International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 10 | Oct - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Applications of Operations Research in Portfolio Optimization in Finance

A Bibliometric Study, 2017–2026

Dr. Richa Saxena[2], Arham Sanghvi[1], Adavya Atri[1], Arnav Bansal[1], Anvesha Gupta[1], Ahaanaa Girish[1]

[1]BBA, NMIMS, Anil Surendra Modi School of Commerce, Mumbai, India

[2]Assistant Professor, NMIMS, Anil Surendra Modi School of Commerce, Mumbai, India

ABSTRACT: Construction of an optimized portfolio is one of the most important problems in finance and operations research, but the vast amount of literature has not been possible to review in narrative form. The present study is a bibliometric analysis of the literature on portfolio optimization in finance from 2017 to 2026, with a total of 761 published documents (726 articles and 35 reviews) finally retrieved from the Scopus database by the query "Portfolio Optimization AND Finance". The study maps the field by means of a keyword co-occurrence analysis with the help of the tools Biblioshiny/Bibliometrix and VOSviewer to provide insights into the publication and citation trends, the leading authors, the sources and countries and the thematic structure of the field. The growth of the output was at a compound annual rate of 18.56 per cent; 2025 and 2026 together account for 42.2 per cent of the corpus. Despite the significant progress in the literature, it continues to be organized around the Markowitz mean–variance framework, which has been extended to dynamic and multi-period models, metaheuristic algorithms under realistic constraints, machine learning and reinforcement learning, and a new strand of quantum computing. The analysis also describes the constraints of the data set, including the absence of cited-reference information, which made co-citation and bibliographic-coupling analyses impossible, and the fact that a large fraction of the documents retrieved is on the outskirts of the topic.

KEYWORDS: Portfolio optimization; Bibliometric analysis; Operations research; Scopus; Science mapping; VOSviewer; Biblioshiny.

1. INTRODUCTION

Researchers searching for portfolio optimization papers often cannot tell how much of the returned set actually concerns asset allocation and how much is simply what the database has indexed under that label. The field is broad enough that the two are easy to mix.

Portfolio optimization sits at the center of finance and operations research. It became practical once Harry Markowitz showed that investors could choose holdings by trading expected return against risk. Variance does not capture every form of risk an investor faces (Aouni et al., 2018), yet the mean-variance model remains the standard against which later methods are tested.

Computing power, larger datasets, machine learning, and quantum methods have changed the practical work around that model. The single-period, two-variable setup has not disappeared, but shrinkage estimators now reduce estimation error and multi-period models let portfolios be revised over time. Transaction costs, cardinality constraints, and minimum investment sizes make the problem harder to solve. Reinforcement learning and large language models treat allocation as a sequence of decisions rather than a single optimization. Quantum and quantum-inspired algorithms are being tried on larger portfolios. Sustainability, liquidity, and assets such as cryptocurrencies have widened the objective function (Kalayci et al., 2019; Orús et al., 2019; Rout et al., 2026; Mantshimuli & Mwamba, 2025).

A conventional literature review struggles once the literature reaches thousands of papers. Which studies matter, and how they connect, depends on the reviewer's judgment. Bibliometric methods supply counts of publications and citations, maps of co-authorship, and keyword networks. These show which papers and topics have drawn attention and where the literature is thin (Donthu et al., 2021; Zupic & Čater, 2015; Mahmoud & Surwanti, 2026; Manogna & Anand, 2024).

This paper applies that approach to portfolio optimization and records the limits of the data used. The records come from Scopus for the query "Portfolio Optimization AND Finance," years 2017 to 2026. Four records dated 2027, probably early-access items, were dropped, leaving 761 documents: 726 articles and 35 reviews, from 400 sources and 2,039 authors. Output grew at an average annual rate of 18.56 percent. Growth was not steady. Publications fell in 2019, recovered through 2024, then rose sharply in 2025.

Two limits matter. The Scopus export contains no cited-reference data, so co-citation and bibliographic coupling are not possible. The analysis therefore uses publication and citation counts, author and source productivity, and keywords. The broad search also returned 170 of the 761 papers that do not contain the word "portfolio" in the title or author keywords. One of the three most-cited papers concerns synthetic data in healthcare. Discussion of selection and allocation methods therefore stays with papers that address those topics or the algorithms used to solve them.

The paper has three aims: to describe publication trends, leading authors, sources, countries, and citation patterns for portfolio optimization research in Scopus from 2017 to 2026; to map the main themes through keyword co-occurrence and relate them to the classical, dynamic, metaheuristic, machine-learning, and quantum strands in the literature; and to state the gaps and data limits that follow from the broad search and the missing reference fields. These aims produce four questions:

RQ1: What are the publication and citation trends between 2017 and 2026, and what do they show about growth and the concentration of impact?

RQ2: Which countries, authors, and sources account for most of the output, and what does the pattern of collaboration indicate about how the work is produced?

RQ3: What thematic clusters appear in the keyword co-occurrence maps, and how do they line up with the classical, dynamic, metaheuristic, machine-learning, and quantum developments already described?

RQ4: What limits in the dataset and the search strategy constrain the conclusions, and how should later bibliometric studies of this topic be designed?

The rest of the paper covers the methods, a literature review that moves from the mean-variance model through dynamic, machine-learning, and quantum approaches, the bibliometric results and their limits, and a short conclusion.

2. RESEARCH METHODOLOGY

The last collection of 761 documents (726 articles and 35 reviews) is then analysed using bibliometrics in order to investigate the evolution and thematic structure of research related to Portfolio Optimization AND Finance. Bibliometric Analysis is a suitable tool for the purpose of the study since it allows for quantitative analysis of a significant amount of literature rather than performing a narrative review. According to Donthu et al. (2021), bibliometric analysis can be used to analyse the performance and the thematic structure of a particular research field in the terms of various quantitative metrics. In a similar vein, Zupic and Čater (2015) suggests that bibliometric analysis provides an objective and systematic insight into the relationships and the thematic structure within a research area.

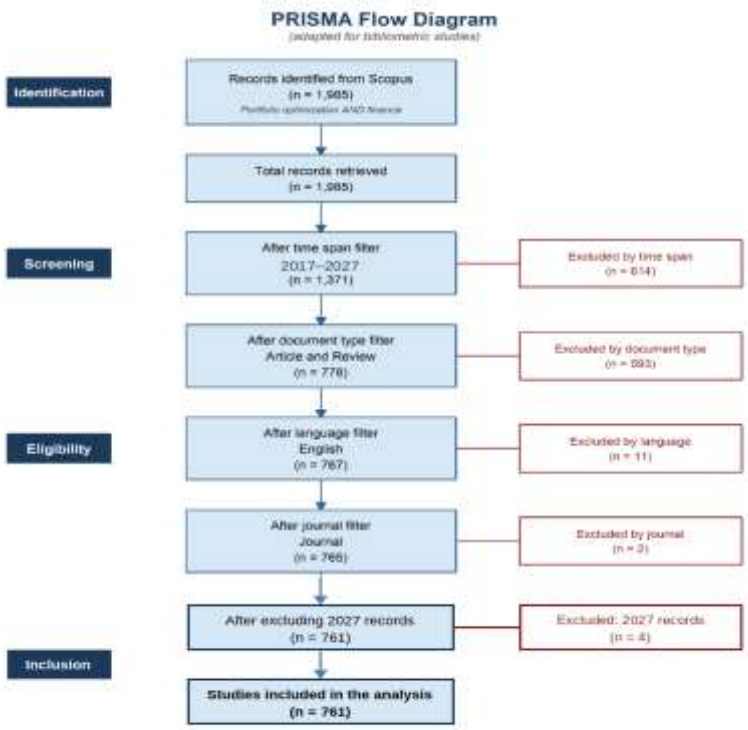
Article Selection and Inclusion and Exclusion Criteria

A search was conducted on 11th September 2026 on Scopus database using the search keywords "Portfolio Optimization AND Finance", and the resulting records were exported on 26th September 2026. The Scopus database has been chosen as the source of bibliographic data since it contains a vast multidisciplinary collection of academic sources as well as bibliographic and citation data for bibliometric analysis (Pranckutė, 2021; Akinlolu et al., 2020). The total number of records found in the Scopus database during the first step was 1985. The records found in this step were filtered according to several parameters such as period of publication, document type, language, and source (Thomé et al., 2016; Ren et al., 2020). The criterion used for time span was from 2017 to 2027. Of the 1985 documents, 1371 documents remained, while 614 documents were removed. Articles and review papers were kept in the study while other documents were filtered out, reducing the total from 1371 to 778 documents and removing 593 documents (Cobo et al., 2018; Goyal & Kumar, 2021). The language filter reduced the records from 778 to 767

documents, excluding 11 documents. The journal filter resulted in 765 publications, excluding 2 documents. In summary, 765 publications met all the required criteria of selection in the research process.

Four records dated 2027, likely early-access items, were then excluded, leaving a final analytical set of 761 documents (726 articles and 35 reviews) covering 2017–2026.

The entire process of article-identification and screening has been represented by the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) flow diagram (Jiang et al., 2019; Ren et al., 2020).



