Generalized Conic Fitting for Interpretable Shape Representation with Applications to Biomedicine

Mathias Oehmen[1], Xiaoyi Jiang[1] *, and Walter G. Kropatsch[2]

[1] Department of Computer Science, University of Münster, Münster, Germany
[2] TU Wien, Vienna, Austria
{moehmen,xjiang}@uni-muenster.de, krw@prip.tuwien.ac.at

Abstract. Generalized conics, defined as continuous level sets of aggregated distance functions, provide a compact and interpretable shape representation using only a small number of focal points and scalar weights. While such multifocal representations have recently found applications in shape analysis, optimization, and robotics, fitting generalized conics to discrete 2D contours remains challenging due to optimization instabilities, often leading to disconnected predictions in practice. To address this limitation, we propose a decoupled two-phase optimization pipeline for generalized conic fitting. Instead of jointly optimizing all parameters directly on the contour, the proposed approach first stabilizes the global geometry through distance-field fitting and subsequently fine-tunes the focal points for improved boundary alignment. We demonstrate a biomedical application of the proposed pipeline on a curated benchmark of isolated *Drosophila melanogaster* larval deformations, showing that the method yields compact parameterizations of biological shapes while reducing topological artifacts during optimization.

Keywords: Generalized Conics · Shape Representation · Optimization · Parameter Decomposition · Biomedicine

1 Introduction

The geometric representation of deformable biological structures, such as bending and coiling *Drosophila* larvae or marine worms [1, 22], requires models capable of capturing continuous shape changes. While classical active contours represent boundaries accurately, they often require many control points and struggle with highly irregular geometries [3, 19]. Generalized conics provide a compact alternative by representing shapes as implicit distance fields defined through weighted sums of distances to multiple focal points [7]. Their focal structure yields interpretable and compact representations of biological postures [8]. However, fitting multi-focal conics to arbitrary deformations remains challenging. While recent methods successfully fit simpler generalized conic families [9, 10],

* The work was supported by the German Research Foundation (DFG: CRC 1450 "inSight" – 431460824). The first author thanks the EU ERASMUS traineeship program supporting his stay at TU Wien from February – April 2026.

unconstrained gradient-based optimization often destabilizes the underlying distance field, leading to fragmented or topologically inconsistent predictions.

To investigate the broader feasibility of generalized conics as compact, interpretable shape descriptors beyond classic 2D contour representation, we present the following contributions in this paper: (1) We introduce a novel gradient-based optimization pipeline for generalized conics fitting. To mitigate optimization instabilities of multi-focal fitting, the proposed pipeline employs a two-phase procedure comprising a global distance field fit, followed by fine-tuning only the focal points directly on the 2D boundary with frozen weights. (2) We demonstrate a biomedical application by evaluating our approach on a curated benchmark of isolated *Drosophila* larvae from the LCDseg dataset [1].

2 Related Work

Explicit vs. Implicit Shape Modeling. Explicit models define boundaries either through global frequency analysis, such as Fourier descriptors [17], or local control points like Active Contours [15, 19] and splines [3, 4], encoding only boundary geometry. Implicit representations like Level Sets [23] embed shapes as zero-level sets on a spatial grid, capturing topological changes at the cost of memory-intensive grid-based embeddings. This motivates continuous parametric models combining regional implicit properties with explicit parameter efficiency.

Generalized Conics. Generalized conics provide compact and interpretable analytic shape representations based on continuous level sets of aggregated distance functions, with applications in shape representation [5, 9], optimization [2], and robotics [25]. While previous work has shown that they can represent complex geometries, including sharp corners and concavities [7, 8], gradient-based fitting to arbitrary shapes remains largely unexplored. Unlike abstract Fourier coefficients or global moment invariants [12], the focal points and weights of generalized conics carry direct geometric meaning, making them interpretable. In this work, we investigate the feasibility of generalized conics as interpretable shape representations rather than aiming for a general-purpose benchmark. Although deep learning has become the dominant paradigm for medical image segmentation, it continues to benefit from advances in classical computer vision. In particular, shape representations can be integrated into deep learning pipelines. For example, [13] performs object segmentation by directly learning the Fourier coefficients of object masks, highlighting the potential of further exploring classical shape representations.

