Julius-Maximilians-
UNIVERSITÄT
WÜRZBURG

Wirtschafts-
wissenschaftliche
Fakultät

An Introduction to Data-driven Operations & Logistics Management

An online course offered by the Chair of Logistics & Quantitative Methods at the University of Würzburg, Germany

Target Audience – Who should participate in this course and continue reading this document?

It all depends on your background and your interests. With this course, we address three groups of students:

- *Basic operations and logistics learners*: You have an interest in operations and logistics management, have attended some basic courses in this field, have basic knowledge in quantitative techniques and statistics and are interested in getting to know how "more and better data" and the use of machine learning techniques will shape the future of operations and logistics management. You do not have extensive skills in developing code and do not know much about machine learning. This course will provide you with an intuitive avenue into Data-driven Operations Management (DDOM) – see below for a description – and will show you the basic concepts of how we can learn to take better decisions based on available data. You will be able to understand the concepts and the underlying mechanisms without getting into the details of the code and the theory of machine learning. As a side effect, however, you will understand the basic logic of certain machine learning techniques (e.g.….)

- *Basic ML-Learners*: You have less or no background in operations and logistics, but you already have a working knowledge in machine learning and consider yourself a future data scientist. You want to understand how your ML-expertise can be used for taking better decisions in companies, in particular, in operations and logistics, which is a big and important field. From this course you will learn how basic and more sophisticated ML-techniques can be employed for better decision making in operations and logistics. In particular, you will see how we can combine ML-techniques and optimization in a practically relevant context and you will understand how these concepts can be used for other decision making problems in operations and logistics. You will be able to see and interpret the code and develop other approaches by yourself.

- *ML-Experts*: You are already acquainted with ML-techniques and know how to implement them in Python or the like, but want to find out how your expertise can make an impact on important management problems. This course will provide you with a real-world problem, data, concepts and the corresponding implementations. You will work through the analytical concepts, the code and – because we share everything with you – develop something new and exciting that will hopefully make decisions even better. We provide you with code to start with and to question, as well as sample results that you can use as benchmarks for you own approaches.

Depending on your assessment of your expertise, you will have to complete a final assignment that is tailored to your level of experience.

Motivation – Why is this this interesting?

Companies face considerable uncertainty when taking important decisions in operations & logistics management, for example, when deciding upon capacities, inventory levels, transportation and production schedules. One important source of uncertainty is, of course, customer demand. Final products are sold and distributed through different channels, and it is extremely difficult to predict the channels' demands for individual products on individual days, weeks or months, making it even more difficult to set the right inventory and capacity levels. Other sources of uncertainty (production and replenishment lead times, for example) further complicate decision making in operations and logistics. Most operations managers would probably argue that uncertainty has increased in the recent past, because there are more products (SKUs), life cycles have become shorter, demand is more fragmented across channels, and for a number of other reasons.

The good news, however, is that companies have unprecedented access to (more and more) data that can potential help taking better decisions under these uncertain and challenging conditions. Data from social media, clickstreams, web searches, weather data, calendar data (holidays, etc.), data scraped from competitors' websites, for example, can potentially be used to better predict customer demand and to take better inventory and capacity decisions. Sensor data, detailed shipping records and weather data can possibly be used to better predict replenishment lead times and to improve ordering decisions. These are just some examples of how access to more and better data may improve operations and logistics management, and there are many more.

Turning this data into better decisions in terms of costs and customer satisfaction lies at the heart of what we call "data-driven operations management" (DDOM). DDOM is mostly concerned with figuring out how we can best leverage a given data set that contains, for example, historical demands and many other explanatory variables, to optimize our inventory and capacity levels, our replenishment decisions and various other decisions in operations and logistics management. DDOM has always been there – managers have always taken decisions based on some data – but today we have access to more and better data and we have more sophisticated methods and computational power. Increasingly, DDOM embraces methods from the machine learning domain and combines these with optimization techniques.

Course contents – What can you expect from this course?

This course gives you a basic hands-on introduction to the new field of DDOM. Based on a real-world example, you will learn step-by-step how we can use simple and more sophisticated techniques to get from data to (better) decisions. To focus on the core concepts of DDOM and for pedagogical reasons we use the most common inventory model (the "Newsvendor Model") to explain important concepts of DDOM. We do, however, also explain how these concepts can be translated to other (more complex) decision making problems in operations and logistics management.

The course is structured around an online tutorial implemented in a Google Colab Notebook in which you are introduced to the real-world problem and the available data, and are guided through its pre-processing and various DDOM techniques. We start with very simple concepts that any participant can relate to, and gradually increase complexity so that, finally, we cover fairly sophisticated machine learning-based techniques for obtaining a decision for the case company. The individual steps in the analyses are explained in an intuitive way and are illustrated based on sample results. Each step of the analyses is accompanied by the corresponding Python code that can be directly executed within the Colab Notebook. The code is explained in an intuitive manner and you have the option to delve deeply into the implementation, but are not required to do so.

Pre-requisites – What do you need to know before starting with this course?

We will provide material for preparation that you will have to work through before starting the course. Your effort for the preparation largely depends on your level of expertise and previous education.

Since this course will focus mainly on data-driven implementations of the well-known Newsvendor model, any participant has to have a solid knowledge of the concepts of the model and how we typically use it. This is the most important pre-requisite.

Moreover, we draw on basic concepts from multivariate statistics (multiple linear regression) and basic optimization techniques that are typically taught in undergraduate programs in business/management, engineering, etc.

Basic knowledge in coding is a plus, but not strictly required

Course mechanics – How are we going to run this course?

We will provide the material for preparation 6 weeks before the start of the course. As described previously, the course is structured around an online tutorial and a Colab Notebook. The course will span 5 weeks (???) and you will have to work through individual chapters of this tutorial each week. For each chapter there will be a live review-session. We will also have a kick-off workshop and a final workshop in which participants present the results of their assignments and where we discuss future applications of DDOM.

Timelines

- Enrolment to the course by May 31, 2021 the latest. This is a hard deadline.
- Distribution of materials for preparation: June 1, 2021
- Kick-off session: July 16
- Review of individual chapters on a weekly basis from July 23 to end of August. Dates will be scheduled according to participants' preferences and availability of instructors.
- Final workshop: Mid of September. Date and time will be scheduled based on participants' preferences and availability of instructors.

Fees and Admission

The course is free of charge. However, you have to register with the University of Würzburg as a student in order to participate and to obtain credits. The registration also comes free of charge, but you have to furnish a set of documents (transcripts, etc.) for being registered. We will provide you with the details.

We admit up to 20 participants and select from applications based on merit.

Credits & Certificates

Being an official student at the University of Würzburg, you will receive 5 credit points according to the European Credit Transfer System (ECTS). Participants outside the European Union have to check with their Registrar's office or their International Office how their respective Universities or Schools account for these credit points.

We can issue a separate certificate of participation if required.

Contact

Andreas Philippi, Researcher at the Chair of Logistics & Quantitative Methods (andreas.philippi@uni-wuerzburg.de)

Dr. Richard Pibernik, Full Professor of Logistics & Quantitative Methods (richard.pibernik@uni-wuerzburg.de)