A Geometric Gaussian Mixture Representation of Curves and Trajectories in Euclidean Space

Benedikt Stratmann[#¶⊥], Emadedin Soufian[⊛⊥], Jürgen Beyerer[⊛⊥] and Ali Darijani[⊛□#¶]

***Abstract*— Curves and trajectories in $\mathbb{R}^d$ are basic objects in tracking and sensor data fusion. *Gaussian Mixture Model*s (GMMs) are widely used as flexible parametric *Probability Density Function*s (PDFs) over the object state space. GMMs are usually obtained by applying *Expectation Maximization* (EM)-based schemes to point samples. When the underlying geometric structure is known deterministically - for example a road centerline in map-assisted tracking, a nominal flight path, or a planned robot trajectory - such iterative fitting is unnecessary, can be slow, and only indirectly reflects the curve or trajectory.**

We propose a purely arithmetic, geometry-driven construction that maps a curve or trajectory in $\mathbb{R}^d$ with user-defined transverse uncertainty directly to a GMM in a single pass, without EM or any iterative optimization.

Starting from a parametrized curve $\alpha\colon I \to \mathbb{R}^d$, we form a *User Defined Probabilistic Polygonal Representation* (UDPP) and assign each segment a normal-subspace standard deviation, obtaining thin curve-aligned primitives encoding position, tangent, normal, and length scale together with transverse uncertainty. Choosing component weights proportional to segment length distributes probability uniformly over arc length, so all GMM parameters follow directly from the discretized curve and uncertainty profile.

The construction applies to smooth, open, closed, non-regular, and self-intersecting curves and trajectories in $\mathbb{R}^d$. It supports adaptive discretization and spatially varying uncertainty and yields differentiable GMMs that directly encode geometric structure. We demonstrate the approach on canonical planar curves and nominal trajectories and include a simulation study in higher-dimensional Euclidean spaces comparing approximation quality and runtime against EM-based fitting. Applications include map- and road-constrained tracking and sensor data fusion, uncertainty-aware trajectory planning, probabilistic obstacle and boundary modeling in robotics, and uncertainty-aware CAD and digital twins.

I. INTRODUCTION

Curves and trajectories in Euclidean state space $\mathbb{R}^d$ are among the simplest geometric objects, yet they are fundamental in tracking and sensor data fusion. In many fusion systems, prior knowledge on object motion and environment structure is represented by one-dimensional manifolds in state space. Related fields of application include computer vision [1], robotics [2], measurement and calibration [3], [4], shape analysis [5], motion planning and tracking [6], and *Computer Aided Design* (CAD) [7], [8].

Classical curve representations (parametric curves, polygons [9], polynomials, splines) are deterministic: they specify where the curve lies in $\mathbb{R}^d$, but do not directly encode uncertainty normal to the curve. In contrast, probabilistic representations are widely used in Bayesian estimation and inference in the fusion community as flexible parametric PDFs over object state vectors and environment geometry. Typical usage fits a GMM to point samples via EM or variants [10], [11], sometimes exploiting underlying mesh geometry [12]. In such settings, curve or surface structure is not imposed a priori; it is only indirectly reflected via sampled data.

In many applications, however, the curve itself is known deterministically (e.g., a CAD edge or nominal path), and one wants to endow it with a tunable uncertainty field, especially in the normal direction. Examples include manufacturing tolerances and wear in CAD and digital twins, obstacle boundary uncertainty in robotics, and probabilistic models of trajectories and shapes.

In this paper we propose a direct mapping from a curve or trajectory in $\mathbb{R}^d$ with a UDPP to a GMM. The mapping is arithmetical, does not require iterative fitting, and preserves local and global geometric structure.

Contributions

The contributions of this paper are twofold:

- We propose and analyze a deterministic, geometry-driven GMM representation of rectifiable curves and trajectories in $\mathbb{R}^d$. Starting from a probabilistic polygonal approximation with segment-wise transverse uncertainty, we derive closed-form Gaussian components whose means, covariances, and weights are obtained directly from the curve geometry.
- We evaluate the resulting representation against standard EM-based GMM fitting, comparing approximation quality via *Kullback–Leibler* (KL) divergence and runtime over different numbers of components and ambient dimensions.

In addition, the construction supports adaptive discretization, spatially varying uncertainty, non-regular and self-intersecting curves, and extensions to higher-dimensional Euclidean state spaces.

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

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

[⊛⊥]Emadedin Soufian is affiliated with DNV, Hamburg, Germany. emadedin.soufian@dnv.com

[⊛⊥]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

[□#¶]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

© 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.

