UNIVERSITY
OF TURKU
Turku School of
Economics

Augmented Analytics in Supply Chain Management

Reducing Decision Latency Through Reliability and Human Trust

Information Systems Science
Bachelor's thesis

Author:
Sebastian Nieminen

Supervisor:
Tarja Matikka

08.05.2026
Turku

Student’s statement regarding the use of Artificial Intelligence (AI) for preparing and/or writing this thesis:

☐ **I have not used any AI-based tools.**

☒ **I have used AI-based tools**. Their use is documented in the Appendix. The AI tools were used in a way that complies with academic integrity guidelines.

Bachelor's thesis

Subject: Information Systems Science
Author: Sebastian Nieminen
Title: Augmented Analytics in Supply Chain Management
Supervisor: Tarja Matikka
Number of pages: 33 pages (+ appendices 3 pages)
Date: 08.05.2026

Abstract

Modern supply chain management operates under structural volatility where disruptions generate information processing demands that exceed organizational capacity. This leads to decision latency that compounds operational damage. This thesis examines how augmented analytics can reduce decision latency for operational supply chain managers. It further identifies the factors that determine the reliability and trustworthiness of AI-generated outputs. The study is conducted as a systematic literature review across information systems and supply chain management.

The research utilizes Organizational Information Processing Theory, which posits that the mismatch between information processing demands and organizational capacity constitutes an information processing gap. The findings indicate that augmented analytics reduces decision latency through two primary mechanisms. These are the compression of analysis latency via automated data preprocessing and the reduction of interpretive effort through natural language technologies that broaden access to analytical insights beyond specialist users.

However, technical capability alone does not ensure organizational value. The research reveals a trust gap that can preserve decision latency even when system outputs are accurate. This gap is driven by algorithmic aversion, representational load, and a tendency of users to follow recommendations without genuine trust in their basis. The thesis identifies Explainable Artificial Intelligence and Human-in-the-Loop design as critical organizational interventions that facilitate the bridging of the trust gap by enhancing transparency. These approaches support human decision authority within the process. Ultimately, the study suggests that the reduction of analysis latency translates most effectively into operational responsiveness when technical speed is supported by human accountability and trust.

Keywords: augmented analytics, decision latency, supply chain management, information processing gap, explainable AI, human-in-the-loop, algorithmic aversion, trust

Kandidaatintutkielma

Oppiaine: Tietojärjestelmätiede
Tekijä: Sebastian Nieminen
Otsikko: Lisätty analytiikka osana toimitusketjujen johtamista
Ohjaaja: Tarja Matikka
Sivumäärä: 33 sivua (+ liitteet 3 sivua)
Päivämäärä: 08.05.2026

Tiivistelmä

Toimitusketjujen hallinta kohtaa jatkuvaa epävakautta, jossa häiriötilanteet synnyttävät organisaation kapasiteetin ylittäviä tiedonkäsittelytarpeita. Tämä johtaa päätöksenteon viiveisiin, jotka syventävät operatiivisia vahinkoja. Tutkielma tarkastelee, miten lisätty analytiikka voi lyhentää päätöksenteon viiveitä operatiivisten toimitusketjupäälliköiden työssä. Lisäksi tunnistetaan tekijöitä, jotka määrittävät tekoälyjärjestelmien tuottamien vastausten luotettavuutta ja käyttäjän luottamusta. Tutkimus on toteutettu systemaattisena kirjallisuuskatsauksena tietojärjestelmätieteen ja toimitusketjujen johtamisen aloilta.

Tutkielma hyödyntää organisaation tiedonkäsittelyteoriaa, jonka mukaan häiriöiden synnyttämien tiedonkäsittelytarpeiden ja organisaation kapasiteetin välinen epäsuhta muodostaa tiedonkäsittelykuilun. Tulokset osoittavat, että lisätty analytiikka lyhentää päätöksenteon viiveitä automatisoimalla datan esikäsittelyä ja hyödyntämällä luonnollisen kielen teknologioita, jotka laajentavat analyyttisten tulosten saatavuutta muille kuin asiantuntijakäyttäjille.

Tekninen suorituskyky ei kuitenkaan yksin takaa organisatorista arvoa. Tutkimus tunnistaa luottamuskuilun, joka voi ylläpitää päätöksenteon viiveitä tarkimmistakin tuloksista huolimatta. Tämä kuilu johtuu algoritmiaversiosta, tulosten esitystapaan liittyvästä kognitiivisesta kuormituksesta sekä siitä, että käyttäjät saattavat noudattaa järjestelmän suosituksia ilman todellista luottamusta niiden perusteisiin. Selitettävä tekoäly ja human-in-the-loop -suunnittelu tarjoavat organisatorisia ratkaisuja, jotka voivat kuroa umpeen luottamuskuilua lisäämällä järjestelmän päättelyn läpinäkyvyyttä. Nämä lähestymistavat tukevat ihmisen päätösvaltaa prosessin aikana. Tutkielman johtopäätöksenä todetaan, että analyysiviiveen lyhenemisen muuntuminen operatiiviseksi reagointikyvyksi näyttää edellyttävän teknisen nopeuden ohella inhimillistä hallintaa ja luottamusta.

Avainsanat: lisätty analytiikka, päätöksenteon viive, toimitusketjujen hallinta, tiedonkäsittelykuilu, selitettävä tekoäly, human-in-the-loop, algoritmiaversio, luottamus

TABLE OF CONTENTS

1 Introduction 7

2 Decision making in volatile supply chains 9

2.1 The information processing gap 9

2.2 Symptoms of overload and decision latency 10

2.3 Increasing capacity and agility 11

2.4 Challenges with reliability and trust 12

3 Augmented Analytics in Business Intelligence 14

3.1 Defining the technical framework of augmented analytics 14

3.2 Natural language processing and generation as interpretive infrastructure 15

3.3 Generative AI and data governance 16

3.4 Increasing information processing capacity and strategic agility 17

4 Human-AI Synergy and Reliability 21

4.1 Factors influencing the reliability of analytical outputs 21

4.2 The paradox of trust and algorithmic aversion 22

4.3 Theoretical frameworks of technology acceptance 23

4.4 Bridging the gap: XAI and Human-in-the-Loop models 24

4.5 Trust as a foundation for resilience 27

5 Conclusions 28

References 30

Appendices 34

Appendix 1 Explanation of the use of AI 34

FIGURES

Figure 1. The augmented analytics decision pipeline 25

TABLES

Table 1. Core technologies in augmented analytics and their relevance to supply chain decision-making 18

1 Introduction

In March 2021, a single grounded container ship blocked the Suez Canal for six days, holding up an estimated $9.6 billion in trade per day and exposing how physical disruptions in highly interconnected global networks can generate spreading knock-on effects across supply chains (Dierker et al., 2024). Such crises expose a critical organizational vulnerability: when companies cannot process disruption signals quickly enough to act, they encounter an information processing gap, which is the point at which incoming data and environmental complexity exceed an organization's analytical capacity (Galbraith, 1974). This informational delay can be just as damaging as the physical disruption itself, since time pressure compounds cognitive overload and narrows the window within which managers can act with sufficient clarity (Eppler & Mengis, 2004; Kulpa et al., 2026). The result is higher transport costs and disrupted operations. Reducing decision latency is therefore vital for maintaining supply chain resilience (Kulpa et al., 2026).

