Kai Hormann, N. Sukumar

Generalized Barycentric Coordinates in Computer Graphics and Computational Mechanics

Contents

CHAPTER 1

Maximum-entropy meshfree coordinates in computational mechanics

Marino Arroyo

LaCàN, Department of Civil and Environmental Engineering
Universitat Politècnica de Catalunya-BarcelonaTech
Carrer Jordi Girona 1-3, 08034 Barcelona, Spain

CONTENTS

CONVENTIONAL BARYCENTRIC COORDINATES are defined relative to vertices defining the boundary of a polytope. Here, we develop barycentric coordinates relative to a cloud of vertices sampling not only the boundary, but also the interior of a region in $\mathbb{R}^d$, with the objective of using these barycentric coordinates as basis functions to approximate partial differential equations or parametrize surfaces. We show that entropy maximization provides a rational way to define smooth barycentric coordinates, but for the resulting basis functions to be localized, and hence lead to sparse matrices in computational mechanics, entropy maximization needs to be biased by a suitable notion of locality. The basis functions that result from this approach are smooth, reproduce polynomials, are localized around their corresponding vertex, and their definition does not rely on

an underlying mesh or grid but rather on less structured neighborhood relations between vertices. Thus, they can be viewed as meshfree basis functions [12, 21, 37]. The theory and practical evaluation behind these basis functions is reviewed next, and selected applications in computational mechanics are presented.

1.1 INTRODUCTION

A standard approach to numerically represent a function u over a domain in $\Omega \subset \mathbb{R}^d$ is to expand it in terms of a finite set of basis functions, that is

$$u^h(\boldsymbol{x}) = \sum_{i=1}^{n} \phi_i(\boldsymbol{x}) u_i, \tag{1.1}$$

where $\phi_i : \bar{\Omega} \longrightarrow \mathbb{R}$ is the $i-$th basis function and u_i its corresponding coefficient. Such a representation is the basis of Galerkin methods to solve partial differential equations (PDEs), see [15] and also Chapter ??. Here, rather than choosing a priori the set of basis functions, such as piecewise polynomials defined over a mesh, we initially leave their definition open, and then motivate and make explicit the specific choices leading to maximum-entropy meshfree basis functions.

To show convergence of numerical solutions obtained by a Galerkin method to the exact solution, the basis functions need to satisfy the so-called reproducibility conditions. These conditions ensure that polynomials up to degree p are exactly represented by the basis functions. For second-order PDEs such as the heat equation and the most common systems arising in fluid and solid mechanics, consistency conditions up to $p = 1$ are required at all points $\boldsymbol{x} \in \Omega$

$$\sum_{i=1}^{n} \phi_i(\boldsymbol{x}) = 1, \tag{1.2a}$$

$$\sum_{i=1}^{n} \phi_i(\boldsymbol{x}) \boldsymbol{v}_i = \boldsymbol{x}, \tag{1.2b}$$

for some vector coefficients $\boldsymbol{v}_i \in \mathbb{R}^d$. If these conditions are met, then any affine function $w(\boldsymbol{x}) = \boldsymbol{a} \cdot \boldsymbol{x} + b$, where $\boldsymbol{a}$ is a vector in $\mathbb{R}^d$ and $\cdot$ denotes the scalar product, can be exactly represented in Ω as $w(\boldsymbol{x}) = \sum_{i=1}^{n} \phi_i(\boldsymbol{x})(\boldsymbol{a} \cdot \boldsymbol{v}_i + b)$.

Interpreting the coefficients $\boldsymbol{v}_i$ as the coordinates of points or vertices, we can naturally associate each basis function to a vertex, and view the set of vertices $V = \{\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_n\}$ as points sampling the domain Ω. With this interpretation, we only need to require the basis functions to be non-negative

$$\phi_i(\boldsymbol{x}) \geq 0 \quad \forall \boldsymbol{x} \in \Omega, \; i = 1, 2, \ldots, n \tag{1.3}$$

and recall (1.2) to realize that the functions $\phi_i(\boldsymbol{x}), \; i = 1, 2, \ldots, n$ define a set of generalized barycentric coordinates relative to V. A difference with respect to previous chapters is that now we allow for interior vertices to the domain, and for multiple vertices along a face of the domain.

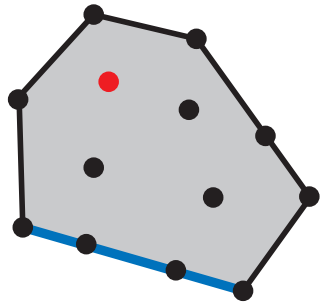

Figure 1.1 Set of vertices V in two dimensions, along with the convex hull P.

Because generalized barycentric coordinates at $\boldsymbol{x}$ are coefficients of a convex combination of the vertices resulting in $\boldsymbol{x}$, see (1.2b), it immediately follows that such basis functions can only be defined in the convex hull of V, and therefore $\Omega \subset \operatorname{conv} V$. Further exploiting elementary facts of convex geometry, it is possible to characterize the behavior of *convex approximation schemes* of V, that is barycentric coordinates defined in the vertex set V, on the boundary of the polytope $P = \operatorname{conv} V$ [4]. In particular, it can be shown that if F denotes a face of P in the sense of convex geometry (in $\mathbb{R}^3$ 0-dimensional faces are extreme points, 1-dimensional faces are edges, and 2-dimensional faces are a proper faces) and $\boldsymbol{v}_i \notin F$, then the corresponding basis function ϕ_i vanishes on F [4].

Immediate consequences of this fact are that: (1) basis functions of interior vertices (red vertex in Figure 1.1) vanish at the boundary of P, (2) the basis functions of extreme points $\boldsymbol{v}_i$ of P satisfy the Kronecker-delta property, $\phi_i(\boldsymbol{v}_j) = \delta_{ij}$, and (3) on a given face F (in blue in Figure 1.1), only the basis functions of vertices lying on F are non-zero in F. As a result of (3), if we denote by I_F the set of indices of nodes lying on F, then (1.2a,1.2b) become $\sum_{i\in I_F} \phi_i(\boldsymbol{x}) = 1$ and $\sum_{i\in I_F} \phi_i(\boldsymbol{x})\boldsymbol{v}_i = \boldsymbol{x}$ for all $\boldsymbol{x} \in F$, that is the basis functions of nodes on a face maintain the full approximation properties on that face. Furthermore, (3) also implies that the basis functions on a face can be defined without reference to information in the higher-dimensional ambient space. These properties can be referred to as a weak Kronecker-delta property, which facilitate the imposition of boundary conditions in the numerical approximation of PDEs.

