INTERNATIONAL JOURNAL OF ADVANCE RESEARCH, IDEAS AND INNOVATIONS IN TECHNOLOGY

ISSN: 2454-132X
Impact Factor: 6.078
(Volume 12, Issue 5 - V12I5-1228)
Available online at: https://www.ijariit.com

Applications of Operations Research in AI for Optimization: A Bibliometric Review

Dr Richa Saxena
richa.saxena@nmims.edu
Anil Surendra Modi School of Commerce, NMIMS, Mumbai, Maharashtra

Dhruv Jain
dhruv.jain411@nmims.in
Anil Surendra Modi School of Commerce, NMIMS, Mumbai, Maharashtra

Dhrushank Thakker
dhrushank.thakker016@nmims.in
Anil Surendra Modi School of Commerce, NMIMS, Mumbai, Maharashtra

Deepti Soneji
deepti.soneji547@nmims.in
Anil Surendra Modi School of Commerce, NMIMS, Mumbai, Maharashtra

Dev Udeshi
dev.udeshi042@nmims.in
Anil Surendra Modi School of Commerce, NMIMS, Mumbai, Maharashtra

Deshna Malu
deshna.malu295@nmims.in
Anil Surendra Modi School of Commerce, NMIMS, Mumbai, Maharashtra

ABSTRACT

Artificial intelligence (AI) is increasingly used to solve optimisation problems across manufacturing, supply chains, energy, healthcare and food, yet existing reviews examine these sectors in isolation. This study maps the field of AI for optimisation through a bibliometric analysis of 1,550 Scopus-indexed journal articles and reviews published between 2016 and 2026, screened using the PRISMA approach and analysed with the bibliometrix R package, Biblioshiny and VOSviewer, alongside a structured review of the literature. The results show a young, rapidly growing and highly collaborative field: 88.3% of the documents were published in the last three years, 5,822 authors contributed, and output is scattered across 894 sources. Thematic mapping identifies analytics-based decision-making fed by internet-of-things data as the motor theme, optimisation and machine learning as the basic theme, clinical applications as a niche theme, and large language models as an emerging theme. The study also identifies five gaps: prediction dominates over prescription, reported gains rely largely on benchmarks and simulations, sectors are studied separately, adoption barriers are rarely examined, and fairness and computing costs are seldom weighed against benefits. A future research agenda of six priorities is proposed to address these gaps.

Keywords: *Artificial Intelligence, Operations Research, Optimisation, Bibliometric Analysis, Prescriptive Analytics, Machine Learning*

1. INTRODUCTION

Operations Research (OR) has always been concerned with finding the best possible decision for a problem with limited resources, and throughout most of its existence this has meant casting the problem as a mathematical programme and solving it by means of an exact algorithm or heuristic. Artificial intelligence (AI) has altered the manner in which such problems may be addressed. Machine learning now provides the forecasts, parameters and learned rules that are supplied to an optimisation model, while optimisation is, in turn, used to train and tune the AI models themselves. Bengio et al. (2021) illustrate this two-way dependency for combinatorial optimisation, where learning may either be used to substitute some costly computations that live within a solver, or may be used to directly learn which decisions are good. Bertsimas and Kallus (2020) go one step further and contend that the value of the data used does not lie in prediction per se but in prescription, or the conversion of a prediction into a decision that can be optimised. These two concepts – learning that serves optimisation, and optimisation that serves learning – are central to what is now most often simply referred to as AI for optimisation.

Its appeal can be attested in applications across many sectors. In manufacturing it is used for monitoring machines, predicting breakdowns and scheduling maintenance. In supply chains and logistics it is used for forecasting demand, planning inventories

and routing vehicles, and more recently, with sustainability criteria in mind (Chen et al., 2024; El Mariami et al., 2025). In energy systems it is used to balance supply, demand, cost and emissions in grids and buildings, while in healthcare it is used to aid diagnosis, treatment planning and hospital-level decisions. An illustrative example of the prescriptive logic underlying such applications is provided by Pessach et al. (2020), where an interpretable machine learning model is used first to predict whether the placement of a new employee will be successful, and a mathematical programming model is used next to plan the hiring process across the entire organisation. Digital twins, or virtual replicas of physical systems that are continuously updated with real-time data, have been used to make the connection between the two sides even tighter, and Sajadieh and Noh (2025) identify a role for AI at each stage of a twin's life cycle, including decision-making and optimisation.

The toolset has also evolved rapidly. Many applications are still based on neural networks and nature-inspired search techniques like particle swarm optimisation and genetic algorithms, appearing in 41 and 36 of the abstracts of the set of records used for this study, respectively. More recently, however, applications draw on deep learning, reinforcement learning, federated learning and edge computing, and since 2025, a separate stream of applications can be found on large language models and generative AI. Explainable AI is another area which has attracted attention recently, as a decision based on a black-box model may be difficult for a manager or clinician to accept without a sense of the reasons behind the model's recommendation. The number of methods at a researcher's disposal has therefore also multiplied, but it is not clear that evidence on what methods work, and where, and at what cost, has grown at a similar pace.

A second motivation to look at the field as a whole is that sophistication and usefulness are not necessarily one and the same. Many applications report impressive gains in terms of accuracy or cost reduction, but such gains are often obtained on benchmark datasets or in simulation, and the settings, data and performance indicators are different from study to study. The presence of barriers such as data availability, interoperability between systems and organisational adoption of techniques are often cited, but seldom examined in detail (El Mariami et al., 2025; Sajadieh & Noh, 2025). The question of whether more sophisticated AI techniques actually make for better decisions, once cost, constraints and operating conditions are considered, remains open therefore.

This question is not addressed in the existing review literature, because it is structured by sector. El Mariami et al. (2025), for example, review OR and AI for supply chain planning, Chen et al. (2024) review AI in logistics optimisation with sustainability criteria, and Sajadieh and Noh (2025) review the integration of AI and digital twins. Bibliometric work is similarly sectorally focused: Jaramillo et al. (2025), for example, assess AI, the internet of things, blockchain and big data, but only in the context of power systems in renewable-energy. Each of these reviews is helpful in its own right and within the boundaries of its focus, but none maps the landscape of AI-based optimisation across manufacturing, supply chains, energy, healthcare and food together. A title-based search of the set of records used for this study for bibliometric reviews also yields only single-sector or single-technology work. To the best of our knowledge, there is therefore no study that would map the whole field and show where OR-style optimisation remains central, where it has been absorbed into machine learning, and which themes are evolving in isolation from each other.

Bibliometric analysis is well suited to such a task. It quantifies the growth of a field, the contributions that authors and sources have made, and the relationships between keywords, and is particularly helpful where a field has become too big to be read paper by paper (Donthu et al., 2021). The data used for this study were retrieved on 10 September 2026 from Scopus on the basis of the search string "AI AND Optimization AND analytics" and were screened, following the PRISMA approach, down to the English-language journal articles and reviews published from 2016 onwards. After removing five records dating from 2027, the resulting set consists of 1,550 documents (1,118 articles and 432 reviews), published in 894 sources by 5,822 authors. The field is very young. Output was negligible until 2018, rose to 50 documents in 2021, and then increased to 147 in 2024, 546 in 2025 and 675 in 2026, which means that 88.3% of all records were published in the last three years. The analysis is carried out using the bibliometrix R package and its Biblioshiny interface (Aria & Cuccurullo, 2017), supported by VOSviewer for the mapping of keywords, and themes are placed on a strategic diagram according to their center of gravity and density (Cobo et al., 2011).

Accordingly, this study has three objectives. The first is to describe how research into AI for optimisation has developed, including the field's growth, authorship, collaboration and sources in which it is published. The second is to identify its thematic and conceptual structure, by distinguishing its motor, basic and niche themes, as well as its emerging or declining themes, and to compare this structure with the concepts identified in a structured reading of the literature. The third is to identify the gaps that future research should address and to turn them into a research agenda. In doing so, the study provides a picture that is useful across sectors, as single-sector reviews of the field cannot offer, not only to researchers in operations research, computer science and management, but also to practitioners whose jobs include deciding where AI-based optimisation is ready to be used. The paper proceeds with the research methodology, followed by the literature review, the bibliometric analysis, the future research agenda and, finally, the conclusion.

