Closed-Form Gaussian Mixture Representations of d-Dimensional Triangulated Domains

Ali Darijani[⊛□#¶⊥], Benedikt Stratmann[#¶⊥], Uwe D. Hanebeck[⊛¶⊥], Ratan Bahadur Thapa[¶⊥], Jürgen Beyerer[⊛⊥], Nehal Afifi[⊛⊥],

***Abstract*— We derive a Gaussian Mixture representation of a full-dimensional subset of Euclidean space with finite volume from a finite simplicial triangulation in closed form. On each simplex, we replace the uniform probability density function with a Gaussian density that minimizes the Kullback–Leibler divergence, arriving at the information projection onto the family of Gaussian distributions. By exploiting affine isomorphisms between simplices, the joint convexity and bijection invariance properties of the Kullback–Leibler divergence, identities for Dirichlet distributions, and a partition-of-unity argument, we establish a dimension-dependent but geometry-independent sharp upper bound on the global Kullback–Leibler divergence between the resulting Gaussian Mixture and the uniform distribution on the original domain. Numerical examples illustrate the construction and demonstrate its potential in practical applications.**

I. Introduction

Many engineering and physical systems are geometric before they are probabilistic. Objects are defined by their shape, environments by spatial maps, and mechanical systems by admissible configurations. At the same time, uncertainty is unavoidable due to noisy measurements, incomplete knowledge, and variability. A natural question is therefore how a known geometric description can be converted directly into a probabilistic representation suitable for estimation, fusion, and decision making.

*Gaussian Mixture Model*s (GMMs) provide a flexible representation of probability densities and positive measures and can be interpreted as basis functions for such representations [1], [2]. However, most existing GMM construction methods solve a different problem: they infer a mixture from samples or a discretized representation of an unknown distribution. Expectation maximization (*Expectation-Maximization* (EM)) [3]–[5], Bayesian approaches, and grid-based density representations are powerful in this setting, but require additional choices concerning samples or grids, initialization, optimization, stopping criteria, and the number of mixture components. Their iterative or discretization-dependent nature also makes it difficult to obtain a direct geometry-to-probability mapping together with an analytical approximation guarantee in an appropriate probabilistic distance measure. In particular, existence results for GMM approximations do not provide a constructive procedure for obtaining the parameters from a given geometric object [1], [2].

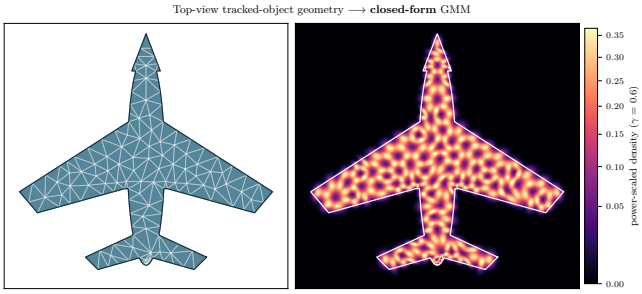


Fig. 1. A fixed-wing planar silhouette is triangulated and directly converted into a Gaussian mixture, with one closed-form component per triangle.

This work considers the complementary, geometry-driven problem: the target distribution is known through a geometric domain, and the objective is to construct its GMM representation directly. Given a uniform distribution over a finite triangulated domain, where each "triangle" is the full-dimensional analogue of a tetrahedron in higher dimensions, we seek a reasonably accurate Gaussian mixture representation. We replace this uniform distribution locally by its Gaussian I-projection and combine the resulting components according to the "triangle" volumes. The Gaussian parameters are obtained directly from the "triangle" vertices, so the complete GMM is constructed in closed form without sampling or iterative optimization. The topological equivalence of "triangles" further allows the local approximation error to be characterized independently of the position, scale, and shape of the "triangles" within the triangulation.

Figure 1 illustrates this direct construction for a planar fixed-wing silhouette.

The resulting representation is particularly useful when geometry is part of the computational model. For example, geometric refinement naturally produces a finer probabilistic representation: in the experiments, barycentric and polyhedral

[⊛]Idea[□]Math[#]Programming[¶]Writing[⊥]Critique

[⊛□#¶]Ali Darijani is affiliated with Fraunhofer IOSB and the Vision and Fusion Laboratory (IES), Karlsruhe Institute of Technology, Karlsruhe, Germany. ali.darijani@rwth-aachen.de

[#¶⊥]Benedikt Stratmann is affiliated with Fraunhofer IOSB, Karlsruhe, Germany. benedikt.stratmann@iosb.fraunhofer.de

[⊛¶⊥]Uwe D. Hanebeck is affiliated with Intelligent Sensor-Actuator-Systems (ISAS), Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany. uwe.hanebeck@kit.edu

[□⊥]Ratan Bahadur Thapa is affiliated with the Institute for Artificial Intelligence, Analytical Computing, University of Stuttgart, Stuttgart, Germany. ratan.thapa@ki.uni-stuttgart.de

[⊛⊥]Jürgen Beyerer is affiliated with Fraunhofer IOSB and the Vision and Fusion Laboratory (IES), Karlsruhe Institute of Technology, Karlsruhe, Germany. juergen.beyerer@iosb.fraunhofer.de

[⊛⊥]Nehal Afifi is affiliated with Institute for Production Engineering (IPEK), Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany. nehal.afifi@kit.edu

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

refinement empirically reduce the global forward *Kullback–Leibler* (KL) divergence. At equal component counts, the geometry-derived construction also avoids the iterative fitting required by EM. In a nut–bolt clearance example, translations of the geometric object directly translate the corresponding Gaussian means, enabling analytic overlap rescoring without refitting the mixture. In a map-aware fusion example, the geometry-derived prior can be combined sequentially with Gaussian sensor likelihoods in closed form. These examples illustrate that the mixture is not merely a statistical approximation, but retains an explicit correspondence to the underlying geometric representation.

The approach is also motivated by the close relationship between geometry and physical modeling. Geometry provides the structure of physical systems, including configurations, constraints, and admissible states, while probability provides a natural representation of uncertainty. Geometric formulations of mechanics and physics illustrate the fundamental role of geometry in describing physical systems [6], [7]. Combining this geometric structure with probabilistic inference and uncertainty-aware control [8] suggests a broader "probabilistic physics" perspective in which geometry, probability, and decision making share a common computational representation.

The main contribution of this work is therefore an explicit geometry-to-GMM construction for uniform distributions on finite full-dimensional triangulated domains. The construction is closed form, uses one Gaussian component per simplex, and requires no numerical optimization. Furthermore, the resulting mixture admits a dimension-dependent upper bound on the global forward KL divergence. The number of components is controlled by the chosen triangulation, providing a natural accuracy–complexity trade-off; empirical refinement behavior suggests further investigation of adaptive triangulation and component reduction. The remainder of this paper develops the construction, establishes the local and global approximation results, and evaluates the resulting representation through geometric refinement, clearance analysis, and probabilistic fusion.

