Service Operations Research Seminar FSS 2026 (OPM 781)

"Current Topics in Service Operations Management Research"

General Information:

1. The goal of this seminar is to introduce participants to conducting applied scientific research in the field of (service) operations management. Also, the seminar aims at practicing presentation skills, such as speaking with clarity, confidence and connection.

2. The master thesis prepares students for writing their M.Sc./Diploma Thesis primarily at the Chair of Service Operations Management, but OPM781 also qualifies you formally for writing a master thesis at any other chair in the Operations Area.

3. The offered topics are presented below and designed to be explored by a single student based on the fundamental literature. Each participant will present his/her findings in a written report (about 20 pages) as well as in an in-class presentation (~20 min), followed by a discussion (~10 min).

4. The **application procedure** for this seminar is combined with those for the seminars of the Chair of Production Management (OPM 761), the Chair of Logistics (OPM 701) and the Chair of Procurement (OPM 791). Students can apply for topics from all chairs by joining the ILIAS application group and completing the online form provided there. Topics labeled with "L" refer to the Chair of Logistics (OPM 701), topics labeled with "P" refer to the Chair of Production Management (OPM 761), topics labeled with "B" refer to the Chair of Procurement (OPM 791) and topics **labeled with "S"** refer to the **Chair of Service Operations Management (OPM 781)**. The **assignment of topics** to students will be preference-based through ILIAS. To better match topic and student background, applicants for OPM 781 may in addition send a CV and official grades overview by post to the chair or by e-mail to sekretariat.som@uni-mannheim.de with subject "OPM 781 Seminar Application".[1]

5. The **application period** starts on **Nov. 8th** and ends on **Nov. 22nd, 2025**.

6. **Admission** to the seminar is **binding** and will be confirmed by E-mail by **Dec. 1st, 2025** at latest.

[1] Data protection: Please note that a breach of confidentiality and the unauthorized access by third parties cannot be excluded when transmitting an unencrypted email. Note on data protection: The submitted documents will be returned only if an envelope with sufficient postage is included. Otherwise they will be destroyed after the application process according to the requirements of the data protection law. Electronic applications will be deleted accordingly.

7. A **kick-off meeting** for all participants will be held on Wednesday, **Dec 3rd at 10:15 till 11:45 in SO318**. During this meeting, general guidelines for conducting scientific work will be discussed.

8. The latest **submission** date for the written report incl. appendices is **Apr. 20th (2026),** For submission, please …

    a. **Upload your report** (Word- / Latex-document and PDF) via Task "Upload of final Thesis & Calculations/Software Output" in the OPM781 ILIAS group. If you have multiple files (e.g. a pdf and some Excel analysis), please upload all in a single zip file.

    b. **Submit a hard copy** at our secretary's office (Mon-Thu before noon) or at your thesis supervisor. Please make an appointment for submitting the hard copy.

9. The **final presentations** of the seminar participants will be held by default in the **regular presentation session on May 8th (2026),** in room SO 318. A **fast-track presentation track** may be offered to students who desire to start with their master thesis early based on their request. Attendance is mandatory for all presentations on your own presentation date.

10. In addition, we will offer an **optional mock-up presentation** session one week before the regular final presentations, i.e., on **Apr. 29th** (in SO 322). Here, participants can practice their final presentation and get tips on structure, content and presentation style risk-free without being graded – if they want to. In the kick-off meeting, we will provide some guidance and resources on how to train your presentation skills upfront, such as speaking with clarity, confidence and connection.

11. Please **upload your final presentation slides** (ppt and PDF) on Task "Upload of Final Presentation" in the ILIAS group **one day before the presentation**, latest by 18:00 pm.

12. The final grade for the seminar is composed of the following components: Written report (60%), presentation (30%), contribution to discussion of your own topic and of potentially other topics presented on the same date (10%).

13. For questions concerning the seminar contact us by email at sekretariat.som@uni-mannheim.de

Seminar topics

Please note:

The amount of recommended literature does NOT indicate more or less workload! Every thesis will contain some literature review, and more recommendations maybe helpful for this. Also, your supervisor may have more recommendations for you, in particular, if the initial list of recommended references for a topic is short.

