INTERNATIONAL JOURNAL OF COMPUTERS COMMUNICATIONS & CONTROL
Online ISSN 1841-9844, ISSN-L 1841-9836, Volume: 18, Issue: 6, Month: December, Year: 2023
Article Number: 5735, https://doi.org/10.15837/ijccc.2023.6.5735

Transforming Financial Decision-Making: The Interplay of AI, Cloud Computing and Advanced Data Management Technologies

S.A. Ionescu, V. Diaconita

Sergiu-Alexandru IONESCU
Department of Economic Informatics and Cybernetics
Bucharest University of Economic Studies
Bucharest, 010374 , Romania
ionescusergiu14@stud.ase.ro

Vlad DIACONITA*
Department of Economic Informatics and Cybernetics,
Bucharest University of Economic Studies
Bucharest, 010374 , Romania
*Corresponding author: diaconita.vlad@ie.ase.ro

Abstract

Financial institutions face many challenges in managing modern financial transactions and vast data volumes. To overcome these challenges, they are increasingly harnessing advanced data management technologies such as artificial intelligence and cloud computing. This paper presents a comprehensive review of how these tools transform financial decision-making in various domains and applications. We analyzed both foundational and recent advancements using a rigorous methodology based on the PRISMA 2020 guideline. Our findings indicate that many major financial institutions are adopting AI-driven solutions to potentially enhance real-time risk assessment, transactional efficiency, and predictive analytics. While they bring benefits like faster decision-making and reduced operational costs, they also pose challenges like data security and integration complexities that require further research and development. Looking ahead, we envision a more integrated, responsive, and secure financial ecosystem that leverages the convergence of AI, cloud computing, and advanced data storage. This synthesis underscores the significance of contemporary data management solutions in shaping the future of data-driven financial services, offering a guideline for stakeholders in this evolving domain.

Keywords: Data Management Technologies, Blockchain, Artificial Intelligence, Decision Support, Cloud Computing

1 Introduction

Financial institutions find themselves in a constant state of adaptation, driven by the essential demand to deliver effective and proficient financial services. At the core of this transformational journey is the requirement for trustworthy and precise data management systems. These institutions have moved beyond basic record-keeping methods to leverage innovative technologies like big data and cloud computing. These advancements not only offer improved data management capabilities but also facilitate better decision-making processes, reflecting the dynamic nature of the financial sector's evolution. The journey continues, with new challenges and opportunities arising from these ever-evolving technologies and their application within the financial sector.

The escalating intricacy of financial transactions, paired with an exponential growth in data, emphasizes the necessity for proficient and accurate data management systems. Key technological breakthroughs during the mid-20th century in integrated circuit technology and data storage capabilities mark a significant milestone in the evolutionary journey of financial institutions. Such advancements not only initiated the automation of manual processes and storage of enormous data volumes, but also sparked the onset of complex data analysis tools and methodologies, thus ushering in a significant transformation in the financial sector.

In the contemporary digital era, financial institutions are faced with the daunting task of managing and interpreting immense data volumes. This challenge has stimulated innovative progress in data analysis, driven by burgeoning technologies such as big data and artificial intelligence. These tools are harnessed to process and analyze data in real-time, thereby fortifying the efficiency of financial operations, promoting informed decision-making, mitigating risks, and elevating customer experiences.

The constant improvement in data storage and processing methods is key for financial institutions that need to meet growing demands for accurate and quick data, which is vital for making informed decisions. These technological improvements have become the base for financial institutions to quickly adapt to a fast-changing financial environment, strengthening their competitive standing and importance in a world focused on data. By using cutting-edge technologies like big data, artificial intelligence, and blockchain, financial institutions can make their operations more efficient, increase the accuracy of predictions, identify new trends, and uncover hidden opportunities. This leads to better management of risk, improved customer experiences, and better overall performance.

This paper offers a comprehensive overview of the evolution and current trends of data storage and processing technologies in the financial sector [1].

At the core of data management, relational databases [2] have served distinct roles in supporting both online transaction processing (OLTP) and online analytical processing (OLAP). OLTP systems prioritize fast and reliable transactional activities, typically found in everyday database operations such as order entry, retail sales, and financial transactions. These databases adeptly manage structured entities like customer details or orders information. In contrast, OLAP systems, often integrated within data warehouses [3], are designed to support complex queries and analytics, consolidating data from various sources, and enabling multidimensional analysis. As a result, data warehouses have become instrumental in serving as the backbone for business intelligence operations, bridging the gap between transactional data and actionable insights.

However, with the explosion of data volume and variety, systems like Hadoop [4] have become indispensable. Hadoop provides a scalable solution for storing and managing vast amounts of diverse data across distributed systems. Complementing Hadoop, Apache Spark [5] offers rapid data processing capabilities, equipped to handle streaming data and machine learning tasks.

As data complexities grow, the concept of Data Lakes has gained traction. These are expansive storage repositories, many times stored in the Cloud, capable of housing diverse and unstructured data, serving as preliminary storage for raw information intended for future analysis [6]. On the other hand, NoSQL databases [7] offer flexible data structures, accommodating both large and varied datasets, thereby becoming a choice for certain big data projects. In the realm of real-time data processing, Kafka [8] stands out. It efficiently moves data, streamlining the process of data transfer and real-time analysis. Additionally, Blockchain, with its immutable records, ensures transactional security and integrity [9]. Once data is written in systems like Blockchain, it becomes unalterable, which is pivotal for areas such as financial auditing and supply chain management [10].

These technologies, from the foundational relational databases to the modern tools like Spark and Kafka, collectively drive the digital transformation in the financial sector. The integration of these tools could be visualized as a workflow: data stored in Hadoop, processed with Spark, transferred via Kafka, and metadata potentially secured using Blockchain.

The current digital era underscores a significant shift towards data-driven decision-making in financial institutions. Although these advancements offer promising avenues for data management and analysis, they also introduce challenges, especially in handling large data volumes. By understanding the historical progression and current trends of these technologies, we aim to provide insights into the future landscape of the financial sector.

Central to our exploration are the research questions: "How are evolving data storage and processing technologies transforming the financial sector, and what are the future directions and implications of these advancements?" and "What are the challenges and potential solutions in implementing these new technologies within the financial sector?".

2 Methodology

To ensure the rigor and comprehensiveness of our review, we adopted the PRISMA 2020 guideline [11], a well-established framework for systematic reviews and meta-analyses. This methodology provided a structured and systematic approach to identify, screen, and include relevant studies and sources.

We began by defining a clear research question aimed at understanding the interplay of AI, Cloud Computing, and Advanced Data Management Technologies in financial decision-making. Considering our research questions, we identified keywords and search terms relevant to the topic. These terms included, but were not limited to, "Data Management Technologies," "Blockchain," "Artificial Intelligence," "Decision Support," and "Cloud Computing."

Using these search terms, we conducted an exhaustive search across multiple databases, including Web of Science, IEEE Xplore, ScienceDirect, Google Scholar, and the ACM Digital Library. Our search spanned articles published mainly from 2000 to 2022 to capture both foundational and recent advancements in the field. However, foundational, and historically significant papers, even if older, were highly valued for their seminal contributions.

After retrieving a preliminary set of articles, we employed a two-step screening process. Initially, we scanned titles and abstracts to exclude articles that were clearly not relevant to our research question. In the second phase, the full texts of the shortlisted articles were reviewed in depth to determine their relevance and quality. Any articles that did not provide substantial insights into the technologies or their implications in financial decision-making were excluded.

For each included study, we extracted relevant data points, such as the technology discussed, its application in the financial sector, benefits, challenges, and prospects. This information was tabulated to facilitate comparative analysis and synthesis.