II. Construction

Definition 1 (Convex Hull [9]): Let $A \subset \mathbb{R}^d$. The convex hull of A is defined by

$$\operatorname{conv}(A) \coloneqq \left\{ \sum_{i=1}^{n} \lambda_i v_i \colon n \in \mathbb{N}_{>0},\ v_i \in A,\ \lambda_i \geq 0,\ \sum_{i=1}^{n} \lambda_i = 1 \right\}. \tag{1}$$

Definition 2 (Standard Simplex [9]): The standard d-simplex $\Delta^d \subset \mathbb{R}^d$ is defined by:

$$\Delta^d \coloneqq \operatorname{conv}\{0, e_1, \ldots, e_d\}, \tag{2}$$

where e_i denotes the i-th canonical basis in $\mathbb{R}^d$.

Definition 3 (Geometrically Independent Points [9]): $\{v_i\}_{i=0,\ldots,d} \subset \mathbb{R}^d$ are geometrically independent iff $\{v_i - v_0\}_{i=1,\ldots,d}$ are linearly independent.

Definition 4 (General Simplex [9]): Let $v_0, \ldots, v_d \in \mathbb{R}^d$ be geometrically independent. The general d-simplex σ^d generated by these points is defined by

$$\sigma^d(v_0, \ldots, v_d) \coloneqq \operatorname{conv}\{v_0, \ldots, v_d\}. \tag{3}$$

Definition 5 (Faces of a Simplex [9]): Let

$$\sigma^d = \operatorname{conv}\{v_0, \ldots, v_d\}$$

be a d-simplex. A *face* of σ^d is a simplex of the form

$$\tau \coloneqq \operatorname{conv}\{v_i : i \in I\}, \qquad \varnothing \neq I \subseteq \{0, \ldots, d\}.$$

If $|I| = r + 1$, then $\dim \tau = r$. In particular, 0-faces, 1-faces, and $(d-1)$-faces are called *vertices*, *edges*, and *facets*, respectively.

Definition 6 (Geometric Simplicial Complex [9]): A *geometric simplicial complex* $\mathcal{S}$ in $\mathbb{R}^d$ is a collection of simplices in $\mathbb{R}^d$ such that:

1) every face of a simplex in $\mathcal{S}$ is also in $\mathcal{S}$;
2) for any $\sigma, \tau \in \mathcal{S}$,
$$\sigma \cap \tau$$
is either empty or a common face of both σ and τ.

Definition 7 (Geometric Realization [10]): The *geometric realization* of a geometric simplicial complex $\mathcal{S}$ is

$$|\mathcal{S}| \coloneqq \bigcup_{\sigma \in \mathcal{S}} \sigma \subseteq \mathbb{R}^d.$$

Definition 8 (Dimension and Purity [9]): Let $\mathcal{S}$ be a finite geometric simplicial complex. Its *dimension* is

$$\dim \mathcal{S} \coloneqq \max_{\sigma \in \mathcal{S}} \dim \sigma.$$

A simplex $\sigma \in \mathcal{S}$ is *maximal* (or a *facet*) if it is not a proper face of any simplex $\tau \in \mathcal{S}$. The complex $\mathcal{S}$ is *pure of dimension* d if every maximal simplex has dimension d.

Definition 9 (Geometric Triangulation [9]): Let

$$\Omega \subseteq \mathbb{R}^d, \qquad \dim \Omega = d, \qquad \operatorname{vol}_d(\Omega) < +\infty.$$

A *finite geometric triangulation* of Ω is a pure d-dimensional geometric simplicial complex $\mathcal{S}$ such that

$$|\mathcal{S}| = \Omega, \qquad \#\mathcal{S} < \infty.$$

We denote its set of d-simplices by

$$\mathcal{S}^{[d]} \coloneqq \{\sigma \in \mathcal{S} : \dim \sigma = d\}.$$

Theorem 1 (Triangulation for 2-Manifolds [11]): Every 2-manifold is homeomorphic to the polyhedron of a 2-dimensional simplicial complex, in which every 1-simplex is a face of exactly two 2-simplices.

Theorem 2 (Triangulation for 3-Manifolds [11]): Every 3-manifold is triangulable.

Theorem 3 (Triangulation for Differentiable Manifolds [12]): Differentiable manifolds are always triangulable.

Proposition 1 (Intersections of maximal simplices [9], [13]): Let $\mathcal{S}$ be a finite, pure, d-dimensional geometric simplicial complex. Then, for distinct $\sigma, \tau \in \mathcal{S}^{[d]}$,

$$\sigma \cap \tau = \varnothing \quad \text{or} \quad \sigma \cap \tau \in \mathcal{S}, \qquad \dim(\sigma \cap \tau) \leq d - 1.$$

Hence

$$\operatorname{vol}_d(\sigma \cap \tau) = 0,$$

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

and $\operatorname{int}(\sigma) \cap \operatorname{int}(\tau) = \varnothing$.

Proof: By the definition of a geometric simplicial complex, $\sigma \cap \tau$ is empty or a common face. Since $\sigma \neq \tau$ and $\dim \sigma = \dim \tau = d$, any non-empty common face has dimension at most $d-1$, hence has d-dimensional Lebesgue measure zero. ∎

Corollary 1 (Volume of a finite triangulation [13]): Let $\mathcal{S}$ be a finite, pure, d-dimensional geometric simplicial complex. Then

$$\operatorname{vol}_d(|\mathcal{S}|) = \sum_{\sigma \in \mathcal{S}^{[d]}} \operatorname{vol}_d(\sigma).$$

Proof: Since

$$|\mathcal{S}| = \bigcup_{\sigma \in \mathcal{S}^{[d]}} \sigma$$

and distinct d-simplices intersect only in sets of d-dimensional measure zero, finite additivity yields the result. ∎

Proposition 2 ($F : \Delta^d \overset{\text{aff}}{\cong} \sigma^d(v_0, \dots, v_d)$): Let Δ^d be the standard d-simplex, and $\sigma^d(v_0, \dots, v_d) = \operatorname{conv}\{v_0, \dots, v_d\}$ be the general d-simplex, Then there exists a unique affine isomorphism

$$F : \Delta^d \overset{\text{aff}}{\cong} \sigma^d(v_0, \dots, v_d), \quad x \mapsto Ax + v_0, \tag{4}$$

where

$$A = \begin{bmatrix} v_1 - v_0 & v_2 - v_0 & \cdots & v_d - v_0 \end{bmatrix}. \tag{5}$$

Proposition 3 ($\operatorname{vol}_d \sigma(v_0, \dots, v_d)$ [14], [15]):

$$\operatorname{vol}_d \sigma(v_0, \dots, v_d) = \frac{1}{d!} \left(\det G\right)^{1/2}, \tag{6}$$

where G is the Gram matrix:

$$G_{ij} \coloneqq (v_i - v_0)^\top (v_j - v_0), \quad i, j = 1 \dots, d. \tag{7}$$

