



InShaDe: Invariant Shape Descriptors for visual 2D and 3D cellular and nuclear shape analysis and classification

Khaled Al-Thelaya^a, Marco Agus^{a,*}, Nauman Ullah Gilal^a, Yin Yang^a, Giovanni Pintore^b, Enrico Gobbetti^b, Corrado Cali^c, Pierre J. Magistretti^d, William Mifsud¹, Jens Schneider^{a,*}

^aCollege of Science and Engineering, Hamad Bin Khalifa University, Qatar Foundation, Doha, Qatar

^bVisual Computing Group, Center for Advanced Studies, Research and Development in Sardinia (CRS4), Cagliari, Italy

^cDepartment of Neuroscience "Rita Levi Montalcini", Neuroscience Institute "Cavalieri Ottolenghi", University of Turin, 10043, Turin, Italy.

^dBiological and Environmental Sciences and Engineering Division, King Abdullah University of Science and Technology, Saudi Arabia.

^eSidra Medicine, Qatar Foundation, Doha, Qatar

ARTICLE INFO

Article history:

Received May 9, 2021

ABSTRACT

We present a shape processing framework for visual exploration of cellular nuclear envelopes extracted from microscopic images arising in histology and neuroscience. The framework is based on a novel shape descriptor of closed contours in 2D and 3D. In 2D, it relies on a geodesically uniform resampling of discrete curves to compute unsigned curvatures at vertices and edges based on discrete differential geometry. Our descriptor is, by design, invariant under translation, rotation, and parameterization. We achieve the latter invariance under parameterization shifts by using elliptic Fourier analysis on the resulting curvature vectors. Uniform scale-invariance is optional and is a result of scaling curvature features to z-scores. We further augment the proposed descriptor with feature coefficients obtained through sparse coding of the extracted cellular structures using K-sparse autoencoders. For the analysis of 3D shapes, we compute mean curvatures based on the Laplace-Beltrami operator on triangular meshes, followed by computing a spherical parameterization through mean curvature flow. Finally, we compute the Spherical Harmonics decomposition to obtain invariant energy coefficients. Our invariant descriptors provide an embedding into a fixed-dimensional feature space that can be used for various applications, e.g., as input features for deep and shallow learning techniques or as input for dimension reduction schemes to provide a visual reference for clustering shape collections. We demonstrate the capabilities of our framework in the context of visual analysis and unsupervised classification of 2D histology images and 3D nuclear envelopes extracted from serial section electron microscopy stacks.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

The last decades have witnessed the rapid improvement and proliferation of high-throughput digital acquisition technology. As a result, high-quality digital representations of real-world

scenes and objects have become commonplace in many application domains. In biology and medicine in particular, the rise of whole-slide scanners and the digitization of traditional, confocal, and electron microscopy has led to both fully digital analysis and the creation of large image databases [1]. Early uses of this technology mostly included tele-pathology, solicitation of second opinions, and education in research and clinical practice, as well as visual ultra-structural analysis in neu-

*Corresponding authors: magus@hbku.edu.qa, jeschneider@hbku.edu.qa.

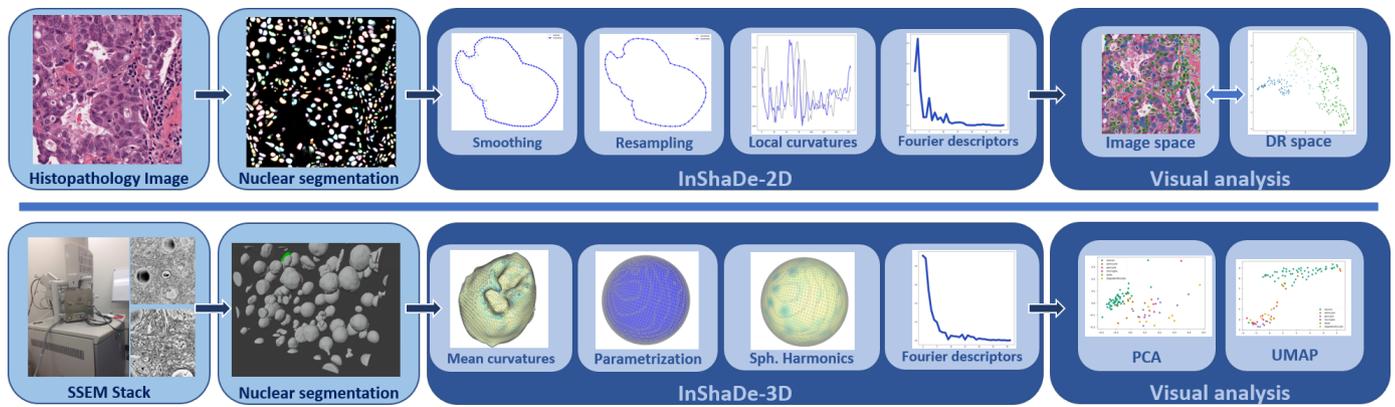


Fig. 1: **InShaDe pipeline**: from cell contours extracted from digital histology images (on top) or 3D shapes reconstructed from Serial Section Electron Microscopy stacks (on bottom), our pipeline computes invariant energy-based Fourier descriptors on top of discrete curvature embeddings. These synthetic descriptors can be used for visual analysis, proof-reading segmentation results, domain-specific clustering and classification according to specific taxonomies.

rosience. In most cases, digital workflows were designed to closely mimic the traditional investigation process. Only recent years saw a shift towards exploiting the large amount of information in the acquired images and collections by means of novel data-driven analysis methods [2]. In this context, a wide array of basic tools are employed, ranging from handcrafted feature descriptors over fully data-driven approaches to combinations of matching of various approaches [3]. Machine learning, and, especially, deep learning approaches have become popular in the context of digital pathology and biology. Their significant success stems from their ability to provide automatic tools for tasks such as segmentation and labeling of cellular entities from large individual microscope images [4, 5] or for supporting connectomics investigations by reconstructing the neural connections in large portions of brain tissue samples [6, 7]. One drawback of the current purely data-driven frameworks, however, is that they often disregard specific domain knowledge and taxonomies [8]. As a result, tools for domain-specific proofreading of segmented images, classification according to taxonomies, filtering, visual exploration, and subsequent computation are still lacking. For this reason, many applications require the use of descriptors that, by design, preserve some domain-specific characteristics. To better exploit the capabilities of novel learning frameworks we thus advocate not to dismiss designing features a priori. Instead, we believe that it is necessary to research the development of methods integrating such designed features into powerful descriptive models. One benefit of taking a hybrid design+data-driven approach is that domain knowledge can be integrated into the design process. This potentially leads to models that are easier to explain, and, thus, result in increased discrimination performance for human analysis. In addition, since features are designed, the training efforts in terms of the amount of data and computational power required can be eased.

Inspired by these considerations, we propose a novel visual analysis pipeline based on discrete differential geometry concepts, whose recent findings provide very powerful theoretical formulations for describing 2D and 3D shapes [9, 10]. The framework is based on a 2D/3D shape descriptor that can be used in various application domains to complement or enhance

generic deep learning networks, such as U-Net [11]. The descriptor, dubbed *InShaDe*, is based on the concept of discrete curvature along closed, resampled contours (see Fig. 1 for an overview of the proposed pipeline). In 2D, discrete curvature is computed using vertex and edge osculating circles. Interleaving the resulting edge and vertex curvatures produces a high-resolution curvature vector. These curvature vectors are naturally invariant under rigid body transformations (translations and rotations). Invariance under parametric shifts is ensured by using energy-based elliptic Fourier descriptors. Cellular shapes in 2D tissue samples are sliced, implying that their apparent size may be smaller than the real radius [12]. We therefore also propose to achieve optional invariance under uniform scaling by replacing components in the curvature vectors by standard (z -)scores. This manuscript is an extended version of the conference contribution [13] recently presented at the Eurographics Workshop on Visual Computing for Biology and Medicine (EG VCBM 2020) held in Tübingen. The conference paper presented:

- i) a robust geometry processing pipeline for computing 2D invariant shape descriptors exploiting shifted linear interpolation and discrete differential geometry schemes, and
- ii) visual mapping schemes from 2D cellular contours to shape descriptor embeddings based on modern dimension reduction schemes such as UMAP [14].

This manuscript extends the original pipeline by:

- iii) extending the shape descriptors to 3D shapes represented by triangle meshes, and
- iv) augmenting the shape descriptors by additional features based on sparse coding to improve analysis and classification performance.

In addition, we study the use of the proposed pipeline for proof-reading and visual, unsupervised classification of various histology images. We also present results for various 2D and 3D data sets stemming from histology and neuroscience.

2. Related Work

Our work is concerned with shape feature extraction from closed contours and surfaces and with the analysis of histopathology images and electron microscopy stacks. These are very broad topics and a full coverage of the state-of-the-art is beyond the scope of this paper. For a comprehensive overview of all related fields, we refer the reader to various surveys on 2D shape analysis [15, 3], digital histopathology analysis [16, 17, 18], and recent deep learning methods for cellular analysis [19].

In the following, we discuss the methods that are most closely related to our approach.

2.1. Shape feature descriptors

During the last two decades, significant research efforts have been carried out on both the theoretical and the practical aspects of the shape-based image retrieval problem [3]. For an overview of the seminal methods for shape-based invariant feature extraction for object recognition, we refer also to Yang et al. [20].

In general, there are two main modeling strategies for representing shapes: region-based methods and boundary-based ones. Region-based techniques use moment descriptors to describe shapes, like geometrical moments [21], Zernike moments [22, 23], Legendre moments [24], and Tchebichef moments [25]. Although region-based approaches are global in nature and can be applied to generic shapes, boundary-based techniques appear to be more efficient for handling objects that can be described by their object contours. In this latter category, a number of boundary-based techniques have been proposed, including Fourier descriptors [26], curvature scale space [27], and wavelet descriptors [28]. Our descriptor combines the features of curvature analysis and Fourier analysis, similarly to the technique proposed by El Ghazal et al. [29, 30]. Differently from them, our method is based on recent findings in discrete differential geometry [31], thus resulting in a more robust formulation with respect to the sampling strategy and better classification results. Moreover, Osjanikov et al. applied the concepts of invariant features in the 3D world to the problem of non-rigid shape search and retrieval in large databases [32].

Complementing advances for the general classification problem, machine learning strategies exploiting the existence of large amounts of data have led to significant advances [33, 34]. Many current efforts attempt to work directly on raw data images [35, 36], by designing deep neural networks in which the modeling is hidden in the network design and training strategy and the feature computation and filtering of information is automatically performed by the network. At the same time and in an attempt to simplify classification and automatic shape generation, techniques to reduce the depth of networks by introducing meaningful parameterizations or embeddings of input shapes are gaining interest, since such parameterizations can simplify the automatic classification or shape generation (model-based or “shallow” learning) and reduce the number of training examples [37]. Our work goes towards that direction, since we propose a simplified contour description that can be used either

for supporting machine learning frameworks or for supervised visual analysis. In this work, we focus on the latter aspect.

2.2. Histology analysis

Digital pathology and microscopy-image analysis is widely used in the biomedical domain for comprehensive studies of cell morphology or tissue structure. In most cases, analysis is carried out through manual assessment, which is labor-intensive and prone to inter-observer variations. Computer-aided systems have recently attracted significant interest since they can dramatically reduce the manual efforts and increase reproducibility [38, 39, 17].

Among the various parts composing a computer-aided diagnostic system, nucleus or cell detection and segmentation play a key role to describe the molecular and morphological information underlying the investigated samples [17, 40]. In the past few decades, many efforts have been devoted to automated nucleus/cell detection and segmentation, and an independent field named computational pathology emerged simultaneously to the rapid proliferation of deep learning (DL) models for quantitative analysis of spatial patterns in digitized whole-slide images (WSIs) of cancerous tissue [41]. To this end, various techniques for the detection, extraction, recognition of pathological patterns at various scales have been recently established [42, 43, 44]. Various medical studies have since demonstrated the potential of DL models in detecting neoplastic tissue and recognizing diagnostically relevant structures [45, 44].

One of the most successful and widely used architecture is U-Net, introduced in 2015 by Ronneberger et al. [11]. U-Nets operate on the entire image and jointly segment and provide per-pixel labels, leading to an improvement in spatial segment and label coherence. The same authors also demonstrate that U-Nets improve accuracy on several bio-image segmentation tasks, even when the data set is relatively small [11]. In the context of nuclear segmentation of histopathology images, Chidester et al. [4] enhance U-Nets by enforcing rotation-equivariance to groups, similar in style to group-equivariant CNNs (GCNNs) [46].

Moreover, in order to attract efforts to particular tasks in medical imaging, various challenge contests and public data sets have been published [47, 41]. However, as of this writing, such methods are still far from being accepted in fully automated clinical workflows [48]. Proofreading efforts from domain scientists are, thus, still required to double-check labeling consistency and segmentation accuracy [8].

Consequently, the work presented herein provides a visual analysis framework that supports digital histologists to efficiently carry out investigations on labeling and segmentation quality. Our input data is the automatic segmentation obtained from networks of the U-Net family [4]. The proposed framework then allows for visual analysis in a reduced parameter space obtained by performing dimension reduction on our Fourier-based contour shape descriptor.

To properly capture the visual variance of nuclear shapes under dimension reduction, autoencoders [49] provide a convenient way to effectively uncover latent feature spaces. It is thus no surprise that their use is increasing in popularity [50, 51]. For instance, Xu et al. propose Stacked Sparse

Autoencoders [50] to learn high-level features from pixel intensities. They are then applied to high resolution breast cancer histopathology images. Hou et al. [51] modify the general autoencoder scheme by applying adaptive convolutional filters to match the size of the nuclei to be represented. In this work, we use k-sparse autoencoders [52] to produce feature vectors that describe the inner visual features of nuclei. We then use these feature vectors to augment our proposed shape descriptor, resulting in a description of the exterior (shape descriptor) and interior (auto-encoder) of each nucleus.

2.3. Shape analysis in neuroscience

Recent advances in imaging technology have led to the availability of 3D sparse and dense reconstructions of brain cells at high resolution. This, in turn, has fueled the development of various methods for shape analysis in the context of automatic classification to aid studying the variability associated with different structures and conditions [53, 54]. Likewise, the availability of high-resolution imaging data has also triggered shape analysis studies of brain structures at the nanometer scale [55, 56, 54]. For instance, Queisser et al. [57] propose a method to reconstruct the 3D view of cell nuclear envelopes from laser scanning confocal microscopy data. Wittmann et al. [58] later use this method to show how synaptic activity induces significant modifications in the geometry of the cell nucleus. To study heterogeneities in nuclear shapes obtained through optical projection tomographic microscopy, Nandukumar et al. [59] use conformal mapping to extract rotation-invariant shape descriptors. Finally, Agus et al. [60, 61] perform classification of nuclear brain cells through implicit and explicit shape representations of cell nuclei obtained from electronic-imaging data. They demonstrate an improvement in terms of classification accuracy over previous approaches based on simple spherical or ellipsoidal fittings.