2. LITERATURE REVIEW

Artificial Intelligence for Optimisation

2.1 Background and Rationale

Artificial intelligence has become closely tied to optimisation research over the past few years, showing up in work on manufacturing and maintenance, supply chains and logistics, energy systems, healthcare, and food and agriculture. Part of the reason is that the two fields help each other: machine learning supplies the forecasts and learned rules that feed into an optimisation problem, while optimisation methods are used to tune and train the AI models themselves. The result is a literature that has grown quickly and spread across engineering, computer science, medicine and operations research venues at once, but that is rarely read as a single field.

A bibliometric approach fits this kind of situation reasonably well. Rather than reading a handful of representative papers, it works from publication patterns, keyword structures and thematic clustering across hundreds of records at once. That is the approach taken here, applied to a Scopus dataset of 1,557 bibliographic records, of which 1,123 are articles and 434 are reviews. The dataset is very recent, with 1,223 records (78.5%) published in 2025 or 2026, so this review is a snapshot of a field that is still moving fast. Keyword maps built in VOSviewer (349 and 485 keywords) each resolve into six clusters and point to the same three large areas: Industry 4.0, the internet of things and energy; healthcare and drug discovery; and food, agriculture and supply chains.

2.2 Existing Reviews and Positioning of This Study

There is a spectrum in what counts as a literature review. A systematic review follows a defined protocol for finding, screening and synthesising studies around a specific question. A meta-analysis goes further and statistically pools results across studies. A bibliometric review is different from both: it works from the publication metadata itself, such as counts, keyword co-occurrence and thematic clusters, to trace how a field has grown and where its structure currently sits. Because this review asks how AI-based optimisation has developed across several areas, rather than answering one narrow question, the bibliometric approach is the more natural fit.

A number of reviews already exist in this space, but they tend to work one area at a time. El Mariami, Mharzi, El Oualidi Moulay and Abadi (2025) reviewed 55 articles on AI in supply chain optimisation and planning, and Chen, Men, Fuster, Osorio and Juan (2024) reviewed AI for logistics optimisation with sustainability criteria. Hector and Panjanathan (2024) surveyed planning models and machine learning for predictive maintenance, and Sajadieh and Noh (2025) reviewed how AI and digital twins are being integrated. Bibliometric work exists too, but within single areas: Al-Thani et al. (2025) combine a bibliometric and systematic review of smart food supply chains, and Jaramillo, Carrión, Muñoz and Tipán (2025) assess AI, IoT, blockchain and big data in renewable-energy power systems. Table 1 lays these six reviews out alongside what each covers and where it stops short of what this study is trying to do.

Table -1: Existing reviews and their scope

Study	Year	Sub-Problem Examined	Contribution / Limitation
El Mariami, Mharzi, El Oualidi Moulay & Abadi	2025	AI for supply chain optimisation and planning, 55 articles from 2020 to 2024	Systematic review; supply chain only
Chen, Men, Fuster, Osorio & Juan	2024	AI in logistics optimisation with sustainability criteria	Covers machine learning and metaheuristics; logistics only
Hector & Panjanathan	2024	Planning models and machine learning for predictive maintenance	Survey; maintenance only
Sajadieh & Noh	2025	Integration of AI and digital twins	Proposes an integration framework; digital twins only
Al-Thani, Hadid, Padmanabhan, Kerbache & Elomri	2025	Smart food supply chain management	Bibliometric and systematic review; food supply chains only
Jaramillo, Carrión, Muñoz & Tipán	2025	AI, IoT, blockchain and big data in renewable-energy power systems	Bibliometric; power systems only

Taken together, these six reviews cover supply chains, logistics, maintenance, digital twins, food systems and power systems, but each in isolation. A title search of the dataset for bibliometric studies returned only work focused on one sector or one technology, and none that maps AI-based optimisation across manufacturing, supply chains, energy, healthcare and food together. That absence is treated here as a real, evidence-based gap rather than an assumption. The claim is deliberately limited to the dataset examined and not to the literature at large, since no search can fully rule out a cross-sector review sitting outside the sources consulted.

2.3 Key Concepts

A few concepts come up often enough in this literature that they are worth explaining before going further. The first is the move from prediction to prescription. A forecast on its own is not a decision, so many studies add an optimisation step after the prediction. Pessach et al. (2020) first predict whether a job placement will succeed, using an interpretable machine-learning model, and then use a mathematical model to plan hiring across the whole organisation. Lepenioti, Bousdekis, Apostolou and Mentzas (2021) note that methods for recommending actions have mostly been built for one specific field, and propose a more general approach based on interactive multi-objective reinforcement learning. Miikkulainen et al. (2021) use predictive models to test many COVID-19 control strategies and an evolutionary search to find the best ones.

The second concept runs the other way, with optimisation algorithms used to improve AI models. Ehteram et al. (2021) hybridise two AI models with the sunflower optimisation algorithm, which finds their best tuning parameters, to predict lake water levels. Kumar, Addula, Lind, Brown and Odion (2026) pair a deep-learning model with Sand Cat Optimisation to identify faults in factory robots. The logic in both cases is that the learning model supplies the prediction while the search method makes that prediction as accurate as it can be.

A digital twin is a virtual copy of a physical system that is kept up to date with live data. Emmert-Streib (2023) argues that the role of AI and machine learning in digital twins has been understudied, and Sajadieh and Noh (2025) place AI at every stage of a twin's life, including decision-making and optimisation.

Explainable AI is the last concept worth flagging. It covers methods that make a model's predictions understandable to a human rather than leaving them as the output of a black box. Parisineni and Pal (2024) propose a method that combines local explanations with SHAP because complex models are hard to trust, and Teixeira, Carvalhais, Pinto and Vale (2025) present an

explainable AI framework for building energy management. This matters in optimisation because an operating decision based on a model's output is hard to act on without some sense of why the model produced it.

Table -2: Keyword groups on the VOSviewer maps

Keyword group	Examples on the maps	Cluster colour (Map 1 / Map 2)	What it tells us
Core AI terms	artificial intelligence, machine learning, artificial neural network, convolutional neural network	Green / Green	The largest circles on both maps. They sit between the sectors and act as the shared toolkit that links manufacturing, health and food research.
Industry 4.0 and smart manufacturing	Industry 4.0, internet of things (IoT), predictive maintenance, cyber-physical systems, embedded systems, fault detection	Red / Red	The manufacturing strand: AI used to monitor, maintain and control machines and production lines.
Digital twins, security and connectivity	digital twin, interoperability, edge computing, cybersecurity, blockchain, federated learning	Red / Red	The conditions connected systems depend on: data must flow between devices and organisations and stay secure. This matches the interoperability barrier raised in the review.
Energy and sustainability	energy efficiency, sustainability, smart power grids, electric vehicles, renewable energies, circular economy, energy management	Red / Red	The energy strand: AI used to balance supply, demand, cost and emissions in grids, buildings and transport.
Decisions and optimisation	decision making, optimization, predictive analytics, model predictive control, reinforcement learning, information management	Red / Red	Where forecasts are turned into actions. These terms sit close to the centre of the maps, which is the prediction-to-prescription link discussed earlier.
Supply chains, big data and forecasting	supply chains, supply chain management, supply chain optimization, big data, data analytics, forecasting, automation	Yellow / Yellow, some red	The supply chain strand: data-driven forecasting and planning for logistics, inventory and operations.
Food and agriculture	food safety, food processing, food traceability, quality control, precision agriculture, precision livestock farming, near infrared spectroscopy, remote sensing, shelf life	Blue / Yellow	The food and farming strand: AI for safe, traceable, good-quality food and efficient use of land and water. The cluster changes colour between maps, but its terms stay together.
Healthcare and personalised medicine	personalized medicine, precision medicine, clinical decision support, electronic health record, diagnostic imaging, image analysis, patient care, clinical outcome	Green / Green	The healthcare strand: AI for diagnosis, treatment planning and hospital decisions.
Drug discovery and pharmaceutics	drug discovery, drug development, drug design, proteomics, multiomics, pharmacokinetics, drug delivery system, nanocarrier	Green / Green, with a small cyan group	The molecule-level end of healthcare. On the second map, the pharmaceutics terms split off into a small cyan cluster.
Generative AI and explainability	large language model, generative AI, explainable AI, natural language processing, knowledge management, strategic planning	Few small nodes / Blue	The newest cluster, most visible on the second map. It shows large language models and explainable AI arriving as a separate topic.
Ethics, privacy and trust	ethics, privacy, data privacy, reproducibility, benchmarking, reliability	Mostly green / Mostly green	Where fairness, privacy and validation appear, mainly in health research. This supports the gap on fairness and real-world validation.
Review and method terms	systematic review, bibliometric analysis, scoping review, review, publishing, simulation	Green and purple / Purple	These describe the type of paper or method, not a research topic. They show how many records are reviews (434 of 1,557).