Definition 10 (Dirichlet Distribution on Δ^d [15]): Let $\Delta^d = \operatorname{conv}\{0, e_1, \dots, e_d\}$. For $x = (x_1, \dots, x_d) \in \Delta^d$, define

$$x_0 \coloneqq 1 - \sum_{i=1}^{d} x_i.$$

Thus $(x_0, \dots, x_d)$ are the barycentric coordinates of x. For $\alpha_\sharp = (\alpha_0, \dots, \alpha_d) \in \mathbb{R}^{d+1}_{++}$, the Dirichlet density on Δ^d is

$$\mathcal{D}_{\Delta^d}(x|\alpha_\sharp) = \frac{1}{B(\alpha_\sharp)} \prod_{i=0}^{d} x_i^{\alpha_i - 1},$$

where

$$B(\alpha_\sharp) = \frac{\prod_{i=0}^{d} \Gamma(\alpha_i)}{\Gamma(|\alpha_\sharp|)},$$

Corollary 2 (Uniform PDF on Δ^d in $\mathbb{R}^d$ [14], [15]): Let $1_{d+1} \coloneqq (1, \dots, 1) \in \mathbb{R}^{d+1}$. The uniform probability density function on Δ^d is the Dirichlet density with parameter 1_{d+1}, i.e.

$$X \sim \mathcal{D}_{\Delta^d}(x|1_{d+1}). \tag{8}$$

Its density is given by

$$\mathcal{U}_{\Delta^d}(x) \coloneqq \mathcal{D}_{\Delta^d}(x|1_{d+1}) = \begin{cases} d! & x \in \Delta^d \\ 0 & x \notin \Delta^d. \end{cases} \tag{9}$$

Proposition 4 ($\mathbb{E}(X), \mathbb{V}(X)$ for $X \sim \mathcal{D}_{\Delta^d}(x|\alpha)$ [15]): Let $X = (X_1, \dots, X_d)$ have a Dirichlet distribution on Δ^d with parameter $\alpha_\sharp = (\alpha_0, \dots, \alpha_d) \in \mathbb{R}^{d+1}_{++}$ and $\alpha_\natural = (\alpha_1, \dots, \alpha_d) \in \mathbb{R}^d_{++}$. Then

$$\begin{aligned} \mathbb{E}(X) &= \frac{\alpha_\natural}{|\alpha_\sharp|}, \\ \mathbb{V}(X) &= \frac{|\alpha_\sharp| \operatorname{diag}(\alpha_\natural) - \alpha_\natural \alpha_\natural^\top}{|\alpha_\sharp|^2 (|\alpha_\sharp| + 1)}. \end{aligned} \tag{10}$$

Definition 11 (KL Divergence): Let P and Q be probability measures on $\Omega = \mathbb{R}^d$. The Kullback–Leibler (KL) divergence of P from Q is defined as

$$D_{\mathrm{KL}}(P : Q) \coloneqq \begin{cases} \displaystyle\int_\Omega \log\left(\frac{\mathrm{d}P}{\mathrm{d}Q}\right) \mathrm{d}P, & P \ll Q, \\ +\infty, & \text{otherwise}, \end{cases} \tag{11}$$

where $P \ll Q$ denotes that P is absolutely continuous with respect to Q.

Remark 1 (KL Divergence for PDFs): Suppose that P and Q admit densities p and q with respect to the Lebesgue measure, then

$$D_{\mathrm{KL}}(P : Q) = \begin{cases} \displaystyle\int_\Omega p(x) \log\left(\frac{p(x)}{q(x)}\right) \mathrm{d}x, & P \ll Q, \\ +\infty, & \text{otherwise}. \end{cases} \tag{12}$$

Definition 12 (I-Projection): Let $\mathcal{Q}$ be a non-empty set of PDFs and let p be a PDF. The I-projection of p onto $\mathcal{Q}$ is defined as

$$\Pi^{\mathrm{I}}_{\mathcal{Q}}(p) \coloneqq \arg\min_{q \in \mathcal{Q}} D_{\mathrm{KL}}(p : q). \tag{13}$$

Remark 2: For example, the I-projection of a PDF $p(x)$ onto the Gaussian family $\mathcal{Q}_{\mathcal{N}} \coloneqq \{\mathcal{N}(x|m, \Sigma) : m \in \mathbb{R}^d, \Sigma \in \mathbb{S}^d_{++}\}$ is the Gaussian PDF that minimizes $D_{\mathrm{KL}}\big(p(x) : \mathcal{N}(x|m, \Sigma)\big)$ over $m \in \mathbb{R}^d$ and $\Sigma \in \mathbb{S}^d_{++}$.

Definition 13 (M-Projection): Let $\mathcal{Q}$ be a non-empty set of PDFs and let p be a PDF. The M-projection of p onto $\mathcal{Q}$ is defined as

$$\Pi^{\mathrm{M}}_{\mathcal{Q}}(p) \coloneqq \arg\min_{q \in \mathcal{Q}} D_{\mathrm{KL}}(q : p). \tag{14}$$

Remark 3: For example, the M-projection of a PDF $p(x)$ onto the Gaussian family $\mathcal{Q}_{\mathcal{N}} \coloneqq \{\mathcal{N}(x|m, \Sigma) : m \in \mathbb{R}^d, \Sigma \in \mathbb{S}^d_{++}\}$ is the Gaussian PDF that minimizes $D_{\mathrm{KL}}\big(\mathcal{N}(x|m, \Sigma) : p(x)\big)$ over $m \in \mathbb{R}^d$ and $\Sigma \in \mathbb{S}^d_{++}$.

Proposition 5 ($\min_{m \in \mathbb{R}^d, \Sigma \in \mathbb{S}^d_{++}} D_{\mathrm{KL}}\left(\mathcal{U}_{\Delta^d}(x) : \mathcal{N}(x|m, \Sigma)\right)$):

$$\begin{aligned} &\min_{m \in \mathbb{R}^d,\ \Sigma \in \mathbb{S}^d_{++}} D_{\mathrm{KL}}\left(\mathcal{U}_{\Delta^d}(x) : \mathcal{N}(x|m, \Sigma)\right) \\ &\qquad = \log\left(\frac{(2\pi e)^{d/2} (d!)}{(d+1)^{(d+1)/2} (d+2)^{d/2}}\right). \end{aligned} \tag{15}$$

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

$$(m^\star, \Sigma^\star) = \underset{m\in\mathbb{R}^d, \Sigma\in\mathbb{S}^d_{++}}{\operatorname{argmin}} D_{\mathrm{KL}}\left(\mathcal{U}_{\Delta^d}(x) : \mathcal{N}(x|m,\Sigma)\right) = \left(\mathbb{E}_{\mathcal{U}_{\Delta^d}}(X), \mathbb{V}_{\mathcal{U}_{\Delta^d}(x)}(X)\right). \tag{16}$$