II. RELATED WORK

We focus on work related to constructing GMM representations of geometry under uncertainty.

A common approach is to fit a GMM to point samples, typically using variants of the EM algorithm. Such methods treat the data as noisy samples from an unknown distribution and recover component means, variances, and weights numerically [10], [11]. When the underlying object is a curve, it is usually first discretized into a point cloud by sampling its parameter domain and pushed through the curve mapping, and then a GMM is fitted to these samples. In this setting, the curve structure is not hard coded in the GMM parameters; it is only indirectly reflected via the sampled data.

A more geometry aware strategy is to exploit a given mesh or surface representation. For instance, [12] constructs GMM on triangular meshes, modifying the EM procedure to respect mesh structure. This couples data driven fitting with an explicit geometric scaffold, but still relies on iterative optimization over GMM parameters.

GMMs have also been used as compact environment models in robotics and autonomous driving. Recent work uses GMMs to represent occupancy and free space directly [13], [14], enabling efficient communication and planning. In these approaches, GMM parameters are typically obtained from sensor data or learned models.

In computer graphics and vision, 3D Gaussian splatting [15] represents scenes by a large number of anisotropic Gaussian "splats" in space, optimized to match image observations. While not focused on curves, this demonstrates the practical value of Gaussian primitives as a differentiable renderer and learning friendly representation of geometry.

In contrast to these fitting based or mesh driven methods, our work starts from a known deterministic curve or trajectory in $\mathbb{R}^d$ (e.g., a CAD edge or nominal path) together with a user defined uncertainty field in the normal direction. We derive GMM component means, variances, and weights in closed form from the UDPP, without EM or numerical optimization. As a result, local tangent, normal, and length scale are encoded directly in the GMM parameters, and uncertainty is introduced in a controlled, interpretable way.

III. BACKGROUND AND PROBLEM FORMULATION

A curve $\alpha : I \to \mathbb{R}^d$ is called rectifiable if it has finite arc length, i.e.,

$$L(\alpha) := \sup_{\{t_i\}} \sum_i \|\alpha(t_i) - \alpha(t_{i-1})\| < \infty,$$

where the supremum is taken over all finite partitions of I. Thus rectifiability excludes curves of infinite length while still allowing non-smooth, non-regular, and self-intersecting curves. For notational simplicity and for visualization, most examples in this paper use $d = 2$, but the construction is defined for every finite $d \geq 2$ and applies directly to higher-dimensional state spaces (e.g., 3D position, position–velocity spaces, or other Euclidean state vectors used in tracking and sensor data fusion).

Definition 1 (Curve and its polygonal approximation): Let $I = [a, b] \subset \mathbb{R}$ be a compact interval and let $\alpha : I \to \mathbb{R}^d$ be a (rectifiable) curve. A partition of I is a finite ordered set

$$a = t_0 < t_1 < \cdots < t_N = b.$$

The curve vertices are $v_i := \alpha(t_i) \in \mathbb{R}^d$, $i = 0, \ldots, N$, and the segments $S_i := (v_{i-1}, v_i)$, $i = 1, \ldots, N$.

Remark 1 (Derivative-free construction): The construction only uses point evaluations $\alpha(t_i)$ at the partition nodes and the induced segments (v_{i-1}, v_i). It does not require an arc-length parametrization or access to derivatives of α. Any continuous parametrization of a rectifiable curve in $\mathbb{R}^d$ is therefore admissible.

The polygonal chain $\{S_i\}$ is a standard approximation of α, and can be refined by choosing a finer partition. Throughout the construction, zero-length segments with $\ell_i = 0$ are discarded, since they do not contribute to arc length and do not define a unique tangent direction. We now attach a transverse uncertainty to each segment.

Definition 2 (Normal-subspace uncertainty): For each segment $S_i = (v_{i-1}, v_i)$ we assign a scalar parameter

$$\tau_i > 0, \quad i = 1, \ldots, N,$$

interpreted as an isotropic standard deviation in the segment's $(d-1)$-dimensional normal subspace.

Definition 3 (UDPP): A *User Defined Probabilistic Polygonal Representation* (UDPP) of α is a collection

$$\mathcal{P} = \{(S_i, \tau_i)\}_{i=1}^N,$$

where $S_i = (v_{i-1}, v_i)$ is a segment of a polygonal approximation of α and $\tau_i > 0$ is the associated normal-subspace uncertainty.

