

VI

CHAPTER SIX

**AI AS AMPLIFIER: ETHICAL GOVERNANCE AND
HUMAN EXPERIENCE IN FASHION RETAIL**

Dr. Joey Pui Ki Chan

Development of Organizational Behavior Management Systems in the Implementation of Anthropomorphic AI Agents in Enterprises

Joey Pui Ki Chan¹

¹ SBS Swiss Business School, Zurich, Switzerland
Corresponding Author: Joey Pui Ki Chan, joeypkchan@hotmail.com

Chapter Information

- **Date of Receipt:** [19/03/2026]
- **Date of Acceptance:** [24/05/2026]
- **JEL Classification Codes :** M15 ; G21 ; O33 ;

Abstract

Artificial intelligence (AI) is advancing at a pace that renders traditional approaches unsustainable for brands and retailers. The critical question is no longer whether AI will reshape shopping, but how it amplifies organizational power and values, and why ethical considerations must guide its application. As technological capability accelerates more rapidly than human judgment, integrity emerges as a defining leadership competency. AI systems do not possess moral agency; they reflect the values and blind spots of those who design and deploy them. This underscores the urgent need to embed responsible AI principles, including safety, security, transparency, and accountability, into the core of business strategy rather than treating them as compliance afterthoughts. This chapter uses a qualitative single-case study and reflexive thematic analysis of interviews with executives and store managers to examine real-world applications of AI within a Fortune 500 fashion retail brand. Drawing on interviews across organizational levels, it captures both strategic intent and operational realities, analyzing immediate opportunities alongside longer-term transformations. A central concern is the potential erosion of human interaction in automated environments. The study explores how trust and credibility can be cultivated in AI-mediated experiences, ensuring brands retain an authentic, distinctively human presence. By foregrounding human experience as a differentiator, this chapter identifies four empirical themes in AI governance and offers practical insights on how the fashion industry can harness AI while preserving human connection and advancing sustainability.

Keywords: *artificial intelligence (AI) in retail, ethical AI, responsible AI, consumer trust, AI governance, sustainability in fashion.*

INTRODUCTION

The rapid acceleration of artificial intelligence (AI) has brought the fashion retail industry to a critical inflection point. For decades, the sector has relied on human intuition, creative direction, and interpersonal connections to build brand loyalty and drive consumption. Today, however, AI systems are increasingly capable of executing tasks that once required human judgment, from predictive inventory management and algorithmic pricing to generative design and personalized customer interactions. While the business case for AI adoption is compelling, with unprecedented efficiency and scale, the human and ethical implications remain profoundly complex.

This chapter examines the ethical, regulatory, and sustainability implications of AI adoption within fashion retail. It moves beyond theoretical debates by presenting an in-depth, qualitative case study of a Fortune 500 fashion brand navigating this transition. Informed by the socio-technical systems perspective, which holds that technology and human organization co-evolve and must be designed together, the research analyses how purpose-driven organizations attempt to scale AI without weakening trust, ethics, or the quality of human connection (Trist, 1981; Afroogh et al., 2024). This study makes an original empirical contribution by examining AI governance in fashion retail through the combined lens of executive strategy and frontline store operations, bridging corporate intent and lived employee experience in a way prior literature has not.

The structure of this chapter is organized as follows: Section 2 reviews the current literature on AI in retail, focusing on global regulatory frameworks, consumer concerns, the unique ethical risks of agentic AI, and the paradox of AI in sustainability. Section 3 details the qualitative, reflexive thematic analysis methodology used to conduct the case study. Section 4 presents the findings across four core themes, illustrating how AI acts as an amplifier of organizational intent. Finally, Section 5 offers strategic recommendations for retail leaders, acknowledges the study's limitations, and proposes directions for future research in responsible AI governance.

LITERATURE REVIEW

This section critically reviews the growing body of literature on artificial intelligence (AI) in retail and fashion, with a particular focus on its human, ethical, and governance implications. AI is widely recognized for optimizing supply chains, personalizing customer experiences, and improving operational efficiency, yet less attention has been given to how it reshapes relationships among consumers, employees, and brands (KPMG, 2025; Grewal et al., 2017; Chen & Prentice, 2025). The literature suggests that AI is a socio-technical force that can reinforce or challenge existing inequalities, biases, and norms within the fashion retail ecosystem (KPMG, 2026; The Business of Fashion, 2026). This reliance also introduces a critical dependence on third-party vendors, shifting strategic control away from retailers and toward technology providers. By synthesizing insights from business, ethics, and governance research, this section sets the stage for a practice-oriented discussion of how retailers can capture AI's benefits while deliberately preserving authentic human experience.

Artificial Intelligence in Retail

Artificial intelligence lies at the center of a major transformation in how consumers shop, how brands engage customers, and how supply chains create value (Bieńkowska, 2025;

Fontaine et al., 2019). Retail executives increasingly view AI adoption as strategically necessary, and many believe firms that embrace it will gain a competitive advantage (KPMG, 2025). Yet this optimism is tempered by concerns about ethics, regulation, and implementation. For example, KPMG (2025) reports that 56% of retail leaders identify ethical challenges as the main barrier to AI adoption, while *The Business of Fashion* (2026) notes that legal, compliance, and brand review processes are the top reasons fashion marketers do not scale AI. These tensions show that retailers need strong technical capability and clear governance frameworks, so AI enhances, rather than weakens, the human experience.

Global AI Regulation: EU Risk Framework and US Ethics Clash

The European Union (EU) has positioned itself as a global leader in protecting citizens' rights through AI-specific legislation. The European Commission proposed the Artificial Intelligence Act in April 2021, and the regulation entered into force on 1 August 2024, making it the world's first comprehensive legal framework on AI (European Commission, 2024). Its provisions apply progressively: prohibitions on the most harmful AI practices, including certain forms of social scoring and biometric mass surveillance, took effect in February 2025, while the main rules for high-risk AI systems and their enforcement apply from 2 August 2026. For fashion retailers operating in European markets, the Act's risk-based classification system is particularly consequential. AI tools used in recruitment, performance management, or credit assessment are classified as high-risk and subject to mandatory audits, transparency requirements, and human oversight obligations. Non-compliance with prohibited practices carries fines of up to €35 million or 7% of global annual turnover, whichever is higher (European Commission, 2024). Compliance, therefore, is not merely a legal obligation; it is a signal of trustworthiness to consumers, employees, and investors alike.