Proof: The derivation of Equation (16) follows from [16].

$$(m^\star, \Sigma^\star) = \left(\frac{1}{d+1}1_d, \frac{(d+1)I_d - 1_d 1_d^\top}{(d+1)^2(d+2)}\right) \qquad \det\Sigma^\star = \frac{1}{(d+1)^{d+1}(d+2)^d}. \tag{17}$$

$$\min_{m\in\mathbb{R}^d,\ \Sigma\in\mathbb{S}^d_{++}} D_{\mathrm{KL}}\left(\mathcal{U}_{\Delta^d}(x) : \mathcal{N}(x|m,\Sigma)\right) = \log\left(\frac{(2\pi e)^{d/2}(d!)}{(d+1)^{(d+1)/2}(d+2)^{d/2}}\right). \tag{18}$$

■

Proposition 6 ($D_{\mathrm{KL}}(p(x) : q(x)$ is $\overset{\text{aff}}{\cong}$ Invariant [17]): Let $p(x)$ and $q(x)$ be PDFs on $\mathbb{R}^d$, and let

$$f\colon \mathbb{R}^d \to \mathbb{R}^d, \quad x \mapsto y, \tag{19}$$

be a diffeomorphism. Let $p'(y)$ and $q'(y)$ denote the PDFs obtained by transforming $p(x)$ and $q(x)$ under f, respectively. Then

$$D_{\mathrm{KL}}\big(p'(y) : q'(y)\big) = D_{\mathrm{KL}}\big(p(x) : q(x)\big). \tag{20}$$

Proposition 7 ($\min_{m\in\mathbb{R}^d,\Sigma\in\mathbb{S}^d_{++}} D_{\mathrm{KL}}\left(\mathcal{U}_{\sigma^d}(y) : \mathcal{N}(y|m,\Sigma)\right)$): For Δ^d and $\sigma^d(v_0,\dots,v_d)$ let:

$$F : \Delta^d \overset{\text{aff}}{\cong} \sigma^d(v_0,\dots,v_d), \quad x \mapsto Ax + v_0 = y, \tag{21}$$

be the affine diffeomorphism, then:

$$\min_{m\in\mathbb{R}^d,\ \Sigma\in\mathbb{S}^d_{++}} D_{\mathrm{KL}}\left(\mathcal{U}_{\sigma^d}(y) : \mathcal{N}(y|m,\Sigma)\right) = \log\left(\frac{(2\pi e)^{d/2}(d!)}{(d+1)^{(d+1)/2}(d+2)^{d/2}}\right), \tag{22}$$

and:

$$(m^\star, \Sigma^\star) = \underset{m\in\mathbb{R}^d, \Sigma\in\mathbb{S}^d_{++}}{\operatorname{argmin}} D_{\mathrm{KL}}\left(\mathcal{U}_{\sigma^d}(y) : \mathcal{N}(y|m,\Sigma)\right) = \left(A\mathbb{E}_{\mathcal{U}_{\Delta^d}}(X) + v_0, A\mathbb{V}_{\mathcal{U}_{\Delta^d}}(X)A^\top\right). \tag{23}$$

Example 2.1 (2-simplex and its Gaussian I-projection): Figure 2 numerically diagnoses the analytical simplex KL value in Proposition 7.

Remark 4 (I-projection vs M-projection): In our setting, we consider the I-projection rather than the M-projection because the corresponding optimization problems have fundamentally different behavior. The I-projection yields a finite minimum,

$$\min_{q\in\mathcal{Q}} D_{\mathrm{KL}}(p : q) < \infty, \tag{24}$$

whereas the reverse divergence appearing in the M-projection is infinite for every $q \in \mathcal{Q}$,

$$D_{\mathrm{KL}}(q : p) = \infty. \tag{25}$$

Consequently, the M-projection does not provide a meaningful finite minimizer in this setting, while the I-projection does.

Proposition 8 (Uniform Measure as a Mixture): Let $\mathcal{S}$ be a finite geometric triangulation of $\Omega \subseteq \mathbb{R}^d \wedge \dim\Omega = d \wedge 0 < \mathrm{vol}_d(\Omega) < +\infty$. Then

$$\mathcal{U}_\Omega = \sum_{\sigma\in\mathcal{S}^{[d]}} \frac{\mathrm{vol}_d(\sigma)}{\mathrm{vol}_d(\Omega)}\mathcal{U}_\sigma \qquad \text{a.e.,}$$

where

$$\mathcal{U}_A(x) \coloneqq \frac{1_A(x)}{\mathrm{vol}_d(A)}.$$

In particular, with

$$w_\sigma \coloneqq \frac{\mathrm{vol}_d(\sigma)}{\mathrm{vol}_d(\Omega)},$$

one has $w_\sigma > 0$ and

$$\sum_{\sigma\in\mathcal{S}^{[d]}} w_\sigma = 1.$$

Proof: By Proposition 1,

$$1_\Omega = \sum_{\sigma\in\mathcal{S}^{[d]}} 1_\sigma \qquad \text{a.e.,}$$

and by Corollary 1,

$$\mathrm{vol}_d(\Omega) = \sum_{\sigma\in\mathcal{S}^{[d]}} \mathrm{vol}_d(\sigma).$$

Therefore

$$\sum_{\sigma\in\mathcal{S}^{[d]}} w_\sigma \mathcal{U}_\sigma = \frac{1}{\mathrm{vol}_d(\Omega)} \sum_{\sigma\in\mathcal{S}^{[d]}} 1_\sigma = \mathcal{U}_\Omega \qquad \text{a.e.,}$$

and $\sum_\sigma w_\sigma = 1$. ■

Proposition 9 (Joint Convexity of $D_{\mathrm{KL}}(p : q)$ [17]): For any PDFs p_1, p_2, q_1, q_2 and any $\lambda \in [0,1]$, we have:

$$D_{\mathrm{KL}}(\lambda p_1 + (1-\lambda)p_2 : \lambda q_1 + (1-\lambda)q_2) \le \lambda D_{\mathrm{KL}}(p_1 : q_1) + (1-\lambda) D_{\mathrm{KL}}(p_2 : q_2). \tag{26}$$

Proposition 10: For PDFs $\{p_i\}_{i=1,\dots,K}$ and $\{q_i\}_{i=1,\dots,K}$, and $\pi \in \mathbb{R}^K_+ \wedge |\pi|_1 = 1$:

$$D_{\mathrm{KL}}(\sum_{i=1}^K \pi_i p_i : \sum_{i=1}^K \pi_i q_i) \le \sum_{i=1}^K \pi_i D_{\mathrm{KL}}(p_i : q_i). \tag{27}$$

