iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** X **Month of publication:** October 2026

DOI: https://doi.org/10.22214/ijraset.2026.85317

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue X Oct 2026- Available at www.ijraset.com

Vehicle Routing Problem Research and the Quick Commerce Gap: A Bibliometric Analysis of Scopus Literature, 2011 to 2026

Dr. Richa Saxena[1], Fenil Gindra[2], Deep Soni[3], Divyan Arora[4], Falguni Gupta[5], Fazal Patel[6]
SY BBA (Division F), Anil Surendra Modi School of Commerce, SVKM's NMIMS, Mumbai, India

Abstract: ***The Vehicle Routing Problem (VRP) is one of the most studied problems in operations research, yet it is unclear whether research has followed the rise of quick commerce and hyperlocal delivery. This paper presents a bibliometric analysis of VRP literature indexed in Scopus. Records were retrieved on 16 September 2026 using the terms vehicle routing, optimization and heuristics, screened under the PRISMA framework, and reduced from 4,364 records to 2,559 English-language journal articles and reviews published from 2011 onwards. The data were analysed with Biblioshiny, VOSviewer and Tableau. Output grew at an annual rate of 14.51%, from 41 documents in 2011 to 313 in 2026, with 3.48 authors per document and 31.62% international co-authorship. Authorship follows a generalised Lotka pattern, sources follow Bradford's Law closely, and China leads in output ahead of the USA and India. Keyword analysis shows core VRP methodology terms far outnumbering last-mile delivery, while quick commerce and hyperlocal delivery barely appear. An input-process-output analysis shows that solution methods receive more attention than service and fairness outcomes. The paper sets out research questions on quick commerce, drone delivery, sustainability trade-offs and emerging-market settings, and discusses implications for theory, management, policy and society.***

Keywords: Vehicle routing problem, bibliometric analysis, Biblioshiny, quick commerce, last-mile delivery, research agenda.

I. INTRODUCTION

Every product that reaches a shop shelf, a hospital ward or a front door has travelled on a vehicle whose route someone had to plan. These decisions are made millions of times a day and their combined effect is visible in fuel consumption, congestion, air quality, prices for deliveries and working hours of drivers. Bad routes cost companies money through wasted fuel and late deliveries and they cost society through avoidable emissions and crowded roads. Good routes reduce the cost of moving goods and reduce the distance driven to deliver the same volume. Improvements to routing are therefore a win for consumers, business and public authorities alike, as transportation is a large part of the cost of any physical product. Operations research (OR) gives a formal framework for these decisions. Each is formulated as an optimisation problem with an objective, a set of constraints and a space of possible solutions that can be searched systematically (Laporte, 2009; Toth & Vigo, 2014).

This is further weighted by environmental pressure. Road transport is a major source of greenhouse gas emissions and urban air pollution and cities with growing numbers of delivery vehicles also have to manage noise, parking and road safety. Routing models that consider emissions and energy consumption as well as distance (Khalid et al., 2026) show a growing appreciation that a good route is not simply the shortest one.

The formal study of this problem began with Dantzig and Ramser (1959) who modelled the dispatch of a fleet of trucks from one centralised hub to petrol stations at least possible distance. The Vehicle Routing Problem (VRP), in its current form, is to determine a set of routes that start and finish at a depot and service all customers at minimum cost. Laporte (2009) surveyed the first fifty years of progress on exact and heuristic methods for it and the problem has since become one of the most studied topics in combinatorial optimisation.

The VRP is unique in OR in that it is easy to formulate and hard to solve. It is a generalisation of the travelling salesman problem. Lenstra and Rinnooy Kan (1981) proved that vehicle routing and scheduling problems are NP-hard. In practice, no algorithm is known to be able to guarantee an optimal solution in a reasonable amount of time when the number of customers is large. Precise methods such as branch-and-cut are viable for small and medium-sized instances, but operational reality exceeds them rapidly. The field has therefore turned to heuristics and metaheuristics which give up the guarantee of optimality in exchange for quickly finding good solutions.

©IJRASET: All Rights are Reserved

Recent work includes genetic algorithms, ant colony optimisation, variable neighbourhood search and guided local search (Du et al., 2026; Rachmawati et al., 2026; Zhou et al., 2026; Khalid et al., 2026). The availability of benchmark instances, such as the Solomon instances for time windows, which allow routine comparison of new algorithms (Du et al., 2026), has also driven progress. This common yardstick allows the results to be compared. It can also highlight problems for which benchmarks already exist, to which we return in the conclusion.

The classical problem has led to a large family of variants. Restrictions on vehicle capacities, customer time windows, multiple depots, pickups and deliveries, mixed fleets, and time-dependent travel speeds all change the problem structure and the methods needed to solve it. Now mathematical programming is mixed with heuristic search (Vadseth et al., 2026). Learning-based approaches such as graph neural networks and deep reinforcement learning have entered the field (Xiao et al., 2026; Zabraoui et al., 2026). Also, recent studies bring in drones, electric vehicles and congestion (Kong et al., 2026; Wang et al., 2026). The VRP is thus a good domain to see how OR methods evolve with the problems posed to them.

The environment in which routing decisions are made has likewise evolved quickly. Online retail has moved much of the buying from stores to people's homes, and consumers increasingly expect delivery within hours, or even minutes. The most obvious example of this shift is fast commerce and hyperlocal delivery. Small fulfilment points in neighbourhoods, often called dark stores, collect short baskets of groceries and daily items which are delivered to customers by riders on two wheelers within a very short promised window. In a short time, the model has become a common feature of urban life in many Indian cities. At the core are routing decisions: which orders to batch, which rider to assign and in which order to deliver.

These services have different routing requirements from traditional parcel delivery. Orders flow continuously and cannot be known in advance, delivery promises are very short, distances are small, vehicles are light and plans have to be constantly revised. Each of these features raises questions about dynamic dispatch, rider workload, congestion and emissions on dense urban streets. Meanwhile publications on VRP have flourished, including work on synchronised vehicle and drone delivery (Kong et al., 2026) and on generative artificial intelligence as a routing tool (Ramingwong & Jintana, 2026). A reasonable question is whether routing research has followed commercial practice into quick commerce and this paper sets out to test that.

This paper is written by a group of undergraduate BBA students who want to know how the VRP literature has developed and where it has gone wrong. Our problem is the disconnect between research and commercial practice. Quick commerce and hyperlocal delivery are a visible and fast-growing reality, but our reading of the field suggests that they barely appear in VRP research. Our keyword counts show 104 mentions for core VRP methodology vs. only 4 for last-mile delivery. What has not been done is to quantify the extent to which these delivery models are represented in the literature. To abstain from it makes the impression that research lags practice. If the pattern is consistent, the field has developed a deep repertoire of methods for problems that are only partially like those delivery businesses now face. This conclusion applies to papers which use the language of routing, optimisation and heuristics. Thus it characterises the VRP literature as defined by that search and excludes research on delivery in other fields.

The reason for the study is recency. The output of VRP is growing at such a rate that a map of its content quickly becomes outdated, and bibliometric review is a suitable way to synthesise a large body of work, detect patterns and gaps (Paul & Criado, 2020). The study has four aims. The first is to describe the growth and structure of VRP research from 2011 to 2027 through bibliometric indicators. The second is to identify authors, journals and countries that produce most of the work. The third is to explore the themes of the field through its keywords. The fourth is to assess the presence of last-mile delivery, quick commerce and hyperlocal delivery and to delineate the implications for theory, management, policy and society.

We searched Scopus on 16 September 2026 with the terms vehicle routing, optimisation and heuristics in titles, abstracts and keywords. We retrieved 4,364 records according to PRISMA framework (Moher et al., 2010), keeping the 3,740 published after 2011, limiting these to articles and reviews (2,766), then to English-language papers (2,578) and finally to journal sources, resulting in 2,559 papers. Data were analysed using the Biblioshiny interface of the bibliometrix package (Aria & Cuccurullo, 2017) supported by VOSviewer and Tableau. Annual production increased from 41 documents in 2011 to 313 in 2026, authorship and journals are concentrated in a small core, and keyword analysis indicates last-mile delivery as marginal and quick commerce and hyperlocal delivery as close to absent.