Record screening flow of the Scopus search, adapted from the PRISMA framework.

Bibliometric Analysis

The bibliometric analysis is concerned with trends in publication, citation behavior, authors who have made a significant impact, journals which have had a strong influence, keywords used in this research area, and collaboration in research. Alongside all these factors, co-authorship analysis and co-word analysis will be used. Citation analysis quantifies a publication's acceptance in academia and identifies the authors, organisations, and countries producing the most influential work on a topic (Khanra et al., 2020). Co-word analysis examines the actual content of publications through co-occurrence of keywords to produce a semantic map. Co-authorship analysis examines the collaborative networks among scholars, institutions, and countries (Donthu et al., 2021). Bibliometric data extracted from Scopus database is analysed using the combination of Bibliometrix/Biblioshiny and VOSviewer tools. Microsoft Excel is used for organising the data gathered from the database in case such necessity arises. Bibliometrix is an R based bibliometric and science mapping tool (Aria & Cuccurullo, 2017). VOSviewer is applied in visualisation of bibliometric networks and enables identification of relationships among authors, keywords, and publications (Strozzi et al., 2017; Van Eck et al., 2010).

3. LITERATURE REVIEW

Portfolio optimization is the process of allocating assets to various assets available to the investor so that he can balance the risk and return of his portfolio in an optimum manner. Since Markowitz introduced his mean–variance model in 1952, portfolio selection has been seen as a return–variance trade-off. It has been the standard model ever

since and although it has been criticized it is still the most widely used. It has been criticized for being unable to consider factors other than expected return and variance (Aouni et al., 2018).

There have been many recent studies that have created numerous methods. Researchers have evolved beyond the mean–variance framework, thanks to advances in computing, the increased amount of data and new methods like machine learning, metaheuristics and quantum computing. Other areas of focus include transaction costs, liquidity, sustainability and personalization. This paper provides descriptive statistics of the bibliographic database, it explores the main research themes, and reviews the different streams of portfolio optimization, with an eye on research gaps and future directions.

The publication has been reviewed based on a Bibliometrix export of 761 documents that are available from 2017 to 2026, of which 726 are research articles and 35 are reviews. The collection contains 10,420 citations, which average out to ~13.69 citations per document, with a median of 3 citations per document, meaning that influence lies in a few highly cited documents. The keyword co-occurrence reveals areas of research that are closely related: the main topic is investments and portfolio optimization, whereas the areas of research are risk, computational methods, data-driven finance, and quantum technologies. These patterns serve as a framework for the discussion to follow.

The papers discussed below were selected not so much for their number of citations as to illustrate each theme, as there was no cited-reference data included in the export for a citation network. The inclusion and hence the selection are thus indicative of the authors' opinion of representative contributions and other relevant studies may not have been included.

This is a substantial area: data-driven methods and fintech, which includes machine learning, deep learning, reinforcement learning, learning systems, electronic trading, and decentralized finance, as well as blockchain, financial markets and systematic reviews, representing the artificial-intelligence and financial-technology part of the field.

Algorithmic and computational optimization is a second area: constrained, multi-objective, and evolutionary optimization, metaheuristic and heuristic methods, swarm intelligence, NP-hard problems, costs, and transaction costs.

Investment decision-making is still a key challenge, as in the classical decision-making theory, related to the choice of investments, regression analysis, downside risk and optimization.

Another constant theme is risk; risk management, risk assessment, risk analysis and risk analysis applications are often taken up together with finance, optimization, artificial intelligence, stochastic systems and others, indicating the significance of managing uncertainty in the decision of the portfolio.

Quantum and heuristic computing (quantum computing, quantum algorithms and theory; computation theory; heuristic algorithms) is a newer and focused field that is expanding but has yet to forge links with the portfolio methods.

Investments, portfolio optimization, decision making, metaheuristic and constrained optimization are keywords that often appear alongside the keywords from machine learning, fintech and risk research; additionally, risk analysis overlaps with optimization, AI, quantum computing and mean-variance optimization. The literature can be seen to have three common themes: classical optimization (mean-variance approach, constraints, metaheuristics), risk analysis, and data-driven or quantum technologies. Quantum computing is still in its infancy, and more frequently mentioned in the context of risk and known portfolio problems. These themes are followed in the sections below.

1. Classical Foundations and the Mean-Variance Framework

Markowitz's mean-variance model is the foundation, with several syntheses of its extensions in the data set. (Kalayci et al., 2019)review 175 operations-research papers on deterministic mean-variance models, covering variants arising from real-life conditions and the solution methods tested. (Gunjan & Bhattacharyya, 2023a, 2023b)widen the view by compiling classical, statistical, intelligent and quantum-inspired techniques in a comparative study. (Aouni et al., 2018) review exact methods for problems with criteria beyond mean and variance, and argue that multiple-criteria decision aid is increasingly important for portfolio selection.

The practical weakness of mean-variance optimization is estimation error, which becomes serious when the number of assets is large relative to the sample. (Bodnar et al., 2018)addresses this with a shrinkage-type estimator of the global minimum variance portfolio in high dimensions, using random matrix theory with only weak assumptions about return distributions.

A related stream combines investor views with market equilibrium. (Kolm et al., 2021) describe Black-Litterman as one of the most popular models in quantitative finance, bringing Bayesian statistics into portfolio construction and letting managers express subjective views in a structured way. (Kocuk & Cornuéjols, 2020) extend the idea by modelling returns as a mixture of normals and building portfolios that minimise conditional value-at-risk (CVaR) while incorporating Black-Litterman views. Classical theory has thus been refined rather than discarded.

2. Multi-Period, Dynamic and Personalised Optimization

A second set of papers relaxes the single-period assumption. (Li et al., 2022a)apply model predictive control to multi-period optimization under mean-variance and risk-parity objectives; in out-of-sample tests over 2006 to 2020, multi-period models outperform single-period ones when market impact costs are present, and risk-parity reaches a Sharpe ratio of 0.97 against 0.64 for mean-variance. These figures come from a single study over one historical window and should be read as indicative rather than general.(Capponi et al., 2022)introduce a robo-advisor that solves an adaptive mean-variance problem as the client's risk profile changes. The optimal strategy has myopic and intertemporal hedging terms, and the advisor must trade off obtaining client information promptly against counteracting behavioural biases in the stated risk profile. This links portfolio optimization to behavioural finance and fintech.

3. Metaheuristics, Multi-Objective Optimization and Real-World Constraints

Once realistic constraints such as cardinality limits, minimum lots, and transaction costs are introduced, the problem becomes NP-hard and exact solution is often impractical, hence the attention to computational methods. (Soler-Dominguez et al., 2018)survey metaheuristics across finance, with applications including rich portfolio optimization, index tracking, enhanced indexation, and credit risk. (Doering et al., 2019)review metaheuristics for NP-hard portfolio optimisation and risk management, propose a unified classification linking the two, and outline likely future trends.

Transaction costs have attracted particular attention.(Meghwani & Thakur, 2018) formulate a tri-objective model (risk, return, transaction cost) with cardinality and self-financing constraints, compare variance, value-at-risk and CVaR as risk measures, and propose a repair algorithm for equality constraints in multi-objective evolutionary algorithms. (Leal et al., 2020) use a bilevel formulation in which the financial intermediary sets transaction costs while the investor chooses the portfolio, and(Czichowsky et al., 2018) study optimisation under proportional transaction costs with fractional Brownian motion. The most recent benchmark in the data set, (Hidalgo-Marín et al., 2026), compares NSGA-II, MOEA/D, SMS-EMOA and SMPSO on real data for over 11,000 U.S.-listed assets and reports clear trade-offs between return and risk and between algorithms as dimensionality and constraint complexity rise. These studies show a movement from stylised problems to benchmarked, high-dimensional, realistically constrained ones.

4. Machine Learning, Reinforcement Learning and Sentiment

Artificial-intelligence methods have become increasingly prominent in portfolio research.(Hambly et al., 2023)survey reinforcement learning in finance, noting that, unlike classical stochastic control, it can exploit large amounts of data with fewer model assumptions.(Ferreira et al., 2021)review artificial intelligence applied to stock trading, and(Sonkavde et al., 2023)systematically review machine- and deep-learning models for forecasting stock prices.

For portfolio allocation, the literature increasingly links reinforcement learning to classical risk-return ideas. (Aboussalah et al., 2022) combine reinforcement learning for dynamic asset allocation with the mean-variance paradigm. (Choudhary et al., 2025) address reward-function design by training three deep reinforcement-learning agents on log returns, differential Sharpe ratio and maximum drawdown, fusing their actions with a convolutional network into one risk-adjusted policy. (Sattar et al., 2025) propose an RMS-driven deep reinforcement-learning framework for portfolio management. A newer strand uses language models: (Mantshimuli & Mwamba, 2025) combine sentiment from three finance-tuned large language models with an LSTM and feed the score into Black-Litterman as investor views, reporting an annualised return of 31.22% against 24.57% for a market-capitalisation-

weighted S&P 500 portfolio. This is a single backtest, and language models trained on broad historical text carry look-ahead risk on past market data, so the margin should be interpreted cautiously.(Weber et al., 2024) add a caution by reviewing explainable artificial intelligence in finance: black-box allocation models raise transparency questions. (Mahmoud & Surwanti, 2026) provide a bibliometric view of artificial intelligence in investment decision-making, consistent with the growth documented here.