In this paper, we improve existing work [61] by introducing feature vectors based on the spherical harmonics spectrum of mean curvatures. This reduces the amount of data serving as the descriptor from three complex vectors to a single real-valued vector, resulting in faster implementation and processing times.

3. Methodology Overview

Fig. 1 schematically summarizes the InShaDe framework. As can be seen, we use two separate yet similar pipelines for 2D and 3D closed nuclear envelopes that use different microscope imaging techniques as data source. The two methods have important similarities:

- they use the curvature signal (planar curvature in the 2D case, mean curvature in the 3D case),
- they involve parametrization (circular parametrization in the 2D case, and spherical parametrization in the 3D case),
- they involve Fourier analysis (Elliptic Fourier Analysis in the 2D case, and Spherical Harmonics decomposition in the 3D case), It is worth noticing that the Spherical Harmonics framework is a 3D generalization of the Elliptic Fourier Analysis [62],

- they use the same strategy for computing energy descriptors according to the harmonic frequencies [63].

Overview of InShaDe 2D. The input to the InShaDe 2D pipeline are segmented nuclear envelopes of cells obtained by applying an U-Net [4] on microscopic histopathology images (see Fig. 1 top). We then extract closed contours (Sec. 4.1) from each segmentation mask and perform the following processing steps (see also Fig. 1, top):

- (1) contour smoothing (Sec. 4.2),
- (2) geodesically uniform resampling (Sec. 4.3),
- (3) discrete curvature computation (Sec. 4.4),
- (4) opt. feature scaling using standard (z-)scores (Sec. 4.5),
- (5) embedding to constant dimensions (Sec. 4.6),
- (6) elliptic Fourier analysis (“EFA”, Sec.4.7).

Contour smoothing serves to reduce pixelation noise, whereas geodesically uniform resampling removes sampling biases and is a pre-requisite to computing discrete curvatures using discrete differential geometry formulations using osculating circles. Embedding the resulting descriptors in constant dimensions helps in removing noise and spurious frequencies during the EFA stage, but is also necessary to allow for easy comparison between shapes using, e.g., cosine or Euclidean metrics. The Fourier analysis is used to remove shift (i.e., choice of origin) from the parameterization of the closed curve. So far the resulting descriptor is invariant under translation and rotation (3) and invariant under parameterization shift (6). In addition, the optional feature scaling step (4) ensures invariance under uniform scaling. In the result section, we furthermore show how the final descriptor can be used in combination with dimension reduction schemes for visualizing clusters of nuclear shapes with similar geometric characteristics.

Overview of InShaDe 3D. The input of the InShaDe 3D pipeline are closed triangular meshes extracted from image stacks obtained through Serial Section Electron Microscopy acquisition of samples from rodent brains [54]. These shapes represent the envelopes of brain cell nuclei and are obtained from images through a processing pipeline involving automatic segmentation tools as well as manual proof-reading tools [64]. We then process the 3D meshes by performing the following operations (also see Fig. 1, bottom):

- (1) discrete mean curvature computation (Sec. 5.1),
- (2) spherical parameterization using Willmore flow (Sec. 5.2),
- (3) spherical harmonics decomposition (Sec. 5.3),
- (4) computation of invariant energy coefficients (Sec. 5.4).

We use discrete mean curvature (1) as the basis of our embedding and to represent the features of 3D shapes. In contrast to previous formulations of shape decomposition [60, 61], the proposed embedding is based on a single scalar- and real-valued

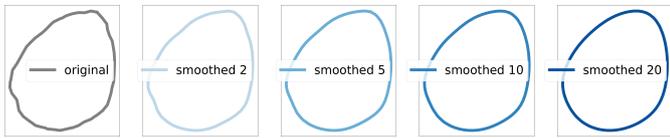


Fig. 2: **Contour smoothing:** we apply iterative contour smoothing to the closed contours (here, $N=271$) extracted from histology images. The higher the number of iteration steps, the smoother the contours: in the example, 2, 5, 10, 20 steps respectively.

function on the spherical domain. The result is a simpler numerical formulation involving only the real part of Spherical Harmonics (SPH) as well as a significantly lower number of coefficients. From the coefficients of a truncated SPH decomposition, we then compute invariant energy coefficients. In the result section (Sec. 6.4), we show how the obtained descriptor can be used for shallow classification of nuclei representing brain cells from different layers of somatosensory cortex of adult rodents.

4. InShaDe 2D

In this section, we provide details for the various processing steps for computing the descriptor for closed shapes extracted from 2D images.

4.1. Contour Extraction & Chordal Parameterization

Given a segmentation mask, we extract a closed contour enveloping each nucleus using iso-contouring (specifically Marching Squares, which is a special case of the Marching Cubes algorithm [65]). We reject open contours (i.e., the nucleus intersects the image boundary) and contours falling into the lowest 5% with respect to their number of samples. Let $\mathcal{C} := \{\mathbf{p}_i\}_1^N$, a closed curve with N vertices \mathbf{p}_i . We let $\Delta_i := \mathbf{p}_{i+1} - \mathbf{p}_i$, the i^{th} edge, consistent with Bobenko [31], and abbreviate $l_i := \|\Delta_i\|_2$ (edge length). We then obtain an initial chordal parameterization $t(\mathcal{C})$ with $t_1 := t(\mathbf{p}_1) = 0$ and $t_{i+1} := t(\mathbf{p}_{i+1}) = \|\Delta_i\|_2 + t_i \forall i > 1$.

4.2. Contour Smoothing

The discrete nature of binary segmentation masks may lead to pixelation artifacts in the extracted contour. To prevent the resulting high spikes in curvature, we pre-smooth contours iteratively, using a superscript $\star^{(k)}$ to denote quantities at iteration k . The process is shown in Fig. 2. Specifically, we replace each vertex with a length-weighted average of the bisector of adjacent edges,

$$\mathbf{p}_i^{(k+1)} = \frac{l_i^{(k)} (\mathbf{p}_{i+1}^{(k)} + \mathbf{p}_i^{(k)}) + l_{i-1}^{(k)} (\mathbf{p}_i^{(k)} + \mathbf{p}_{i-1}^{(k)})}{2(l_i^{(k)} + l_{i-1}^{(k)})}. \quad (1)$$

As shown by Gottschalk [66], this sum of length-weighted edge bisectors computes the barycenter of the points on the piecewise linear curve segment $\mathbf{p}_{i-1}, \mathbf{p}_i, \mathbf{p}_{i+1}$. Since it is a 2-stage convex combination of $\mathbf{p}_{i-1}, \mathbf{p}_i, \mathbf{p}_{i+1}$, it is numerically stable and robust. Similar to virtually all smoothing operators, this

does not yet preserve area. We therefore compute the area $a^{(0)}$ enclosed by the curve prior to smoothing and the area $a^{(k)}$ after each iteration. We then scale the curve by

$$\mathbf{p}_i^{(k)} \leftarrow \mathbf{p}_i^{(k)} \sqrt{\frac{a^{(0)}}{a^{(k)}}}. \quad (2)$$

4.3. Geodesically Uniform Resampling

In order to remove sampling bias and to employ discrete differential geometry formulations for vertex and edge curvature, we perform geodesically uniform resampling. We do so by placing equidistant samples $\tilde{\mathbf{p}}$ on the piece-wise linear curve \mathcal{C} , thereby yielding a new piece-wise linear curve $\tilde{\mathcal{C}}$ that is Arc-length parameterized with respect to a unit scale u . Starting at a point $\mathbf{p}_1 = \tilde{\mathbf{p}}_1$ and $u = 1$, we intersect the edges of $\tilde{\mathcal{C}}$ with a unit circle around \mathbf{p}_1 . This yields between zero and two intersections. If we find two intersections, we select one intersection as $\tilde{\mathbf{p}}_2$ and keep track of the last edge, $\tilde{\Delta}_1 = \tilde{\mathbf{p}}_2 - \tilde{\mathbf{p}}_1$. We then continue intersecting linear segments with unit spheres, but when deciding on $\tilde{\mathbf{p}}_i$, we chose the intersection that maximizes $\langle \tilde{\Delta}_i, \tilde{\Delta}_{i-1} \rangle$, with $\tilde{\Delta}_i$ defined analogously to $\tilde{\Delta}_1$. This enforces progress along the curve and prevents jumping back and forth on the curve. For our data, we did not encounter the case of finding less than two circle-curve intersections. A total of zero intersections would correspond to extremely small contours that cover less than a few pixels after processing; and we remove the bottom 5% shortest curves. One crossing would arise if part of the contour degenerates into a double line segment; Marching Squares does not extract such pathological curves.

Once the best intersection $\tilde{\mathbf{p}}_{\tilde{N}}$ “laps” past $\tilde{\mathbf{p}}_1$, we use $\tilde{\mathbf{p}}_{\tilde{N}} = \tilde{\mathbf{p}}_1$ instead to close the loop. This means that the last edge $\tilde{\Delta}_{\tilde{N}-1}$ may be shorter than unit length. In order to resolve this issue, we now calculate the length L of the curve. Knowing that $\|\tilde{\Delta}_i\|_2 = u$ for all but the last edge, we have $L = (\tilde{N} - 2)u + \|\tilde{\Delta}_{\tilde{N}}\|_2$. To obtain an u for which $u^{-1}L$ is approximately integral, we round $u^{-1}L$ to the nearest integer L' and update $u \leftarrow L'^{-1}L$.

We then revert to placing samples along the original curve \mathcal{C} with the updated spacing u . We repeat this process until the rounding error $\rho = |u^{-1}L - L'|$ (using the old u and the updated L') becomes negligibly small. While we do not have a proof of convergence of this heuristic at the time, we note that the rest of our method is orthogonal to this Arc-length parameterization. This means that, in the future, as more robust methods become available, this step can be exchanged. In all of our experiments, three to five iterations reduced ρ to less than 10^{-4} . Given any number $x \in \mathbb{R}^+$, rounding to the nearest integer changes x by 0.25 on average. We therefore expect that $|1 - u| \approx 0.25L^{-1}$, which we see confirmed in our experiments with typical contour lengths of more than 100 pixel widths (for reference, $L = 100 \rightarrow |1 - u| \approx 2.5 \times 10^{-3}$). The result of this step is a new piece-wise linear curve $\tilde{\mathcal{C}}$ that is Arc-length parameterized with respect to a close-to-unit scale u .

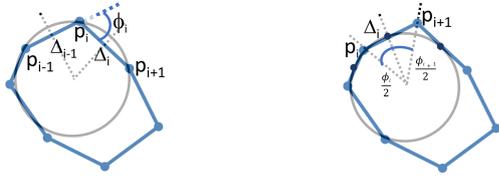


Fig. 3: **Discrete curvatures:** following discrete differential geometry [31] we compute discrete curvatures by considering vertex osculating circles (left), and edge osculating circles (right).

4.4. Discrete Curvatures

For a discrete Arc-length parameterized curve, there are two definitions of discrete curvature based on osculating circles [31] (Sec. 2.3 therein). By defining the turning angle at vertex \mathbf{p}_i as

$$\phi_i \equiv \arccos \langle \Delta_i, \Delta_{i-1} \rangle, \quad (3)$$

and by embedding the planar curve in the $z = 0$ plane (see also Fig. 3), we obtain, assuming for now an Arc-length parameterization with $\|\Delta_i\| = 1$ for all i , the (unsigned) **vertex curvature:**

$$\kappa_v = 2 \frac{|\sin \phi_i|}{\|\mathbf{p}_{i+1} - \mathbf{p}_{i-1}\|_2} = 2 \frac{\|\Delta_i \times \Delta_{i-1}\|_2}{\|\Delta_i + \Delta_{i-1}\|_2}. \quad (4)$$

For the **edge curvature** we use the standard equation [31]:

$$\kappa_e = \tan \frac{\phi_i}{2} + \tan \frac{\phi_{i+1}}{2}. \quad (5)$$

The choice to use unsigned vertex curvature was made to be consistent with the unsigned edge curvature. Using such a discrete differential geometry approach results in much more robust and stable curvature estimates than by using an intermediate interpolating spline. A reason may be that splines tend to over- and undershoot near vertices, and are thus not representative of the curvature in these points. Since one of our goals for the final shape descriptor is *optional* scale-invariance, we still have to scale curvatures back from our arbitrary unit length u to $u = 1$ in case scale-invariance is not desired. This is achieved by dividing each κ_v and κ_e by u^2 . Finally, we interleave vertex and edge curvatures to obtain a high-resolution, coherent descriptor. After this step, we also abandon the notion of curvature “living” on vertices and edges and transition to the notion that the shape descriptor computed so far is a vector in a high-dimensional vector space. We also adopt the notion that this vector represents a 1D periodic signal on a uniform grid on the 2D circle. This interpretation is crucially supported by the fact that all edges have the same length prior to computing curvature. The descriptor computed so far is invariant under translation and rotation, but neither parametric shift nor scale. We now establish the optional scale-invariance followed by shift-invariance.

4.5. Feature Scaling

Given a sequence of curvatures, $\{\kappa_i\}_{i=1}^{2\tilde{N}}$, we compute standard scores (also called z-scores) by mapping

$$\kappa_i \leftarrow \frac{\kappa_i - \mu_\kappa}{\sigma_\kappa}, \quad (6)$$

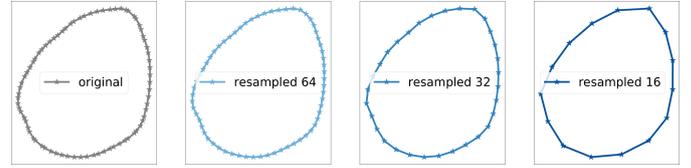


Fig. 4: **Resampling to constant dimension:** to reduce noise and spurious frequencies during the Fourier analysis and to enforce constant dimensionality of our descriptor, we apply uniform resampling through shifted linear interpolation [67]. In this example, we show resampling with 64, 32 and 16 points respectively.

where

$$\mu_\kappa = \frac{1}{2\tilde{N}} \sum_{i=1}^{2\tilde{N}} \kappa_i \quad \text{and} \quad \sigma_\kappa = \frac{1}{2\tilde{N}-1} \sum_{i=1}^{2\tilde{N}} (\kappa_i - \mu_\kappa)^2 \quad (7)$$

are the empiric mean and variance, respectively. Such a scaling is commonly employed in statistics as well as in training convolutional neural networks. However, normally standard scores are computed using global moments derived from the entire data set. This, in turn, does not provide full scale-invariance, since vectors with pre-dominantly small components will stay small. In contrast, by computing individual standard scores we enforce the optional scale-invariance of our descriptor. Assuming that the curvature components of each vector are normal-distributed results in the expectation that all but 0.2% of the data is represented by z-scores in the range $[-3, 3]$.