Remark 2 (Non-regular and self-intersecting curves): The curve α need not be regular and may exhibit singularities, loops, or self-intersections. Since the construction uses only the vertices $\alpha(t_i)$ and the segment-wise tangents derived from (v_{i-1}, v_i), it applies to any rectifiable parametrized curve for which a suitable partition can be chosen. Different branches or disconnected components are simply represented by disjoint sets of segments in the resulting mixture.

Each pair (S_i, τ_i) is interpreted as an idealized "probabilistic segment": probability mass is spread uniformly along the tangent direction of S_i and Gaussian in its normal subspace, with standard deviation τ_i.

IV. GAUSSIAN MIXTURE REPRESENTATION

Figure 1 summarizes the construction. A known curve or trajectory is first sampled at partition nodes, giving a polygonal approximation. Each segment is then equipped with a transverse uncertainty parameter. From the segment midpoint, tangent, normal subspace, length, and uncertainty, the corresponding Gaussian mean, covariance, and mixture weight are obtained directly. Thus the final GMM is constructed without sampling-based fitting or iterative optimization.

© 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.

© 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.

(a) Partitioned curve α (b) Uniform tangent, Gaussian normal (c) Moment-matched GMM

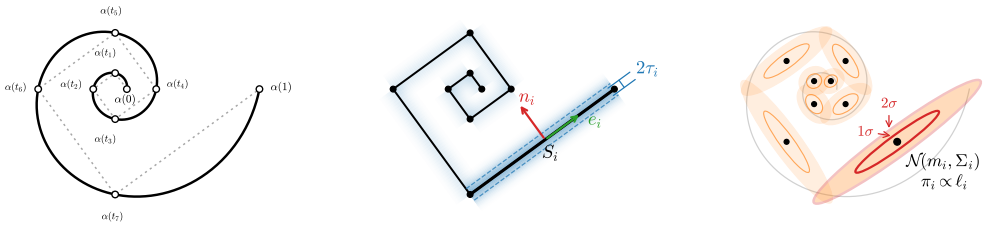


Fig. 1. Schematic construction of the geometric GMM. (a) A known curve or trajectory $\alpha(t)$ is given. (b) The curve is replaced by a UDPP; each segment S_i has tangent direction e_i, normal direction n_i, and transverse uncertainty τ_i. (c) Each segment is mapped to a Gaussian component $\mathcal{N}(x|m_i, \Sigma_i)$ with length-proportional mixture weight $\pi_i \propto \ell_i$.

We next define a *Random Variable* (RV) associated with each probabilistic segment and derive a Gaussian component that matches its first and second central moments.

Remark 3: The construction below holds for every $d \geq 2$. The scalar τ_i specifies isotropic uncertainty in the $(d-1)$-dimensional normal subspace. For $d = 2$, this subspace is spanned by the unique normal n_i up to sign, as shown in Fig. 1.

A. Per segment RV

For $S_i = (v_{i-1}, v_i)$, define

$$d_i := v_i - v_{i-1}, \qquad \ell_i := \|d_i\|, \qquad e_i := \frac{d_i}{\ell_i} \quad (\ell_i > 0),$$

and let the columns of $N_i \in \mathbb{R}^{d\times(d-1)}$ form an orthonormal basis of $e_i^{\perp}$, so that

$$N_i^{\top} N_i = I_{d-1}, \qquad N_i N_i^{\top} = I_d - e_i e_i^{\top}.$$

For $d = 2$, we write $N_i = n_i = Re_i$, where R is the $90°$ rotation matrix

$$R = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix},$$

and n_i is the unit normal of the segment.

Definition 4 (Per segment RV): Let

$$U_i \sim \mathcal{U}(u_i|0,1), \qquad Z_i \sim \mathcal{N}(0, I_{d-1}), \qquad U_i \perp Z_i.$$

The RV associated with (S_i, τ_i) is

$$X_i := v_{i-1} + \ell_i U_i\, e_i + \tau_i N_i Z_i. \tag{1}$$

Thus U_i distributes mass uniformly along the segment in tangent direction, and $\tau_i N_i Z_i$ introduces isotropic Gaussian uncertainty in the normal subspace.