Theorem 4 ($D_{\mathrm{KL}}\big(\mathcal{U}_\Omega(x) : \mathcal{G}_\mathcal{S}(x|g)\big) \le C(d)$): Let $\Omega \subseteq \mathbb{R}^d \wedge \dim\Omega = d \wedge 0 < \mathrm{vol}_d(\Omega) < +\infty$ admit a finite geometric triangulation $\mathcal{S}$. For each $\sigma = \mathrm{conv}\{v_0,\dots,v_d\} \in \mathcal{S}^{[d]}$, define

$$w_\sigma \coloneqq \frac{\mathrm{vol}_d(\sigma)}{\mathrm{vol}_d(\Omega)} \qquad (m^\star_\sigma, \Sigma^\star_\sigma) = \underset{m\in\mathbb{R}^d,\ \Sigma\in\mathbb{S}^d_{++}}{\operatorname{argmin}} D_{\mathrm{KL}}\big(\mathcal{U}_\sigma(x) : \mathcal{N}(x|m,\Sigma)\big). \tag{28}$$

Then

$$\mathcal{G}_\mathcal{S}(x|g) \coloneqq \sum_{\sigma\in\mathcal{S}^{[d]}} w_\sigma \mathcal{N}(x|m^\star_\sigma, \Sigma^\star_\sigma)$$

is a Gaussian mixture model with $\#\mathcal{S}^{[d]}$ components satis-

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

fying

$$D_{\mathrm{KL}}\big(\mathcal{U}_\Omega(x) : \mathcal{G}_{\mathcal{S}}(x|g)\big) \le \log\left(\frac{(2\pi e)^{d/2} d!}{(d+1)^{(d+1)/2}(d+2)^{d/2}}\right).$$

Proof: By Proposition 8, $\mathcal{U}_\Omega = \sum_{\sigma\in\mathcal{S}^{[d]}} w_\sigma \mathcal{U}_\sigma$ a.e.; hence, by the joint convexity of D_{KL},

$$D_{\mathrm{KL}}\big(\mathcal{U}_\Omega(x) : \mathcal{G}_{\mathcal{S}}(x|g)\big) \le \sum_{\sigma\in\mathcal{S}^{[d]}} w_\sigma D_{\mathrm{KL}}\big(\mathcal{U}_\sigma : \mathcal{N}(x|m_\sigma^\star, \Sigma_\sigma^\star)\big).$$

By Proposition 7,

$$D_{\mathrm{KL}}\big(\mathcal{U}_\sigma(x) : \mathcal{N}(x|m_\sigma^\star, \Sigma_\sigma^\star)\big) = \log\left(\frac{(2\pi e)^{d/2} d!}{(d+1)^{(d+1)/2}(d+2)^{d/2}}\right) \quad \forall \sigma \in \mathcal{S}^{[d]} \tag{29}$$

Since $\sum_{\sigma\in\mathcal{S}^{[d]}} w_\sigma = 1$, the claim follows. ∎

III. Numerical Examples and Applications

Example 3.1 ($\Pi^{\mathrm{I}}_{\mathcal{Q}_{\mathcal{N}}}(\mathcal{U}_{\sigma^2}(x))$): Consider the standard 2-simplex (an isosceles right triangle):

$$\Delta^2 = \mathrm{conv}\{0, e_1, e_2\} \subset \mathbb{R}^2 \tag{30}$$

with uniform density $\mathcal{U}_{\Delta^2}(x)$. From Corollary 2 and Proposition 4, its mean and covariance are

$$\begin{aligned} m_{\Delta^2}^\star &= \mathbb{E}_{\mathcal{U}_{\Delta^2}}(X) = \frac{1}{3}\begin{pmatrix}1\\1\end{pmatrix} \\ \Sigma_{\Delta^2}^\star &= \mathbb{V}_{\mathcal{U}_{\Delta^2}}(X) = \frac{1}{36}\begin{pmatrix}2 & -1\\ -1 & 2\end{pmatrix}. \end{aligned} \tag{31}$$

Hence, by Proposition 5, the Gaussian I-projection of $\mathcal{U}_{\Delta^2}(x)$ is

$$\mathcal{N}\left(x|m_{\Delta^2}^\star, \Sigma_{\Delta^2}^\star\right), \tag{32}$$

and the corresponding minimum KL divergence is

$$D_{\mathrm{KL}}\left(\mathcal{U}_{\Delta^2}(x) : \mathcal{N}(x|m_{\Delta^2}^\star, \Sigma_{\Delta^2}^\star)\right) = \log\left(\frac{\pi e}{3\sqrt{3}}\right). \tag{33}$$

Now consider an arbitrary nondegenerate triangle (a simplex with nonvanishing area)

$$\sigma^2(v_0, v_1, v_2) = \mathrm{conv}\{v_0, v_1, v_2\} \subset \mathbb{R}^2. \tag{34}$$

The corresponding affine isomorphism from the standard simplex Δ^2 to σ^2 is

$$F : \Delta^2 \overset{\mathrm{aff}}{\cong} \sigma^2, \qquad x \mapsto Ax + v_0 = y, \tag{35}$$

where

$$A = \begin{bmatrix} v_1 - v_0 & v_2 - v_0 \end{bmatrix}. \tag{36}$$

Since σ^2 is nondegenerate, A is invertible.

Under this affine transformation, the uniform density $\mathcal{U}_{\Delta^2}$ on Δ^2 is mapped to the uniform density $\mathcal{U}_{\sigma^2}$ on σ^2. Furthermore, the Gaussian I-projection transforms according to:

$$m_{\sigma^2}^\star = A m_{\Delta^2}^\star + v_0, \qquad \Sigma_{\sigma^2}^\star = A \Sigma_{\Delta^2}^\star A^\top. \tag{37}$$

Consequently,

$$m_{\sigma^2}^\star = \frac{v_0 + v_1 + v_2}{3}. \tag{38}$$

Moreover,

$$\Sigma_{\sigma^2}^\star = \frac{1}{36} A \begin{pmatrix}2 & -1\\ -1 & 2\end{pmatrix} A^\top = \frac{1}{12}\sum_{i=0}^{2} (v_i - m_\sigma^\star)(v_i - m_\sigma^\star)^\top. \tag{39}$$

Finally, by the affine invariance of the KL divergence in Proposition 6,

$$\begin{aligned} &D_{\mathrm{KL}}\left(\mathcal{U}_{\sigma^2}(y) : \mathcal{N}(y|m_{\sigma^2}^\star, \Sigma_{\sigma^2}^\star)\right) \\ &= D_{\mathrm{KL}}\left(\mathcal{U}_{\Delta^2}(x) : \mathcal{N}(x|m_{\Delta^2}^\star, \Sigma_{\Delta^2}^\star)\right) = \log\left(\frac{\pi e}{3\sqrt{3}}\right). \end{aligned} \tag{40}$$

Thus, every nondegenerate triangle has the same minimum KL divergence to its Gaussian I-projection, independent of its position, size, and shape. Figure 2 shows random 2-simplices σ^2 and Monte Carlo estimates consistent with the analytical minimum of the KL divergence.