5. Quantum Computing

Quantum computing is a distinctive and emerging theme. The most cited directly relevant document is(Orús et al., 2019), who explain how quantum annealers can optimise portfolios and quantum amplitude estimation could speed up Monte Carlo sampling for pricing and risk analysis.(Mugel et al., 2022)test dynamic portfolio optimization on real daily prices for 52 assets over eight years, comparing classical solvers, D-Wave hybrid quantum annealing, IBM-Q variational quantum eigensolvers and, for the first time here, a quantum-inspired tensor-network optimizer on Sharpe ratios, profits and computing times.(Alcazar et al., 2020)compare classical restricted Boltzmann machines with quantum circuit Born machines on scenarios from a probabilistic portfolio problem.(Khan et al., 2021)propose Quantum Beetle Antennae Search, a metaheuristic for constrained portfolio selection, merging quantum ideas with the metaheuristic tradition.(Naik et al., 2025)systematically review quantum computing in finance, covering portfolio optimization, fraud detection, derivative pricing, and blockchain effects. The evidence is promising but largely experimental; practical advantage on realistic portfolio sizes is not yet established.

The investable universe and objective function are also being redefined. On cryptocurrencies, (Trimborn et al., 2020) propose a liquidity-bounded risk-return approach and find that adding cryptocurrencies in tradable amounts improves the risk-return trade-off of portfolios of stocks, bonds and commodities. (Ma et al., 2020) similarly report that adding five cryptocurrencies to four traditional portfolios raised returns in most cases and reduced volatility in all of them. (Ali et al., 2024) extend the diversification question to green cryptocurrencies.

Sustainability enters through both constraints and objectives. (Gasser et al., 2017) modify Markowitz's theory with a social-responsibility measure, producing a three-dimensional capital-allocation surface across return, risk and responsibility. (Anquetin et al., 2022) add a carbon penalty to constrained mean-variance optimization and find emission intensity can be cut by at least half with virtually no loss in Sharpe ratio. (Akhtaruzzaman et al., 2023) find that a portfolio including a green asset performs better than the benchmark over short horizons, and(Lalwani, 2024) finds that portfolios combining green asset ETFs with broad indices achieve higher returns and Sharpe ratios, even after transaction costs and excluding COVID-19. These studies suggest sustainability-aware optimization need not be costly, though results depend on period, data provider and horizon.

Several conclusions emerge. First, mean-variance optimization remains the organising framework, but nearly every stream modifies it: robust or shrinkage estimators, Bayesian views, additional objectives, realistic constraints, multi-period dynamics or new risk measures such as CVaR. Second, the computational frontier has three waves: metaheuristics, deep reinforcement learning and language models, and quantum or quantum-inspired solvers. Third, optimization objectives are expanding to include liquidity, transaction costs, carbon exposure and social responsibility.

Promising avenues include standardised, open benchmarks with transaction costs and constraints, explainable and risk-aware reinforcement learning allocators, hybrid classical-quantum solvers tested on realistic dimensions, robust estimation for high-dimensional portfolios, and models that treat sustainability, liquidity and investor behaviour as first-class elements of the optimization problem rather than afterthoughts.

The portfolio optimization literature is growing, with output about four times higher in 2025 than in 2017, and is becoming more multidisciplinary. Its core remains the Markowitz risk-return trade-off, expanded through constraints, dynamics, artificial intelligence, quantum computing, and sustainability, with investments and portfolio optimization at the center of a wider conversation spanning risk, computational optimization, fintech, machine learning, and quantum computing. The next frontier likely involves not only better optimizers but also more rigorous testing and closer attention to real-world investor considerations.

4. DATA ANALYSIS AND RESULTS

The results below report the performance of the corpus and the structure of its keyword network. Screening, the Scopus query "Portfolio Optimization AND Finance" and the PRISMA counts are set out in the methods. The analysed tables use the Biblioshiny extract of that screened set after four 2027 records are dropped, leaving 761 documents from 2017 to 2026 (Scopus file dated 26 September 2026). Interpretation is limited to this extract.

Main Information about the Data

The indicators in Table 1 follow the staged science-mapping workflow of Zupic and Čater (2015), Aria and Cuccurullo (2017, 2026) and Aria, Cuccurullo, D'Aniello and Spano (2026), with performance analysis preceding science mapping (Donthu et al., 2021). Co-citation and bibliographic coupling cannot be run because every cited-reference field is empty, an export setting rather than a limit of Scopus. The set holds 761 English-language documents from 400 sources by 2,039 authors: 726 articles and 35 reviews, with 34 records tagged article in press (12 from 2025, 22 from 2026).

At 1.90 documents per source the journal base is broad rather than concentrated; the European Journal of Operational Research leads with 38 documents (5.0 per cent). Biblioshiny's 2.98 co-authors per document is the mean number of authors. Only 107 documents (14.1 per cent) are single-authored, and international co-authorship is 30.75 per cent.

Table 1. Main information about the filtered dataset (761 of 765 documents) as reported in Biblioshiny

Description	Results
Timespan	2017:2026
Sources	400
Documents	761
Annual growth rate %	18.56
Average citations per document	13.69
Cited references in the export	0
Index keywords (ID)	2,650
Author keywords (DE)	2,395
Authors	2,039
Co-authors per document	2.98
International co-authorships %	30.75
Records lacking index keywords	388 (51.0%)

Average citations (13.69) are incoming Scopus counts, unadjusted for time: the mean falls from 35.1 in 2017 to 0.9 in 2026. Index keywords are missing on 388 of 761 records (51.0 per cent), stably across years (40 per cent in 2017, 51 per cent in 2026), so co-occurrence maps built on merged keywords rest on an incomplete indexer field.

Citations are skewed: the median is 3 and 161 documents (21.2 per cent) have none. The three most-cited records are a 2019 review of quantum computing for finance (612 citations; Orús et al., 2019), a 2023 paper on synthetic data in healthcare (370; Giuffrè & Shung, 2023) and a 2020 study of stock-market reaction to the Paris Agreement (268; Monasterolo & de Angelis, 2020). The second has no portfolio content in its title or author keywords, and 170 records (22.3 per cent) contain no "portfolio" string in either. Because the query joined portfolio optimisation to finance with no further topical screen, the citation core is broader than the research problem, partly as a property of the search rather than of the field.

Annual Scientific Production

Table 2. Annual scientific production and mean incoming citations, 2017 to 2026

Year	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026
Documents	35	41	32	41	50	68	76	97	159	162
Mean citations	35.1	18.7	35.3	28.3	17.8	21.2	26.0	10.5	4.2	0.9

Output grows from 35 documents in 2017 to 162 in 2026, a compound annual growth rate of 18.56 per cent, $(162/35)^{1/9} - 1$. The only contraction is in 2019 (nine documents, 22.0 per cent) and the largest gain in 2025 (62 documents, 63.9 per cent). 2026 is incomplete, since the file was built in late September, yet it already matches 2025. Dropping 2027 as incomplete while keeping 2026 is inconsistent.

A log-linear OLS model on the ten annual counts gives a year coefficient of 0.1878 (SE 0.0210, $p < 0.001$; $R^2 = 0.909$), an implied annual growth of $\exp(0.1878) - 1 = 20.7$ per cent, which differs from the CAGR. Durbin–Watson is 1.17, so residual autocorrelation is present, and Poisson dispersion of 2.22 breaks the variance assumption; a quasi-Poisson fit gives 0.206. Dropping 2026 as well gives a slope of 0.181 ($R^2 = 0.875$) and a CAGR of 20.8 per cent. With $n = 10$ and a 2025 step, these fits only describe the observed series.