1.2 SELECTING BARYCENTRIC COORDINATES THROUGH ENTROPY MAXIMIZATION

On examining (1.2a,1.3), it is clear that for each point $\boldsymbol{x} \in P$, the barycentric coordinates $\{\phi_1(\boldsymbol{x}), \phi_2(\boldsymbol{x}), \ldots, \phi_n(\boldsymbol{x})\}$ can be viewed as a discrete probability distribution for n events. Defining the barycentric coordinates then becomes a problem of statistical inference. In the absence of additional information and if one tries to avoid any bias, Laplace's principle of insufficient reason would suggest selecting

$\phi_i(\boldsymbol{x}) = 1/n$ for $i = 1, \ldots, n$. Yet, we do have additional information, since (1.2b) states that the average of the vertex positions is $\boldsymbol{x}$. How to incorporate this piece of information, while making a choice of the probabilities (barycentric coordinates) as unbiased as possible?

One classical answer to this question in information theory is Jaynes' principle of maximum entropy [17]. Let us first introduce the information entropy associated to a finite probability distribution by considering a coin toss. In this case, there are two events (heads or tails) that for a perfectly balanced coin have equal probabilities, $p_1 = p_2 = 1/2$. Instead, we could conceive an extremely unbalanced coin with probabilities $q_1 = 0.99, q_2 = 0.01$. A fundamental question in information theory is how to quantify the uncertainty of a probability distribution, or in other words, how to quantify the amount of information gained by realizing a coin toss. It is obvious that the balanced coin leads to a much more uncertain outcome, whereas, for the unbalanced one, we will learn very little by throwing it since we know that the outcome will most likely be heads. It can be shown that a canonical measure of uncertainty or information is given by the Shannon entropy function

$$H(p_1, p_2, \ldots, p_n) = -\sum_{i=1}^{n} p_i \log p_i, \tag{1.4}$$

where the function $x \log x$ is extended by continuity at zero, i.e. $0 \log 0 = 0$. See [18] for a detailed motivation of this measure of information and its properties. By evaluating Shannon entropy on the previous two probability distributions, we find that the first coin is about 12 times more uncertain than the second one. In fact, it can be shown that the function in (1.4) is strictly concave with a maximum at the uniform distribution $p_i = 1/n,\ i = 1, 2, \ldots, n$. This last fact shows that maximizing the entropy is consistent with Laplace's principle of insufficient reason.

Let us now introduce Jaynes' principle using another simple example, that of a dice with six faces, $A_1, A_2, \ldots, A_6$. Consider the random function that assigns to a given face its numerical value, $f(A_i) = i$. For a perfectly balanced dice, the average of f is $\sum_{i=1}^{6} i/6 = 7/2$. Suppose, however, that an accurate sample mean is computed for this quantity after a large number of trials, giving a result of 4. Clearly, this new piece of information is inconsistent with the uniform probabilities that we expect for a perfectly balanced dice. We would like to incorporate this information in the inference of the probabilities $p_1, p_2, \ldots, p_6$ associated to each face of the dice, but clearly it is not sufficient to uniquely determine these probabilities. Jaynes' recipe is to consider the most uncertain probability distribution, that is the one that maximizes entropy, while being consistent with the known partial information [17]. Mathematically, this statement can be formulated as an optimization program with constraints

$$\begin{aligned} \max_{p_1,\ldots,p_6} \quad & -\sum_{i=1}^{6} p_i \log p_i \\ \text{subject to} \quad & p_i \geq 0, \ \sum_{i=1}^{6} p_i = 1, \ \sum_{i=1}^{6} i\ p_i = 4, \end{aligned}$$

where the last constraint encodes the additional information about the system. The solution to this constrained optimization problem is $p_1 = 0.103$, $p_2 = 0.123$, $p_3 = 0.146$, $p_4 = 0.174$, $p_5 = 0.207$, and $p_6 = 0.247$, showing that the higher expected value for f biases the probabilities towards faces with a higher numerical value.

Going back to the problem of inferring the generalized barycentric coordinates, it is now clear that application of Jaynes' principle of maximum entropy leads to an optimization program at each point $\boldsymbol{x} \in P$ and given by

$$\begin{aligned} \max_{\phi_1,\ldots,\phi_n} \quad & -\sum_{i=1}^{n} \phi_i \log \phi_i \\ \text{subject to} \quad & \phi_i \geq 0, \;\; \sum_{i=1}^{n} \phi_i = 1, \;\; \sum_{i=1}^{n} \phi_i \; \boldsymbol{v}_i = \boldsymbol{x}, \end{aligned} \tag{1.5}$$

where the dependence on $\boldsymbol{x}$ only enters through the last vectorial constraint. The solution to the program is the value of the barycentric coordinates selected by entropy maximization at $\boldsymbol{x}$, $\phi_i(\boldsymbol{x})$, $i = 1, 2, \ldots, n$. Because the entropy function is strictly concave, this program has a unique solution if and only if it is feasible, that is whenever $\boldsymbol{x} \in \bar{P}$ [8]. This method, reviewed in Section **??**, was introduced in [35] to generate basis functions in polygons.

Figure 1.2 shows the basis functions defined by entropy maximization for a set of vertices that not only defines a polygon, but also samples its interior and its faces. It can be observed that in agreement with the general results for convex approximation schemes mentioned earlier, the basis function associated to the extreme vertex $\boldsymbol{v}_1$ is 1 at $\boldsymbol{v}_1$, and therefore all other basis functions are zero at this point. Also in agreement with the general results, the basis function of the interior vertex $\boldsymbol{v}_3$ vanishes on the boundary of P, and the basis function ϕ_2 does not satisfy the Kronecker-delta property because $\boldsymbol{v}_2$ is on the boundary but is not an extreme point. Yet, this set of basis functions presents a significant disadvantage when used to approximate functions in general, or solutions of PDEs in particular: the basis functions are completely nonlocal. The most dramatic case is that of basis functions of interior nodes, such as ϕ_3, which spread through the entire domain and are as uniform as allowed by the convex structure of the optimization program. This is not a desirable feature to approximate functions over P because with such a basis it is not possible to define a localized feature, say around vertex $\boldsymbol{v}_i$, and correlate this feature with the nodal value u_i in (1.1). Besides the lack of local resolution, such a global set of basis functions leads to a dense stiffness matrix when used in a Galerkin method to approximate partial differential equations [15].