4.6. Constant Dimensionality

Resampling the contour to a constant dimensionality as depicted in Fig. 4 allows us to control the number of elliptic harmonics in our Elliptic Fourier Analysis in a way to agree with the Nyquist sampling constraint. It is also a pre-requisite for easy comparison of shape descriptors using, e.g., cosine and Euclidean metrics. As an added side-benefit, it also allows us to eliminate remaining traces of noise on the curve. In this paper, we perform this resampling step based on shifted-linear interpolation [67] for the following reasons: (i) shifted linear interpolation achieves performances that compare favourably to cubic interpolation at a much lower computational cost, (ii) shifted-linear interpolation is still convex, albeit with respect to shifted samples. It is thus free of oscillations and the amount of foreign frequencies introduced by resampling can be computed easily.

The basic idea of shifted-linear interpolation is to sample the original signal at positions other than the original underlying sampling grid, followed by standard linear interpolation. Blu et al. prove, somewhat surprisingly, that there is a data-independent and thus constant shift $\tau \approx 0.21$ that results in L^2 -optimal reconstruction of the unknown original signal given only the known samples [67]. Samples κ'_i at shifted positions t'_i are obtained using the infinite impulse response scheme described by Blu et al.,

$$\kappa'_i = -\frac{\tau}{1-\tau} \kappa'_{i-1} + \frac{1}{1-\tau} \kappa_i. \quad (8)$$

It should be noted, however, that linear interpolation on κ' is literally shifted “to the right” by τ , meaning that a sample $\kappa'(t)$

corresponds to $\kappa(t - \tau)$. The resulting interpolation thus becomes a shifted discrete convolution of the hat kernel

$$\Lambda(t) := \begin{cases} 1 - |t| & \text{if } |t| < 1 \\ 0 & \text{otherwise,} \end{cases} \quad (9)$$

with the shifted discrete signal κ' :

$$\kappa(t) = \sum_i \kappa'_i \Lambda(t - t_i - \tau). \quad (10)$$

4.7. Elliptic Fourier Analysis (EFA)

To achieve shift-invariance (i.e., invariance under choice of parametric origin), we consider the Fourier spectrum of each given curve. In particular, we compute elliptic Fourier descriptors [62], similarly to what was proposed by Khazhdan et al. [63] and what has been successfully used in various applications [60, 61, 12].

For a piecewise linear, periodic function $\kappa(t) \in [0, 2\pi]$ representing the curvature of a contour, its Fourier elliptic expansion is obtained through linear combination of elliptic harmonics functions which provide a complete orthonormal basis for the decomposition

$$\kappa(t) = a_0 + \sum_{n=1}^{\infty} \left(a_n \cos\left(\frac{2\pi nt}{T}\right) + b_n \sin\left(\frac{2\pi nt}{T}\right) \right). \quad (11)$$

In order to compute the coefficients for the curvature function $\kappa(t)$ representing closed contours, we normalize the parameterization t to the interval $[0, 2\pi]$. As we are concerned with closed contours, the assumption of periodicity, $t = 0 \equiv 2\pi$ is naturally supported. We then consider the classic method proposed by Kuhl and Giardina [68] for piecewise linear contours. This method essentially equates the discrete time derivative of Eqn. (11), at locations \mathbf{p}_i ,

$$\dot{\kappa}_i := \left. \frac{\partial \kappa}{\partial t} \right|_{t_i}, \text{ thus,} \quad (12)$$

$$\dot{\kappa}_i = \sum_{n=1}^{\infty} \left(-a_n \frac{2\pi n}{T} \sin\left(\frac{2\pi nt_i}{T}\right) + b_n \frac{2\pi n}{T} \cos\left(\frac{2\pi nt_i}{T}\right) \right),$$

with a Fourier expansion of the time derivative of the curvature,

$$\dot{\kappa}_i = \sum_{n=1}^{\infty} \left(\alpha_n \cos\left(\frac{2\pi nt_i}{T}\right) + \beta_n \sin\left(\frac{2\pi nt_i}{T}\right) \right). \quad (13)$$

Noting that in Eqn. (13), the coefficients α_n and β_n can be computed as

$$\alpha_n = \frac{2}{T} \sum_{i=1}^N \dot{\kappa}_i \left(\sin\left(\frac{2\pi nt_i}{T}\right) - \sin\left(\frac{2\pi nt_{i-1}}{T}\right) \right) \text{ and}$$

$$\beta_n = \frac{2}{T} \sum_{i=1}^N \dot{\kappa}_i \left(\cos\left(\frac{2\pi nt_i}{T}\right) - \cos\left(\frac{2\pi nt_{i-1}}{T}\right) \right), \quad (14)$$

Kuhl and Giardina derive the following for the n^{th} harmonic, by equating the two different derivative expressions in Eqn. (13)

and (12):

$$a_n = -\frac{1}{\pi n} \sum_{i=1}^N \dot{\kappa}_i \left(\cos\left(\frac{2\pi nt_i}{T}\right) - \cos\left(\frac{2\pi nt_{i-1}}{T}\right) \right),$$

$$b_n = \frac{1}{\pi n} \sum_{i=1}^N \dot{\kappa}_i \left(\sin\left(\frac{2\pi nt_i}{T}\right) - \sin\left(\frac{2\pi nt_{i-1}}{T}\right) \right). \quad (15)$$

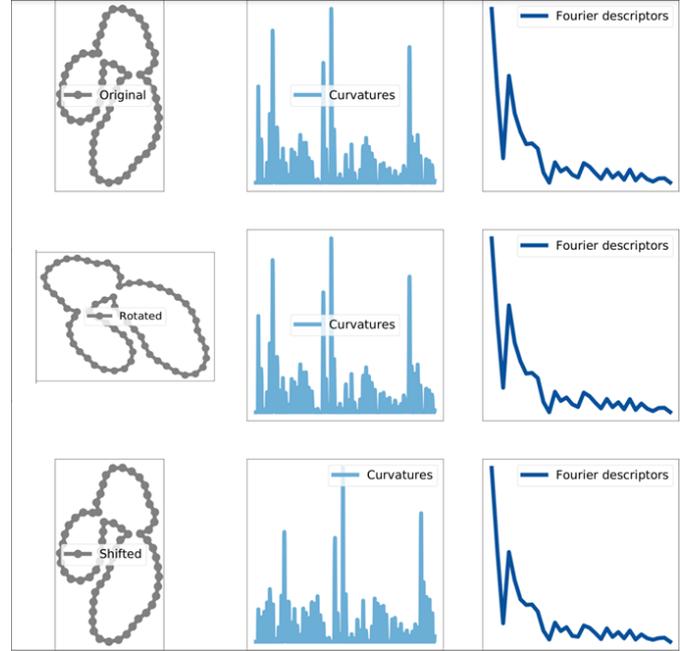


Fig. 5: **Invariant descriptor:** the discrete curvature formulation of InShaDe descriptor is by design invariant to rotations (middle row), while the derive frequency-based energy descriptors are invariant also with respect to shift (bottom row).

We would like to remind here that, according to the Nyquist theorem, the number N_s of contour samples after smoothing and resampling limits the number N_h of harmonics necessary to reconstruct the contour curvature without adding noise ($N_h \leq \frac{N_s}{2}$). Finally, in order to obtain shift-invariance, we compute harmonic energies through the Euclidean norm of the harmonic coefficients [63], resulting in the following Curvature Fourier Descriptor K with

$$K_n = \sqrt{a_n^2 + b_n^2}, \quad (16)$$

which provides a vector of shape features that can be used for various machine learning applications. Like the more commonly employed traditional Fourier transform, the elliptic Fourier transform results in a space-agnostic spectrum, thereby making our descriptor invariant under parameter shift (translation of the underlying domain). In this paper, we chose the elliptic Fourier transform over the traditional Fourier transform since its additional expressiveness resulted in better results. Fig. 5 demonstrates both rotation- and shift-invariance.

Sorted curvatures. We also consider another, much simpler scheme for obtaining shift-invariance, namely, to sort the individual (unsigned) curvatures from highest to lowest (see also

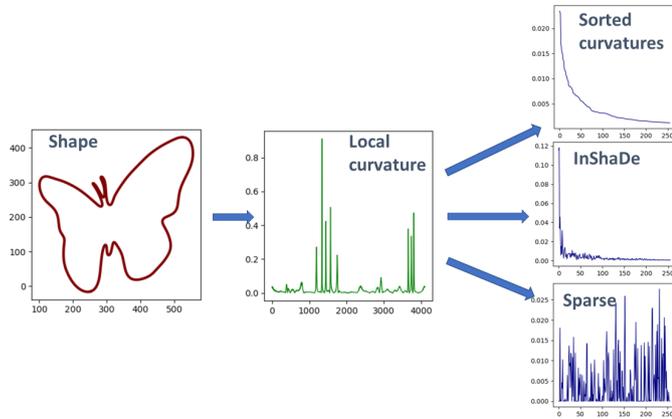


Fig. 6: **Synthetic descriptors:** In addition to the proposed InShaDe energy-based shape descriptors, we considered other representations based on sorted local curvatures (top right) and sparse coding of the foreground segment against a black background (bottom right).

Fig. 6). The feature vector obtained in this way can be further augmented with energy-based coefficients to obtain a composite feature vector. We would like to note that, in our experience, sorted descriptors are outperformed by spectral descriptors if used in stand-alone fashion.

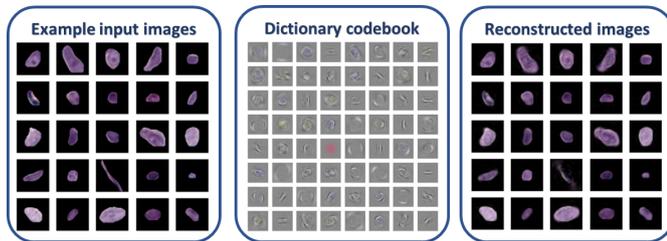


Fig. 7: **Sparse coding:** given images of segmented cell nuclei against a black background (on the left some examples), we use k-sparse autoencoders to find a dictionary codebook (center) that can be used for computing a compressed representation of the original images (on the right some examples).

Sparse-coding based image descriptors. Finally, we consider a compact descriptor of images representing cellular structures obtained through sparse coding. To provide additional shape cues for this case, we place the segmented nuclei on a black background, as shown in Fig. 7, left. Sparse coding methods are typically composed of two steps. Firstly, an offline learning process for finding a dictionary W that sparsely represents the image data $\{I_i\}_{i=0}^N$, and, secondly, an encoding step that maps a given input image I to a compressed feature vector \hat{x} using W , normally through a pursuit algorithm for minimizing the constrained least squares problem

$$\hat{x} = \arg \min_x \|I - Wx\|_2^2, \text{ s.t. } \|x\|_0 < k. \quad (17)$$

For obtaining the codebook (see also Fig. 7, center) and creating the approximated sparse representation of nuclei images (Fig. 7, right), we use the K-sparse autoencoder proposed by Makhzani and Frey [52]. The technique uses linear activation functions and tied weights. In contrast to other autoencoders, only the k largest codes are used while the others are set to zero. The

resulting code k-vector gives us additional cues that, albeit not rotation invariant, can be combined with the InShaDe descriptor (see also Fig. 6, right). Section 6 evaluates various descriptors obtained by composing the three different feature vectors: sorted curvatures, energy coefficient and sparse coding weights.

5. InShaDe 3D

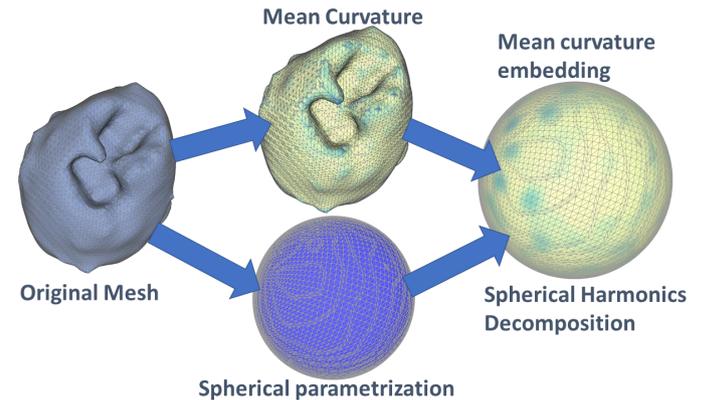


Fig. 8: **InShaDe 3D processing:** starting with a triangular mesh of a closed object, we compute mean curvature H , and we use Willmore flow to obtain a conformal spherical parametrization of the original mesh. The result is a scalar function $H(\theta, \phi)$ over the spherical domain that we further decompose using spherical harmonics.

Our 3D pipeline is a natural adaptation of the 2D case. From a mathematical perspective, the Laplace-Beltrami operator on either 1- or 2-manifold induces a Fourier space. That is, the eigenfunctions of the Laplace-Beltrami operator constitute the "classic" Fourier space in the 2D case and Spherical Harmonics for the 3D spherical case. Albeit we use the elliptic Fourier transform for its superior performance in the 2D case [12], the two pipelines share the same mathematical foundations. We then compute discrete differential geometric attributes of the manifold and express them in this Fourier space. It is imaginable to generalize this even further by utilizing manifold harmonics [69], at the likely expense of higher computational complexity. The bottom half of Fig. 1 depicts the pipeline schematically for easy comparison with the 2D pipeline (Fig. 1, top). The details of the processing steps are provided in this section.

5.1. Mean curvature

According to the differential geometry theory of surfaces, for every twice-differentiable surface we can find the tangent plane for a point on the surface. We can then proceed to define a quadratic form, that is a polynomial containing only terms of degree two, using the two tangent directions x, y [70, chapter 19 therein]. This quadratic form, sometimes called the *shape tensor* describes extrinsic invariants of the surface, such as principal curvatures, at the point where manifold and tangent plane touch. This form is called the *second fundamental form* Π ,

$$f(x, y) \approx \frac{1}{2} [dx \ dy] \Pi \begin{bmatrix} dx \\ dy \end{bmatrix}. \quad (18)$$

The second fundamental form approximates the surface $z = f(x, y)$ with $z = 0$ the plane tangent to the surface (informally, z is the “height over tangent plane”) in a neighborhood around the touching point. Therefore, the idea of the second fundamental form is to measure, in \mathbb{R}^3 , how a surface curves away from its tangent plane at a given point. The eigenvectors of the 2×2 matrix \mathbf{II} are called *principal directions*, and the eigenvalues are called *principal curvatures*, denoted κ_1, κ_2 [9]. Given the principal curvatures, the mean curvature $H = \frac{\kappa_1 + \kappa_2}{2}$ provides a meaningful and natural description of 3D surfaces, and it can be computed on triangular mesh using a discretization of the Laplace-Beltrami operator [9]. In this work, we compute the mean curvature at each vertex of a closed triangle mesh. Given a parameterization (s, t) of the triangle mesh, we thus obtain a discrete, scalar function $H(s, t)$ (see also Fig. 8).