Keyword group	Examples on the maps	Cluster colour (Map 1 / Map 2)	What it tells us
Database index tags	human, humans, nonhuman, animal, female, procedures, article	Green / Green	Labels added by biomedical databases, not research themes. They look large because many health records carry them, so they should not be read as topics.

2.4 Evidence Gaps in Prior Studies

Beyond the issues mentioned above, there are a few gaps that could be noticed in the way the evidence was produced for this field:

- While several articles address prediction and forecasting, relatively fewer contain a prescriptive element, i.e., follow up on predicted outcomes with a suggested course of action
- The research areas of manufacturing, supply chains, energy, healthcare, and food production are addressed separately, with none of the articles synthesizing the interconnections between them
- Comparability between the articles is limited, as they often use different settings, data, and performance metrics, while the abstracts often mention comparisons to operations-management practices, without specifying the details of such comparisons
- Data availability, interoperability, and organizational adoption barriers are identified as areas of concern (El Mariami et al., 2025; Sajadieh & Noh, 2025), but often without going into specifics
- Fairness concerns and costs of adopting solutions, particularly in terms of resources spent on AI training and deployment are rarely balanced out against the benefits, while the topics are often only discussed in one area, such as medicine (Hussain, Bresnahan & Zhuang, 2025).

With that said, it is important to note that several of the articles mentioned above are recent ones, and while time is a factor in research, it is not a determining factor. In other words, simply being recent does not make an article a gap in the evidence, unless it specifically limits the perspective. For instance, Hang, Yu, Morabito and Tan (2024), and Chen et al. (2025) are both recent because they address LLMs, which are recent additions to the field of AI. As such, a gap would be represented by an article that specifically limits its scope to one industry, one data set or methodology, while a recent article does not necessarily do that.

2.5 Literature Review Table

Table 3 pulls together the studies discussed above with their research focus, key findings, and the specific gap or limitation each one leaves open.

Table -3: Literature review table

Author(s) & Year	Research Focus	Key Findings	Research Gap
Pessach et al. (2020)	Hiring: ML prediction plus mathematical programming	Interpretable model predicts placement success; optimisation plans hiring across the organisation	Single application (HR)
Lepenioti et al. (2021)	Prescriptive analytics with interactive multi-objective RL	Proposes a general method adaptable to human preferences	Method proposal
Miikkulainen et al. (2021)	Evolutionary surrogate-assisted prescription for COVID-19 interventions	Predictive models score candidate strategies; evolutionary search finds the best	Early-stage experiments
Ehteram et al. (2021)	AI models tuned with sunflower optimisation for lake water level	ANFIS-SO performs best; RMSE 1.89 m (training), 1.92 m (testing)	Single lake
Kumar et al. (2026)	Attention GRU with Sand Cat Optimisation for robot maintenance	Improves fault identification	Evaluated on a benchmark dataset
Lim & Jeong (2023)	Deep RL for storage devices in factory simulation	RL used to optimise storage in Industry 4.0 and digital-twin setting	Simulation-based
Sajadieh & Noh (2025)	AI and digital twin integration	Framework placing AI across the twin life cycle	Interoperability, scalability and security remain open
Iseri et al. (2025)	Forecasting for reverse logistics of solar panels	Statistical and deep-learning forecasts support closed-loop supply chain optimisation	Single product
Arvind et al. (2025)	ML and IoT for warehouse inventory	Forecast error 32% lower than ARIMA; costs 35% lower	Single framework and dataset
Skoropad et al. (2025)	RL for traffic signal control in Belgrade	Deep Q-Learning and PPO control signals using real-time data	Simulated network
Saini, Thakur & Yadav (2025)	AI in drug discovery and pharmaceutical operations	Speeds discovery; supports trial design and recruitment	Review-level evidence

Author(s) & Year	Research Focus	Key Findings	Research Gap
Hussain, Bresnahan & Zhuang (2025)	Bias in AI health algorithms	Racial and ethnic bias found; oversight and governance examined	Scoping review of 23 sources

3. RESEARCH METHODOLOGY

A number of techniques have been developed to understand the growth, development and emerging trends within a particular area of academic research. Among these techniques, Systematic Literature Review (SLR) and bibliometric analysis are more often used to organise and synthesise large body of existing literature. Systematic Literature Review is helpful in understanding the prior knowledge and synthesis in the previous outcomes to understand the gaps for future research (Paul & Rialp, 2020). It provides a structured approach to identification, screening, eligibility and inclusion. The systematic nature of such a review also improves transparency in the research process by clearly specifying how studies were identified and selected. This research study utilizes one such technique with an objective to identify and analyse the research trends in the application of AI FOR OPTIMIZATION in academic and business research. The PRISMA framework provides a diagrammatic view of the article selection process followed in this study. It shows how the research data were identified, screened and finally selected on the basis of the predefined inclusion and exclusion criteria. This helps to clearly explain the basis for selecting the final set of articles and make the overall selection process more transparent and systematic (Donthu, Kumar, Mukherjee, Pandey, & Lim, 2021).

3.1 Article selection, inclusion, and exclusion criteria

A search was conducted on 10th September 2026 on the Elsevier's database Scopus. Scopus provides extensive coverage and allows researchers to retrieve important bibliographic information such as publication details, authors, affiliations, abstracts, keywords and citation-related information. The search was conducted using the search keywords "AI AND Optimization AND analytics". After the search further refining was done using the filter option in the Scopus such as Time span, document type, subject criteria and many more, this is used to note the number of excluded articles using filtering criteria, which was further used while making the PRISMA diagram.

After collecting the database, the article selection was further structured using the PRISMA flow diagram available through Biblioshiny, a web-based interface of the Bibliometrix package (Aria & Cuccurullo, bibliometrix: An R-tool for comprehensive science mapping analysis, 2017). The process began with the identification stage where the source of the data was specified as the Scopus database, along with the search query "AI AND Optimization AND analytics" and the total number of records retrieved from the search. The retrieved bibliographic data were than exported from Scopus and imported into Biblioshiny for further screening and analysis. Biblioshiny provides a systematic filtering criteria which can be applied sequentially to refine the initial database (Aria, Cuccurullo, D'Aniello, & Spano, 2026).

In the screening and eligibility stage, predefined criteria were applied to make the dataset more relevant to the objective of the study. The first criterion was the publication time span, where publications before 2015 were excluded, though there were publications of 2012, 2014 and 2015, the study focused on literature published from 2015 onwards only. The next criterion was the document type, where the study focused on articles and reviews while excluding conference papers, book chapters, conference reviews, books, retracted, notes, erratum, editorials, short surveys and data papers. Further, the next criterion was subject areas where publication from Computer science, Engineering, Mathematics, Business, Management and Accounting, Energy, Decision Sciences and other related fields were included.

In the next criterion publications written in languages other than English were excluded and only English language articles and reviews were retained. This criterion was adopted to minimise language related barriers during the subsequent analysis and interpretation of the data, as publications in multiple languages like Chinese, Russian, German etc. may create difficulties in consistently understanding, comparing and synthesising the findings (Pieper & Puljak, 2021). In the last, filtering was done on the basis of source type in the database. It showed three categories namely journal, trade journal and book series from which only journal source type was included for further analysis.

After all the criterion details filled the PRISMA diagram was generated which showed all the stages of systematic study and also provided us the set of publications that remained after applying the predefined screening and eligibility criteria. The PRISMA diagram therefore provides a clear diagrammatic representation of the entire selection process, showing how the initial records obtained from Scopus, were progressively screened and refined until the final dataset was obtained. It also shows what data were selected and why particular records were excluded at different stages (Page, et al., 2021).