This discussion suggests that in defining the generalized coordinates by mere entropy maximization, we have been "too unbiased". We have ignored the fundamental requirement of locality, by which the approximation in (1.1) at a given point $\boldsymbol{x}$ should more strongly depend on coefficients u_i associated to vertices close to $\boldsymbol{x}$. Local basis functions should concentrate their "mass" in the vicinity of their associate vertex, where their value should be closer to 1 than for instance ϕ_3 in Figure 1.2.

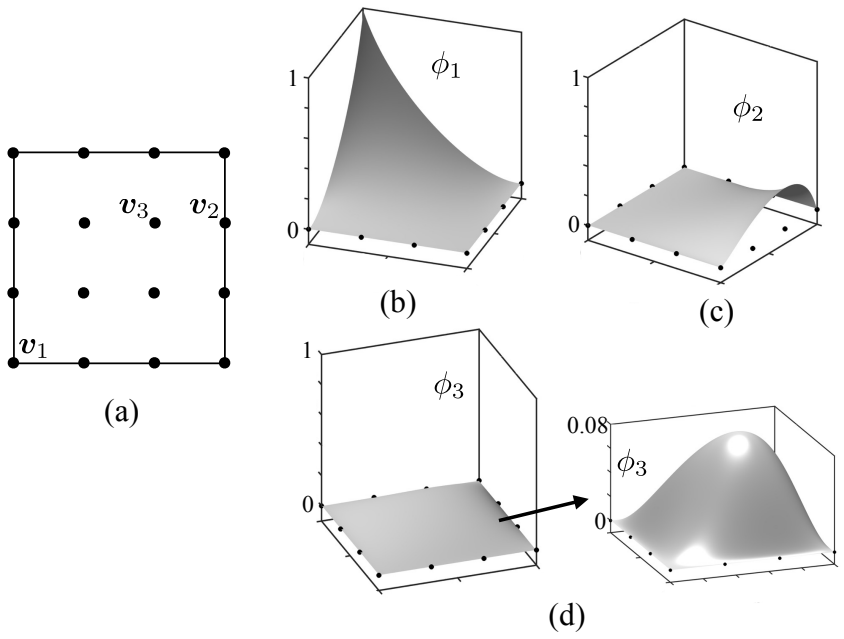

Figure 1.2 Basis functions selected by entropy maximization for the set of vertices shown in (a). Selected basis functions for vertices at extreme positions of P (b), along a face of P (c), and in the interior of P (d).

1.3 INTRODUCING LOCALITY: LOCAL MAXIMUM-ENTROPY APPROXIMANTS

Clearly, locality and entropy maximization are conflicting objectives; at any given point $\boldsymbol{x}$ the former tries that only a few basis functions, those of vertices closer to $\boldsymbol{x}$, have values significantly larger than 0 while all others are close to 0, while the latter tries to have as uniform values of $\phi_i(\boldsymbol{x})$ as possible. Interestingly, there is another family of generalized barycentric coordinates associated to V that can be characterized through an optimization program analogous to that in (1.5): the piecewise linear basis functions associated to the Delaunay triangulation of V. Indeed, if the vertices in V are in *general positions* (there are no $d+1$ cospherical vertices in V), then there is a unique Delaunay triangulation with vertices V and the solution to the constrained optimization problem

$$\begin{aligned} \min_{\phi_1,\dots,\phi_n} \quad & \sum_{i=1}^{n} \phi_i |\boldsymbol{x}-\boldsymbol{v}_i|^2 \\ \text{subject to} \quad & \phi_i \geq 0, \;\; \sum_{i=1}^{n} \phi_i = 1, \;\; \sum_{i=1}^{n} \phi_i \, \boldsymbol{v}_i = \boldsymbol{x}, \end{aligned} \tag{1.6}$$

are the barycentric coordinates of the simplex in the d-dimensional Delaunay triangulation where $\boldsymbol{x}$ belongs [30]. Thus, at each point there are at most $d+1$ non-zero basis functions. See [30, 4] for a discussion about the situation in which the vertices are not in general positions. Arguably, these are the most local generalized barycentric coordinates associated to V.

Comparison of (1.5) and (1.6) suggests combining both objective functions to define the following optimization problem

$$\begin{aligned} \min_{\phi_1,\dots,\phi_n} \quad & \beta \sum_{i=1}^{n} \phi_i |\boldsymbol{x}-\boldsymbol{v}_i|^2 + \sum_{i=1}^{n} \phi_i \log \phi_i \\ \text{subject to} \quad & \phi_i \geq 0, \;\; \sum_{i=1}^{n} \phi_i = 1, \;\; \sum_{i=1}^{n} \phi_i \, \boldsymbol{v}_i = \boldsymbol{x}, \end{aligned} \tag{1.7}$$

which defines Pareto optima harmonizing the two conflicting objectives: information theory optimality and locality. Depending on β, or its non-dimensional counterpart $\gamma = \beta h^2$ where h is the typical spacing between vertices, the basis functions will more closely resemble nonlocal approximants such as those shown in Figure 1.2 or highly local piecewise affine barycentric coordinates supported on a Delaunay triangulation. In fact, the locality parameter β, controlling the aspect ratio of the basis functions, can be defined locally by considering the following objective function in (1.7)

$$\sum_{i=1}^{n} \beta_i \phi_i |\boldsymbol{x}-\boldsymbol{v}_i|^2 + \sum_{i=1}^{n} \phi_i \log \phi_i,$$

where β_i is the aspect ratio parameter associated to vertex i. We call the set of basis

