



Autonomous Databases and Artificial Intelligence

Architectures, Optimization, and Governance

Shashipurna Kurapati

● DeepScience
,

Autonomous Databases and Artificial Intelligence: Architectures, Optimization, and Governance

Shashipurna Kurapati

Artificial Intelligence, Data Management



DeepScience

Published, marketed, and distributed by:

Deep Science Publishing, 2025
USA | UK | India | Turkey
Reg. No. MH-33-0523625
www.deepscienceresearch.com
editor@deepscienceresearch.com
WhatsApp: +91 7977171947

ISBN: 978-93-7185-667-6

E-ISBN: 978-93-7185-652-2

<https://doi.org/10.70593/978-93-7185-652-2>

Copyright © Shashipurna Kurapati, 2025.

Citation: Kurapati, S. (2025). *Autonomous Databases and Artificial Intelligence: Architectures, Optimization, and Governance*. Deep Science Publishing. <https://doi.org/10.70593/978-93-7185-652-2>

This book is published online under a fully open access program and is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0). This open access license allows third parties to copy and redistribute the material in any medium or format, provided that proper attribution is given to the author(s) and the published source. The publishers, authors, and editors are not responsible for errors or omissions, or for any consequences arising from the application of the information presented in this book, and make no warranty, express or implied, regarding the content of this publication. Although the publisher, authors, and editors have made every effort to ensure that the content is not misleading or false, they do not represent or warrant that the information-particularly regarding verification by third parties-has been verified. The publisher is neutral with regard to jurisdictional claims in published maps and institutional affiliations. The authors and publishers have made every effort to contact all copyright holders of the material reproduced in this publication and apologize to anyone we may have been unable to reach. If any copyright material has not been acknowledged, please write to us so we can correct it in a future reprint.

Preface

Artificial Intelligence-native databases are currently at the forefront of the rapidly evolving data management landscape. The book examines how database systems are changing to satisfy the needs of real-time, intelligent decision-making in different industries. The transition from traditional relational models to AI-driven architectures, cloud integration, optimization, and new developments like automation, explainability, and security are all covered in the chapters.

This book's writing has involved both a thorough examination of contemporary data technology and a contemplation of the field's continuing opportunities and challenges. I want professionals, students, and anybody else interested in the future of databases to be able to understand both basic and advanced topics. I hope it encourages readers to welcome innovation and investigate the wise opportunities that lie ahead.

I want to express my gratitude to my parents for their unwavering support during my journey, as well as to my peers, fellow researchers, and everyone else who has helped and inspired me. Their guidance and collaboration have been invaluable in shaping this book.

Shashipurna Kurapati

Table of Contents

Chapter 1: The Transformation of Database Technologies: From Relational to AI-Enhanced Systems

- 1. Introduction to Database Technologies.....1
- 2. Overview of Relational Database Systems2
 - 2.1. History of Relational Databases2
 - 2.2. Key Features of Relational Databases2
 - 2.3. Limitations of Relational Databases.....3
- 3. Emergence of Non-Relational Database Systems.....3
 - 3.1. NoSQL Databases4
 - 3.2. NewSQL Databases.....4
 - 3.3. Comparison with Relational Databases5
- 4. Introduction to AI-Enhanced Database Systems5
 - 4.1. Definition and Scope5
 - 4.2. Technologies Driving AI in Databases6
- 5. Applications of AI-Enhanced Database Systems6
 - 5.1. Healthcare7
 - 5.2. Finance7
 - 5.3. Retail7
 - 5.4. Manufacturing8
 - 5.5. Telecommunications.....8
- 6. Case Studies of AI-Enhanced Systems9
 - 6.1. Case Study 1: AI in Healthcare Databases9
 - 6.2. Case Study 2: AI in Financial Services.....9
 - 6.3. Case Study 3: AI in Retail Operations.....10
- 7. Challenges and Ethical Considerations10
 - 7.1. Data Privacy Issues11
 - 7.2. Bias in AI Algorithms.....12
 - 7.3. Regulatory Compliance.....12
- 8. Future Trends in Database Technologies13
 - 8.1. Integration of AI and Machine Learning.....13

8.2. The Role of Cloud Computing	14
8.3. Emerging Technologies	14
9. Conclusion	14
Chapter 2: Architectural Strategies for Managing Databases in AI Environments.....	17
1. Introduction to AI and Database Architecture	17
2. Architecting Databases for AI Workloads.....	18
2.1. Data Lakehouses	18
2.2. Multimodal Storage	19
2.3. Unstructured Data Handling	20
3. Serving Machine Learning Models from Database Systems	20
3.1. Integration of ML Models within Databases	20
3.2. Performance Optimization Techniques.....	21
4. Real-time Feature Stores and Streaming Architectures	21
4.1. Designing Real-time Feature Stores.....	22
4.2. Implementing Streaming Architectures	22
5. Data Management Techniques for AI.....	23
5.1. Data Governance and Compliance.....	23
5.2. Data Quality and Validation	24
6. Scalability Considerations in AI Database Architectures	24
6.1. Horizontal vs. Vertical Scaling.....	25
6.2. Load Balancing Strategies	25
7. Security and Privacy in AI Database Management.....	25
7.1. Data Encryption Techniques	26
7.2. Access Control Mechanisms	26
8. Case Studies in AI Database Architectures	27
8.1. Industry Applications of Data Lakehouses.....	27
8.2. Real-world Examples of ML Model Serving.....	28
9. Future Trends in Database Management for AI	28
9.1. Emerging Technologies and Innovations	29
9.2. Predictions for AI Database Architectures	29
10. Conclusion.....	30

Chapter 3: Exploring the Impact of AI on Query Optimization and Database Performance Tuning

.....

32

1. Introduction to AI in Database Management

32

2. Fundamentals of Query Optimization.....

33

3. AI-Powered Query Optimization and Tuning

33

3.1. Reinforcement Learning in Query Planning

34

3.2. Cost Models Driven by Machine Learning

34

3.3. Self-Tuning Databases

34

4. Case Study: Oracle's AI Capabilities

35

5. Case Study: Azure SQL Hyperscale.....

35

6. Case Study: Snowflake's AI Features

36

7. Comparative Analysis of Database Systems.....

36

8. Challenges in AI-Driven Query Optimization

37

9. Future Trends in AI and Database Performance.....

37

10. Ethical Considerations in AI Implementation

38

11. Performance Metrics for AI-Driven Systems.....

38

12. User Experience and AI in Databases

39

13. Integration of AI with Traditional Optimization Techniques

39

14. Impact of Cloud Computing on Database Performance.....

39

15. AI for Predictive Analytics in Databases.....

40

16. Real-World Applications of AI in Database Management

40

17. The Role of Data Quality in AI Optimization

41

18. Security Implications of AI in Databases

41

19. Regulatory Compliance in AI-Enhanced Systems

42

20. User Training and AI Systems.....

42

21. Collaboration Between IT and Data Science Teams

42

22. Cost-Benefit Analysis of AI Implementation.....

43

23. Conclusion.....

43

Chapter 4: Embedding Intelligence into Data Pipelines: Exploring the Intersection of MLOps and DataOps for Enhanced Automation and Quality Assessment

.....

46

1. Introduction to MLOps and DataOps

46

2. The Convergence of MLOps and DataOps.....	46
2.1. Historical Context	47
2.2. Key Principles of MLOps.....	47
2.3. Key Principles of DataOps	48
2.4. Benefits of Integration	48
3. Automating ETL/ELT with AI.....	48
3.1. Overview of ETL/ELT Processes	49
3.2. Role of AI in Automation.....	50
3.3. Tools and Technologies for Automation.....	50
3.4. Case Studies of AI-Driven ETL/ELT	51
4. Data Quality Assessment.....	51
4.1. Importance of Data Quality	52
4.2. Traditional vs. AI-Driven Approaches	52
4.3. Frameworks for Quality Assessment.....	52
5. Anomaly Detection with Machine Learning.....	53
5.1. Understanding Anomalies in Data	53
5.2. Machine Learning Techniques for Anomaly Detection.....	54
5.3. Implementation Strategies	54
5.4. Real-World Applications.....	55
6. Challenges and Solutions in MLOps and DataOps Integration.....	55
6.1. Cultural and Organizational Barriers.....	56
6.2. Technical Challenges.....	56
6.3. Best Practices for Overcoming Challenges	56
7. Future Trends in MLOps and DataOps.....	57
7.1. Emerging Technologies	57
7.2. Predictions for the Next Decade	58
8. Conclusion	58
 Chapter 5: The Age of Vector and Graph Databases: Foundations for Advanced Information Retrieval and Reasoning	
1 Introduction.....	61
2 The Age of Vector and Graph Databases	62

3 Vector Embeddings and Semantic Search.....	63
4 Integrating with LLMs for Retrieval-Augmented Generation (RAG)	64
5 Knowledge Graphs and Reasoning Engines	65
6 Synthesis and Future Directions	66
7 Conclusion	67
References	68
 Chapter 6: Exploring Security, Governance, and Explainability in AI Systems	70
1. Introduction.....	70
2. Understanding AI Systems	70
2.1. Definition of AI	71
2.2. Types of AI Systems.....	71
2.3. Applications of AI.....	72
3. Security in AI Systems.....	72
3.1. Threats to AI Security	73
3.2. Mitigation Strategies.....	73
3.3. Case Studies in AI Security Breaches	74
4. Governance of AI Systems	74
4.1. Frameworks for AI Governance	75
4.2. Roles and Responsibilities.....	75
4.3. Policy Development and Implementation	76
5. Explainability in AI	76
5.1. Importance of Explainability	77
5.2. Techniques for Explainable AI.....	77
5.3. Challenges in Achieving Explainability.....	78
6. Anomaly Detection in AI Systems.....	78
6.1. Methods of Anomaly Detection	79
6.2. Applications of Anomaly Detection.....	79
6.3. Impact on Security and Governance.....	79
7. Data Lineage in AI Systems.....	80
7.1. Understanding Data Lineage.....	80
7.2. Tools for Tracking Data Lineage.....	81

- 7.3. Importance for Compliance and Governance81
- 8. Compliance with Explainable AI82
 - 8.1. Regulatory Requirements82
 - 8.2. Best Practices for Compliance82
 - 8.3. Case Studies on Compliance Issues83
- 9. Integration of Security, Governance, and Explainability83
 - 9.1. Holistic Approach to AI Management.....84
 - 9.2. Interdependencies Among Security, Governance, and Explainability84
- 10. Future Directions in AI Systems84
 - 10.1. Emerging Trends85
 - 10.2. Research Opportunities85
 - 10.3. Ethical Considerations86
- 11. Conclusion.....86

Chapter 7: Exploring Case Studies, Industry Implementations, and Future Research Directions in AI, Big Data, and Blockchain Technologies89

- 1. Introduction89
- 2. Overview of AI Technologies89
 - 2.1. Machine Learning90
 - 2.2. Natural Language Processing.....90
 - 2.3. Computer Vision91
- 3. Big Data Fundamentals91
 - 3.1. Data Collection Techniques.....91
 - 3.2. Data Storage Solutions.....92
 - 3.3. Data Analysis Methods92
- 4. Blockchain Technology Overview93
 - 4.1. Decentralization93
 - 4.2. Smart Contracts94
 - 4.3. Consensus Mechanisms94
- 5. Case Studies in AI94
 - 5.1. Healthcare Applications95
 - 5.2. Financial Services96

5.3. Retail Innovations	96
6. Industry Implementations of Big Data	96
6.1. Telecommunications	97
6.2. Manufacturing	97
6.3. Marketing Analytics	98
7. Blockchain in Various Industries	98
7.1. Supply Chain Management	99
7.2. Real Estate Transactions	99
7.3. Voting Systems	100
8. Challenges in AI Implementation	100
8.1. Ethical Considerations	101
8.2. Data Privacy Issues	101
8.3. Algorithmic Bias	102
9. Big Data Challenges	102
9.1. Data Quality Management	103
9.2. Scalability Issues	103
9.3. Integration with Legacy Systems	104
10. Blockchain Challenges	104
10.1. Scalability	105
10.2. Regulatory Compliance	106
10.3. Interoperability	106
11. Future Research Directions in AI	107
11.1. Explainable AI	108
11.2. AI in Climate Change	108
11.3. AI for Cybersecurity	109
12. Future Research Directions in Big Data	109
12.1. Real-time Data Processing	110
12.2. Data Ethics and Governance	110
12.3. Predictive Analytics	111
13. Future Research Directions in Blockchain	111
13.1. Blockchain for Social Good	112
13.2. Integration with IoT	113

13.3. Cross-chain Solutions	113
14. Interdisciplinary Approaches.....	114
14.1. Collaboration between AI, Big Data, and Blockchain.....	115
14.2. Case Studies of Interdisciplinary Projects	115
15. Conclusion.....	116

Chapter 1: The Transformation of Database Technologies: From Relational to AI-Enhanced Systems

1. Introduction to Database Technologies

Database technologies are an integral part of the global economy. They present an attractive area of investment for venture capitalists, are a fundamental component of most software projects, and moderate our personal, social, and financial lives with their representations of individuals and the everyday. However, the hype that surrounds them is a far cry from their humble roots as a simple means of recording and querying progress on budget and hiring tasks. They eventually were reinterpreted as the key to information management and enabling business administrations to be competitive in increasingly digital economies.




Fig 1. The Transformation of Database Technologies: From Relational to AI-Enhanced Systems

They were subsequently understood as the catalyst for creating better products and services in almost every walk of life, culture, and pursuit.

2. Overview of Relational Database Systems

Relational Database Management Systems (RDBMS) have long been the foundation for storing and retrieving data for businesses across all industries worldwide, having been introduced in the 1970s. These systems are based on a rigid structure—the data model, formally defined through the relational model—that defines data objects and the relationship between them. This approach offers great flexibility; data relationships can be created, updated, and deleted easily, queries are built based on the data model rather than the physical model, and the databases are highly scalable and fast. However, it also features many rigid characteristics, with inflexible schema, less support for complex queries, and less support for other types of data (such as graph, spatial, or object data). The relational model has provided a highly established and proven approach for storing and retrieving data. Transactional systems largely use relational databases with OLTP for applications such as core banking, manufacturing, billing, and so on.

2.1. History of Relational Databases

Database technologies have undergone dramatic transformations since the foundations of relational databases were laid in the 1970s. In the early years, relational database management systems (RDBMSs) completely overshadowed all other approaches to data management. However, as modern applications became increasingly complex and database designs were targeted to serve a broader range of services, the limitations of the relational technology started to surface[1]. Two classes of non-relational database systems emerged at the beginning of the new millennium as an alternative approach for data management in the cloud, for Big Data analytics, and for large-scale data streaming. NoSQL systems enabled the manipulation of semi-structured and structured databases, whereas NewSQL systems supported structured and relational data.

Today, a new wave of database systems is on the rise. Next-generation systems, known collectively as AI-Enhanced Database (AI-DB) Systems, are enriched with AI capabilities that are designed to handle one or more stages of the data pipeline. These new systems leverage the extensive use of Artificial Intelligence (AI) and Machine Learning (ML). Progress in ML has also led to an increased interest in the physical design of modern workloads, mainly because of the impact that an optimal physical design and choice of configuration parameters can have on the price and performance.

2.2. Key Features of Relational Databases

Relational databases organize data in tables composed of rows and columns. Each table has a key column that uniquely identifies a relationship inside the table and acts as a key in other tables to

form relationships across the database. Tables can be combined with operations like joins to create new tables for query results. Referential integrity constraints control how tables can be modified without breaking relationships. No-sql databases do not enforce these constraints. Relational databases use the SQL language for data query and programming. In the last decades, non-relational database systems appeared with models like no-SQL and new-SQL. Non-relational databases have an edge over relational databases in terms of scalability, availability, and fault tolerance. However, non-relational databases should not be considered as an opposition to the relational database model.

Rather, the emergence of non-relational database systems fills the gap of relational database systems by addressing the current modern business requirements, such as the cloud, the large size of data, and the intensive methods of extracting meaning from data, driven mainly by artificial intelligence. The ability of XAI databases to explain their outputs results in a reduction of the number of processing steps that are currently done outside the database and at a higher level. Such ability also leads to the creation of new applications in areas where explainability of AI is crucial, such as healthcare, financial services, manufacturing, telecommunications, and retail.

2.3. Limitations of Relational Databases

Even though relational databases have served the world well, their atomic grouping of data into rows is not always desirable. Relational databases store data at the atomic level in tables, a two-dimensional structure of rows and columns. Each entry in a table must be atomic, indivisible. This constraint prevents a single cell from holding more than one value. For instance, storing a customer's phone numbers becomes challenging: each number requires a separate row, or a fixed maximum number of columns must be allocated, though the actual number of phone numbers varies between customers [1-3]. When the data model is not normalized to a third-normal form, the data cannot be stored in an RDBMS. If the model is not normalized for a business report, a NoSQL database with hierarchical storage is ideal.

Postgres has overcome this limitation by allowing non-atomic groups to be stored in a cell. Postgres is a great RDBMS for geospatial queries. Another disadvantage of RDBMSs is their difficulty in scaling across distributed computing architectures like clusters or cloud-scale virtual private servers that evolve and add more machines over time.

3. Emergence of Non-Relational Database Systems

Non-relational databases—based on so-called NoSQL (not only SQL) and NewSQL models—have evolved to address many of the limitations embedded in the relational model. NoSQL databases are characterized by flexible schemas, simplified design, and horizontal scaling, all of which contribute to efficient handling of large volumes of diverse, semi-structured, and unstructured data without the rigidity imposed by schemas and relations. Consequently, NoSQL supports iterative and agile development, enabling rapid data extraction and querying.

On the other hand, NewSQL databases preserve the anterior relational models and SQL querying standards, including ACID transactional properties, yet overcome drawbacks through non-locking concurrency control mechanisms, distributed architectures, and elastic resource allocation. Both NoSQL and NewSQL represent a response to the expansion of database-related challenges not amenable to solutions offered by commercial relational systems.

3.1. NoSQL Databases

NoSQL, an abbreviation for “not only SQL,” designates a broad range of data management models that do not adhere to the relational model and SQL query language. NoSQL models acquired popularity in the first decade of the 2nd millennium thanks to their demonstrated capability to manage Big Data and to scale on machine clusters.

Using a different terminology, one might state that NoSQL models typically relax one or more of the ACID principles that underlie relational databases. The motivation behind such relaxations is their reconciliation with the CAP theorem, which establishes that a distributed database cannot simultaneously provide strong consistency, availability, and high tolerance to network partitions (due to node failures). NoSQL models sacrifice some degree of either consistency or availability to preserve their ability to cope with node failures. The main NoSQL categories with representative examples are as follows: document (MongoDB), key–value (Dynamo), column (Big Table), and graph (Neo4J).

3.2. NewSQL Databases

The limitations of NoSQL databases have resulted in the emergence of various NewSQL databases, exemplified by MemSQL (now SingleStore). As described by designer.io, MemSQL, a distributed, scalable, relational database supporting SQL queries, combines the benefits of traditional RDBMS and NoSQL. This fusion enables transactional applications while delivering high performance, scalability, and simplicity for real-time analytics solutions. MemSQL achieves this through a distributed, shared-nothing architecture, capturing SQL workflows and ACID semantics of transactional databases, while providing NoSQL-like scalability and operability. Consequently, it supports the command of SQL for data definition, manipulation, query, and transaction, and offers comprehensive relational database services with consistent low latency and high throughput. MemSQL delivers one unified database service for operational analytics, delivering rich data insights on real-time transactional data.

Thus, the DBMS market has developed multiple alternative database models and technologies to overcome the limitations of relational database systems. NoSQL databases represent a family of nonrelational databases designed to address the challenges of managing large volumes of distributed, semi-structured, and unstructured data while supporting the rigidity of relational databases. The eight prominent NoSQL databases outlined above—from columnar CQL through to graph Gremlin—offer drastically improved storage management and query processing capabilities for novel application varieties, including web, IoT, and big data. Nevertheless,

limitations remain. NewSQL databases seek to reconcile the benefits of SQL support and ACID semantics with the scalability and low latency of NoSQL, delivering a unified database service for operational analytics.

3.3. Comparison with Relational Databases

Relational databases have been dominant since IBM published the paper by Codd (1970). Yet certain shortcomings have limited their wide acceptance. First, they are generally not good at handling very large amounts of unstructured data, such as images and video. Second, scaling up relational databases is very costly, since it requires more powerful CPUs and large amounts of RAM—a practice known as vertical scaling. Third, the dependency on the rigid schema tends to complicate the ETL (Extract, Transform, Load) processes to feed the data into the database. Finally, performing efficient large-scale analytics for data stored in relational databases tends to be challenging.

Over the years, these limitations have resulted in the emergence of non-relational database technologies, popularly labelled as NoSQL. Subsequently, the NoSQL approach was complemented by NewSQL technologies. The term NoSQL was first coined in late 1998 for a lightweight open-source relational database that did not expose the standard SQL interface found in traditional databases. However, the term NoSQL gained wide popularity only in 2009, when Johan Oskarsson used it for a meetup discussing open-source distributed databases. The meetup included databases such as Redis, Cassandra, and Neo4j. EM Codd’s 1970 paper continues to shape data management.

4. Introduction to AI-Enhanced Database Systems

The term “AI-enhanced database” and the concept of “AI in databases” describe the growing integration of artificial intelligence into database technologies and database management systems. This development utilizes new technologies that support artificial intelligence to enhance or supplement existing data management processes.

In recent years, several industries—healthcare, financial services, retail, manufacturing, and telecommunications—have augmented or replaced traditional databases with AI-enhanced alternatives. These systems support new applications that incorporate artificial intelligence and machine learning, enabling more rapid decision-making with access to a broader range of unstructured data. The trend toward AI-enhanced systems aims to improve current machine-learning capabilities by seamlessly integrating user intent into the database, thereby facilitating advanced analytics and operational efficiency.

4.1. Definition and Scope

Database technologies are the backbone of modern data management; they enable organizations to efficiently store, manage and analyze large amounts of information. This support of decision-

making processes is crucial in a digital environment where business success is often directly linked to the quality of the decisions made.

Relational database management systems account for a large portion of industrial and commercial deployments in the last three decades [3]. Commercial implementations such as Oracle Database and Microsoft SQL Server are two of the twenty most powerful computer programs of all time. However, relational databases also present limitations and therefore NoSQL / NewSQL non-relational models have emerged to address these shortcomings.

4.2. Technologies Driving AI in Databases

The integration of AI in database technology encompasses a range of discrete yet interrelated developments. For instance, Turbo-VCM combines probabilistic data models with machine-learned components, automating the inference of group-by location for contextual visualizations. Data requirements analysis, as performed by AIPlanner, utilizes AI planning to construct detailed plans that match user-supplied questions. Additionally, machine-learned components comprising AI-EDLC are designed to alleviate the complexity of data wrangling, cleaning, and integration.

Beyond AI-enhanced database systems, advancements in cloud computing have catalyzed the rise of specialized data ecosystem services—including popular cloud analytics engines such as Snowflake and BigQuery. The services encompass data-messaging, storage, security, and governance. Furthermore, even within the domain of relational database systems, interest in the transformative impact of AI and machine learning has intensified.

5. Applications of AI-Enhanced Database Systems

A wide range of industries rank among the early AI adopters. As AI is embraced and applied, the volume and scope of available data is rapidly increasing, triggering a rising-bar effect on the underlying databases and database-management systems. This cause—effect relationship highlights the potential of AI-enhanced database systems to contribute to innovative applications across multiple traditional and new AI sectors.

The currently evolving applications of AI-enhanced database systems address emerging business needs, as illustrated by examples across various industries. Enterprises in healthcare, financial services, retail, manufacturing, and telecommunications face common challenges in using data for increased revenue growth, cost reduction, and risk minimization. However, AI's actual impact goes much "deeper," producing a transformational effect on the selected use-case category, the corresponding sector, and, quite often,--amplified through the supply chain or related sectors--the overall economy. Many of these applications echo the key trends identified in other sectors, and their lessons and principles can be extended to other areas. Three examples further illustrate the stage and breadth of AI application in database systems.

5.1. Healthcare

The healthcare industry, with its vast data volumes, faces clinical data management challenges that traditional systems struggle to address. AI-enhanced healthcare systems offer solutions that lower costs, improve patient safety, and deliver quality medical care. The advantages of an AI-enhanced database can be seen in its ability to categorize, discriminate, transform, forecast, and prescribe, assisting in decision-making processes. For example, associating disease patterns with X-ray images aids in making accurate, quick decisions; banking fraud detection guards against account hacking; customer income prediction supports financial consultancy; and manufacturing winner selection guides marketing and manufacturing strategies. Although the key database technology employs a Data Warehouse Engine, Artificial Intelligence is instrumental in ensuring the safety and quality of human life.

5.2. Finance

Artificial intelligence (AI) applications are transforming banking and the broader financial services industry. AI technologies are being deployed for a wide range of tasks, from setting credit and insurance policy rates to determining which applications for loans, mortgages, and insurance benefits to approve [2,4]. AI is also highly beneficial for detecting credit card fraud, with most major credit card providers employing machine learning algorithms specialized for that purpose.

AI applications facilitate cross-selling for banks and hedge funds, analyze the effectiveness of advertising campaigns in financial services, and recommend preferred shares and bonds for individual investors based on their risk tolerance. Additionally, investment management firms utilize AI-powered chatbot assistants and document analysis tools to help investors make informed decisions. In treasury departments, interest rate forecasting applications employ AI techniques to minimize risk. Before the pandemic, many financial services firms were already deploying artificial intelligence. Enhanced customer experience, greater efficiency, reduced costs, and improved operational control were the top benefits driving adoption. During the COVID-19 crisis, these benefits became even more apparent, leading more institutions to implement AI.

5.3. Retail

AI-enhanced database systems are deployed in the retail and e-commerce sectors to improve the customer experience through real-time responses to questions or issues during the decision-making process while shopping online, at home, on their phone or computer, or in the store. The promise of AI is that data can be analyzed in real-time and insights delivered to customers in the form of personalized experiences and recommendations based on their spending habits. Real-time targeted advertising can be presented in price-reduced coupons through emails or ads on social media sites. Physical stores can reduce operating costs using AI to monitor security, loss prevention, customer movement and products they pick up but don't purchase.

The retail industry uses the vast amount of customer data it collects to manage inventory, change marketing strategies in real time, and identify products to suggest to each customer during their shopping experience by leveraging AI technologies. The data as well as the predictions can be accessed on mobile devices by the decision makers of the company. Both online and brick-and-mortar stores have increased their use of robotics to replenish shelves, fetch items in the store and fulfill online orders. Chatbots acting as virtual sales associates guide shoppers through the stores, highlighting promotions and advising on complementary products while improved geographic information systems (GIS) map customer shopping patterns for the retailer.

5.4. Manufacturing

The manufacturing sector is adopting AI-driven database technologies to achieve increased production efficiency and production delivery. By integrating AI-driven systems with IoT equipment, real-time connection to production lines generates extensive data, which is then cleaned, processed, stored, and analyzed continuously by AI-enhanced database systems. The identified challenges in production process management can be addressed by AI-integrated solutions; however, adequate data security is essential to mitigate potential leakage during the analysis phase.

Manufacturing enterprises seek continuous improvements in production efficiency, cost reduction, and overall increased product quality and related services. Challenges involving the complex interactions between physical and logical devices can have significant impacts when overlooked [5-8]. An intelligent human-machine environment supported by AI enables manufacturers to efficiently schedule production, dynamically manage the supply chain, and assist in machine maintenance optimization. Real-time and effective management information services for production lines are critical to the development of smart factories.

5.5. Telecommunications

The transformation of telecommunications is sustained if not enabled by advances in data collection and data analysis. The Big Data challenge is considerable: A Cisco report forecasted in 2017 a quadrupling of global mobile data traffic, underpinned by a 10-fold increase in the number of mobile devices [53]; and a McKinsey analysis for the European Union concluded that within Europe the telecom operators would be responsible for managing a share of Big Data responsible for 50 to 80 exabytes of new data annually, and that they are well-positioned to help other sectors manage their Big Data [54]. Public networks are evolving towards 5G, which for the first time has been designed to deliver highly reliable and low latency connectivity that supports flexible network slicing and an increasing number of devices per cell. This enables new services as well as new business and revenue models for the network operators.

Adoption of AI-enhanced database technology in telco markets requires a flexible platform that can process large training sets, prepare and cleanse the data, and apply machine learning. In a broad cross-industry McKinsey analysis, three applications for telcos were highlighted: real-time

multimedia translation and analysis, cybersecurity, and fraud detection [51]. Dialpad, a provider of cloud telephony and voice, video, and conferencing solutions, is using AI to create customer-facing applications with speech performance and emotion recognition, in addition to assisting support agents. In Hungary, Telekom is working with Accenture to reduce customer churn, improve marketing campaign effectiveness, and enhance the overall customer experience using AI.

6. Case Studies of AI-Enhanced Systems

Examples of AI-enhanced database systems are emerging in many industries, including health care, financial services, and retail. Customer support in all industries is being transformed by AI chatbots, including banking, airline travel, and supply chain management. Surveillance systems in manufacturing, logistics, and telecommunications are enriched by image recognition systems.

Health care institutions are increasing database performance and security with AI and implementing AI applications. These include disease monitoring, remote patient monitoring, image analysis, and prediction of potential epidemics. Banks and other financial companies are using AI-based database services to improve power, performance, and risk analysis, as well as fraud detection. Several companies use AI-driven database services to support vibration analysis, image recognition, trash detection, facial recognition, and natural language processing.

6.1. Case Study 1: AI in Healthcare Databases

AI-enabled databases have begun to transform many real-world applications, including healthcare, financial services, retail, manufacturing, and telecommunications. In each case, incorporating AI capabilities into the database provides a competitive advantage. The following examples illustrate three of these applications.

An exponentially growing population, along with advancements in treatment methods and healthcare facilities, is generating unprecedented volumes of healthcare-related data. This data explosion renders the efficient assessment and diagnosis of patients more difficult. By utilizing AI-enabled databases, healthcare organizations can support physicians with advanced decision support systems. These systems analyze vast amounts of data to detect, interpret, and predict trends in order to establish a connection between patient scenarios and known outcomes. Hospitals can then apply these techniques to develop advanced Intelligent Patient Assessment systems that rapidly analyze patient data and conditions for swift and accurate diagnosis.

6.2. Case Study 2: AI in Financial Services

Information is the foundation of every business. Companies that use their data effectively can sharply improve decision-making and deliver faster turnaround times. The principle applies to every individual in a company and at all levels—from entry-level agents to executives, from providing answers for customers, to making the right call when investing, to managing a company in a manner that achieves maximized value for its shareholders [6,9]. The growth of

artificial-intelligence (AI) tools is accelerating. As these tools gain market penetration, the impact on the database market is significant. AI-enhanced database systems are those that use various database and other technologies in conjunction with AI technology. The term “AI-enhanced” refers to the fact the database is strengthened in some manner by the integration of AI and associated technologies.

Chatbots are a foundational AI use case and continue to evolve toward smarter answers as machine-learning and natural-language-processing techniques advance. Even relatively unsophisticated AI applications can reduce costs and improve customer satisfaction significantly. Major financial institutions depend on AI-assisted technologies such as optical character recognition (OCR) for simple document handling and authentication tasks such as mortgage applications and credit card applications. AI tools help regional banks and credit unions battle fraud, comply with increasingly complex regulations, and generate more business from cross-selling and up-selling. Other applications include virtual financial planners, analysis of investment possibilities, and fraud detection, identification, and prevention. Case studies confirm that AI, in large and small banks alike, delivers better efficiencies, improves quality, reduces workload, and generates higher profits.

6.3. Case Study 3: AI in Retail Operations

Rapid developments in artificial intelligence (AI) have affected different levels of database management systems (DBMSs) and transformed applications through enhanced data intelligence [10]. AI engines now contribute to closing the gap between data and value by enhancing understanding and facilitating the extraction of value. Businesses adopting AI-enhanced DBMSs and exploiting the enhanced intelligence across DBMS components showcase significant advantages across diverse sectors.

Retailing is a prime example of where AI has revolutionized corporate activities. Data from internal vertical chains—suppliers, distributors, warehouses, retail stores—and historical activities serve as a major application area. Given the growing importance of retailing, especially in developed countries, retail data analysis is increasingly recognized as a strategic element for corporate success. AI's application in retailing is dedicated to making corporate activities more responsive to end-user needs through data intelligence. It contributes to process innovation and functional management, ultimately enhancing value to customers. Deployment areas include sales forecasting, customer profiling, inventory planning and management, cloud resource optimization, customer relationship management (CRM), and strike detection, illustrating the broad integration of AI into retail operations.

7. Challenges and Ethical Considerations

AI-enhanced databases are increasingly addressing key issues faced by industry verticals, including database security provisioning, real-time data insights, and regulatory compliance with GDPR, Basel III, among others. The management of access privilege is critical to ensure data

confidentiality and prevent unauthorized access or data breaches. Automated database activity monitoring and the application of AI can provide real-time alerts regarding abnormal or suspicious database activities, particularly during events such as terrorist attacks or bank robberies.

However, recent research has highlighted concerns regarding bias in AI algorithms. Examples include facial recognition systems that predominantly identify white males and the targeting of black men by online crime prediction tools. Accordingly, various institutions have issued ethical guidelines for AI and related applications. AI-enhanced database systems are penetrating diverse industry verticals, offering smarter and more efficient services. In healthcare, AI analyzes patient data to determine disease susceptibility and optimize treatment recommendations. The financial sector employs AI to detect suspicious transactions and prevent fraud. Retail companies utilize AI algorithms to tailor products and services to customer preferences, while the manufacturing industry leverages AI for quality control, fault prediction, and inventory management. Customer churn prediction also benefits from these intelligent systems. Telecom service providers use AI to detect fraudulent calls and spam.

