

The rise of artificial intelligence in algorithmic trading: a bibliometric study

European Journal
of Management
and Business
Economics

Aref Mahdavi Ardekani and Damien Dupré

Business School, Dublin City University, Dublin, Ireland, and

María Luisa Medrano-García and Raúl Gómez-Martínez

Rey Juan Carlos University – Madrid Campus, Madrid, Spain

Received 30 January 2026

Revised 3 June 2026

21 July 2026

Accepted 28 August 2026

Abstract

Purpose – This paper aims to provide a comprehensive bibliometric mapping of research on artificial intelligence in algorithmic trading by analyzing 417 publications from 1993 to 2025. It seeks to identify publication trends, influential journals and authors, collaboration patterns, and the field's main thematic structures, in order to clarify the intellectual foundations, emerging research fronts, and future directions of this rapidly evolving interdisciplinary domain.

Design/methodology/approach – The study employs a bibliometric research design based on 417 publications retrieved from the Web of Science Core Collection (1993–2025). Descriptive indicators are used to analyze publication trends and source productivity. Network-based techniques, implemented through VOSviewer, are applied to conduct citation analysis, co-authorship mapping, and term co-occurrence analysis. These methods enable the identification of influential contributions, collaboration structures, thematic clusters, and emerging research fronts within artificial intelligence-driven algorithmic trading.

Findings – The results show a sharp growth in publications since 2018, indicating the rapid consolidation of AI-driven algorithmic trading as a research field. A small number of highly cited studies act as intellectual anchors, while most authors and institutions remain weakly connected, revealing fragmented collaboration networks. Thematic analysis identifies coherent clusters centered on methodological innovation, trading strategy applications, time-series forecasting, and ethical and governance issues. Emerging research fronts, particularly transformer-based models and explainable AI, signal a shift toward more interpretable and robust trading systems.

Originality/value – This study provides the first large-scale, longitudinal bibliometric mapping of artificial intelligence research in algorithmic trading, covering more than three decades of scientific production. It uniquely integrates publication trends, citation impact, collaboration networks, and thematic clustering within a single analytical framework. By revealing both the intellectual consolidation of dominant AI paradigms and the fragmentation of collaboration structures, the paper offers novel insights into the field's evolution. It delivers strategic value for scholars and practitioners by identifying emerging research fronts and underexplored areas that shape future theoretical development and practical innovation.

Keywords Artificial intelligence, Algorithmic trading, Machine learning, Deep learning, Bibliometric analysis, Co-occurrence networks, Financial decision-making, FinTech innovation, AI governance, Quantitative finance, Algorithmic decision systems, Explainable artificial intelligence

Paper type Research article

1. Introduction

Artificial intelligence (AI) is transforming how trading strategies are generated, evaluated, and executed in financial markets. The adoption of machine learning (ML), deep learning (DL), and reinforcement learning (RL) has shifted algorithmic trading from traditional rule-based systems towards adaptive, data-driven systems that learn nonlinear relationships, detect structural market patterns, and automate decision-making at scale.

© Aref Mahdavi Ardekani, Damien Dupré, María Luisa Medrano-García and Raúl Gómez-Martínez. Published in *European Journal of Management and Business Economics*. Published by Emerald Publishing Limited. This article is published under the Creative Commons Attribution (CC BY 4.0) licence. Anyone may reproduce, distribute, translate and create derivative works of this article (for both commercial and non-commercial purposes), subject to full attribution to the original publication and authors. The full terms of this licence may be seen at <http://creativecommons.org/licences/by/4.0/>



European Journal of Management and
Business Economics
Emerald Publishing Limited
e-ISSN: 2444-8494
p-ISSN: 2444-8451
DOI 10.1108/EJMBE-01-2026-0047

Beyond its technical applications, the adoption of AI in financial markets is part of a broader process of digital transformation that is reshaping organisational decision-making, innovation processes, and competitive dynamics in the financial sector. AI-driven trading systems influence how investment firms allocate resources, manage risk, process information, and respond to changing market conditions. Algorithmic decision systems have consequently evolved from specialised analytical tools into strategic assets that affect organisational performance, market efficiency, and the structure of financial intermediation. Understanding the evolution of research in this area is therefore relevant not only from a technological perspective but also from management, innovation, and business-economics perspectives.

As data availability, computational capacity, and market microstructure complexity have expanded, AI models have moved from peripheral forecasting tools to central components of quantitative trading and market-microstructure research. Despite this expansion, the literature remains fragmented and conceptually under-integrated. Prior studies have examined specific algorithms, asset classes, or performance metrics, but no data-driven synthesis connects the evolution of methodological approaches, the structure of scientific collaboration, and the emerging research fronts within a unified analytical framework. In particular, little is known about how the intellectual foundations of AI-driven trading have developed over time, whether the field is converging towards a dominant paradigm or remains methodologically dispersed, and what these patterns imply for future theoretical development and managerial practice. This absence of an integrative perspective limits both academic understanding and the capacity of researchers, practitioners, and policymakers to assess the field's likely direction. To address this gap, the present study employs bibliometric techniques to map the evolution, impact, and thematic structure of AI research in algorithmic trading, offering a structured foundation for theory building, practical decision-making, and future research agendas.

By mapping more than three decades of scholarship, this study clarifies how the field has developed, where its core knowledge structures lie, and how emerging paradigms are shaping its trajectory.

2. Theoretical framework

The growth of AI-driven trading has been partly motivated by the limitations of traditional financial models in handling high-dimensional, nonlinear, and non-stationary market behaviour. Early academic contributions showed the potential of ML methods relative to classical econometric approaches, particularly their capacity to capture complex temporal structures and extract predictive signals from noisy data. Previous studies have reported that AI-based systems can outperform conventional linear benchmarks across diverse markets and trading horizons (Chen *et al.*, 2016; e.g. Li *et al.*, 2016; Sezer and Ozbayoglu, 2018; Shynkevich *et al.*, 2017). These studies created a methodological foundation that subsequent research refined through increasingly sophisticated architectures.

