

Maximum Entropy Coordinates for Arbitrary Polytopes

K. Hormann¹ and N. Sukumar²

¹Department of Informatics, Clausthal University of Technology, Germany

²Department of Civil & Environmental Engineering, University of California, Davis, USA

Abstract

Barycentric coordinates can be used to express any point inside a triangle as a unique convex combination of the triangle's vertices, and they provide a convenient way to linearly interpolate data that is given at the vertices of a triangle. In recent years, the ideas of barycentric coordinates and barycentric interpolation have been extended to arbitrary polygons in the plane and general polytopes in higher dimensions, which in turn has led to novel solutions in applications like mesh parameterization, image warping, and mesh deformation. In this paper we introduce a new generalization of barycentric coordinates that stems from the maximum entropy principle. The coordinates are guaranteed to be positive inside any planar polygon, can be evaluated efficiently by solving a convex optimization problem with Newton's method, and experimental evidence indicates that they are smooth inside the domain. Moreover, the construction of these coordinates can be extended to arbitrary polyhedra and higher-dimensional polytopes.

Categories and Subject Descriptors (according to ACM CCS): G.1.1 [Numerical Analysis]: Interpolation formulas G.1.6 [Numerical Analysis]: Constrained optimization I.3.5 [Computer Graphics]: Geometric algorithms

1. Introduction

Barycentric coordinates were first introduced by Möbius as a special kind of homogeneous coordinates with respect to the vertices of a simplex [Möb27]. While unique for simplices, they can be generalized in several ways to arbitrary polygons [Wac75, Flo03, MLD05, FHK06], polyhedra [FKR05, JSW05, LBS06], higher dimensional polytopes [War96, JLW07], and even curves [Bel06, SJW07, DF08].

Let $\Omega \subset \mathbb{R}^d$ be an arbitrary polytope (e.g., a polygon in \mathbb{R}^2 or a polyhedron in \mathbb{R}^3), with vertices v_1, \dots, v_n . The functions $b_i: \Omega \rightarrow \mathbb{R}$, $i = 1, \dots, n$ are called *barycentric coordinates* with respect to Ω if they form a *partition of unity*,

$$\sum_{i=1}^n b_i(v) = 1, \quad (1)$$

allow to write any point $v \in \Omega$ as an *affine combination* of the vertices,

$$\sum_{i=1}^n b_i(v)v_i = v, \quad (2)$$

and satisfy the *Lagrange property*

$$b_i(v_j) = \delta_{ij}. \quad (3)$$

Due to the property in (3), barycentric coordinates can be used as basis functions for *barycentric interpolation*: indeed, it is clear that the function

$$f(v) = \sum_{i=1}^n b_i(v)f_i \quad (4)$$

interpolates the data f_i at the vertices v_i for $i = 1, \dots, n$, and properties (1) and (2) further guarantee the reproduction of affine functions by barycentric interpolation. Obviously, the interpolant f inherits the smoothness from the functions b_i and if the b_i can be evaluated efficiently, then so can f . Barycentric interpolation has many useful applications, ranging from Gouraud and Phong shading, rendering of quadrilaterals [HT04], image warping [HF06, WSHD07], and mesh deformation [JSW05, LKCOL07, JMD*07, LS08] to generalized Bézier surfaces [LD89, LS07] and finite element applications [AO06, SM06, WBG07, MP07, TS08].

Many of these applications require or at least benefit from the barycentric coordinates being *non-negative*,

$$b_i(v) \geq 0, \quad (5)$$

so that (2) and (4) become *convex combinations* and so $f(v)$ is guaranteed to lie inside the convex hull of the data f_i .

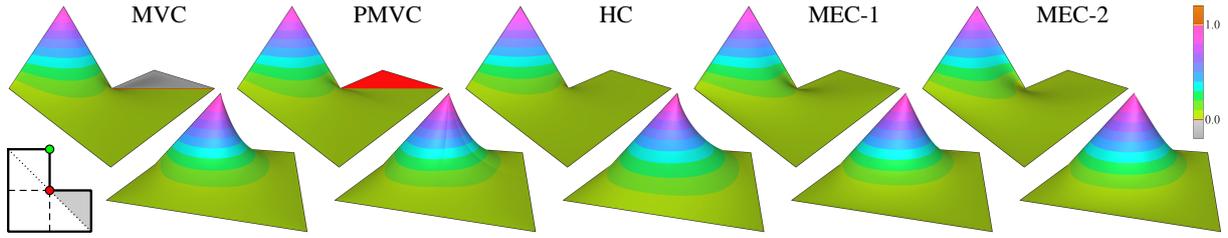


Figure 1: Barycentric coordinates for the green (top row) and the red vertex (bottom row) of an L-shaped polygon. Note that the MVC for the green vertex is negative inside the grey region of the polygon and that the PMVC for that vertex is constant zero there. Moreover, PMVC are only C^0 along the dashed lines and C^1 along the dotted lines.

1.1. Related Work

If Ω is a simplex (e.g., a triangle in \mathbb{R}^2 or a tetrahedron in \mathbb{R}^3), then the barycentric coordinates are uniquely determined by conditions (1) and (2), and they automatically satisfy properties (3) and (5). For general polytopes, this uniqueness breaks down and the non-negativity is not always guaranteed by the several constructions that have been proposed in the past.

