1. Discipline:

Operations Management/Operations Research

2. Title:

Data-driven Operations Management

3. Lecturer:

Dr. Richard Pibernik, Professor of Logistics and Quantitative Methods, Julius-Maximilians-Universität Würzburg, Germany

Dr. Christoph M. Flath, Chair of Information Systems and Management, Julius-Maximilians-Universität Würzburg, Germany

4. Date and Location

Würzburg, October 7-10, 2019

5. Course Description

5.1 Abstract and Learning Objectives

This course explores when and how the availability of large amounts of relevant data ("big data") affects decision making in operations management. As such, the course combines elements of artificial intelligence (AI)/machine learning (ML) with traditional approaches in operations management. Thereby, it addresses the pertinent question of how new models in operations and supply chain management will evolve, or traditional models have to be modified, to leverage extensive auxiliary data. The course will first introduce participants to the most relevant AI/ML techniques for operations management. After an in-depth discussion of the traditional paradigm of operations management ("sequential estimation and optimization"), its critical assumptions and potential shortfalls, participants will experience and discuss how decision making in operations and supply chain management may change when "big data" is available. The latter will be based on a number of new and relevant publications in the field of data-driven operations management (see references below), as well as selected practical cases and datasets that are currently used in research projects of the organizers of the course. While the primary focus of the course lies on recent developments in operations and supply chain management, and not the development of hands-on skills in implementing ML techniques, the course will include a number of lab session to illustrate how "novel" models in data-driven operations management leveraging ML techniques can or should be implemented and evaluated. The course will also include contributions of international guest lecturers who will share their view of how data-driven operations management will evolve, and how it will shape the research agenda in the field.

5.2 Content

The course will cover the following topics:

- Fundamentals of machine learning and artificial intelligence for operations management
- From predictive to prescriptive analytics in operations management
- Data-driven inventory management
- Data-driven capacity management
- Data-driven logistics operations
- Data-driven pricing and revenue management
- Discussion of a future research agenda in data-driven operations management

5.3 Tentative Schedule

Day 1: Fundamentals: A primer in ML and AI for operations management

- The foundations, frameworks and applications of the emerging field of advanced analytics
- Design, implement, and evaluate the core algorithms underlying an end-to-end data science workflow, including data import, analysis, and presentation of information
- Leverage Jupyter Notebook application and data infrastructure that supports data acquisition, manipulation, visualization and modeling
- Implementation and execution skills for descriptive and predictive analytics

Day 2: From predictive to prescriptive analytics in operations management

- Combining Optimization and Machine Learning
- Sequential Estimation and Optimization (SEO) vs. Joint Estimation and Optimization (JEO)
- Machine learning with operational costs
- Lab session & demos

Day 3: Data-driven inventory & capacity management

- The Big-data Newsvendor and its extensions
- Data-driven multi-period inventory management
- Inventory management and pricing for new products based on online data
- Capacity management in service operations – from queuing theory to data-driven approaches
- Data-driven multi-period capacity management
- Lab session & demos

Day 4: Further applications of data-driven methodologies in operations management & discussion

- The machine learning and traveling repairman problem
- Sales forces routing with uplift information
- Thinking ahead – challenges and opportunities of an emerging research area
- Which journals to target?
- How to strike the balance between OM and ML?

5.4 Course format

The course will comprise lectures and presentations addressing the fundamental theories and applications of the different topics described above, as well as the presentation and discussion of important research papers by the participants. Short lab sessions will complement the presentations and discussions.

To make this short course productive, each participant is required to study the fundamental literature (see below) and to prepare at least one presentation of a research paper (either individually or in a team of two, depending on the number of participants).

5.5 Prerequisites

Solid quantitative knowledge in operations management and logistics (basic inventory models, basic statistics, economic analysis, fundamentals of optimization). Some basic programming skills are beneficial.

5.6 Essential Reading Material

- Bertsimas/Kallus (2018): From Predictive to Prescriptive Analytics. Forthcoming in *Management Science* https://arxiv.org/pdf/1402.5481.pdf
- Tulabandhula/ Rudin (2013): Machine learning with operational costs. *The Journal of Machine Learning Research,* 14.1: 1989-2028.
- Ban/Rudin (2018): The Big Data Newsvendor: Practical Insights from Machine Learning, Forthcoming in *Operations Research*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2559116
- Ramamurthy/Shanthikumar/Shen (2012): Inventory Policy with Parametric Demand: Operational Statistics, Linear Correction, and Regression. *Production and Operations Management*, Vol. 21 (2): 291-308.
- Ferreira/Lee/Simchi-Levi (2016): Analytics for an Online Retailer: Demand Forecasting and Price Optimization. *Manufacturing & Service Operations Management*, Vol. 18 (1): 69-88.
- Ban/Keskin (2017): Personalized dynamic pricing with machine learning. Working Paper. London Business School. https://www.london.edu/faculty-and-research/academic-research/p/personalized-dynamic-pricing

5.7 Additional Reading Material

Will be distributed prior to the course.

5.8 To prepare

Participants will have to study – in detail – the essential readings (see 5.6) and prepare a presentation on a selected research paper. The set of research papers to be discussed in class will be distributed and assigned prior to the course.

6. Administration

6.1 Max. number of participants

20

6.2 Assignments

Participants will have to present at least one research paper (see 5.8).

In addition, all participants have to provide a written report of approximately 15-20 pages in which they review the relevant literature pertaining to the research paper they were assigned, describe the most important research findings of the paper, and outline future research directions.

6.3 Exam

The final grade will be composed of the grade for the paper presentation (weight: 1/3), and the written report (2/3).

6.4 Credits

6 ECTS.

7. Working Hours

Working Hours	**Stunden**
Preparation	60 h
Class participation	30 h
Preparation of presentation	20 h
Preparation of written report	70 h
SUMME	**180 h**
ECTS: 6	