The remainder of this paper is organised as follows. The literature review presents the main methods and variants of problems described in recent research of VRP. The research methodology details the search, screening and analytical procedure. The data analysis shows the bibliometric results in terms of production, citations, authors, journals and countries. The future research agenda presents directions that are drawn from the identified gaps. The conclusion sums up the findings, implications and limitations of the study.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

II. LITERATURE REVIEW

The Vehicle Routing Problem (VRP) remains one of the central optimization problems in logistics because it connects a basic operational decision (how a fleet should visit a set of customers) with broader concerns such as cost, travel time, service quality, vehicle capacity, environmental impact, and increasingly complex delivery technologies. The literature reviewed here shows that VRP research is moving well beyond the classical formulation toward problems that combine time windows, multiple depots, pickup and delivery requirements, heterogeneous fleets, drones, electric vehicles, congestion, and sustainability objectives. At the same time, solution methods are also changing. Traditional heuristics and metaheuristics continue to play an important role, but they are increasingly being combined with mathematical programming, deep reinforcement learning, attention-based models, graph neural networks, and generative artificial intelligence. This development suggests that contemporary VRP research is less about replacing one method with another and more about finding ways to combine complementary strengths in order to solve increasingly realistic routing problems.

A useful starting point is the continued importance of heuristic and metaheuristic search. The review by Kong, Wei, and Jiang (2026) on vehicle-drone delivery routing shows how the growing complexity of last-mile delivery has created several different synchronization settings between vehicles and drones. Their classification distinguishes non-synchronized, low-synchronization, and high-synchronization scenarios, and compares the associated exact and heuristic approaches. This perspective is important because it shows that routing difficulty is not determined only by the number of customers or vehicles; the relationships between delivery resources can fundamentally change the structure of the problem. Similar complexity can be seen in the work of Duan et al. (2026), who study a time-dependent vehicle routing problem with drones, simultaneous pickup and delivery, time windows, energy consumption, differentiated service times, and multiple objectives. Their multi-objective hyper-heuristic memetic algorithm combines several low-level heuristics with estimation-of-distribution and decomposition-based local search. The study illustrates how hybrid search can be designed around the specific constraints of a routing problem rather than applying a generic metaheuristic without adaptation.

The continuing relevance of constructive and population-based methods is also visible in research on the Vehicle Routing Problem with Time Windows (VRPTW). Du, Tanimizu, and Watanabe (2026) propose a Hybrid Constructive Genetic Algorithm that integrates Solomon's insertion heuristic with a sequential constructive crossover operator. The underlying idea is that good initial individuals can strengthen the evolutionary search and improve the balance between global exploration and local refinement. Their experiments on the Solomon100 benchmark report best-known solutions for a number of instances while remaining close to the best-known values on the others. In another application, Rachmawati, Lohil, and Tarigan (2026) develop an enhanced Ant Colony Optimization approach for the capacitated VRP with time windows in franchise distribution. These studies demonstrate that classical metaheuristic ideas, namely constructive heuristics, genetic search, pheromone-based learning, and local improvement, remain useful when they are carefully adapted to the structure of the routing constraints.

Another important direction is the integration of mathematical programming with heuristic search. Vadseth et al. (2026) introduce route-modifying improvement models that use mixed-integer programming to improve solutions generated by local search and metaheuristics. Their approach is evaluated on several VRP variants, including the multi-trip VRP, pickup and delivery with time windows, and VRPTW. Rather than treating exact optimization and heuristics as competing alternatives, the study uses mathematical programming as an improvement and post-optimization mechanism. This matheuristic perspective is particularly relevant for large routing problems, where obtaining an exact optimum may be computationally difficult but where a high-quality heuristic solution can provide an effective starting point. The work therefore reinforces a broader trend in the literature: strong VRP solvers increasingly combine fast heuristic exploration with targeted exact optimization in order to improve solution quality without sacrificing practicality.

Artificial intelligence is becoming another major component of VRP research. Zabraoui, Chafi, and Alami (2026) discuss the integration of graph neural networks and reinforcement learning for VRP and related variants such as the capacitated VRP and VRP with time windows. Their work emphasizes the ability of graph-based learning to represent transportation networks while reinforcement learning frames routing as a sequential decision process. This direction is extended by Xiao et al. (2026), who investigate deep reinforcement learning for the vehicle routing problem with route balancing. Their Equity Attention Model is designed for a multi-objective setting in which route workload balance is considered alongside travel cost and fleet constraints. The proposed two-stage learning strategy is also intended to improve scalability to larger instances. These studies indicate that learning-based routing is increasingly addressing practical limitations that were often simplified in earlier neural approaches, particularly fleet size, multiple objectives, and generalization across problem instances.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

The same shift toward learned improvement can be seen in Yang et al. (2026), who propose a Collaborative Adaptive Fusion Transformer for the Capacitated Vehicle Routing Problem. Instead of only constructing a route from scratch, the method iteratively improves complete solutions. Its architecture combines node information with visit-index information and uses a task-adaptive reinforcement learning framework with curriculum learning. The emphasis on iterative improvement is significant because it brings learning-based methods closer to the logic of traditional heuristics, where a feasible solution is repeatedly modified through local changes. Wang, Zhang, and Wang (2026) similarly combine heuristic information with attention mechanisms for the Electric Vehicle Routing Problem with Time Windows. Their approach explicitly incorporates time-window information and dynamic heuristic information into the attention process to address the interaction between customer service requirements and battery limitations. Together, these studies suggest that future AI-based routing methods may benefit most when they incorporate problem-specific heuristic knowledge rather than relying on purely data-driven representations.

The expansion of VRP into drone-supported and time-dependent logistics provides another important research direction. Zhou et al. (2026) study a truck-drone routing problem under time-dependent traffic congestion and load-dependent drone endurance. Their variable neighborhood search metaheuristic uses parallel exploration, structured restart, and decomposition-based re-optimization to handle large-scale instances. The problem is particularly relevant because truck travel times can vary with congestion while drone energy requirements can depend on payload. Such interactions make simple distance minimization inadequate for realistic routing decisions. Duan et al. (2026) reach a similar conclusion from a multi-objective perspective, showing that synchronization, waiting, energy use, service times, and time-dependent roads must be considered together when vehicle-drone systems are evaluated. The literature therefore increasingly treats routing as a coordinated system rather than as an isolated vehicle sequencing problem.

Sustainability is also becoming a stronger component of VRP research. Khalid et al. (2026) develop an emission-capacitated vehicle routing model for urban waste collection and use a hybrid guided local search to address the resulting optimization problem. The study reflects a broader movement from purely distance-based routing toward objectives that include emissions and environmental constraints. Electric vehicle routing research extends this challenge further because route feasibility depends not only on customer locations and vehicle capacity but also on energy consumption, charging requirements, and service times. Wang et al. (2026) show how these factors create additional interactions between time-window feasibility and battery capacity. As logistics systems become more sustainability-oriented, routing models are therefore expected to account for environmental and energy-related variables without losing the computational efficiency required for operational use.

A particularly recent development in the reviewed literature is the use of generative AI for routing. Ramingwong and Jintana (2026) investigate whether a generative AI system can produce competitive solutions for a real-world vehicle routing case involving 80 service locations. Their comparison includes a linear-programming spreadsheet approach, a commercial heuristic, and a genetic algorithm. The results reported in the study show that the genetic algorithm produced the shortest best trial, while the generative AI approach also produced competitive solutions with relatively low setup requirements. This is an interesting shift because it frames VRP not only as an algorithmic optimization problem but also as an accessibility problem. Conventional optimization tools can require specialist modelling knowledge, whereas generative AI may lower the technical barrier for smaller organizations. At the same time, the probabilistic nature of generative systems means that consistency, reproducibility, constraint handling, and verification remain important considerations before such methods can be treated as dependable operational solvers.

Overall, the reviewed literature presents VRP research as a field moving toward hybridization. Classical heuristics remain valuable because they provide fast, interpretable, and problem-specific search mechanisms. Metaheuristics such as genetic algorithms, ant colony optimization, variable neighborhood search, guided local search, and hyper-heuristics continue to perform important roles, especially when routing constraints become too complex for simple constructive methods. Mathematical programming adds a mechanism for exact improvement, while reinforcement learning, graph neural networks, and transformer-based architectures introduce adaptive learning and scalable solution construction or improvement. Drone routing, electric vehicle routing, time-dependent traffic, route balancing, pickup and delivery, and sustainability further expand the problem space. The literature therefore suggests that the main challenge is no longer simply finding a route, but designing solution frameworks that remain effective when several operational constraints interact simultaneously. Future research is likely to place greater emphasis on generalization across problem distributions, transparent hybrid AI methods, real-time adaptation, energy and emissions considerations, and reliable integration of learned policies with established optimization techniques. In this sense, the evolution of the VRP literature points toward increasingly integrated systems in which heuristics, exact methods, and artificial intelligence complement rather than replace one another.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

III. RESEARCH METHODOLOGY

Systematic Literature Review (SLR) is an all-encompassing technique that aims in the understanding and gaining valuable insights in a subject area by synthesising the exact literature. SLR techniques are used to identify gaps for future research by reviewing and understanding the past trends and prior knowledge on the specific area of study. (Paul & Criado, 2020) This current study employs a qualitative methodology to critically evaluate past trends and developments in the area of Vehicle Routing Problem, using data obtained through keywords and using citations analysis in peer-reviewed journal articles and review papers.

