Report No.
UCB/SEMM-2011/08

Structural Engineering
Mechanics and Materials

Convergence of an Efficient Local Least-Squares Fitting Method for Bases with Compact Support

By

Sanjay Govindjee, John Strain,
Toby J. Mitchell and Robert L. Taylor

June 2011

Department of Civil and Environmental Engineering
University of California, Berkeley

Convergence of an Efficient Local Least-Squares Fitting Method for Bases with Compact Support

Sanjay Govindjee[a,*], John Strain[b], Toby J. Mitchell[a], Robert L. Taylor[a]

[a]*Structural Engineering, Mechanics, and Materials*
Department of Civil and Environmental Engineering
University of California, Berkeley, CA 94720 USA
[b]*Department of Mathematics*
University of California, Berkeley, CA 94720 USA

Abstract

The least-squares projection procedure appears frequently in mathematics, science, and engineering. It possesses the well-known property that a least-squares approximation (formed via orthogonal projection) to a given data set provides an optimal fit in the chosen norm. The orthogonal projection of the data onto a finite basis is typically approached by the inversion of a Gram matrix involving the inner products of the basis functions. Even if the basis functions have compact support, so that the Gram matrix is sparse, its inverse will be dense. Thus computing the orthogonal projection is expensive.

An efficient local least-squares algorithm for non-orthogonal projection onto smooth piecewise-polynomial basis functions is analyzed. The algorithm runs in optimal time and delivers the same order of accuracy as the standard orthogonal projection. Numerical results indicate that in many computational situations, the new algorithm offers an effective alternative to global least-squares approximation.

Keywords: Least Squares, Isogeometric Analysis, Dirichlet Boundary Conditions

*Corresponding author
Email addresses: s_g@berkeley.edu (Sanjay Govindjee), strain@math.berkeley.edu (John Strain), tobymitchell@berkeley.edu (Toby J. Mitchell), rlt@ce.berkeley.edu (Robert L. Taylor)

1. Introduction

In a recent publication [1], an efficient approximate method for the imposition of Dirichlet boundary conditions in isogeometric analysis was proposed. Isogeometric analysis [2] uses non-interpolatory basis functions for both geometry and analysis, and adjusts "control point" values to impose Dirichlet boundary conditions. Standard methods for finding the control point values solve a least-squares (LSQ) problem [3]. This provides optimal accuracy but increases the computational cost, especially when dealing with transient Dirichlet boundary conditions, and in certain classes of non-linear problems. The local least-squares (LLSQ) method of [1] approximates the global least-squares problem by a collection of decoupled local least-squares problems, greatly diminishing the computational cost. Examples in [1] demonstrate the efficiency and accuracy of the method. However, the theoretical properties of the LLSQ method were not addressed and thus the true limitations of the method are unknown. In this work, we carefully examine the LLSQ method and elucidate its theoretical properties. In particular, we analyze the error in the LLSQ approximation, prove convergence to the global LSQ solution for model problems under appropriate conditions, and make positive statements about the overall impact of using LLSQ solutions in place of LSQ solutions.

In brief, the LLSQ method exploits the finite element concepts of local basis functions and global basis functions. When we speak of a global LSQ solution, we mean the projection of a given function onto the span of the global basis functions. This projection is an orthogonal projection in the L^2 norm which is computationally expensive. The LLSQ solution proceeds by first computing the orthogonal projection of the given function onto the span of the local basis functions. This inexpensive local computation yields a discontinuous approximation. The LLSQ solution then projects the discontinuous approximation non-orthogonally into the space spanned by the global basis functions, in a way that is not equivalent to the global LSQ projection. The primary theoretical question is to clearly ascertain the relation between the resulting approximations. We answer this question by factorizing the LSQ and LLSQ projections, identifying the common factors, and bounding the differences.

2. The Least-Squares Method

The least-squares method [4] is a well-established scheme for computing the orthogonal projection which best approximates given data from a finite dimensional subspace; see e.g. [5, §6.2-6.5] for an elementary introduction and [6, §5.3] for a more in-depth presentation and further references. In summary, consider a given vector-valued data function $u \in H^s$ with values $u(x)$ in the state space $\mathbb{R}^r$, for x in a subset Ω of Euclidean space $\mathbb{R}^d$, where the dimension d is 1, 2, or 3 and r is typically $O(1)$. Here H^s is the usual Sobolev space consisting of functions with distributional derivatives up to order s in L^2. Suppose one wishes to project this data onto a finite dimensional subspace $\mathcal{F}_n = \text{span}\{f_1(x), f_2(x), \cdots, f_n(x)\}$ of $L^2(\Omega)$. Here $f_i(x)$ are given linearly independent (but not usually orthogonal) scalar basis functions, and we seek a projection which is orthogonal in the standard L^2 inner product $\langle u, v\rangle = \int_\Omega u(x)^* v(x)\,\mathrm{d}x$. The global least-squares projection $u_F \in \mathcal{F}_n$ of the data u minimizes the squared residual norm

$$\|u - \sum_{i=1}^{n} u_i f_i\|^2 = \langle u - \sum_{i=1}^{n} u_i f_i, u - \sum_{i=1}^{n} u_i f_i\rangle \tag{1}$$

with respect to the parameters $u_i \in \mathbb{R}^r$. Employing an obvious abuse of notation, the nr-vector of minimizing components is also denoted $u_F \in \mathbb{R}^{nr}$ and it satisfies the well-known normal equations

$$G u_F = p\,, \tag{2}$$

where the $nr \times nr$ Gram matrix has $r \times r$ block matrix-valued components

$$G_{ij} = \langle f_i, f_j\rangle\, I \tag{3}$$

and the right hand side has vector-valued components

$$p_i = \langle f_i, u\rangle\,. \tag{4}$$

Here I is the $r \times r$ identity matrix.

Remarks:.

1. Standard solution methods for the normal equations (independent of their assembly) exploit the uncoupled nature of the r components of each vector coefficient u_i, but typically do not exploit sparsity of the Gram matrix. Thus they require $O(rn^3)$ floating point operations.

2. The numerical stability of the least-squares solution process in finite precision arithmetic requires the residual norm in (1) to be small (consistency of the projection) and the condition number of the Gram matrix to be controlled (stability, which is equivalent to uniform linear independence of the basis functions); see [6, §5.3].
3. The least-squares solution u_F is characterized by the orthogonality of its error: $\langle u - u_F, v\rangle = 0$ for any $v \in \mathcal{F}_n$. It can also be written as $u_F = Fu$ where F is the orthogonal projection operator from H^s onto $\mathcal{F}_n$. The operator F satisfies $F^2 = F$ because it is a projection, and F is selfadjoint, $F = F^*$, because it is an orthogonal projection.
4. In order to avoid cumbersome notation and terminology, from here on, we assume without loss of generality that the state space is one-dimensional; i.e. $r = 1$.