The key barrier to rapid responsiveness is the analytical bottleneck in traditional business intelligence systems, where manual data processing cannot keep up with environmental volatility. This bottleneck occurs at the point at which the volume and complexity of incoming information exceed the organization's capacity to interpret it in time for effective action (Kania & Mehta, 2026; Ivanovic et al., 2025). Augmented analytics addresses this challenge by using artificial intelligence and natural language processing to automate insight generation. It works as an analytical support system, meaning a system that augments rather than replaces human reasoning capacity, and helps organizations respond more agilely in volatile conditions (Kania & Mehta, 2026).

While augmented analytics offers a technical pathway to close the information processing gap, its organizational value is not guaranteed. Integrating AI-driven systems into human decision-making raises fundamental challenges around reliability and trust. If managers cannot understand or verify the basis of an automated recommendation, they may hesitate to act on it, preserving the very decision latency the system was designed to eliminate.

This research is situated at the intersection of Information Systems Science and Supply Chain Management. From an information systems perspective, augmented analytics is examined as an analytical support system that enables the processing, interpretation, and communication of vast data volumes to close the information processing gap. From a supply chain perspective, the focus is on operational performance, specifically the reduction of decision latency and the enhancement of resilience in volatile markets. This bachelor's thesis investigates, based on peer-reviewed scientific articles, how augmented analytics can reduce decision latency in operational supply chain management and how the reliability and trustworthiness of AI-generated outputs shape their organizational value.

The thesis aims to answer the following research questions:

- How can augmented analytics reduce decision latency for operational managers in supply chain management?
- What are the key factors influencing the reliability of analytical outputs and user trust in augmented decision-making?

The study focuses on large-scale supply chain operators, as the existing literature on augmented analytics and information processing capacity is predominantly grounded in complex, high-volume operational contexts where the information processing gap is most acute. Studies addressing SME contexts or less interconnected supply networks fall outside this scope. The research does not examine the technical implementation of AI models but focuses on the organizational and cognitive factors, specifically the information processing gap, decision latency, and user trust, that determine whether augmented analytics delivers operational value.

The structure of the thesis is as follows: Chapter 2 establishes the theoretical framework by examining decision-making in volatile supply chains through the lens of Organizational Information Processing Theory and its relationship to decision latency. Chapter 3 examines the technological foundations of augmented analytics, focusing on machine learning, natural language processing, and natural language generation as components of cognitive infrastructure. Chapter 4 analyzes the challenges of reliability and trust in human-AI decision-making, including transparency, explainability, and the conditions for effective adoption. Chapter 5 presents the conclusions and synthesizes the findings in relation to the research questions.

2 Decision making in volatile supply chains

This chapter establishes the theoretical framework for the study by examining the relationship between environmental uncertainty and organizational decision-making. It begins by introducing Organizational Information Processing Theory as a lens for understanding how information processing gaps emerge in volatile environments, then examines how the symptoms of this gap manifest as decision latency and operational exposure and finally discusses how technological interventions can increase organizational processing capacity and strategic agility.

Modern supply chain management operates in an environment where volatility has become a structural condition rather than an exception. Major systemic events such as the 2021 Suez Canal blockage and ongoing trade frictions between global economies demonstrate how physical disruptions quickly spread through interconnected networks (Kulpa et al., 2026; Ding et al., 2026). These disruptions are rarely isolated events. Research increasingly documents that organizations face concurrent, overlapping shocks that test their resilience and recovery capabilities simultaneously, compounding the difficulty of predicting the final impact on operations (Ivanovic et al., 2025). To understand how firms navigate these challenges, it is necessary to examine the relationship between environmental uncertainty and the internal ability of the organization to process information.

2.1 The information processing gap

Organizational Information Processing Theory (OIPT) provides a comprehensive framework for analyzing how companies maintain performance under high levels of uncertainty. OIPT conceptualizes the organization as an information processing system whose capacity to maintain performance depends on its ability to process sufficient information to coordinate interdependent tasks under conditions of environmental uncertainty (Galbraith, 1974). Environmental uncertainty increases the demand for information processing because managers require constant, reliable insights to reconfigure logistics and production schedules during a crisis (Galbraith, 1974; Ivanovic et al., 2025).

When the volume of incoming data and the complexity of the disruption exceed the existing processing capabilities of the organization, an information processing gap is created (Bag et al., 2026). This gap is driven by a specific cognitive mechanism known as information overload. Eppler and Mengis (2004) define information overload as occurring when the processing demands of a task exceed an individual's cognitive capacity. This relationship is captured by the inverted U-curve model: decision accuracy and reasoning quality initially improve as more information becomes available, but beyond a critical threshold, additional information begins to actively degrade performance (Eppler & Mengis, 2004).

Supply chain disruptions are particularly damaging from an information processing perspective because they trigger multiple overload causes simultaneously rather than in isolation. Eppler and Mengis (2004) identify five such causes: the volume and ambiguity of incoming information, the limited cognitive capacity of the individual manager, the extreme time pressure imposed by the crisis, the breakdown of routine organizational coordination, and the inability of existing information systems to keep pace with the data flow. Consider a port closure triggered by extreme weather: within hours, a logistics manager may be receiving conflicting updates from carriers, customs authorities, and warehouse operators across multiple time zones, all while being expected to make rerouting decisions that affect weeks of downstream production. When a major disruption occurs, all five overload causes converge at once, overwhelming the organization's capacity to make sense of what is happening. The result is what Eppler and Mengis (2004) describe as analytical paralysis, a state where the firm possesses sufficient raw data but lacks the cognitive and technical capacity to transform it into actionable intelligence.

2.2 Symptoms of overload and decision latency

The consequences of analytical paralysis manifest in two specific ways that are directly relevant to supply chain performance. First, managers delay or avoid decisions entirely because they cannot integrate incoming information fast enough to act with confidence. Second, they accept lower-quality choices because the cognitive cost of continued analysis exceeds what the situation allows. Eppler and Mengis (2004) identify these as the paralysis and delay of decisions, along with a greater tolerance of error, the primary behavioral symptoms of information overload. These symptoms represent the practical reality of the information processing gap: when overwhelmed, managers spend excessive time attempting to interpret incoming data, producing decision latency, or accept lower-quality choices because they can no longer integrate new information effectively (Eppler & Mengis, 2004).

Decision latency refers to the delay between the emergence of a disruption signal and the execution of a managerial response. In the post-disruption phase, supply chain managers must first process information to become aware of a disruption and identify its scope before corrective action is possible (Kulpa et al., 2026). This means that informational delays directly translate into operational exposure. As the response window contracts, the costs of inaction accumulate in the form of higher transaction costs, lost sales, and disrupted coordination across the network (Kulpa et al., 2026). In volatile supply chain environments, the speed of information processing is therefore not merely a technical concern but a strategic one, since the value of a correct decision diminishes rapidly as the window for its execution closes (Kulpa et al., 2026).

2.3 Increasing capacity and agility

Augmented analytics addresses the information processing gap directly by using artificial intelligence to automate the identification of patterns and insights within large data volumes, thereby reducing the analytical burden on human decision-makers (Kania & Mehta, 2026). This technology allows a firm to achieve supply chain ambidexterity, which refers to the dual organizational ability to remain agile in responding to immediate disruption shocks while remaining adaptable to longer-term strategic shifts (Al Mamun et al., 2025). By shifting the computational burden of data processing from human analysts to AI-driven systems, the organization can process information at a speed that more closely matches the pace of environmental volatility (Al Mamun et al., 2025).