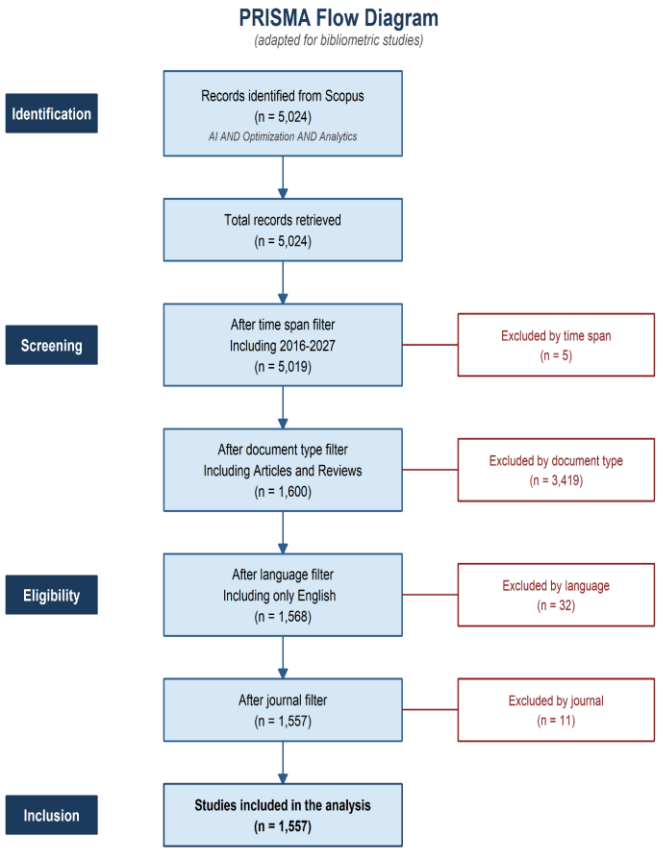

Figure -1: PRISMA diagram

3.2 Bibliometric analysis

The major advantage of bibliometric analysis is to systematically and methodically quantify, evaluate and study the data pertaining to the review articles of this paper. In practical application, bibliometric analysis constitutes a quantitative inquiry into the realm of academic publishing, encompassing research articles, periodicals, citation patterns, and associated bibliographic attributes. The data has been obtained through Elsevier's database Scopus using the keyword search and introducing the above-mentioned filters. The selected review articles were downloaded and exported in an .ris file from the database and imported into Mendeley Desktop which is reference manager. Further the papers were exported from Scopus in a BibTex file (.bib), the file formal to be used in further analysis using biblioshiny. The study utilised biblioshiny package of RStudio. Biblioshiny converted the data in an excel format (.xlxs), which was generated for the data cleaning and can also be used for further analysis using VosViewer and Tableau analysis. The first research objective investigates the performance of the existing literature, including publication details, sources, authors and citations. The second objective employs thematic mapping and co-occurrence analysis to identify the relationships between the keywords. Finally, future research areas are identified for further analysis and studies (Dr Richa Saxena, 2024).

4. DATA ANALYSIS

4.1 Overview

The 1,550 documents were published between 2016 and 2026 in 894 different sources, with the five records dated 2027 removed as they lie beyond the data-collection period. The annual growth rate of 78.99% indicates rapid expansion of the literature, although this figure should be interpreted with caution because it is calculated from a very small starting base (two documents in 2016). In terms of authorship, the dataset involves 5,822 authors, of whom only 159 (2.7%) have written single-authored documents, which suggests that research in this area is predominantly collaborative. This is supported by an average of 4.21 co-authors per document and an international co-authorship rate of 32.65%, meaning that roughly one in three documents involves authors from more than one country. Finally, the 1,550 documents contain 5,515 distinct author keywords, which points to a wide and varied vocabulary and is consistent with the thematic diversity observed in the conceptual structure and thematic maps.

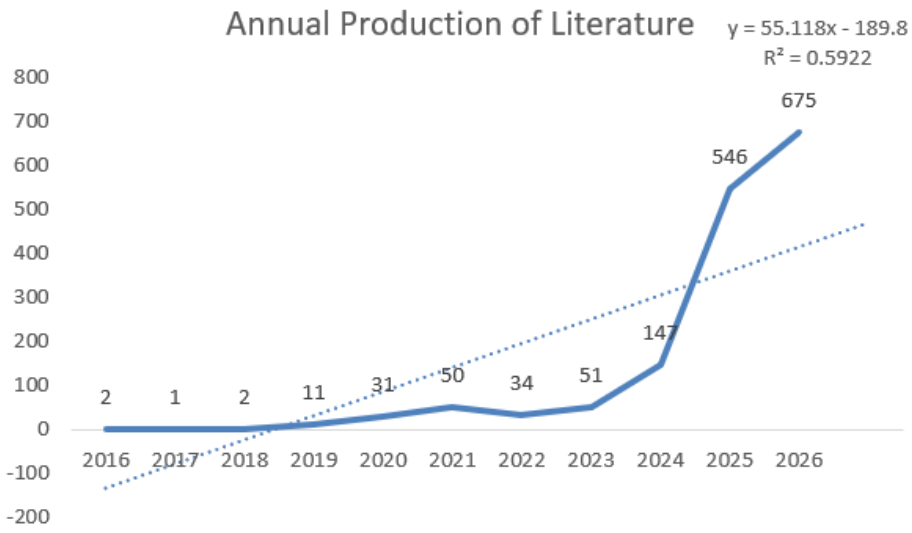

Diagram -1.1: Annual production of literature

Diagram 1.1 presents the annual production of literature. Output was negligible between 2016 and 2018 (five documents in total), began to grow in 2019 (11 documents) and reached 50 documents in 2021, after which it fluctuated at a similar level (34 in 2022 and 51 in 2023). A marked acceleration is evident from 2024, when 147 documents were published, followed by 546 in 2025 (a 3.7-fold increase) and 675 in 2026. It is noteworthy that the last three years together account for 88.3% of the dataset (1,368 documents), and that 2026 is already 23.6% higher than 2025 despite being a partial year. The linear trend line ($y = 55.118x - 189.8$) indicates an average increase of around 55 documents per year; however, its R^2 value of 0.5922 shows that a linear model explains the pattern only moderately well, as growth is concentrated in the most recent years rather than being steady. This surge may reflect the growing attention that AI-enabled optimisation and analytics have received in recent years, although the results for 2026 should be interpreted with caution because the year is incomplete (data retrieved in September 2026).

Table -1.2: Most relevant sources

Sources	Articles
IEEE ACCESS	35
APPLIED SCIENCES (SWITZERLAND)	30
SCIENTIFIC REPORTS	30
SUSTAINABILITY (SWITZERLAND)	28
DISCOVER ARTIFICIAL INTELLIGENCE	16
SENSORS	14
ENERGIES	13
RESULTS IN ENGINEERING	13
SHODHKOSH: JOURNAL OF VISUAL AND PERFORMING ARTS	13
IEEE INTERNET OF THINGS JOURNAL	12
ELECTRONICS (SWITZERLAND)	11
INTERNATIONAL JOURNAL OF ADVANCED COMPUTER SCIENCE AND APPLICATIONS	11
BUILDINGS	10
INTERNATIONAL JOURNAL OF ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING	10
JOURNAL OF INTELLIGENT DECISION MAKING AND INFORMATION SCIENCE	10
COMPUTERS, MATERIALS AND CONTINUA	9
INTERNATIONAL JOURNAL OF INFORMATION AND COMMUNICATION TECHNOLOGY	9
SYSTEMS	9
ALGORITHMS	8
IEEE TRANSACTIONS ON CONSUMER ELECTRONICS	8
INTERNATIONAL JOURNAL OF HYDROGEN ENERGY	7
PROCESSES	7

Table 1.2 presents the most relevant sources by number of articles published. The 1,550 documents appeared across 894 sources, which indicates a highly dispersed literature: 670 sources (74.7%) contributed only one document, while the ten most productive sources account for just 204 documents (13.1%). This pattern is broadly consistent with Bradford's law of scattering (Bradford, 1934). As can be seen from Table 1.4, IEEE Access leads with 35 articles, followed by Applied Sciences and Scientific Reports (30 each) and Sustainability (28). A closer inspection reveals that the leading sources span engineering and computing, sustainability and energy, and AI-specific journals, which suggests that research on AI for optimization and analytics is published mainly in applied, interdisciplinary venues. It is worth noting that productivity does not always indicate influence: Sensors averages 62.9 citations per article, well above the dataset mean of 18.5, whereas Scientific Reports averages 9.2. These findings should, however, be interpreted with caution, as Shodhkosh: Journal of Visual and Performing Arts (13 articles, two citations in total) lies outside the engineering core of this study.

