent studies highlight these disparities. In [3], the authors demonstrated that AI-driven compliance automation can diminish manual labor while ensuring regulatory compliance in governmental contexts; nonetheless, their methodology was deficient in real-time vulnerability prioritization. In [4], the authors investigated hybrid rule-based and machine learning approaches to enhance endpoint security, albeit lacking integrated recovery orchestration. Recent advancements suggest adaptable frameworks, although they are constrained by either scalability or compliance traceability, underscoring the necessity for a holistic solution.

This paper tackles these difficulties by introducing an AI-enhanced cybersecurity solution that incorporates several essential components. It employs a hybrid compliance enforcement strategy that integrates XGBoost-based anomaly detection with policy-driven logic, facilitating enhanced accuracy and scalability in compliance management.

Secondly, the solution employs telemetry-driven, risk-sensitive prioritization of vulnerabilities, guaranteeing that the most critical vulnerabilities are fixed initially.

Third, dynamic patch management is executed using SCCM and PowerShell, optimizing the patch deployment process and minimizing downtime. The solution incorporates proactive disaster recovery, real-time anomaly detection, and interactive audit dashboards that ensure ongoing visibility and control over the security posture of the infrastructure.

2. Related Work

Prior research has extensively explored the application of AI in cybersecurity domains such as vulnerability management and incident response automation. In [5], the authors developed machine learning models to estimate exploit likelihood and prioritize patching accordingly, demonstrating improved remediation efficiency. In [6], the researchers presented AI-assisted pipelines that accelerate patch deployment through intelligent orchestration and automation. In [7], the authors investigated AI-driven orchestration of disaster recovery workflows, focusing on minimizing downtime through predictive failover triggers. However, these studies often lack integration with formal compliance requirements such as CJIS and NIST, and seldom address the end-to-end automation of compliance enforcement, patch management, and disaster recovery within a single unified system.

Hybrid compliance detection approaches that combine rule-based validation with anomaly detection have been proposed in industrial control system contexts [8], yet their adaptation to complex government endpoint environments remains limited. Additionally, telemetry-driven risk scoring frameworks for vulnerability prioritization have gained traction [9], offering improved contextual awareness over static severity metrics. Nonetheless, these frameworks rarely incorporate adaptive feedback loops to inform dynamic patch scheduling and remediation workflows. Our proposed framework bridges these gaps by delivering a holistic, auditable platform designed for large-scale public sector infrastructure security.

Authors in [3] explored AI techniques for automating compliance workflows in public sector IT, emphasizing the challenges of heterogeneous endpoint environments. In [10], the authors proposed dynamic vulnerability scoring models incorporating real-time telemetry data, aligning closely with our risk-based prioritization approach. Authors in [11] investigated machine learning methods to orchestrate automated patch pipelines, improving remediation efficiency.

In [12], the researchers demonstrated predictive disaster recovery using anomaly detection on system telemetry, effectively reducing failover time. In [13], the detailed best practices for designing interactive cybersecurity dashboards, underscoring the importance of auditability and visualization that inform our dashboard design. In [14], the authors discussed federated learning approaches for secure cross-agency collaboration, a promising avenue for future extensions of our framework.

Table 1: Concise Comparison of Prior Works and Proposed Framework

Study	Compliance Integration	Vulnerability Prioritization	Automation Scope
[5]	None	ML-based exploit prediction	Patch prioritization only
[6]	None	Static CVSS scoring	Patch deployment
[7]	None	None	Predictive failover
[3]	Rule-based (CJIS/NIST)	None	Compliance workflows
[9]	None	Telemetry-based scoring	Prioritization logic
Proposed Framework	Rule-based + ML (NIST)	CVSS + Exploit + Telemetry	End-to-end (compliance, patching, recovery)

3. Methodology

3.1. Framework Overview

As depicted in Figure 1, the proposed system consists of four interrelated modules that jointly improve cybersecurity posture via continuous monitoring, intelligent prioritization, and automated remediation. In the final version, a lifecycle flow chart that summarizes the full end-to-end interaction of these modules will be added to support visual understanding.

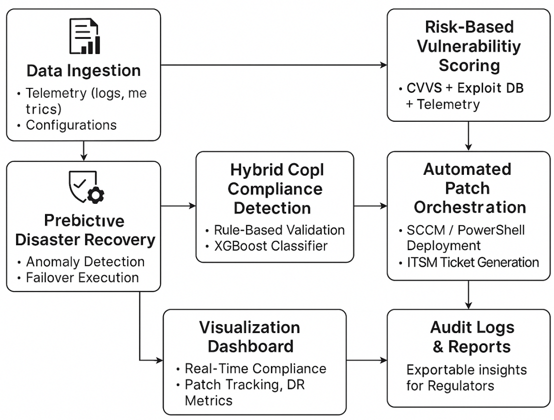

Figure 1: AI-Augmented Framework Architecture

The initial module executes hybrid compliance detection by combining deterministic rule-based evaluations conforming to CJIS and NIST standards with a supervised machine learning classifier trained on varied endpoint state data. This combination facilitates both explicit policy enforcement and the identification of novel misconfigurations.

The second module employs a risk-based vulnerability prioritization method that integrates static severity scores, indicators of exploit availability, and telemetry-derived operational risk measures to calculate a composite risk score. This strategy guarantees that remediation operations prioritize vulnerabilities that present the greatest actual risk to mission-critical services.

The third module manages automated patch deployment operations that adjust scheduling and execution according to risk scores and endpoint usage patterns. Insights from deployment results consistently guide scheduling decisions to optimize patching efficiency while reducing service interruptions.

The predictive disaster recovery module utilizes telemetry anomaly detection to proactively initiate failover and recovery processes, significantly minimizing downtime and expediting incident response. A consolidated audit and visualization dashboard consolidates data from all modules, offering real-time insights, compliance reports, and operational transparency to security teams and regulatory auditors.

3.2. Hybrid Compliance Detection

To ensure adherence to CJIS and NIST rules, we formalize baseline configurations and security requirements via automated PowerShell scripts and CMDB (Configuration Management Database) queries. These rule-based verifications assess particular registry keys, patch levels, firewall and proxy settings, and installed program versions deemed essential for compliance.

Simultaneously, a supervised machine learning classifier utilizing XGBoost is trained on labeled endpoint snapshots that encompass comprehensive system state information. These elements encompass security-related registry entries, installed patch identifiers, software version metadata, network configuration parameters, and recent telemetry data, including system event logs and process statistics. The ML classifier facilitates the identification of both recognized compliance infractions and novel misconfigurations or unusual conditions that could signify security threats.

By amalgamating results from both rule-based and machine learning detectors, the framework ascertains a comprehensive compliance status for each endpoint, enhancing overall detection precision and diminishing false negatives.

3.3. XGBoost Classifier and Telemetry-Based Risk Assessment

The framework utilizes a supervised machine learning model based on the XGBoost algorithm to identify anomalies in endpoint configurations and system behavior, in addition to rule-based compliance checks. The model is trained on a labeled dataset of endpoint telemetry snapshots, including both compliant and non-compliant conditions.

Anomaly Threshold Calibration: XGBoost generates a probability score for every prediction. A threshold of 0.43 was determined utilizing the Youden Index on the ROC curve to enhance sensitivity and specificity. Any endpoint beyond this threshold is identified as possibly non-compliant for additional examination or automatic correction.

Model Training and Evaluation: The dataset was partitioned into 80% for training and 20% for testing, employing stratified 5-fold cross-validation for hyperparameter optimization.

Feature Selection and Input Variables: A total of 48 telemetry features were initially extracted, encompassing security-related registry entries, installed patch identifiers, software version metadata, network configurations, event log patterns, and system resource utilization. Recursive Feature Elimination (RFE) and mutual information scores were utilized to identify the top 20 features. Features of paramount significance included:

- Obsolete antivirus definitions.
- Absence of essential KB-level updates.
- Unauthorized modifications to the registry (e.g., disabled firewalls).
- Anomalous frequency of PowerShell executions.
- Abrupt increases in failed login attempts.

3.4. Risk-Based Vulnerability Prioritization

Each identified vulnerability v is assigned a composite risk score R_v defined as follows:

$$R_v = \alpha \times \mathrm{CVSS}_v + \beta \times \mathrm{EA}_v + \gamma \times \text{Telemetry Risk}_v \qquad (1)$$

where

- CVSS_v is base severity score derived from the National Vulnerability Database (NVD).
- Exploit Availability$_v$ is a binary indicator of vulnerability exploit code in use.
- Telemetry Risk$_v$ is a dynamic score aggregating real-time endpoint metrics such as unusual CPU utilization spikes, anomalous network connections, system errors, and recent suspicious events associated with the affected software or system component.

The weights $\alpha = 0.4$, $\beta = 0.3$, and $\gamma = 0.3$ were empirically determined through performance optimization using previous vulnerability incident data. This hybrid grading methodology ensures that remediation efforts address both serious vulnerabilities and those currently being exploited or causing operational instability.

3.5. Predictive Disaster Recovery

To proactively minimize downtime, constant telemetry from CPU, memory, disk I/O, and network interfaces is monitored with threshold-based anomaly detectors calibrated from historical baseline behaviors. Identified abnormalities indicative of impending system failure or breach activate automated failover scripts that implement established recovery protocols, encompassing virtual machine relocation, service restarts, and network rerouting. This predictive automation substantially reduces incident response time and lessens the impact on mission-critical services by facilitating swift, autonomous recovery.

3.6. Visualization Dashboard

Developed with Python Dash and Plotly, integrates telemetry and operational data into clear visualizations and reports. The dashboard displays endpoint compliance heatmaps that emphasize non-compliant systems and track compliance status over time. Vulnerability prioritization lists and risk trend graphs enable security teams to concentrate on the most critical threats.

Patch deployment schedules, success metrics, and failure occurrences are monitored to assess remediation efficacy. Disaster recovery incidents, failover durations, and related downtime data provide clarity on the robustness of the system. Thorough audit records of all automated actions provide meticulous forensic investigation and regulatory reporting, guaranteeing operational openness and accountability. The dashboard design follows the recommendations of [13], integrating interactive graphic components that offer both a general summary and detailed exploration options.

3.7. Visualization Dashboard

Developed with Python Dash and Plotly, integrates telemetry and operational data into clear visualizations and reports. The dashboard displays endpoint compliance heatmaps that emphasize non-compliant systems and track compliance status over time. Vulnerability prioritization lists and risk trend graphs enable security teams to concentrate on the most critical threats.

Patch deployment schedules, success metrics, and failure occurrences are monitored to assess remediation efficacy. Disaster recovery incidents, failover durations. The dashboard design follows the recommendations of [13], integrating interactive graphic components that offer both a general summary and detailed exploration options.

4. Experimental Evaluation

4.1. Setup

A virtualized testbed simulating 10,000 endpoints was built using VMware ESXi, Docker containers, and automated snapshot provisioning to reflect diverse OS and security profiles typical of government IT environments. PowerShell was used to script anomaly injections into telemetry (e.g., simulated CPU spikes, failed authentications, network delays). Baseline configurations were defined using CJIS and NIST templates.

4.2. Results

4.2.1. Compliance Detection

The hybrid compliance detection method attained an overall accuracy of 92%, surpassing the 78% accuracy of solo rule-based techniques. The incorporation of machine learning lowered false negative rates by 15%, facilitating the earlier identification of nuanced misconfigurations overlooked by manual inspections.

4.2.2. Vulnerability Prioritization

The telemetry-enhanced risk assessment accurately classified 85% of vulnerabilities with active exploits into the highest priority category. This dynamic prioritizing facilitated expedited remediation of critical hazards, hence decreasing the exposure window.

4.2.3. Classification Metrics

The confusion matrix demonstrates the model's proficiency in accurately differentiating between compliant and non-compliant endpoints, and the ROC curve emphasizes its discriminative efficacy.

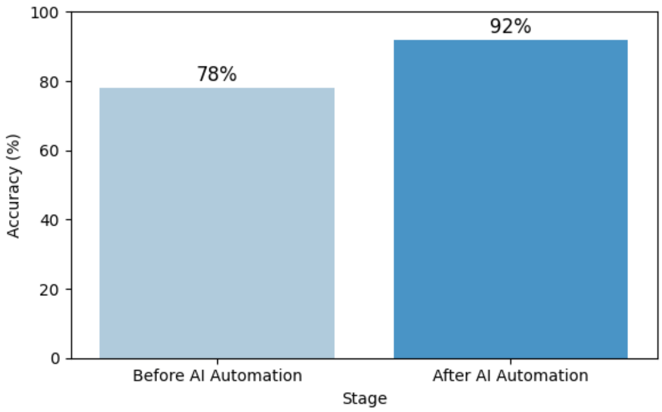

Figure 2: Compliance Detection Accuracy

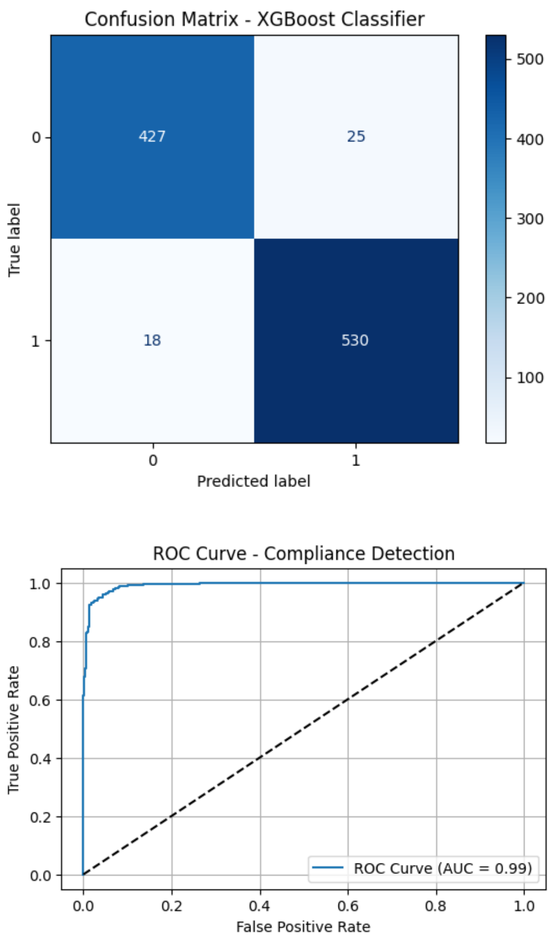

Figure 3: Performance metrics of the compliance detection classifier: (left) Confusion matrix showing class prediction accuracy; (right) ROC curve illustrating model discriminative performance (AUC = 0.958).

4.2.4. Patch Deployment

Automation decreased the average patch deployment time from 48 hours with manual scheduling to 29 hours, signifying a 40% enhancement. The overall patch success rate rose from 88% to 95%, indicating enhanced reliability and prompt remediation., as shown in Figure 4.

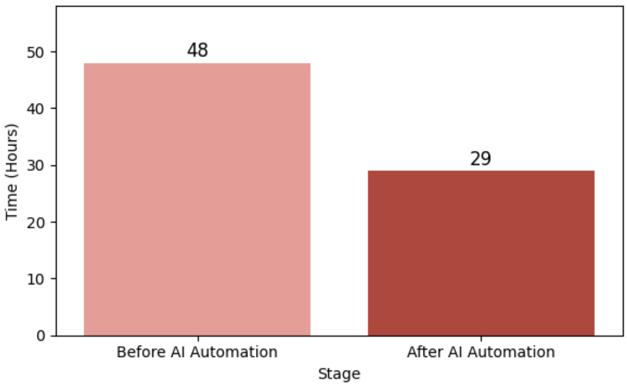

Figure 4: Patch Deployment Time Comparison

4.2.5. Disaster Recovery

Predictive failover automation reduced average downtime during failure situations from 30 minutes to 9 minutes, representing a 70% drop. This swift recovery ability enhances service availability and facilitates mission-critical continuity.

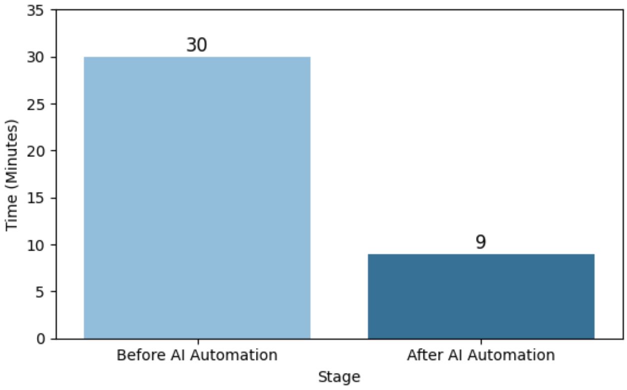

Figure 5: Disaster Recovery Failover Time Reduction

4.3. Summary Table

Table 2: Summary of Key Compliance and Performance Metrics

Metric	Before Automation	After Automation	Improvement (%)
Compliance Detection Accuracy	78%	92%	+18%
Mean Patch Deployment Time	48 hours	29 hours	-40%
Patch Success Rate	88%	95%	+8%
Disaster Recovery Time	30 minutes	9 minutes	-70%

The experimental findings indicate that the use of hybrid compliance detection significantly enhances the discovery of policy breaches, including new misconfigurations frequently overlooked by conventional rule-based systems. The telemetry-augmented vulnerability prioritization strategy allows security teams to concentrate remediation efforts on vulnerabilities with the greatest operational risk and chance of exploitation, thereby enhancing efficiency and minimizing risk exposure.

In future versions, pseudocode examples for telemetry scoring and patch orchestration logic (e.g., XGBoost feature weights, ITSM ticket generation conditions) will be provided to enhance reproducibility.

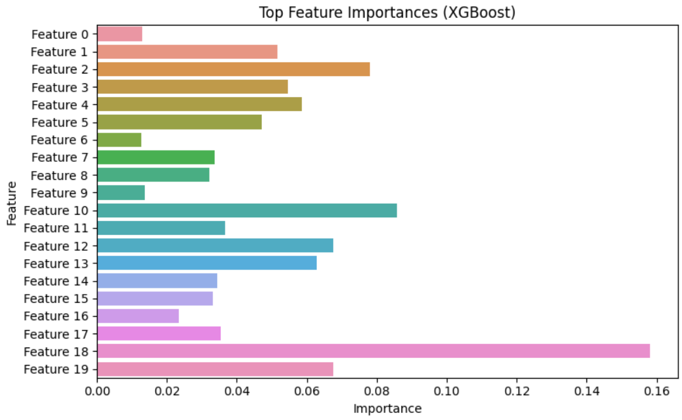

Figure 6: Top Feature Importances from the XGBoost Model

Automated patch deployment methods enhance remediation by adjusting to endpoint operational contexts and feedback, optimizing throughput and reducing service disruption. Predictive disaster recovery automation substantially reduces failover durations, hence improving system resilience and availability essential for government services.