B. Moment matched Gaussian component

Proposition 1 (Mean and covariance of X_i): Let X_i be defined by (1). Then

$$\mathbb{E}[X_i] = \frac{v_{i-1} + v_i}{2}, \tag{2}$$

$$\mathbb{V}(X_i) = \frac{\ell_i^2}{12}\, e_i e_i^{\top} + \tau_i^2 \big(I_d - e_i e_i^{\top}\big). \tag{3}$$

Proof: [Proof sketch] Using independence of U_i and Z_i and the known moments of the uniform and standard normal distributions,

$$\mathbb{E}[U_i] = \tfrac{1}{2}, \quad \mathbb{V}(U_i) = \tfrac{1}{12}, \quad \mathbb{E}[Z_i] = 0, \quad \mathbb{V}(Z_i) = I_{d-1},$$

one obtains (2) by linearity of expectation, and (3) by decomposing the variance into tangential and normal contributions. A detailed derivation is provided in the technical report [16]. ■

The mean $m_i := \mathbb{E}[X_i]$ is the segment midpoint, and the covariance decomposes into a tangential term proportional to ℓ_i^2 and an isotropic normal-subspace term proportional to τ_i^2. Its eigenspaces are aligned with the segment's tangent and normal subspace.

Remark 4 (Gaussian approximation and KL projection): The distribution of X_i is uniform along the segment tangent and Gaussian in the normal subspace. A natural Gaussian approximation is obtained by matching its first and second central moments, leading to a Gaussian

$$\mathcal{N}(x|m_i, \Sigma_i), \quad m_i = \tfrac{v_{i-1}+v_i}{2}$$

$$\Sigma_i = \tfrac{\ell_i^2}{12}\, e_i e_i^{\top} + \tau_i^2 \big(I_d - e_i e_i^{\top}\big).$$

This Gaussian can also be characterized information-theoretically: it is the optimal Gaussian approximation in the KL sense. If p_i denotes the density of X_i and we consider

$$\min_{m,\Sigma} D_{\mathrm{KL}}\big(p_i \,\|\, \mathcal{N}(\cdot|m, \Sigma)\big),$$

then the unique minimizer over all non-degenerate Gaussians is given by $m = \mathbb{E}[X_i]$ and $\Sigma = \mathbb{V}(X_i)$. Thus the moment matched Gaussian is precisely the KL-optimal Gaussian approximation to p_i.

For $d = 1$, if U is uniform on $[-a/2, a/2]$ and we approximate U by a Gaussian in the same KL sense, the minimizer is $\mathcal{N}(x|0, a^2/12)$, i.e., the Gaussian with matching mean and variance, and the minimal KL divergence is:

$$D_{\mathrm{KL}}\big(\mathcal{U}(x|-a/2, a/2)\big\|\mathcal{N}(x|0, a^2/12)\big) = \tfrac{1}{2}\log\big(\tfrac{\pi}{6}\big) + \tfrac{1}{2}.$$

C. Mixture along the curve

Assuming probability mass is distributed uniformly along arc length, the total probability assigned to segment S_i is proportional to ℓ_i. This motivates the following choice of mixture weights.

Definition 5 (Mixture weights): By rectifiability of the curve, every polygonal approximation has finite total length. For a nonconstant curve, after removing possible zero-length segments, we have

$$0 < \sum_{j=1}^{N} \ell_j < \infty.$$

We define the length-based mixture weights by

$$\pi_i := \frac{\ell_i}{\sum_{j=1}^{N} \ell_j}, \qquad i = 1, \dots, N.$$

Definition 6 (GMM representation): Given a UDPP $\mathcal{P} = \{(S_i, \tau_i)\}_{i=1}^{N}$ of a rectifiable curve, the associated GMM PDF is

$$p(x) = \sum_{i=1}^{N} \pi_i \mathcal{N}(x | m_i, \Sigma_i), \tag{4}$$

with m_i and Σ_i as above and π_i the length based weights.

Remark 5 (Zero normal-subspace uncertainty): If the model is relaxed to permit $\tau_i = 0$, then Σ_i has rank one. Geometrically this corresponds to uncertainty only along the segment tangent, consistent with a purely one dimensional curve segment. Numerically, such components are degenerate and may require regularization (e.g., adding a small positive multiple of the identity matrix) in algorithms that invert covariance matrices.

Remark 6 (Adaptive discretization and automatic choice of τ_i): The construction permits non-uniform segment lengths and segment dependent τ_i, which enables local refinement in regions of high curvature or near singularities, and varying uncertainty along the curve (e.g., encoding spatially varying tolerances or sensor noise). In practice, segment lengths can be chosen adaptively based on geometric criteria (such as an estimate of curvature or a user specified sampling density), and the values τ_i can be set automatically as functions of local curvature, estimated measurement noise, or tolerance classes taken from a CAD or other specification. Designing principled, fully automated schemes for joint discretization and τ_i selection is an interesting direction for future work and lies beyond the scope of this paper.