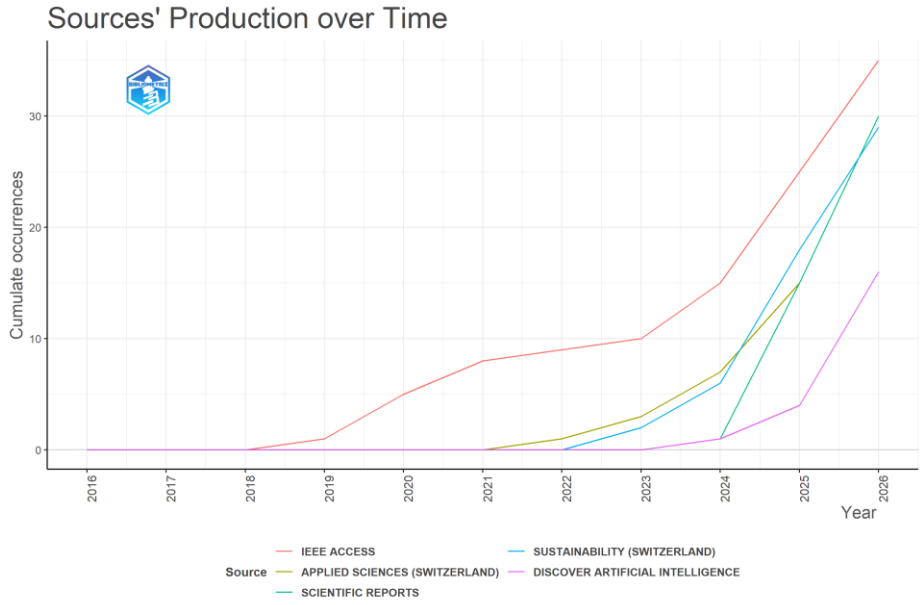

Diagram -1.3: Sources' production over time

Diagram 1.3 illustrates the cumulative production of the five most productive sources between 2016 and 2026. As can be seen from the graph, none of the sources published relevant documents before 2019, when IEEE Access became the first to appear. It remained the only source with output before 2022, reaching 10 cumulative articles by 2023. A notable feature of the graph is the sharp increase in production from 2024 onwards. Scientific Reports rose from a single article in 2024 to 30 by 2026, Discover Artificial Intelligence from 1 to 16, and Sustainability and Applied Sciences from 6 and 7 to 29 and 30 respectively, so that these three journals had almost caught up with IEEE Access (35 articles) by 2026. This pattern suggests a shift from engineering-oriented outlets towards multidisciplinary and sustainability-focused journals, and it is consistent with the recent growth of the literature, as roughly 79% of the articles in these five sources were published in 2025 and 2026. These results should, however, be interpreted with caution, as 2026 is a partial year and the lines show cumulative counts, which can only rise.

Diagram -1.4: Thematic map

Diagram 1.4 presents the thematic map of the 1,550 documents, constructed from the co-word network of keywords and plotted according to centrality and density (Callon et al., 1991; Cobo et al., 2011). Centrality, shown on the horizontal axis, measures the strength of a theme's links with other themes and thus its importance to the field, whereas density, shown on the vertical axis, measures the internal cohesion of a theme and thus its degree of development (Cobo et al., 2011). The map is divided into four quadrants, and the size of each bubble is proportional to the keyword occurrences within the cluster. The motor themes quadrant (upper right) contains a single cluster comprising predictive analytics, decision making and internet of things, which is

characterised by both high centrality and high density. This suggests that analytics-driven decision support, fed by IoT data, is the best developed and most influential strand of the literature, and it corresponds directly to the "analytics" dimension of the present study. The basic themes quadrant (lower right) contains a cluster of learning systems, machine-learning and optimization. Because this cluster is central but lies close to the median lines, it appears to represent foundational, transversal topics that underpin many studies. It is noteworthy that optimization, the second key concept of this study, falls here, which indicates that it acts as a link between the motor theme and the rest of the field. In contrast, the niche themes quadrant (upper left) is occupied by a cluster of artificial intelligence, machine learning and human, which is internally well developed but only loosely connected to the other themes. A possible explanation for this is that the cluster captures AI applications in human-centred, largely clinical research, which is consistent with the medical terms observed on the right-hand side of the conceptual structure map. Finally, the emerging or declining themes quadrant (lower left) holds a small cluster of data integration, large language model and language model. Its low centrality and density, together with the recent publication dates of its documents (over 90% of the records mentioning large language models appear from 2025 onwards), suggest an emerging rather than a declining theme. These results should, however, be interpreted with caution, as variants such as machine learning and machine-learning are assigned to different clusters, which may distort the position of the AI theme.

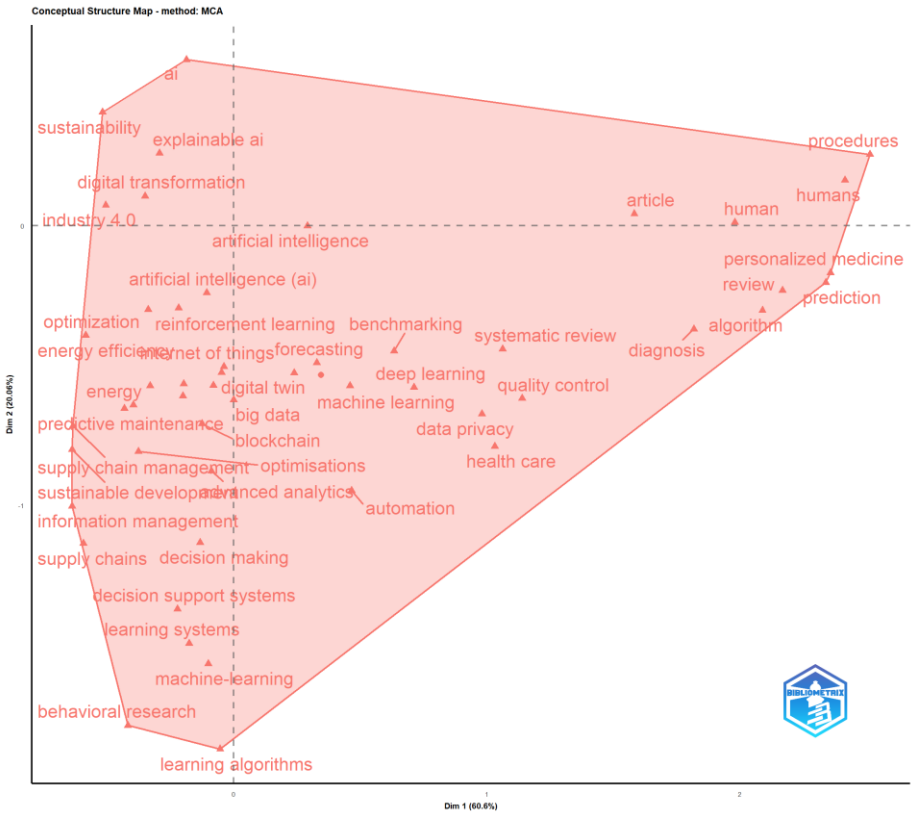

Diagram -1.5: Conceptual structure map (factorial map)

Diagram 1.5 presents the conceptual structure map generated through multiple correspondence analysis (MCA) of the keywords in the 1,550 articles and reviews retained for analysis. As can be seen from the map, the first dimension accounts for 60.6% of the variance and the second for 20.06%, which together capture approximately 80.7% of the variation in the keyword structure. A closer inspection of Dimension 1 reveals a clear separation between two bodies of literature. Keywords such as industry 4.0, supply chain management, energy efficiency, predictive maintenance and optimization are located on the left, indicating that these concepts frequently co-occur in studies applying AI-driven optimization to operational, industrial and sustainability problems. In contrast, diagnosis, personalized medicine, health care and prediction are positioned on the right, suggesting a distinct stream of clinical applications. It is interesting to note that general-purpose techniques, namely machine learning, deep learning, big data and digital twin, lie close to the origin, which implies that they are shared across both domains. Dimension 2 appears to distinguish strategic and conceptual terms (e.g., explainable ai, digital transformation) from methodological terms (e.g., learning algorithms, decision support systems, decision making), although this distinction is less pronounced. These findings should, however, be interpreted with caution. First, all keywords fall within a single cluster, so the map indicates the spread of themes rather than discrete thematic groups. Second, terms such as humans, procedures and article are database indexing terms rather than substantive concepts, and spelling variants such as optimization and optimisations, or artificial intelligence and ai, appear as separate points. Further research is therefore needed to harmonise synonyms before the analysis is repeated.