In contrast, the United States has taken a more fragmented, ethics-driven approach that often clashes with commercial interests. A notable illustration occurred in early 2026 when Anthropic, a leading AI developer, refused Pentagon contracts involving autonomous lethal warfare and mass surveillance of U.S. citizens, citing internal ethical red lines (CBS News, 2026). In response, the Trump administration characterised Anthropic as a "supply chain risk" and attempted to ban federal agencies from using its Claude AI model, though a federal judge temporarily blocked this ban in March 2026 (NBC News, 2026). This episode highlights a broader dilemma: if specialized AI developers face institutional friction over ethical boundaries, fashion retailers will likely encounter significant challenges in holding third-party vendors accountable to brand-specific values. The implication is that AI adoption will continue to accelerate, but its long-term sustainability depends on rigorous governance. Auditable decision logs, human-in-the-loop controls over pricing and inventory decisions, and independent ethics audits are increasingly recognized as essential compliance mechanisms (PwC, 2024; Mokander & Floridi, 2021; OECD, 2021).

Building AI Trust: User Concerns and Industry Response

Despite the advancing capabilities of AI, consumer trust remains fragile and, in some respects, declining. A 2026 survey of 250 readers by *Vogue Business*, conducted specifically within the fashion and beauty sector, found that fewer than one in four consumers (24%) trust the recommendations and summaries made by AI chatbots, while 55% actively express distrust (O'Brien & Shoaib, 2026). This sector-specific finding is more alarming than broader cross-industry figures, which show 54% of respondents are wary of trusting AI (Gillespie et al., 2025).

The distrust is structural: consumers are willing to share low-risk data such as dress size, but 72% refuse to share card details, and 46% will not share browsing history with AI systems. Only 31% would outsource a purchase to an AI agent, even one that knows their taste and history (O'Brien & Shoaib, 2026).

Employees, meanwhile, fear that AI will threaten their economic and professional well-being. A 2023 study by EY found that 72% of workers are concerned AI will negatively affect their salary or pay, 67% worry about losing promotions if they lack AI-related skills, and 66% fear falling behind peers who adopt AI tools at work (EY, 2023). These anxieties underscore the need for organisations to pair AI efficiency with human oversight and empathy, particularly in retail contexts where customer relationships and employee morale are central to brand value. The implication for fashion retailers is clear. Convenience alone is insufficient to secure consumer confidence; trust must be cultivated through transparency, security, and demonstrated respect for privacy and personal boundaries.

Fashion-retail leaders have responded by developing structured responsible-AI frameworks. H&M Group, for example, introduced a Responsible AI Framework built on nine principles, including fairness, transparency, and security, operationalised through a "Checklist for Responsible AI" applied to individual projects (H&M Group, 2021). The company also established an Ethical AI Debate Club, an internal forum where employees discuss fictional scenarios and ethical dilemmas, thereby preparing the organisation for future developments that may reshape the fashion industry (Gomes, 2025). In parallel, Stella McCartney collaborates with AI while using blockchain to document the sustainable materials in its products, reinforcing its ethical commitment through verifiable, transparent records (Kaur, 2026). Together, these initiatives illustrate concrete strategies for building trust by aligning AI deployment with human-centric values and governance structures.

Ethics of AI

Although the business-case rationale for AI in retail is compelling, the human-case dimension demands careful scrutiny. The way AI systems are designed, trained, and deployed significantly influences how they shape individuals' everyday experiences and broader social structures. If optimised solely for efficiency and productivity, AI can function as a tool for exploitation, reinforcing existing inequalities and eroding trust (Manchiraju, 2025). In an industry built on relationships, aesthetics, and emotional resonance, robust ethical standards are therefore essential to ensure that the future of fashion is both sustainable and socially equitable (Manchiraju, 2025).

A commonly used distinction in the literature is between "ethics of AI," which concerns the moral principles governing AI design and deployment, and "ethical AI," which refers to systems that actively embody those principles in practice (Siau & Wang, 2020). This two-level view is relevant to fashion retail, where the gap between stated values and operational reality is a recurring theme. In *The Oxford Handbook of Ethics of AI*, Dignum (2020) argues that the core of AI development must prioritise collective public benefit, reflecting what is commonly described as an "AI for Good" and "AI for All" ethos (ITU, 2026; AI4ALL, n.d.). This perspective calls for AI systems that benefit not just shareholders but also customers, employees, communities, and the environment. Historian and philosopher Yuval Noah Harari warns that AI can amplify human frailties, such as bias and shortsightedness, by encoding them into autonomous systems with near-divine power, thereby escalating the risks of unintended harm (Harari, 2016; Harari, 2024). In response, robust regulatory and ethical frameworks are needed to mitigate these existential and societal risks (Harari, 2024). AI does not act in a moral vacuum; its design

choices carry long-term consequences for fairness, inclusion, and democratic values.

The ethical stakes are not hypothetical; leading AI developers have publicly acknowledged the profound ethical dilemmas inherent in large-scale model deployment, ranging from life-or-death moderation choices to the protection of deeply sensitive user data (CNBC, 2025). What this means in practice is that the international AI-governance landscape is uneven, and many fashion retailers are trying to navigate ethical questions in a patchwork of local rules and emerging best practices.

Evolving Ethics in the Age of Agentic AI

The ethical landscape is further complicated by the rapid emergence of agentic AI, systems designed not only to execute predefined tasks but to autonomously identify problems, formulate strategies, and execute actions with minimal human intervention (Burnstine, 2025). In fashion retail, this may translate to systems that autonomously adjust pricing, manage inventory, and personalise customer interactions. However, this autonomy introduces significant risks, notably opaque decision-making and the erosion of accountability, which are well recognized in governance and risk management research (Kalia & Gill, 2023).

Agentic AI systems can operate without constant human attention, but responsible deployment does not require surrendering oversight. It requires designing oversight in from the start. By embedding explainability, audit trails, and human-in-the-loop controls into the AI architecture, organisations can maintain confidence in both the outcomes and the underlying reasoning, as highlighted in recent analyses of agentic AI within the fashion ecosystem (Burnstine, 2025). The EU AI Act formalises this expectation for high-risk systems, requiring that operators establish “appropriate human oversight measures” that are technically feasible before deployment (European Commission, 2024). In practice, this means fashion retailers must treat explainability not as a technical afterthought but as a design-time commitment, one that protects both the consumer and the brand.

The cybersecurity dimension of agentic AI is equally pressing. As AI systems gain the ability to access databases, execute transactions, and communicate with customers on behalf of the brand, the attack surface expands dramatically. A 2026 Microsoft report indicates that generative AI tools are involved in 32% of organisations’ data security incidents (Microsoft, 2026). In the fashion sector, this risk is compounded by the high value of intellectual property and the reputational sensitivity of customer data. The hacking of Gucci, Balenciaga, and Alexander McQueen in 2025 demonstrated that luxury and premium fashion brands are high-value targets (BBC News, 2025; Walsh et al., 2026). When brands shift from static e-commerce to autonomous, agentic AI systems that possess read/write access to customer databases, an identical breach does not merely leak historical data; it allows malicious actors to potentially hijack active customer-facing AI agents.

