Computers and Mathematics with Applications 71 (2016) 2449–2465

Contents lists available at ScienceDirect

Computers and Mathematics with Applications

journal homepage: www.elsevier.com/locate/camwa

A dynamically adaptive sparse grids method for quasi-optimal interpolation of multidimensional functions

Miroslav K. Stoyanov *, Clayton G. Webster

Department of Computational and Applied Mathematics, Oak Ridge National Laboratory, Oak Ridge, TN 37831, United States

ARTICLE INFO

Article history:
Available online 1 February 2016

Keywords:
Sparse grids
Quasi-optimal approximation
Multidimensional interpolation

ABSTRACT

In this work we develop a dynamically adaptive sparse grids (SG) method for quasi-optimal interpolation of multidimensional analytic functions defined over a product of one dimensional bounded domains. The goal of such approach is to construct an interpolant in space that corresponds to the "best M-terms" based on sharp a priori estimate of polynomial coefficients. In the past, SG methods have been successful in achieving this, with a traditional construction that relies on the solution to a Knapsack problem: only the most profitable hierarchical surpluses are added to the SG. However, this approach requires additional sharp estimates related to the size of the analytic region and the norm of the interpolation operator, i.e., the Lebesgue constant. Instead, we present an iterative SG procedure that adaptively refines an estimate of the region and accounts for the effects of the Lebesgue constant. Our approach does not require any a priori knowledge of the analyticity or operator norm, is easily generalized to both affine and non-affine analytic functions, and can be applied to sparse grids built from one dimensional rules with arbitrary growth of the number of nodes. In several numerical examples, we utilize our dynamically adaptive SG to interpolate quantities of interest related to the solutions of parametrized elliptic and hyperbolic PDEs, and compare the performance of our quasi-optimal interpolant to several alternative SG schemes.

1. Introduction

This paper considers constructing an approximation to multidimensional analytic functions defined over a product of one dimensional bounded domains. The main challenge facing all methods in this context is the *curse of dimensionality*, i.e., the computational complexity of approximation techniques increases exponentially with the number of dimensions. To alleviate the curse, methods have been proposed that reduce the dimensionality of the problem [1,2], reduce the complexity of the target function [3,4], or approximate the function in an optimal polynomial subspace [5–7]. We take the latter approach and we build upon the recent results in best M-terms approximation [7–9], where the function is projected onto the polynomial space associated with the dominant coefficients of either a Taylor or Legendre expansion. In implementation, finding the optimal space is intractable and instead sharp a priori estimates of the expansion coefficients are used to select a quasi-optimal space. Such approach can achieve sub-exponential convergence rate in the context of both projection, e.g., [10–12]

* Corresponding author.
E-mail address: stoyanovmk@ornl.gov (M.K. Stoyanov).

http://dx.doi.org/10.1016/j.camwa.2015.12.045
0898-1221/

and interpolation, e.g., [9,13], however, the quasi-optimal methods rely heavily on a priori estimates of the size of the region of analyticity of the function and sharp estimates are available only in few special cases.

Given a suitable polynomial space, orthogonal projection results in the best L^2 approximation, however, the projection approach often times comes at a heavy computational cost [10–12]. In contrast, sampling based techniques require only the values of the function at a set of nodes, e.g., Monte Carlo random sampling for computing the statistical moments of a function [14,15], and Sparse Grids (SG) method for high order polynomial approximation [5,6,16], which is the focus on this work. SG sampling does not result in best approximation in the associated polynomial space and the error is magnified by the norm of the SG operator, a.k.a., the Lebesgue constant. However, sampling tends to be computationally cheaper than projection as well as more susceptible to parallelization which usually offsets the moderate increase of the error. In addition, sampling procedures can wrap around simulation software that computes single realization of the function, which simplifies the implementation and allows the use of legacy and third party code.

