

Special Section on 3DOR2026

SHREC2026: Human perceived visual complexity of 3D shapes[☆]

Roberto M. Dyke^a,^{*}, Yang Deng^b, Charlotte Brassey^c, Yu-Kun Lai^b, Paul L. Rosin^b,^{*}, James Gardiner^c, Atsushi Tatsuma^d, Christoforos Vlachos^e, Anthony Belessis^e, Ioannis Romanelis^e, Vlassis Fotis^e, Konstantinos Moustakas^e, Ihar Volkau^f, Sergei Krasovskii^g, Marco Guerra^h, Ulderico Fugacci^h, Simona Cutruzzulà^{h,i}, Michela Mortara^h, Michela Spagnuolo^{h,1}, Ivan Sipiran^j, Isaac Aguirre^j

^a Univ Toulouse, CNRS, Toulouse INP, IRIT (CMR 5505) Toulouse, France

^b Cardiff University, Abacws, Senghennydd Road, Cathays, CF24 4AG, Cardiff, Wales, United Kingdom

^c Manchester Metropolitan University, John Dalton Building, Chester Street, Manchester, M1 5GD, United Kingdom

^d National Institute of Technology, Tsuyama College, 624-1 Numa, Tsuyama, Okayama, Japan

^e Department of Electrical and Computer Engineering, University of Patras, Patras, Greece

^f Alfa Algorithms, Singapore

^g Belarussian State University, Minsk, Belarus

^h Istituto di Matematica Applicata e Tecnologie Informatiche, Consiglio Nazionale delle Ricerche (CNR-IMATI), Genoa, Italy

ⁱ Department of Mathematics, University of Genoa, Genoa, Italy

^j Department of Computer Science, University of Chile, Santiago, Chile

ARTICLE INFO

Keywords:

3D visual complexity
Perceptual complexity
Shape analysis
Comparative study

ABSTRACT

The estimation of visual complexity in 3D remains a relatively under-studied problem. An aspect of this deficiency is due to the lack of suitable publicly available benchmarks. In this work, three datasets are used to investigate this problem from different angles: (1) a subset of the ABC dataset of CAD shapes, augmented with perceived visual complexity rankings; (2) an existing primate tooth dataset, with complexity rankings derived from established domain attributes; and (3) a novel dataset of fractal shapes with perceived visual complexity rankings. Ground truth rankings for (1) and (3) were derived from human judgments obtained from two studies in which lay individuals performed two-alternative forced-choice comparisons between pairs of shapes. We then invited researchers to submit results for algorithms designed to model shape complexity. We find that most methods that operate directly in 3D yield low, positive correlations; whereas—despite working on 2D projections derived from the original geometry—the best results were obtained by a subset of 2D methods.

1. Introduction

Determining the visual complexity of 3D shapes is a sophisticated problem. While there exists a corpus of literature that examines 2D shape complexity [e.g., 1–4], 3D shape complexity is still relatively under-studied. Furthermore, our understanding of how humans intuitively evaluate complexity in 3D remains limited [5]. Generally, the perceptual complexity of a shape is considered to be a multi-faceted

construct involving quantitative factors (e.g., number of elements) and structural factors (e.g., symmetries) [6].

The subject of complexity has been investigated in a broad range of domains. As with other constructs, such as saliency and similarity, the meaning of complexity changes depending on the context. The reliable operationalization of complexity using geometry-based metrics is desirable in a variety of domains (e.g., manufacturing [7], orthography [8], marine ecology [9], marketing [10], etc.).

[☆] This article is part of a Special issue entitled: 'CAG_SS_3DOR2026' published in Computers & Graphics.

^{*} Corresponding authors.

E-mail addresses: roberto-marco.dyke@irit.fr (R.M. Dyke), dengy29@cardiff.ac.uk (Y. Deng), c.brassey@mmu.ac.uk (C. Brassey), laiy4@cardiff.ac.uk (Y.-K. Lai), RosinPL@cardiff.ac.uk (P.L. Rosin).

URLs: <http://www.robertodyke.com> (R.M. Dyke), <https://isyangdeng.github.io/> (Y. Deng), <https://www.mmu.ac.uk/staff/profile/dr-charlotte-brassey> (C. Brassey), <https://yukunlai.github.io/> (Y.-K. Lai), <https://rosinpl.github.io/> (P.L. Rosin).

¹ Given his/her/their role as Associate Editor, Michela Spagnuolo had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to another journal editor.

A comparative study of complexity measures for 3D shapes was recently conducted in [11]. Rather than performing human-based studies, proxy tasks such as counting the number of parts/segments a shape is composed of were used. Here, we employ two datasets that are grounded by human annotation, as well as a third dataset that is based on an established technique [12] for assessing the complexity of primate teeth.

We observe that the lack of publicly available benchmarks with suitable ground truths may be limiting progress on this topic. Inspired by prior studies [11,13], we seek to establish a set of benchmark datasets that are publicly available for future studies. Due to the multifaceted nature of complexity, we examine the problem in multiple contexts by utilizing three thematically different datasets. We then invited researchers to collaborate on this work to contribute results of methods that are designed to estimate complexity.

The main contributions of this work are as follows:

- A novel fractal shape dataset with perceived complexity ground truth, and perceived complexity rankings contributed to two existing datasets (ABC shapes; primate teeth), all made publicly available.
- The presentation of several novel methods for estimating shape complexity.
- A systematic evaluation of various methods to the problem of measuring shape complexity in 3D.

In the following sections: we discuss previous benchmarking datasets and other works that are related to the problem of visual complexity in Section 2. In Section 3, we describe the curated datasets, which are used to evaluate the complexity estimation methods described in Section 4. We define a set of metrics for evaluating the methods and discuss the resulting evaluations in Sections 5 and 6.

2. Prior works

While much empirical research on 2D visual complexity has involved the collection of human-grounded data [e.g., 1,4], in this section, we focus on notable works that involve the use of 3D benchmark datasets.

Our work extends the prior SHREC track on quantifying shape complexity [11]. Since few datasets for this problem existed at the time, the track proposed three new datasets, two of which consisted of synthetic shapes. The authors hypothesized that shape complexity may be characterized by: (1) adding noise; (2) estimating the number of segments found via the proxy task of mesh segmentation; or (3) composing simple shapes. However, (1) and (2) assume that the complexity of a shape transitively corresponds to the parametric (or simple statistical) representation. While the labels for (2) were manually created by eleven individuals, the small number of annotators and subjective nature of the problem may limit reliability.

Other studies have examined specific factors that contribute to perceived 3D shape complexity. Sukumar et al. [13] conducted a user study to analyze the effects of surface variation, symmetry, part count, part decomposability, intricate details, and topology on complexity. This work relied on several datasets created by the authors; unfortunately, the data is not publicly available, limiting its use for evaluating future methods.

Ben Slama et al. [7] recently conducted a human-based study examining complexity in additive and subtractive manufacturing. Fifty experts were asked to assign complexity scores to digital models of 25 workpieces, from which a ranking was derived.

A common limitation in previous work is the lack of suitable human-grounded benchmarks [14,15]. In fact, Saleem et al. [15] highlight the potential utility of a user study to establish a shape complexity database as part of their discussions.

3. Dataset

We employ three 3D-shape datasets: a subset of the ABC dataset [16] augmented with visual complexity rankings, the primate tooth dataset [12] with complexity derived from existing domain attributes, and a novel fractal shape dataset with visual complexity rankings. These datasets cover distinct notions of complexity. While an algorithm may align well with human judgments on one dataset, this may not hold in general; therefore, relying on a single dataset would give an incomplete picture.

The *ABC (subset) dataset* is a large (500 shape) dataset that contains many CAD shapes of varying complexity (described in Section 3.1). The ABC dataset's scale and diversity of engineering geometries provide a context for assessing generalization across real design practices. This complements the domain-specific (tooth) and carefully controlled (fractal) contexts by grounding complexity measurement in a multi-featured setting. The *primate tooth dataset* contains 107 samples that have been attributed with an expected level of complexity by an expert in zoology (detailed in Section 3.2). This dataset is particularly interesting because it represents an established real-world application of complexity for understanding a biological phenomenon, namely primate diet. The specimens were collected in a single survey and scanned using μ CT scanning of casts of 25 species of platyrrhine (New World) primate lower second molar teeth. Finally, the *fractal dataset* (Section 3.3) consists of 24 diverse models generated using Mandelbulb 3D. This final dataset was chosen to mitigate familiarity bias (discussed in 17) that may be present in the other two datasets as these shapes are devoid of familiar structures.

All the data is available online from the track's webpage <http://complexity.robertodyke.com/shrec>, or is alternatively available from the Cardiff University Open Data repository (DOI: [10.17035/cardiff.32113027](https://doi.org/10.17035/cardiff.32113027)). In the case of the ABC and primate tooth datasets, scripts are provided to obtain the meshes from their respective repositories.

Ground truth collection. Excluding the primate tooth dataset, we ran two independent user studies on the ABC and fractal datasets. In each study, users were shown pairs of samples from the same dataset and asked to choose which shape was more visually complex using the *two-alternative forced choice* (2AFC) method, with no option for equality. The study was made available via a web-based application. Participants were non-expert annotators. We avoid domain expertise to prevent bias toward non-visual conceptions of complexity. Since a single view image may occlude details, an interactive renderer was developed to allow users to rotate shapes. Eye-dome lighting [18] was used to enhance depth perception (Figs. 1, 2 & 3). Furthermore, to ensure that participants thoroughly examined the shapes, users could not select an option until at least 60% of both shapes were visible from sampled camera positions.

To ensure participants understood the task, a series of practice pairs with clear differences and immediate feedback were presented at the start. Due to the large size of the datasets, it was not practical for users to evaluate all combinations of shape pairs. Therefore, the collected pairwise choices represent a sparse subset of the datasets.