Given the nature of our review and the computer science focus, the quality of referenced articles was evaluated based on several key criteria that ensure relevance, depth, and rigor:

- **Source Authenticity and Prestige:** Source Authenticity and Prestige: The origin of the publication plays an important role in determining its reliability. Studies published by renowned publishers, journals, or conferences (like ACM, IEEE, Springer, etc.) were given precedence due to their rigorous peer-review processes. Articles in journals published directly by universities or research centers were also favored.

- **Depth and Rigor of Content:** We evaluated the depth of the content in each article. Articles that provided comprehensive insights, rigorous methodologies, and substantive discussions on topics were deemed to be of higher quality. For instance, foundational papers like those by E. F. Codd on relational models or discussions on SQL's history by D. D. Chamberlin were considered vital due to their depth and historical significance.

- **Relevance to the Research Question:** We assessed each paper based on its direct relevance to our overarching theme. Papers that directly addressed the interplay of AI, Cloud Computing, and Advanced Data Management Technologies in financial decision-making, such as those discussing NoSQL and NewSQL data stores or the evolution of data warehousing, were given higher weight.

- **Recency and Timeliness:** Research is constantly evolving, so it is important to stay up-to-date on the latest findings. For this reason, recent papers are often given more weight in research reviews and other scholarly publications.
- **Citations and Impact:** Citations and Impact: We also considered the impact of the paper in the community. Articles that were widely cited and had significantly influenced the field were preferred. publications.
- **Reproducibility and Open Source:** Studies were evaluated based on the availability of their source code, datasets, and other resources. Those that provided open access or clear guidelines for reproduction were deemed of higher quality. In this context, the ACM's initiative of introducing Reproducibility Badges provided a valuable benchmark. For instance, papers awarded badges like "Artifacts Available" or "Results Replicated" indicated a strong commitment by authors to ensure verifiable results, elevating their credibility and quality.
- **Validation and Testing:** The extent and rigor of validation were assessed. Studies that showcased comprehensive testing, either on benchmark datasets, through simulations, or real-world applications, were favored.

Using these criteria, we critically assessed each article in the bibliography. Those that did not sufficiently meet these benchmarks were carefully weighed before inclusion in our discussions to maintain the robustness and integrity of our review.

3 Thematic Breakdown of Reviewed Works

To provide clarity on the scope of our research review, we have categorized the principal papers by distinct topic areas. Table 1 offers a concise breakdown of these categories, accompanied by brief descriptions and key references identified by the first author and year of publication. For a more in-depth exploration, you can easily locate these studies in our comprehensive bibliography.

Table 1: Thematic classification of reviewed literature highlighting main contributors.

Topic Area	**Description**	***Key References***
Historical Perspectives on Computing	Insights into the evolution and foundational concepts of computing.	Copeland et al. (2015)[12], Ceruzzi & Aspray (2003) [13], Oakley & Kenneth (1990) [14]
Database Management and SQL	Development, significance, and early history of SQL and relational databases.	Codd (1983) [2], Chamberlin (2012) [15], Chamberlin & Boyce (1976) [16], Elmasri et al. (2013) [17], Inmon (2008) [3]
NoSQL and NewSQL Databases	Emergence, characteristics, and importance of NoSQL and NewSQL data storage systems.	Grolinger et al. (2013)[7], Stonebraker (2010) [18], Fudivada et al. (2014) [19]
Big Data and Distributed Systems	Analysis and exploration of big data technologies, Spark, distributed storage, process mining and processing frameworks.	Harter et al. (2014) [20], Dean & Ghemawat (2008), Zaharia et al. (2016) [21], Gandomi & Haider (2015) [22] Hicham & Anis [23]
Cloud Computing	Concepts, definitions, and advancements in cloud computing technologies, including Edge Computing, Fog Computing and Cloudlet.	Varghese & Buyya (2018) [24], Red et al. (2020) [25], Chen et al. (2021) [26]
Artificial Intelligence and Decision Making	Role of AI in decision-making, advancements, and prospects.	Filip (2021) [27], Filip (2022) [28], Luger (2023) [29], Duan et al. (2019) [30], Collins et al. (2021) [31], Singh & al. (2019)[32]
Blockchain and Digital Assets	Exploration of blockchain technology, its applications, and significance in finance.	Dinh et al. (2018) [9], Ullah et al. (2022) [33], Castonguay & Smith (2020) [10]
Financial Sector and Technology	The role of modern technologies in shaping the financial sector, including advancements in data-driven decision-making.	Nauta & al. (2019) [34], Javaid et al. (2022) [35], Mirestean et al. (2022) [36], Gul & Al-Faryan (2023) [37]

The references listed in Table 1 serve as a foundational starting point for each respective topic area. It is important to note that these references are not exhaustive, and not all the cited articles in the comprehensive review are represented in this table. Readers should consult the main text for a broader range of references that provide a more comprehensive understanding of each topic.

4 The Evolution of Data Management

The progression of data recording and processing technologies has transitioned from simple mechanisms to more sophisticated and effective ones. Over the years, these systems have undergone continuous advancement, with each iteration improving upon its predecessor, leading us to the contemporary tools and methods utilized for handling and analyzing data.

Pivotal moments in the digital revolution occurred with the advent of innovative tools like the Electronic Numerical Integrator and Computer (ENIAC). Finished in 1945, ENIAC was the inaugural programmable, electronic, all-purpose digital computer. While preceding computers encompassed some of these characteristics, ENIAC consolidated all these attributes into a single machine. It was Turing-complete and could solve "a large class of numerical problems" through reprogramming and was a significant step in the evolution of computing, marking the transition from manual programming to the stored-program concept [12]. Its development laid the groundwork for the modern computing era.

The development of Integrated Circuits (ICs) and Microprocessors between 1958 and 1971 revolutionized the field of computing and played a crucial role in the evolution of data storage and processing. Prior to the invention of ICs, computers relied on bulky and unreliable vacuum tubes to store and process data [13]. ICs were smaller, faster, and more reliable, leading to the creation of smaller and more powerful computers.

Magnetic storage gradually replaced punch cards as the principal data storage medium by the mid-1960s. As personal computers and magnetic disk storage became more accessible and affordable, punch cards fell out of favor.

A major leap forward was the introduction of the hard disk drive. This innovation reshaped the landscape of data storage by enabling the storage of large volumes of data on a single device. It facilitated direct data access and dramatically increased data storage capacity.

Fast forward to today, we are witnessing substantial enhancements in hard disk drive technology. The advent of solid-state drives has refined the performance of traditional hard drives, offering faster and more reliable alternatives. Although they typically come with smaller storage capacities and are more expensive, the ongoing advancements, such as the NVM Express (NVMe) logical interface are pushing the boundaries of data storage and processing capabilities.

SSDs, replacing traditional Hard Disk Drives, ushered in a new era of speed and durability in storage. But the hardware revolution was only a facet of the transformation. The surge in data volumes necessitated decentralized storage and processing solutions. Distributed databases emerged, offering scalability and data integrity across vast networks. The Cloud didn't just offer storage; it presented a multitude of services from on-demand infrastructure to machine learning capabilities.

Big Data technologies like Hadoop and Spark further revolutionized the field, democratizing data processing and allowing for real-time insights from massive datasets. As we advance into the modern era, it's clear that the evolution of data management has been a precursor to the current innovative trends shaping the industry.

While the past has set the stage, the present and future of data management are brimming with innovations that have a profound impact on various sectors, particularly financial institutions. The tools and methodologies birthed in recent years aren't just evolutionary—they're revolutionary. As we transition to modern trends and technologies, it's essential to understand their significance and the transformative potential they hold, especially in areas like financial decision-making.