3. Compact Support

We now assume the functions $f_i(x)$ have compact support, so the Gram matrix elements $\langle f_i, f_j\rangle$ vanish except for $O(n)$ index pairs (i, j). This assumption holds for NURBS, B-spline, and standard Lagrange polynomial finite element basis functions, and so forth. A useful construct for bases with compact support is the notion of elements, as in finite element analysis. A typical finite element structure partitions the domain Ω into elements Ω_e. On each Ω_e, each global basis function f_j restricts to a local basis function f^e_m defined on Ω_e and vanishing outside Ω_e. In this setting, the rectangular assembly operator A is the Boolean matrix that maps between local element degrees of freedom v_{em}, representing any function $v = \sum_e \sum_m v_{em} f^e_m$, and global degrees of freedom v_j, where $v = \sum_j v_j f_j$. Each column of A contains a single unity-valued entry; all other entries are zero. Furthermore, due to our compact support assumption, the rows of A are sparse. This allows us to write

$$G = AG_EA^* \tag{5}$$
$$p = Ap_E\,, \tag{6}$$

where G_E is a block diagonal matrix whose blocks are computed via integration over individual elements, and likewise for p_E. In other words, G_E and p_E represent local (unassembled) element contributions to the global vectors and

matrices and the operator A assembles them. In this notation, the normal equations for the global LSQ problem become

$$AG_E A^* u_F = A p_E \,. \tag{7}$$

Form (7) of the normal equations is the point of departure for the approximate solution scheme proposed in [1].

4. The Local Least-Squares Method

The LLSQ method of [1] rewrites (7) in the form

$$G_E A^* u_F = p_E + q \,, \tag{8}$$

where $q \in \ker(A)$. Here q typically represents highly oscillatory data not present in $\mathcal{F}_n$, and measures the consistency of the given data u with the chosen function subspace $\mathcal{F}_n$. The LLSQ method ignores q and solves the relation $G_E A^* u_S = p_E$. This generates an approximation

$$u_S = (AA^*)^{-1} A G_E^{-1} p_E \tag{9}$$

to the LSQ solution u_F. Approximation (9) is extremely efficient to compute because AA^* is a diagonal matrix and the inversion of G_E is done element-by-element.

Concretely, the LLSQ method solves a local least-squares problem on each element with G_E^{-1} to obtain a globally discontinuous solution, $u_E = G_E^{-1} p_E$; it then smooths the discontinuous solution with the operator $(AA^*)^{-1}A$. The diagonal matrix AA^* counts the number of local degrees of freedom corresponding to each global degree of freedom; so $(AA^*)^{-1}A$ sets shared degrees of freedom to the average of the corresponding local degrees of freedom. The result is a function $u_S \in \mathcal{F}_n$, which enjoys the smoothness of the global basis. Note $u_S \neq u_F$ in general.

The geometric situation is depicted in Fig. 1. The given data is $u \in H^s = \{u \in L^2 \mid u^{(k)} \in L^2 \text{ for } 0 \leq k \leq s\}$. The span of the local functions f_m^e is $\mathcal{E}_n$, a linear space of discontinuous functions. The combination of the local least-square solutions on each element generates u_E, the orthogonal projection of u onto $\mathcal{E}_n$. The LLSQ method of [1] projects this function non-orthogonally onto $\mathcal{F}_n = \mathcal{E}_n \cap H^s$, the span of the global basis functions. The global LSQ solution u_F is the orthogonal projection of both u and u_E onto $\mathcal{F}_n$.

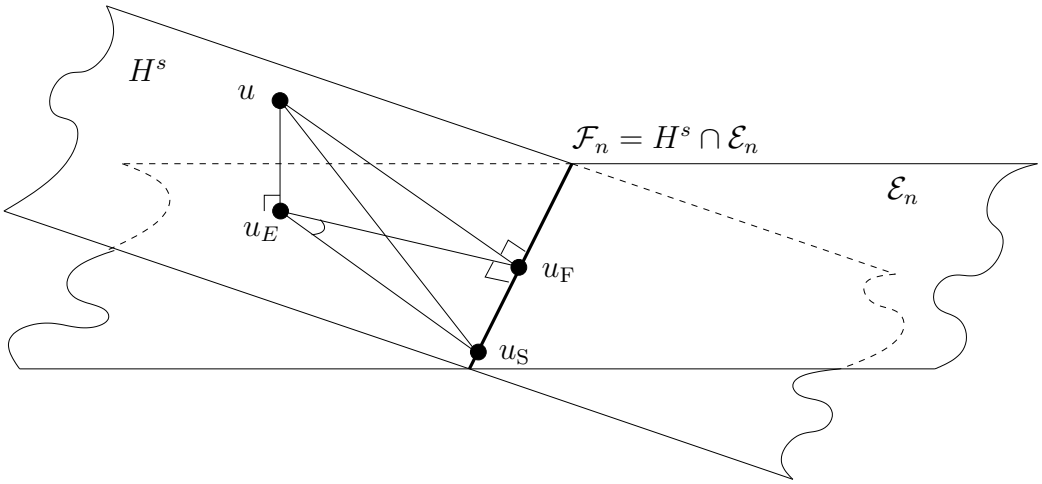

Figure 1: Geometry of orthogonal and skew projections u_F and u_S of $u \in H^s$ onto $\mathcal{F}_n$.

4.1. Illustration

To concretely illustrate the LLSQ method in a simple setting, consider given data $u(x)$ over an interval $\Omega = [0, L]$. We desire to approximate $u(x)$ by a function in $\mathcal{F}_3 = \mathrm{span}\{f_1(x), f_2(x), f_3(x)\}$, where $f_j(x)$ are known functions as depicted in Fig. 2. The classical global least square solution to this problem generates a function $u_F(x) \in \mathcal{F}_3$ which is the orthogonal projection of $u(x)$ onto $\mathcal{F}_3$ (in the L^2 norm). In this example, $u_F \in C^0$. The LLSQ method of [1] is a two step process. In step one, one first breaks the domain up into elements and orthogonally projects $u(x)$ onto $\mathcal{E}_3 = \mathrm{span}\{f_1^1(x), f_2^1(x), f_1^2(x), f_2^2(x)\}$, where $f_m^e(x)$, the local functions over element e, are the restrictions of the global functions $f_j(x)$; the mapping between the global index j and the local index m of element e is given by the geometry of the problem and the definition of the global functions. In finite element parlance, this relation is given by the assembly operator A. The elements of the basis for $\mathcal{E}_3$ are illustrated in Fig. 3. The result of this step is a function

$$u_E(x) = u_1^1 f_1^1(x) + u_1^2 f_1^2(x) + u_1^2 f_1^2(x) + u_2^2 f_2^2(x) \in \mathcal{E}_3.$$

In this example, $u_E \in C^{-1}$. Step two of the LLSQ method takes $u_E(x)$ and non-orthogonally projects it onto $\mathcal{F}_3$. This last step is non-orthogonal for reasons of efficiency – an orthogonal projection would simply generate $u_F(x)$ with all its attendant costs (should one consider a basis of non-trivial size). The final result is a function $u_S(x) \in \mathcal{F}_3$ which we call the skew projection

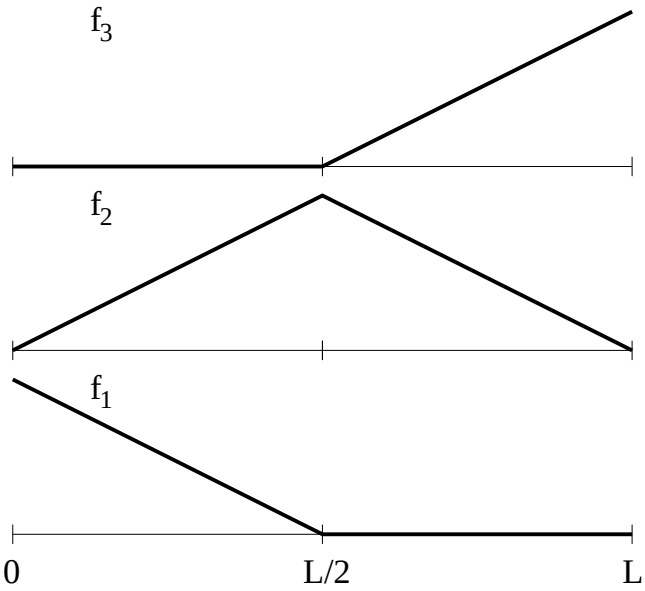

Figure 2: Example global basis with 3 functions $f_j(x)$.