5. FUTURE RESEARCH AGENDA

The bibliometric data reveal AI for optimisation as a rapidly expanding, cross-sectoral and widely scattered research domain. Its development seems, however, to have far outpaced its conceptual and empirical foundations, with the bulk of research ending with prediction, claims of performance generally relying on simulation or benchmark data, each research sector largely pursuing its growth independently of each other and the human, ethical and cost implications of AI-based decisions hardly quantified. The primary research question for future research, then, is whether the gains reported for AI-based optimisation are also generated in real organisations, at a reasonable cost and risk.

Table 1 identifies six research priorities. Priorities 1 to 5 correspond to evidence gaps in substantive areas, and Priority 6 provides the methodological and conceptual infrastructure on which the others build.

Table -1: Summary of research priorities

Priority	Observation from the corpus	Core research question	Proposed approach	Expected contribution
1. From prediction to prescription	"Predictive" occurs in 837 of 1,550 abstracts; "prescriptive" in only 35	Under what circumstances does the addition of an optimisation step to an AI prediction make for a better decision?	Head-to-head benchmarking of predict-then-optimise, decision-focused learning and classical OR	Evidence of when AI adds prescriptive value
2. From benchmark to real-world performance	Gains reported typically depend on benchmark datasets, simulation or single case study	Do reported gains hold up under real operating constraints?	Meta-analysis by evaluation design; field pilots with real-world partners	Comparable, realistic performance estimates
3. Cross-sector integration	The largest separation dimension (60.6%) cleanly divides the industrial and clinical streams of the research field; no existing review combines multiple sectors	Under what conditions do methods transfer from one sector to another, and which do?	Taxonomy of OR problem types; transfer studies on equivalent problems in different sectors	A connected field rather than parallel silos
4. Explainability, trust and adoption	Explainable AI is an author keyword in 57 records; barriers to adoption listed but rarely studied	When do mangers understand, trust and act upon AI-recommended decisions?	Decision experiments; case studies of implementation	Design principles for systems that can be adopted
5. Fairness and the cost of AI	"Fairness" occurs in 32 abstracts, most within health; costs of computer-based work seldom reported	Do the benefits outweigh the risks in terms of fairness and the energy and financial costs of training and running the models?	Standard reporting for cost and fairness; net-benefit studies	A balanced account of net value
6. Conceptual coherence and evidence infrastructure	"Artificial intelligence" occurs under at least three labels; LLM work only beginning	What counts as AI for optimisation, and how much of its growth reflects real, substantive expansion?	A shared thesaurus; multiple-database living review	A cumulative, reproducible evidence base

5.1 Priority 1: From Prediction to Prescription

Gap: Predictive analytics occurs as an author keyword 140 times, the third most common term (out of 1,550 abstracts), and is located in the motor theme. Only 35 of 1,550 abstracts mention prescriptive analytics as an autor keyword, and classical OR methods like linear programming (13 abstracts) and mixed-integer programming (6) also occur seldom. Most studies thus conclude with a forecast but fail to go on to a decision. Bertsimas and Kallus (2020) assert that the connection between prediction and decision should be explicitly modelled, and Elmachtoub and Grigas (2022) demonstrate that a model that is trained to reduce prediction error does not necessarily yield the best decision.

Research questions: To what extent does an explicit step of optimisation add to the quality of a decision as opposed to just making a prediction? Under what conditions does training models for decision quality outperform the two-stage approach, predict-then-optimise?

Proposed approach: Benchmark studies of common problems like inventory, maintenance scheduling and energy dispatch would be asked to compare prediction with rule-based decisions, predict-then-optimise, decision-focused learning and classical stochastic or robust OR models. The same data should be used for all studies, and the cost of a decision the most important output for comparison.

Contribution: This would demonstrate when AI adds real prescriptive value and help reinvestigate the field's roots in OR.

5.2 Priority 2: From Benchmark to Real-World Performance

Gap: Many of the gains reported in the literature result from benchmark datasets, simulated environments or one-off case studies. For example, Arvind et al. (2025) reports a 32% lower forecast error compared with ARIMA and 35% lower inventory cost, and Skoropad et al. (2025) is able to control traffic signals with a reinforcement learning model; however, both studies rely on a single decision framework or are tested on a simulated network. Most studies use different data, baselines and metrics, so the results of each are incomparable.

Research questions: Do the reported gains hold up under the real-world operating constraints, data-quality issues and change? Which evaluation designs lead to the largest reported gains, and the least reliable results?

Proposed approach: First, a meta-analysis should categorise studies according to their evaluation design and extract effect sizes (benchmark, simulation, historical back-test or field deployment). For this, effect sizes, baselines and outcomes should be made comparable. Where multiple effects are reported per study, a robust variance estimation technique should be used to accommodate this. (Hedges et al., 2010.) Second, field pilots with real-industry partners should be undertaken. The incumbent decision process should be compared to an AI-based decision process for a primary outcome agreed on in advance.

Contribution: This would replace performance claims with comparable estimates of how AI is expected to perform in practice.

5.3 Priority 3: Cross-Sector Integration

Gap: The conceptual structure map cleanly divides the industrial and sustainability applications from the clinical applications along a first dimension that explains 60.6% of the variance between them, and then general techniques like machine learning and digital twins lie at the centre of the map. Existing reviews typically cover only one of these sectors, and none of the 100-odd records in the dataset makes any attempt to integrate the manufacturing, supply chains, energy, healthcare and food. A number of similar problems like scheduling, resource allocation and inventory show up under different names in different research communities.

Research questions: Which AI-based optimisation methods do transfer across sectors, and under what circumstances is this transfer limited? Is hospital scheduling a different problem to production scheduling simply because the problems themselves are not the same, or because the research communities are not the same?

Proposed approach: A taxonomy should categorise studies by the underlying OR problem type rather than by sector. Then transfer studies can test a method found in one sector but validated on an equivalent problem in another. Digital twins, as defined by Sajadieh and Noh (2025), span the entire life cycle of a system and are well placed to test the transfer of solutions in this way.

Contribution: This will transform the parallel sector-based literatures into a connected field in which lessons and methods move freely.

5.4 Priority 4: Explainability, Trust and Adoption

Gap: Explainable AI is an author keyword in 57 records. The focus is mostly on the accuracy of predictions or optimal solutions, without reference to whether the people making decisions are able to understand them or have confidence in and will act upon them. This is partly due to well-known barriers of interoperability, data and organisational barriers, but these are mentioned only rarely if at all (Sajadieh & Noh, 2025). The growing research on algorithm aversion reveals that people often downweight advice from algorithms even when they are provably correct (Burton et al., 2020), but an excessive reliance on AI is no less of a risk.

Research questions: What kinds of explanation bring with them appropriate levels of reliance without people simply accepting or rejecting the prediction/decision without scrutiny? What organisational conditions determine whether AI-based optimisation is used, over-ridden or abandoned?

Proposed approach: Experiments in decision-making should be used with operation managers, supply chain managers and clinical managers. These will test how different types of explanation for a given recommendation and different rules regarding rights and responsibilities to override the recommendation, affect the quality of the decision made, how frequently the decision is over-ridden and how well individuals' trust is calibrated. A number of multi-site case studies will be used to trace how each AI-based optimisation implementation works its way through the decision-making process, succeeding or failing.

Contribution: This will bring the design of models back to organisational reality, and produce principles for systems that can be relied upon and used by people.

5.5 Priority 5: Fairness and the Cost of AI

Gap: Fairness occurs in 32 abstracts, almost exclusively for health, but Hussain et al. (2025) find evidence of racial and ethnic bias in AI algorithms applied to healthcare. Sustainability appears 70 times as a keyword, and yet the energy and financial cost of training and running AI models is seldom reported in comparison with the efficiency gains they are claimed to create, although these are not inconsequential, as reported in, for example, Strubell et al. (2019).

Research questions: Do the benefits of AI-based optimisation outweigh the risks of unfairness and the computing costs of deployment? How are benefits of the optimised allocations distributed among different groups?

Proposed approach: A full-text synthesis should code the outcome measures as firm-level or stakeholder level, such as cost, accuracy, fairness, access and welfare. A standard template for reporting should be defined that includes the time and energy used in training and making decisions with these models as well as standard measures for fairness. Net-benefit studies should compare the benefits of simpler OR models with the large AI models.

Contribution: This will produce a balanced account of whether AI-based optimisation creates net value, and for whom.