As DL matured, researchers widely adopted Long Short-Term Memory (LSTM) networks, Convolutional Neural Networks (CNNs), and hybrid DL models to analyse financial time series, identify market regimes, and generate trading signals. Several studies reported advantages of these approaches in capturing long-range dependencies and uncovering latent market structures. LSTM-based prediction and policy-learning frameworks, hybrid convolutional-recurrent architectures, and attention-enhanced models, for example, reported strong forecasting and trading performance across equities, foreign-exchange, and cryptocurrency markets (Gómez-Martínez *et al.*, 2022; Qing *et al.*, 2022; e.g. Wu *et al.*, 2019; Ye and Schuller, 2023). Collectively, this body of work positioned DL as a central methodological pillar in AI-based trading, often extending or outperforming traditional time-series approaches in predictive accuracy and adaptability.

RL expanded this paradigm further by conceptualising trading as a sequential decision-making problem in which an autonomous agent interacts with the market, receives rewards, and learns optimal policies. RL has been particularly influential in the development of end-to-

end trading systems capable of self-optimisation. A growing body of research, from early RL portfolio frameworks for cryptocurrency markets to more recent studies incorporating macroeconomic indicators, multi-agent interactions, and derivatives, illustrates the versatility of RL-driven trading (Conegundes and Pereira, 2024; e.g. Jiang and Liang, 2017; Lin *et al.*, 2022; Park *et al.*, 2024; Sang *et al.*, 2025). These contributions point to RL's capacity to integrate heterogeneous information sources, adapt to regime shifts, and optimise long-horizon risk-adjusted performance.

AI-driven trading research has also expanded beyond price prediction and strategy generation to encompass portfolio optimisation, trend-following, explainability, behavioural modelling, and volatility forecasting. Studies report that ML-enhanced allocation models can outperform traditional portfolio construction approaches, particularly under shifting market conditions or in emerging markets characterised by structural inefficiencies (e.g. Kawy *et al.*, 2021; Ye and Schuller, 2023). AI-based trend-following systems have shown robustness during periods of extreme volatility, such as the COVID-19 crisis, outperforming rule-based strategies by adjusting dynamically to changing momentum structures (e.g. Gómez-Martínez *et al.*, 2022). The integration of behavioural and sentiment-based data adds a further dimension. Studies incorporating investor sentiment, macroeconomic narratives, and volatility indicators into ML models have reported performance improvements and offer insights into market psychology. Examples include ML classifiers applied to sentiment-enhanced datasets for forecasting equity indices and models predicting volatility indices in markets with limited historical data (He *et al.*, 2021; e.g. Wang and Xie, 2024). These contributions mark a convergence between behavioural finance and AI, expanding the informational scope of algorithmic trading.

Beyond their predictive capabilities, AI-based trading systems can also be interpreted through theoretical perspectives on market adaptation, behavioural finance, and strategic decision-making. From a market-adaptation perspective, AI models are mechanisms through which financial actors adjust to evolving information environments and competitive conditions. Behavioural finance suggests that market outcomes are influenced by cognitive biases, sentiment dynamics, and information-processing limitations, factors that AI systems increasingly seek to identify, exploit, or mitigate. At the organisational level, AI-driven trading platforms function as strategic decision-support systems, enabling firms to process large volumes of information, automate complex decisions, and respond quickly under uncertainty. The evolution of AI in algorithmic trading therefore reflects not only technological progress but also broader changes in how financial organisations generate, interpret, and act upon market knowledge.

Cryptocurrency markets, characterised by extreme volatility, decentralised structure, and rapid innovation, have served as a major proving ground for AI-driven trading research. DL and RL models have been widely deployed to capture rapid price swings, detect structural anomalies, and optimise cryptocurrency portfolios. The development of explainable-AI tools for cryptocurrency trading reflects a broader shift towards interpretability and transparency, addressing concerns about model opacity in markets with high volatility and substantial speculative activity (e.g. Fior *et al.*, 2022; Jiang and Liang, 2017; Qing *et al.*, 2022).

As the literature has matured, scholars have increasingly recognised the limitations of static learning systems in nonstationary and adversarial markets. This recognition has motivated meta-learning frameworks and cognitive-theoretic RL models designed to improve adaptability. Research in this developing area suggests that meta-RL and game-theoretic learning mechanisms can improve strategy robustness in markets characterised by structural shifts and interactions among competing agents (Lin *et al.*, 2022; Park *et al.*, 2024; e.g. Shen and Huang, 2025).

More recently, transformer architectures and large language models (LLMs) have opened a new stage in AI-driven financial decision-making. Earlier research focused primarily on predictive performance using ML, DL, and RL techniques. Current developments increasingly emphasise AI systems that integrate structured and unstructured information, support human

decision-making, and operate within explainable and auditable environments. Recent research shows that LLM-based architectures are being developed as transparent, domain-adaptive components that support data management, semantic reasoning, regulatory compliance, and organisational governance rather than merely improving prediction accuracy (Karras *et al.*, 2025). Similarly, agentic AI architectures integrating LLMs, retrieval-augmented generation (RAG), multimodal reasoning, and human-in-the-loop mechanisms reflect a transition towards intelligent financial systems that prioritise explainability, accountability, and adaptive decision support (Ahi *et al.*, 2025). These developments suggest that AI in finance is moving beyond predictive modelling towards trustworthy, governance-oriented decision systems, a context relevant for interpreting the recent thematic trends identified in the present bibliometric analysis.

Despite the breadth of methodological innovation and domain expansion, the research remains fragmented. Studies are distributed across finance, engineering, computer science, and AI journals. Co-authorship networks exhibit weak global integration, and the field's intellectual structure is anchored by a relatively small number of highly cited works. These characteristics call for a structured, data-driven analysis of the field's development and of its thematic and collaborative architecture. Bibliometric methods are well suited to this task: they map citation structures, visualise thematic clusters, quantify collaboration networks, and track temporal research dynamics. Prior bibliometric work in finance and technology shows how science-mapping techniques capture the diffusion of innovations, the emergence of research fronts, and the evolution of interdisciplinary fields (Bhuiyan *et al.*, 2025; Chen *et al.*, 2025; e.g. Horobet *et al.*, 2024). Applying these methods to AI in algorithmic trading therefore offers a framework for understanding how the field has evolved and where it is headed.

