Title: Algorithmic and Computational Statistics

Instructor: Robert S. Maier (Professor, Mathematics; Professor, Physics; Affiliate Member, Statistics Program)

Textbook: G. H. Givens and J. Hoeting, "Computational Statistics", second edition, Wiley, 2013.

Context:

This Topics course has been taught several times and is now proposed for Fall 2025. (I last taught it five years ago, successfully.) It will be a higher-level counterpart to the existing course STAT 675 ("Statistical Computing"). Unlike STAT 675, it will not provide training in the use of statistical software, such as R. Instead it will focus on theory: the mathematics of advanced computational statistics, including algorithms other than ones for machine learning.

The course should appeal to students in the Statistics & Data Science Ph.D. program, and to students in the Mathematics and Applied Mathematics graduate programs who would like to see statistical applications of analysis and numerical analysis. It will not compete with any of the existing core or non-core graduate courses.

Prerequisites: Any of the core courses MATH 523, 584, 589, or 564/566; or a comparably advanced course taught by the Systems & Industrial Engineering (SIE) Department.

Approximate Schedule:

1. Aspects of maximum likelihood and Bayesian inference (review). 1.5 weeks.

2. Numerical optimization and the numerical solution of nonlinear equations (Newton-Raphson and more advanced multivariate schemes; statistical applications). 1.5 weeks.

3. Simulated annealing. 1.5 weeks.

4. Numerical integration, with applications to Bayesian inference and simulation. Related mathematical topics, such as orthogonal polynomials of many types; how orthogonal polynomials are related to standard distributions appearing in probability and statistics. 3 weeks.

5. The issue of limited machine precision: the IEEE format for floating point numbers, and how it affects the accuracy of numerical computations. 1 week.

6. Sampling from a Bayesian posterior distribution: MCMC (Markov chain Monte Carlo), etc. Also, much background information on Markov chains and their structure. 2.5 weeks.

7. Smoothing and nonparametric density estimation; mathematical aspects. 1.5 weeks.

8. Alternative information criteria for model selection, i.e., deciding between fitted statistical models on the basis of data; connections with information theory. 1 week.

9. Brief remarks on quantum information and computation. 1.5 weeks.

Expected Learning Outcomes:

The following will be taught and learned.

1. How to interpret statistical inference function-theoretically.

2. How to exploit numerical optimization methods in a statistical context (e.g., while performing parameter estimation).

3. How to exploit simulated annealing techniques in any optimization context.

4. How to identify the standard families of orthogonal polynomials, and explain how they arise in applied probability and also numerical integration (whether or not in the context of Bayesian inference).

5. How to resolve floating-point issues in numerical and statistical computation (i.e., limited machine precision, overflow and underflow, and NaN [not-a-number] problems).

6. How to analyse the structure of a Markov chain, whether or not it arises in MCMC (Markov Chain Monte Carlo).

7. How to analyse nonparametric density estimation, theoretically and also practically (e.g., how to shrink histogram bins as a sample size increases).

8. How to apply alternative information criteria (AIC, BIC, and others), in model selection, and how to justify them using asymptotic analysis.