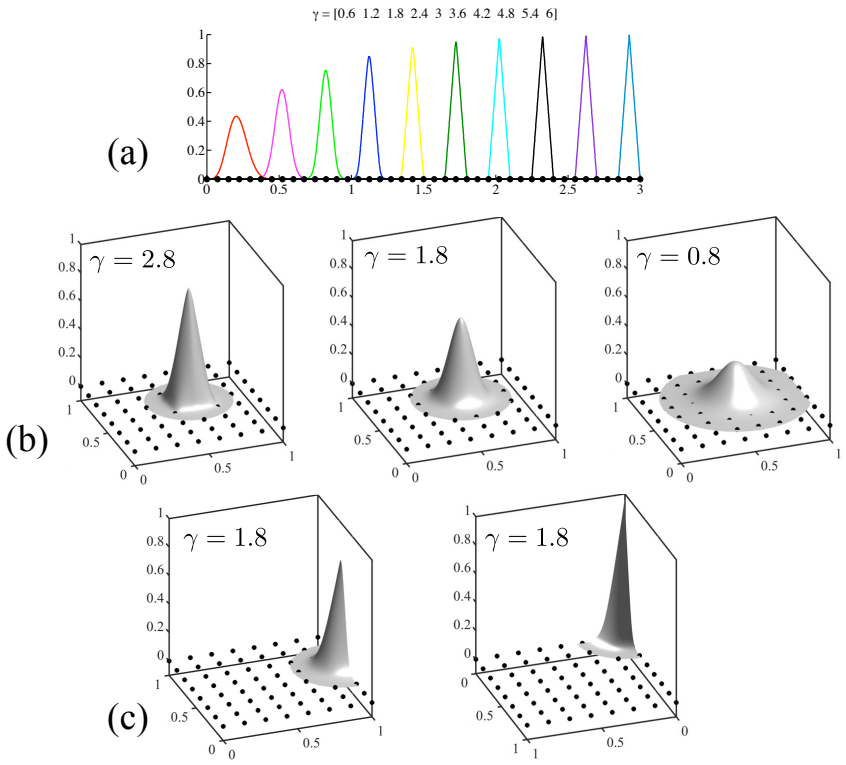

Figure 1.3 Local maximum-entropy (LME) basis functions. (a) Selected basis functions for a one-dimensional vertex set, with varying nondimensional aspect ratio parameter $\gamma_i = \beta_i h^2$, see main text. (b) Basis functions for an interior node in a two-dimensional set of vertices and different aspect ratio parameters, and (c) representative basis functions of boundary nodes.

functions obtained through the optimization program (1.7), $\phi_i(\boldsymbol{x})$, $i = 1, 2, \ldots, n$, *local maximum-entropy* (LME) approximants. Figure 1.3 illustrates the LME basis functions in one and two dimensions, although the formulation can be implemented in principle in any dimension. The smooth basis functions closely resemble other meshfree basis functions, such as those generated by the moving least-squares method [37]. The crucial difference, however, is that such alternative methods do not produce in general generalized barycentric coordinates because the basis functions can be negative. The fact that LME functions are non-negative, and thus convex approximants with a weak Kronecker-delta property, facilitates the imposition of essential boundary conditions in the numerical solution of PDEs.

Note that while (1.6) is a linear program with unique solution only when vertices are in general positions, the one-parameter family of optimization programs in (1.7)

is nonlinear and strictly convex as a result of the entropy function. Thus, since it is feasible for all $\boldsymbol{x} \in \bar{P}$, it has a unique solution. Thus, the entropy functional regularizes in some sense the linear program in (1.6). In fact, as shown in [4], the entropy regularization leads to a unique well-defined set of generalized Delaunay barycentric coordinates in the limit $\beta \to +\infty$.

Another important consequence of the entropy functional in (1.7) is the smoothness of the resulting LME barycentric coordinates, as compared to those resulting from (1.6). While the Delaunay barycentric coordinates are only C^0 in P, since they have discontinuous derivatives along the faces of the Delaunay simplices, it can be shown using the implicit function theorem that the LME approximants are C^∞ in P with respect to $\boldsymbol{x}$ (also with respect to β_i).

The LME basis functions have been defined so far implicitly through a convex constrained optimization program with as many unknowns as vertices in V. Through standard duality methods, however, it is possible to obtain an almost explicit form of the basis functions. For this, let us first define the partition function

$$Z(\boldsymbol{x}, \boldsymbol{\lambda}) = \sum_{i=1}^{n} \exp\left[-\beta_i |\boldsymbol{x} - \boldsymbol{v}_i|^2 + \boldsymbol{\lambda} \cdot (\boldsymbol{x} - \boldsymbol{v}_i)\right]. \tag{1.8}$$

Then, the LME basis functions can be expressed at any point $\boldsymbol{x} \in P$ as

$$\phi_i(\boldsymbol{x}) = \frac{1}{Z\left(\boldsymbol{x}, \boldsymbol{\lambda}^*(\boldsymbol{x})\right)} \exp\left[-\beta_i |\boldsymbol{x} - \boldsymbol{v}_i|^2 + \boldsymbol{\lambda}^*(\boldsymbol{x}) \cdot (\boldsymbol{x} - \boldsymbol{v}_i)\right]. \tag{1.9}$$

This expression is not fully explicit because to compute $\boldsymbol{\lambda}^*(\boldsymbol{x})$ an unconstrained optimization must be solved: the mininization of the strictly convex function $\log Z(\boldsymbol{x}, \boldsymbol{\lambda})$. It can be observed that the functions can be viewed as the product of a Gaussian function around vertex $\boldsymbol{v}_i$, a normalizing factor (Z) so that all functions add up to one, and an exponential factor that depends on the Lagrange multiplier $\boldsymbol{\lambda}^*$ enforcing the linear reproducing condition. Examination of (1.9) shows that, strictly speaking, these basis functions have global support. However, thanks to the fast decay of the Gaussian function, for all practical purposes they can be considered as compactly supported. Furthermore, the sum in calculation of Z in (1.8) effectively involves only a handful of terms corresponding to vertices in the neighborhood of $\boldsymbol{x}$, which makes the evaluation of the basis functions efficient.

We emphasize the dependence of $\boldsymbol{\lambda}^*$ on position to highlight the fact that at each evaluation point, an unconstrained nonlinear optimization problem must be solved. This problem, however, is only d-dimensional, has a unique solution, and can be efficiently solved numerically using Newton's method. See [4, 32, 23] for mathematical and computational details, as well as for the explicit expressions to compute first and second spatial derivatives of the basis functions, and see [13] for the calculation of the gradients of the basis functions on the boundary of P.