Ethical AI in Fashion Retail

The integration of AI into fashion retail, from design and production to marketing and customer service, raises serious ethical issues that demand urgent attention (Giovanola et al., 2023). These developments force the industry to confront profound moral questions about consumer privacy, environmental impact, algorithmic bias, the protection of human creativity, and the stability of employment across the sector (Giovanola et al., 2023). Table 2.1 summarises key concerns, underlying issues, and illustrative examples.

Table 2.1*Ethical Concerns Associated With AI in Fashion Retail*

Concern	Key Issue	Example
Consumer Privacy	Unauthorised collection and exposure of purchase and personal data undermine consumer trust and consent (Siau & Wang, 2020).	Gucci, Balenciaga, and Alexander McQueen were hacked in 2025, exposing millions of customers' personal data (BBC News, 2025).
Environmental Footprint	AI training and inference carry significant energy and water costs: a single ChatGPT query consumes approximately five times more electricity than a standard web search, and data centres may consume up to 1,000 TWh of electricity globally by 2026 (World Economic Forum, 2025b).	Google and Microsoft have both disclosed year-on-year increases in emissions from AI workloads, warning that current growth trajectories challenge their own climate commitments (World Economic Forum, 2025b).
Algorithmic Bias	AI models trained on non-representative data can propagate unfair discrimination in hiring, marketing, and product recommendations (European Union Agency for Fundamental Rights, 2022).	Amazon's AI résumé-screening tool was found to systematically downrank applications containing women-linked terms, amplifying gender bias at scale (Dastin, 2018).
Human Creativity	Generative AI tools trained on existing creative works risk eroding design originality and undermining intellectual property rights (Niyompatama & Lapatoura, 2025).	Aerie explicitly rejected AI-generated bodies in its 2026 campaign, foregrounding authentic human representation as a deliberate brand response to AI-generated imagery (O'Brien, 2026).
Employment Disruptions	Automation technologies can displace frontline and routine roles more rapidly than new opportunities emerge, with 92 million jobs projected to be displaced globally by 2030 (World Economic Forum, 2025a).	The WEF estimates that 40% of employers expect to reduce their workforce where AI can automate tasks, with retail among the most exposed sectors (World Economic Forum, 2025a).

These examples illustrate that ethical AI in fashion retail requires more than technical compliance; it demands a deliberate effort to align AI deployment with human dignity, environmental sustainability, and social justice (Giovanola et al., 2023). Retailers that integrate these concerns into their governance frameworks are better positioned to preserve the human touch while still harnessing the transformative potential of AI (Giovanola et al., 2023).

Furthermore, the use of generative AI in fashion design raises significant intellectual property concerns. Generative AI tools frequently leverage copyrighted patterns and existing creative works without explicit design attribution or compensation to the original creators. This practice challenges traditional frameworks of copyright ownership and necessitates new legal and ethical norms to protect human creativity while fostering innovation (Niyompatama & Lapatoura, 2025).

As the fashion industry seeks to leverage AI in pursuit of sustainability, the intersection of AI and sustainability presents a risk of AI-enabled greenwashing. While AI can enhance supply chain transparency, it can also be deployed to generate sophisticated but unverifiable sustainability claims. This risk is increasingly targeted by regulators in the EU and UK under

emerging green claims legislation, highlighting the need for AI to be used to substantiate, rather than merely assert, environmental credentials (Guglyuvatyy, 2026).

Gaps in Literature

A central gap in the literature is the absence of fashion-specific governance frameworks for AI, especially for agentic systems that design, recommend, and interact with customers autonomously (Burnstine, 2025). While current work stresses the importance of data protection, bias mitigation, and consumer consent, it tends to offer high-level principles without concrete, sector-specific guidelines for designers, marketers, and store managers (Giovanola et al., 2023; Burnstine, 2025). This gap is particularly acute given the growing divergence between AI governance principles articulated at the national level and their practical implementation within individual organisations (Stanford Institute for Human-Centered Artificial Intelligence, 2026). The present study addresses this gap by presenting findings from frontline store managers (Section 4.1), illustrating how high-level corporate intent is operationalised in daily workflows.

A second gap is the weak integration of ethics into human-centric AI design in fashion retail. Existing discussions acknowledge that AI reshapes aesthetics, cultural representation, and creative work, yet they rarely show how ethical reflection can be systematically embedded into the design and deployment of AI tools (Giovanola et al., 2023; Burnstine, 2025). In particular, agentic AI poses new questions about responsibility and explainability (Mokander & Floridi, 2021). This research links this gap directly to the strategic choice of implementing human-in-the-loop overrides, as discussed in Theme 3 and the final recommendations (Section 5.1).

A third gap lies in the under-explored intersection of sustainability and intellectual property risk. While the literature acknowledges AI's potential to replicate creative work and its capacity for supply chain optimisation, it offers few empirically grounded pathways for how retailers can experiment with different governance models to mitigate IP infringement and AI-enabled greenwashing while protecting human experience (Guglyuvatyy, 2026; OECD, 2026). Future research directions addressing this gap are proposed in Section 5.3. By filling these gaps, future research can help fashion retailers use AI not only to grow profits but also to deepen trust, inclusion, and meaningful human experiences.

Research Objectives and Questions

This study aims to address these gaps by examining the ethical, governance, and sustainability implications of AI adoption in the fashion industry through a comprehensive case study. The research seeks to answer the following questions:

- Research Question 1 (RQ1):
How does AI reshape retail and sustainability strategy at its core?

- Research Question 2 (RQ2):
What governance practices help maintain ethical standards as AI scales?

- Research Question 3 (RQ3):
How can brands preserve human connection in AI-mediated retail?

Through a real-world case study, this research explores the human experience and ethical considerations from top executives to frontline employees in the deployment of AI within the fashion industry. The following sections present the research framework, analyse the case study findings, and examine research gaps, implications, and limitations.

METHODOLOGY AND FRAMEWORK

To deeply investigate the ethical, regulatory, and human-centric challenges of scaling artificial intelligence within fashion retail, this research adopts a qualitative single-case study design (Yin, 2018). A qualitative paradigm was chosen because the core constructs of this investigation, including moral judgment, brand authenticity, and the human touch, are subjective, contextual, and not easily captured through quantitative measures (Creswell & Poth, 2018). The single-case design is appropriate because it allows for a close examination of a bounded, real-world organizational setting in which AI adoption is unfolding under strong commercial and ethical pressures.

The chosen case is a prominent Fortune 500 fashion brand, selected for its analytical richness rather than statistical representativeness. It offers a useful setting for exploring how a purpose-driven organization attempts to scale AI while preserving trust, employee judgment, and customer connection. The study is therefore designed to generate conceptual insights into responsible AI governance in fashion retail, rather than to produce generalizable claims about the wider industry.