The contribution of this study lies in providing an integrative bibliometric assessment of research on artificial intelligence in algorithmic trading. Recent studies have examined related topics, including algorithmic trading, deep learning applications, and sentiment analysis in financial markets, but existing reviews tend to focus on specific technologies, narrower thematic domains, or shorter observation periods. This study instead examines the evolution of AI-driven algorithmic trading research over more than three decades (1993–2025) and combines publication trends, citation impact analysis, collaboration networks, and thematic mapping within a single analytical framework.

Unlike previous bibliometric and review studies, which have addressed algorithmic trading in general, specific AI techniques (e.g. deep learning), or particular financial applications such as sentiment analysis, this study offers the first bibliometric assessment devoted exclusively to artificial intelligence in algorithmic trading. Integrating publication trends, citation impact, collaboration networks, and thematic evolution over the period 1993–2025 yields a broader and more cohesive account of the intellectual development of this interdisciplinary research domain.

By integrating these complementary dimensions, the study describes how the field has evolved, identifies the publications and authors that have shaped its development, reveals patterns of scientific collaboration, and uncovers emerging research fronts. Rather than focusing exclusively on methodological advances, the analysis also examines the interdisciplinary convergence between artificial intelligence, quantitative finance, and managerial decision-making. In doing so, the study contributes to a better understanding of the intellectual evolution of AI-based trading research and provides a structured foundation for future theoretical, methodological, and practical developments.

Although several review studies have examined related areas, their scope differs substantially from that of the present work. Horobet *et al.* (2024) mapped the broader literature on algorithmic trading without focusing specifically on artificial intelligence techniques. Bhuiyan *et al.* (2025) conducted a systematic review of deep learning applications for algorithmic trading, emphasising predictive models and optimisation strategies rather than the overall intellectual structure of the field. Chen *et al.* (2025) analysed the bibliometric evolution of sentiment analysis in stock market research, concentrating on a specific AI

application rather than the wider AI-driven trading ecosystem. In contrast, the present study provides the first bibliometric mapping focused specifically on artificial intelligence in algorithmic trading, integrating machine learning, deep learning, reinforcement learning, citation structures, collaboration networks, and thematic evolution over more than three decades (1993–2025). This broader perspective allows us to identify not only publication trends but also the intellectual foundations, interdisciplinary connections, and emerging research fronts of the domain.

Table 1 summarises the positioning of the present study relative to previous reviews.

Specifically, the study has four objectives: (1) to examine the temporal evolution of research output and identify key publication trends; (2) to assess the scientific influence of the journals, authors, and articles shaping the field; (3) to analyse collaboration patterns through co-authorship networks and evaluate the degree of interdisciplinary integration; and (4) to identify the main thematic structures and emerging research fronts through term co-occurrence analysis. By addressing these objectives, the study seeks to clarify the intellectual foundations of AI-driven algorithmic trading and to offer a structured basis for theoretical development, managerial insight, and future research agendas in finance, management, and business economics.

Given the rapid expansion, interdisciplinary nature, and methodological diversity of research on AI in algorithmic trading, a bibliometric approach is well suited to examining the field. Unlike traditional literature reviews, which typically synthesise findings from a limited number of studies, bibliometric methods enable the systematic analysis of large bodies of literature, revealing publication trends, influential contributions, collaboration patterns, and emerging thematic structures. The search strategy and data-collection procedures used to construct the bibliometric dataset are described in the Methodology section.

3. Methodology

Bibliometric analysis has become a central method of research evaluation. Since Garfield’s introduction of citation indexing in the 1960s (Garfield, 1964), bibliometric approaches have developed into tools for mapping the production, dissemination, and influence of knowledge. Supported by databases and software such as Web of Science, Scopus, VOSviewer, and bibliometrix, bibliometric analysis is particularly valuable in rapidly developing fields with high publication volumes (Aria and Cuccurullo, 2017). Applied to AI in algorithmic trading, it provides both a historical overview of the field and an assessment of its current research frontiers. When applied to algorithmic trading, bibliometric analysis can identify distinct streams of research. Recent studies point to clusters focused on ML and DL methods for stock

Table 1. Comparison of the scope, method, and main contribution of recent reviews of AI-related trading research (Horobet *et al.*, 2024; Bhuiyan *et al.*, 2025; Chen *et al.*, 2025) and the gap addressed by the present study

Study	Type	Scope	Main contribution	Gap addressed by this study
Horobet <i>et al.</i> (2024)	Bibliometric	Algorithmic trading	Evolution of algorithmic trading literature	Does not focus specifically on AI
Bhuiyan <i>et al.</i> (2025)	Systematic review	Deep learning in algorithmic trading	Predictive DL models	Limited to deep learning
Chen <i>et al.</i> (2025)	Bibliometric	Sentiment analysis in finance	AI-based sentiment research	Limited thematic scope
This study	Bibliometric	AI in algorithmic trading	Integrated mapping of publications, citations, collaboration and themes (1993–2025)	Comprehensive interdisciplinary perspective

price prediction and time series forecasting (Bhuiyan *et al.*, 2025), as well as RL and optimisation techniques that support automated decision-making in trading strategies (Horobet *et al.*, 2024). Other notable clusters include the integration of sentiment analysis and behavioural indicators into trading systems, where unstructured data such as news and social media play a growing role (Chen *et al.*, 2025), and the examination of risk management and market microstructure, with particular attention to liquidity, volatility, and the stability of algorithmic systems (Demir *et al.*, 2022).

To investigate the research landscape of AI in algorithmic trading, we conducted a bibliometric analysis based on publications indexed in the Web of Science (WoS) Core Collection (Clarivate Analytics, 2025), a widely used multidisciplinary database for scholarly output. The Web of Science Core Collection was selected for its rigorous journal selection criteria, standardised bibliographic records, and widespread use in bibliometric research, which support the reliability and reproducibility of science-mapping analyses. Broader databases such as Scopus or Dimensions include additional conference papers and regional publications; whether their inclusion would materially alter the publication patterns or intellectual structure reported here remains an empirical question.

The search strategy was designed to capture the intersection between AI and algorithmic trading. The search query applied was:

$$\begin{aligned} \text{TS} &= (\text{"algo * trading"} \text{ OR } \text{"quant * trading"}) \text{ AND TS} \\ &= (\text{"artificial intelligence"} \text{ OR } \text{"machine learning"} \text{ OR } \text{"AI"}) \end{aligned}$$

Here, TS denotes the Topic Search field, which includes title, abstract, and author keywords. The search was conducted within the Web of Science Core Collection.