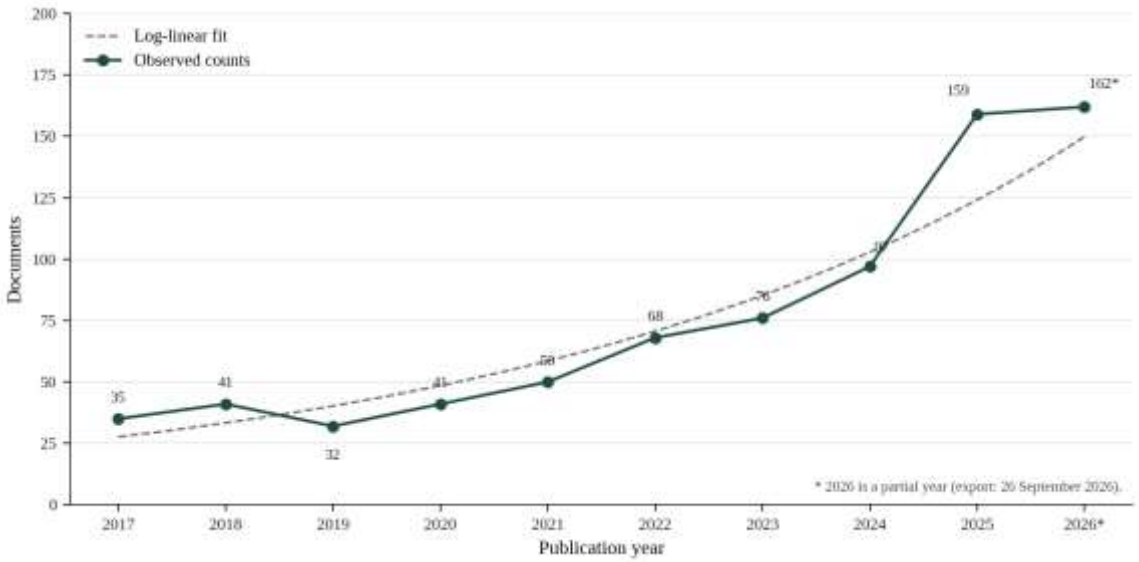


Figure 1. Annual scientific production, 2017–2026, with observed counts and fitted log-linear trend. 2026 is a partial year.

Two papers from the first three years top a young corpus: Metawa, Hassan and Elhoseny (2017), genetic algorithms for bank lending (224 citations), and Kalayci, Ertenlice and Akbay (2019), deterministic mean-variance models (162). Singh et al. (2022), a reinforcement-learning review (225), is not early and is not shown to supply the learning vocabulary of the later maps.

Three-Field Plot

Figure 2 maps authors, author countries and merged keywords with the threeFieldsPlot function in Biblioshiny (Aria & Cuccurullo, 2017, 2026). Each author is joined to the countries on their documents, and each country to those documents' keywords; rectangle size follows frequency and line thickness association strength. Author names are as exported, homonyms were not disambiguated, overlapping labels were not redrawn, and "machine learning" and "machine-learning" were not merged.

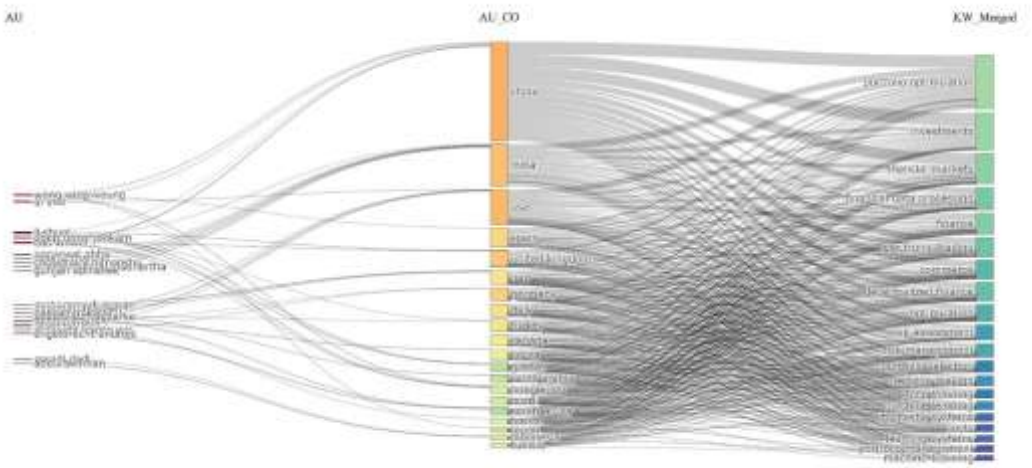


Figure 2. Three-field plot of authors (left), author countries (centre) and merged keywords (right).

The left field is less dispersed than it appears: Li Shuai (eight documents) and Cao Xinwei (seven) share affiliations across China, Hong Kong, Finland, the United Kingdom, Greece and Serbia, and Qi Yue (seven) publishes from Nankai University, so several names count teams as well as individuals.

Table 3. Leading authors in the extract. Countries are those listed on the papers

Author	Documents	Countries on their papers
Li Shuai	8	China, Hong Kong, Finland, UK, Greece, Serbia
Cao Xinwei	7	China, Hong Kong, Finland, UK, Greece, Serbia
Qi Yue	7	China; United States
Mohammadi Emran	6	Iran; Viet Nam
Oben Remy Jonkam	5	Turkey, India, Oman, Pakistan

Connection load in the middle field is heaviest at China, then India and the United States. As keywords are reachable only through the country layer, the thematic structure in Section 4 is sorted by geography before subject.

Country Scientific Production

Table 4 reports Biblioshiny's country production: 78 countries, frequencies summing to 1,708. This is not a tally of the 1,756 C1 address tokens or a count of unique countries per paper. A paper is multi-country when the last token of two or more C1 strings differs, giving 237 multi-country and 524 single-country documents. Hong Kong appears in Table 3 but is folded into China in Table 4.

Table 4. Leading countries, 2017 to 2026 (author-affiliation appearances)

Rank	Country	Appearances	Share of 1,708
1	China	300	17.6%
2	United States	185	10.8%
3	India	135	7.9%
4	United Kingdom	72	4.2%
5	Germany	57	3.3%
6	Spain	55	3.2%

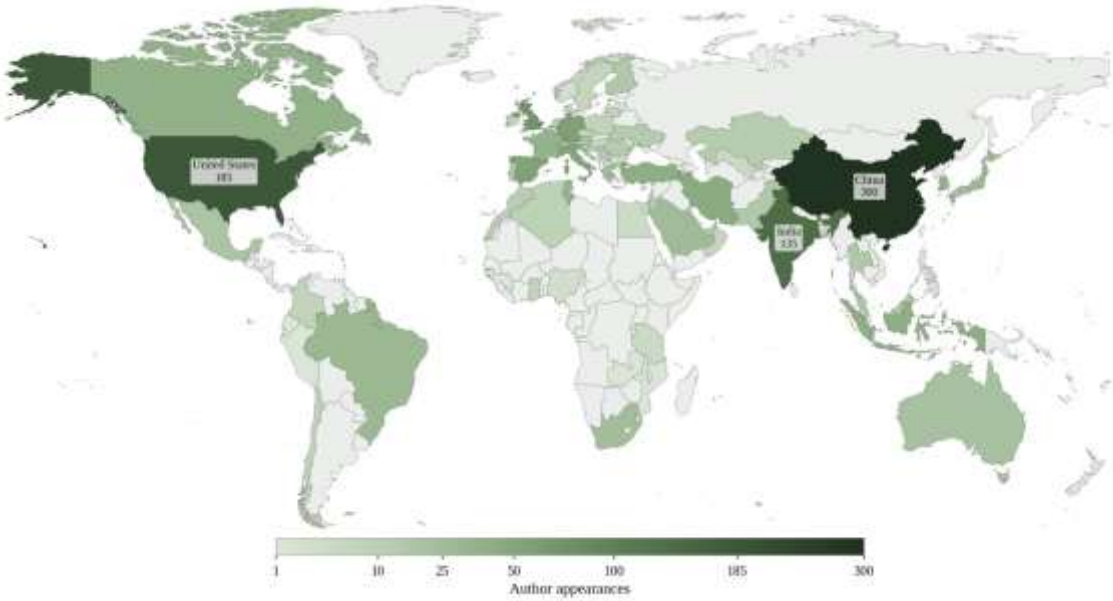


Figure 3. Country scientific production by author appearances, 2017 to 2026, on a power scale (γ = 0.50). Countries with no documents are unshaded.

China accounts for 300 appearances (17.6 per cent), then the United States (185) and India (135); the three hold 36.3 per cent, the top five 43.9 and the top ten 57.5. Multi-country papers (237, 31.1 per cent) average 22.5 citations (median 5) and hold 51.1 per cent of all citations; single-country papers (524) average 9.7 (median 3). Mann–Whitney tests give $p = 1.2 \times 10^{-4}$ on raw counts and 1.1×10^{-9} on within-year percentile ranks. Reviews confound this, averaging 55.6 citations against 11.7 for articles, so the advantage is neither only a year nor only a review effect, though counts remain unnormalised.

Keyword Co-occurrence

Three VOSviewer maps (van Eck & Waltman, 2010) were built at minimum occurrences of 5, 7 and 9 on merged keywords (author plus index) with no thesaurus. Author keywords alone split “portfolio optimization” (249 documents) from “portfolio optimisation” (30). A recount gives 250, 160 and 117 merged terms at the thresholds, and 84, 46 and 29 author-keyword terms. Software version and clustering parameters were not recorded. Rectangles are keywords sized by frequency and coloured by software-assigned cluster; clusters at one threshold are not the same objects as at another, and overlap was not computed.

