INTRODUCTION

The advancement of artificial intelligence (AI) and machine learning (ML) has substantially influenced multiple sectors, with the finance sector undergoing a particularly notable transformation. AI refers to computational systems designed to perform tasks typically associated with human cognition, including decision-making, perception, language comprehension, and reasoning (Vernon & Furlong, 2007; Lieto et al., 2018; Korteling et al., 2021).

Within this domain, ML enables systems to improve performance through data exposure, without the need for explicitly programmed instructions (Jordan & Mitchell, 2015). These algorithms are well-suited for identifying patterns, generating predictions, and supporting decision-making in real time, particularly in contexts involving high-frequency and multidimensional data (Heaton et al., 2017).

In financial contexts, these technologies have transitioned from experimental automation tools to integral components of operational and strategic decision systems. Their applications span areas such as portfolio management, trading, credit evaluation, risk assessment, and sustainability-oriented investment strategies (El Hajj & Hammoud, 2023).

Initial implementations in finance can be traced to automated credit assessment models from the 1960s and 1970s, which relied on statistical techniques like linear and logistic regression. The following decade saw the emergence of expert systems designed to emulate decision rules, though these approaches proved limited in adaptability (Newell, 1982). During the 1990s, more flexible machine learning methods, such as decision trees and support vector machines (SVMs), gained prominence.

The current landscape has been shaped by developments in computational infrastructure, data availability, and software scalability. Contemporary learning systems are capable of analysing complex datasets in real time, supporting predictive insights and dynamic responsiveness to market fluctuations (Thayyib et al., 2023). These capabilities underpin a growing reliance on data-driven processes in financial institutions.

Whereas early systems were primarily used for task automation or basic econometric forecasting, current applications extend to autonomous decision systems. Multi-layer neural networks facilitate the modelling of non-linear phenomena in forecasting, text analysis, and anomaly detection (LeCun et al., 2015; Markauskaite et al., 2022; Russell & Norvig, 2010). Structured and sequential financial data can be processed through convolutional and recurrent network architectures, including long short-term memory (LSTM) networks, for applications such as pricing prediction and anomaly detection in transaction streams.

Decision-making in dynamic financial environments is also supported by methods that learn through iterative interaction. These approaches have been adopted in trading strategy development and portfolio allocation processes that adjust in response to continuously updated market data (Sutton & Barto, 2018). These techniques introduce flexibility not present in traditional optimisation models and enable financial agents to operate under uncertainty with increased adaptability.

Integration across financial services is extensive. In investment management, personalised advisory services are now automated through algorithmic systems capable of tailoring portfolios to individual risk profiles and financial goals (Sironi, 2016; Belanche et al., 2014). These tools contribute to increased accessibility and operational scalability.

In market operations, ML techniques support high-frequency strategies and adaptive trading algorithms capable of modelling short-term inefficiencies. These models continuously learn from new information and adjust execution strategies accordingly (Lu, 2019).

In credit evaluation, alternative sources of information, including behavioural and digital interaction data, are increasingly incorporated into modelling approaches, supplementing conventional financial indicators (Biju et al., 2024). These methods enhance predictive accuracy and inclusion, particularly for individuals outside the scope of traditional banking systems.

In fraud detection, real-time anomaly identification is enabled through pattern recognition algorithms trained on historical transaction data. These systems detect irregular activity without predefined labels, employing unsupervised learning methods well suited to the detection of novel fraud patterns (Lakhchini et al., 2022).

The methodologies underpinning these developments include a spectrum of learning approaches. Supervised models are used extensively for forecasting and classification tasks, including credit scoring and volatility prediction. Unsupervised models assist in segmentation and pattern identification. Techniques based on limited labelled data, such as semi-supervised and self-supervised methods, are increasingly important in financial domains where comprehensive annotation is unavailable (Wang et al., 2022).

Foundational to system accuracy is the preprocessing of input data. Procedures such as cleaning, normalisation, and dimensionality reduction contribute to improved model reliability. In textual data applications, processing steps such as stemming and term filtering facilitate improved model performance (Albahra et al., 2023). Data preparation is essential for ensuring robustness and consistency in financial prediction.

In settings characterised by uncertainty, interdependence, and nonlinearity, traditional statistical frameworks often prove insufficient. Learning-based approaches support risk assessment and forecasting under such conditions by identifying latent dependencies and adapting to shifting distributions. These methods are particularly useful for asset pricing, systemic risk evaluation, and the detection of operational threats.

ML techniques also demonstrate efficacy in handling irregular distributions and time series instability, both common in financial data. This makes them particularly suitable for volatility forecasting, scenario testing, and derivatives valuation. Ensemble models, transfer learning methods, and Bayesian optimisation provide further enhancements in prediction accuracy and uncertainty calibration.

An increasingly relevant domain is sustainability-oriented finance. Learning algorithms are employed to assess environmental, social, and governance (ESG) related indicators by analysing both structured and unstructured data from multiple sources, including disclosures, textual content, and sensor data (Chen et al., 2021). These techniques support the construction of sustainable investment portfolios and inform risk assessments sensitive to sustainability metrics.

The analysis of unstructured data through language processing tools enables the extraction of relevant signals from financial statements, regulatory texts, and public discourse. Sentiment tracking and reputational assessment benefit from such models, particularly in the context of ESG-sensitive investment. In addition, model interpretability frameworks have been proposed to address the opacity of predictive scoring systems.

Despite these advances, significant concerns remain. These include fairness, privacy, and transparency, particularly in models where decision logic is difficult to reconstruct. In response, new methodological directions have emerged, including interpretability frameworks, privacy-preserving learning protocols, and novel computational architectures. These developments aim to balance predictive performance with principles of accountability and control (Hassija et al., 2024). Distributed model training across institutions may enhance data protection, while emerging computational paradigms may address scale and complexity constraints (Herman et al., 2022). These directions indicate a broader shift towards responsible innovation in financial ML.

The book's objectives are threefold: i) to outline the core methodologies and technological enablers that power AI systems in finance; ii) to explore their applications in areas such as forecasting, trading, portfolio construction, fraud detection, credit evaluation, and sustainable finance; iii) to discuss the ethical, regulatory, and methodological challenges that will shape the future of AI-driven financial ecosystems.

The structure of this book unfolds across four chapters. Chapter 1 provides the technical foundation of financial AI systems, laying the groundwork for understanding the complex methodologies involved. Chapter 2 delves into predictive modeling and risk-oriented applications, exploring how AI and ML are used to address real-world challenges in finance. In Chapter 3, the focus shifts to AI's role in financial strategies, including market behavior and the integration of ESG principles. Finally, Chapter 4 tackles open challenges and emerging trends in the field, including explainable AI and new paradigms in ethical ML. The conclusion synthesizes key insights and offers strategic recommendations for investors, policymakers, and researchers.

Through this exploration, this book aims to shed light on the transformative potential of AI and ML in finance, providing readers with the knowledge and tools to navigate the future of this rapidly evolving field. By equipping practitioners and scholars with a deeper understanding of both the power and the limitations of these technologies, the book also aims to encourage the development of AI-driven financial ecosystems that are not only efficient and profitable but also ethical, resilient, and aligned with broader societal goals.