The lecture "TMA4265: Stochastic Processes" is roughly divided into five sections: discrete-time Markov chains, Poisson processes, continuous-time Markov chains, queuing theory, Brownian motion. At the end of the first course segment the students should be able to

- explain what a discrete-time Markov chain is. You should be able to formulate a Markov chain and derive the corresponding transition graph linking the states of the state space.
- compute state probabilities and classify the states into equivalence classes and determine important class properties, such as transience, recurrence (positive and null) or periodicity.
- apply first-step analysis to answer questions such as in which state the Markov chain is ultimately trapped and how long, on average, it takes to reach this state.
- investigate long-term behavior of a Markov chain and distinguish between stationary and limiting probabilities.
- compute the expected time a Markov chain spends in transient states
- recognize and describe a branching process, determine if it converges to infinity or dies out, and compute certain quantities of interest.
- derive the transition probabilities of the reversed Markov chain, decide whether the process is time-reversible and exploit time-reversibility to compute limiting probabilities.
- simulate a Markov chain and check key-quantities derived theoretically using simulations. You should have a rough idea about Markov chain Monte Carlo procedures.

By the end of the second part the students will be able to

- define and illustrate a Poisson process, and compute key quantities such as arrival or inter-arrival times.
- work with the Poisson and exponential distribution, exploit their properties in the computation of conditional probabilities, and evaluate the superposition and intersection of several Poisson processes.

At the end of the third part the student will be able to

- explain a birth-death process and state its key properties. Further birth and death rates can be determined and visualized in a graph.
- compute limiting probabilities and expected values such as the mean time until the process reaches a certain state for the first time. Furthermore, the student should be able to derive the probability that a departure happens before a new arrival.

- set up and apply Kolmogorov's backward and forward equations.

By the end of the fourth part the student will be able to

- set-up or recognize a queueing model and describe it completely.
- name cost equations and understand the PASTA principle.
- work with exponential server models using a single or multiple server with both finite and infinite capacity. The student will be able to compute the average number of customers in the system and in the waiting line. Further the times spent in the system or in the waiting queue can be computed for different processes.
- understand more general queueing models and the concept of busy periods.

By the end of the fifth part the students is able to

- define a Brownian motion (Wiener) process.
- compute hitting times and probabilities for the maximum value the process attains in a certain time interval.
- extend the Brownian motion to include drift components or to the geometric version.