or alternately the LLSQ solution. In our current example, this final step is executed by averaging the computed coefficients for the shared degrees of freedom $f_2^1(x)$ and $f_1^2(x)$ and assigning the result to the coefficient of $f_2(x)$:

$$u_S(x) = u_1^1 f_1(x) + \frac{1}{2}\left(u_1^2 + u_1^2\right) f_2(x) + u_2^2 f_3(x)\,.$$

5. General Analysis

The LLSQ method computes a "skew projection" $u_S \in \mathcal{F}_n$ by a two-step procedure: First, project u orthogonally onto $\mathcal{E}_n$ to get the best approximation u_E to u from $\mathcal{E}_n$. Second, project u_E non-orthogonally onto $\mathcal{F}_n$ by setting shared degrees of freedom controlling the first s derivatives of u_E to their average values. The second step yields a projection S from $\mathcal{E}_n$ to $\mathcal{F}_n$ defined by $u_S = Su_E$. Because repeating the second step makes no further change in u_S, we have $S^2 = S$. However, S is not orthogonal since shared degrees of freedom do not correspond to coefficients of a global orthonormal basis: thus $S^* \neq S$. Indeed, if S were the orthogonal projection from $\mathcal{E}_n$ onto $\mathcal{F}_n \subset \mathcal{E}_n$, then u_S would be exactly the global solution u_F. Figure 1

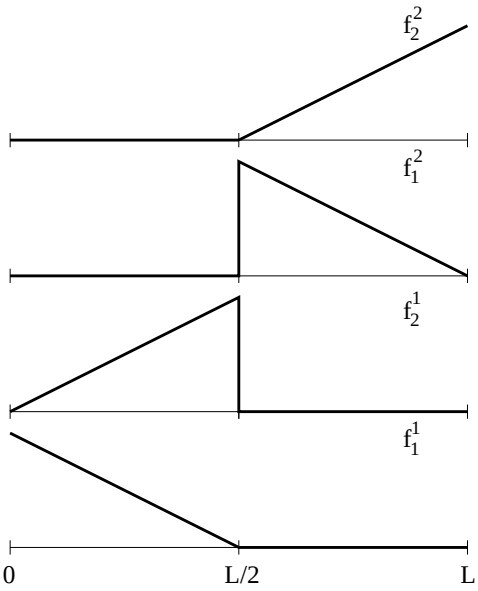

Figure 3: Four local basis functions corresponding to the global basis in Fig. 2.

exhibits the geometry of the three projections E, F and S and suggests three observations that simplify the analysis:

Factorization: The orthogonal projection F onto $\mathcal{F}_n$ could in principle be implemented by projecting first orthogonally onto $\mathcal{E}_n$ and then orthogonally from $\mathcal{E}_n$ to $\mathcal{F}_n$. Symbolically, $F = FE$. Hence the only difference between $u_F = Fu = FEu$ and $u_S = SEu$ is the route from $u_E \in \mathcal{E}_n$ to the final result in $\mathcal{F}_n$.

Projection: Both S and F are projections onto $\mathcal{F}_n$, so $Sv = Fv = v$ for all $v \in \mathcal{F}_n$.

Optimality: The error $\|u - u_E\|$ in orthogonal projection onto the space $\mathcal{E}_n$ of discontinuous functions is never larger than the error $\|u - u_F\|$ in orthogonal projection onto the space $\mathcal{F}_n$ of continuous functions, because $\mathcal{F}_n$ is a subspace of $\mathcal{E}_n$ and orthogonal projection finds the best approximation.

Applying these three observations in succession shows that the error $\|u - u_S\|$ in the LLSQ solution u_S is controlled by the error $\|u - u_F\|$ in the global least-squares solution u_F, multiplied by a bounded constant $C \geq 1$: By

factorization and projection,

$$\begin{aligned} u - u_S &= u - u_F + u_F - u_S \\ &= u - u_F + (F - S)u_E \\ &= u - u_F + (F - S)(u_E - u_F) \\ &= u - u_F + (F - S)(u_E - u + u - u_F). \end{aligned}$$

Applying norms of functions and operators, using the triangle inequality, and applying optimality gives

$$\begin{aligned} \|u - u_S\| &= \|u - u_F + (F - S)(u_E - u + u - u_F)\| \\ &\leq \|u - u_F\| + \|F - S\| \left(\|u_E - u\| + \|u - u_F\|\right) \\ &\leq (1 + 2\|F - S\|) \|u - u_F\| \\ &\leq C\|u - u_F\|. \end{aligned}$$

Thus $C = 1+2\|F-S\|$ is proportional to the operator norm of the difference between orthogonal and skew projections from $\mathcal{E}_n$ onto $\mathcal{F}_n$. If C is bounded independent of the element size, this result shows that the LLSQ method delivers the same order of accuracy as the global least-squares solution with considerably less computational effort.

6. Examples

The two step recipe elaborated upon in the prior section defines the skew projection procedure of the LLSQ method. To understand its properties, and in particular compute the error constant $C = 1 + 2\|F - S\|$, we will look at two particular cases. We will start with Lagrange polynomial finite elements of variable polynomial order and smoothness and examine the behavior of the LLSQ method as well as one variant that we will define later. Next we will look at the case of B-spline bases of variable polynomial order. Our tool of choice, in both cases, will be von Neumann analysis.

6.1. Lagrange Polynomial Finite Elements

We prove the order-$p+1$ convergence of the LLSQ method for projection of smooth data u from the order-s Sobolev space

$$H^s = \{u \in L^2 | u^{(k)} \in L^2 \quad \text{for} \quad 0 \leq k \leq s\}$$

onto standard degree-p H^s polynomial finite elements in dimension $d = 1$. It is known that the global least-squares approximation u_F enjoys order-$p+1$ convergence for $p + 1 \le s$:

$$\|u - u_F\| \le Kh^{p+1}\|u\|_{H^s} = Kh^{p+1}\sqrt{\sum_{k=0}^{s} \|u^{(k)}\|^2}\,,$$

where $\|v\|$ is the L^2 norm of v and K is independent of the element size h. We will show that the LLSQ approximation maintains the same order of convergence. For simplicity, we assume state space dimension $r = 1$, and elements with uniform size $2h$, degree p and smoothness s, and employ classical von Neumann analysis. The structure and conclusions of the analysis carry over to related basis functions in multidimensional geometry, but more sophisticated tools are required.