We end this section by presenting an alternative method to introduce locality into the maximum-entropy framework, which in addition provides a means to define strictly compactly supported basis functions. It is very easy to define a set of non-negative smooth basis functions $\psi_i(\boldsymbol{x})$ relative to V, localized around the

corresponding vertex, and satisfying only the zeroth-order reproducing condition, $\sum_{i=1}^{n} \psi_i(\boldsymbol{x}) = 1$. Indeed, consider for instance a scalar, non-negative and smooth window function $w : [0, +\infty) \longrightarrow \mathbb{R}$, which decays when the argument is large compared to 1 and satisfies that all odd derivatives up to k at 0 vanish (to guarantee that the even extension of w to $\mathbb{R}$ is C^k). Then, for a large enough ρ the functions

$$\psi_i(\boldsymbol{x}) = \frac{w\left(|\boldsymbol{x} - \boldsymbol{v}_i|/\rho\right)}{\sum_{j=1}^{n} w\left(|\boldsymbol{x} - \boldsymbol{v}_j|/\rho\right)}$$

are clearly non-negative, localized around their corresponding vertex $\boldsymbol{v}_i$ with typical width ℓ, smooth, and add up to one. Furthermore if w is compactly supported, then the functions $\psi_i(\boldsymbol{x})$ are also compactly supported. Thus, at each point, the numbers $\psi_1(\boldsymbol{x}), \psi_2(\boldsymbol{x}), \ldots, \psi_n(\boldsymbol{x})$ define a discrete probability distribution, which contains information about locality ($\psi_i(\boldsymbol{x})$ is larger if $\boldsymbol{x}$ is closer to $\boldsymbol{v}_i$) but does not satisfy the first-order reproducing constraint (1.2b). This constraint can be recovered using the concept of *relative entropy*, also called Kullback-Leibler distance, which measures the amount of information required to obtain a discrete probability distribution from a prior distribution. Viewing $\psi_i(\boldsymbol{x}),\ i = 1, \ldots, n$ as a prior distribution, the relative entropy can be written as

$$D(\boldsymbol{\phi}|\boldsymbol{\psi}(\boldsymbol{x})) = \sum_{i=1}^{n} \phi_i \log \frac{\phi_i}{\psi_i(\boldsymbol{x})}. \tag{1.10}$$

Minimization of this function with respect to $\phi_i,\ i = 1, \ldots, n$, subject to (1.2) and (1.3) will produce the closest set of barycentric coordinates to the prior from an information theoretical viewpoint, i.e. introducing the least extra information [5, 37]. A direct calculation using duality methods shows that the resulting basis function ϕ_i can be written as the product of ψ_i and another factor enforcing the linear reproducing condition, and thus the prior and the relative maximum-entropy functions have the same support. In particular, if $w(r) = e^{-r^2}$ and $\rho = 1/\sqrt{\beta}$, then these functions coincide with the LME basis functions.

1.4 FURTHER EXTENSIONS

We summarize next extensions to the meshfree and convex basis functions generated by the LME optimization program. An obvious question is whether these approximants can be extended to higher-order, that is, if in addition to exactly reproducing affine functions due to (1.2), they can be designed to exactly reproduce higher-order polynomials. This would lead to a higher rate of convergence when approximating PDEs. Focusing in 1D for simplicity, a naive extension would be to maximize the LME objective function subject to

$$\sum_{i=1}^{n} \phi_i = 1, \quad \sum_{i=1}^{n} \phi_i v_i = x, \quad \sum_{i=1}^{n} \phi_i v_i^2 = x^2, \tag{1.11}$$

a set of conditions met for instance by quadratic finite element spaces. However, as identified in [4], conditions in (1.11) together with the non-negativity constraint

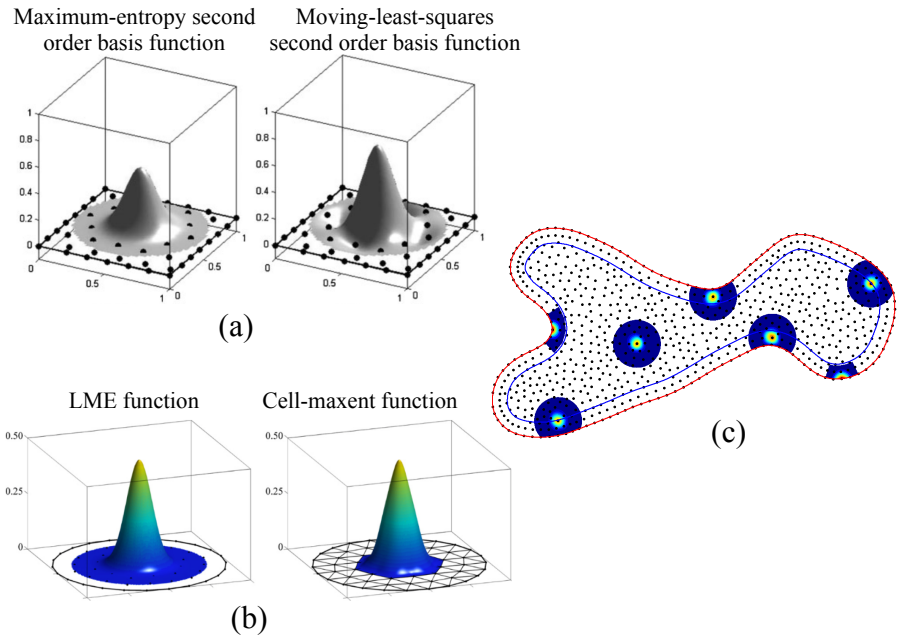

Figure 1.4 Extensions of maximum-entropy approximations. (a) Representative second-order maximum entropy basis function, compared to a moving-least-squares function, which is wiggly and negative at some points. (b) Comparison of a standard LME basis function and a cell-based maximum entropy basis function supported on a triangulation, which results in a structured adjacency relation and more efficient calculations to approximate PDEs for a comparable accuracy. (c) Combination of LME approximants with a boundary representation based on B-splines.