5.6 Priority 6: Conceptual Coherence and Evidence Infrastructure

Gap: The same concept appears under different labels. For example, there are 478 records where the phrase "artificial intelligence" appears, and 100 where the phrase "artificial intelligence (AI)" appears. There are 55 that use "AI" as an author keyword. There are 50 records with the term "digital twin" and 34 with the term "digital twins". There are three different variants of the internet of things. This discrepancy, and others, distort the thematic map. The term "operations research" appears in two abstracts. Large language models, particularly in AI and manufacturing, are only just appearing as research and are probably under-represented. These models often first appear as the main contribution to conference papers and preprints.

Research questions: What is counted as AI for optimisation? How much of the field's apparent growth represents true, substantive growth?

Proposed approach: First a shared thesaurus to map onto the same concepts. Second, searches should include Scopus, the Web of Science and conference proceedings. Third, reporting should be undertaken following PRISMA 2020 guidelines (Page et al., 2021) and an independent double coding process implemented. Fourth, the work should be updated in a living review process at regular intervals.

Contribution: This will increase comparability and reproducibility of results and provide the base upon which the previous five priorities are built.

5.7 Conclusion

The research agenda is focused on one question relevant to practice: under what circumstances does AI-based optimisation produce a better decision in practice, at a cost and risk that organisations and society are willing to accept? Conceptual coherence and a common infrastructure of shared evidence for these constitute the bottom. Prescriptive design and the test in real-world settings can prove whether there is value being created or not. A cross-sector transfer and adoption, combined with studies of fairness and cost, determine whether that value can be sustained and distributed fairly. If pursued in concert, these research prioritises will help transform the field's rapid growth into a cumulative research programme that managers, engineers and policy makers can rely upon.

6. CONCLUSION

The study set out to map how artificial intelligence is being applied to optimisation problems, and how the resulting body of research is organised. By using a bibliometric approach on 1,550 Scopus-indexed articles and reviews published between 2016 and 2026 in 894 sources and a structured reading of the literature, it finds a field that is growing very quickly, and spreading across a large number of sectors, but whose parts remain weakly connected and whose evidence base is yet to catch up with its methods. This section summarizes its main findings, discusses what they tell us and sets out the limitations of the data.

The results address all three of the objectives stated in the introduction. As to the field's development, the field is very young, highly collaborative and scattered across a large number of journals. As to the field's structure, analytics-based decision-making makes up the motor of the field, while optimisation and machine learning are its foundation, and healthcare and generative AI sit on its periphery. As to the gaps that it leaves in the field, the findings are those of a shortage of prescriptive work, limited real-world validation, weak integration between sectors, and inconsistent terminology.

Until 2018 output was negligible, just five documents in total, before it rose to 11 in 2019 and 50 in 2021. It then fluctuated (34 in 2022 and 51 in 2023) before accelerating sharply to 147 in 2024, 546 in 2025 and 675 in 2026. The last three years make up 1,368 documents, or 88.3% of the dataset as a whole, and 2026 already exceeds 2025 by 23.6% despite running only part of a year. The 78.99% annual growth rate reported is inflated by the very small base of two documents in 2016, and the linear trend line only fits moderately ($R^2 = 0.59$), which confirms that growth is concentrated in the most recent years, and not steady. This timing also coincides with the spread of generative AI and with wider interest in digital transformation and sustainability, although the bibliometric data is not fine-grained enough to separate these causes.

The research into this area is overwhelming team-based. Only 159 of 5,822 authors (2.7%) have ever written single-authored documents, documents have a 4.21 average number of co-authors, and 32.65% of all documents feature international co-authorship. Author affiliation from India is the most frequently represented, followed by those from China and the United States. The literature is also highly dispersed: 670 of the 894 sources (74.7%) have each contributed just one document to the field and the ten sources that publish the most documents each account for only 13.1% of output, a patten consistent with Bradford's (1934) law of scattering. The source leading the field is IEEE Access, with 35 articles, followed by Applied Sciences and Scientific Reports, with 30 apiece and Sustainability, with 28. Before 2022 only IEEE Access published a paper on a regular basis in this field, but from 2024 the journals that focus on multi- or inter-disciplinarity and sustainability have quickly caught up, and around 79% of the articles published in those five leading sources were written in 2025 and 2026. Productivity is not a measure of influence, however: Sensors reports an average of 62.9 citations per article against a dataset mean of 18.5, and Scientific Reports an average of 9.2. The 5,516 distinct author keywords represent a diverse vocabulary that cuts across sectors.

The thematic map comprises four clusters, located by centrality and density (Cobo et al., 2011). The motor theme is made up of predictive analytics, decision making and internet of things, and is both central and well developed, which indicates that analytics-based decision support drawing on data from the internet of things is the most mature strand of the field.

Learning systems and machine learning and optimisation form a basic theme; they are central, but lie close to the median lines, which is to say that they constitute the shared foundation that links the motor theme to the remainder of the literature. The cluster of terms made up of the words artificial intelligence, machine learning and human makes up a niche, that is internally cohesive but weakly connected, and most probably captures the applications of the field in clinical settings. The emerged or declining cluster is made up of data integration and the large language model cluster; given that more than 90% of the records on large language models date from 2025 onwards, the theme is in the former rather than the latter category. The conceptual structure map supports this interpretation. Its first dimension, which accounts for 60.6% of the variance, separates an industrial and sustainability stream (Industry 4.0, supply chain management, energy efficiency and predictive maintenance) on the one side from a clinical stream (diagnosis, personalised medicine and health care) on the other, while general purpose techniques like machine learning, deep learning, big data and digital twins are found close to the origin and are shared by both.

The structured reading of the literature offers a consistent picture. There are two logics running through the field. In the first a decision-making process is followed by an optimisation step and a prediction becomes a decision (Lepenioti et al., 2021; Miikkulaisen et al., 2021). In the second, optimisation algorithms are used to train and tune the AI models themselves. Digital twins and explainable AI form the links between the two and are increasingly treated as preconditions for using AI in real operations (Sajadieh & Noh, 2025; Teixeira et al., 2025).

Both sources indicate the same overall picture, that is, that the development of the field's methods have advanced much faster than the evidence on how these methods perform in real circumstances. The five gaps that stand out from the data are the following. First, prediction dominates over prescription. The word "predictive" occurs in 837 of the abstracts, but "prescriptive" occurs in just 35, and the classical OR terms like linear or mixed-integer programming are but seldom found. Second, many reported gains result from a benchmark dataset, a simulation or a single case study; the 32% reduction in forecast error and 35% reduction in cost reported in Arvind et al. (2025), for example, are based on just one framework and dataset, and results like these are seldom compared in the same conditions. Third, sectors are studied individually, but no review in the set of records connecting manufacturing, supply chains, energy, healthcare and food. Fourth, barriers such as data availability, interoperability between systems and organisational adoption of techniques are often mentioned, but seldom studied in any detail (El Mariami et al., 2025). Fifth, fairness and the costs of training and deploying AI are seldom considered against its advantages and are discussed mainly in the healthcare (Hussain et al., 2025). The gaps mentioned are based on an analysis of the abstracts and keywords and would require confirmation by a more in-depth coding of the full texts.

The main contribution of this study is its picture of a cross-sector field. Previous reviews looked at a single sector, supply chains, logistics, maintenance, digital twins, food or power systems, or at a single technology, like AI or optimization. Bring them together in the same picture reveals how optimisation is the linking thread between the motor theme of analytics and the sectoral clusters, how healthcare constitutes a largely independent stream and, finally, that large language models are a new addition, only beginning to enter the field.

The next useful step might then be less to add another single-sector algorithm to the field, but to draw together what already exists: to compare AI-based methods and classical OR approaches with the same data, to move beyond forecasts to decisions and, finally, to test if methods transfer across sectors. For managers the message might be that headline performance numbers should be treated with caution until they have been tested under realistic data, constraints and costs conditions, and that attention should be paid to whether the people who have to act on the recommendation of a model can understand and trust those recommendations.

The findings are subject to several limitations. Only Scopus was searched, so work that has been indexed elsewhere may have been left out, as well as conference papers, which are an important outlet in the computer science world. The search string used ("AI AND Optimization AND analytics") may have overlooked relevant studies using other terminology, and may have included some oddities, such as those from Shodhkosh: Journal of Visual and Performing Arts. Because 2026 is but a partial year, the counts and citations in recent years are underexpressed. Variants of the key terms (optimization and optimisation, AI and artificial intelligence, digital twin and digital twins) and database indexing terms, like human and article, were not fullly harmonized, which affects the network maps (Aria & Cuccurullo, 2017). Some of the interpretations, for example linking the niche cluster to clinical research, are based on a read of the maps themselves rather than by direct measurement.