To obtain a ranking for each set of pairwise comparisons, we fit a Bradley-Terry model [19]. A key limitation—albeit necessary—of our approach is the assumption that visual complexity behaves as a transitive property. We contend that focusing on narrowly defined shape classes (e.g., CAD parts, teeth, fractals) helps mitigate the impact of this assumption, as shapes within a narrow domain tend to vary along fewer, more consistent complexity-driving factors, reducing the likelihood of intransitive judgments.

3.1. ABC dataset

The ABC subset contains shapes selected from [16] (Fig. 1). The original dataset includes duplicates and many minor variations of




Fig. 1. Example shapes from the ABC dataset ordered by visual complexity from (a)–(e).




Fig. 2. Example shapes from the tooth dataset ordered by ground-truth complexity from (a)–(e).

simple primitives. We prune these systematically to retain only visually unique, non-planar single-component shapes.

We performed a two-stage user study with 65 participants in total. One participant was excluded as their results did not pass the control tests. In the first stage, 35 participants each evaluated 100 pairs of shapes from a collection of 600 curated shapes.

We pruned the 600 curated shapes to a 500-shape benchmark for evaluation, keeping the remaining 100 for training. This procedure was intended to improve the distribution of complexities amongst the shapes. We removed shapes iteratively as to minimize the increase in the Cramér–Rao lower bound; equivalently, this preserves the Fisher information of the Bradley–Terry model parameters in the remaining set. The average parameter-wise estimation deviation decreased from 1.499 to 1.033, indicating the removed shapes were weakly connected in the comparison graph—thus degrading the overall estimation quality. In the second stage, an additional 31 participants were invited to compare 100 pairs of shapes each on the now smaller collection.

3.2. Primate tooth dataset

Dental topographic analysis has established that an animal's diet imposes functional demands on the surface of its teeth, and that these demands manifest as differences in measurable shape complexity (e.g., small-scale surface features and overall roughness) [12]. We therefore use diet as a domain-grounded proxy for the expected visual complexity of a tooth.

To generate an *a priori* 'expected tooth complexity' ordering from diet, each species was assigned to a broad dietary class (as per Winchester et al. [12]) and classes were assigned an order reflecting our expectations of functional demands (folivory ranked highest *i.e.*, most 'complex', followed by insectivory/omnivory, then frugivory, then hard-object (*i.e.*, seeds, nuts) feeding). In geometric terms, these dietary categories correspond to expected differences in surface complexity, with the shearing-driven action of folivory producing highly crenulated, high-relief surfaces, insectivory/omnivory intermediate, and more crushing-driven of frugivory and hard-object feeding associated with smoother, lower-relief surfaces. To split ties within each broad class, we then calculated a 'dietary contribution' score from the percentage contribution of each food component, as per EltonTraits 1.0 [20].

$$\text{Dietary contribution score} = (1 \times \text{PlantO}) + (0.7 \times A) \\ + (0.4 \times \text{Fruit \& Nect}) + (0.1 \times \text{Seed}),$$

where PlantO, plant other = EltonTraits category 'Plant Other', *i.e.*, leaves/foilage; A, animal = the sum of EltonTraits categories 'Inv'




Fig. 3. Example shapes from the fractal dataset ordered by visual complexity from (a)–(e).

(invertebrates), 'Vend' (vertebrate endotherm), 'Vect' (vertebrate ectotherm), 'Vfish' (vertebrate fish), 'Vunk' (vertebrate unknown) and 'Scav' (scavenged carcass).

Species were then ranked by (i) broad class, and then within each class by (ii) dietary contribution. If species remained tied after (ii), tied ranks were retained. Where subspecies lacked diet percentages from EltonTraits, we assigned the corresponding species-level diet composition.

3.3. Fractal dataset

From a collection of reconstructed fractals, a set of 24 shapes was chosen. Combinations of functions, such as Integer Power, Bul-box, and IQ-bulb, were used to vary the geometry. We limited our selection to relatively compact shapes. The shapes are visually abstract, which makes it difficult to attribute notions of complexity based on non-visual qualities, such as 'manufacturing complexity' (see examples in Fig. 3).

To construct the ground truths, a similar approach to that of the ABC dataset was taken. In this case, however, the user study was performed in a single stage. Due to the small size of the dataset, only ten individuals were asked to each evaluate 65 pairs of shapes. Small components were removed and all fractals were uniformly simplified to contain 200,000 faces. This was chosen as a reasonable balance between visual fidelity and mesh complexity.

4. Methods

Excluding the organizers, six international research groups participated in the track. From the announcement of the challenge, participants had 11 weeks to submit their results. In this section, we describe the methods evaluated.

4.1. Entropy-based shape complexity estimation from local and global point cloud features (Atsushi Tatsuma)

Method name: Tatsuma

To capture shape complexity, we propose a method based on local and global features extracted from point cloud representations.

Since the dataset consists of 3D mesh models, we sample 30,000 surface points from each 3D mesh using the method of Osada et al. [21], where Sobol's quasi-random sequence [22] is used for sampling.

From the point cloud, we extract *Signature of Histograms of Orientations* (SHOT) descriptors [23] as local features. To summarize multiple local features into a scalar measure of shape complexity, we use eigenvalues of their covariance matrix. These eigenvalues represent the variance of the principal components of the features. We assume that complex shapes scatter features across many principal components, while simple shapes concentrate them in fewer components. Therefore, we use the normalized information entropy of the eigenvalues as a measure of shape complexity. Let $\lambda_1, \lambda_2, \dots, \lambda_N$ be the eigenvalues of




Fig. 4. Overview of the SCI-FIVE method.

the covariance matrix. We normalize them as $p_i = \lambda_i / \sum_{j=1}^N \lambda_j$, and compute the normalized entropy as

$$H = -\frac{1}{\log_2 N} \sum_{i=1}^N p_i \log_2 p_i,$$

where N is the dimensionality of the descriptor.

For global features, we use extended *Fast Point Feature Histograms* (FPFH) [24]. The extended FPFH consists of sub-histograms that represent the geometric features of the point cloud. For each sub-histogram, we normalize the bin values and compute the normalized information entropy. We adopt the mean of these sub-histogram entropies as a measure of global shape complexity.

Finally, we adopt the mean of the local and global entropies as a measure of shape complexity.

4.2. SCI-FIVE: Shape complexity indicator using feature integration and view evaluation (Christoforos Vlachos, Anthony Belessis, Ioannis Romanelis, Vlassis Fotis and Konstantinos Moustakas)

Method name: SCI-FIVE

An overview of the applied method may be seen on Fig. 4. Each object is transformed by Poisson-Disk Sampling into a point cloud containing 10,000 points. Each point cloud is then segmented using a standard SED-Net [25] and the resulting parts are extracted, after which, a set of common complexity features is calculated for each part and normalized appropriately. These features are combined in a weighted sum, yielding an initial, per part complexity score. To obtain a rough complexity score for the entire object, the per part complexity scores are aggregated, weighted by the fraction of points represented by each part, and normalized by the number of parts in the object.

At the same time, different views are obtained both from the initial objects, and each part, and their respective view similarity scores are calculated to encode (approximate) symmetry—which is closely linked to low object complexity. Lastly, these are used to refine the initial complexity score.

Finally, the number of recognized segments is itself used as a measure of complexity, to refine the obtained complexity score on each object.

The overall complexity score of a given object is given by:

$$\text{Complexity score} = \frac{1}{n_{\text{parts}}} \sum_{i=1}^{n_{\text{parts}}} \left(s_p \sum_i w_i f_i \right) + \left(w_{\text{vs}_{\text{parts}}} \times \text{vs}_{\text{parts}} \right)$$

$$+ \left(w_{\text{vs}_{\text{full}}} \times \text{vs}_{\text{full}} \right) + \left(w_{n_{\text{parts}}} \times n_{\text{parts}} \right),$$

where $f_i \in F$ (set of features) and w_i is the respective weight (discussed later). s_p is the number of points in each part (analogous to surface area, due to the Poisson Disk Sampling), vs_{parts} , vs_{full} are the view similarity penalty scores, and n_{parts} is the number of parts segmented.

Complexity features. Commonly adopted complexity features were employed in order to adequately judge the perceived complexity of each object. These may be divided into two classes: *Principal Component Analysis* (PCA)- and curvature-based.

For each part, PCA yields eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$. The following features are then calculated:

$$\text{Linearity} = (\lambda_1 - \lambda_2) / \lambda_1, \quad \text{Entropy} = -\sum_{i=1}^3 \lambda_i \log \lambda_i,$$

$$\text{Omnivariance} = (\lambda_1 \lambda_2 \lambda_3)^{1/3}.$$

The principal curvatures k_1 and k_2 are also computed, along with the mean curvature $H = \frac{k_1 + k_2}{2}$ and the Gaussian curvature $K = k_1 k_2$. Then, a number of relevant statistics are computed: the mean $\mu_{|H|}$, the standard deviation $\sigma_{|H|}$, the 90th percentile $P90_{|H|}$, the mean μ_K , and the Shannon entropies $\mathbb{H}(H)$ and $\mathbb{H}(K)$. We tested additional statistics but removed those that correlated less with perceived complexity.

View similarity. A major factor affecting perceptual object complexity is perceived symmetry [6]. A notable difference between true symmetry and its subjective interpretation is that our method should assign low scores to both symmetric and approximately symmetric objects; in other words, symmetry should be treated as a continuous score rather than a binary property.

To quantify symmetry, we estimate similarity across the object's views. Similar to Saleem et al. [15], we capture 42 grayscale views of each object and derive a view-similarity matrix, which we then project to two dimensions using Multi-Dimensional Scaling. Because sparsely distributed points in the projected similarity space correspond to dissimilar views, we use the convex hull area of these points as the similarity metric. Instead of computing the similarity matrix with Contour-to-Centroid Triangulation on the view boundary contours, as in the baseline method, we use a pretrained vision model [26] and compute pairwise cosine similarity across all views.