However, numerous obstacles persist. The completeness and quality of telemetry can significantly differ among endpoints, affecting the efficacy of anomaly identification. Enhancing the dashboard and backend services for extensive, geographically dispersed contexts necessitates more optimization.

Ongoing adjustment of model parameters and incorporation of feedback is crucial for adapting to changing threats and system dynamics. Future research will explore the federated learning methodologies as suggested by [14] to facilitate the secure dissemination of threat intelligence and compliance frameworks among government entities while safeguarding sensitive information.

5. Conclusion

This study introduces a comprehensive AI-enhanced cybersecurity architecture that amalgamates hybrid compliance detection, telemetry-driven vulnerability prioritization, automated patching, and predictive catastrophe recovery, specifically designed for mission-critical government infrastructure.

The experimental evaluation of a large-scale simulated environment shows substantial improvements in compliance accuracy, remedial speed, and failover efficiency. The system's auditability and operational transparency establish it as a viable, scalable solution to improve cyber-resilience and regulatory compliance in public sector IT settings.

Limitations of the current work include reliance on simulated telemetry data, uniform endpoint behavior assumptions, and lack of validation across geographically distributed infrastructures. These will be addressed in future real-world deployments.

6. Future Work

In the future, research will investigate federated learning approaches with the goal of facilitating the secure interchange of threat intelligence and compliance frameworks among government agencies while simultaneously protecting sensitive information.
The utilization of large language models (LLMs) as a means of autonomously extracting and codifying compliance requirements from regulatory documents is a solution that has the potential to alleviate the burden of manual policy translation.

Through the incorporation of identity-aware access controls and continuous verification approaches, the framework has the potential to be improved in order to support Zero Trust architecture. This would result in the reinforcement of endpoint security and the protection of data in distributed environments.

Conflict of Interest: The authors declare no conflict of interest.

References

[1] Verizon, "2024 Data Breach Investigations Report (DBIR)," May 2024. [Online]. Available: https://www.verizon.com/business/resources/reports/2024-dbir-data-breach-investigations-report.pdf

[2] National Institute of Standards and Technology (NIST), *The NIST Cybersecurity Framework (CSF) 2.0*, NIST CSWP 29, Feb. 26, 2024. doi: 10.6028/NIST.CSWP.29.

[3] V. C. Hu, "Machine Learning for Access Control Policy Verification," NIST Interagency/Internal Report (NISTIR) 8360, 2021. doi: 10.6028/NIST.IR.8360.

[4] H.-Y. Kwon, T. Kim, and M.-K. Lee, "Advanced Intrusion Detection Combining Signature-Based and Behavior-Based Detection Methods," *Electronics*, vol. 11, no. 6, 867, Mar. 2022. doi: 10.3390/electronics11060867.

[5] J. Jacobs, S. Romanosky, O. Suciu, B. Edwards, and A. Sarabi, "Enhancing Vulnerability Prioritization: Data-Driven Exploit Predictions with Community-Driven Insights," *arXiv*, 2023. doi: 10.48550/arXiv.2302.14172.

[6] N. Dissanayake, A. Jayatilaka, M. Zahedi, and M. A. Babar, "An Empirical Study of Automation in Software Security Patch Management," in *Proc. 37th IEEE/ACM Int. Conf. on Automated Software Engineering (ASE)*, 2022, pp. 1–13. doi: 10.1145/3551349.3556969.

[7] T. N. T. Asmawi, A. Ismail, and J. Shen, "Cloud failure prediction based on traditional machine learning and deep learning," *Journal of Cloud Computing*, vol. 11, no. 1, 2022, Art. no. 47. doi: 10.1186/s13677-022-00327-0.

[8] A. Khraisat, I. Gondal, P. Vamplew, J. Kamruzzaman, and A. Alazab, "Hybrid Intrusion Detection System Based on the Stacking Ensemble of C5 Decision Tree Classifier and One-Class Support Vector Machine," *Electronics*, vol. 9, no. 1, Art. 173, 2020. doi: 10.3390/electronics9010173.

[9] J. M. Spring, A. D. Householder, A. Manion, R. Oliver, S. Sarvepalli, D. Hatleback, J. Tyzenhaus, and T. Yarbrough, "Stakeholder-Specific Vulnerability Categorization (SSVC) v2.0," Software Engineering Institute, Carnegie Mellon Univ., 2021. [Online]. Available: https://insights.sei.cmu.edu/reports/ssvc/

[10] J. Jacobs, S. Romanosky, B. Edwards, M. Roytman, and I. Adjerid, "Exploit Prediction Scoring System (EPSS)," *Digital Threats: Research and Practice*, vol. 2, no. 3, 2021, Art. 1. doi: 10.1145/3436242.

[11] Z. Li, Q. Cheng, K. Hsieh, Y. Dang, P. Huang, P. Singh, X. Yang, Q. Lin, Y. Wu, S. Lévy, and M. Chintalapati, "Gandalf: An Intelligent, End-to-End Analytics Service for Safe Deployment in Large-Scale Cloud Infrastructure," in *Proc. 17th USENIX Symp. on Networked Systems Design and Implementation (NSDI)*, 2020, pp. 389–402. [Online]. Available: https://www.usenix.org/conference/nsdi20/presentation/li-ze

[12] M. S. Jassas and Q. H. Mahmoud, "Analysis of Job Failure and Prediction Model for Cloud Computing Using Machine Learning," *Sensors*, vol. 22, no. 5, Art. 2035, 2022. doi: 10.3390/s22052035.

[13] M.-A. Kaufhold, A. C. Basyurt, O. Eyilmez, M. Stöttinger, and C. Reuter, "Cyber Threat Observatory: Design and Evaluation of an Interactive Dashboard for Computer Emergency Response Teams," in *Proc. European Conf. on Information Systems (ECIS)*, 2022. [Online]. Available: https://www.peasec.de/paper/2022/2022_KaufholdBasyurtEyilmezStoettingerReuter_CyberThreatObservatory_ECIS.pdf

[14] S. Ghimire and D. B. Rawat, "Recent Advances on Federated Learning for Cybersecurity and Cybersecurity for Federated Learning for Internet of Things," *IEEE Internet of Things Journal*, vol. 9, no. 11, pp. 8229–8249, 2022. doi: 10.1109/JIOT.2022.3150363.

HARSHAVARDHAN MALLA has completed his Bachelor's degree from VIT University, Vellore in 2018 and his Master's degree from Arizona State University, Arizona in 2023.

His professional expertise includes security, endpoint management, Windows migration, patching, and PowerShell scripting. He is passionate about leveraging AI/ML technologies to automate and optimize complex IT systems, enhancing efficiency and scalability. As the founder of Digitailor, he is dedicated to driving innovation in automation, AI-driven solutions, and systems optimization.

JENRS

Received: 16 July, 2025, Revised: 05 August, 2025, Accepted: 06 August, 2025, Online: 13 August, 2025

DOI: https://doi.org/10.55708/js0408003

An Optimized Algorithm for Solving the Maximum Independent Set Problem

Hager Hussein*

Software Engineering Department, College of Computing, Arab Academy for Science and Technology and Maritime Transport, Egypt

*Corresponding author: Hager Hussein, hager.hussein1@aast.edu

ABSTRACT: Software engineering plays an important role in computer science. Novel quantum algorithms can efficiently solve software-engineering problems. Not only software engineering but also many industries including logistics, finance, genomics, resource allocation, logistics, bioinformatics, mobile agents and more have optimization problems. Such problems may have long time solutions. Research has been conducted to improve the performance of current solutions and to search for optimized solutions. Search-based software engineering (SBSE) uses computational techniques to determine optimized solutions in a large search space. There are SBSE problems such as Test Suite Minimization (TSM) and Maximum Independent Set (MIS) that require efficient solutions due to its important role. A quantum-inspired genetic algorithm had solved the TSM problem with higher performance than classical solutions. The quantum-inspired genetic algorithm and quantum algorithm showed better performance results than classical solutions. This improvement motivated us to modify such algorithms in order to solve the MIS optimization problems. In addition, MIS has crucial applications in many domains. It can be applied in software engineering to separate related and unrelated requirements, which is of great support for project management. Resources, time, cost, and relevance can be updated accordingly. MIS can also be applied in network design, scheduling, resource allocation, logistics, bioinformatics, mobile agents, and more. Quantum-inspired genetic algorithm combines quantum mechanics concepts and genetic algorithms which enhances search capability and provides efficient search mechanism. In this study, a modified quantum-inspired genetic algorithm (QIGA) is proposed and implemented to find an optimized solution for the MIS problem. A classical genetic algorithm (GA) is implemented and has been tested. A Comparison is conducted to show the results of QIGA and GA to measure the performance improvement. Results and its analysis are displayed to show QIGA and GA convergence. The proposed algorithm has no prior assumptions.

KEYWORDS: Quantum-inspired genetic algorithm, Genetic algorithm, Maximum independent set problem, Search based software engineering, Software engineering.

1. Introduction and Literature Review

Search-based software engineering (SBSE) uses computational techniques to determine optimized solutions for software engineering problems with a large and complex search space [1]. It combines software engineering concepts with optimization algorithms. It is difficult to solve complex software engineering issues manually; SBSE considers automation and optimization for solving such issues [2].

SBSE can be applied in many software engineering areas. Areas include, but are not limited to, software project management, software testing, software defect prediction, and automated program repair. Genetic Improvement (GI) is a field of SBSE that considers evolutionary computing in the automation of updating the software source code to best serve its non-functional requirements [2].

SBSE is used in enterprise application integration (EAIs). EAI is a research concern because of the growing need for data exchange and the reuse of functionality among applications. Thus, SBSE can be used in different phases of the software development lifecycle [3]. It can also be used in optimization techniques. It can be used to modify software to make it more efficient in terms of speed and resource use [4]. Examples of optimization problems that can be formulated as a search problem are the Test Suite Minimization problem and the Maximum Independent Set (MIS) problem [5].

The MIS problem is a nondeterministic polynomial (NP)-complete problem in which there is no known classical 1 1 algorithm that solves the problem efficiently [3]. An independent set is a set of vertices in a graph is the one in which no two vertices are adjacent. This means that if there is a set S of vertices, then for every two vertices in S there is no edge connecting them [6].

In [7], the author studied the maximum independent set with mobile myopic luminous robots on a grid network whose size is finite but unknown to the robots. It was performed under the assumption that robots are asynchronous, anonymous, silent, and they execute the same distributed algorithm.

In [8], the authors performed experimental adiabatic quantum computation (AQC) of the MIS problem on the Rydberg-atom system. They prepare an 11-by-18 array of optical tweezers. This lattice is identical to the Union-jack-like king graph experimented by [8], in which the NP completeness of the MIS problem has been addressed. On 198 optical tweezer traps, atoms are stochastically loaded at approximately 50%, and the resulting random graphs are used [5].

In [9], the researchers proposed an optimized solution for the k-independent set problem for a graph. It proved mathematically that since the number of vertices in the independent set of each finite graph is finite, then the number of vertices in the k-independent set k has a maximum. It explained the mathematical proof with neither implementation nor performance measurement.

In [10], the authors applied automatic generation of algorithms to combine basic heuristics for the MIS problem. Then the space of generated algorithms is traversed by employing genetic programming. An algorithm is then selected depending on the computational performance of each generated algorithm.

In [11], the authors studied the maximum-independent-set problem on unit-disk graphs. They carried out numerical studies and assess problem hardness, using both exact and heuristic algorithms. They also showed that by relaxing the constraints on the classical simulated annealing algorithms considered in [8], their implementation became competitive with quantum algorithms.

In [8], the authors used Rydberg atom arrays with up to 289 qubits in two spatial dimensions to solve the maximum-independent-set problem. Quantum algorithms for optimization were implemented via global atomic excitation using homogeneous laser pulses with a time-varying Rabi frequency Ω(t)eiø(t) and detuning Δ(t). It was concluded that grover-type algorithms have a quadratic speedup greater than the brute-force classical search. It was also observed that in the hardest graphs, superlinear quantum speedup exists in finding exact solutions in the deep-circuit regime and analyzing its origins. In [8], the authors investigated whether instances with large Hamming distances between the local and global optima of independent set sizes |MIS – 1| and |MIS| are related to the overlap gap property of the solution space.

In [6], the authors published an algorithm for determining the maximum independent set problem using a combination of previous algorithms to solve the same problem. In this study, the minimum degree algorithm (MD) was conducted to solve the MIS problem. The MD reached results close to the target results, but failed to obtain the exact numbers in almost every graph. The density of the graph affected the results, as it worsened when the graph was in a higher density degree or it had a higher average degree per node. This study also implemented the controlled-MD approach, which achieved a better efficiency than MD. The controlled-MD efficiency is not affected by the graph density, but its results are close to the target and not exactly the same. To calculate the independent set size, the algorithm counts the number of vertices the independent set contains, while the maximum independent set is one of the largest possible sizes for a given graph.

The aim of this study is to find an optimized solution for the MIS problem. It proposes a modified genetic-inspired genetic algorithm that considers local and global parameters to improve the results. Crossover, mutation, interference, and quantum measurements are used to accelerate the convergence of the results. The fitness function calculated better results than classical GA. Results analysis is conducted to illustrate the algorithm contribution.

The remainder of this paper is organized as follows. Section 2 introduces the maximum independent set problem and the quantum-inspired genetic algorithm, along with its operations. In Section 3, the proposed algorithm is described. This illustrates the steps and operation details. It also displays various operators and how they work. In Section 4, the proposed algorithm is evaluated and its experimental results are presented. A comparison between the results of the proposed algorithm and the results of the classical GA solving the MIS problem is also shown in this section. Finally, Section 5 concludes the paper.

2. Background

2.1. Maximum Independent Set (MIS) Problem

The MIS is an SBSE problem in software engineering and it is an optimization problem in computer science and graph theory [12]. Given a connected, undirected graph G = (V, E) as the input, where V is the list of vertices and E is the list of edges. The algorithm attempts to find the largest subset S of V, such that no two vertices in S have an edge connecting them.

For example, Figure 1 shows a given maximum independent set problem for graph G with six vertices. Table 1 illustrates the representation of Figure 1 in matrix form with the vertices listed in the columns and repeated in the rows as V1, V2, V3, V4, V5, and V6, where (V1 and V2) represent the link between V1 and V2. If the cell value is 1, there is an edge connecting the two vertices that intersect in that cell. If the cell value is zero, then there is no edge between the two vertices intersecting in that cell. The solution is the vertex set {V2, V3, V5, V6}, which is not difficult to find but becomes more complicated with large datasets [13].

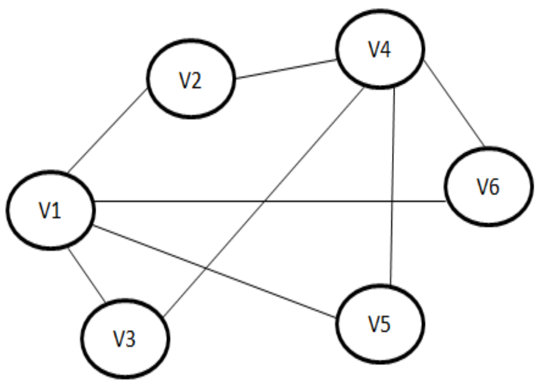

Figure 1: A maximum independent set problem example graph G with 6 vertices

Table 1: Example for maximum independent set problem

Vertices	V1	V2	V3	V4	V5	V6
V1	0	1	1	0	1	1
V2	1	0	0	1	0	0
V3	1	0	0	1	0	0
V4	0	1	1	0	1	1
V5	1	0	0	1	0	0
V6	1	0	0	1	0	0

2.2. Quantum Inspired Genetic Algorithm (QIGA)

2.2.1. Quantum Basics

Classical computers perform n operations simultaneously using n bits, while quantum computers perform 2^n operations in n qubits simultaneously [14]. This relationship is exponential. Qubits can be in the superposition of |0> and |1>, such that $\alpha|0>+\beta|1>$, where α and β are complex numbers with

$$|\alpha 2> + |\beta 2> = 1. \tag{1}$$

Here are some unitary logic gates' effects (Hadmard gate [15]).

$$H.|0> = \frac{1}{\sqrt{2}}\begin{bmatrix}1 & 1\\ 1 & -1\end{bmatrix}\begin{bmatrix}1\\0\end{bmatrix} = \frac{1}{\sqrt{2}}\begin{bmatrix}1\\1\end{bmatrix}. \tag{2}$$

$$H.|1> = \frac{1}{\sqrt{2}}\begin{bmatrix}1 & 1\\ 1 & -1\end{bmatrix}\begin{bmatrix}0\\1\end{bmatrix} = \frac{1}{\sqrt{2}}\begin{bmatrix}1\\-1\end{bmatrix}. \tag{3}$$

The X-gate, which is the NOT gate in classical computers, can have the following effect [15]:

$$X.|0> = \frac{1}{\sqrt{2}}\begin{bmatrix}0 & 1\\ 1 & 0\end{bmatrix}\begin{bmatrix}1\\0\end{bmatrix} = \frac{1}{\sqrt{2}}\begin{bmatrix}0\\1\end{bmatrix} = |1>. \tag{4}$$

$$X.|1> = \frac{1}{\sqrt{2}}\begin{bmatrix}0 & 1\\ 1 & 0\end{bmatrix}\begin{bmatrix}0\\1\end{bmatrix} = \frac{1}{\sqrt{2}}\begin{bmatrix}1\\0\end{bmatrix} = |0>. \tag{5}$$

Basic quantum logic gates are used form quantum circuits.

2.2.2. QIGA Operations

The Quantum-Inspired Genetic Algorithm (QIGA) builds its operations in a qubit concept representation [16].