Wachspress [Wac75] was the first to come up with a generalization of barycentric coordinates for finite element applications. These *Wachspress coordinates* as well as their higher-dimensional extensions [War96, WSHD07] satisfy (5) as long as Ω is *convex* and can further be evaluated efficiently [MLBD02]. For planar polygons, alternative generalizations are the *discrete harmonic coordinates* [PP93, EDD*95] and the *metric coordinates* [MLD05, SM06], but they both can be negative even inside a convex polygon.

A major breakthrough came with the advent of *mean value coordinates* (MVC) that were discovered by Floater [Flo03] in the context of mesh parameterization and later generalized to polyhedra [FKR05, JSW05, LBS07]. In contrast to Wachspress coordinates, they are well-defined even if Ω is non-convex [HF06], but they can take on negative values then (see Figure 1).

By modifying the transfinite description of mean value coordinates [JSW05, Bel06], Lipman et al. [LKCOL07] were able to overcome this drawback, which significantly improves the results in applications like mesh deformation. Their *positive mean value coordinates* (PMVC) no longer have a simple closed form, but they can still be evaluated efficiently with the GPU. However, these coordinates are only piecewise smooth (see Figure 1).

Up to now, the only known barycentric coordinates that are smooth and non-negative for arbitrary polytopes are the *harmonic coordinates* (HC), which have first been mentioned by Floater et al. [FHK06] and later realized by Joshi et al. [JMD*07] for animating characters. The drawback of harmonic coordinates is that they are rather costly to evaluate because they require to compute the solution of Laplace’s equation subject to suitable Dirichlet boundary conditions.

Another approach for constructing generalized barycentric coordinates that has been suggested independently by Sukumar [Suk04] and Arroyo and Ortiz [AO06] is based on Jaynes’s principle of maximum entropy [Jay57]. By construction, these coordinates are non-negative and always satisfy conditions (1) and (2), but the Lagrange property (3) holds only if Ω is *strictly convex* (see Figure 2).

1.2. Contribution

In this paper we show how to adapt the maximum entropy approach in order to get non-negative barycentric coordinates for arbitrary polygons that also satisfy the Lagrange property and can thus be used for barycentric interpolation. In contrast to PMVC, these new coordinates are smooth and unlike HC, they can be evaluated directly.

After a brief introduction to informational entropic measures and the maximum entropy formalism (Section 2), we review how to derive barycentric coordinates from *prior estimates* by maximizing the Shannon-Jaynes entropy with respect to linear constraints (Section 3). We then present two choices of appropriate prior functions (Section 4) that yield *maximum entropy coordinates* (MEC) with all the desired properties for arbitrary polygons (see Figure 1) and explain how to extend the construction to higher dimensions. We further describe how Newton’s method can be used to efficiently evaluate these coordinates (Section 5). Finally, we compare MEC with previous constructions (MVC, PMVC, and HC) by studying some application examples (Section 6) and discuss their limitations as well as interesting open questions for future research (Section 7).

2. Principle of Maximum Entropy

Shannon [Sha48] introduced the concept of entropy as a measure of uncertainty in information theory, with an eye on its applications in communication theory. The *Shannon entropy* of a discrete probability distribution is

$$H(p) = \langle -\ln p \rangle = - \sum_{i=1}^n p_i \ln p_i, \quad (6)$$

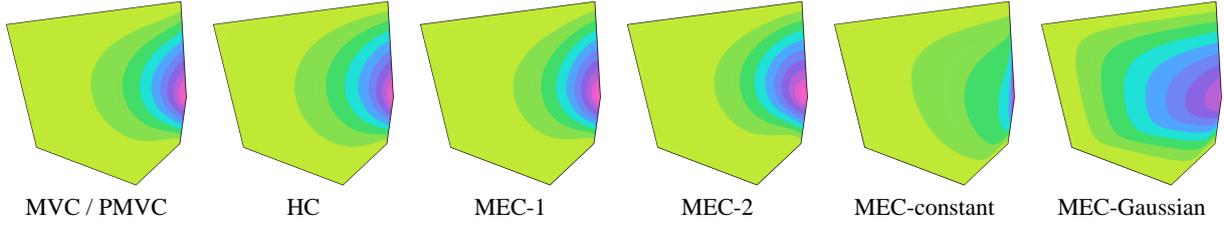


Figure 2: Barycentric coordinates for the rightmost vertex of this convex polygon. For convex polygons, MVC and PMVC are the same, and the maximum entropy coordinates based on constant [Suk04] and Gaussian priors [AO06] (with $\beta = 5$ in this example) satisfy the Lagrange property. Note that the coordinate derived from constant priors is very steep near the vertex and would lose the Lagrange property if the polygon was weakly convex and the vertex and its neighbours were collinear.

where $\langle \cdot \rangle$ is the expectation operator, $p_i = p(x_i)$ is the probability of the occurrence of the event x_i , and $p_i \ln p_i \doteq 0$ if $p_i = 0$. Note that the above form of H satisfies the axiomatic requirements of an uncertainty measure [Khi57].

As a means for least-biased statistical inference in the presence of testable (known) constraints, Jaynes used the Shannon entropy to propose the *principle of maximum entropy* [Jay57]. While Jaynes’s initial emphasis was on applications in statistical mechanics, the principle has broader appeal and can be applied to any ill-posed problem that requires inductive inference [Jay03]. In a nutshell, maximizing entropy provides the *least-biased* statistical inference solution when insufficient information is available.

To illustrate this statement, consider a coin toss experiment and let p_H and p_T be the unknown probabilities of heads and tails, respectively. Here, the only known constraint is $p_H + p_T = 1$, which involves two unknowns but only one equation. But if we regularize the problem by maximizing the Shannon entropy (6) subject to this constraint, then the unique maximum entropy solution is $p_H = p_T = 1/2$, which is consistent with our expectations for an unbiased coin.