The view similarity metric is computed both globally for every object, and for each segmented part, in order to take into account both the quantitative and structural dimensions of the object complexity.

4.3. CORE: Complexity via reconstruction via guided point cloud reconstruction (Christoforos Vlachos, Anthony Belessis, Ioannis Romanelis, Vlassis Fotis and Konstantinos Moustakas)

Method name: CORE

We propose a generative method to estimate 3D shape complexity. Our method builds on the hypothesis—previously explored in the 2D context [27]—that perceived complexity is directly linked to geometric complexity, i.e., visually complex shapes tend to exhibit similarly intricate 3D geometry. A generative model trained on a collection of 3D shapes is expected to reconstruct shapes near the center of the training distribution with higher fidelity than outliers that deviate from the learned prior. Additionally, constraining the capacity and training budget of the generative model limits its ability to encode high-frequency geometric details, which are typically more prevalent in complex shapes. Following this hypothesis, we link the complexity of a 3D object to the reconstruction loss of our generative model.

As our generative backbone we select SPVD [28], due to its high efficiency and low computational requirements. Because SPVD operates on point clouds, we sample 2048 points on the surface of each mesh using farthest point sampling. As our task is reconstruction, we extend




Fig. 5. The DtH method: a satellite sphere with Fibonacci observation directions. The normalized shape S lies within a unit sphere. The satellite sphere at radius $R = 1 + \epsilon$ surrounds it and carries N_H quasi-uniform Fibonacci directions h_i . The observer at each point Rh_i gazes inward along $-h_i$ to the center of S ; the scalar field value $F(h_i)$ is computed via Eq. (1).

the model with point cloud-based conditional guidance following the paradigm of G-SPVD [29], replacing the image-based vision transformer with a PointNet [30] encoder. The lightweight PointNet encoder further reinforces higher reconstruction loss on complex objects, as it cannot efficiently encode salient geometric details. Finally, in place of the diffusion pipeline used in the original SPVD, we adopt a rectified flow formulation [31]. To measure reconstruction loss between the input and the generated shape we use Chamfer Distance.

Since this method is fully unsupervised and requires no human annotation, it can estimate complexity in any sufficiently large set of 3D shapes. To demonstrate its effectiveness, we train a variant exclusively on the ABC subset, and one on the primate tooth dataset, without any pretraining on larger datasets.

4.4. Distance-to-heaven (DtH) field with dual-scale complexity metrics (Igor Volkov and Sergei Krasovskii)

Method name: Volkov

We propose a method based on the hypothesis that perceived complexity reflects how surface details remain visible across multiple viewing scales. Each input shape S is represented as an oriented triangulated surface mesh, where each triangle T_i carries an outward unit normal $\mathbf{n}(T_i)$ and an area $\text{area}(T_i)$. As a pre-processing step, each mesh is translated so that its area-weighted centroid lies at the origin and scaled so the surface fits within the unit sphere. This defines a canonical coordinate frame that is unique up to rotation. Meshes with more than 20,000 faces are decimated to 20,000 faces using quadratic error simplification [32] before normalization.

Satellite sphere. We introduce the *satellite sphere* of radius $R = 1 + \epsilon$, centered at the origin. Each point $R\mathbf{h}$ on the satellite sphere serves as a vantage point; the observer gazes inward along $-\mathbf{h}$, and the visibility gate $\langle \mathbf{n}(T_i), \mathbf{h} \rangle_+$ in Eq. (1) selects triangles whose outward normal faces toward the observer, analogous to an orbiting satellite scanning a planet from a fixed altitude ϵ . An illustration of the satellite sphere with the Fibonacci observation directions is shown in Fig. 5.

The distance-to-heaven (DtH) field. The DtH scalar field $F(\mathbf{h})$ is an area-weighted, visibility-gated, kernel-smoothed projection of the normalized surface onto the satellite sphere:

$$F(\mathbf{h}; R, \sigma) = \sum_{T_i} \text{area}(T_i) \cdot \langle \mathbf{n}(T_i), \mathbf{h} \rangle_+ \cdot K(\|\mathbf{b}(T_i) - R\mathbf{h}\|; \sigma) \quad (1)$$

where $\langle \cdot, \cdot \rangle_+ = \max(0, \langle \cdot, \cdot \rangle)$ denotes the positive part and $\mathbf{b}(T_i)$ is the barycenter of triangle T_i . The three multiplicative factors serve distinct roles: $\text{area}(T_i)$ ensures the sum approximates the surface area integral $\int_S \dots dA$, making $F(\mathbf{h})$ independent of mesh tessellation density. $\langle \mathbf{n}(T_i), \mathbf{h} \rangle_+$ equals zero for back-facing triangles, approximately modeling a *visibility gate*: only surface area that an observer at $R\mathbf{h}$ could see contributes to $F(\mathbf{h})$. $K(\|\mathbf{b}(T_i) - R\mathbf{h}\|; \sigma)$ is a spatial smoothing kernel centered at $R\mathbf{h}$, weighting each triangle's contribution by its proximity to the observer.

Spatial scale-difference metric. Our first descriptor D_1 uses the Wendland C^2 kernel (K_W) [33], a compactly supported positive-definite radial basis function. From (1), we use K_W in place of K , which we shall refer to as F_W .

F_W is calculated over two scales: $R_{\text{fine}} = 1.01$ and $R_{\text{coarse}} = 1.05$. We denote the standard deviation of $F_W(h)$ over all $N_H = 1000$ observation directions at radius R as $\text{std}(F_W | R)$. The D_1 metric is:

$$D_1 = \text{std}(F_W | R_{\text{fine}}) - \text{std}(F_W | R_{\text{coarse}}).$$

At the fine scale, the kernel support cone is narrow, making std sensitive to localized features. At the coarse scale, the cone is wider and suppresses fine details. D_1 as a difference of std at fine and coarse scales is therefore a *scale-selective detector of fine geometric detail*, i.e., positive for shapes with fine-scale features, and near zero for geometrically homogeneous shapes (e.g., a sphere).

When $D_1 > 0$, it may be physically interpreted as indicating fine surface details (e.g., thin protrusions, sharp edges, high-frequency textures, etc.). For smooth shapes such as a sphere or box, $\text{std}(\cdot)$ changes little between scales and $D_1 \approx 0$.

Spectral scale-ratio metric. Our second descriptor D_2 examines the *spherical harmonics* (SH) coefficients [34] at a fine and coarse scale. A Gaussian kernel K_G is used in place of K in F from (1). For the fine scale $\sigma = 0.30$ while for the coarse scale $\sigma = 1.20$.

At each scale, the SH power spectrum $\{P_\ell\}$ is computed. Three summary statistics are derived from the high-frequency bands ($\ell_0 = 5, \dots, L_{\text{max}}$ where $L_{\text{max}} = 10$):

$$\underbrace{hf_{\text{ratio}} = \frac{\sum_{\ell=\ell_0}^{L_{\text{max}}} P_\ell}{\sum_{\ell=0}^{L_{\text{max}}} P_\ell}}_{\text{Fraction of spectral energy}} \quad \underbrace{\mu_{\text{norm}} = \frac{\sum_{\ell=\ell_0}^{L_{\text{max}}} P_\ell}{L_{\text{max}} \sum_{\ell=\ell_0}^{L_{\text{max}}} P_\ell}}_{\text{Energy-weighted mean } \mu_{\text{norm}} \in [0,1]}$$

$$\underbrace{H_{\text{norm}} = \frac{-\sum_{\ell=\ell_0}^{L_{\text{max}}} P_\ell \ln P_\ell}{\ln(L_{\text{max}} - \ell_0 + 1)}, \quad P_\ell = \frac{P_\ell}{\sum_{\ell=\ell_0}^{L_{\text{max}}} P_\ell}}_{\text{Normalized spectral entropy}}$$

The three statistics capture complementary properties of the high-frequency spectral distribution. hf_{ratio} answers *how much?*—the fraction of total spectral energy in the high-frequency bands. μ_{norm} answers *where?*—where within those bands the energy is concentrated (low value: energy sits at the lowest high-frequency bands; high value: energy reaches the finest spatial scales near $\ell = L_{\text{max}}$). H_{norm} answers *how spread?*—whether the energy is concentrated in one band or distributed across many, indicating a multi-scale texture.

The composite spectral complexity index M_6 at scale ϵ is structured as amount $\times$ quality:

$$M_6(\epsilon) = hf_{\text{ratio}} \cdot \sqrt{\mu_{\text{norm}} \cdot H_{\text{norm}}},$$

where hf_{ratio} acts as a gate. If there is no high-frequency energy, the distributional statistics μ_{norm} and H_{norm} are meaningless, and M_6 collapses to zero. The factor $\sqrt{\mu_{\text{norm}} \cdot H_{\text{norm}}}$ is the geometric mean of the two distributional properties. A surface with all energy concentrated in a single band has $H_{\text{norm}} = 0$, which drives the mean to zero regardless of μ_{norm} .

The resulting descriptor is the combination of the log ratios at the two scales:

$$D_2 = \frac{\log_{10} M_6(0.20) - \log_{10} M_6(0.80)}{\log_{10}(0.80) - \log_{10}(0.20)}.$$

D_2 is structurally analogous to a two-scale box-counting slope, measuring how the high-frequency spectral content of $F(\mathbf{h})$ changes between a finer and a coarser observation scale. Surfaces with richer angular complexity obtain a higher M_6 even at the coarser scale, resulting in a larger D_2 .

Ranking. Both D_1 and D_2 are converted to per-dataset ordinal ranks. The Borda score [35] for shape S_j is:

$$\text{Borda}(S_j) = \text{rank}_{D_1}(S_j) + \text{rank}_{D_2}(S_j).$$

Final ranking is by ascending Borda score. Ties are broken by D_1 rank.

