Training in Advanced Low-Energy Nuclear Theory
(see TALENT website at nucleartalent.org)

Course 11. Learning from Data: Bayesian Methods and Machine Learning

A three-week TALENT course on Bayesian inference, focused on its applications in nuclear physics, will take place at the University of York in June 2019 (tentatively June 10-28).

Introduction
In recent years there has been an explosion of interest in the use of Bayesian methods in nuclear physics. These methods are being used to quantify the uncertainties in theoretical work on topics ranging from the NN force to high-energy heavy-ion collisions, to develop more reliable extrapolants of nuclear-energy-density functionals towards the dripline, to predict the impact that future NICER observations may have on the equation of state of neutron matter, and to determine whether or not nucleon resonances are present in experimental data. Meanwhile machine learning is gaining increased currency as a method for identifying interesting signals in both experiments and simulations.

While most nuclear-physics Ph.D. students are taught some standard (frequentist) statistics as part of their course work, very few encounter Bayesian methods until they are engaged in research. But Bayesian methods provide a coherent and compelling framework to think about inference, and so can be applied to many important questions in nuclear physics. The overall learning goal of this school is to take students who have had no previous exposure to Bayes' theorem and show them how it can be applied to problems of parameter estimation, model selection, and machine learning.

Personnel:
Lecturers:

- Christian Forssén (Chalmers University, Sweden)
- Dick Furnstahl (Ohio State University, USA)
- Daniel Phillips (Ohio University, USA)
- Ian Vernon (Durham University, UK)

Forssén, Furnstahl, and Phillips all have experience applying Bayesian methods to problems in low-energy nuclear physics. Vernon is a statistician who has applied Bayesian methods to a range of problems including galaxy formation and epidemiology. They are all involved in organizing the "ISNET" (Information and Statistics in Nuclear Experiment and Theory) series of workshops. Alessandro Pastore (York University, UK) will act as the local liaison and work with the lecturers to help students complete the exercises. Additional facilitators may be added later.

Audience: advanced graduate students and early postdocs in all subfields of low-energy nuclear physics (broadly construed), including both theorists and experimentalists. Students will be selected by a competitive process, based on CVs and letters from their advisor.

Format:
The course material delivered over three weeks. Each day will begin with lectures and "in-class" exercises in the mornings, and then students work in groups on more extended analytical and computational exercises in the afternoons. Solutions, implications, and extensions of these exercises will be discussed by everyone at special wrap-up sessions at the end of each day. We will use Python/Jupyter notebooks to supplement lecture material and provide computational exercises. Total lecture and exercise hours will be comparable to previous TALENT courses, with a

total of 45 hours of lectures and directed exercises, and about 75 hours devoted to out-of-class analytical and computational exercises (out of this 60 hours will be supervised). York University will provide individual certificates to attest to students' participation in the course.

Topics (tentative)**:**

- Basics of Bayesian statistics
  - Bayes theorem, the laws of probability, and what it does and doesn't do for you
  - Bayesian updating 101: the coin-toss example; lighthouse problem
  - What pdfs for parameters mean: 68%, 95%, etc. intervals, correlations, error ellipses, etc.
- Bayesian parameter estimation: recovering least-squares regression and standard-deviation estimators
- Why Bayes is better
  - I: dealing with outliers
  - II: fitting in the presence of higher-order terms
- How to choose a prior
  - subtleties of non-informative priors
  - maximum entropy, Jeffrey's prior, etc.
- Modeling systematic errors
  - general principles
  - example: how to deal with experimental normalization errors (including a discussion of the "d'Agostini bias" problem)
- Sampling
  - MCMC: how it works
  - more sophisticated MC (HMC, NUTS)
- Model selection
  - model evidence: what it is, why it's useful
  - LOO cross-validation metrics
  - example: How many peaks?
  - example: What order is the polynomial?
- Experimental planning
- Gaussian processes
  - Introduction and examples
  - Gaussian processes as machine learning
  - Emulators for expensive computations
- Other topics as time permits
  - Dimensional reduction and sloppy parameters
  - Bayesian neural networks
  - Maximum entropy methods for inversion
  - Non-parametric bootstrap
  - Bayesian optimization
- Survey of applications in low-energy nuclear physics, integrated as appropriate

Preparation:

Introductory materials and tutorial Jupyter notebooks will be made available in advance of the course so that all students arrive appropriately prepared.