Table of Contents

Topic S01: Regulating Environmental Emissions

Environmental regulators worldwide have implemented a broad range of policies to mitigate the harmful externalities associated with industrial emissions, including acid rain from sulfur oxides, particulate pollution, and the greenhouse effect driven by carbon dioxide. Among these policy instruments, tradable emission permits have emerged as one of the most widely adopted and economically grounded regulatory tools. Under such systems, regulated entities are required to acquire and surrender permits that correspond to their verified emissions.

Programs such as the European Union Emissions Trading System (EU ETS), the U.S. Acid Rain Program (ARP), the NOx Budget Trading Program (NBP), and the Hunter River Salinity Trading Scheme exemplify the large-scale and durable implementation of permit-based regulation. These systems illustrate that market-based control can produce meaningful abatement incentives while preserving flexibility for firms.

Xu et al. (2025) examine emission-permit designs with varying temporal flexibility across multiple compliance periods and discuss the results of extreme regimes as well as hybrids (e.g., transfer caps/discounts, permit-tax mixes).

The primary objectives of this thesis are to:

- provide an overview of regulatory instruments governing emissions,
- present selected components of the model of Xu et al. (2025), discuss the results and relate them to the literature,
- identify future research opportunities.

Selected Literature Recommendations:

Fu, Xingyu, Ying-Ju Chen, Guillermo Gallego, Pin Gao, Mengqian Lu (2025). Designing Emission Permits Regulation for Multiple Compliance Periods. Manufacturing & Service Operations Management 27(5):1587-1603.

Gronwald, M., & Hintermann, B. (2016). Emissions trading as a policy instrument: evaluation and prospects. Cambridge, Massachusetts.

Yuxuan Cao, Wanyu Ren, Li Yue (2024). Environmental regulation and carbon emissions: New mechanisms in game theory. Cities. 149 (104945).

Topic S02: Carbon-Aware Computing at Google: Optimizing Datacenter Emissions

Data centers are critical infrastructure for the digital economy, and their rapid expansion has made greenhouse-gas emissions a material concern. Emissions arise in operations, electricity for computation and cooling, and in embodied sources from construction and hardware manufacturing. Cloud growth, AI loads, and high availability intensify both.

Decarbonization in operations therefore relies on a combination of energy-efficiency gains across servers, cooling, and software; carbon-aware computing aligned with low-emissions supply; 24/7 and location-matched procurement of carbon-free electricity.

Within this landscape, Google is a pivotal actor. Their work on carbon-aware computing is especially consequential. The approach shifts flexible workloads across time and location to coincide with lower grid-carbon intensity, using short-term forecasts and locational signals to route or defer jobs while respecting performance and data-locality constraints.

The primary objectives of this thesis are to:

- motivate the importance of emissions of data centers,
- review the literature on approaches aiming to reduce emissions from data centers,
- focus on the approach of Google and discuss its design and the results,
- identify future research opportunities.

Selected Literature Recommendations:

Arshi, Praneet and Joel Miller (2025). Our approach to carbon-aware data centers: Central data center fleet management. Google.

Panda, S. K., Srivastav, A., Singh, A., and Li, K.-C. (2025). A renewable energy-based virtual machine placement algorithm for managing energy and carbon in geographically distributed datacenters. Cluster Computing, 28(10), 666.

Radovanovic, Ana, Bokan Chen, Saurav Talukdar, Binz Roy, Alexandre Duarte and Mahya Shahbazi (2021). Power Modeling for Effective Datacenter Planning and Compute Management. IEEE Transactions on Smart Grid.

Radovanovic, Ana, Ross Koningstein, Ian Schneider, Bokan Chen, Alexandre Duarte, Binz Roy, Diyue Xiao, Maya Haridasan, Patrick Hung, Nick Care, Saurav Talukdar, Eric Mullen, Kendal Smith, Mariellen Cottman and Walfredo Cirne (2022). Carbon-Aware Computing for Data-centers. IEEE Transactions on Power Systems.