5.2. Spherical parametrization

As hinted at in the last section, our 3D pipeline relies on a surface parameterization. We use a spherical parameterization, which, despite its apparent simplicity, is much harder to obtain than circular parameterizations of planar shapes. From a geometrical point of view, the second fundamental form can be used for classifying surface points according to the signs and values of principal curvatures κ_1 and κ_2 . Of particular interest are so-called umbilic (“locally spherical”) points for which $\kappa_1 = \kappa_2$. A measure for “local sphericity” can thus be defined based on κ_1, κ_2 , such as the Willmore energy of a surface \mathcal{S} ,

$$E = \frac{1}{4} \int_{\mathcal{S}} (\kappa_1 - \kappa_2)^2 dA. \quad (19)$$

The geometric flow associated with this energy,

$$\dot{\mathcal{S}} = -\nabla_{\mathcal{S}} E(\mathcal{S}), \quad (20)$$

will evolve any genus-0 surface \mathcal{S} to a sphere, providing a way to obtain a spherical parameterization. In this work, we use the discrete Willmore flow formulation proposed by Crane et al. [71] that was also previously used by Agus et al. [61]. This spherical parameterization maps each vertex of the 3D input shape to a corresponding point (θ, ϕ) on the unit sphere \mathcal{S}^2 , thereby also providing a parameterization for the discrete scalar function $H(\theta, \phi)$.

5.3. Spherical Harmonics decomposition

The Spherical harmonic basis provides a Fourier basis for functions defined over a sphere. We can thus approximate a generic function defined over a closed surface as a finite linear combination of spherical harmonics $Y_l^m(\theta, \phi)$ up to a given maximum frequency L :

$$F(\theta, \phi) \approx \sum_{l=0}^L \sum_{m=-l}^l w_l^m Y_l^m(\theta, \phi), \quad (21)$$

where the weights w_l^m can be found through least-square error minimization with respect to the samples computed on the original 3D shape. To this end, we used a method similar to [61], with the main difference that, since the mean curvature signal is scalar, we only consider the real part of the

spherical harmonics. As a result, the weight coefficients are also real. Specifically, given the real part of spherical harmonics $R_l^m(\theta, \phi) = \Re(Y_l^m(\theta, \phi))$, a spherical parameterization of the surface \mathcal{S} , $\Sigma_{\mathcal{S}} = \{(\theta_i, \phi_i) = (\theta(v_i), \phi(v_i)) \in \mathcal{S}^2, \forall v_i \in \mathcal{S}\}$, and the mean curvature values computed across the surface $\mathbf{H}_{\mathcal{S}} = \{H_i := H(p_i), \forall p_i \in \mathcal{S}\}$, the spherical harmonic decomposition is obtained by computing the coefficients $\mathbf{w} = \{w_l^m, 0 \leq l \leq L, -l \leq m \leq l\}$ that minimize the square error:

$$\mathbf{w} = \arg \min_{\mathbf{w}} \sum_i \left\| H_i - \sum_l \sum_m w_l^m R_l^m(\theta_i, \phi_i) \right\|_2^2, \quad (22)$$

leading to the linear system

$$\mathbf{R}^T \mathbf{R} \cdot \mathbf{w} = \mathbf{R}^T \mathbf{H}_{\mathcal{S}}. \quad (23)$$

We solve this system using LDL^T factorization, a robust, symmetric pivoting variant of the Cholesky decomposition [72, chapter 4.1.7 therein], in combination with Tikhonov regularization [61]. This regularization add a diagonal matrix \mathbf{T} depending on the Spherical Harmonics order and weighted by a small numeric value ν (see [61] for details). Hence, the final linear system has the following form.

$$(\mathbf{R}^T \mathbf{R} + \nu \mathbf{T}) \cdot \mathbf{w} = \mathbf{R}^T \mathbf{H}_{\mathcal{S}}. \quad (24)$$

We set the Tikhonov regularization weight to $\nu = 10^{-5}$ in all experiments reported in this paper.

5.4. Energy coefficients

Similarly to the 2D formulation we use the harmonic energies, to gain rotation-invariance. Energies are defined as the Euclidean norm of the Spherical Harmonic coefficients \mathbf{w} for each harmonic frequency separately. Specifically, the 3D curvature Fourier descriptor Φ is defined by

$$\Phi_l := \sqrt{\sum_{m=-l}^l (w_l^m)^2}. \quad (25)$$

These coefficients provide a compact descriptor of genus-0 shapes, and can be used for the analysis of nuclear envelopes extracted from Serial Section Electron Microscopy stacks.

6. Results

We implemented our general shape descriptor pipelines and tested them on several challenging use cases. In this section, we first provide details on our implementation (Sec. 6.1) and then provide an evaluation on general shape analysis, on the analysis of histopathological images, and nuclear shapes extracted from Serial Section Electron Microscopy (SSEM) stacks. We separate the evaluation of the two pipelines: for the 2D framework, we report on consistency evaluation performed on classic shape collections commonly used in literature for testing shape retrieval methods (Sec. 6.2). We also provide results obtained with our pipeline on various histology samples for medical diagnostics and neuroscience investigations (Sec. 6.3). For the 3D

framework, we report on the usage of the pipeline for the classification of neural cells reconstructed from a rodent brain sample (Sec. 6.4). In both evaluations, we involve expert domain scientists, for providing a qualitative evaluation of the framework, and for getting suggestions for designing a full visual analytics framework for histology images.

6.1. Implementation notes

The code used to generate the results presented in this paper for the InShaDe 2D pipeline is available in GitHub¹(Python scripts & Jupyter notebooks). After further testing and cleaning, we plan to also release the C++ code for the InShaDe 3D pipeline.

InShaDe 2D. We implemented the 2D geometry processing pipeline in Python using the following building blocks & modules: G-U-Net [4] for automatic segmentation, sklearn, skimage for contour processing and dimension reduction, interactive matplotlib for visualization. For testing the pipeline, we developed simple interactive widgets in which users can compare the clustering visualization in the parameter space to the reconstructed cellular shapes in the histology images. In order to attenuate the amplitudes of high frequencies, we used a frequency equalization scheme weighting of the InShaDe coefficients according to the square root of their order ($w(k) = \sqrt{k}$). Our geometry processing pipeline can be used in combination with different dimension reduction schemes and clustering methods. In this work we use the recent Uniform Manifold Approximation and Projection (UMAP) method, which is based on Riemannian geometry and algebraic topology [14]. For clustering, we used HDBSCAN [73] or k-Means [74, 75, 76] (depending on the case).

InShaDe 3D. We implemented the InShaDe 3D pipeline in C++, using conformal curvature flow based on spin transformations². We furthermore used the Eigen library [77] for the LDL^T solver arising in the SPH least squares optimization. For the SPH functions we used boost and libigl for geometry processing. We then fed the parameters derived from the SPH decomposition to standard machine learning methods (e.g., support vector machines, SVMs) using Python’s sklearn to classify the nuclei.

6.2. Consistency validation

We first performed a consistency validation of the InShaDe descriptor. For this, we used MPEG-7 and Animals [78], which are among the most popular data sets for evaluating and comparing the accuracy of shape retrieval methods [78]. The MPEG-7 shape collection is composed of 1,400 binary images containing objects of 70 different classes [79] (see Fig. 9 left), while the Animals shape collection (see Fig. 9 right) is an even more challenging data set containing 2,000 binary images grouped in 20 classes of 100 animals each one. To test

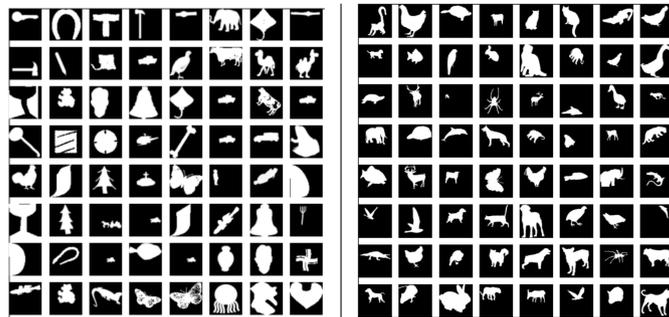


Fig. 9: **Shape retrieval experiments:** we evaluate the InShaDe 2D pipeline through classical shape collections commonly used for testing shape retrieval methods: MPEG-7 (left), and Animals (right).

the InShaDe 2D descriptor we considered three different assessment criteria:

- (1) **Retrieval accuracy** of a basic SVM scheme, trained on an augmented data set. We triple the size of the input data set by adding randomly rotated and shifted copies of original images. Moreover, we use a hyperparameter optimization scheme to find the best SVM linear parameters with respect to cross-correlation accuracy, and we test the obtained model over the original collection. We also show the accuracy of shape retrieval in the form of a confusion matrix to highlight the accuracy differences between classes.
- (2) **Bull’s Eye accuracy**, which is commonly used to score shape retrieval tasks when the number of objects is limited: First, a similarity distance between objects represented by feature vectors is defined. For this, we use the L¹ norm (Manhattan Distance), $d(\mathbf{x}, \mathbf{y}) = \sum_i |x_i - y_i|$. For any object O of class C , find a given K_O nearest neighbors with respect to d ($K_O = 4$ in our case). After that, count how many objects N_O in the set of K_O nearest neighbors share the same class C_O of object O . Finally, the Bull’s Eye score is defined as $B = \sum_O \frac{N_O}{K_O}$.
- (3) **Qualitative visual assessment** of the reduced parameter space obtained by projecting the feature vectors on a 2D plane through dimension reduction techniques and observing how objects cluster together. In our experiments we used UMAP [14].

For the composition of the feature vector we consider three different contributions: the sorted local curvature signal Γ , the elliptic Fourier analysis energy coefficients Φ , and the weight values of k-sparse autoencoding [52] (see the examples in Fig. 10). The purpose of using sparse coding features is not to precisely reconstruct the original image, but rather to extract the important information using very few parameters. The reconstructed images, using only 256 coefficients, clearly represent the main contents of the original images.

We investigate all possible composition permutations: sorted local curvatures alone (Γ), Fourier energy coefficients alone (Φ), sparse coding weights alone (Ω), local curvatures and energy coefficients ($\Gamma\Phi$), local curvatures and sparse coding

¹<https://github.com/HBKUVisCommunity/inshade/>

²<https://github.com/nitronoid/flo>

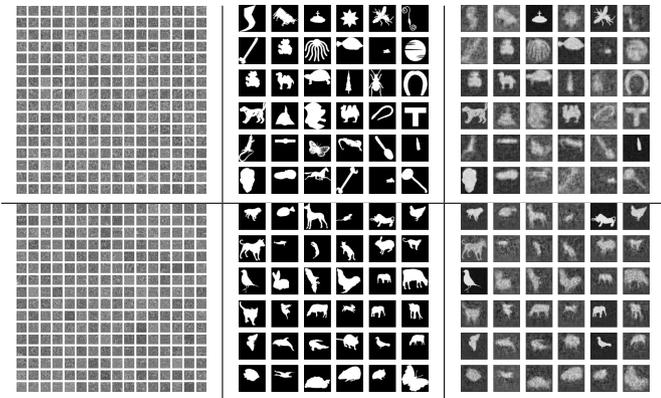


Fig. 10: **Sparse coding:** we used k -sparse autoencoders for encoding MPEG-7 (top row) and Animals (bottom row) shape images. Left: dictionary W . Center: shape examples. Right: corresponding images reconstructed from sparse codes with $k = 256$ coefficients. The reconstructed images clearly represent the main contents of the original images.

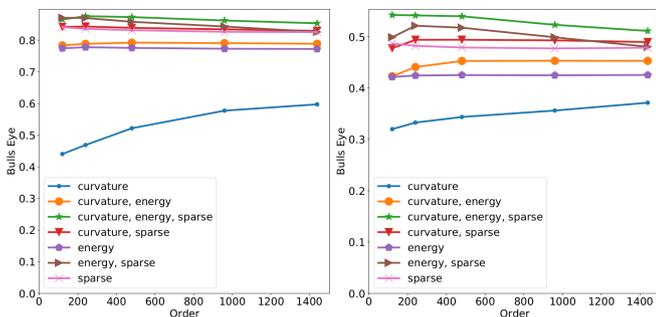


Fig. 11: **Bull's Eye score for shape collections:** various descriptors are compared with respect to the Bull's Eye accuracy on top of 4 retrievals for varying number of coefficient. On the left, MPEG7 results, while on the right Animals results. We compare simple and composed descriptors based on sorted curvatures, elliptic Fourier energy coefficients, and sparse coefficients. The highest Bull's Eye accuracy scores are 0.54 for Animals and 0.87 for MPEG7, and they are obtained for the composed descriptor $\Gamma\Phi\Omega$ containing sorted curvatures, energy coefficients, and sparse coefficients.

weights ($\Gamma\Omega$), energy coefficients and sparse coding weights ($\Phi\Omega$), local curvatures together with energy coefficients and sparse coding weights ($\Gamma\Phi\Omega$). For merging heterogeneous feature vectors, we perform pre-normalization of the various feature vectors. Fig. 11 shows the Bull's Eye score obtained for both shape collections on top of four retrievals for various feature vector compositions at varying number of coefficients. The highest accuracy was obtained for 240 coefficients and with the descriptor $\Gamma\Phi\Omega$ (composing sorted curvatures, energy coefficients, and sparse coefficients). Considering the feature coefficients alone, sparse coding descriptors outperform sorted curvatures and energy descriptors: we suspect that this is due to the fact that sparse coding is able to provide a valid description of both the boundary shape and the inner part of the objects. Nonetheless, incorporating the energy coefficients proved to be beneficial since the composed descriptors can take into account transformations like rotations and shifts. The obtained values are in line with current state of the art methods (see tables in [78]): for example the obtained Bull's Eye score on MPEG-7 data set is 0.87 (versus 0.863 for Hierarchical String Cats [80] and 0.876 for Fourier Transform Group Feature [81]),

while for the Animals data set the Bull's Eye rate is 0.54 (versus 0.436 Hierarchical String Cats [80]). Better performances can be obtained through post-processing retrieval schemes that are orthogonal to our method and can be incorporated successively: for example the Online to Offline O2O scheme [78] applied on top of different descriptors can achieve Bull's Eye score up to 0.99 for the MPEG-7 data set and 0.66 for the Animals data set. In our experiments, we also noticed a slight degradation of accuracy for feature dimensions higher than 500, indicating that the curse of dimensionality can affect the proposed descriptor.