It was later recognized that for H to be invariant under invertible mappings of the continuous random variable x , the general form of the continuous entropy should be

$$H(p, m) = - \int p(x) \ln \left(\frac{p(x)}{m(x)} \right) dx,$$

where m is called a *p-estimate* or *prior distribution* [KL51, Jay63, SJ80]. In this paper we use the discrete version of the *Shannon-Jaynes entropy functional*:

$$H(p, m) = - \sum_{i=1}^n p_i \ln \left(\frac{p_i}{m_i} \right). \quad (7)$$

In the literature, the quantity $D(p \| m) = -H(p, m)$ is known as the Kullback-Leibler (KL) distance. If the KL-distance is adopted as the objective functional, the variational principle is known as the principle of minimum relative entropy [SJ80]. Obviously, maximizing the Shannon-Jaynes entropy functional $H(p, m)$ is equivalent to minimizing the relative entropy functional $D(p \| m)$.

3. Maximum Entropy Coordinates

Historically, discrete probability measures have been seen as weights and hence their association with the construction of barycentric coordinates is natural. As in condition (1), discrete probability measures sum to one, and condition (2) is the counterpart of the expectation value of the first moment (or mean) of a discrete probability distribution being known.

Sukumar [Suk04] adopted the Shannon entropy (6) to construct non-negative barycentric coordinates for strictly convex polygons, whereas Arroyo and Ortiz [AO06] used a modified entropy functional in the variational principle to derive basis functions for meshfree methods. The modified entropy chosen in [AO06] is a linear combination of Rajan’s functional [Raj94] and the Shannon entropy, and the solution of the variational problem provides a smooth transition from Delaunay interpolation as a limiting case at one end to global maximum entropy approximation at the other end of the spectrum.

Sukumar and Wright [SW07b] later realized that both constructions can be described in a unifying framework that uses the Shannon-Jaynes entropy functional with a prior (7). The variational formulation for maximum entropy coordinates in general then is: find $b = (b_1, \dots, b_n) : \Omega \rightarrow \mathbb{R}_+^n$ as the solution of the constrained optimization problem

$$\max_{b(v) \in \mathbb{R}_+^n} H(b, m), \quad H(b, m) = - \sum_{i=1}^n b_i(v) \ln \left(\frac{b_i(v)}{m_i(v)} \right) \quad (8a)$$

subject to the linear precision conditions

$$\sum_{i=1}^n b_i(v) = 1, \quad (8b)$$

$$\sum_{i=1}^n b_i(v)(v_i - v) = 0 \quad (8c)$$

for any $v \in \Omega$. In (8a), \mathbb{R}_+^n is the non-negative orthant and $m_i : \Omega \rightarrow \mathbb{R}_+$ is a prior estimate for b_i . Note that if Ω is a simplex in \mathbb{R}^d , then $n = d + 1$ and the linear constraints yield a unique solution for the b_i . For $n > d + 1$, which is the case of interest in this paper, the linear constraints form an under-determined system.

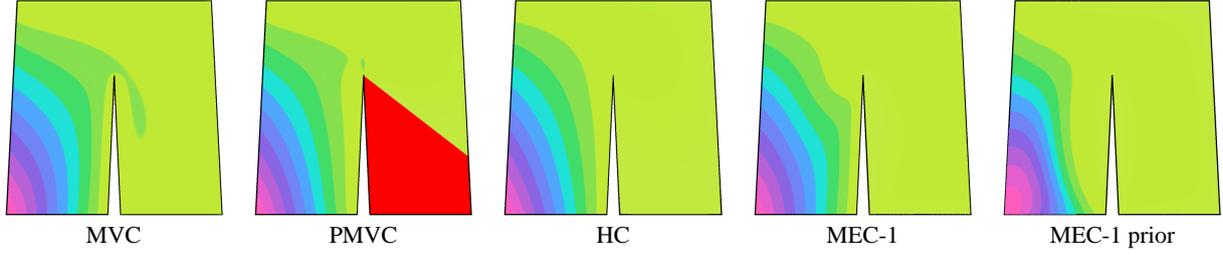


Figure 3: Barycentric coordinates for the lower left vertex of this non-convex polygon and prior estimate of MEC-1. In this example, all functions are between zero and one. Note that PMVC is constant zero over the region shaded in red and also attains a local maximum just above the vertex in the centre.

In the present context, the prior functions can be seen as weight functions associated with each vertex v_i , and the variational principle provides a *correction* that modifies the weight functions m_i in a minimal (i.e., least-biased) way to form coordinates b_i that satisfy the constant and linear precision conditions. For non-negative m_i , the objective functional $H(b, m)$ is strictly concave, and therefore the problem posed in (8) admits a unique non-negative solution for b_i . Note that this construction of barycentric coordinates is also valid for points v inside the convex hull $\Omega_c = [v_1, \dots, v_n]$ of the polytope as long as the priors are well-defined over Ω_c .

It now turns out that the coordinates of Sukumar [Suk04] are the solutions of (8) for constant priors $m_i(v) \equiv 1$ and that Gaussian prior functions $m_i(v) = \exp(-\beta \|v_i - v\|^2)$ result in the coordinates proposed by Arroyo and Ortiz [AO06] (see Figure 2). We further note that one obtains $b_i(v) = m_i(v)$ for $i = 1, \dots, n$ if the priors m_i a priori satisfy the linear constraints. Moreover, if we omit the linear precision condition (8c) and use just the constant precision condition (8b) as a constraint, then the solution of (8) is $b_i(v) = m_i(v) / \sum_{j=1}^n m_j(v)$, which we recognize as having the same form as the Shepard function [She68].