A. Article Selection, Inclusion and Exclusion

The database search was conducted on 16th September, 2026 on the Elsevier's Scopus database, using search keywords appearing in the title, abstract and keywords using the PRISMA framework, which enables data collection as an iterative process which is well-documented and replicable. (Moher et al., 2010)

The flowchart for the PRISMA is given in Figure 1. The keywords "vehicle routing" AND "optimization" AND "heuristics" were used to conduct the search. The database search revealed publications ranging all the way back in time from 1974. However, the study excluded the articles published before 2011 and it showed 3740 documents in the past 15 years before filtering. It should be noted that the authors considered whether there were any skip years during the data collection. Skip years could have created gaps in the data which could have caused disrupted consistency. However, it was found that no skip years occurred within the time period (2011-2027) covered by the data used by the authors. If the whole data ranging from 1974 was to be taken, it was not only bound to have skip years but more importantly there was a small number of papers to accurately identify the trend before 2011.

It was observed that the search results for the past 15 years showed that papers were written in various subject areas, and after careful consideration it was finalised that all the subject areas will be included in the paper as they were within the scope of the research. The research areas included computer science, engineering, mathematics, decision sciences as well as more unconventional areas such as medicine and psychology. No filtering was done in subject areas.

It was also realized that it was essential to leave out any conference papers, book chapters, conference reviews, books, editorials, erratum, notes, retracted items, articles in press and data papers. This made the sample space contain only reviews and articles. This step was very important, as it made the bibliometric analysis for the topic much simpler. The other datatypes would not be useful for a bibliometric study as we need metadata for it. To be noted that 974 documents were removed in this step as they were deemed irrelevant. After the datatype filtering was done, we were left with 2766 results.

The authors noticed that the papers that were remaining were written in multiple different languages. This would have made it harder to do any analysis of the data. So, a decision was made apply a filter for language as an exclusion criterion which excludes languages, like Russian, Chinese, French and Persian. It was decided to include only English Language papers as English is accepted all over the world. It also indicates that the majority of the papers written on the selected topics are written in English Language, as 2578 papers were left after the filtering, resulting in the exclusion of 188 papers, mainly consisting of those in Chinese.

The authors added the next filter as the source type, in the database. The source types showed three categories: journal, trade journal and book series. The authors decided to exclude trade journals and book series and only include journals as the source type to be utilised in further study. This step revealed that most of the papers were journals as only 19 papers were excluded by the end of the process, leading to a final number of 2559 papers.

Scopus also had filters for country/territory, with China standing out as the country with the most papers published, followed by the likes of USA, Iran, India and Canada. No filtering was done in Country of Publishing as it is an irrelevant factor, as long as the data used is in English, it was taken for the bibliometric study, no matter the country.

Majority of the documents were open access (996) followed by those which were Green, Gold, Hybrid Gold and Bronze. Again, no filtering was done in this area as it does not matter if the paper is behind a paywall or openly accessible for everyone, as even the paid papers provide with meaningful metadata. There's no methodological reason to exclude paywalled records.

©IJRASET: All Rights are Reserved

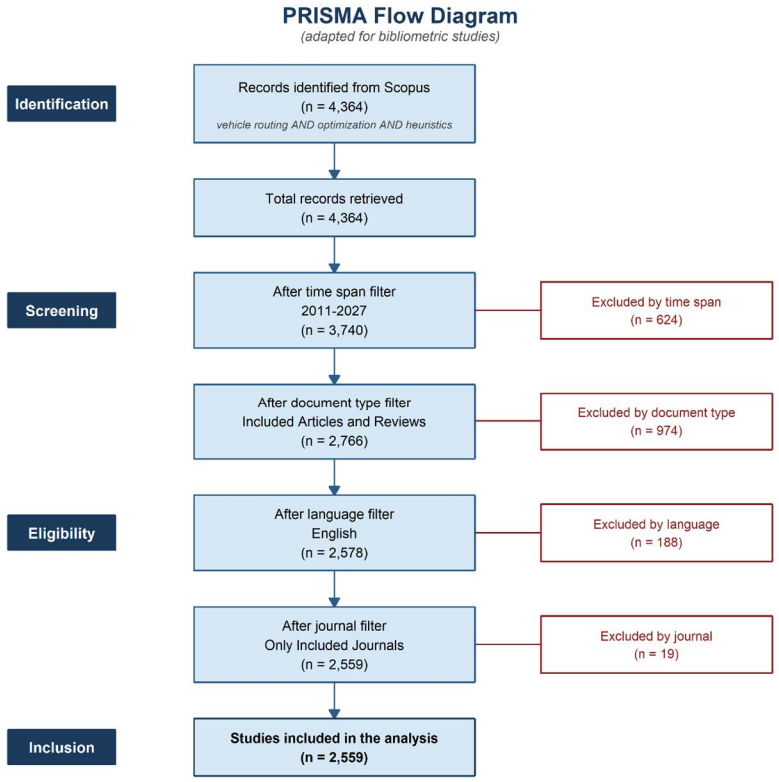


Figure 1. PRISMA flow diagram of the article selection process.

The authors have made a PRISMA diagram. They documented every inclusion and exclusion criterion used to narrow down the sample space. After that they used Biblioshiny (Aria & Cuccurullo, 2017), a package, in RStudio to build the PRISMA diagram. This diagram follows the guidelines set by the PRISMA 2020 Statement.

B. Bibliometric Analysis

Bibliometric reviews analyse a massive amount of published research using statistical tools. This helps to find trends and citations or co-citations for a theme. The analysis is done by year, country, authors, journal, method, theory and research problem (Paul & Criado, 2020). The authors have used analysis and review for the same purpose for the topic chosen in this paper. After doing a keyword search in Elsevier's Scopus database and adding the filters mentioned above the papers were exported in a RIS file from the database. Brought into Mendeley Desktop, then the papers were also exported from Scopus in a BibTex file (.bib) as the file format is utilised to further analyse the data.

The authors used RStudio to use a package inside the RStudio library called "bibliometrix". The authors used "Biblioshiny" from the bibliometrix package. Brought the BibTex file into Biblioshiny. Biblioshiny turned the data into an Excel file to format and analyse the details and insights. The data was also analysed using tools like VOSviewer, Tableau. Used thematic mapping and co-occurrence analysis to find or see a relationship, between the keywords that were searched.

IV. DATA ANALYSIS

This section presents six of the main outputs generated in Biblioshiny for the Vehicle Routing Problem dataset: 2,242 documents from 2011 to 2026, retrieved from Scopus. Each subsection pairs one chart or table with a short reading of what it shows and why it matters for the wider study. The figures are taken directly from each Biblioshiny screen, not estimated, and are cross-checked with a separate Excel trendline where the two overlap. Together, the six views give a reasonably complete picture of the shape of the field.

A. Summary of Important Information

The collection's vital-signs panel is the Main Information dashboard (Figure 2). The 2,242 documents come from 607 different sources over the period 2011 to 2026, with a strong annual growth rate of 14.51%, well above the typical 8 to 10% seen in bibliometrics as a whole. That growth rate alone points to an actively growing, still-young research area, not a mature or plateauing one.



Figure 2. Main information of the dataset, as reported by Biblioshiny.

Measures of authorship tell the same story. There are 4,931 unique authors in the corpus but only 105 who published a single-authored paper. There are only 114 single-authored documents in total, a small fraction of 2,242 total. This, combined with the average of 3.48 coauthors per paper, suggests a very collaborative field, which is standard for applied optimisation and engineering research, as opposed to a theory-heavy or solo-author discipline. The international co-authorship of 31.62% is well above the 30% threshold usually taken as evidence of global research networks, meaning that around one in three papers already cross borders.