4.5. Object complexity measured from different viewpoints (Simona Cutruzzola, Michela Mortara and Michela Spagnuolo)

Method name: Cutruzzola

We evaluate the visual complexity of each mesh by performing measurements from different viewpoints. Initially, the scale of each mesh is normalized with respect to its bounding box diagonal, and translated to place its bounding box center at the origin $O = (0, 0, 0)$. This allows the method to be scale and translation invariance.

Curvature estimation. We compute the principal curvatures k_1, k_2 using Meyer et al. [36]. To reduce noise that would alter the results, we perform 3 iterations of a modified Laplacian smoothing on the curvature values.

The value of k_j (for $k \in \{1, 2\}$) at vertex v after iteration $(i + 1)$ is given by:

$$k_{j,v}^{(i+1)} = \frac{k_{j,v}^{(i)} + \sum_{u \in N(v)} w_{uv} k_{j,u}^{(i)}}{1 + \sum_{u \in N(v)} w_{uv}},$$

where $N(v)$ is the 1-ring neighborhood of v and the weights w_{uv} are defined as:

$$w_{uv} = \frac{\cot \alpha_{uv} + \cot \beta_{uv}}{2} \cdot e^{-\frac{(k_{j,u} - k_{j,v})^2}{2m^2}}.$$

Here, α_{uv} and β_{uv} are the angles opposite edge (u, v) in the adjacent triangles, and m is the median of $|k_{j,u} - k_{j,v}|$ computed over all neighboring vertices in the mesh. Weights are thus the product between the cotangent weights and an edge-preserving function, added to enhance smoothing across neighboring regions with similar curvature while preventing excessive rounding of edges and sharp features.

Subsequently, each vertex is classified by its shape index as either *planar*, *cup*, *trough*, *saddle*, *ridge*, or *dome* following the scheme of Koenderink and van Doorn [37].

Quantifying visual clutter. We sample 12 viewpoints at the vertices of a regular icosahedron, each at a distance of 1.5 from O . For each view, we use ray tracing and region growing to select all visible mesh faces, defining a face as “visible” if its barycenter is visible.

The selected faces are extracted to form a mesh of the visible portion of the model. The vertices of this copy are then projected onto a unit sphere at the viewpoint’s position, yielding a realistic approximation of the perceived view.

Our goal is to quantify the visual clutter from every point of view. We do this by computing the entropy of the distribution of shape index classes within the neighborhood of each vertex on the spherical projection of the visible mesh. At the same time, we want to control the perspective-included distortion that might yield higher entropy in regions with skewed elements squeezed by the projection rather than in correspondence with truly complex structural details.

To this end, given the subject vertex v , and its neighbors $u_1, \dots, u_n$ on the spherical projection within a fixed angular distance α (for our




Fig. 6. Vertex v over the unit sphere centered at the observer’s position with its neighbors u_1 and u_2 , their counterparts in the three-dimensional environment v' , u'_1 and u'_2 , as well as the projections u''_1 , u''_2 onto the sphere intersecting v' .

experiments $\alpha = 2^\circ$), found using a k -d tree. The first part of the computation of local visual clutter in v involves Shannon’s entropy formula:

$$H_v^{(\alpha)} = \frac{1}{\log_2 6} \sum_{c=1}^6 -p_c \log_2 p_c,$$

where p_c is the proportion of neighboring vertices (including v) belonging to class c . This value is then scaled by a term $R_v^{(\alpha)}$, which accounts for perspective distortion. Specifically, let $v', u'_1, \dots, u'_n$ be the counterparts of $v, u_1, \dots, u_n$ in the three-dimensional setting, and let $u''_1, \dots, u''_n$ be the projections of $u'_1, \dots, u'_n$ onto the sphere centered at the observer’s position and intersecting v' (see Fig. 6). $R_v^{(\alpha)}$ is then defined as:

$$R_v^{(\alpha)} = \frac{1}{n} \cdot \sum_{i=1}^n \frac{\text{dist}(v', u'_i)}{\text{dist}(v', u_i)} = \text{dist}(Ob, v') \cdot \frac{1}{n} \sum_{i=1}^n \frac{\text{dist}(v, u_i)}{\text{dist}(v', u_i)}.$$

This term compares the distances in the 3D domain with those measured on the sphere centered at the observer and passing through v' .

Finally, the local visual entropy values $E_v^{(\alpha)} = H_v^{(\alpha)} \cdot R_v^{(\alpha)}$ are averaged over all visible vertices to obtain the overall scene entropy $E^{(\alpha)}$. This procedure is repeated for each viewpoint, and the 12 values of $E^{(\alpha)}$ are averaged together to obtain the complexity index of the entire mesh.

4.6. Topology-based complexity scoring (Marco Guerra, Ulderico Fugacci, Silvia Biasotti)

Method name: Guerra

Shape descriptors. We adopt an array of topological and geometrical descriptors to describe properties of the mesh associated with complexity. Several basic features are used such as the number of vertices, number of faces, circumradius, and total surface area. We also compute the Betti numbers b_0 , b_1 and b_2 , corresponding to the number of connected components, the number of tunnels, and the number of cavities respectively. Next, we compute the persistence diagrams for the three axis-aligned filtrations in dimension 0 and 1, and measure the total persistence of bars whose lengths lie between two thresholds defined as fractions of the circumradius. The same descriptor was computed for the Alpha complex [38] filtration.

Finally, to account for a more geometrical description we computed a multi-scale curvature descriptor by sampling points uniformly inside the mesh. For each point, a ball of radius r is constructed at the point. The intersection of the mesh with the ball is considered. We computed face normal vectors and their covariance matrix. The ratio of the smallest eigenvalue to the sum of eigenvalues is recorded. This procedure is repeated at three different scales for three different values of r .

Data preparation. To obtain a small ground-truth set to validate our descriptors, we used a two-way ranking procedure based on the Elo algorithm [39]. For the ABC and fractal datasets, we used shapes selected for training. For the primate dataset, we derived rankings from a subset of 25 test shapes due to limited additional data. For each dataset, three individuals performed repeated pairwise comparisons on shapes. The comparison histories were then interleaved to obtain a combined Elo based on all three individuals.

Scoring and ranking. The selected shape descriptors were then used as input to an ℓ_1 -regularized linear model, using the ground truth subset to tune parameters. Due to the aggressive regularization, many coefficients are pushed to zero, highlighting features of particular importance for the complexity classification. With the linear function thus approximated, we have then computed a score for each mesh, and constructed the final ranking.

4.7. Qwen & Gemini

Method name: Qwen, Gemini

Large *Vision Language Models* (VLMs) have become pervasive for certain tasks in recent years. They perform well in zero-shot settings, but their utility for complexity estimation is unknown. For this task we used the Qwen3 VL 235B Thinking (open-weight) model and a recent Google Gemini (closed-weight) model. Because these datasets require many queries, we used the more economical Google Gemini 3.1 Flash Lite Preview variant.

Data preparation. VLMs are well suited to image data, we therefore render each mesh from four viewpoints at the corners of a tetrahedron. Images were rendered at 256×256 pixels an isometric camera view and using `vtkLightKit` lighting. Initial experiments showed that the number of views and image resolution were sufficient.

Simulation analysis. We conduct a Monte Carlo simulation to estimate how many pairwise comparisons Rank Centrality requires for reliable global ranking recovery. For a dataset of n items, we generate a ground-truth ranking, randomly sample m pairwise comparisons with replacement, and simulate outcomes with a fixed upset probability $p = 0.1$. The recovered ranking is evaluated using Spearman correlation ρ . For the ABC dataset ($n = 500$), performance increases from $\rho = 0.23$ at $m = 100$ to $\rho = 0.87$ at $m = 10,000$, reaching $\rho = 0.91$ at $m = 20,000$ and saturating around $\rho = 0.95$ at $m = 50,000$. For the Fractal dataset ($n = 24$), ρ improves from 0.35 at $m = 10$ to 0.73 at $m = 100$, exceeding 0.90 at $m = 800$ and reaching 0.95 at $m = 2000$.

Method. Since asking a VLM to directly attribute a complexity score is unlikely to result in a consistent score, we therefore choose to perform a series of two-way comparisons, as this is more likely to lead to a more reliable output. We adopt the Rank Centrality method [40] and obtain $n \lceil \log_2 n \rceil$ samples (where n is the number of shapes in a dataset) as suggested. For each shape, $k = \lceil \log_2 n \rceil$ other shapes were randomly selected.

The VLMs were provided with the prompt below followed by the encoded images interleaved with labels indicating which images corresponded to shape A and shape B:

You are given two 3D shapes (Shape A and Shape B)
represented by multiple rendered views.

Task:
Determine which shape is more visually complex.

Answer with either 'A' or 'B'.

For each dataset of each VLM, Rank Centrality was applied to the responses to obtain an ordering.

4.8. Alpha-shape complexity (Charlotte Brassey and James Gardiner)

Method name: Gardiner et al. [41]

Any models that were not watertight were preprocessed prior to analysis, including hole filling and the fixing of non-manifold edges performed in Autodesk MeshMixer. If more extensive repairing was required, 3D alpha wrapping was applied using CGAL [42].

Once watertight meshes were obtained, model volume and a reference length (used as a metric of scale) were calculated as per Gardiner et al. [41]. Point clouds were then generated for each model by sampling a spatially random distribution of 100,000 points within the mesh volume.

Following generation of point clouds, alpha-shape complexity analysis, which was originally developed for analyzing biological datasets lacking homologous landmarks, was applied. For a more detailed description of the methodology, see Gardiner et al. [41].