The key ingredient in the maximum entropy formulation is the selection of the prior functions m_i , and this flexibility presents the possibility of designing tailored barycentric coordinates. The next section describes the specific choices for prior functions that we make to construct non-negative barycentric coordinates for arbitrary polygons and polyhedra.

4. Prior Functions

To obtain non-negative barycentric coordinates for an arbitrary polygon Ω , we first construct non-negative priors m_i that possess the desired boundary properties. We then use these within the entropy formulation to obtain functions b_i that are non-negative, interpolate on the boundary of Ω , and are linearly precise.

To this end, let e_i be the edge of the polygon between v_i and v_{i+1} and consider the *edge weight function*

$$\rho_i(v) = \|v - v_i\| + \|v - v_{i+1}\| - \|v_{i+1} - v_i\| \quad (9)$$

that vanishes along e_i and is positive elsewhere, by virtue of the triangle inequality. Note that ρ_i is also used in the construction of metric coordinates [MLD05, SM06]. It then follows that the product

$$\pi_i(v) = \prod_{j \neq i-1, i} \rho_j(v)$$

is non-negative and vanishes on all edges that are not adjacent to vertex v_i . The same clearly holds for the prior functions m_i that we derive from the π_i through normalization,

$$m_i(v) = \frac{\pi_i(v)}{\sum_{j=1}^n \pi_j(v)}.$$

Note that after dividing both the numerator and the denominator by the product of all ρ_j , we get the equivalent form

$$m_i(v) = \frac{\tilde{\pi}_i(v)}{\sum_{j=1}^n \tilde{\pi}_j(v)} \quad \text{with} \quad \tilde{\pi}_i(v) = \frac{1}{\rho_{i-1}(v)\rho_i(v)}, \quad (10)$$

which renders it amenable to stable numerical computations.

Now solving the optimization problem (8) with these m_i as input yields functions b_i with all the desired properties. Due to the linear constraints (8b) and (8c), the b_i clearly satisfy conditions (1) and (2) and as the m_i are non-negative, then so are the b_i . It remains to be shown that the b_i inherit the Lagrange property from the prior functions m_i . If $m_i(v^*) = 0$ for some $i \in \{1, \dots, n\}$ and $v^* \in \Omega$, then solving (8) gives $b_i(v^*) = 0$ because

$$\lim_{v \rightarrow v^*} b_i(v) \ln \left(\frac{b_i(v)}{m_i(v)} \right)$$

is zero if $b_i(v^*) = 0$ and diverges otherwise. Therefore, the functions b_i have the same zero sets as their prior estimates m_i , and the only function that does not vanish at vertex v_i is b_i . The Lagrange property (3) now follows because of (1) and we similarly conclude that the b_i are linear on the edges of Ω . Hence, the functions b_i are non-negative barycentric coordinates. We call them ‘‘MEC-1’’ and Figure 3 shows an example of such a function b_i and its prior estimate m_i .

We would like to point out that this construction of barycentric coordinates is rather general and also works for any other choice of edge weight functions as long as ρ_i is

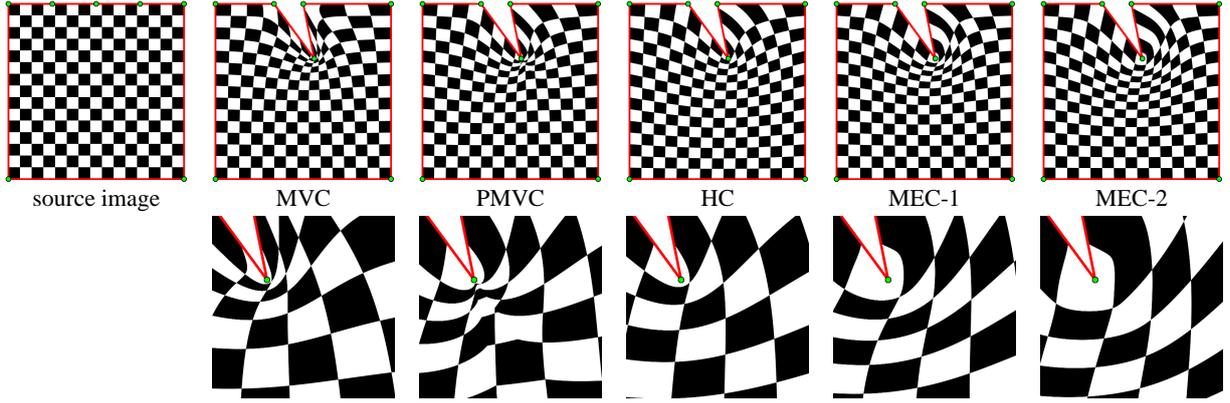


Figure 4: Image warping using different barycentric coordinates and backward mapping (cf. [HF06]). The bottom row shows a close-up of the region around the concave vertex of the target polygon. Note that the result using PMVC is only C^0 .

non-negative and vanishes along e_i . For example, using the alternative edge weight functions

$$\rho_i(v) = \|v - v_i\| \cdot \|v - v_{i+1}\| + (v - v_i) \cdot (v - v_{i+1})$$

in (10) gives the barycentric coordinates that we refer to as “MEC-2” in our examples.