The average document age of 5.02 years and average citations per document of 30.94 point to a weighting toward recent work but with meaningful citation history, consistent with a field that has been building momentum for over a decade without reaching saturation.

Two other figures from the table view are worth noting. The References count reads 0, a known Biblioshiny behaviour when cited-reference lists are not included in the Scopus export, so citation-network analyses such as co-citation or bibliographic coupling could not be run on this dataset. Keywords Plus also provides a total of 8,157 terms compared with 5,241 author keywords, giving two independent, cross-checkable sets of terms for thematic analysis.

B. Scientific Production per Year

Output grew from 41 articles in 2011 to 313 in 2026 with almost no sustained down year, providing a textbook picture of an accelerating research field (Figure 3). The curve is not a straight line: production is rather flat through 2012 to 2015 (41 to 77 articles), then takes a notable step up from 2018 on, with the steepest climb after 2021 (160, 184, 216, 225, 269, 313). The late acceleration is consistent with the broader adoption of machine learning and LLM-adjacent methods in operations research, and marks an inflection point for the field.

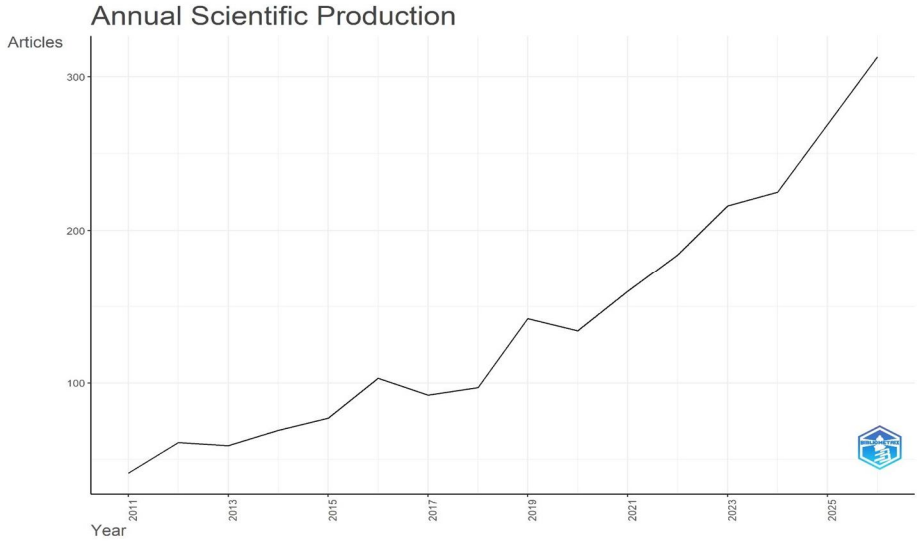


Figure 3. Annual scientific production, 2011 to 2026.

If a linear trendline is fitted to the counts year by year (fitted separately in Excel), the slope is about 16.4 additional articles per year, the R^2 is 0.92 and the CAGR is 14.51%, exactly matching the dashboard's growth-rate figure, so the two tools agree. An R^2 this close to 1 means the growth is unusually consistent rather than driven by one or two spike years, so extrapolating forward a couple of years is reasonably safe. The main caveat is the final 2026 bar, which is very likely an underestimate because the dataset was collected in September 2026 and indexing is still catching up, so it should not be read as a peak or plateau.

Worth noting separately is the fact that the raw export already contains 9 records dated 2027, most likely early-online or in-press articles that Scopus has pre-assigned a forward publication year. They are not included in the table of Annual Scientific Production in Biblioshiny, which lists 2,242 documents from 2011 to 2026 (16 years), and they do not skew the trendline above.

C. Average Number of Citations per Year

This chart shows the average number of citations per year for documents published in each year, and not the raw counts of citations, so it is a rate and not a total (Figure 4). The early years (2011 to 2013) are around 3 to 3.7, then the rate begins a steady upward climb to a first peak near 6.5 around 2015, dipping to about 5 in 2016, and then reaching its overall high point around 2017 to 2018 (roughly 6.7). From 2019 it stays in a fairly tight 5.8 to 6.3 band through 2023, which reads as the field's mature citation velocity. Once a paper in this area is a few years old, it tends to collect citations at a stable, predictable pace, which is a reassuring sign for anybody citing older work from this corpus.

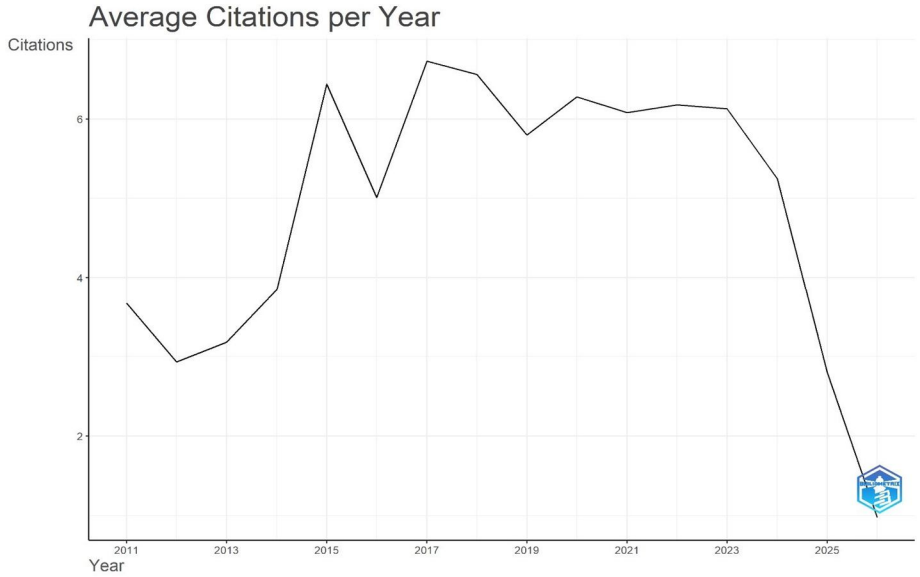


Figure 4. Average citations per year, by year of publication.

The sharp decline after 2023, to under 3 for 2025 and about 1 for 2026, is citation lag, not declining quality or relevance. A paper published in 2025 or 2026 simply hasn't had time to pick up citations, and the CitableYears column in the table view reflects this directly, with citable years shrinking from 16 for the 2011 cohort to single digits for the most recent cohorts. That is the typical shape of any citations-per-year chart, and should not be taken to imply that recent papers in this dataset are declining in impact. The fair interpretation is that only cohorts from 2023 and earlier have a reliable citations-per-year rate to report at this point.

It is worth differentiating MeanTCperArt and MeanTCperYear in the Biblioshiny table view. The former is total citations averaged per document regardless of age, and varies from roughly 44 to 77 across the cohorts. The chart above tracks the latter, citations normalised by the number of citable years for each cohort. The N column grows steadily from 41 documents in 2011 to 142 in 2019, showing that later cohorts are larger and, given enough time, as well cited as earlier cohorts.

D. Lotka's Law: Productivity of Authors

Lotka's Law describes the distribution of research productivity over authors: a small number publish many papers, and a large majority publish only one (Figure 5). The fitted curve here gives an estimated Beta of 2.3, as opposed to the classical Lotka value of Beta = 2, with a fitted C of 0.416 and a strong overall fit, $R^2 = 0.9182$, meaning the underlying power-law shape genuinely holds for this dataset even before comparing it to the textbook case.

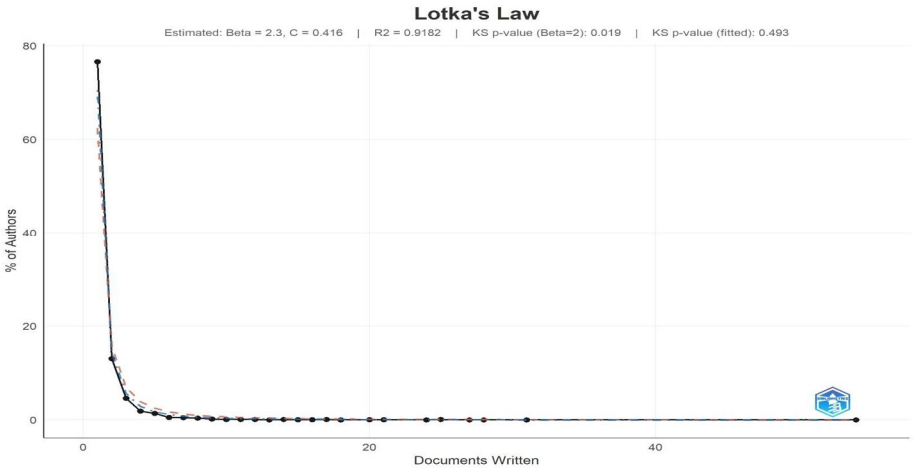


Figure 5. Lotka's Law: distribution of author productivity.