In this setting, the domain $\Omega = \mathbb{R}$ is the whole real line and the elements $\Omega_j = [x_j - h, x_j + h]$ are intervals of equal size $2h$ centered about points $x_j = 2jh$ for all integers $j \in \mathbb{Z}$. Each element supports polynomials of degree $p \ge 1$, expressed for convenience in a local orthonormal basis

$$p_{jm}(x) = \frac{1}{\sqrt{h}} P_m\left(\frac{x - x_j}{h}\right) \qquad |x - x_j| \le h, \qquad 0 \le m \le p. \tag{10}$$

Here we define $p_{jm}(x) = 0$ for $|x - x_j| > h$, and P_m are orthonormal Legendre polynomials on the interval $[-1, 1]$:

$$P_0(x) = \frac{1}{\sqrt{2}}, \qquad P_1(x) = \frac{\sqrt{3}}{\sqrt{2}}x,$$

$$P_{m+1}(x) = \frac{2m+1}{m+1}\sqrt{\frac{2m+3}{2m+1}}xP_m(x) - \frac{m}{m+1}\sqrt{\frac{2m+3}{2m-1}}P_{m-1}(x).$$

Let $\mathcal{E}_n = \text{span}\{p_{jm}\}$ be the span of these basis functions, so $\mathcal{E}_n$ is the space of piecewise degree-p polynomials over the elements Ω_j. Since $p_{jm}(x)p_{j'm'}(x) = 0$ for $j \ne j'$ and $\langle P_m, P_{m'}\rangle = 0$ in $L^2(-1, 1)$ for $m \ne m'$, the p_{jm}'s form an orthonormal basis for their span. Hence any $v \in \mathcal{E}_n$ has the form

$$v(x) = \sum_{j\in\mathbb{Z}} \sum_{m=0}^{p} v_{jm}p_{jm}(x)\,, \tag{11}$$

where each coefficient v_{jm} is an inner product

$$v_{jm} = \langle v, p_{jm} \rangle = \int_{x_j - h}^{x_j + h} v(x) p_{jm}(x) \, \mathrm{d}x$$

and for any given element Ω_J, only the $p+1$ terms in Eq. (11) with $j = J$ are nonzero. The coefficients give an isometry since

$$\|v\|^2 = \int_{-\infty}^{\infty} |v(x)|^2 \, \mathrm{d}x = \sum_{j \in \mathbb{Z}} \sum_{m=0}^{p} |v_{jm}|^2.$$

Functions in $\mathcal{E}_n$ are typically discontinuous, because no interelement continuity or smoothness conditions have been imposed on the basis functions (10). In compensation, the process of computing the orthogonal projection u_E onto $\mathcal{E}_n$ of an arbitrary function $u \in H^s$ is completely local:

$$u_E(x) = \sum_{j \in \mathbb{Z}} \sum_{m=0}^{p} u_{jm} p_{jm}(x) \, ,$$

where

$$u_{jm} = \langle u, p_{jm} \rangle = \int_{x_j - h}^{x_j + h} u(x) p_{jm}(x) \, \mathrm{d}x.$$

Because the p_{jm}'s form an orthonormal basis for $\mathcal{E}_n$, $u_E = Eu$ is the closest function in $\mathcal{E}_n$ to the data u, and defines an orthogonal projection operator: $E = E^2 = E^*$.

The global least-squares problem of Eq. (1) projects u orthogonally onto the subspace $\mathcal{F}_n = \mathcal{E}_n \cap H^s = \mathcal{E}_n \cap C^{s-1}$ of $\mathcal{E}_n$ consisting of piecewise degree-p polynomials v which satisfy s additional interelement smoothness conditions

$$\lim_{\epsilon \downarrow 0} v^{(k)}(x_j + h - \epsilon) = v^{(k)}(x_j + h_-) = v^{(k)}(x_{j+1} - h_+) = \lim_{\epsilon \downarrow 0} v^{(k)}(x_{j+1} - h + \epsilon)$$

for derivatives of order $0 \le k \le s-1$. The result of this projection is the closest function $v = u_F = Fu \in \mathcal{F}_n$ to the data u, and defines an orthogonal projection operator: $F = F^2 = F^*$.

In this concrete setting, the skew projection S which implements the LLSQ method takes the discontinuous degree-p piecewise polynomial

$$u_E(x) = \sum_{j \in \mathbb{Z}} \sum_{m=0}^{p} u_{jm} p_{jm}(x) \tag{12}$$

and maps the coefficients u_{jm} to values v_{jm} such that

$$u_S(x) = \sum_{j\in\mathbb{Z}} \sum_{m=0}^{p} v_{jm} p_{jm}(x) = Su(x)$$

satisfies the conditions

$$u_S^{(k)}(x_j + h_-) = u_S^{(k)}(x_{j+1} - h_+), \qquad 0 \le k \le s-1, \qquad j \in \mathbb{Z},$$

on jumps across each interelement interface $x_j + h_- = x_{j+1} - h_+$. Since only $p+1$ terms in the sum defining u_S are nonzero for each x, these conditions can be written

$$\sum_{m=0}^{p} P_m^{(k)}(1) v_{jm} = \sum_{m=0}^{p} P_m^{(k)}(-1) v_{j+1,m}, \qquad 0 \le k \le s-1, \qquad j \in \mathbb{Z},$$

where we have scaled out a common power of h. Since Legendre polynomials are alternately even and odd, $P_m^{(k)}(-1) = (-1)^{m+k} r_{km}$ where $r_{km} = P_m^{(k)}(1)$. Thus these conditions simplify to

$$Rv_j - Lv_{j+1} = 0, \qquad j \in \mathbb{Z},$$

where each v_j is an $(p+1)$-vector $(v_{j0}, v_{j1}, \ldots, v_{jp})^T \in \mathbb{R}^{p+1}$ and R and L are $s \times (p+1)$ matrices with elements r_{km} and $(-1)^{m+k} r_{km}$ respectively. Row k of R defines a shared degree of freedom controlling the kth derivative of u_S at the right interval endpoint, with the rows of L playing the same roles at the left interval endpoint. Note that R and L are independent of the mesh size h.

Specifically, the LLSQ method sets the shared values for the skew projection v to their averages from $u_E(x)$:

$$Rv_j = Lv_{j+1} = \frac{1}{2}\left(Ru_j + Lu_{j+1}\right)$$

for each $j \in \mathbb{Z}$; here as with v_j, $u_j \in \mathbb{R}^{p+1}$. Shifting the index $j+1$ gives an infinite block tridiagonal system consisting of two linear systems for each coefficient vector v_j:

$$Rv_j = \frac{1}{2}\left(Ru_j + Lu_{j+1}\right)$$

and

$$Lv_j = \frac{1}{2}\left(Ru_{j-1} + Lu_j\right) .$$

Thus each v_j satisfies a block linear system,

$$\begin{bmatrix} R \\ L \end{bmatrix} v_j = \frac{1}{2}\begin{bmatrix} 0 \\ R \end{bmatrix} u_{j-1} + \frac{1}{2}\begin{bmatrix} R \\ L \end{bmatrix} u_j + \frac{1}{2}\begin{bmatrix} L \\ 0 \end{bmatrix} u_{j+1} \tag{13}$$

for $j \in \mathbb{Z}$. Block j of the system consists of $2s$ equations for the $p+1$ coefficients in v_j, so we expect a solution if $2s \leq p+1$. (The left-hand block matrix has full row rank because Hermite interpolation of s values and derivatives at interval endpoints by a degree-p polynomial has at least one solution for $2s \leq p+1$.) Equality $2s = p+1 = 2$ holds for continuous (H^1) piecewise-linear basis functions, but we cannot expect a solution for H^3 cubic polynomials with $2s = 6 > p+1 = 4$. Instead, quintic polynomials are required to achieve H^3 or C^2 smoothness. The LLSQ method sets shared degrees of freedom to specified average values rather than simply matching them between elements, so it requires about twice the polynomial degree to achieve a specified level of smoothness vis-a-vis the global least-squares computation.