The search strategy prioritised precision over breadth by focusing on terms directly associated with algorithmic and quantitative trading. Although this approach reduced the inclusion of peripheral studies, it may have excluded relevant contributions using alternative terminology, such as automated trading, trading systems, or AI-based financial decision-making.

The dataset comprised 417 documents published between 1993 and 2025 (Figure 1). The dataset consisted primarily of journal articles (70%), followed by conference proceedings (29%), and a small proportion of book chapters and reviews. For each record, we collected the available metadata, including title, authors, source, year of publication, abstracts, keywords, and citation counts. Unlike systematic literature reviews, which typically involve manual screening and exclusion stages, the present bibliometric study relied on a targeted search strategy. All 417 records retrieved from the Web of Science Core Collection were therefore retained for analysis, and no additional exclusion criteria were applied.

In accordance with established bibliometric protocols, the study used VOSviewer to conduct several network analyses. Descriptive statistics and figures were derived from the Web of Science export and VOSviewer outputs (Van Eck and Waltman, 2010). Co-authorship analysis was used to examine collaboration patterns, whereas term co-occurrence analysis based on titles and abstracts was used to identify the field's main thematic clusters. Citation analysis identified the most influential publications and contributors. Although export constraints prevented complete co-citation and keyword-network analyses, the available term co-occurrence networks provided useful evidence on the field's thematic organisation and research trends.

4. Results

4.1 The evolution of publications

The evolution of publications in this field reveals a clear trajectory of growth. The earliest publication in the dataset appeared in 1993, yet research activity remained limited for more




Figure 1. Flow diagram of the search and selection strategy applied to the Web of Science core collection, yielding the 417 publications (January 1993 to September 2025) analysed in this study

than two decades. A marked expansion began around 2018, coinciding with the growing adoption of DL and RL in financial applications. Publication output reached notable peaks in 2022 with 65 documents and in 2024 with 90, whereas the dataset contains 44 publications for 2025 as of 4 September, meaning that the final annual count is incomplete (Figure 2).

This trajectory reflects the emerging and rapidly expanding nature of the field.

Journal articles constitute approximately 70% of the dataset, reflecting the increasing consolidation of the field in peer-reviewed scholarly outlets. Conference proceedings account for approximately 29%, indicating the field's continuing links to computer science and engineering communities, in which preliminary findings are often disseminated through conferences. Reviews and book chapters constitute only a small proportion of the dataset. This acceleration in publication volume indicates that AI in algorithmic trading has moved from an experimental research niche towards an established interdisciplinary domain. The post-2018 surge coincides with wider access to computational resources (GPU/TPU), open-source frameworks (TensorFlow, PyTorch), and the mainstreaming of DL and RL. The temporal pattern also matches the diffusion curve typical of emerging technological paradigms: an initial exploratory phase (pre-2010), rapid expansion during methodological innovation (2018–2024), and a current phase of consolidation marked by increased journal specialisation. The predominance of journal articles over conference proceedings is a further indication that AI-based trading research is moving from technical feasibility studies towards theoretical and applied maturity.

The most relevant journals by publication count are presented in Table 2.

These journals illustrate the interdisciplinary nature of the field, spanning AI, computer science, and finance.

Table 3 shows the most productive authors by publication count.




Figure 2. Number of publications on AI in algorithmic trading indexed in the Web of Science core collection, 1993–2025. The 2025 count (44 documents) covers publications up to September and is therefore incomplete

Table 2. The 15 journals with the highest number of publications in the dataset ($n = 417$), led by expert systems with applications (30), IEEE access (21), and applied soft computing (12)

Journal	Publications
EXPERT SYSTEMS WITH APPLICATIONS	30
IEEE ACCESS	21
APPLIED SOFT COMPUTING	12
APPLIED INTELLIGENCE	11
MATHEMATICS	9
COMPUTATIONAL ECONOMICS	9
APPLIED SCIENCES-BASEL	7
QUANTITATIVE FINANCE	7
JOURNAL OF INTELLIGENT & FUZZY SYSTEMS	4
IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING	4
IEEE INTELLIGENT SYSTEMS	3
PATTERN RECOGNITION	3
ELECTRONICS	3
AI	3
NEUROCOMPUTING	3

Several other authors have 2 or 3 publications, reflecting a relatively dispersed author network without a single dominant research group.

4.2 Network analysis

Beyond descriptive indicators of productivity and impact, bibliometric methods allow for the exploration of the relational structure of a research field. Network analysis, applied to bibliographic data, captures how knowledge is created and disseminated through collaborations, citations, and thematic linkages. Mapping these relationships makes it possible to identify the intellectual structure of a domain, the main communities of researchers, and the emerging research fronts.

Table 3. The 15 most productive authors in the dataset by number of publications (range: 3–7), illustrating the absence of a single dominant research group

Author	Publications
Huang, Szu-Hao	7
Cohen, Gil	6
Cagliero, Luca	5
Huang, Yuling	5
Cliff, Dave	5
Gomez-Martinez, Raul	4
Chen, Chiao-Ting	4
Garza, Paolo	4
Wu, Mu-En	4
Sun, Shuo	4
An, Bo	4
Attanasio, Giuseppe	4
Lu, Xiaoping	4
Wu, Junjie	3
Sang, Chuan-Yun	3

4.2.1 Co-authorship networks. The co-authorship analysis indicates that collaboration within this field remains limited and dispersed. A few small groups of researchers co-publish, but most studies are produced either by individual authors or by teams that rarely collaborate beyond their immediate circles. The largest connected components contain only three to five authors, and many researchers appear as isolated nodes. This pattern reflects the emerging and interdisciplinary character of the field, in which expertise is spread across computer science, finance, and engineering departments. Among the most productive contributors are Gil Cohen and Yuling Huang, yet no single research hub or leading collaborative network has formed.

Figure 3 presents the co-authorship network for the 417 documents analysed. The map shows highly fragmented collaboration patterns: the majority of authors appear as isolated nodes or form very small clusters. Only a few weakly connected groups can be observed, and there is no large cohesive network of researchers working systematically together. This pattern is consistent with an emerging field in which contributions remain dispersed across institutions and disciplines, and most articles are authored by small independent teams.