Example 3.2 (Triangle Barycentric Subdivision): Starting from a single triangle, repeated barycentric subdivision [9] produces increasingly fine triangulations of the same domain and, consequently, increasingly detailed Gaussian mixture representations. At each of the four finite measured levels in Figure 3, the sampled global forward KL divergence and leakage decrease relative to the preceding shown level. These finite observations establish neither general monotonicity under barycentric subdivision nor asymptotic convergence of the KL divergence.

Example 3.3 (Polyhedral Refinement): Starting from a polyhedral domain, increasingly fine triangulations yield Gaussian mixture representations with increasing numbers of components. In the displayed run in Figure 4, finer meshes yield lower sampled global forward KL divergence for the geometry-derived construction. The held-out forward KL of the EM fit approaches near zero and then turns upward beginning at $K = 256$. Because held-out forward KL differs from held-out negative log likelihood only by the target entropy constant, this upturn means that held-out negative log likelihood worsens rather than improves. The behavior is consistent with finite-sample overfitting and/or sensitivity to optimization and regularization, but it does not diagnose a cause. In contrast, the geometry-derived construction has no fitted training sample and hence no train–test overfitting mechanism, although it still incurs approximation and discretization error.

At equal component counts, the geometry-derived parameters are obtained directly from the supplied mesh without iterative optimization. The comparison nevertheless has intentionally limited scope: EM is one fixed seeded fit using finite training and held-out draws, whereas the geometry timing starts from a supplied mesh and excludes meshing. It is neither a global-optimality claim nor a full benchmark. The finite refinement results also do not establish monotonic decrease or asymptotic convergence in general.

Example 3.4 (Map-Aware Fusion): Let $q_\Omega(x)$ be the geometry-derived GMM prior, and let $\ell_r(x)$ and $\ell_c(x)$ denote the robot and camera Gaussian likelihoods. Under conditional independence of the two measurements given x, the

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

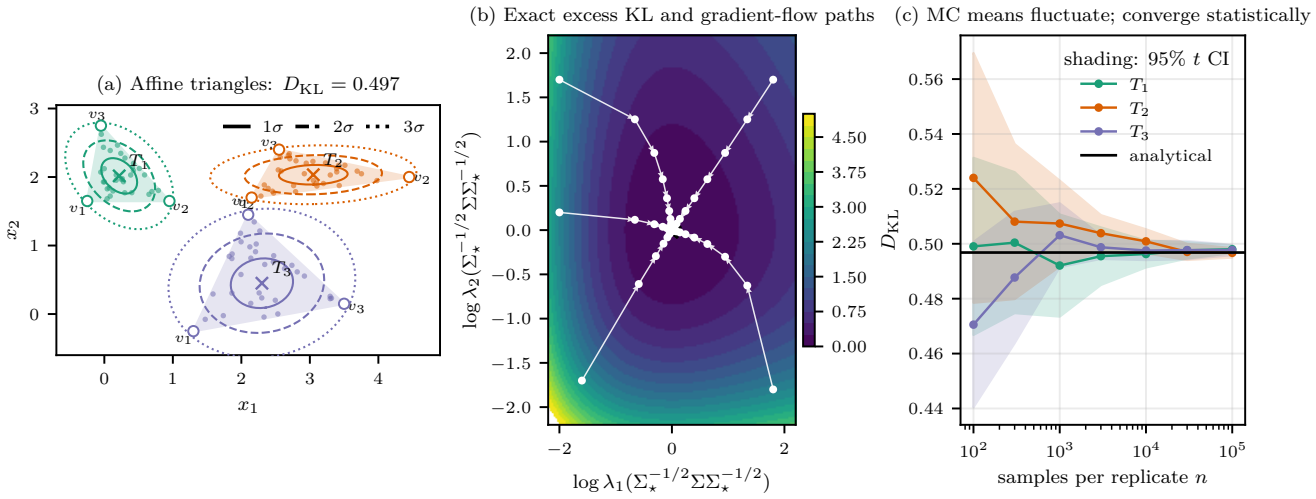


Fig. 2. Affine triangles share the proposed analytical simplex based KL. Panel b shows the exact excess-KL landscape with gradient-flow paths; panel c shows repeated Monte Carlo estimates and 95% intervals approaching the analytical value.

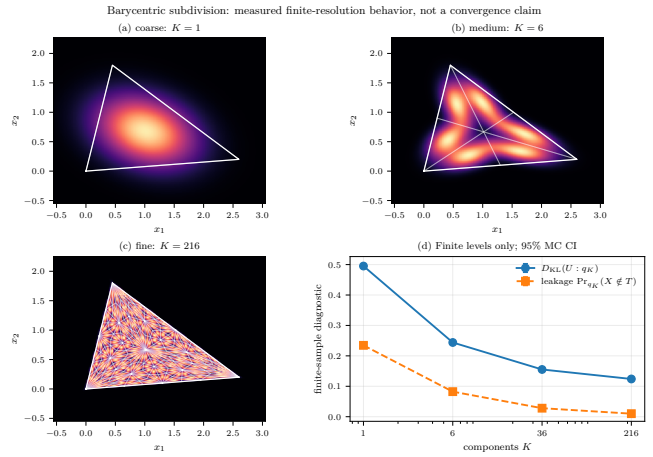


Fig. 3. Four finite barycentric-subdivision levels show decreasing sampled forward-KL and leakage diagnostics at each displayed level; they do not establish general monotonicity or asymptotic convergence.

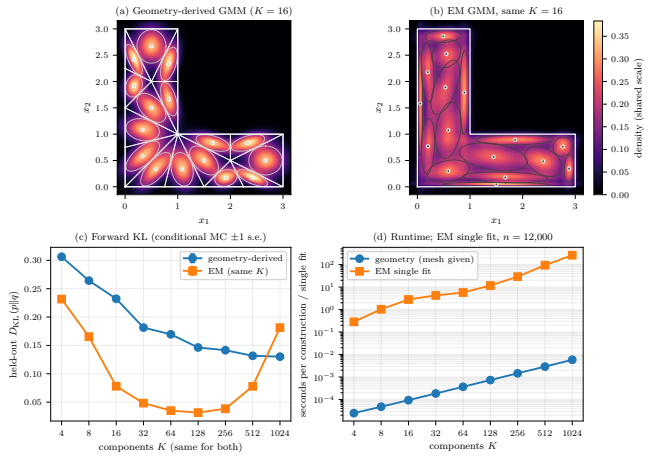


Fig. 4. Same-K comparison of the geometry-derived GMM and one fixed seeded EM fit on finite training and held-out draws. The EM held-out forward KL turns upward from $K = 256$ in this run. Geometry timing starts from the supplied mesh and excludes meshing; the comparison is not a global-optimality or full-benchmark claim.

normalized posterior is

$$q(x \mid z_r, z_c) = \frac{q_\Omega(x)\,\ell_r(x)\,\ell_c(x)}{Z}, \qquad Z = \int q_\Omega(u)\,\ell_r(u)\,\ell_c(u)\,\mathrm{d}u. \tag{41}$$

Multiplying each Gaussian component by the two Gaussian likelihoods in sequence gives the exact posterior update for the constructed GMM. The constrained uniform-domain field shown as a contour in Figure 5 is evaluated on a finite grid and serves only as a numerical reference. If the sensor correlation is unknown, the conditional independence factorization is invalid and naive multiplication is not justified.