$\phi_i \geq 0$ are unfeasible for almost all points in P. This may seem contradictory at first sight, since there are examples of non-negative approximants that can reproduce quadratic and higher-order functions, notably B-splines. One option is the give up non-negativity and consider a notion of entropy for signed probability distributions, which however destroys the convex structure of the approximation scheme, does not produce generalized barycentric coordinates, and can lead to wiggly basis functions [36, 6]. Insisting on non-negativity, the apparent contradiction referred to above is resolved by noting that for the second-order consistency condition to hold, the coefficients that multiply the generalized barycentric coordinates do not need to be v_i^2. By designing appropriate coefficients, it is possible to restore feasibility, and thus define second-order maximum entropy approximants in multiple dimensions [9, 33], see Figure 1.4(a). The construction of feasible constraints for higher-order reproducing conditions, however, is an open question.

When applied to the numerical solution of PDEs with Galerkin methods, one aspect complicating the implementation and diminishing the efficiency of LME approximations is the "uncontrolled" support of the basis functions. This is important because any two basis functions with overlapping support will generate an entry in the stiffness matrix resulting from the Galerkin method. Because the LME basis functions have completely unstructured support, the underlying adjacency structure setting the matrix structure is quite dense, but with many very small entries [28]. To alleviate this, a modification of the LME basis functions was proposed, and called cell-based maximum entropy approximations, in which prior functions ($\psi_i(\boldsymbol{x})$ in the previous section) were designed to be supported on the a k-ring of simplices around $\boldsymbol{v}_i$ of a triangulation, see Figure 1.4(b). In this way, the adjacency structure of the stiffness matrix is given directly by the topology of the triangulation and is much sparser, leading to more efficient calculations without nearly any degradation in accuracy [25]. In this reference, the smooth and compactly supported prior functions were computed using approximate distance functions and R-functions, but other techniques may result in better priors.

We finally report on a method that uses LME approximants to represent domains with complex geometry and smooth boundaries. As described up to now, the natural domain where the LME approximants are defined is the convex hull of the P of the set of vertices, or subsets of P, but in this case the approximants do not satisfy the weak delta-Kronecker on the boundary. However, there is an increasing awareness on the importance of geometric fidelity in engineering calculations [16]. In fact, it is quite natural to blend maximum-entropy approximations with any other convex approximation scheme, such as most approximation methods used in computer graphics. It is sufficient to consider an optimization program such as that in (1.7), in which some of the barycentric coordinates are taken as data (the known convex scheme) and the rest are the unknowns of the problem [31]. For instance, in Figure 1.4(c) one layer of tensor product B-spline basis functions extruded from a B-spline boundary representation was blended with LME functions.

1.5 APPLICATIONS

The LME approximants and their extensions have been adopted as trial and test functions in Galerkin methods to solve various PDEs, including linear and nonlinear elasticity [31, 4, 32, 33], incompressible solids [27], and fluid flow problems [19, 26, 29]. In these problems, which often exhibit smooth solutions, LME approximations provide highly accurate numerical approximations with relatively coarse sets of points, as compared to standard C^0 finite elements. Furthermore, the meshfree basis functions can accommodate very large grid distortions in Lagrangian large deformation simulations [19, 29]. However, some applications truly benefit from the smoothness of the LME approximants for general unstructured sets of points. We outline next two example: the numerical approximation of high-order PDEs and the approximation of manifolds defined by scattered sets of points.

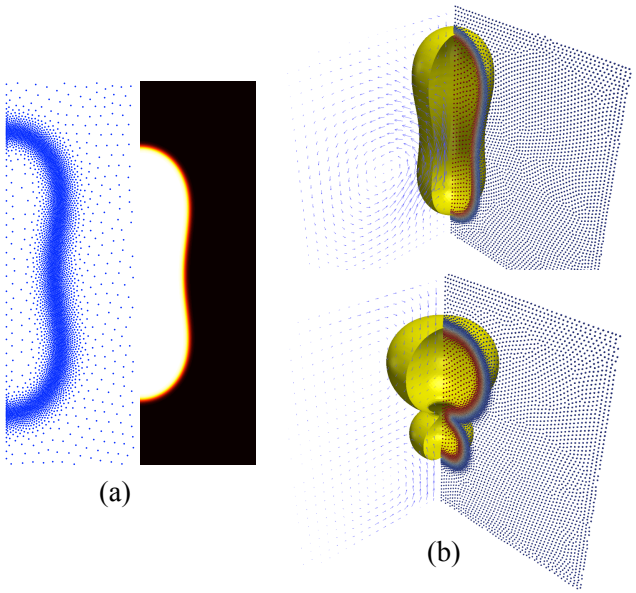

Figure 1.5 Application of the LME approximants to high-order PDEs. (a) In phase-field models of biomembranes, the unknown field defining the shape of a moving interface is governed by a fourth-order PDE. Therefore, a Galerkin method requires smooth basis functions (with square integrable second derivatives). The phase field develops sharp gradients, calling for an adaptive solution and an approximations scheme devoid of Gibbs phenomenon. All these conditions are met by LME approximations. (b) shows two snapshots in a dynamical simulation tracking the shape of a deforming vesicle made out of a fluid membrane with curvature elasticity, together with the hydrodynamics of the surrounding fluid. The resolution of the set of vertices follows the sharp variations of the phase-field.

1.5.1 High-order partial differential equations

Many physical phenomena of interest in science and engineering are modeled using higher-order PDEs. For instance, the Kirchhoff-Love equations of thin shells are fourth-order PDEs. Similarly, many problems involving evolving discontinuities can be modeled using phase-field models, where the discontinuity is smeared out and tracked by a field governed by a PDE. In many cases, such a PDE is of fourth order, including models for crack propagation [7], for crystal growth [38], or for the mechanics of fluid membranes [10]. The numerical treatment of such problems with C^0 finite elements is cumbersome and involves interpolating independently the field and its gradient [11]. In contrast, a Galerkin approach is straightforward if smooth basis functions are used. LME are not only smooth, but also easily amenable to adaptive methods and have monotonicity properties that help them accurately describe the sharp variations characteristic of phase-field solutions. See Figure 1.5 for an illustration of LME approximants applied to a phase-field model of vesicle dynamics [34, 29]. LME approximants have also been applied to phase-field models of fracture [20], and to other high-order PDEs such as those arising in the mechanics of thin shells [24] or in the coupled electromechanics of flexoelectric materials [2, 1].