Shape complexity was quantified as the 'tightness' of alpha shape with a volume equal to that of the original watertight mesh, with more complex geometries requiring tighter fits to achieve equivalence. Alpha-shapes were fitted using the `alphavol` function [Jonas Lundgren, [MATLAB File Exchange](#)], and all analyses were performed in MATLAB R2020b.

4.9. Nearest neighbor and topological complexity descriptor (Ivan Sipiran and Isaac Aguirre)

Method names: Sipiran NN and Sipiran TC

To capture the spatial distribution of features over the surface, we introduce two additional complexity measures that operate on the per-vertex feature representation. These feature vectors are obtained via the enhanced back-projection pipeline of Aguirre and Sipiran [43]: the object is rendered from multiple viewpoints sampled on a Fibonacci sphere, and patch-level descriptors extracted by DINOv2 [26] are projected back onto each visible vertex and averaged across views. To reduce computational cost, the resulting 384-dimensional per-vertex descriptors are compressed to 3 dimensions via PCA prior to complexity estimation.

The first measure estimates the *nearest-neighbor dispersion* of the feature distribution. For each vertex, we compute the Euclidean distance to its nearest neighbor in feature space, and report the mean of these distances over all vertices as the complexity score. The rationale is that geometrically simple objects (such as planes or spheres) produce highly concentrated feature distributions in which neighboring points in feature space are densely packed, whereas complex objects with varied surface geometry yield more dispersed distributions where even the closest feature-space neighbors remain relatively distant from one another.

The second measure quantifies the *topological complexity* of the feature signal as a scalar function defined over the 3D surface. For each of the d dimensions of the PCA-reduced feature representation, we treat the feature values as a scalar field over the point set and count the number of local maxima (vertices whose feature value strictly exceeds that of all their $k = 10$ nearest geometric neighbors). This count is normalized by the total number of vertices and averaged across all d dimensions to yield a single score. The underlying intuition is that a geometrically simple object produces a smooth and slowly varying feature field with few local extrema, while a complex object generates a rugged feature landscape with many isolated peaks.

4.10. 2D multi-view based shape convexity and complexity (Paul Rosin)

In SHREC 2026 we have used the identical set of 2D boundary-based methods that were used in SHREC 2021 [11]. Therefore, only a brief outline is provided here, and a fuller description can be found in [11]. With the exception of the last measure (C_{PC}) all measures are applied




Fig. 7. Top row: silhouettes of 2D views. Middle row: silhouettes with cuts to each hole. Note that for clearer visualization the cuts have been drawn thicker than normal. Bottom row: extracted contours with arclength mapped to hue.

independently to each 2D view, and the combined result is the average of the measure across the views.

These methods can be split into convexity or fractal dimension measures (indirectly measuring complexity) and direct complexity measures. The convexity measures consist of: Rosin's convexification [44], Žunić and Rosin's measure based on the perimeter of the minimum area bounding rectangle, and two standard convexity measures based on the area and perimeter of the convex hull denoted as C_1 and C_2 . The fractal measures consist of: Vasselle and Giraudon's k-regularity [45] based on a ratio between distances between points on the curve at different scales, Hayward et al.'s slope of the regression line for log versus log plots of ruler length versus perimeter [46], and Borkowski's [47] box counting method. The direct complexity measures are Brinkhoff et al.'s measure [48] that is a linear combination of the number of notches (non-convex vertices) normalized to be in the range [0,1], and two terms similar to C_1 and C_2 , Page et al.'s measure [49] which uses the entropy of boundary turning angles, also a version denoted C_{CRE} using cumulative residual entropy [50], The measure of Matsumoto et al. [51] using the sum of absolute Gaussian curvature, and a version denoted C_σ using the standard deviation of signed curvature. Finally, the multi-view method, denoted as C_{PC} , is our simplified version of Saleem et al. [15]. We perform a pairwise comparison of the boundaries across all the views for a given model, and return the mean score across all comparisons.

The 2D boundary-based methods in SHREC 2021 [11] were applied to the largest area single contour extracted from each binary image. This approach was generally sufficient for the three datasets used in SHREC 2021. However, many of the shapes in the three datasets used in SHREC 2026 contain holes, and ignoring these holes would neglect an important aspect of complexity. In particular the ABC dataset includes many shapes with one or multiple holes, and these can be large or small in area. The fractal dataset also has shapes with holes, but these holes tend to be small. Finally, the number and size of holes in the tooth dataset is negligible. Therefore a different approach was taken for SHREC 2026, which (for shapes containing holes) essentially consisted of connecting the holes to each other and/or to the background in order to split the object in such a way that the outer contour would now capture all the internal details of such shapes.

Examples of the results of this process are shown in Fig. 7. After cutting the initial 2D silhouettes (top row) the holes have all been removed, and so the new silhouettes (middle row) have a single (exterior) boundary. This boundary can be seen more clearly when it is plotted with arclength mapped to hue (bottom row).

The cutting process for a silhouette starts by extracting the hole regions and thinning them. This set of points C forms part of the cutting curve, and they are all interior to the shape. A further point on the exterior boundary is added to C in order that the interior holes can



Fig. 8. Left: cutting with full MST; middle: cutting with thresholded MST; right: second iteration of cutting with thresholded MST. Note that for clearer visualization the cuts have been drawn thicker than normal.

be removed by connecting them to the exterior. This exterior point is chosen as the point on the exterior boundary that is closest (i.e. has the smallest Euclidean distance d_{min}) to C . The minimal spanning tree (MST) of C provides a cut connecting the holes to the exterior, as shown in Fig. 8 (left).

However, it would be preferable to minimize the length of the cutting curve, since the cutting curve itself is an added artefact and not part of the original data. In this example, it can be seen that there is a long MST edge connecting the large holes at the top of the shape to the small hole at the bottom. Therefore, a modified iterative approach is employed. To prevent excessively long MST edges, all edges longer than $d_{min} + \epsilon$ are deleted. If any such edges exist, and are removed, the cutting consequently does not remove all the holes, as shown in Fig. 8 (middle). Further iterations are applied to the partially cut silhouette until the result contains a single component. At each iteration C contains the thinned points of the remaining holes and the new exterior point, and d_{min} is updated. For the example in Fig. 8 two iterations are sufficient.

4.11. Other methods

We also include results from some other methods in the literature. We briefly describe them here and refer the reader to the corresponding papers for details.

Genctav et al. [52] and Genctav and Tari [53] are PDE-based shape descriptors that transform geometry into entropy-summarized scalar fields, emphasizing boundary sensitivity and deviation from roundness respectively. Both are further described in [11].

Reichert et al. [54] employs fractal dimensionality as a proxy for shape complexity. The experiments were run using the available binaries [54]. To generalize the method, a small modification was made to the protocol where the fractal dimensionality was estimated by the slope of the line of best fit, which is the common practice.

Machado and Cardoso [55] estimate an image's complexity by measuring its compressibility. They encode the image using JPEG compression at a fixed quality factor quality; we used 50. For more complex images, compression will incur a high error, and also yield a low file size compression ratio. Therefore the ratio of these two terms is an estimate of the complexity of the original image.

Rothgänger et al. [27] propose a self-supervised approach employing two variational autoencoders with independent parameters and latent dimensions of 16 and 64, compressing the input at two levels. This can be interpreted as a contemporary version of prior methods that approximate Kolmogorov complexity [e.g., 56]. To extend the method to 3D shapes, we compute the mean complexity score across a set of 2D silhouettes corresponding to each shape. Since the method is self-supervised and thus does not require labels, we train the method for 200 epochs on each respective dataset.

5. Evaluation metrics

To quantify the performance of the aforementioned methods, we compare the predicted rankings to the ground-truth rankings for each dataset individually. We adopt several metrics that quantify different aspects of the performance of these methods. The most fundamental

statistics are the rank correlation coefficients obtained using Kendall's τ and Spearman's ρ .

Given a set of shapes S where the number of shapes is $n = |S|$. Let $r^* : S \mapsto \{1, \dots, n\}$ and $r : S \mapsto \{1, \dots, n\}$ be the ground-truth and predicted ranking respectively. The Kendall distance may be defined as

$$d_K = \# \{ (i, j) \in S \times S : i < j, (r^*(i) < r^*(j)) \neq (r(i) < r(j)) \}.$$

Kendall's correlation statistic can then be calculated as follows:

$$\tau = 1 - \frac{4d_K}{n(n-1)}.$$

Kendall's τ coefficient has the advantage of being insensitive to the magnitude of individual errors and exhibits better robustness to datasets that are not distributed uniformly.

Unlike Kendall's τ , the Spearman's rank correlation coefficient takes into account the distance between the predicted and ground-truth ranking. We define the sum of squared rank differences as

$$d_S = \sum_{i \in S} (r(i) - r^*(i))^2.$$

In the case of tied ranks, we measure the distance using the mid-rank. Spearman's rank correlation coefficient is then given by

$$\rho = 1 - \frac{6d_S}{n(n^2 - 1)}.$$

To illustrate the differences in performance we also report a cumulative mass function that is based on the absolute distance,

$$\text{CMF}(k) = \frac{1}{n} \# \{ i : |r(i) - r^*(i)| < k \}, \quad \forall k = 0, \dots, n$$

This can also be summarized quantitatively by the area under the curve, *i.e.*,

$$\text{AUC} = \sum_{k=0}^{n-1} \text{CMF}(k),$$

which can be normalized by $1/n$. Unlike precision and recall at k , this formulation does not bias items near the start of the ranking. The resulting statistic may be interpreted similarly to other related statistics such as the cumulative mass characteristic, where a larger AUC score is better.

6. Results

Quantitative results using the aforementioned metrics are reported for each dataset presented in Table 1. The results have been arranged to help distinguish between the approaches that operate in 2D and those that operate in 3D directly. Illustrations of the shape ordering for each method are shown in Figs. 9(a), 10(a) & 11(a). Cumulative error curves that complement the tables are given in Figs. 9(b), 10(b) & 11(b). To probe for systematic directional biases, we examined the shapes with the greatest per-shape mean ranking error and assessed whether methods within each category consistently over- or underestimated complexity. The results of individual methods have been ordered consistently from simple to complex; ideally, the correlations should be positive.