7.1. Data Privacy Issues

Data security plays a vital role in establishing AI-enhanced database systems for industrial applications. Privacy, confidentiality, integrity, and security remain the top concerns for data repositories in both relational and non-relational database management systems. The recent data breaches at leading corporations such as Facebook and Amazon have highlighted the criticality of protecting user data stored in AI-enhanced information systems. These risks can be addressed by adopting a procedure of data anonymization or deidentification to ensure individual privacy during the data analysis process [10-12]. The disclosure of sensitive data can generate potentially devastating effects on the privacy of individuals during a retrieval or data-analysis process.

In the early 1970s, researchers at the Federal Trade Commission in the USA recognized that the disclosure of sensitive and confidential information might lead to potentially devastating effects on individual privacy. One of the reasons for potential privacy leakage is the association between individual information in the published data and their identities. The release of personal information of individuals for conducting data analysis has been acknowledged as a privacy disclosure vulnerability. Credit card companies, governmental organisations, hospitals, banks, and many other institutions collect a lot of information about their customers, and their privacy is vital for a safe and fair society. Indeed, most of the practices of the companies that handle the data of their customers fall under data regulation. Organisations use the information of their customers to support their business models and strategies. Sensitive and confidential data are protected against unauthorised individuals or organisations; however, these data can, in principle, be accessed by parties internal to the company for use in various activities. Thus, the exploitation of these data did not balance privacy concerns and interests of the public.

7.2. Bias in AI Algorithms

The emergence of powerful AI algorithms has led to the integration of cognitive capabilities into traditional large-scale data management systems, giving rise to the AI-enhanced Database Technology (AI-DB) research area. Databases serve as the backbone of many organizations, underpinning critical business applications across sectors such as financial services, healthcare, retail, manufacturing, telecommunications, and government agencies. AI-DB systems leverage AI techniques to extract meaningful insights from vast and varied data repositories. Through the inclusion of AI modules and components, these systems can analyze, interpret, infer, and even make decisions based on their analysis.

Research has explored the deployment of AI-DB in several application domains. Healthcare organizations can employ AI-DB solutions to detect and prevent diseases using medical records; for example, systems can monitor blood pressure data to identify health deterioration. Bias in AI algorithms, however, poses significant concerns that cannot be overlooked. Pattern analysis in the financial services industry can identify fraudulent transactions; bias may lead to either false accusations or missed fraudulent activities. Retailers analyze POS and inventory data using AI-DB to forecast sales trends that influence operational and promotional activities. Task scheduling during production can be optimized to reduce manufacturing costs through pattern analysis. Additionally, pattern analysis of network data plays a critical role in enhancing customer satisfaction ratings in the telecommunications sector. The implementation of AI-enhanced capabilities in modern database systems empowers organizations to optimally allocate resources. Ethical issues warrant serious consideration, especially when AI-DB assists in regulatory areas such as child protection, teacher monitoring, debt management, credit control, and arrest warrant decisions.

7.3. Regulatory Compliance

When an AI system makes a decision, the decision may be traceable to a database that is no longer in compliance with the regulations in place at the time or place of the decision. For example, a machine learning model trained on data obtained under GDPR may be used to enable a decision that is always made outside the scope of GDPR, thereby creating potential liability for the data controller. Probes emerge in the context of an audit and are designed to test the machine learning model for compliance with GDPR principles and regulations. Regulatory compliance probes are typically implemented using black-box testing techniques, which involve providing inputs to a system and observing the outputs, without any knowledge of the internal structure, logic, or handling of data.

Regulatory compliance is critical because an AI system that is either in or out of regulatory compliance cannot be described as fair or ethical. Furthermore, regulatory compliance plays a significant role since AI-related jobs are increasingly regulated and governed by entities such as the EU, IEEE Standards Association, the United Kingdom, the United States, and the Organization for Economic Cooperation and Development (OECD). Financial regulatory bodies,

while currently concentrated on the financial services sector, are beginning to investigate and assess other industries that utilize AI, focusing on AI credit decisioning, AI recommendations, and AI pricing.

8. Future Trends in Database Technologies

The ongoing integration of artificial intelligence (AI) and machine learning (ML) with advanced database systems and the ever-expanding capabilities of cloud computing are expected to transform the future of database technologies. Continuous research is focused on developing more intelligent and highly automated database systems, including systems with DBMS-internal AI capabilities. Driven by the hybrid edge cloud paradigm, the rise of geodistributed, multicloud, and polymorphic database services also points to a future of database services that are more universally accessible, flexible, and elastic, thus catering to business and user needs more in spirit with the meaning of Database-as-a-Service.

Interest in trend-setting new database technologies remains as strong as ever. AI- and ML-enhanced (or empowered) database and data-intensive decision-making systems will play a key role in the digital transformation of many industries. Recent extraordinary advances in AI, powered by deep learning and large neural network models, point to the prospect of greatly enhanced, semantically oriented natural language and multimodal (e.g., speech and vision) capabilities that promise to address the persistent data management challenges of extracting rich business intelligence and knowledge from unstructured data, as well as from structured and text data combined. Demonstrating the cross-industry impact of AI-enabled data management, the applications range from managing electronic health records and enabling COVID-19 research for improved patient diagnosis, treatments, and policy recommendations, to financial decision-making, personalized retail, production plans in manufacturing, and network optimization in telecommunications.

8.1. Integration of AI and Machine Learning

Over recent years, a host of novel types of database systems—often classified as NoSQL and NewSQL models—have emerged to overcome some of the many limitations of traditional relational database systems. The addition of artificial intelligence (AI) and machine learning (ML) technologies constitutes the most recent phase of this transformative evolution; the resulting AI-enhanced systems are already delivering advanced capabilities for data and information management.

An AI-enhanced database system represents any database product that incorporates fundamental AI elements—such as natural language processing, knowledge representation, automated reasoning, machine learning, and computer vision—or relies on AI-based technologies to provide support for its development and operational phases. The broad range of AI techniques applied within database-supportive products continues to grow, and their applications span multiple industries.

8.2. The Role of Cloud Computing

Cloud computing is an architectural model for delivering shared pools of computing resources scalable on-demand. Public cloud databases capitalise on the elasticity and scalability of their underlying cloud computing environment. At the same time, cloud providers share the responsibility of security with their customers and implement rigorous compliance controls and regular independent audits. Cloud providers such as Amazon, Microsoft, and Google built data centres in many countries to accommodate customer requests for legal boundaries on data.

8.3. Emerging Technologies

Subsections 8.3.1, Emerging technologies, and 8.3.2, Applications, focus on new approaches to and uses of database technology. Database technology has come a long way since early systems such as IBM's IMS and Relational Database Systems, notably Oracle. NoSQL and NewSQL systems exploit different storage and processing models in order to meet the diverse requirements of modern applications. Artificial Intelligence (AI)—more precisely, Deep Learning combined with Cloud Computing—plays a leading role in the transformation of database technologies into AI-Enhanced Database Systems, which are similarly being applied in a variety of contexts.

8.3.1. Emerging technologies. Artificial Intelligence has long been a popular research topic. Breakthroughs in Deep Learning, combined with the growth of Cloud Computing and Cloud Storage, now enable practical applications in many fields. 8.3.2. Applications. The integration of Artificial Intelligence technologies is transforming many industries. Recent studies and reviews indicate that AI has achieved remarkable results in areas such as clinical healthcare, financial services, and retail [7,13-15]. In the healthcare sector, AI tools enable doctors to analyse clinical data for early identification of patients at risk of critical conditions, while equally important applications include supporting clinical decisions. In the financial industry, AI-based algorithms assist users with financial advice and offer services such as fraud detection. Other industries seeking to optimize their operations with a digital transformation approach also increasingly rely on AI technologies. Manufacturing businesses use AI for quality control and preventive maintenance, whereas the telecom sector exploits AI to enhance the customer experience. The integration of Artificial Intelligence technologies is thus transforming many aspects of everyday life.

9. Conclusion

Database technologies have come a long way since the early 1970s, when Edgar Frank Codd of IBM proposed the relational database model. The proliferation of data sources, the increasing demand for data storage associated with new emerging areas such as big data and the Internet of Things, and the need for faster processing of read and write transactions have highlighted many limiting characteristics of traditional relational database implementations. These limitations have led to the development of non-relational (NoSQL) database management systems, which differ

from traditional relational databases in various ways. NewSQL databases, which are fully relational but implement different mechanisms to achieve scale-out, have also appeared, offering solutions to the scale of certain databases. Artificial intelligence is beginning to have an impact on the database world and may become one of the main strategic areas for the future of database development.

Over the past decade, more and more industries have started to see how AI solutions can bring value to their products and services. AI-enhanced database management technologies can simplify the creation and maintenance of industrial databases, while intelligent data management systems can improve both the quality of the stored information and the automatic generation of business intelligence. AI can optimise manufacturing processes and detect fraudulent actions. As a visionary research topic, the integration of AI and database technologies is making great strides and will have a profound impact in the coming years across all industrial domains, including software development, healthcare, finance, retail, manufacturing and telecommunications.

References:

- [1] Gadde H. AI-Augmented Database Management Systems for Real-Time Data Analytics. *Revista de Inteligencia Artificial en Medicina*. 2024;15(1):616-49.
- [2] Ojika FU, Owobu O, Abieba OA, Esan OJ, Daraojimba AI, Ubamadu BC. A conceptual framework for AI-driven digital transformation: Leveraging NLP and machine learning for enhanced data flow in retail operations. *IRE Journals*. 2021 Mar;4(9).
- [3] Muppala, M. . (2025). *Digital Oceans: Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience*. Deep Science Publishing. 2025; doi:10.70593/978-93-7185-787-1
- [4] Konet SB. Artificial Intelligence-Powered Finance Algorithms, Analytics, and Automation for the Next Financial Revolution. *Deep Science*. 2025; doi:10.70593/978-93-7185-613-3
- [5] Panda SP. Artificial Intelligence Across Borders: Transforming Industries Through Intelligent Innovation. *Deep Science Publishing*. 2025; doi:10.70593/978-93-49910-25-6
- [6] Eboigbe EO, Farayola OA, Olatoye FO, Nnabugwu OC, Daraojimba C. Business intelligence transformation through AI and data analytics. *Engineering Science & Technology Journal*. 2023 Nov 29;4(5):285-307.
- [7] Mohapatra PS. Intelligent Assurance Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. *Deep Science Publishing*. 2025; doi:10.70593/978-93-7185-046-9
- [8] Ojika FU, Owobu WO, Abieba OA, Esan OJ, Ubamadu BC, Daraojimba AI. Transforming cloud computing education: Leveraging AI and data science for enhanced access and collaboration in academic environments. *Journal name and details missing*. 2023 Jan
- [9] Panda SP, Padhy A. Business Intelligence with Power BI and Tableau: Cloud-Based Data Warehousing, Predictive Analytics, and Artificial Intelligence-Driven Decision Support. *Deep Science Publishing*; 2025 Aug 15.
- [10] Konet SB. Microstructure, and Risk Optimization Models. *Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution*. 2025 Aug 12:17.
- [11] Mohapatra PS. Artificial Intelligence-Driven Test Case Generation in Software Development. *Intelligent Assurance: Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle*. 2025 Jul 27;4:38.
- [12] Konet SB. Artificial Intelligence in Financial Systems: Digital Transformation, and Machine Learning Applications. Available at SSRN 5401202. 2025 Aug 12.
- [13] Muppala M. SQL Database Mastery: Relational Architectures, Optimization Techniques, and Cloud-Based Applications. *Deep Science Publishing*; 2025 Jul 27.
- [14] Muppala M. Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience. *Digital Oceans: Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience* | Deep Science Publishing. 2025 Jul 8.

- [15] Panda SP, Koneti SB, Muppala M. Benefits of Site Reliability Engineering (SRE) in Modern Technology Environments. Available at SSRN 5285768. 2025 May 1.

Chapter 2: Architectural Strategies for Managing Databases in AI Environments

1. Introduction to AI and Database Architecture

Artificial intelligence (AI) workloads differ markedly from those of traditional business intelligence, primarily because they incorporate the use of non-structured data like images, audio, and text. Furthermore, it demands real-time access to features that are feeding models for inference, typically via a real-time feature store, in order to be able to respond to requests with minimal latency. Databases have historically been considered non-natural hosts for machine learning (ML) models, which was often solved by building architecture that extracts data from the database and loads it on the serving infrastructure. However, certain techniques can be leveraged such that ML models can reside natively within the database itself.




Fig1. Architectural Strategies for Managing Databases in AI Environments

The continuous integration of automation into everyday tasks hinges on the real-time extraction of knowledge from data, a capability provided by ML models. These models, encapsulated in a production-ready state, interact with external systems to deliver predictions. The external systems, in turn, generate feature data based on live and historical events. While the data retrieved is non-real-time, it must remain fresh and comply with stringent service-level agreements. As AI workloads become widespread and address diverse use-cases, organizations across sectors such as automotive, financial services, manufacturing, retail, and technology are revisiting database architecture and design principles.

2. Architecting Databases for AI Workloads

AI (artificial intelligence) workloads pose distinctive challenges; these involve training ML models on massive volumes of data or deploying previously trained models in live applications. The insights generated are then presented alongside traditional data, in dashboards or standard reports. Databases have long addressed the basic requirements of AI workloads [1-2]. Can they do more? Perhaps a broader ecosystem is required, exploiting the efficiency and features of databases where feasible, but going beyond when necessary.

The broad architecture of an AI system that exploits databases differs greatly from the more familiar data warehouse or summary reporting environment. Three core capabilities are required. First, a “real-time feature store” must make suitable inputs to a model available for querying while respecting constraints on data freshness, low latency and consistency. Second, the streaming architecture must efficiently forward events to one or more live models, observing the varying tolerances of different use cases to dropped events. Third, ML models must be served to live applications under strict performance requirements. Specific techniques for serving models from the database kernel promise environmental benefits by consolidating infrastructure and a compelling performance advantage for large-scale use.

2.1. Data Lakehouses

Artificial Intelligence (AI) workloads consist of tasks that require algorithms to identify patterns and generate forecasts from data, as well as to modify system behavior based on past experience. Such workloads are typically categorized at a high level as being either predictive or mostly automated. Predictive AI focuses on anticipating outcomes, while mostly automated AI encompasses routine decisions made by autonomous systems. The complexity of AI workloads has profound implications on the engineering and database infrastructure needed to support them. Most notably, native AI is requiring enterprises to rethink their data management systems and architectures. More generally, it is becoming a best practice to position Artificial Intelligence as an independent workload within a company, supported by a specialized branch of the information technology (IT) organization that is distinct from business intelligence (BI) and data engineering.

The advent of AI workloads has in turn created a set of new requirements for database management systems (DBMSs). For one thing, AI workloads tend to require the handling of much larger volumes and more diverse modalities of data than most business intelligence and operational analytics workloads, many of which are in unstructured format or simply textual format. This diverse corpus of unstructured data has to be stored and indexed inside the database for efficient retrieval of relevant context during inference. The issues of scale latency and performance also apply to operationalized model serving, real-time feature stores, and streaming architectures [2]. Data management challenges arising from legal, regulatory, and business governance frameworks also require attention. A prominent example is the demand for explainability and fairness in model predictions, which underscores the importance of model monitoring and validation in production. Security and privacy for AI data is a growing concern for all organizations, especially during inferences involving sensitive personal information. Broadly, optimizing DBMS performance, ensuring availability, scalability, and security for AI-related data, models, and inference responses concurrently and at scale are critical architectural considerations within modern DBMSs and data platforms.

Due to mixed workloads, enterprises rarely adopt a single database approach. Instead, of the different AI use-case categories, data lakehouses form a foundational capability because organizations have the lowest degree of control over currently-existing data sources. Hence, “many data scientists focus on building predictive AI applications on top of native AI capabilities. These predictive AI applications have fairly lax latency requirements in terms of inference response times, ranging from subsecond to a few minutes.”

2.2. Multimodal Storage

AI workloads require managing data of diverse types and modalities. Images, text, video, audio, time series, and other unstructured sources constitute a growing share of the volume of data ingested and processed by organizations, as well as of the features associated with use cases supported by AI. In fact, a substantial number of AI use cases rely on principles of multimodality. Consider, for instance, Alexa or Siri, which can synthesize speech and images. Interpretations at the foundation of vision-language models are deeply multimodal and combine images and text data. Consequently, multimodal storage capabilities represent another essential requirement for an AI database architecture.

Because every use case relies on its own data types, it is impractical to craft custom pipelines to convert a single feature store to support new target modalities when addressing a new scenario. Instead, the ability to natively support all the relevant modalities within a single feature store emerges as a more natural API to present to AI systems. The different components of data within a single AI workload may have very different requirements. A database suitably architected for these applications must be able to satisfy the specific needs of each data category. For that reason, multimodal storage and processing are typically associated with specialized submodules or subengines optimized for each task. For example, unstructured features—such as images, time

series, or text—are mainly consumed by AI models and therefore exhibit much lower query rates with very different access patterns and query types from the main relational table of the system, which is optimized for fast, random reads and writes, focused mainly on serving enrichment features for model queries.

2.3. Unstructured Data Handling

Artificial Intelligence workloads present new challenges for data management. Large-scale pretrained models require query, analysis, and scanning of vast amounts of training data to understand their behavior and bias. Online inference with these models requires retrieval of a fresh set of features for each request. A second wave of AI adoption in enterprise applications involves synthetic data generation and incorporating AI-generated content into transactions and communications. The size and rarity of such AI-generated content inhibit traditional storing and loading that large language models (LLMs) utilize. Consequently, database management must adapt to serve AI applications at scale.

Two core architectural aspects enable data management for these setting requirements. Firstly, data lakehouses combine data warehouses with data lakes to support unstructured AI training data while retaining enterprise capabilities such as governance and ACID consistency [2-4]. Secondly, multimodal storage, exemplified by real-time feature-store and streaming architectures, supports the time-sensitive nature of AI inference prediction. These approaches influence strategies for serving machine learning (ML) models within database systems, where tightly integrating a model's architecture with the database engine is crucial for efficient serving and is orthogonal to the challenge of accessing features during prediction.

3. Serving Machine Learning Models from Database Systems

Artificial intelligence (AI) is transforming almost every aspect of life, and with that transformation arise new requirements for many of the underlying systems that support the technology. An increasing number of AI workloads are hitting database management systems, but database management systems are not optimized for these kinds of AI workloads. The storage of data of different types, including unstructured data, and the serving of machine learning models for inference are two particular challenges.

Machine learning models are typically created with an ML framework such as TensorFlow or PyTorch, but inference issues occur because the creation processes within the frameworks are disconnected from the real-time processes needed for production. Feeding models at production into a database system and serving them from that system might solve both the incompatibility and performance issues, but doing so requires new techniques.

3.1. Integration of ML Models within Databases

Architectural Strategies for Managing Databases in AI Environments

Artificial intelligence (AI) workloads present a profound challenge in designing the systems and architectures of enterprises and software companies that provide AI services. The data involved in AI workloads is diverse, often stored in various database management systems (DBMSs) in use within the company. Data organizations must retrieve, integrate, clean, and prepare data from all these different systems for executing model inferences and training, with structures capable of responding to requests in real time; these architectures are often referred to a real-time feature stores. Graphical information often accompanies data tables to support the AI inference or training. Managing this information demands a streaming architecture, which can capture and serve sequences of past events, enabling a machine learning model in production to access its historical inputs and predictions.

Model inferences produce representations or embeddings that encode semantic information about the concept the embedding captures (a joke, a human face, the meaning of a sentence, the characteristics of an image, etc.). These embeddings represent the different entities involved in a problem (the topics, the profile of a new product, or the preferences of a customer) [5-6]. The role of databases, and how these are architected, lies not only in storing and managing the large amounts of information used to train the models but also in giving support to the operations involved in the organization of inputs and outputs, both during the execution of inferences and of model training.

3.2. Performance Optimization Techniques

Machine learning (ML) is considered a key technology for the new era of Artificial Intelligence (AI). Databases have been central to business for many decades, providing capabilities such as storage, access, management, protection, and security of company data. ML models form the core of many AI applications, and these applications evolve with changes to the models themselves.

ML models and their implementations are frequently deployed outside of data management systems. However, serving ML models inside databases can yield significant benefits, including efficiency gains, incremental model updating, execution of pipelines inside the database, and seamless integration with existing pipelines. These advantages help reduce inference latency and enable real-time loading of information into models. A major challenge in performing ML inference inside databases is the computationally expensive nature of matrix operations required by neural networks. Existing work has demonstrated that hardware acceleration through GPUs and FPGAs noticeably improves prediction speed. Nevertheless, to effectively support serving ML models inside databases, further optimizations are essential.

4. Real-time Feature Stores and Streaming Architectures

Multi-model databases can play a pivotal role not only during model training but also at serving time, when the ML model is used to perform inference. The straightforward approach is to separate training and serving. When it's time to perform inference, the model is integrated into

your service's source code, making queries to your main PSQL database to fetch all required features. While this approach scales horizontally with the rest of the application, it requires special code for logging features and model inferences, and burdens the main database with ML-related traffic.

Real-time feature stores address this challenge by decoupling serving from feature creation and model training. Features are generated and materialized within the feature store, a dedicated, distributed, read-optimized storage layer that supports very low latency and up-to-date feature retrieval, ensuring consistency among all features required by the model. To close the real-time feature data pipeline, streaming architectures are used.

4.1. Designing Real-time Feature Stores

Feature stores are data structures that store the features that are consumed by machine-learning (ML) models. Both AI and traditional ML models can benefit from feature stores. In batch-model building, feature stores provide saved historical features. Many models support transactions and require data windows; feature stores provide the data needed to support these windows. In model life cycles, feature stores oversee both the development of new models and the deployment of online (real-time) models — including fraud detection in banking, product recommendation in e-commerce, and loan approval. The data that supports these models changes as individuals perform different actions on the bank or e-commerce platforms.

Although feature-store design applications cover a broad range of feature-store details, one key component is a real-time serving capability: supporting low latency and consistently retrieved features in live ML models. Supporting real-time features is of minimal value unless a model can immediately utilise these features in its inference process.

4.2. Implementing Streaming Architectures

Streaming architectures enable continuous data flow and processing, crucial for real-time AI applications such as recommendation engines and anomaly detection. Different subsystems involved in streaming architectures include data ingestion, data stream processing, and serving ML models from streaming [7,8]. Data ingestion continuously captures events in the machine learning workflow and the application. Data stream processing transforms these events and calculations over historical data into useful real-time signals. Finally, a subsystem is responsible for serving these ML models from streaming—making models available for consumption as part of a stream.

Data ingestion systems continuously capture signals generated in the application and ML workflow—such as user profiles, real-time events like clicks, login timestamps, or ML closest neighbours—at low latency. Historical sensors feed the batch systems but have high latency (e.g., 24 hours). Real-time sensors can reduce latency to a few minutes or even seconds. As the number and variety of signals grow, organizations investing in scalable real-time capabilities

often adopt a standard data ingestion pipeline. However, this introduces significant complexity to the streaming architecture.

5. Data Management Techniques for AI

Data management for AI methods covers database auditability, compliance, data quality, and data-validation/generation methods. Auditability and compliance can be tackled with provenance data management methods and cryptographic data verification schemes, both areas that have an established literature with techniques now integrated into commercial-grade relational database engines. Data quality for multimodal AI training data can be addressed with integration of data profiling and data cleaning with ML theory so as to derive, for example, the best selection of training datasets that minimizes downstream model error-index. Similarly, methods from adversarial machine learning can be used to perform semantic data validation, for example, by generating the sets of training data that make a given ML model achieve optimal performance. Techniques in this area also look at the trade-off of data-quantity and data quality in noisy training datasets in order to achieve optimized investment on data-generation and data-curation for enhancing model quality.

Architecture principles to optimize machine-learning model serving inside the database engine have been discussed in the previous section. Other complementary techniques in this space involve (1) methods of model serving that achieve low latency while guaranteeing model accuracy, (2) approaches that use data structures analogous to indexes to serve high-dimensional data such as neural network embeddings, (3) real-time feature stores that provide the feature vectors used during prediction by taking the value of the relevant attributes at the query time, (4) streaming architectures that continuously feed events of real life into prediction framework for near-real time responses, and (5) approaching model serving as an event streaming problem for handling large influx of prediction requests. Database architectures in AI also present scalability-related concerns, such as comparing the use case of horizontal versus vertical scaling as well as investigating optimal load balancing strategies. Lastly, given the sensitivity of the data involved in AI workloads, security concerns require methods and systems for encryption, access control, and auditing.

5.1. Data Governance and Compliance

The quality and quantity of data used to train machine learning models determine model performance. The training data must be minimally curated and rich in information, with a test set accurately representing the real-world data distribution. Many organizations face obstacles such as compliance and regulatory requirements, which are particularly challenging in scalable training-data preparation. Training data often contains sensitive personal details that require governance to comply with regulatory and corporate rules [9-12]. Large organizations sometimes house training data in general-purpose data management systems. These industries project rapid growth in training-data preparation, necessitating the addressing of the outlined challenges.

Despite significant individual progress in data governance, data compliance, data-quality checking, and training-data validation, there is a lack of synergy to realize scalable, parallel training-data-preparation pipelines. Optimal utilization of underlying big data frameworks is also missing, resulting in subpar scalability and efficiency.

5.2. Data Quality and Validation

Data quality and validation comprise an active area of research and development within AI data management. The myriad of techniques employed for quality management, primarily developed for OLTP (Online Transaction Processing) systems, confront challenges when applied to AI data. These challenges are especially evident with unstructured data, where the absence of a defined schema hampers the enforcement of data quality assurances through traditional methods. The sector of AI data management dedicated to these issues is poised for an increasingly prominent role.

The issue of quality management has grown even more difficult due to the growing interest in self-driving or autonomous systems. Within such environments, an AI model—such as an autonomous car—constantly collects new data as it interacts with the physical world. The data generated must then be used to retrain and update different models. As in a continuous integration environment for software, the data produced must be checked to verify whether it contains any anomalies, errors, or gaps that could potentially degrade the quality of the models trained for the system.

6. Scalability Considerations in AI Database Architectures

The rapidly growing volume of data is the primary cause for the evolution of the AI database architecture. The ability to handle increasing amounts of data, both in terms of storage and computational needs, is crucial for AI applications since large-scale datasets are essential for training accurate models [7,13-15]. Databases designed for AI workloads must scale effectively, ensuring consistent performance even as data volumes surge. Scalability is often categorized into two dimensions: horizontal scaling, which involves adding more machines to a system through sharding or replication to distribute load and increase capacity, and vertical scaling, which focuses on enhancing a single machine's resources to manage more substantial workloads. Both require efficient load balancing strategies to optimize resource utilization and minimize response times.

In the future, databases tailored for sensitive AI applications will likely employ enhanced security features, incorporating detailed encryption and access control mechanisms for stored data. Such measures are indispensable for maintaining data privacy and protecting sensitive information. Moreover, business or mission-critical AI workloads often include stringent compliance requirements related to data governance, making comprehensive data management an essential aspect for AI databases.

6.1. Horizontal vs. Vertical Scaling

Modern databases for AI need to build an efficient execution framework and flexible data model to process all the different data types without losing performance. Storage systems supporting AI should store all the different data types together and should be integrated closely to the processing engine. In this way, it is possible to quickly combine large amounts of meta-data and unstructured data. AI database management systems should also provide real-time responses, helping to get insights from specific events.

Traditional centralized data systems do not satisfy all these requirements, especially regarding scalability. Scaling up typically requires expensive and complex hardware upgrades, such as faster CPUs, more RAM, or higher I/O disks, which can quickly deplete available resources. Additionally, the processing capacity might be bounded by a sequential workflow, causing the system to become a bottleneck that restricts overall performance. A distributed architecture removes these constraints, allowing for seamless scalability by adding more machines with off-the-shelf hardware. Furthermore, data can be partitioned across different nodes and processed locally, combining the results to deliver a global answer efficiently.

6.2. Load Balancing Strategies

AI workloads are often distributed across fixed clusters, and balanced load distribution is crucial to maximize resource utilization and maintain high throughput. Model serving engines adopt various strategies to assign workloads among stored model replicas. One straightforward method is the round robin approach, assigning requests to replicas sequentially in a loop. For stateful streaming architectures, more sophisticated techniques ensure that all related events of a given stream are directed to the same process, preserving state consistency. Recent work optimizes model serving by dynamically adjusting the number of replicas at runtime based on the real-time workload.

Real-time feature stores require low-latency feature retrieval to serve live model inference. Their architecture resembles that of serving engines, necessitating mechanisms to distribute querying load evenly across servers to achieve low latency and efficient resource use. A simple approach employs round robin distribution [16]. Alternatively, assigning requests for the same entity key to a specific replica enables that replica to cache features for the entity, boosting cache hit ratio. These strategies parallel those applied in model serving engines.

7. Security and Privacy in AI Database Management

Security and privacy are paramount in managing databases for AI, particularly through encryption and role-based access control. Underlying data for AI workloads can be highly sensitive, such as industrial sensor data in predictive maintenance scenarios or data sources prone to attacks in autonomous vehicle systems. Sensitive information also emerges from stored

embeddings used for similarity searches, semantics, and data retrieval. Data breaches have driven the adoption of encryption methods applied to both data at rest and data in transit.

Data at rest can benefit from file-system encryption that protects static data on disk, but practical considerations can lead to IT staff disabling this feature. Consequently, a more robust defense is offered by database layer encryption. Data in transit requires an additional layer of encryption to implement TLS between distributed systems and end users, securing data as it moves across different nodes in the system. Implementing these measures alongside role-based access control mechanisms ensures that unencrypted data, whether in storage or in motion, remains accessible only to authorized personnel who require it for their work.

7.1. Data Encryption Techniques

Artificial Intelligence (AI) has been around for many decades but interest in the field has soared recently, with new applications like ChatGPT using foundation models to enable systems that can perform a variety of tasks, such as answering questions, summarizing texts, creating new content and more, and doing so with high quality and low latency. This high demand for interactive AI systems requires a rethinking of the underlying database architecture. Modern AI systems impose new requirements on database management systems, such as support for data lakehouses that combine the efficiency of data warehouses with the low-cost storage and flexibility of data lakes, the ability to support multimodal storage of data in the same database, support for unstructured data like images and text, and the ability to serve ML models inside the database for low-latency, high-throughput processing.

In addition, a wide variety of techniques are necessary to support AI systems. Real-time feature stores enable the retrieval of ML features in a consistent, low-latency, highly available manner, and streaming architectures enable the continuous flow of events. Privacy, security, compliance, data management and governance techniques help produce trustworthy AI systems. Finally, the scalability of the system ensures that low latency and high throughput are maintained, and encryption techniques protect data both at rest and in flight.

7.2. Access Control Mechanisms

The management of databases used in artificial intelligence (AI) workloads benefits from the general security measures used in conventional environments. In addition, AI data repositories require privacy-preserving modeling techniques such as differential privacy. The latter protects privacy during model training, ensuring that no subset of data samples—either as individuals or as a group—can exert an outsized influence on the output of a model. AI databases that contain information like names, addresses, and credit limits often have specific governance and compliance requirements that must be enforced. The databases themselves must provide access control policies, governance, auditing, and the ability to iterate on data (e.g., remove users) for compliance reasons. It is crucial to build in fine-grained access control on data that undergoes

heavy processing before final storage. Public data can be stored in cheaper tiers, whereas sensitive data can be stored in expensive but encrypted storage.

AI databases often contain highly sensitive data that must be protected, either owing to the nature of the data or the nature of the model training process. In traditional sensitive-data scenarios such as financial and medical applications, encrypting the data is both an organizational and regulatory requirement. One of the most expensive phases of machine learning is the training phase, which needs to be repeated each time the underlying data changes by even a single record. If such encrypted data is outsourced to the cloud for performing training, then the question becomes: How much does the database knowledge that it is training on encrypted data provide an advisory role in strengthening model training privacy while simultaneously being cognizant of the computational overhead? The standard input/output primitives of a database system support privacy preservation during inference as well. For example, during the inference phase, the database system can query a trained model over an encrypted dataset [9,16-18]. The output of the model can also be encrypted so that it provides inference without leakage of information.

8. Case Studies in AI Database Architectures

Several real-world case studies illustrate how database management tools have been scaled up to meet the demands of artificial intelligence. Data lakehouses, combining data lakes and data warehouses architectures, provide unified storage for historical, operational, and machine-learning data. As AI databases must support multiple data types—'multimodal' data including structured tables, time series, images, videos, hypertext, speech, and audio—such companies as UnifyID store the data in a single database instead of housing different data types in separate repositories. The demand for speedy live inference (serving predictions in response to queries in real time) has led to the creation of dynamically updated feature tables containing the inputs for machine-learning models, together with streaming architectures that feed a flow of live data into the models.

Architectural aspects also influence the integration of machine-learning models within the database engine by considering embeddings, transformers, bloom-filters, multiplexers, and tokenizers. An architecture that performs horizontal scaling—distributing the models across multiple nodes in a cluster—enables the serving of more model inferences in parallel. Vertical scaling entails the selective allocation of model components to specialized hardware accelerators such as GPUs. Furthermore, model inference serving can be accelerated by the deployment of load balancers to distribute incoming HTTP inference requests efficiently, and by the implementation of caching mechanisms to store and quickly retrieve results of frequently accessed model inferences.