2.2.3. Quantum Mutation

It defines a mutation rate to randomly pick a mutation point and change the chromosome value by replacing that randomly picked point with another randomly chosen point. This is performed as follows [17]:

$$P = \begin{bmatrix}\alpha 1\ \alpha 2\ \alpha 3\ \ldots\ \alpha q\\ \beta 1\ \beta 2\ \beta 3\ \ldots\ \beta q\end{bmatrix}. \tag{6}$$

The new chromosome will become:

$$P' = \begin{bmatrix}\alpha' 1\ \alpha 2\ \alpha 3\ \ldots\ \alpha q\\ \beta' 1\ \beta 2\ \beta 3\ \ldots\ \beta q\end{bmatrix}, \tag{7}$$

where

$$|\alpha' 2> + |\beta' 2> = 1. \tag{8}$$

2.2.4. Quantum Crossover

A crossover point is chosen randomly in two different chromosomes according to the crossover rate, and the operation is applied as follows [18]:

$$P1 = \begin{bmatrix}\alpha 1\ \alpha 2\ \alpha 3\ \ldots\ \alpha q\\ \beta 1\ \beta 2\ \beta 3\ \ldots\ \beta q\end{bmatrix}, \tag{9}$$

$$P2 = \begin{bmatrix}\alpha' 1\ \alpha' 2\ \alpha' 3\ \ldots\ \alpha' q\\ \beta' 1\ \beta' 2\ \beta' 3\ \ldots\ \beta' q\end{bmatrix}. \tag{10}$$

After applying the crossover, the chromosomes will be as the following:

$$P'1 = \begin{bmatrix}\alpha 1\ \alpha' 2\ \alpha' 3\ \ldots\ \alpha' q\\ \beta 1\ \beta' 2\ \beta' 3\ \ldots\ \beta' q\end{bmatrix}, \tag{11}$$

$$P'2 = \begin{bmatrix}\alpha' 1\ \alpha 2\ \alpha 3\ \ldots\ \alpha q\\ \beta' 1\ \beta 2\ \beta 3\ \ldots\ \beta q\end{bmatrix}. \tag{12}$$

2.2.5. Interference

The interference or rotation operator can be applied as follows [19]:

$$U(\theta)\ |\psi_t\rangle = |\psi_{t+1}\rangle = \begin{bmatrix}\cos(\theta)\alpha t - \sin(\theta)\beta t \\ \sin(\theta)\,\alpha t + \cos(\theta)\beta t\end{bmatrix}. \quad (13)$$

3. The Proposed Technique

3.1. Problem Representation

The graph of the MIS problem is composed of nodes, and the edges between the nodes link these nodes. If there is no edge, the two nodes are not connected. This problem is represented in the proposed technique as a table with a list of nodes shown in rows and the same nodes shown in columns. The intersection between the node in the column and the node in the row shows whether there is an edge connecting them or not. The 0 value is for the no existing edge between the nodes while the 1 value is for the existing edge connecting the nodes. For example, if there is a link between two nodes V1 and V2, then the intersection of column V1 and row V2 takes a value of "1" and similarly the intersection of row V1 and column V2. "0" is put otherwise. An example is presented in Table 1. This table is then represented in a 2D matrix form. Table 1 can be represented as follows.

$$\text{MIS matrix}=\begin{bmatrix}0\,1\,1\,0\,1\,1\\1\,0\,0\,1\,0\,0\\1\,0\,0\,1\,0\,0\\0\,1\,1\,0\,1\,1\\1\,0\,0\,1\,0\,0\\1\,0\,0\,1\,0\,0\end{bmatrix}. \quad (14)$$

As shown in Figure 2 and Algorithm 1, the algorithm first generates or reads the MIS matrix. For each node a global α a global β are calculated as using (15) and (16). These global α and global β are not updated later on. Each matrix has various tests that use the global α and β. After selecting the chromosomes from the MIS matrix, local α and local β are calculated for each chromosome in the population. Local α and β are calculated using the global α and β. They are updated from one generation to another. Table 2 shows an example of the global α and global β for each row as calculated using (15) and (16). Table 3 gives an example of the local α and β for each chromosome in population "p" assuming a population size of 3.

Table 2: Global α and global β example.

Row Number	Global Values
V1	α_1, β_1
V2	α_2, β_2
V3	α_3, β_3
V4	α_4, β_4
V5	α_5, β_5
V6	α_6, β_6

Table 3: Local α and local β example.

Chromosome Number	Chromosome(population, chromosome, row)		
1	V1 $\alpha_{p11}, \beta_{p11}$	V3 $\alpha_{p13}, \beta_{p13}$	V5 $\alpha_{p15}, \beta_{p15}$
2	V2 $\alpha_{p22}, \beta_{p22}$	V4 $\alpha_{p24}, \beta_{p24}$	V5 $\alpha_{p25}, \beta_{p25}$
3	V3 $\alpha_{p33}, \beta_{p33}$	V5 $\alpha_{p35}, \beta_{p35}$	V6 $\alpha_{p36}, \beta_{p36}$

$$\text{global\_}\beta=\sqrt{1.0 - global\_\beta^2}. \quad (15)$$

$$\text{global\_}\alpha=\sqrt{sumOfZeros/sumOfAllZeros}. \quad (16)$$

The matrix values are then considered individually in each row to calculate sumOfZeros and sumOfAllZeros. sumOfZeros is calculated when the intersection is zero and the adjacent node is zero, as shown in (17). If the intersection is zero, then sumOfAllZeros is calculated, as shown in (18). The fitness value can then be calculated, as illustrated in (19).

$$\text{sumOfZeros} = \sum_{k=0}^{colSize}((\text{colSize} - (\text{k})) * (k+1) \quad (17)$$

$$\text{sumOfAllZeros} = \sum_{k=0}^{colSize}((\text{colSize} - (\text{k})) \quad (18)$$

$$\text{fitnessValue} = \text{sumOfZeros / sumofAllZeros * 100} \quad (19)$$

4. Experimental Results

All the experiments are conducted in a laptop with Intel® Core™ i7 processor and 64-bit Windows 11 operating system.

A summary of the proposed technique is presented in the pseudocode of Algorithm 1. Figure 2 shows a flowchart for the proposed technique for visual clarity. The pseudocode and flowchart illustrates the QIGA process as it starts with the MIS matrix itself, then it generates the initial population and initializes the parameters including the local α and local β. Then a loop starts with applying interference operation. The fitness function is measured after the interference operation. Based on this measurement, parameters are updated and a population is selected. Crossover operation is applied on the selected population and then mutation operation is performed. The fitness function and the average fitness are then calculated. Then it updates the loop counter to go for the next iteration. These steps are performed as long as the predetermined number of iterations is not yet reached or the average fitness is less than 100. When this stopping condition becomes false, that means the MIS is solved and the algorithm ends. Table 4 lists the Genetic Algorithm (GA) parameters used to measure the technique. The experiments were performed on 200 × 200 matrices to represent a graph of 200 nodes. The experimental results were applied to three different types of matrices. The sparse that contains 80% of zeros and 20% of ones, the dense that contains 80% of ones and 20% of zeros, and the 50-50 that contains 50% of zeros and 50% of ones. The matrices are formed with randomly chosen values, but they follow each matrix type constraint.

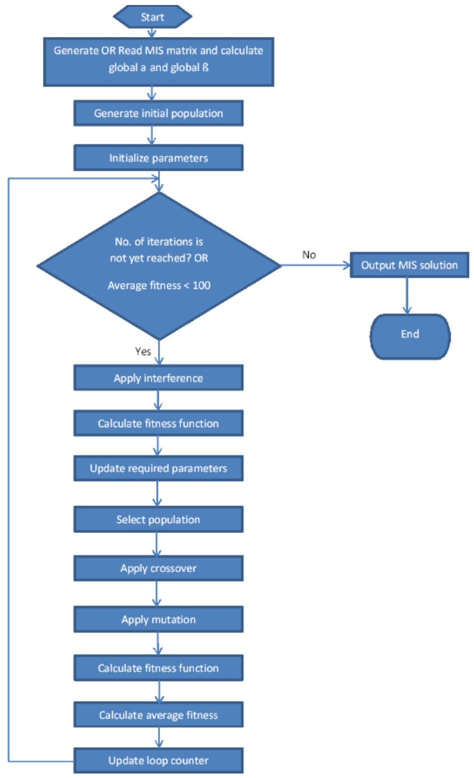

Figure 2: Flowchart summarizing the QIGA process.

Algorithm 1: Pseudocode for the proposed technique

Read MIS matrix m × n.
Calculate global α and global β for each node.
Calculate sumOfZeros and sumOfAllZeros as shown in (17) and (18).
Choose a population size.
Generate the initial population.
Give initial values to the local α and local β.
While number of iterations is not yet reached OR average fitness == 100 **do**
Apply interference and measure the fitness results using the fitness function in (19).
Update the local α and local β accordingly.
Select from the population using Roulette wheel.
Apply crossover with 90% .
Apply mutation with a mutation rate 1%
Measure the fitness results.
Calculate the average fitness.
Update loop counter.
end while
Print the MIS solution.

The maximum independent set problem was solved using GA and QIGA. Both the algorithms were tested using the same parameters. A total of 500 iterations were performed ten times to measure the average of the results.

Table 4: GA Parameters for the proposed technique

GA Parameter	Value
Population Size	500
Crossover	Single-point
Crossover Rate	%90
Mutation Rate	%1
Selection	Roulette Wheel
$\theta_{initial}$	π
$\delta\theta s$	A random number between 0 and 1

Figure 3, 4, and 5 show that QIGA achieved faster convergence than classical GA in the three matrix types. In addition to convergence, the QIGA fitness value results were higher than the classical GA results. Figure 3 shows the faster convergence and higher fitness values of sparse matrices, whereas Figure 4 shows the same successful results for balanced matrices. Figure 5 illustrates the convergence and results achieved for dense matrices.

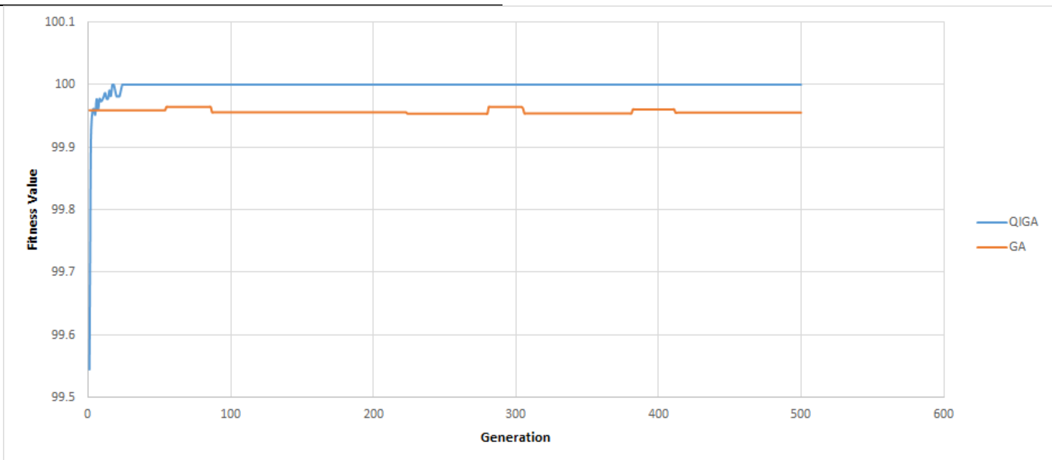

Figure 3: A comparison between QIGA and GA that shows the number of generations with the output fitness value for each generation, in case of balanced matrixes.

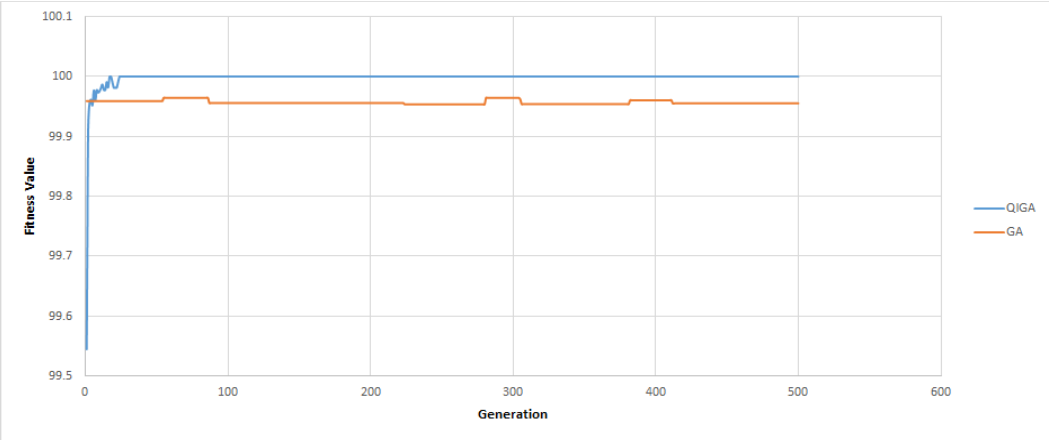

Figure 4: A comparison between QIGA and GA that shows the number of generations with the output fitness value for each generation, in case of sparse matrixes.

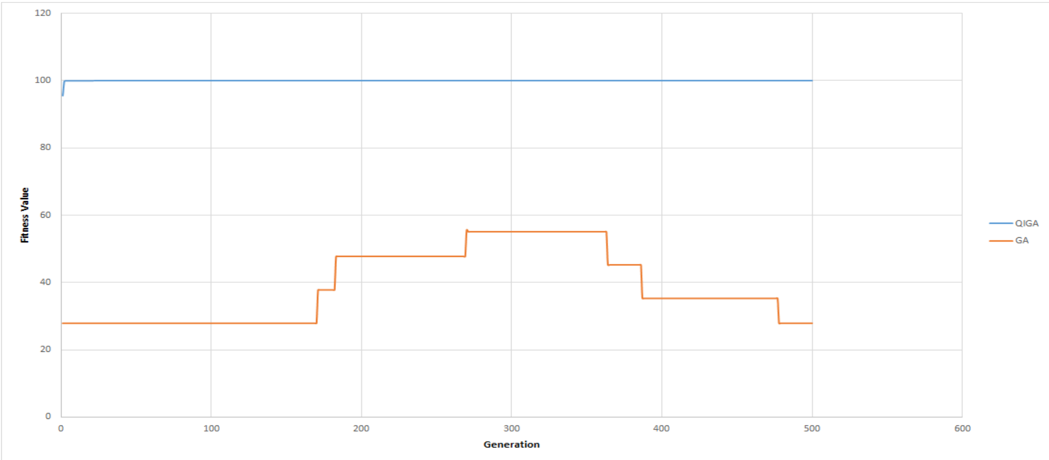

Figure 5: A comparison between QIGA and GA that shows the number of generations with the output fitness value for each generation, in case of dense matrixes.

5. Conclusion and Future Work

In this paper, a modified quantum-inspired genetic algorithm (QIGA) is proposed to solve the maximum independent problem, where quantum superposition has been used in the encoding of the chromosome to increase the size of the search space over approximately the same physical space. Quantum gates, such as crossover, mutation, and interference gates, have been used to achieve better and faster results. The experimental results have been shown for sparse, balanced, and dense test cases. The results show that QIGA performed faster and better than classical GA. It converges more rapidly and it achieved higher fitness values. This solution can be used in many domains such as software engineering to separate related requirements from unrelated requirements, time management, cost management, resource management, network design, scheduling, resource allocation, logistics, bioinformatics, mobile agents, and more [20]. Future work will be held on creating more fitness functions to give better results. Other problems will be considered to be solved using QIGA. Future application to the proposed technique can be performed on other domains.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The author received no direct funding for this research.

Funding

No financial support for this research.

Availability of data and materials

Data was generated and available upon request.

References

[1] M. Harman, Y. Jia, J. Krinke, W. B. Langdon, J. Petke, and Y. Zhang, "Search based software engineering for software product line engineering," Sep. 2014, doi: 10.1145/2648511.2648513.

[2] F. Sarro, "Search-Based Software Engineering in the Era of Modern Software Systems," IEEE 31st International Requirements Engineering Conference (RE), 2023.

[3] A. Mazzonetto, R. Z. Frantz, F. Roos-Frantz, C. Molina-Jimenez, and S. Sawicki, "A Systematic Mapping Study of Search-Based Software Engineering for Enterprise Application Integration,"

International Journal of Software Engineering and Knowledge Engineering, 2022.

[4] S. Memeti, S. Pllana, A. Binotto, J. Kołodziej, and I. Brandic, “Using meta-heuristics and machine learning for software optimization of parallel computing systems: a systematic literature review,” Computing, vol. 101, no. 8, pp. 893–936, Apr. 2018, doi: 10.1007/s00607-018-0614-9.

[5] K. Kim, M. Kim, J. Park, A. Byun, and J. Ahn, “Quantum computing dataset of maximum independent set problem on king lattice of over hundred Rydberg atoms,” Scientific Data, vol. 11, no. 1, Jan. 2024, doi: 10.1038/s41597-024-02926-9.

[6] A. W. Chaudhry, “A New Algorithm for Solving the Maximum Independent Set Problem,” Australia, May 2019.

[7] S. Kamei and S. Tixeuil, “An Asynchronous Maximum Independent Set Algorithm By Myopic Luminous Robots On Grids,” The Computer Journal, vol. 67, no. 1, pp. 57–77, Nov. 2024, doi: 10.1093/comjnl/bxac158

[8] H. Pichler, S. Wang, L. Zhou, S. Choi, and M. D. Lukin, “Quantum optimization of maximum independent set using Rydberg atom arrays,” Science, vol. 376, no. 6598, pp. 1209–1215, Jun. 2022, doi: 10.1126/science.abo6587.

[9] J. Luo and S. Ding, “Solving the k-Independent Sets Problem of Graphs by Gröbner Bases,” Open Journal of Discrete Mathematics, vol. 13, no. 03, pp. 86–94, 2023, doi: 10.4236/ojdm.2023.133008.

[10] M. Silva-Muñoz, C. Contreras-Bolton, C. Rey, and V. Parada, “Automatic generation of a hybrid algorithm for the maximum independent set problem using genetic programming,” Applied Soft Computing, vol. 144, pp. 110474–110474, Jun. 2023, doi: 10.1016/j.asoc.2023.110474.

[11] R. S. Andrist et al., “Hardness of the maximum-independent-set problem on unit-disk graphs and prospects for quantum speedups,” Physical review research, vol. 5, no. 4, Dec. 2023, doi: 10.1103/physrevresearch.5.043277.

[12] Y. Dong, A. Goldberg, A. Noe, N. Parotsidis, M. Resende, and Q. Spaen, A Local Search Algorithm for Large Maximum Weight Independent Set Problems, 30th Annual European Symposium on Algorithms (ESA 2022), 2022.

[13] Z. Wang, J. Tan, L. Zhu, and W. Huang, “Solving the Maximum Independent Set Problem based on Molecule Parallel Supercomputing,” Applied Mathematics & Information Sciences, vol. 8, no. 5, pp. 2361–2366, Sep. 2014, doi: 10.12785/amis/080531.

[14] S. Hakemi, M. Houshmand, S. A. Hosseini, and X. Zhou, “A Modified Quantum-Inspired Genetic Algorithm Using Lengthening Chromosome Size and an Adaptive Look-Up Table to Avoid Local Optima,” Axioms, vol. 12, no. 10, p. 978, Oct. 2023, doi: 10.3390/axioms12100978.

[15] P. S. Menon and M. Ritwik, “A Comprehensive but not Complicated Survey on Quantum Computing,” IERI Procedia, vol. 10, pp. 144–152, 2014, doi: 10.1016/j.ieri.2014.09.069.