5 Modern Trends and Technologies in Data Storage and Processing

As we see a surge in data availability, discussions on efficient data storage and retrieval have intensified. In earlier times, the terms "Data Storage" and "memory" were often used interchangeably. However, our contemporary understanding of these terms has evolved to distinguish between them. Memory is generally viewed as a form of volatile short-term storage, while the concept of data storage is commonly associated with long-term storage. The multitude of available data storage options paved the way for advancements in relational databases, data warehouses, NoSQL databases and Big Data frameworks such as Hadoop or Spark.

5.1 The Advent of Modern Databases, Data Warehouses, and NoSQL

E.F. (Ted) Codd, a computer scientist at IBM's San Jose Research Laboratory, proposed a new way of organizing data that he called the "relational data model." This was a significant shift in the field of computer science in the early 1970s, particularly in the development of systems and languages for handling persistent data. Inspired by Codd's work, Don Chamberlin and Ray Boyce created the Structured Query Language (SQL), which is the de-facto standard for accessing data [2].

SQL implementations, like Oracle and DB2, emerged in the late 1970s and early 1980s. By 1986, the ANSI and ISO standards groups formally adopted a standard language definition, "Database Language SQL" [15]. These systems are called Relational Database Management Systems (RDBMS) and allow for data to be organized in a structured way, making it easier to search and analyze. This made it possible to perform complex data analysis and improve the efficiency of financial operations. The widespread adoption of relational databases paved the way for the development of new technologies and tools, such as data mining and predictive analytics, that transformed the financial sector [16].

The primary benefit of relational databases is their ability to store and manage large amounts of structured data in a well-organized and efficient manner and to enforce relationships between different data elements, which helps to maintain data integrity and consistency. Relational databases management software also provides robust security features, such as user authentication and access control, that help to protect sensitive financial data [17].

Data warehouses, introduced in the late 1980s, emerged as a pivotal solution for the consolidation, organization, and analysis of vast amounts of data from various sources. These systems collect and store data from diverse operational systems within an organization, facilitating access and comprehensive analysis. Whether housed in an organization's mainframe server or in cloud-based storage, they amass data from multiple Online Transaction Processing (OLTP) applications and other sources, serving business intelligence, decision support, and user query responses. Not only does this architecture support Online Analytical Processing (OLAP), enabling historical analysis and reporting for enlightened decision-making, but the flexibility of their data storage structures also allows for efficiency and versatility. While many maintain a relational structure analogous to databases, the modality of data storage can vary; some use row-based storage, others employ column-based storage, and certain systems can interchange between the two depending on the task at hand [3]. This adaptability has enabled financial institutions to conduct complex data analysis, thereby enhancing the accuracy and reliability of financial operations and enabling informed decision-making.

However, there are also some disadvantages to using relational databases and data warehouses. One of the main challenges is the limitations of the relational model in handling large volumes of unstructured data, such as text, images, and videos. This difficulty arises partly from the schema-on-write approach employed by relational databases, which requires data to conform to a predefined schema before it's written into the database. This approach can be restrictive and time-consuming when dealing with diverse and unstructured data types. These challenges can lead to performance bottlenecks and make it difficult for financial institutions to analyze and make sense of big data.

Another disadvantage of relational databases is that they can be inflexible and difficult to scale, which can make it challenging for financial institutions to respond to changing business needs and demands. This has led to the development of new data storage solutions, such as NoSQL databases, that are better suited to handling big data and can provide greater scalability and flexibility [7]. The volume, velocity, and variety of data increased dramatically, pushing the limits of relational databases. In response, in the late 2000s, NoSQL databases emerged, addressing the challenges posed by traditional relational databases in dealing with extensive, distributed, and high-speed data ecosystems [18]. The term "NoSQL" was first used around 2009 by Carlo Strozzi, who developed a lightweight database system that did not depend on SQL [38].

CouchDB, one of the earliest NoSQL databases, was launched in 2005, utilizing a document-based data model [39]. MongoDB, another well-known NoSQL database, was introduced in 2009 and employed a similar document-based data model [40], [41]. Multidimensional data warehouses can be implemented in column-oriented NoSQL systems such as HBase [20] or Cassandra [42]. One approach is to consolidate all related dimension attributes and all measures into a single table and column family for each fact. While this method is straightforward, it may not be the most efficient in terms of storage and querying. Another strategy involves storing all attributes of a single dimension in a dedicated column family for each fact, and all fact attributes (measures) in a separate column family. This method offers a more organized structure but may still encounter inefficiencies. A different technique involves allocating data in dedicated tables, with one for each dimension and another for the fact table. The fact table contains references to the dimension tables. This approach is beneficial in terms of reducing storage space usage, but it can potentially slow down querying due to the challenges of performing joins in NoSQL systems [43].

While Data Lakes and Big Data are part of the same data ecosystem, their functions are unique but synergistic. Big Data covers a wide array of tasks related to data processing and management. On the other hand, Data Lakes focus predominantly on the storage of raw data, which can encompass structured, semi-structured, and unstructured forms. Although Data Lakes can accommodate NoSQL databases, their principal emphasis is not on hosting such databases, but on serving as reservoirs for raw data [44]. This specific focus promotes a more streamlined and efficient data storage process, which can be pivotal in situations where swift data ingestion and storage are essential.

5.2 Technological Innovations: In-Memory Computing, Cloud Computing, Big Data and Blockchain

Simultaneously, advances in hardware technology led to the development of in-memory computing which built upon the strengths of NoSQL while pushing the boundaries of performance, speed, and real-time processing capabilities.

In-memory storage technology, which emerged in the 2000s, has improved data processing and analysis. By storing data directly in the main memory (RAM) rather than on disk, in-memory storage has considerably decreased data access times, enabling real-time analysis and decision-making across various sectors, including finance. In-memory storage allows for real-time analysis of large data volumes, supporting applications such as fraud detection, risk management, and high-frequency trading found in the financial industry [45].

The advancement of in-memory storage technology has also resulted in the development of in-memory databases and analytics platforms. SAP HANA is one example of an in-memory database platform that integrates database capabilities with advanced analytics tools like machine learning and predictive analytics[46]. Other well-known in-memory databases include VoltDB, MemSQL, and Apache Ignite[47].

Capitalizing on the progress of NoSQL and in-memory storage, two significant technological developments emerged in the 2000s - cloud computing and Big Data. These innovations significantly transformed the approach to managing and analyzing financial data.

Cloud Computing, by providing computing services through the internet[48], introduces a scalable, flexible platform capable of processing and storing massive data sets. This has effectively reduced the requirement for large-scale on-premises infrastructure [49]. Despite the challenges of security, privacy, and the potential for vendor lock-in, financial institutions have widely embraced cloud computing, drawn by the multitude of benefits it brings to the table.

It's essential to note that the key players in the cloud computing arena - AWS, Azure, and GCP - each provide unique advantages to the financial sector. AWS's strength lies in its scalability and flexibility, making it particularly adept at handling large volumes of financial data. Azure, with its seamless integration with other Microsoft software products, simplifies the migration process for financial institutions transitioning their existing applications to the cloud. Meanwhile, GCP excels at big data and analytics, making it an ideal choice for financial data analysis [50]. A SWOT (Strengths, Weaknesses, Opportunities, and Threats) analysis of the leading cloud service providers - AWS, OCI (Oracle Cloud Infrastructure), GCP (Google Cloud Platform), and Azure is presented in Table 2. The practical consequences for individual financial institutions may vary significantly depending on a multitude of factors such as existing infrastructure, particular application requirements, cost considerations, among others. Consequently, it is crucial for financial institutions to undertake a comprehensive assessment tailored to their unique needs prior to selecting a cloud service provider.

Table 2: SWOT analysis of cloud providers.

