



Friday, April 7 • 2:00 PM – BPB, Rm. 130

## **Uncertainty Quantification using Embarrassingly Small Numbers of Simulations and Experiments**

Ilias Bilonis, Assistant Professor  
School of Mechanical Engineering, Purdue University

**Abstract:** Increasing modeling detail is not necessarily correlated with increasing predictive ability. Setting modeling and numerical discretization errors aside, the more detailed a model gets, the larger the number of parameters required to accurately specify its initial/boundary conditions, constitutive laws, external forcing, object geometries, etc. To be predictive, we need to quantify this high-dimensional parametric uncertainty by combining our prior physical knowledge with noisy experimental data obtained from various heterogeneous sources. Once we have quantified this uncertainty, all we need to do is propagate it through the model and obtain predictive error bars for any quantity of interest. This program sounds reasonable, but after a little bit of thinking, we realize that it is completely impractical! How many simulations or experiments do you think are required to bring such a program to fruition? Tens, hundreds, thousands, or, perhaps, millions? It all depends on the number of uncertain parameters and the smoothness of the response surface. Unfortunately, for general response surfaces the number of required information source queries grows exponentially with increasing number of parameters (the curse of dimensionality). The situation seems hopeless for both a cutting-edge modeler, with simulation models taking a week or more to run on a supercomputer and an experimentalist with a high-dollar-cost per experiment. Questions like "What is the mean and the variance of a quantity of interest?" cannot be answered without incurring an unreasonably high cost of information acquisition. Therefore, these are not the right questions to ask. Instead, we should be asking more practical questions like "What do we know about the mean and the variance of a quantity of interest if all we have is ten simulations/experiments?", or "If we can afford ten simulations/experiments, which ones give us the greatest amount of information about the mean and the variance of a quantity of interest?" In this talk, I demonstrate how the Bayesian formalism can be used to lay the foundations of a unifying mathematical paradigm for posing such questions. Specifically, I show how we can capture the epistemic uncertainty induced by the small number of simulations/experiments and exploit it to construct task-specific information acquisition policies. To demonstrate the unifying nature of these ideas, I will present results from uncertainty propagation, model calibration, and optimization under uncertainty. The application examples I will consider are 1) the stochastic multi-objective optimization of an FEM-modeled wire-drawing process (selecting simulations) and 2) the optimization of high-quality graphene manufacturing using a plasma enhanced chemical vapor deposition reactor (selecting experiments).

**Biographical Sketch:** Dr. Ilias Bilonis is an Assistant Professor at the School of Mechanical Engineering, Purdue University leading the Predictive Science Laboratory. His research focuses on uncertainty quantification, design optimization under uncertainty, model calibration, and socio-technical modeling. Prior to his appointment at Purdue, he was a Postdoctoral Researcher at the Mathematics and Computer Science Division (MCS), Argonne National Laboratory. He received his Ph.D. in Applied Mathematics from Cornell University in 2013 and his Diploma in Applied Mathematics from the National Technical University of Athens in 2008.

For additional information, please contact Prof. Ying Li at (860) 486-7110, [yingli@engr.uconn.edu](mailto:yingli@engr.uconn.edu) or Laurie Hockla at (860) 486-2189, [hockla@engr.uconn.edu](mailto:hockla@engr.uconn.edu)