3 Generalized Conics and Optimization Challenges

Let $x \in \mathbb{R}^2$ and $\mathcal{F} = \{F_1, \ldots, F_k\}$ be a set of focal points. A generalized conic is defined by a scalar distance field aggregating weighted Euclidean distances to the foci [5, 7]:

$$D(x) = \sum_{i=1}^{k} w_i ||x - F_i||_2 \tag{1}$$

where $w_i \in \mathbb{R}$ are scalar weights and $\delta(x, F_i) = ||x - F_i||_2$ denotes the Euclidean distance. The resulting shape is defined as the implicit level set $L = \{x \in \mathbb{R}^2 \mid D(x) = c\}$. Positive weights generate multifocal ellipses, while negative weights introduce hyperbolic components. We refer to [5, 7, 9] for further details.

Given a set of discrete contour points $C = \{x_1, \ldots, x_n\}$, our goal is to estimate the optimal parameters $\theta = \{F_i, w_i\}$ such that the continuous level set L approximates the target geometry. Unlike traditional Level Set Methods [23], which evolve an implicit boundary on a memory-intensive spatial grid (Eulerian representation), our approach relies on an analytical, parametric representation. However, gradient-based optimization of this parameter space faces two practical optimization challenges.

Focal Singularity: Gradient-based fitting becomes singular when a focal point coincides with the contour, since $\partial\delta(\mathbf{x}, F_i)/\partial F_i$ is undefined for $\mathbf{x} = F_i$. However, sharp corners and concavities in generalized conics naturally require focal points on the boundary [5, 8], driving the optimizer toward this singularity. We therefore regularize the distance as $\delta(\mathbf{x}, F_i) = \sqrt{|\mathbf{x} - F_i|^2 + \epsilon}$ with $\epsilon = 10^{-8}$. While this removes the singularity, optimization near the boundary may still remain numerically unstable.

Scale Invariance: The generalized conic equation is scale-invariant. Multiplying both the target level c and the weights w_i by a constant yields the same geometric shape. While existing literature constrains this parameter space by normalizing the scalar weights against their maximum value [5], our continuous optimization pipeline reduces drift along equivalent scalar multiples by internally normalizing the target boundary level to $c = 1.0$ for all computations.

4 Methods

4.1 Two-Phase Optimization Pipeline

For improved optimization robustness, we propose a two-phase optimization pipeline consisting of a global spatial approximation (Phase 1) followed by a local boundary fine-tuning (Phase 2).

Phase 1 (Global Distance Field Optimization): Phase 1 approximates the global shape geometry over a discrete 2D spatial grid $\mathcal{G}$, corresponding to the pixel coordinates of the target image domain, with $p \in \mathcal{G}$ denoting individual pixel locations. To ensure numerical stability, this grid is normalized by centering it at the object's centroid and scaling it to a unit range. Following classical level-set embedding [23] and discrete distance fields for generalized conics [5, 6], we extract the Signed Distance Field (SDF) (positive outside, negative inside) from the binary segmentation mask. To align the object boundary with the target contour level $c = 1.0$ and ensure scale invariance, we define the normalized target field $D_{target}(x) = 1.0 + SDF(x)/D_{\max}$, where $D_{\max}$ denotes the maximum Euclidean distance from the shape centroid. We model the shape as a Confocal Multifocal Elliptic Field [5], parameterized by focal points $F_i = (x_i, y_i)$ and

scalar weights w_i. The global Distance Field Loss is defined as:

$$\mathcal{L}_{DF}(F, w) = \frac{1}{|\mathcal{G}|} \sum_{p \in \mathcal{G}} \left(\sum_{i=1}^{k} w_i \delta(p, F_i) - D_{target}(p) \right)^2 \tag{2}$$

Direct optimization over all parameters results in a nonlinear problem with $3k$ degrees of freedom ($2k$ spatial coordinates and k scalar weights). We follow the parameter decomposition framework of [14], which has proven effective in practical application scenarios [18]. The parameter space is split into spatial variables $\theta_F \in \mathbb{R}^{2k}$ and linear weights $\theta_w \in \mathbb{R}^k$. For fixed focal positions F, we construct the evaluation matrix $A \in \mathbb{R}^{|\mathcal{G}| \times k}$ with entries $A_{j,i} = \delta(p_j, F_i)$ and the vectorized target field $D_{target} \in \mathbb{R}^{|\mathcal{G}|}$. The loss reduces to a linear least-squares problem $\mathcal{L}_{DF}(w) = \frac{1}{|\mathcal{G}|} \|Aw - D_{target}\|_2^2$. Expanding this objective yields:

$$\mathcal{L}_{DF}(w) = \frac{1}{|\mathcal{G}|} (w^T A^T A w - 2 D_{target}^T A w + D_{target}^T D_{target}) \tag{3}$$