Attributes	AWS	OCI	GCP	Azure
Strengths	Broad range of services, largest market share, high scalability, robust security	Strong hybrid capabilities, optimized for Oracle databases, competitive pricing	Excels in big data processing and analytics, strong machine learning capabilities	integration with Microsoft products, extensive compliance coverage, strong hybrid capabilities
Weaknesses	Costs can add up with scale, complex pricing model	Less established in the cloud market, machine learning and AI services still developing	Less market share compared to AWS and Azure, some services less mature	Complexity and learning curve for non-Microsoft users
Opportunities	Continued dominance in the cloud market, expanding AI and machine learning services	Growth in the cloud market, leveraging existing Oracle database customers	Growth in the cloud market, expanding AI and machine learning capabilities	Leverage existing base of Microsoft product users, expanding AI and machine learning services
Threats	Growing competition from other cloud providers, risk of vendor lock-in	Competition from established cloud providers, risk of vendor lock-in	Competition from established cloud providers, risk of vendor lock-in	Growing competition from other cloud providers, risk of vendor lock-in

The integration of cloud computing within financial operations has become indispensable, yet it is not lacking complexities. Foremost among these concerns is the imperative of security and privacy. Financial entities play the pivotal role of guardians of sensitive financial information. Compounding this is the issue of vendor entrapment, which potentially amplifies costs and complications when transitioning between cloud providers or contemplating a return to an on-site system.

However, the incentives for embedding cloud technologies in the realm of finance are compelling. Through the incorporation of cloud infrastructures, financial institutions can achieve substantial cost savings, attain a higher degree of scalability and nimbleness, fortify their security stance and regulatory adherence, and fast-track technological and service innovations.

An essential application of cloud computing in finance is in the domain of data analytics. The prominence of Big Data analytics in modern financial transactions is irrefutable. Cloud solutions furnish the requisite infrastructure and tools for the swift and adept processing of vast data sets. As an example, a global banking corporation, leveraging cloud capabilities for storing and computing massive daily transactional data, can utilize Big Data analytical tools on this cloud-resident data. This facilitates a nuanced understanding of consumer behaviors, real-time detection of suspicious activities, and the customization of offerings tailored to individual clients. This synthesis between cloud technologies and Big Data not only enhances the operational acumen of banking entities but also vastly improves client experiences.

In summary, while cloud computing emerges as a transformative catalyst enabling financial entities to elevate their operational proficiency, optimize costs, and usher in innovations, it's imperative for these organizations to critically evaluate the intertwined advantages and challenges. Ensuring robust security protocols and strict adherence to relevant regulations is of the utmost importance.

The emergence of Big Data is attributable to several groundbreaking advancements in data storage and processing technologies. Distributed Computing, for instance, allows for the simultaneous processing of vast datasets across multiple computing devices, drastically enhancing performance. Cloud Technologies provide access to flexible and expansive computing resources, ideal for processing and analyzing large data volumes. Additionally, specialized Big Data Analytical Instruments, designed to process and interpret vast and intricate datasets, utilize a spectrum of methods, from machine learning to artificial intelligence, to derive insights.

The convergence of these innovations has catalyzed a paradigm where organizations are empowered to gather, archive, and critically analyze vast data collections both economically and efficiently. Consequently, Big Data has instigated profound transformations across numerous sectors, with the financial sector standing out prominently.

Traditional data management systems struggled to keep pace with this meteoric rise in data, highlighting the need for more sophisticated technologies. In response to this demand, solutions like Hadoop, Spark, and NoSQL databases emerged. These technologies brought with them increased scalability, flexibility, and performance [5], [22], [51], [52].

In this context, it's important to discuss the concept of schema-on-read versus schema-on-write.

Traditional relational databases, which follow a schema on write approach, require data to adhere to a predefined schema before being written into the database (i.e, tables must be created before data can be inserted). This method can be rigid and inefficient when dealing with Big Data, which often encompasses varied and complex data structures.

In contrast, the schema on read approach, typically adopted by Big Data technologies such as Hadoop and NoSQL databases, allows data to be stored in its raw form and the schema to be defined only when it's read for analysis. This provides the flexibility necessary for handling the diverse and rapidly changing data structures associated with Big Data.

The schema on read approach enables real-time processing and analysis of large data sets, a feature that is paramount in the dynamic environment of financial institutions. Hence, technologies that employ schema on read offer a compelling solution to the challenges posed by the influx of Big Data.

Hadoop, born out of Google's need to tackle large-scale data collection challenges, enabled scalability. However, the early versions of Hadoop's MapReduce component heavily depended on the Hadoop Distributed File System (HDFS), thereby curtailing the use of scalable cloud storage solutions [53], [4].

Most importantly, Apache Spark, launched in 2014, delivered notable performance enhancements by optimizing memory usage, revolutionizing the way big data processing was handled. Spark introduced the DataFrame API in 2016, assisting users transitioning from data exploration to harnessing Spark's scalable MapReduce engine. Spark 2.0 marked the integration of Spark SQL into the same optimization engine used for the DataFrame API [5], [54]

Other significant Big Data related technologies, such as Apache Kafka, also emerged, each with distinct advantages and challenges, including ingesting voluminous data sets and safeguarding sensitive financial data [8].

The key takeaway is that the intersection of Cloud Computing and Big Data technologies fundamentally rests on the premise that cloud computing provides the infrastructure and computational power indispensable to meet the demands of Big Data. The scalability of cloud services facilitates the storage, processing, and analysis of massive datasets, aligning with the essential requirements of Big Data. Big Data analytics tools, in turn, can be deployed on the cloud, leveraging the scalability and distributed processing capabilities of the cloud. This symbiotic relationship empowers businesses to improve data-driven decision-making, real-time analytics, and predictive modelling, while concurrently reducing operational costs and augmenting efficiency. Hence, the relationship between cloud computing and Big Data is not merely opportunistic, but rather a necessity, forming an integral link in the chain of digital transformation.

As the volume and complexity of data skyrocketed over time, a more flexible and scalable approach to data storage and management became imperative, resulting in the conception of Data Lakes. Emerging in the 2010s, Data Lakes represent a significant evolution from traditional data warehouses [55], serving as centralized repositories capable of storing raw data in its native format. This development constitutes a key component within the Big Data ecosystem, ushering in a fundamental shift in data management strategies. Contrary to traditional systems where data schemas are defined on-write, Data Lakes facilitate defining schemas on read, enhancing the flexibility for data exploration and discovery. As a result, Data Lakes prove particularly beneficial for exploratory data analysis, data science, and machine learning applications, where the data requirements might not be pre-defined.

The integration of blockchain technology and smart contracts into financial systems can be attributed to several critical factors. Notably, the inherent transparency and trustworthiness of blockchain's distributed ledgers can significantly enhance the clarity of financial transactions and prevent fraud without the need of banks or other intermediaries [35]. Therefore, it has been seen more as a threat to the well-established financial institutions rather than an opportunity to advance the traditional banking system, especially in emerging economies by reducing overhead [33]. Moreover, technical challenges such as scalability, security, flexibility, reliability, and interoperability must be overcome for effective implementation. If there are flaws in the programming languages used for smart contracts or in the classes implemented, there's a significant risk of large-scale fraudulent activities [10].

To determine whether blockchain is an apt solution for specific financial issues, a robust evaluation method is necessary. This includes establishing design principles for blockchain applications, considering technological, contextual, and organizational aspects to cater to financial system needs. Adoption of blockchain can lead to organizational transformation, involving changes to strategies, processes, operational structures, and risk management. The unmodifiable nature of data on a blockchain may be problematic for firms that regularly need to alter stored data. Concurrently, an appropriate legal framework is essential to accommodate the wide adoption of blockchain technologies, especially given the regulatory nature of financial systems. This demands collaboration between different financial institutions, regulators, and technology providers to balance regulatory measures and the opportunities provided by blockchain technology.