The fragmentation observed in the co-authorship network suggests that knowledge diffusion in AI-based trading occurs through parallel innovation streams rather than centralised research hubs. Such structural dispersion is typical of rapidly evolving technological fields, where expertise develops simultaneously across institutions and geographies. This limits cross-fertilisation of ideas, but it also fosters methodological diversity, which may explain the rapid proliferation of heterogeneous models (Convolutional Neural Networks, Long Short-Term Memory, RL agents) in the literature. From an innovation-theory standpoint, this network topology indicates a pre-paradigmatic phase in which localised expertise precedes the formation of global epistemic communities. Strengthening cross-institutional collaboration could therefore accelerate the standardisation of methods and the translation of AI innovations into financial practice.

4.2.2 Term co-occurrence analysis. A co-occurrence analysis of terms drawn from titles and abstracts revealed three main thematic clusters within the field. The first, centred on methodological and experimental approaches, includes terms such as analysis, novel, propose, system, models, and accuracy, reflecting research devoted to the development and testing of new ML and AI techniques. The second cluster focuses on applications to financial trading, with frequent terms such as strategies, trading, stock, returns, financial, and results, capturing work on the design of trading strategies, the assessment of financial performance, and the practical integration of AI into trading systems. The third cluster, related to time-series and



forecasting, is smaller but notable; it includes terms such as time-series, long-term, technique, and success, reflecting research that applies AI to time-series prediction and long-horizon financial forecasting.

Figure 4 presents the co-occurrence map of terms extracted from titles and abstracts, produced using VOSviewer. The visualisation identifies four principal clusters: methodological and experimental approaches (blue and green), applications to trading and financial markets (green), studies focused on accuracy, datasets, and forecasting methods (yellow), and conceptual or ethical discussions surrounding artificial intelligence (red). The map illustrates the field's dual focus on technical innovation and financial application, while also revealing emerging trends such as the growing interest in transformer-based models, represented by the purple cluster.

4.2.3 Citation analysis. The publications in the dataset had received 4,605 citations in total as of the data-extraction date, with an average of 11 citations per publication. The most frequently cited work is that of [Sezer and Ozbayoglu \(2018\)](#), titled *Algorithmic financial trading with deep convolutional neural networks: Time series to image conversion approach*, published in *Applied Soft Computing*, which has received 253 citations. This study introduced a method that converts financial time series into two-dimensional images derived from technical indicators, thereby enabling the use of convolutional neural networks originally developed for image recognition in financial data analysis. The proposed CNN model classifies market patterns into buy, sell, or hold signals and outperformed the reported traditional benchmarks, including Buy-and-Hold and other standard trading systems, in out-

 VOSviewer

Several highly cited contributions, each with more than 100 citations, have shaped research at the intersection of AI and algorithmic trading. [Chen *et al.* \(2016\)](#) provided early evidence of the applicability of deep CNNs to noisy financial time-series data. [Jiang and Liang \(2017\)](#) introduced a model-free deep-RL framework for cryptocurrency portfolio management and reported strong back-tested performance relative to conventional strategies. [Shynkevich *et al.* \(2017\)](#) showed that forecasting performance depends materially on the relationship between the input-window length and the forecast horizon. [Théate and Ernst \(2021\)](#) subsequently introduced the Trading Deep Q-Network and a structured methodology for evaluating risk-adjusted trading performance. Together, these studies illustrate the growing influence of DL and RL on algorithmic-trading research and on standards of empirical performance evaluation.

into coherent theoretical frameworks that bridge financial modelling, ML, and behavioural finance.

5. Discussion

The bibliometric analysis shows the rapid growth and diversification of research on AI in algorithmic trading over the past decade. The descriptive indicators confirm that the field has experienced a sharp increase in scientific output since 2018, coinciding with the widespread adoption of DL and RL techniques in finance. Journals such as *Expert Systems with Applications* and *IEEE Access* have emerged as key publication venues, while contributions remain relatively dispersed across other outlets in computer science and applied intelligence.

5.1 Overview of the results

Citation metrics show that a small group of highly cited studies serves as the field's methodological foundation. The mean of approximately 11 citations per publication indicates growing scholarly visibility, although citation counts are sensitive to publication age. This concentration may encourage cumulative development around established architectures, but it may also narrow theoretical diversity if subsequent studies focus primarily on incremental extensions.

The co-authorship network shows that collaboration remains weak and fragmented. Most authors appear as isolated nodes or in very small clusters, with no dominant research hub or cohesive international network. This pattern is typical of emerging interdisciplinary domains, where scholars from finance, computer science, and engineering work in parallel but have not yet established strong collaborative ties. Strengthening co-authorship networks could improve knowledge integration and accelerate methodological innovation in the field.

The fragmented collaboration network suggests that AI-based trading research is still characterised by dispersed communities operating across finance, computer science, and engineering. This fragmentation encourages methodological diversity and experimentation, but it may also slow the development of common research standards, reduce knowledge transfer across disciplines, and delay the emergence of shared theoretical frameworks.

The co-occurrence map of terms reveals a more structured intellectual organisation within the field. Four main thematic clusters can be distinguished. The first, centred on methodological development, focuses on the creation of new models and the evaluation of their performance. The second, devoted to applications in trading strategies, examines stock prediction, financial returns, and automated trading systems. The third cluster, related to forecasting and time-series modelling, explores comparisons between traditional and AI-based predictive methods. The fourth, encompassing conceptual and ethical perspectives, addresses the broader implications and risks associated with the use of AI in financial markets. A smaller emerging cluster indicates growing interest in transformer-based architectures, a potential new front for research.

Viewed through the lens of technological diffusion and paradigm formation, the results indicate that AI in algorithmic trading is moving from a pre-paradigmatic stage towards early consolidation. The initial phase (pre-2018) was characterised by dispersed experimentation and heterogeneous methodologies, while the subsequent expansion phase shows increasing convergence towards DL and RL paradigms. Thematic clusters and citation hierarchies now function as epistemic anchors, structuring how new entrants to the field frame problems and methods. Yet the fragmented co-authorship network reveals that institutional diffusion lags behind intellectual diffusion: methodological ideas spread quickly, while cross-institutional collaboration has yet to catch up.