8.1. Industry Applications of Data Lakehouses

The increased use of artificial intelligence in the industry is not new, but the team responsible for designing database management systems that support AI workloads faces several new

challenges. Storage engines commonly used in a data lake type architecture to store unstructured data from text, images, and video are combined with the low latency, high throughput, and efficient update and delete operations of the data warehouse storage engine in the so-called data lakehouse architecture. As a result, the database management system operation is up to an order of magnitude more expensive because it was not designed for this scenario. For example, a feature store that supports serving features for an ML model inferences at online prediction time must scale horizontally to support thousands of requests per second and process events in real-time for very low latency and consistent retrieval of features. Another example is the need to perform real-time event streaming in a streaming architecture.

8.2. Real-world Examples of ML Model Serving

Databricks Lakehouse Platform embodies the data lakehouse architecture by merging the scalability of data lakes with the management and tuning capabilities of data warehouses to provide a single source of truth for AI-driven analytics and decision-making. The platform employs Apache Spark™ for distributed processing of large datasets and Delta Lake for reliable streaming and batch data pipelines, delivering consistent high-quality data at scale on cloud object storage. Its low-latency optical caching capability allows users to cache frequently accessed data, significantly boosting performance. Databricks also supports a wide range of machine learning models, from classification and regression to natural-language processing and computer vision, enhancing security and governance across the platform's extensive analytics ecosystem.

EaseML offers a declarative abstraction system for rapid development and deployment of ML services. Built on top of PostgreSQL, it allows users to easily implement ML services within the popular open-source database, utilizing shared-disk elasticity for deploying models on fewer machines and eliminating redundancy. Unlike typical AI databases, EaseML provides a declarative interface that decouples ML services from the underlying implementation of storage and model serving. Developers can quickly deploy state-of-the-art models with ease, and graduate to more specialized tasks without abandoning the simplicity of the declarative model. EaseML also serves as a reusable foundation for implementing ML-serving infrastructures within a database[2,19-20].

9. Future Trends in Database Management for AI

Artificial Intelligence (AI) workloads differ significantly from traditional database workloads in the direction, structure, and velocity of data movement. Consequently, databases supporting AI applications require architectural features that depart from the organization of transactional and traditional analytical processing. These architectural considerations have given rise to the concepts of the data lakehouse and the serving of machine-learning models from the database. Real-time feature stores further extend databases to support low-latency retrieval of training and live inference data for ML models.

The data lakehouse architectural model combines the best features of data lakes and data warehouses in a single database and is particularly suitable for AI workloads. Data lakes provide highly scalable, low-cost multimodal storage that supports structured, semi-structured, and unstructured data in their native formats. However, the schema-on-read approach of a data lake often complicates data governance, performance, and compliance. Data warehouses, by contrast, improve governance and performance through the imposition of a structured, relational schema on highly curated data, but they only support a limited number of data modalities. The tightly bound transport, schema, and storage of the data warehouse also inhibit the performance achievable for much higher data velocities demanded by AI that operate in all directions — training, validation, and live inference.

9.1. Emerging Technologies and Innovations

Artificial Intelligence is giving rise to various emerging technologies that require data management and back-end infrastructures at scale. Artificial Intelligence Interactive Applications like ChatGPT demonstrate the complex and partially unresolved challenges encountered when managing databases. While feature stores and real-time data streams support lightning-fast AI inference for interactive applications, scaling the training set for the entire underlying AI model(s) poses immense challenges in a rapidly evolving environment.

Several emerging areas illustrate the challenges of managing AI databases. First, database architectures must scale effectively—horizontally, vertically, or both across clusters—while also incorporating load balancing strategies to optimize operational performance. Second, governance, compliance, and data quality processes are essential to ensure appropriate data use at scale. Third, the management of sensitive data requires privacy mechanisms, such as encryption and access control models. New cutting-edge architectures are evolving to address these challenges.

9.2. Predictions for AI Database Architectures

The architecture of future data management systems will be driven by Artificial Intelligence. What is special about AI workloads? While training large neural networks requires tensors encompassing both model parameters and their corresponding gradients, AI is about more than just tensors. As AI continues to advance, unstructured data in the form of images, audio, video, and text gain significance. Therefore, AI databases need to be multimodal universities that support multiple types of data and enable machines to learn. Since AI models operate on data but cannot generate predictions in a vacuum, AI databases must support real-time feature stores—repositories where we can look up key-based, high-dimensional feature vectors for live inference requests. Additionally, real-time event streaming architectures are essential to channel streaming events into the models. Furthermore, AI models operate on data. Given the dependence of AI model predictions on input data, enhancement models are necessary to improve AI data quality and support auditing and compliance related to AI models and their predictions. The diverse requirements of the data translate to a multitude of design alternatives.

Currently, two architectures dominate: lakehouse architectures aimed solely at training, and serving models from database systems that support both training and serving. The future likely holds a convergence of these architectures to deliver the best of both worlds.

10. Conclusion

Data management plays an important role in artificial intelligence (AI) workloads. AI generates large amounts of data, such as the objects that a self-driving car passes on the street or a patient's color image database for cancer diagnosis. This data can be either structured or unstructured. Structured data can be organized into rows and columns within a database, while unstructured data can take forms such as text, color images, and videos. Due to the large size of unstructured data, storing it in SQL databases can lead to very large tables. Another important part of AI data is the model and its weights, which are critical components of any AI application. New techniques aim to store machine-learning models, their weights, and inferencing functions within database systems.

AI data is quickly becoming one of the largest data domains. Databases are moving the goalpost by supporting new AI workloads. Among the different types of AI databases, data lakehouses offer the ability to manage data of different modalities and provide data management services such as governance, compliance, security, privacy, data validation, and data quality—all essential components for handling AI data. Models can be stored in a multimodal database along with features from different data modes and the data lakehouse that stores the ground-truth data. Real-time machine-learning inferencing requires features to be retrieved with exceptionally low latency, while online machine-learning training demands consistency. Both needs can be addressed by implementing feature stores either as part of or closely integrated with the URI. Moreover, streaming architectures are necessary to deliver upstream events to the online machine-learning training pipeline.

References:

- [1] Koneti SB. Artificial intelligence Applications in Retail and Investment Banking: Personalization, Robo-Advisory and Behavioral Analytics. *Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution*. 2025;4:72.
- [2] Muppala M. Architectures in relational databases: An analytical study of SQL-based data models and ACID principles. *database*.;2:4.
- [3] Bentahar J. A Survey on Explainable Artificial Intelligence for Network Cybersecurity. *arXiv (Cornell University)*. 2023 Mar 7.
- [4] Gadde H. AI-Assisted Decision-Making in Database Normalization and Optimization. *International Journal of Machine Learning Research in Cybersecurity and Artificial Intelligence*. 2020;11(1):230-59.
- [5] Koneti SB. Algorithmic Trading and Quantitative Finance Strategies: High-Frequency Trading, Market Microstructure, and Risk Optimization Models. *Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution*. 2025;4:17.
- [6] Panda SP. The Evolution and Defense Against Social Engineering and Phishing Attacks. *International Journal of Science and Research (IJSR)*. 2025 Jan 1.

- [7] Mohapatra PS. Artificial Intelligence and Machine Learning for Test Engineers: Concepts in Software Quality Assurance. Intelligent Assurance: Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. 2025 Jul 27:17.
- [8] Konet SB. Analysis, Predictive Analytics, and Macroeconomic. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025 Aug 12:90.
- [9] Panda S. Scalable Artificial Intelligence Systems: Cloud-Native, Edge-AI, MLOps, and Governance for Real-World Deployment. Deep Science Publishing; 2025 Jul 28.
- [10] Shivadekar S, Halem M, Yeah Y, Vibhute S. Edge AI cosmos blockchain distributed network for precise ablh detection. Multimedia tools and applications. 2024 Aug;83(27):69083-109
- [11] Gadde H. AI-Enhanced Adaptive Resource Allocation in Cloud-Native Databases. Revista de Inteligencia Artificial en Medicina. 2022 Oct 18;13(1):443-70.
- [12] Konet SB. Microstructure, and Risk Optimization Models. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025 Aug 12:17.
- [13] Panda SP. Augmented and Virtual Reality in Intelligent Systems. Available at SSRN. 2021 Apr 16. Mohapatra PS. Artificial Intelligence-Driven Test Case Generation in Software Development. Intelligent Assurance: Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. 2025 Jul 27:38.
- [14] Eboigbe EO, Farayola OA, Olatoye FO, Nnabugwu OC, Daraojimba C. Business intelligence transformation through AI and data analytics. Engineering Science & Technology Journal. 2023 Nov 29;4(5):285-307.
- [15] Mohapatra PS. Intelligent Assurance Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. Deep Science Publishing. 2025; doi:10.70593/978-93-7185-046-9
- [16] Ojika FU, Owobu WO, Abieba OA, Esan OJ, Ubamadu BC, Daraojimba AI. Transforming cloud computing education: Leveraging AI and data science for enhanced access and collaboration in academic environments. Journal name and details missing. 2023 Jan
- [17] Panda SP, Padhy A. Business Intelligence with Power BI and Tableau: Cloud-Based Data Warehousing, Predictive Analytics, and Artificial Intelligence-Driven Decision Support. Deep Science Publishing; 2025 Aug 15.
- [18] Konet SB. Microstructure, and Risk Optimization Models. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025 Aug 12:17.
- [19] Mohapatra PS. Artificial Intelligence-Driven Test Case Generation in Software Development. Intelligent Assurance: Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. 2025 Jul 27:4:38.
- [20] Konet SB. Artificial Intelligence in Financial Systems: Digital Transformation, and Machine Learning Applications. Available at SSRN 5401202. 2025 Aug 12.

Chapter 3: Exploring the Impact of AI on Query Optimization and Database Performance Tuning

1. Introduction to AI in Database Management

Artificial intelligence (AI), defined as computer algorithms that undertake tasks normally requiring human intelligence, is an enabler for self-tuning databases and improvements in query optimization. By drawing on specialized machine learning techniques such as reinforcement learning and deep learning, database management systems can incorporate AI to optimize critical-performance-determining features. With query plan selection having such a significant impact on performance, the prospect of being able to dynamically choose plans within a single query provides exciting possibilities for query optimization.




Fig1. Exploring the Impact of AI on Query Optimization and Database Performance Tuning

Databases have always needed to be tuned, and in response, engineers have created tools that assist with or automate the tuning process. Oracle, Azure SQL Hyperscale, and Snowflake offer

feature sets that are self-tuning in varying degrees. Self-tuning involves eliminating, removing, or at least reducing the amount of manual intervention required to tune the database, by incorporating AI techniques. Tuning operations that are particularly painful or difficult from a human perspective are usually the primary candidates for AI-driven self-tuning.

2. Fundamentals of Query Optimization

To optimize a query, a database system must decide on an efficient plan for executing the query. The database does so by first deriving an algebraic representation of the query and then creating a tree representation for it. The system transforms the tree and assigns data access methods to create different query plans [1]. Cost models estimate the execution costs of these different plans, determining an apparently optimal query plan. In some implementations, these models are driven by machine learning, which predicts query execution costs. Its focus is on reinforcement learning, which can be used to create dynamic query plans.

In the earliest phases of the database revolution, development teams created specialized engineering tools. As technologies matured, the team could build self-scaling and self-healing mechanisms, incorporating more automated algorithms. Businesses soon demanded self-tuning technologies, which could address performance bottlenecks—such as long-duration queries—without human intervention. Self-tuning is now recognized as an indispensable component of a database environment as data drain increases. Certain modern databases, including Oracle, Azure SQL Hyperscale, and Snowflake, collect query execution data and proactively tune the system to improve performance for repetitive queries, thereby minimizing customer pain points.

3. AI-Powered Query Optimization and Tuning

Common query optimization strategies adopted in database management system implementations can be supplemented with reinforcement learning to enable dynamic query planning decisions. Because cost models constitute the core of query optimization, machine learning methods can be employed to enhance the accuracy of predicted query execution costs. However, there are also ways in which databases can be made largely self-tuning through the application of AI techniques.

Several major database vendors offer cloud-based databases that harness artificial intelligence for this purpose. Oracle, for example, provides a self-tuning feature for its autonomous database. Microsoft Azure offers a self-tuning capability for the hyperscale tier on Azure SQL Database, while Snowflake has incorporated features that render its service largely self-tuning. Each of these implementations leverages aspects of AI to automate and optimize performance tuning tasks.

3.1. Reinforcement Learning in Query Planning

Advanced technologies employ AI techniques, such as reinforcement learning, to generate optimized query plans dynamically. Strategies using reinforcement learning demonstrate the resulting potential of such an approach.

Reinforcement learning provides a foundational approach to self-tuning databases that can be extended successfully by including cost models based on machine learning. These advanced cost models are also implemented using machine learning techniques.

3.2. Cost Models Driven by Machine Learning

Cost models can be transparently enhanced by leveraging Machine Learning techniques, as they excel in uncovering and leveraging hidden correlations within datasets. Models driven by Machine Learning are capable of adapting to a myriad of objectives, including query execution duration, resource consumption, monetary expenditures, and the number of input/output requests. They adeptly accommodate diverse implementations of identical data operations. Unlike traditional cost models anchored in database statistics, estimations generated through Machine Learning are resilient to the pitfalls of cardinality estimation errors, given that stock cost parameters inherently absorb errors across preceding processing stages. Incorporating the cost of the reward function itself remains a non-trivial endeavor, as certain cost entities might not be pertinent to all queries; for instance, the CPU cost of an index scan is zero.

Machine Learning techniques excel at discovering hidden correlations within diverse datasets. Models erected on the foundations of Machine Learning can adapt to various targets, such as the duration of query execution, consumption of resources, monetary costs, or the tally of input/output requests [1-2]. They gracefully accommodate implementations of the same data operation in its different achievable variants and are less vulnerable to the effects of cardinality estimation mistakes. This robustness stems from the fact that, contrary to cost models relying on conventional database statistics, cost model parameters are frequently integrated within the Machine Learning data annotations. The estimation procedure for the reward function remains challenging, as not all cost categories are relevant to every query; for example, the CPU cost associated with an index scan does not apply.

3.3. Self-Tuning Databases

Databases incorporate self-tuning features to enhance query performance and adapt to evolving workloads. Operators can configure these mechanisms to allocate more resources to heavy workloads or query groups, reducing execution time. In query planning, plans can be dynamically adjusted using feedback during execution or runtime information. Cost models may rely on machine learning predictions of query performance, improving over heuristics. Operational tuning leverages activity and timing metrics to rebalance workloads and manage resources, with alternatives including the addition of virtual nodes for scalability.

Recent trends also show major database vendors incorporating artificial intelligence (AI) into their products. Self-tuning through machine learning lowers administration costs and enhances user-friendliness by reducing the need for manual parameter settings and role assignments. It is anticipated that AI automates all tuning operations in the future, leading to fully autonomous databases. Such developments are pixelating the traditional roles of DBA and race analyst, as AI in the cloud handles the bulk of operational decisions. Making operations smart is nearly synonymous with incorporating AI at some level.

4. Case Study: Oracle's AI Capabilities

Artificial intelligence (AI) plays an increasingly important role in database management and optimization. Database vendors integrate AI automation into self-tuning features, aiming to remove tedious and error-prone manual interventions in query optimization. Oracle, for example, refers to its Autonomous Database features as the "Future of Databases."

The following examples demonstrate the benefits of Oracle Autonomous Database for the user, highlighting how AI autobot businesses are revolutionizing database management. As AI-based optimization algorithms mature, more advanced approaches appear beyond traditional optimization heuristics and query planning strategies. Reinforcement learning, for instance, can optimize live query plans. It adapts the query planning strategy dynamically to limit execution times and enhance user experience. Other approaches employ cost models driven by machine learning to estimate query costs better. Clear advantages over traditional cost modes enable model-driven optimizers that estimate costs based on richer representations than individual plan operator costs.

5. Case Study: Azure SQL Hyperscale

AI-enabled tools for database performance optimization—specifically quasi self-tuning databases—have now been introduced at all of the major database vendors. The following case study on Azure SQL Hyperscale highlights how Microsoft incorporates AI to tune performance.

The Hyperscale service tier for Azure SQL Database enables rapid scaling to hundreds of terabytes for single databases in the cloud. The data is stored in page servers and transaction logs are stored and managed separately by the log service. Another component, the Log Replay Service, is responsible for applying transaction log records to page servers. It plays a key role in database creation, scaling, backup, and restoring [3-5]. Log Replay must be highly performant to meet scaled-out logging requirements. Recently, artificial intelligence techniques have been applied to optimize Log Replay. By automating tuning of configuration parameters and improving the commit rate, AI integration has markedly enhanced update performance of Log Replay.

In SQL Server 2019, a new, lightweight, and scalable architecture for batch-mode query processing was introduced. Operating on a columnar batch of rows, this new engine now lets

more data-processing tasks benefit from the performance of batch mode. AI Red Optimizer technology extends the benefits of the batch-mode processing engine to rowstore data—without requiring any indexes or materialized query tables to be created. By predicting the efficiency of various batch-mode query plans with greater accuracy, and guiding SQL Server to choose better batch-mode plans, AI Red Optimizer enables better plan selection for both rowstore and columnstore data formats. An experiment on TPC-H Q9 using AI Red Optimizer yields a 4× improvement in overall query runtime, and a TPC-DS Q67 workload with joins between narrow rowstore, wide rowstore, and columnstore tables runs 3.5× faster with AI Red Optimizer enabled.

6. Case Study: Snowflake's AI Features

Modern database systems are incorporating increasing degrees of self-tuning query optimization. Snowflake's cloud data warehouse has recently launched a range of features leveraging machine learning to recognize operational patterns on their platform and optimize customer workloads accordingly. Their Query Acceleration Service dynamically determines the number of resources utilized to process a query, impacting both response time and run-time cost. Machine learning is employed to identify and deliver the optimal provisioning for a given query that offers the best trade-off for the user, taking into account user preferences. The service supports various query types, such as SDR pre-provisioned queues and ad hoc queries.

Snowflake also uses artificial intelligence in their automatic clustering service. Each table is examined to find the single best key to re-cluster on, allocating the compute cluster in accordance with the recommended system cost for such re-clustering. The metadata is analyzed to monitor and detect performance degradation caused by existing clustering keys, while workload information is incorporated to assign benefit scores to partition-level activity within current keys. By examining the distribution of data access by clustering key, it ensures that only required partitions are reclustered, thereby minimizing additional costs.

7. Comparative Analysis of Database Systems

Artificial intelligence enables databases to automate or partially automate complex operations across the entire data lifecycle, encompassing collection, storage, analysis, and protection. Additionally, AI optimizes cost-effectiveness by dynamically allocating resources in support of an organization's digital transformation. Prominent cloud vendors such as Oracle, Microsoft, and Snowflake are currently advancing these capabilities, offering clients products that facilitate the processing of increasingly complex and demanding workloads with minimal supervision and reduced tuning overhead.

AI-powered query optimization emerges as a fundamental element in the development of cost-effective and scalable database systems. The application of reinforcement learning to dynamic query planning enhances scalability and adaptability, while the incorporation of machine

learning augments the optimizer's cost model, resulting in more accurate query cost estimations. Fully automated tuning remains one of the most captivating aspects of query optimization; for instance, Oracle provides an autonomous cloud that leverages AI to administer workload management, resource allocation, and tuning decisions. Similarly, Azure SQL Database Hyperscale exploits AI in its autoscaling mechanisms to support highly scalable and flexible workloads. Snowflake integrates AI-driven elastic scaling to address the demands of compute-intensive workloads, enabling clients to perform analyses with greater efficiency and speed.

8. Challenges in AI-Driven Query Optimization

AI techniques, particularly those making database management systems self-tuning, hold great promise. They can transform operations: as less human intervention is required, fewer errors occur, flexibility increases, workloads scale more easily, and data-driven decision making and pattern recognition become quicker and more reliable. Nevertheless, numerous technical difficulties remain as the technology matures.

The size, complexity and sophistication of modern database management systems call for machine learning that can scale and adapt. Most engine components already adjust themselves automatically, but these systems mimic adjustment by means of thresholds applied to meticulously crafted metrics [6-8]. As workloads constantly evolve, these thresholds must be recalibrated—preferably in an automated way, with as little human input as possible. Adapting to new and changing configurations requires learned models to consider operations that lie beyond their current experience—such as adding, removing or relocating indexes, changing memory parameters or updating the concurrency control algorithm.

9. Future Trends in AI and Database Performance

Machine learning (ML) techniques are becoming increasingly popular within the database community. For example, reinforcement learning has recently been applied to query planning, allowing for the dynamic selection of optimal query plans as the plan executes. Moreover, ML can be used to build enhancements that leverage the latest cloud technology. Given the central importance of cost models in traditional query optimizers, ML models trained to provide accurate query cost estimates also hold great promise for adapting query processing to evolving hardware architectures.

One obvious candidate for ML techniques is self-tuning databases, which have been a research focus since the 1990s. Database vendors are beginning to incorporate these concepts into commercial platforms. For instance, Oracle has branded part of its Autonomous Database as self-driving. Self-scaling and self-tuning features are central to Azure SQL Hyperscale and Snowflake. Specifically, two key aspects of self-tuning—the ability to automatically scale virtual machines (VMs) or compute nodes and the use of AI techniques to optimize query performance—are expected to become ubiquitous in cloud databases. Self Scaling has been

addressed previously. Self Optimization, the application of AI and ML techniques to improve query optimization, is outlined below.

Recent research demonstrates that query optimization can be improved by caching the execution of sub-plans within the same query and using this information to produce better query plans for subsequent sub-plans. In addition, a dynamic query optimization approach based on reinforcement learning adapts the execution strategy based on the current resource availability of the distributed system and the progress of the query execution, significantly reducing query latency and optimizing resource consumption across multiple nodes.

10. Ethical Considerations in AI Implementation

Ethics becomes paramount when AI query optimization functions are deployed in production, as the stakes are high. Not only can they vastly influence the monthly operating costs of a company, but wrongful implementations can cause catastrophic outages or even seriously damage a company's reputation. Consequently, companies internally demand the highest level of confidence in those systems, partly because the opaque nature of AI makes its predictions less explainable than those of rule-based heuristics. The costs associated with manual validation, however, often complicate matters.

Moreover, problems of fairness and bias appear in similar polymorphic forms. It is not reassuring, for example, if a subgroup of users receives a less-efficient query plan because of socioeconomic or racial data inferred from the query or users' historical query patterns. Despite the challenge—since certain kinds of optimizations are specifically targeted in boosting performance in high-frequency queries—such considerations must be taken into account in future developments. These are just a few of the many ethical concerns that arise and intensify as the research field moves in the direction of AI-powered database management [9].

11. Performance Metrics for AI-Driven Systems

Appropriate metrics guide AI tools and quickly assess gains. Performance metrics include 11.1 Average Query Runtime (seconds) and 11.2 Average Cost Metric as an Abstract Value.

Query runtime is the prime performance indicator, the one for which database optimization exists. What else is? AI and machine learning typically evaluate and compare performance improvements in terms of cost functions. While the cost function is a value in the optimization model, it does not always translate directly into a scalable, measurable unit like seconds or milliseconds. The cost metric is an abstract, dimensionless value, designed to serve as a proxy for cost or run time. Query planning aims to minimize the cost metric.

12. User Experience and AI in Databases

Experiences differ between setups and scale. For data warehouses, a simple query to read a million rows from a basic table should run flawlessly in any vendor's ecosystem. Yet, anomalies often arise that confuse even seasoned developers. These disruptions stem from which optimization features are toggled, affecting plan search strategies, index selection, auto-parallelization, distributed query plans, and many other aspects. The issue's root is that automatic parameter setting is necessary but not sufficient. Autonomous systems must accurately understand the workload. Achieving 100% reliable SQL plan stability remains improbable—and might never be attainable—due to the unpredictability of future queries.

Data warehouses represent the low-hanging fruit for AI in automation. Yet, similar self-tuning capabilities are emerging in OLTP database engines. Oracle Autonomous, for instance, incorporates self-tuning features, while the Azure Hyperscale variant introduces a horizontally scaled SQL Server engine. Snowflake leverages AI to optimize data handling within CSV and JSON semi-structured files.

13. Integration of AI with Traditional Optimization Techniques

Artificial Intelligence and machine learning have received increasing attention in database query optimization and tuning. A recent line of research exploits diverse AI techniques, including reinforcement learning, deep learning, and learned cost models, to complement and overcome the limitations of traditional query optimizers. Following this trend, the focus here is on their application to self-tuning mechanisms.

Protecting Database Performance with AI—Why It Matters While SQL performance is determined by several factors—such as thin client speed, network condition, latency, query execution plan, and database resources—the consistency of query performance especially depends upon the database part. As the volume of data grows exponentially and the workload varies dynamically, the database execution plans keep starving for new information in order to come up with an effective query plan. This makes AI models well-suited to query optimization. Further three cases examine how Artificial Intelligence integrates with traditional optimization techniques in Oracle Database, Azure SQL Hyperscale, and Snowflake.

14. Impact of Cloud Computing on Database Performance

Cloud providers use many techniques to improve the performance of cloud database services. One example is the hyperscaler database architectures used by Microsoft and Oracle. These are different from a traditional database.

The hyperscale architecture is primarily designed to support very large database sizes (petabytes of data). The unique feature of the hyperscale architecture is the separation of the compute layer from the storage layer. The compute layer is a database engine that supports classic database

operations (queries, data manipulation/statements) and a remote cloud storage layer (e.g. Azure Storage or Oracle Autonomous Data), which is a fully managed, tabular, auto-scaling, highly reliable, and performant distributed storage service for structured relational data in the cloud.

15. AI for Predictive Analytics in Databases

Just as AI has been incorporated into a variety of different database management functions, query optimization is another database operation that is highly influenced by AI technology. Traditionally, a query optimizer calculated the best access path to the raw data for a query being processed. Query plan optimization techniques such as cost-based optimization attempt to find the query plan with the lowest cost to return the results from a query. AI such as reinforcement learning has been used to dynamically generate a query plan as the query fetches its results. AI can also be used to create a cost model with machine learning that is used by the optimizer to calculate the cost of a query plan and choose the plan with the best cost. Self-tuning technologies can also be embedded in a database. Oracle or Azure SQL Hyperscale SQL Server are two examples [7,9-10]. Wisdom gained from tuning procedures used by DBAs or the engine can be introduced as heuristics to be applied when working on a potential slowdown, scaling or concurrency problem. Snowflake can also be added to these as some of the aspects of its public API address predictive analytics at scale.

In the most demanding cases for DBMSs, some queries start clogging the system. AI is ready to use its knowledge to prevent this scenario coming from resources saturations such as CPU, RAM or IOPS. One DBA's first reaction is to run either a tuning stored procedure or a dynamic management view (DMV) SQL Server query. These queries will usually bring the most expensive items within the system. AI goes further by creating a baseline of historic performance metrics on the main resources, linked also to factors behind the user workload. If it is out of norm, AI can react and use the information that it had collected earlier to address the on-going demand. The AI procedure can also warn about an excessive number of concurrent users for a T-SQL or stored procedure compiled object. Scalability issues can be tackled proactively during a peak time for OLTP or OLAP work. For recent cloud-connected databases such as Snowflake, a SQL API can be triggered to provide additional scalable resources, followed by a subsequent API call to remove it once the activity is back to normal.

16. Real-World Applications of AI in Database Management

Artificial intelligence (AI) enhances modern database management by automating performance tuning and query optimization, areas traditionally addressed through manual configuration and heuristics. Emerging techniques in reinforcement learning, machine learning, and deep learning enable the creation of AI-augmented tools that optimize query plans dynamically and generate more effective physical design strategies.

AI-powered tuning features have become a core component of many cloud database services and hyperscale architectures, providing automated workload-aware scalability and self-tuning capabilities. For example, Oracle's self-driving database leverages machine learning to recommend optimal configurations and analyze query performance across operational workloads. Azure SQL Hyperscale introduces an architecture capable of auto-scaling and auto-tuning to maximize performance. Snowflake employs AI-driven optimization to adjust its internal structures continually, enhancing query execution speed automatically.

17. The Role of Data Quality in AI Optimization

Credit costs for a given query during the process of trade-off interaction. Expedia partnered with Snowflake to develop an ML-based cost model for its SQL queries within Snowflake. In the domain of self-tuning, automated or semi-automated tuning of database parameters and structures is a well-known and desired concept in database management systems. Oracle, Microsoft, and Snowflake incorporate such technologies to enhance query optimization and database performance.

Trading-off last-level-cache (LLC) misses against prefetching opportunities can substantially influence query performance, which can be addressed by dynamically choosing the best tuning operation. Controlling the large scale of databases in the cloud is challenging, yet Azure Hyperscale enables dynamic scaling of compute and storage independently [1,11-14]. While cgroups can assign a profile to each database container, finding the best profile is non-trivial, especially when workload patterns change over time. Snowflake developed a controller that dynamically scales the resources of a virtual warehouse to optimize performance during periods of peak workload intensity. The achieved return on investment ranges from 20 to 40%, depending on the current workload. Credit costs for queries change during these trade-off adjustments.

18. Security Implications of AI in Databases

AI-powered SQL generation also implies the AI can perform queries that might compromise security. Any machine learning system still fails if trained with bad data. Even when trained properly, training data can always be hacked through poisoning and backdoors. Thus if care is not taken, AI can output hard-to-explain queries that will execute on the database. To overcome this limitation, any practitioner would therefore narrow down the amount of access the AI can do to just a small part of the database [13,15-17].

In general, the use of AI in databases introduces new security concerns, such as data privacy risks from automated indexing and schema management, the risk of data leaks via generated queries, and vulnerabilities to data poisoning and model backdoors. Addressing these challenges is both necessary and difficult, requiring a balanced approach between the benefits of AI-driven optimization and the imperative of maintaining robust security.

19. Regulatory Compliance in AI-Enhanced Systems

Regulatory compliance represents a crucial consideration in the implementation of services like a Self-tuning Azure SQL Database that employs AI and machine learning techniques for auto-scaling. Such organizations must comply with relevant laws and regulations regarding software application development and delivery, the collection, processing, storage, and transfer of personal data, consumer protection, and contractual and tort liability. The level of compliance responsibility will depend on the nature and location of the organization's operations, the data managed, and how the service is designed and delivered. It is the organization's responsibility to ensure that the use of Azure SQL Database complies with all relevant laws and regulations. Additionally, national security regulations might also apply.

The organization must adopt any measures necessary to prevent the unauthorized transfer of services or data outside regions or countries where their residence or operation may trigger additional compliance requirements. Microsoft is positioned to provide the Artificial Intelligence capability needed to support the Self-tuning Azure SQL Database, yet it is currently the organization's responsibility to comply with all relevant laws and regulations concerning the use of Azure SQL Database. These considerations extend beyond compliance alone, encompassing the perceived impact that outsourcing core components of an organization's business operations to third-party providers may have on external reputation and investor confidence[18-20].

20. User Training and AI Systems

A particularly novel application of reinforcement learning to query optimization is the resulting prospect of a self-tuning database. Major database vendors have been applying machine learning (ML) methods to automation in self-tuning. Google demonstrated the value of ML in a DBMS Autotuner. Oracle has applied ML methods in a self-tuning version of its database. Azure SQL Hyperscale supports a self-tuning feature that scales compute nodes up and down based on actual workload demands. Snowflake added support for AI-powered tuning.

Self-tuning involves applying AI in the broader context of query optimization and query planning. During query planning, the DBMS breaks complex queries into simpler steps and structures the plan to execute steps in an order tailored to the size and clustering of input query results. Reinforcement learning may guide the database system to choose an execution strategy at each step.

21. Collaboration Between IT and Data Science Teams

The rapid growth of unstructured data, such as images, documents, and sensor recordings, is transforming data-driven applications in modern trade and commerce. This has led to an increasing focus on AI technologies that allow databases to understand and process unstructured information. Despite the recognized benefits of integrating AI with databases, practitioners face several challenges in implementing and adopting these capabilities within real-world databases.

Given the complexity and diversity of database systems across various domains, a one-size-fits-all approach is impractical. Instead, targeted community efforts should consider the needs of different groups of users and database engines [19,21-22].

Effective collaboration between IT professionals and data scientists is critical to realizing AI's full potential in query tuning and database performance optimization. To build trust in AI-powered solutions, vendors have developed systems that are not only powerful but also easy to use. For example, Oracle Database's self-tuning technologies leverage AI at various levels, enabling automatic tuning and workload management in single-node and multi-node settings, as well as seamless workload scale-up and scale-down in the Autonomous Data Warehouse. Similarly, AI-driven features in Azure SQL Hyperscale and Snowflake are significantly changing the way organizations optimize query performance [23,25].

22. Cost-Benefit Analysis of AI Implementation

The challenges faced by human experts attempting to manually optimize the query plan can be highlighted by considering the growth of the search space for an SQL query as described by Surajit Chaudhuri and Gerhard Weikum. Queries can be joined in different orders, creating a complex space of hundreds of millions of logically equivalent plans for just a few tables. The problem arises when the cost model is inaccurate for the specific hardware and workload being run. This inaccuracy causes the query optimizer to pick a local optimal plan when a globally optimal plan exists but cannot be detected due to the cost model, thus causing performance degradation. Recently, the reinforcement learning technique has been used to optimize the query planning stage itself, finding the optimal sequence of join operations without using a cost model [26,27].

Machine learning has made a significant impact on the database community. Traditionally, machine learning has improved the cost model of query optimization using classification, regression, and learning to rank, which still use, at their core, a cost model but train it with real query data. With the increased interest shown in database tuning, vendors have been developing self-tuning databases that enable automated database tuning. Although database management systems such as Microsoft SQL Server, Google Cloud Spanner, and SAP HANA offer automated tuning capabilities, Oracle, Azure SQL Hyperscale, and Snowflake have pushed the concept of an autonomous, self-tuning database furthest.