Example 3.5 (Nut–Bolt Clearance): Let e_r be the upward radial unit vector. A rigid signed radial translation by δ, positive upward, transforms the nut density according to

$$q_{\mathrm{nut},\delta}(x) = q_{\mathrm{nut}}(x - \delta e_r), \tag{42}$$

so every nut-component mean shifts by $+\delta e_r$, while its weight and covariance remain unchanged. The overlap score is

$$S(\delta) = \int_{\mathbb{R}^2} q_{\mathrm{nut}}(x - \delta e_r) q_{\mathrm{bolt}}(x)\,\mathrm{d}x. \tag{43}$$

Writing the two mixtures with parameters $(w_i^n, \mu_i^n, \Sigma_i^n)$ and $(w_j^b, \mu_j^b, \Sigma_j^b)$ gives the exact pairwise Gaussian sum

$$S(\delta) = \sum_{i,j} w_i^n w_j^b \mathcal{N}\big(\mu_i^n + \delta e_r \big| \mu_j^b, \Sigma_i^n + \Sigma_j^b\big)\,. \tag{44}$$

These translation and overlap identities are analytic and exact for the resulting GMMs. The displayed normalized

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

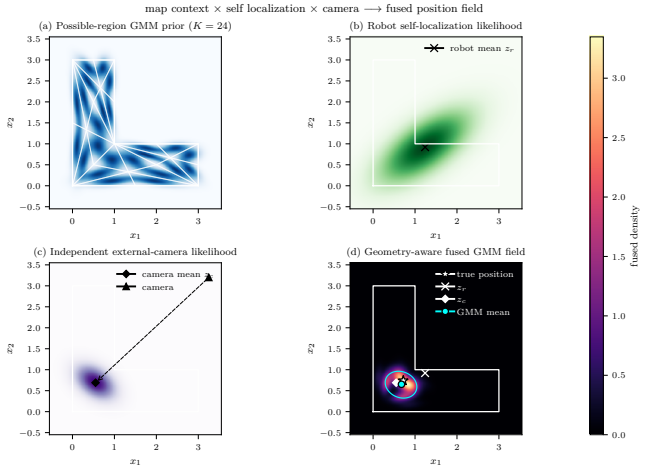


Fig. 5. A map prior times conditionally independent robot and camera Gaussian likelihoods gives a sequential closed-form update. The contour is a finite-grid constrained reference; no validity is claimed under unknown sensor correlation.

score is relative and dimensionless, but remains only a soft overlap/interference surrogate: it is not a collision probability, contact model, ISO analysis, or certification result. A geometry edit requires remeshing and regeneration of the affected GMM; a rigid translation changes only its component means and can be rescored analytically.

IV. Conclusion

We considered the problem of representing a uniform distribution on a geometrically described domain $\Omega \subset \mathbb{R}^d$ by a Gaussian mixture model (GMM). Given a finite triangulation of Ω, we proposed a construction that replaces the uniform distribution on each d-simplex by its KL-optimal Gaussian and combines the resulting components according to the volumes of the simplices.

We showed that the Gaussian I-projection of the uniform distribution on a simplex is given in closed form by matching its mean and covariance. Furthermore, exploiting affine equivalence between simplices and the invariance of the KL divergence under bijections, we showed that the corresponding minimum KL divergence is independent of the position, size, and shape of the simplex and depends only on the dimension d.

By combining the local Gaussian approximations with the volume-weighted representation of the uniform distribution on Ω and using the joint convexity of the KL divergence, we obtained a global approximation guarantee. In particular, the resulting GMM $\mathcal{G}_{\mathcal{S}}(x|g_{\mathcal{S}})$ satisfies:

$$D_{\mathrm{KL}}\left(\mathcal{U}_{\Omega} : \mathcal{G}_{\mathcal{S}}\right) \leq \log\left(\frac{(2\pi e)^{d/2} d!}{(d+1)^{(d+1)/2}(d+2)^{d/2}}\right). \quad (45)$$

Thus, the bound is independent of the particular geometry and triangulation and depends only on the dimension of the manifold.

The main advantage of the proposed construction is that it provides an explicit geometry-to-probability mapping without requiring numerical optimization and in closed form. The Gaussian parameters are obtained directly from the vertices of the simplices, while the resulting mixture preserves the global structure of the original domain. At the same time, the number of mixture components is determined by the number of d-simplices in the triangulation, so fine triangulations may lead to large GMMs. Moreover, the dimension-dependent bound does not quantify the approximation quality of a particular triangulation. The finite-level diagnostics in Figure 3 decrease, but do not establish asymptotic convergence.

The numerical examples separate validation from application: Figures 2 to 4 examine the local result, finite refinement behavior, and the deterministic construction relative to a fitted baseline, while Figures 5 and 6 demonstrate closed-form Gaussian fusion, rigid translation, and analytic overlap evaluation using the resulting mixtures.

This work provides an explicit construction and a dimension-dependent upper bound in the KL sense for GMM representations of triangulated domains. Future work could investigate adaptive triangulations and component reduction methods to obtain more compact representations while controlling the approximation error. Extensions to more general geometric objects, including lower-dimensional manifolds embedded in higher-dimensional spaces, are also of interest.

On the application side, we envision "Probabilistic Physics Using GMM" as a natural next step in the GMM based geometric formulation of physical systems. Geometry has provided a remarkably successful language for formulating physical laws and capturing the underlying structure of physical systems, from geometric formulations of classical mechanics to modern theories of dynamics and field physics [6]. Probability, in turn, provides a principled language for representing uncertainty, variability, and incomplete knowledge, forming the foundation of modern state estimation and sensor fusion. Building on this probabilistic representation, stochastic and optimal control provide a natural framework for reasoning and decision-making under uncertainty [8], opening the possibility of methods that jointly learn, model, estimate, and control complex physical systems.

V. Declaration on the Use of Artificial Intelligence

GPT-6 Astra (Pro Version) was used to assist in the preparation of the examples and their visualizations.