Topic S03: Data-Driven Revenue Management – Recent Advances and Selected Models

Revenue management (RM) has become a vital strategy for firms across various industries, aiming to predict consumer behavior and optimize product availability and prices to maximize revenue or profit. With the advent of data-driven decision-making and the emergence of new industries such as internet advertising, cloud computing, and e-commerce, there has been a growing demand for innovative RM models, frameworks, algorithms, and policies. In response to this demand, researchers have been focusing on advancing RM methodologies, particularly in accommodating RM to emerging industries and leveraging data-driven approaches. This includes the development of advanced personalized pricing and assortment models, exploration of RM strategies in industries with reusable resources, and the proposal of network RM and end-to-end decision-making techniques, etc.

The primary objectives of this thesis are to:

- review and classify existing data-driven approaches in RM and specify the recent advances, published in leading academic journals, such as INFORMS journals, European Journal of Operational Research, Production and Operations Management, Journal of Operations Management, Journal of Revenue and Pricing Management,
- delve into several model-free approaches, which are in a subset of data-driven approaches and do not assume any specific distribution of consumer valuation, with the focus on their model, advantages, limitations, and applications,
- identify future research opportunities for model-free approach in revenue management.

Selected Literature Recommendations:

Chen, N., Cire, A. A., Hu, M., & Lagzi, S. (2023). Model-free assortment pricing with transaction data. Management Science, 69(10), 5830-5847.

Chen, N., & Hu, M. (2023). Frontiers in Service Science: Data-Driven Revenue Management: The Interplay of Data, Model, and Decisions. Service Science, 15(2), 79-91.

Topic S04: Applications of Machine Learning in Pricing and Assortment Optimization

In today's highly competitive and data-driven marketplace, firms face increasing pressure to make smart and timely decisions about what products to offer and at what prices. Two of the most critical levers in driving profitability, customer satisfaction, and operational efficiency are pricing and assortment optimization. Assortment optimization focuses on selecting the right mix of products to offer, balancing customer preferences, inventory constraints, and business goals. Traditional demand models for pricing and assortment optimization often rely on strong assumptions and offer limited flexibility when modeling real-world consumer behavior. However, real-world decision-making involves large assortments with many interrelated products, heterogeneous and dynamic customer preferences, and frequent changes in market conditions, inventory, and competition, etc. Machine learning demand models offer powerful alternatives by capturing nonlinear and complex patterns in customer choice, handling high-dimensional data, enabling real-time and adaptive pricing and assortment decisions, and improving accuracy in forecasting and personalization, etc.

The primary objectives of this thesis are to:

- review the state-of-the-art applications of machine learning in pricing and assortment optimization as published in leading academic journals, such as INFORMS journals, European Journal of Operational Research, Production and Operations Management, Journal of Operations Management, Journal of Revenue and Pricing Management,
- analyze selected applications in detail, focusing on the data and contexts, models, methodologies, key findings, managerial implications, and limitations,
- identify future research opportunities for machine learning in pricing and assortment optimization.

Selected Literature Recommendations:

Li, T., Wang, C., Wang, Y., Tang, S., & Chen, N. (2024). Deep reinforcement learning for online assortment customization: A data-driven approach. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4870298.

Topic S05: Bridging Prediction and Decision: A Review of Integrated Learning and Optimization Approaches

Many operational decisions—from inventory control to pricing and logistics—depend on uncertain parameters that must be estimated from data. Traditionally, such problems are addressed in a two-step framework: first, parameters are predicted using machine learning or statistical models, and second, an optimization model is solved using these predictions. This predict-then-optimize approach is intuitive and modular but may lead to suboptimal decisions when the predictive model does not account for how its outputs affect the optimization task. In recent years, a new stream of research has emerged to integrate prediction and optimization more tightly, aiming to improve decision quality rather than just prediction accuracy. This seminar provides an overview of the main approaches in this area:

- Predict-then-Optimize with Decision-Focused Learning:
  Methods like Smart Predict-then-Optimize (SPO) introduce loss functions that reflect downstream decision performance rather than traditional prediction error.
- One-Step Integrated Models:
  Recent work proposes to embed learning directly into the optimization problem, removing the boundary between prediction and decision-making (e.g., Ban & Rudin, 2019). These models often use machine learning (including deep learning) to learn directly from data what decisions to make.