Sparse grids algorithms construct multidimensional function approximation from a linear combination of tensors of one dimensional interpolation rules. Quasi-optimal SG are traditionally constructed as the solution to a Knapsack problem [17,13], where the selected set of tensors is associated with the largest profit index that is derived from an a priori estimate of the hierarchical surplus, the Lebesgue constant, and the number of samples in a tensor. In the case when the one dimensional rules grow by one node at a time, a near optimal greedy procedure using the Taylor coefficients of the function can construct a suitable approximation [18,9], however, without a priori assumptions, selecting the optimal set of coefficients comes at a very high computational cost.

In this work, we present an iterative procedure for constructing a sequence of SG interpolants with increasing number of nodes and accuracy, that does not require a priori estimates of the region of analyticity. We focus our attention to the nested SG case, where all nodes associated with one grid are also utilized by the next grid in the sequence, thus reusing all available samples. We review popular one dimensional nested rules such as Clenshaw–Curtis [19] and Leja [20,21] and we present several new rules based on greedy minimization of operator norms. In addition, for any chosen rule and any arbitrary lower (i.e., admissible [17]) polynomial space, we present a strategy for selecting the minimal set of tensors that yields an interpolant in that space. Every interpolant in the sequence is constructed using this strategy, which circumvents the Knapsack problem and allows us to restrict our attention to the selection of the optimal polynomial spaces.

The quasi-optimal polynomial space associated with Legendre coefficients is a total degree space with a small logarithmic correction [10,8]. However, while the Legendre space is optimal with respect to projection, in the context of interpolation, the quasi-optimal estimate does not account for the effect of the Lebesgue constant. Using estimates of the operator norm of the one dimensional rules, we add a strong correction to the total degree space to arrive at an estimate for the quasi-optimal interpolation space. Our estimate is parametrized by two vectors associated with the size of the analytic region of the function and the growth of the Lebesgue constant of the interpolation rules.

In order to keep our approach free from a priori assumptions, we present a procedure for dynamically estimating the two vector parameters. For each interpolant in the sequence, we consider the orthogonal decomposition of the interpolant into a linear combination of multivariate Legendre polynomials. Then, we seek the vectors that give the best fit of our quasi-optimal estimate to the decay rate of the Legendre coefficients, i.e., using least-squares approach. The polynomial space used for the construction of the next interpolant in the sequence is optimal with respect to the parameters inferred from the previous interpolant. The number of additional nodes in each interpolant can be chosen arbitrarily, however, few nodes result in more frequent update of the parameter vectors which leads to better accuracy, while larger number of nodes allows for greater parallelization.

The procedure for estimating the quasi-optimal polynomial space can be coupled with any approximation strategy that satisfies a mild assumption regarding the growth of the Lebesgue constant. One potential alternative is to use interpolation based on Fekete points, however, even in moderate dimensions, finding those points involves an ill-conditioned and prohibitively expensive problem. Other popular alternatives are the optimization based methods that construct an approximation based on minimization of ℓ^2 (e.g., least-squares [22]) or ℓ^1 (e.g., compressed sensing [23]) norms. Those methods can be applied to sets of random samples, however, the number of samples needed to construct the approximation always exceeds the cardinality of the optimal polynomial space. We assume that we can choose the abscissas for each samples and we want to exploit the fact that the range of an interpolation operator has exactly the same degrees of freedom as the number of interpolation nodes. Thus, the sparse grids interpolants are best suited for our context.

The rest of the paper is organized as follows, in Section 2 we derive an estimate of the quasi-optimal interpolation space and we present an iterative procedure for generating a sequence of quasi-optimal polynomial spaces. In Section 3, we present a strategy for constructing sparse grids operators with minimal number of nodes and we present several one dimensional interpolation rules. In Section 4 we present several numerical examples.

2. Quasi-optimal polynomial space

We consider the problem of approximating a multivariate function $f(\boldsymbol{y}) : \Gamma \to \mathbb{R}$, where $\Gamma \subset \mathbb{R}^d$ is a d-dimensional hypercube, i.e., $\Gamma = \bigotimes_{k=1}^{d} \Gamma_k$ and without loss of generality we let $\Gamma_k = [-1, 1]$. We assume that $f(\boldsymbol{y})$ admits holomorphic extension to a poly-ellipse in complex plane, i.e.,