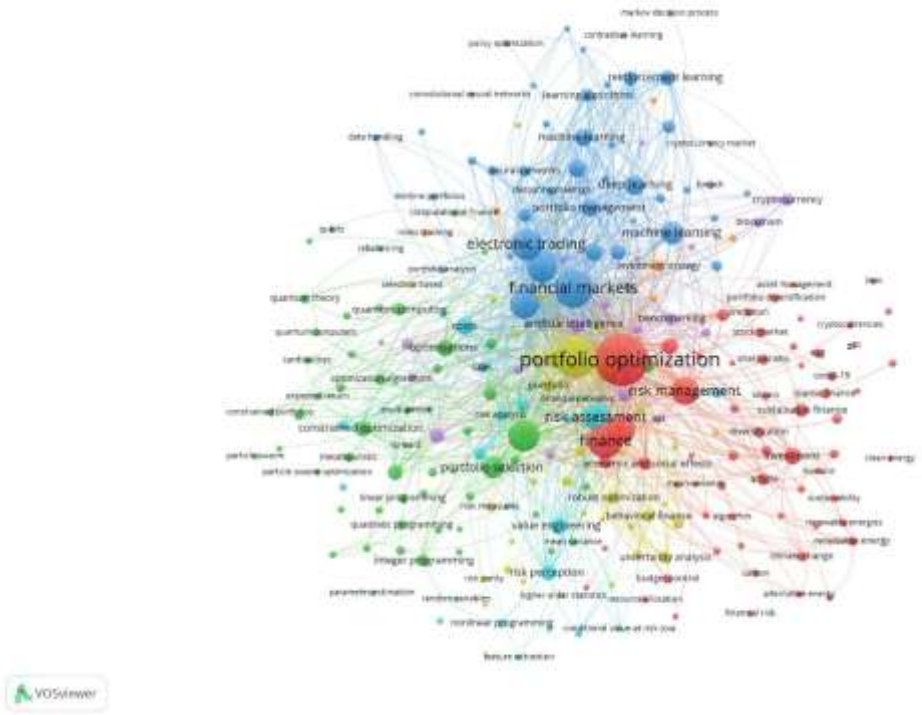


Figure 4. Keyword co-occurrence network, minimum five occurrences (252 terms, seven clusters).

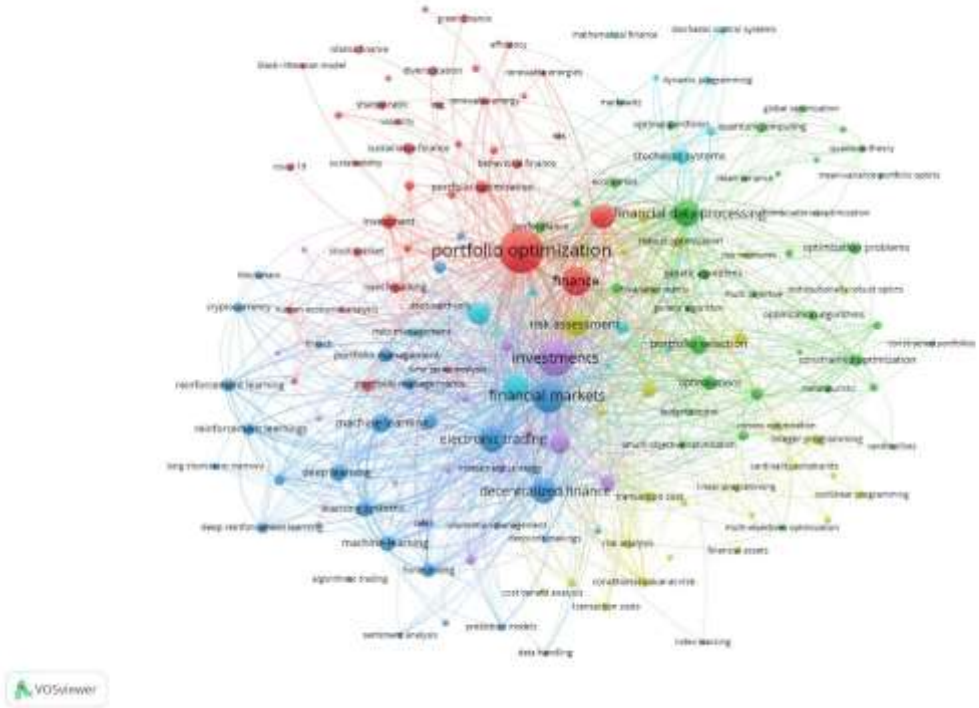


Figure 5. Keyword co-occurrence network, minimum seven occurrences (161 terms, six clusters).

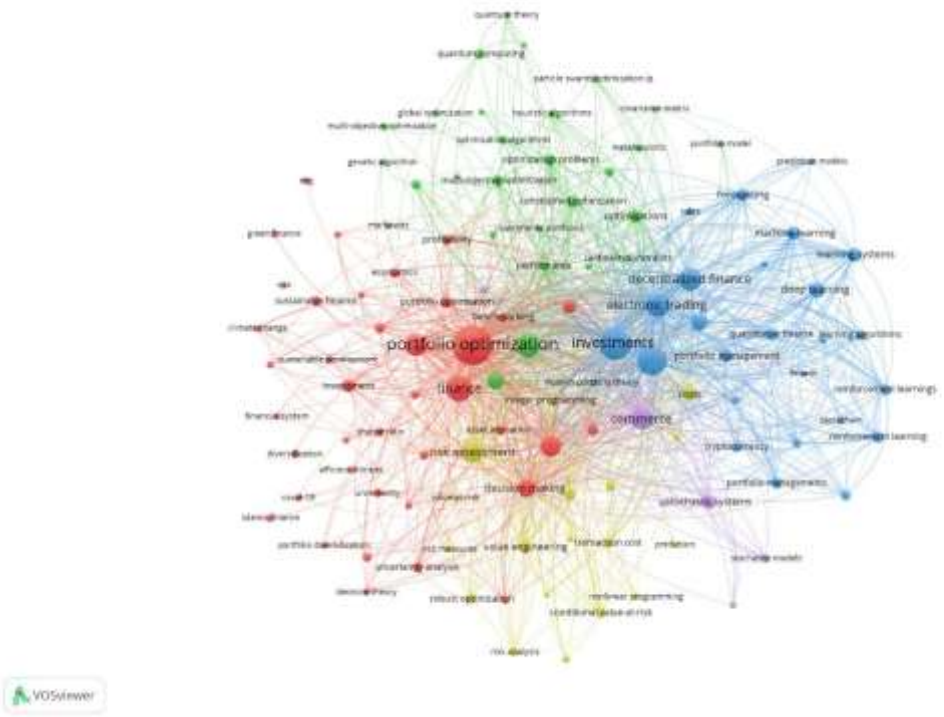


Figure 6. Keyword co-occurrence network, minimum nine occurrences (118 terms, five clusters).

Table 5. Cluster sizes at three co-occurrence thresholds

Minimum occurrences	Terms	Clusters	Cluster sizes, descending
5	252	7	63, 56, 46, 36, 22, 21, 8
7	161	6	38, 37, 30, 27, 16, 13
9	118	5	39, 28, 26, 19, 6

At five occurrences the seven groups are: a core of portfolio optimisation, finance, risk management and sustainable finance (63 terms); operations research, with constrained optimisation, genetic algorithms and quantum computing (56); markets and learning (46); decision and uncertainty (36); artificial intelligence and multi-objective methods, with cryptocurrency (22); risk measurement (21); and investment strategy (8).

Portfolio optimization is the top keyword at 321 documents in every map, a product of the search string that puts it at the hub of any network from this export; removing it would be the cleaner test. Indexer strings (article, commerce, financial data processing, electronic trading) also survive every cut but are absent from half the records.

Index keywords assign “decentralized finance” to 100 documents, 99 from 2024–2026, yet only seven titles and 12 author-keyword records mention DeFi, and the overlap with DeFi, blockchain or crypto mentions is 10. A paper can enter the 2025.4 mean-year figure without discussing DeFi, so the post-2024 learning-and-market block is partly tagging.

Operations-research vocabulary (constrained optimisation, genetic algorithms, particle swarm) has an earlier mean publication year than learning vocabulary (machine, deep and reinforcement learning). Sustainability terms sit inside the core cluster, and country names almost never become keywords despite skewed production. These describe how this export is tagged, not whether studies use local assets or a global book, and they do not survive a rebuild on author keywords alone.

Re-exporting with cited references would allow co-citation and bibliographic coupling (Donthu et al., 2021), and rebuilding the maps on an author-keyword thesaurus without the query term would show how much of the structure is an indexing artefact.

5. FUTURE RESEARCH AGENDA

The purpose of the present review was to trace the developments of the field of portfolio optimization research, the authors of this research and the thematic structure of this field under a general finance-type search. The evidence from the 761 documents, read alongside the qualitative literature, shows a field that is expanding quickly but unevenly. Output grew at an annual rate of 18.56% between 2017 and 2026, yet 2025 and 2026 alone account for 42.2% of the corpus, and the citation impact is concentrated in a small number of older papers. Notwithstanding this growth, several gaps remain, some arising from the field itself and some from the way the field is captured in bibliographic databases. The agenda below is therefore organized into six areas, beginning with the methodological and ending with the substantive.