6.1. ABC dataset

Of the learning-based methods, Gemini obtained the best result. It was observed—on occasion—that the VLM would verbally describe the scene such as decomposing the shape into parts (*e.g.*, ‘a cylinder and an arc’) and textures (*i.e.*, ‘bumpy’ vs. ‘smooth’) when weighing its decision. The heuristic method, SCI-FIVE, performs competitively. Unfortunately, the method performed poorly on the other datasets.

The top heuristic method was C_{CRE} , which is a variation of Page et al. [49]—whose original formulation performs notably worse. C_{CRE} replaces Shannon entropy with the cumulative residual entropy, which

offers more stable estimations by avoiding explicitly estimating the probability density function.

Examining the results in Fig. 9(a), we see that while the top methods broadly obtain the correct order, there are many outliers where a very simple or very complex shape is ranked near the opposite extreme. This manifests again in the cumulative error plot (Fig. 9(b)) where we see that the scores converge slowly to 100%. Among the most erroneous shapes, 3D methods show a tendency to overestimate complexity, whereas 2D methods exhibit no strong directional bias. This systematic overestimation indicates a sensitivity in 3D methods to geometric features that are uncorrelated with human complexity judgments.

We additionally hypothesize that self-occluding geometry may disadvantage 3D methods: since human raters could not inspect shape interiors, 3D methods risk attending to signals that were imperceptible to annotators. Pruning shapes with self-occlusions yields marginal improvements for most 3D methods, supporting this hypothesis.

6.2. Primate tooth dataset

The best performing method on the dataset was Genctav and Tari [53]—a 3D-based method. Despite this, most 2D methods outperformed those operating directly in 3D. Silhouette-based methods perform surprisingly well despite occluding concave surface features, suggesting that prominent surface relief is a principal complexity signal. Fig. 10(a) illustrates that most 2D methods produce similarly ordered results, with rankings that are comparably erroneous and clustered in the same regions. Unlike on the ABC dataset, 3D methods here tended to underestimate complexity among the most erroneous shapes, while 2D methods showed no clear directional bias.

This dataset is particularly challenging as the shapes are reconstructed from an isosurface using marching cubes, which causes stair-stepping artifacts. A robust method must therefore be insensitive to such repeating high-frequency patterns. Relevantly, Arslan et al. [11] attributed degraded performance under increased perturbation to local high-frequency changes rather than the count of perturbations. Indeed, despite performing well on the ABC dataset, Qwen and Gemini failed to produce good results. It was observed that the VLM focused particularly on these artifacts of the reconstruction.

Of the 3D methods, Genctav and Tari [53] and Tatsuma achieve markedly better results than methods that have previously been investigated for estimating tooth complexity [41,54]. In the case of Tatsuma, the down-sampling of shapes to 30,000 points likely lessens the influence of stair-stepping artifacts. While Arslan et al. [11] note that down-sampling necessarily discards shape information, we find this can, in some cases, reduce sensitivity to artifacts. The use of entropy on the computed descriptors then quantifies the surface variability or regularity of features.

6.3. Fractal dataset

This dataset proved to be the most challenging for the 3D methods, with half of the participating methods being uncorrelated ranking. Consistent with the self-occlusion hypothesis, removing such shapes yields more pronounced improvements for most 3D methods than observed on the ABC dataset. Of the 3D methods, Gardiner et al. [41] performed the best. Broadly, the ranking of the fractal dataset is dominated by the degree of concavity of the shapes (see Fig. 3). The method's use of α -shape volume ‘tightness’ captures this property particularly well.

In comparison with the 3D methods, many of the 2D methods perform significantly better, with Machado and Cardoso [55] performing the best. One possible explanation is that many of the fractal shapes exhibit repeated structure and partial symmetry, so a single-view projection may still capture much of the relevant visual organization. Of the top ten methods visualized in Fig. 11(a) the majority are 2D methods. We observe that there are few outlier rankings, and the greatest absolute error was 14 places, as can be seen in Fig. 11(b). Neither 2D nor 3D methods exhibit a consistent directional bias on this dataset.

Table 1

Quantitative results on the three datasets. For each dataset, the best performing method of the 3D and 2D methods, respectively, are highlighted in **bold** and the second best in **bold italics**. The overall best method for each dataset is highlighted in **bold red**.

Method	ABC			Primate tooth			Fractal		
	τ	ρ	AUC	τ	ρ	AUC	τ	ρ	AUC
3D methods									
Gardiner et al. [41]	0.3125	0.4519	0.7664	0.3441	0.4554	0.7652	0.6304	0.8035	0.8611
Sipiran TC	0.0914	0.1336	0.6946	0.4750	0.6198	0.8022	0.4710	0.6330	0.8194
Tatsuma	0.2826	0.4128	0.7545	0.5521	0.7184	0.8300	0.1884	0.2426	0.7257
CORE ^a	0.1934	0.2872	0.7256	0.4573	0.6088	0.7932	–	–	–
Genctav and Tari [53]	0.0604	0.0875	0.6820	0.6442	0.8235	0.8605	0.0725	0.1078	0.6979
Reichert et al. [54]	0.1087	0.1650	0.6999	0.3648	0.4719	0.7525	0.3623	0.5061	0.7813
Guerra ^b	0.2588	0.3804	0.7451	0.3645	0.5116	0.7825	0.0435	0.0991	0.6875
Genctav et al. [52]	−0.1447	−0.2109	0.6299	−0.1871	−0.2444	0.6261	0.7029	0.8835	0.8924
Cutruzzola	0.1637	0.2369	0.7225	–	–	–	−0.0797	−0.1383	0.6562
Volkov	0.1029	0.1525	0.6967	0.0785	0.1026	0.6750	0.0725	0.1591	0.6840
Sipiran NN	0.0391	0.0581	0.6806	0.2535	0.3437	0.7331	−0.1449	−0.2183	0.6389
SCI-FIVE	0.3244	0.4735	0.7690	−0.2243	−0.2691	0.5940	−0.1739	−0.2313	0.6215
2D methods									
Gemini ^c	0.4522	0.6360	0.8085	−0.3452	−0.4682	0.5697	0.5145	0.7009	0.8160
Qwen ^c	0.3984	0.5661	0.7934	−0.0374	−0.0637	0.6335	0.4348	0.5843	0.8125
Machado and Cardoso [55]	0.1635	0.2427	0.7130	0.1057	0.1469	0.6943	0.7391	0.8852	0.8993
Rothgänger et al. [27] ^a	0.1336	0.1964	0.7083	0.5275	0.6943	0.8225	0.1594	0.2461	0.7083
C_{CRE}	0.3749	0.5347	0.7871	−0.1403	−0.2046	0.6456	0.4348	0.5817	0.7847
C_e	0.2511	0.3643	0.7444	0.5268	0.6947	0.8216	0.4203	0.5626	0.7778
Matsumoto et al. [51]	0.3278	0.4667	0.7709	0.5701	0.7315	0.8361	0.5290	0.7009	0.8194
Brinkhoff et al. [48]	0.1555	0.2288	0.7159	0.5632	0.7409	0.8266	0.6014	0.7643	0.8611
Page et al. [49]	0.2206	0.3237	0.7367	−0.0965	−0.1402	0.6452	0.4783	0.6235	0.8125
C_{PC}	0.2113	0.3094	0.7321	0.4680	0.6300	0.8051	0.2609	0.3704	0.7257
C_1	0.1217	0.1801	0.7044	0.5698	0.7586	0.8354	0.2971	0.4148	0.7535
Rosin [44]-1	0.1866	0.2776	0.7222	0.5237	0.7002	0.8188	0.5797	0.7522	0.8576
Rosin [44]-2	0.1906	0.2809	0.7230	0.5429	0.7227	0.8248	0.5797	0.7522	0.8576
Žunić and Rosin [57]	0.2131	0.3125	0.7296	0.5471	0.7393	0.8246	0.6522	0.8061	0.8750
C_2	0.2130	0.3134	0.7289	0.5529	0.7230	0.8234	0.6522	0.8000	0.8715
Borkowski [47]	0.1349	0.2016	0.7064	0.3851	0.5268	0.7805	0.5072	0.6765	0.8125
Hayward et al. [46]	0.2010	0.2970	0.7282	0.1173	0.1724	0.6939	0.2319	0.2635	0.7292
Vasselle and Giraudon [45]	0.2783	0.4055	0.7561	−0.1307	−0.1890	0.6429	0.1159	0.1861	0.7153
2D methods using a simple outer contour (SHREC2021)									
C_{CRE}	0.3080	0.4488	0.7640	−0.1445	−0.2105	0.6442	0.3188	0.3870	0.7500
C_e	0.2578	0.3759	0.7448	0.5264	0.6945	0.8201	0.0870	0.0817	0.6597
Matsumoto et al. [51]	0.3583	0.5130	0.7798	0.5682	0.7321	0.8347	0.3261	0.4435	0.7361
Brinkhoff et al. [48]	0.1328	0.1942	0.7085	0.5640	0.7451	0.8275	0.3986	0.5765	0.7951
Page et al. [49]	0.2704	0.3966	0.7513	−0.0969	−0.1418	0.6445	0.2029	0.2591	0.7326
C_{PC}	0.1977	0.2927	0.7279	0.4677	0.6297	0.8062	0.0362	−0.0183	0.6632
C_1	0.0965	0.1415	0.6959	0.5701	0.7592	0.8355	0.2174	0.3183	0.7222
Rosin [44]-1	0.1263	0.1874	0.7048	0.5222	0.7017	0.8188	0.3333	0.4635	0.7743
Rosin [44]-2	0.1338	0.1981	0.7077	0.5425	0.7235	0.8246	0.3551	0.4896	0.7778
Žunić and Rosin [57]	0.1981	0.2919	0.7308	0.5452	0.7394	0.8246	0.4348	0.6270	0.8090
C_2	0.2072	0.3057	0.7320	0.5525	0.7258	0.8237	0.4058	0.5800	0.8021
Borkowski [47]	0.0576	0.0830	0.6860	0.3894	0.5292	0.7814	0.4275	0.6070	0.8021
Hayward et al. [46]	0.1750	0.2550	0.7241	0.1169	0.1722	0.6939	0.0797	0.0930	0.7049
Vasselle and Giraudon [45]	0.2198	0.3240	0.7367	−0.1357	−0.1950	0.6422	0.0435	0.0983	0.7049

^a Self-supervised.