V. EXAMPLES AND APPLICATIONS

We illustrate how the GMM representation behaves on canonical curves and discuss several application scenarios.

A. Canonical curves

The construction applies uniformly to:

- *Smooth open curves* (e.g., parabola, measurement curve, logarithmic and Archimedean spirals), which can be interpreted as nominal trajectories in state space or as road centerlines.
- *Smooth closed curves* (e.g., circle, ellipse, superellipse): the mixture forms a closed "probabilistic tube" tightly aligned with the curve.
- *Curves with isolated singularities* (e.g., semicubical cusp, astroid, cardioid): away from singular points, behavior is as above; near cusps, neighboring components reflect the ambiguity in tangent direction.
- *Self-intersecting curves* (e.g., lemniscate, Lissajous, rose curves): different branches correspond to disjoint sets of segments, and at self-intersections the local PDF is a superposition of components associated with distinct branches.
- *Composite and disconnected curves* (e.g., Yin–Yang like curves, disjoint circles and semicircles): components from different pieces simply add in the mixture.

Figure 2 shows a "complex gallery" of curves where only the final GMM PDF is displayed; per-curve split views and convergence studies are provided in the technical report [16].

At isolated cusps or other singular points, a classical tangent direction may be undefined or nonunique. Our construction does not assign a single tangent or normal direction to such a point. Instead, each adjacent nonzero polygonal segment contributes its own tangent and normal direction. Thus, the "ambiguity" near a cusp is represented by several nearby Gaussian components with different principal directions, corresponding to the incoming and outgoing local branches of the curve.

At self-intersections, "superposition" means that the mixture density is the sum of the Gaussian components associated with all branches passing through the same spatial region. The components are not merged by the construction; hence branch information remains available at the component level, although the scalar density may be higher near the intersection.

These effects are usually not problematic for density representation itself. However, they can matter in downstream tasks such as tracking, gating, or path following, where one may need to distinguish between branches. Coarse discretization or a normal uncertainty τ_i that is large compared with the distance between nearby branches may also blur distinct parts of the curve and create artificial connections. In practice, these issues can be mitigated by adaptive refinement near cusps and intersections, by choosing smaller local uncertainty values, and by retaining component or branch labels when topological information is required.

B. Influence of the number of components and dimension

To put the proposed deterministic construction into context, we compare it against standard EM-based GMM fitting with the same number of components. We report both *accuracy* (via a Monte Carlo estimate of the KL divergence to a reference curve distribution) and *runtime*.

We consider a family of smooth Lissajous-type curves in $\mathbb{R}^d$ for $d \in \{2, 5, 10, 20, 50\}$. For each d and each $N \in \{4, 8, 16, 32, 64, 128, 256, 512\}$ we construct:

- a *geometric* GMM p_{geom} by applying the deterministic mapping (4) to the curve, and

© 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.

© 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.

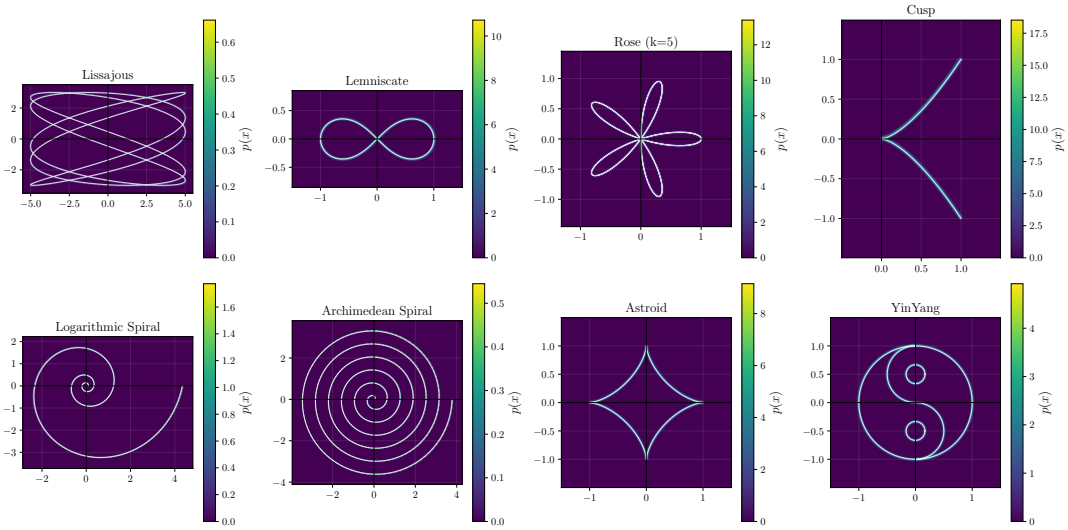


Fig. 2. Complex gallery of curve types (Lissajous, lemniscate, rose curve, Yin–Yang curve, logarithmic and Archimedean spirals, astroid, cycloid, semicubical cusp).