The primary objectives of this thesis are to:

- review recent methods integrating prediction and optimization, focusing on their conceptual differences, solution quality, and computational complexity,
- explain the Smart Predict-then-Optimize (SPO) framework, highlighting the decision-focused loss function and its improvements over traditional approaches,
- compare performance trade-offs between two-step and integrated approaches, especially under limited data or model misspecification,
- evaluate the challenges of solving these models, considering non-convexity, and data availability and so on.

Selected Literature Recommendations:

Ban, G.-Y., & Rudin, C. (2019). The big data newsvendor: Practical insights from machine learning. Operations Research.

Bertsimas, D., & Kallus, N. (2020). From predictive to prescriptive analytics. Management Science.

Elmachtoub, A. N., & Grigas, P. (2022). Smart “Predict, then Optimize”. Management Science.

Liu, M. (2024). Decision-focused learning and its applications in operations management (Doctoral dissertation, UC Berkeley).

Qi, M., & Shen, Z.-J. (2022). Integrating prediction/estimation and optimization. INFORMS Tutorials.

Topic S06: Evaluating the Performance of MNL and MMNL Choice Models in Assortment Optimization and Product Line Selection

Assortment optimization and product line selection are crucial strategic decisions for firms aiming to align product offerings with consumer preferences. These decisions are typically modeled using discrete choice frameworks such as the Multinomial Logit (MNL) and Mixed Multinomial Logit (MMNL) models, each capturing customer behavior with varying degrees of complexity and realism.

Under the MNL model, optimal assortments can often be computed efficiently using convex optimization techniques. A notable example is provided by Chen and Hausman (2000), who examined a joint pricing and assortment selection problem. Their formulation as a nonlinear binary integer program could be relaxed without loss of optimality due to the problem's favorable structural properties, allowing the relaxed solution to remain integral and globally optimal. However, the MNL model assumes homogeneous consumer preferences, which may not accurately reflect real-world behavior. The MMNL model addresses this by incorporating preference heterogeneity, most probably offering a better empirical fit. Despite its advantages, the MMNL-based assortment problem is computationally intractable and can only be solved exactly for small problem sizes using mixed-integer linear programming (MILP) formulations. This seminar examine how much solution quality is lost when an MNL model is used as an approximation for an MMNL model?

The primary objectives of this thesis are to:

- explore recent work on assortment optimization and product line selection under MNL and MMNL models,
- utilize parameter estimates from a relevant empirical study (e.g., Keane & Wasi, 2013) to formulate and solve the problem under both MNL and MMNL frameworks,
- quantify the performance bias introduced by using a single-segment MNL model instead of MMNL.

Selected Literature Recommendations:

Chen, K. D., & Hausman, W. H. (2000). Mathematical properties of the optimal product line selection problem using choice-based conjoint analysis. Management Science, 46(2), 327–332.

Keane, M., & Wasi, N. (2013). Comparing alternative models of heterogeneity in consumer choice behavior. Journal of Applied Econometrics, 28(6), 1018–1045.

Schön, C. (2010). On the optimal product line selection problem with price discrimination. Management Science, 56(5), 896–902.

Topic S07: Aligning User Preferences and Profits: Exploring the Intersection of Recommender Systems and Choice-Based Optimization

With digital services becoming ever more embedded in daily life, users face an overwhelming abundance of options — whether selecting a restaurant on Lieferando.de, booking a hotel on Booking.com, or finding a new mobile phone contract on Check24.de. Recommender Systems (RS) have become vital tools for helping users navigate this abundance by filtering and ranking options based on inferred preferences.