Classical von Neumann analysis applies Fourier series analysis to block-diagonalize Eq. (13). For any sequence of vector coefficients $u_j \in \mathbb{R}^{p+1}$, we consider the associated Fourier series

$$\hat{u}(\theta) = \frac{1}{\sqrt{2\pi}} \sum_{j\in\mathbb{Z}} u_j \mathrm{e}^{\mathrm{i}j\theta}$$

where $\mathrm{i} = \sqrt{-1}$ and the Fourier coefficients $u_j \in \mathbb{R}^{p+1}$ are given by the standard orthogonality formula

$$u_j = \frac{1}{\sqrt{2\pi}} \int_{-\pi}^{\pi} \mathrm{e}^{-\mathrm{i}j\theta} \hat{u}(\theta)\, \mathrm{d}\theta.$$

Because both orthogonal Legendre coefficients and the Fourier transform are isometries of the L^2 inner product, the L^2 norm of our original piecewise polynomial $u_E \in \mathcal{E}_n$ is the same as the discrete l^2 norm of its vector coefficients and the $L^2(-\pi,\pi)$ norm of $\hat{u}$:

$$\int_{-\infty}^{\infty} |u_E(x)|^2 \mathrm{d}x = \sum_{j\in\mathbb{Z}} \|u_j\|^2 = \int_{-\pi}^{\pi} \|\hat{u}(\theta)\|^2 \,\mathrm{d}\theta\,,$$

where $\|u\|$ is the standard Euclidean norm of $u \in \mathbb{R}^{p+1}$. Similarly, inner product relationships such as adjoints and norms of operators are preserved.

The shifted sequence u_{j+1} corresponds to multiplying the Fourier series corresponding to u_j by $\mathrm{e}^{-\mathrm{i}\theta}$, so the Fourier transform of Eq. (13) yields

$$\begin{aligned}\begin{bmatrix} R \\ L \end{bmatrix} \hat{v}(\theta) &= \left(\frac{1}{2} \begin{bmatrix} 0 \\ R \end{bmatrix} \mathrm{e}^{\mathrm{i}\theta} + \frac{1}{2} \begin{bmatrix} R \\ L \end{bmatrix} + \frac{1}{2} \begin{bmatrix} L \\ 0 \end{bmatrix} \mathrm{e}^{-\mathrm{i}\theta} \right) \hat{u}(\theta) \\ &= \frac{1}{2} \begin{bmatrix} R + \mathrm{e}^{-\mathrm{i}\theta} L \\ L + \mathrm{e}^{\mathrm{i}\theta} R \end{bmatrix} \hat{u}(\theta).\end{aligned}$$

6.1.1. LLSQ skew projection case

If the block matrix on the left is square and invertible, then the skew projection S is uniquely determined by these equations and

$$\hat{v}(\theta) = \frac{1}{2} \begin{bmatrix} R \\ L \end{bmatrix}^{-1} \begin{bmatrix} R + \mathrm{e}^{-\mathrm{i}\theta} L \\ L + \mathrm{e}^{\mathrm{i}\theta} R \end{bmatrix} \hat{u}(\theta) = \hat{S}(\theta)\hat{u}(\theta) = \widehat{Su}(\theta)$$

where $\hat{S}$ is the $(p+1) \times (p+1)$ matrix symbol of S. If $2s < p+1$ then S is underdetermined and several possible skew projections exist.

Among these a natural choice (in the context of orthogonality) involves the pseudoinverse or right inverse defined by $B^\dagger = B^*(BB^*)^{-1}$ for a matrix B with full row rank:

$$\hat{S}(\theta) = \frac{1}{2} \begin{bmatrix} R \\ L \end{bmatrix}^\dagger \begin{bmatrix} R + \mathrm{e}^{-\mathrm{i}\theta} L \\ L + \mathrm{e}^{\mathrm{i}\theta} R \end{bmatrix}.$$

Back in real space, this choice defines S by

$$Su_j = \begin{bmatrix} R \\ L \end{bmatrix}^\dagger \left(\frac{1}{2} \begin{bmatrix} 0 \\ R \end{bmatrix} u_{j-1} + \frac{1}{2} \begin{bmatrix} R \\ L \end{bmatrix} u_j + \frac{1}{2} \begin{bmatrix} L \\ 0 \end{bmatrix} u_{j+1} \right). \tag{14}$$

6.1.2. Alternate skew projection

Another natural choice would be to set a subvector of v_j equal to the corresponding entries of u_j, and use the remaining degrees of freedom in v_j to satisfy the smoothness conditions. Setting the low-order entries of v_j to those of u_j and satisfying smoothness conditions with the high-order entries gives another skew projection T with symbol

$$\hat{T}(\theta) = \begin{bmatrix} R \\ J \\ L \end{bmatrix}^{-1} \begin{bmatrix} \frac{1}{2}(R + \mathrm{e}^{-\mathrm{i}\theta} L) \\ J \\ \frac{1}{2}(L + \mathrm{e}^{\mathrm{i}\theta} R) \end{bmatrix},$$

where J is the $(p+1-2s) \times (p+1)$ matrix consisting of the top $p+1-2s$ rows of the $(p+1) \times (p+1)$ identity matrix.

6.1.3. Explicit Computations

For continuous (H^1) piecewise-linear polynomials with $s = p = 1$, the skew projection S is uniquely determined by Eq. (13). Using the values

$$r_{00} = P_0(1) = \frac{1}{\sqrt{2}}, \qquad r_{01} = P_1(1) = \frac{\sqrt{3}}{\sqrt{2}},$$

gives

$$\begin{bmatrix} R \\ L \end{bmatrix} = \begin{bmatrix} r_{00} & r_{01} \\ r_{00} & -r_{01} \end{bmatrix} = \begin{bmatrix} P_0(1) & P_1(1) \\ P_0(-1) & P_1(-1) \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & \sqrt{3} \\ 1 & -\sqrt{3} \end{bmatrix}$$

and

$$\hat{S}(\theta) = \frac{1}{2\sqrt{3}} \begin{bmatrix} \sqrt{3}(1+\cos\theta) & 3\mathrm{i}\sin\theta \\ -\mathrm{i}\sin\theta & \sqrt{3}(1-\cos\theta) \end{bmatrix}.$$

Clearly $\hat{S}(\theta)^2 = \hat{S}(\theta)$, and $\hat{S}(\theta)^* \neq \hat{S}(\theta)$. Hence S is a non-orthogonal projection. Since $\det \hat{S}(\theta) = 0$ and trace $\hat{S}(\theta) = 1$, its eigenvalues are 1 and 0, with 1 corresponding to the invariant subspace $\mathcal{F}_n$. In Fourier space, it is straightforward to compute the orthogonal projection F from $\mathcal{E}_n$ onto the range of S (which is $\mathcal{F}_n$). Indeed, $F = S(S^*S)^{-1}S^* = F^2 = F^*$ is the orthogonal projection operator with the same range as S. Since the Fourier transform preserves inner products and therefore adjoints, F is represented by the matrix symbol $\hat{F} = \hat{S}(\hat{S}^*\hat{S})^{-1}\hat{S}^*$. A brief calculation shows that

$$\hat{F} = \frac{1}{2(2+\cos\theta)} \begin{bmatrix} 3(1+\cos\theta) & \sqrt{3}\mathrm{i}\sin\theta \\ -\sqrt{3}\mathrm{i}\sin\theta & 1-\cos\theta \end{bmatrix}.$$

Clearly $\hat{F}(\theta)^2 = \hat{F}(\theta)$, $\det \hat{F}(\theta) = 0$, trace $\hat{F}(\theta) = 1$, and $\hat{F}(\theta)^* = \hat{F}(\theta)$. Hence F is an orthogonal projection. Since $\hat{F}$ is a rational function rather than a trigonometric polynomial, it is the symbol of a nonlocal operator F which will be expensive to apply. We see that the error constant $C = 1 + 2\|F - S\|$ is independent of the element size. A tedious calculation shows that

$$(\hat{S} - \hat{F})(\hat{S} - \hat{F})^* = \frac{\sin^2\theta}{6(2+\cos\theta)} \begin{bmatrix} 3(1+\cos\theta) & \sqrt{3}\mathrm{i}\sin\theta \\ -\sqrt{3}\mathrm{i}\sin\theta & 1-\cos\theta \end{bmatrix} = \frac{\sin^2(\theta)}{3}\hat{F}$$

has eigenvalues 0 and $\sin^2(\theta)/3$, so $\|\hat{S} - \hat{F}\| = |\sin\theta|/\sqrt{3} \leq 1/\sqrt{3}$ and

$$C = 1 + 2\|F - S\| = 1 + 2\max_{|\theta|\leq\pi} \|\hat{F}(\theta) - \hat{S}(\theta)\|_2 = 1 + 2/\sqrt{3} \leq 3.$$

Table 1: Error bounds $\lceil C \rceil$ for skew projection S onto piecewise polynomials of degrees p and smoothness H^s.