Participant Selection and Data Gathering

The primary data source comprises ten semi-structured, individual interviews and one focus group with six assistant store managers at the target organization. One store manager (P5) participated in both the individual interview and the focus group; therefore, the total number of unique participants in the study was 16. Purposeful sampling was used to recruit participants whose roles directly intersect with AI implementation, corporate governance, cybersecurity, sustainability, and frontline retail experience. This functional spread is critical because it captures both the strategic imperatives driving AI adoption at the executive level and the operational realities experienced on the shop floor.

Data collection occurred between January and May 2026. Individual interviews lasted approximately 30 to 45 minutes. To accommodate participants and logistical needs, four interviews (P1, P2, P8, P9) were conducted via video conferencing (Zoom), while the remaining six individual interviews and the focus group were conducted in person. Interviews followed a flexible protocol organized around three central topics: how AI is reshaping retail strategy, how the organization manages ethical guardrails, and how it preserves human connection. The semi-structured format allowed participants to raise unanticipated issues, including shadow AI use, internal governance gaps, and the emotional effects of automation.

Where interviewees felt comfortable, three sessions (P5, P6, P10) were recorded and fully transcribed. For the remaining sessions, detailed verbatim and paraphrased notes were taken. This approach was a deliberate methodological choice to protect participant anonymity and encourage corporate candour, given the sensitive discussions. All data were subsequently verified for analytical readiness.

To protect identities and encourage candor, both the participants and the company have been anonymized. Participants are identified as P1 to P16 throughout the chapter. The composition of the participant sample is detailed in Table 3.1.

Table 3.1*Overview of Research Participants*

Identifier	Job Title / Functional Focus	Format
P1	Cybersecurity and Technology	Individual Interview (Zoom)
P2	General Counsel, Legal	Individual Interview (Zoom)
P3	Creative Direction	Individual Interview (In Person)
P4	Information and Security	Individual Interview (In Person)
P5	Store Manager	Individual Interview (In Person, Recorded)
P6	Regional Manager	Individual Interview (In Person, Recorded)
P7	Senior Architect	Individual Interview (In Person)
P8	Sustainability	Individual Interview (Zoom)
P9	Data, AI Engineering and AI Research	Individual Interview (Zoom)
P10	AI and Technology	Individual Interview (In Person, Recorded)
P11	Assistant Store Manager	Focus Group (In Person)
P12	Assistant Store Manager	Focus Group (In Person)
P13	Assistant Store Manager	Focus Group (In Person)
P14	Assistant Store Manager	Focus Group (In Person)
P15	Assistant Store Manager	Focus Group (In Person)
P16	Assistant Store Manager	Focus Group (In Person)

Note. P5 participated in both an individual interview and the focus group. Zoom interviews: P1, P2, P8, P9. Recorded with transcription: P5, P6, P10.

Analytical Approach: Reflexive Thematic Analysis

The data were analysed using Reflexive Thematic Analysis (RTA), following Braun and Clarke's (2019, 2021) six-phase approach. RTA was selected because it treats analysis as an interpretive and reflexive process rather than a purely mechanical coding exercise. This was appropriate for a study in which meaning, judgment, and organisational culture are central to the research question.

The analysis proceeded in six stages. First, the researcher immersed themselves in the dataset through repeated reading and note-taking. Second, the data were coded in an iterative, organic process where codes evolved to capture recurring ideas, tensions, and patterns relevant to the research questions, rather than being constrained by a fixed rubric. Third, candidate themes were developed by clustering related codes. Fourth, these themes were reviewed against the full dataset to ensure coherence and explanatory value. Fifth, the themes were named and defined to clarify their analytical contribution. Sixth, the findings were written up through a process of interpretation, comparison with the literature, and integration of anonymised participant quotations. Because RTA explicitly rejects mechanical codebooks or inter-rater reliability metrics, the analysis relied on the researcher's active, subjective interpretation of the data. Thematic saturation was reached across the dataset, with later interviews confirming rather than extending the emerging themes (Braun & Clarke, 2021).

The researcher's professional background in retail was treated as a source of insight rather than a bias to be eliminated. Reflexivity was maintained through ongoing memo-writing and critical self-reflection during the analysis process. For the focus group data, the analysis

explicitly accounted for group dynamics, capturing moments of consensus and disagreement among participants to contextualise their collective experience. To account for the overlap of P5, their individual interview data was analysed first to establish a baseline perspective before examining how their views interacted with the focus group dynamic. This approach helped surface assumptions while preserving the contextual depth needed for a credible interpretation of the data.

Trustworthiness and Ethics

The study was designed to strengthen trustworthiness through specific procedural safeguards based on Lincoln and Guba’s (1985) evaluative criteria. Credibility was supported by prolonged engagement with the case, careful questioning during interviews, and the use of rich, descriptive accounts in the reporting of findings (Lincoln & Guba, 1985). Dependability was supported through reflexive memo-writing and transparent documentation of the analytic process. Confirmability was strengthened through reflexive journaling and systematic comparison of codes and themes across the dataset (Braun & Clarke, 2021; Lincoln & Guba, 1985). Transferability was supported by providing sufficient contextual detail for readers to judge relevance to similar fashion retail settings (Lincoln & Guba, 1985).

Rigorous ethical standards were maintained throughout the research process. All participants were informed of the study’s purpose and provided verbal or written consent prior to participation. They were told they could withdraw at any time without penalty. Because the study addresses internal governance practices and strategic concerns, careful anonymisation was essential to protect both individual participants and the organisation. The research therefore prioritised confidentiality, respect, and non-maleficence at every stage.

DISCUSSION AND FINDINGS

This section presents the findings from the Reflexive Thematic Analysis (RTA) of ten semi-structured interviews and one focus group conducted with executives and store managers at a leading Fortune 500 fashion retail brand. The analysis was guided by three research questions: how AI reshapes retail and sustainability strategy at its core (RQ1); what governance practices help maintain ethical standards as AI scales (RQ2); and how brands can preserve human connection in AI-mediated retail (RQ3).

The RTA process generated four overarching themes that together capture the central tension of this study: the immense efficiency potential of AI set against the ethical, cultural, and relational imperatives of a purpose-driven brand.