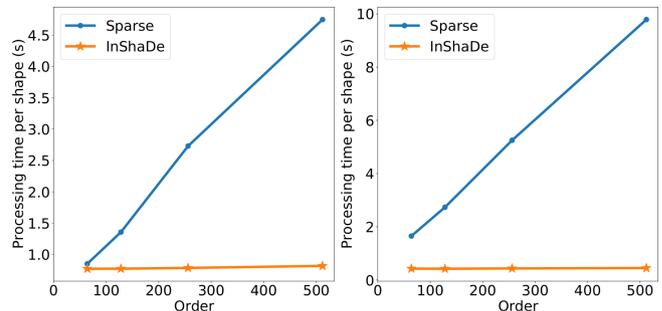


Fig. 12: **Processing time for descriptors:** we compared the CPU processing time of InShaDe 2D energy descriptor with the sparse coding scheme computed through k -sparse autoencoders [52] with respect to the number of coefficients. We report the processing time per shape for MPEG-7 (on the left), and Animals (on the right).

We also compared the CPU computation times per shape between the various descriptors on the MPEG-7 dataset (the average number of vertices per shape is 1917) and Animals (the average number of vertices per shape is 961). For timing measurements, we used a workstation equipped with an Intel i9-9900 CPU (8 cores, 3.1 Ghz), 64 GB of RAM, and an Nvidia RTX 2080 GPU with 8GB RAM. In Fig. 12, we report the processing times per shape as function of number of coefficients, and for both datasets (left: MPEG-7, right: Animals). For the k -sparse autoencoder, we report the training time for 500 epochs using the total images of the dataset for training but averaging the times reported by the number of images. The input image resolution is 256×256 for MPEG-7 and 640×432 for Animals. It is worth noticing that, according to the number of coefficients, the memory resources needed for using k -sparse autoencoder are proportional to the size of input images and the size of output sparse descriptor and can easily reach the limits of available RAM in many systems. Moreover, the InShaDe pipeline can be further accelerated through GPU-friendly implementations that would be able to manage batches of images in parallel.

Fig. 13 and 14 show the UMAP projection of the composed descriptor for the various shapes for both MPEG-7 and Animals data sets as a visual reference (separated in groups of at most 12 labels to reduce clutter).

It appears evident that the proposed descriptor is, in most cases, able to discriminate the shapes of MPEG-7 data set. The clusters for the Animals data set appear more confused, thus confirming the retrieval rates in this work and prior literature. Fig. 15 depicts a typical failure case of our scheme on the Animals data set. A leopard is considered very similar to a cow, a cat and another cow. The middle row shows the local curvature

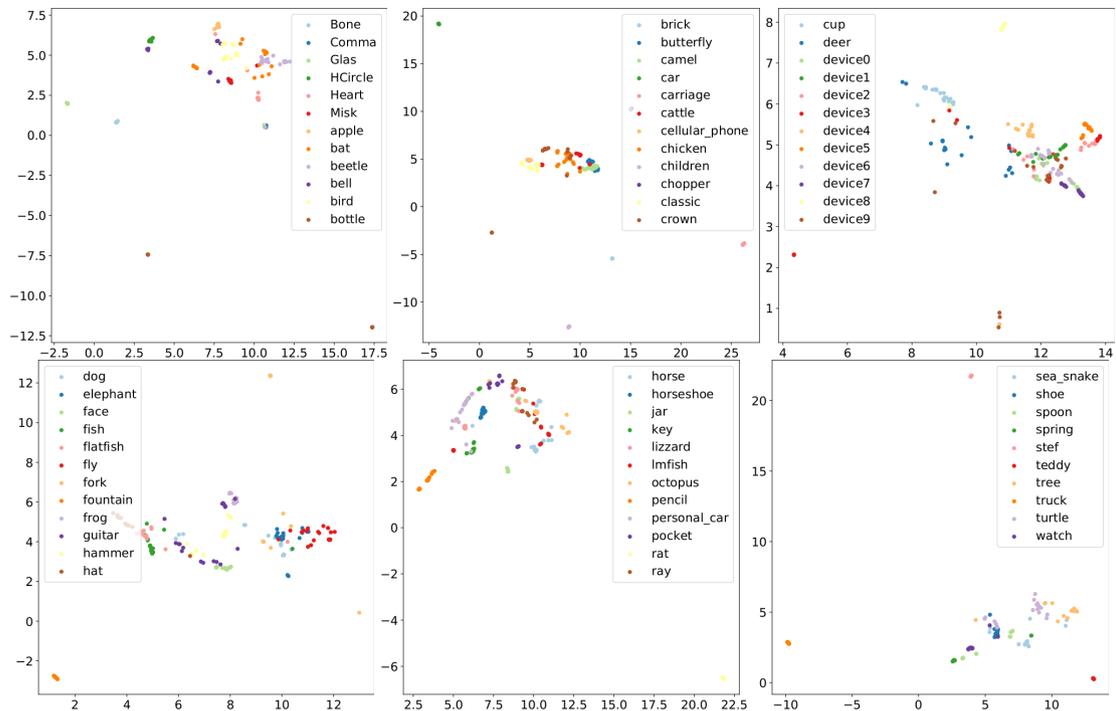


Fig. 13: **MPEG-7 UMAP clustering:** we test UMAP dimension reduction on the composed descriptor $\Gamma\Phi\Omega$ for MPEG-7 collections (256 feature elements per descriptor for a total of 768 features). The Bull's Eye score obtained for this collection was 0.898. In order to reduce visual cluttering, the various shapes are separated in 6 groups of maximum 12 labels.

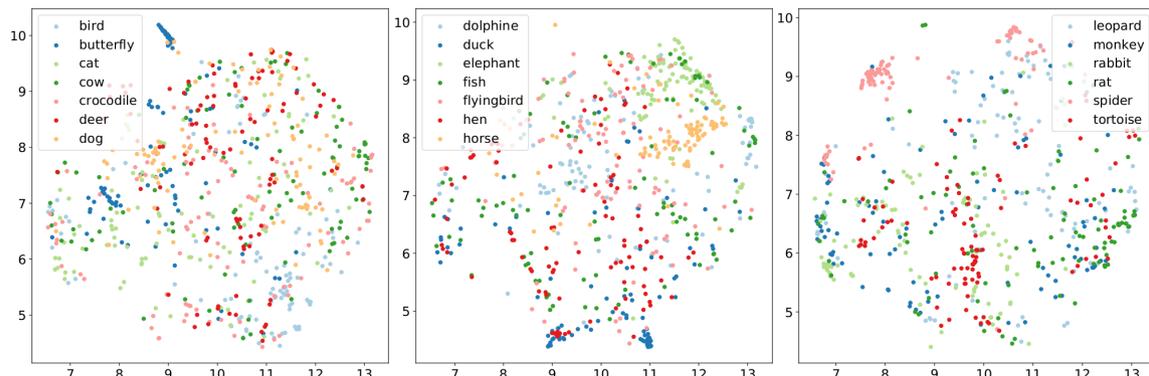


Fig. 14: **Animals UMAP clustering:** we test UMAP dimension reduction on the composed descriptor $\Gamma\Phi\Omega$ for Animals collection. The Bull's Eye score obtained for this collection was 0.543. In order to reduce visual cluttering, the various shapes are separated in 3 groups of maximum 7 labels.

signals computed over the shape contour, while the bottom row shows the composed feature vector containing the sorted curvatures (left), the energy coefficients (middle), and the sparse coefficients (right).

Fig. 16 shows the confusion matrix for linear SVM classification obtained for our composite descriptor $\Gamma\Phi\Omega$. In terms of accuracy, we obtain results aligned with state of the art methods (88.4% for our descriptor on the MPEG-7 data set versus 66% for Curvature based Fourier descriptor [29], 78% for blurred shape models [82], 78% Morphological Pattern Spectrum [83] and 90% for Zernicke moments with geometric features [79]). It is important here to note that the composition of shape features and image features significantly improved retrieval accuracy in agreement with prior work [79] (88.4% for the composed descriptor versus 87.5% for the sparse coding descriptor and 78% of the simple energy shape descriptor proposed in the

conference paper [13]). Given that the results show the proposed descriptor to be consistent for classifying shapes of natural objects, we will now proceed to analyze its performance for the analysis of biomedical images.

6.3. Histopathology analysis

For the analysis of histopathology images, we use public domain data from the MoNuSeg contest [47], and the very recent PanNuke data set [41]. The former contains 30 images from seven organs with unclassified annotations of more than 20k individual nuclei. The latter contains more than 220K labeled nuclei from 19 different tissue types and, as of writing, is the largest open pan-cancer histology data set for nuclei instance segmentation and classification. Finally, we apply the pipeline to a whole slide image representing a paediatric appendix specimen.

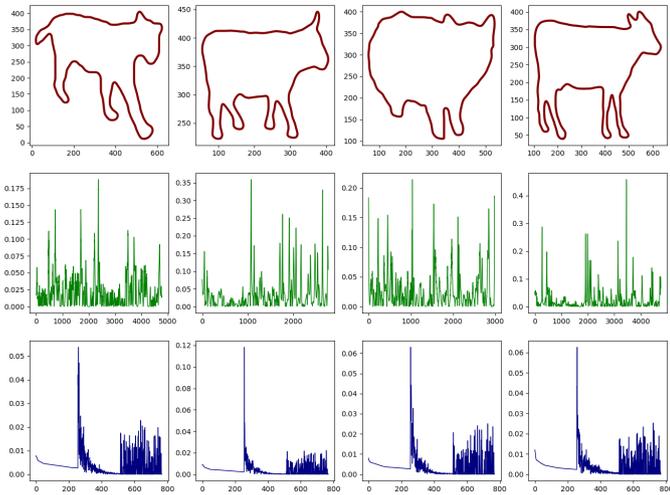


Fig. 15: **Bull's Eye testing:** Here, we show a typical failure case using the L1 Bull's Eye score for the descriptor $\Gamma\Phi\omega$ on the Animals data set. Top row, left to right: a leopard is considered similar to a cow, a cat, and another cow. We also show the curvature signal of each curve (middle row) and composite descriptor (bottom row).

Fig. 17 shows examples of images from the MoNuSeg data set [47] classified by our framework: we reduce the dimensions of the feature descriptors using UMAP, followed by k-Means for clustering. We then color-code contours by cluster. We notice that cells recognized as having similar shape features according to UMAP do not necessarily relate to intuitive discrimination through visible attributes such as length or thickness or smoothness. However, they do not only form feature clusters (same color) but also tend to form spatial clusters. The latter fact can provide additional visual information to digital pathologists for diagnosis through spatial aggregation of such clusters. While further investigation is needed to understand and evaluate the clinical value and to find explainable taxonomies, initial feedback from pathologists confirmed that in many cases nuclear features and clusters can provide decisive information for recognizing specific conditions.

Histopathology shallow classification. We also tested whether our descriptor could be used for shallow classification of nuclear cells for diagnostic purposes. To this end, we tested various composed descriptors (Φ , Ω , $\Phi\Omega$, and $\Gamma\Phi\Omega$) with varying number of features (64, 128, 256), and we trained a linear SVM classifier on PanNuke data set [41] for discriminating between three classes of nuclei: neoplastic cells, inflammatory cells, and others. Fig. 18, left, shows the accuracy performance of the various descriptors considered in this work. The highest accuracy (0.601) is obtained for the composite descriptor $\Gamma\Phi\Omega$ with 256 features per component. However, the improvement of the composed descriptor with respect to the sparse coding descriptor is almost imperceptible (0.578 for the descriptor Ω). In this case, sparse coding captures not just the shape but also texture information, which might be helpful for classification. Thus, the improvement brought by InShaDe is less pronounced compared to Fig. 11, where the input images only contain shape and not texture. In this case, the proposed shape descriptor cannot adequately discriminate the various cell classes accord-

ing to the proposed taxonomy. It is still far from being sufficiently accurate for reliable classification of individual nuclei. However, this preliminary accuracy performance was obtained with a simple SVM classifier, and it can be improved by considering more sophisticated classifiers, like ANNs. In general, the dimension reduction plots show that cells of same type do not cluster together when using the InShaDe descriptor (see in Fig. 18 right some examples). Nonetheless, the presence of outliers in the parameter space can provide pathologists visual hints for proofreading the labeling of nuclei or evaluating the accuracy of contours (see an example in Fig. 18 right, in which a group of images is processed together to obtain a parametric scatter plot to be used for proof-reading patches). Fig. 19 shows an example for the visual analysis of whole slide images (WSIs). We trained an SVM model on InShaDe feature vectors derived from PanNuke data. Then, we used the SVM to classify nuclei in a large-scale, $80,986 \times 99,328$ -pixel WSI of a paediatric appendix specimen. All nuclei are classified as either neoplastic (red), inflammatory (blue), or other (green). Inflamed nuclei cluster together, providing a clear indication of specific affected areas. Neoplastic nuclei are very rare and do not form structured clusters. They are therefore considered classification errors by the domain scientists. Histopathologists can use the processing pipeline for preliminary analysis targeted at the individuation of inflammatory areas. We believe (a hypothesis supported by the domain scientists in our team) that spatial aggregation of classes, i.e., density estimations of the nuclei distribution in space (also see Fig. 19), could become a valuable diagnostic tool in discriminating and individuating different tissue regions. Since a full study into the usefulness of such spatial density estimates is beyond the scope of this work, it is left as a future research direction.

Qualitative evaluation. We tested our InShaDe processing pipeline also on histology images of rodent brain samples stemming from neuroscience. Two expert neuroscientists aided this study by providing a qualitative evaluation of the framework as applied to images obtained with different staining techniques. As a general outcome, the domain scientists particularly appreciated the fact that they could try to map specific features in the shape features space to specific patterns in the histology images.