1. Strengthening the bibliometric foundation

The first gap concerns the evidence base of this review and of similar reviews. Because every cited- reference field in the Scopus export was empty, co-citation analysis and bibliographic coupling (Small, 1973; Kessler, 1963) could not be carried out, and the thematic structure of the field was inferred only from keywords. Future researchers should re-export the data with cited references, so that co-citation can reveal the foundational works and bibliographic coupling can map the current research front (Donthu et al., 2021; Zupic & Čater, 2015). Further, the search string "Portfolio Optimization AND Finance" retrieved 170 documents with no portfolio term in the title or author keywords,. A more focused query, restricted to title, abstract and author keywords, would give a cleaner picture of the field properly.

2. Collaboration and geographic diversity

China (300 appearances), the United States (185) and India (135) together account for 36.3% of author- affiliation appearances, and the leading authors appear to belong to recurring transnational teams rather than to work independently. Multi-country papers carry a mean of 22.5 citations against 9.7 for single- country papers, a pattern consistent with the wider literature on international co-authorship (Glänzel & Schubert, 2001). The research on emerging and non-Western markets is limited, such as the research on Islamic and conventional indices (Abu-Alkheil et al., 2017) and on financial sector assets such as fintechs in China (Li & Wang, 2026). Future research could therefore pursue collaborations between leading and emerging contributors, and test whether portfolio models validated on United States data hold in markets with different liquidity, regulation and investor behaviour.

3. Reinforcement learning, language models and explainability

The operations research vocabulary has a later mean publication year than the learning vocabulary (machine learning, deep learning, reinforcement learning) as the latter is a separate block in the keyword maps. From reinforcement learning as a trading device (Hambly et al., 2023), to risk-aware allocation such as multi-reward agents (Choudhary et al., 2025), CVaR-constrained safe agents with action repair (Choudhary et al., 2026) to hybrid of RL and evolutionary computation (Yadav et al., 2025). The next is language models which are used to generate Black–Litterman views directly from multi-model sentiment

(Mantshimuli & Mwamba, 2025), to evolve strategies (Zeng et al., 2025), or to directly produce the portfolio weights (Mantshimuli & Muteba Mwamba, 2026).

Despite this advancement, it is recognized that there is a lack of a strict comparison. Deep reinforcement learning studies are reported to be diverse in data (Rezaei & Nezamabadi-pour, 2025), reward design and evaluation (Rezaei & Nezamabadi-pour, 2025), and systematic studies to have highlighted the lack of attention to interpretability (Sutiene et al., 2024; Weber et al., 2024). Future studies should, therefore, (i) perform a comprehensive evaluation of learning-based allocators net of transaction costs and on out-of- sample windows of several market regimes, (ii) compare with the classical mean-variance and Black- Litterman benchmarks on the same set of data (Kolm & Ritter, 2025), and (iii)

incorporate explainability into the allocation model itself rather than as a post-hoc analysis (Gunjan & Bhattacharyya, 2025).

4. Quantum and quantum-inspired optimization

In this corpus, quantum computing is an emerging theme, in so far as it is part of a theme cluster that is shunted to the periphery at higher keyword thresholds and is only linked with risk and mean-variance at wider thresholds. The available evidence is encouraging but mainly experimental, focusing on quantum annealers and variational circuits on real data (Mugel et al., 2022), on a practical guide for real quantum devices (Buonaiuto et al., 2023), on counterdiabatic algorithms (Hegade et al., 2022), on one quantum algorithm with analyzed data input/output (Rebentrost & Lloyd, 2024), and on quantum-inspired metaheuristics (Khan et al., 2021; Gunjan & Bhattacharyya, 2024). Future research should therefore test hybrid classical–quantum solvers at realistic dimensions with cardinality and transaction-cost constraints, report time-to-solution and solution quality against strong classical baselines, and clarify where quantum- inspired methods, which run on conventional hardware, offer the more immediate gain.

5. Benchmarking metaheuristics under realistic constraints

Once cardinality limits, minimum lots and transaction costs are introduced, the problem becomes very hard, which explains the strong operations research cluster of constrained optimization and genetic algorithms (Chang et al., 2000; Doering et al., 2019; Meghwani & Thakur, 2018). Recent work shows a movement from stylized to benchmarked problems, such as the comparison of four multi-objective algorithms on more than 11,000 listed assets (Hidalgo-Marín et al., 2026a) and its supporting open-source tool (Hidalgo-Marín et al., 2026b). Hence, future studies should adopt shared, openly available datasets and standard metrics, so that algorithms can be compared across the same constraints and should add frictions such as taxes and rebalancing rules that are seldom modelled together.

6. Widening the objective function: sustainability

Sustainability terms sit inside the core portfolio cluster rather than forming a separate theme, indicating that they have entered optimization as constraints and objectives. Studies add carbon penalties (Anquetin et al., 2022), green assets (Akhtaruzzaman et al., 2023; Lalwani, 2024), social-responsibility dimensions (Gasser et al., 2017), fuzzy treatments of scarce sustainability data (Yadav et al., 2023) and robust handling of ESG uncertainty (Li et al., 2022b).

6. CONCLUSION AND LIMITATIONS

The analysis covers 761 of the 765 records retrieved from Scopus on portfolio optimization between 2017 and 2026. Publications rose from 35 in 2017 to 162 in 2026, an average annual growth of 18.56 percent, though the path was uneven. For RQ1, output dropped 22 percent in 2019, recovered through 2024, and jumped 63.9 percent in 2025 to 159 papers. The years 2025 and 2026 together hold 42.2 percent of the corpus. 2026 is incomplete; 22 of its papers remain in press. Mean citations fall from 35.1 per paper in 2017 to 0.9 in 2026 because recent papers have had less time to be cited. The median is 3 citations. 161 papers (21.2 percent) have none, so impact is concentrated in a smaller set of older work.

Of the 761 papers, 170 do not contain “portfolio” in the title or author keywords. One of the three most-cited items is a 2023 healthcare paper on synthetic data (370 citations) that is not about portfolio optimization. The other two are a 2019 review of quantum computing in finance (612 citations) and a 2020 stock-market study (268 citations). Both touch related topics but are not central to portfolio selection. Citation totals therefore describe the retrieved set, not portfolio optimization alone. Comparisons with studies that used narrower searches should be made with that difference in mind.

For RQ2, the work is spread across many countries yet concentrated in a few. The set contains 2,039 authors and 400 sources, with 2.98 authors per paper on average. Only 14.1 percent of papers are single-authored; 30.75 percent have international co-authors. China accounts for 300 of 1,708 author-affiliation appearances (17.6 percent), the United States for 185 (10.8 percent), and India for 135 (7.9 percent). These three countries make up 36.3 percent of appearances; the top ten make up 57.5 percent. Li Shuai (eight papers), Cao Xinwei, and Qi Yue (seven each) appear in recurring international teams. Multi-country papers (237, or 31.1 percent) average 22.5 citations against 9.7 for

single-country papers, a statistically significant difference. Reviews average more citations (55.6 versus 11.7), but that difference does not fully account for the collaboration effect.

For RQ3, keyword co-occurrence at minimum frequencies of 5, 7, and 9 produced similar themes, with the number of clusters falling from seven to five as the threshold rose. The largest clusters center on portfolio optimization, finance, risk management, and sustainable finance; on operations-research methods such as constrained optimization, genetic algorithms, and quantum computing; and on financial markets, electronic trading, and machine learning. Smaller clusters cover decision-making, artificial intelligence, multi-objective optimization, risk measurement, and investment strategies. These map onto the literature review: mean-variance optimization remains the base, with later work adding optimization under uncertainty, metaheuristics, reinforcement learning, and quantum computing. Quantum computing still forms a distinct cluster at the stricter thresholds but links to risk and mean-variance topics at the broader ones (Cui et al., 2022; Soler-Dominguez et al., 2018; Orús et al., 2019; Singh et al., 2022).

“Decentralized finance” appears as an index keyword in 100 papers, 99 of them from 2024 to 2026. Only 10 of those papers mention DeFi, blockchain, or cryptocurrency in the title or author keywords. Part of the recent rise in that term may therefore come from indexing rather than from a shift in research focus.

For RQ4, the main limits come from the export. All cited-reference fields are empty, so co-citation analysis and bibliographic coupling cannot be done. A re-export that includes references would open those methods. Index keywords are missing for 388 papers (51 percent), so the keyword maps cover only about half the set. The broad search also means the results do not map cleanly onto portfolio optimization. Restricting the analysis to author keywords and dropping the main search term would tighten the focus. The treatment of incomplete years is inconsistent: four 2027 records were removed, yet all 162 papers from 2026 were kept, including 22 still in press. Papers from 2025 and 2026 have had less time to accumulate citations, so their lower counts do not by themselves indicate lower quality.

The counts and maps still show how the literature under this search has grown, which countries and authors contribute most, and how the keyword clusters track the methodological shifts described in the literature review. Later bibliometric work on the topic would gain from a narrower query, an export that includes cited references, and conclusions that stay within what the chosen data can support. The present study acknowledges these limits.