Increasing information processing capacity is also fundamentally linked to the democratization of data access. As organizations implement AI-driven analytics tools, the technical barriers to insight generation are lowered, enabling a broader range of operational managers to access and act on analytical outputs without specialist data expertise (Desai & Desai, 2025). Passlick et al. (2023) describe this structural shift in terms of self-service business intelligence and analytics, identifying the emergence of distinct user roles, including what the literature terms citizen data scientists, who can perform analytical tasks without belonging to a dedicated data function. This broadening of the analytical user base is not merely a technical capability but an organizational one: it redistributes the capacity for insight generation across the firm rather than concentrating it in specialist departments (Passlick et al., 2023; Alghamdi and Al-Baity, 2022). A critical enabler of this increased capacity is end-to-end supply chain visibility, which refers to a real-time, network-wide view of inventory positions, shipment status, and demand signals across all tiers of the supply chain (Ivanov, 2024). Without such visibility, organizations cannot feed timely and accurate data into analytical systems fast enough to generate actionable insights during a disruption (Ivanov, 2024).

Technology alone, however, does not close the information processing gap. The organizational culture surrounding data use is an equally important determinant of whether analytical tools deliver value in practice (Cadden et al., 2022). If an organization does not cultivate a culture in which managers trust and act on data-driven insights, even technically sophisticated augmented analytics systems will fail to reduce decision latency. The goal of augmented analytics is therefore not simply to accelerate computation, but to reduce the analytical burden sufficiently that managers can make sense of complex, rapidly evolving situations and act on that understanding before the response window closes.

2.4 Challenges with reliability and trust

Augmented analytics introduces new risks to the decision-making process even as it increases the speed of analysis. Simply accelerating the volume of information does not always lead to better outcomes if the sensemaking capacity of the firm is overwhelmed. In a crisis, a manager receiving a high frequency of automated alerts may experience a new form of information overload, one generated by the very systems designed to reduce it, making it harder to distinguish critical signals from noise (Ivanovic et al., 2025). This condition can freeze the decision-making process and create what might be described as technical latency, a delay produced not by insufficient data but by insufficient interpretive clarity (Ivanovic et al., 2025).

There are also concerns about the accuracy of AI-generated outputs. Generative AI systems can produce outputs that are fluent and internally consistent but factually incorrect, a phenomenon that directly undermines the reliability of analytical recommendations and the willingness of managers to act on them (Desai & Desai, 2025). A related challenge is the opacity of many AI models, which provide recommendations without revealing the reasoning behind them. If a manager is instructed to cancel a major shipment without any explanation of why, they are likely to hesitate regardless of whether the recommendation is correct (Arrieta et al., 2020). This hesitation represents a form of decision latency that originates not in the information processing gap but in a trust gap between the manager and the system. When a manager doubts the machine, the speed advantage of augmented analytics is lost to human uncertainty (Dietvorst et al., 2015).

These challenges suggest that the relationship between augmented analytics and decision latency is not straightforwardly positive. The technology can compress the time required to generate an insight, but if that insight is not trusted, understood, or acted upon, the information processing gap remains open. Closing that gap therefore requires not only technical capability but also the organizational and psychological conditions under which managers can confidently rely on AI-generated outputs (Dietvorst et al., 2015). Research confirms that decision-maker characteristics shape whether analytical capabilities translate into resilience outcomes in practice, meaning that the human dimension of adoption is as consequential as the technical one (Asgari et al., 2025).

3 Augmented Analytics in Business Intelligence

This chapter examines the technological infrastructure that constitutes augmented analytics and connects it to the information processing gap established in Chapter 2. It first defines the technical framework of augmented analytics, focusing on the foundational role of machine learning and its relationship to natural language technologies. It then explores how natural language processing and natural language generation reduce interpretative latency and democratize data access. Following this, it examines how the integration of generative AI extends these capabilities while introducing new requirements for data governance. Lastly, it analyzes how the combination of these technologies increases organizational information processing capacity and strategic agility.

3.1 Defining the technical framework of augmented analytics

Traditional analytical methods are no longer sufficient to handle the scale and speed of modern data environments. Augmented analytics represents a significant shift in how organizations apply data science. Kania and Mehta (2026) define augmented analytics as an interactive, data-driven system that combines personal judgment with advanced statistical analysis to improve decision-making. At its core, this is what the literature calls Augmented Data Science, a framework that removes preprocessing bottlenecks and empowers business users rather than concentrating analysis in specialist roles (Kania & Mehta, 2026). This empowerment is operationalized through the emergence of what Alghamdi and Al-Baity (2022) and Passlick et al. (2023) term the citizen data scientist, a user who can perform advanced analytical tasks including predictive modeling and pattern recognition without a primary background in statistics or data science. Passlick et al. (2023) identify this user category as one of three distinct role types within self-service analytics environments, alongside information consumers and information producers, each placing different demands on the system design and the analytical infrastructure supporting it.

The technical evolution from structured data analytics toward machine learning and natural language processing as core business intelligence capabilities has been well documented in the information systems literature (Chen et al., 2012). Passlick et al. (2023) situate this evolution within a broader trajectory of self-service business intelligence and analytics, identifying that contemporary platforms have progressively shifted analytical capability from IT specialists to business users across successive generations of tooling. Alghamdi and Al-Baity (2022) similarly document this generational progression, characterizing augmented analytics as the third generation of business intelligence in which AI automates the analytics cycle that earlier generations required users to perform manually. It is this accumulated evolution that forms the foundational research landscape on which modern augmented analytics systems are built (Alghamdi and Al-Baity, 2022; Passlick et al., 2023).

Machine learning, broadly defined as the discipline of building systems that improve automatically through experience, has progressed from a laboratory curiosity into a practical technology in widespread commercial use across sectors including manufacturing and financial modeling (Jordan & Mitchell, 2015). Machine learning serves as the fundamental engine within augmented analytics. Its primary role within the augmented analytics pipeline is to reduce the time-consuming steps of data preprocessing and feature development that have traditionally represented significant bottlenecks for human analysts (Kania & Mehta, 2026). By utilizing algorithms such as recurrent neural networks and long short-term memory networks, the system can learn from historical data to identify and predict trends in supply chain behavior. This evolution has moved artificial intelligence from simple rule-based expert systems toward cognitive computing, where the emphasis is on enhancing the quality of human decision-making rather than replacing human judgment entirely (Duan et al., 2019). In a business context, this shift allows firms to move from descriptive analytics, which explains what has already happened, toward predictive and prescriptive analytics, which suggest what should be done (Lepenioti et al., 2020). This transition is critical for supply chain resilience because it provides the foresight needed to mitigate risks before they propagate through the global network (Duan et al., 2019).

3.2 Natural language processing and generation as interpretive infrastructure

Natural language processing, referred to here as NLP, is the field of computer science concerned with using computational techniques to learn, understand, and produce human language content, enabling seamless interaction between human decision-makers and technical data systems (Hirschberg & Manning, 2015). Through this capability, managers can interact with complex datasets using everyday language rather than requiring specialized knowledge of query languages such as SQL (Kania & Mehta, 2026). A logistics manager who needs to know the current inventory status of a critical component can, for instance, type a plain-language question into the system and receive an immediate answer, rather than waiting for a data analyst to run a query and prepare a report. The practical consequence of this is the democratization of data access, ensuring that insights are available across all levels of the organization rather than being restricted to data specialists (Desai & Desai, 2025). This democratization is particularly valuable during a supply chain crisis because it allows frontline managers to obtain immediate answers regarding shipment statuses or inventory levels, directly reducing the initial phase of decision latency identified in Chapter 2 (Desai & Desai, 2025).