Table 4.1

Overview of Themes, Research Questions, and Key Participants

Theme	Primary RQ Addressed	Key Participant Perspectives
Theme 1: AI as an Amplifier of Intent, Not a Substitute for Strategy	RQ1, RQ3	P5, P6, P11–P16
Theme 2: The Governance Gap and the Double Standard of Machine Bias	RQ2	P1, P2, P4, P7, P9, P10

Theme 3: Redefining the Human Touch from Transaction to Connection	RQ3	P3, P5, P11–P16
Theme 4: The Sustainability Paradox: Efficiency vs. Overconsumption	RQ1, RQ2	P3, P8, P9

Theme 1: AI as an Amplifier of Intent, Not a Substitute for Strategy

The most consistent finding across the dataset is that AI has not yet fundamentally reshaped the company’s core strategy. Instead, participants described AI as a tool that reveals and amplifies existing intent, provided that intent is clear and purposeful to begin with. This finding aligns with Brynjolfsson, Rock, and Syverson (2019) and Davenport and Mittal (2023), who argue that AI creates value not by replacing strategic thinking, but by removing the friction that prevents people from doing their best thinking. The most articulate expression of this idea came from a store manager with direct experience of AI adoption at the frontline:

“AI reshaped us by allowing us to return to our purpose by removing the distractions... It’s not a structured tool like Canva or Photoshop where a company decides what’s possible. It’s open, and what it produces is entirely down to the person curating it. If your thinking is purposeful, AI amplifies that. If it isn’t, AI just makes the average faster.” (P5)

This framing of AI as a mirror of human intention is a significant departure from the dominant discourse, which tends to frame AI as either a disruptive threat or a transformative force in its own right.

At the store level, the focus group participants confirmed this view. AI is being used for scheduling, drafting communications, translating guest emails into local language, and synthesising performance data, but always with a human in the loop. A critical barrier to realizing this potential is what participants described as a lack of “technological mindset” within the organization. P6 observed that “bringing in technical people does not mean you are going to change the company. You need to convince everyone else to be more technical and embrace the tools.” The proposed solution is a bottom-up adoption strategy: start with small teams solving small problems, demonstrate the value, and let the behaviour spread organically. This approach resonates with the “lighthouse-project” model discussed in the World Economic Forum’s Global Lighthouse Network, in which early, high-impact AI pilots are used to build capability and buy-in for broader transformation (World Economic Forum, 2023).

In short, the returns on AI investment are inseparable from the clarity of the human intent behind it: organisations that articulate a purposeful strategy before deploying AI will consistently outperform those that treat it as a productivity shortcut.

Theme 2: The Governance Gap and the Double Standard of Machine Bias

If Theme 1 addresses what AI can do for the enterprise, Theme 2 establishes the ethical and regulatory guardrails under which it must be deployed. Across multiple organizational levels, participants described a stark institutional reality: technological capability is outpacing internal governance. The core challenge facing the organization is not an absence of AI enthusiasm, but a deficit in governance readiness. The most concrete evidence of this gap came from a data and AI engineering leader, who acknowledged that the company is “not good at guardrails” and cited a regulatory fine paid to authorities for a failure in customer communications that, while not AI-related, illustrated the real-world cost of inadequate oversight (P9). This

finding is consistent with the growing body of regulatory and industry literature warning that companies are deploying AI faster than they are building the governance structures to manage it (European Commission, 2021; KPMG, 2023).

This governance vacuum directly manifests as “shadow AI,” an operational risk where employees leverage unsanctioned, public AI tools on personal devices because corporate alternatives are too restrictive to be commercially useful. As one regional manager observed: “The company gives us this tool that people do not even want to use. Eventually, people put it on their own device, which compromises data security” (P6). The cybersecurity lead confirmed this risk, noting that employees at comparable retailers have been observed inputting confidential data into public AI systems (P1). This dynamic creates a sharp irony: risk-averse governance policies designed to protect corporate intellectual property inadvertently generate the exact security vulnerabilities they were engineered to prevent. This aligns with recent warnings that frontline employee adoption routinely bypasses slow-moving corporate compliance protocols, widening institutional vulnerabilities (Puthal et al., 2025; Walsh et al., 2026).

Table 4.2

The Anatomy of Corporate AI Friction: Rhetoric vs. Operational Reality

Domain	Corporate Rhetoric / Policy	Operational Reality (Employee Experience)
Tool Adoption	Employees should use sanctioned, secure corporate AI tools (e.g., Copilot).	Sanctioned tools are overly restricted; employees use “shadow AI” (ChatGPT, Claude) on personal devices.
Security	Strict data governance protects intellectual property and customer privacy.	Overly strict governance drives behaviour underground, creating unmonitored data vulnerabilities.
Performance	AI should be used to increase efficiency and productivity, with employees expected to review and own the output.	Employees who use AI and review the output may still face disciplinary action if the AI-generated language is not adapted to corporate terminology, conflating tool use with a failure of diligence.

A second and more philosophically significant finding within this theme is what P10 termed the “double standard” applied to AI systems. Drawing on his background in machine learning, P10 argued that human bias is fundamentally unmeasurable, difficult to correct, and nearly impossible to audit, yet society tolerates it. AI bias, by contrast, is mathematically measurable, correctable, and auditable, yet it is held to a standard of perfection that no human system meets. While P10’s concern about unequal scrutiny is worth taking seriously, this observation reveals a profound corporate opportunity. Because machine bias is auditable and measurable, retailers have a unique capability to use AI as a diagnostic tool to identify and correct systemic human biases within their operations. Rather than viewing the standard of perfection as merely a “double standard,” responsible retailers can leverage AI’s transparency to build more equitable systems than human judgment alone could achieve (Binns, 2018; Raji et al., 2020).

The governance gap is also a cultural problem. P7’s account illustrates this with precision: having used AI to research a problem, P7 reviewed the output before submitting it to a manager, but did not translate the AI-generated response from generic industry terminology into the company’s specific corporate language. The manager interpreted this as evidence that P7 had neither done the work nor checked it, and P7 received a performance improvement plan as a result. The incident reveals a deeper institutional failure: in the absence of clear guidance on how AI-assisted work should be presented, employees are penalized not for negligence, but

for the visibility of their tool use. “Don’t want to pretend I’m not using AI,” P7 said. “Someone should announce: use AI and check the work.” The absence of psychologically safe, explicit norms around AI use is itself an ethical failure, one that conflates AI use with a lack of diligence and forces employees to choose between efficiency and the performance of effort.

The governance gap is therefore not merely a compliance risk; it is a cultural and ethical failure that, left unaddressed, will erode both employee trust and the brand values that differentiate a purpose-driven retailer from its competitors.

Theme 3: Redefining the Human Touch from Transaction to Connection

As AI becomes capable of handling an increasing share of routine tasks, the question of what remains distinctively human becomes both more urgent and more answerable. The findings from this study suggest a clear and consistent answer: human connection is not transactional. It is relational, grounded in intention, community, and lived experience.

P3, the company’s creative director, articulated the strategic stakes with characteristic directness: “GenZ is allergic to AI. They don’t believe it and can’t tell what’s true. Real physical touch is more valuable now. Can’t take away what a chef can do. Not live music, comedy, theatre. We must premiumise human experience.” This observation is supported by recent consumer research suggesting that authenticity and human connection are becoming premium attributes in a world saturated with AI-generated content (Kirkby et al., 2023). The company’s decision not to use AI-generated imagery for e-commerce photography is a concrete expression of this principle.