Econometric choice models, while generally simpler and more theory-driven than data-driven approaches, such as RS, have long served as powerful tools for predicting consumer behaviour. Naturally, they are regularly used to solve pricing, product design, and assortment planning problems. Both recommender systems (RS) and choice-based optimization address a selection problem: identifying the optimal subset of items. In choice-based optimization, this process is typically guided by explicit business objectives (e.g., maximizing expected revenue) and constrained by operational limits such as inventory, budget, or shelf space. Recommender Systems focus on user-oriented objectives such as engagement, satisfaction, and retention, which may not always align directly with a company's strategic or financial goals. In addition, RS frequently integrate design considerations like diversity, novelty, and catalogue coverage to improve the overall user experience. These aspects, however, are typically embedded implicitly within the model training process rather than being formulated as explicit optimization constraints. Despite these conceptual parallels between RS and choice-based optimization, a systematic investigation of their relationship is largely missing from the existing literature. This presents a valuable opportunity to integrate theoretical insights from choice modelling with data-driven approaches from recommender systems, ultimately leading to more effective, user-aware, and business-aligned decision-making frameworks.

This thesis will critically examine the relationship between RS and choice-based optimization, with a particular focus on how these fields can complement and inform one another. To this end, the thesis will begin by introducing the fundamentals of discrete choice models and choice-based optimization, alongside a comprehensive overview of RS, including collaborative filtering, content-based filtering, and hybrid approaches. Literature at the intersection of RS and choice-based optimization will be critically reviewed to assess the conceptual and methodological relationships between the two fields, briefly evaluating their respective strengths and weaknesses, and areas of divergence and convergence. Finally, the thesis will identify open research gaps at the intersection of these domains and outline promising directions for future work.

The primary objectives of this thesis are to:

- introduce the theoretical foundations of RS and choice-based optimization,
- critically review literature at the intersection of RS and choice-based optimization,
- evaluate the strengths and weaknesses of RS and choice-based optimization in the context of managing company operations,
- identify open research gaps and outline future research directions.

Selected Literature Recommendations:

Feldman, J., Zhang, D. J., Liu, X., & Zhang, N. (2022). Customer Choice Models vs. Machine Learning: Finding Optimal Product Displays on Alibaba. Operations Research, 70(1), 309–328.

Kallus, N., & Udell, M. (2020). Dynamic Assortment Personalization in High Dimensions. Operations Research, 68(4), 1020–1037.

Jiang, H., Qi, X., & Sun, H. (2014). Choice-Based Recommender Systems: A Unified Approach to Achieving Relevancy and Diversity. Operations Research, 62(5), 973–993.

Ricci F, Rokach L, Shapira B (2011). Introduction to recommender systems handbook. Ricci F, Rokach L, Shapira B, Kantor PB, eds. Recommender Systems Handbook (Springer, Boston), 1–35.

Roy, D., & Dutta, M. (2022). A systematic review and research perspective on recommender systems. Journal of Big Data, 9(1), 59.

Topic S08: Learning to Optimize: Machine Learning for Combinatorial Optimization

Combinatorial optimization (CO) problems are ubiquitous in real-world decision-making and industrial planning. Companies such as Amazon, UPS, and FedEx must solve large-scale vehicle routing and delivery scheduling problems daily to minimize transportation costs and delivery times. Likewise, airlines face complex crew scheduling and aircraft assignment problems, while retail and e-commerce, firms like Walmart and Zara tackle assortment optimization and product line design problems, where the goal is to select an optimal subset of products or configurations that maximize expected profit or customer satisfaction subject to production, space, or cannibalization constraints. Each of these settings can be formalized as a combinatorial optimization problem, where decisions are discrete and interdependent, and the feasible set grows exponentially with the problem size.

Traditionally, such problems have been addressed using operations research (OR) techniques such as branch-and-bound algorithms and various meta heuristics. These approaches rely heavily on handcrafted rules and expert knowledge to guide the search process. However, recent research has shown that machine learning (ML) methods can learn optimal or near-optimal decision policies for combinatorial optimization problems by leveraging data from past instances, for example, historical delivery routes or previous product configurations. Nevertheless, other important classes of combinatorial optimization problems, unrelated to transportation, remain relatively underexplored, offering promising avenues for advancing learning-based optimization schemes.