Natural language generation, referred to here as NLG, represents the output stage of the analytical pipeline, where the system translates mathematical findings into human-readable narratives. While traditional analytics often presents users with a series of complex charts that require independent interpretation, NLG produces explanatory text that communicates the significance of findings in plain language (Kania & Mehta, 2026). Rather than presenting a manager with a logistics delay probability distribution across multiple carrier routes, an NLG system might instead generate a sentence stating that carrier reliability on the Rotterdam-Hamburg corridor has declined by 12 percent over the past 14 days and that rerouting via Antwerp is advisable based on current lead time data. As stated, the strategic role of NLG is to bridge the gap between raw data and actionable insights, enabling managers to grasp the implications of a disruption without requiring specialist data expertise (Kania & Mehta, 2026).

Taken together, NLP and NLG function as an interpretive infrastructure that sits between the computational power of machine learning and the judgment of the human decision-maker. NLP translates human intent into machine-readable queries; NLG translates machine outputs into human-readable insights (Kania & Mehta, 2026). This bidirectional translation capability is what allows augmented analytics to function as a cognitive infrastructure in the sense defined in Chapter 1, supporting rather than replacing human reasoning capacity.

3.3 Generative AI and data governance

A significant recent development in augmented analytics is the integration of generative artificial intelligence and large language models (Desai & Desai, 2025). This integration enhances the effectiveness of human-system interaction by providing conversational interfaces capable of handling unstructured and messy data that traditional analytics pipelines struggle to process. Generative AI enables faster democratization of data access across organizational domains including finance, retail, and logistics by lowering the technical threshold required to query and interpret complex datasets (Desai & Desai, 2025).

The use of generative AI also introduces specific and consequential challenges regarding the accuracy and reliability of analytical outputs. Generative AI models generate the most probable response to a prompt rather than necessarily the correct one, a phenomenon known as hallucination, where outputs are semantically plausible but factually incorrect and typically not easily verifiable (Feuerriegel et al., 2024; Desai & Desai, 2025). This risk represents a primary concern for data governance in organizational contexts, and to address it, organizations must implement robust validation frameworks that compare AI-generated outputs against real-world datasets and establish clear protocols for human review of high-stakes recommendations. Successful integration of generative AI into augmented analytics therefore requires that the generated insights are not only produced rapidly but are also contextually relevant and factually verifiable, ensuring that speed does not come at the cost of reliability (Desai & Desai, 2025).

3.4 Increasing information processing capacity and strategic agility

When machine learning, NLP, NLG, and generative AI are integrated into a unified augmented analytics platform, the combined effect is a substantial increase in the organization's information processing capacity. According to OIPT, a firm maintains resilience only when its internal processing capabilities are sufficient to match the level of environmental uncertainty it faces (Galbraith, 1974). In modern logistics environments, where disruptions propagate rapidly through global networks, the volume and velocity of incoming data routinely exceed the limits of manual analysis, creating the processing gap described in Chapter 2. Galbraith (1974) identifies investment in vertical information systems, meaning systems that collect, process, and distribute information across organizational levels, as a primary organizational strategy for increasing processing capacity without expanding human headcount. Following this logic, augmented analytics addresses the processing gap by automating the preprocessing and feature engineering stages that have historically represented the most significant analytical bottlenecks, allowing organizations to generate insights at a speed that manual analysis cannot match (Kania & Mehta, 2026).

A central outcome of this increased capacity is the reduction of analysis latency, which refers to the time required to transform raw data into actionable intelligence. Kania and Mehta (2026) identify data preprocessing and feature development as the most significant time-consuming bottlenecks in the traditional analytical process. Empirical support for this claim is provided by Alghamdi and Al-Baity (2022), whose comparative study of traditional BI and augmented analytics platforms found that automation of data preparation and model building measurably compressed the time required at each pipeline stage relative to manual approaches. When these steps are automated, the period between data collection and insight generation is substantially compressed. It should be noted that Alghamdi and Al-Baity's study (2022) was conducted in a retail context using a single company dataset, and the magnitude of time savings will vary by organizational context and data complexity.

The direction of the effect is nonetheless consistent across both their empirical findings and the theoretical predictions of OIPT established in Chapter 2: automation reduces analysis latency regardless of the organizational context in which it is measured (Alghamdi and Al-Baity, 2022; Kania & Mehta, 2026).

Increased processing capacity also supports supply chain ambidexterity, which Chapter 2 defined as the dual organizational ability to remain agile in responding to immediate shocks while remaining adaptable to longer-term strategic shifts. Machine learning contributes to this by providing predictive and prescriptive insights that go beyond describing past events. By identifying patterns in historical data, the technology provides the foresight needed to prepare alternative logistics routes or reconfigure production schedules in advance of a disruption (Duan et al., 2019; Lepenioti et al., 2020). This predictive capability transforms the firm from a reactive to a proactive system, which is a fundamental shift in the nature of supply chain decision-making (Duan et al., 2019). The core technologies that contribute to this transformation, along with their function in augmented analytics and their relevance to supply chain decision-making, are summarized in Table 1.

Table 1. Core technologies in augmented analytics and their relevance to supply chain decision-making

Technology	**Function in Augmented Analytics**	**Relevance to Supply Chain Decision-Making**	**Source**
Machine Learning (ML)	Automates data preprocessing, feature development, and pattern recognition from historical data.	Reduces analysis latency by identifying supply chain disruption signals and generating predictive and prescriptive insights without manual intervention.	Jordan & Mitchell (2015); Duan et al. (2019); Kania & Mehta (2026)
Natural Language Processing (NLP)	Enables plain-language querying of complex datasets, translating human intent into machine-readable queries.	Allows frontline managers to interrogate supply chain data directly without specialist data expertise.	Kania & Mehta (2026); Desai & Desai (2025)
Natural Language Generation (NLG)	Translates mathematical findings and model outputs into human-readable narrative summaries.	Communicates disruption signals and recommendations in plain language, enabling managers to grasp	Kania & Mehta (2026)

		implications without requiring data expertise.	
Large Language Models (LLMs)	Provides conversational interfaces capable of handling unstructured data and generating contextually responsive outputs.	Lowers the technical threshold for querying analytical systems across organizational domains including logistics and finance, while introducing hallucination risk requiring governance controls.	Feuerriegel et al. (2024); Desai & Desai (2025)
Explainable AI (XAI)	Makes the reasoning behind AI recommendations interpretable to human users through local interpretability mechanisms.	Reduces epistemic discomfort and algorithmic aversion by communicating which variables drove a recommendation and with what weight, enabling managers to know when to trust and when to override outputs.	Gunning & Aha (2019); Arrieta et al. (2020)
Human-In-The-Loop (HITL)	Preserves human decision authority within the automated analytics pipeline by assigning computational tasks to AI and judgment tasks to the manager.	Ensures accountability and contextual intelligence are retained while augmented analytics accelerates the preceding analytical stages.	Alghamdi & Al-Baity (2022); Kania & Mehta (2026); Arrieta et al. (2020)

Increasing processing capacity does not, however, automatically translate into better decisions. The final judgment still relies on the expertise and contextual understanding of the human manager. Augmented analytics provides the analytical support that allows managers to make sense of complex situations more rapidly and with greater confidence, but it does so in service of human decision-making rather than as a substitute for it (Kania & Mehta, 2026). The value of the technology is therefore conditional: it closes the information processing gap on the computational side, but as established in Chapter 2, a trust gap between the manager and the system can preserve decision latency even when the analytical output is accurate. This conditional relationship between technical capability and organizational value is the central concern of Chapter 4 (Kania & Mehta, 2026).