23. Conclusion

AI plays an important role in modern database management systems. Query optimization aims to change a SQL statement into an efficient execution plan that accesses the data and computes the result more quickly. In a self-tuning database, machine learning is used in various tuning areas, including indexing, resource provisioning, and concurrency control. Reinforcement learning can

help improve query performance by selecting different query plan operations. A cost model powered by machine learning offers more accurate query optimization.

Oracle databases incorporate artificial intelligence to automate query optimization and performance tuning. Azure SQL Hyperscale leverages machine learning to automatically optimize resources and services for extreme scalability in the cloud. Snowflake also applies AI techniques for query optimization and performance tuning.

References:

- [1] Reis J, Housley M. Fundamentals of data engineering. " O'Reilly Media, Inc."; 2022 Jun 22.
- [2] Muppala, M. . (2025). Digital Oceans: Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience. Deep Science Publishing. 2025; doi:10.70593/978-93-7185-787-1
- [3] Konet SB. Artificial Intelligence in Financial Systems: Digital Transformation, and Machine Learning Applications. Available at SSRN 5401202. 2025 Aug 12.
- [4] Panda SP, Konet SB, Muppala M. Benefits of Site Reliability Engineering (SRE) in Modern Technology Environments. Available at SSRN 5285768. 2025 May 1
- [5] Gadde H. AI-Augmented Database Management Systems for Real-Time Data Analytics. *Revista de Inteligencia Artificial en Medicina*. 2024;15(1):616-49.
- [6] Eboigbe EO, Farayola OA, Olatoye FO, Nnabugwu OC, Daraojimba C. Business intelligence transformation through AI and data analytics. *Engineering Science & Technology Journal*. 2023 Nov 29;4(5):285-307
- [7] Muppala M. Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience. *Digital Oceans: Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience*| Deep Science Publishing. 2025 Jul 8.
- [8] Ojika FU, Owobu O, Abieba OA, Esan OJ, Daraojimba AI, Ubamadu BC. A conceptual framework for AI-driven digital transformation: Leveraging NLP and machine learning for enhanced data flow in retail operations. *IRE Journals*. 2021 Mar;4(9).
- [9] Konet SB. Artificial Intelligence-Powered Finance Algorithms, Analytics, and Automation for the Next Financial Revolution. Deep Science. 2025; doi:10.70593/978-93-7185-613-3
- [10] Panda SP. Artificial Intelligence Across Borders: Transforming Industries Through Intelligent Innovation. Deep Science Publishing. 2025; doi:10.70593/978-93-49910-25-6
- [11] Mohapatra PS. Intelligent Assurance Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. Deep Science Publishing. 2025; doi:10.70593/978-93-7185-046-9
- [12] Ojika FU, Owobu WO, Abieba OA, Esan OJ, Ubamadu BC, Daraojimba AI. Transforming cloud computing education: Leveraging AI and data science for enhanced access and collaboration in academic environments. *Journal name and details missing*. 2023 Jan
- [13] Panda SP, Padhy A. Business Intelligence with Power BI and Tableau: Cloud-Based Data Warehousing, Predictive Analytics, and Artificial Intelligence-Driven Decision Support. Deep Science Publishing; 2025 Aug 15.
- [14] Konet SB. Microstructure, and Risk Optimization Models. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025 Aug 12:17.
- [15] Mohapatra PS. Artificial Intelligence-Driven Test Case Generation in Software Development. *Intelligent Assurance: Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle*. 2025 Jul 27;4:38.
- [16] Muppala M. SQL Database Mastery: Relational Architectures, Optimization Techniques, and Cloud-Based Applications. Deep Science Publishing; 2025 Jul 27.
- [17] Enemosah A. Enhancing DevOps efficiency through AI-driven predictive models for continuous integration and deployment pipelines. *International Journal of Research Publication and Reviews*. 2025 Jan;6(1):871-87.
- [18] Ivanov SH, Webster C. Adoption of robots, artificial intelligence and service automation by travel, tourism and hospitality companies—a cost-benefit analysis. *Artificial intelligence and service automation by travel, tourism and hospitality companies—a cost-benefit analysis*. 2017.

- [19]Ramadhan M, Naseeb A. The cost benefit analysis of implementing photovoltaic solar system in the state of Kuwait. *Renewable energy*. 2011 Apr 1;36(4):1272-6.
- [20]Cordes JJ. Using cost-benefit analysis and social return on investment to evaluate the impact of social enterprise: Promises, implementation, and limitations. *Evaluation and program planning*. 2017 Oct 1;64:98-104.
- [21]Dykes PC, Curtin-Bowen M, Lipsitz S, Franz C, Adelman J, Adkison L, Bogaisky M, Carroll D, Carter E, Herlihy L, Lindros ME. Cost of inpatient falls and cost-benefit analysis of implementation of an evidence-based fall prevention program. *InJAMA Health Forum* 2023 Jan 6 (Vol. 4, No. 1, pp. e225125-e225125). American Medical Association.
- [22]Gadde H. AI-Enhanced Adaptive Resource Allocation in Cloud-Native Databases. *Revista de Inteligencia Artificial en Medicina*. 2022 Oct 18;13(1):443-70.
- [23]Koneti SB. Microstructure, and Risk Optimization Models. *Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution*. 2025 Aug 12:17.
- [24]Panda SP. Augmented and Virtual Reality in Intelligent Systems. Available at SSRN. 2021 Apr 16. Mohapatra PS. Artificial Intelligence-Driven Test Case Generation in Software Development. *Intelligent Assurance: Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle*. 2025 Jul 27:38.
- [25]Eboigbe EO, Farayola OA, Olatoye FO, Nnabugwu OC, Daraojimba C. Business intelligence transformation through AI and data analytics. *Engineering Science & Technology Journal*. 2023 Nov 29;4(5):285-307.
- [26]Mohapatra PS. Intelligent Assurance Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. *Deep Science Publishing*. 2025; doi:10.70593/978-93-7185-046-9
- [27]Ojika FU, Owobu WO, Abieba OA, Esan OJ, Ubamadu BC, Daraojimba AI. Transforming cloud computing education: Leveraging AI and data science for enhanced access and collaboration in academic environments. *Journal name and details missing*. 2023 Jan

Chapter 4: Embedding Intelligence into Data Pipelines: Exploring the Intersection of MLOps and DataOps for Enhanced Automation and Quality Assessment

1. Introduction to MLOps and DataOps

MLOps—the intersection of Machine Learning, DevOps, and Data Engineering—introduces continuous creation and operationalization of machine learning workflows. MLOps enables rapid experimentation and model leveraging for business enhancement. DataOps improves automation, monitoring, and quality of data. Embedding intelligence into data pipelines allows automation of many steps in ETL/ELT workflows beyond what is achievable with traditional DevOps. Artificial Intelligence methods can enhance data quality assessment and anomaly detection in data pipelines.

ETL (Extract, Transform, Load) operations are among the most manually intensive business processes. Organizations strive to operationalize and automate ETL steps with continuous monitoring and alerting for breaks or delays, but monitoring data quality is usually not implemented. As data volumes grow powerfully, they deliver greater value in rapid analyses but also generate actions requiring control for possible risks. Ensuring high quality is critical for making business decisions. The intersections between MLOps and DataOps enable this intelligent layer within data pipelines, providing mechanisms to connect data quality assessment and anomaly detection with monitoring, alerting, and auditing systems, supporting feedback loops and continuous development.

2. The Convergence of MLOps and DataOps

The DevOps revolution has paved the way for modern transformations in the fields of machine learning (Machine Learning Operations, MLOps) and data (Data Operations, DataOps). MLOps is a set of cultural and technical practices that enables the deployment and maintenance of machine learning models in production reliably and efficiently. Its aim is to increase automation and improve the quality of production models while focusing on business and regulatory

requirements. DataOps is a collaborative data management practice focused on improving the communication, integration, and automation of data flows between data managers and data consumers across an organization. These practices address against the pitfalls of building and maintaining a data pipeline.

When MLOps meets DataOps, the potentials for a fully automated, adaptable, and scalable integrated data pipeline solution arise. Today, business needs require combining data-driven decision services, ML predictive services, and forecasting services into a single product. The convergence stimulates the automation of ETL/ELT pipelines through AI techniques and models, thus giving birth to a new research and technology domain. These issues are discussed in section 3, where the focus is mainly on the automation of ETL/ELT pipelines using AI techniques, and the associated products that can be developed and delivered [1-3].

2.1. Historical Context

Both DataOps and MLOps appeared almost simultaneously and succeeded DevOps principles, long established. DataOps appeared in 2014 as the new practice of Agile analytics, aiming to reduce the cycle time of data analytics. MLOps appeared in 2015 as a variant of DevOps necessary to deploy and maintain machine learning applications in production. The continuous integration and continuous delivery/deployment (CI/CD) in the MLOps approach are extended to the integration of machine learning models. Both disciplines focus on bringing teams working with data closer by shortening feedback cycles, but they emphasize different dimensions. DataOps managers are responsible for the complete data lifecycle in an enterprise, including data preparation and governance functions, as well as data delivery. MLOps managers focus primarily on the phases around training machine learning models and deploying them in production, including business metrics and key performance indicators evaluation.

2.2. Key Principles of MLOps

The key principles of MLOps are cross-functional collaboration and automation. Cross-functional collaboration ensures that teams across different departments work together, use common tools, and share a unified purpose. Automation leverages recurring processes and patterns to establish a trustworthy, repeatable, and disaster-resilient deployment pipeline. Such pipelines not only deploy trained models to production but also handle big data traffic flows serving the data related to these models [2].

Given the scale of data processed by operations, DataOps practices are crucial for assuring data quality. Machine learning and deep learning techniques, which are themselves data-driven, depend heavily on data quality. Recent research explores how DataOps and MLOps can complement each other in providing automated quality assessment for data transformations in extract, transform, and load operations (ETL). MLOps tools and mechanisms can learn from DataOps anomalies and error patterns, fueling automation in data pipelines. Although Electrical,

Electronics and Communication Engineering was not originally centered on Machine Learning and Artificial Intelligence, advances in these fields have naturally steered them in that direction.

2.3. Key Principles of DataOps

DataOps emerged from experiences at companies such as Netflix, Facebook, and Spotify, evolving over time and drawing inspiration from the principles of Agile, Scrum, and DevOps. Agile introduced Developer Operations to continuous software development and deployment, Scrum provided a framework for continuous team collaboration and sprint structuring, and DevOps integrated development and operations teams to rapidly deploy executable software code to production. These concepts directly influenced the birth of DataOps, which extends the same DevOps principles to manage and maintain stable and scalable data pipelines.

Scrum, Agile, and DevOps were integrated through automation and collaboration, bringing together cultural and organizational philosophies that supported their principles. DataOps shares these foundational elements, guided by a set of automation and collaboration practices as reflected in "The DataOps Manifesto." The resulting continuous cycle advances collaboration between users, developers, and operations personnel, cutting across disparate workflows from various companies and increasing the automation rate of data pipeline phases.

2.4. Benefits of Integration

About a decade ago, the ground broke and a growing community of machine learning (ML) engineers came together to form a new discipline called MLOps, or machine learning operations [2,4,5]. The premise is simple: embedding machine learning intelligence into the data pipeline. It works with data operations and management because it focuses on the data used by machine learning, especially for predictions and ranking. Industry analysts proposed the complementary term DataOps, emphasizing the need for a culture change that brings data-first thinking with automation and processes like continuous integration/continuous delivery (CI/CD) for data pipelines.

Embedding intelligence into the data pipeline enables a new level of automation allowing organizations to work faster, develop more pipelines and models and ultimately derive more value from their data. This new level of automation optimizes operations by blending AI with the operational aspect of data management. It helps businesses avoid the risks and compliance issues caused by problems in the underlying data structures by analyzing metadata and engine logs from all components. A meta machine learning system effectively enables Artificial ETL (extract, transform and load): an AI system that automatically performs the ETL/ELT embedding intelligence into the data pipeline.

3. Automating ETL/ELT with AI

Modern MLOps pipelines increasingly feature ELT (Extract, Load, Transform) operations to organize data for AI model training, testing, and validation. Similarly, DataOps pipelines rely on

ETL (Extract, Transform, Load) steps that prepare data before analytics and visualization. Especially in DataOps, these workflows process data related to the business, the product, pricing, or the resultant revenue. Most of the time, organizations develop these processes based on business decisions and data-management staff placed in two separate silos. At least in principle, AI can help handle some of this organizational complexity but also provide guidance about the potential impact of decisions related to these processes.

Recent trends in machine-learning operations deliver new business values based on managing business metadata, business metrics, and data quality/observability. These advances open a completely new perspective for DataOps approaches [6-8]. The General Data Protection Regulation (GDPR)-related requirements for data transparency, and the general data-protection and data-governance objectives of enterprises, are additional drivers for accelerating the automation of metadata management and data quality assessment. Another trend in automating data-preparation pipelines is using ML techniques to detect data anomalies.

3.1. Overview of ETL/ELT Processes

MLOps meets DataOps: Embedding intelligence into the data pipeline Automation in the extract, transform and load (ETL) or extract, load and transform (ELT) process plays an important role in order to provide timely, consistent and accurate data for business intelligence. MLOps meets DataOps: Embedding intelligence into the data pipeline Automation in the extract, transform and load (ETL) or extract, load and transform (ELT) process plays an important role in order to provide timely, consistent and accurate data for business intelligence. Extract, transform and load (ETL) describes the process of extracting the data from the source systems, applying transformations to the data and loading the transformed data into its destination table in the target system. Extract, load and transform (ELT), on the other hand, loads the extracted data into dedicated tables in the target system without applying any transformations. Subsequently, the transformations are applied on the data in the target system and the data is moved into its final destination table. This principle allows the database engine of the target system to perform the transformations more efficiently by using native commands and parallel processing.

Both ETL and ELT processes are usually constructed manually. For instance, in the case of the Zeppelin Movet project, team members are required to design the workflows in Apache Zeppelin and continue with a manual follow-up and maintenance. The ETL/ELT automation, however, also includes the automatic inferral of dependencies for workflows used to extract the data from the source, transform the data and load the data into the warehouse. For most businesses, data is a key asset and it is essential that data quality is assured and continuously monitored in order to provide consistent, accurate preferably up-to-date data for the business intelligence analysts, business users and company executives. Anomalies in the data can indicate potential problems and appropriate actions can be taken to reduce the likelihood of having wrong data in the data warehouse. Data quality functions have become a standard feature of demand management systems and most vendors offer a number of predefined data quality functions. Artificial

intelligence (AI) can play a decisive role in the assessment of the data quality of the ETL/ELT process.

3.2. Role of AI in Automation

Modern software development is undergoing a transition towards the Automation Age. Rapid growth in software, data generated by applications, and the development of artificial intelligence (AI) methods and tools present software engineers with abundant automation options. Indeed, new AI tools related to Software Engineering (SE) tasks appear rapidly, often disseminated through social media. Owing to the availability and ease of use of AI tools—particularly code generative large language models (LLMs) such as Copilot and ChatGPT—developers are beginning to adopt AI-generated code. These AI tools can be integrated into existing software development automation practices (e.g., DevOps, MLOps) to address other aspects of the SDLC [9,10]. For example, incorporating AI-generated code, automated analysis methods, and AI-augmented tools can accelerate and automate aspects of Data Engineering in DataOps pipelines. DataOps, an emerging discipline that combines Agile software development methods with continuous delivery aspects from DevOps, focuses on data pipelines—particularly the move, process, and transformation of data. Controlling the quality and consistency of data in these pipelines remains a significant challenge.

3.3. Tools and Technologies for Automation

A variety of existing tools integrate artificial intelligence with ELT and ETL processes. Knowledge Graphs and Knowledge Bases form an intelligent layer above ETL technologies by extracting, structuring, enriching, storing, reusing, and sharing enterprise semantics and related knowledge. Machine Learning Operations ecosystem addresses the automation of data pre-processing and quality assessment [11-13]. Quality assurance tools allow the user to define ETL quality rules and monitor data quality, while data anomaly detection solutions are based on Machine Learning. Finally, executing analysis models in an online environment involves orchestrating different components with an automatic scheduler.

Embedding intelligence into the data pipeline is an emerging area of artificial intelligence. It is leveraged by MLOps and DataOps integration. Applying artificial intelligence to automation breaks simple data ingestion processes into microservices within the larger mutation phase of a typical ELT workflow—Extract, Load, and Transform. Integrating MLOps and DataOps transforms the traditional Extract, Transform, and Load data pipeline, establishing governance, orchestration, monitoring, and quality controls across the full data lifecycle, from ingestion to reporting. The confluence of MLOps and DataOps facilitates embedding intelligent capabilities into ETL/ELT processes. In particular, artificial intelligence supports automation at three levels: automating ETL workflow construction, automating data quality assessment, and automating data anomaly detection.

3.4. Case Studies of AI-Driven ETL/ELT

The automation of data extraction, transformation, and loading processes can be achieved by harnessing different artificial intelligence methods. The following cases illustrate the effective integration of machine learning and AI methodologies in ETL/ELT pipelines.

There's one specific DataOps pipeline that is focused on data quality (DQ) verification using machine learning to analyze records as they enter the data lake. Low quality is detected, triggering a quick reaction from data engineers and reducing reprocessing costs [2,14-17]. Machine learning based anomaly detection and predictive data validation is used by big fintech companies to identify root causes of data quality issues there are more accurate and earlier in the ETL process. These are just a few examples of how the intelligence built in to DataOps workflows drive automation and improve data quality.

4. Data Quality Assessment

CheKing as an Analytical Reasoning Tool for Data Quality Assessment Decision making is important task in decision life cycle of any organization particularly when organizations want to have competitive edge through business intelligent. Preprocessing and manipulating the data is thus critical. Various methods are designed for automatically assessing data quality through artificial intelligence which deal with problems of practical applications to automate AI. A joint solution of both supervised and unsupervised learning is in the direction with good potential. The right algorithm or ensemble algorithm to use is important and also related with the business context. Signs, such as vessel size or ETA that may influence container release and payment process should be constantly monitored. According to the model of European Foundation for Quality Management (EFQM) with "radicalism", a customer-oriented, process driven, and integrated approach is intensified toward excellence and continuous improvement. Data quality is a vital part of the EFQM model and contributes to a business's reputation for achieving business excellence.

Machine learning offers robust forecasts of vessel arrivals, supports the identification of anomalies, and highlights where data quality should be prioritized. Machine learning techniques can identify and compute the anomalies detected within a dataset. Subsequently, decision trees can categorize these anomalies into clusters within the vicinity of the anomaly. Anomalies can be categorized as organisational, catastrophic, or suspected anomalies, and labeled with the associated cause in the dataset. This annotated dataset then serves as a foundation for further classification tasks. In these projects, assessment was performed within a DataOps culture, where DataOps tools and workflows are integrated into the actual work culture of the organization [9,18-21].

4.1. Importance of Data Quality

Modern data pipelines do more than collect, transform, and pre-select data—they also monitor and analyze incoming and output data to assess its overall trustworthiness and suitability for intended downstream tasks. Data quality evaluation remains critical, with data subjected to comprehensive validation before use in models, reports, or dashboards. Moreover, data quality assessment demands continuous temporal monitoring to detect and address technical glitches or natural business fluctuations and seasonality.

Today, organisations can leverage artificial intelligence to enhance the automation and quality of traditional extract, transform, and load (ETL), or more recently, extract, load, and transform (ELT) processes. Organizations have developed MLOps and DataOps methodologies to assist data practitioners in confronting cultural, organizational, and technical challenges associated with adopting and operationalizing these new AI tools across the enterprise.

4.2. Traditional vs. AI-Driven Approaches

The growth of data economies promotes the recognition that data is a product represented by its Quality, and Data Quality is recognized as a critical element. Data quality and associated control procedures have been developed, like Six Sigma and Analytical Quality Control, to manage the Quality of the Product.

Data quality assessment can be designed in various ways depending on the use case. Nevertheless, data quality gates or rules can be automatically built with AI methods that augment data engineers' skills. It is common to start by implementing quality checks for source data. Data quality gates can then be established for each step in the ETL/ELT process. These gates perform control checks to ensure that the data transform performed on a dataset is as expected and does not create unexpected data anomalies or biases.

4.3. Frameworks for Quality Assessment

Notable efforts have been made to approach data quality assessment systematically.^{1–5} Although the techniques differ, the majority fuse scored indicator metrics to provide an overall data quality assessment and visualization. Machine learning has been recent applied to assess data quality,⁶ but unsupervised anomaly detection has proven more effective — in essence, training an ML model on fresh data for discovering data anomalies or outliers should reflect most data-quality issues, rather than relying on a traditional trained model.

Data quality assessment is intimately related to anomaly detection. Implementation considerations arise, however, for organizations that have yet to establish a complete MLOps environment. Since DataOps and MLOps are mutually dependent, the integration challenge deserves particular attention. Therefore, the preceding methods maximize automation for ETL workflows, work in both supervised and unsupervised scenarios, and require only a minimal level of MLOps support.

5. Anomaly Detection with Machine Learning

Data anomalies represent deviations or inconsistencies in data that diverge from an expected or typical state, presenting challenges to proper data handling and business operations. Detecting these anomalies is crucial for preserving data quality and enhancing security. Machine Learning techniques provide robust means for identifying unusual samples within datasets, applicable across diverse domains such as credit card fraud, network intrusion, product defect identification, and disease diagnosis. Typical approaches include Clustering, Neural Networks, Support Vector Machines (SVMs), Statistical methods, and Ensemble Learning. Clustering assesses groupings and detects points distant from defined clusters as anomalies. Neural Networks, particularly Self-Organizing Maps, leverage clustering and visualization to identify atypical samples. SVMs employ hyperplanes to segregate normal from abnormal data, assigning labels accordingly. Statistical methods analyze attribute distributions to flag data points falling outside normal boundary values [22,23]. Ensemble Learning techniques integrate multiple models to compute anomaly scores, facilitating outlier detection.

Anomaly detection—considering temporal, spatial, relational, and other contexts—is a fundamental application in domains including Cyber-Physical Systems, Environment, Transportation, Finance, and Healthcare. Evaluation metrics such as the F1 score, Accuracy, Precision, and Recall quantify the performance of detection systems.

5.1. Understanding Anomalies in Data

MLOps and DataOps can be viewed as ways to environmentalize AI or, in other words, as ways to operationalize AI-based workflows that either actively make use of AI or represent a necessary prerequisite for successfully applying AI in other stages of the data flow. Once the underlying processes of the data flow have been environmentalized and automated, the integration of AI is a natural next step to embed intelligence into the different steps of the data flow. In the context of DataOps, this can be seen as a means to transition from merely managing the data flow to enabling a self-managing data flow. The overall benefits of this integration are clear: it leverages the potential of AI for existing tasks, improves the efficiency of data flows, and reduces the manual workload of DataOps teams.

Extract–Transform–Load (ETL) or Extract–Load–Transform (ELT) serve as crucial methods for automatising data transformation and preparation. Apart from the other dimensions, automatic ways of determining and evaluating quality of the transformed data are equally important. A potential application for artificial intelligence in this domain would be to assist a company's DataOps team when data quality is assessed, for example, by detecting anomalies within transformed data tables [24-26]. Anomalies can be in the forms of anomalies, collective data distribution and lost periods in short. By using machine learning methods, a data-driven anomaly detection method is developed that provides significantly higher capability versus traditionally defined rule-based approaches.

5.2. Machine Learning Techniques for Anomaly Detection

The discipline of Operations in Machine Learning (MLOps)—also called Automatic Machine Learning or Tools of Artificial Intelligence—consists of tools and techniques that introduce automation into building Machine Learning models.

Data pipelines consist of sequences of processes that extract and load information into other systems. Traditionally, these pipelines have served as the backbone of Modern Data Engineering, enabling the generation of reports and business intelligence. They have become the need of the hour for data-driven companies that drive growth and compete better in the market.

The rich and varied family of Artificial Intelligence techniques could also be exploited to automate other steps of the conventional Data pipeline (such as the extract-transform-load step, or operations related to Data Quality Assessment, such as these performed in DataOps operations).

The parallel presence of Machine Learning oriented pipelines and the classical data pipeline paves a way for an Operational Model for Data. This model would combine aspects of MLOps and DataOps, incorporating intelligence into data pipelines for improved automation and quality assurance.

5.3. Implementation Strategies

Implementing Data Quality Control in Production Once the phase of Data Quality Control Design is completed, the entire working regime, including monitoring, alerting and troubleshooting, must be implemented, preferably in an automated and integrated way. Partial or primitive forms of automation can be supported by popular workflow managers such as Apache Airflow, but full automation is frequently attained through the employment of MLOps tools, which enable the integration of ML workflows into broader data pipelines. Automation Code Once data quality KPIs and Anomaly Detection models become operational, the corresponding monitoring and alerting processes are systematically automated to achieve full automation of data quality control. [27,28] Given the mature DevOps infrastructure that supports ETL/ELT tooling, including CI/CD pipelines, containerization and orchestration, the deployment of new data quality workflows on prem or in the cloud can be accomplished with minimal effort and maximal reliability in the shortest possible timeframe. Monitoring Integration into various dashboards is also straightforward through Loggers, Cerberus, etc. Culture and Mindset The effective implementation of the concept of embedding intelligence into data pipelines necessitates concerted efforts to harmonize MLOps and DataOps both culturally and organizationally. A common understanding of the objectives must be established, along with a clear recognition that the creation and maintenance of AI elements in all processes is an everyday job that demands continuous attention and adjustment.

5.4. Real-World Applications

Implementations of machine-learning techniques for data-quality assessment and anomaly detection are steadily growing in number and increasingly apply to various real-world data sets. Companies such as Hume and Soda revolve these concepts around the more modern DataOps discipline, which is evolving into the MLOps framework successfully used within artificial intelligence (AI) teams. Many tools and integrations support this evolution and allow for the use of data-quality checkers built on AI methods, such as Deepchecks, Great Expectations, or Datafold. These data-quality checkers offer automated, intelligent support in writing expectations, analyzing data, and detecting anomalies, thus not only ensuring clean data but also assisting in model selection and the definition of model boundaries.

Within the AWS ecosystem, a service named Deequ automates data-quality checks using ML methods. Deequ processes raw data and computes metrics and constraint checks, which it then evaluates. The results determine whether a data row passes or fails the quality criteria, enabling the definition of guard rails. Deequ's automated support facilitates the generation of constraints by analyzing historical data and identifying anomalies. Combining the functionality of guard rails with the automatic generation of constraints inspired by Concerns-Less Data Engineering fosters cooperative evolution between MLOps and DataOps, thereby embedding intelligence in data pipelines.

6. Challenges and Solutions in MLOps and DataOps Integration

The fields of MLOps and DataOps are growing very fast but there are several cultural, organizational and technical challenges to overcome. Fairly distinct skills are required to tackle MLOps and DataOps, with DataOps engineers being typically skilled more on synthetic languages like Python or Scala and MLOps engineers focusing more on system or shells programming languages like C or GO. Both actions are performed by completely separate teams with very little crossover in high level strategic technical decisions [19,29-31].

Where we most need to shift is the operationalization of machine learning. There are some big barriers there, such as lack of architecture and methodologies integrated with operations systems and parts. From there on out it is all about data AND not enough people that can combine those skills, roles, responsibilities etc. (Selected quotes: 1) In reality (and common sense), without data you cannot feed any intelligence at all. DevOps automation system is more slowly being adopted because of lacking the common MLOps architecture and guidance and lack of cross skilled manpower. Organizations move at different paces in synchronizing DevOps traditional stream with new MLOps stream and police forcefulness towards the latter can hinder progress. To overcome the challenges and realize the benefits requires an approach that embeds intelligence into existing data pipelines, i.e., an approach where MLOps meets DataOps.

6.1. Cultural and Organizational Barriers

As with any new methodology, there are several challenges when applying MLOps and DataOps conceptually. First, DataOps emphasizes an iterative and continuous approach to the design and development of data analytics, which requires supervision motivation and business support, all of which may be lacking in traditional siloed organizations. Second, DataOps calls for the automation of a multitude of steps and activities associated with data analytics, not only the retrieval and storage of data but also profiling, cleansing, integration, and transformation. The lack of end-to-end automation and data pipelines often interferes with the ability to address changing market demands.

These problems have been solved in the cloud business intelligence (BI) area, offering easy access from anywhere and a low level of knowledge to exploit their capabilities. However, it still remains a considerable challenge to construct DataOps that can automate the entire process of web presentation, periodical use, and scripts scheduling along with data visualization.

6.2. Technical Challenges

MLOps and DataOps together ensure the maintainability, replicability, reliability, security, and quality control of data pipelines with embedded intelligence. Technical challenges in integration arise from the nature of these two disciplines. DataOps encompasses technical practice, particularly Julia Language web development, data transformation pipelines, and Apache Kafka data streaming. MLOps focuses more on machine-learning topics such as hyperparameter searching, architecture evaluation, and model training/rendering. Designing, implementing, and maintaining extensive pipelines that rely on MLOps and DataOps live in practice within two development teams with distinct primary skills and tools.

The different goals of DataOps and MLOps represent an additional obstacle. While DataOps aims to create quality data in a reliable and automated manner, any machine-learning method introduced must act as an assistant. Consequently, MLOps techniques should provide robust quality assessment and anomaly detection support that supports domain expert decisions, enable data-preparation automation that improves efficiency, and ensure a seamless, fail-proof train-predict-iterate circle.

6.3. Best Practices for Overcoming Challenges

The convergence of data and machine-learning pipelines brings intelligence to well-known DataOps problems such as ETL automation, data-quality assessment, and anomaly detection. Typically, these problems are tackled by dedicated teams with a DataOps mindset. The emergence of MLOps establishes a similar mindset, organization, culture, and tools for machine-learning pipelines. Combining DataOps and MLOps practices is a logical evolution that can significantly increase automation and intelligence in production data and machine-learning pipelines.

Surveyed practitioners have observed that for effective collaboration between DataOps and MLOps teams on end-to-end data pipelines, "a cultural shift and organizational mindset supported by a conducive structure" is essential. Realizing the vision of data and machine-learning pipelines as a single intelligence embedding pipeline reveals several challenges, including knowledge gaps, duplicated tasks, poor pipeline quality, conflicting objectives, focused projects over shared outcomes, and siloed domain knowledge. Practitioners have shared recommended practice patterns to address these issues and capitalize on the benefits of integrated operations.

7. Future Trends in MLOps and DataOps

MLOps and DataOps have followed an analogous trajectory, shaped by analogous forces and driven by analogous forces—first the cloud, followed by containers and orchestration, and now AI. For the same reasons, the two domains will also converge and merge into a single practice, drawing upon complementary core competencies.

The automation of ETL and ELT workflows, data quality assessment, and data anomaly detection has traditionally been carried out using rule-based algorithms and heuristics. The incorporation of AI can produce better results and enable the automation of more aspects of an analyst's or a data engineer's work. Indeed, the differentiation of the future AI orchestration platforms currently under development will be their focus on automation or AI-first support for these critical processes.

7.1. Emerging Technologies

The integration of MLOps and DataOps supports the deployment of machine learning models into production and engineering workflows for analytics, reporting and decision making. Cloud providers as well as open-source projects have produced numerous supporting tools and managed services. For example, Amazon Web Services has announced the general availability of code, notebooks, pipelines, evaluation reports and templates for training, tuning, endpoint deployment and data capture alongside an orchestrating workflow for monitoring data drift and report generation. Data quality frameworks augment batch and streaming data pipelines to produce detailed reports. Services such as Amazon SageMaker Model Monitor support baseline creation and data profile monitor deployment for batch and real-time inference within active endpoints.

Embedded intelligence in data pipelines offers clear benefits through improved monitoring, decision-making, quality and automation. Further, decision-making automation has been brought out even to traditional ETL/ELT workflows which are almost fully automated based on AI and have expert users empowered with or replaced by the value of AI. Data quality assessment models use validations based on rules (RBVs), which can be improved integrating AI techniques to avoid the drawbacks of fixed schemas and plain heuristics. Just as classification-based approaches are replacing rule and threshold-based methods that analyze data anomalies through

machine learning, anomaly detection based on supervised models is subsuming classical anomaly scoring techniques.

7.2. Predictions for the Next Decade

Automatic of ETL/ELT workflows with AI is a hot area in research. The basic idea is to push intelligence into the data pipelines to improve efficiency and reduce human effort to handle data, especially in anomaly detection which may be present in machine learning models. These activities have found new forms in the recently introduced paradigms of MLOps-from-the-Data-Quality-Perspective and DataOps-paradigms, working to bring DevOps into data-science-and-respectively data-engineering-operations.

The integration of machine-learning methods addresses two main challenges. There is a demand for deploying machine-learning models that determine the possible causes of quality failures and for implementing models that anticipate these causes. This dual capability would enable data-engineering and data-science teams to be proactive, providing them with triggers for data failure and difficulties, and to respond promptly to data failures. However, attempts to meet this demand have failed because many of the more mature MLOps and DataOps frameworks and tools emphasize the operationalization of the MLOps lifecycle only after release or focus on establishing Continuous Integration and Continuous Delivery pipelines for machine-learning projects. Typically, MLOps and DataOps do not explicitly consider the operationalization of the data-preparation pipelines that make datasets appropriate for training machine-learning models.

8. Conclusion

MLOps and DataOps are the two faces of the automated data process cycle. While MLOps deals with exposing and automating the intelligence of the data in the pipeline, DataOps focuses on the pipeline itself. The two are seldom considered together, but many can benefit from the synergies of embracing both philosophies. Having intelligence embedded inside the pipeline can help in automating much of the ETL/ELT process as well as providing a quality assessment of the data.

Embedding intelligence inside the pipeline automates much of the ETL/ELT process and provides a quality assessment of the data. A specific popular aspect of quality assessment— anomaly detection—is chosen to demonstrate the potential of introducing Machine Learning models into the data pipeline. An overview of how MLOps and DataOps complement each other is presented, followed by the various types of data pipelines. Finally, using univariate forecasting of data-behaviour as an example, an end-to-end deployment is conducted to reveal one approach to combining these two philosophies for a common purpose.