This thesis will critically review the existing literature on the application of ML to CO problems. It will provide an overview of the different types of combinatorial problems and their defining characteristics, examine the range of ML schemes employed in the literature, and analyse how these learning-based approaches operate and perform in comparison to traditional optimization methods. Furthermore, the thesis will briefly discuss the strengths and limitations of current approaches, highlighting their practical implications and theoretical underpinnings. Finally, it will identify open research gaps at the intersection of machine learning and combinatorial optimization and outline promising directions for future research in this rapidly evolving field.

The primary objectives of this thesis are to:

- introduce the field of ML for CO,
- explain the principal classes of combinatorial optimization problems along with their defining characteristics and elaborate on the theoretical foundations that underpin the various optimization schemes analyzed in this thesis,
- review the range of ML schemes employed in the literature to solve CO problems and analyze how these learning-based approaches operate and perform in comparison to traditional optimization methods,
- discuss the strengths and limitations of ML for solving CO problems,
- identify open research gaps and outline future research directions,

Selected Literature Recommendations:

Bello, I., Pham, H., Le, Q. V., Norouzi, M., & Bengio, S. (2017). Neural Combinatorial Optimization with Reinforcement Learning (No. arXiv:1611.09940). arXiv.

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.

He, H., Daumé, H., & Eisner, J. (2014). Learning to Search in Branch and Bound Algorithms. Advances in Neural Information Processing Systems, 27.

Zhou, F., Lischka, A., Kulcsar, B., Wu, J., Chehreghani, M. H., & Laporte, G. (2025). Learning for routing: A guided review of recent developments and future directions. Transportation Research Part E: Logistics and Transportation Review, 202, 104278.

Topic S09: Price Optimization in Omnichannel Retailing

Omnichannel retailing integrates a retailer's physical stores, online shops, mobile apps, and social media platforms into a unified and seamless customer experience. This integration enables customers to switch effortlessly between channels — for instance, by browsing products online on IKEA.de, and locating them via the mobile app inside the physical store before purchasing them.

A key advantage of omnichannel retailing from an operations management perspective lies in the wealth of data generated across multiple customer touchpoints, which can inform better and more responsive pricing decisions. However, coordinating prices across online and offline channels remains one of the central challenges for omnichannel retailers. Consumers frequently compare prices across channels, engage in showrooming and webrooming behaviours, and expect price consistency and fairness — all while firms strive to maintain profitability and competitiveness. Traditional channel-specific pricing practices often fail to exploit cross-channel substitution effects and can lead to revenue loss or brand inconsistency. In response, recent research has proposed advanced pricing frameworks that integrate data-driven demand estimation, cross-channel elasticity modelling, and optimization techniques to coordinate prices across channels and locations.

This thesis will introduce the fundamental concepts of omnichannel pricing and explain how pricing decisions interact with demand, inventory, and service considerations. It will review state-of-the-art approaches to price optimization and critically discuss findings from this review, with particular emphasis on how data is leveraged to inform and enhance pricing decisions. Finally, the thesis will identify open research gaps and outline promising directions for future work.

The primary objectives of this thesis are to:

- introduce the concept of omnichannel retailing and its implications for price optimization,
- explain the theoretical fundamentals of the most relevant optimization approaches,
- critically review state-of-the-art approaches to price optimization,
- discuss key findings and insights from the literature, with particular emphasis on how data is leveraged to inform and enhance pricing decisions, while highlighting the challenges, limitations, and managerial implications of current approaches,
- identify open research gaps and outline future research directions.

Selected Literature Recommendations:

Gupta, V. K., Ting, Q. U., & Tiwari, M. K. (2019). Multi-period price optimization problem for omnichannel retailers accounting for customer heterogeneity. International Journal of Production Economics, 212, 155–167.

Harsha, P., Subramanian, S., & Ettl, M. (2019). A Practical Price Optimization Approach for Omnichannel Retailing. INFORMS Journal on Optimization, 1(3), 241–264.

Jiang, Y., Liu, L., & Lim, A. (2020). Optimal pricing decisions for an omni-channel supply chain with retail service. International Transactions in Operational Research, 27(6), 2927–2948.