The estimated Beta of 2.3 is larger than 2, indicating a greater concentration of author productivity in this dataset than the classical Lotka pattern would predict, with a steeper drop-off in the proportion of highly prolific authors than in an average field. This is a pattern often seen in speciality or fast-moving niches, which is a reasonable fit for a still-emerging subfield like ML-driven vehicle routing. The two Kolmogorov-Smirnov tests tell slightly different stories, which are worth flagging separately. The test against the classical Beta = 2 model yields a p-value of 0.019, i.e. the data statistically deviates from the textbook Lotka curve, while the test against the fitted Beta = 2.3 model yields a p-value of 0.493, i.e. the fitted, generalised model is a good and statistically acceptable fit. To summarise, this field does not strictly obey Lotka's classical law but a well fitted generalised inverse power law with a higher concentration of authors than predicted by the textbook case.

The curve visually and directly confirms this concentration, with over 75% of authors at exactly one document, falling to single digits by three documents, with only a handful of authors beyond the 20-document mark. For a targeted citation search this means that a small set of prolific authors will emerge with a large share of the field's core contributors. The long tail of single-paper authors is better reached through keyword and venue searches instead.

E. Bradford's Law: Source Concentration

Bradford's Law divides journals into three zones whose total output is roughly equal and demonstrates the concentration of literature in a few core outlets (Figure 6). The Core zone has 14 sources, Zone 2 has 84 sources and Zone 3, the long tail, has 509 sources, a ratio of approximately 1 : 6 : 36 due to a Bradford multiplier of k = 6.03. This implies that each subsequent zone demands roughly 6 times as many journals to deliver the same volume of articles, a classic sign of high concentration around a small set of home venues for the field, and a trend that tends to persist across most engineering and operations-research literatures.

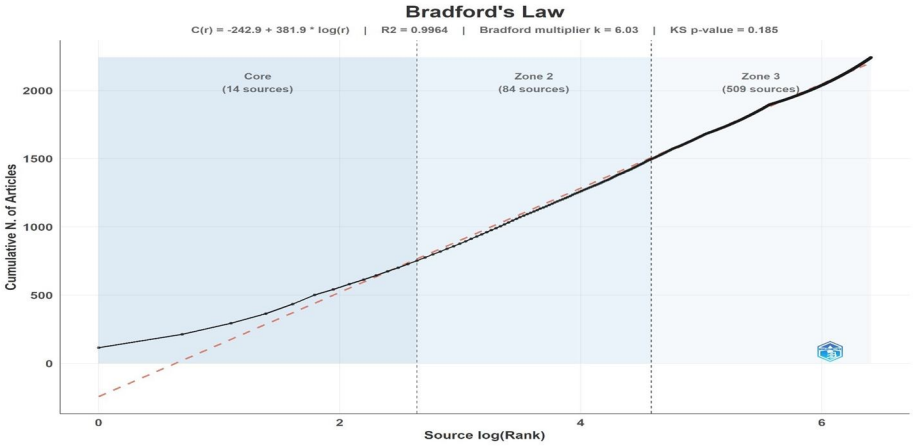


Figure 6. Bradford's Law: concentration of articles across sources.

The fit of the model is very good: $R^2 = 0.9964$ and the p-value of the KS goodness-of-fit test (0.185) is much higher than the usual 0.05 threshold, showing that the empirical source distribution is consistent with the theoretical Bradford curve and not deviating from it. In practice, this justifies cutting off a most-relevant-sources table at about 10 to 14 journals without materially under-representing the field. Those 14 core journals already carry a disproportionate share of the relevant literature, and the returns diminish rapidly after that point. The other 509 peripheral journals contribute one or two papers each, and are not worth pursuing for a literature review individually.

The plotted curve also shows the classic Bradford shape: an initial concave region through the core zone, settling into the near-linear region that the R^2 and trendline above are measuring. The empirical curve in Zone 2 and Zone 3 closely follows the fitted line, not just in the core. This is what gives the 0.9964 R^2 its practical meaning; the pattern of concentration is consistent across the entire source list, not just the top few journals.

F. Country Scientific Output

Research output is highly concentrated geographically. China is way out in front with 1,498 documents, more than three times the next highest country, followed by USA (442), India (281), Iran (233), Turkey (190), Canada (186), France (178), Brazil (164), UK (157) and Italy (151). These ten countries alone account for the vast majority of output among the 83 countries represented in the corpus, with a long tail of countries contributing only a handful of papers each.

This is also visually evident in the map view (Figure 7), where research activity is clustered around East Asia (dominated by China), North America and parts of Europe and South Asia. Much of Africa and Central Asia is only minimally shaded, indicating genuinely low publication volumes from those regions, not a display artefact. China's dominance is specifically a reflection of the overall scale of engineering and logistics research output in China, and the strong presence of vehicle-routing and supply-chain optimisation groups. These groups recur throughout the author-affiliation data pulled from the raw export.

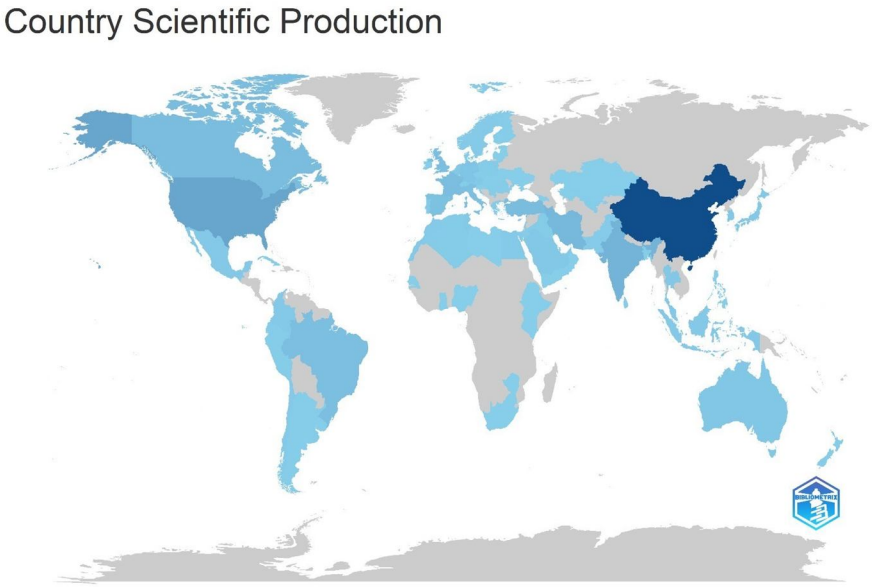


Figure 7. Country scientific production.

This concentration matters: any claim about global trends in this literature is in effect describing a small set of research hubs. Cross-country collaboration patterns, namely single versus multiple country publications by corresponding author's country, were not analysed here and would show how internationalised the field is beyond raw volume.

The gap between first and second place is worth stating plainly: China's 1,498 documents are more than 3 to 1 ahead of the USA's 442, and the numbers fall off quickly after the top ten, with Italy in tenth place at 151. Overall the data represent 83 countries, but the vast majority of these have only a handful of papers each.

V. FUTURE RESEARCH AGENDA: AN INPUT-PROCESS-OUTPUT ANALYSIS

A. Introduction to Framework

Keyword co-occurrence networks demonstrate the scope of a field under investigation, but say little about the efficacy of a study's components. To reveal this, the updated corpus of 304 records (all published in 2026) is projected onto an Input-Process-Output (IPO) frame. Inputs are the conditions defining a routing problem: fleet, demand, network and delivery setting. Processes are the solution methods: metaheuristics, exact methods, machine learning and hybrids. Outputs are the performance dimensions on which solutions are judged: cost, sustainability and service quality. Research advances through pairings, so a field can be prolific in each category while leaving specific input-process-output combinations unexplored. This matters because a field better instrumented to solve routing problems than to define what a good solution means will continue optimizing the wrong things. Because the exported keyword field is truncated at 50 characters, concepts are counted as the number of records mentioning them in the title or abstract.

B. Terminology

Table 1 lists the key terms used in this agenda.