At the store level, the focus group data reveals a more nuanced picture. Participants expressed genuine enthusiasm for AI tools that remove administrative burden, precisely because they free up time for the human interactions that matter most. “Before I wouldn’t reply to a guest in local language by email, and now I get so confident to do it,” one assistant store manager noted, describing how AI-assisted translation had expanded her ability to connect authentically with customers (P11). This is authentic human connection redefined: not the absence of AI, but the presence of a more capable, more confident human, enabled by AI.

The most sophisticated articulation of this theme came from P5, who proposed that the company’s human advantage lies not in any single interaction, but in the accumulated, irreplaceable depth of its community and institutional knowledge. “What they can’t copy is why we make the product, who we make it for, and the intention behind all of it. That’s where our community originates from. Many years of that. You can’t duplicate it.” The proposed solution to transform the company’s internal knowledge platform into a “shared brain” that preserves and distributes this intent is a practical expression of the principle that AI should serve human connection, not replace it.

This theme also surfaces an important concern about the long-term erosion of foundational human skills. P2, the company’s legal counsel, warned that junior staff who rely on AI for research may never develop the critical reasoning and advocacy skills that define excellence in their field. This concern is not unique to law; it reflects a broader risk of AI-induced deskilling, where overreliance on automated systems may gradually erode critical reasoning and domain expertise over time (Ferdman, 2026).

The strategic implication is clear: fashion retailers that invest in AI as a tool for human enablement, rather than human replacement, will build the kind of authentic brand equity that no algorithm can replicate.

Theme 4: The Sustainability Paradox: Efficiency vs. Overconsumption

The intersection of AI and sustainability in the fashion industry presents a genuine paradox. AI offers powerful tools for reducing waste, accelerating material innovation, and improving supply chain transparency, yet the same efficiency gains that reduce operational waste can also create conditions that encourage higher overall consumption, reflecting a sustainability paradox that is particularly acute in fashion retail: AI can improve efficiency while also enabling the systemic conditions that sustain higher overall consumption (Alcott, 2005; Bieñkowska, 2025).

On the positive side, the findings reveal several concrete applications of AI in the company's sustainability strategy. P8, the sustainability leader, described the use of AI to analyse many years of supply chain audit data and predict where compliance problems are likely to emerge before they occur. AI is also being used to generate novel enzyme molecules for the enzymatic recycling of nylon and polyester, a breakthrough that could fundamentally change the economics of circular fashion (P8). These applications align with the broader literature showing that AI can support circular economy transitions in fashion and textiles by improving design, process optimisation, and recyclability assessment (Nisa et al., 2025).

However, the findings also surface a more uncomfortable truth. P9 acknowledged that while AI can optimise operations, it can equally be used to drive more sales: "Presumptuous for us to say, in the business of generating more revenue... The role of AI is helping make the right choice." The tension is structural: a company whose business model depends on selling more products cannot fully resolve the contradiction between growth and sustainability through efficiency alone.

This finding exemplifies Jevons' Paradox, an economic principle positing that technological advancements which increase resource efficiency tend to increase, rather than decrease, the overall consumption of that resource (Alcott, 2005). In the context of AI in fashion retail, efficiency gains that lower the relative cost of production or transaction can inadvertently fuel aggregate consumption, offsetting the environmental benefits of optimized supply chains. P3 offered a more radical reframing to address this: "Most sustainable product is what we use the most. Highly functional, so you will use it. Not a regenerative cotton couture dress worn once." This shifts the sustainability metric from material inputs to product longevity and utility, a framing that AI-powered personalization could actively support.

The resale programme illustrates both the promise and the tension of this theme. P8 acknowledged that the fashion business model "creates more sales, which is at odds with encouraging" reduced consumption. Yet circularity offers an end-of-life solution: it generates storytelling opportunities, connects customers to the brand's circular values, and gives used products a second life. The moral compass here requires the company to ask not just whether AI makes the business model more efficient, but whether it makes the business model more honest, more aligned with the brand's stated commitment to doing less harm.

P8's framing of the ethical test applies equally to AI governance and sustainability strategy: "Ask yourself: what do you stand for? [It requires a] gut check. Is it consistent with the values we have as a brand?... Lots of brands said sustainability was important to them, when times get tough or politics changes, it's not important anymore. Have to constantly, as a senior leader, reaffirm what the values are."

Ultimately, the sustainability paradox can only be resolved through deliberate governance: retailers must set explicit environmental guardrails on their AI systems and hold themselves accountable to those commitments with the same rigour they apply to financial targets.

CONCLUSION AND RECOMMENDATIONS

The integration of artificial intelligence into fashion retail represents a socio-technical transformation that challenges foundational assumptions about creativity, efficiency, and the nature of human connection. This study set out to explore how a major fashion brand navigates the ethical, governance, and sustainability implications of this shift. The findings reveal that AI is not a substitute for strategic vision or moral judgment; rather, it is a powerful amplifier of existing organizational intent. While AI offers remarkable tools to reduce friction and accelerate sustainability initiatives, its unchecked deployment risks widening governance gaps, eroding consumer trust, and exacerbating the paradox of efficiency-driven overconsumption. To harness the benefits of AI while preserving authentic human experience, fashion retailers must move beyond treating ethics as a compliance exercise. Responsible AI requires deliberate design choices that prioritize transparency, human oversight, and genuine relational integrity.

Strategic Recommendations

Based on the literature review and case study findings, the following recommendations are proposed for retail leaders.

First, organizations should close the governance gap with “safe sandboxes.” To address the shadow AI problem identified in Theme 2, organizations must provide employees with secure, internally sanctioned AI environments. These tools must be genuinely useful rather than overly restrictive but must be enclosed within corporate firewalls to protect proprietary data and consumer privacy.

Second, human-in-the-loop controls should be implemented as a core design principle. To preserve the relational quality highlighted in Theme 3, AI should be deployed to augment human decision-making rather than replace it. Store managers and creative teams must retain the authority to override algorithmic recommendations, ensuring that the brand’s moral compass remains in human hands.

Third, retailers should create visible governance and oversight structures. To build trust with both consumers and employees, retailers must establish clear, auditable decision logs and independent ethics review boards. The governance structure must be transparent, demonstrating a commitment to fairness and accountability that aligns with emerging regulations like the EU AI Act.

Fourth, guardrails must be set for recommendation engines and demand shaping. To navigate the sustainability paradox identified in Theme 4, retailers must evaluate AI deployments against environmental metrics, not just revenue growth. AI should be directed toward solving complex circularity challenges, such as supply chain transparency and material recycling, and strict guardrails should be placed on recommendation algorithms to prevent them from blindly driving overconsumption.