5.3 The Future of Data Management: AI and Machine Learning

Artificial Intelligence (AI) and Machine Learning (ML) have risen to prominence in the digital revolution as key tools for interpreting and utilizing the extensive volumes of data that are currently available. The advancements in big data and cloud computing, coupled with the rapid growth of computational power, have opened new horizons for AI and ML. The transformation of Artificial Intelligence (AI) into both a reputable scientific field and an influential set of tools has been remarkable. Reflecting Engelbart's vision, AI's growth over the years has served to amplify human intelligence [56]. This has empowered individuals and organizations to tackle and resolve intricate problems in ways that would have once been considered beyond reach. The financial sector, much like other industries, has been quick to seize the opportunities these technologies offer [57].

Artificial Intelligence (AI) is a broad concept that encapsulates the idea of machines performing tasks in ways that humans would consider "smart." This could range from a computer program engaging in a game of chess to a voice-recognition system like Amazon's Alexa or to a chatbot as ChatGPT. A subset of AI, Machine Learning (ML), is predicated on the notion that we should be able to provide machines with data and allow them to learn independently.

Deep learning, a further specialization within ML, is a central theme in recent research. This research emphasizes the use of neural language models and recurrent neural networks, which are innovative tools in the field of AI. These models are designed to learn how to convert word symbols into word vectors or embeddings, thereby predicting the next word in a sequence. They can even 'translate' the meaning of an image into an English sentence, demonstrating the remarkable versatility and potential of these models [58].

Moreover, the potential of unsupervised learning, a type of machine learning where the system learns by observing the world, much like humans and animals do, is underscored. This method of learning doesn't rely on explicit instructions but instead discovers the structure of the world through observation. It is anticipated that unsupervised learning will play a significant role in the future of AI, further expanding the capabilities of smart machines.

From automating mundane tasks to predicting trends, AI and ML offer a myriad of benefits. In banking, AI-driven chatbots are being used to automate customer service, while ML algorithms are being employed to detect fraudulent transactions and predict customer behavior. In insurance, AI and ML are used to automate claims processing and risk assessment.

AI and ML algorithms require vast amounts of data to learn and improve, and the growth of big data has fueled their advancement. The availability of large, diverse datasets has enabled the development of complex ML models that can make accurate predictions and decisions.

The ability to store and process these large datasets has been made possible by cloud computing. The immense computational resources provided by the cloud enable the training and execution of complex ML models that would otherwise be prohibitively resource intensive. Cloud platforms like Amazon Web Services, Google Cloud, and Microsoft Azure offer specialized services for AI and ML, providing the necessary tools and infrastructure for developing, training, and deploying ML models.

Despite these advancements, technology should not be seen as a panacea for all the problems facing the financial sector. Technology should be used as a tool to enhance the services offered by financial institutions and improve their operational efficiency, rather than as a replacement for the core values and principles that underpin the financial sector.

5.4 Advancements in AI based Decision Support in the Financial Sector

As an addition to data management technologies, Decision Support Systems (DSS) have become vital to the financial sector. These systems consist of diverse software applications, such as data analysis tools and machine learning algorithms, that assist in the organization, analysis, and management of unstructured data, thereby supporting the decision-making process. Consequently, the integration of DSS with data warehouses is critical, and the collaboration between them can drive unparalleled efficiency and effectiveness in data-driven decision making.

By design, a data warehouse is an optimal solution for integrating data from various sources, consolidating it into a unified and coherent format that is conducive to further analysis. When integrated with a DSS, the data warehouse becomes a powerful tool, providing the necessary data in a structured and accessible format that can be leveraged to generate insights and support strategic decision making[59]. Moreover, the integration of DSS with data warehousing enables real-time decision-making. With data continually being updated and processed in the data warehouse, the DSS can provide up-to-date insights that reflect the current state of the market, thus facilitating prompt and informed decisions.

The DSS can employ data mining techniques to extract patterns and trends, predictive analytics to forecast future scenarios, and machine learning algorithms to learn from historical data and make informed predictions. The outcome is a comprehensive and dynamic system that not only stores data but also uses it to inform strategic decisions, providing a competitive edge in the rapidly evolving financial landscape. DSSs can be designed to incorporate advanced visualization tools, which can transform complex data into easy-to-understand graphs, charts, and other visual representations. These visualizations can elucidate trends, patterns, and relationships in the data, thereby improving the overall decision-making process.

AI has been incorporated into decision support systems for a long time. During the early days of AI, "expert systems" commanded a lot of attention. They were usually built on rules and tried to copy how a human expert makes decisions [60]. But after some projects didn't work out, people realized these systems were more like helpers for human decision-makers, not experts telling humans what to do. Because of this, the term "expert systems" became less popular.

The term "knowledge-based system" (KBS) started to get attention when the UK Government launched the Alvey Programme for IT research in the '80s [14]. Some people thought expert systems were just a part of KBS, while others thought they were the same thing. At the start of the 21st century, businesses and managers started to prefer the term "knowledge-based systems" over "expert systems". This was because they started to focus more on "knowledge" rather than "the expert", especially with the rise of knowledge management in the late 1990s. Ontology and Social Networks were identified as crucial subjects associated with the evolution of KBS in the Big Data era [61].

In the 2010s, AI became a popular term again[29]. Within AI, people started to use the terms "machine learning" and "data mining" a lot more [31]. This change shows the progress in technology and the success of machine learning systems that we often hear about in the news. "Data mining" is a newer term compared to the others. It was first used in economics to describe the process of trying different models to find the best fit for the data. It wasn't until the mid-1990s that the term started to be used more broadly.

Addressing the challenges of AI requires a careful and balanced approach, involving regulatory oversight, ethical considerations, including the interaction between AI and humans, the impact of AI on job roles and skills, and technological solutions [30], [36]. Questions have been raised about the future of decision-making activities, specifically whether AI-based tools will evolve to become collaborative partners and active contributors, or if humans will simply serve as sources of information for these algorithms [28].

5.5 Comparison in the context of decision making

The blend of DSSs with data warehouses, Big Data, augmented by AI, signifies a major leap forward in data management within the financial sector. This combination not only enhances the storage and organization of data but also revolutionizes how this data is utilized, making it an indispensable tool for strategic decision-making. As the volume and complexity of financial data continue to grow, the significance of efficient data management systems like these, augmented by AI, will only become more prominent. The previously discussed technologies are compared in Table 3. This table focuses on several key attributes that differentiate these technologies.

Table 3: Technology comparison.

Technology	Primary Use	Data Model	Performance	Scalability	Flexibility	Decision Support
Relational Databases	Structured data storage and complex queries	Table-based (Rows and Columns)	High for structured data and complex queries; lower for unstructured data	Moderate; vertical scaling, distributed systems can be complex	Low; predefined schema required, less flexible with unstructured data	Moderate; supports structured data-based decision making
Hadoop	Batch processing of large data sets	File-based (HDFS)	High for large, batch jobs; lower for real-time analytics	Excellent; built for distributed processing	Moderate; built for structured and semi-structured data	Moderate;Suitable for large-scale data analysis and mining; Less efficient for real-time decision making
Spark	Fast, in-memory data processing and machine learning	In-memory / Disk-based	High for both batch and real-time analytics	Excellent; can handle large data sets in a distributed manner	High; can process structured and unstructured data	High; supports both batch and real-time decision making
Data Lakes	Storage of raw data in its native format	File-based and object-based	Varies depending on the processing engine used	Excellent; can store massive amounts of data	High; can handle any type of data (structured, semi-structured, unstructured)	Excellent for AI/ML applications due to diverse data storage; Requires efficient data management
Kafka	Real-time streaming data	Message-based	High for real-time data streaming	Excellent; can handle high-volume real-time data	Moderate; built for streaming data, not ideal for static data	Perfect to accompany real-time decision support systems; Less efficient for batch processing
NoSQL Databases	Handling large volumes of structured and unstructured data	Varies (document, key-value, column, graph)	High; designed for fast read/write operations	Excellent; built for horizontal scalability	High; schema-less design allows for high flexibility	High; Suitable for diverse and complex data types; Good for AI/ML applications
Blockchain	Decentralized and secure record keeping	Chain of blocks	Lower; transactions take longer to process	Moderate; adding new nodes increases network strength but slows transaction time	Low; transactions are immutable, and the structure is predefined	Low; Not ideal for traditional DSS due to transaction speed and immutability

It is necessary to point out that the comparison from Table 3 serves as a generalized overview. The practical performance, scalability, and adaptability of each technology can considerably diverge depending on the specific circumstances and implementation nuances.