The bibliometric evidence also points to the interdisciplinary nature of this research ecosystem. AI-driven trading sits at the intersection of computational science, quantitative finance, and behavioural economics, domains that differ in their epistemic assumptions and

validation standards. The observed thematic clusters reflect these tensions: one cluster prioritises algorithmic precision and model optimisation, while another focuses on interpretability, ethics, and market implications. Such epistemic plurality fragments collaboration networks in the short term, but it also fuels innovation by maintaining methodological diversity. In the long term, bibliometric trends suggest gradual epistemic integration, in which shared methodological norms (benchmarking, reproducibility, and hybrid models) may support closer links between AI research and financial-market theory.

Overall, these findings indicate that the field is simultaneously maturing and diversifying: intellectual structures are beginning to consolidate around a few methodological and applied clusters, while the social structure of research remains fragmented. The combination of strong thematic coherence and weak collaboration networks suggests that AI-enabled algorithmic trading is a developing research area with considerable potential for integration across disciplines and geographies.

The rapid growth of publications on AI in algorithmic trading can be largely explained by the convergence of technological, economic, and institutional drivers. Advances in computational power, the widespread availability of financial data, and the maturation of ML frameworks have substantially lowered entry barriers, enabling researchers from diverse disciplines to contribute to the field. At the same time, the economic relevance of algorithmic trading and the strategic value of AI-driven decision systems have intensified academic and professional interest, accelerating publication output since 2018. This expansion has, however, occurred in a fragmented manner. The bibliometric evidence shows that research has developed through connected streams across finance, computer science, and engineering, reflecting differences in epistemic traditions, validation standards, and publication outlets. Such fragmentation limits cumulative knowledge building and interdisciplinary collaboration, but it has also fostered methodological diversity and rapid experimentation. Looking ahead, the field is likely to enter a phase of consolidation characterised by stronger theoretical integration, increased cross-disciplinary collaboration, and greater emphasis on the explainability, governance, and robustness of AI-based trading systems. As regulatory scrutiny intensifies and market environments become more complex, future research is expected to move beyond performance optimisation towards integrative frameworks that align technological innovation with managerial decision-making and financial stability.

Taken together, the bibliometric patterns identified in this study suggest that AI in algorithmic trading is following a trajectory similar to that observed in other emerging technology-driven fields. The coexistence of strong thematic convergence and weak collaboration networks indicates that methodological consolidation is occurring more rapidly than institutional integration. Researchers are increasingly converging on a shared set of tools and analytical approaches, particularly deep learning and reinforcement learning, even though collaboration remains fragmented across disciplines and research communities. This pattern is characteristic of fields undergoing paradigm formation, in which intellectual convergence precedes the emergence of stable collaborative ecosystems. If this trend continues, future research is likely to become more interdisciplinary, integrating advances in artificial intelligence with broader concerns related to financial decision-making, governance, explainability, and market resilience.

5.2 Theoretical and practical implications

From a theoretical perspective, this study contributes to the literature on management and business economics by clarifying the intellectual structure and evolutionary dynamics of AI in algorithmic trading. The mapping of citation patterns and thematic clusters shows that the field is converging towards a limited set of dominant methodological paradigms, particularly DL and RL, while remaining fragmented in terms of interdisciplinary integration and collaborative structures. This tension between methodological convergence and institutional dispersion points to a theoretical gap in how AI-driven financial innovation is conceptualised. The

findings suggest opportunities for finance and management theory to better account for adaptive, data-driven decision systems, which complicate conventional assumptions regarding market efficiency, rationality, and model stability. By positioning AI-based trading as an emerging technological paradigm rather than a collection of isolated techniques, this study provides a conceptual foundation for future theory building that integrates financial economics, strategic management, and AI research.

The identification of emerging research fronts, such as transformer-based architectures, explainable AI, and meta-learning frameworks, extends current theoretical debates by signalling a shift from performance-oriented modelling towards concerns related to interpretability, robustness, and governance. These trends indicate that future theoretical development in algorithmic trading will need to move beyond predictive accuracy and incorporate organisational, ethical, and regulatory dimensions, reinforcing the relevance of management-oriented perspectives in AI-driven financial research.

For practitioners, including asset managers, quantitative trading firms, and fintech organisations, the findings offer insights into the strategic landscape of AI-driven algorithmic trading. The concentration of academic citations around a limited number of model architectures may indicate methodological homogeneity within the research literature; bibliometric evidence alone, however, cannot establish the degree of model homogeneity in industry practice. Practitioners should nevertheless consider diversity in model design, data sources, validation methods, and learning frameworks to improve resilience across changing market conditions.

The fragmented collaboration networks identified in this study also carry implications for organisational strategy. Although bibliometric collaboration patterns cannot be mapped directly onto organisational structures, the findings point to the potential value of interdisciplinary teams when developing AI-based trading systems. Encouraging collaboration between quantitative analysts, data scientists, and risk managers can improve model design, deployment, and oversight. In this regard, the growing prominence of explainable-AI tools represents a practical opportunity to connect technical performance with managerial accountability, supporting better communication between technical and non-technical stakeholders.

Finally, the emergence of ethical and governance-related research themes indicates the importance of responsible AI adoption in financial markets. Regulators and practitioners alike can use the structured overview provided by this study to identify issues likely to attract regulatory attention, including transparency, model risk, and market stability. By aligning trading innovation with explainability, robustness, and ethical considerations, organisations can improve performance while strengthening trust and legitimacy in increasingly automated financial environments.

The findings also have broader implications for financial institutions, regulators, and policymakers responsible for overseeing AI-enabled financial markets. Financial institutions should strengthen governance frameworks for AI-based trading systems by incorporating model validation, continuous monitoring, and explainability mechanisms to mitigate operational and model risks. For regulators and policymakers, the growing complexity of AI-driven trading points to the need for standards that balance innovation with transparency, accountability, and market stability. The bibliometric evidence presented in this study suggests that governance-related research is becoming an increasingly important component of the field, supporting the development of regulatory approaches capable of addressing the systemic implications of autonomous financial decision-making.

5.3 Limitations and future research agenda

This study is subject to certain limitations. First, it relies exclusively on the Web of Science Core Collection, which, while authoritative, may omit relevant contributions indexed in Scopus, IEEE Xplore, or SSRN. Second, due to database export constraints, complete co-

citation and author-keyword network analyses could not be conducted, limiting the scope of the intellectual structure captured. Finally, the field itself is still in a state of flux, meaning that bibliometric patterns will likely evolve rapidly in the coming years.

