Sparse Pseudospectral Approximation: The Proper Use of Sparse Grids in Polynomial Chaos

Mathematical and Computational Tools for Predictive Simulation of Complex Coupled Systems Under Uncertainty

Paul G. Constantine, Michael S. Eldred, Eric T. Phipps
Sandia National Laboratories
PO Box 5800, Albuquerque, NM 87185

Abstract

Polynomial approximation (e.g., polynomial chaos and stochastic collocation methods) provides a useful alternative to random sampling (e.g., Monte Carlo and quasi Monte Carlo methods) for the forward propagation of uncertainty in physical models with stochastic inputs. The *non-intrusive* polynomial chaos method computes a polynomial approximation of a quantity of interest using only evaluations of the quantity at specified values the stochastic inputs; the function evaluations are used to numerically integrate the coefficients of a Fourier-type expansion in orthogonal polynomials. For quantities that vary smoothly as a function of the inputs, this global polynomial is known to converge rapidly as higher order basis functions are added to the representation. However, polynomial methods often suffer from the so-called *curse of dimensionality*, where the work required to compute the approximation – measured in the number of function evaluations – increases exponentially as the dimension of the inputs increases.

Sparse grid methods alleviate this curse by using specially chosen point sets to evaluate the quantity of interest. They are most effective for smooth functions with bounded mixed partial derivatives, promising comparable accuracy to standard tensor product approximation with many fewer function evaluations. Sparse grid integration methods are consequently used to approximate the coefficients of a the polynomial chaos expansion. Using an integration rule with n nodes, one expects to obtain accurate approximations of n coefficients; this is well established for tensor product type approximation. However, it has been observed in practice that the sparse grid integration rules with n nodes accurately computes far fewer than n coefficients; the coefficients associated with higher degree polynomials often contain an unacceptable errors.

In this work, we revisit Smolyak's original algorithm to construct a *sparse pseudospectral approximation*, which uses the point set of a sparse grid integration rule to accurately compute the coefficients of a multivariate polynomial approximation that naturally corresponds to the given point set. We observe that the error associated with the approximate coefficients decays at the rate comparable to the truncation error, which is optimal. We also prove that the sparse pseudospectral approximation is point-wise equivalent to the stochastic collocation response surface constructed on the given sparse grid.

We compare the sparse pseudospectral approximation method to the standard sparse grid integration method for the coefficients of the polynomial chaos expansion for a series of test problems. The compelling results suggest that the sparse pseudospectral approximation is the proper way to build a polynomial approximation from the evaluations of the quantity of interest at the points of the sparse grid.