[16] A. M. Mohammed, N. A. Elhefnawy, M. M. El-Sherbiny, and M. M. Hadhoud, “Quantum crossover-based quantum genetic algorithm for solving non-linear programming,” International Conference on Informatics and Systems, May 2012.

[17] N. Toronto and D. Ventura, “Learning Quantum Operators From Quantum State Pairs,” IEEE International Conference on Evolutionary Computation, vol. 3103, pp. 2607–2612, Sep. 2006, doi: 10.1109/cec.2006.1688634.

[18] M. Ridha, “Reversible Logic Synthesis Methodologies with Application to Quantum Computing,” Springer International Publishing, 2015. doi: 0.1007/978-3-319-23479-3.

[19] H. Hussein, A. Younes, and W. Abdelmoez, “Quantum-Inspired Genetic Algorithm for Solving the Test Suite Minimization Problem,” WSEAS TRANSACTIONS ON COMPUTERS, vol. 19, pp. 143–155, Aug. 2020, doi: 10.37394/23205.2020.19.20.

[20] A. M. Salman and A. S. Al-Jilawi, “Applications of maximum independent set,” AIP conference proceedings, Jan. 2022, doi: 10.1063/5.0093375.

DR. HAGER HUSSEIN had her PhD degree in computer science from Alexandria University in 2021. She received her MSc degree in software engineering from CCIT in AASTMT in 2012.

She is an assistant professor at the Software Engineering Department, College of Computing and Information Technology (CCIT), Arab Academy for Science, Technology, and Maritime Transport (AASTMT). Her research interests include quantum algorithms and software engineering.

Dr. Hager Hussein had published a research paper with Cogent Engineering journal in Taylor & Francis group. She acted as a Scientific Committee member in the 34th International Conference on Computer Theory and Applications.

Special Issue on Multidisciplinary Sciences & Advanced Technology 2025

JENRS

Received: 06 June 2025, Revised: 14 June 2025, Accepted: 07 August 2025, Online: 20 August 2025

DOI: https://dx.doi.org/10.55708/js0408004

Magnetic AI Explainability: Retrofit Agents for Post-Hoc Transparency in Deployed Machine-Learning Systems

Maikel Leon*

Department of Business Technology, Miami Herbert Business School, University of Miami, Miami, Florida, USA

*Corresponding author. Email: mleon@miami.edu

ABSTRACT: Artificial intelligence already influences credit allocation, medical diagnosis, and staff recruitment, yet most deployed models remain opaque to decision makers, regulators, and the citizens they affect. A new wave of transparency mandates across multiple jurisdictions will soon require organizations to justify automated decisions without disrupting tightly coupled production pipelines that have evolved over the years. We advance a conceptual proposal to address this tension: the magnetic AI agent. This external, attachable software layer learns a faithful surrogate of any target model, delivering audience-tailored explanations on demand. The paper first synthesizes fragmented scholarship on post-hoc explainability, sociotechnical alignment, and model governance, revealing an unmet need for lightweight retrofits that minimize downtime. It then creates a basic framework based on design principles, explaining methods for data collection, ongoing learning processes, and user-friendly explanation tools. A plan for evaluation lists both numerical and descriptive measures, including how closely a model matches reality and how much extra time it takes, as well as the mental effort required and how well policies work, which users can adjust for different fields like credit scoring, medical imaging, and predictive maintenance. Overall, the work contributes a roadmap for upgrading the installed base of black-box systems while aligning with emergent regulatory frameworks and ethical guidelines for trustworthy AI.

KEYWORDS: Magnetic AI, Explainable Artificial Intelligence, Agentic AI, Retrofit Transparency, Design-Science Research, Policy Compliance.

1. Introduction

Artificial Intelligence (AI) systems that once resided in research labs now power high-stakes finance, health care, logistics, national security, and public administration decisions. These models deliver unprecedented speed and predictive accuracy, yet they rarely reveal the internal logic that drives their outputs. This asymmetry between performance and interpretability poses reputational, operational, and legal risks for organizations that rely on opaque algorithms. Recent incidents—such as biased credit approvals, flawed recidivism predictions, and inconsistent medical triage decisions—demonstrate how opacity can erode stakeholder trust and invite regulatory scrutiny [1].

Last century, AI research surged on the back of expert systems, decision trees, and the early "neural nets" revival. Success was measured almost entirely by how precisely these models could predict outcomes, whether diagnosing disease, flagging credit risk, or recognizing handwritten digits. Researchers fine-tuned rule bases or tweaked hidden-layer weights to squeeze out a few extra percentage points of accuracy, and industry adopters celebrated any gains that outperformed human benchmarks. Yet this accuracy-first mindset treated the models as opaque black boxes: engineers rarely asked why a particular rule fired or a neuron activated, and users seldom demanded a justification. As a result, explainability remained an afterthought; the momentum and funding of the era were channeled into sharpening predictive performance, not into opening the "black box" so stakeholders could trust and understand the reasoning inside it.

Across major jurisdictions, regulation is converging on a common requirement that AI systems be explainable: the European Union's AI Act, recent U.S. executive directives, and China's updated generative-AI rules all mandate that high-impact models provide meaningful information about how they reach their outputs. This amounts to an emerging right for everyday users to demand clear, human-readable reasons for automated predictions or decisions, even when those decisions come from complex neural networks. Anticipating audits, fines, and reputational risks, companies are building explanation layers into their products—dashboards that visualize feature contributions, surrogate models that translate deep-learning logic into plain language, and customer portals that show "what-if" scenarios—because meeting this new transparency baseline is becoming less a nice-to-have and more a competitive necessity.

Societal expectations for transparency have accelerated. Policymakers on both sides of the Atlantic have enacted or proposed frameworks that place the burden of justification on automated decision-makers. The European Union's AI Act, the United Kingdom's Algorithmic Transparency Standard, and various U.S. proposals such as the Algorithmic Accountability Act collectively signal a shift from self-regulation to explicit accountability. These initiatives often focus on two intertwined requirements: the ability to generate human-understandable explanations and the capacity to audit models throughout their life cycle. Organizations, therefore, face the dual challenge of upgrading legacy AI assets and operationalizing governance processes at scale.

Despite rapid advances in post-hoc interpretability techniques, most production environments cannot easily accommodate invasive code changes, extensive retraining cycles, or computational overhead that might jeopardize service-level agreements. Enterprise Machine Learning pipelines typically integrate proprietary libraries, tightly coupled microservices, and third-party APIs that preclude direct intervention. A non-disruptive alternative is to attach an explanatory agent to the outside of an existing pipeline, much like a magnetic device that snaps onto the surface of a machine without changing its internal workings. We label this solution the magnetic AI agent. The magnetic analogy underscores three salient properties: passive attachment, minimal friction, and continuous real-time learning [2].

While the concept of attaching post-hoc interpretability layers has precedent in techniques such as shadow models, knowledge distillation, and wrapper-based surrogates, the magnetic AI agent diverges in critical ways. Unlike shadow models that mimic predictions for evaluation purposes or distillation methods that compress complex models into simpler ones, the magnetic agent is designed to operate continuously alongside the original model without approximation or replacement [3]. Its emphasis is not only on interpretability but also on modular deployment, governance integration, and lifecycle adaptability in real-world production systems. The magnetic metaphor is not a rhetorical flourish—it reflects an architectural philosophy: to enable passive but intelligent observability without disrupting the core model's functioning or retraining requirements.

The remainder of the paper deepens the conceptual foundation, formalizes the design space, and proposes an actionable evaluation pathway for magnetic AI. While empirical results are not presented here, this absence is by design: the work is intended as a conceptual proposal that lays the groundwork for future implementation and experimentation. Its primary aim is to contribute a structured framework, design rationale, and deployment blueprint that researchers and practitioners can build upon. First, Section 2 surveys the multidisciplinary literature on explainable AI and model-agnostic wrappers, identifying persistent gaps that motivate a new approach. Section 3 introduces the conceptual framework that positions the retrofit agent within sociotechnological constraints and elaborates design principles, reference architecture, and governance interfaces. Section 4 describes a design-science research strategy and methodological considerations for constructing and refining the artifact. Section 5 details an evaluation blueprint that organizations can replicate or adapt in their domains. Section 6 discusses operational, ethical, and societal implications, mapping the proposal onto current regulatory trends. Section 7 concludes by summarizing contributions, delineating limitations, and articulating a future research agenda that includes full-scale prototypes, multimodal extensions, and integration with next-generation foundation models.

2. Related Work

Research on explainability spans multiple disciplines, each supplying partial answers to how automated systems should justify their outputs. Algorithmic contributions range from ante-hoc transparent models to post-hoc attribution methods such as LIME, SHAP, and integrated gradients to compression techniques that create interpretable surrogates. Human-computer interaction studies examine the cognitive load of different explanation formats, user mental-model accuracy, and the conditions under which explanations raise or erode calibrated trust. Work in organizational behavior documents how power dynamics, siloed incentives, and technical debt shape whether explanations are acted upon or ignored. Legal scholarship and policy analyses frame transparency as a right, exploring liability, due-process entitlements, and the evolving notion of algorithmic accountability [4].

This review weaves the strands together, pinpointing where they fall short and how they complement one another. Algorithmic methods often optimize fidelity or sparsity but rarely address maintenance overhead once a model is in production. HCI experiments illuminate user comprehension in laboratory settings, yet evidence remains sparse on sustained behavior change in real workflows. Organizational case studies highlight governance bottlenecks but seldom tie them to concrete design artifacts. Legal work identifies transparency duties but leaves practitioners with little guidance on technical implementation. Magnetic AI draws on the strengths of each field while addressing their gaps: a passive attachment strategy respects intellectual-property boundaries emphasized in law, continuous fidelity auditing answers organizational concerns about drift and technical debt, and explanation pluralism accommodates the heterogeneous user needs documented in HCI research [5].

Key takeaways that inform the design are as follows:

- Algorithmic insight: incremental surrogates balance fidelity with latency, enabling explanations at line speed without altering the primary model. They learn from a sliding window of recent requests, refresh continuously without full retraining, and respect the intellectual-property boundaries of closed models, making them suitable for third-party APIs and in-house stacks.
- HCI insight: multiple discourse formats—ranked feature tables, layered saliency maps, natural-language counterfactual narratives, and compliance-ready audit summaries—are necessary because data scientists, end users, and regulators each privilege different cues. Adaptive rendering lets the same evidence flow into analyst dashboards, tooltips for consumers, or machine-readable JSON for supervisory authorities.
- Organizational insight: modular deployment decouples the four layers—interception, surrogate learning, explanation rendering, and fidelity auditing—so firms can adopt only the components they lack. This bolt-on architecture avoids rewriting brittle legacy code,

shortens change-management cycles, and reduces the blast radius of defects to a single microservice rather than the full model pipeline.

- Legal insight: persistent audit logs, role-based explanation access, and optional differential-privacy noise satisfy both transparency duties and data-protection rules. The same artifacts can populate internal risk registers, respond to freedom-of-information requests, or demonstrate compliance during external audits, aligning technical controls with emerging statutes such as the EU AI Act and national consumer-protection guidelines [6].

By fusing these lessons, magnetic AI offers a coherent blueprint that advances beyond silo-specific approaches toward an integrated, production-ready solution for trustworthy machine learning.

2.1. Post-Hoc Explainable AI

Early work on interpretability concentrated on "glass-box" algorithms—decision trees, linear or logistic regressions, and simple rule lists—whose parameters and splits can be read like prose. As deep learning's opaque layers dominated predictive accuracy, researchers shifted toward post-hoc techniques that wrap explanations around otherwise black-box models [7].

The most influential of these are LIME and SHAP. Both build local surrogate models that mimic the original model's behavior near a single instance, then report feature attributions: LIME perturbs inputs and fits a sparse linear model, whereas SHAP samples coalitions of features to compute Shapley values that satisfy additivity and consistency. Their appeal lies in domain-agnostic deployment—data scientists can drop in a few lines of code and hand users a ranked list of "which variables mattered most"—yet the price is high computational overhead, sensitivity to sampling noise, and explanations that change when the same point is probed twice [8].

Beyond LIME and SHAP, gradient-based saliency maps track the partial derivatives of a convolutional network to highlight the pixels that nudge an image score upward or downward; attention visualizations in transformer models color the tokens that capture a language model's gaze; counterfactual methods search the input space for the most minor tweak that flips the prediction, offering an actionable "what would need to change?"; and prototype- or example-based explanations surface representative cases that anchor abstract probability scores in concrete, human-readable examples. Each broadens the explanatory toolbox, yet each inherits its drawbacks: saliency maps blur under adversarial noise, attention plots do not always align with causal importance, counterfactuals become infeasible in high-dimensional data, and prototype selection can reinforce majority-class bias [9].

Across the board, explanation strength often comes at the cost of latency, stability, or hardware resources. Empirical studies still debate whether richer explanations meaningfully boost user trust or downstream decision quality, highlighting an unsolved interpretability-accuracy-usability triangle.

2.2. Wrapper and Surrogate Paradigms

Building a simpler model that imitates a complex one is hardly new. In the 1980s, credit bureaus built "shadow" logistic regressions to track the decisions of proprietary loan scoring engines, and in the 1990s, speech-recognition teams used teacher–student pairs to shrink large hidden-Markov networks so they could run on low-power chips. These ideas matured into what is now called knowledge distillation, where an extensive teacher network produces soft targets—probability distributions rather than hard labels—that guide a smaller student network. The result is a faster, lighter model that often matches the teacher's top-line accuracy but may blur fine-grained decision boundaries, especially in rare or ambiguous cases.

Modern workflows try to close that gap by performing distillation continuously. An online student receives a stream of teacher outputs and updates its weights on the fly, or it joins a replay buffer that mixes new observations with old exemplars to resist catastrophic forgetting. Continual-learning variants add regularizers that anchor key teacher activations so the student does not drift when the data distribution shifts. Yet experiments on non-stationary benchmarks show that even these advanced students struggle with concept drift and are highly sensitive to mislabeled or adversarially perturbed examples [10].

A parallel line of work forgoes access to internal weights altogether. Instead, engineers wrap the black-box service with a data interceptor that logs inputs and outputs, then train a surrogate, often a decision tree or gradient-boosted ensemble, purely from those pairs. This wrapper strategy sidesteps intellectual-property barriers and can be swapped before any commercial API. Still, it introduces fresh privacy challenges: synthetic or cached query data must be stored outside the original security perimeter, and reconstruction attacks can expose sensitive attributes if the wrapper is breached [11].

Taken together, today's surrogate models fall into two camps. Static snapshots captured once during development grow stale as the real world evolves, while dynamic surrogates that retrain or distill online demand constant monitoring, a computation budget, and careful privacy safeguards. Neither camp fully resolves the tension between efficiency, fidelity, and maintainability in production environments that change by the hour.

2.3. Regulatory and Business Context

Across regions, lawmakers and standard-setters are locking into a shared vocabulary—transparency, accountability, fairness, and meaningful human oversight—and turning it into binding or quasi-binding rules. In Europe, the AI Act labels credit scoring, hiring, medical diagnosis, and other "high-risk" applications. It forces them to generate understandable explanations, document data provenance, and pass third-party conformity assessments before entering the market.

In the United States, the Federal Trade Commission, Consumer Financial Protection Bureau, Department of Justice, and other agencies have warned that undisclosed bias, dark-pattern interfaces, or the sale of inscrutable models can trigger enforcement actions under existing consumer-protection and civil-rights statutes. At the same time, the White House blueprint for an AI Bill of Rights and the NIST AI Risk-Management Framework give regulators a benchmark for what "reasonable" governance should look like. China's updated Interim Measures on generative AI require providers to watermark outputs, publish model cards, and supply "interpretive" summaries on demand; Canada's forthcoming AI and Data Act mandates impact assessments and real-time monitoring; Brazil and India are drafting parallel bills; and the G7's Hiroshima Process is pressing multinationals to align with these norms wherever they operate.

Industry bodies reinforce the trend: the Partnership on AI, the OECD, the ISO/IEC 42001 management-system standard, and voluntary procurement checklists now ask vendors to show audit logs, bias tests, and plain-language explanations as a condition of sale. Non-compliance can mean multimillion-euro fines, exclusion from public-sector tenders, investor divestment, and reputational damage that stalls digital-transformation roadmaps. Yet most enterprises run on entrenched code bases, brittle data pipelines, and overlapping legacy models; ripping and replacing them is rarely feasible. This clash between external pressure and internal technical debt drives demand for retrofit solutions—lightweight layers that bolt onto existing systems, capture inputs and outputs, monitor drift, and surface user-friendly explanations—so firms can satisfy new governance obligations without rebuilding their entire machine-learning stack [12].

2.4. Gap Analysis

Table 1 contrasts prevailing approaches against operational requirements and spotlights the unresolved disconnect between research prototypes and production realities. While the literature offers algorithmic sophistication, it rarely addresses day-two concerns such as deployment pipelines, monitoring infrastructure, and heterogeneous stakeholder needs. The magnetic AI proposal aims to bridge this gap by integrating passive attachment, continuous fidelity auditing, and human-centered explanation delivery into a unified artifact.

There seems to be a clear trade-off pattern: methods that are easiest to bolt onto any model (LIME, SHAP, Anchors) suffer from high inference latency or heavy sampling, while techniques that are fast enough for production (knowledge-distilled surrogates, ante-hoc interpretable models) often under-fit or drift from the source model without constant retraining. Vision-specific tools like Grad-CAM are efficient but narrow in scope, and counterfactual or prototype-based approaches provide the most human-friendly "what-if" stories yet demand large compute budgets and carefully curated instance libraries [13].

In short, no single technique simultaneously delivers low latency, high fidelity, and broad stakeholder usability. This operational gap motivates a hybrid solution, such as the proposed magnetic AI artifact, that couples passive attachment for real-time capture with continuous fidelity auditing and layered explanation modes tuned to different audiences.

Table 1: Operational gap between explainability techniques and production requirements

Approach	Strengths	Limitations
LIME / SHAP	Model-agnostic; easy to add	High latency in production; explanations local
Knowledge distillation	Compact, fast surrogates	Needs labelled outputs; surrogate drift
Counterfactuals	Actionable "what-if" paths	Heavy compute; plausibility issues
Magnetic AI (proposed)	Passive attachment; continuous learning	Concept stage; governance pending
Integrated Gradients	Faithful to deep nets; low single-call overhead	Requires differentiable model; noisy for saturated neurons
Grad-CAM	Intuitive heat-maps for vision CNNs; real-time on GPU	Vision-only; coarse spatial resolution
Anchors	Sparse, high-precision rules; human-readable	Sampling-intensive; struggles with high-dimensional mixes
Partial Dependence / ICE	Global feature-effect trends; offline computation	Assumes feature independence; stale in changing data
Prototype & Criticism	Example-based, domain-relatable explanations	Needs large representative set; weak in very sparse spaces
Ante-hoc interpretable mdl.	Transparency built-in (e.g., GAMs, monotonic GBMs); low latency	May under-fit complex tasks; restricted model choices

3. Magnetic AI Conceptual Framework

The magnetic AI framework delineates the core constructs, operational boundaries, and design guidelines necessary to retrofit explainability into black-box systems. Building on sociotechnical theory, the framework positions the agent as an intermediary that negotiates between opaque algorithms and heterogeneous human audiences [14].