Since $A^T A$ is positive semidefinite for any matrix A, the objective is convex in w. Any minimizer satisfies the normal equations $A^T A w = A^T D_{target}$. To ensure numerical stability and avoid ill-conditioning caused by correlated distance evaluations, we apply Ridge Regression [11]:

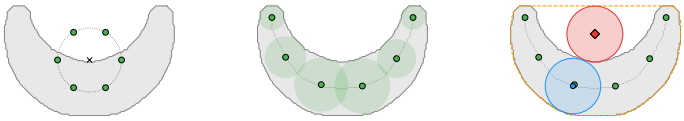
$$\hat{w} = (A^T A + \lambda I)^{-1} A^T D_{target} \tag{4}$$

Then, with $\hat{w}$ fixed, the focal points F_i are updated via gradient descent on $\mathcal{L}_{DF}$. This alternating optimization separates the linear subproblem from the nonlinear optimization, reducing the effective optimization space from $3k$ to $2k$ dimensions and improving optimization robustness.

Alternative regularizers like Lasso [24] or Elastic Net [26] remain interesting for future work, as they promote sparse weight estimates, effectively driving some weights to zero, to support adaptive model complexity with a variable number of active foci. However, they address a fundamentally different modeling objective than mitigating the linear subproblem's ill-conditioning. Furthermore, even for model adaptivity, weight sparsity alone is insufficient to identify geometrically redundant foci, since redundancy depends jointly on focal positions and weights. We therefore do not expect these methods to meaningfully improve upon Ridge Regression here.

Phase 2 (Boundary Fine-Tuning): While Phase 1 establishes the global geometry, Phase 2 refines boundary alignment using only the discrete contour points $x_j \in C$. Restricting optimization to the target isoline removes many global constraints of the distance field, making the problem under-constrained and potentially destabilizing the generalized conic into disconnected level sets [5, 7]. Therefore, the scalar weights remain fixed to the least-squares solution $\hat{w}$ obtained in Phase 1. The boundary refinement loss minimizes the squared error between the contour evaluations and the target isoline level $c = 1.0$:

$$\mathcal{L}_{Boundary}(F) = \frac{1}{n} \sum_{j=1}^{n} \left(\sum_{i=1}^{k} \hat{w}_i \delta(x_j, F_i) - 1 \right)^2 \tag{5}$$



(a) Circular (b) Spine Sampling (c) Hybrid

Fig. 1. Initialization strategies ($k = 6$). (a) Circular: Foci (green) on a radius (dashed) around the centroid (cross). (b) Spine: Foci along medial axis; weights match radii of the largest inscribed circles (green rings). (c) Hybrid: Adds a negative focus (red diamond) at the center of the largest inscribed circle (red ring) between the convex hull (orange) and the shape, scaled by max shape thickness (blue ring).

By updating only the focal points F_i via gradient descent on $\mathcal{L}_{Boundary}$, the optimization operates in the reduced $2k$-dimensional spatial parameter space. This helps preserve the coarse global topology established in Phase 1.

4.2 Boundary-Only Baseline

For comparison we will also evaluate standalone Phase 2 (*P2 Only*) as a baseline, applying the alternating decomposition exclusively on the discrete contour C with unfrozen weights. Unlike the sequential pipeline we initialize the weights to better control the early optimization trajectories. The weighting is done in a way that reflects the natural thickness variation along the body of an elongated biological structure [5]. Foci at wider body regions receive higher weights to account for the greater local extent of the shape, while foci near the tapered tips receive lower weights. Note that the weight initialization discussed in Section 4.3 is solely dedicated to *P2 Only*.

Given the initialization, each iteration first updates focal positions F_i via gradient descent on $\mathcal{L}_{Boundary}$. Weights are subsequently updated as:

$$\hat{w} = (A^\top A + \lambda I)^{-1} A^\top \mathbf{1} \tag{6}$$

where matrix $A \in \mathbb{R}^{n \times k}$ has entries $A_{j,i} = \delta(x_j, F_i)$ and $\mathbf{1} \in \mathbb{R}^n$ is the target level vector.

4.3 Initialization Strategies

To navigate the loss landscape of $\mathcal{L}_{DF}$ and guide optimization towards geometrically plausible solutions, we define three initialization strategies (see Fig. 1). Note that weight initialization is only needed for the Boundary-Only Baseline.