Liu, S., Wang, J., Wang, R., Zhang, Y., Song, Y., & Xing, L. (2024). Data-driven dynamic pricing and inventory management of an omni-channel retailer in an uncertain demand environment. Expert Systems with Applications, 244, 122948.

Ram, B., Lai, K. K., & Xiao, F. (2024). Retail reimagined: Pricing optimization in single channel, multi-channel and omni-channel retailing. Annals of Operations Research.

Topic S10: Driving Change: Leveraging Analytics for Sustainable Freight Transportation

The global population, currently estimated at 7.6 billion, is projected to increase by roughly 50% by 2100, as forecasted by the United Nations Department of Economic and Social Affairs. This growth will amplify the demand for freight transport, which is fundamental not only for economic growth but also for ensuring social welfare. Freight, moved by multiple modes of transport—road, rail, sea, and air—serves as a critical link in global supply chains. Despite advances in efficiency and distribution, the environmental impact of freight transportation remains substantial. Greenhouse gas emissions, fuel consumption, and pollution associated with transport have spurred an urgent need for more sustainable logistics practices. Here, analytics has a significant role to play, offering powerful tools for enhancing both the sustainability and efficiency of freight transportation.

The primary objectives of this thesis are to:

- identify and analyze the primary challenges faced by the freight transportation sector in transitioning to greener practices, including issues related to emissions reduction, regulatory compliance, infrastructure limitations, and cost implications;
- review and classify the literature for models, tools, and technologies available for optimizing green freight logistics;
- discuss the role that analytics has played and may play in the future towards greener freight transportation, along with its limitations;
- discuss future research opportunities and practice needs.

Selected Literature Recommendations:

Bektaş, T., Ehmke, J. F., Psaraftis, H. N., & Puchinger, J. (2019). The role of operational research in green freight transportation. European Journal of Operational Research, 274(3), 807-823.

Moghdani, R., Salimifard, K., Demir, E., & Benyettou, A. (2021). The green vehicle routing problem: A systematic literature review. Journal of Cleaner Production, 279, 123691.

McKinsey & Co. (2024). Decarbonizing logistics: Charting the path ahead. https://www.mckinsey.com/capabilities/operations/our-insights/decarbonizing-logistics-charting-the-path-ahead

Topic S11: Airline Schedule Development – Overview, Selected Optimization Models and Applications

The planning of an airline's schedule poses significant challenges, as it involves multiple interdependent, customer-oriented operational decisions ranging from strategic to short-term planning: which origin-destination (OD) markets to serve, at what frequencies and departure times, over which routes, with which fleet types and crew, and under which prices and fare conditions. These elements are tightly coupled: fleet and schedule decisions largely determine fixed and variable operational costs, while pricing and schedule characteristics strongly influence customer demand and revenue. Due to the complexity of real-world networks, airlines rely on decision support tools that traditionally decompose the planning process into sub-problems – such as schedule design, fleet assignment, maintenance routing, crew scheduling, and revenue management – which are solved sequentially. However, this sequential approach often leads to suboptimal system-wide outcomes. In contrast, recent research has proposed integrated planning models that combine two or more sub-problems (e.g., schedule design, fleet assignment, and pricing) to achieve better overall solutions. These integrated models are computationally challenging and require advanced solution methods capable of solving large-scale instances efficiently.

The focus of this seminar thesis is on integrated optimization models to jointly determine the schedule and fleet assignment with choice-based passenger demand, resulting in more realistic and profitable solutions. In particular, the paper by Yan et al. (2022) introduces the so-called CSD-FAM model, which incorporates schedule-based demand into a combined schedule design and fleet assignment problem. To handle the combinatorial complexity, they use a Sales-Based Linear Programming (SBLP) revenue approximation to estimate passenger choice and revenue. This SBLP is itself an approximation to a true discrete choice model such as Multinomial Logit (MNL) or discrete mixed multinomial logit (MMNL).

The primary objectives of this thesis are to:

- provide a comprehensive literature review of optimization models for airline schedule development with a particular focus on schedule design and fleet assignment, including both, separate and integrated models,
- explain the model by Yan et al. (2022) in detail, including limitations and improvement opportunities. You do not need to discuss the decomposition approach to solving the problem,
- compare the model theoretically to alternative approaches for integrated schedule design and fleet assignment,
- identify future research opportunities in strategic airline schedule planning.

