1157-90-630 **Raghu Pasupathy*** (pasupath@purdue.edu), Dept. of Statistics, HAAS Hall, West Lafayette, IN 47907, and **Aaron Yip** and **David Newton**. *Adaptive Sequential SAA as a Universal Paradigm for Stochastic Optimization.*

The Stochastic Approximation (SA) recursion, also known as Stochastic Gradient Descent (SGD), is the "workhorse" recursion for solving stochastic optimization problems arising in machine learning and elsewhere. For smooth, strongly convex objectives, it is well-known that SA iterates achieve the information theoretic Cramer-Rao (CR) lower bound when the step size is chosen correctly. Such correct choice, however, depends on unknown curvature constants; moreover, the CR bound as a measure of performance ignores transient evolution and is thus not fully reflective of finite-time performance. This explains why SGD is rarely implemented in its native form. Instead, heuristic variations, e.g., ADAM, have come to be widely preferred (see, for example, TensorFlow). In this talk, I will argue for an alternative to SGD and its variants. Specifically, I will re-introduce a decades-old idea and rigorize it in a way that allows well-established deterministic optimization methods. The power of the framework is that it naturally retains the structure that is often inherent to sample-paths, and is expressly equipped to exploit deterministic optimization methods. I will demonstrate that the proposed framework retains the CR lower bound to within a factor in the convex smooth case. (Received February 04, 2020)