Moreover, the construction can be extended to higher dimensions. For example, if Ω is a polyhedron in \mathbb{R}^3 with convex planar faces, then we first define for each face f_i a function ρ_i that vanishes on f_i and is positive elsewhere. If f_i has k vertices v_{i_1}, \dots, v_{i_k} , then the analogue of the edge weight function (9) is the *face weight function*

$$\rho_i(v) = \sum_{j=1}^k A(v_{i_j}, v_{i_{j+1}}, v) - A(v_{i_1}, \dots, v_{i_k}),$$

where $A(v_1, \dots, v_m)$ denotes the area of the polygon with vertices v_1, \dots, v_m . The prior function m_i for any vertex v_i of Ω is then computed as in (10) with the product in the denominator of $\tilde{\pi}_i$ ranging over all faces adjacent to v_i .

5. Numerical Algorithm

In order to use the proposed maximum entropy coordinates for practical applications, it is essential to be able to efficiently solve the constrained optimization problem in (8). We resort to the method of Lagrange multipliers to first convert (8) into an unconstrained problem and then use Newton’s method to solve it.

Let $\lambda_0 \in \mathbb{R}$ be the Lagrange multiplier for the constraint (8b) and $\lambda \in \mathbb{R}^d$ be the Lagrange multipliers for the d constraints (8c). Then the Lagrangian for problem (8) is

$$\begin{aligned} \mathcal{L}(b; \lambda_0, \lambda) = & \sum_{i=1}^n -b_i(v) \ln \left(\frac{b_i(v)}{m_i(v)} \right) - \lambda_0 \left(\sum_{i=1}^n b_i(v) - 1 \right) \\ & - \lambda \cdot \left(\sum_{i=1}^n b_i(v) \tilde{v}_i \right), \end{aligned}$$

where $\tilde{v}_i = v_i - v$. On setting the first variation of \mathcal{L} to zero, namely $\delta \mathcal{L}(b; \lambda_0, \lambda) = 0$, we obtain

$$\left[-1 - \ln \left(\frac{b_i(v)}{m_i(v)} \right) - \lambda_0 - \lambda \cdot \tilde{v}_i \right] \delta b_i(v) = 0$$

for $i = 1, \dots, n$, and since the variation $\delta b_i(v)$ is arbitrary, the term within the bracket must be identically equal to zero:

$$-1 - \ln \left(\frac{b_i(v)}{m_i(v)} \right) - \lambda_0 - \lambda \cdot \tilde{v}_i = 0.$$

Therefore,

$$b_i(v) = \frac{m_i(v) \exp(-\lambda \cdot \tilde{v}_i)}{Z},$$

where the substitution $\ln Z = 1 + \lambda_0$ has been made (Z is known as the *partition function* in statistical mechanics). Now, on using condition (8b), we obtain [SW07b]

$$\begin{aligned} b_i(v) &= \frac{Z_i(\lambda)}{Z(\lambda)}, \\ Z_i(\lambda) &= m_i(v) \exp(-\lambda \cdot \tilde{v}_i), \\ Z(\lambda) &= \sum_{j=1}^n Z_j(\lambda). \end{aligned} \quad (11)$$

Note that λ , $Z_i(\lambda)$, and $Z(\lambda)$ implicitly depend on v and that once the Lagrange multipliers $\lambda = (\lambda_1, \dots, \lambda_d)$ are determined, then $b_i(v)$ can be obtained from (11). To compute λ we note that the $b_i(v)$ in (11) must satisfy (8c), which leads to the d non-linear equations

$$\frac{1}{Z(\lambda)} \sum_{i=1}^n Z_i(\lambda) \tilde{v}_i = 0 \quad (12)$$

in the d unknowns $\lambda_1, \dots, \lambda_d$. The solution of (12) is equivalent to solving the dual unconstrained optimization problem [BV04]

$$\lambda^* = \operatorname{argmin} F(\lambda), \quad F(\lambda) = \ln Z(\lambda),$$

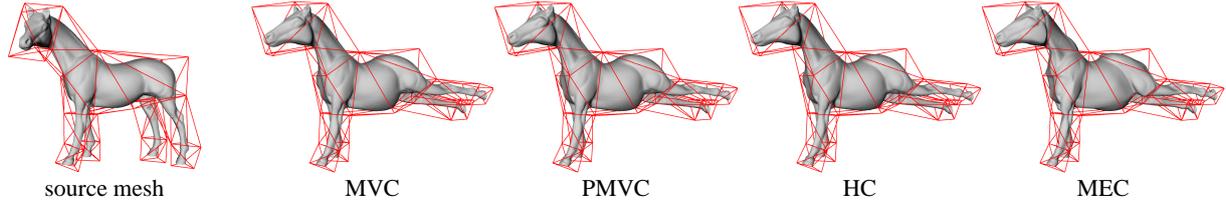


Figure 5: Mesh deformation using different barycentric coordinates and a control polyhedron with triangular faces. The results for MVC, PMVC, and HC are taken from [LKCOL07].

where λ^* is the optimal solution. Since F is strictly convex in Ω , Newton’s method is the natural choice. The steps in the Newton algorithm to compute the maximum entropy coordinates $b_i(v)$ of any $v \in \Omega$ are as follows:

1. For given v , compute and store $\tilde{v}_i = v_i - v$; also functions to compute the prior functions $m_i(v)$ are available;
2. Start with iteration counter $k = 0$, the initial guess $\lambda^0 = 0$, and let ϵ be the convergence tolerance. The convergence tolerance dictates to what accuracy the linear precision conditions are satisfied. A value of ϵ in the range 10^{-3} to 10^{-10} is suitable;
3. Compute $g^k = \nabla_{\lambda} F(\lambda^k)$ and $H^k = \nabla_{\lambda} \nabla_{\lambda} F(\lambda^k)$, which are the gradient and Hessian of F , respectively;
4. Determine Newton search direction $\Delta\lambda^k = -(H^k)^{-1} g^k$;
5. Update: $\lambda^{k+1} = \lambda^k + \alpha \Delta\lambda^k$, where α is the step size. For Newton’s method (*damped* or *guarded*), a line search algorithm [BF04] is used to determine α if the error is greater than 10^{-4} , otherwise α is set to unity;
6. Check convergence: if $\|g^{k+1}\| > \epsilon$, then increment the iteration counter k and goto 3, else continue;
7. Set $\lambda^* = \lambda^{k+1}$ and compute $b_i(v)$ from Equation (11).