^b Primate tooth results were trained on independent annotations of the test set.

^c Large vision language model.




Fig. 9. Qualitative results on the ABC dataset of the top ten methods. (a) ranking consistency heatmap with Kendall's τ reported in parentheses, (b) CMF the AUC is reported in parentheses.




Fig. 10. Qualitative results on the tooth dataset of the top ten methods. Note that due to ties, the colors of the heatmap in (a) correspond to subsets of shapes that share the same ground truth rank. (b) illustrates the CMF and the AUC is reported in parentheses.




Fig. 11. Qualitative results of the top ten methods on the fractal dataset. (a) ranking consistency heatmap with Kendall's τ reported in parentheses, (b) CMF the AUC is reported in parentheses.

7. Conclusion

In this work, we curated three benchmarks for evaluating visual complexity estimation in 3D, enabling a comparative analysis of a range of methods. Across all three datasets, the strongest performing approaches were predominantly 2D-based—despite operating on projections rather than the full 3D geometry. We find that performance

was dataset-dependent and that no single algorithm dominated, consistent with Arslan et al. [11]; we further confirm on human-grounded benchmarks their finding that projection-based methods outperform native 3D methods. While 2D projections sacrifice geometric detail, 3D methods may attend to self-occluded regions that were invisible to human raters—an apparent advantage that becomes a source of noise against human-grounded ground truth. These results suggest that current computational approaches only partially capture human judgments

of visual complexity. A key limitation of this work is the reduction of complexity to a single scalar ranking, which may not fully reflect its multi-faceted nature. Overcoming this limitation will likely require a substantially larger and more nuanced survey. We leave this as future work. Finally, by making the ground-truth data publicly available, we hope this work will support future research on human-grounded visual complexity estimation, including the study of the domain-specific complexity of primate tooth morphology.

CRedit authorship contribution statement

Roberto M. Dyke: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Data curation, Conceptualization. **Yang Deng:** Writing – original draft, Validation, Resources, Project administration, Formal analysis, Data curation, Conceptualization. **Charlotte Brassey:** Writing – review & editing, Writing – original draft, Supervision, Resources, Methodology, Data curation, Conceptualization. **Yu-Kun Lai:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Conceptualization. **Paul L. Rosin:** Writing – review & editing, Writing – original draft, Supervision, Software, Methodology, Investigation, Data curation, Conceptualization. **James Gardiner:** Writing – original draft, Methodology. **Atsushi Tatsuma:** Writing – original draft, Methodology. **Christoforos Vlachos:** Writing – original draft, Methodology. **Anthony Belessis:** Writing – original draft, Methodology. **Ioannis Romanelis:** Writing – original draft, Methodology. **Vlassis Fotis:** Writing – original draft, Methodology. **Konstantinos Moustakas:** Writing – original draft, Methodology. **Ihar Volkau:** Writing – original draft, Methodology. **Sergei Krasovskii:** Writing – original draft, Methodology. **Marco Guerra:** Writing – original draft, Methodology. **Ulderico Fugacci:** Writing – original draft, Methodology. **Simona Cutruzzulà:** Writing – original draft, Methodology. **Michela Mortara:** Writing – original draft, Methodology. **Michela Spagnuolo:** Writing – original draft, Methodology. **Ivan Sipiran:** Writing – original draft, Methodology. **Isaac Aguirre:** Writing – original draft, Methodology.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Michela Spagnuolo is a symposium chair for 3DOR 2026. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made publicly available via a university repository. The evaluation code will also be made available online.

References

- Atneave F. Physical determinants of the judged complexity of shapes. *J Exp Psychol* 1957;53(4):221–7. <http://dx.doi.org/10.1037/h0043921>.
- Bazazian D, Magland B, Grimm C, Chambers E, Leonard K. Perceptually grounded quantification of 2D shape complexity. *Vis Comput* 2022;38(9):3351–63. <http://dx.doi.org/10.1007/s00371-022-02634-8>.
- Donderi DC. Visual complexity: a review. *Psychol Bull* 2006;132(1):73–97. <http://dx.doi.org/10.1037/0033-2909.132.1.73>.
- Snodgrass JG, Vanderwart M. A standardized set of 260 pictures: norms for name agreement, image agreement, familiarity, and visual complexity. *J Exp Psychol Hum Learn* 1980;6(2):174–215. <http://dx.doi.org/10.1037/0278-7393.6.2.174>.
- Pizlo Z, Sawada T, Li Y, Kropatsch WG, Steinman RM. New approach to the perception of 3D shape based on veridicality, complexity, symmetry and volume. *Vis Res* 2010;50(1):1–11. <http://dx.doi.org/10.1016/j.visres.2009.09.024>.
- Gartus A, Leder H. Predicting perceived visual complexity of abstract patterns using computational measures: The influence of mirror symmetry on complexity perception. *PLoS One* 2017;12(11):e0185276. <http://dx.doi.org/10.1371/journal.pone.0185276>.
- Ben Slama M, Chatti S, Louhichi B. Novel method for shape complexity evaluation: a threshold from machining to additive manufacturing in the early design phase. *Res Eng Des* 2024;35(2):191–214. <http://dx.doi.org/10.1007/s00163-023-00429-z>.
- Chang L-Y, Plaut DC, Perfetti CA. Visual complexity in orthographic learning: Modeling learning across writing system variations. *Sci Stud Read* 2016;20(1):64–85. <http://dx.doi.org/10.1080/1088438.2015.1104688>.
- Ferrari R, McKinnon D, He H, Smith RN, Corke P, González-Rivero M, et al. Quantifying multiscale habitat structural complexity: A cost-effective framework for underwater 3D modelling. *Remote Sens* 2016;8(2). <http://dx.doi.org/10.3390/rs8020113>.
- Pieters R, Wedel M, Batra R. The stopping power of advertising: Measures and effects of visual complexity. *J Mark* 2010;74(5):48–60. <http://dx.doi.org/10.1509/jmk.74.5.048>.
- Arslan JM, Haridis A, Rosin PL, Tari S, Brassey C, Gardiner JD, et al. SHREC'21: Quantifying shape complexity. *Comput Graph* 2022;102:144–53. <http://dx.doi.org/10.1016/j.cag.2021.09.005>.
- Winchester JM, Boyer DM, St. Clair EM, Gosselin-Ildari AD, Cooke SB, Ledogar JA. Dental topography of platyrrhines and prosimians: convergence and contrasts. *Am J Phys Anthr* 2014;153(1):29–44. <http://dx.doi.org/10.1002/ajpa.22398>.
- Sukumar SR, Page DL, Koschan AF, Abidi MA. Towards understanding what makes 3D objects appear simple or complex. In: 2008 IEEE computer society conference on computer vision and pattern recognition workshops. 2008, p. 1–8. <http://dx.doi.org/10.1109/CVPRW.2008.4562975>.
- Rossignac J. Shape complexity. *Vis Comput* 2005;12(01;21(12):985–96. <http://dx.doi.org/10.1007/s00371-005-0362-7>.
- Saleem W, Belyaev A, Wang D, Seidel H-P. On visual complexity of 3D shapes. *Comput Graph* 2011;35(3):580–5. <http://dx.doi.org/10.1016/j.cag.2011.03.006>.
- Koch S, Matveev A, Jiang Z, Williams F, Artemov A, Burnaev E, et al. ABC: A big CAD model dataset for geometric deep learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 2019, p. 9593–603. <http://dx.doi.org/10.1109/CVPR.2019.00983>.
- Forsythe A. Visual complexity: Is that all there is? In: Harris D, editor. Engineering psychology and cognitive ergonomics. Berlin, Heidelberg: Springer Berlin Heidelberg; 2009, p. 158–66. http://dx.doi.org/10.1007/978-3-642-02728-4_17.
- Ribes A, Boucheny C. Eye-dome lighting: a non-photorealistic shading technique. 2011, <https://www.kitware.com/eye-dome-lighting-a-non-photorealistic-shading-technique/>. [Accessed: 2025-11-06].
- Bradley RA, Terry ME. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika* 1952;39(3/4):324–45. <http://dx.doi.org/10.2307/2334029>.
- Wilman H, Belmaker J, Simpson J, De La Rosa C, Rivadeneira MM, Jetz W. EltonTraits 1.0: Species-level foraging attributes of the world's birds and mammals. *Ecological archives* E095-178. *Ecol* 2014;95(7):2027. <http://dx.doi.org/10.1890/13-1917.1>.
- Osada R, Funkhouser T, Chazelle B, Dobkin D. Shape distributions. *ACM Trans Graph* 2002;21(4):807–32. <http://dx.doi.org/10.1145/571647.571648>.
- Joe S, Kuo FY. Remark on algorithm 659: Implementing sobol's quasirandom sequence generator. *ACM Trans Math Software* 2003;29(1):49–57. <http://dx.doi.org/10.1145/641876.641879>.
- Tombari F, Salti S, Di Stefano L. Unique signatures of histograms for local surface description. In: Computer vision – ECCV 2010. ECCV '10, Berlin, Heidelberg: Springer-Verlag; 2010, p. 356–69. http://dx.doi.org/10.1007/978-3-642-15558-1_26.
- Rusu RB, Bradski G, Thibaux R, Hsu J. Fast 3D recognition and pose using the Viewpoint Feature Histogram. In: 2010 IEEE/RSJ international conference on intelligent robots and systems. 2010, p. 2155–62. <http://dx.doi.org/10.1109/IROS.2010.5651280>.
- Li Y, Liu S, Yang X, Guo J, Guo J, Guo Y. Surface and edge detection for primitive fitting of point clouds. In: ACM SIGGRAPH 2023 conference proceedings. SIGGRAPH '23, New York, NY, USA: Association for Computing Machinery; 2023, p. 1–10. <http://dx.doi.org/10.1145/3588432.3591522>.
- Oquab M, Darcet T, Moutakanni T, et al. DINOv2: Learning robust visual features without supervision. In: TMLR. 2024, p. 1–32.
- Rothgänger M, Melnik A, Ritter H. Shape complexity estimation using VAE. In: Arai K, editor. Intelligent systems and applications. Cham: Springer Nature Switzerland; 2024, p. 35–45. http://dx.doi.org/10.1007/978-3-031-47715-7_3.
- Romanelis I, Fotis V, Kalogeras A, Alexakos C, Munteanu A, Moustakas K. Efficient and scalable point cloud generation with sparse point-voxel diffusion models. *IEEE Trans Neural Netw Learn Syst* 2025;1–12. <http://dx.doi.org/10.1109/TNNLS.2025.3636409>.
- Romanelis I, Fotis V, Munteanu A, Moustakas K. G-SPVD: Image and sketch guided point cloud generation with sparse point-voxel diffusion models. In: 2025 international conference on visual communications and image processing. VCIP, 2025, p. 1–5. <http://dx.doi.org/10.1109/VCIP67698.2025.11396860>.
- Charles RQ, Su H, Kaichun M, Guibas LJ. PointNet: Deep learning on point sets for 3D classification and segmentation. In: 2017 IEEE conference on computer vision and pattern recognition. CVPR, 2017, p. 77–85. <http://dx.doi.org/10.1109/CVPR.2017.16>.