References:

- [1] HUSSAIN, Fatima; HUSSAIN, Rasheed; HOSSAIN, Ekram. Explainable artificial intelligence (XAI): An engineering perspective. arXiv preprint arXiv:2101.03613, 2021.
- [2] Koneti SB. Artificial Intelligence in Financial Systems: Digital Transformation, and Machine Learning Applications. Available at SSRN 5401202. 2025 Aug 12.
- [3] Muppala M. SQL Database Mastery: Relational Architectures, Optimization Techniques, and Cloud-Based Applications. Deep Science Publishing; 2025 Jul 27.
- [4] Muppala M. Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience. Digital Oceans: Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience| Deep Science Publishing. 2025 Jul 8.
- [5] Panda SP, Koneti SB, Muppala M. Benefits of Site Reliability Engineering (SRE) in Modern Technology Environments. Available at SSRN 5285768. 2025 May 1.
- [6] Koneti SB. Artificial intelligence Applications in Retail and Investment Banking: Personalization, Robo-Advisory and Behavioral Analytics. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025;4:72.
- [7] Muppala M. Architectures in relational databases: An analytical study of SQL-based data models and ACID principles. database.;2:4.
- [8] Bentahar J. A Survey on Explainable Artificial Intelligence for Network Cybersecurity. arXiv (Cornell University). 2023 Mar 7.
- [9] Gadde H. AI-Assisted Decision-Making in Database Normalization and Optimization. International Journal of Machine Learning Research in Cybersecurity and Artificial Intelligence. 2020;11(1):230-59.
- [10] Koneti SB. Algorithmic Trading and Quantitative Finance Strategies: High-Frequency Trading, Market Microstructure, and Risk Optimization Models. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025;4:17.
- [11] Panda SP. The Evolution and Defense Against Social Engineering and Phishing Attacks. International Journal of Science and Research (IJSR). 2025 Jan 1.
- [12] Mohapatra PS. Artificial Intelligence and Machine Learning for Test Engineers: Concepts in Software Quality Assurance. Intelligent Assurance: Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. 2025 Jul 27:17.
- [13] Koneti SB. Analysis, Predictive Analytics, and Macroeconomic. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025 Aug 12:90.
- [14] Panda S. Scalable Artificial Intelligence Systems: Cloud-Native, Edge-AI, MLOps, and Governance for Real-World Deployment. Deep Science Publishing; 2025 Jul 28.
- [15] Shivadekar S, Halem M, Yeah Y, Vibhute S. Edge AI cosmos blockchain distributed network for precise ablh detection. Multimedia tools and applications. 2024 Aug;83(27):69083-109
- [16] Reis J, Housley M. Fundamentals of data engineering. " O'Reilly Media, Inc."; 2022 Jun 22.
- [17] Ivanov SH, Webster C. Adoption of robots, artificial intelligence and service automation by travel, tourism and hospitality companies—a cost-benefit analysis. Artificial intelligence and service automation by travel, tourism and hospitality companies—a cost-benefit analysis. 2017.
- [18] Ramadhan M, Naseeb A. The cost benefit analysis of implementing photovoltaic solar system in the state of Kuwait. Renewable energy. 2011 Apr 1;36(4):1272-6.
- [19] Cordes JJ. Using cost-benefit analysis and social return on investment to evaluate the impact of social enterprise: Promises, implementation, and limitations. Evaluation and program planning. 2017 Oct 1;64:98-104.
- [20] Dykes PC, Curtin-Bowen M, Lipsitz S, Franz C, Adelman J, Adkison L, Bogaisky M, Carroll D, Carter E, Herlihy L, Lindros ME. Cost of inpatient falls and cost-benefit analysis of implementation of an evidence-based fall prevention program. InJAMA Health Forum 2023 Jan 6 (Vol. 4, No. 1, pp. e225125-e225125). American Medical Association.

- [21] Frank J. Artificial intelligence and intrusion detection: Current and future directions. In Proceedings of the 17th national computer security conference 1994 Oct 11 (Vol. 10, pp. 1-12).
- [22] Wang F, Preininger A. AI in health: state of the art, challenges, and future directions. Yearbook of medical informatics. 2019 Aug;28(01):016-26.
- [23] Lu Y. Artificial intelligence: a survey on evolution, models, applications and future trends. Journal of management analytics. 2019 Jan 2;6(1):1-29.
- [24] Cordes JJ. Using cost-benefit analysis and social return on investment to evaluate the impact of social enterprise: Promises, implementation, and limitations. Evaluation and program planning. 2017 Oct 1;64:98-104.
- [25] Dykes PC, Curtin-Bowen M, Lipsitz S, Franz C, Adelman J, Adkison L, Bogaisky M, Carroll D, Carter E, Herlihy L, Lindros ME. Cost of inpatient falls and cost-benefit analysis of implementation of an evidence-based fall prevention program. In JAMA Health Forum 2023 Jan 6 (Vol. 4, No. 1, pp. e225125-e225125). American Medical Association.
- [26] Gadde H. AI-Enhanced Adaptive Resource Allocation in Cloud-Native Databases. Revista de Inteligencia Artificial en Medicina. 2022 Oct 18;13(1):443-70.
- [27] Konecni SB. Microstructure, and Risk Optimization Models. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025 Aug 12:17.
- [28] Panda SP. Augmented and Virtual Reality in Intelligent Systems. Available at SSRN. 2021 Apr 16. Mohapatra PS. Artificial Intelligence-Driven Test Case Generation in Software Development. Intelligent Assurance: Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. 2025 Jul 27:38.
- [29] Eboigbe EO, Farayola OA, Olatoye FO, Nnabugwu OC, Daraojimba C. Business intelligence transformation through AI and data analytics. Engineering Science & Technology Journal. 2023 Nov 29;4(5):285-307.
- [30] Mohapatra PS. Intelligent Assurance Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. Deep Science Publishing. 2025; doi:10.70593/978-93-7185-046-9
- [31] Ojika FU, Owobu WO, Abieba OA, Esan OJ, Ubamadu BC, Daraojimba AI. Transforming cloud computing education: Leveraging AI and data science for enhanced access and collaboration in academic environments. Journal name and details missing. 2023 Jan

Chapter 5: The Age of Vector and Graph Databases: Foundations for Advanced Information Retrieval and Reasoning

1 Introduction

The explosive growth of unstructured data and the escalating complexity of AI applications have begun to reveal fundamental flaws in traditional relational database systems. These systems are intended to be used with structured data and predefined schemas, making them ill-equipped to handle the kinds of semantic relationships and complex interdependencies that now frequently crop up in today's modern types of data: text, images, knowledge graphs. This technology mismatch has encouraged several special-purpose database paradigms that work well for certain types of data and operations.

Two such complementary approaches are vector and graph databases. Vector-based databases store items as mathematical vectors in a high-dimensional space which permits to perform efficient retrieval based on similarity that captures semantics, instead of exact equivalence. In contrast, graph databases use node-edge-node structures to directly express relationships between entities on the data keeping level allowing for advanced processing and analysis of linked data. The combination of these technologies with more advanced AI methods has led to powerful knowledge-based application frameworks.

This chapter explores how these database technologies underpin today's AI systems. In the following, we first compare vector and graph databases focusing on their different data model, queries as well as potential usage. We also walk through the foundational idea of vector embeddings and how they support semantic search, as well as an investigation into how retrieval systems are combined with large language models using Retrieval-Augmented Generation (RAG) architectures. Finally, we are considering knowledge graphs and reasoning engines, bringing explainable inference to AI systems. We highlight throughout the practical applications of these technologies and their synergistic effects in fostering advanced, efficient, and transparent intelligent systems..

2 The Age of Vector and Graph Databases

The abstraction between vector and graph databases is the underlying data model and what kinds of relationships are most important. The aim of this review is to highlight these differences and show the reasons for choosing the right technology, and also to understand their complimentary potentials in hybrid architectures [1-3].

2.1 Data Representation Models

Vector databases convert raw data of any sort — text, images, audio — into a numerical representation in high-dimensional space called the embedding. These embeddings map the data points into a continuous vector space in which distances correspond to degrees of semantic similarity. For instance, in an e-commerce use case like online shopping items' descriptions are converted to vectors so that similar items falls on same set of points(not only close but also harder for recommendation system to find interested items if different vendors decide to offer the same item under a slightly different name / description) [2].

In contrast, graph databases explicitly model complex networks using a structure formed of nodes (entities or concepts) and edges (relations). This model is a natural representation for connected domains such as social networks, where the vertices are users and edges represent connecting pair of them or fraud detection systems, where patterns of relationships among entities can indicate to anomalous activities.

Table 1: Comparative Analysis of Vector and Graph Databases

Feature	Vector Databases	Graph Databases
Data Representation	Points in multi-dimensional space based on semantic similarity	Nodes (entities) and edges (relationships) forming interconnected networks
Primary Query Method	Similarity search using metrics like cosine similarity or Euclidean distance	Graph traversal algorithms (e.g., breadth-first search) to navigate relationships
Optimal Use Cases	Recommendation systems, semantic search, anomaly detection	Social network analysis, fraud detection, knowledge representation
Scalability Considerations	Generally handle large-scale similarity searches well with horizontal scaling	Can face performance challenges with highly complex queries over massive graphs

2.2 Complementary Strengths and Convergence

Despite their different modeling approaches, vector and graph databases are strong in different types of scenarios. Vector-based systems are useful for similarity search in unstructured data, while graph-based databases can address the complicated relationships and paths from proximal data..

Curiously, these paradigms are now more and more merging in modern AI systems. Vector-like capabilities are being integrated to graph databases endowing them with the ability to perform similarity searches on node properties and not just an nested relationship traversal. Conversely, some vector databases are integrating graph-like relationships to enhance their similarity metrics with contextual information [2,4,5]. This Now the intersected items of two facets represent an entity's both semantic attributes and its position relations in a knowledge graph.

3 Vector Embeddings and Semantic Search

Vector embeddings are the mathematical basis for making sense of meaning between words in AI systems. These techniques allow automatic derivation of continuous vector space representations of discrete data, in which the similarity between elements can be queried to retrieve semantically related information.

3.1 The Embedding Generation Process

Vector embeddings are compact numerical representations that encode relevant properties and semantics of data into vectors in a very high-dimensional space. This is based on the fundamental framework of neural networks, in particular specialized paradigms such as Word2Vec for text variations or Convolutional Neural Networks (CNNs) for images, that learn to translate raw data into vectors through representation learning. In training, these networks try to move semantically similar items to places close together in the vector [6-8]. For example, words that have similar meanings (e.g. king and queen) have very close vector representations in the semantic space, whereas unrelated words like car are placed further.

The high dimensionality of these embeddings (usually a few hundred to a thousand dimensions) affords the capacity necessary to encode such complex semantic relationships. Modern embedding models such as OpenAI's text-embedding-ada-002 convert variable-length text into fixed-dimensional vectors, providing a uniform input size for processing. This transformation opens the door to all kinds of data, from text and images, through audio and other sensations up to molecular structures, being processed in one single mathematical space where dependencies between any kinds of inputs might be quantified systematically.

3.2 Similarity Metrics and Search Implementation

In a simplex embedding, query and document representations are projected into this embedded space and their cosine similarity in such a space is measured. Unlike keyword-based search, which uses exact matches, semantic search grasps the meaning of context, and is able to serve up relevant results despite variations in terminology across queries and documents.

The performance of semantic search is tightly bound to similarity metric selected. Some of the widely used measures are as follows:

Cosine similarity: Computes the cosine of the angle between vectors a and b (the vectors must have length 1 for this to be a proper distance metric). That makes it especially useful for text tasks where we need frequency-independent semantic similarity [9,10] .

The (dis)similarity measure is euclidean distance, which measures the straight-line distance between points in a vector space; it is easy to interpret in terms of absolute separation.

In practice, semantic search systems use Approximate Nearest Neighbor (ANN) algorithms (e.g., Facebook's FAISS and Google's ScaNN), to perform efficient searches in high-dimensional vector spaces. These methods allow fast similarity search over even billion-scale vector databases, so semantic search is possibly for real-time scenarios.

4 Integrating with LLMs for Retrieval-Augmented Generation (RAG)

Cross-linking retrieval systems and Large Language Models (LLMs) via Retrieval-Augmented Generation (RAG) is a crucial step towards addressing the limitations of LLMs such as static knowledge truncation, hallucination and non-expertise in domain. RAG architectures effectively bridge the gap between the generative capabilities of LLMs and the dynamic, verifiable knowledge stored in external databases.

4.1 RAG Architecture and Workflow

A typical RAG system follows a structured pipeline that combines information retrieval with contextual generation [11-13]. The process begins with an indexing phase, where domain-specific documents are chunked into manageable segments, converted into vector embeddings, and stored in a vector database . At inference time, when a user submits a query, the system embeds this query using the same model and performs a similarity search against the vector database to retrieve the most relevant contextual documents .

These retrieved documents are then concatenated with the original query and fed to an LLM, which generates a response grounded in the provided context . This approach significantly

enhances the factuality, accuracy, and timeliness of LLM outputs while providing explainability through source attribution. For example, a customer support chatbot utilizing RAG can retrieve relevant information from a constantly updated knowledge base and generate responses that reflect the most current policies and procedures.

4.2 Evolution of RAG Frameworks

RAG designs have moved from simplistic naive implementations to abstract and modular structures. Naive RAG Naive Retrieval-Generation techniques extracting replies by simply retrieving a candidate context and estimating its likelihood as a reply to the current turn have been already evaluated using metrics like retriever precision, in which was found naive New state-of-the-art results on how AdvRAG New SOTA Applying improved pre-retrieval optimization (e.g. better chunking strategies, the enrichment of metadata), enhanced retrieval techniques (e.g. fine-tuned embeddings) and post-retrieval refinements (e.g. reranking or prompt compression may improve these results) addresses this~"{}" This is also brought to question by new experiments with SingleSeqStruc with more suitable architecture and a better pre-training method [2,14-17].

The most adaptable position, which accepts additional functional blocks, is Modular RAG that includes search modules, memory mechanisms and routers that adapt the retrieval policy according to the query. This modularity facilitates the use of advanced methods as recursive retrieval, that begins with smaller semantic chunks and goes higher with respect to larger context semantics, or query Decomposition which decomposes complex questions relative to simple sub-queries able to address different sources of data.

5 Knowledge Graphs and Reasoning Engines

Although vector databases are efficient at similarity-driven retrieval, KGs and their reasoning capabilities offer complementary strengths for representing structured knowledge and conducting logical inference. Such technologies add explainability, relational understanding and deductive power to AI systems, especially in the setting of complex decision making based on facts interconnected [9,18-21].

5.1 Knowledge Representation and Inference Mechanisms

A knowledge graph is an architectural way of representing knowledge as a linked-network of entities, where the nodes are concepts or objects and the links are their relationships. This framework fitly contains complex relationship (i.e., hierarchical classification, spatial relations and temporal order). Contrast to vector representation where semantic closeness is encoded in a continuous space implicitly, Knowledge graphs encode relationship in an explicit and symbolic fashion

LogicaJ reasoning engines apply logical rules on these explicit representations to derive new knowledge (I). For instance, given facts "Person X works for Company Y" & "Company Y is headquartered in City Z", it should be possible to provide RSETs with a reasoning engine which can deduce the implicit fact that Person X is based on City Z. The techniques which advanced reasoning engines such as RDFox use are incremental in the sense that real-time inferences can be updated when new data comes in without having to recompute the entire knowledge base. They also support negation as failure (by being able to infer based on the lack of some data), and aggregation, permitting numerical calculations over groups of data in the graph [22,23].

5.2 Applications and Integration with AI Systems

Domain-specific regulatory requirements and the ability of knowledge graphs to represent relationships explicitly and perform logical inference, which makes them particularly promising in regulated industries and complex decision-making domains. On health, knowledge graphs could depict relations among symptoms and diseases or treatments (thereby allowing reasoning engines to propose the most probably diagnosis according to a symptom pattern). For example in finance, they can discover complex fraud cases that have uncommon relations between entities and do not get discovered through isolated similarity searches[24-26].

Integrated with LLMs, KGs serve as sources of structured knowledge grounding, which contributes to improved factual and relational correctness in the generated responses. Such integration can operate in both directions: knowledge graphs may provide verifiable information to LLMs, and in reverse, LLMs may help fill out and update knowledge graphs by filling it with structured information extracted from free-form text. This nexus of statistical learning from LLMs and symbolic reasoning with knowledge graphs offers a promising route to more powerful and reliable AI systems

6 Synthesis and Future Directions

The technologies that we have seen in this chapter—vector and graph databases, semantic search, RAG, and knowledge graph reasoning—are additional pieces of a modern AI stack. Instead of just being its own separate choice, they form the integrated parts in an advanced AI design. There is a need to understand the synergies between these views and their possible trajectories for future intelligent systems.

6.1 Hybrid Architectures and Emerging Convergences

The most critical trend impacting database infrastructure for AI is the merging of vector and graph techniques. Graph databases are adding vector capabilities in their systems, allowing for similarity searches on node properties; whereas vector databases are embracing the addition of graph-like relationships to improve its semantic matching with surrounding contexts. This hybridization makes for powerful synergies: for example, a recommendation system could use

vector similarity to find similar products and graph traversals to discover items liked by users with similar network profiles.

For RAG architectures, such convergence enables advanced retrieval strategies [27,28]. Hybrid search The hybrid search approaches fuse the semantic vector search with traditional keyword-based as well graph-based retrieval in order to handle different query types and information requirements. For multi-perspective questions, taking different point of view systems can break down the question into suitable sub-queries in which each sub-query is posed to a different database system and by combining the returned information they attempt to produce a coherent answer. This multi-modality retrieval method leads to much stronger and richer AI systems..

6.2 Research Challenges and Future Outlook

Although much progress has been made, many challenges in scaling and improving these techniques still puzzle researchers. Computational efficiency remains an issue as the size of datasets increase to billions of vectors and trillions of graph relationships. This scaling is being challenged considerably by more efficient indexing structures, approximate algorithms and hardware acceleration. Moreover, bias in embedding mitigation is a continuing reminder before embeddings that comes with continuing neglect as the pre-trained vectors may amplify social biases in training data.

In the future, we expect greater emphasis to be put on the real-time reasoning capabilities that could leverage the strengths of the pattern recognition in neural methods and other forms with explainability from symbolic AI [19,29-31]. Such technologies can unleash edge-based implementations that will make AI applications more responsive as well as more privacy-conscious. Moreover, when AI systems come to play an increasingly important role in high-stakes decisions (e.g., about my health or wealth), the accountability and explainability provided by tools like knowledge graphs will be necessary rather than just highly desirable [32].

7 Conclusion

In this chapter, we discussed the role of vector database systems and graph database systems as well as technologies related to these newfound types in transforming AI. We showed how these dedicated data management systems overcome basic limits in managing unstructured information and relationships, making it possible to support more advanced AI applications.

The era of vector and graph databases is a time when “one size fits all” data management is replaced by specialized offerings developed for certain variety of the data and the universe of use cases that surround it. Vector databases are especially effective to capture semantic similarities and deliver context-sensitive retrieval due to the embedded-based approach. Graph databases are extremely powerful at modelling and querying for complex relationships. When incorporated via interfaces such as RAG, they enrich large language models with external knowledge to reduce

hallucinations and support domain task applications. In addition, knowledge graphs and reasoning engines deliver explainable, logic-based inference to AI systems — critical in regulated sectors and in complex decision scenarios.

As these technologies merge and progress, they promise to become ever more central in the construction of AI systems that are but not only, capable, but also transparent, trustworthy and aligned with human modes of reasoning. The future is not to decide between these paradigms but to exploit the complementary strengths of each other with hybrid approaches that combine a statistical power of neural methods and structured reasoning in symbolic AI.

References

- [1] Ojika FU, Owobu WO, Abieba OA, Esan OJ, Ubamadu BC, Daraojimba AI. Transforming cloud computing education: Leveraging AI and data science for enhanced access and collaboration in academic environments. Journal name and details missing. 2023 Jan
- [2] Panda SP, Padhy A. Business Intelligence with Power BI and Tableau: Cloud-Based Data Warehousing, Predictive Analytics, and Artificial Intelligence-Driven Decision Support. Deep Science Publishing; 2025 Aug 15.
- [3] Koneti SB. Microstructure, and Risk Optimization Models. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025 Aug 12:17.
- [4] Mohapatra PS. Artificial Intelligence-Driven Test Case Generation in Software Development. Intelligent Assurance: Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. 2025 Jul 27;4:38.
- [5] Koneti SB. Artificial Intelligence in Financial Systems: Digital Transformation, and Machine Learning Applications. Available at SSRN 5401202. 2025 Aug 12.
- [6] Muppala M. SQL Database Mastery: Relational Architectures, Optimization Techniques, and Cloud-Based Applications. Deep Science Publishing; 2025 Jul 27.
- [7] Muppala M. Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience. Digital Oceans: Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience| Deep Science Publishing. 2025 Jul 8.
- [8] Panda SP, Koneti SB, Muppala M. Benefits of Site Reliability Engineering (SRE) in Modern Technology Environments. Available at SSRN 5285768. 2025 May 1.
- [9] Koneti SB. Artificial intelligence Applications in Retail and Investment Banking: Personalization, Robo-Advisory and Behavioral Analytics. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025;4:72.
- [10] Muppala M. Architectures in relational databases: An analytical study of SQL-based data models and ACID principles. database.;2:4.
- [11] Bentahar J. A Survey on Explainable Artificial Intelligence for Network Cybersecurity. arXiv (Cornell University). 2023 Mar 7.
- [12] Gadde H. AI-Assisted Decision-Making in Database Normalization and Optimization. International Journal of Machine Learning Research in Cybersecurity and Artificial Intelligence. 2020;11(1):230-59.
- [13] Koneti SB. Algorithmic Trading and Quantitative Finance Strategies: High-Frequency Trading, Market Microstructure, and Risk Optimization Models. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025;4:17.
- [14] HUSSAIN, Fatima; HUSSAIN, Rasheed; HOSSAIN, Ekram. Explainable artificial intelligence (XAI): An engineering perspective. arXiv preprint arXiv:2101.03613, 2021.

- [15] deployment pipelines. *International Journal of Research Publication and Reviews*. 2025 Jan;6(1):871-87.
- [16] Reis J, Housley M. *Fundamentals of data engineering*. " O'Reilly Media, Inc."; 2022 Jun 22.
- [17] Ivanov SH, Webster C. Adoption of robots, artificial intelligence and service automation by travel, tourism and hospitality companies—a cost-benefit analysis. *Artificial intelligence and service automation by travel, tourism and hospitality companies—a cost-benefit analysis*. 2017.
- [18] Ramadhan M, Naseeb A. The cost benefit analysis of implementing photovoltaic solar system in the state of Kuwait. *Renewable energy*. 2011 Apr 1;36(4):1272-6.
- [19] Cordes JJ. Using cost-benefit analysis and social return on investment to evaluate the impact of social enterprise: Promises, implementation, and limitations. *Evaluation and program planning*. 2017 Oct 1;64:98-104.
- [20] Dykes PC, Curtin-Bowen M, Lipsitz S, Franz C, Adelman J, Adkison L, Bogaisky M, Carroll D, Carter E, Herlihy L, Lindros ME. Cost of inpatient falls and cost-benefit analysis of implementation of an evidence-based fall prevention program. *InJAMA Health Forum* 2023 Jan 6 (Vol. 4, No. 1, pp. e225125-e225125). American Medical Association
- [21] Chu H. *Information representation and retrieval in the digital age*. Information Today, Inc.; 2003.
- [22] Dominich S. *Mathematical foundations of information retrieval*. Springer Science & Business Media; 2001 Mar 31.
- [23] Reinanda R, Meij E, de Rijke M. Knowledge graphs: An information retrieval perspective. *Foundations and Trends® in Information Retrieval*. 2020 Oct 14;14(4):289-444.
- [24] O'Leary DE. *Artificial intelligence and big data*. *IEEE intelligent systems*. 2013 Jun 27;28(2):96-9.
- [25] Gadde H. AI-Augmented Database Management Systems for Real-Time Data Analytics. *Revista de Inteligencia Artificial en Medicina*. 2024;15(1):616-49.
- [26] Ojika FU, Owobu O, Abieba OA, Esan OJ, Daraojimba AI, Ubamadu BC. A conceptual framework for AI-driven digital transformation: Leveraging NLP and machine learning for enhanced data flow in retail operations. *IRE Journals*. 2021 Mar;4(9).
- [27] Muppala, M. . (2025). Digital Oceans: Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience. *Deep Science Publishing*. 2025; doi:10.70593/978-93-7185-787-1
- [28] Koneti SB. Artificial Intelligence-Powered Finance Algorithms, Analytics, and Automation for the Next Financial Revolution. *Deep Science*. 2025; doi:10.70593/978-93-7185-613-3
- [29] Panda SP. *Artificial Intelligence Across Borders: Transforming Industries Through Intelligent Innovation*. Deep Science Publishing. 2025; doi:10.70593/978-93-49910-25-6
- [30] Eboigbe EO, Farayola OA, Olatoye FO, Nnabugwu OC, Daraojimba C. Business intelligence transformation through AI and data analytics. *Engineering Science & Technology Journal*. 2023 Nov 29;4(5):285-307.
- [31] Mohapatra PS. *Intelligent Assurance Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle*. Deep Science Publishing. 2025; doi:10.70593/978-93-7185-046-9
- [32] Ojika FU, Owobu WO, Abieba OA, Esan OJ, Ubamadu BC, Daraojimba AI. Transforming cloud computing education: Leveraging AI and data science for enhanced access and collaboration in academic environments. *Journal name and details missing*. 2023 Jan

Chapter 6: Exploring Security, Governance, and Explainability in AI Systems

1. Introduction

There are numerous hurdles in AI Systems too ranging from security of the models to model governance to model explainability. Anomaly Detection of AI Systems, Data Lineage of AI Systems and Compliance with eXplainable_AI. Each challenge comes with a specific problem statement, solutions and references to real world use cases along with libraries, tools or services.

AI is the science and engineering of making intelligent machines (intelligence in such sense, enables learning, based on experience, user provided knowledge reasoning, and adapting to new environment s). AI learns training data to address similar problems. The veracity of predictions or classifications is dependent on the assimilation and learning procedure. The AI world divides into—narrow AI labs building systems for very specific tasks; artificial superintelligence work aiming to exceed human intelligence; and artificial general intelligence trying to match the human ability to reason in a variety of domains. It is safe to say that end-use AI applications are limitless, encompassing from disease diagnosis in medicine to intelligent gaming models, voice-activated assistants or self-driving cars and the application of machine created art.

2. Understanding AI Systems

Foundations Artificial Intelligence (AI) is the science and engineering of making intelligent machines. Weak AI does specific tasks (think Siri or Alexa), Strong AI hopes to do all that humans can as we wander our ways through the world. Applications of AI include Decision tree regression, AdaBoost, Support Vector Machines (SVM), Naive Bayesian classification, K-nearest neighbours (KNN), Neural Networks and others. IBM's Deep Blue (Chess) are some of the AI applications which shows how we can apply AI to wide variety of areas. For example, Liquibase is a database schema change management tool which designs to keep track of the schema change history so as to detect anomaly based on data lineage[1].

Anomaly Mining for AI Systems The security considerations in AI systems are not mature and are likely to evolve with the increase of risk exploiting data and model, when being sourced through multi-party collaborations, e.g., data annotation & training. Sensitive applications,

especially in law enforcement, financial analytics, or credit risk analysis, require corporate governance policies for AI systems to satisfy regulations like the General Data Protection Regulation or the Equal Credit Opportunity Act (ECOA). Regulatory bodies such as the Consumer Financial Protection Bureau oversee these areas, highlighting the importance of comprehensible models for outcome explanation. Best practice governance policies therefore leverage anomaly detection to ensure AI models perform as anticipated.

2.1. Definition of AI

A working definition of artificial intelligence (AI) can be constructed by considering several aspects. Firstly, AI is intelligence exhibited by machines. Secondly, considering the distinctions between artificial and natural, absent and existent, nonbiological and biological, inorganic and organic, unnatural and natural, AI is also unnatural intelligence that may be manifested naturally or artificially [1-2]. Thirdly, considering distinctions between machines and organisms, unnatural intelligence exhibited by machines is artificial. Fourthly, considering different types of intelligence, AI corresponds to the performance of activities associated with human thinking or intelligent behaviour. Finally, considering that designing machines with AI consists of making them perform activities requiring experiences associated with human intelligence, AI can be defined as the discipline of designing machines that exhibit unnatural intelligence.

The definition, in its broadest sense, does not account for the major differences between various levels of AI. For example, AI may merely emulate human behaviour, as is the case in the game of chess, or it may emulate the human mind, producing behaviour that is hardly distinguishable from human behaviour, such as determining a diagnosis or recommending a course of action. Four types of AI have been proposed: reactive AI (such as Deep Blue); limited-memory AI (such as Siri); theory of mind AI (advanced psychological systems under development); and self-aware AI (capable of self-awareness). Understanding AI solely as a discipline of designing machines capable of exhibiting unnatural intelligence greatly broadens its scope and application, enabling it to penetrate virtually every aspect of daily life.

2.2. Types of AI Systems

AI systems currently in use may be categorized into narrow AI, also called weak AI, and AGI, also called strong AI, human-level AI or deep AI. Narrow AI is designed and trained to perform specific tasks. Voice recognition systems, image recognition software, recommendation engines and self-driving cars are all examples of narrow AI applications. Artificial General Intelligence is an AI system with general-purpose understanding and reasoning abilities that is similar to or exceeds those of humans. AGI research is evolving rapidly and includes applications in gaming, voice recognition, image recognition, robotics, text generation and many others.

Within the AI category, there are several sub-areas including Machine Learning, a technique that enables computers to learn from new data inputs or experiences; Machine Reasoning, focused on reasoning processes; Robotic Process Automation; Supervised Learning, using labeled data to

train models; Unsupervised Learning, processing information without labeled examples; and Natural Language Processing, which deals with text and voice recognition. Although each subfield serves different purposes, they are all focused on using machines to simulate elements of higher intelligence for specific tasks [3-5]. Examples of AI applications include recommendation systems, healthcare diagnostics, autonomous vehicles, social media, and smart manufacturing.

2.3. Applications of AI

The ability of AI to learn operations or strategies from training data or an environment means it can perform a wide range of tasks, as long as the task can be described well and input data are available. Modern AI systems use deep learning to mimic the human brain, enabling them to create associations and learnable patterns between any kind of data. Reinforcement learning is an iterative method in which an AI agent learns to solve a task by exploring different steps in a potentially very complex landscape. In the following subsection, some important use cases of AI are listed, which help in understanding the scope of AI systems.

Personal assistants like Siri, Google Now, and Cortana provide users with help by sending messages, searching for information, providing traffic information, and so on. Few mobile phones currently offer AI features as an upgrade, aiming to assist the user even more. Many games utilize AI to provide players with believable and challenging opponents. AI bots can be used in these games for social media marketing purposes, to send automatically generated messages to a large amount of users with the same content or topic or to provide answers to frequently asked questions. AI systems also play a major role in other industries. In the banking sector, credit scoring is applied to assess a customer's creditworthiness. AI is also present in fraud detection. Data of new customers are matched with the crime data to see if there is coincidence between the two. Audit trails could be analyzed by AI technology, the data in contractor forwarding and generation of reports that could detect discrepancies.

3. Security in AI Systems

AI Security studies how AI systems react to arbitrary inputs that induce incorrect behavior and investigates mechanisms of attacks or vandalism on these systems. It includes the identification, resistance and recovery from such adversarial events. Critical threats include model extraction, evasion and poisoning [6-8]. Model stealing attacks want to copy the model's capabilities perfectly or just approximately and do so without authorization; they can result in proprietary models leak or service integrity compromise. Evasion attacks utilize perturbed data that is true under normal evaluations and cause a misclassification during runtime (adversarial example) which have been demonstrated to fool real-world systems including self-driving cars. Backdoor attacks consist in poisoning training or testing data with perturbed examples, or abusing the feature extraction pipeline in those phases, for degrading model performance, biasing inputs toward certain classes or facilitating target evasion at inference time.

Credit allayance tactics range from regulatory participation, platform design to model-level resuscitation. These regulations aspire to create normative security model, compare with existing standards and clarify who is responsible and liable. In terms of platform, their concerns are about defending against attacks on process and data – e.g., by sanitizing the training data to prevent poisoning or filtering at runtime to resist evasion – so that they can enforce model confidentiality, integrity and availability with as low operational overhead [9]. However, The practical use of these methods is very limited, especially in the evasion attacks. Formulating model defense methods aim at achieving recovery of original performance by excluding backdoor functionality or improving resistance with respect to input perturbation. Examples of incidents that make this point include the 2016 trainwreck when Microsoft’s Tay chatbot was manipulated to create racist and obscene content; a 2017 Amazon Echo proof-of-concept procedure shattered so-called cryptographically secure secrets where signals could be sent out from an in-pocket device to record what supposed-to-be-private sounds were around; and the latest Capital One data breach which occurred because of a AWS cloud computing vulnerability, which wound up exposing more than 100 million customer accounts and credit card applications

3.1. Threats to AI Security

Artificial intelligence comes with security issues like other computing technologies, which can lead to very undesirable consequences. In recent years deep-learning models have been exposed to attacks such as data poisoning, evasion, or model-inversion attacks. These actions are addressed by adversarial machine-learning techniques. There is a need to identify other probable threats across the various parts of the system (AI model, training dataset, results, and underlying infrastructure), determine how these can be carried out, and develop mitigating controls.