3.1. Definition and Scope

A magnetic AI agent functions as a sidecar or proxy service that eavesdrops on every request–response pair flowing to and from a production model. As each new interaction arrives, the agent adds it to a sliding window buffer—say the most recent ten thousand cases—and updates an online surrogate such as an incremental gradient-boosted tree or a compact transformer fine-tuned with parameter-efficient adapters. This continual refresh allows the surrogate to track concept drift without incurring the full cost of retraining. Because the agent learns only from observable inputs and outputs, it can attach to black-box APIs, commercial SaaS endpoints, or legacy binaries without source code or training data. Once the surrogate reaches a configurable fidelity threshold, the agent can emit different explanation "dialects" on demand: concise ranked feature lists for customer-service representatives, multi-layer saliency maps for data scientists, counterfactual recourse suggestions for end users, or timestamped audit reports that regulators can archive. A governance layer encrypts the buffered data, records model-to-surrogate agreement scores, triggers alerts when fidelity degrades, and exposes REST or gRPC endpoints so downstream dashboards can pull explanations in real time [15].

Deployment is lightweight—often a Docker container or Kubernetes sidecar—so platform teams can roll it out with minimal changes to existing pipelines. Because the agent never touches proprietary weights or training sets, intellectual-property boundaries remain intact, and privacy can be reinforced with hashing or differential-privacy noise in the captured feature vectors. This combination of passive attachment, incremental learning, and audience-specific explanation formats positions magnetic AI agents as a practical retrofit for organizations that must meet new transparency rules without redesigning their entire machine-learning stack.

3.2. Design Principles

Design principles serve as invariant heuristics that guide implementation choices across contexts:

- Plug-and-play attachment via standardized data taps that conform to common message-queue or REST interfaces, minimizing engineering overhead.
- Model and domain agnosticism that enables deployment across tabular, image, NLP, time-series, and multimodal pipelines.
- Continuous auditing that monitors surrogate fidelity over time using drifting-window statistical tests and triggers automatic recalibration when thresholds are breached.
- Explanation pluralism that tailors output modalities to stakeholder expertise, regulatory requirements, and situational constraints, thereby enhancing relevance and comprehension.
- Privacy-preserving learning that supports on-device distillation, differential privacy budgets, and federated aggregation when data sovereignty is paramount.

3.3. Reference Architecture

The architecture is divided into four loosely coupled layers. The data interception layer attaches to message brokers, REST gateways, or in-process hooks to duplicate each input–output pair with millisecond-level delay. Captured data is written to an encrypted sliding-window buffer sized to the latency budget. The surrogate learning layer ingests this stream and updates an incremental model such as an online gradient-boosted tree, streaming k-nearest neighbors, or a partial-fit neural network.

A fading factor emphasizes recent samples so the surrogate can track concept drift without unbounded memory growth. The explanation rendering layer queries the current surrogate to extract local and global importance signals, then converts them into human-readable artifacts by combining a template engine with natural-language generation. Supported formats include ranked feature lists, layered saliency maps, counterfactual recourse narratives, and compliance-oriented audit summaries.

The fidelity auditing layer compares surrogate outputs with the target model on a hold-back stream slice, records agreement statistics, raises drift alerts when error thresholds are exceeded, and exposes metrics to governance dashboards through an HTTP endpoint. The modular design permits selective adoption, so an organization may activate only the components that fill existing gaps:

- Data interception choices: sidecar proxy, service-mesh filter, or Kafka consumer
- Surrogate learning supports pluggable incremental algorithms and optional ensembling
- Explanation rendering exports Markdown, JSON, PDF, or SVG artefacts for integration with existing portals
- Fidelity auditing pushes metrics to Prometheus or OpenTelemetry and routes alerts to Slack or PagerDuty

Figure 1 illustrates the magnetic AI agent operating across four loosely coupled layers.

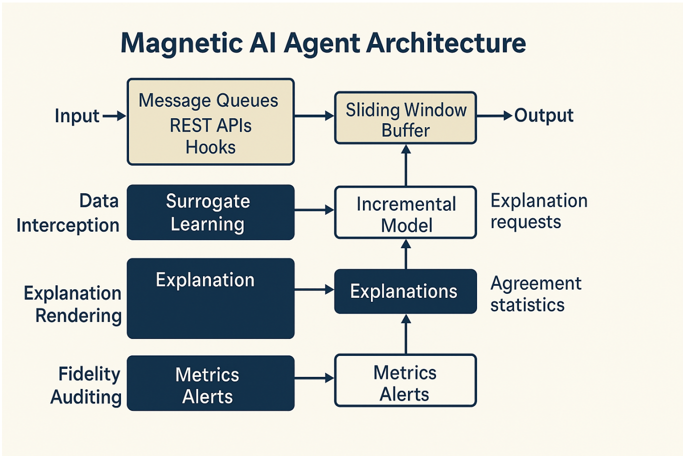

Figure 1: Magnetic AI Reference Architecture: A four-layer system that retrofits explainability into black-box models using passive data interception, online surrogate learning, audience-specific rendering, and continuous fidelity auditing.

4. Research Design and Methodology

Table 2 summarizes the guiding questions. Rigorous methodological scaffolding is essential to transform a design idea into an evaluable artifact. We adopt a design-science paradigm that iteratively synthesizes knowledge through constructing and assessing purposeful artifacts.

4.1. Artifact Construction Strategy

The construction strategy unfolds in three stages. Stage 1 employs synthetic benchmarks such as tabular classification tasks from the UCI repository to validate algorithmic viability under controlled conditions. Stage 2 transitions to semirealistic testbeds—for example, open medical-imaging datasets—where data sensitivity approximates production scenarios. Stage 3 involves shadow deployments within partner organizations, embedding the agent in parallel with live systems to observe operational impacts without influencing decision outcomes. Each stage employs a build-measure-learn loop, refining data-tap APIs, surrogate hyperparameters, and explanation formats based on empirical feedback.

Table 2: Guiding questions for magnetic AI research design

Research question	Section
What functions must a retrofit agent perform to satisfy transparency mandates?	Framework
How can fidelity be maintained as underlying models drift?	Methodology
Which usability metrics best capture explanation quality across domains?	Evaluation
What governance processes are necessary to embed magnetic agents responsibly?	Discussion

4.2. Proposed Evaluation Metrics

Comprehensive evaluation encompasses technical fidelity, human factors, and organizational fit.

- Surrogate fidelity quantified by macro-averaged agreement, calibration error, and local explanation stability across perturbed inputs.
- Latency overhead measured as the delta between baseline prediction response time and pipeline response time with the agent attached, segmented by cold-start and steady-state conditions [16].
- Cognitive burden assessed via the NASA-TLX workload instrument and validated comprehension quizzes administered to diverse user cohorts.
- Policy sufficiency mapped to ISO-based checklists and jurisdiction-specific compliance rubrics, with binary pass/fail indicators and narrative justifications.
- Maintenance complexity captured through engineer-reported setup time, mean time to detection, and time to repair when drift alarms are triggered.

5. Evaluation Blueprint

A structured evaluation helps an organization transition from proof of concept to full roll-out without losing sight of risk, cost, or stakeholder value. Below, we will break the adoption into four incremental phases, each with its entry criteria, success indicators, and decision gates. Escalation to the next phase occurs only when the previous one meets predefined thresholds, reducing the likelihood of expensive rework later in the project. As shown in Figure 2, the evaluation progresses through four structured phases.

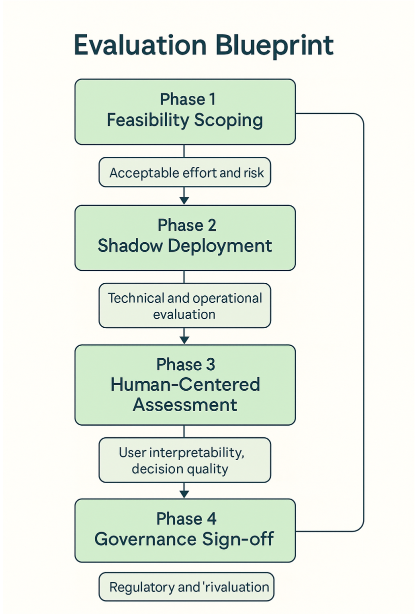

Figure 2: Evaluation Blueprint: A four-phase process guiding the deployment of magnetic AI agents from feasibility scoping to governance sign-off.

5.1. Phase 1: Feasibility Scoping

The objective is to decide whether a magnetic agent can attach to existing systems with acceptable effort and risk. A cross-functional team—product owners, data engineers, legal counsel, and compliance officers—maps the technical and organizational landscape before a single line of code is written.

- Catalog candidate models, including version numbers, input modalities, and traffic volumes.
- Identify data-tap points such as message queues, microservice gateways, or in-process hooks.
- Segment explanation audiences: internal analysts, external customers, and regulators.
- Run a one-week pilot that captures a small sample of input–output pairs to confirm data visibility, latency overhead, and encryption requirements.
- Document legal constraints on data copying, retention, and cross-border transfer.

A green light to Phase 2 requires evidence that data taps are technically feasible, that no show-stopper legal barriers exist, and that the surrogate can be trained within the latency budget on a representative sample.

5.2. Phase 2: Shadow Deployment

The magnetic agent now runs parallel with the production model but remains invisible to end users. The aim is to measure technical fidelity and operational impact without altering business outcomes.

- Stream live input–output pairs to the surrogate and store them in a ring buffer sized to the retention policy.
- Generate explanations, drift graphs, confusion matrices, and saliency heat maps; push them to a read-only dashboard.
- Track surrogate-to-model agreement, memory growth, and compute cost hourly.
- Stress-test the agent under peak traffic loads to verify scaling rules and auto-healing scripts [17].
- Perform red-team exercises to probe for model inversion and data leakage vectors.

Promotion to Phase 3 requires that fidelity metrics reach a predefined threshold, that resource consumption stay within budget, and that no critical security vulnerabilities remain open.

5.3. Phase 3: Human-Centered Assessment

With technical soundness established, the focus shifts to human interpretability and decision quality. Explanations are shown to real users in a sandbox or pilot workflow.

- Recruit subject-matter experts—credit underwriters, fraud analysts, radiologists—for structured review sessions.
- Present a stratified sample of explanations, including edge-case and adversarial examples.
- Collect quantitative scores using metrics from Section 4 and qualitative feedback on clarity, usefulness, and domain language.
- Run A/B trials where some users receive explanations and others do not, measuring changes in decision time, error rate, and confidence calibration.
- Iterate on templates, terminology, and granularity until user-acceptance criteria are met.

Advancement to Phase 4 depends on demonstrable gains in user understanding or workflow efficiency and the absence of new cognitive or fairness concerns.

5.4. Phase 4: Governance Sign-off

The final checkpoint aligns the deployment with corporate risk appetite and external regulatory obligations. A multidisciplinary committee reviews evidence accumulated in earlier phases.

- Audit logs: fidelity trends, drift alerts, red-team findings, and remediation actions.
- Human-factor reports: focus-group transcripts, A/B test statistics, and user-acceptance sign-offs.
- Compliance dossier: data-protection impact assessment, model card, explanation samples mapped to regulatory articles.
- Operational playbook: on-call rotation, retraining schedule, rollback triggers, and key performance indicators.

Once approved, the magnetic agent's explanation endpoints are activated in consumer portals, internal tools, or regulator-facing audit trails. Post-deployment, a quarterly review loop checks for concept drift, escalating to retraining or policy revision when thresholds are breached.

6. Discussion

The empirical and design insights above converge on a central theme: explainability is no longer a research luxury but an operational requirement that influences competitive advantage, regulatory posture, and societal trust. Deploying a magnetic agent transforms transparency from an expensive, one-off retrofit into a continuous service layer that scales with business growth [18]. This shift prompts decision makers to treat explainability as a cross-cutting capability, like security or observability, rather than a bolt-on feature. It carries strategic implications at three levels.

First, at the enterprise level, magnetic AI offers a cost–benefit inflection point. Faster compliance approvals, reduced litigation risk, and new value propositions, such as premium data-lineage services for high-stakes customers, offset the marginal expense of streaming surrogates and auditing dashboards. Firms adopting early may shape industry standards and lock in reputational capital that late movers struggle to match.

Second, at the ecosystem level, widespread passive-attachment architectures could generate large, anonymized corpora of model–surrogate disagreement events. These data could be shared under federated learning or secure multiparty protocols, catalyzing sector-wide benchmarks for robustness and enabling collaborative defense against adversarial attacks and systemic bias.

Third, granular yet comprehensible explanations at the societal level recalibrate the power balance between institutions and individuals. Users gain procedural recourse, auditors gain verifiable artifacts, and policymakers gain a practical blueprint for enforcement. The trade-off, however, is a thicker layer of governance overhead and an expanded attack surface that demands ongoing vigilance [19].

Against this backdrop, executive sponsors should treat magnetic AI deployment as a phased capability-maturity journey. Early milestones include establishing a data-tap inventory, codifying explanation-quality metrics, and funding interdisciplinary training programs so that engineers, risk officers, and product managers share a common vocabulary. Later stages focus on automating drift remediation, integrating feedback loops into agile release cycles, and participating in cross-industry consortia that set open standards for explanation fidelity and fairness. Organizations can navigate tightening regulations and rising public expectations by internalizing these priorities without sacrificing innovation velocity [20].

6.1. Prototype Model Demonstration

To illustrate the feasibility and behavior of the magnetic AI agent in a controlled environment, we implemented a toy model scenario. This lightweight empirical demonstration, while not intended as a comprehensive validation, serves to ground the concept in observable mechanics and provide an early proof of plausibility.

We used the classic Iris dataset and trained a black-box model using a random forest classifier. The magnetic agent was simulated as a proxy service that intercepted each input–output interaction and updated an online logistic regression model as its surrogate. The surrogate was constrained to observe only the request–response pairs, without access to feature importances, decision paths, or model internals.

Explanations were then generated by querying the logistic surrogate for each prediction and mapping the coefficients to ranked features. A fidelity audit compared surrogate predictions to the random forest decisions over a sliding window of 150 samples. Surrogate agreement stabilized at approximately 92%, and drift detection flagged one period where surrogate performance dropped due to a change in the class distribution, prompting automatic retraining.

Latency benchmarks were also recorded. On a commodity laptop (2.4 GHz, 8 GB RAM), average inference time per sample for the surrogate was under 3 milliseconds, including update and explanation rendering. This suggests that passive learning and auditing are feasible in near-real-time scenarios with moderate throughput. The latency–fidelity trade-off was observed to be tunable: larger sliding windows and ensemble surrogates marginally improved fidelity (up to 95%) but increased inference latency to 7–9 milliseconds per sample.

Input and output interfaces were defined as JSON over HTTP, simulating a REST-based production API. The surrogate processed flattened tabular features of fixed-length float vectors (4 dimensions for Iris), and the agent operated asynchronously in a sidecar thread. All components were implemented in Python using scikit-learn, Flask, and asyncio.

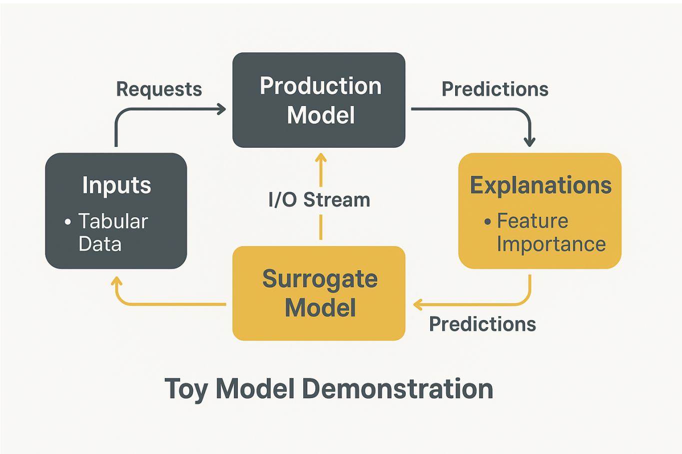

Figure 3: Toy Model Setup: The magnetic AI agent observes request–response pairs from a black-box random forest classifier trained on the Iris dataset. It trains a surrogate logistic regression model in real time, generates explanations, and audits fidelity in a sliding window.

6.2. Operational Considerations

Deploying a magnetic-AI layer replaces the usual pain of rewriting core models with the more manageable task of tapping live data streams. In companies that route traffic through Kafka, Kinesis, or a service-mesh sidecar, engineers can expose the request and response topics, spin up an agent container, and reach baseline fidelity in a morning.

By contrast, firms that still rely on tightly coupled middleware or batch ETL pipelines have to interpose a shim: a wrapper script that logs function calls or a lightweight message broker that mirrors production payloads without breaking the original code path. Once the tap is in place, the dominant cost moves from development time to compute cycles. Surrogate training scales almost linearly with input volume, so high-traffic applications—think personalized advertising or fraud detection at the millisecond level—can drive up cloud bills. Most teams blunt the cost curve by batching updates, down-sampling low-value events, or letting the agent burst to spot GPUs only during load spikes. Role clarity is essential to keep the system maintainable.

Data engineers own the interception code and service orchestration; data scientists tune the surrogate's learning rate, curate explanation templates, and validate fidelity thresholds; and compliance officers monitor the audit metrics, approve threshold changes, and archive drift reports for regulators. Without that three-way handshake, incremental tweaks in one area can silently break obligations in another, turning a retrofit to reduce risk into a new source of operational debt [21].

6.3. Ethical and Societal Dimensions

Agentic explainability shifts control from the system to the individual: a user can probe why their loan application was declined, inspect which pixels persuaded a vision model to flag an X-ray as malignant, or test what-if scenarios to see how a recommendation would change if inputs were different. This new transparency fosters autonomy and contestability and cracks open fresh attack surfaces.

Detailed feature-importance scores can reveal sensitive correlations that a company regards as trade secrets; if queried repeatedly, counterfactual examples let adversaries approximate the decision boundary and reconstruct private training data. To balance empowerment with protection, platform teams typically combine three defenses: rate-limiting caps the number of explanation calls per user or session, and throttling brute-force inversion attempts. Second, tiered access gates fine-grained explanation modes—local SHAP values, raw probability vectors, and full counterfactual paths—behind roles, entitlements, or paywalls, so casual consumers see only high-level summaries.