$s \backslash p$	1	3	5	7	9	11	13	15
1	3	4	4	4	4	4	4	4
2		10	7	6	6	6	6	6
3			56	33	27	25	23	22
4				383	199	153	132	121
5					2844	1318	953	791
6						21869	9176	6250
7							171924	65915
8								1372945

Since this is independent of the element size h, the LLSQ projection S delivers the same order of accuracy as the global least-squares projection F.

It may be worth noting that the LLSQ method is related to some recent investigations [7] into banded matrices with banded inverses. Indeed, the LSQ method solves a banded system with a non-banded inverse, while the LLSQ method approximates the non-banded inverse by a banded matrix, with an accuracy independent of the matrix size.

The general case of arbitrary polynomial order p and smoothness s with $2s \leq p + 1$ proceeds similarly. Numerically computed upper bounds for the constant C, covering polynomial degrees $p = 1, 3, 5, \ldots, 15$ and smoothness $0 \leq s \leq \lfloor (p + 1)/2 \rfloor$ are reported in Tables 1 and 2 for the skew projections S and T based on pseudoinversion and high-order modification, respectively. The error constants are independent of the element size h, and remain small as long as the level of smoothness s remains well below its maximum. Along each row of values with fixed smoothness s and increasing polynomial degree p, the error constants rapidly approach an asymptote of moderate value. Along table diagonals where a fixed percentage of the available degrees of freedom are devoted to smoothness, the error constants grow factorially, but the rapid convergence of approximations constructed from high-order smooth basis functions provides considerable compensation. The skew projection S based on pseudoinversion is slightly better than the alternative T, as expected.

Table 2: Error bounds $\lceil C \rceil$ for skew projection T onto piecewise polynomials of degrees p and smoothness H^s.

$s\backslash p$	1	3	5	7	9	11	13	15
1	3	4	4	5	5	5	5	6
2		10	8	10	13	16	20	24
3			56	40	51	68	92	122
4				383	243	305	424	599
5					2844	1621	1962	2737
6						21869	11311	13102
7							171924	81359
8								1372945

6.2. B-Spline Basis

As a second example, we now consider the case where we wish to project onto the space of functions spanned by a B-spline basis defined on a uniform knot vector. Our approach will be to construct a map between the B-spline basis and our Legendre basis and then to leverage our prior analysis. The primary difference in the B-spline setting is that the averaging operation of the LLSQ method can no longer be interpreted as a simple setting of inter-element jump values. Instead, it corresponds to an averaging operator with a stencil extending over $p+1$ elements.

We begin by noting that the skew projection can be written as:

$$w(x) = u_S(x) = \sum_{j \in \mathbb{Z}} w_j b_p(x - 2jh) = Su(x)$$

but is formally computed by first creating the discontinuous piecewise polynomial given by Eq. (12) and then projecting onto the H^p piecewise degree-p B-spline basis $\{b_p(x-2jh) \mid j \in \mathbb{Z}\}$. This basis is defined by

$$b_p(x) = \frac{p+1}{p} \sum_{i=0}^{p+1} \gamma_{ip} (x - t_i)_+^p,$$

where $t_i = (2i-1)h$,

$$\gamma_{ip} = \prod_{\substack{0 \le j \le p+1 \\ j \neq i}} \frac{1}{t_j - t_i},$$

and $(t)_+ = \max(t, 0)$. It consists of translates of a single function $b_p(x)$, which is a degree-p polynomial on each interval $t_j < x < t_{j+1}$ and vanishes identically for $x < t_0$ and $x > t_{p+1}$.

Since the piecewise Legendre polynomials p_{jm} form an orthonormal basis for piecewise degree-p polynomials, the coefficients b_{jm} such that

$$b_p(x) = \sum_{m=0}^{p} b_{jm} p_{jm}(x) \qquad \text{for} \quad t_j \le x \le t_{j+1}$$

must be given by the formula

$$b_{jm} = \int_{t_j}^{t_{j+1}} p_{jm}(x) b_p(x)\,\mathrm{d}x = \sqrt{h} \int_{-1}^{1} P_m(\zeta) b_p((2j+\zeta)h)\,\mathrm{d}\zeta.$$

Since p_{jm} and b_p are polynomials of degree $\le p$ on the interval of integration, the q-point Gauss-Legendre integration formula

$$\int_{-1}^{1} f(\zeta)\,\mathrm{d}\zeta = \sum_{k=1}^{q} \alpha_k f(\zeta_k)$$

(where α_k and ζ_k are tabulated in [8]) is exact whenever $q \ge p+1$. Hence

$$b_{jm} = \sqrt{h} \sum_{k=1}^{q} \alpha_k P_m(\zeta_k) b_p((2j+\zeta_k)h).$$

Using these coefficients and the characteristic function $\chi_j(x)$ which is 1 for $t_j \le x \le t_{j+1}$ and 0 otherwise, we have

$$b_p(x) = \sum_{j=0}^{p} \sum_{m=0}^{p} b_{jm} p_{jm}(x) \chi_j(x),$$

and inverting the relationship expresses the Legendre polynomials in terms of the b_p restricted to each interval of support:

$$p_{jm}(x) = \sum_{k=0}^{p} \beta_{mk} b_p(x + 2(k-j)h) \chi_j(x)$$

where $\beta = \{\beta_{kj}\}$ is the $(p+1)\times(p+1)$ inverse matrix to $\{b_{jk}\}$ and $|x-x_j| \le h$. Thus a discontinuous polynomial $u_E(t)$ is transformed to the segmented B-spline basis by

$$\begin{aligned} u_E(x) &= \sum_{j\in\mathbb{Z}}\sum_{m=0}^{p} u_{jm}p_{jm}(x) \\ &= \sum_{j\in\mathbb{Z}}\sum_{m=0}^{p} u_{jm}\sum_{k=0}^{p}\beta_{mk}b_p(x-x_j+x_k)\chi_j(x) \\ &= \sum_{j\in\mathbb{Z}}\sum_{k=0}^{p}\left(\sum_{m=0}^{p}\beta_{mk}u_{j+k,m}\chi_{j+k}(x)\right)b_p(x-x_j) \\ &= \sum_{j\in\mathbb{Z}}\left(\sum_{k=0}^{p} w_{jk}\chi_{j+k}(x)\right)b_p(x-x_j). \end{aligned}$$

The segmented B-spline basis functions are $\chi_{k+j}(x)b_p(x-x_j)$ for $j\in\mathbb{Z}$ and $0\le k\le p$, and the coefficients of u_E in this basis are

$$w_{jk} = \sum_{m=0}^{p}\beta_{mk}u_{j+k,m}.$$

The LLSQ method proceeds to average and spread the coefficients over the support of each smooth unsegmented basis function $b_p(x-x_j)$:

$$w_{jk}\chi_{j+k}(x) \to w_j = \left(\frac{1}{p+1}\sum_{k=0}^{p} w_{jk}\right)\left(\sum_{k=0}^{p}\chi_{j+k}(x)\right).$$