Circular (Baseline): This shape-agnostic strategy distributes a predefined number of foci k around the center of mass $\mu = \frac{1}{n}\sum_{j=1}^{n} x_j$ of the discrete contour points. To keep foci inside the contour, we set radius $r = \frac{1}{2n}\sum_{j=1}^{n} ||x_j - \mu||_2$. The initial focal coordinates are established as $F_i^{(0)} = \mu + r(\cos(\alpha_i), \sin(\alpha_i))^T$ with $\alpha_i = 2\pi(i-1)/k$ for $i = 1, \ldots, k$. Scalar weights are initialized inversely

proportional to their minimum Euclidean distance to the contour C, so foci closer to the boundary receive larger weights. Letting $d(F_i^{(0)}, C) = \min_{x \in C} ||F_i^{(0)} - x||_2$ and defining $\bar{d}$ as the mean of these distances, the weights are formalized as $w_i^{(0)} = \bar{d}/(d(F_i^{(0)}, C) + \epsilon)$, where $\epsilon = 0.1$ prevents numerical singularities.

Spine Sampling: This strategy exploits the medial axis of the biological structure. Let $B \subset \mathbb{Z}^2$ be the binary mask of the target shape. We compute the topological skeleton $S = \text{skel}(B)$ and extract a path P through S starting from a morphological endpoint of S. Assuming $k \geq 2$, foci are sampled at uniform index intervals along the discrete path P: $F_i^{(0)} = P(s_i)$ with $s_i = (L-1) \cdot (i-1)/(k-1)$ for $i = 1, \ldots, k$, where L denotes the total number of skeleton points. The initial weights reflect the local body thickness at each focus, measured by the radius of the largest inscribed circle fitting inside the shape at that point, i.e. $\tilde{w}_i = \text{EDT}_B(F_i^{(0)})$ (see Fig. 1(b)). They are normalized with an exponential decay to emphasize the narrower anatomical extremities: $w_i^{(0)} = (\tilde{w}_i / \max_l \tilde{w}_l)^{1.5}$. If skel($B$) yields no connected path, we default to Circular initialization (note that for the curated benchmark evaluated in Section 5, this fallback was never triggered).

Hybrid: Building on spine sampling, this method extends the focal set to capture deep morphological concavities using hyperbolic components ($w_i < 0$). Retaining the k positive medial foci, we compute the convex hull $H(B)$ and define the set of external bays as $B_{bays} = H(B) \setminus B$. To suppress discretization noise, the connected components of B_{bays} are filtered, keeping only M valid bays $V = \{v_1, \ldots, v_M\}$ that exhibit a maximum inscribed circle radius ≥ 5 pixels or an area ≥ 40 pixels. For each bay $v_j \in V$, an additional focus F_{k+j} is placed at the center of its largest inscribed circle (point inside v_j with maximum distance to the bay boundary ∂v_j): $F_{k+j}^{(0)} = \arg\max_{x \in v_j}(\min_{y \in \partial v_j} ||x - y||_2)$. Following [9], the negative weight scales with the inscribed circle radius at that point $d_j = \text{EDT}_{B_{\text{bays}}}(F_{k+j}^{(0)})$, relative to the maximum inscribed circle radius of the larva body $T_{\max} = \max_{x \in B} \text{EDT}_B(x)$, via $w_{k+j}^{(0)} = -\text{clip}(d_j/T_{\max}, 0.2, 0.9)$. This limits the influence of negative weights, expanding the focal set to $k_{\text{total}} = k + M$.

5 Experiments

We implement the optimization pipeline in PyTorch. To maintain numerical stability during optimization, we pre-normalize all spatial coordinates into a zero-centered $[-1, 1]$ unit space. For reproducibility, all experiments use fixed random seeds. For all initialization methods, we set $k = 6$, matching the six medial tracking points used in [1]. Spatial coordinates were optimized using Adam [16] for up to 10,000 epochs per phase with an early stopping patience of 20. Analytical weights were solved via Ridge Regression [11] using a fixed penalty of $\lambda = 10^{-5}$. Phase 1 utilized an Adam learning rate of 0.01 and a 10^{-6} loss tolerance, while Phase 2 applied a 0.005 learning rate and a 10^{-4} tolerance.