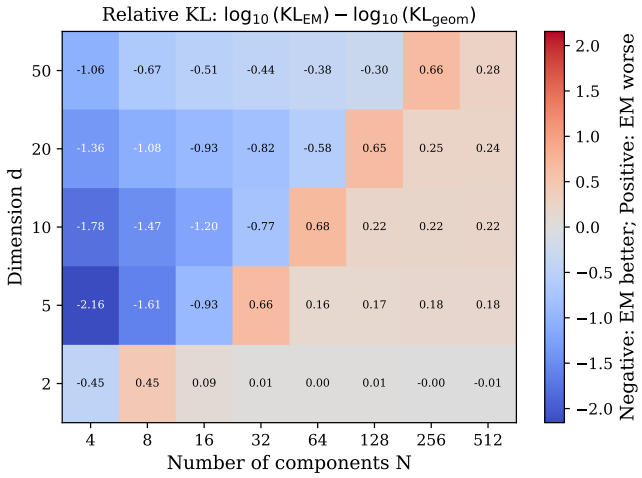


Fig. 3. Relative KL divergence for Lissajous-type curves in $\mathbb{R}^d$: $\Delta_{\mathrm{KL}}(d, N) = \log_{10} D_{\mathrm{KL}}(p_{\text{true}} \,\|\, p_{\text{EM}}) - \log_{10} D_{\mathrm{KL}}(p_{\text{true}} \,\|\, p_{\text{geom}})$. Positive values (red) indicate higher KL divergence for EM; negative values (blue) indicate higher KL divergence for the geometric method.

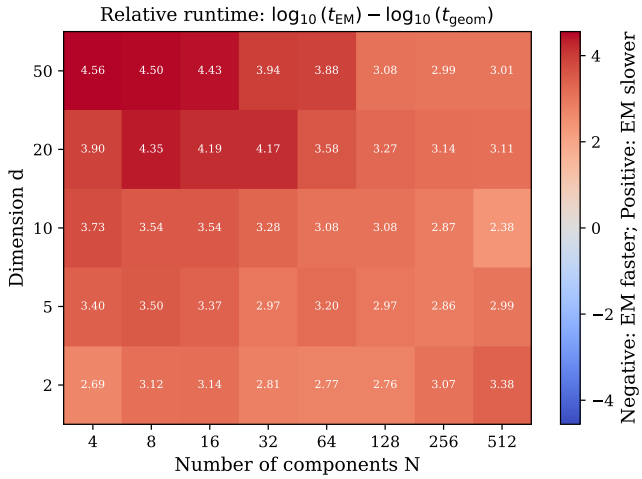


Fig. 4. Relative runtime: $\Delta_t(d, N) = \log_{10} t_{\text{EM}} - \log_{10} t_{\text{geom}}$. Positive values mean EM is slower.

- an *EM* GMM p_{EM} with N components fitted to samples drawn from the same underlying curve-with-normal distribution.

For both methods we measure the wall-clock runtime t and estimate the KL divergence

$$D_{\mathrm{KL}}(p_{\text{true}} \,\|\, p)$$

by Monte Carlo. To visualize relative performance, we plot the differences

$$\Delta_{\mathrm{KL}}(d, N) = \log_{10} D_{\mathrm{KL}}(p_{\text{true}} \,\|\, p_{\text{EM}}) - \log_{10} D_{\mathrm{KL}}(p_{\text{true}} \,\|\, p_{\text{geom}}),$$

and

$$\Delta_t(d, N) = \log_{10} t_{\text{EM}} - \log_{10} t_{\text{geom}},$$

as heatmaps in Fig. 3 and Fig. 4.

The KL heatmap in Fig. 3 shows three qualitative effects:

- For low dimensions ($d = 2$) and moderate to large N ($N \geq 32$), Δ_{KL} is close to zero, meaning both methods achieve very similar approximation quality.
- For higher dimensions and small N, Δ_{KL} is negative with magnitudes up to about 2, so EM attains a smaller KL (up to roughly two orders of magnitude in KL).
- As N increases, the differences rapidly shrink in all dimensions. Beyond a method- and dimension-dependent crossover (around $N \approx 32$ for $d = 5$, $N \approx 64$ for $d = 10$, and somewhat larger N for $d = 20$ and $d = 50$), Δ_{KL} becomes small and often positive, indicating that the geometric mixture matches or slightly improves on EM in KL while using the same number of components.