Relational Databases play an integral role in decision support systems owing to their capacity for structured data storage and complex queries. They are designed with a table-based data model, using rows and columns to facilitate easy comprehension and detailed analysis. This makes relational databases highly valuable for decision-making processes that rely on structured data and complex analysis. Nevertheless, their performance and scalability are contingent on the nature of the data. While they are high performing for structured data and complex queries, their efficiency shrinks with unstructured data. The moderate scalability associated with relational databases is through vertical scaling, though distributed systems can pose complexity. Furthermore, flexibility is relatively low as a predefined schema is necessary, thus curtailing their ability to handle unstructured data fluidly. Within the realm of relational databases, for instance, there exist multiple options such as MySQL, PostgreSQL, Oracle, and more, each with its own unique attributes and advantages.

Hadoop is a robust tool for batch processing of voluminous data sets. The decision-making utility of Hadoop stems from its capability to process large batches of data. It is especially beneficial in scenarios where decision-making isn't time-critical and can be based on comprehensive historical data analysis.

Its file-based data model (Hadoop Distributed File System - HDFS) is designed to handle structured and semi-structured data. However, when it comes to real-time analytics, Hadoop's performance may not be as efficient, somewhat limiting its decision support in real-time scenarios.

park is distinguished by its capacity for swift, in-memory data processing, and machine learning. As such, Spark is a versatile tool that facilitates both batch and real-time data analytics, covering a wide range of decision-making scenarios. Its high-performance nature, coupled with the ability to handle large data sets in a distributed manner (it can also use HDFS), empowers decision-makers with accurate insights both from historical and real-time data. Moreover, Spark's capacity to process structured and unstructured data further enhances its decision support capabilities.

Data Lakes provides a platform for the storage of raw data in its native format. This means they store data as it is, without the need for a predetermined structure or schema. The flexibility in data storage opens the door for a variety of analyses and decision support. Data Lakes are particularly advantageous in complex decision-making processes that involve diverse data types - structured, semi-structured, and unstructured. The performance of a data lake, however, hinges on the processing engine used, influencing its effectiveness in decision support.

Kafka is a system designed for real-time streaming data. Its architecture makes it well-suited for decisions that necessitate real-time data streaming and instantaneous decision-making. Kafka's capability to handle high-volume real-time data empowers decision-makers in dynamic and fast-paced environments where swift decision-making based on real-time insights is critical. Nevertheless, Kafka's emphasis on streaming data makes it less ideal for static data, thus somewhat limiting its use in decision scenarios that demand thorough historical data analysis.

NoSQL Databases cater to the need for handling substantial volumes of structured and unstructured data. Their versatility in data handling offers an unprecedented level of flexibility in decision support. The schema-less design of NoSQL databases not only enables the accommodation of varied data types but also allows for adaptability to evolving decision-making environments. The inherent architecture of NoSQL databases, primed for fast read/write operations, empowers swift decision-making based on up-to-date data. However, the umbrella of NoSQL databases covers an array of technologies, each designed to fulfill the demands of large-scale data storage and retrieval with distinct capabilities. This category spans document databases like MongoDB, key-value stores like Redis, wide-column stores like Cassandra[42], and graph databases like Neo4j. Therefore, a more detailed understanding of each database type and its specific application is vital when assessing these systems.

Blockchain, with its unique attribute of decentralized and secure record-keeping, can aid decision support in scenarios requiring data veracity. The immutable nature of transactions within a blockchain and its predefined structure can, however, restrict flexibility in decision support. Moreover, scalability becomes a concern in blockchain-based systems as the transaction processing time increases with the addition of new nodes, potentially impacting real-time decision-making. The field of blockchain technologies presents a wide spectrum of options, each with unique features. Prominent examples include Bitcoin, Ethereum, BNB and many others, each possessing distinctive capabilities and attributes.

While certain technologies such as Apache Kafka and Blockchain may seem different, given that the former is a distributed event streaming platform and the latter a distributed ledger technology, they do have some fundamental similarities that significantly influence their ability to support decision-making processes:

- **Distributed Architecture:** Both Kafka and Blockchain are designed with a distributed architecture. This feature allows for reliable data sharing across multiple systems in Kafka, and resilient data storage across several nodes in Blockchain. In terms of decision support, this architecture offers a robust and reliable system that can handle high volumes of data and ensure data availability, key factors in data-driven decision-making processes.

- **Immutable Records:** Both technologies treat data as immutable once it's written. In decision-making scenarios, this immutability ensures the veracity of data. In Kafka, once an event is published, it remains unchanged, providing a reliable source of historical data for decision-making. Similarly, in Blockchain, the inability to alter a block once added to the chain ensures an accurate and reliable history of transactions, critical for decision-making processes especially in areas like financial auditing and supply chain management.

- **Sequential Data Storage:** The storage of data in a sequential manner in both Kafka and Blockchain helps maintain data integrity. In decision support, this feature assists in chronological analysis of data and identification of trends over time, thereby enhancing decision accuracy.
- **Fault Tolerance:** Both technologies are designed to be fault-tolerant, ensuring that data is available and durable, even in the event of system failures. This characteristic guarantees data availability for decision-making at all times. In the case of Kafka, data is replicated across multiple brokers, while in Blockchain, the ledger is replicated across all nodes, reinforcing system reliability.
- **Real-Time Processing:** Kafka is renowned for its capability to facilitate real-time data processing, a feature that is critical for many use-cases, including financial transactions where decisions need to be made promptly based on real-time data. The real-time processing capability varies with different Blockchain types and consensus mechanisms, but when available, it provides up-to-date data for timely decisions.

This comparison provides an initial understanding of the presented technology's' capabilities, while the rapidly evolving technology landscape suggests that their impacts will continue to diversify and probably deepen.

6 Conclusions and Future Directions

The continuous evolution of data storage and processing technologies has been pivotal in shaping the financial sector. Our research journey, initiated by an intrigue into the transformative role of these technologies, revealed their underlying impact on financial decision making.

From the foundational developments in data organization tools to the sophisticated capabilities of relational databases, data warehouses, and NoSQL databases, the trajectory of technological adoption has been driven by the financial sector's quest for efficiency, accuracy, and security. As we detailed, the emergence of AI and cloud computing has further accelerated data accessibility and processing speeds, directly influencing real-time financial decisions.

Our findings underscore the profound influence of big data technologies like Hadoop and Spark, which have redefined risk assessment and decision-making in financial institutions. Their ability to process vast datasets has facilitated more informed and timely decisions, empowering institutions to navigate the complexities of modern financial ecosystems.

However, the journey is not without challenges. The integration of AI with big data, while promising transformative interactions with clients, raises concerns related to biases, the interpretability of AI models, and data security. Addressing these challenges is vital to maintaining trust and ensuring ethical AI-driven financial decision making.

Looking forward, the promise of technologies like blockchain—with its potential for decentralized and transparent transactions—is being profoundly observed, though its maturity and standardization remain areas of development. Similarly, the nascent potential of DNA data storage suggests intriguing possibilities for the future of financial decision-making. The successful adoption and integration of these technologies will play a crucial role in shaping the agility and robustness of future financial systems, but it is essential to approach them with a balanced perspective, acknowledging their current limitations.