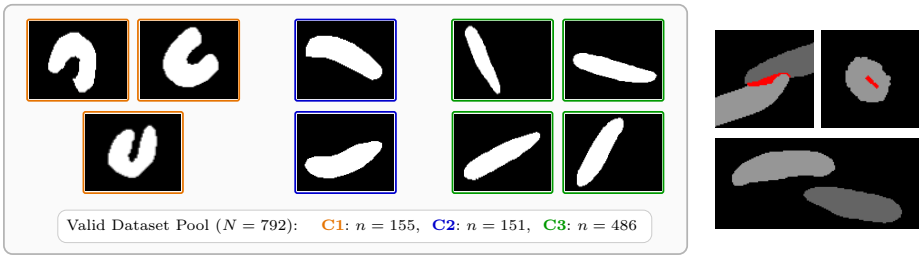


Fig. 2. Dataset curation and shape distribution. **Left:** The valid dataset clustered by shape (C1: Severely Bent, C2: Slightly Bent, C3: Straight). **Right Top:** Excluded invalid source frames (occlusions and holes). **Right Bottom:** Source frame with two valid shapes.

5.1 Dataset Curation: Drosophila Larvae

We evaluate our pipeline on the LCDseg dataset [1], which contains manually segmented masks of interacting *Drosophila* larvae recorded using a Frustrated Total Internal Reflection imaging setup [21]. This imaging modality enables high-contrast acquisition of larval body deformations and has enabled advanced biological studies [20]. We processed the complete set of 60 two-larvae collision sequences (1,989 frames), yielding an initial pool of 3,978 larval instances.

Curation was guided by the assumptions of the optimization pipeline. 3,160 instances containing touching or overlapping larvae were excluded. Additionally, 26 instances containing interior holes or contour artifacts were removed to favor simply connected target shapes for stable Signed Distance Field computation (Fig. 2, right top). This filtering yields a final evaluation pool of 792 valid larval instances across 60 sequences, each represented as an isolated binary mask. We partitioned the dataset into three morphological groups using K-Means clustering ($k = 3$) over two geometric descriptors: solidity (mask area over convex hull area) and aspect ratio (of the minimum bounding rectangle). After manually reassigning five outliers from the slightly-bent to severely-bent cluster, we obtained the final three posture classes: Severely Bent (C1, $n = 155$), Slightly Bent (C2, $n = 151$), and Straight (C3, $n = 486$) (Fig. 2, left).

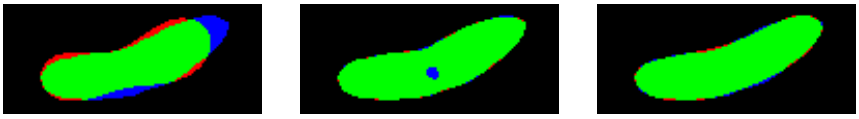
Because LCDseg provides pre-cropped sequences acquired under a fixed imaging setup, the spatial resolution remains constant across all samples. All optimization steps therefore operate directly on the native image resolution, yielding an average evaluation grid $\mathcal{G}$ of 160×173 pixels during dense global distance-field optimization (Phase 1).

5.2 Performance Evaluation

Evaluation metrics include the Jaccard Index (IoU), the Dice Coefficient, and the Maximum Hausdorff Distance (H_{Max}), following established protocols for geometric boundary fitting of implicit shape representations [10]. For pixel-wise evaluation, the continuous generalized conic field is evaluated at discrete

Table 1. Quantitative performance evaluation. Overall metrics are shown alongside Mean IoU broken down by shape complexity: Severely Bent (C1), Slightly Bent (C2), and Straight (C3). Failure rate: instances with IoU < 0.85. Topo. Invalid: predictions with Euler number $\neq 1$, split into holes and fragments (disconnected level sets). **Bold**: best per column.