Malicious hackers can take several approaches either individually or in combination to breach an ML system. They may perform a data poisoning attack by corrupting the training data, which will result in faulty training. They might attempt an inference attack on the training model to gain information about the training data or the model. The adversaries may also try to launch a reverse engineering attack to build a model similar to the training model. A successful system breach may allow them to alter or destroy the training model or its results[7,9-10]. External entities may breach the underlying infrastructure which hosts the ML system, initiate a denial-of-service attack, or take control of the system for command and control. They may also possess legitimate access rights to the model’s output and attempt to encrypt or sell it to untrusted parties.

3.2. Mitigation Strategies

While holistic security of capitals remains a demanding task, mitigating known attack types helps minimize the attack surface. Each attack category can be assessed based on damage, implementation level, target target capital, and attack direction. Although some attacks can have catastrophic damage, others allow limited inferences on sensitive information. User profiling data at the inference level can lead to a range of attacks if disclosed. Automatized Privacy

Setting Attack and the Feature Reconstruction Attack typically impact Symbiotic Capital, while the Hard Label Attack at the inference level affects all four identified capitals (Data, Model, Users, and Services). Model Extraction attacks generally bear high damage as they rely on compromised sensitive information.

Implementing mitigation strategies at the affected level(s) significantly curbs attack capabilities. For example, attacks taking place at the training data level are limited if no privileged access is granted to the model. Conversely, substantial damage emerges if attacks commence at the training privileges level or at the training model level. Similarly, the Sneaky Metadata Attribution attack requires metadata access to evade detection. Most model-level attacks interact directly with the user, necessitating their implementation within the user's interaction level. More broadly, three fundamental strategies underpin most mitigation measures:

- * Feature/Model Hardening: during ML/AI training phase, data sanitization is conducted to remove noise from data or strengthen the ML/AI model against manipulation.
- * Detection: implemented on the ML/AI system to identify anomalous model responses (testing/inference phase) or recognize manipulated data.
- * Response: upon detection of anomalies, a specialized response mechanism is triggered to counter the threat.

3.3. Case Studies in AI Security Breaches

Information exchange-no request and response relationship breaches involve sharing sensitive data with unauthorized parties [1,11-14]. Examples include a return flight booking email sent to an incorrect recipient, an investment risk evaluation questionnaire filled by an uninvolved participant, and disclosure of a confidential secure meeting notice to external vendors.

Denial-of-service (DOS) breaches can render network services useless, for instance, when an attack blocks a trading system, an irrecoverable error arises due to no reserve price data, or a shortage of seats results.

4. Governance of AI Systems

The growing adoption of artificial intelligence (AI) in a wide array of applications has created a number of concerns that have an impact on the decision-making process.

AI governance can be designed around the threats to and risks of an AI system, to the roles and responsibilities of those involved in the AI lifecycle, and to the policies, practices, and arrangement of controls for effectively directing and managing AI systems in order to achieve the organization's strategies and business objectives. AI governance addresses the management of the organization's AI by providing direction and continuous monitoring through the evaluation of AI models, parameters, sources, and targets, as well as through the clear assignment of roles and responsibilities. Ensuring holistic AI governance requires a mechanism to provide transparency into AI systems through explainability, security, compliance, fairness, and performance.

4.1. Frameworks for AI Governance

Artificial intelligence continues to advance and shape society. For when AI models are employed in decision-making, it is important to be able to get a sense of how they use data on an ongoing basis during their several-year lifecycle as a means of detecting abnormal behavior and as feedback mechanism for AI governance policies. Various approaches exist to secure and govern AI systems, including anomaly detection, data lineage and explainable AI. [13,15-17].

AI anomaly detection can be used to detect abnormal activities throughout the entire AI lifecycle. Anomaly-based approaches help users discover unusual behaviors related with a model or training data— some attacks that fall under this category are poisoning the input data, and identifying an unanticipated surge in the number of training records from a particular city. Data lineage enables to track the flow of information across the AI lifecycle. It provides traceability between the input sources, the operations executed upon them and its result. Data lineage allows for ensuring that AI governance policies are accomplished and it constitutes an important factor when it comes to explainable AI. Regulations such as the GDPR demand that humans execute decisions rather than relying fully on AI, thus when making AI-supported decisions it is fundamental to understand how AI arrives at a certain conclusion or recommendation.

4.2. Roles and Responsibilities

AI Governance is the process and policies used to govern an AI system. Its purpose is to assign clear and unambiguous responsibility for the tasks involved in designing, training, tuning, testing, deploying, monitoring, maintaining, updating, and decommissioning an AI system. The roles involved in governing an AI system and the associated policies differ for each phase and thus are described separately.

The testing procedure determines who is responsible for development and execution and what the expected outcome is by defining the governance test objectives, governance test setup, source of data, and pass/fail criteria or assessment methods. The monitoring process specifies who is responsible for the monitoring, when to monitor, and how to monitor the AI system and evolving operating environment; what to monitor (including measurement of KPIs and detection of anomalies in model and/or operating environment); what are the thresholds indicating an anomaly; and how to respond in case of an anomaly (e.g., raise a warning, launch an investigation, retrain or adjust the model, or deactivate the AI system). Operational procedures state the required governance aspects for the deployment, maintenance, and decommissioning of an AI system. These include the labeling requirements for the human consumers, limitations in web and/or mobile usage, frequency and amount of carbon footprint allowed, cost of service (subscription or per use), and any other policy–rule–process–regulation applicable in each phase.

The governance policies generate requirements for the verification, validation, and test activities for the AI system. In addition, they contribute to the configuration management of the AI System Artifacts – the documented and implemented information for each phase of the life cycle.

4.3. Policy Development and Implementation

Following the assignment of roles and responsibilities, the process of developing and publishing AI policies begins. The AI governance tool provides support with templates to accelerate policies creation, for example, from top-level rules to operational policies and procedures. The publication of policies signals the initial step toward formal AI governance. When policies are enacted on an AI system, governance controls, opinions, and anomaly detection become operational. Subsequently, the governance capability analyzes the policy content to identify obligations and prohibitions, and it determines the artefacts within the AI system to which these policies correspond. Finally, it enforces control mechanisms to verify that the policy requirements are properly addressed.

Operational Roles consolidate the AI security role, providing a comprehensive service framework for addressing all security-related requirements. A police Detection verifies that the final outcomes of an AI system, when combined with anomaly detection, can be used to measure the accuracy of the AI system results. Policy compliance policing checks how much the policies are satisfied and detects violations for further investigation [18-20].

5. Explainability in AI

Explainability is a critical design factor to make AI systems trustworthy and is closely associated with the capability to detect abnormal events by AI systems and to establish data lineage. 4Reliable AI develops trust in an autonomous system and enables human understanding, control, prediction and response to the behaviour of adversarial or unforeseen AIs. Explainability also refers to ethical and societal impacts of AI decision-making, and is useful for stakeholders who want to know: what the system aims to do; why it comes to a certain conclusion or prediction..

Explainable Artificial Intelligence (XAI) is important for various reasons: it enables the end-users and the stakeholders to comprehend the outputs of an AI, mitigates any worries on errors or biases that might occur in an AI prediction, builds Trust on decision-making by AI, satisfies legal and regulatory requirements such as GDPR, harness ethical implementation and fairness of AI-powered applications; finally helps detecting anomalous events driven by attacks. Explainable AI techniques have been separated into data, design, and post-hoc interpretability methods, and challenges for achieving explainability in the context of AI are data quality issues, algorithmic complexity high-level methods such as decision trees solving problems through analogue grey-box models under difficult conditions and negative consequences associated with over-interpretation. Moreover, the requirement of explainability is critical for compliance in AI environments which can help organizations comply with regulations like GDPR and inject transparency in AI results.




Fig. 1. Stakeholders in XAI

5.1. Importance of Explainability

It is well known that many artificial intelligence (AI) systems are being used to assist data-based decision-making and knowledge discovery in diverse domains. Yet a lack of knowledge about how their predictions are arrived at has limited the range of scenarios in which AI's can be deployed [19,21-22]. AI anomaly detection deals with unstructured data that is distinct from training data. Data lineage involves tracking the flow, changes, and possible challenges affecting the data. Compliance with explainable AI meets the requirements defined for deploying models within organizations and projects.

AI can assist in detecting security threats to AI systems and alert the appropriate personnel. Responsible AI governance encompasses the policies and practices for managing, monitoring, governing, and maintaining AI models. The increasing use of AI in areas such as education, healthcare, enterprise network security, financial fraud detection, self-driving cars, supply chain management, and customer analytics underscores the significance of "explainable AI," which offers the rationale behind prediction decisions, enhancing human trust in AI models.

5.2. Techniques for Explainable AI

Explainability techniques can be categorized by scope within a system, stage in the systems lifecycle, the time-frame in which the explanation is delivered and, relatedly, the "target" or intended audience. A useful taxonomy distinguishes pre-model, in-model and post-model explainability. Pre-model explainability methods seek to clarify input data and the mechanics of the model building process in order to help users understand the inner workings of the prediction model, for example by explaining why certain training data was selected, visualizing data inputs or even elucidating the technical steps in the model building phase. In-model explainability is achievable with some types of models – in particular white-box approaches such as rule-based

learning, use of decision trees or logistic regression – where the resulting model is sufficiently simple that some transparency can be achieved.

In contrast, black-box models such as deep neural networks do not admit an inherent explanation of logic. Post-model explainability methods seek to shed light on predictions made by such black-box models, at a local (i.e., a single prediction) or global level. Typical examples include: visualizations, explanations by example, feature relevance, and outcome explanations. Explainability is a cornerstone of many governance strategies; however, recent work has established a strong relationship between explainability and security. Moreover, the ongoing development of perturbation-based explainability methods has demonstrated the utility of such techniques as an anomaly detection approach in their own right.

5.3. Challenges in Achieving Explainability

Producing adequate and appropriate explanations in AI systems is not straightforward. Current techniques support only specific array of activities, roles, filtering, etc., resulting in only partial coverage of the framework requirements [11,23-25]. Transparent or interpretable methods are explainable by design; however, this characteristic often tends to negatively impact the predictive performance of the model. Furthermore, the explainability of a method is sometimes subjective; that is, what one particular stakeholder finds acceptable may not satisfy another. For example, customers use different types of explanations to justify and assign responsibility for AI decisions. Several explanations that accompany models are precise and accurate but not always understandable, given the cognitive limitations and mental models of human users.

6. Anomaly Detection in AI Systems

In the AI and machine learning landscapes, anomaly detection denotes the process wherein AI systems automatically classify and identify suspicious activities. The term "anomaly detection" arises from the nature of the data being analyzed; some types of data, such as medical images, bank transactions, and time series, can be used to detect abnormal situations. A prevailing category of AI security threats are the adversarial data attacks that can deceive or deceive production-trained AI models. Over the last two decades, multiple defense techniques have been presented to successfully build robust AI models through algorithmic adaptation.

Anomaly detection also involves recognizing abnormal operations and behaviors within AI systems themselves. AI auditing aims to detect AI-related issues and risks within production AI systems by examining system behaviors and accessing AI system operation logs. An AI system breach, by definition, undermines the goal and reliability of an AI system, which is to generate trustworthy outputs and intelligence. Anomaly detection models continuously monitor AI system conditions as an ongoing defensive procedure, analyzing AI system operational data using AI/ML techniques. Intrusion-detection techniques and existing anomaly detection models can be applied. Alerts generated by the anomaly detection model can serve as early-warning indicators to preempt potential AI security breaches.

6.1. Methods of Anomaly Detection

AI systems face an array of anomalies, ranging from attacks and system faults to policy violations and unexpected business events. While anomalies such as attacks and faults compromise the AI system's security, privacy, or safety, other types of anomalies underscore the importance of governance. Regardless of an anomaly's nature, prompt detection and explanation are crucial. Anomaly detection techniques identify instances or sequences of AI operation deviating from expected normal behavior, while anomalous event prediction methods forecast pairs of time points that may mark the ranging boundaries of potentially anomalous operation.

In AI systems governance, anomaly detection methods contribute across all roles and activities. They verify the adequate operational status of policy controls defined for implemented policies; penalize violators; derive model adjustment actions when necessary to ensure policy adherence; explain anomalies; and report the state of AI system operation to responsible parties.

6.2. Applications of Anomaly Detection

Anomaly detection can help harden AI systems against failure and data breaches. Examples illustrate how AI anomaly detection constraints result in specific use cases. A company trained a language model on their technical architecture documents. The service used production data to generate diagrams if the user input syntax was correct and contained only entities present in the architecture design documents. The data contained no sensitive information. The system failed when users asked for diagrams that required entity relationships independent of the company data, because the language model produced diagrams without those relationships. Another company used a Bayesian model to assign confidence scores for a sensitive classification use case. The service returned only high-confidence labels. However, a side-channel attack revealed sensitive data. Bayesian labels minimized leakage but did not mitigate the threat as a whole.

Anomaly detection in AI systems enables broader risk-detection frameworks. Detecting AI anomalies can prevent or reduce consequences of AI failures, data breaches, and misuse. Examine practices for anomaly detection. Languages for rule-based detection can express metrics governing AI in production. Operators can specify and evaluate AI constraints and detect possible anomalies. Practices for constraint programming in AI systems are outlined, exploring reasons for constraint creation and the impact of missing constraints.

6.3. Impact on Security and Governance

While developers and business managers of AI have security objectives, AI systems are also increasingly exploited as vectors of attack. Anomaly detection in input data, logic, and output both enables sophisticated new systems and mitigates threats and misuse. In addition, the explosion in AI system design components, styles, and vendors has created a “composition problem” in that it is difficult to specify and enforce control policies for each component of a complex AI system. Data lineage enables the tracking and tracing of AI data through its

lifecycle, so that governance policies can be developed and executed on the data being used at each point.

Security and governance objectives are heavily influenced by the trend toward higher levels of explainability for AI systems. Explainable AI (XAI) is the third pillar of responsible AI together with bias reduction and robustness. Techniques for achieving XAI provide real-time detection of AI attacks and result in governance policies that accelerate audit support and exploit lesson learning. Explainability mitigates liability threats by revealing bias and gaps in AI training and inference data. Regulatory requirement with XAI is very important in finance, health and public sector domain. Balancing between how well you are governed, secured, and explained will ultimately maximize productivity and minimize costs..

7. Data Lineage in AI Systems

In particular, tracing the input data that pass through an AI process can offer significant knowledge of the decision made by AI systems. The outputs of the AI system are typically influenced by two sources: the current input to which the system gives a reference and makes a decision, and the dataset on which it was trained. Each of the input and training data sets yields elements vital to establish the 'right' qualityliness and reliability of the decision.

Data lineage is a term familiar to anyone involved in data management, and it is concerned with exposing the original source of a dataset, as well as tracking how that dataset has changed or been moved around over time. Data Lineage in an AI Environment In the context of AI, data lineage is the ability to track and trace how both training and input data change over time. The capability to record and explain the data within AI processes as it changes over time is a key part of any governance framework, which in turn enables adherence to local regulations for use of AI or internal corporate policies.”.

7.1. Understanding Data Lineage

Data lineage is the life cycle of data: from where it came, how it moves through an organization and what becomes of it. Which is all cataloging, curating and visualizing flows of data. Data lineage is used to track the movement of data around an AI infrastructure, following it from its raw source location all the way through pre-processing, analysis and modelling to reports produced [26-28].. The information recorded in a data lineage can contain many details; for example, the time when the model or report was generated, the geographical origin of a dataset or model, the physical and organisational environment in which the AI system has been trained and executed, or the range of phenomena represented in the data.

The increasing reliance of businesses on AI systems has led managers and regulators to pay special attention to the explainability and governance of AI. Several regulations for the Explainable AI highlight the need to record and supervise the flows of data from and to both the

AI models and reports. Ensuring the availability of detailed data lineage information becomes then the first step towards the compliance with the ExAI regulations.

7.2. Tools for Tracking Data Lineage

Tools for tracking data lineage are software systems that maintain, trace, and manage the origin and path of data assets throughout an enterprise environment. Data lineage focuses on the flow of data on a more granular level than data provenance, yet it remains a closely related use case that involves tracking the exact source of data sets. In the context of AI, it is important to trace and explain the lineage of data fed into, flowing inside, and powering data-driven models. Defining explainability in terms of the data utilized by AI systems emerges as a promising approach. However, responsible AI practices demand more comprehensive governance that covers all components of an AI pipeline.

Data architectures and job scheduling are fundamental for constructing data pipelines, and related open source tools like Apache Nifi provide data lineage tracking. Many open source libraries, such as Pandera and OpenLineage, can also provide end-to-end data lineage for AI pipelines. Enterprise-level data warehouses often build data lineage frameworks upon OpenLineage. Additionally, data flow modes within AI systems can be uncovered through general dependency analysis—controlling data flow and usage is a key security practice associated with Artificial Intelligence of Things Compliance (AIoTC). Tools that allow querying data dependencies within AI pipelines can thus bolster overall security and governance.

7.3. Importance for Compliance and Governance

Explainability has become a critical governance requirement, with several jurisdictions mandating organizations capable of making automated decisions to provide meaningful information about how those decisions were made. Currently, emerging regulations include the European AI Act, the European GDPR, and the New York City Automated Decision Systems Law. Data lineage is a core practice for broader governance, where it is used to trace the flow of data from decision output back through the transformation steps and ultimately to the origin of the data. By combining data lineage, governance, and anomaly detection, organizations can address a powerful use case for compliance and governance with data.

The European AI Act: The future of AI applications and opportunities in Europe (and globally)
The European Commission has recently released the potential regulatory act on AI, known as a prop ASS. On 21st April 2021. The bill, which was released in 2021, examines how A.I. algorithms are created and deployed in detail. In particular, it categorizes AI uses based on the risk to rights and freedoms which they may entail and imposes specific obligations and requirements in relation to high-risk AI applications with a view to ensuring trustworthiness. The nyc automated decision systems law: It applies to all city agencies, including those in control of a contract with the city or created by government entities. It sets responsibilities for how automated decision systems (ADS) are used, including public policy applications, human

override, primary and secondary use.. Similarly, The US Federal Trade Commission has recommended that organizations employ explainable AI to avoid risks associated with non-explainability. It also suggests that using explainable AI can improve algorithmic auditability, assist with bias detection, and support ongoing compliance.

8. Compliance with Explainable AI

Explainability AI tools are critical to meeting governance principles as well as requirements for responsible and ethical AI. Explainability methods allow compliance by making transparent and revealing how an AI system comes to a decision. In areas where explainable AI is required by law, such as credit risk assessment, explainability is a regulatory compliance enabler that can be a part of AI governance [29].

Explainability can be selected as a governance focus for an AI workflow or a governance profile can specify a focus on explainability. Compliance patterns for explainability determine the explainability techniques and modes to be used periodically throughout the AI lifecycle. Explainability is commonly discussed in the context of a specific use case, for example, in a healthcare application or assessing credit risk. These use case-specific compliance patterns can create an audit trail with the activities in an AI lifecycle.

8.1. Regulatory Requirements

Explainability is an increasingly important subject as AI systems are regulated. This section looks at the regulatory environment for AI, the requirements of explainability in AI compliance, and considerations for the development of a compliance program.

The global regulatory landscape comprises many regulations and initiatives that impose varying forms of governance requirements on the development and implementation of AI systems. Although the regulations are for the most part different, they all contain similar fundamental principles that encourage principles of transparency, fairness, ethics, explicability, mitigation of implicit bias, protection of fundamental rights, the ability to challenge decisions, consideration of health and safety, data protection and privacy, and data quality and well-being. These principles, which are also among AI principles defined in the IEEE Standard, should be incorporated within an organization's overarching AI governance program, which will typically be developed at an organizational level and then applied per jurisdiction.

8.2. Best Practices for Compliance

Under various compliance regulations, organizations that use artificial intelligence technologies must demonstrate integral explainability during the development and implementation phases. Considerable thought and attention are therefore required when designing explainability capabilities for these organizations. Organizations can begin by establishing specific objectives for the Explainable AI (XAI) capabilities most pertinent to their AI solutions, noting that the objectives chosen directly influence the selection of suitable XAI techniques.

For example, if the primary concern is for security and governance staff to receive information that aids in the identification and analysis of anomalous AI behaviour, then they are likely to require different explanations than if their interest lies in ensuring the AI systems are fair, ethical, and compliant with internal and external regulations. In these latter cases, the explanation information needs to be suitable for business stakeholders (perhaps even the customers impacted by the decisions) rather than focusing solely on the technical description of the AI systems. The explainability framework proposed by Moore and Anderson presents one approach that organizations can adopt to integrate XAI objectives into the broader governance controls associated with AI deployment.

8.3. Case Studies on Compliance Issues

Several compliance issues encountered while developing real-world explainability frameworks call for attention. Adversarial scenarios must be considered, and countermeasures integrated into deployable systems. For example, a framework that detects anomalous data inputs entering AI models governing crime rates alleviates the risk of spurious predictions. Correlated anomaly scores also help in tracking the lineage of the suspicious data back to the responsible data source/operator, and thus assist in enforcing policies for trusted AI.

More compliance-centric questions arise in context of datasets and their distributions, such as "how trustworthy is the policy explanation for a particular request?" Regulations like GDPR mandate the use of only non-sensitive or anonymized data during model development. When a policy is requested for sensitive information, either within the training dataset or the data distributions underpinning the dataset, it is essential to verify whether the model predictions are indeed reliable for the specific policy generated [30]. Likewise, when a set of rules explaining a policy are generated for a sub-region of the data distribution, the trustworthiness of those rules should be confirmed. Existing tools like the Health Learner Angle (healthLENS) enable such compliance evaluations on an Adverse Impact analysis basis.

9. Integration of Security, Governance, and Explainability

AI is a complex of technologies that allows computers and software to imitate human intellectual abilities electronically. Depending on its functionality, it can be divided into reactive machines, limited memory, theory of mind, and self-conscious. The bedrock domains of AI are computer vision, speech recognition and synthesis, natural language processing, robotics, machine learning, and deep learning. The potential of AI can be deployed in the fields of education, military, finance, governments, healthcare, and so on. AI needs proper safeguards. Ignorance or lack of AI foresight can lead the world into a dystopian future.

The design and development of AI technologies has a severe impact on modern society in the form of security, governance, explainability, and ethics. Governing AI is concerned with how society decides what its development should look like in light of its broader societal implications for progress and development. AI policy is the practice of developing appropriate rules and

regulations for AI development and deployment. Anomaly detection, the identification of unexpected items or events in datasets. Data lineage is the ability to track where data comes from, what becomes of it and where it goes. Explainability is interested in the data explain and decision making levels. Compliance tackles the how, why, what, when and who of AI decision making.

9.1. Holistic Approach to AI Management

Rise of AI-generated content has created a need for tools that can control origin (provenance) and application rules (classification) ruling out the usage for training AI systems copyrighted or sensitive data. With the escalating use of AI around the world, governments around the world have published an array of AI governance policies to mitigate risks such as bias, opacity, job displacement, disinformation and more. AI Explainability (XAI) is an emerging field that aims to let AI decision be explainable to different stakeholders in order to mitigate risks in various domains.

Deep learning models are susceptible to different attacks, and the diverse behaviors of deep-learning models lead them into being challenging to be trusted in the context of security for their deployment. To solve these challenges, the AI community is concentrating on Anomaly Detection in AI Systems, Data Lineage in AI Systems and Compliant Explainable AI.

9.2. Interdependencies Among Security, Governance, and Explainability

The aforementioned three pillars are all about the discipline of making it work and successfully managing and running AI powered systems. Each of these pillars is designed to advance a particular aspect in order to equip AI practitioners with the requisite means to address the associated challenges. Hence, an ambitious approach that integrates all three perspectives will naturally contribute significantly to the advancement of AI as well.

10. Future Directions in AI Systems

Governance: It needs to be the hot topic in AI if researchers want to diminish blackbox models. Let's talk about roles and duties, governance policies, anomalies (e.g. anomaly detection in deep neural networks) etc. Besides, the explainability in AI requires attention to compliance considerations thus leading us to explore regulations, guidelines for ensuring compliance and how lineage can help with AAI compliance. It is important that integrating of these new topics in the AI literature can be achieved.

AI scrutiny has stepped up as use case deployments have pushed into virtually every industry. "AI technologies are a set of techniques, tools and methods that can enable you to do the sort of things that we would normally associate with human intelligence. In the last decade a fast-growing process was observed in this field due to deep learning methods, that make possible work directly with raw data from physical world and also to train highly complex systems while simulating human-brain mechanisms. Nowadays, AI can categorize objects, make sense of

images or text without relying on people at all and it's a break from Artificial Narrow Intelligence followers. The extreme level of automation and decision-making inherent in AI systems amplifies the importance of security; even minor faults, errors, or attacks can produce significant impact, jeopardizing AI projects and their alignment with business objectives.

10.1. Emerging Trends

With the growing intertwining of artificial intelligence (AI) in our daily lives, new trends emerge that raises hopes and concerns for the current governance levels. Recent work Reimers and Gurevych (2019) on large language models donors' ability to mimic truthfulness of human behavior, with LLMs pretrained using large amounts of natural language. LLMs are multimodal models, which have the ability to comprehend and generate text, they encode human-like fine-grained meaning of written language by binding with image or code. What's more, an active research direction concerns adapting base LLMs to the interests of users. Reinforcement learning with human feedback (RLHF) optimizes LLMs for human preferences; then models harness their language generation abilities to provide self-collected feedback in future iterations. But other than that there are some worrying signs: data breaches, attacks via prompts, malicious jailbreaks and politicking chatbot responses.

Future research must address trustworthy AI, focusing on governance, control, and alignment. Emerging trends in AI continue to facilitate control over the model's generated content. For example, textual inversion enables the embedding of synthesized subjects in text-to-image diffusion models. Finegrained captioning and grounding generate detailed captions with spatial relationships, while textual guidance establishes spatial properties for generated images. Moreover, developments in generative AI and explainable AI enhance model transparency and user understanding. Notwithstanding these innovations, comprehensive actions and guidance for trustworthy AI remain limited, and research community interest in governance and control risks diminishing if subsequent issues remain unaddressed.

10.2. Research Opportunities

Anomaly detection techniques are very useful in many applications, and several studies show that no one method or model can detect all kinds of anomalies. Although currently non-explainable AI techniques, such as deep learning models, achieve very accurate results, the main drawback is that such models hide the internal decision-making process and cannot explain why a sample is marked as an anomaly. This is why explainable AI techniques for anomaly detection should be further explored. In particular, methods should be developed to not only explain the reasons why a model detects anomalies but also explain why samples are classified as specific types of anomalies. This level of explanation would significantly assist practitioners and analysts during the denoising and cleansing process, helping them identify the root causes behind abnormal instances. Data lineage is one of the most important aspects surrounding AI governance. It offers detailed information about the flow of data, including details about the transformation or creation of data, data science models, and software items along pipelines. The

literature does not address data lineage and its use in AI governance, although it does describe how to trace data elements or explain different levels of lineage. Therefore, lineage information across different pipeline components should be maintained and used for governance, supervision, and control purposes. Supporting full tracing across pipelines would help the auditing process of AI systems by granting control, supervision, and anomaly detection of all pipeline elements. These AI security, governance and explainable AI research fields have tremendous potential to be deployed in real advanced AI systems.

10.3. Ethical Considerations

Ethical challenges are emerging in the AI era and rapid growth of artificial intelligence (AI) make these concerns even more urgent to govern for potential harm to prevent. Three major issues are the spread of bogus news and bias or disinformation — such as the notorious “deepfakes”: highly realistic images, sound, video or text generated by AI techniques that can be used to produce lies and defamation; the automation of warfare through lethal autonomous machines and other weapons systems; and how human work will be displaced by AI as employment is lost to automation [31]. Policy makers also must anticipate and address risks of this kind, along with more familiar dangers such as bias amplified by AI decision-making.

Yet in all these areas, AI, when used judiciously, can offer substantial advantages. For example, it can make a contribution to solving the problems caused by climate changes: new material or energy systems using AI eventually called Tfi (Technology for the Imagination), improve people’s health as a result of accurate medical diagnosis and therapy, human rights in terms of detecting contemporary slavery by processing visual information concerning hate risk attack. Regulating AI thus necessitates us to take a nuanced approach and actively monitor for potentially positive outcomes while regulating to moderate and prevent negative outcomes.

11. Conclusion

We believe that governance, security, and explainability of AI systems are key to shaping an enduring AI economy. With AI systems taking over decision making for various firms and organizations, the issue of governance becomes apparent. Re-imagine the types of harms that an AI system can incur in cases where it is biased, or vulnerable to manipulation and attack. Maintaining governance in AI systems demands the ability to query and interrogate the provided results and decisions of an AI system. They use their approach for anomaly detection to scrutinize an AI system in operation, determining if the outcomes are based largely on biases or adversarial attacks. The data lineage of a AI system, including queries on training data to train an model, the model that is trained and the resulting decision provides important aspects at different levels of the how its processed by an AI system. Anomalies can then be reported for critical parts as disturbances, etc. Moreover, compliance with explainable AI, such as the EU directive for AI and other proposed legislation, requires explainability of AI. Explainable AI techniques can

explain how an AI system came to a decision or result, thereby enabling scrutiny, fault detection, and correction of AI systems.

References:

- [1] HUSSAIN, Fatima; HUSSAIN, Rasheed; HOSSAIN, Ekram. Explainable artificial intelligence (XAI): An engineering perspective. arXiv preprint arXiv:2101.03613, 2021.
- [2] Koneti SB. Artificial Intelligence in Financial Systems: Digital Transformation, and Machine Learning Applications. Available at SSRN 5401202. 2025 Aug 12.
- [3] Muppala M. SQL Database Mastery: Relational Architectures, Optimization Techniques, and Cloud-Based Applications. Deep Science Publishing; 2025 Jul 27.
- [4] Muppala M. Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience. Digital Oceans: Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience| Deep Science Publishing. 2025 Jul 8.
- [5] Panda SP, Koneti SB, Muppala M. Benefits of Site Reliability Engineering (SRE) in Modern Technology Environments. Available at SSRN 5285768. 2025 May 1.
- [6] Koneti SB. Artificial intelligence Applications in Retail and Investment Banking: Personalization, Robo-Advisory and Behavioral Analytics. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025;4:72.
- [7] Muppala M. Architectures in relational databases: An analytical study of SQL-based data models and ACID principles. database.;2:4.
- [8] Bentahar J. A Survey on Explainable Artificial Intelligence for Network Cybersecurity. arXiv (Cornell University). 2023 Mar 7.
- [9] Gadde H. AI-Assisted Decision-Making in Database Normalization and Optimization. International Journal of Machine Learning Research in Cybersecurity and Artificial Intelligence. 2020;11(1):230-59.
- [10] Koneti SB. Algorithmic Trading and Quantitative Finance Strategies: High-Frequency Trading, Market Microstructure, and Risk Optimization Models. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025;4:17.
- [11] Panda SP. The Evolution and Defense Against Social Engineering and Phishing Attacks. International Journal of Science and Research (IJSR). 2025 Jan 1.
- [12] Mohapatra PS. Artificial Intelligence and Machine Learning for Test Engineers: Concepts in Software Quality Assurance. Intelligent Assurance: Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. 2025 Jul 27:17.
- [13] Koneti SB. Analysis, Predictive Analytics, and Macroeconomic. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025 Aug 12:90.
- [14] Panda S. Scalable Artificial Intelligence Systems: Cloud-Native, Edge-AI, MLOps, and Governance for Real-World Deployment. Deep Science Publishing; 2025 Jul 28.
- [15] Shivadekar S, Halem M, Yeah Y, Vibhute S. Edge AI cosmos blockchain distributed network for precise ablh detection. Multimedia tools and applications. 2024 Aug;83(27):69083-109
- [16] Reis J, Housley M. Fundamentals of data engineering. " O'Reilly Media, Inc."; 2022 Jun 22.
- [17] Ivanov SH, Webster C. Adoption of robots, artificial intelligence and service automation by travel, tourism and hospitality companies—a cost-benefit analysis. Artificial intelligence and service automation by travel, tourism and hospitality companies—a cost-benefit analysis. 2017.
- [18] Ramadhan M, Naseeb A. The cost benefit analysis of implementing photovoltaic solar system in the state of Kuwait. Renewable energy. 2011 Apr 1;36(4):1272-6.
- [19] Cordes JJ. Using cost-benefit analysis and social return on investment to evaluate the impact of social enterprise: Promises, implementation, and limitations. Evaluation and program planning. 2017 Oct 1;64:98-104.
- [20] Dykes PC, Curtin-Bowen M, Lipsitz S, Franz C, Adelman J, Adkison L, Bogaisky M, Carroll D, Carter E, Herlihy L, Lindros ME. Cost of inpatient falls and cost-benefit analysis of implementation of an evidence-based

- fall prevention program. InJAMA Health Forum 2023 Jan 6 (Vol. 4, No. 1, pp. e225125-e225125). American Medical Association.
- [21] Frank J. Artificial intelligence and intrusion detection: Current and future directions. InProceedings of the 17th national computer security conference 1994 Oct 11 (Vol. 10, pp. 1-12).
- [22] Wang F, Preininger A. AI in health: state of the art, challenges, and future directions. Yearbook of medical informatics. 2019 Aug;28(01):016-26.
- [23] Lu Y. Artificial intelligence: a survey on evolution, models, applications and future trends. Journal of management analytics. 2019 Jan 2;6(1):1-29.
- [24] Cordes JJ. Using cost-benefit analysis and social return on investment to evaluate the impact of social enterprise: Promises, implementation, and limitations. Evaluation and program planning. 2017 Oct 1;64:98-104.
- [25] Dykes PC, Curtin-Bowen M, Lipsitz S, Franz C, Adelman J, Adkison L, Bogaisky M, Carroll D, Carter E, Herlihy L, Lindros ME. Cost of inpatient falls and cost-benefit analysis of implementation of an evidence-based fall prevention program. InJAMA Health Forum 2023 Jan 6 (Vol. 4, No. 1, pp. e225125-e225125). American Medical Association.
- [26] Gadde H. AI-Enhanced Adaptive Resource Allocation in Cloud-Native Databases. Revista de Inteligencia Artificial en Medicina. 2022 Oct 18;13(1):443-70.
- [27] Konet SB. Microstructure, and Risk Optimization Models. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025 Aug 12:17.
- [28] Panda SP. Augmented and Virtual Reality in Intelligent Systems. Available at SSRN. 2021 Apr 16.
- Mohapatra PS. Artificial Intelligence-Driven Test Case Generation in Software Development. Intelligent Assurance: Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. 2025 Jul 27:38.
- [29] Eboigbe EO, Farayola OA, Olatoye FO, Nnabugwu OC, Daraojimba C. Business intelligence transformation through AI and data analytics. Engineering Science & Technology Journal. 2023 Nov 29;4(5):285-307.
- [30] Mohapatra PS. Intelligent Assurance Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. Deep Science Publishing. 2025; doi:10.70593/978-93-7185-046-9
- [31] Ojika FU, Owobu WO, Abieba OA, Esan OJ, Ubamadu BC, Daraojimba AI. Transforming cloud computing education: Leveraging AI and data science for enhanced access and collaboration in academic environments. Journal name and details missing. 2023 Jan

Chapter 7: Exploring Case Studies, Industry Implementations, and Future Research Directions in AI, Big Data, and Blockchain Technologies

1. Introduction

While the terms “Artificial Intelligence” (AI), “Big Data,” and “Blockchain” constitute three distinct disciplines, recent developments centered on the implementation of these technologies demonstrate their strong mutual interaction. AI in its many manifestations—Machine Learning, Natural Language Processing, and Computer Vision—creates the data that Big Data then analyzes to detect latent patterns. Other disciplines also build on Big Data’s outputs; Blockchain is one such area. Properly designed, deployed, and administered, Blockchain technology introduces an infrastructure extracted from the norms of computer science, cryptography, and economics. The purpose of the present contribution is precisely to provide case studies and industry implementations for these three disciplines, address the challenges encountered in their deployment, and finally indicate promising directions for future research.