In sum, our research emphasizes the intertwined relationship between data technologies and financial decision making. As we progress, it is vital that the financial sector remains adaptive, ensuring that technology serves as a tool for enhancement and not an unchecked force. The goal should be a symbiotic relationship between technology and human expertise, leading to decisions that are not only efficient but also ethically sound and beneficial for all stakeholders.

Funding

This paper was co-financed by The Bucharest University of Economic Studies during the PhD program

Author contributions

The authors contributed equally to this work.

Conflict of interest

The authors declare no conflict of interest.

References

[1] Hasan, Md.M.; Popp, J.; Oláh, J. (2020). Current landscape and influence of big data on finance, *J Big Dat*, doi: 10.1186/s40537-020-00291-z, vol. 7, no. 1, p. 21, 2020.

[2] Codd, E.F. (1983). A relational model of data for large shared data banks, *Commun ACM*, doi: 10.1145/357980.358007, vol. 26, no. 1, pp. 64–69, 1983.

[3] Inmon, W. (2008). Building the Data Warehouse 3rd Edition, *Wiley*.

[4] White, T. (2012). Hadoop: The Definitive Guide, 3rd edition *Commun ACM*, doi: 10.1002/mus.23546.

[5] Zaharia, M. et al., (2016). Apache Spark *Commun ACM*, doi: 10.1145/2934664, vol. 59, no. 11, pp. 56–65, 2016.

[6] Weintraub, G. ;Gudes, E.; Dolev, S. (2021). Indexing cloud data lakes within the lakes *Proceedings of the 14th ACM International Conference on Systems and Storage*, doi: 10.1145/3456727.3463828, vol. 59, no. 11, pp. 56–65, New York, NY, USA: ACM, 2021, pp. 1–1.

[7] Grolinger, K. ;Higashino, W.A.; Tiwari, A.; Capretz M.A. (2013). Data management in cloud environments: NoSQL and NewSQL data stores *Journal of Cloud Computing: Advances, Systems and Applications*, doi: 10.1186/2192-113X-2-22, vol. 2, no. 1, p. 22, 2013.

[8] Kato, K. ;Takefusa, A. ; Tiwari, A.; Nakada H.; Oguchi M. (2018). A Study of a Scalable Distributed Stream Processing Infrastructure Using Ray and Apache Kafka *International Conference on Big Data (Big Data), IEEE*, doi: 10.1109/BigData.2018.8622415, pp. 5351–5353, 2018.

[9] Dinh, T. T. A., Liu, R., Zhang, M., Chen, G., Ooi, B. C., Wang, J. (2018). Untangling blockchain: A data processing view of blockchain systems. *IEEE transactions on knowledge and data engineering*, doi: 10.1109/TKDE.2017.2781227, vol. 30, no. 7, pp. 1366–1385, 2018

[10] Castonguay, J. J., Stein Smith, S. (2020). Digital Assets and Blockchain: Hackable, Fraudulent, or Just Misunderstood. *IEEE transactions on knowledge and data engineering*, doi: 10.1111/1911-3838.12242, vol. 19, no. 4, pp. 363–387, 2020

[11] Page, M. J., et al. (2021). The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *BMJ*, doi: 10.1136/bmj.n71, p. n71, 2021

[12] Copeland, B. J.,Sommaruga, G. (2015). The Stored-Program Universal Computer: Did Zuse Anticipate Turing and von Neumann? *Turing's Revolution, Cham: Springer International Publishing*, doi: 10.1007/978-3-319-22156-4_3, pp. 43–101, 2015

[13] Ceruzzi, P., Aspray W. (2003). A History of Modern Computing, second edition (History of Computing) *MIT Press*, 2003

[14] Oakley, B., Kenneth, O. (1990). Britain's Strategic Computing Initiative. *MIT Press*, 1990

[15] Chamberlin, D. D. (2012). Early history of SQL *IEEE Annals of the History of Computing*, doi: 10.1109/MAHC.2012.61, vol. 34, no. 4, pp. 78–82, 2012

[16] Chamberlin, D. D., Boyce, R. F. (1976). SEQUEL: A structured English query language *Proceedings of the 1976 ACM SIGFIDET (now SIGMOD) workshop on Data description, access and control*, doi: 10.1145/800296.811515, pp. 249–264., 1976

[17] Elmasri, R., Navathe, S.B.,Baydaoui J. (2013). Fundamental of Database Systems Seventh Edition. 2013.

[18] Stonebraker, M. (2010). SQL databases v. NoSQL databases *Communications of the ACM*, doi: 10.1145/1721654.1721659, vol. 53, no. 4, pp. 10–11, 2010

[19] Gudivada, V. N., Rao, D., Raghavan, V. V. (2014) NoSQL Systems for Big Data Management *EEE World Congress on Services*, doi: 10.1109/SERVICES.2014.42, pp. 190–197, 2014

[20] Harter, T., et al. (2014) Analysis of HDFS under HBase: A Facebook messages case study *Proceedings of the 12th USENIX Conference on File and Storage Technologies*, FAST 2014

[21] Zaharia, M., Chowdhury, M.,Stoica, I. (2012) Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing *In 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, doi: 10.1111/j.1095-8649.2005.00662.x, pp. 2–2, 2012

[22] Gandomi, A., Haider, M. (2015) Beyond the hype: Big data concepts, methods, and analytics *International journal of information management*, doi: 10.1016/j.ijinfomgt.2014.10.007, vol. 35, no. 2, pp. 137–144, 2015

[23] Hicham, R., Anis, B. M. (2022) Processes meet Big Data: Scaling process discovery algorithms in Big Data environment *Journal of King Saud University - Computer and Information Sciences*, doi: 10.1016/j.jksuci.2021.02.008, vol. 34, no. 10, pp. 8478–8489, 2022

[24] Varghese, B., Buyya, R. (2018) Next generation cloud computing: New trends and research directions *Future Generation Computer Systems*, doi: 10.1016/j.future.2017.09.020, vol. 79, pp. 849–861, 2018

[25] Ren, J., Zhang, D., He, S., Zhang, Y., Li, T. (2020). A Survey on End-Edge-Cloud Orchestrated Network Computing Paradigms *ACM Computing Surveys (CSUR)*, doi: 10.1145/3362031, vol. 52, no. 6, pp. 1–36, 2020

[26] Chen, X., Song, H., et al. (2021). Achieving low tail-latency and high scalability for serializable transactions in edge computing *in Proceedings of the Sixteenth European Conference on Computer Systems*, doi: 10.1145/3447786.3456238, pp. 210–227, 2021

[27] Filip, F.G. (2021). Automation and computers and their contribution to human well-being and resilience. *Studies in Informatics and Control*, doi: 10.24846/v30i4y202101, vol. 30, no. 4, pp. 5–18, 2021

[28] Filip, F.G. (2022). Collaborative Decision-Making: Concepts and Supporting Information and Communication Technology Tools and Systems *International Journal of Computers Communications and Control*, doi: 10.15837/ijccc.2022.2.4732, vol. 17, no. 2, 2022

[29] Luger, G.F. (2023). A Brief History and Foundations for Modern Artificial Intelligence *International Journal of Semantic Computing*, doi: 10.1142/S1793351X22500076, vol. 17, no. 01, pp. 143–170, 2023

[30] Duan, Y., Edwards, J.S., Dwivedi, Y.K. (2019) Artificial intelligence for decision making in the era of Big Data – evolution, challenges and research agenda *International journal of information management*, doi: 10.1016/j.ijinfomgt.2019.01.021, vol. 48, pp. 63–71, 2019