Future research should expand bibliometric mapping by incorporating multiple databases and by applying complementary techniques such as bibliographic coupling or text mining of full articles. As the field matures, it will also be important to study how interdisciplinary collaborations evolve, how new AI architectures (e.g. transformers, generative models) reshape trading research, and how ethical, regulatory, and societal concerns influence both academic and practical developments in AI-driven finance.

6. Conclusion

This study maps the development of research on AI in algorithmic trading between January 1993 and September 2025. Publication output increased sharply after 2018, and a limited number of journals and highly cited studies now serve as important publication outlets and intellectual anchors. Deep learning and reinforcement learning dominate the methodological landscape, while transformer-based, explainable, and governance-oriented approaches are emerging.

The network analyses reveal a contrast between the field's intellectual and social structures. Term co-occurrence patterns show increasingly coherent thematic groupings around methodological innovation, trading applications, forecasting, and ethical or governance concerns. By contrast, the co-authorship network remains fragmented, with limited cross-institutional and cross-disciplinary integration. The field therefore appears to be moving towards early intellectual consolidation without a correspondingly integrated collaborative structure.

These findings have implications for both research and practice. Scholars should strengthen the links between AI methods, financial theory, and rigorous empirical validation, with particular attention to reproducibility, transaction costs, model drift, explainability, and robustness across market regimes. Practitioners and policymakers should view AI-enabled trading not only as a source of predictive and operational benefits but also as a source of model, governance, and market-stability risks.

The study is limited by its reliance on the Web of Science Core Collection, the incomplete 2025 observation period, time-sensitive citation counts, and export constraints that restricted co-citation and keyword analyses. Future research should use multiple databases, apply bibliographic coupling and full-text analysis, and examine whether the field's growing methodological coherence is accompanied by stronger interdisciplinary collaboration. Such work will help clarify how emerging AI architectures can be integrated into transparent, robust, and responsible financial decision systems.

Disclaimer

Generative artificial intelligence tools were used solely for language support, including grammatical and stylistic revision, readability improvements, translation of selected passages, and assistance with manuscript organisation. They were not used to generate research data, conduct bibliometric analyses, select references, interpret findings, develop theoretical arguments, or formulate conclusions. All methodological decisions, data-collection procedures, analyses, interpretations, and conclusions were undertaken and validated by the authors, who assume full responsibility for the accuracy, originality, and integrity of the manuscript.

References

Ahi, K., Hsieh, C.-H. and Fenger, G. (2025), "LLMs and LVMs for agentic AI: a GPU-accelerated multimodal system architecture for RAG-grounded, explainable, and adaptive intelligence", in

- Aria, M. and Cuccurullo, C. (2017), "Bibliometrix: an R-tool for comprehensive science mapping analysis", *Journal of Informetrics*, Vol. 11 No. 4, pp. 959-975, doi: [10.1016/j.joi.2017.08.007](https://doi.org/10.1016/j.joi.2017.08.007).
- Bhuiyan, M.S.M., Rafi, M.A., Rodrigues, G.N., Mir, M.N.H., Ishraq, A., Mridha, M.F. and Shin, J. (2025), "Deep learning for algorithmic trading: a systematic review of predictive models and optimization strategies", *Array*, Vol. 26, 100390, doi: [10.1016/j.array.2025.100390](https://doi.org/10.1016/j.array.2025.100390).
- Chen, J.-F., Chen, W.-L., Huang, C.-P., Huang, S.-H. and Chen, A.-P. (2016), "Financial time-series data analysis using deep convolutional neural networks", *2016 7th International Conference on Cloud Computing and Big Data (CCBD)*, IEEE, pp. 87-92, doi: [10.1109/CCBD.2016.027](https://doi.org/10.1109/CCBD.2016.027).
- Chen, X., Xie, H., Li, Z., Zhang, H., Tao, X. and Wang, F.L. (2025), "Sentiment analysis for stock market research: a bibliometric study", *Natural Language Processing Journal*, Vol. 10, 100125, doi: [10.1016/j.nlp.2025.100125](https://doi.org/10.1016/j.nlp.2025.100125).
- Clarivate Analytics (2025), available at: <https://webofscience.com>
- Conegundes, L. and Pereira, A.C.M. (2024), "A methodology for developing deep reinforcement learning trading strategies: a case study in the futures market", *2024 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics (CIFEr)*, IEEE, pp. 1-10, doi: [10.1109/CIFEr62890.2024.10772851](https://doi.org/10.1109/CIFEr62890.2024.10772851).
- Demir, S., Stappers, B., Kok, K. and Paterakis, N.G. (2022), "Statistical arbitrage trading on the intraday market using the asynchronous advantage actor-critic method", *Applied Energy*, Vol. 314, 118912, doi: [10.1016/j.apenergy.2022.118912](https://doi.org/10.1016/j.apenergy.2022.118912).
- Fior, J., Cagliero, L. and Garza, P. (2022), "Leveraging explainable AI to support cryptocurrency investors", *Future Internet*, Vol. 14 No. 9, p. 251, doi: [10.3390/fi14090251](https://doi.org/10.3390/fi14090251).
- Garfield, E. (1964), "Science Citation Index—a new dimension in indexing", *Science*, Vol. 144 No. 3619, pp. 649-654, doi: [10.1126/science.144.3619.649](https://doi.org/10.1126/science.144.3619.649).
- Gómez-Martínez, R., Orden-Cruz, C. and Medrano-García, M.L. (2022), "Quantitative trading using artificial intelligence on trend-following indicators: an example in 2020", *Journal of Investing*, Vol. 32, pp. 35-49, doi: [10.3905/joi.2022.1.235](https://doi.org/10.3905/joi.2022.1.235).
- He, C., Lin, M. and Wang, N. (2021), "Trading on online social mood: a machine learning strategy based on Twitter sentiment", *International Journal of Financial Engineering*, Vol. 8 No. 4, 2141011, doi: [10.1142/S2424786321410115](https://doi.org/10.1142/S2424786321410115).
- Horobet, A., Boubaker, S., Belascu, L., Negreanu, C.C. and Dinca, Z. (2024), "Technology-driven advancements: mapping the landscape of algorithmic trading literature", *Technological Forecasting and Social Change*, Vol. 209, 123746, doi: [10.1016/j.techfore.2024.123746](https://doi.org/10.1016/j.techfore.2024.123746).
- Jiang, Z. and Liang, J. (2017), "Cryptocurrency portfolio management with deep reinforcement learning", *2017 Intelligent Systems Conference (IntelliSys)*, IEEE, pp. 905-913, doi: [10.1109/IntelliSys.2017.8324237](https://doi.org/10.1109/IntelliSys.2017.8324237).
- Karras, A., Theodorakopoulos, L., Karras, C., Krimpas, G.A., Giannaros, A. and Bakalis, C.-P. (2025), "LLM-driven big data management across digital governance, marketing, and accounting: a Spark-orchestrated framework", *Algorithms*, Vol. 18 No. 12, p. 791, doi: [10.3390/a18120791](https://doi.org/10.3390/a18120791).
- Kawy, R.A., Abdelmoez, W.M. and Shoukry, A. (2021), "Financial portfolio construction for quantitative trading using deep learning technique", in *Artificial Intelligence and Soft Computing (ICAISC 2021)*, *Lecture Notes in Computer Science*, Springer, pp. 3-14, doi: [10.1007/978-3-030-87986-0_1](https://doi.org/10.1007/978-3-030-87986-0_1).
- Li, Y., Wu, J. and Bu, H. (2016), "When quantitative trading meets machine learning: a pilot survey", *2016 13th International Conference on Service Systems and Service Management (ICSSSM)*, IEEE, pp. 1-6, doi: [10.1109/ICSSSM.2016.7538632](https://doi.org/10.1109/ICSSSM.2016.7538632).
- Lin, Y.-C., Chen, C.-T., Sang, C.-Y. and Huang, S.-H. (2022), "Multiagent-based deep reinforcement learning for risk-shifting portfolio management", *Applied Soft Computing*, Vol. 123, 108894, doi: [10.1016/j.asoc.2022.108894](https://doi.org/10.1016/j.asoc.2022.108894).