The runtime heatmap in Fig. 4 is uniformly positive, with values between approximately 2.4 and 4.5. Thus EM is

consistently between about $10^{2.4} \approx 250$ and $10^{4.5} \approx 30\,000$ times slower than the geometric construction across all tested (d, N). This behavior is consistent with the asymptotic scaling: for the geometric construction, forming N components in $\mathbb{R}^d$ requires computing N means and N covariance matrices, which costs

$$O(Nd^2)$$

floating-point operations. In contrast, an EM algorithm with $K = N$ components fitted to M samples over T iterations incurs

$$O\big(TMNd^2 + TNd^3\big)$$

operations: the first term from evaluating responsibilities in the E-step, the second from updating and (at least implicitly) inverting $d \times d$ covariance matrices in the M-step. For typical choices of M and T, this predicts an additional factor of order TM between EM and the geometric method, which is in line with the empirical runtime ratios of about $10^{2.4}$–$10^{4.5}$ observed in Fig. 4.

Overall, the scaling study indicates that:

- EM can reach a smaller KL in the regime of very few components and high dimension, but only at the cost of this much higher computational effort.
- For moderate and larger N, the geometric method is competitive or slightly better in KL, while remaining several orders of magnitude faster and completely deterministic.

In curve- and trajectory-based applications where moderate N and real-time operation are more relevant than minimizing KL at extremely small N, this leads to a more favorable accuracy–runtime tradeoff for the geometric construction.

C. Circle example

We illustrate the construction on a circle $\alpha(t) = r(\cos t, \sin t)$, $t \in [0, 2\pi]$, with radius $r > 0$. We choose N equally spaced parameters $t_i = 2\pi i/N$, $i = 0, \ldots, N$, set $v_i = \alpha(t_i)$, and form segments $S_i = (v_{i-1}, v_i)$ of length $\ell_i \approx 2\pi r/N$. For simplicity we take a constant normal uncertainty $\tau_i = \tau > 0$ for all segments.

For each S_i the associated Gaussian component has mean at the segment midpoint and covariance

$$\Sigma_i = \frac{\ell_i^2}{12}\, e_i e_i^\top + \tau^2 n_i n_i^\top,$$

where e_i and n_i are the unit tangent and normal of S_i, respectively. Length-based weights $\pi_i = \ell_i / \sum_{j=1}^{N} \ell_j = 1/N$ distribute probability uniformly over arc length.

As N increases, the component means form a regular polygon converging to the circle, and the principal axes of Σ_i remain aligned with the circle tangent and normal. For fixed τ, the resulting mixture approximates a probabilistic tubular neighborhood of the circle with essentially constant thickness; see Fig. 5. In a trajectory-planning context, this can be interpreted as a safety corridor of tunable thickness around a nominal circular path.

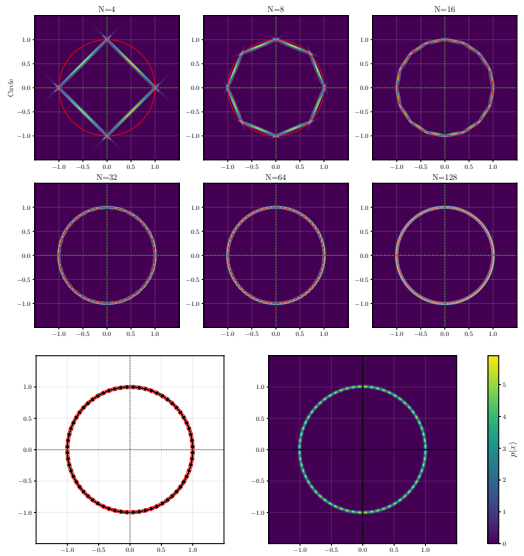


Fig. 5. Circle example. Top: convergence of the GMM PDF as the number of components N increases. Bottom: split view with component covariance ellipses overlaid on the circle (top subpanel) and the corresponding GMM PDF heatmap (bottom subpanel).

VI. DISCUSSION AND LIMITATIONS

A. Convergence and consistency

As the polygonal approximation is refined, segment midpoints converge to the curve and tangents to the curve's derivative where it exists, so the mixture components' means and principal axes track local geometry more closely. Under suitable regularity assumptions and appropriate refinement strategies, one expects the sequence of mixtures to approximate a probabilistic tubular neighborhood of the curve in a weak sense.