Table 1. Key terms used in this agenda

Term	Definition used here
Vehicle routing problem (VRP)	Designing least-cost routes for a fleet serving dispersed customers.
Capacitated VRP (CVRP)	VRP in which each vehicle has a maximum load.
VRP with time windows	VRP requiring each customer to be served within a set interval.
Capacitated arc routing (CARP)	Routing where service occurs on road segments, such as street coverage, not at customer nodes.
Last-mile delivery	The final leg moving goods from a hub to the customer.
Quick commerce	Ten-to-twenty-minute delivery of groceries and essentials.
Dark store	Micro-warehouse closed to the public, used only to fulfil online orders.
Drone routing	Routing of drones, alone or with trucks, to serve customers.
Green / electric VRP	Routing that limits emissions or energy use, including battery fleets with charging constraints.
Metaheuristic	General search framework, such as a genetic algorithm, giving good but not provably optimal solutions.
Hybrid algorithm	Method combining exact mathematical programming with heuristic search.
Deep reinforcement learning	Neural policy trained by reward to construct routes quickly.
Multi-objective optimization	Simultaneous optimization of conflicting objectives, yielding trade-off solutions.

C. IPO Classification

Cost dominates outputs, mentioned in 150 of 307 output mentions. In each concept group, outputs attract fewer records than processes (on average 61 vs. 89), and service quality (23) and fairness (11) are thin. The asymmetry persists but has shifted: sustainability (88 records; 29%) has entered the mainstream, while service quality, equity and context-specific inputs lag. Inputs are fragmented across many small concept groups, suggesting that isolated context-specific settings are being studied rather systematically than compared to one another. Each question below therefore links an under-studied input or process to one of these thin outputs, so that the agenda targets combinations rather than isolated topics (Table 2).

Table 2. Input-Process-Output classification of corpus concepts

Category	Concept groups	Record mentions	Leading concepts (records)
Inputs	13	475	Time windows (66), drone routing (62), urban logistics (60)
Processes	5	446	Metaheuristics (148), exact methods (98), machine learning (86)
Outputs	5	307	Cost (150), sustainability (88), delay violation (35)

Mentions exceed 304 because a record can mention several concepts. Runtime (71) is treated as a process attribute.

D. Quick Commerce

The most notable limitation of this research is the absence of a certain node. Only one record contains information on quick commerce or dark stores, while seven entries describe on-demand grocery or food delivery services. Quick commerce is different in terms of delivery timeframes, micro-warehouse fulfilment, and demand volatility through apps, which are not always reflected in conventional studies. The closest substitutes were small-scale and peripheral, but no single article or chapter focuses on these aspects, which is why quick commerce appears to be ahead of its time and may require further observations to form a distinct cluster.

The following research questions can be considered for future study:

- How does the density of the dark-store network impact the reliability of deliveries within shortened timeframes through dynamic reoptimization?
- In what ways can the volatility of demand from mobile applications affect the utilization rates and idle time of couriers under reinforcement learning?
- How do ten to twenty-minute delivery windows impact the balance between cost and service efficiency under hybrid metaheuristic-exact methods?
- How do time-sensitive and temperature constraints of dark-store inventories impact the accuracy of deliveries under capacitated dynamic routing?

E. Autonomous and Drone Delivery

Drone routing has now overtaken the second-largest input theme (62 records, 20%), while sustainability links have increased (20 records), with 17 records addressing emergency or humanitarian settings, building on the previous finding. There remains a gap on the service-side, with only three drone records addressing service quality, and one concerning health logistics, despite drones being proposed for medical logistics frequently. Drone routing is therefore predominantly considered as a novel optimization problem to be tested against service and public-health outcomes, rather than as a lever to be trialed.

Future research could include:

- Could constraints inherent in drone batteries and payloads (under multi-objective optimization) impact on service achievement at the last-mile?
- How do cold-chain requirements in drone routing impact on delivery reliability and cost for medical supplies?
- How do urban airspace and no-fly-zone constraints (considering dynamic routing under constraints) impact on delivery-time reliability?

F. Machine Learning Methods

The machine learning area is now a major process family (86 records, 28%), with 53 on reinforcement learning, indicating a shift toward fast learned solvers, including for capacitated arc routing. Yet few of these studies connect to outcomes beyond cost: only six address service quality, 19 address sustainability, and none address emerging-market settings. Whether learned policies trained on benchmark instances are reliable under noisy, real operating data is largely untested, and the new corpus offers little head-to-head evidence against hybrid metaheuristics.

Future work could include:

- How do reinforcement-learning routing policies trained on benchmark instances perform on real, noisy demand and travel-time data, compared with hybrid metaheuristics, in terms of service reliability?
- How does embedding sustainability and service-level targets in the reward function of deep reinforcement learning routing change the cost-emissions-service trade-off?

G. Sustainability Trade-offs

The problem of transportation sustainability is not isolated since 33 of the 55 electric-vehicle records address it. However, it has a weak association with service-related outcomes. Multi-objective studies are combined with sustainability in 18 records and service quality in 5 while electric-vehicle studies with service quality in only 2. The speed vs. emissions trade-off is mostly unquantified. Hence, the opportunity cost in terms of speed and/or service quality of a more sustainable routing is rarely explicitly stated to the decision-maker.

Future research could focus on:

- How do emissions-constrained multi-objective metaheuristics affect the trade-off between delivery speed and carbon output in urban last-mile delivery?
- How do the availability of charging stations and the time that it takes to charge electric vehicles affect customer service levels and total cost under electric-fleet routing?

H. Geographical and Sectoral Scope

Healthcare logistics was embedded in 15 records and humanitarian or emergency logistics in 31 records. Explicit emerging-market contexts were mentioned in only five records, none of which employed machine learning, as the heterogeneous and densely populated traffic conditions are fundamentally different from the orderly lanes and reliable addresses usually postulated in traditional vehicle routing problems, let alone the use of standardized commercial fleets. Healthcare was combined with sustainability in two records and with service quality in one.

The following research questions could be pursued:

- In what way do the challenges posed by mixed and non-lane disciplines traffic impact the feasibility and delivery time performance of adaptive metaheuristic routing in emerging-market contexts?
- In what way do the limitations of informal address and geocoding systems affect the ability to achieve successful delivery outcomes in emerging-market e-commerce through machine learning-assisted routing?
- In what way do two-wheeler and informal fleets impact the cost-efficiency of last-mile delivery when routings are optimized with respect to both capacity and mode of transportation?
- In what way can routing be leveraged to simultaneously optimize sustainability and service performance in healthcare logistics?

I. Standardizing Terminology

The terms last-mile delivery, e-commerce delivery, food delivery and city logistics are separate and only loosely connected labels, despite substantial overlap, and the new corpus adds variants such as on-demand delivery. The fragmentation allows co-occurrence networks to underestimate the centrality of a concept and systematic searches to miss work indexed elsewhere. This study proposes 'vehicle routing problem' and 'last-mile delivery' as anchor terms for the methodology and application layers, with the specific terms in Table 1 as descriptors, not substitutes. This would improve the discoverability and the reliability of future bibliometric mapping.

J. Priorities for Future Work

Three priorities emerge from these fifteen questions. First, logistics solutions specific to such environments as quick commerce and emerging-market settings are needed, given their limited attention in the literature. Second, a more thorough analysis of the trade-offs between learning-based and hybrid approaches is required, focused on service quality and sustainability rather than time and cost. Third, studies concerning drone and electric-fleet logistics should explicitly consider service performance. If pursued, these priorities will shift the field's focus from defining a successful logistics solution (e.g., the shortest route) to determining a properly functioning one (e.g., an affordable route with acceptable sustainability).

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

VI. CONCLUSION

This paper examined the Vehicle Routing Problem literature through a bibliometric lens, using 2,559 journal articles and reviews retrieved from Scopus and analysed with Biblioshiny. It set out to describe how the field has grown, who produces it and where, and which themes it covers. It also tested a specific concern, namely that quick commerce and hyperlocal delivery, now a visible part of urban commerce, receive little attention from routing researchers. The evidence supports that concern. The field is large, productive and growing quickly, but its attention sits with established problem families and with newer technical themes, and the delivery model that shapes daily life in many cities is barely present in it.