4 Human-AI Synergy and Reliability

This chapter analyzes what might be called the last mile of the augmented analytics pipeline: the interaction between the human expert and the automated system. The preceding chapters established that augmented analytics can technically close the information processing gap by automating insight generation and reducing analysis latency. This chapter examines whether that technical capability translates into organizational value in practice. To answer the second research question, the chapter first evaluates the technical factors that determine whether an analytical output is objectively reliable, then examines the psychological factors that determine whether a manager trusts and acts on that output, and lastly introduces Explainable AI and Human-in-the-Loop models as the primary organizational responses to these challenges.

4.1 Factors influencing the reliability of analytical outputs

Before trust can be established between a manager and an analytical system, the system must produce outputs that are objectively reliable. Three primary factors threaten reliability in the context of augmented analytics: first, data quality problems in the underlying sources; second, the risk of hallucinations in generative models; and third, the limited contextual awareness of automated systems. These challenges reflect broader responsible AI concerns around data governance, transparency, and the boundaries of human oversight that have been identified as central to the safe organizational deployment of AI systems (Mikalef et al., 2022).

The first threat is data quality. A system is only as reliable as the data it consumes, and if the underlying data sources contain fragmented, duplicated, or outdated information, the resulting prediction will be logically flawed regardless of the sophistication of the model producing it. Organizations across sectors frequently encounter challenges in maintaining data quality and consistency when integrating information from multiple internal systems, and these structural data problems represent a foundational threat to the reliability of augmented analytics outputs (Desai & Desai, 2025; Passlick et al., 2023). Passlick et al. (2023) treat data governance as an independently significant dimension of self-service analytics deployment, finding through both tool analysis and a company case study that the reliability of analytical outputs depends critically on whether data modeling is controlled by IT or delegated to business users. When business users perform their own data modeling, quality is not necessarily guaranteed, and the degree of risk varies with the complexity of the transformations required. This finding is directly relevant to augmented analytics in supply chain settings, where data inputs frequently originate from multiple heterogeneous systems across supplier and logistics networks, compounding the data quality challenges that Desai and Desai (2025) identify at the system level (Passlick et al., 2023).

The second threat is hallucination. The integration of large language models introduces a distinct reliability risk in which the system generates plausible-sounding but factually incorrect summaries or recommendations. These errors occur because large language models are designed to predict linguistically probable outputs based on statistical patterns in training data rather than to verify the factual accuracy of what they produce (Feuerriegel et al., 2024; Desai & Desai, 2025). In a supply chain context, a system might report that a key supplier has sufficient buffer stock to cover a two-week delay when the underlying data actually reflects a different supplier's inventory, leading to a response strategy built on false assumptions. Without robust data governance protocols and systematic output validation, these inaccuracies make managerial skepticism a rational response to genuine technical limitations rather than an irrational bias (Desai & Desai, 2025). This distinction matters because the appropriate organizational response to rational skepticism is improved system reliability, whereas the appropriate response to irrational aversion is improved transparency and communication, a point developed in section 4.2 (Desai & Desai, 2025).

The third threat is limited contextual awareness. Machine learning models are trained on historical data and optimize for pattern recognition within that data. They may therefore fail to account for qualitative developments that fall outside their training distribution. A model trained on years of stable supplier performance data will not independently register that a key supplier relationship is deteriorating due to a contract dispute, or that a new trade restriction has just been announced. Human judgment remains essential precisely because it can incorporate this kind of contextual intelligence that automated systems cannot reliably capture (Kania & Mehta, 2026).

4.2 The paradox of trust and algorithmic aversion

Even when a system is technically reliable, managers may still reject or discount its recommendations. Algorithmic aversion refers to this tendency of decision-makers to discount automated recommendations even when the system demonstrably outperforms human judgment (Dietvorst et al., 2015). Understanding this phenomenon is essential because it means that improving technical reliability alone is insufficient to ensure adoption (Dietvorst et al., 2015).

Algorithmic aversion originates in what the literature describes as epistemic discomfort, a sense of professional vulnerability that arises when a manager cannot evaluate the reasoning behind a recommendation. When a manager receives advice from a human colleague, they can ask follow-up questions, probe the underlying logic, and assess the credibility of the source. An opaque automated system offers none of this interaction. Because the manager cannot verify the reasoning, they are unable to take intellectual ownership of the decision, which creates discomfort that often manifests as preference for their own judgment even when that judgment is objectively inferior (Arrieta et al., 2020; Kania & Mehta, 2026).

The consequences of algorithmic aversion for decision latency are significant. If a manager spends time doubting or overriding an accurate AI recommendation, the temporal advantage gained by automated insight generation is consumed by human hesitation. The information processing gap is technically closed but practically remains open. Algorithmic aversion is therefore the mechanism by which a technically solved problem remains operationally unsolved, and it represents the central challenge this chapter addresses (Dietvorst et al., 2015).

4.3 Theoretical frameworks of technology acceptance

Two theoretical frameworks help explain how behavioral barriers to augmented analytics adoption operate at the organizational level: the Technology Acceptance Model and the Unified Theory of Acceptance and Use of Technology.

The first is the Technology Acceptance Model, established by Davis et al. (1989), which identified perceived usefulness and perceived ease of use as the two primary determinants of technology adoption intention. In a supply chain context, perceived usefulness translates to whether a manager believes the augmented analytics system actually improves their ability to respond to disruptions. Perceived ease of use translates to whether interacting with the system feels manageable under the time pressure of a crisis. While TAM remains influential, its explanatory power is limited in complex organizational contexts where adoption involves not just individual preference but institutional pressure, task characteristics, and prior experience (Davis et al., 1989).

The second framework is the Unified Theory of Acceptance and Use of Technology, known as UTAUT, developed by Venkatesh et al. (2003) to address these limitations by synthesizing eight competing adoption models. UTAUT identifies four core determinants of adoption intention: performance expectancy, which is the belief that the technology improves work outcomes; effort expectancy, which is the perceived ease of using the system; social influence, which is the degree to which peers and supervisors are seen to support use; and facilitating conditions, which refers to the organizational infrastructure available to support the user. Together, these four determinants explain approximately 70 percent of the variance in adoption intention, a substantial improvement over prior models (Venkatesh et al., 2003).

In large-scale supply chain operations, the social influence dimension of UTAUT carries particular significance. Adoption in these contexts is typically mandatory rather than voluntary, and Venkatesh et al. (2003) found that in mandatory contexts, social influence tends to produce compliance without genuine trust. A manager may use the system because institutional expectations require it while continuing to discount its outputs during high-stakes decisions. This compliance-without-trust dynamic is especially dangerous in crisis conditions, precisely the situations where augmented analytics is most needed, because it preserves decision latency at the moment when reducing it is most critical (Venkatesh et al., 2003).

Recent empirical work by Burnay et al. (2024) adds a further dimension to this framework by identifying representational load as a significant adoption barrier. Representational load refers to the cognitive effort required to interpret an analytical output. A dashboard displaying seventeen simultaneously updating KPI indicators, several probability distributions, and a heat map of supplier risk scores may technically contain all the information a manager needs, but the effort required to parse it under crisis conditions can be so high that the manager disengages from the system entirely. When outputs are difficult to parse quickly, effort expectancy increases and adoption intention falls directly (Burnay et al., 2024). This finding is particularly relevant for augmented analytics, where the volume and complexity of outputs can inadvertently recreate the information overload conditions the technology was designed to eliminate.