Limitations

This study is subject to several limitations. First, as a single-case study focused on one Fortune 500 brand, the findings may not be entirely generalisable to smaller retailers, luxury heritage houses, or fast-fashion conglomerates, which operate under different economic and cultural pressures. Second, the sample is heavily weighted toward executive and managerial perspectives, introducing possible social desirability bias and omitting the direct, lived experiences of the end consumers interacting with these AI systems. Finally, the rapid pace of AI

development means that the technological capabilities and regulatory environments discussed are continuously evolving.

Future Research

Future research should address these limitations by exploring the following areas. Longitudinal studies tracking how fashion retailers adapt their internal AI policies over time, in response to the enforcement of the EU AI Act and emerging global regulations, would provide valuable insight into governance maturity. Quantitative research examining how fashion consumers respond to autonomous AI agents handling customer service or personal styling would help identify the threshold at which “personalisation” becomes “intrusive.” Investigations into how AI-driven compliance and auditing tools affect the labor rights and working conditions of garment workers in producing countries would extend the ethical focus beyond the Western retail floor. Finally, the legal and ethical analysis of how fashion brands can establish fair copy-right practices when using generative AI models trained on existing creative works remains an urgent, under-examined priority.

The fashion industry has long served as a cultural barometer, reflecting the values, anxieties, and aspirations of the societies it inhabits. As artificial intelligence becomes embedded in how brands design, sell, and engage with their customers, the ethical stakes of that reflection intensify. The findings of this study suggest that the defining question is no longer whether AI will transform fashion retail, but whether organisations possess the strategic clarity, governance maturity, and relational commitment to ensure that transformation remains oriented toward human flourishing rather than operational convenience. AI does not determine what a brand stands for; it amplifies what is already present. The responsibility for what gets amplified remains, as it always has, entirely human.

REFERENCES

- Afroogh, S., Akbari, A., Malone, E., Kargar, M., & Alambeigi, H. (2024). Trust in AI: Progress, challenges, and future directions. *Humanities and Social Sciences Communications*, *11*(1), Article 1568. <https://doi.org/10.1057/s41599024040448>
- AI4ALL. (n.d.). *AI4ALL – Making AI accessible and inclusive for all*. <https://aifourall.org>
- Alcott, B. (2005). Jevons' paradox. *Ecological Economics*, *54*(1), 9–21. <https://doi.org/10.1016/j.ecolecon.2005.03.020>
- BBC News. (2025, September 15). Gucci, Balenciaga and Alexander McQueen hacked in cyber-attack. BBC News. <https://www.bbc.com/news/articles/cr15j8ld615o>
- Bieńkowska, J. (2025). The effects of artificial intelligence on the fashion industry: Opportunities and challenges for sustainable transformation. *Sustainable Development*, *33*(3), 3774–3790. <https://doi.org/10.1002/sd.3312>
- Binns, R. (2018). Fairness in machine learning: Lessons from political philosophy. In *Proceedings of the Conference on Fairness, Accountability and Transparency* (pp. 149–159). PMLR.
- Braun, V., & Clarke, V. (2019). Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health*, *11*(4), 589–597. <https://doi.org/10.1080/2159676X.2019.1628806>
- Braun, V., & Clarke, V. (2021). One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative Research in Psychology*, *18*(3), 328–352. <https://doi.org/10.1080/14780887.2020.1769238>
- Brynjolfsson, E., Rock, D., & Syverson, C. (2019). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 23–60). University of Chicago Press. <https://doi.org/10.7208/chicago/9780226613475.003.0001>
- Burnstine, A. (2025). Autonomous intelligence in fashion: A comprehensive analysis of agentic AI across the fashion ecosystem. *Asian Business Research Journal*, *10*(4), 31–37. <https://doi.org/10.55220/25766759.405>
- CBS News. (2026, February 23). Hegseth demands full military access to Anthropic's AI model Claude and sets deadline for end of week. CBS News. <https://www.cbsnews.com/news/hegseth-anthropic-full-access-claude-ai-model/>
- Chen, Y., & Prentice, C. (2025). Integrating artificial intelligence and customer experience: A conceptual framework of AI applications in customer experience. *Australasian Marketing Journal*, *33*(2), 141–153. <https://doi.org/10.1177/14413582241252904>
- CNBC. (2025, September 15). Why is Sam Altman losing sleep? OpenAI CEO addresses controversies in sweeping interview. CNBC. <https://www.cnbc.com/2025/09/15/sam-altman-losing-sleep-open-ai-ceo-addresses-controversies-interview.html>
- Creswell, J. W., & Poth, C. N. (2018). *Qualitative inquiry and research design: Choosing among five approaches* (4th ed.). Sage.
- Dastin, J. (2018, October 9). *Amazon scraps secret AI recruiting tool that showed bias against women*. Reuters. <https://www.reuters.com/article/world/insight-amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK0AG/>
- Davenport, T., & Mittal, N. (2023). *All in on AI: How smart companies win big with artificial intelligence*. Harvard Business Review Press.
- Dignum, V. (2020). Responsibility and artificial intelligence. In M. Dubber, F. Pasquale, & S. Das (Eds.), *The Oxford handbook of ethics of AI* (pp. 214–231). Oxford University Press.