Due to the quadratic convergence of Newton’s method, only 3 to 7 iterations are needed to obtain an accuracy of 10^{-10} .

6. Results and Discussion

The advantages of MEC over previous barycentric coordinates are threefold and we present several examples to illustrate them.

Firstly, MEC are non-negative for arbitrary polytopes, which is important as it guarantees the convex hull property of the barycentric interpolation (4). PMVC and HC share this property with MEC, but MVC can be negative inside a non-convex domain (see Figure 1). In some applications, this can lead to undesired results: for example, the interpolation of colour values can yield values outside the valid range [HF06] and it can lead to severe artefacts in mesh deformation in some extreme cases [LKCOL07]. However, in many situations, the influence of negative weights is not noticeable (see Figure 4).

Secondly, like MVC and HC, MEC are smooth inside the domain, which in turn leads to well-behaved results in

m	MVC	PMVC	HC	MEC	
				$\epsilon = 10^{-5}$	$\epsilon = 10^{-10}$
50 K	0.07	0.33	2.44	4.52	5.09
100 K	0.15	0.67	6.75	8.99	10.5
200 K	0.32	1.42	17.3	18.7	24.3
400 K	0.59	2.61	195	34.6	39.7

Table 1: Timings (in sec.) for evaluating the barycentric interpolant at m points inside the L-shaped polygon with $n = 6$ vertices from Figure 1.

applications that build on barycentric interpolation. PMVC, however, are only piecewise smooth and have discontinuous derivatives along certain lines inside a non-convex polygon (see Figure 1), and similarly inside non-convex polyhedra. In particular, this happens along the lines defined by a concave vertex and its two neighbours and can lead to visible “kinks” if used for image warping (for example, see Figure 4).

Thirdly, MEC can be evaluated *locally* at any point $v \in \Omega$, though not as efficiently as MVC or PMVC, and the cost scales reasonably with the desired accuracy. In contrast, the evaluation of HC requires the solution of a *global* approximation problem, and the computational cost substantially increases if high accuracy is desired.

For planar polygons, we implemented all methods considered in this paper in the following way. For MVC, we use the pseudo code suggested in [HF06], and for PMVC we first determine the segments of Ω that are “visible” from some point $v \in \Omega$ and then use the MVC formulas, restricted to these segments. In order to evaluate HC, we use *triangle* [She02] to triangulate Ω with m vertices and *TAUCS* [Tol03] to solve the linear system arising from the piecewise linear finite element discretization of the Laplacian. The details of our MEC implementation are described in Section 5.

The timings in Table 1 were measured on a 2 GHz Intel Pentium M with 1 GB of RAM and confirm that the cost for computing MVC, PMVC, and MEC at m points grows linearly with m and slightly worse for HC, because the latter requires to solve a linear system with a sparse system matrix of size $m \times m$. The significantly larger constant of proportionality for HC in the last row is due to the fact that the linear system solver required more memory than available as RAM and thus started swapping data to the hard disk.

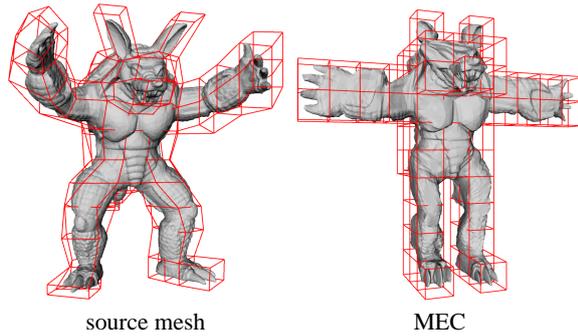


Figure 6: Mesh deformation using a control polyhedron with quadrilateral faces.

However, we point out that HC can be computed more efficiently by using boundary element methods [Rus08], which requires solving a linear system with a dense $k \times k$ system matrix (with $n < k \ll m$, where n and m are the numbers of polygon vertices and interior points, respectively). The timings also show that it is not very costly to increase the accuracy of MEC due to the quadratic convergence rate of Newton's method.

We also implemented MEC for polyhedra in \mathbb{R}^3 and Figures 5 and 6 show the results of using them for mesh deformation. It took about 100 sec. to compute the $n = 51$ MEC of the $m = 48485$ vertices of the horse, and about 140 sec. for the $n = 218$ MEC of the $m = 15002$ vertices of the armadillo, confirming again that the evaluation cost depends linearly both on n and m . Note that MEC can also be used if the control polyhedron consists of quadrilateral faces, as shown in Figure 6. Moreover, the mesh to be deformed is not required to be inside the control polyhedron but only inside its convex hull (see Section 3), which is the case in this example.

