Applying the Sequential Monte Carlo Method of Particle Filtering with Dynamic Models: Theory, Implementation and Best Practices

While dynamic models growing use in health and social decision making, such models suffer from an important shortcoming: Given the inevitable limitations on model scope and the presence of stochastics, even the best such models give increasingly stale projections as time passes since their creation. By contrast, the application of the SMC methods of particle filtering (PF) to a dynamic model allows for recurrently regrounding estimates of the current system state in the model against observed data. This approach allows for a model that learns as new data becomes available, including by ongoing estimation of the evolving current state of the system. Such a model can also be used to support probabilistic projections, and probabilistic assessment of intervention tradeoffs from the current point -- considering all the evidence available to this point.

This tutorial -- taught by a practitioner who has led dozens of projects successfully applying PF with dynamic models -- will provide an introduction to the theory and of PF with dynamic models, as well as guidelines for providing these techniques effectively. Contents include discussion of the Bayesian updating underlying PF, the theory of sequential importance sampling that provides the basis for sampling from such distributions via particles, and the way in which this theory is captured in practice within model implementation. We further will emphasize the central role that model stochastics play within particle filtered models, and the essential need to balance and tune the magnitude of such stochastics so as to achieve suitable performance of the particle filter. Additional discussion of model tuning will focus on the role of the initial model state, the count of particles to be used, and the temporal granularity of the empirical data. An important element of the tutorial will discuss the critical role of the likelihood function, the distributional assumptions associated with the likelihood function, its use in integrating multiple lines of evidence. We will also note shortcomings & added needs associated with certain likelihood functions.

As time allows, we will discuss practical mechanisms that capitalize on the natural alignment between PF and streaming high-velocity data to create constantly updating dynamic models, and transforming model projects from a focus on delivering discrete products to a service offering ongoing value. We may also cover the additional power afforded by the use of Particle Markov Chain Monte Carlo techniques, and the textured challenges and opportunities afforded by use of PF with agent-based models.

Participants will be provided and experiment with a complete example of PF for an aggregate (System Dynamics/compartmental) model -- including visualizations that allow for understanding the detailed dynamics of that PF -- and detailed guidelines that allow this implementation to additionally serve as a template for other PF models that attendees may wish to build.

Expected Audience: Practitioners familiar with dynamic (simulation) modeling & comfortable sitting through some mathematical explanations using Bayesian probability & small Java code snippets.

Instructor: Nathaniel Osgood (osgood@cs.usask.ca) serves as Professor in the Department of Computer Science and Assoc. Faculty in the Department of Community Health & Epidemiology at U. Saskatchewan. His research is focused on advancing and cross-linking system science, data science, computational science and applied math to inform understanding and decision making in health and health care. Prior to joining the U of S faculty, he graduated from MIT with a PhD in Computer Science in 1999, and served as a Senior Lecturer in MIT’s School of Engineering.