Given the averaged coefficients, the skew projection $w = Su$ can be reassembled in the piecewise Legendre basis by

$$\begin{aligned} w(x) &= \sum_{j\in\mathbb{Z}} w_j b_p(x-x_j) \\ &= \sum_{j\in\mathbb{Z}} w_j \sum_{k=0}^{p}\sum_{m=0}^{p} b_{km}p_{km}(x-x_j) \\ &= \sum_{j\in\mathbb{Z}}\sum_{m=0}^{p}\left(\sum_{k=0}^{p} b_{km}w_{j-k}\right)p_{jm}(x) \\ &= \sum_{j\in\mathbb{Z}}\sum_{m=0}^{p}\left(\sum_{k=0}^{p} b_{km}\frac{1}{p+1}\sum_{l=0}^{p}\sum_{n=0}^{p}\beta_{nl}u_{j-k+l,n}\right)p_{jm}(x). \end{aligned}$$

Here we have interchanged summation over j, shifted the infinite sum by k, and interchanged the sums again. Fourier analysis of this last result yields the matrix symbol $\hat{S}(\theta)$ of the skew projection such that

$$\widehat{Su}(\theta)_m = \sum_{n=0}^{p} \hat{S}(\theta)_{mn} \hat{u}(\theta)_n$$

where

$$\begin{aligned} \hat{S}(\theta)_{mn} &= \sum_{k=0}^{p} \sum_{l=0}^{p} b_{km} \frac{1}{p+1} \beta_{nl} \mathrm{e}^{\mathrm{i}(k-l)\theta} \\ &= \frac{1}{p+1} \left(\sum_{k=0}^{p} b_{km} \mathrm{e}^{\mathrm{i}k\theta} \right) \left(\sum_{l=0}^{p} \beta_{nl} \mathrm{e}^{-\mathrm{i}l\theta} \right). \end{aligned}$$

The separated form of $\hat{S}$ is a natural consequence of the low-rank structure of averaging and spreading.

6.2.1. Numerical Computations

In the first row of Table 3 we report the error constants $C = 1 + 2\|S - F\|$ for B-spline polynomial degrees 1 through 6. These are computed from the norms of the operators derived in the previous section. The values are larger than the error constants of Table 2 in compensation for the lower degree of the B-spline basis functions, but still remain moderate in size and independent of the element size. Thus the local skew projection provides the same order of accuracy as the global projection onto the space of B-splines. The error constants are relatively tight. To illustrate this point, in the second row of Table 3 we show computed values of C for specially chosen data functions $u(x)$. These functions were determined by the right singular vector corresponding to the maximum singular value of $\hat{S}(\theta) - \hat{F}(\theta)$ for each value of p. Fig. 4 shows a graph of one such function used for the case $p = 3$ on the interval $[0, 10]$; the knots t_j are uniformly spaced 1 unit apart.

As a second numerical B-spline demonstration, we consider the smooth data function $u(x) = \cos(6\pi x/L)$ over the interval $[0, L]$ and project it onto the B-spline basis for $p = 1, 2, 3, 4$. Figure 5 shows the errors $\|u - u_F\|$ and $\|u - u_S\|$ versus the reciprocal of the number of knot spacings (i.e. h). For each value of p, one observes that the rate of convergence for the skew projection (solid line) and the full least square projection (dashed line) is

Table 3: (First row) Error bounds $\lceil C \rceil$ for skew projection S onto B-splines of degree p. (Second row) Computed values.

$p = 1$	2	3	4	5	6
3	10	67	854	12205	253587
1.3	8.2	50	657	8113	166991

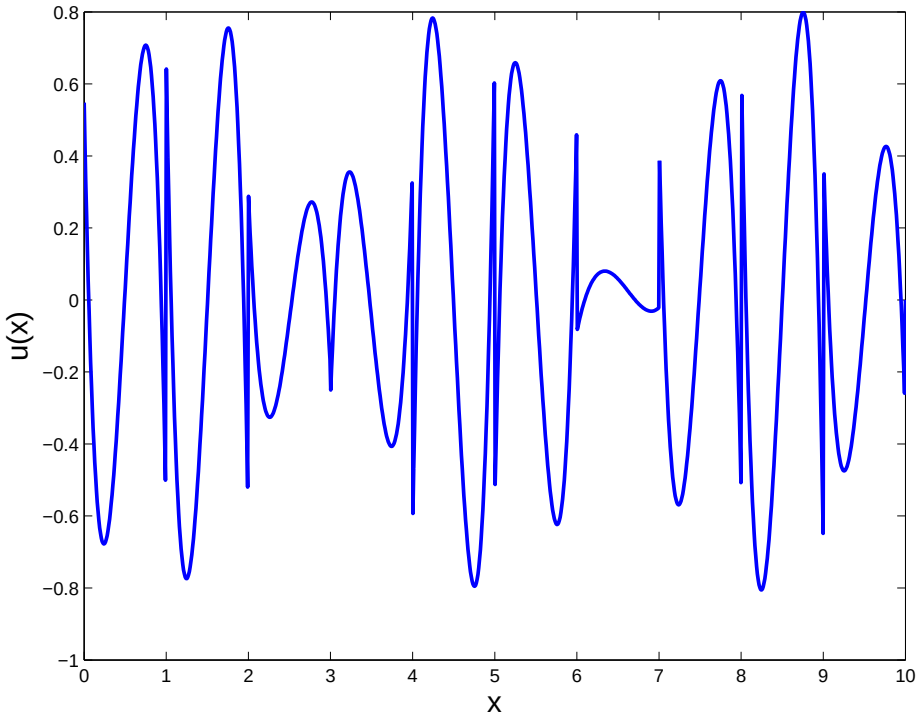

Figure 4: Example data function $u(x)$ for the case $p = 3$ which results in a near maximal value for C.

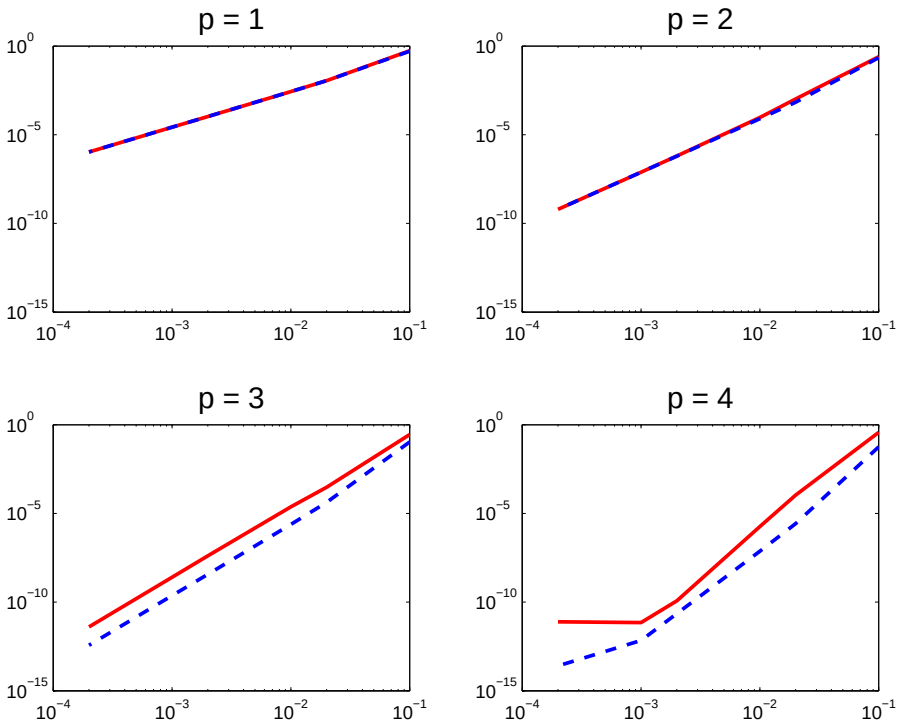

Figure 5: Convergence curves for $\|u - u_F\|$ (dashed curves) and $\|u - u_S\|$ (solid curves) versus the reciprocal of the number of knot spacings (i.e. h).