VI. Acknowledgment

This work was funded by Project SFB 1574 "A Circular Factory for the Perpetual Product", Deutsche Forschungsgemeinschaft (DFG), Project ID 471687386.

References

[1] Vladimir Gilelevich Maz'ya and Gunther Schmidt, *Approximate Approximations*. American Mathematical Society, 2007.

[2] Yulong Lu and Andrew Stuart and Hendrik Weber, "Gaussian Approximations for Probability Measures on $\mathbb{R}^d$," *SIAM/ASA Journal on Uncertainty Quantification*, vol. 5, pp. 1136–1165, 2017.

[3] Arthur Pentland Dempster and Nan McKenzie Laird and Donald Bruce Rubin, "Maximum Likelihood from Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 39, no. 1, pp. 1–22, 1977.

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

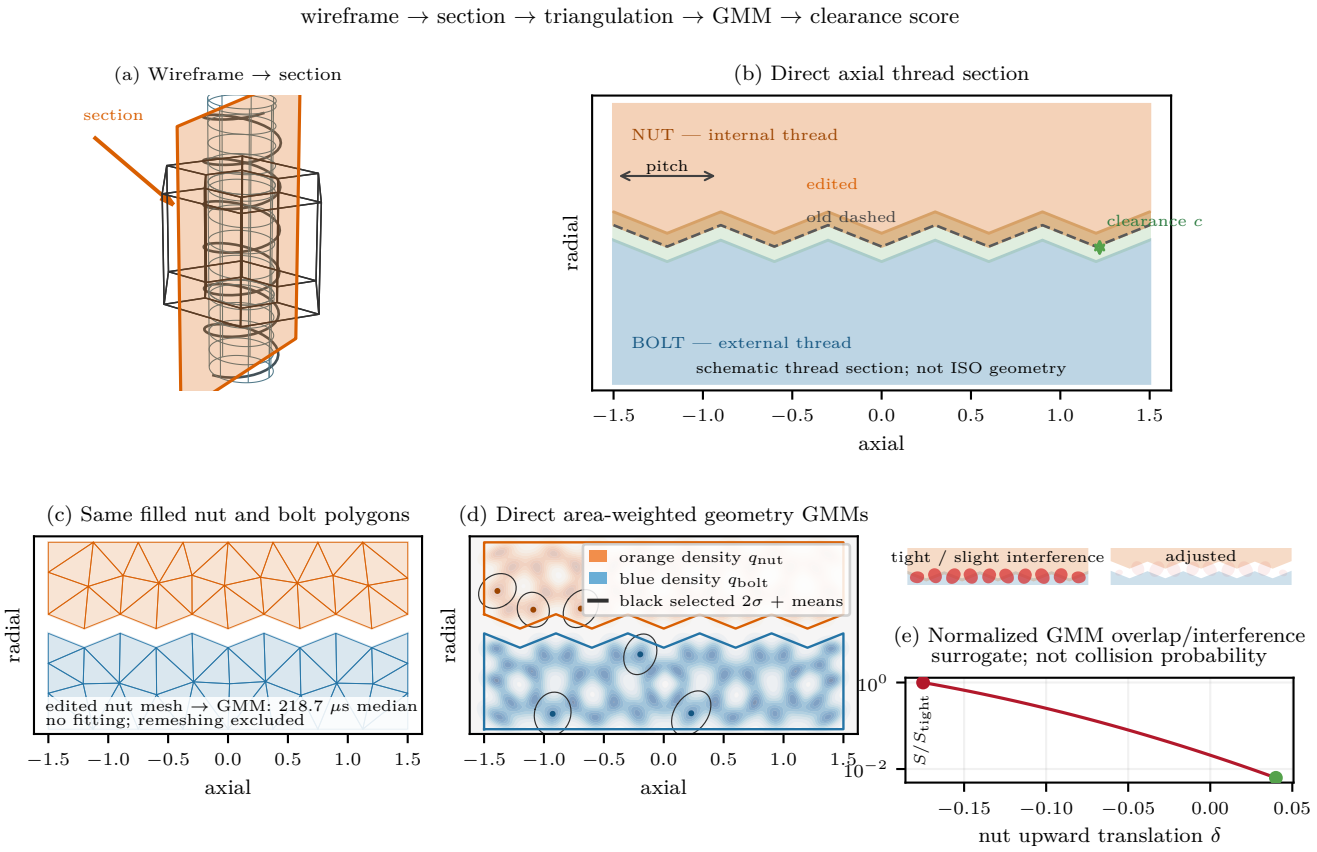


Fig. 6. A schematic wireframe is sectioned; the same bounded full-dimensional nut and bolt material polygons are directly triangulated and converted to area-weighted geometry GMMs. Nut-clearance translations update component means and allow analytic overlap rescoring. Normalized overlap/interference is a soft surrogate, not a collision probability; this is not an ISO, contact, or certification analysis.

[4] Chien-Fu Jeff Wu, "On the Convergence Properties of the EM Algorithm," *Annals of Statistics*, vol. 11, no. 1, pp. 95–103, 1983.
[5] Geoffrey John McLachlan and Thriyambakam Krishnan, *The EM Algorithm and Extensions*, 2nd ed. Wiley, 2007.
[6] Vladimir Igorevich Arnold, *Mathematical Methods of Classical Mechanics*. Springer, 1989.
[7] Theodore Frankel, *The Geometry of Physics: An Introduction*, 3rd ed. Cambridge University Press, 2012.
[8] Wendell Fleming and Raymond Rishel, *Deterministic and Stochastic Optimal Control*. Springer, 1975.
[9] James R. Munkres, Steven G. Krantz, and Harold R. Parks, *Elements of Algebraic Topology*, 2nd ed. CRC Press, 2025.
[10] Dimitry Kozlov, *Combinatorial Algebraic Topology*. Springer, 2008.
[11] John Marshall Lee, *Introduction to Topological Manifolds*, 2nd ed. Springer, 2011.
[12] John Henry Constantine Whitehead, "On C^1-complexes," *Annals of Mathematics*, vol. 41, no. 4, pp. 809–824, 1940.
[13] Terence Tao, *An Introduction to Measure Theory*. American Mathematical Society, 2011.
[14] P. Stein, "A Note on the Volume of a Simplex," *The American Mathematical Monthly*, vol. 73, no. 3, pp. 299–301, 1966.
[15] Kai Wang Ng, Guo-Liang Tian, and Man-Lai Tang, *Dirichlet and Related Distributions*. Wiley, 2011.
[16] Daniel Sanz-Alonso,, Andrew Stuart, and Armeen Taeb, *Inverse Problems and Data Assimilation*. Cambridge University Press, 2023.
[17] Nihat Ay, Jürgen Jost, Hông Vân Lê, and Lorenz Schwachhöfer, *Information Geometry*. Springer, 2017.

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.