Initialization	Pipeline	IoU ↑				Dice ↑	H_{Max} ↓	Fail % ↓	Topo. Invalid % ↓		
		Overall (Mean ± Std)	C1	C2	C3	(Mean ± Std)	(Mean ± Std)		Total	Holes	Frags.
Circular (Baseline)	P1 Only	0.691 ± 0.176	0.670	0.731	0.685	0.802 ± 0.143	17.28 ± 16.83	81.8%	24.9%	0.3%	24.2%
	P2 Only	0.643 ± 0.411	0.499	0.602	0.701	0.683 ± 0.400	24.38 ± 31.01	41.4%	35.4%	15.9%	19.9%
	P1 → P2	0.873 ± 0.156	0.788	0.889	0.896	0.923 ± 0.117	9.81 ± 17.52	22.6%	6.9%	3.3%	3.9%
Spine Sampling	P1 Only	0.702 ± 0.171	0.684	0.750	0.693	0.811 ± 0.140	16.32 ± 15.88	81.4%	19.2%	0.1%	18.8%
	P2 Only	0.921 ± 0.138	0.849	**0.943**	**0.938**	0.951 ± 0.110	6.00 ± 12.48	11.5%	15.0%	9.7%	5.3%
	P1 → P2	0.887 ± 0.140	0.791	0.907	0.912	0.933 ± 0.104	8.74 ± 17.29	19.8%	**4.2%**	1.5%	**2.7%**
Hybrid	P1 Only	0.761 ± 0.137	0.744	0.810	0.751	0.856 ± 0.108	11.86 ± 12.59	72.3%	11.7%	**0.0%**	11.6%
	P2 Only	**0.926 ± 0.137**	**0.895**	0.930	0.935	**0.954 ± 0.104**	**5.15 ± 10.75**	**10.7%**	17.3%	11.0%	6.8%
	P1 → P2	0.901 ± 0.132	0.811	0.916	0.925	0.941 ± 0.098	7.97 ± 14.89	16.9%	6.7%	2.3%	4.5%



P1 (IoU: 0.800) P2 (IoU: 0.964) P1→P2 (IoU: 0.953)

Fig. 3. Topology–metric trade-off. **Left:** P1 preserves coarse shape. **Middle:** P2 maximizes IoU but introduces topology-breaking artifacts. **Right:** P1→P2 preserves structural consistency. TP (green), FP (red), FN (blue).

pixel centers of the image grid. Pixels satisfying $D(x) \leq 1.0$ are classified as foreground, all others as background. A segmentation is considered a failure if IoU < 0.85. Table 1 reports quantitative results across initialization strategies and morphological shape clusters (C1–C3), comparing global-only (*P1 Only*), boundary-only (*P2 Only*), and sequential (*P1→P2*) optimization.

Table 1 reveals a consistent trade-off between structural stability and peak performance across initialization strategies and shape complexities (C1–C3).

P1 Only consistently yields the lowest contour accuracy. Since Phase 1 evaluates the dense target field rather than the contour directly, it remains comparatively insensitive to initialization, producing similarly smoothed results across all three strategies. It captures coarse geometry but fails on fine boundary details, particularly for Severely Bent (C1) shapes. Topologically, *P1 Only* is virtually hole-free across all 792 instances ($\leq$ 0.3% holes), with invalid predictions almost exclusively taking the form of fragments (up to 24.2%), consistent with the smooth monotonic structure of the SDF target suppressing interior field maxima.

P2 Only is highly sensitive to initialization and shape complexity. Under Hybrid initialization, it yields the highest precision metrics, particularly on severely bent (C1) shapes. However, paired with Circular initialization, it fails to converge, resulting in a 41.4% IoU failure rate and a 35.4% topological failure rate, the highest across all configurations, suggesting that many failures are associated with topological artifacts. Even under favorable initializations, high quantitative performance does not necessarily reflect structural correctness, as illustrated in

Fig. 3: under Spine initialization *P2 Only* achieves high IoU overall, yet introduces a clear internal hole. The two phases exhibit structurally distinct failure signatures: *P1 Only* produces primarily fragments ($\leq 0.3\%$ holes, up to 24.2% fragments), while *P2 Only* introduces holes (up to 11.0% under Hybrid, 9.7% under Spine). The sequential pipeline suppresses both, reducing total topological failures to 4.2%–6.9%.

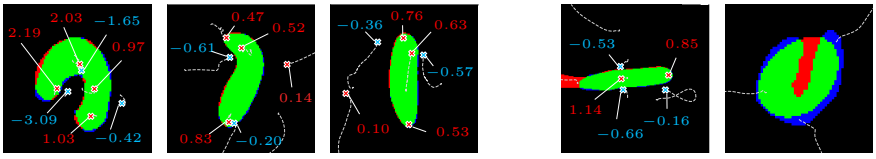
As detailed in Table 1, the two-phase approach achieves the most consistent overall behavior across the entire dataset, balancing accuracy and stability across all morphological clusters (C1-C3). It improves the performance of the unstable Circular baseline to a competitive mean IoU of 0.873 and yields stable performance for both Spine and Hybrid initializations. By combining global and local objectives, the sequential pipeline suppresses both failure modes characteristic of the individual phases, namely the fragmentation of global-only fitting, and the internal holes of unconstrained boundary optimization.