the same and that the separation of the error curves remains nearly constant for all values of h. For $p = 1, 2$ this constant is very close to 1 so the two error curves are not distinguishable from each other. For $p = 3$ the value is approximately 20 and for $p = 4$ it is approximately 50 – all consistent with the error bounds in Table 3. The deviation for the LLSQ method on the finest knot spacing for $p = 4$ is due to round-off errors.

6.3. NURBS Example

As a final example we consider the case of a non-rational B-spline basis (NURBS) for a surface embedded in $\mathbb{R}^3$. This case is also directly applicable to the original motivation for the development of the LLSQ, viz., isogeometric analysis. Due to the technical complexity associated with providing proofs in this case we consider the problem from a purely numerical standpoint.

Consider the projection of the data function

$$u(x, y) = \sin(3\pi x/\sqrt{2}R)\sin(2\pi y/L)$$

onto $\mathcal{F}$ and $\mathcal{E}$, where these spaces are defined by a tensor product of NURBS functions that exactly map one-quarter of a cylindrical shell as shown in Fig. 6(upper left). The radius of curvature of the shell is $R = 5$ units and the length of the shell is $L = \pi R/2$ units. Also shown in Fig. 6 are the convergence curves for the cases of a uniformly refined NURBS basis in both directions of orders $p = 2, 3, 4$. In each case, one observes that the rate of convergence for the skew projection (solid line) and the full least square projection (dashed line) match each other and that the separation of the curves, for each order, remains essentially constant for all values of h. The separation ratios are further noted to be compatible with the theoretical B-spline constants C – in this case roughly 1.1, 10, and 35 for $p = 2$, 3, and 4, respectively. Given the intimate relation between the two bases, this is not unexpected. Again, we see a the deviation of the LLSQ method on the finest knot spacing for $p = 4$ due to round-off errors.

7. Conclusions

The LLSQ method was introduced in [1] and shown to be effective on a series of problems arising in isogeometric analysis of solids but without rigorous justification. In the present work we have examined the methodology in detail and shown in general that the LLSQ method possesses an error that is bounded by a constant times the global least square error. For two special cases in one dimension we have shown that the constant is bounded and independent of the element size used to define the local basis functions at the heart of the LLSQ method. Using similar tools but with far more involved algebra the results shown carry over to higher dimensions. As a simple illustration, we have provided a numerical example employing a NURBS basis in two-dimensions demonstrating good behavior. Considering the original motivation of isogeometric analysis in [1], we can also conclude that using the LLSQ method to enforce Dirichlet boundary conditions will not change the expected rates of global convergence for an isogeometric analysis. Since in higher dimensions the LLSQ method is more efficient than the LSQ method, we view it as a practical method for this purpose.

8. Acknowledgments

John Strain's work was sponsored by the Air Force Office of Scientific Research, USAF, under grant number FA9550-08-1-0131, and by the Division

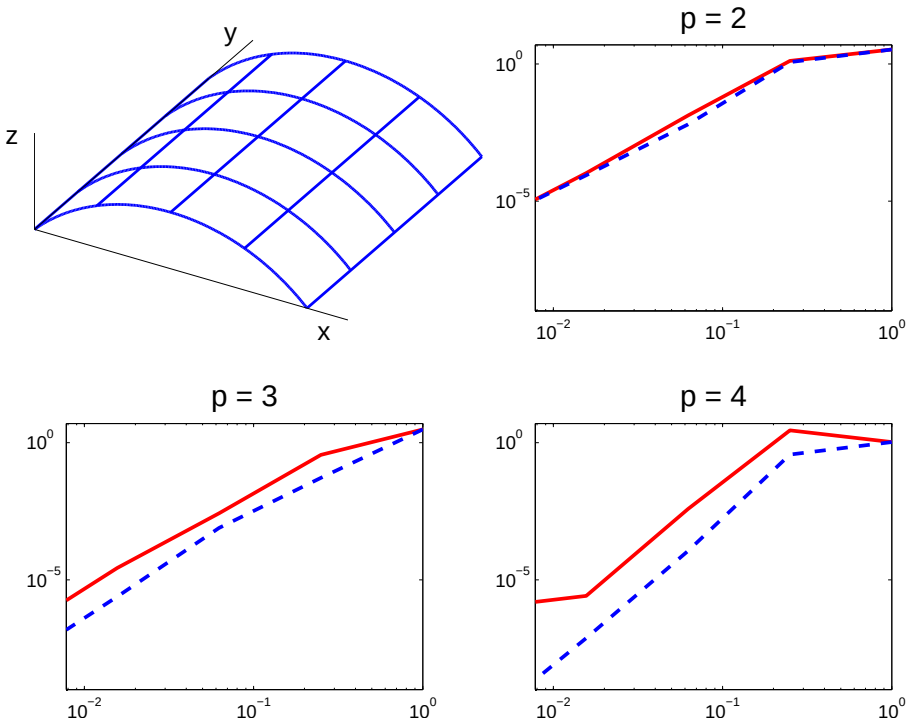

Figure 6: (Upper left) One-quarter cylindrical shell domain over which the projection is performed. (Upper right and lower row) Convergence curves for $\|u - u_F\|$ (dashed curves) and $\|u - u_S\|$ (solid curves) versus normalized knot spacing (i.e. h/L) for a two-dimensional NURBS example.

of Mathematical Sciences, National Science Foundation, under grant number DMS-0913695.

References

[1] T. J. Mitchell, S. Govindjee, R. L. Taylor, A method for enforcement of Dirichlet boundary conditions in isogeometric analysis, in: D. Mueller-Hoeppe, S. Loehnert, S. Reese (Eds.), Recent Developments and Innovative Applications in Computational Mechanics, Springer-Verlag, 2011, pp. 283–293.

[2] T. Hughes, J. Cottrell, Y. Bazilevs, Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement, Computer Methods in Applied Mechanics and Engineering 194 (2005) 4135–4195.

[3] J. A. Cottrell, T. J. R. Hughes, Y. Bazilevs, Isogeometric Analysis: Towards Integration of CAD and FEA, John Wiley and Sons, 2009.

[4] C. Gauss, Theoria Combinationis Observationum: Erroribus Minimis Obnoxiae, Apud Henricum Dieterich, 1825.

[5] S. D. Conte, C. de Boor, Elementary Numerical Analysis: An Algorithmic Approach, McGraw-Hill, 1980.

[6] G. H. Golub, C. F. Van Loan, Matrix Computations, Johns Hopkins University Press, 3rd edition, 1996.

[7] G. Strang, Groups of banded matrices with banded inverses, Proceedings of the American Mathematical Society doi: 10.1090/S0002-9939-2011-10959-6 (2011) 1–10.

[8] M. Abramowitz, I. A. Stegun, Handbook of Mathematical Functions, Dover, 1965.