A. Findings and Results

VRP research has grown steadily. Annual output rose from 41 documents in 2011 to 313 in 2026, an annual growth rate of 14.51%, well above the 8 to 10% typical of bibliometric datasets. A linear trend fitted to the yearly counts adds about 16 documents a year with an R^2 of 0.92, and the climb is steepest after 2021. Because the data were collected in September 2026, the 2026 figure should be read as incomplete. Average citations per document reach 30.94, and the citations-per-year rate settles at a stable level for cohorts up to 2023, while later cohorts show lower rates because they have had less time to be cited.

The field is collaborative, with 3.48 authors per document and international co-authorship of 31.62%. Authorship follows a generalised Lotka pattern ($\beta = 2.3$, $R^2 = 0.918$), in which a small group of prolific authors sits above a long tail of single-paper contributors. Sources follow Bradford's Law closely ($R^2 = 0.9964$), with 14 core journals, 84 in the second zone and 509 in the third. Geographically, China contributed 1,498 documents, more than three times the 442 from the USA, followed by India (281), Iran (233) and Turkey (190), so claims about global trends in the literature describe a small set of research hubs.

The keyword analysis gives the central result. Core VRP methodology terms were recorded 104 times, against only 4 mentions of last-mile delivery. Quick commerce and hyperlocal delivery barely register at all. The literature review points the same way, since the recent studies we reviewed centre on metaheuristics, hybrid methods, drones, electric vehicles, emissions and learning-based solvers (Du et al., 2026; Kong et al., 2026; Xiao et al., 2026). Research effort has followed methodological and technological opportunity, while the delivery model reshaping city logistics has received little direct attention.

B. Theoretical Implications

For theory, the study offers a quantified map of a field that is usually described through narrative reviews. The close fit of Bradford's Law and of a generalised Lotka distribution shows that VRP research follows the regular concentration patterns of an established field even though its output is still rising fast. The keyword results suggest something specific to routing: new variants enter the literature when a technology or constraint becomes tractable, as with drones, electric vehicles and learned solvers, and they enter more slowly when they originate in a business model. This indicates that theory building in the VRP may be driven more by methodological supply than by commercial demand, an interpretation that further studies should test. For quick commerce, the gap points to a need for formulations that combine dynamic order arrival, very short delivery promises, batching and rider assignment, which a static capacitated model with one extra constraint cannot capture. The study also shows how bibliometric methods (Paul & Criado, 2020) can test a claim about research relevance with counts.

C. Managerial Implications

Logistics managers can draw on a mature toolkit. Hybrid heuristics and matheuristics that use mathematical programming to improve heuristic solutions (Vadseth et al., 2026) are well documented and can serve as a baseline for firms that still plan routes by hand or by spreadsheet. Quick commerce operators should not expect off-the-shelf methods built for their setting, since the literature gives them little. They are better placed to test established time-window and dynamic routing methods on their own order data and to treat the results as a starting point. Generative artificial intelligence may lower the technical barrier for smaller firms, but reproducibility and constraint handling need checking before it is trusted in operations (Ramingwong & Jintana, 2026). Managers who want to follow the research can concentrate on the 14 core journals identified in the Bradford analysis, which carry a large share of the field's output.

D. Policy Implications

Cities that host fast-delivery fleets face congestion, kerbside parking pressure and emissions, yet the research that could inform rules on dark store location, rider numbers or delivery windows is thin.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

Public research funders could treat dense, short-window urban delivery as a priority area and support work that links routing models to emission and congestion outcomes, building on emission-aware formulations such as Khalid et al. (2026). Regulators could also encourage operators to share anonymised delivery data with universities, since the shortage of public data is a likely barrier to research on these services.

E. Social Implications

Better routing shortens the distance driven for each delivery, which lowers fuel use and local air pollution. Route-balancing models (Xiao et al., 2026) show that workload can be shared more evenly among drivers, a concern that applies strongly to riders working under tight time promises. Research output is also concentrated in a few countries, led by China, so the findings that shape delivery practice come from a narrow set of research settings. Wider participation, including from the cities where quick commerce is growing fastest, would make future work more representative of how these services affect workers and neighbourhoods.

F. Limitations

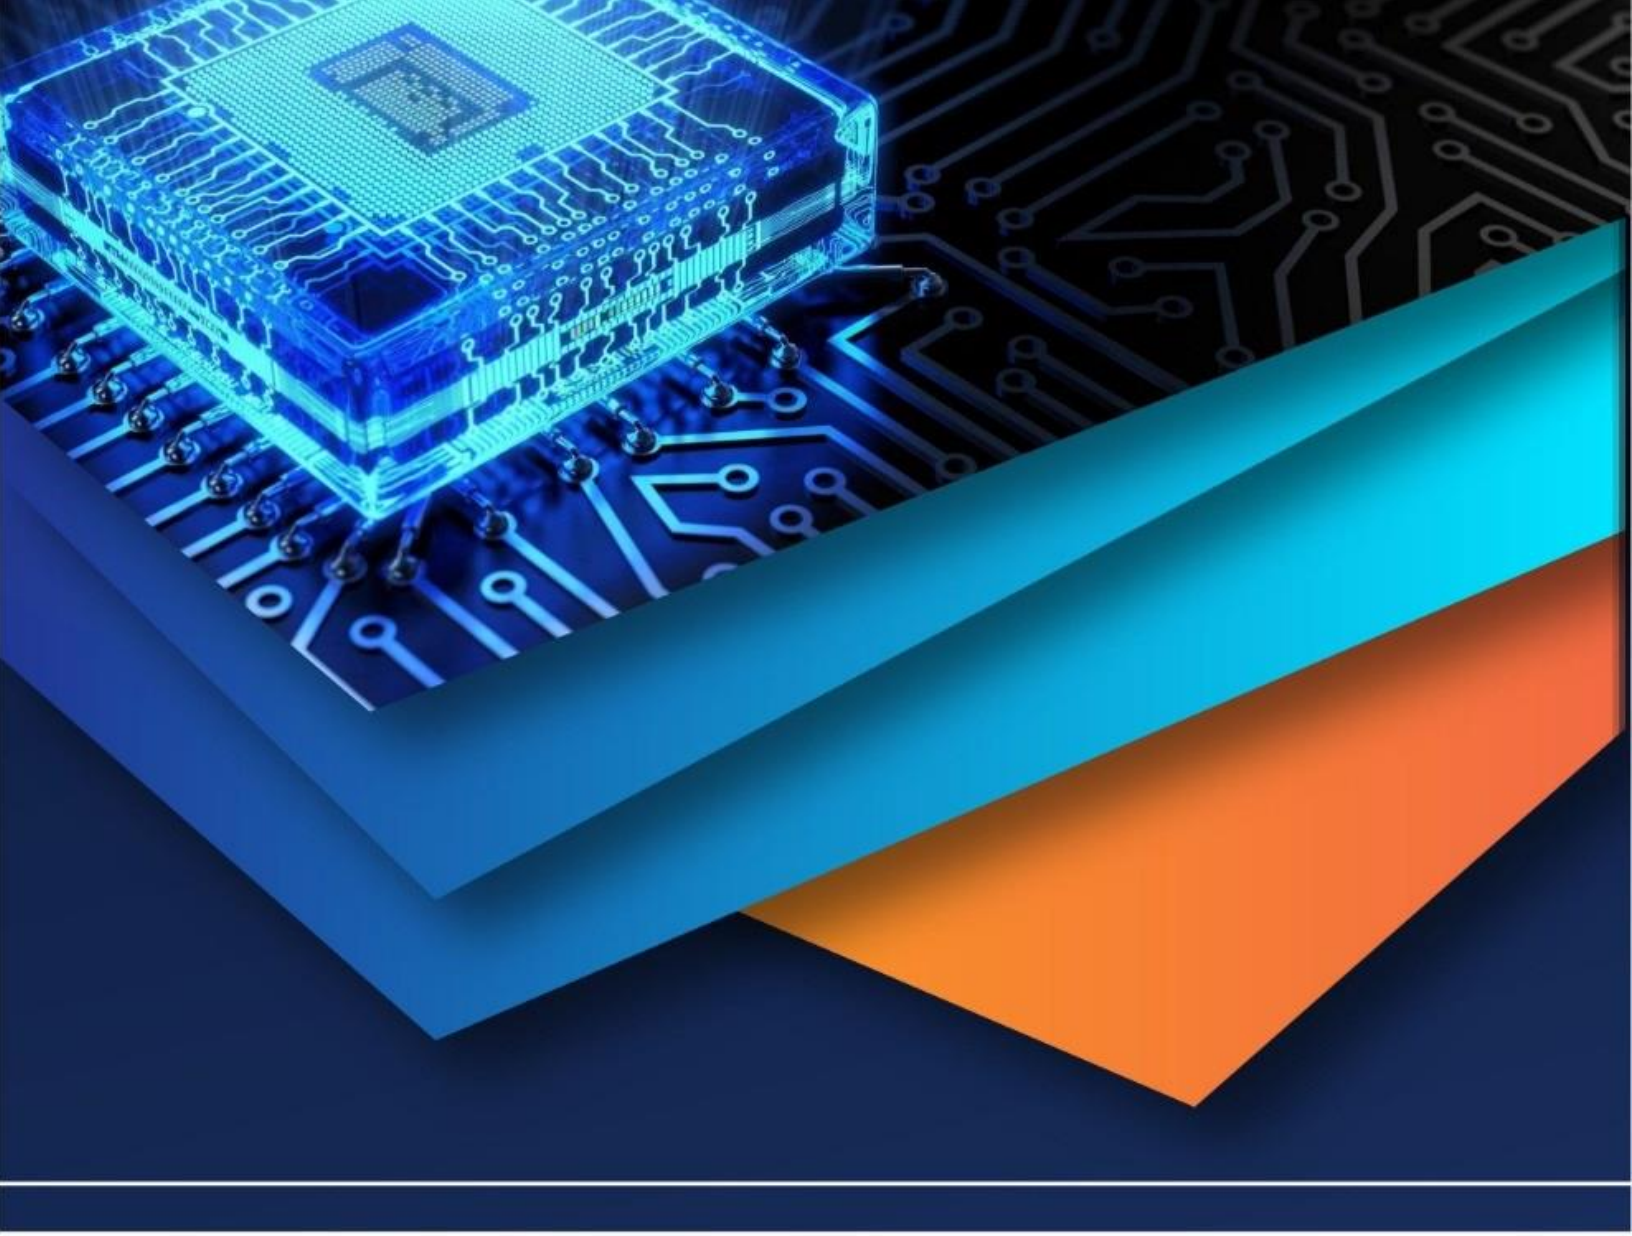
Several limits apply. The search used the terms vehicle routing, optimization and heuristics, so the dataset captures the VRP literature as defined by that vocabulary. Papers on quick commerce or last-mile delivery that use other terms were never retrieved, and a broader search string might show more coverage. The data come from Scopus alone and cover English-language journal articles and reviews, which excludes conference papers, where much computing research appears. Keyword counts do not read full texts, so a paper could address quick commerce without using the term. The 2026 data are incomplete, and bibliometric indicators describe volume and visibility, not the quality of individual studies. The Biblioshiny export also lacked cited-reference data, which ruled out co-citation and bibliographic coupling analysis. These limits do not weaken the main observation but they define how far it can be generalised.