Specifically, Fig. 20, left, shows the outcomes of Nissl staining of mice brain sections. clustering was obtained with k-Means. The Nissl staining is not specific for particular cell types and is commonly used for cell counting, since it provides an excellent contrast between the cellular and extracellular space. On the other hand, it does not provide a very good contrast between the cytoplasm and the cell nucleus. In the example reported, the contrast allowed the automated algorithm to efficiently segment cell profiles, but only few nuclei were segmented correctly (mostly in light blue, some of them highlighted with blue arrows). In this case, the usage of the parameter space for highlighting the contour shapes in the image space provides visual hints for recognizing particular features, like blurred segmentations of soma mixed with dendrites, appearing as irregular and elongated shapes (see red arrows in Fig. 20 top right). So far, neuroscientists consider the framework potentially useful for proofreading the quality of the staining, and

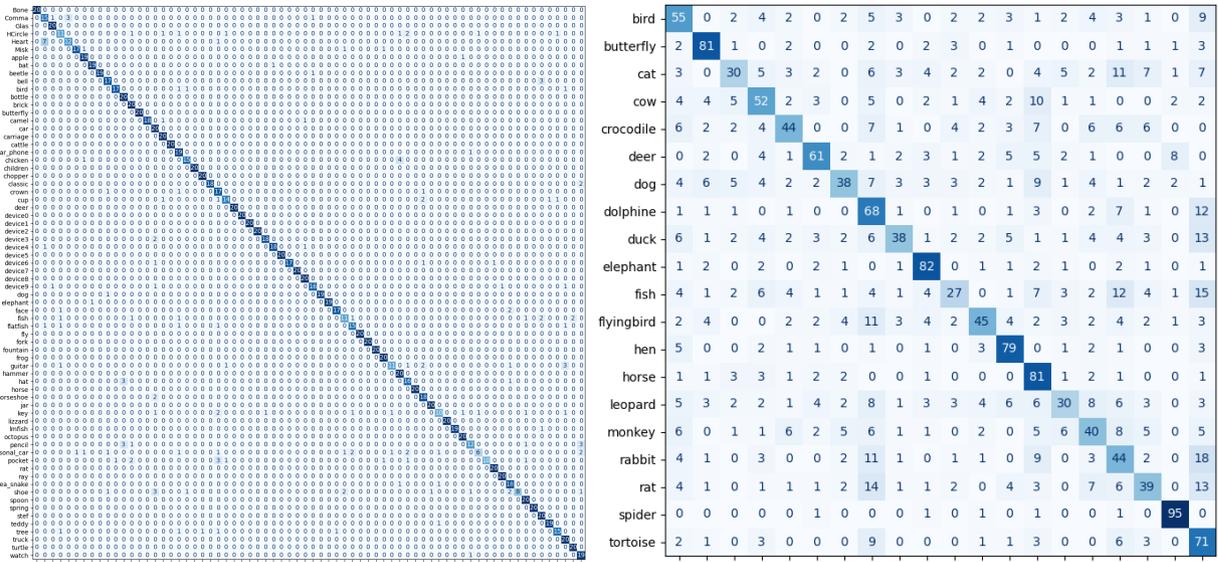


Fig. 16: **Shape retrieval experiments:** a simple Support Vector Machine classifier using our descriptor is able to obtain classification accuracy on par with standard geometry-based classification methods (88.4% for the MPEG-7 and 55% for Animals over the complete shape collection). We also show the full confusion matrix obtained on the testing data (left: MPEG-7, right: Animals collection).

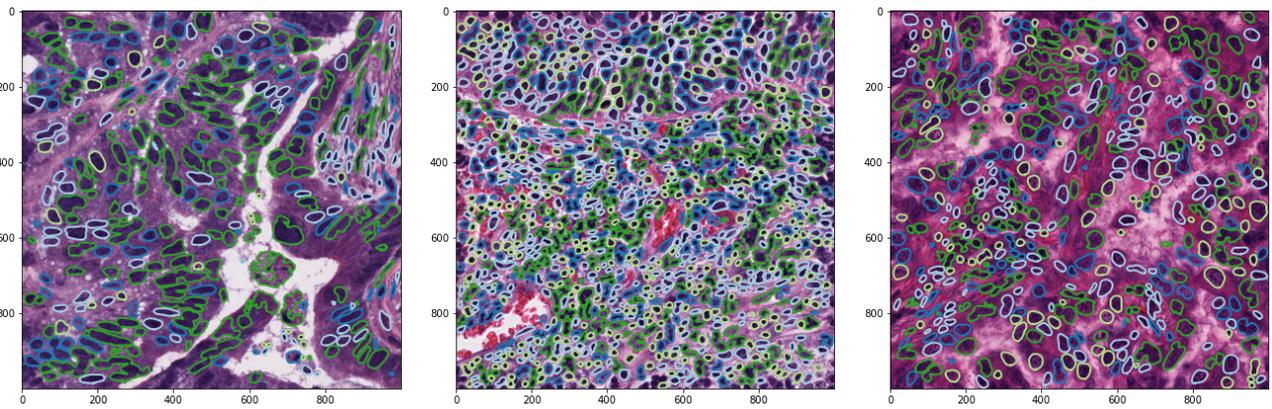


Fig. 17: **Examples of combining the InShAde pipeline with dimension reduction and clustering for visual classification in histology:** Color-coding of shape clusters in the MoNuSeg data set [47] result in recognizable spatial patterns. Similar shape features according to UMAP do not necessarily relate to intuitive discrimination through visible attributes such as length or thickness or smoothness.

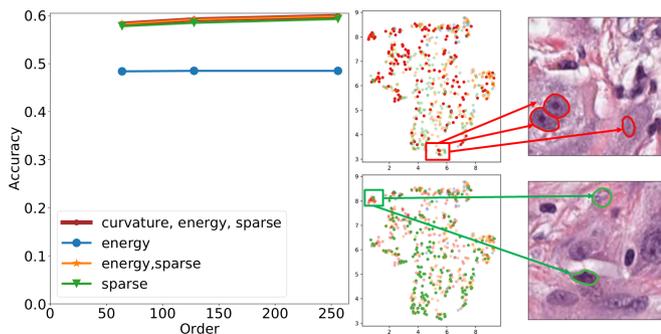


Fig. 18: **Accuracy on PanNuke data set:** we trained a linear SVM model on our descriptor and we compared the accuracy with respect to different feature vectors obtained by composing energy coefficients and/or sparse coding coefficients (left). The maximum obtained accuracy over 3-classes is 0.601. PanNuke represents the largest open pan-cancer histology data set for nuclei instance segmentation and classification [41]. InShAde can be used for proof-reading the quality and the accuracy of labelling (right).

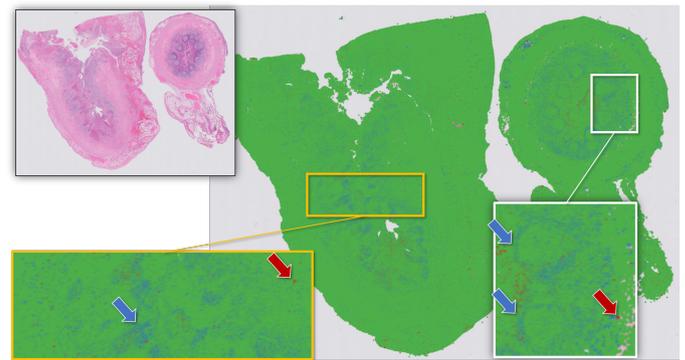


Fig. 19: **Visual analysis of WSI:** the InShAde processing pipeline is applied to the analysis of Whole Slide Images (WSI). A paediatric appendix specimen (top left inset) is analyzed by a linear SVM model trained on InShAde features from PanNuke data. Histopathologists can identify inflammatory areas (blue) against irrelevant background (green). Spurious cancer cells (red) do not form structures and are correctly interpreted as classification error.

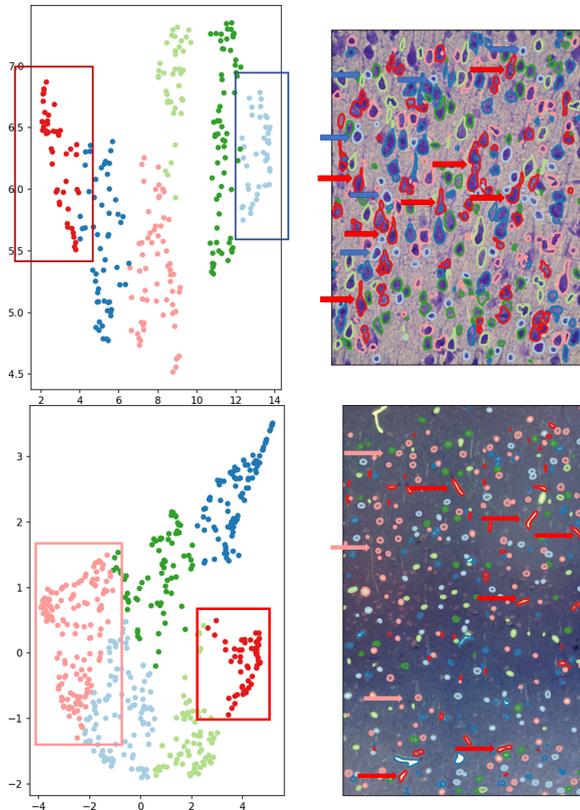


Fig. 20: **Visual analysis of mouse brain sections:** our visual analysis pipeline is used for a neuroscience investigation. Top: a brain section fixed with paraformaldehyde is stained with Cresyl Violet, which highlights Nissl substance in the cytoplasm of neurons. Only few nuclei are segmented correctly (in light blue and highlighted with blue arrows on the right) and in various cases soma are mixed with dendrites (in red and highlighted with red arrows). Bottom: toluidine blue is used in an attempt to discriminate pyramidal neurons nuclei (in pink and highlighted by pink arrows) from blood vessels (in red and highlighted with red arrows) and artifacts.

filtering some information even in case of wrong staining.

Finally, Fig. 20, right, shows a portion of somatosensory cortex from an ultrastructural work on ageing [84, 85, 86]. Nuclei were stained with toluidine blue on semithin sections prepared for electron microscopy in order to count cells. The extracted contours were clustered through k-Means. In this case, the shape feature space enabled scientists to distinguish immediately between blood vessels (in red, with some of them highlighted by arrows), and nuclei from different kind of neurons (pyramidal neurons mostly in pink, highlighted by arrows and with different distribution according to the layer).

6.4. Evaluation of 3D pipeline

For the evaluation of the InShaDe 3D pipeline, we used two collections of 3D reconstructions of brain cells nuclei, extracted from reconstructions of nanometric scale electron microscopy stacks, obtained after imaging a volume of brain parenchyma from layer II/III (see Fig. 21 left) and layer VI (see Fig. 21 right) somatosensory cortex of a P14 rat [54]. The nuclear shapes were manually assigned to known cell types, namely neurons, astrocytes, microglia, pericytes, unknown cells (most likely oligodendrocytes), and endothelium cells for both collections. We used InShaDe 3D feature vectors as the input for

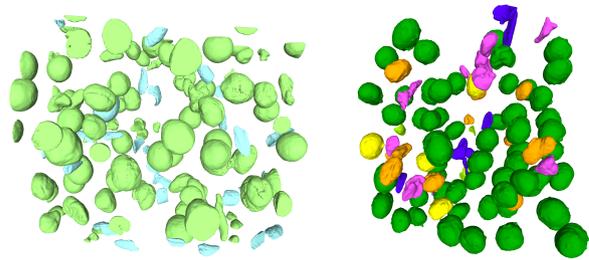


Fig. 21: **Data sets:** we tested the InShaDe 3D pipeline on two collections of brain cells nuclei extracted from layer II/III (left) and layer VI (right) of somatosensory cortex of a P14 rat.

a kernel SVM with radial basis functions. To assess the classification performance, we considered four cases for SPH decomposition with order $L_{max} = 8, 16, 24, 32$, corresponding to the number of rotation-invariant energy descriptors (see also Sec. 5.1). For each case, we performed a grid-search to configure the two hyperparameters in the SVM model: the constant γ of the Gaussian radial basis function, and the weight C for the soft margin regularization function. We chose a grid logarithm-

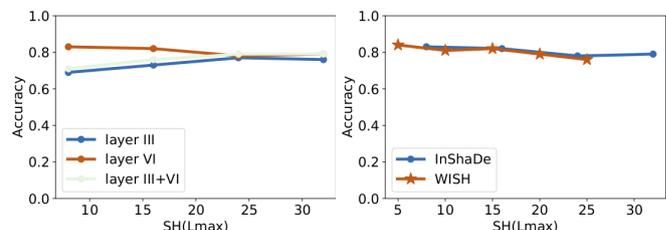


Fig. 22: **Accuracy for 3D nuclei classification:** we trained SVMs with InShaDe 3D features on nuclei shape collections from somatosensory cortex of a juvenile rat in layer III and layer VI. Left: cross accuracy of the SVM model for layer III collection, layer VI collection and full collection. Right: we compare the accuracy performance of InShaDe with respect to WISH [61] on the same data. Despite the simplified formulation, the accuracy is similar.

mic in C (ranging from 10^{-2} to 10^{10}) and γ (ranging from 10^{-9} to 10^3). We then trained the model on 95 nuclear shapes for the layer VI shape collection, and 82 shapes for the layer II/III shape collection. We performed a 5-fold cross-validation using sklearn's StratifiedShuffleSplit function. This partitions the input data into five image sets while maintaining the relative ratio of classes in each set. Four sets were used for training and validation (using an 80/20 split) and the remaining set was used for a blind test. Fig. 22, left, summarizes the best cross-accuracy among the five folds for the SVM models trained on different shape collections and varying feature dimensions.

We also compared the performances of this new formulation with respect to our previous framework WISH [61] (also see Fig. 22 right). The InShaDe accuracy is similar to WISH, with a best score of 83% versus 84% for the layer VI data set. It is worth noting that this is *despite* the new formulation proposed here containing fewer coefficients (one real-valued vector here as opposed to three complex vectors [61]) and is easier to compute numerically. In Fig. 23 we report the average processing times of the Spherical Harmonics coefficients for varying orders of components. It is evident that the simplified formulation results in a dramatic reduction of computation times. These tim-

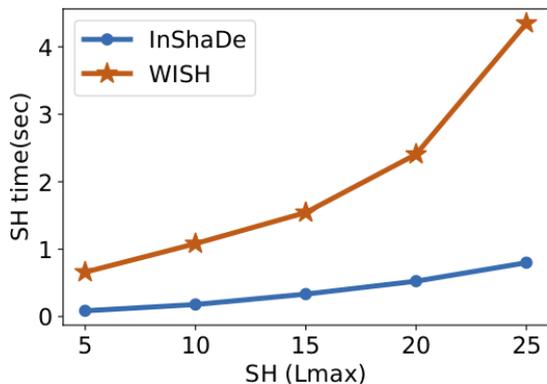


Fig. 23: **Processing time for SH computation:** we compare the processing time for coefficients computation between InShaDe 3D and WISH [61]. The simplified formulation results in a dramatic reduction of processing times.