At the same time, regulators or auditors can request deeper details under non-disclosure constraints. Third, an adversarial-testing regime injects synthetic queries that mimic hostile behavior and flags the agent if leakage thresholds are exceeded.

Technical safeguards alone are insufficient because the audience's ability to parse explanatory artifacts is uneven. A compliance officer versed in statistics might understand the caveats of partial-dependence plots, whereas a consumer reading a heat map could misinterpret bright red pixels as causal rather than correlative. Organizations supplement the raw output with plain-language tooltips, short videos, or interactive walk-throughs that coach users on what the colors or numbers mean and, equally important, what they do not guarantee. Regulators are starting to codify such practices, requiring that explanations be available and comprehensible to a layperson in the decision context [22].

Lastly, equity audits need to extend beyond prediction fairness to explanation parity. A system may produce identical acceptance rates for two demographic groups, yet still describe its reasoning in more detailed or actionable ways for one group than the other. Auditors should measure the consistency of feature rankings, saliency intensities, and counterfactual suggestions across protected attributes. They should verify that any differences can be justified by legitimate factors rather than reflecting hidden bias. Without such checks, well-intentioned transparency can entrench inequities by giving some users a more straightforward path to recourse while leaving others in the dark.

7. Conclusions

This paper positions magnetic AI as a practical, scalable strategy for injecting explainability into the countless black-box models influencing credit decisions, hiring, medical triage, and other facets of economic and social life. Rather than requiring expensive retraining or code rewrites, the magnetic approach attaches passively to existing data flows, learns a lightweight surrogate in real time, and delivers multiple explanation formats that can satisfy data scientists, end users, auditors, and regulators alike. We first synthesize decades of research on interpretability, model compression, and drift detection to ground the proposal in established theory. We then distill that literature into concrete design principles: non-intrusiveness, continual fidelity auditing, modular deployment, and explanation pluralism tailored to stakeholder needs.

Building on these principles, we outline an evaluation blueprint that cuts across three dimensions. The technical track measures surrogate accuracy, latency overhead, and drift-detection sensitivity. The human track uses controlled studies and field pilots to gauge whether different user groups understand and act on the explanations. The regulatory track maps the agent's outputs to statutory requirements such as the EU AI Act's transparency duty, U.S. consumer protection guidelines, and industry standards like ISO 42001. By integrating these perspectives, the paper provides a holistic roadmap for retrofitting trustworthy AI capabilities into existing machine-learning stacks without disrupting production workflows. Ultimately, magnetic AI extends the idea of surrogate modeling from a one-off snapshot to a living, continuously audited companion, positioning organizations to meet emerging policy mandates and rising public expectations for transparency and accountability.

7.1. Limitations

The magnetic-AI framework is, at present, a theoretical blueprint. It has not yet been stress-tested on production traffic in banking, retail, health care, or public-sector settings, where data rates, latency budgets, and privacy constraints differ sharply. Field trials are needed to reveal whether the surrogate can keep pace with high-volume streams, whether passive interception introduces unacceptable delay, and which sectors face unique regulatory or contractual hurdles.

These deployments will also expose weak security points, such as opportunities for adversaries to infer proprietary decision logic or poison the surrogate's sliding-window buffer. In addition, the current design assumes a supervised task with stable labels—credit approval, fraud detection, or image classification—leaving open how a magnetic agent would operate in unsupervised anomaly detection, continuous exploratory reinforcement learning, or free-form generative applications where outputs are text, images, or code snippets rather than class scores. Each paradigm raises new questions about what counts as a faithful surrogate, how to define drift or fidelity, and which explanation formats are meaningful to users. Therefore, comprehensive empirical studies across these settings are essential before the approach can be considered production-ready.

7.2. Future Work

Future research must move the magnetic-AI concept from controlled prototypes into live production pipelines. Pilot deployments in banking, e-commerce, and telemedicine sectors would reveal practical limits on throughput, latency, and privacy while showing how easily the agent can be co-containerized, versioned, and rolled back under real traffic. Once embedded, the surrogate-learning engine should evolve from periodic mini-batch updates to accurate streaming operation, digesting continuous flows of tabular events, log sequences, sensor signals, and even raw audiovisual frames without halting for retraining. Handling these multimodal inputs will require hybrid learners that combine gradient-boosted trees for structured features, lightweight convolutional backbones for images, and adapter-based mini-transformers for text, all coordinated by a reservoir buffer that prioritizes the most recent or conceptually novel samples.

A second avenue involves deeper integration with large foundation models that have chain-of-thought capabilities. Instead of treating the surrogate purely as a predictive mimic, an agent could query a frozen language model for self-rationalizing traces, then cross-check those traces against feature-importance scores to generate richer, more coherent explanations. This hybrid could also let users ask follow-up questions in natural language—Why did age matter more than income?—and receive conversational clarifications grounded in statistical evidence and domain policy.

Finally, the community needs shared benchmarks that evaluate explanation quality across domains rather than in narrow, single-task silos. A standard suite might pair representative workloads—credit risk, dermatology imaging, autonomous-vehicle perception—with crowdsourced judgment tests, cognitive-load surveys, and perturbation-based robustness checks. Metrics would cover fidelity, sparsity, stability under re-queries, resistance to inversion attacks, and user comprehension measured through decision-making tasks. Establishing such benchmarks would allow researchers to compare methods rigorously, accelerate regulatory acceptance, and guide practitioners toward solutions whose benefits generalize beyond any industry.

References

[1] M. Leon, "Ai-driven digital transformation: Challenges and opportunities", *Journal of Engineering Research and Sciences*, vol. 4, no. 4, p. 8–19, 2025, doi:10.55708/js0404002.

[2] M. Leon, "Generative artificial intelligence and prompt engineering: A comprehensive guide to models, methods, and best practices", *Advances in Science, Technology and Engineering Systems Journal*, vol. 10, no. 02, p. 01–11, 2025, doi:10.25046/aj100201.

[3] M. Leon, B. Depaire, K. Vanhoof, "Fuzzy cognitive maps with rough concepts", "Artificial Intelligence Applications and Innovations: 9th IFIP WG 12.5 International Conference, AIAI 2013, Paphos, Cyprus, September 30–October 2, 2013, Proceedings 9", pp. 527–536, Springer Berlin Heidelberg, 2013.

[4] H. DeSimone, M. Leon, "Leveraging explainable ai in business and further", "2024 IEEE Opportunity Research Scholars Symposium (ORSS)", p. 1–6, IEEE, 2024, doi:10.1109/orss62274.2024.10697961.

[5] M. Leon, H. DeSimone, "Advancements in explainable artificial intelligence for enhanced transparency and interpretability across business applications", *Advances in Science, Technology and Engineering Systems Journal*, vol. 9, no. 5, p. 9–20, 2024, doi:10.25046/aj090502.

[6] M. Velmurugan, C. Ouyang, R. Sindhgatta, C. Moreira, "Through the looking glass: evaluating post hoc explanations using transparent models", *International Journal of Data Science and Analytics*, 2023, doi:10.1007/s41060-023-00445-1.

[7] S. Jesus, C. Belém, V. Balayan, J. Bento, P. Saleiro, P. Bizarro, J. Gama, "How can i choose an explainer?: An application-grounded evaluation of post-hoc explanations", "Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency", FAccT '21, p. 805–815, ACM, 2021, doi:10.1145/3442188.3445941.

[8] M. Leon, N. M. Sanchez, Z. G. Valdivia, R. B. Perez, "Concept maps combined with case-based reasoning in order to elaborate intelligent teaching/learning systems", "Seventh International Conference on Intelligent Systems Design and Applications (ISDA 2007)", pp. 205–210, IEEE, 2007.

[9] H. DeSimone, M. Leon, "Explainable ai: The quest for transparency in business and beyond", "2024 7th International Conference on Information and Computer Technologies (ICICT)", p. 532–538, IEEE, 2024, doi:10.1109/icict62343.2024.00093.

[10] M. Leon, G. Nápoles, M. M. García, R. Bello, K. Vanhoof, "Two steps individuals travel behavior modeling through fuzzy cognitive maps pre-definition and learning", "Advances in Soft Computing: 10th Mexican International Conference on Artificial Intelligence, MICAI 2011, Puebla, Mexico, November 26-December 4, 2011, Proceedings, Part II 10", pp. 82–94, Springer Berlin Heidelberg, 2011.

[11] G. Nápoles, F. Hoitsma, A. Knoben, A. Jastrzebska, M. Leon, "Prolog-based agnostic explanation module for structured pattern classification", *Information Sciences*, vol. 622, p. 1196–1227, 2023, doi:10.1016/j.ins.2022.12.012.

[12] V. Hassija, V. Chamola, A. Mahapatra, A. Singal, D. Goel, K. Huang, S. Scardapane, I. Spinelli, M. Mahmud, A. Hussain, "Interpreting black-box models: A review on explainable artificial intelligence", *Cognitive Computation*, vol. 16, no. 1, p. 45–74, 2023, doi:10.1007/s12559-023-10179-8.

[13] M. Leon, "Gail: Enhancing student engagement and productivity", *The International FLAIRS Conference Proceedings*, vol. 38, 2025, doi:10.32473/flairs.38.1.138689.

[14] S. Bordt, M. Finck, E. Raidl, U. von Luxburg, "Post-hoc explanations fail to achieve their purpose in adversarial contexts", "2022 ACM Conference on Fairness Accountability and Transparency", FAccT '22, p. 891–905, ACM, 2022, doi:10.1145/3531146.3533153.

[15] W. J. von Eschenbach, "Transparency and the black box problem: Why we do not trust ai", *Philosophy & Technology*, vol. 34, no. 4, p. 1607–1622, 2021, doi:10.1007/s13347-021-00477-0.

[16] S. Hosseini, H. Seilani, "The role of agentic ai in shaping a smart future: A systematic review", *Array*, vol. 26, p. 100399, 2025, doi:10.1016/j.array.2025.100399.

[17] D. B. Acharya, K. Kuppan, B. Divya, "Agentic ai: Autonomous intelligence for complex goals—a comprehensive survey", *IEEE Access*, vol. 13, pp. 18912–18936, 2025, doi:10.1109/ACCESS.2025.3532853.

[18] N. Karunanayake, "Next-generation agentic ai for transforming healthcare", *Informatics and Health*, vol. 2, no. 2, p. 73–83, 2025, doi:10.1016/j.infoh.2025.03.001.

[19] M. Leon, "The escalating ai's energy demands and the imperative need for sustainable solutions", *WSEAS Transactions on Systems*, vol. 23, pp. 444–457, 2024.

[20] U. Ehsan, P. Wintersberger, Q. V. Liao, E. A. Watkins, C. Manger, H. Daumé III, A. Riener, M. O. Riedl, "Human-centered explainable ai (hcxai): Beyond opening the black-box of ai", "Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems", CHI EA '22, Association for Computing Machinery, New York, NY, USA, 2022, doi:10.1145/3491101.3503727.

[21] S. Nyawa, C. Gnekpe, D. Tchuente, "Transparent machine learning models for predicting decisions to undertake energy retrofits in residential buildings", *Annals of Operations Research*, 2023, doi:10.1007/s10479-023-05217-5.

[22] J. Mökander, "Auditing of ai: Legal, ethical and technical approaches", *Digital Society*, vol. 2, no. 3, 2023, doi:10.1007/s44206-023-00074-y.

Biography

Dr. Maikel Leon is interested in Artificial Intelligence (AI), Generative AI (GenAI), and Machine Learning (ML). His work bridges intelligent systems' theoretical foundations and practical applications, particularly emphasizing explainability, hybrid models, and educational innovation. Dr. Leon has published in high-impact journals such as IEEE Transactions on Cybernetics, Information Sciences, Knowledge and Information Systems, International Journal on Artificial Intelligence Tools, Intelligent Decision Technologies, and International Journal of Learning, Teaching and Educational Research. His research explores cutting-edge topics, including prompt engineering, sustainable AI, personalized tutoring via generative models, hybrid fuzzy systems, and large language model benchmarking. He was awarded the Cuban Academy of Sciences National Award for the Most Relevant Research in Computer Science. Dr. Leon obtained his PhD in Computer Science at Hasselt University (Belgium), an MSc and a BSc in Computation from Central University of Las Villas (Cuba), and currently serves as Associate Professor of Practice in Business Technology at the Miami Herbert Business School, University of Miami (Florida, USA).

Received: 14 June 2025, Revised: 02 August 2025, Accepted: 04 August 2025, Online: 21 August 2025

DOI: https://dx.doi.org/10.55708/js0408005

AI-Powered Decision Support in SAP: Elevating Purchase Order Approvals for Optimized Life Sciences Supply Chain Performance

Vinil Apelagunta[*,1], **Vishnuvardhan Reddy Tatavandla**[2]

[1]Senior SAP Consultant, AIFA Labs, Frisco, 75033, USA

[2]Senior Business System Engineer-SAP Development

[*]Corresponding author: Vinil Apelagunta, Belle Mead, NJ, USA & Email: apelaguntavinil@gmail.com

ABSTRACT: Resilient and compliant supply chains, while essential to the Life Sciences, depend heavily upon SAP systems to manage the complexities involved. The standard Purchase Order (PO) approval process in SAP is an important upstream control point in the supply chain, but seldom has the required intelligence needed to manage endorsed compliance (e.g., GxP) or to be proactive in supply chain risk mitigation. This paper offers an introduction to a proof of concept that demonstrates how an AI enabled, decision support solution that embeds into SAP processes and workflows can provide opportunities to transform this critical process and improve overall performance within the supply chain. Beginning with the evolution of SAP's approval workflows, the paper updates the concepts around AI/ML applications for improving various supply chain functions, and situates intelligent automation as part of the strategic digital transformation landscape for Life Sciences. The paper establishes constructs to improve PO approvals through the embedding of AI contextually based insight to build in performance trend analysis of suppliers (e.g., delivery, quality) and contextually relevant compliance checks as part of the decision process. These and other safeguards can move the PO approval process from being predominately procedural to a more strategic control point, increasing supply chain visibility, resilience, compliance assurance, and operational performance relevant to the Life Sciences. .

KEYWORDS: SAP, S/4 HANA, Artificial Intelligence (AI), Purchase Order Approval, Procurement Controls, Supply Chain Management (SCM), Life Sciences, Trend Analysis, GxP Compliance, Regulatory Compliance, Supply Chain Resilience, Supply Chain Optimization, Digital Transformation

1. Introduction

In today's unpredictable global markets, especially under regulatory requirements (GxP) in the Life Sciences sector, where patient safety, regulation, and product efficacy are responsibilities we cannot overlook [1], to enable resilient, compliant, and efficient supply chain operations is not only necessary. The critical upstream supply chain processes of procurement and the subsequent approval of purchase orders (POs) are crucial control points that have a significant impact on downstream performance. The performance of such downstream processes could extend to manufacturing continuity, inventory levels, quality of finished product, and most importantly, on-time delivery of finished product [2], all of which involves approval of purchase orders (POs). We rely increasingly on Enterprise Resource Planning systems (predominantly SAP S/4HANA) and often embedded with specialized systems such as those offers for Integrated Business Planning (IBP), Quality Management and Business Network, most of which are interconnected by the same database, to manage the flow of such processes throughout the supply chain. Nevertheless, the traditional PO approval mechanisms in SAP, despite efforts to enhance flexibility, traditionally remain procedure-focused [3]. They usually lack the situational intelligence to determine where suppliers might present a supply chain risk— such as reliability declines, quality declines— or the ability to rigorously enforce compliance with any stringent Life Sciences quality agreement or GxP requirements [4, 5]. This analytical gap at an upstream control point can introduce risks throughout the supply chain and open the door to costly disruptions, serious quality failures that impact patient safety, or compliance failures.

To tackle many supply chain issues harnessing the transformative potential of Business Process Automation (BPA), Digital Transformation and advanced technologies (Artificial Intelligence (AI) and Machine Learning (ML)), across the end-to-end value chain must be a key part of the solution [2, 6]. This paper will examine these trends, in the context of SAP purchase order (PO) approvals. More specifically, this paper will:

- Track the historical evolution of SAP PO approval workflows, identifying their historical weakness in enabling dynamic supply chain risk assessment.
- Review prevalent applications of AI/ML now impacting supply chain activities that are managed through SAP and set the context for advanced analytics.
- Consider the use cases of AI/ML techniques in the

context of supply chain digitalization and a BPA strategy.

- Provide an aspirational two future states, embedding advanced analytics (specifically supplier performance trend analysis and context relevant quality/compliance analytics) in the PO approval decision process, to strengthen upstream controls, for improved supply chain effectiveness.

The aim is to shine a light on the continuum from basic workflow automation to real intelligent decision support, then find examples of how we can take an essential administrative action and conceptualize it as a control point that is more informed, efficient, and strategic in the world of contemporary procurement, particularly the Life Sciences supply chain.

2. Literature Review

2.1. *Evolution of SAP Purchase Order Approval Workflow*

The approval of purchase orders (POs) in SAP's Materials Management (MM) module is a significant control within the overall supply chain's Procure-to-Pay cycle. Over the years, SAP systems have leveraged configurable, however, inflexible, rule-based methods called 'Release Strategies' (ME28/ME29N) to route POs for approval based on predetermined criteria like value, material group or plant. As summarized in Figure 1, while helpful for basic policy enforcement, these strategies are limited in authority and focus on authorizations limited to factual invariant conditions. More recently, SAP has the 'Flexible Workflow' architecture in S/4HANA, which has greater flexibility in specifying multi-step approval processes, dynamic approval authorities (based on roles or logic specified through BAdIs) and time frames [3, 7]. Yet even with more flexible frameworks, the standard decision-making functions of these workflows still rely heavily on procedural notification to approvers whose documents are routed according to configured rules, rather than providing contextual analysis, or data-driven insights to support the approvers' decision-making process [3].

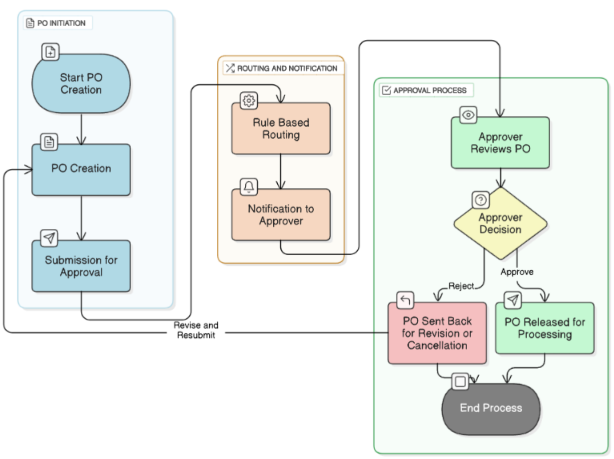

Figure 1: Standard Rule-Based SAP Purchase Order Approval Workflow (As-Is Process)

2.2. *Leveraging AI and Machine Learning in SAP-Driven Life Sciences Supply Chains*

Organizations are not only taking advantage of typical workflow capabilities, but also increasingly utilizing Artificial Intelligence (AI) and Machine Learning (ML) to build efficiency and predictive capabilities across critical supply chain processes handled in SAP systems, which is important in a stressful Life Sciences environment. One critical area is improving supplier management, where AI/ML can go well, beyond simple financial checks into evaluating historical delivery performance (e.g. OTIF rates), quality records (e.g. batch acceptance rates, audit outcomes), GxP standards, regulatory compliance certifications, and accepted historical rates to establish dynamic risk scores related to possible impacts on supply chain performance or compliance failures [2, 8].