REFERENCES

[1] Aria, M., & Cuccurullo, C. (2017). bibliometrix: An R-tool for comprehensive science mapping analysis. Journal of Informetrics, 11(4), 959–975.
[2] Bogyrbayeva, A., Meraliyev, M., Mustakhov, T., & Dauletbayev, B. (2024). Machine learning to solve vehicle routing problems: A survey. IEEE Transactions on Intelligent Transportation Systems, 25(6), 4754–4772.
[3] Chen, J., Wang, X., & Lim, M. K. (2017). Bibliometric analysis of the emerging trends in green logistics. Journal of Cleaner Production.
[4] Dantzig, G. B., & Ramser, J. H. (1959). The truck dispatching problem. Management Science, 6(1), 80–91.
[5] Du, H., Tanimizu, Y., & Watanabe, R. (2026). A hybrid constructive genetic algorithm (HCGA) for the vehicle routing problem with time windows in logistics optimization. IEEE Access, 14, 40532–40549. https://doi.org/10.1109/ACCESS.2026.3672634
[6] Duan, H., Li, X., Feng, Y., Zhang, G., Lu, Q., & Meng, X. (2026). Multi-objective hyper heuristic memetic algorithm for time-dependent vehicle routing problem with drone considering simultaneous pickup and delivery. Expert Systems with Applications, 297. https://doi.org/10.1016/j.eswa.2025.129196
[7] Erdem, E., Erkayman, B., & Aydin, T. (2026). Advances in AI-based solutions for the multi-depot vehicle routing problem: A review of recent trends and future directions. Archives of Computational Methods in Engineering, 33(2), 1997–2035. https://doi.org/10.1007/s11831-025-10356-y
[8] Kahiya, E. T. (2018). Five decades of research on export barriers: Review and future directions. International Business Review, 27(6), 1172–1188.
[9] Khalid, Q. S., Maqsood, S., Mumtaz, J., & Qureshi, S. M. (2026). An emission-capacitated vehicle routing model for sustainable urban waste collection using hybrid guided local search. Scientific Reports, 16(1). https://doi.org/10.1038/s41598-026-38829-5
[10] Kong, J., Wei, L., & Jiang, X. (2026). A literature review of vehicle and drone delivery routing problems in different synchronization level scenarios. Drones, 10(3). https://doi.org/10.3390/drones10030206
[11] Laporte, G. (2009). Fifty years of vehicle routing. Transportation Science, 43(4), 408–416.
[12] Lenstra, J. K., & Rinnooy Kan, A. H. G. (1981). Complexity of vehicle routing and scheduling problems. Networks, 11(2), 221–227.
[13] Moher, D., Liberati, A., Tetzlaff, J., & Altman, D. G. (2010). Preferred reporting items for systematic reviews and meta-analyses: The PRISMA statement. International Journal of Surgery, 8(5), 336–341.
[14] Paul, J., & Criado, A. R. (2020). The art of writing literature review: What do we know and what do we need to know? International Business Review, 29(4), 101717.
[15] Rachmawati, D., Lohil, T., & Tarigan, J. T. (2026). Enhanced ant colony optimization for capacitated vehicle routing problem with time windows in franchise distribution. International Journal of Advanced Computer Science and Applications, 17(2), 1045–1058. https://doi.org/10.14569/IJACSA.2026.01702103
[16] Ramingwong, S., & Jintana, J. (2026). Investigating the potential and performance of generative AI for a vehicle routing problem. Logistics, 10(6). https://doi.org/10.3390/logistics10060120
[17] Randhawa, K., Wilden, R., & Hohberger, J. (2016). A bibliometric review of open innovation: Setting a research agenda. Journal of Product Innovation Management, 33(6), 750–772.
[18] Ren, Y. S., Ma, C. Q., Chen, X. Q., Lu, Z. N., & Wang, Y. R. (2020). Sustainable finance and blockchain: A systematic review and research agenda. Research in International Business and Finance.
[19] Toth, P., & Vigo, D. (Eds.). (2014). Vehicle routing: Problems, methods, and applications (2nd ed.). SIAM.
[20] Vadseth, S. T., Andersson, H., Cordeau, J.-F., & Stålhane, M. (2026). Mixed integer programs to improve solutions of vehicle routing problems with intra-route constraints. Journal of Heuristics, 32(2). https://doi.org/10.1007/s10732-026-09583-8

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

[21] Wang, C., Zhang, R., & Wang, H. (2026). Heuristic-augmented attentions for the electric vehicle routing problem with time windows. IEEE Transactions on Vehicular Technology, 75(7), 12707–12719. https://doi.org/10.1109/TVT.2026.3665950
[22] Xiao, J., Kong, D., Cao, Z., & Zhao, J. (2026). Deep reinforcement learning for the vehicle routing problem with route balancing. Transportation Research Part E: Logistics and Transportation Review, 208. https://doi.org/10.1016/j.tre.2025.104632
[23] Yang, Z., Li, J., Wang, Y., & Zhang, H. (2026). Solving the capacitated vehicle routing problem via a collaborative adaptive fusion transformer. Applied Soft Computing, 202. https://doi.org/10.1016/j.asoc.2026.115866
[24] Zabraoui, O., Chafi, A., & Alami, S. K. (2026). AI-driven optimization of vehicle routing problems in supply chain: integrating graph neural networks and reinforcement learning. Flexible Services and Manufacturing Journal. https://doi.org/10.1007/s10696-025-09653-2
[25] Zhou, S., Liu, X., Chen, J., Wu, L., Xu, Z., & Zhao, M. (2026). Truck-drone routing problem under time-dependent traffic congestion: A novel variable neighborhood search metaheuristic. Transportation Research Part E: Logistics and Transportation Review, 216. https://doi.org/10.1016/j.tre.2026.105167



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)