- Park, J.-H., Kim, J.-H. and Huh, J.-H. (2024), “Deep reinforcement learning robots for algorithmic trading: considering stock market conditions and U.S. Interest rates”, *IEEE Access*, Vol. 12, pp. 20705-20725, doi: [10.1109/ACCESS.2024.3361035](https://doi.org/10.1109/ACCESS.2024.3361035).
- Qing, Y., Sun, J., Kong, Y. and Lin, J. (2022), “Fundamental multi-factor deep-learning strategy for cryptocurrency trading”, *2022 IEEE 20th International Conference on Industrial Informatics (INDIN)*, IEEE, pp. 674-680, doi: [10.1109/INDIN51773.2022.9976116](https://doi.org/10.1109/INDIN51773.2022.9976116).
- Sang, C.-Y., Huang, S.-H., Chen, C.-T. and Chang, H.-T. (2025), “Portfolio management using online reinforcement learning with adaptive exploration and multi-task self-supervised representation”, *Applied Soft Computing*, Vol. 172, 112846, doi: [10.1016/j.asoc.2025.112846](https://doi.org/10.1016/j.asoc.2025.112846).
- Sezer, O.B. and Ozbayoglu, A.M. (2018), “Algorithmic financial trading with deep convolutional neural networks: time series to image conversion approach”, *Applied Soft Computing*, Vol. 70, pp. 525-538, doi: [10.1016/j.asoc.2018.04.024](https://doi.org/10.1016/j.asoc.2018.04.024).
- Shen, Z. and Huang, H. (2025), “An adaptive quantitative trading strategy optimization framework based on meta reinforcement learning and cognitive game theory”, *Applied Intelligence*, Vol. 55 No. 10, p. 689, doi: [10.1007/s10489-025-06423-3](https://doi.org/10.1007/s10489-025-06423-3).
- Shynkevich, Y., McGinnity, T.M., Coleman, S.A., Belatreche, A. and Li, Y. (2017), “Forecasting price movements using technical indicators: investigating the impact of varying input window length”, *Neurocomputing*, Vol. 264, pp. 71-88, doi: [10.1016/j.neucom.2016.11.095](https://doi.org/10.1016/j.neucom.2016.11.095).
- Théate, T. and Ernst, D. (2021), “An application of deep reinforcement learning to algorithmic trading”, *Expert Systems with Applications*, Vol. 173, 114632, doi: [10.1016/j.eswa.2021.114632](https://doi.org/10.1016/j.eswa.2021.114632).
- Van Eck, N.J. and Waltman, L. (2010), “Software survey: VOSviewer, a computer program for bibliometric mapping”, *Scientometrics*, Vol. 84 No. 2, pp. 523-538, doi: [10.1007/s11192-009-0146-3](https://doi.org/10.1007/s11192-009-0146-3).
- Wang, H. and Xie, D. (2024), “Optimal profit-making strategies in stock market with algorithmic trading”, *Quantitative Finance and Economics*, Vol. 8 No. 3, pp. 546-572, doi: [10.3934/QFE.2024021](https://doi.org/10.3934/QFE.2024021).
- Wu, J., Wang, C., Xiong, L. and Sun, H. (2019), “Quantitative trading on stock market based on deep reinforcement learning”, *2019 International Joint Conference on Neural Networks (IJCNN)*, IEEE, pp. 1-8, doi: [10.1109/IJCNN.2019.8851831](https://doi.org/10.1109/IJCNN.2019.8851831).
- Ye, Z.J. and Schuller, B.W. (2023), “Human-aligned trading by imitative multi-loss reinforcement learning”, *Expert Systems with Applications*, Vol. 234, 120939, doi: [10.1016/j.eswa.2023.120939](https://doi.org/10.1016/j.eswa.2023.120939).

Further reading

- Cohen, G. (2022), “Algorithmic trading and financial forecasting using advanced artificial intelligence methodologies”, *Mathematics*, Vol. 10 No. 18, p. 3302, doi: [10.3390/math10183302](https://doi.org/10.3390/math10183302).

Corresponding author

Raúl Gómez-Martínez can be contacted at: raul.gomez.martinez@urjc.es