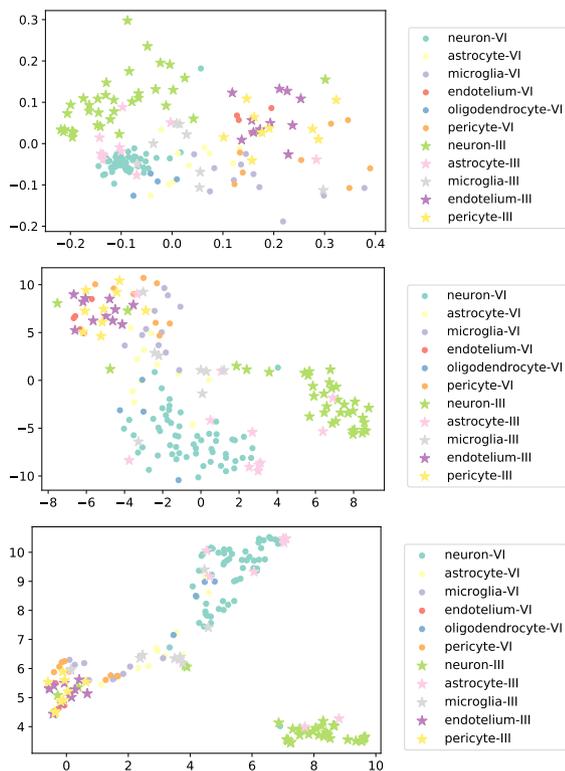


Fig. 24: **Dimension reduction:** dimension reduction projections of InShaDe 3D features for all collection shapes. From top to bottom: PCA, t-SNE, and UMAP. Neurons from different layers tend to form clearly separate clusters.

ings were measured on a Razor Stealth laptop equipped with an Intel i7-8565U CPU (4 cores, 1.8GHz), 16GB RAM and connected through USB-C to an e-GPU NVIDIA Titan RTX with 24GB RAM. Moreover, the spherical parametrization step is identical between the two pipelines InShaDe and WISH, and the mean curvature signal is obtained as free by-product from the usage of Willmore Flow [87]. Given the availability of shape collections extracted from different layers, we tested whether models trained on one collection could be generalized for inference on the shapes of another collection. The resulting performance was poor, in particular for neurons (below 60%). This confirms the hypothesis from domain scientists [54] that nu-

clear envelopes exhibit different shape features according to the layer from which they are extracted. To confirm this point, Fig. 24 shows the full shape collection under different dimension reduction schemes. It appears that it is not possible to cluster together cells extracted from different layers (II & III vs VI). From these preliminary results, it appears evident that:

- it is difficult to find models using InShaDe 3D descriptors that can generalize the classification of cell types regardless of the layers from which they are extracted;
- neurons of different layers appear to form separate clusters, suggesting a shape variability depending on the area from which they are extracted [88]. This is an interesting hypothesis worthy of further investigations which we plan to carry out in the future.

6.5. Discussion

We summarize the main outcomes of this study as follows.

- **Relationships between shape parameter space and image space:** in various cases we notice that spatial clusters of cells exhibit closer shape features in the reduced parameter space. Further investigation is needed to understand whether and in which cases spatial patterns or clusters in the image space correspond to patterns or clusters in the parameter space, and to associate shape clusters to specific taxonomies. In this context, we would like to remind that performing clustering on parameter space obtained after dimension reduction is still considered a complex task prone to producing unreliable results [89]. Therefore, we plan to explore different automatic and manual dimension reduction techniques to support domain scientists during their analysis.
- **Coupling with image descriptors:** we integrated the InShaDe 2D with sparse coding for decomposing the inner part of nuclei as function of specific texture patterns with different physical and molecular characteristics. Sparse coding features outperform the proposed descriptor (we suspect that the reason is related to the fact that they can also describe shape boundary features), while composed descriptors provide slightly better discrimination capabilities for standard evaluation datasets because they recover invariance with respect to transformations. However, the composed descriptors are not reliable yet for fine-grain diagnosis on histopathology images. Spatial aggregation (i.e., class density estimation) could be used to alleviate this problem, but future research is needed. We also suspect that we have hit a performance wall for engineered and model-based descriptors. Therefore, we plan to integrate model-based descriptors into more general deep learning architectures.
- **Caveats due to staining techniques:** depending on the structure to be identified within a cell, or the type of tissue, a large plethora of immunohistochemical staining techniques are available. The proposed analysis framework can provide effective proof-reading tools for checking the

quality of staining methods and semi-automatically individuating the structures of interest.

- **Taxonomy-based visual analytics system:** a real challenge in the analysis of histology images is the difficulty to individuate correct taxonomies of nuclei in order to simplify understanding and diagnosis. A visual analytics framework incorporating contour analysis, image analysis, and expert domain knowledge would help digital pathologists in labeling and proof-reading, and would provide fast ways for creating labeled data for more sophisticated artificial intelligence frameworks. To this end, our processing pipeline provides encouraging results and can be easily integrated in such systems.
- **2D Arc-length parameterization:** while we have yet to observe our Arc-length parameterization algorithm to diverge, we do not have a formal proof of convergence at the time of writing. We believe it works so well since changes in u happen very gradually and the original curve remains untouched. Each reparameterization attempt therefore slides vertices around the input curve. While formal analysis is hindered by the fact that our method is discontinuous at original vertices, we believe a full treatise to be an interesting direction for future work.
- **Benefits of scale-invariance:** for the MPEG-7 and animals data sets, utilizing the optional scale-invariance boosts performance by up to 10% in many of our experiments. This should come as no surprise, since, e.g., the outline of a butterfly stays the outline of a butterfly under magnification and minification. Consequently, deep learning-based pipelines have made rescaling a main step of their data augmentation stage, in an attempt to achieve de-facto rather than by-design scale-invariance. What is surprising, however, is that scale-invariance added only marginal and in many cases statistically insignificant improvements for cell nuclei classification. We believe that this is due to the “apparent size” problem, in which cell nuclei always appear smaller than the original size due to slicing. It seems that having plenty of slices under different angles at the classifiers disposal is more important than to remove scale-variance. A full analysis of this problem is beyond the scope of this paper and offers an interesting future research direction.
- **Limitations of 3D pipeline:** the encouraging results obtained with our 3D formulation are counterbalanced by two important limiting bottlenecks. Firstly, the process for producing nuclear surfaces from electron microscopy image stacks is still time consuming and requiring highly specialized human efforts. Even though important progresses in automatic segmentation of EM stacks has recently been made [19], custom models for automatic extraction of nuclei are not available to our knowledge. We plan to focus future efforts towards this direction. Secondly, the spherical parameterization task is complex and can be unstable. One of its limitations is that it cannot be applied to arbitrary closed shapes but only genus-0 and (if the flow

is appropriately regularized) genus-1 surface (that is, surfaces either homeomorphic to spheres or torii with at most 1 hole). To overcome these limitations, we plan to investigate more general invariant formulations based on manifold harmonics [90].

7. Conclusion

We have presented a general shape processing framework rooted in a novel differential-geometry-based descriptor of closed contours and surfaces. Our descriptor provides an embedding into a fixed-dimensional feature space that can be utilized for various applications, which range from serving as input feature for deep and shallow learning techniques to supporting dimension reduction schemes for providing a visual reference for clustering collection of shapes. While our methods are of general use, our work is motivated by the study of cellular nuclear envelopes extracted from histopathological images and serial section electron microscopy stacks. In this context, we have shown the capabilities of the proposed framework for visual analysis and unsupervised classification. Our results are very encouraging and we identify several major areas of future work in the previous section. In particular, we plan to develop, on top of our pipeline, a taxonomy-based visual analytics system to simplify study and diagnosis.

Acknowledgments. This work was supported by the College of Science and Engineering at Hamad Bin Khalifa University. It was partially supported by Sardinian Regional Authorities under projects VIGELAB and DIFRA, and by the KAUST-EPFL Alliance for Integrative Modeling of Brain Energy Metabolism <https://www.kaust.edu.sa/en> under KAUST CRG6 Grant No. 2313.

References

- [1] Boges, DJ, Agus, M, Magistretti, PJ, Cali, C. Forget about electron micrographs: A novel guide for using 3d models for quantitative analysis of dense reconstructions. In: Volume Microscopy. 2020, p. 263–304. doi:10.1007/978-1-0716-0691-9_14.
- [2] Madabhushi, A, Lee, G. Image analysis and machine learning in digital pathology: Challenges and opportunities. *Medical Image Analysis* 2016;33:170–175. doi:10.1016/j.media.2016.06.037.
- [3] Kurnianggoro, L, Wahyono, SK, Jo, KH. A survey of 2D shape representation: Methods, evaluations, and future research directions. *Neurocomp* 2018;300:1–16. doi:10.1016/j.neucom.2018.02.093.
- [4] Chidester, B, Ton, TV, Tran, MT, Ma, J, Do, MN. Enhanced rotation-equivariant U-Net for nuclear segmentation. In: Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition Workshops (CVPRW). 2019, p. 0. doi:10.1109/CVPRW.2019.00143.
- [5] Gamper, J, Koohbanani, NA, Benet, K, Khuram, A, Rajpoot, N. PanNuke: An open pan-cancer histology dataset for nuclei instance segmentation and classification. In: Proc. European Congress on Digital Pathology. 2019, p. 11–19. doi:10.1007/978-3-030-23937-4_2.
- [6] Scheffer, LK, Xu, CS, Januszewski, M, Lu, Z, Takemura, Sy, Hayworth, KJ, et al. A connectome and analysis of the adult drosophila central brain. *Elife* 2020;9:e57443.
- [7] Motta, A, Berning, M, Boergens, KM, Staffler, B, Beining, M, Loomba, S, et al. Dense connectomic reconstruction in layer 4 of the somatosensory cortex. *Science* 2019;366(6469).
- [8] Tajbakhsh, N, Jeyaseelan, L, Li, Q, Chiang, JN, Wu, Z, Ding, X. Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation. *Medical Image Analysis* 2020;:101693doi:10.1016/j.media.2020.101693.

- [9] Bobenko, AI, Sullivan, JM, Schröder, P, Ziegler, G. Discrete differential geometry. Springer; 2008. URL: <https://www.springer.com/gp/book/9783764386207>; ISBN 978-3-7643-8621-4.
- [10] Crane, K, Wardetzky, M. A glimpse into discrete differential geometry. Notices of the American Mathematical Society 2017;64(10). doi:10.1090/noti1578.
- [11] Ronneberger, O, Fischer, P, Brox, T. U-Net: Convolutional networks for biomedical image segmentation. In: Int'l Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI). 2015, p. 234–241. doi:10.1007/978-3-319-24574-4_28.
- [12] Diaz, G, Zuccarelli, A, Pelligra, I, Ghiani, A. Elliptic Fourier analysis of cell and nuclear shapes. Computers and Biomedical Research 1989;22(5):405–414. doi:10.1016/0010-4809(89)90034-7.
- [13] Agus, M, Al-Thelaya, K, Cali, C, Boido, M, Yang, Y, Pintore, G, et al. InShaDe: Invariant shape descriptors for visual analysis of histology 2d cellular and nuclear shapes. In: Proc. Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM). 2020, p. 61–70. doi:10.2312/vcbm20201173.
- [14] McInnes, L, Healy, J, Melville, J. UMAP: Uniform manifold approximation and projection for dimension reduction. arXiv preprint; 2018. URL: <https://arxiv.org/abs/1802.03426>.
- [15] Mingqiang, Y, Kidiyo, K, Joseph, R. A survey of shape feature extraction techniques. Pattern Recognition 2008;15(7):43–90. doi:10.5772/6237.
- [16] Irshad, H, Veillard, A, Roux, L, Racoceanu, D. Methods for nuclei detection, segmentation, and classification in digital histopathology: A review—Current status and future potential. IEEE Reviews in Biomedical Engineering 2013;7:97–114. doi:10.1109/RBME.2013.2295804.
- [17] Xing, F, Yang, L. Robust nucleus/cell detection and segmentation in digital pathology and microscopy images: A comprehensive review. IEEE Reviews in Biomedical Engineering 2016;9:234–263. doi:10.1109/RBME.2016.2515127.
- [18] Jothi, JAA, Rajam, VMA. A survey on automated cancer diagnosis from histopathology images. Artificial Intelligence Review 2017;48(1):31–81. doi:10.1007/s10462-016-9494-6.
- [19] Moen, E, Bannon, D, Kudo, T, Graf, W, Covert, M, Van Valen, D. Deep learning for cellular image analysis. Nature methods 2019;16(12):1233–1246.
- [20] Yang, M, Kpalma, K, Ronsin, J. Shape-based invariant feature extraction for object recognition. In: Advances in Reasoning-Based Image Processing Intelligent Systems. 2012, p. 255–314. doi:10.1007/978-3-642-24693-7_9.
- [21] Xu, D, Li, H. Geometric moment invariants. Pattern Recognition 2008;41(1):240–249. doi:10.1016/j.patcog.2007.05.001.
- [22] Khotanzad, A, Hong, YH. Invariant image recognition by Zernike moments. IEEE Transactions on Pattern Analysis and Machine Intelligence 1990;12(5):489–497. doi:10.1109/34.55109.
- [23] Novotni, M, Klein, R. Shape retrieval using 3D Zernike descriptors. Computer Aided Design 2004;36(11):1047–1062. doi:10.1016/j.cad.2004.01.005.
- [24] Shu, H, Luo, L, Bao, X, Yu, W, Han, G. An efficient method for computation of Legendre moments. Graphical Models 2000;62(4):237–262. doi:10.1006/gmod.2000.0523.
- [25] Mukundan, R, Ong, S, Lee, PA. Image analysis by Tchebichef moments. IEEE Transactions on image Processing 2001;10(9):1357–1364. doi:10.1109/83.941859.
- [26] Zhang, D, Lu, G. A comparative study of curvature scale space and fourier descriptors for shape-based image retrieval. Journal of Visual Communication and Image Representation 2003;14(1):39–57. doi:10.1016/S1047-3203(03)00003-8.
- [27] Mokhtarian, F. Silhouette-based isolated object recognition through curvature scale space. IEEE Transactions on Pattern Analysis and Machine Intelligence 1995;17(5):539–544. doi:10.1109/34.391387.
- [28] Laga, H, Takahashi, H, Nakajima, M. Spherical wavelet descriptors for content-based 3d model retrieval. In: IEEE Int'l Conf. on Shape Modeling and Applications. 2006, p. 15–15. doi:10.1109/SMI.2006.39.
- [29] El-Ghazal, A, Basir, O, Belkasim, S. A novel curvature-based shape fourier descriptor. In: Proc. IEEE Int'l Conf. on Image Processing. 2008, p. 953–956. doi:10.1109/ICIP.2008.4711914.
- [30] El-Ghazal, A, Basir, O, Belkasim, S. Invariant curvature-based Fourier shape descriptors. Journ of Visual Comm and Image Representation 2012;23(4):622–633. doi:10.1016/j.jvcir.2012.01.011.
- [31] Bobenko, AI. Geometry II: Discrete differential geometry. Lecture Notes, Math Dept., TU Berlin, Germany; 2015.
- [32] Ovsjanikov, M, Bronstein, AM, Bronstein, MM, Guibas, LJ. Shape google: a computer vision approach to isometry invariant shape retrieval. In: 2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops. IEEE; 2009, p. 320–327.
- [33] Chang, AX, Funkhouser, T, Guibas, L, Hanrahan, P, Huang, Q, Li, Z, et al. Shapenet: An information-rich 3D model repository. arXiv preprint; 2015. URL: <https://arxiv.org/abs/1512.03012>.
- [34] Wu, Z, Song, S, Khosla, A, Yu, F, Zhang, L, Tang, X, et al. 3D shapenets: A deep representation for volumetric shapes. In: Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR). 2015, p. 1912–1920. doi:10.1109/CVPR.2015.7298801.
- [35] Wang, PS, Liu, Y, Guo, YX, Sun, CY, Tong, X. O-CNN: Octree-based convolutional neural networks for 3D shape analysis. ACM Transactions on Graphics 2017;36(4):72. doi:10.1145/3072959.3073608.
- [36] Yi, L, Su, H, Guo, X, Guibas, LJ. SyncSpecCNN: Synchronized spectral CNN for 3D shape segmentation. In: Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR). 2017, p. 2282–2290. doi:10.1109/CVPR.2017.697.
- [37] Sinha, A, Bai, J, Ramani, K. Deep learning 3D shape surfaces using geometry images. In: Proc. ECCV. 2016, p. 223–240. doi:10.1007/978-3-319-46466-4_14.
- [38] Jeong, WK, Schneider, J, Turney, S, Faulkner-Jones, BE, Meyer, D, Westermann, R, et al. Interactive histology of large-scale biomedical image stacks. IEEE Transactions on Visualization and Computer Graphics 2010;16(6):1386–1395. doi:10.1109/TVCG.2010.168.
- [39] Jeong, W, Schneider, J, Hansen, A, Lee, M, Turney, SG, Faulkner-Jones, BE, et al. A collaborative digital pathology system for multi-touch mobile and desktop computing platforms. Computer Graphics Forum 2013;32:227–242.
- [40] Rohde, GK, Ribeiro, AJ, Dahl, KN, Murphy, RF. Deformation-based nuclear morphometry: Capturing nuclear shape variation in HeLa cells. Cytometry Part A: The Journal of the International Society for Analytical Cytology 2008;73(4):341–350. doi:10.1002/cyto.a.20506.
- [41] Gamper, J, Koohbanani, NA, Graham, S, Jahanifar, M, Khurram, SA, Azam, A, et al. PanNuke dataset extension, insights and baselines. arXiv preprint; 2020. URL: <https://arxiv.org/abs/2003.10778>.
- [42] Xing, F, Xie, Y, Su, H, Liu, F, Yang, L. Deep learning in microscopy image analysis: A survey. IEEE Transactions on Neural Networks and Learning Systems 2017;29(10):4550–4568. doi:10.1109/TNNLS.2017.2766168.
- [43] Zheng, Y, Jiang, Z, Xie, F, Zhang, H, Ma, Y, Shi, H, et al. Feature extraction from histopathological images based on nucleus-guided convolutional neural network for breast lesion classification. Pattern Recognition 2017;71:14–25. doi:10.1016/j.patcog.2017.05.010.
- [44] Al-Milaji, Z, Ersoy, I, Hafiane, A, Palaniappan, K, Bunyak, F. Integrating segmentation with deep learning for enhanced classification of epithelial and stromal tissues in H&E images. Pattern Recognition Letters 2019;119:214–221. doi:10.1016/j.patrec.2017.09.015.
- [45] Nagpal, K, Foote, D, Liu, Y, Chen, PHC, Wulczyn, E, Tan, F, et al. Development and validation of a deep learning algorithm for improving gleason scoring of prostate cancer. NPJ Digital Medicine 2019;2(1):1–10. doi:10.1038/s41746-019-0112-2.
- [46] Cohen, T, Welling, M. Group equivariant convolutional networks. In: Proc. Int'l Conf. on Machine Learning. 2016, p. 2990–2999. URL: <http://proceedings.mlr.press/v48/cohen16.html>.
- [47] Kumar, N, Verma, R, Anand, D, Zhou, Y, Onder, OF, Tsougenis, E, et al. A multi-organ nucleus segmentation challenge. IEEE Transactions on Medical Imaging 2019;39(5):1380–1391. doi:10.1109/TMI.2019.2947628.
- [48] Schneider, J, Agus, M. Reflections on the clinical acceptance of artificial intelligence. In: Househ, M, Kushniruk, A, Borycki, E, editors. Multiple Perspectives on Artificial Intelligence in Healthcare: From Challenges to Opportunities. SpringerBriefs; 2021, p. 0. doi:10.1007/978-3-030-67303-1; to appear.
- [49] Kingma, DP, Welling, M. Auto-encoding variational Bayes. arXiv preprint; 2013. URL: <https://arxiv.org/abs/1312.6114>.
- [50] Xu, J, Xiang, L, Liu, Q, Gilmore, H, Wu, J, Tang, J, et al. Stacked sparse autoencoder (SSAE) for nuclei detection on breast cancer histopathology images. IEEE Transactions on Medical Imaging 2015;35(1):119–130. doi:10.1109/TMI.2015.2458702.