Computational Cost. The computational cost of the proposed pipeline is dominated by Phase 1, whose per-epoch cost scales linearly with the grid size $\mathcal{G}$ for any fixed number of foci k. Evaluated on an 8-Core AMD Ryzen 7 7800X3D and 32 GB RAM the average runtime of the complete *P1→P2* pipeline was 17.03 s for Circular initialization, 12.55 s for Spine Sampling, and 12.02 s for Hybrid. While all initialization strategies achieve comparable segmentation quality after Phase 1, shape-aware initializations satisfy the fixed early stopping criterion substantially sooner. Hybrid and Spine required only 1232 and 1412 epochs, respectively, compared to 2536 for Circular.

Hyperparameter Robustness. We also assessed hyperparameter robustness via a sensitivity analysis on a stratified subset of 15 representative shapes (5 per cluster). Varying the Ridge penalty λ across six orders of magnitude (10^{-7} to 10^{-2}) yielded a negligible mean IoU variation (≤ 0.005) and identical failure rates for Spine Sampling and Hybrid initialization. The less stable Circular baseline showed higher sensitivity to λ (33.3%–46.7% failure rates), though this variation was concentrated strictly on borderline instances near the IoU failure threshold. Scaling the heuristic geometric thresholds of the Hybrid initialization (minimum inscribed radius and minimum bay area, Section 4.3) by $\pm 50\%$ similarly produced only a minor IoU variation of 0.006 without affecting the failure rate. Overall, these results support that segmentation quality primarily depends on the initialization strategy rather than the precise hyperparameters.

5.3 Qualitative Analysis and Structural Integrity

We qualitatively analyze the Hybrid *P1→P2* configuration, which achieves the most robust performance across all shape clusters. Fig. 4(a–c) shows successful fitting of severely bent (C1), slightly bent (C2), and elongated (C3) larval shapes, including frozen weights from Phase 1 and focal trajectories during Phase 2. The remaining failure cases (Fig. 4d–e) reveal two limitations caused by fixed scalar weights during Phase 2. **Overshoot:** At tapered extremities (Fig. 4d), fixed scalar weights may mismatch the required local curvature, causing focal



(a) C1; IoU: 0.926 (b) C2; IoU: 0.967 (c) C3; IoU: 0.981 (d) Overshoot (e) Bridging

Fig. 4. Qualitative results (Hybrid initialization). **(a–c):** Successful abstractions. **(d–e):** Limitations (overshoot and bridging). Dotted lines indicate focal trajectories during Phase 2; numbers denote frozen scalar weights. Pixels: TP (green), FP (red), FN (blue).

points to move outside the contour and the predicted shape to overshoot the biological boundary. **Bridging:** For deep, narrow concavities (Fig. 4e), the isoline may bridge the gap instead of tracking the indentation. To preserve the global shape, Phase 1 generally establishes a balanced weight distribution. However, forming sharp concavities requires strong local adaptation of negative weights [5, 8]. Manual inspection of the 134 failure cases (IoU < 0.85, 16.9%) reveals overshoot as the dominant failure mode (86/134), followed by bridging (23/134), fragments (14/134), general imprecision (6/134), and holes (5/134).

6 Conclusion

This paper presents a gradient-based framework for fitting generalized conics as compact and interpretable shape representations. To mitigate optimization instabilities in multi-focal fitting, we introduce a two-phase strategy separating global distance-field regression from local boundary refinement. Experiments on the LCDseg dataset show that the proposed pipeline substantially improves topological stability, reducing failure rates from up to 35.4% to 4.2%–6.9% across all initializations while preserving connected single-component shapes. Future work will investigate adaptive focal structures, multi-contour decompositions, comparisons against alternative contour descriptors or learned shape representations, and broader geometric applications.