It should lastly be noted that UTAUT was validated primarily in routine workplace technology contexts. The high-pressure, time-constrained conditions of supply chain disruptions differ meaningfully from this baseline, and the framework's predictions may not translate without qualification to crisis decision-making environments. This is a limitation that future research in this area will need to address (Venkatesh et al., 2003).

4.4 Bridging the gap: XAI and Human-in-the-Loop models

Explainable AI and Human-in-the-Loop models represent the primary organizational and design responses to the barriers identified in the preceding sections. They address the reliability and trust challenges from complementary angles: XAI improves the transparency of system outputs, while

HITL preserves meaningful human agency within the decision-making process. Figure 1 illustrates how XAI and HITL operate as distinct intervention points within the augmented analytics pipeline, translating automated outputs into trusted, human-verified decisions.

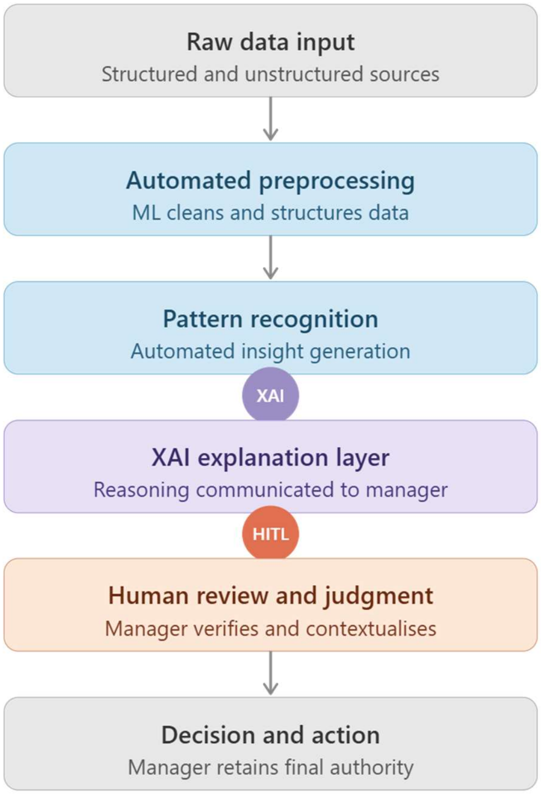

Figure 1. The augmented analytics decision pipeline

Blue boxes indicate ML-driven stages. The XAI circle marks where system reasoning is communicated to the manager. The HITL circle marks where human judgment enters as the governing principle. Gray boxes indicate human-owned stages. Adapted from Kania and Mehta (2026), Arrieta et al. (2020), and Alghamdi and Al-Baity (2022).

As shown in Figure 1, the XAI explanation layer translates automated insights into reasoning that the manager can interrogate before judgment is applied. Explainable AI, commonly abbreviated as XAI, is a field of research and design practice dedicated to making the reasoning of AI systems interpretable to human users, enabling them to understand, appropriately trust, and effectively use AI-generated outputs (Gunning & Aha, 2019). A central feature of XAI systems is local interpretability, which is the capacity to explain the specific reasoning behind an individual recommendation rather than only describing the general behavior of the model (Arrieta et al., 2020). In a supply chain context, this means the system does not only flag a shipping route as high-risk but also communicates which specific variables contributed to that assessment and with what weight. A manager might, for instance, see that the system's recommendation to reroute a shipment is based 60 percent on port congestion data from the past 72 hours, 25 percent on a declining carrier on-time delivery score, and 15 percent on adverse weather forecasts, giving them the information needed to either trust the recommendation or override it based on knowledge the system does not have. The goal of explanation is not merely to describe system behavior but to support the user's ability to know when to trust and when to doubt an automated recommendation (Gunning & Aha, 2019). By providing the reasoning behind the recommendation, XAI reduces the epistemic discomfort that drives algorithmic aversion and lowers the representational load that suppresses adoption. It also allows the manager to identify cases where the system's reasoning is based on incomplete or contextually inappropriate data, which directly addresses the reliability challenges identified in section 4.1 (Arrieta et al., 2020; Kania & Mehta, 2026).

The Human-in-the-Loop model, in turn, ensures that the efficiency gains of automated analytics are not pursued at the cost of human judgment and accountability. As Figure 1 shows, HITL occupies the position between the explanation layer and the final decision and action stage. In this model, the AI system performs the computationally intensive tasks of data processing, pattern recognition, and insight generation, while the human manager serves as the final decision authority. This division of labor reflects a fundamental characteristic of augmented analytics that its proponents and empirical evaluators consistently affirm: the technology facilitates but does not replace human analytical work (Arrieta et al., 2020; Kania & Mehta, 2026). Alghamdi and Al-Baity (2022) state this conclusion explicitly in their comparative study, observing that augmented analytics platforms accelerate and support the analytical process but that the primary analytical workload remains with the human analyst, who is responsible for identifying which automatically generated insights are operationally meaningful and which are not. Not all automatically generated insights are useful to the specific business goal in question, and decision-makers retain an essential role in translating outputs into action. This empirically grounded observation reinforces the theoretical case for HITL design made by Kania and Mehta (2026) and Arrieta et al. (2020), and grounds it in the practical reality of how augmented analytics platforms perform when applied to actual organizational data (Alghamdi and Al-Baity, 2022; Kania & Mehta, 2026).

In short, the HITL model does not slow decision-making when implemented well. Rather, it concentrates human attention on the judgment call rather than the preceding analytical work, which is the most efficient possible use of managerial cognitive capacity. Together, XAI and HITL transform the role of augmented analytics from an autonomous decision system into a decision support system, a distinction that has significant implications for both reliability and trust. The manager retains authority and visibility; the system provides speed and analytical breadth. This is the organizational form that the cognitive infrastructure concept introduced in Chapter 1 ultimately describes (Kania & Mehta, 2026; Arrieta et al., 2020).

4.5 Trust as a foundation for resilience

This chapter has shown that the two research questions of this thesis are not independent but interdependent. Closing the technical information processing gap, as Chapters 2 and 3 examined, is necessary, but it is not sufficient for reducing decision latency. Managers must also find the outputs both reliable and trustworthy enough to act on (Kania & Mehta, 2026).

The frameworks examined in this chapter reveal why satisfying both conditions is difficult in practice. First, TAM and UTAUT demonstrate that adoption intention is shaped by perceived usefulness, cognitive effort, and social context in ways that can produce compliance without genuine trust. Second, Burnay et al. (2024) show that representational load can recreate information overload within the very systems designed to eliminate it. Third, Dietvorst et al. (2015) show that algorithmic aversion can persist even when a system is demonstrably accurate, because the barrier is psychological rather than technical. Taken together, these findings confirm that the organizational challenge of augmented analytics is as much a human problem as a technological one (Dietvorst et al., 2015).

XAI and HITL models address this challenge by making the system interpretable and keeping the human in a position of genuine authority. When these design principles are applied, the speed of automated analytics and the judgment of the human manager are complementary rather than competing. The information processing gap is closed on the technical side; the trust gap is closed on the human side. Only when both conditions are satisfied does the reduction in analysis latency translate into the reduction in decision latency that this thesis identifies as the central value proposition of augmented analytics in volatile supply chain environments (Arrieta et al., 2020; Kania & Mehta, 2026).