7. Conclusions

Maximum entropy coordinates offer a new way of generalizing barycentric coordinates to arbitrary polytopes. By construction, MEC are non-negative and have affine precision, and we have shown that they also satisfy the Lagrange property as long as the corresponding prior functions have the correct zero set. Due to this flexibility, there is a lot of room for improvement in designing "good" prior functions. For example, it might be possible to manipulate the derivative of the barycentric coordinates in a prescribed way by modifying the priors so as to allow for Hermite interpolation as in [LS08].

A good prior will probably need to account for the local geometry of the polytope, and it should also have local support. The construction in (10) seems to be a good recipe for constructing priors and in essence, any function that measures the distance to an edge (or a face in \mathbb{R}^3) in some way

can be used as a weight function ρ_i . Due to the close connection between distance fields and level set functions, we believe that it is worth investigating the possible use of level sets for the design of prior functions in future work.

Although confirmed by the many numerical results that we ran, it also remains to prove the smoothness of MEC. For the MEC, based on Gaussian prior functions, Arroyo and Ortiz [AO06] could prove smoothness, and it is generally assumed that the MEC are as smooth as the prior functions. A first step towards a proof has been taken by Sukumar and Wets [SW07a] who established C^0 -continuity of MEC for any set of C^k prior functions, $k \geq 0$.

We would finally like to mention again that MEC are only defined inside the convex hull of the polytope Ω and not everywhere in \mathbb{R}^d like MVC and HC (see [Rus08]), because the optimization problem (8) does not have a feasible non-negative solution at points outside the convex hull.

Acknowledgements

The authors would like to thank Johannes Kopf for kindly providing us with the MVC, PMVC, and HC data for the example in Figure 5. This work was partially supported by the US NSF through contract grant CMMI-0626481.

References

- [AO06] ARROYO M., ORTIZ M.: Local *maximum-entropy* approximation schemes: a seamless bridge between finite elements and meshfree methods. *Int. J. Numer. Methods Eng.* 65, 13 (Mar. 2006), 2167–2202.
- [Bel06] BELYAEV A.: On transfinite barycentric coordinates. In *Geometry Processing* (2006), Eurographics Symposium Proceedings, Eurographics Association, pp. 89–99.
- [BF04] BURDEN R. L., FAIRES J. D.: *Numerical Analysis*, eighth ed. Brooks/Cole, Belmont, CA, 2004.
- [BV04] BOYD S., VANDENBERGHE L.: *Convex Optimization*. Cambridge University Press, Cambridge, UK, 2004.
- [DF08] DYKEN C., FLOATER M. S.: Transfinite mean value interpolation. *Comput. Aided Geom. Des.* (2008). To appear.
- [EDD*95] ECK M., DE ROSE T., DUCHAMP T., HOPPE H., LOUNSBERRY M., STUETZLE W.: Multiresolution analysis of arbitrary meshes. In *Proceedings of SIGGRAPH 95* (1995), ACM Press, pp. 173–182.
- [FHK06] FLOATER M. S., HORMANN K., KÓS G.: A general construction of barycentric coordinates over convex polygons. *Adv. Comput. Math.* 24, 1–4 (Jan. 2006), 311–331.
- [FKR05] FLOATER M. S., KÓS G., REIMERS M.: Mean value coordinates in 3D. *Comput. Aided Geom. Des.* 22, 7 (Oct. 2005), 623–631.
- [Flo03] FLOATER M. S.: Mean value coordinates. *Comput. Aided Geom. Des.* 20, 1 (Mar. 2003), 19–27.
- [HF06] HORMANN K., FLOATER M. S.: Mean value coordinates for arbitrary planar polygons. *ACM Trans. Graph.* 25, 4 (Oct. 2006), 1424–1441.