In addition, AI/ML applications are being adopted throughout the wider Life Sciences supply chain: increasing the accuracy of demand forecasting in platforms like SAP IBP in order to assure product availability [2]; optimizing inventory levels, especially for materials requiring special storage conditions (e.g., temperature) or products with short shelf-life [2]; enabling predictive quality analytics by combining data from manufacturing, quality management (QM) and laboratory information management system (LIMS) in order to proactively predict future batch deviations [9]; improving logistics agility, including route optimization, or the prediction of transport delays and risks [4, 10]; and AI/ML can address substantial compliance requirements including analyzing traceability data (e.g., serialization) or the automation of some aspects of regulatory documentation management . Many academic surveys report on AI applications used in supply chain management, but usually focus on using resilience for creating risk management strategies that are important in Life Sciences sector [4].

AI/ML provides supply chain efficiency across the Procure-to-Pay (P2P) cycle by automating functions like invoice matching [11] and utilizing complex multi-way matching rules (PO, Goods Receipt, Quality Inspection) [11] to identify anomalies related to errors or fraud [11]. Typically used in conjunction with these intelligent applications, Robotic Process Automation (RPA) technology provides solutions for repetitive, high-volume, rule-based functions [10]. RPA sometimes blends into AI technologies, enabling Intelligent Process Automation (IPA) [10]. Although AI/ML applications enhance visibility, efficiency, and predictive capabilities in the supply chain, many are limited to specific functional areas and largely rely on structured data [12]. Integration of these outputs without displacing them, expanding AI logic capabilities to support complex, cross-functional supply chain decisions, such as the PO approval situation described later, is still work in progress and opportunity.

2.3. *The Strategic Role of Digital Transformation in Life Sciences Supply Chains*

The journey on established workflows within SAP (which are discussed in "Evolution of SAP Purchase Order Approval Workflows") and the usage of specific AI/ML applications (which are discussed in "Leveraging AI and Machine Learning in SAP-Driven Life Sciences Supply Chains") are important building blocks on a larger strategic objective: Business Process Automation (BPA) and Digital Transformation fundamentally changing Life Sciences supply chain operations. This transformational ambition is beyond just incremental improvements, but redesigning supply chain processes from end-to-end - planning, sourcing, manufacturing, quality, and logistics - to support better operational efficiency, maintain strong GxP compliance, foster greater end-to-end visibility and traceability, enhance supply chain resilience from disruptions, better manage costs, and to optimize the overall function of the supply chain to generate strategic value, all while preserving product quality and patient safety [1, 4, 5].

In this context of supply chain transformation, technologies marketed by vendors like SAP, such as their 'Flexible Workflow', are designed to offer flexible infrastructures to enable coordination of complex processes with many internal or external participants. Robotic Process Automation (RPA) tackles high-volume, rule-based tasks in various activities [10]; specific AI/ML solutions (described above) provide predictive when the use of pattern recognition is needed to better manage supply chain risk with forecasting accuracy, quality control, or logistics [2, 4, 9]. The overall aim is increased integrated, intelligent, data-driven supply chain workflows to remove manual bottlenHosseinOrdibazar2025, ecks, secure the same data accuracy necessary for regulation oversight, and exploit real-time-analytics to quickly inform decisions across the supply chain life-cycle [13].

However, in order for the full promise of digital transformation in the complex and highly regulated Life Sciences value chain to be realized, a number of considerable barriers must be addressed, such as orchestrating together the many disparate systems (e.g., ERP, MES, QMS, LIMS, logistics interface), ensuring consistent, governed data quality across the value chain, executing organizational change management across siloed functions, and scaling from incremental improvements demonstrated during pilot projects to comprehensive, end-to-end intelligent supply chain automation that embodies advanced levels of reasoning and adaptive management [14, 15, 16].

It is essential to recognize that there are many commercial offerings, such as SAP's own Ariba and Spend Control Tower, that are now including AI as part of providing high-level spend analytics and supplier risk scorecards. These are effective solutions for purposes like strategic spend categorization, and identifying high-level supplier risks. But the framework which is put forward in this paper has an advantage: it puts highly contextual, real-time trend analysis directly into the transactional PO approval workflow. Unlike high-level dashboards that analytics providers have developed, we provide not just situational or landscape awareness, but micro intelligence (e.g., "Is the quality of this specific material from this supplier trending down in the last month?") , and most importantly, we provide it to the decision-maker exactly at the moment that they use it at a GxP-relevant control point. Operationalizing intelligence for front-line decision making and not just strategic decision making is a key feature identifying our contribution.

Operationalizing intelligence for key decisions at the frontline, and not exclusively at the strategic level, is proprietary to our offering. This distinction mirrors the character of Procurement 4.0, moving from the procurement of duel to procuring in a more intelligent way and ultimately for the transformative process. [17].

3. A Proposed Framework: AI-Powered Contextual Intelligence for Strategic PO Approvals

The purpose of the strategic supply chain transformation previously outlined is to more than simply provide a procedural checkpoint with respect to the SAP purchase order approval process (PO). The intention is to create a knowledgeable, data-driven point of decision that enables effective management of Life Sciences supply chain resilience, compliance, and performance [18]. This means inserting capabilities that use advanced analytics — AI/ML based — into the workflow so that approvers receive actionable intelligence related to the PO, and the supply chain context surrounding it.

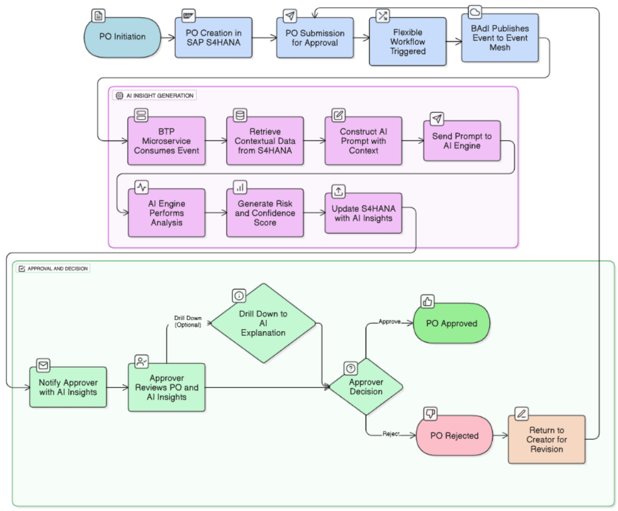

Figure 2: Proposed AI-Powered Purchase Order Approval Workflow with Event-Driven Architecture (To-Be Process)

Advanced decision support should provide performance data in context and trends [13, 19]. For example, if the system examines the historical delivery data of the items on the current PO, it could highlight if, while recently restrictions on single items were only recent, the delays appear to be a sign of a generally worsening trend that has ingress to disruption of timelines in critical production or supplies of clinical trials. Just as the system should examine trends in product quality, what tends are we seeing in rejection rates, returns and notifications/CAPAs associated with Quality Management (QM) for these specific materials? Are the trends arising from a concerned review of data about even possible lapses in GxP compliance or risks to the integrity of a batch, based on the recent performance of this supplier? Most critically, the System will recognize and track initial quality issues to specific products/batches that appear on this PO, as well as determining if this supplier and material combinations are consistent with current validation statuses and if quality agreements are indicated where this is particularly essential in Life Sciences [5, 9]. This type of dynamic trend analysis will provide intelligence not typically achieved from scorecards of static supplier performance in places.

3.1. Conceptual Technical Architecture and Implementation Approach (Event-Driven)

In order to appropriately deliver the AI-powered PO approvals insights in a decoupled and resilient architecture, an event-based solution that uses SAP S/4HANA and the SAP Business Technology Platform (BTP) is a fundamentally sound approach aligned to clean core principles [20]. The example 'To-Be' process flow is presented in Figure 2. The process would initiate from within the SAP S/4HANA Purchase Order approval workflow (e.g., Flexible Workflow). At some determinate step (e.g., at the approval step being ready to send assignment, or just prior to triggering a notification), we would implement a Business Add-In (BAdI) or enhancement spot within the workflow. This ABAP logic would be responsible for:

- **Event Publication:** An example for discussion is publishing a business event (e.g., "POApprovalInsightsRequired") using any topic on SAP Event Mesh (running on BTP) or other enterprise messaging queue. The event payload would include some critical identifiable information for the Purchase Order (e.g. PO number, company code, etc.). Thus, the S/4HANA core process remains standard, with the extension point being unambiguous and upgrade safe.

A microservice, built and hosted on SAP BTP (or any other cloud platform of your choosing), would subscribe to the "POApprovalInsightsRequired" event. On encountering the event, this BTP service would first conduct the following steps:

- **Contextual Data Retrieval:** The BTP service would call secure APIs (e.g., OData services backed by pre-built Core Data Services (CDS) views) back to S/4HANA to retrieve the full breadth of data for context. The BTP service would retrieve historical supplier performance, quality management data (QM notifications, CAPA summaries), compliance status, batch info, and relevant document info for the PO items and supplier.
- **Prompt Development & AI Engagement:** The service would develop an optimized prompt for the designated AI engine (ideally a company-native or private instance of an LLM for data security). This prompt would combine the retrieved contextual data and instruct the AI to conduct risk assessments, trend analysis and generate the score, sub-metrics and rationale. A significant item in design consideration is the selection of the AI model. While traditional machine learning models (e.g. regression or time-series) could be very powerful and computationally efficient for performing analysis of structured performance data, the proposed framework is based on a Large Language Model (LLM) for a specific purpose: synthesis & explanation. The goal is to have a hybrid approach whereby a LLM acts as the reasoning engine using & interpreting many data types, including structured outputs of the traditional machine learning models, unstructured text in Quality Notification long-text fields, & codified GxP compliance rules [5, 21]. The uniqueness of LLM models comes from their ability to synthesize these data types into a cohesive narrative (that is human-readable) and actionable recommendation, and solve the "black box" question by providing clarity into why it provided the risk score.
- **Asynchronous Contextual Insight Update:** After receiving the AI analysis, the BTP service would then asynchronously update the relevant context in S/4HANA. This could happen by hitting a custom API exposed by S/4HANA (to update custom fields on the PO, custom tables or append to workflow attachments) or by saving those insights in a data store on BTP, which could be queried by S/4HANA or Fiori apps when the approval task was presented.

The PO approval notification (email or Fiori 'My Inbox' item) would then be customized to retrieve and visualize AI-based insights, which could include a minor conditional branch in the notification generation or Fiori UI to check for existing insights before rendering, or perhaps a secondary notification/update if the insights are received after the original task is created.

For the risk analysis report, a similar event-driven, or direct API call, could be established. A user action (i.e., clicking a link in Fiori app) may create a new event for the BTP service, or directly call it to request the full report from the AI engine, which will ultimately be passed back to the user (i.e., email, Fiori UI, etc.).

To make the information consumable and actionable, the vision is to present insights directly in the approver's standard interface (i.e., SAP Fiori app, notification email, etc.). This could be a summary "PO Risk/Confidence Score" [4], along with some key contributing sub-metrics (i.e., 'Recent On-Time Delivery Trend', 'Quality/Compliance Flags (PO Items)'). For transparency, and to allow for exploratory dives, approvers would be able to quickly drill-down to a detailed report that explains how the score was derived, visualizes performance trends, and cites specific supporting data points or events (e.g., recent delivery delays affecting key production lines, specific QM notifications against PO items) [4, 13].

With such capabilities, the approval task could be quantitatively transformed from primarily an administrative check, to an informed risk-based decision. It would allow approvers to make decisions in more-informed, fast, and confident manners, while allowing the explicate identification of upcoming supply chain interruptions or compliance failures, and complete an important contribution to overall Life Sciences supply chain integrity, agility, and performance, and in doing so, the appropriate safeguarding of patient safety, and continuity of care [1, 5, 19].

3.2. Key Considerations for this Event-Driven Architecture

- **Asynchronous Nature:** Despite promoting decoupling and resilience, this architecture may contribute to latent time to insights. A formal Service Level Agreement (SLA) for insights needs to be defined (e.g., 95% of insights are received in 60 seconds of PO submission), in order to mitigate delays/models that may stall the approval process. The UI, e.g., the Fiori "My Inbox" application, needs to anticipate this and provide a status such as "AI insights are being generated..." to establish user expectations effectively.

- **Complexity:** The complexities of managing distributed transactions, ensuring you have eventual consistency, and monitoring event flows can be more complicated than using a synchronous model.

- **Resilience & Scalability:** This architecture is typically more resilient to failures in individual components and can scale components in the solution independently.

- **Clean Core:** With the use of BAdIs publishing events from S/4HANA, and using BTP for extensions, you are staying true to SAP's clean core.

4. Simulated Case Study: AI-Powered Decision Support at 'Innovida Life Sciences'

To demonstrate the applied impact of the proposed framework and what it means in practice, this section presents a fictitious case study of a fictitious company, "Innovida Life Sciences".

Scenario: Innovida Life Sciences is a mid-sized biopharmaceutical company that has a leading product, "Gerocept," a temperature-sensitive biologic drug used in critical patient therapies. Innovida's manufacturing process of Gerocept is based on the availability and quality of a critical raw material, in this case, "Stabilizer-7", and associated tight quality and delivery processes as a GxP compliant product. Unfortunately, there have been sporadic delays in manufacture and delivery of Gerocept-linked to the performance of the company's primary supplier of Stabilizer-7, "Global Bio-Reagents Inc" [4]. Such challenges are common in complex pharmaceutical supply chains, where AI and big data analytics are increasingly being leveraged for greater efficiency and risk mitigation [22].

4.1. The 'As-Is' Process: A High-Risk PO Approval

A new Purchase Order (PO #4500012345) is generated in SAP S/4HANA for a transactional shipment of 1000L of Stabilizer-7. The PO now enters the standard flexible workflow for approval, as shown in general terms in Figure 1. The approver who is a procurement manager will review the PO using the standard criteria: right material, right quantity, right price, and right cost center; and will see that the vendor, Global Bio-Reagents Inc., is on the approved vendor list. At this stage it looks pretty normal.

But there are important risk signals buried in separate SAP data modules. In the next section, we provide a table (Table 1) of the risk signals, which and in the 'As-Is' situation are not presented to the approver in a consolidated, contextual way.

Table 1: Supplier Performance Data: Global Bio-Reagents Inc.

Metric	Data Point	Observation/ Trend
Supplier Delivery Performance	On-Time-In-Full (OTIF) rate has fallen from 95% to 70% over the last 6 months.	A consistent and sharp negative trend, indicating deteriorating reliability.
Quality Management Data	3 new Quality Notifications (QNs #80012, #80015, #80019) in the last quarter.	A recurring GxP compliance issue specifically related to temperature deviations for "Stabilizer-7."
Goods Receipt Data	Batch acceptance rate at goods receipt has dropped from 100% to 92% in the last 2 months.	Increasing number of shipments failing initial quality checks, suggesting a potential systemic issue.

Outcome of the 'As-Is' Process: Missing a unified, contextual view of these deteriorating trends, the procurement manager approves the PO. The shipment from Global Bio-Reagents arrives two days late. More critically, it arrives with another temperature excursion. This requires a mandatory quality investigation and corrective and preventive action (CAPA) plan, which takes time. The two-week delay and the investigation delayed and complicated the production of a vital Gerocept batch. This had an estimated opportunity cost of lost sales of US $1.2 million, plus an estimated $50,000 in internal costs to manage the CAPA. This event also created a high compliance risk that may be flagged in a global regulatory audit in the future.

4.2. The 'To-Be' Process: AI-Powered Risk Mitigation

Now, think about the same PO #4500012345 being processed with the proposed AI-enabled framework which is intended to inject contextual intelligence into the approval workflow.

1. **PO Submission & Event Trigger:** The PO is sent for approval in SAP S/4HANA. A Business Add-In (BAdI) developed in the Flexible Workflow sends a "POApprovalInsightsRequired" event to the SAP Event Mesh. The event-driven model supports sidecar architecture and enables independence of the core SAP process.

2. **AI Insight Generation:** A microservice on the SAP Business Technology Platform (BTP) listens for the event. When the assistant receives this notification, it collects all historical delivery, quality, and goods receipt information from S/4HANA using pre-defined Core Data Services (CDS) views. This context is then forwarded to a secure AI engine for a risk assessment and trend analysis to be carried out

3. **Actionable Insights Provided:** The approver opens SAP Fiori's 'My Inbox' and selects the PO approval task. The approver sees more than just the basic PO data, they can engage with the AI generated summary located on the approval screen (as shown in Table 2).

Now the approval screen is being used as an active decision point.

Table 2: Simulated AI Insights in Fiori App

Insight Component	Details & Justification
Overall Confidence Score	35% **(High Risk)**
Key Risk Factor: Delivery Trend	Supplier's on-time delivery for this material has degraded by 25% over the last six months, indicating a high probability of a schedule-impacting delay.
Key Risk Factor: Quality Lapses	Recurring GxP compliance issues (3 QNs) for this specific material due to temperature deviations. High risk of repeat quality failure.
AI Recommendation	**Action: Reject PO**. The combination of worsening delivery and repeated quality failures presents a significant and immediate risk to manufacturing continuity and compliance.
Suggested Next Steps	Initiate an expedited order with the qualified secondary supplier and trigger a formal performance review for Global Bio-Reagents Inc.

Quantified Benefits: The AI-based intervention has transformed approval from a procedural task to a strategic risk-reduced approval stage. By preventing the high-risk purchase order (PO), Innovida Life Sciences:

- Avoids $1.2 million in lost revenue as the Gerocept batch is able to be on time for production and avoid disruption to existing production operations.
- Avoids the $50,000 cost of the CAPA investigation.
- Obtain improved supply chain resiliency by correcting a weak-link in its supply chain which will generate systemic improvements in performance and patient safety.

This pseudo case study has indicated that as a result of integrating contextual relevant AI-based intelligence, directly into the approving users workflow in the SAP PO module to transform reactive problem solving into proactive risk management, organizations have created value, both meaningfully and quantifiably.