5 Conclusions

This thesis examined how augmented analytics can reduce decision latency in operational supply chain management and under what conditions AI-generated outputs are sufficiently reliable and trusted to deliver organizational value. The study was motivated by the observation that supply chain disruptions generate not only physical and logistical consequences but also an information processing gap, a condition in which the volume and complexity of incoming disruption signals exceed the analytical capacity of the organization. The central argument pursued across the chapters is that augmented analytics offers a workable response to this gap, but that its organizational value is conditional on factors that extend beyond the technology.

The first research question asked how augmented analytics can reduce decision latency for operational managers in supply chain management. The answer developed across Chapters 2 and 3 is that it does so through two complementary mechanisms. First, by automating the most time-consuming stages of the analytical pipeline, specifically data preprocessing, feature development, and insight generation, augmented analytics compresses the period between disruption onset and actionable intelligence. Second, by enabling natural language querying and natural language generation, the technology democratizes access to analytical outputs, allowing frontline managers to interrogate data directly without routing requests through specialist functions. Both mechanisms address different phases of decision latency: the first reduces analysis latency at the computational level, and the second reduces interpretive latency at the human level. Together, they close the information processing gap that Organizational Information Processing Theory identifies as the primary barrier to effective disruption response.

The second research question asked what factors determine whether analytical outputs are reliable and trusted. Chapter 4 established that this is where the organizational challenge becomes most acute. Technical reliability is threatened by data quality problems, hallucinations in generative AI systems, and the limited contextual awareness of models trained on historical data. Even when outputs are objectively accurate, managerial trust is not guaranteed. Algorithmic aversion, representational load, and compliance without genuine trust are persistent barriers that can preserve decision latency even when the analytical infrastructure performs as intended. The conclusion drawn is that closing the information processing gap on the computational side is a necessary but insufficient condition for reducing decision latency. The sufficient condition requires that outputs are both reliable and trusted, and that managers retain genuine decision authority within the process.

Explainable AI and Human-in-the-Loop design are identified as the primary organizational responses to this challenge, because they address the trust gap by making system reasoning transparent and keeping the human manager in a position of meaningful accountability.

It is important to note that this thesis was not intended as a technical evaluation of business intelligence architectures. The focus throughout has been on the organizational and managerial dimension of augmented analytics adoption: how the technology affects the decision-making capacity of supply chain managers, what risks it introduces into that process, and what conditions must be satisfied for it to deliver operational value. This matters because the obstacles to realizing value from augmented analytics are largely human and organizational, not technical. Organizations that treat augmented analytics as a purely infrastructural investment, without attending to trust, interpretability, and data governance, are unlikely to achieve the reductions in decision latency that technology theoretically enables.

Several limitations of this study should be acknowledged. The analysis is grounded in peer-reviewed literature rather than primary empirical data, which means the findings reflect theoretical arguments and existing research rather than direct organizational observation. Additionally, the field is evolving rapidly, and sources from 2025 and 2026 represent early empirical work whose findings may be revised as the technology matures.

Future research should address several open questions. Empirical studies examining how augmented analytics adoption affects decision latency in real supply chain environments would provide the kind of evidence that literature-based analysis cannot. The role of organizational culture in mediating the relationship between analytical capability and decision quality warrants dedicated investigation, as the existing literature treats it as a moderating factor without fully theorizing the mechanism. Lastly, the continued development of agentic AI systems, in which AI agents act autonomously across multi-step tasks rather than simply generating recommendations for human review, represents a significant extension of the augmented analytics paradigm. As these systems become more prevalent in supply chain contexts, questions of accountability, interpretability, and human oversight will become considerably more complex and will require dedicated theoretical and empirical attention.

References

Al Mamun, A., Reza, M. N. H., Yang, Q., & Aziz, N. A. (2025). Dynamic capabilities in action: The synergy of big data analytics, supply chain ambidexterity, green supply chain and firm performance. Journal of Enterprise Information Management, 38(2), 636–659. https://doi.org/10.1108/JEIM-08-2024-0441

Alghamdi, A., & Al-Baity, H. H. (2022). Toward augmented data scientists: An empirical study for evaluating the effect of automation on data science workflow performance. Applied Sciences, 12(17), 8443. https://doi.org/10.3390/app12178443

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82–115. https://doi.org/10.1016/j.inffus.2019.12.012

Asgari, A., Alaeddini, M., & Reaidy, P. (2025). The role of big data and predictive analytics in developing resilient supply chains: A resource-based perspective. IFAC-PapersOnLine, 59(10), 2874–2879. https://doi.org/10.1016/j.ifacol.2025.09.483

Bag, S., Srivastava, G., Routray, S., & Chiarini, A. (2026). Generative AI, ESG sensemaking, and environmental performance: An OIPT perspective. Business Strategy and the Environment. https://doi.org/10.1002/bse.70520

Burnay, C., Jureta, I. J., Linden, I., & Faulkner, S. (2024). Representational load and its impact on the adoption of business intelligence systems. Information Systems Journal, 34(1), 102–138. https://doi.org/10.1111/isj.12432

Cadden, T., McIvor, R., Cao, G., Treacy, R., Yang, Y., Gupta, M., & Onofrei, G. (2022). Unlocking supply chain agility and supply chain performance through the development of intangible supply chain analytical capabilities. International Journal of Operations & Production Management, 42(9), 1329–1355. https://doi.org/10.1108/IJOPM-06-2021-0383

Chen, H., Chiang, R. H. L., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. MIS Quarterly, 36(4), 1165–1188. https://doi.org/10.2307/41703503

Davis, F. D., Bagozzi, R. P., & Warshaw, P. R. (1989). User acceptance of computer technology: A comparison of two theoretical models. Management Science, 35(8), 982–1003. https://doi.org/10.1287/mnsc.35.8.982

Desai, D., & Desai, A. (2025). Integrating generative AI in business intelligence: A practical framework for enhancing augmented analytics. International Journal of Mathematical, Engineering and Management Sciences, 10(3), 704–728. https://doi.org/10.33889/IJMEMS.2025.10.3.036

Dierker, D., Lopez, N., Murnane, J., Saxon, S., Szigethi, J., & Zampelas, A. (2024, January 19). How could Panama Canal restrictions affect supply chains? McKinsey & Company. https://www.mckinsey.com/industries/travel-logistics-and-infrastructure/our-insights/how-could-panama-canal-restrictions-affect-supply-chains

Dietvorst, B. J., Logg, J. M., & Soll, J. B. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. Journal of Experimental Psychology: General, 144(1), 114–126. https://doi.org/10.1037/xge0000033

Ding, H., Huang, D., Li, Z., & Lu, X. (2026). Digital safeguards in trade wars: Assessing the impact of China's CBEC pilot zone on global supply chain resilience. Journal of International Money and Finance, 162. https://doi.org/10.1016/j.jimonfin.2026.103529

Duan, Y., Edwards, J. S., & Dwivedi, Y. K. (2019). Artificial intelligence for decision making in the era of big data – evolution, challenges and research agenda. International Journal of Information Management, 48, 63–71. https://doi.org/10.1016/j.ijinfomgt.2019.01.021

Eppler, M. J., & Mengis, J. (2004). The concept of information overload: A review of literature from organization science, accounting, marketing, MIS, and related disciplines. Information Society, 20(5), 325–344. https://doi.org/10.1080/01972240490507974

Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. Business & Information Systems Engineering, 66(1), 111–126. https://doi.org/10.1007/s12599-023-00834-7

Galbraith, J. R. (1974). Organization design: An information processing view. Interfaces, 4(3), 28–36. https://doi.org/10.1287/inte.4.3.28

Gunning, D., & Aha, D. (2019). DARPA's Explainable Artificial Intelligence (XAI) program. AI Magazine, 40(2), 44–58. https://doi.org/10.1609/aimag.v40i2.2850

Hirschberg, J., & Manning, C. D. (2015). Advances in natural language processing. Science, 349(6245), 261–266. https://doi.org/10.1126/science.aaa8685

Ivanov, D. (2024). Digital supply chain management and technology to enhance resilience by building and using end-to-end visibility during the COVID-19 pandemic. IEEE Transactions on Engineering Management, 71, 10485–10495. https://doi.org/10.1109/TEM.2021.3095193

Ivanovic, N., de Vries, T., van der Vegt, G., & van Donk, D. P. (2025). Handling disruption concurrence: The importance of inter- and intra-departmental communication for critical infrastructure resilience. Journal of Supply Chain Management, 61(3), 55–76. https://doi.org/10.1111/jscm.12346

Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. Science, 349(6245), 255–260. https://doi.org/10.1126/science.aaa8415

Kania, S. S., & Mehta, Y. (2026). Augmented analytics: Data interpretation and natural language generation. In T. Senjyu et al. (Eds.), Smart trends in computing and communications (Lecture Notes in Networks and Systems, Vol. 1465, pp. 167–175). Springer. https://doi.org/10.1007/978-981-96-7517-3_15

Kulpa, T., Miller, J. W., Darby, J. L., Thomas, R., & Waller, M. (2026). Flexible responses to global supply chain disruptions: Empirical evidence from international logistics. Transportation Journal, 65(1), e70016. https://doi.org/10.1002/tjo3.70016

Lepenioti, K., Bousdekis, A., Apostolou, D., & Mentzas, G. (2020). Prescriptive analytics: Literature review and research challenges. International Journal of Information Management, 50, 57–70. https://doi.org/10.1016/j.ijinfomgt.2019.04.003

Mikalef, P., Conboy, K., Lundström, J. E., & Popovič, A. (2022). Thinking responsibly about responsible AI and 'the dark side' of AI. European Journal of Information Systems, 31(3), 257–268. https://doi.org/10.1080/0960085X.2022.2026621

Passlick, J., Guhr, N., Lebek, B., & Breitner, M. H. (2023). Explaining the ambiguous nature of self-service business intelligence and analytics systems' influence on information overload. Information Systems Journal, 33(3), 548–582. https://doi.org/10.1111/isj.12394

Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. MIS Quarterly, 27(3), 425–478. https://doi.org/10.2307/30036540

Appendices

Appendix 1 Explanation of the use of AI

In this bachelor’s thesis, generative artificial intelligence tools (ChatGPT, Gemini, Scopus AI, and Microsoft Word’s automated spell check) have been utilized for tasks such as outlining the thesis structure, searching for scientific sources, translating and clarifying terminology, and assisting in the translation process. AI has not been used to generate the final text, analyze data, or answer the research questions. In accordance with good academic practices, AI has been used openly and transparently. AI served as a tool to support the process but does not replace scientific analysis or independent thinking.

OpenAI ChatGPT (5.3) and Google Gemini (3) were used in the following ways:

1. **Research Design and Structuring:**
    - **Purpose:** Used to ensure a logical flow, assist in narrowing the research scope, and organize the key elements of the introduction and conclusion.
        - **Example Prompt 1 (15.01.2026):** "Suggest structural improvements for a chapter discussing augmented analytics in the context of decision latency in supply chain management."
        - **Example Prompt 2 (20.01.2026):** "Suggest improvements to the scope of my thesis taking the delivered ideas into consideration. Can you suggest three different angles that would fit a Bachelor's level study on organizational decision-making?"
        - **Example Prompt 3 (10.02.2026):** "Review my current contents. Is there a logical gap between the theoretical framework in Chapter 2 and the methodology in Chapter 3?"
    - **Verification:** The final structure, research questions, and scope were independently determined. The AI's suggestions served only as a springboard for independent structural planning.

2. **Literature Search and Source Retrieval:**
    - **Purpose:** Generating keywords and constructing search strings for academic databases (e.g., Scopus) to ensure systematic coverage of topics like XAI and decision latency.
        - **Example Prompt 1 (02.03.2026):** "Generate a Scopus-compatible search string using Boolean operators for research focusing on the intersection of Explainable AI and decision latency in organizational contexts. Provide synonyms for the key terms."
        - **Example Prompt 2 (02.03.2026):** "Can you find cited researches and researchers regarding Organizational Information Processing Theory? Can you suggest search words?"
        - **Example Prompt 3 (15.03.2026):** "Suggest searchwords for academic papers that could link Explainable AI to managerial trust or human-AI collaboration."
    - **Verification:** All search strings were manually tested. Sources were retrieved directly from Scopus and Web of Science, and their credibility was assessed using the Finnish Publication Forum (JUFO) levels.

3. **Terminology Clarification and Translation:**
    - **Purpose:** Clarifying complex theoretical constructs and ensuring accurate translation of domain-specific terminology between Finnish and English.
        - **Example Prompt 1 (12.02.2026):** "Explain the core components of OIPT (uncertainty, equivocality, and information processing capacity) in simple terms as if explaining them to a professional outside the field."
        - **Example Prompt 2 (15.02.2026):** "Explain the word decision latency and how is it seen?"
        - **Example Prompt 3 (20.02.2026):** "Translate and explain these terms XAI, black-box theory, interpretability in Finnish from your perspective."
    - **Verification:** Translated terms were cross-referenced with Finnish academic dictionaries and existing theses in the same field to ensure consistency with academic standards.

4. **Textual Refinement and Language Checking**

- **Purpose:** Improving transitions between paragraphs, ensuring an academic tone, and identifying unnecessary repetition in the draft.
    - **Example Prompt 1 (08.04.2026):** "Review the following paragraph for academic tone and clarity. Ensure the transition between OIPT and XAI feels fluid and logical."
    - **Example Prompt 2 (08.04.2026):** "Check chapters 2.2 and 2.3 for redundancy. Are there points where I am repeating the same definition of XAI needlessly?"
    - **Example Prompt 3 (09.04.2026):** "Identify three sentences in my conclusion that are too informal for a Bachelor's thesis and suggest more formal, objective alternatives."
- **Verification:** AI was used as a sophisticated editor. Every sentence was reviewed, and the final wording remains the author's own choice. No text was generated by AI without significant manual revision.

5. **Visualizations and Diagrams**
    - **Purpose:** Designing and rendering conceptual frameworks and process flows into a visual format.
        - **Example Prompt 1 (09.04.2026):** "I have a theoretical model where XAI influences trust, which in turn reduces decision latency. Suggest a layout for a diagram that clearly illustrates these relationships."
        - **Example Prompt 2 (09.04.2026):** "Provide the Mermaid.js code for a vertical flowchart. From top to bottom: raw data, preprocessing, pattern recognition, explanation layer, human review, decision."
        - **Example Prompt 3 (09.04.2026):** "Critique this diagram: Does it accurately reflect the standard way of representing these variables in a conceptual model?"
    - **Verification:** The logic and relationships depicted in all diagrams were verified against the primary literature. The AI assisted only in the technical layout and rendering of the visual elements.