A rigorous analysis of convergence rates and consistency, including the impact of adaptive discretization and segments with $\tau_i \to 0$, requires additional work and is beyond the scope of this paper. A detailed theoretical treatment, including convergence results and further examples, is given in [16].

B. Beyond planar curves and the need for a calculus

Although we have visualized planar examples, the construction is defined for curves in $\mathbb{R}^d$ with arbitrary finite d. The extension to space curves and, conceptually, to surfaces in $\mathbb{R}^3$ is immediate, where components would encode local tangent spaces and normal direction uncertainty. However, once curves or surfaces are represented as GMMs in the ambient state space, classical differential geometric tools (e.g., connections, curvature) are no longer directly available in their standard form.

This suggests the need for a "calculus" on GMM representations that plays a role analogous to differential geometry on smooth manifolds, allowing geometric quantities to be expressed directly in terms of GMM parameters. Developing such a framework is an interesting direction for future work.

© 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.

VII. CONCLUSION

We introduced a geometric GMM representation of curves and trajectories in $\mathbb{R}^d$ based on a UDPP. Each segment is equipped with normal-subspace uncertainty, and an associated RV distributes mass uniformly along the tangent and Gaussian in the normal subspace. By matching the first and second central moments, we derive Gaussian components with means at segment midpoints and covariances aligned with local tangent and normal subspaces. Length based weighting then yields a GMM that encodes both the curve and its transverse uncertainty field.

The construction is arithmetical, does not require iterative fitting, and applies to smooth, non-regular, and self-intersecting curves. Experiments on canonical examples indicate that the GMMs capture local and global shape faithfully, while providing a probabilistic, differentiable representation suitable for uncertainty aware modeling in CAD and digital twins, robotics, and trajectory and shape analysis.

ACKNOWLEDGMENT

Funded by: Project SFB 1574 “A Circular Factory for the Perpetual Product”, Deutsche Forschungsgemeinschaft (DFG), Project ID 471687386.

REFERENCES

[1] R. Szeliski, *Computer Vision: Algorithms and Applications*, ser. Texts in Computer Science. Springer International Publishing, 2022.

[2] T. Nanayakkara, *Handbook on Soft Robotics*. Springer Nature Switzerland, 2024.

[3] G. Gerlach, K. D. Sommer, and C. Scholz, *Measurement Uncertainty*. De Gruyter, 2026.

[4] A. Darijani, J. Beyerer, Z. S. H. Nasrollah, L. Hoffmann, and M. Heizmann, “Comprehensive description of uncertainty in measurement for representation and propagation with scalable precision,” 2026. [Online]. Available: https://arxiv.org/abs/2603.20365

[5] I. L. Dryden and K. V. Mardia, *Statistical Shape Analysis*, ser. Wiley Series in Probability and Statistics. Wiley, 1998.

[6] S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.

[7] J. J. Risler, *Mathematical Methods for CAD*, ser. Recherches en mathématiques appliquées. Cambridge University Press, 1992.

[8] G. Farin, *Curves and Surfaces for CAGD: A Practical Guide*, ser. Computer graphics and geometric modeling. Elsevier Science, 2002.

[9] K. Tapp, *Differential Geometry of Curves and Surfaces*. Springer International Publishing, 2016.

[10] B. Eckart, K. Kim, A. Troccoli, A. Kelly, and J. Kautz, “Accelerated generative models for 3d point cloud data,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 5497–5505.

[11] W. Tabib, C. O’Meadhra, and N. Michael, “On-manifold gmm registration,” *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3805–3812, 2018.

[12] L. Keselman and M. Hebert, “Direct fitting of gaussian mixture models,” in *2019 16th Conference on Computer and Robot Vision (CRV)*, 2019, pp. 25–32.

[13] K. Goel and W. Tabib, “Gira: Gaussian mixture models for inference and robot autonomy,” 2024. [Online]. Available: https://arxiv.org/abs/2307.00071

[14] Y. Yuan, J. Liu, W. Chi, G. Chen, and L. Sun, “A gaussian mixture model based fast motion planning method through online environmental feature learning,” *IEEE Transactions on Industrial Electronics*, vol. 70, no. 4, pp. 3955–3965, 2023.

[15] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, “3d gaussian splatting for real-time radiance field rendering,” *ACM Transactions on Graphics*, vol. 42, no. 4, 2023.

[16] A. Darijani, B. Stratmann, and J. Beyerer, “A geometric gaussian mixture representation of plane curves,” 2026. [Online]. Available: https://arxiv.org/abs/2606.06505

© 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.