4.3. *Implementation Considerations: Cost-Benefit Analysis and ROI*

While an accurate return on investment (ROI) is reliant on organization-specific elements, we can estimate a directional cost-benefit analysis that is rooted in the real-world application of the "Innovida Life Sciences" case study. This is also critical to make a business case for the upfront investment necessary for this kind of strategic implementation [2, 14].

Estimated Costs The implementation costs are a composite of technology, development, and ongoing maintenance.

- **Technology & Infrastructure Costs:**
  - **SAP BTP Services:** The subscription costs associated with the main BTP services (SAP Event Mesh for the event-driven architecture, Cloud Foundry, or Kyma runtime environment for the microservice, SAP AI Core to manage the AI/ML model life cycle).
  - **AI/LLM Services:** A large window for operational costs. Organizations can either use the Commercial pay-per-use Large Language Model (LLM) API and incur usage costs each time the LLM is used, or take the hit on the upfront investment to host a private company-native LLM that is similar to a LLM API but provides data security and privacy.
- **Development & Implementation Costs:**
  - **Personnel:** One-time project cost of engaging a number of skilled people, including SAP BTP/ABAP developers to create and plan the integration, AI/ML engineering resources to develop prompt engineering, and model tuning.
  - **Data Integration:** A substantial amount of working to develop and validate the Core Data Services (CDS) views to extract clean, reliable data from source systems (ERP, QMS, LIMS).
- **Ongoing Maintenance & Governance:**
  - **Model Monitoring:** Considerable ongoing effort is required to monitor the performance of the LLM-enabled AI model to ensure the model is not drifting or 'hallucinating', and to help maintain accuracy.
  - **Change Management & Training:** The upfront costs associated with training end users as well as the ongoing costs associated with managing the organizational shift to AI-assisted decision making are of considerable size, and in consideration of the importance of actually adopting the systems, have program affect changes in the organization too.

Projected Benefits & Value Proposition The benefits that are offered by the proposed solution go beyond straightforward automation and will provide considerable financial and strategic value too.

- **Quantifiable (Direct) Benefits:**
  - **Avoidance of Disruption Costs:** As we illustrated in the case study, proactively avoiding just one high-risk PO can avoid major losses. The $1.2 million in lost revenue was avoided by stopping one production from stopping; this is a compelling rationale.
  - **Reduction in Compliance & Quality Costs:** The framework will allow organizations to avoid direct costs due to quality failures, similar to the estimated $50,000 to investigate a CAPA, cited in the case study. As well, the level of manual work necessary for quality reviews will also be decreased.

  - **Increased Operational Efficiency:** By automating the data gathering and analysis that an expert approver would do manually, the workflow system could shorten approval cycle time and release procurement professionals to spend more time on negotiations and on managing and developing suppliers.
- **Qualitative (Strategic) Benefits:**
  - **Enhanced Supply Chain Resilience:** The value in moving from a reactive risk management approach to a proactive approach, will ultimately drive a more resilient and agile supply chain.
  - **Strengthened GxP Compliance:** The system provides an automated, auditable control point in order to enforce GxP upstream in the procurement process, to protect product quality and patient safety.

5. Limitations of the Model Proposed

This model has considerable promise for improving supply chain management across the Life Sciences Value Chain, but it is also important to highlight the limitations and challenges of implementation.

5.1. Dependency on Data

The effectiveness of the AI analysis is dependent not only on the volume and granularity of data from potentially many data sources in SAP S/4HANA and other systems, but also on the availability, accessibility and quality of that data. Achieving a holistic view of a product's lifecycle often challenges data silos and the need for quality data (which must be done with data integrity because of GxP) and data governance. In particular, the approach requires the integration of diverse systems such as the ERP, Quality Management Systems (QMS), Laboratory Information Management Systems (LIMS) and Manufacturing Execution Systems (MES), with effective data pipelines and validated data lineage that allows for an audit trail for everything the AI model ingests. Establishing this level of data integrity is substantial governance and technical challenge in and of itself.

5.2. AI Model Limitations

- **Explainability and Trust:** The "black box" aspect of some complex AI models can be a serious obstacle. To realize the complete transparency and auditability required for AI-generated risk scores within a GxP-regulated framework will need additional enhancements in the development and incorporation of Explainable AI (XAI) techniques.
- **Bias and Accuracy:** AI models can perpetuate biases in their training data, which, if not carefully addressed, can result in an unintentional and unfair supplier evaluation. In addition to addressing the uncertainty of the AI's output, and preventing model drift, or "hallucinations", will require continuous monitoring, and validation processes.

5.3. Implementation and Governance Risks

- **Data Privacy within a Proprietary Context:** Managing sensitive supplier data happens within proprietary contexts. The use of public AI services inevitably exposes this sensitive data to some risk. Therefore, this architecture proposes a strong recommendation for a proprietary or company-native LLM hosted in a secure setting to keep data confidential. All data transmissions must be encrypted, and access to the underlying data sources must be strictly governed based on roles [23].
- **Regulatory Validation with a GxP Context:** Within the Life Sciences industry, any system that can impact/affect quality decisions must undergo Computer Systems Validation (CSV) as part of any framework such as GAMP 5. This will require comprehensive documentation and auditability of the AI Model and decision-making logic [24]. For the AI recommendations to be validated, it will be crucial to be able to trace recommendations back to the physical data points that affect each recommendation because of the requirement for data lineage and integrity to prove GxP compliance. Changing the model would require a change control process.
- **Complexity of Change Management:** Implementing this system involves more than a technology change; it will involve changing core business processes, and possibly causing a change in the roles of people working with the system. Within the highly regulated, risk-averse atmosphere of the pharmaceutical sector, dealing with user adoption will be a significant roadblock. Apart from user adoption, one of the greatest risks faced, will be user resistance to AI recommendations. Resistance needs to be considered in the context that trust must be established between the user and the AI System. To construct this level of trust, a change management strategy must promote the AI technology as a decision support system instead of a substitute for the decision being made [3, 16].

5.4. Technical Challenge and Cost

The implementation of this proposed event-based architecture, which will require technical use development, deployment, and ongoing stewardship, involves many technical challenges and associated expenses. Additionally, to ensure there is no excessive delay introduced into the time-sensitive PO approval process, the performance of the system must be managed carefully.

6. Future Opportunities

Can AI-based approaches to PO approval evolve and this framework for PO approval enhanced? The answer is yes with many possibilities for development and research to further build intelligent automation in Life Sciences SCLs.

Short-Term: Prototyping and Pilot Testing The next step in this project is the development of a prototype or pilot program in a controlled Life Sciences environment. In this environment, it will be possible to assess technical feasibility, refine the AI models and prompt, assess the usability of the interface, and collect empirical data on the systemic impacts the framework has on approval times, risk mitigation, and compliance adherence. Establishing robust methods and KPI's to quantify and measure the ROI, and other tangible benefits gained using the intelligent approval system is also an important task. This aligns with the broader academic push to move from conceptual frameworks for AI in the pharmaceutical supply chain towards validated, real-world applications [25].

Mid-Term: Building on The AI Core and Explainability In future versions it would be possible to have access to a wider range of internal and external data sources (i.e. broader market intelligence, feeds on geopolitical risk, sustainability data, even more detailed IoT sensor data based on logistics/manufacturing) that would lend themselves to richer contextual analysis [4]. Additional investment in Explainable AI (XAI) methods inside the decision support interface, will be important for user buy-in and to facilitate regulatory audit requirements [26, 27]. There are many developing robust XAI capabilities, as an ongoing area of research, and which can make decision-making more effective in complex contexts such as SCL [13]. Indeed, there are views from other scientific disciplines regarding the challenges associated with the process of delivering trust in XAI methods in relation to dealing with complexity [13, 28]. Furthermore, designing into the AI system, the ability for the AI to learn from the outcomes associated with approved POs and any feedback from users on its recommendations, will enhance its predictive accuracy over time.

Long-Term: Strategic Deployment and Governance With the fundamental principles of AI contextual intelligence, it would be possible to expand and apply the concepts to other high-importance decision points with the SCM area of the SAP ecosystem, (e.g. supplier qualification and onboarding, contract life cycle management, proactive quality event management). In deploying these concepts it will be important to have in mind the ongoing development and implementation of strong ethical guidelines, oversight frameworks, and management processes for AI execution in high-importance decisions in SCL, underpinned by the concepts of fairness, accountability, transparency, and bias mitigation [23, 24, 29].

7. Conclusion

The development of Purchase Order approval processes in SAP systems, moving from typical rule-based decisions to the more advanced, intelligent, contextually aware decision support systems discussed in this paper represents an important technological and strategic change that is necessary for modern Life Sciences supply chains. By combining the principles of Business Process Automation with the advanced analytical capabilities of AI and Machine Learning, especially when we can use this technology to embed contextual intelligence and trend analysis into critical upstream controls like Purchase Order approvals, we can advance supply chain standards well beyond that of daily operational execution.

This shift in thinking represents more than just automating processes—it is about giving supply chain professionals advanced insight to gain insight into timely actionable information at critical decision points in order to make better strategic decisions, to proactively address supply chain risks in relation to supply chain performance and quality, to adequately demonstrate GxP compliance, and ultimately to develop more robust and flexible Life Sciences supply chains. As intelligent technologies continue to evolve and be integrated across the SAP landscape (connecting the ERP with QM, IBP, MES, and outside partner data sources on networks), we can imagine a world where adaptive, predictive, and intelligent supply chain operations exist. This revolutionizes the way in which Life Sciences organizations can navigate the complexities of their global supply chain networks to manage reliable, compliant, and efficient material flows that are most critical for patient safety, regulatory compliance and sustainability in innovation.

References

[1] J. Y. Ma, L. Shi, T. W. Kang, "The effect of digital transformation on the pharmaceutical sustainable supply chain performance: The mediating role of information sharing and traceability using structural equation modeling", *Sustainability 2023, Vol. 15, Page 649*, vol. 15, p. 649, 2022, doi:10.3390/SU15010649.

[2] R. Toorajipour, V. Sohrabpour, A. Nazarpour, P. Oghazi, M. Fischl, "Artificial intelligence in supply chain management: A systematic literature review", *Journal of Business Research*, vol. 122, pp. 502–517, 2021, doi:10.1016/J.JBUSRES.2020.09.009.

[3] J. Idogawa, F. S. Bizarrias, R. Câmara, "Critical success factors for change management in business process management", *Business Process Management Journal*, vol. 29, pp. 2009–2033, 2023, doi:10.1108/BPMJ-11-2022-0625.

[4] A. H. Ordibazar, O. K. Hussain, R. K. Chakrabortty, E. Irannezhad, M. Saberi, "Artificial intelligence applications for supply chain risk management considering interconnectivity, external events exposures and transparency: a systematic literature review", *Modern Supply Chain Research and Applications*, 2025, doi:10.1108/MSCRA-10-2024-0041.

[5] J. Vellanki, "Discuss how ai-driven solutions can improve compliance with good practice (gxp) regulations in the biotechnology industry.-a comparative study", *International Journal of Research in Management Studies*, 2024.

[6] A. Daios, N. Kladovasilakis, A. Kelemis, I. Kostavelis, "Ai applications in supply chain management: A survey", *Applied Sciences 2025, Vol. 15, Page 2775*, vol. 15, p. 2775, 2025, doi:10.3390/APP15052775.

[7] A. Vaid, C. Sharma, "Pioneering digital transformation initiatives with cutting-edge sap s/4hana solutions", *World Journal of Advanced Engineering Technology and Sciences*, vol. 3, pp. 109–121, 2021, doi:10.30574/wjaets.2021.3.2.0075.

[8] G. Baryannis, S. Validi, S. Dani, G. Antoniou, "Supply chain risk management and artificial intelligence: state of the art and future research directions", *International journal of production research*, vol. 57, no. 7, pp. 2179–2202, 2019.

[9] S. Mundhra, K. S. Kumar, T. Prashant, "Harnessing ai and machine learning in pharmaceutical quality assurance", *Journal of Pharmaceutical Quality Assurance and Quality Control*, vol. 6, pp. 19–29, 2024.

[10] A. K. Percherla, "Effortless logistics automation - sap logistics process automation with uipath rpa", *Journal of Artificial Intelligence & Cloud Computing*, pp. 1–4, 2024, doi:10.47363/JAICC/2024(3)262.

[11] M. R. Kunchala, "Sap finance and management accounting with integration of ai and ml", *International Journal of Innovative Research in Engineering & Multidisciplinary Physical Sciences*, vol. 10, no. 3, pp. 1–5, 2022, doi:10.37082/IJIRMPS.v10.i3.231962.

[12] M. Brylowski, M. Schroeder, S. Lodemann, W. Kersten, "Machine learning in supply chain management: A scoping review", "Hamburg International Conference of Logistics (HICL) 2021", pp. 377–406, epubli, 2021.

[13] F. Olan, K. Spanaki, W. Ahmed, G. Zhao, "Enabling explainable artificial intelligence capabilities in supply chain decision support making", *Production Planning and Control*, vol. 36, pp. 808–819, 2025, doi:10.1080/09537287.2024.2313514.

[14] M. S. Attiany, S. A. Al-Kharabsheh, L. S. Al-Makhariz, M. A. Abed-Qader, S. I. S. Al-Hawary, A. A. Mohammad, A. A. A. Rahamneh, "Barriers to adopt industry 4.0 in supply chains using interpretive structural modeling", *Uncertain Supply Chain Management*, vol. 11, pp. 299–306, 2023, doi:10.5267/J.USCM.2022.9.013.

[15] Y. A. Bena, R. Ibrahim, J. Mahmood, "Current challenges of big data quality management in big data governance: A literature review", *Lecture Notes on Data Engineering and Communications Technologies*, vol. 210, pp. 160–172, 2024, doi:10.1007/978-3-031-59711-4_15.

[16] C. Kewalramani, S. Neema, "Generative ai in change management: A new framework for organizational transformation", *International Journal of Research and Analytical Reviews*, p. 838, 2024.

[17] D. K. Vaka, "Procurement 4.0: Leveraging technology for transformative processes", *Journal of Scientific and Engineering Research*, vol. 11, no. 3, pp. 278–282, 2024.

[18] C. M. Sakala, S. M. Bwalya, "The role of artificial intelligence in optimizing supply chain performance", *Journal of Procurement and Supply Chain Management*, vol. 2, no. 1, p. 1–14, 2023.

[19] N. Y. Hussain, P. A. Adepoju, A. I. Afolabi, B. Austin-Gabriel, "Ai and predictive modeling for pharmaceutical supply chain optimization and market analysis", *International Journal Of Engineering Research And Development*, vol. 20, pp. 191–197, 2024.

[20] J. Singh, "Event-driven architecture for real-time analytics in cloud crm platforms", *European Journal of Computer Science and Information Technology*, vol. 13, pp. 24–34, 2025, doi:10.37745/ejcsit.2013/vol13n422434.

[21] M. C. Nezianya, A. O. Adebayo, P. Ezeliora, "A critical review of machine learning applications in supply chain risk management", *World Journal of Advanced Research and Reviews*, vol. 23, pp. 1554–1567, 2024, doi:10.30574/WJARR.2024.23.3.2760.

[22] N. Bhatt, "Ai and big data analytics in pharmaceutical supply chain management", *International Journal of Scientific Research & Engineering Trends*, pp. 599–601, 2023.

[23] T. Birkstedt, M. Minkkinen, A. Tandon, M. Mäntymäki, "Ai governance: themes, knowledge gaps and future agendas", *Internet Research*, vol. 33, no. 7, pp. 133–167, 2023, doi:10.1108/INTR-01-2022-0042.

[24] M. Sendak, M. C. Elish, M. Gao, J. Futoma, W. Ratliff, M. Nichols, A. Bedoya, S. Balu, C. O'Brien, ""the human body is a black box": supporting clinical decision-making with deep learning", "Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency", FAT* '20, p. 99–109, Association for Computing Machinery, New York, NY, USA, 2020, doi:10.1145/3351095.3372827.

[25] A. D. Adekola, S. A. Dada, "Optimizing pharmaceutical supply chain management through ai-driven predictive analytics: A conceptual framework", *Computer Science & IT Research Journal*, vol. 5, pp. 2580–2593, 2024, doi:10.51594/csitrj.v5i11.1709.

[26] O. . Shaughnessy, P. Couto, M. Saarela, V. Podgorelec, "Recent applications of explainable ai (xai): A systematic literature review", *Applied Sciences 2024, Vol. 14, Page 8884*, vol. 14, p. 8884, 2024, doi:10.3390/APP14198884.

[27] V. Hassija, V. Chamola, A. Mahapatra, A. Singal, D. Goel, K. Huang, S. Scardapane, I. Spinelli, M. Mahmud, A. Hussain, "Interpreting black-box models: A review on explainable artificial intelligence", *Cognitive Computation*, vol. 16, pp. 45–74, 2024, doi:10.1007/s12559-023-10179-8.

[28] R. J. O'Loughlin, D. Li, R. Neale, T. A. O'Brien, "Moving beyond post hoc explainable artificial intelligence: a perspective paper on lessons learned from dynamical climate modeling", *Geoscientific Model Development*, vol. 18, no. 3, pp. 787–802, 2025, doi:10.5194/gmd-18-787-2025.

[29] D. B. Patel, C. Author, "Ethical ai: Addressing bias and fairness in machine learning models for decision-making", *Journal of Computer Science and Technology Studies*, vol. 3, pp. 13–17, 2021, doi:10.32996/JCSTS.2021.3.1.3.

Vinil Apelagunta received his Bachelor of Technology degree from Osmania University College of Technology in 2010. He earned his Master of Science from the Illinois Institute of Technology in 2012 and later completed his Master of Business Administration (MBA) in Digital Innovation from HEC Paris in 2021.

Vinil Apelagunta is a results-driven Digital Transformation Leader and SAP Project Manager with extensive experience in modernizing legacy systems within traditional industries. He specializes in leading the full product lifecycle, from ideation to enterprise-wide adoption, with a focus on enhancing operational efficiency and customer satisfaction. His work involves designing and implementing intuitive, data-driven solutions like inventory management platforms and customer-centric delivery systems that replace manual processes and drive significant business impact.

Vishnuvardhan Reddy Tatavandla has done his bachelors degree from Osmania University in 2009. He has done his Masters from Southern Illinois University Edwardsville in 2012.

Vishnuvardhan Reddy Tatavandla is an accomplished SAP Enterprise Architect with over 14 years of experience across the industrial spectrum and deep expertise in the SAP Order To Cash process. His career has been dedicated to mastering SAP ecosystems, with a strong focus on modules such as SD (Sales and Distribution), MM (Materials Management), CRM (Customer Relationship Management), ABAP (Advanced Business Application Programming), and extending into advanced areas like SAP BTP, Fiori and S/4HANA.