REFERENCES

1. Aboussalah, A. M., Xu, Z., & Lee, C.-G. (2022). What is the value of the cross-sectional approach to deep reinforcement learning? Quantitative Finance, 22(6), 1091–1111. https://doi.org/10.1080/14697688.2021.2001032

2. Abu-Alkheil, A., Khan, W. A., Parikh, B., & Mohanty, S. K. (2017). Dynamic co-integration and portfolio diversification of Islamic and conventional indices: Global evidence. Quarterly Review of Economics and Finance, 66, 212–224. https://doi.org/10.1016/j.qref.2017.02.005

3. Akhtaruzzaman, M., Banerjee, A. K., Boubaker, S., & Moussa, F. (2023). Does green improve portfolio optimisation? Energy Economics, 124, 106831. https://doi.org/10.1016/j.eneco.2023.106831

4. Alcazar, J., Leyton-Ortega, V., & Perdomo-Ortiz, A. (2020). Classical versus quantum models in machine learning: Insights from a finance application. Machine Learning: Science and Technology, 1(3). https://doi.org/10.1088/2632-2153/AB9009

5. Ali, F., Khurram, M. U., Sensoy, A., & Vo, X. V. (2024). Green cryptocurrencies and portfolio diversification in the era of greener paths. Renewable and Sustainable Energy Reviews, 191, 114137. https://doi.org/10.1016/j.rser.2023.114137

6. Anquetin, T., Coqueret, G., Tavin, B., & Welgryn, L. (2022). Scopes of carbon emissions and their impact on green portfolios. Economic Modelling, 115, 105951. https://doi.org/10.1016/j.econmod.2022.105951

7. Aouni, B., Doumpos, M., Pérez-Gladish, B., & Steuer, R. E. (2018). On the increasing importance of multiple criteria decision aid methods for portfolio selection. Journal of the Operational Research Society, 69(10), 1525–1542. https://doi.org/10.1080/01605682.2018.1475118

8. Aria, M., & Cuccurullo, C. (2017). bibliometrix: An R-tool for comprehensive science mapping analysis. Journal of Informetrics, 11(4), 959–975. https://doi.org/10.1016/j.joi.2017.08.007

9. Aria, M., & Cuccurullo, C. (2026). Science mapping analysis: A primer with Biblioshiny. McGraw-Hill Education. ISBN 978-88-386-2297-7.

10. Aria, M., Cuccurullo, C., D'Aniello, L., & Spano, M. (2026). Biblioshiny and the SAAS workflow: An integrated framework for transparent and reproducible science mapping. A demonstration through the replication of a study. Journal of Informetrics. https://doi.org/10.1016/j.joi.2026.101837

11. Bodnar, T., Parolya, N., & Schmid, W. (2018). Estimation of the global minimum variance portfolio in high dimensions. European Journal of Operational Research, 266(1), 371–390. https://doi.org/10.1016/j.ejor.2017.09.028

12. Buonaiuto, G., Gargiulo, F., De Pietro, G., Esposito, M., & Pota, M. (2023). Best practices for portfolio optimization by quantum computing, experimented on real quantum devices. Scientific Reports, 13, 19434. https://doi.org/10.1038/s41598-023-45392-w

13. Capponi, A., Ólafsson, S., & Zariphopoulou, T. (2022). Personalized robo-advising: Enhancing investment through client interaction. Management Science, 68(4), 2485–2512. https://doi.org/10.1287/mnsc.2021.4014

14. Chang, T.-J., Meade, N., Beasley, J. E., & Sharaiha, Y. M. (2000). Heuristics for cardinality constrained portfolio optimisation. Computers & Operations Research, 27(13), 1271–1302. https://doi.org/10.1016/S0305-0548(99)00074-X

15. Choudhary, H., Orra, A., Sahoo, K., & Thakur, M. (2025). Risk-adjusted deep reinforcement learning for portfolio optimization: A multi-reward approach. International Journal of Computational Intelligence Systems, 18(1), 126. https://doi.org/10.1007/s44196-025-00875-8

16. Choudhary, H., Orra, A., Thakur, M., Gao, X.-Z., & Sahu, P. K. (2026). A CVaR-constrained safe reinforcement learning framework with action repair for practical portfolio optimization. IEEE Transactions on Artificial Intelligence. https://doi.org/10.1109/TAI.2026.3686783

17. Cui, X.-Y., Gao, J.-J., Li, X., & Shi, Y. (2022). Survey on multi-period mean–variance portfolio selection model. Journal of the Operations Research Society of China, 10(3), 599–622. https://doi.org/10.1007/s40305-022-00397-6

18. Czichowsky, C., Peyre, R., Schachermayer, W., & Yang, J. (2018). Shadow prices, fractional Brownian motion, and portfolio optimisation under transaction costs. Finance and Stochastics, 22(1), 161–180. https://doi.org/10.1007/s00780-017-0351-5

19. Doering, J., Kizys, R., Juan, A. A., Fitó, À., & Polat, O. (2019). Metaheuristics for rich portfolio optimisation and risk management: Current state and future trends. Operations Research Perspectives, 6, 100121. https://doi.org/10.1016/j.orp.2019.100121

20. Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines. Journal of Business Research, 133, 285–296. https://doi.org/10.1016/j.jbusres.2021.04.070

21. Ferreira, F. G. D. C., Gandomi, A. H., & Cardoso, R. T. N. (2021). Artificial intelligence applied to stock market trading: A review. IEEE Access, 9, 30898–30917. https://doi.org/10.1109/ACCESS.2021.3058133

22. Gasser, S. M., Rammerstorfer, M., & Weinmayer, K. (2017). Markowitz revisited: Social portfolio engineering. European Journal of Operational Research, 258(3), 1181–1190. https://doi.org/10.1016/j.ejor.2016.10.043

23. Giuffrè, M., & Shung, D. L. (2023). Harnessing the power of synthetic data in healthcare: Innovation, application, and privacy. npj Digital Medicine, 6, 186. https://doi.org/10.1038/s41746-023-00927-3

24. Glänzel, W., & Schubert, A. (2001). Double effort = double impact? A critical view at international co- authorship in chemistry. Scientometrics, 50(2), 199–214. https://doi.org/10.1023/A:1010561321723

25. Goyal, K., & Kumar, S. (2021). Financial literacy: A systematic review and bibliometric analysis. International Journal of Consumer Studies, 45(1), 80–105.

26. Gunjan, A., & Bhattacharyya, S. (2024). Quantum-inspired meta-heuristic approaches for a constrained portfolio optimization problem. Evolutionary Intelligence, 17(4), 3061–3100. https://doi.org/10.1007/s12065-024-00929-4

27. Gunjan, A., & Bhattacharyya, S. (2025). Explainable artificial intelligence enhanced quantum-inspired spider monkey optimization for a constrained portfolio optimization problem. Quantum Machine Intelligence, 7(2), 108. https://doi.org/10.1007/s42484-025-00338-5

28. Hambly, B., Xu, R., & Yang, H. (2023). Recent advances in reinforcement learning in finance. Mathematical Finance, 33(3), 437–503. https://doi.org/10.1111/mafi.12382

29. Hegade, N. N., Chandarana, P., Paul, K., Chen, X., Albarrán-Arriagada, F., & Solano, E. (2022). Portfolio optimization with digitized counterdiabatic quantum algorithms. Physical Review Research, 4(4), 043204. https://doi.org/10.1103/PhysRevResearch.4.043204

30. Hidalgo-Marín, A. J., Nebro, A. J., & García-Nieto, J. (2026a). Efficient selection of investment portfolios in real-world markets: A multi-objective optimization approach. Algorithms, 19(1), 20. https://doi.org/10.3390/a19010020

31. Hidalgo-Marín, A. J., Nebro, A. J., & García-Nieto, J. (2026b). ParetoInvest: Integrating real-time financial data and multi-objective meta-heuristics for portfolio optimization. SoftwareX, 33, 102469. https://doi.org/10.1016/j.softx.2025.102469

32. Kalayci, C. B., Ertenlice, O., & Akbay, M. A. (2019). A comprehensive review of deterministic models and applications for mean-variance portfolio optimization. Expert Systems with Applications, 125, 345–368. https://doi.org/10.1016/j.eswa.2019.02.011

33. Kessler, M. M. (1963). Bibliographic coupling between scientific papers. American Documentation, 14(1), 10–25. https://doi.org/10.1002/asi.5090140103

34. Khan, A. T., Cao, X., Li, S., Hu, B., & Katsikis, V. N. (2021). Quantum beetle antennae search: A novel technique for the constrained portfolio optimization problem. Science China Information Sciences, 64(5), 152204. https://doi.org/10.1007/s11432-020-2894-9

35. Kocuk, B., & Cornuéjols, G. (2020). Incorporating Black-Litterman views in portfolio construction when stock returns are a mixture of normals. Omega, 91, 102876. https://doi.org/10.1016/j.omega.2018.11.017

36. Kolm, P. N., & Ritter, G. (2025). Reinforcement learning for asset and portfolio management. Journal of Portfolio Management, 52(2), 81–95. https://doi.org/10.3905/jpm.2025.1.782