Advanced technologies are frequently underestimated—even undervalued—because their ultimate implementation is sometimes partially lost in translation. Practical applications and real-world implementation often represent the final destination of a discovery, so the journey should be recorded for the sake of future research. Practitioners need to be convinced that the technology they are using or plan to implement—described in so many written contributions—should indeed be implemented. The first portfolio supports this approach. It offers a series of examples of AI, Big Data, and Blockchain implementation across different industries. Real-life AI application is shown in healthcare, finance, and retail; Big Data is depicted in telecommunications, manufacturing, and marketing analytics; and Blockchain is demonstrated in the supply chain, real estate, and voting. Grounded in a wide range of different applied contexts, these cases provide complementary insights into the uses of the technologies discussed.

2. Overview of AI Technologies

Artificial Intelligence (AI) refers to the simulation of human intelligence in machines, allowing them to perform tasks that typically require human intellect. It encompasses Planning, Learning,

Reasoning, Problem-solving, and Language Processing skills. The section explores the roles of three subfields in contemporary AI applications [1].

Machine Learning (ML) involves training algorithms on data to make predictions without explicit programming. Supervised, Unsupervised, and Reinforcement Learning are key types. Recent progress in Deep Learning, leveraging Deep Neural Networks (DNNs), has led to breakthroughs in image and speech recognition. Natural Language Processing (NLP) equips machines with the ability to understand and respond to human language, enabling sentiment analysis, chatbots, and language translation. Computer Vision enables machines to analyze visual inputs, facilitating applications like face recognition, image detection, and Optical Character Recognition (OCR).

2.1. Machine Learning

Machine learning (ML) is a branch of artificial intelligence (AI) that studies computer algorithms that improve automatically through experience and the use of data. These algorithms operate by building mathematical models from sample data—also known as training data—to make predictions [1-2]. Some ML algorithms make decisions and provide answers in response to data entered by users. Machine learning includes both supervised learning, where the algorithm is presented with example inputs and outputs in order to learn a general rule, and unsupervised learning, where only example inputs are given, and the model attempts to find patterns directly in the data.

ML algorithms are used in a wide variety of applications, such as speech recognition, computer vision, medical diagnosis, and forecasting. Machine learning can be applied anywhere where tasks are accomplished through analysis of data. ML applies the principles of statistics, data mining, and knowledge discovery in databases (KDD).

2.2. Natural Language Processing

Natural Language Processing (NLP) is a subfield of AI that studies the interaction between computers and human natural languages. The ultimate objective of NLP is to read, decipher, understand, and make sense of human languages in a manner that is valuable. NLP is significant because it enables computers to perform useful tasks, such as text translation, sentiment analysis, chatbot interactions, and market intelligence. Numerous NLP applications are found in fan fiction recommendation systems, fake news detection, question answerers, part of speech taggers, and space domain search engines.

Advanced Language Models (ALMs) refer to a collection of pre-trained text generation models, each boasting billions of input parameters. These models employ transformer architecture and attention mechanisms to improve learning accuracy and speed. ALMs exhibit the ability to generate creative content without explicit input from humans, writing lyrics, scripts, business plans, stories, and source code. The most popular ALMs are GPT-4, LaMBDA, BERT, LLaMa,

and ChatGPT, which are impacting society in areas such as content creation, language translation, question answering, and sentiment analysis.

2.3. Computer Vision

Machine vision and computer vision are closely related but distinct fields. Machine vision entails imaging-based automatic inspection and analysis and places particular emphasis on the control aspect of the system. Machine vision systems are often designed to replace human visual inspection during operations such as manufacturing, assembly, or printing.

On the other hand, computer vision is concerned with the theory behind artificial systems that extract information from images. Its overarching goal is to create an artificial system with visual sensing capabilities that parallel or surpass those of a human observer. In practical terms, it deals with the automatic construction of realistic or abstract description of objects, scenes, and events from images. Despite these differences, both disciplines share the use of computers to automate tasks that the human visual system can do.

3. Big Data Fundamentals

The term “Big Data” refers to the volume of structured and unstructured data being generated and stored, forming the foundation of data-driven enterprises and business processes [7]. Big Data analysis involves searching through vast amounts of data to identify hidden patterns or correlations, thereby enhancing organizational team efficiency. Big Data methods process large datasets that traditional databases cannot, enabling companies to predict trends and customer preferences [3-5].

The Internet of Things (IoT) significantly contributes to the generation of Big Data, coupled with decreasing storage costs and continuous improvements in data warehousing and analytics technologies. Over recent years, organizations have increasingly invested in Big Data analysis to optimize their business strategies. Unlike conventional data processing techniques, the new capabilities offered by Big Data analytics facilitate the processing of unstructured data at high velocities. Meantime, the traditional data processing steps of collection, transmission, storage, and analysis are no longer adequate for real-time Big Data requirements. The growing volume of digital data worldwide brings significant benefits but also introduces complex risks and hidden vulnerabilities.

3.1. Data Collection Techniques

Organizations collect Big Data from a variety of sources, including social media platforms such as Twitter, Facebook, and Instagram; web platforms like Google and YouTube; commercial sources that provide financial information; health-care sources that publish records and scan results; investment sources; weather repositories; multimedia data; and scientific projects page publications. The proprietary dataset of a company depends on the nature of its business. The

objective of data-collection techniques is to identify such sources and select the right sources on the basis of the business objective.

A number of techniques are implemented when collecting Big Data, especially information/web scraping, application programming interface (API) based techniques, scraping the Google trends website, web scraping on Medium, and Reddit web scraping. Web scraping techniques create a small program to collect information from a web page. These techniques extract specific information, like Liu numbers, followers, and total blogs from different websites.

3.2. Data Storage Solutions

Major technological advancements create the need for efficient data storage solutions, which help companies store, access, and manage data effectively. Data storage solutions range from traditional on-premise storage to modern cloud-based storage. Innovations in cloud computing have led to cloud storage solutions, enabling users to save and backup data on the cloud. Cloud service providers manage data centres and offer shared resources, allowing users to store data and access it whenever needed. In addition to cloud storage, companies also rely on data warehouses and data lakes, which can be hosted on-premise or on the cloud. Data warehouses use a structured format to store data, while data lakes can handle unstructured data. Important considerations in choosing a data storage solution include cost, scalability, data accessibility, and security.

Once data has been collected and stored, companies can analyse it to gain valuable insights. The goal of big data analytics is to provide actionable intelligence to decision-makers in a timely and cost-effective manner. Big data analytics solutions vary widely and include business intelligence, data mining, and data science. Business intelligence requires structured data as input and usually generates reports with rich visualisations for executives [6-8]. Data mining aims to discover patterns and trends, enabling organisations to support business objectives. Data science focuses on creating predictive models for automation and improving efficiency across different areas of a business.

3.3. Data Analysis Methods

Data analysis is basically the process of transforming useful data into valuable, insightful, and valuable information. Using a variety of analysis techniques, the data collected is subjected to an examination. It makes it possible to draw effective conclusions. Data analysis is the most difficult step in decision-making. It aids in evaluation, planning series of actions, controls that implementation.

Data analysis can be accomplished in a variety of ways. An effective technique is chosen based on the data collection. Data structure, goal of analysis, and many other elements. The explained techniques are appropriate in relation to decision-making issues in the telecom and

manufacturing industries. Decision tree and market basket analysis are two frequently employed techniques. Additionally, data analysis can fall into one of these three categories:

Descriptive Analysis: This analysis reconstructs and describes the history of analyzed data: what, when, where, how many, and how often something happened.

Diagnostic Analysis: The analysis takes a deeper dive into the data to figure out why something happened.

Predictive Analysis: Although it requires a ton of data, an extensive predictive analysis sheds light on what will probably happen in the future.

4. Blockchain Technology Overview

Bitcoin, a cryptocurrency in which transactions are verified and records maintained by a decentralized system using cryptography rather than by a centralized authority, was proposed in 2008 by Satoshi Nakamoto. Subsequently, Nakamoto also proposed a consensus mechanism called Proof-of-Work (PoW). PoW has also been applied on several other cryptocurrency protocols like Ethereum, Litecoin, and Bitcoin Cash. During the last decade or so, Nakamoto's Bitcoin blockchain was extended to incorporate a feature called smart contract, and the new blockchain was called Ethereum [9]. A smart contract contains a specific set of instructions that trigger contract terms and conditions.

Several consensus mechanisms have been proposed in the new blockchain. Apart from PoW, others include Proof-of-Stake (PoS), Delegated-(D)PoS, Proof-of-Elapsed Time (PoET), Raft, Practical Byzantine Fault Tolerance (PBFT), and Federated Byzantine Agreement (FBA). PoS is used on board cryptocurrencies like Cardano, Nxt, and Peercoin. DPoS, the delegated version of PoS, is used on Steem and BitShares. The new version of the PoW, Litecoin, uses the proof-of-capacity (PoC), which is based on mining cluster capacity. Quorum Bitcoin employs Raft and Istanbul BFT. The SEC blockchain uses PBFT, and Ripple employs FBA.

4.1. Decentralization

The most prominent feature of blockchain technology is decentralized consensus. Decentralized storage Data is spread out in a decentralized way and stored on multiple devices of the participants that belong to different parts of the world. This decentralized structure is intended to reduce the ability for any one entity to dictate rate data, and in thus doing accelerates fair order execution across the board by avoiding acting parties tampering said data. The lack of a central authority implies that all members have the same standing in the network. Decentralized systems are fault-tolerant as well; if one computer goes down, the others chug on. Also, with decentralization, there is no single point of failure [7,9-10].

Other benefits of decentralized systems include greater security, self-executing smart contracts, and the lack of reliance on trusted third parties. Blockchain is preferred over others because it

can provide the additional benefits of immutable database, distributed consensus and enhanced security even though there are known weaknesses that can be exploited. This technology forms the foundation of such categories as crypto-currency and digital currency. Smart contracts implemented by blockchain enable authorized transactions automatically, and its token creation capability allows enterprises to design and launch their own tokens for various applications.

4.2. Smart Contracts

Smart contracts are computer programs or digital protocols which manage automatically, control and enforce the performance of specific kind of agreements (or any other conditional statements) on a blockchain. Leveraging blockchain's unique properties — namely, decentralization, autonomy, transparency, immutability and trust — they allow for parties in a transaction to be able to forge trust without the need for an intermediary. After 2013, smart contracts are used in domains including finance, healthcare, government, crisis management and IoT.

Smart contracts on a blockchain involve three parts: the decentralized ledger, the infrastructure behind it (the ledgers and support systems etc., which must all be running seamlessly) and lastly the business rules that manage industries and enterprises. The Ethereum ecosystem and its Solidity programming language is the most popular for writing contracts. It is also necessary to implement the business rules, in the decentralized context, into Solidity language in order to develop the smart contract. Consensus mechanism guarantees that business rules are correctly enforced despite of the complexity and possible mischief of the underlying network gridding. A summary of smart contract applications existing in these different industrial reviewing sectors is presented in the table.

4.3. Consensus Mechanisms

Consensus algorithms including PoW, PoS, BFT and PBFT are essential part of a blockchain system [1,11-14]. They are a processing module which is responsible for transactions, record-keeping, credit and asset control in the economy. Fault tolerance, Sybil attack resistance, and double-spending attacks prevention are some of the main properties we hope to achieve with a consensus mechanism.

PoW demands that all miners "race" to solve a mathematical challenge and the first one to find an answer is entitled to create the next block in a given blockchain. PBFT and BFT adopt primary replica selection for achieving consensus for permissioned blockchains. PoS does not heavily rely on computational power; rather, the probability of a node being elected for block generation increases with the amount of digital-assets it possesses.

5. Case Studies in AI

Artificial Intelligence (AI) is a domain of Computer Science—offering an application that enables machines to intelligently act like humans—training and building Computers and software for performing various operations such as recognition, learning, reasoning, and problem

solving. A broad branch of AI is Machine Learning (ML), which performs pattern recognition through detailed data analysis. Natural Language Processing (NLP) is employed for training the Machine with languages trained by humans, so it can recognize and understand human-written or spoken text. Furthermore, Computer Vision (CV) is used to train a Computer to view the enactment of an environment [13,15-17].

AI is more recognized for its distribution than any other system and has gradually spread to every business and market. Different companies are performing the implementation of artificial intelligence technology to fulfill business needs and their requirements to handle and detect several things in daily life. AI Technology in different service industries, such as Healthcare, Finance, and Retail is also being used for several purposes; for example, AI in Healthcare helps facilitate the early prediction and detection of cancer and also helps in the detection of anomalies; AI in Finance is widely used by Banks in several operations such as customer care, operational risk management, fraud detection, and several other operations; AI in Retail is being used specifically for customer behavior analysis.

5.1. Healthcare Applications

The healthcare industry has always been at the forefront of the AI revolution. The rapid adoption of AI-enabled systems is evident, with the online AI in Healthcare Market Growth Report 2022, which predicted that AI in healthcare will continue to be a lucrative business, revealing the steep increase in the market size—from US \$8.23 billion in 2020 to a projected US \$120.2 billion by 2028, at an extraordinary CAGR of 45.8% (Fortune Business Insights 2022). Some specific instances of those use cases that are already growing and in research include using AI for emergency response (TimelySense2021), patient data, patient risk and safety, patient engagement, diagnosing conditions, aftercare. Healthcare organizations can harness the power of AI to transform and enhance the quality of care and improve relationships and communication with patients and clients. This will soon contribute to the aim of Revolutionizing Human Health through Artificial Intelligence.

AI systems are the key to effective and efficient delivery of diverse health and wellness services. Voice recognition algorithms help doctors in setting reminders and taking notes on the go, enabling better coordination of care, improved staff productivity, better patient monitoring, and faster response in emergency situations and personalized service delivery. Varieties of algorithms generate hyper-personalized experiences for patients and customers and simulate high-level thinking and many human cognitive functions. During the COVID-19 pandemic, AI technologies deployed across the globe to make response and mitigation more effective through discovery of new drugs and treatment models, creating a distributed ledger of diagnostic reports, providing clinical decision support to doctors, identifying high-risk patients, and averting disease transmission through scanning, reporting, and analysing incoming data and information [18-20].

5.2. Financial Services

Financial institutions are turning to AI to develop advanced trading algorithms and predictive models capable of extracting patterns in historical data to prevent fraud, conduct risk assessments, forecast market trends, and manage hedging strategies. These functions lead to remarkable improvements in efficiency, enabling more accurate predictions, faster execution, and substantial cost reductions.

Additionally, AI is being applied in retail banking to enhance customer experience. Chatbots, for instance, offer instant support, affordability and 24/7 availability. These sophisticated virtual assistants communicate with customers through voice and/ or text, in real time (over the phone and chat). Through NLP, users can receive help for common services or the newly revised tax and law knowledge. Personalize information to help customers The understanding of customers contexts and moods via Deep Learning, bedeudet Personal Advice for the best banking products.

5.3. Retail Innovations

Data insights and algorithms are to be had, and being used more, in the retail industry decision-making process. Use cases run the gamut, from prediction and anomaly detection to sentiment analysis and customer profiling. Tactically leveraging AI-powered solutions helps retailers derive drive significant business value by evolving according to the constantly changing preferences of their shoppers.

The ability of AI to pick out, categorize and interpret images — and even facial expressions -- is already having an increasing use case in retail. But in addition to enhancing an in-store experience – including but not limited to brand-based deployments and customer interactions – the role of those tools have transformed how returns and complaints are managed in e-commerce, driving substantial improvements in customer satisfaction [19,21-22]. What's more, AI ensures that customers are engaged with personally relevant interactions, engenders loyalty, and supports cross-seeling and up-selling by effectively identifying well-defined customer segments with offers tailored to those needs. The future with AI technologies is even more exciting – leveraging large volumes of data to get a 360 degree view of all aspects of the business, front or back-end..

6. Industry Implementations of Big Data

Organizations collect terabytes of data every minute from customers, partners, and the business environment. It is essential to store, process, organize, and analyze this huge amount of data to extract useful information, which is vital for making good business decisions. The telecom industry is a pioneer in using Big Data for storing and processing millions of transactions daily. The objective is to realize, identify, leverage, and forecast customer needs.

Analyzing industrial manufacturing Big Data has the potential to improve efficiency, decrease costs, and predict and prevent downtime. Supporting data scientists and stakeholders in gaining

high-quality insights into manufacturing processes requires a solid Big Data infrastructure. Industry 4.0 functions in marketing and sales benefit from Big Data. Market strategists can develop modern and innovative products to gain complete customer satisfaction by utilizing Big Data techniques such as data warehousing, data mining, and Customer Relationship Management Software. The rapid growth of online business draws the attention of both customers and companies.

6.1. Telecommunications

With new data generation mechanisms, such as the Internet of Things (IoT), Big Data in the telecommunications industry has grown exponentially. The deployment of smart meters in power grid systems produces massive data, referred to as smart meter data (SMD), which is difficult to recognize, analyze, utilize, and protect. The diverse formats of generated data include the documents and files generated by smart meters themselves, remote control commands of smart meters, and multimedia data (schematic diagrams and photos of meters, multimedia messages containing customer information). A method for storing these diverse types of SMD and related security safeguards has been proposed.

The application of big data in operations, administration, and maintenance (OA&M) in the telecommunications sector has been proposed. Traditional OA&M systems usually handle data generated at each stage sequentially, which is time-consuming, operator-dependent, and incomplete. Big-data-based OA&M systems analyze and process data generated in each stage comprehensively, enabling automatic generation of the optimal plan that meets operations requirements through automatic analysis and judgment, supporting a comprehensive forecast of the entire life cycle of network elements. Several other applications of big data in OA&M have been proposed. Customer analytics has a key role in retaining existing customers and acquiring new ones [11,23-25]. Big-data analytics is widely used to detect the customer's profile, behavior, and frauds, and thereby achieve enhanced customer satisfaction. Campaign management analyses a customer's response toward a specific campaign, and telecommunication companies can identify whether the response is positive or negative; on the basis of this, necessary and suitable actions are taken.

6.2. Manufacturing

Manufacturing is an area that benefits significantly from Big Data analysis. Besides optimizing production, Big Data supports supply chain management, product quality, customer feedback, and data collected from sensors. Consistent and rapid production line operation necessitates constant monitoring. Early identification of potential production errors ensures a fully operational factory and optimal resources. Enabling predictive maintenance in factories and providing operators and engineers with instructions can reduce breakdown times and extend machinery lifespan. Advanced defect detection processes, superior to human inspection, improve the quality of manufactured products.

A novel Big Data logistics-based approach ensures transparency, efficiency, and agility of manufacturing processes throughout the plant's supply chain. A Big Data framework for the real-time condition analysis of automated manufacturing systems helps decision-makers enhance production processes and implement corrective actions. Big Data analytics in automotive manufacturing improves engine speed, vehicle idling, application prioritization, and job scheduling. Optimizing operations planning through Big Data analysis involves collecting and processing datasets related to operations, plant layouts, and machinery status. The big data repository integrates Big Data and Internet of Things (IoT) concepts to provide continuous real-time information on production and machinery health. Predictive analytics enhances just-in-time inventory strategies by forecasting future demand, replenishment requirements, and warehouse stock levels. Data from the production process is employed to analyze, improve, and shape planned product maintenance activities and schedules. In the end, a BD framework able to offer that support to manufacturers in terms of analysis of usage patterns on which their products are used collects such customer-related data while being utilized by them whenever required for this use-case.

6.3. Marketing Analytics

Marketing is now heavily infused with big data technology, as companies constantly collect and analyze a metric ton of consumer data in order to inform effective business/marketing strategies. Marketing analysis results in more effective customer segmentation and profiling, the ability to design targeted communications, increased sales rates, higher customer retention and so much more. With Big Data analytics, marketers can tailor campaigns and offers to suit the needs of individual customers, creating a personal experience that customers appreciate and giving the company a distinct competitive advantage.

Predictive analytics enable marketers to forecast customer demand, avoid stock-outs, reduce costs, and fulfil customers' needs in a more timely fashion. Marketing practitioners can also use sentiment analysis to assess the proportion of positive and negative feedback on their products during new launches, thereby optimizing the marketing mix [26-28]. Furthermore, marketers are increasingly employing clickstream analysis to identify their customers' motives. The harvesting of information agents' clickstream trails, recording their browsing behaviour and examining their favoured websites, assists marketers in the understanding of consumers' search and purchase behaviours. Specific research studies of clickstream trails have enabled the development of predictive models of consumer behaviour parameters including click-through rates, probability of purchase forecasting, consumer segmentation, and frequency of visits, among others.

7. Blockchain in Various Industries

Blockchain Models, Smart Contract Models, and Consensus Mechanism Models, along with their applications in real-world scenarios, illustrate the transformative potential of decentralized technologies. Leveraging these case studies can demonstrate the practical advantages in different

sectors, facilitating an understanding of their wide applications. Supply Chain Management Case Studies elucidate the use of Blockchain to enhance provenance tracking and inventory management. Real Estate lends itself to improvements in property transactions and title searches. Meanwhile, the Voting domain ensures voter verification and protects ballot legitimacy. Experiences in these areas provide numerous avenues for further exploration and development.

Despite the compelling advantages indicated by the various studies, Blockchain faces its share of challenges, particularly in terms of Scalable Storage, Economic Model construction, Security Systems, and Regulatory Frameworks. Therefore, Future Research Directions should include the deployment of lightweight storage models, the establishment of incentive-compatible stakeholder games, the design of advanced permission mechanisms for blockchains, and the development of regulatory systems that promote healthy ecosystem growth. Additional meaningful research areas encompass Blockchain for Social Good, the Blockchain-Internet of Things intersection, and Cross-chain Technology. Furthermore, the synergy between Artificial Intelligence, Blockchain, and Big Data is of great interest. In this regard, the potential of blockchain is explored both as an independent research agenda and as a complementary technology that enables and enhances the capabilities of the other two technologies.

7.1. Supply Chain Management

Today's users are very demanding and do not want to wait. Before buying a product (for example, a car), they search for information about it on the Internet and evaluate comments made by other users, either in the online store of the company or on social networks such as Facebook or Twitter. The question that must arise in the mind of any user before buying a product is, "How was it made?", "Is it original?", "Is this product manufactured by a reliable company?"

Blockchain could be the answer to these questions. A blockchain-based system could provide users with a safe way to understand how their product has been manufactured and whether it is original [29-32]. The fact that every transaction made by companies can be registered and checked by users would offer them transparency about their purchase, adding value to the product and, more important, to the company delivering it. SupChain is a blockchain framework applied to supply chain management systems using a decentralised platform that integrates the Ethereum public blockchain and an internal private blockchain with smart contracts transmitted in a digitalised manner.

7.2. Real Estate Transactions

Several efforts have been made to address the challenges of the real estate industry. Real estate transactions need to collect, organize, and analyze a massive number of documents and transaction records that are normally stored by banks and real estate agents. Blockchain helps collect, track, and record real estate transactions and property information securely and transparently. Concurrently, applications of AI technologies to the real estate industry have also started. One example is Umbra, which leverages a collection of economic and social factors,

such as demographics and life expectancy. The machine learning model captures a property's potential and generates an annual risk and value score. A risk and value score report is provided in a business intelligence format. Real estate investors can make intelligent decisions based on the property risk and value scores.

Real estate investors highly rely on property history information. To make more informed investment decisions, it is essential to dig out the potential investment risks, such as property foreclosures, previous damages, and thefts, from the properties' historical records. Recent advances in blockchain have enabled the property record network to be fully developed with the capacity to gain the public's trust because of the immutability of blockchain records.

7.3. Voting Systems

Prior to the 2016 United States presidential election, presidential candidate Vladimir Putin suggested the implementation of blockchain voting for the election. However, the need for electoral security, voter privacy, and other potential problems, such as "digital divides" and a lack of a paper ballot for recounts, prompted the Russian government to decline the proposal. The Russian government later promoted other usages of blockchain technology and eventually implemented a blockchain voting system for the 2020 Russian constitutional referendum.

Beyond Russia, many countries have tested or implemented blockchain voting systems, including the United States, Switzerland, Estonia, Ukraine, Australia, Canada, and the United Arab Emirates. The Emirates Blockchain Strategy 2021 aims to move 50% of government transactions to blockchain by 2021. In Australia, the city of Fremantle in Western Australia enabled online voting with a blockchain-based voting system, in collaboration between Horizon State and Voatz. Some voting systems support not only voting but also identity verification and voting result detection. Societies are exploring the possibility of implementing national election voting and referendum voting using blockchain voting systems; for instance, Switzerland allowed citizens to test blockchain e-voting during the 2018 political voting season in the city of Zug.

8. Challenges in AI Implementation

The groundbreaking capabilities demonstrated by AI in sectors such as healthcare and climate change have been documented. The most successful AI systems thrive on intelligence produced by large datasets. As AI grows, concerns become more urgent about the security of sensitive data, protection of privacy, ethics, and possible biases. These cutting-edge issues should be addressed through public education and the development of improved data protection legislation.

Example applications include disease diagnosis and treatment recommendations through analysis of medical images, lab data, and genetic information. Aiding COVID-19 diagnosis and prognosis through evaluation of imaging studies, cell count and enzyme levels, age, sex—women, chickenpox—linked to databases of signs and symptoms, can help diagnose new diseases,

evaluating signs, and symptoms of patients and provide diagnostic suspicion and decision support systems. AI naturally complements data-intensive fields such as cyberspace and big data. Areas of application in the field of cybersecurity include the development of IDS in both physical and cyber-cognitive domains, malware, anomaly threat detection and classification as well as situational awareness, and alert prioritization.

8.1. Ethical Considerations

Ethical challenges are prevalent in Artificial Intelligence (AI) implementation due to its ability to extract in-depth knowledge about individuals from large datasets. Potential issues surround concerns on privacy, automated decision making, transparency, and bias. Individuals retain a right to privacy, and the use of AI must indirectly overpower the right to privacy in order to achieve social or individual benefits. However, such issues can be managed with legal and ethical measures on individual and societal levels by adapting data sharing and processing policies; however, the discussion on the appropriate level required for the management of these problems continues. That is the reason that the most current analyses admit explicit ethical agreements so that individuals and states are aware of the generated compromises and therefore of the acceptance of the consequences."

Data bias can considerably impact the accuracy of AI models. It is essential to carefully select the source data to ensure it is neutral and free from gender, regional, or age-related biases. The quality of source data largely determines the quality of the model produced. Ethical issues also arise when robotic agents replace humans. Despite the removal of menial and repetitive tasks from the job market, the displacement of personnel from existing positions raises ethical considerations. Transparency of decision-making processes, especially in sensitive areas such as insurance and private credit scoring, is also crucial in the acceptance of AI technology. Recent developments in explainable AI also contribute to overcoming this challenge. Towards a multi-modal framework, it enables the integration of complementary information beyond pure textual and linguistic metadata, such as temporal, spatial, and emotional metadata. The literature closely linked to the establishment of corpora of AIMFs consists of two complementary research lines [31,33-35]. The first involves interdisciplinary areas such as digital humanities, social sciences, and political sciences; the second is grounded in the work related to the design of resources for automatic multilingual processing of social media texts in general.

8.2. Data Privacy Issues

Artificial Intelligence Models use vast amounts of personal information during training, resulting in data privacy concerns. These AI Models require extensive amounts of training data to perform their acceptable role. Usually, the training dataset consists of users' personal information such as ID number, phone number, bank account number, email address, location information, and other sensitive or confidential information. When the trained model is directly exposed to the external environment, privacy problems arise. Therefore, research on the training process without disclosing private training data has been actively conducted in various AI fields.

The use of data degrades privacy. People get thickly involved in various AI services and require their data to be safe from privacy issues. For example, data sales or data hacking, etc., are threatening the privacy of data. Buyers of the data take advantage of it in unfair ways, such as discrimination and increased risk in insurance companies, insurance applications, bank credits and job advertisements. At this stage, Methodologies connected with Big Data raise user interest and alert in protecting their private data from misuse. Therefore, there is an interest in using Big Data to create machine learning models without violating the training data privacy. Privacy concerns for both individuals and businesses have been addressed explicitly in an efficient manner.

8.3. Algorithmic Bias

This subsection showcases how the tech industry tackles algorithmic bias in AI. Firstly, it highlights Google's approach to making trade-offs between bias and accuracy. Secondly, it discusses Facebook's problem of biased classifiers and explains why Facebook cannot sidestep this issue. The following content is based on the interview with Kate Crawford.

Because Google has fast access to extremely large dataset, it produces better classifier. Better classifiers tend to be less biased for race, gender and thus it helps reduce algorithmic bias. However, as Kate points out, algorithmic bias not always goes hand in hand with accuracy. According to Kate, Google is aware of that and trying to find a way to make trade-offs between bias and accuracy. With better classification accuracy, the algorithm becomes more "equal" with less bias embedded. However, when the topic comes to Facebook, it is about bias reduction. Facebook uses facial recognition to give people a better user experience. Yet bad classifiers tend to be more biased in terms of race and gender because most of the training data is from the United States with predominance of white people. The biased face algorithm leads to cognisant of the unfair treatment of discrimination to people from developing countries. As a result, Facebook needs to debug the facial recognition but in the meantime it cannot not use the algorithm.

9. Big Data Challenges

Big Data has been implemented throughout many different areas. Even simple industries, such as telecommunications and manufacturing, can benefit largely. In manufacturing, Big Data can be used to forecast when equipment is likely to wear and fail, allowing for just-in-time maintenance, thereby reducing downtime, lowering maintenance costs, and preventing missed delivery deadlines. Telecommunications companies utilize Big Data to collect information about their offerings and customers, enabling the creation of products tailored to customer needs while significantly reducing provider churn rate [36-38].

At the other end of the spectrum is marketing analytics, an aspect of Big Data that involves collecting and analyzing customer feedback directed to market products and services in a way that promotes brand loyalty and entices a larger market share. The collection of these data points

enables an organization to identify who is buying their products, determine the most effective channels for marketing to that audience, and enhance customer experience. Despite these practical applications, businesses may be underutilizing, mismanaging, or missing key aspects of Big Data. Challenges such as poor data quality, real-time data collection, integration with legacy systems, scalability, and the reuse of data across multiple departments or business functions remain impediments to Big Data's full potential.

9.1. Data Quality Management

Research into data quality management in large data has been presented. Google BigQuery was used to test big data algorithms and assess the quality of big data. This service allows users to perform SQL-like queries on vast amounts of data. Quality assessment is performed in the solution stages: data gathering, disorganized patterns, viewpoint, preparation for analysis and processing, modeling, and transformation.

The description of stored data must be clear and easy to understand. Milanovic et al. investigated the method of establishing metadata, whereas Zhang et al. adapted the OLAP method in big data. Additionally, the semantic method offers numerous possibilities because it describes the content of stored data, their properties, and the inferring ability.