Two directions seem to be especially fruitful: the validation of AI-based optimisation in real operating conditions relative to the classical OR baselines, as well as linking the distinct streams in sectoral studies through shared terminology and problems. Repeating the analysis with Web of Science, cited references and a complete 2026 would paint a more solid picture. In either case the next stage of progress in AI for optimisation is likely to be based less on new algorithms than on showing which existing algorithms lead to better decisions, in practice

REFERENCES

[1] Al-Thani, Hadid, Padmanabhan, Kerbache, & Elomri. (2025). Smart food supply chain management: A bibliometric and systematic review. Food and Humanity. https://doi.org/10.1016/j.foohum.2025.100736

[2] 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

[3] Aria, M., Cuccurullo, C., D'Aniello, L., & Spano, M. (2026). biblioshiny and the SAAS workflow. journal of informetrics.

[4] Arvind, V. R., Shrinidhi, R. M., Deepa, T., & Maheedhar, M. (2025). Intelligent warehousing: A machine learning and IoT framework for precision inventory optimization. IEEE Access, 13, 169381–169414. https://doi.org/10.1109/ACCESS.2025.3614679

[5] Bengio, Y., Lodi, A., & Prouvost, A. (2021). Machine learning for combinatorial optimization: A methodological tour d'horizon. European Journal of Operational Research, 290(2), 405–421. https://doi.org/10.1016/j.ejor.2020.07.063

[6] Bertsimas, D., & Kallus, N. (2020). From predictive to prescriptive analytics. Management Science, 66(3), 1025–1044. https://doi.org/10.1287/mnsc.2018.3253

[7] Bradford, S. C. (1934). Sources of information on specific subjects. Engineering, 137, 85–86.

[8] Burton, J. W., Stein, M.-K., & Jensen, T. B. (2020). A systematic review of algorithm aversion in augmented decision making. Journal of Behavioral Decision Making, 33(2), 220–239. https://doi.org/10.1002/bdm.2155

[9] Callon, M., Courtial, J.-P., & Laville, F. (1991). Co-word analysis as a tool for describing the network of interactions between basic and technological research: The case of polymer chemistry. Scientometrics, 22(1), 155-205. https://doi.org/10.1007/BF02019280

[10] Chen, W., Men, Y., Fuster, N., Osorio, C., & Juan, A. A. (2024). Artificial intelligence in logistics optimization with sustainable criteria: A review. Sustainability, 16(21), 9145. https://doi.org/10.3390/su16219145

[11] Cobo, M. J., López-Herrera, A. G., Herrera-Viedma, E., & Herrera, F. (2011). An approach for detecting, quantifying, and visualizing the evolution of a research field: A practical application to the Fuzzy Sets Theory field. Journal of Informetrics, 5(1), 146–166. https://doi.org/10.1016/j.joi.2010.10.002

[12] 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

[13] Dr Richa Saxena, A. S. (2024). A BIBLIOMETRIC REVIEW OF SUSTAINABLE SUPPLY CHAIN MANAGEMENT: AN OR PERSPECTIVE. International Journal of Business Management and Economic Review.

[14] Ehteram et al. (2021). Alexandria Engineering Journal. https://doi.org/10.1016/j.aej.2020.12.034

[15] El Mariami, I., Mharzi, H., El Oualidi, M. A., & Abadi, M. (2025). Operations research and artificial intelligence for supply chain planning: A systematic literature review. International Journal of Supply and Operations Management, 12(4), 563–600. https://doi.org/10.22034/ijsom.2025.110642.3282

[16] Elmachtoub, A. N., & Grigas, P. (2022). Smart "predict, then optimize". Management Science, 68(1), 9–26. https://doi.org/10.1287/mnsc.2020.3922

[17] Emmert-Streib, F. (2023). What is the role of AI for digital twins? AI, 4(3), 721–728. https://doi.org/10.3390/ai4030038

[18] Hector, & Panjanathan. (2024). Predictive maintenance in Industry 4.0: A survey of planning models and machine learning techniques. PeerJ Computer Science, 10, e2016. https://doi.org/10.7717/peerj-cs.2016

[19] Hedges, L. V., Tipton, E., & Johnson, M. C. (2010). Robust variance estimation in meta-regression with dependent effect size estimates. Research Synthesis Methods, 1(1), 39–65. https://doi.org/10.1002/jrsm.5

[20] Hussain, S. A., Bresnahan, M., & Zhuang, J. (2025). The bias algorithm: How AI in healthcare exacerbates ethnic and racial disparities – A scoping review. Ethnicity & Health, 30(2), 197–214. https://doi.org/10.1080/13557858.2024.2422848

[21] Iseri et al. (2025). Journal of Global Optimization. https://doi.org/10.1007/s10898-025-01509-1

[22] Jaramillo, M., Carrión, D., Muñoz, J., & Tipán, L. (2025). A bibliometric assessment of AI, IoT, blockchain, and big data in renewable energy-oriented power systems. Energies, 18(12), 3067. https://doi.org/10.3390/en18123067

[23] Kumar, D., Addula, S. R., Lind, M., Brown, S., & Odion, S. (2026). AI-driven hybrid deep learning and swarm intelligence for predictive maintenance of smart manufacturing robots in Industry 4.0. Electronics, 15(3), 715. https://doi.org/10.3390/electronics15030715

[24] Lepenioti, K., Bousdekis, A., Apostolou, D., & Mentzas, G. (2021). Human-augmented prescriptive analytics with interactive multi-objective reinforcement learning. IEEE Access, 9, 100677–100693. https://doi.org/10.1109/ACCESS.2021.3096662

[25] Lim, J.-B., & Jeong, J. (2023). Factory simulation of optimization techniques based on deep reinforcement learning for storage devices. Applied Sciences, 13(17), 9690. https://doi.org/10.3390/app13179690

[26] Miikkulainen, R., Francon, O., Meyerson, E., Qiu, X., Sargent, D., Canzani, E., & Hodjat, B. (2021). From prediction to prescription: Evolutionary optimization of nonpharmaceutical interventions in the COVID-19 pandemic. IEEE Transactions on Evolutionary Computation, 25(2), 386–401. https://doi.org/10.1109/TEVC.2021.3063217

[27] Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71. https://doi.org/10.1136/bmj.n71

[28] Parisineni, & Pal. (2024). International Journal of Data Science and Analytics. https://doi.org/10.1007/s41060-023-00458-w

[29] Paul, J., & Rialp, A. (2020). The art of writing literature review : What do we know and what do we need to know ?

[30] Pessach, D., Singer, G., Avrahami, D., Chalutz Ben-Gal, H., Shmueli, E., & Ben-Gal, I. (2020). Employees recruitment: A prescriptive analytics approach via machine learning and mathematical programming. Decision Support Systems, 134, 113290. https://doi.org/10.1016/j.dss.2020.113290

[31] Pieper, D., & Puljak, L. (2021). Language restrictions in systematic reviews should not be imposed in the search strategy but in the eligibility criteria if necessary. Journal of Clinical Epidemiology.

[32] Saini, Thakur, & Yadav. (2025). RSC Pharmaceutics. https://doi.org/10.1039/d4pm00323c

[33] Sajadieh, S. M. M., & Noh, S. D. (2025). From simulation to autonomy: Reviews of the integration of artificial intelligence and digital twins. International Journal of Precision Engineering and Manufacturing-Green Technology, 12(5), 1597–1628. https://doi.org/10.1007/s40684-025-00750-z

[34] Skoropad, V. N., Deđanski, S., Pantović, V., Injac, Z., & Vujičić, S. (2025). Dynamic traffic flow optimization using reinforcement learning and predictive analytics: A sustainable approach to improving urban mobility in the city of Belgrade. Sustainability, 17(8), 3383. https://doi.org/10.3390/su17083383

[35] Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (pp. 3645–3650). https://doi.org/10.18653/v1/P19-1355

[36] Teixeira, B., Carvalhais, L., Pinto, T., & Vale, Z. (2025). Explainable AI framework for reliable and transparent automated energy management in buildings. Energy and Buildings, 347, 116246. https://doi.org/10.1016/j.enbuild.2025.116246