- [31] Liu X, Gong C, Liu Q. Flow straight and fast: Learning to generate and transfer data with rectified flow. 2022, arXiv preprint [arXiv:2209.03003](https://arxiv.org/abs/2209.03003).
- [32] Garland M, Heckbert PS. Surface simplification using quadric error metrics. In: Proceedings of the 24th annual conference on computer graphics and interactive techniques. SIGGRAPH '97, USA: ACM Press/Addison-Wesley Publishing Co.; 1997, p. 209–16. [http://dx.doi.org/10.1145/258734.258849](https://doi.org/10.1145/258734.258849).
- [33] Wendland H. Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Adv Comput Math* 1995;4(1):389–96. [http://dx.doi.org/10.1007/BF02123482](https://doi.org/10.1007/BF02123482).
- [34] Kazhdan M, Funkhouser T, Rusinkiewicz S. Rotation invariant spherical harmonic representation of 3D shape descriptors. In: Proceedings of the 2003 eurographics/ACM SIGGRAPH symposium on geometry processing. SGP '03, Goslar, DEU: Eurographics Association; 2003, p. 156–64. [http://dx.doi.org/10.2312/SGP/SGP03/156-165](https://doi.org/10.2312/SGP/SGP03/156-165).
- [35] de Borda J-C. Mémoire sur les élections au scrutin. *Hist de L'Académie R Des Sci* 1781;102:657–65.
- [36] Meyer M, Desbrun M, Schröder P, Barr AH. Discrete differential-geometry operators for triangulated 2-manifolds. In: Hege H-C, Polthier K, editors. Visualization and mathematics III. Berlin, Heidelberg: Springer Berlin Heidelberg; 2003, p. 35–57. [http://dx.doi.org/10.1007/978-3-662-05105-4_2](https://doi.org/10.1007/978-3-662-05105-4_2).
- [37] Koenderink JJ, van Doorn AJ. Surface shape and curvature scales. *Image Vis Comput* 1992;10(8):557–64. [http://dx.doi.org/10.1016/0262-8856\(92\)90076-F](https://doi.org/10.1016/0262-8856(92)90076-F).
- [38] Edelsbrunner H, Mücke EP. Three-dimensional alpha shapes. *ACM Trans Graph* 1994;13(1):43–72. [http://dx.doi.org/10.1145/174462.156635](https://doi.org/10.1145/174462.156635).
- [39] Elo AE. *The rating of chessplayers, past and present*. New York: Arco Pub.; 1978.
- [40] Negahban S, Oh S, Shah D. Rank centrality: Ranking from pairwise comparisons. *Oper Res* 2017;65(1):266–87. [http://dx.doi.org/10.1287/opre.2016.1534](https://doi.org/10.1287/opre.2016.1534).
- [41] Gardiner JD, Behnsen J, Brassey CA. Alpha shapes: determining 3D shape complexity across morphologically diverse structures. *BMC Evol Biol* 2018;18(1):184. [http://dx.doi.org/10.1186/s12862-018-1305-z](https://doi.org/10.1186/s12862-018-1305-z).
- [42] Alliez P, Cohen-Steiner D, Hemmer M, Portaneri C, Rouxel-Labbé M. 3D alpha wrapping. In: CGAL user and reference manual. 6.1.1 th ed. CGAL Editorial Board; 2026, p. 1, URL <https://doc.cgal.org/6.1.1/Manual/packages.html#PkgAlphaWrap3>.
- [43] Aguirre I, Sipiran I. Enhanced back-projection of vision features for 3D symmetry detection. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. WACV, 2026, p. 66–76. [http://dx.doi.org/10.1109/WACV61042.2026.00015](https://doi.org/10.1109/WACV61042.2026.00015).
- [44] Rosin PL. Classification of pathological shapes using convexity measures. *Pattern Recognit* 2009;30(5):570–8. [http://dx.doi.org/10.1016/j.patrec.2008.12.001](https://doi.org/10.1016/j.patrec.2008.12.001).
- [45] Vasselle B, Giraudon G. 2-d digital curve analysis: A regularity measure. In: International conference on computer vision. 1993, p. 556–61. [http://dx.doi.org/10.1109/ICCV.1993.378162](https://doi.org/10.1109/ICCV.1993.378162).
- [46] Hayward J, Orford JD, Whalley WB. Three implementations of fractal analysis of particle outlines. *Comput Geosci* 1989;15(2):199–207. [http://dx.doi.org/10.1016/0098-3004\(89\)90034-4](https://doi.org/10.1016/0098-3004(89)90034-4).
- [47] Borkowski W. Fractal dimension based features are useful descriptors of leaf complexity and shape. *Can J Res* 1999;29(9):1301–10. [http://dx.doi.org/10.1139/x99-112](https://doi.org/10.1139/x99-112).
- [48] Brinkhoff T, Kriegl H-P, Schneider R, Braun A. Measuring the complexity of polygonal objects. In: *International workshop on advances in GIS*. 109–117, 1995, p. 9.
- [49] Page DL, Koschan AF, Sukumar SR, Roui-Abidi B, Abidi MA. Shape analysis algorithm based on information theory. In: International conference on image processing. vol. 1, 2003, p. 229–32. [http://dx.doi.org/10.1109/ICIP.2003.1246940](https://doi.org/10.1109/ICIP.2003.1246940).
- [50] Wang F, Vemuri BC, Rao M, Chen Y. Cumulative residual entropy, a new measure of information & its application to image alignment. In: International conference on computer vision. vol. 2, 2003, p. 548–548. [http://dx.doi.org/10.1109/ICCV.2003.1238395](https://doi.org/10.1109/ICCV.2003.1238395).
- [51] Matsumoto T, Sato K, Matsuoka Y, Kato T. Quantification of “complexity” in curved surface shape using total absolute curvature. *Comput Graph* 2019;78:108–15. [http://dx.doi.org/10.1016/j.cag.2018.10.009](https://doi.org/10.1016/j.cag.2018.10.009).
- [52] Genctav M, Genctav A, Tari S. NonLocal via local—NonLinear via linear: A new part-coding distance field via screened Poisson equation. *J Math Imaging Vision* 2016;55(2):242–52. [http://dx.doi.org/10.1007/s10851-015-0614-8](https://doi.org/10.1007/s10851-015-0614-8).
- [53] Genctav A, Tari S. Discrepancy: Local/global shape characterization with a roundness bias. *J Math Imaging Vision* 2019;61(1):160–71. [http://dx.doi.org/10.1007/s10851-018-0851-8](https://doi.org/10.1007/s10851-018-0851-8).
- [54] Reichert J, Backes AR, Schubert P, Wilke T. The power of 3D fractal dimensions for comparative shape and structural complexity analyses of irregularly shaped organisms. *Methods Ecol Evol* 2017;8(12):1650–8. [http://dx.doi.org/10.1111/2041-210X.12829](https://doi.org/10.1111/2041-210X.12829).
- [55] Machado P, Cardoso A. Computing aesthetics. In: *Brazilian symposium on artificial intelligence*. 1998, p. 219–28.
- [56] Chen Y, Sundaram H. Estimating complexity of 2D shapes. In: 2005 IEEE 7th workshop on multimedia signal processing. 2005, p. 1–4. [http://dx.doi.org/10.1109/MMSP.2005.248668](https://doi.org/10.1109/MMSP.2005.248668).
- [57] Žunić J, Rosin PL. A new convexity measurement for polygons. *IEEE Trans Pattern Anal Mach Intell* 2004;26(7):923–34. [http://dx.doi.org/10.1109/TPAMI.2004.19](https://doi.org/10.1109/TPAMI.2004.19).