9.2. Scalability Issues

Addressing scalability in blockchain technology is essential for future innovation and widespread adoption across various sectors. The first concern is that increased block size and faster block interval times can challenge the capacity of hand-held devices to serve as full nodes. If a full node requires too much storage space, it suffers from limited distribution, resulting in the loss of decentralized advantages. Moreover, it can lead to the centralization of mining powers and the weakening of network security. Faster block times can amplify the proof of processes, hindering adaptation to the Internet of Things and the latencies of smart contracts. The second aspect relates to the design choice of different consensus algorithms in different blockchain architectures. The selection from proof of work, proof of stake, proof of storage, and other algorithms is typically related to the demands for throughput, response time, and finality. Furthermore, when the health care and IoT industries merge with blockchain, blockchain applications involve multiple communities, such as users, businesses, and service providers. To address the issue, the innovative cliques—namely, data sharding, transaction sharding, state sharding, and functional sharding—have been proposed. The third aspect pertains to the increase in scalability issues associated with the growth of the underlying block size. The size of the blockchain, which is currently approximately 300 GB, continues to grow every day, making it impractical to download and verify at regular intervals. A similar issue arises due to inefficient consensus algorithms. Although transactions are adequately validated by high-powered devices, the process remains inefficient for mobile nodes with limited computation power and battery life.

9.3. Integration with Legacy Systems

Despite challenges with legacy systems—such as data growth, tightly coupled technologies, security vulnerabilities, and process inefficiencies—large organizations continue to generate comprehensive Big Data. Many remain reluctant to store new data in cloud environments, opting instead to place a Big Data layer on top of their mainframe system, effectively extending automatic mainframe support with a NoSQL environment that stores Big Data outside the core database. Legacy systems can receive mainframe data via MQ or FTP and send it as files through a file system. However, this approach introduces additional files that require specific data management and maintenance. Moreover, an FTP system stores Big Data outside the firewall, further complicating security requirements.

Data quality issues—such as missing values, outliers, and imbalanced data—further complicate the ability to generate accurate insights. Incorporating the entire data pool might not be feasible for real-time decision-making and action. These problems are exacerbated in cumulative analyses, leading to progressively erroneous outcomes. A potential solution involves automatic data cleansing, which identifies the nature of inaccuracies and performs necessary formatting changes, including the removal of inappropriate data elements. Solutions might incorporate basic statistical functions: measures like mean, median, and mode address missing values; percentile calculations correct outliers; and implementing boundary points with upper and lower limits converts imbalanced values into correct forms. The classification and cleaning are based on specific business rules.

10. Blockchain Challenges

Blockchain has recently attracted lots of attention among academicians, industry experts, and governments because of the inherent features of blockchain technology, such as transparency, resilience, trust among dynamically forming consortium components, immutability, and decouple trust from a centralized third party. However, despite these advantages, the blockchain technology is yet to reach the next level of practical implementations and market boom. Recent research has identified some of the outstanding issues with the blockchain technology, including scalability, regulations, cyber-security, privacy and data protection, cybercrime, governance, and interoperability.

This paragraph continues with specific scholarly challenges. Scaling of blockchain is one of the major issues as it prevents a high throughput rate for processing transactions. For example, wang2021evaluating have pointed out several cost barriers in deploying cryptocurrencies by using blockchain technology, and argued that only 20% of the countries are favorable in their environment for adopting cryptocurrency, which reduces the potential of revenue generation through the transaction fees. As a result, the supply and demand gap tends to remain wide. Moreover, with growing users and transaction rates, legacy cryptocurrencies lack the capability to beat the traditional payment providers such as Visa or Mastercard in terms of speed and

efficiency. The study by eenigha2020regulatory have suggested that there should be some regulations for dealing with cryptocurrencies, otherwise, it may lead to dirty money in the system. In addition, cyber-security is another important concern, especially in the use of private key for access and retrieval of bitcoin wallet. Hence, according to issam2020blockchain there is a lack of appropriate preparedness toward malware that affects the existing cryptosystem. User privacy and anonymize quality of cryptocurrencies also raises doubts. tatarnar2018cryptocurrency have suggested that in Europe, the General Data Protection Regulation (GDPR) may conflict with the core characteristics of blockchain technology. Besides, as the blockchain technology is still in nascent stage, it has started attracting the criminals by maliciously utilizing it for money for ransom.

The preceding paragraph continues to supply detail from the literature. As argued by lee2019blockchain, the absence of governance framework in blockchain and lack of government supervision may provide opportunities to misuse blockchain for illegal activities, as criminals can easily exchange money from one cryptocurrency to another. Moreover, banks face serious challenges due to the presence of cryptocurrencies as they provide service without any regulations. A famous rumor says that Nigeria has experienced the second largest drop in bank account holding, and the highest increase of mobile wallet users in the world, because about 15 million Nigerians trade cryptocurrencies on various websites. Again, the degree of interoperability among various blockchain platforms is also crucial for its effective operation, and proper implementation of these issues could lead to exploring a new arena in the field of blockchain technology.

Research directions are also presented. Within the emerging domain of Blockchain Technology, a domain that has recently attracted a lot of attention due to its intrinsic characteristics such as transparency, resilience, trust among dynamically forming consortium members, immutability, and the decoupling of trust from a centralized third party, several open issues have been identified. These include scalability, regulations, cyber-security, privacy and data protection, cybercrime, governance, and interoperability. In line with the challenges previously highlighted, promising future research opportunities exist for addressing these areas, all of which are crucial for advancing the field towards broader practical adoption and market success.

10.1. Scalability

Blockchain technology offers a decentralized and distributed ecosystem where various parties can transact online without a central authority. The integrity, transparency, immutability, privacy, and security of data within the blockchain network are maintained by utilizing multiple technologies in the underlying layers, including cryptography, consensus mechanisms, and computer networks [1,39-40]. However, blockchain systems face scalability challenges similar to other distributed systems. Such challenges are especially pronounced when the number of techniques and services implemented over the blockchain grows at a high rate. The limited processing speed of the system results in delays in validating transactions. Therefore, before

implementing any blockchain service or technique, it is advisable to consider the scalability aspects to avoid performance-level challenges.

Currently, various solutions—for instance, Lightning Network, sidechains, Plasma, Raiden Network, sharding, and off-chain state channels for Bitcoin and Ethereum-based blockchain systems—have been proposed and implemented to alleviate scalability issues. Although these solutions can effectively tackle blockchain scalability, potential issues such as latency, transaction costs, and multi-hop payments may arise. Moreover, not all solutions are suitable for scalability issues in specific blockchain platforms. For example, Plasma and sharding have been primarily implemented for Ethereum. Consequently, a well-defined scalability solution for scalable blockchain ecosystems remains an open research problem.

10.2. Regulatory Compliance

Maintaining regulatory compliance is often challenging when deploying applications in a public cloud infrastructure. Issues related to customer data ownership can arise because, physically, the data may be stored in any part of the world. This situation creates security concerns because the applicable law depends on the location of the data. Certain regulations require that the data cannot leave the country or its borders. In some scenarios, regulations specify the country of incorporation for the cloud provider.

Currently, no single public cloud provider can meet the above-mentioned compliance requirements alone; however, a federation of public cloud providers can do so. In such a case, the customer data is stored in a public cloud depending upon the conditions specified by the governing regulations for that particular industry, although these conditions can be satisfied one after another in a single cloud setup or simultaneously in a multi-cloud setup, depending upon the data-storage requirements. Public cloud providers do not support these compliance requirements natively. Therefore, the customer is responsible for ensuring regulatory compliance before outsourcing its applications and data, which increases management overhead and deployment complexity for application developers.

10.3. Interoperability

With numerous active blockchain projects operating on different platforms, each using different blockchains with different protocols, communication between these individual blockchains offers a great hurdle. The lack of a method or protocol that enables different blockchains to communicate and share data with one another leads to difficulties in scaling the technology. Interoperability in blockchain technology refers to the ability to exchange and develop information or utility without restriction and effectively perform cross-chain transactions.

New solutions for blockchain interoperability Methodological approaches for blockchain interoperability can be largely categorized into three groups: explorer solutions, notary schemes and relays. An explorer, such as BitMetrics, scans the blockchain for all relevant transactions and

their information in order to display it to the user in a specified form. Notaries have the ability and the submission of the transactions that need to be forwarded to the other chain. They act as trusted third parties that accomplish cross-blockchain proof and accompanying functions. Examples of equistructural notary schemes include the Pegged Sidechains, the Liquid sidechain and Sidechains, the Wrapped Tokens as well as the Drivechain Protocol. They essentially represent two-way pegs that enable certification of assets between two networks in both directions. Single-structural notary schemes like Atomic Cross-Chain Swaps on the other hand work without any links established on one side of the system and be accompanied by the disclosure of the private key that controls the funds in exchange for the funds on the other blockchain.

Another type of notary scheme operates on the basis of central notaries: the Custodians. As a third-party holding the assets involved in a transaction, Custodians play the role of a traditional bank. XAPO and BitGo are representative examples of this approach. Relays are smart contracts that verify the validity of transactions on other blockchains and inform the communicating chain accordingly. Blockchains are able to check the authenticity of the underlying records directly via a relay. Both chains therefore must offer the ability to access their own data resources in order to implement a relay. To reduce the blockchain interoperability problem, the use of standardization and regulation should be considered. Common rules, norms and language for the entire sector can enable side-to-side discourse. Blockchain protocols like ERC-20 or ERC-721 allow different projects to develop within the Ethereum network. The scalability of Ethereum and Bitcoin's networks, which play an essential part in overall blockchain adoption, also remains a key factor to achieve the interoperability of blockchains.

11. Future Research Directions in AI

The phenomenal growth of artificial intelligence in recent years has initiated a wave of disruptive innovation, creating a unique momentum and flood of interest in its application across numerous sectors. Present-day work on large-scale, diverse-curated-data inputs has provided novel avenues for addressing the AI open challenges. Many recent investigations have demonstrated AI applications for use in preventing general diseases, COVID-19, and in forecasting diseases. One of the promising future directions for AI can be in climate actions. Explainable artificial intelligence (XAI) is another aspect receiving significant research attention, aimed at enhancing the trustworthiness of machine-based decisions among humans. Additionally, issues related to privacy, security, and cyber-attack detection remain open topics for further exploration.

Future studies may also target the development of advanced, transformative AI methods that are highly interpretable, transparent, robust, cost-effective, and privacy-preserving. The COVID-19 pandemic has revealed that reassuring the public of the reliability and trustworthiness of AI results is a critical concern. Research into human-centric AI, focusing on integrating human expertise with machine intelligence through human–AI interaction, can enhance the overall quality of machine learning models. Cybersecurity, being a serious challenge, requires the

implementation of more sophisticated AI algorithms capable of detecting all types of network-based cyber-attacks.

11.1. Explainable AI

Explainable AI (XAI) endeavors to make the operations of AI-based systems more transparent, intelligible, and thus more understandable and explainable to users. Despite the proliferation of AI applications, there is a conspicuous knowledge gap among many users regarding the mechanisms, functions, and decision processes of these systems. This lack of understanding can engender a sense of distrust or fear toward AI technologies. The inability of complex AI or machine learning models to provide rationale for their decisions, often identified as the “black box” issue, constitutes a significant representation problem that challenges user acceptance.

Recent scholarly investigations have incorporated analysis of content published by industry-leading companies and individuals tasked with marketing AI solutions to public users. The findings reveal that while a substantial proportion of corporate content emphasizes AI's opportunities and potential benefits, comparatively fewer communications address the dangers or ethical concerns associated with artificial intelligence.

11.2. AI in Climate Change

Societal and concerns are becoming important topics in AI. The application of AI for climate change is illustrated here. AI can contribute to climate change mitigation and adaptation in several ways. AI models can be very useful for the prediction of future effects of global warming and support in engineering and developing carbon mitigation technologies, as well as analyzing and developing novel renewable energy sources such as wind and solar energy. Finally, AI can help to reduce and optimize energy consumption in homes, industries, energy networks, and transport and distribution systems. The following case study of the United Nations demonstrates that Machine Learning models contribute to the prediction tasks based on climate change data and focusing on greenhouse gas emissions, fossil fuel production, and reforestation. These tasks can efficiently support decision making and policies for climate change.

An outlined approach to understand the current climate situation and perform a prediction of what might happen in the future is proposed. Emissions dataset for coal, gas, oil, peat combustion, non-energy use by the U.S.A. states, and fossil fuels production dataset for the U.S.A. states during the last decades are merged and analyzed through Exploratory Data Analysis. Daily temperature and CO₂ concentration affecting the reforestation process in the U.S.A. during 10 years have also been considered. Finally, predictions of U.S.A. states' greenhouse gas emissions and fossil fuel production are made up to the year 2050 through Machine Learning models such as Artificial Neural Networks and Gradient Boosting Machine.

11.3. AI for Cybersecurity

Cybersecurity remains one of the major challenges in the information systems field and AI techniques are being used to provide new solutions. Always-on monitoring generates huge amounts of data that can be used for Security Information and Event Management (SIEM). However, compared to other industries, cyber-attacks generate, by definition, very modest amounts of data. The Cyber Battlefield is still «a primacy of the attack». When attacker data will be more abundant, other Cybersecurity strategies will emerge. It is inevitable that a lot of research will be needed on issues related to the CIARDS that will have to protect industrial, commercial, financial Big Data.

AI and ML infrastructures are already being used for cyber-attack detection, identification and response, including spam detection, malware categorization, intrusion detection, traffic analysis, attack analysis and mitigation, and vulnerability assessment. Hacking techniques that use Big Data analysis to target individuals or to identify the services accessed by Home/Business Users are also developing.

12. Future Research Directions in Big Data

Concluding the exploration of Big Data, the discussion now focuses on future research directions. Big Data refers to the voluminous information that inundates organisations on a daily basis. It is principled on the ‘four V’s’: velocity, variety, volume and veracity—in other words, the generation of data in real time of myriad types—stemming from millions and billions of sources worldwide—such that it clearly enables new and innovative capitalisation of the data sets created. Modern organisations are making use of this information to identify risks and gain insights that can be utilised to predict and solve problems. The deluge of data, however, seldom comes in completely uniform, neatly organised and easily digestible form, warning researchers to be cautious of the real value of Big Data.

A massive explosion in computing devices, which are able to collect evermore complex data, as well as increasingly sophisticated algorithms to derive different types of knowledge, requires the utilisation of existing and emerging techniques, as well as the development of new algorithms for real-time processing of Big Data. In many domains, Big Data has the potential to transform existing functions by delivering richer, faster and deeper insights. However, there are technical and non-technical limitations related to Big Data. Therefore, innovative ideas and methodologies are essential to further progress. These requirements need focused research attention and effort. To cultivate a better understanding of these issues, scholars have examined the effect of Big Data on individuals and society in general, e.g., by addressing the Big Data ethical problems of privacy and security. The research agenda is wide and extensive, ranging from frameworks for Big Data checklists to support marketing decision-making to the utilisation of advanced Big Data predictive analytics to identify potential suppliers, and product–market opportunities.

12.1. Real-time Data Processing

Providing a real-time streaming analysis for big data in a Hadoop framework is a challenge; Hadoop has an extremely high access cost and is not fit for real-time analysis. The use of the Superconducting QUantum Interference Device (SQUID) allows the transformation of a superconducting quantum bit to detect real time analysis in the Hadoop framework. Using Spark Streaming, a high-performance platform for real-time data processing and real-time data analysis in all industries, including real-time sensor analysis, is examined. Machine-to-machine communication has become an area of focus in recent years. Emotion is one event in communication that has vital importance for human beings, and Dempster–Shafer Evidence Theory was used for question selecting.

New trends in artificial intelligence techniques have proven that machine intelligence offers promising solutions in the decision-making process for a diverse range of applications. AI has thus been recognized as a support tool that provides human decision-makers with relevant information and enhances their reasoning ability in complex environments. Focus areas for future research in AI include explainable AI, using AI to tackle climate change and protect biodiversity, and applying AI for cybersecurity. Future research directions in big data should be considered in model processing and streaming analytics; ethical and legal concerns related to planned and operational models; and predictive modeling with big data. Employing AI and big data to serve environmental protection presents the challenges and opportunities in the field.

12.2. Data Ethics and Governance

Within the field of Big Data, ethical considerations are broadly defined and include aspects related to the nature and personal information content of the data, cultural elements, and the processing and manipulation of such data [126]. As an example, a survey-based study confirmed that data collected for a particular purpose should not be used for another purpose, even in the presence of anonymization, due to residual concerns about sensitive data [127]. These issues can be addressed by formulating and implementing appropriate data privacy policies. Techniques such as abstraction and aggregation can enforce distributed responsibility, incentivize users, and limit the side effects of broad data accessibility [128]. To address data-related concerns, strategies such as formalizing a data life cycle for collecting and managing personal information; building an organizational framework to regulate data collection, organization, analysis, and exhibition; and developing contextual integrity interpretations can be employed [129].

Data governance relates to a set of responsibilities and practices designed to ensure the quality, availability, usability, consistency, auditability, and security of the data employed in an organization. In China and the United States, the Privacy Protection Model for Information Management (PM4IM) method systematically collects user requirements to help organizations improve information privacy and governance by considering the technical, management, strategic, and social aspects of data privacy protection [130]. However, a significant challenge identified is that organizations attempt to minimize risk and related costs rather than establish an

efficient and transparent information-privacy infrastructure. An information management model applies Big Data in the product development field to enhance operational effectiveness, taking into account regulatory and privacy constraints, competitive threats, and the need for organizations to protect their actions from competitors. Organizations must be long-term oriented and have the capacity to manage and control the risks involved with Big Data.

12.3. Predictive Analytics

The emphasis on real-time analytics and AI has resulted in myriad solutions that purport to support the prediction of future trends, be it in the context of weather forecasting, earthquakes, or diseases. These represent significant paradigms in data-driven analyses. Predictive analytics exercises its influence across a plethora of applications, from the forecasting of natural disasters and at-risk public health environments to crime prevention. The reported results highlight the ability to provide substantial forecasting within these areas. However, the ever-increasing availability of multimodal data, both in terms of structural properties and semantic content, dictates a novel challenge for predictive analytics, which will centre upon the integration of knowledge across multiple heterogeneous data sources.

As an investigation in ethics, the researchers explore the undesirable or potentially harmful side-effects that may arise from instructions contained in product manuals, as a predictive feature designed to reduce the occurrence of erroneous user behaviour. Such errors can lead to adverse impacts upon the system that they support, both in a physical and a cyber context. This line of inquiry reflects the importance of ethical considerations within technologies that employ predictive analytics and serves as an appendix to the broader exposition of ethical issues concerning privacy and surveillance.

13. Future Research Directions in Blockchain

Multiple challenges still need to be solved in blockchain performance, regulation, and business models, namely scalability, sustainable regulation, and integration with legacy systems. Pursuing blockchain for social good also represents an important area of future study. Topics investigated include the integration of blockchain with the Internet of Things, cross-chain and side-chain mechanisms, and solutions that safeguard privacy.

Decentralized applications offer benefits but face limitations regarding privacy, operational costs, and throughput. Challenges arise when transgender individuals attempt to participate in blockchain-based voting or finance systems without disclosure of their transgender identity. While various mechanisms have been proposed for combating fake news, implementing these in a decentralized manner remains an open issue. The popularity of peer-to-peer systems coupled with the global rise in artificial intelligence provides opportunities for collaboration between P2P and AI technologies; creative proposals along these lines are welcome.




Figure 1: Generic applications of blockchain

13.1. Blockchain for Social Good

Blockchain technology can be used to address important social issues such as poverty, inequality, and climate change. Examples of Blockchain for Social Good initiatives include platforms that enable transparent charitable donations, microfinance loans for underserved communities, and decentralized renewable energy projects. Despite its potential, there is still debate about whether Blockchain is the most adequate technology for achieving these social goals.

Currently, challenges associated with the implementation of Blockchain-based initiatives for social good include scalability issues, the lack of regulatory clarity, and the need for interoperability between different platforms. Several questions remain unanswered. Is Blockchain really necessary for addressing social challenges? Do existing applications suffer from data fragmentation? Additionally, how do FinTech and DeFi contribute to the central banking system during a crisis period? Furthermore, how could the integration of Blockchain with other technologies, such as the Internet of Things (IoT), support social initiatives?

13.2. Integration with IoT

Blockchain technology serves as a natural solution to many challenges inherent in the Internet of Things (IoT) industry by providing decentralization, autonomy, and robust security mechanisms to its devices. Recently, blockchain was designated as a top investment area in IoT. Industry players are harnessing blockchain’s unique properties to resolve multiple complex issues, including tamper-proof data storage and uniform, robust infrastructural connectivity across the globe




Figure 2: Taxonomy of IoT security and privacy

The problem areas that can be addressed by combining blockchain and IoT solutions. Juniper Research predicts that, by 2023, over 700,000 active blockchain networks will be managing billions of connected devices, with the combined IoT–blockchain market set to exceed \$12 billion in value. Nevertheless, as with many novel technologies, the IoT–blockchain integration remains in its nascent stages, with many promising avenues yet to be explored and developed.

13.3. Cross-chain Solutions

Many blockchain platforms can be forked and modified, resulting in separate blockchains with mostly similar internal structures but carrying some differences. These distinct blockchains coexist simultaneously, and users can collaborate and transact across platforms that can provide complementary features and/or services. Cross-chain techniques enable interoperability across multiple blockchain platforms through cross-chain transactions, cross-chain routing, and cross-chain communication protocols by connecting independent blockchains. A set of business

applications that leverage two or more blockchains' characteristics and lending each other's help to support the operation or execution of a transaction is known as cross-chain business. A notable example of a cross-chain business is the use of Bitcoin for payment on the Ethereum blockchain.

Some examples of cross-chain protocols are the J.P. Morgan Interbank Information Network, measuring transactions and trades between banks and corporations; Zilliqa, a high-throughput blockchain supporting cross-chain communications; Hadera, offering a decentralized public ledger for the Internet of Things via the hashgraph consensus algorithm; and Tesla, working with Dogecoin to permit transactions in the company's retail operations. By leveraging the properties of these blockchains, emerging blockchain-based applications unlock new value dimensions across different industries, such as finance, insurance, manufacturing, and supply chain management [2,3].

14. Interdisciplinary Approaches

Implementation of the AI, Big Data, and Blockchain technologies are considered very important throughout many industries. Challenges in the implementation of these technologies are explored and analysed in sections 8, 9, and 10, respectively. These challenges are then related to future research directions in sections 11, 12, and 13. In the activities of interdisciplinary research, AI can be combined with those of Big Data and Blockchain. For example, from the GPT series of ChatGPT (Chat Generative Pre-trained Transformer), research can apply the technology to COVID-19-related problems with COVID-19 textual big data (such as research disease spreading models) and publishing records on coin-based blockchain (such as a decentralized publishing model). Exploring case studies and practical application examples of AI, Big Data, and Blockchain technologies yields practical implications in various application areas. Consideration of compelling challenges facilitates the identification of emerging research directions.

AI can be denoted as intelligence exhibited by machines or software. AI subfields include predicate logic, search algorithms, knowledge deduction, neural networks, expert systems, Machine Learning, Natural Language Processing (NLP), and Computer Vision. Machine Learning focuses on teaching machines to learn from historical data. NLP enables computers to understand human language. Computer Vision supports extraction of meaningful information from images and videos. Big Data involves the processes of data collection, data storage, and data analysis. Since Big Data can be collected and stored in large quantities, its analysis often requires advanced analysis techniques such as Artificial Intelligence. Blockchain ensures secure transactional lookup without information tampering by distributing ledgers across a network, with the nodes making collective pairwise decisions on new blocks of information added to the chain.

14.1. Collaboration between AI, Big Data, and Blockchain

Artificial Intelligence, Big Data, and Blockchain are among the most popular technologies of our time, with numerous studies exploring their technical characteristics and industrial applications. Combination also helps reveal emerging challenges and upcoming research directions in these fields. Examples of joint integration are demonstrated by several implementations and associated challenges across the three technologies. A recent study presents case studies, industry implementations, challenges, and future research directions related to artificial intelligence, big data, and blockchain.

Artificial intelligence provides the foundation and technology for creating intelligent robots that mimic human brains. Machine learning, natural language processing, and computer vision are fundamental branches of artificial intelligence. Healthcare, finance, and retail services industries have adopted machine-learning algorithms in their daily operations. Big data involves gathering huge amounts of data and drawing meaningful inferences from raw data. Big data implementation areas include telecommunication companies, manufacturing industry, and customer behavior analysis. Blockchain technology allows a decentralized system to eliminate central authority by implementing distributed ledgers and ledgers. Smart contracts, consensus mechanisms, and transaction blocks are foundational components of blockchain technology. Blockchain has been integrated into various industries, including supply chain management, land registration, and e-voting systems.

14.2. Case Studies of Interdisciplinary Projects

Infrastructure management companies are exploring the application of AI, Big Data, and Blockchain technologies in various sectors. An approach that combines public administration with AI and Big Data analyses was developed at a Dutch land registry company—it automatically determines the priority of client requests. In the Netherlands, AI assists education providers in adapting offerings to learners' wishes and needs. The Public Utilities Department in Toronto addresses electricity outages with an AI model that uses data on weather and past interruptions for prediction. Quality assessment of customers' experiences is also performed with NLP.

Big Data analyses applied to public administration in the Netherlands involve the use of advanced graphical representations and visualization of complex relationships of clients and requests. In the United States, Big Data analysis powers Telecom providers in resolving issues and planning services for rarer types of catastrophes. Integrating customer calls and transactional data generates insights on customer experience, which, when combined with modeling and data visualization, support operations, marketing, and advocacy, as in a Belgian bank. A manufacturing company in Italy applies Big Data analysis to monitor and predict machine status, including maintenance warnings. And in Australia, Big Data enables providers of home and community care services to obtain a better understanding of client needs for improved care planning.




Fig 3.Future directions

15. Conclusion

The three technologies discussed here—AI, Big Data, and Blockchain—are complementing one another and witnessing exponential growth in new research and innovations. Many new and specialized interdisciplinary directions are being constructed at the intersection of these technologies. Fresh sets of challenges emerge for these technologies whenever they become intertwined. The interdisciplinary topics of “opportunities and challenges” and “future research directions” of these three technologies, AI, Big Data, and Blockchain, are indeed fascinating.

The final section of this study selected and briefly presented possible future research directions and challenges of these three technologies. For the detailed analysis of these remaining topics, readers are encouraged to consult other specialized sources. This approach allows a focus on practical applications and emerging trends, in line with the preference for including case studies and proposed research directions.

References:

- [1] Ojika FU, Owobu WO, Abieba OA, Esan OJ, Ubamadu BC, Daraojimba AI. Transforming cloud computing education: Leveraging AI and data science for enhanced access and collaboration in academic environments. Journal name and details missing. 2023 Jan
- [2] Panda SP, Padhy A. Business Intelligence with Power BI and Tableau: Cloud-Based Data Warehousing, Predictive Analytics, and Artificial Intelligence-Driven Decision Support. Deep Science Publishing; 2025 Aug 15.

- [3] Koneti SB. Microstructure, and Risk Optimization Models. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025 Aug 12:17.
- [4] Mohapatra PS. Artificial Intelligence-Driven Test Case Generation in Software Development. Intelligent Assurance: Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. 2025 Jul 27;4:38.
- [5] Koneti SB. Artificial Intelligence in Financial Systems: Digital Transformation, and Machine Learning Applications. Available at SSRN 5401202. 2025 Aug 12.
- [6] Muppala M. SQL Database Mastery: Relational Architectures, Optimization Techniques, and Cloud-Based Applications. Deep Science Publishing; 2025 Jul 27.
- [7] Muppala M. Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience. Digital Oceans: Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience| Deep Science Publishing. 2025 Jul 8.
- [8] Panda SP, Koneti SB, Muppala M. Benefits of Site Reliability Engineering (SRE) in Modern Technology Environments. Available at SSRN 5285768. 2025 May 1.
- [9] Koneti SB. Artificial intelligence Applications in Retail and Investment Banking: Personalization, Robo-Advisory and Behavioral Analytics. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025;4:72.
- [10] Muppala M. Architectures in relational databases: An analytical study of SQL-based data models and ACID principles. database.;2:4.
- [11] Bentahar J. A Survey on Explainable Artificial Intelligence for Network Cybersecurity. arXiv (Cornell University). 2023 Mar 7.
- [12] Gadde H. AI-Assisted Decision-Making in Database Normalization and Optimization. International Journal of Machine Learning Research in Cybersecurity and Artificial Intelligence. 2020;11(1):230-59.
- [13] Koneti SB. Algorithmic Trading and Quantitative Finance Strategies: High-Frequency Trading, Market Microstructure, and Risk Optimization Models. Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution. 2025;4:17.
- [14] HUSSAIN, Fatima; HUSSAIN, Rasheed; HOSSAIN, Ekram. Explainable artificial intelligence (XAI): An engineering perspective. arXiv preprint arXiv:2101.03613, 2021.
- [15] deployment pipelines. International Journal of Research Publication and Reviews. 2025 Jan;6(1):871-87.
- [16] Reis J, Housley M. Fundamentals of data engineering. " O'Reilly Media, Inc."; 2022 Jun 22.
- [17] Ivanov SH, Webster C. Adoption of robots, artificial intelligence and service automation by travel, tourism and hospitality companies—a cost-benefit analysis. Artificial intelligence and service automation by travel, tourism and hospitality companies—a cost-benefit analysis. 2017.
- [18] Ramadhan M, Naseeb A. The cost benefit analysis of implementing photovoltaic solar system in the state of Kuwait. Renewable energy. 2011 Apr 1;36(4):1272-6.
- [19] Cordes JJ. Using cost-benefit analysis and social return on investment to evaluate the impact of social enterprise: Promises, implementation, and limitations. Evaluation and program planning. 2017 Oct 1;64:98-104.
- [20] Dykes PC, Curtin-Bowen M, Lipsitz S, Franz C, Adelman J, Adkison L, Bogaisky M, Carroll D, Carter E, Herlihy L, Lindros ME. Cost of inpatient falls and cost-benefit analysis of implementation of an evidence-based fall prevention program. InJAMA Health Forum 2023 Jan 6 (Vol. 4, No. 1, pp. e225125-e225125). American Medical Association
- [21] Chu H. Information representation and retrieval in the digital age. Information Today, Inc.; 2003.
- [22] Dominich S. Mathematical foundations of information retrieval. Springer Science & Business Media; 2001 Mar 31.
- [23] Reinanda R, Meij E, de Rijke M. Knowledge graphs: An information retrieval perspective. Foundations and Trends® in Information Retrieval. 2020 Oct 14;14(4):289-444.
- [24] O'Leary DE. Artificial intelligence and big data. IEEE intelligent systems. 2013 Jun 27;28(2):96-9.
- [25] Gadde H. AI-Augmented Database Management Systems for Real-Time Data Analytics. Revista de Inteligencia Artificial en Medicina. 2024;15(1):616-49.
- [26] Ojika FU, Owobu O, Abieba OA, Esan OJ, Daraojimba AI, Ubamadu BC. A conceptual framework for AI-driven digital transformation: Leveraging NLP and machine learning for enhanced data flow in retail operations. IRE Journals. 2021 Mar;4(9).
- [27] Muppala, M. . (2025). Digital Oceans: Artificial Intelligence, IoT, and Sensor Technologies for Marine Monitoring and Climate Resilience. Deep Science Publishing. 2025; doi:10.70593/978-93-7185-787-1
- [28] Koneti SB. Artificial Intelligence-Powered Finance Algorithms, Analytics, and Automation for the Next Financial Revolution. Deep Science. 2025; doi:10.70593/978-93-7185-613-3

- [29] Panda SP. Artificial Intelligence Across Borders: Transforming Industries Through Intelligent Innovation. Deep Science Publishing. 2025; doi:10.70593/978-93-49910-25-6
- [30] Eboigbe EO, Farayola OA, Olatoye FO, Nnabugwu OC, Daraojimba C. Business intelligence transformation through AI and data analytics. *Engineering Science & Technology Journal*. 2023 Nov 29;4(5):285-307.
- [31] Mohapatra PS. Intelligent Assurance Artificial Intelligence-Powered Software Testing in the Modern Development Lifecycle. Deep Science Publishing. 2025; doi:10.70593/978-93-7185-046-9
- [32] Ojika FU, Owobu WO, Abieba OA, Esan OJ, Ubamadu BC, Daraojimba AI. Transforming cloud computing education: Leveraging AI and data science for enhanced access and collaboration in academic environments. *Journal name and details missing*. 2023 Jan
- [33] Chattu VK. A review of artificial intelligence, big data, and blockchain technology applications in medicine and global health. *Big Data and Cognitive Computing*. 2021 Sep;5(3):41.
- [34] Pablo RG, Roberto DP, Victor SU, Isabel GR, Paul C, Elizabeth OR. Big data in the healthcare system: a synergy with artificial intelligence and blockchain technology. *Journal of integrative bioinformatics*. 2022 Mar 29;19(1):20200035.
- [35] Yu H, Yang Z, Sinnott RO. Decentralized big data auditing for smart city environments leveraging blockchain technology. *IEEE Access*. 2018 Dec 20;7:6288-96.
- [36] Paramesha M, Rane N, Rane J. Big data analytics, artificial intelligence, machine learning, internet of things, and blockchain for enhanced business intelligence. *Artificial Intelligence, Machine Learning, Internet of Things, and Blockchain for Enhanced Business Intelligence* (June 6, 2024). 2024 Jun 6.
- [37] Chowdhury RH. Blockchain and AI: Driving the future of data security and business intelligence. *World Journal of Advanced Research and Reviews*. 2024 Jul;23(1):2559-70.
- [38] Hassani H, Huang X, Silva E. Big-crypto: Big data, blockchain and cryptocurrency. *Big Data and Cognitive Computing*. 2018 Oct 19;2(4):34.
- [39] Han H, Shiwakoti RK, Jarvis R, Mordi C, Botchie D. Accounting and auditing with blockchain technology and artificial Intelligence: A literature review. *International Journal of Accounting Information Systems*. 2023 Mar 1;48:100598.
- [40] Kend M, Nguyen LA. Big data analytics and other emerging technologies: the impact on the Australian audit and assurance profession. *Australian Accounting Review*. 2020 Dec;30(4):269-82.