- [HT04] HORMANN K., TARINI M.: A quadrilateral rendering primitive. In *Graphics Hardware* (2004), Eurographics Symposium Proceedings, Eurographics Association, pp. 7–14.
- [Jay57] JAYNES E. T.: Information theory and statistical mechanics. *Phys. Rev.* 106, 4 (May 1957), 620–630.
- [Jay63] JAYNES E. T.: Information theory and statistical mechanics. In *Statistical Physics*, Ford K. W., (Ed.), vol. 3 of *1962 Brandeis Lectures in Theoretical Physics*. W. A. Benjamin, New York, 1963, pp. 181–218.
- [Jay03] JAYNES E. T.: *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, UK, 2003.
- [JLW07] JU T., LIEPA P., WARREN J.: A general geometric construction of coordinates in a convex simplicial polytope. *Comput. Aided Geom. Des.* 24, 3 (Apr. 2007), 161–178.
- [JMD*07] JOSHI P., MEYER M., DEROSE T., GREEN B., SANOCKI T.: Harmonic coordinates for character articulation. *ACM Trans. Graph.* 26, 3 (July 2007), 71/1–71/9. Proceedings of ACM SIGGRAPH 2007.
- [JSW05] JU T., SCHAEFER S., WARREN J.: Mean value coordinates for closed triangular meshes. *ACM Trans. Graph.* 24, 3 (July 2005), 561–566. Proceedings of ACM SIGGRAPH 2005.
- [Khi57] KHINCHIN A.: *Mathematical Foundations of Information Theory*. Dover, New York, NY, 1957.
- [KL51] KULLBACK S., LEIBLER R. A.: On information and sufficiency. *Ann. Math. Statist.* 22, 1 (Mar. 1951), 79–86.
- [LBS06] LANGER T., BELYAEV A., SEIDEL H.-P.: Spherical barycentric coordinates. In *Geometry Processing* (2006), Eurographics Symposium Proceedings, Eurographics Association, pp. 81–88.
- [LBS07] LANGER T., BELYAEV A., SEIDEL H.-P.: Mean value coordinates for arbitrary spherical polygons and polyhedra in \mathbb{R}^3 . In *Curve and Surface Design: Avignon 2006*, Chenin P., Lyche T., Schumaker L. L., (Eds.), Modern Methods in Applied Mathematics. Nashboro Press, Brentwood, TN, 2007, pp. 193–202.
- [LD89] LOOP C. T., DEROSE T. D.: A multisided generalization of Bézier surfaces. *ACM Trans. Graph.* 8, 3 (July 1989), 204–234.
- [LKCOL07] LIPMAN Y., KOPF J., COHEN-OR D., LEVIN D.: GPU-assisted positive mean value coordinates for mesh deformations. In *Geometry Processing* (2007), Eurographics Symposium Proceedings, Eurographics Association, pp. 117–123.
- [LS07] LANGER T., SEIDEL H.-P.: Mean value Bézier surfaces. In *Mathematics of Surfaces XII*, Martin R., Sabin M., Winkler J., (Eds.), vol. 4647 of *Lecture Notes in Computer Science*. Springer, 2007, pp. 263–274. Proceedings of the 12th IMA International Conference.
- [LS08] LANGER T., SEIDEL H.-P.: Higher order barycentric coordinates. *Comput. Graph. Forum* 27, 2 (Apr. 2008), 459–466. Proceedings of Eurographics 2008.
- [MLBD02] MEYER M., LEE H., BARR A., DESBRUN M.: Generalized barycentric coordinates on irregular polygons. *J. Graph. Tools* 7, 1 (2002), 13–22.
- [MLD05] MALSCH E. A., LIN J. J., DASGUPTA G.: Smooth two dimensional interpolants: a recipe for all polygons. *J. Graph. Tools* 10, 2 (2005), 27–39.
- [Möb27] MÖBIUS A. F.: *Der barycentrische Calcul*. Johann Ambrosius Barth, Leipzig, 1827.
- [MP07] MILBRADT P., PICK T.: Polytope finite elements. *Int. J. Numer. Methods Eng.* 73, 12 (July 2007), 1811–1835.
- [PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Exp. Math.* 2, 1 (1993), 15–36.
- [Raj94] RAJAN V. T.: Optimality by the Delaunay triangulation in \mathbb{R}^d . *Discrete Comput. Geom.* 12, 1 (Dec. 1994), 189–202.
- [Rus08] RUSTAMOV R. M.: *Boundary Element Formulation of Harmonic Coordinates*. Tech. rep., Department of Mathematics, Purdue University, 2008.
- [Sha48] SHANNON C. E.: A mathematical theory of communication. *Bell Systems Technical Journal* 27 (July, Oct. 1948), 379–423, 623–656.
- [She68] SHEPARD D.: A two-dimensional interpolation function for irregularly-spaced data. In *Proceedings of the 23rd ACM National Conference* (1968), ACM Press, pp. 517–524.
- [She02] SHEWCHUK J. R.: Delaunay refinement algorithms for triangular mesh generation. *Comput. Geom.* 22, 1–3 (May 2002), 21–74.
- [SJ80] SHORE J. E., JOHNSON R. W.: Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Trans. Inf. Theory* 26, 1 (Jan. 1980), 26–37.
- [SJW07] SCHAEFER S., JU T., WARREN J.: A unified, integral construction for coordinates over closed curves. *Comput. Aided Geom. Des.* 24, 8–9 (Nov. & Dec. 2007), 481–493.
- [SM06] SUKUMAR N., MALSCH E. A.: Recent advances in the construction of polygonal finite element interpolants. *Arch. Comput. Methods Eng.* 13, 1 (Mar. 2006), 129–163.
- [Suk04] SUKUMAR N.: Construction of polygonal interpolants: a maximum entropy approach. *Int. J. Numer. Methods Eng.* 61, 12 (Nov. 2004), 2159–2181.
- [SW07a] SUKUMAR N., WETS R. J.-B.: Deriving the continuity of maximum-entropy basis functions via variational analysis. *SIAM J. Optim.* 18, 3 (2007), 914–925.
- [SW07b] SUKUMAR N., WRIGHT R. W.: Overview and construction of meshfree basis functions: from moving least squares to entropy approximants. *Int. J. Numer. Methods Eng.* 70, 2 (Apr. 2007), 181–205.
- [Tol03] TOLEDO S.: TAUCS: A library of sparse linear solvers. <http://www.tau.ac.il/~stoledo/taucs/>, 2003.
- [TS08] TABARRAEI A., SUKUMAR N.: Extended finite element method on polygonal and quadtree meshes. *Comput. Meth. Appl. Mech. Eng.* 197, 5 (Jan. 2008), 425–438.
- [Wac75] WACHSPRESS E. L.: *A Rational Finite Element Basis*, vol. 114 of *Mathematics in Science and Engineering*. Academic Press, 1975.
- [War96] WARREN J.: Barycentric coordinates for convex polytopes. *Adv. Comput. Math.* 6, 2 (1996), 97–108.
- [WBG07] WICKE M., BOTSCH M., GROSS M.: A finite element method on convex polyhedra. *Comput. Graph. Forum* 26, 3 (Sept. 2007), 355–364. Proceedings of Eurographics 2007.
- [WSHD07] WARREN J., SCHAEFER S., HIRANI A. N., DESBRUN M.: Barycentric coordinates for convex sets. *Adv. Comput. Math.* 27, 3 (Oct. 2007), 319–338.