1.5.2 Manifold approximation

Finally, LME approximations have been used to smoothly parametrize manifolds sampled by clouds of points in an automated way and without the need for a mesh, using an atlas of partially overlapping charts [24]. See [22] for recent work along these lines, with a review of similar approaches in computer graphics. Figure 1.6(a) illustrates the procedure, in which the point-set d-dimensional manifold is systematically partitioned until each partition is homeomorphic to an open set in $\mathbb{R}^d$. Then, each partition is embedded in $\mathbb{R}^d$ using nonlinear dimensionality reduction methods, and this embedding is used as a parametric domain for a local parametrization of the partition. Parametrizations of adjacent partitions are "glued" with partition-of-unity functions. This natural but powerful procedure has been used to model thin shells of complex geometry Figure 1.6(b), to parametrize the gait of microscopic swimming cells [3], or to define collective variables for molecular systems based on parametrizing the so-called intrinsic manifold underlying molecular flexibility [14], see Figure 1.6(b).

1.6 OUTLOOK

In summary, the principle of maximum entropy is a conceptually appealing and computationally practical approach to define generalized barycentric coordinates for clouds of points. The resulting barycentric coordinates can be viewed as meshfree basis functions and used in computational mechanics, exploiting their smoothness and the ease to implement adaptive strategies. The fact that these basis functions are barycentric coordinates makes it easier to impose boundary conditions, but also provides a natural framework to coupled maximum entropy basis functions

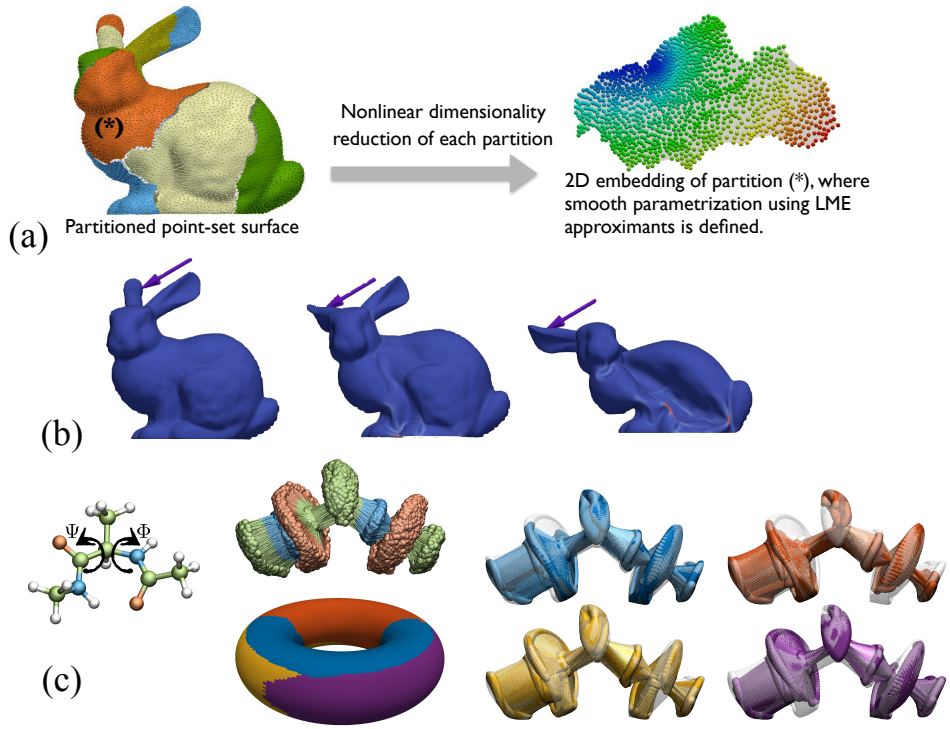

Figure 1.6 Approximation of manifolds using an atlas of smooth LME parametrizations. (a) Methodology to define local parametrizations mapping partitions of a surface defined by a set of points. (b) Application of this method to solve the high-order PDE for a nonlinear Kirchhoff-Love shell. (c) Application of the method to describe the molecular conformations of alanine dipeptide, a small molecule. An ensemble of conformations of this molecule samples an underlying configurational manifold homeomorphic to a torus. With the method presented in (a), this manifold is partitioned into four pieces, each of which is parametrized with LME approximants. This method allows us to define smooth collective variables for this system, required to enhance sampling in molecular dynamics simulations.

with multivariate spline techniques. We have demonstrated this by blending maximum entropy approximations to B-spline boundary representations, but another attractive idea is using such blending to resolve the issues associated with irregular vertices in spline patches [22]. Open issues include the systematic formulation of feasible higher-order consistency conditions in multiple dimensions (see the discussion following (1.11)), and the accurate and efficient numerical quadrature of the integrals appearing in Galerkin methods, which involve products of derivatives of the non-polynomial LME basis functions. Finally, the applicability of the LME approach in any space dimension has not yet been exploited in applications.

Bibliography

[1] A. Abdollahi, D. Millán, C. Peco, M. Arroyo, and I. Arias. Revisiting pyramid compression to quantify flexoelectricity: A three-dimensional simulation study. *Phys. Rev. B*, 91:104103, Mar 2015.

[2] A. Abdollahi, C. Peco, D. Millán, M. Arroyo, G. Catalan, and I. Arias. Fracture toughening and toughness asymmetry induced by flexoelectricity. *Phys. Rev. B*, 92:094101, Sep 2015.

[3] M. Arroyo, L. Heltai, D. Millán, and A. DeSimone. Reverse engineering the euglenoid movement. *Proc Natl Acad Sci USA*, 109:17874–17879, 2012.

[4] M. Arroyo and M. Ortiz. Local maximum-entropy approximation schemes: a seamless bridge between finite elements and meshfree methods. *International Journal for Numerical Methods in Engineering*, 65(13):2167–2202, 2006.

[5] M. Arroyo and M. Ortiz. *Meshfree Methods for Partial Differential Equations III*, volume 57 of *Lecture Notes in Computational Science and Engineering*, chapter Local Maximum-Entropy Approximation Schemes, pages 1–16. Springer, 2007.