[31] Collins, C., Dennehy, D., Conboy, K., Mikalef P., (2021) Artificial intelligence in information systems research: A systematic literature review and research agenda *International Journal of Information Management*, doi: 10.1016/j.ijinfomgt.2021.102383, vol. 60, p. 102383, 2021

[32] Singh, G., Gehr, T., Püschel, M., Vechev, M. (2019) An abstract domain for certifying neural networks *Proceedings of the ACM on Programming Languages*, doi: 10.1145/3290354, vol. 3, no. POPL, 2019

[33] Ullah, N., Al-Rahmi, W. M., Alfarraj, O., Alalwan, N., Alzahrani, A. I., Ramayah, T., Kumar, V. (2022). Hybridizing cost saving with trust for blockchain technology adoption by financial institutions *Telematics and Informatics Reports*, doi: 10.1016/j.teler.2022.100008, vol. 6, p. 100008, 2022

[34] Nauta, M., Bucur, D., Seifert, C. (2019). Causal Discovery with Attention-Based Convolutional Neural Networks *Machine Learning and Knowledge Extraction*, doi: 10.3390/make1010019, vol. 1, no. 1, pp. 312–340, 2019

[35] Javaid, M., Haleem, A., Singh, R. P., Suman, R., Khan, S. (2022) A review of Blockchain Technology applications for financial services *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, doi: 10.1016/j.tbench.2022.100073, vol. 2, no. 3, p. 100073, 2022

[36] Mirestean, A., et al. (2021) Powering the Digital Economy: Opportunities and Risks of Artificial Intelligence in Finance *Departmental Papers*, vol. 2021, no. 024, p. 1, 2021

[37] Gul, R., Al-Faryan, M. A. S. (2023) From insights to impact: leveraging data analytics for data-driven decision-making and productivity in banking sector *Humanities and Social Sciences Communications*, doi: 10.1057/s41599-023-02122-x, vol. 10, no. 1, p. 660, 2023

[38] Sadalage, P., Fowler, M. (2012) NoSQL Distilled: A Brief Guide to the Emerging World of Polyglot Persistence *Vasa*, 2012.

[39] Anderson, J.C., Lehnardt, J., Slater, N. (2010) CouchDB: The Definitive Guide *O'Reilly Media, Inc*, 2010.

[40] Kabakus, A. T, Kara, R. (2017). A performance evaluation of in-memory databases *Journal of King Saud University-Computer and Information Sciences*, doi: 10.1016/j.jksuci.2016.06.007, vol. 29, no. 4, pp. 520–525, 2017

[41] Bradshaw, S., Brazil, E., Chodorow, K. (2019). MongoDB: The Definitive Guide: Powerful and Scalabl *Data Storage 3rd Edition*, 2019

[42] Lakshman, A., Malik, P. (2010). Cassandra: a decentralized structured storage system *ACM SIGOPS operating systems review*, doi: 10.1145/1773912.1773922, pp. 1–6, 2010

[43] Chevalier, M., Malki, M. E., Kopliku, A., Teste, O., Tournier, R. (2015). Implementation of Multidimensional Databases in Column-Oriented NoSQL Systems *In Advances in Databases and Information Systems: 19th East European Conference*, doi: 10.1007/978-3-319-23135-8_6, pp. 79–91, 2015

[44] Halevy, A., Korn, F., Noy, N. F., Olston, C., Polyzotis, N., Roy, S., Whang, S. E. (2016) Goods: Organizing Google's Datasets *in Proceedings of the 2016 International Conference on Management of Data*, doi: 10.1145/2882903.2903730, pp. 795–806, 2016

[45] Plattner, H. (2009, June) A common database approach for OLTP and OLAP using an in-memory column database *in Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*, doi: 10.1145/1559845.1559846, pp. 1–2, 2009

[46] Sikka, V., Färber, F., Goel, A., Lehner, W. (2013) SAP HANA: The evolution from a modern main-memory data platform to an enterprise application platform *Proceedings of the VLDB Endowment*, doi: 10.14778/2536222.2536251, vol. 6, no. 11, pp. 1184–1185, 2013

[47] Kaur, K., Sachdeva, M. (2017). Performance evaluation of NewSQL databases *International Conference on Inventive Systems and Control (ICISC)*, doi: 10.1109/ICISC.2017.8068585, pp. 1–5, 2017

[48] Liu, F., et al. (2011). NIST cloud computing reference architecture *NIST Special Publication*, 2011

[49] Mell, P., Grance, T. (2011). The NIST definition of cloud computing *NIST Special Publication*, doi: 10.6028/NIST.SP.800-145, 2011

[50] Sultan, N. (2014). Making use of cloud computing for healthcare provision: Opportunities and challenges *International Journal of Information Management*, doi: 10.1016/j.ijinfomgt.2013.12.011, vol. 34, no. 2, pp. 177–184, 2014

[51] Zhang, Q., Cheng, L., Boutaba, R. (2010). Cloud computing: state-of-the-art and research challenges *Journal of Internet Services and Applications*, doi: 10.1007/s13174-010-0007-6, vol. 1, no. 1, pp. 7–18, 2010

[52] Sivarajah, U., Kamal, M.M., Irani, Z., Weerakkody, V. (2017). Critical analysis of Big Data challenges and analytical methods *Journal of business research*, doi: 10.1016/j.jbusres.2016.08.001, vol. 70, pp. 263–286, 2017

[53] Dean, J., Ghemawat, S. (2008) MapReduce: simplified data processing on large clusters *Communications of the ACM*, doi: 10.1145/1327452.1327492, vol. 51, no. 1, pp. 107–113, 2008

[54] Armbrust, M., et al. (2015) Spark SQL: Relational Data Processing in Spark *in Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, doi: 10.1145/2723372.2742797, pp. 1383–1394., 2015

[55] Miloslavskaya, N., Tolstoy, A. (2016). Big Data, Fast Data and Data Lake Concepts *Procedia Computer Science*, doi: 10.1016/j.procs.2016.07.439, vol. 88, pp. 300–305, 2016

[56] Filip, F.G. (2021) AI vs AI (Augmenting [Human] Intellect vs Artificial Intelligence *in 2021 IEEE 15th International Symposium on Applied Computational Intelligence and Informatics (SACI)*, doi: 10.1109/SACI51354.2021.9465578, vp. 000011–000012, 2021

[57] Khargonekar, P.P., Dahleh, M.A. (2018). Advancing systems and control research in the era of ML and AI *Annual Reviews in Control*, doi: 10.1016/j.arcontrol.2018.04.001, vol. 45, pp. 1–4, 2018

[58] LeCun, Y., Bengio, Y., Hinton, G. (2015) Deep learning *Nature*, doi: 10.1038/nature14539, vol. 521, no. 7553, pp. 436–444, 2015

[59] March, S.T., Hevner, A.R. (2007) Integrated decision support systems: A data warehousing perspective *Decision support systems*, doi: 10.1016/j.dss.2005.05.029, vol. 43, no. 3, pp. 1031–1043, 2007

[60] Giarratano, J., Riley, G. (1998) Expert Systems: Principles and Programming *Third Edition. Course Technology*, 1998

[61] Cobo, M.J., Martínez, M.Á., Gutiérrez-Salcedo, M., Fujita, H., Herrera-Viedma, E. (2015) 25 years at knowledge-based systems: a bibliometric analysis *Knowledge-based systems*, doi: 10.1016/j.knosys.2014.12.035, vol. 80, pp. 3–13, 2015

Cite this paper as:

Ionescu, S.A.; Diaconita, V. (2023), Transforming Financial Decision-Making: The Interplay of AI, Cloud Computing and Advanced Data Management Technologies, *International Journal of Computers Communications & Control*, 18(6), 5735, 2023.
https://doi.org/10.15837/ijccc.2023.6.5735