

DATA SCIENCES (DS)

DS 540: Multimedia Analytics

3 Credits

This course covers all aspects of multimedia analytics including the foundations, representations, implementations, and applications of extracting patterns from images, video, audio, and related context information for the purpose of prediction, automated discovery, and decision-making. Specific topics to be covered in the course include an overview of multimedia systems, multimedia data representation and compression, image and video analysis, audio analysis, multimedia information retrieval, and advanced topics in multimedia analytics. Emphasis will be placed on the integrative approach to multimedia analytics, in which the systems can benefit from the synergy of leveraging multiple sensing modalities, including audio, image, video, text, and more. The objective of the course is to provide students with a wide range of models and applications that are relevant to their research and development in modern multimedia analytics systems.

recommended prep: IST 510

DS 560: Principles of Causal Inference

3 Credits

The course will give students a comprehensive coverage of the theoretical underpinnings, and practice of causal inference from observational and experimental data. Topics to be covered include: pitfalls of standard machine learning algorithms when applied to observational data; causal inference in the absence of randomized control trials; causal effects and counterfactuals; eliciting causal effects from observations; the Causal Bayesian Network framework for causal inference - do-calculus, identifiability of causal effects from observations and experiments; the Potential Outcomes framework for causal inference - matching and propensity score-based methods and their advanced variants for counterfactual inference; the relationship between the Potential Outcomes and causal Bayes Networks; and learning causal models from observations and experiments. The course will give a principled treatment to confounders as well as practical approaches to cope with them. Additional topics to be covered include mediation analysis; advanced machine learning methods for causal effect estimation; causal transportability; selection bias; and meta-analysis. Finally, the course will include a laboratory component to provide students with hands-on experience with applications of causal inference to problems from several domains. Course projects will focus on applications of causal models and causal inference e.g., in science, public policy, and health. Recommended preparation for the course include basic proficiency in programming, elements of probability theory and statistics, discrete mathematics, and machine learning.

Recommended Preparations: IST 510, Probability and statistics, differential calculus, linear algebra, programming proficiency (e.g., in Python), and machine learning