[6] A. Bompadre, L. E. Perotti, C. J. Cyron, and M. Ortiz. Convergent meshfree approximation schemes of arbitrary order and smoothness. *Computer Methods in Applied Mechanics and Engineering*, 221–222:83–103, 2012.

[7] M. J. Borden, T. J. R. Hughes, C. M. Landis, and C. V. Verhoosel. A higher-order phase-field model for brittle fracture: Formulation and analysis within the isogeometric analysis framework. *Computer Methods in Applied Mechanics and Engineering*, 273:100–118, 2014.

[8] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press: Cambridge, UK, 2004.

[9] C. Cyron, M. Arroyo, and M. Ortiz. Smooth, second order, non-negative meshfree approximants selected by maximum entropy. *International Journal for Numerical Methods in Engineering*, 79(13):1605–1632, 2009.

[10] Q. Du, C. Liu, and X. Wang. A phase field approach in the numerical study of the elastic bending energy for vesicle membranes. *Journal of Computational Physics*, 198:450–468, 2004.

[11] Q. Du and L. Zhu. Analysis of a mixed finite element method for a phase field bending elasticity model of vesicle membrane deformation. *Journal of Computational Mathematics*, 24:265–280, 2006.

[12] T. Fries and H. Matthies. Classification and overview of meshfree methods. Technical report, Institute of Scientific Computing, Technical University Braunschweig, Germany, July 2004.

[13] F. Greco and N. Sukumar. Derivatives of maximum-entropy basis functions on the boundary: Theory and computations. *International Journal for Numerical Methods in Engineering*, 94:1123–1149, 2013.

[14] B. Hashemian, D. Millán, and M. Arroyo. Charting molecular free-energy landscapes with an atlas of collective variables. *The Journal of Chemical Physics*, 145(17), 2016.

[15] T. Hughes. *The Finite Element Method: Lineal Static and Dynamic Finite Element Analysis*. Dover Publications, 2000.

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

[17] E. Jaynes. Information theory and statistical mechanics. *Physics Reviews*, 106:620–630, 1957.

[18] A. Khinchin. *Mathematical Foundations of Information Theory*. Dover, New York, 1957.

[19] B. Li, F. Habbal, and M. Ortiz. Optimal transportation meshfree approximation schemes for fluid and plastic flows. *International Journal for Numerical Methods in Engineering*, 83(12):1541–1579, 2010.

[20] B. Li, C. Peco, D. Millán, I. Arias, and M. Arroyo. Phase-field modeling and simulation of fracture in brittle materials with strongly anisotropic surface energy. *International Journal for Numerical Methods in Engineering*, 102(3-4):711–727, 2015.

[21] S. Li and W. Liu. Meshfree and particle methods and their applications. *Applied Mechanics Reviews*, 55(1):1–34, 2002.

[22] M. Majeed and F. Cirak. Isogeometric analysis using manifold-based smooth basis functions. *Computer Methods in Applied Mechanics and Engineering*, pages –, 2016.

[23] D. Millán, A. Rosolen, and M. Arroyo. Thin shell analysis from scattered points with maximum-entropy approximants. *International Journal for Numerical Methods in Engineering*, 85(6):723–751, 2011.

[24] D. Millán, A. Rosolen, and M. Arroyo. Nonlinear manifold learning for meshfree finite deformation thin shell analysis. *International Journal for Numerical Methods in Engineering*, 93(7):685–713, 2013.

[25] D. Millan, N. Sukumar, and M. Arroyo. Cell-based maximum-entropy approximants. *Computer Methods in Applied Mechanics and Engineering*, 284:712–731, 2015.

[26] K. Nissen, C. J. Cyron, V. Gravemeier, and W. A. Wall. Information-flux method: a meshfree maximum-entropy petrov–galerkin method including stabilised finite element methods. *Computer Methods in Applied Mechanics and Engineering*, 241–244:225 – 237, 2012.

[27] A. Ortiz, M. Puso, and N. Sukumar. Maximum-entropy meshfree method for compressible and near-incompressible elasticity. *Computer Methods in Applied Mechanics and Engineering*, 199(25–28):1859–1871, 2010.

[28] C. Peco, D. Millan, A. Rosolen, and M. Arroyo. Efficient implementation of galerkin meshfree methods for large-scale problems with an emphasis on maximum entropy approximants. *Computers and Structures*, 150:52–62, 2015.

[29] C. Peco, A. Rosolen, and M. Arroyo. An adaptive meshfree method for phase-field models of biomembranes. Part II: a Lagrangian approach for membranes in viscous fluids. *Journal of Computational Physics*, 249:320–336, 2013.

[30] V. Rajan. Optimality of the Delaunay triangulation in R^d. *Discrete and Computational Geometry*, 12(2):189–202, 1994.

[31] A. Rosolen and M. Arroyo. Blending isogeometric analysis and maximum entropy meshfree approximants. *Computer Methods in Applied Mechanics and Engineering*, 264:95–107, 2013.

[32] A. Rosolen, D. Millán, and M. Arroyo. On the optimum support size in meshfree methods: a variational adaptivity approach with maximum entropy approximants. *International Journal for Numerical Methods in Engineering*, 82(7):868–895, 2010.

[33] A. Rosolen, D. Millán, and M. Arroyo. Second order convex *maximum entropy* approximants with applications to high order PDE. *International Journal for Numerical Methods in Engineering*, 94(2):150–182, 2013.

[34] A. Rosolen, C. Peco, and M. Arroyo. An adaptive meshfree method for phase-field models of biomembranes. Part I: approximation with maximum-entropy approximants. *Journal of Computational Physics*, 249:303–319, 2013.

[35] N. Sukumar. Construction of polygonal interpolants: a maximum entropy approach. *International Journal for Numerical Methods in Engineering*, 61(12):2159–2181, 2004.

[36] N. Sukumar. Quadratic maximum-entropy serendipity shape functions for arbitrary planar polygons. *Computer Methods in Applied Mechanics and Engineering*, 263:27–41, 2013.

[37] N. Sukumar and R. Wright. Overview and construction of meshfree basis functions: From moving least squares to entropy approximants. *International Journal for Numerical Methods in Engineering*, 70(2):181–205, 2007.

[38] S. Torabi and J. Lowengrub. Simulating interfacial anisotropy in thin-film growth using an extended Cahn-Hilliard model. *Physical Review E*, 85(4):041603, 2012.