References

1. Bian, A., Jiang, X., Berh, D., Risse, B.: Resolving colliding larvae by fitting ASM to random walker-based pre-segmentations. IEEE/ACM Transactions on Computational Biology and Bioinformatics **18**(3), 1184–1194 (2021)
2. Blanco, V., Puerto, J.: Covering problems with polyellipsoids: A location analysis perspective. European Journal of Operational Research **289**(1), 44–58 (2021)
3. Brigger, P., Hoeg, J., Unser, M.: B-spline snakes: a flexible tool for parametric contour detection. IEEE Transactions on Image Processing **9**(9), 1484–1496 (2000)
4. Delgado-Gonzalo, R., Unser, M.: Spline-based framework for interactive segmentation in biomedical imaging. IRBM **34**(3), 235–243 (2013)
5. Gabdulkhakova, A.: From circles to generalized conics:: Enriching the properties of 2D shape representation and description. Ph.D. thesis, Technische Universität Wien (2022)

6. Gabdulkhakova, A., Kropatsch, W.G.: Confocal ellipse-based distance and confocal elliptical field for polygonal shapes. In: ICPR. pp. 3025–3030 (2018)
7. Gabdulkhakova, A., Kropatsch, W.G.: Generalized conics: Properties and applications. In: ICPR. pp. 10728–10735 (2020)
8. Gabdulkhakova, A., Kropatsch, W.G.: Generalized conics with the sharp corners. In: Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications: 25th Iberoamerican Congress CIARP. pp. 419–429 (2021)
9. Gabdulkhakova, A., Kropatsch, W.G.: Geometry-based approach to find the egg-shape parameters. In: ICPR (2026), to appear
10. Hladůvka, J., Kropatsch, W.G.: Fitting egg-shapes to discretized object boundaries. In: Discrete Geometry and Mathematical Morphology. pp. 107–119 (2024)
11. Hoerl, A.E., Kennard, R.W.: Ridge regression: Biased estimation for nonorthogonal problems. Technometrics **12**(1), 55–67 (1970)
12. Hu, M.K.: Visual pattern recognition by moment invariants. IRE Trans. Inf. Theory **8**(2), 179–187 (1962)
13. Jeon, Y.S., Yang, H., Feng, M.: FCSN: global context aware segmentation by learning the Fourier coefficients of objects in medical images. IEEE Journal of Biomedical and Health Informatics **28**(3), 1195–1206 (2024)
14. Jiang, X.: A decomposition approach to geometric fitting. In: IAPR Conference on Machine Vision Applications (MVA). pp. 467–470 (2000)
15. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. International Journal of Computer Vision **1**(4), 321–331 (1988)
16. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
17. Kurnianggoro, L., Wahyono, Jo, K.: A survey of 2d shape representation: Methods, evaluations, and future research directions. Neurocomputing **300**, 1–16 (2018)
18. Mazzini, F., Kettler, D.T., Guerrero, J., Dubowsky, S.: Tactile robotic mapping of unknown surfaces, with application to oil wells. IEEE Trans. Instrum. Meas. **60**(2), 420–429 (2011)
19. McInerney, T., Terzopoulos, D.: Deformable models in medical image analysis. In: Workshop on Mathematical Methods in Biomedical Image Analysis. pp. 171–180 (1996)
20. Otto, N., Marelja, Z., Schoofs, A., Kranenburg, H., Bittern, J., Yildirim, K., Berh, D., Bethke, M., Thomas, S., Rode, S., Risse, B., Jiang, X., Pankratz, M., Leimkühler, S., Klämbt, C.: The sulfite oxidase shopper controls neuronal activity by regulating glutamate homeostasis in drosophila ensheathing glia. Nature Communications **9**(1), 3514 (2018)
21. Risse, B., Otto, N., Berh, D., Jiang, X., Kiel, M., Klämbt, C.: FIM^{2c}: Multicolor, multipurpose imaging system to manipulate and analyze animal behavior. IEEE Transactions on Biomedical Engineering **64**(3), 610–620 (2017)
22. Sansone, C., Pucher, D., Artner, N.M., Kropatsch, W.G., Saggese, A., Vento, M.: Shape normalizing and tracking dancing worms. In: S+SSPR. pp. 390–400 (2016)
23. Sethian, J., Malladi, R., Vemuri, B.: Shape modeling with front propagation: A level set approach. IEEE TPAMI **17**(2), 158–175 (1995)
24. Tibshirani, R.: Regression shrinkage and selection via the lasso. J. R. Stat. Soc. Ser. B **58**(1), 267–288 (1996)
25. Tsalik, R., Kedarisetty, S., Shima, T.: N-point polyelliptic guidance. Journal of Guidance, Control, and Dynamics **48**(2), 212–225 (2025)
26. Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. J. R. Stat. Soc. Ser. B **67**(2), 301–320 (2005)