37. Kolm, P. N., Ritter, G., & Simonian, J. (2021). Black-Litterman and beyond: The Bayesian paradigm in investment management. Journal of Portfolio Management, 47(5), 91–113. https://doi.org/10.3905/JPM.2021.1.222

38. Lalwani, V. (2024). Incorporating green assets in equity portfolios. Finance Research Letters, 59, 104815. https://doi.org/10.1016/j.frl.2023.104815

39. Leal, M., Ponce, D., & Puerto, J. (2020). Portfolio problems with two levels decision-makers: Optimal portfolio selection with pricing decisions on transaction costs. European Journal of Operational Research, 284(2), 712–727. https://doi.org/10.1016/j.ejor.2019.12.039

40. Li, X., & Wang, K.-H. (2026). Does investment in fintech assets enhance performance in China's financial sector? Evidence from multiple investment strategies. Electronic Commerce Research, 26(2), 1489–1528. https://doi.org/10.1007/s10660-025-09951-9

41. Li, X., Uysal, A. S., & Mulvey, J. M. (2022a). Multi-period portfolio optimization using model predictive control with mean-variance and risk parity frameworks. European Journal of Operational Research, 299(3), 1158–1176. https://doi.org/10.1016/j.ejor.2021.10.002

42. Li, X., Xu, F., & Jing, K. (2022b). Robust enhanced indexation with ESG: An empirical study in the Chinese stock market. Economic Modelling, 107, 105711. https://doi.org/10.1016/j.econmod.2021.105711

43. Ma, Y., Ahmad, F., Liu, M., & Wang, Z. (2020). Portfolio optimization in the era of digital financialization using cryptocurrencies. Technological Forecasting and Social Change, 161, 120265. https://doi.org/10.1016/j.techfore.2020.120265

44. Mahmoud, M. I., & Surwanti, A. (2026). A bibliometric study on the role of artificial intelligence in enhancing investment decision-making processes. Multidisciplinary Reviews, 9(2). https://doi.org/10.31893/multirev.2026092

45. Manogna, R. L., & Anand, A. (2024). A bibliometric analysis on the application of deep learning in finance: status, development and future directions. Kybernetes, 53(12), 5951–5971. https://doi.org/10.1108/K-04-2023-0637

46. Mantshimuli, L. A., & Muteba Mwamba, J. W. (2026). Few-shot portfolio optimization: Can large language models outperform quantitative portfolio optimization? A comparative study of LLMs and optimized portfolio allocators. Journal of Risk and Financial Management, 19(5), 320. https://doi.org/10.3390/jrfm19050320

47. Mantshimuli, L. A., & Mwamba, J. W. M. (2025). Enhancing portfolio optimization with multi-LLM sentiment aggregation: A Black–Litterman integration approach. Investment Management and Financial Innovations, 22(3), 213–226. https://doi.org/10.21511/imfi.22(3).2025.16

48. Meghwani, S. S., & Thakur, M. (2018). Multi-objective heuristic algorithms for practical portfolio optimization and rebalancing with transaction cost. Applied Soft Computing, 67, 865–894. https://doi.org/10.1016/j.asoc.2017.09.025

49. Metawa, N., Hassan, M. K., & Elhoseny, M. (2017). Genetic algorithm based model for optimizing bank lending decisions. Expert Systems with Applications, 80, 75–82. https://doi.org/10.1016/j.eswa.2017.03.021

50. Monasterolo, I., & de Angelis, L. (2020). Blind to carbon risk? An analysis of stock market reaction to the Paris Agreement. Ecological Economics, 170, 106571. https://doi.org/10.1016/j.ecolecon.2019.106571

51. Mugel, S., Kuchkovsky, C., Sánchez, E., Fernández-Lorenzo, S., Luis-Hita, J., Lizaso, E., & Orús, R. (2022). Dynamic portfolio optimization with real datasets using quantum processors and quantum-inspired tensor networks. Physical Review Research, 4(1), 013006. https://doi.org/10.1103/PhysRevResearch.4.013006

52. Naik, A. S., Yeniaras, E., Hellstern, G., Prasad, G., & Vishwakarma, S. K. L. P. (2025). From portfolio optimization to quantum blockchain and security: A systematic review of quantum computing in finance. Financial Innovation, 11(1). https://doi.org/10.1186/s40854-025-00751-6

53. Orús, R., Mugel, S., & Lizaso, E. (2019). Quantum computing for finance: Overview and prospects. Reviews in Physics, 4, 100028. https://doi.org/10.1016/j.revip.2019.100028

54. Pranckutė, R. (2021). Web of Science (WoS) and Scopus: The titans of bibliographic information in today's academic world. Publications, 9(1), 12. https://doi.org/10.3390/publications9010012

55. Rebentrost, P., & Lloyd, S. (2024). Quantum computational finance: Quantum algorithm for portfolio optimization. KI – Künstliche Intelligenz, 38(4), 327–338. https://doi.org/10.1007/s13218-024-00870-9

56. Rezaei, M., & Nezamabadi-pour, H. (2025). A taxonomy of literature reviews and experimental study of deep reinforcement learning in portfolio management. Artificial Intelligence Review, 58(3), 94. https://doi.org/10.1007/s10462-024-11066-w

57. Rout, A. K., Acharya, U. D., Aithal, P. K., & Sharma, A. (2026). Systematic review of reinforcement learning for automated equity portfolio management from single agent to multi agent systems. Discover Computing, 29(1). https://doi.org/10.1007/s10791-026-10336-1

58. Sattar, A., Sarwar, A., Gillani, S., Bukhari, M., Rho, S., & Faseeh, M. (2025). A novel RMS-driven deep reinforcement learning for optimized portfolio management in stock trading. IEEE Access, 13, 42813–42835. https://doi.org/10.1109/ACCESS.2025.3546099

59. Singh, V., Chen, S.-S., Singhania, M., Nanavati, B., Kar, A. K., & Gupta, A. (2022). How are reinforcement learning and deep learning algorithms used for big data based decision making in financial industries: A review and research agenda. International Journal of Information Management Data Insights, 2(2), 100094. https://doi.org/10.1016/j.jjimei.2022.100094

60. Small, H. (1973). Co-citation in the scientific literature: A new measure of the relationship between two documents. Journal of the American Society for Information Science, 24(4), 265–269. https://doi.org/10.1002/asi.4630240406

61. Soler-Dominguez, A., Juan, A. A., & Kizys, R. (2018). A survey on financial applications of metaheuristics. ACM Computing Surveys, 50(1). https://doi.org/10.1145/3054133

62. Sonkavde, G., Dharrao, D. S., Bongale, A. M., Deokate, S. T., Doreswamy, D., & Bhat, S. K. (2023). Forecasting stock market prices using machine learning and deep learning models: A systematic review, performance analysis and discussion of implications. International Journal of Financial Studies, 11(3), 94. https://doi.org/10.3390/ijfs11030094

63. Sutiene, K., Schwendner, P., Sipos, C., Lorenzo, L., Mirchev, M., Lameski, P., Kabasinskas, A., Tidjani, C., Ozturkkal, B., & Cerneviciene, J. (2024). Enhancing portfolio management using artificial intelligence: Literature review. Frontiers in Artificial Intelligence, 7, 1371502. https://doi.org/10.3389/frai.2024.1371502

64. Thomé, A. M. T., Scavarda, L. F., & Scavarda, A. J. (2016). Conducting systematic literature review in operations management. Production Planning & Control, 27(5), 408–420. https://doi.org/10.1080/09537287.2015.1129464

65. Trimborn, S., Li, M., & Härdle, W. K. (2020). Investing with cryptocurrencies: A liquidity constrained investment approach. Journal of Financial Econometrics, 18(2), 280–306. https://doi.org/10.1093/jjfinec/nbz016

66. van Eck, N. J., & Waltman, L. (2010). Software survey: VOSviewer, a computer program for bibliometric mapping. Scientometrics, 84(2), 523–538. https://doi.org/10.1007/s11192-009-0146-3

67. Weber, P., Carl, K. V., & Hinz, O. (2024). Applications of explainable artificial intelligence in finance—a systematic review of finance, information systems, and computer science literature. Management Review Quarterly, 74(2), 867–907. https://doi.org/10.1007/s11301-023-00320-0

68. Yadav, S., Kumar, A., Mehlawat, M. K., Gupta, P., & Charles, V. (2023). A multi-objective sustainable financial portfolio selection approach under an intuitionistic fuzzy framework. Information Sciences, 646, 119379. https://doi.org/10.1016/j.ins.2023.119379

69. Yadav, S., Ravi, V., & Kalyanakrishnan, S. (2025). Hybrids of reinforcement learning and evolutionary computation in finance: A survey. ACM Computing Surveys, 57(10), 247. https://doi.org/10.1145/3728634

70. Zeng, X., Chen, C., Wang, Y., & Liang, Y. (2025). ELfolio: Strategy evolution via large language models for portfolio optimization. Intelligent Computing, 4, 0176. https://doi.org/10.34133/icomputing.0176

71. Zupic, I., & Čater, T. (2015). Bibliometric methods in management and organization. Organizational Research Methods, 18(3), 429–472. https://doi.org/10.1177/1094428114562629