- [51] Hou, L, Nguyen, V, Kanevsky, AB, Samaras, D, Kurc, TM, Zhao, T, et al. Sparse autoencoder for unsupervised nucleus detection and representation in histopathology images. *Pattern recognition* 2019;86:188–200. doi:10.1016/j.patcog.2018.09.007.
- [52] Makhzani, A, Frey, B. K-sparse autoencoders. arXiv preprint; 2013. URL: <https://arxiv.org/abs/1312.5663>.
- [53] Kalinin, AA, Allyn-Feuer, A, Ade, A, Fon, GV, Meixner, W, Dilworth, D, et al. 3D shape modeling for cell nuclear morphological analysis and classification. *Scientific reports* 2018;8. doi:10.1038/s41598-018-31924-2.
- [54] Cali, C, Agus, M, Kare, K, Boges, D, Lehvaslaiho, H, Hadwiger, M, et al. 3D cellular reconstruction of cortical glia and parenchymal morphometric analysis from serial block-face electron microscopy of juvenile rat. *Progress in Neurobiology* 2019;:101696doi:j.pneurobio.2019.101696.
- [55] Härtel, S, Jara, J, Lemus, C, Concha, M. 3D morphotopological analysis of asymmetric neuronal morphogenesis in developing zebrafish. *Computational modelling of objects represented in images Fundamentals, methods and applications* 2018;6:215–220.
- [56] Cali, C, Baghabra, J, Boges, DJ, Holst, GR, Kreshuk, A, Hamprecht, FA, et al. Three-dimensional immersive virtual reality for studying cellular compartments in 3D models from EM preparations of neural tissues. *Journal of Computational Neurology* 2016;524(1):23–38. doi:10.1002/cne.23852.
- [57] Queisser, SG, Wittmann, M, Bading, H, Wittum, G. Filtering, reconstruction, and measurement of the geometry of nuclei from hippocampal neurons based on confocal microscopy data. *Journal of Biomedical Optics* 2008;13(1):014009. doi:10.1117/1.2829773.
- [58] Wittmann, M, Queisser, G, Eder, A, Wiegert, JS, Bengtson, CP, Hellwig, A, et al. Synaptic activity induces dramatic changes in the geom. of the cell nucleus: interplay between nuclear structure, histone H3 phosphorylation, and nuclear calcium signaling. *Journal of Neuroscience* 2009;29(47):14687–14700. doi:10.1523/JNEUROSCI.1160-09.2009.
- [59] Nandakumar, V, An, X, Wang, Y, Johnson, R, Meldrum, D. Conformal mapping of nuclei in 3D tomographic cell images to assess shape heterogeneity. In: *Proc. IEEE Int'l Symp. on Biomedical Imaging*. 2012, p. 222–225. doi:10.1109/ISBI.2012.6235524.
- [60] Agus, M, Veloz Castillo, M, Garnica Molina, JF, Gobetti, E, Lehvaslaiho, H, Morales Tapia, A, et al. Shape analysis of 3D nanoscale reconstructions of brain cell nuclear envelopes by implicit and explicit parametric representations. *Computers & Graphics* 2019;doi:10.1016/j.cagx.2019.100004.
- [61] Agus, M, Gobetti, E, Pintore, G, Cali, C, Schneider, J. WISH: Efficient 3D biological shape classification through Willmore flow and spherical harmonics decomposition. In: *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 2020, p. 972–973. doi:10.1109/CVPRW50498.2020.00494.
- [62] Shen, L, Makedon, F. Spherical mapping for processing of 3D closed surfaces. *Image and Vision Computing* 2006;24(7):743–761. doi:10.1016/j.imavis.2006.01.011.
- [63] Kazhdan, M, Funkhouser, T, Rusinkiewicz, S. Rotation invariant spherical harmonic representation of 3D shape descriptors. In: *Proc. EG Symp. Geometry Processing*. 2003, p. 156–164.
- [64] Coggan, JS, Cali, C, Keller, D, Agus, M, Boges, D, Abdellah, M, et al. A process for digitizing and simulating biologically realistic oligocellular networks demonstrated for the neuro-glio-vascular ensemble. *Frontiers in Neuroscience* 2018;12. doi:10.3389/fnins.2018.00664.
- [65] Lorensen, WE, Cline, HE. Marching cubes: A high resolution 3d surface construction algorithm. In: *Proc. ACM SIGGRAPH*. 1987, p. 163–169. doi:10.1145/37401.37422.
- [66] Gottschalk, S. Collision queries using oriented bounding boxes. PhD Thesis, University of North Carolina at Chapel Hill; 2000.
- [67] Blu, T, Thévenaz, P, Unser, M. Linear interpolation revitalized. *IEEE Transactions on Image Processing* 2004;13(5):710–719. doi:10.1109/TIP.2004.826093.
- [68] Kuhl, FP, Giardina, CR. Elliptic fourier features of a closed contour. *Computer Graphics and Image Processing* 1982;18(3):236–258. doi:10.1016/0146-664X(82)90034-X.
- [69] Vallet, B, Lévy, B. Spectral geometry processing with manifold harmonics. *Computer Graphics Forum* 2008;27(2):251–260. doi:10.1111/j.1467-8659.2008.01122.x.
- [70] Farin, G. *Curves and Surfaces for CAGD—A Practical Guide*. 5th ed.; Morgan Kaufman; 2001. ISBN 978-1-558-60737-8.
- [71] Crane, K, Pinkall, U, Schröder, P. Robust fairing via conformal curvature flow. *ACM Transactions on Graphics* 2013;32(4):61:1–61:10. doi:10.1145/2461912.2461986.
- [72] Golum, GH, Van Loan, CF. *Matrix Computations*. 4th ed.; The Johns Hopkins University Press; 2013. ISBN 978-1-4214-0859-0.
- [73] Campello, RJ, Moulavi, D, Sander, J. Density-based clustering based on hierarchical density estimates. In: *Pacific-Asia Conf. on Knowledge Discovery and Data Mining (PAKDD)*. 2013, p. 160–172. doi:10.1007/978-3-642-37456-2_14.
- [74] Arthur, D, Vassilvitskii, S. k-means++: The advantages of careful seeding. In: *Proc. ACM/SIAM Symp. on Discrete Algorithms (SODA)*. 2007, p. 1027–1035. URL: <http://ilpubs.stanford.edu:8090/778/1/2006-13.pdf>.
- [75] Linde, Y, Buzo, A, Gray, R. An algorithm for vector quantizer design. *IEEE Transactions on Communications* 1980;28(1):84–95. doi:10.1109/TCOM.1980.1094577.
- [76] Lloyd, S. Least squares quantization in PCM. *IEEE Transactions on Information Theory* 1982;28(2):129–137. doi:10.1109/TIT.1982.1056489.
- [77] Guennebaud, G, Jacob, B, Avery, P, Bachrach, A, Barthelemy, S, et al., . *Eigen*. <http://eigen.tuxfamily.org>; 2010.
- [78] Zheng, Y, Guo, B, Yan, Y, He, W. O2O method for fast 2D shape retrieval. *IEEE Transactions on Image Processing* 2019;28(11):5366–5378. doi:10.1109/TIP.2019.2919195.
- [79] Abbas, S, Farhan, S, Fahiem, MA, Tauseef, H. Efficient shape classification using zernike moments and geometrical features on mpeg-7 dataset. *Advances in Electrical and Computer Engineering* 2019;19(1):45–51. doi:10.4316/AECE.2019.01006.
- [80] Wang, B, Gao, Y. Hierarchical string cuts: A translation, rotation, scale, and mirror invariant descriptor for fast shape retrieval. *IEEE Transactions on Image Processing* 2014;23(9):4101–4111. doi:10.1109/TIP.2014.2343457.
- [81] Zheng, Y, Meng, F, Liu, J, Guo, B, Song, Y, Zhang, X, et al. Fourier transform to group feature on generated coarser contours for fast 2d shape matching. *IEEE Access* 2020;8:90141–90152.
- [82] Escalera, S, Fornés, A, Pujol, O, Radeva, P, Sánchez, G, Lladós, J. Blurred shape model for binary and grey-level symbol recognition. *Pattern Recognition Letters* 2009;30(15):1424–1433. doi:10.1016/j.patrec.2009.08.001.
- [83] Shekar, B, Pilar, B. Shape representation and classification through pattern spectrum and local binary pattern—a decision level fusion approach. In: *Proc. Int'l Conf. on Signal and Image Processing*. 2014, p. 218–224. doi:10.1109/ICSIP.2014.41.
- [84] Cali, C, Wawrzyniak, M, Becker, C, Maco, B, Cantoni, M, Jorstad, A, et al. The effects of aging on neuropil structure in mouse somatosensory cortex—A 3D electron microscopy analysis of layer I. *Plos One* 2018;doi:10.1371/journal.pone.0198131.
- [85] Agus, M, Boges, D, Gagnon, N, Magistretti, PJ, Hadwiger, M, Cali, C. GLAM: Glycogen-derived Lactate Absorption Map for visual analysis of dense and sparse surface reconstructions of rodent brain structures on desktop systems and virtual environments. *Computers & Graphics* 2018;74:85–98. doi:10.1016/j.cag.2018.04.007.
- [86] Agus, M, Cali, C, Al-Awami, A, Gobetti, E, Magistretti, P, Hadwiger, M. Interactive volumetric visual analysis of glycogen-derived energy absorption in nanometric brain structures. *Computer Graphics Forum* 2019;38(3):427–439. doi:10.1111/cgf.13700.
- [87] Willmore, T. A survey on Willmore immersions. In: *Geometry and Topology of Submanifolds IV*. 1992, p. 11–16. doi:10.1142/1723.
- [88] Tomassy, GS, Berger, DR, Chen, HH, Kasthuri, N, Hayworth, KJ, Vercelli, A, et al. Distinct profiles of myelin distribution along single axons of pyramidal neurons in the neocortex. *Science* 2014;344(6181):319–324.
- [89] Sedlmair, M, Tatu, A, Munzner, T, Tory, M. A taxonomy of visual cluster separation factors. *Computer Graphics Forum* 2012;31(3pt4):1335–1344. doi:10.1111/j.1467-8659.2012.03125.x.
- [90] Vallet, B, Lévy, B. Spectral geometry processing with manifold harmonics. *Computer Graphics Forum* 2008;27(2):251–260. doi:10.1111/j.1467-8659.2008.01122.x.