- European Commission. (2021). Proposal for a regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>
- European Commission. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union. <https://artificialintelligenceact.eu>
- European Union Agency for Fundamental Rights. (2022). *Bias in algorithms – Artificial intelligence and discrimination*. Publications Office of the European Union. <https://fra.europa.eu/en/publication/2022/bias-algorithm>
- EY. (2023). AI anxiety in business survey. <https://www.ey.com/content/dam/ey-unified-site/ey-com/en-us/campaigns/consulting/documents/ey-us-anxiety-pdf-v7-final.pdf>
- Ferdman, A. (2026). *AI deskilling is a structural problem*. *AI & Society*. <https://doi.org/10.1007/s00146-025-02686-z>
- Fontaine, T., McCarthy, B., & Saleh, T. (2019, July–August). Building the AI-powered organisation. *Harvard Business Review*, 97(4), 62–73. <https://hbr.org/2019/07/building-the-ai-powered-organisation>
- Giovanola, B., Tiribelli, S., Frontoni, E., & Paolanti, M. (2023). Ethical implication of artificial intelligence in the fashion industry: A comprehensive analysis. *Fashion Highlight*, (2), 22–28. <https://doi.org/10.36253/fh-2497>
- Gomes, G. (2025, November 20). H&M's AI playbook: The tech strategy behind its transformation. *CTO Magazine*. <https://ctomagazine.com/ai-at-hm-the-tech-strategy-behind-its-transformation/>
- Grewal, D., Roggeveen, A. L., & Nordfält, J. (2017). The future of retailing. *Journal of Retailing*, 93(1), 1–6. <https://doi.org/10.1016/j.jretai.2016.12.008>
- H&M Group. (2021). Responsible AI. <https://hmgroup.com/our-stories/responsible-ai-is-better-ai/>
- Harari, Y. N. (2016, August 27). *We are quickly acquiring powers that were always thought to be divine*. *The Guardian*.
- Harari, Y. N. (2024). *Nexus: A brief history of information networks from the Stone Age to AI*. Penguin Random House.
- International Telecommunication Union. (2026). *AI for Good*. <https://aiforgood.itu.int/>
- Guglyuvatyy, E. (2026). Artificial integrity? Legal, ethical and technological perspectives on AI greenwashing. *World Development Sustainability*, 8, Article 100286. <https://doi.org/10.1016/j.wds.2026.100286>
- Kalia, A., & Gill, S. (2023). Corporate governance and risk management: A systematic review and synthesis for future research. *Journal of Advances in Management Research*, 20(3), 409–461.
- Kaur, R. (2026). Artificial intelligence for sustainable fashion design: Optimizing processes and reducing waste in the supply chain. In *Proceedings of the International Conference on Creativity, Innovation & Design (ICCID 2025)* (pp. 69–83). Atlantis Press. https://doi.org/10.2991/978-2-38476-535-5_7
- Kirkby, A., Baumgarth, C., & Henseler, J. (2023). *To disclose or not disclose, is no longer the question: Effect of AI-disclosed brand voice on brand authenticity and attitude*. *Journal of Product & Brand Management*, 32(7), 1108–1122. <https://doi.org/10.1108/JPBM-02-2022-3864>
- KPMG. (2023). Thriving in an AI world: KPMG global tech report 2023. <https://kpmg.com/xx/en/home/insights/2023/09/kpmg-global-tech-report-2023.html>
- KPMG. (2025). 2025 global CEO outlook: Consumer and retail. <https://assets.kpmg.com/con>

- tent/dam/kpmgsites/uk/pdf/2026/02/global-ceo-outlook-report.pdf
- Gillespie, N., Lockey, S., Ward, T., Macdade, A., & Hased, G. (2025). *Trust, attitudes and use of artificial intelligence: A global study 2025*. The University of Melbourne and KPMG. <https://doi.org/10.26188/28822919>
- KPMG. (2026). AI in retail: Global lessons from strategy to storefront. <https://assets.kpmg.com/content/dam/kpmgsites/uk/pdf/2026/01/ai-in-retail.pdf>
- Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. Sage.
- Manchiraju, S. (2025). The ethical thread: AI's role in the tapestry of fashion. *AI & Society*, 40(5), 4029–4034. <https://doi.org/10.1007/s00146-024-02151-3>
- Microsoft. (2026). 2026 data security index: Unifying data protection and AI innovation. Microsoft. <https://marketingassets.microsoft.com/gdc/gdccYW33N/original>
- Mokander, J., & Floridi, L. (2021). Ethicsbased auditing to develop trustworthy AI. *Minds and Machines*, 31(3), 323–327. <https://doi.org/10.1007/s11023021095578>
- NBC News. (2026, March 26). Judge blocks Trump administration from limiting Anthropic's government contracts. NBC News. <https://www.nbcnews.com/news/us-news/anthropic-trump-national-security-rcna265399>
- Nisa, H., Van Amber, R., English, J., Islam, S., McCorkill, G., & Alavi, A. (2025). A systematic review of reimagining fashion and textiles sustainability with AI: A circular economy approach. *Applied Sciences*, 15(10), 5691. <https://doi.org/10.3390/app15105691>
- Niyompatama, M. J., & Lapatoura, I. (2025). Generative AI in fashion design creation: A copyright analysis of AI-assisted designs. *Journal of Intellectual Property Law & Practice*, 20(10), 654–666. <https://doi.org/10.1093/jiplp/jpaf045>
- O'Brien, A. (2026, March 26). Aerie takes a stand against AI marketing with Pamela Anderson. *Vogue*. <https://www.vogue.com/article/aerie-takes-a-stand-against-ai-marketing-with-pamela-anderson>
- Organisation for Economic Co-operation and Development. (2021). *Principles on artificial intelligence*. OECD AI Policy Observatory. <https://oecd.ai/en/aiprinciples>
- Organisation for Economic Cooperation and Development. (2026). *Due diligence guidance for responsible AI*. OECD AI Policy Observatory. <https://oecd.ai/en/aipolicyobservatory/duediligenceguidanceforresponsibleai>
- Puthal, D., Mishra, A. K., Mohanty, S. P., Longo, A., & Yeun, C. Y. (2025). Shadow AI: Cyber security implications, opportunities and challenges in the unseen frontier. *SN Computer Science*, 6(5), 405.
- PwC. (2024, May 13). The EU AI Act and its impact on businesses. <https://www.pwc.ie/services/audit-assurance/insights/eu-ai-act-business-impact.html>
- Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 33–44). <https://doi.org/10.1145/3351095.3372873>
- Siau, K., & Wang, W. (2020). Artificial intelligence (AI) ethics: Ethics of AI and ethical AI. *Journal of Database Management*, 31(2), 74–87. <https://doi.org/10.4018/JDM.2020040105>
- Stanford Institute for Human-Centered Artificial Intelligence. (2026, January 12). Translating centralized AI principles into localized practice. Stanford HAI. <https://hai.stanford.edu/news/translating-centralized-ai-principles-into-localized-practice>
- The Business of Fashion. (2026, April 12). Case study: The fashion marketer's guide to AI. <https://www.businessoffashion.com/case-studies/marketing-pr/fashion-marketing-ai-guide-brand/>

- Trist, E. (1981). The evolution of socio-technical systems: A conceptual framework and an action research program. Ontario Quality of Working Life Centre.
- Walsh, D., Doshi, N., & Clemente, A. (2026, February). Navigating security for AI: The view from the front lines [Webinar]. Harvard Business Review Analytic Services. <https://hbr.org/sponsored/2026/02/navigating-security-for-ai-the-view-from-the-front-lines-webinar>
- World Economic Forum. (2023). *Global Lighthouse Network: Adopting AI at speed and scale*. World Economic Forum. https://www3.weforum.org/docs/WEF_Global_Lighthouse_Network_Adopting_AI_at_Speed_and_Scale_2023.pdf
- World Economic Forum. (2025a). The future of jobs report 2025. https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
- World Economic Forum. (2025b, June 1). How to cut the environmental impact of your company's AI use. <https://www.weforum.org/stories/2025/06/how-ai-use-impacts-the-environment/>
- Yin, R. K. (2018). Case study research and applications: Design and methods (6th ed.). Sage.