Selected Literature Recommendations:

Krömer, M. M., Topchishvili, D., & Schön, C. (2024). Sustainable airline planning and scheduling. Journal of Cleaner Production, 434, 139986.

Lohatepanont, M., & Barnhart, C. (2004). Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment. Transportation science, 38(1), 19-32.

Schön, C. (2008). Integrated airline schedule design, fleet assignment and pricing. DSOR-Beiträge zur Wirtschaftsinformatik, DSOR Contributions to Information Systems 5, 73–88.

Xu, Y., Adler, N., Wandelt, S., & Sun, X. (2024). Competitive integrated airline schedule design and fleet assignment. European Journal of Operational Research, 314(1), 32-50.

Yan, C., Barnhart, C., & Vaze, V. (2022). Choice-based airline schedule design and fleet assignment: A decomposition approach. Transportation Science, 56(6), 1410-1431.

Topic S12: Reinforcement Learning Approaches for Dynamic Revenue Management

Revenue management (RM) plays a critical role in industries such as airlines, hospitality, and retail, where firms must optimally set prices and allocate limited capacity to uncertain and time-dependent demand. Traditional RM systems rely on dynamic programming or stochastic optimization models that assume known demand distributions and customer behaviors. However, in practice, these assumptions are often violated: demand functions are unknown, customer preferences evolve, and competitive environments are highly dynamic. Consequently, the ability to learn optimal pricing and allocation policies directly from interaction with the environment – without requiring precise model specifications – has motivated growing interest in reinforcement learning (RL) as a powerful alternative.

In recent years, RL has shown promise in addressing several limitations of classical RM models. By treating pricing and inventory control as sequential decision-making problems under uncertainty, RL can adaptively balance exploration (learning demand patterns) and exploitation (maximizing revenue based on current knowledge). Moreover, RL frameworks naturally handle high-dimensional state spaces, dynamic customer segments, and complex competitive settings where closed-form solutions are infeasible. Despite these advantages, applying RL to RM remains challenging due to issues such as sparse feedback, non-stationary environments, and the need for explainability in decision-making systems that affect pricing fairness and customer trust.

This seminar thesis focuses on the intersection of reinforcement learning and revenue management, exploring how modern RL methods can enhance dynamic pricing and capacity allocation under uncertainty. In particular, it builds upon recent work that applies RL to model dynamic customer behavior, learn pricing strategies from data, and optimize revenues across multiple time horizons and market conditions.

The objectives of this thesis are to:

- provide a comprehensive literature review of reinforcement learning approaches for revenue management,
- explain the underlying key methods and assumptions of RL models for dynamic pricing and capacity control theoretically and with academic examples,
- discuss a selected approach in detail, including contribution, model and method, limitations and improvement opportunities,
- theoretically compare RL-based RM approaches with classical stochastic and optimization-based methods in terms of scalability, adaptability, and robustness,
- identify future research opportunities at the interface of machine learning and revenue management.

Selected Literature Recommendations

Bondoux, N., Nguyen, A.Q., Fiig, T. et al. Reinforcement learning applied to airline revenue management. J Revenue Pricing Manag 19, 332–348 (2020). https://doi.org/10.1057/s41272-020-00228-4

Lange, F., & Schlosser, R. (2025). Dynamic pricing with waiting and price-anticipating customers. Operations Research Perspectives, 14, 1-20.

Ilan Lobel (2020) Revenue Management and the Rise of the Algorithmic Economy. Management Science 67(9):5389-5398. https://doi.org/10.1287/mnsc.2020.3712

Meng, H., Chen, N., & Gao, X. (2024). Reinforcement learning for intensity control: An application to choice-based network revenue management. arXiv preprint arXiv:2406.05358.

Shihab, S.A.M., Wei, P. A deep reinforcement learning approach to seat inventory control for airline revenue management. J Revenue Pricing Manag 21, 183–199 (2022). https://doi.org/10.1057/s41272-021-00281-7