

Modelling natural objects with α -complexes

Bart H.M. Gerritsen^{a,*}, Klaas van der Werff^b, Remco C. Veltkamp^c

^aTNO Netherlands Institute of Applied Geoscience, PO Box 6012, NL-2600 JA Delft, The Netherlands

^bDepartment of Design, Engineering and Production, Delft University of Technology, The Netherlands

^cDepartment of Computer Science, Utrecht University, The Netherlands

Abstract

The shape of natural objects can be so complicated that only a sampling point set can accurately represent them. Analytic descriptions are too complicated or impossible. Natural objects can be vague and rough with many holes. For this kind of modelling, α -complexes offer advantages over triangulations and hulls at only little extra cost. Geometric and topological descriptions are well formalised, with the flexibility to capture holes and separations. Careful layout of the input sampling point set and the attachment of weights make ‘special modelling effects’ possible. Our approach assumes modelling-by-example: start with a sampling point set from a physical example and take the modelling from there. We explore in this paper the merits of geometric modelling with α -complexes, with emphasis on practical values. We discuss the α -complex as a model description and as a representation scheme. We also show how to run FEM computations on α -complexes. © 2002 Elsevier Science Ltd. All rights reserved.

Keywords: CAD; Geometric modelling; Alpha shapes; Representation scheme

1. Introduction

Over the past couple of decades, engineers and designers have come to understand how to unambiguously describe physical objects in terms of boundaries and volumes. Today, using predefined primitives, sweeps, extrusion and part-whole descriptions, fairly complex objects and assemblies can be created. Practical implementations exploiting these capabilities, however, target primarily *engineering* objects and assemblies, with regular mathematical properties. As soon as *natural* objects are to be modelled, many such practical approaches fall short. The more so if modelled objects of this type are to be submitted to some numerical analysis tool, e.g. for stress or fluid flow analysis. Natural objects may become too complicated to find an analytical description: a sampling data point set is often the only feasible form of representation. Moreover, such a point set often serves as a *seed* for the model and modelling from scratch is seen as the exception rather than the rule. There is a large gap, however, between a point set representation and a solid model. Secondly, modelling the variety of holes and separations present in natural objects is complicated if the topology must explicitly be described. Scattered objects, such as fires and flames, clouds, immiscible fluids, eroded regions, sand depositions, etc. can be tedious to describe in a topo-

logical sense. Thirdly, many of today’s modelling approaches are incapable of handling topologies that change over time, which hinders the definition of ‘living’ natural object models, ruled by internal or external erosional or depositional processes, birth–life–death models, renewal processes, etc. Fractals can in part alleviate these difficulties, but many natural objects cannot be described by a single fractal and fractals are generally hard to constrain with observed data. These problems are not exclusive to the modelling of natural objects: similar problems are encountered in *motion planning*, *visibility analysis*, *percolation problems*, *handwriting recognition*, and a wealth of other geometry- and pattern-dominated problems in CAD/CAM.

In all of these fields, α -complexes have come up-front as an alternative geometric modelling approach. The notion of α -complexes grew out of the study of the shape of point sets, over the past two decades. Like a triangulation and a hull, it takes a point set on input and creates a polyhedral solid out of it. However, in contrast with triangulation and hull, the distance over which neighbouring points may connect can be limited by a single non-negative real parameter α . An α -complex is constructed of a collection of *k*-faces: in this case *k*-dimensional simplices, i.e. points, edges, triangles, tetrahedra, etc. An arbitrary α -complex need not be convex, or connected and it may have topological singularities (or: *singular faces*: see Table 1) and holes. By variation of α , a degree of coarseness can be

* Corresponding author.

E-mail address: b.gerritsen@nitg.tno.nl (B.H.M. Gerritsen).

Table 1

Classification of the k -faces $\{\sigma^{(k)}\}$ of an α -complex in terms of singular, regular and internal faces. In general, a k -face goes through the intervals as indicated in the table. Faces need not become singular, faces may become regular or interior straight away upon becoming first α -exposed. For α sufficiently large, all faces have become either regular or interior

Class	In $\partial\mathcal{C}_\alpha$	In Int \mathcal{C}_α	$\exists\sigma^{(k+1)} : \sigma^{(k)} \in \sigma^{(k+1)}$
Singular	✓		No
Regular	✓		Yes
Interior		✓	$k < d$

$[0, \alpha_0)$	$[\alpha_0, \alpha_1)$	$[\alpha_1, \alpha_2)$	$[\alpha_2, \infty)$
<small>not exposed</small>	<small>singular</small>	<small>regular</small>	<small>interior</small>
$\xrightarrow{\alpha}$			

inflicted on the resulting shape: a small value of α rules out all but the smallest faces, a higher value of α permits the inclusion of (almost) all faces also found in the corresponding triangulation of the same point set, see Section 2.3. Fig. 1 shows an example.

The current concept has a number of disadvantages, however. Most and for all, there is only one α for the entire α -complex preventing it from *locally* adapting α to the *local* spatial distribution of the point set. In Fig. 1, for example, no single α -value exists that fits the details of face, neck and shoulders simultaneously. To introduce local adaption, *weighted* α -complexes have been proposed [10]. A weighted point is a point to which a real-valued *weight* has been assigned. A weighted point covers a region, rather than a single spot. In a neighbourhood of points with higher weights, the α -complex tends to develop at lower α -values, whereas negative weights discourage the α -complex to develop. Local weights on top of a global α -value, indeed, allows for local adaption of α -complexes, but the design of an appropriate weight set is far from trivial.

With its *implicit* topology, holes and separations are naturally introduced by point and weight distributions. α -Complex modelling can be used for both *free-form shaping* (*design*, say) and *shape reconstruction*. The hypothesis is now that natural objects with their abundant complexity are typically not modelled using a mere free-form shaping, but using a combination of free-form shaping and shape reconstruction. Most likely, shaping will be feature- or knowledge-based, with geometry and topology resulting from variational geometry. α -Complexes are natural candidates for that purpose. Applications are

sparse so far, and little is known about the α -complex as a formal representation and its use in numerical modelling.

1.1. Goal and motivation

The goal of this paper is to evaluate the merits of modelling with α -complexes, with regard to their practical values for industrial designers and the engineering community. The central question will be as to whether modelling of certain classes of natural objects after initial reconstruction from an observed *physical* example becomes feasible with α -complexes, statically as well as dynamically. This question is inspired by the idea that the ability to (re)produce objects with a more natural ‘look-and-feel’ is of great benefit to the CAD/CAM community, industry, and consumers alike. Therefore, we bring existing developments on α -complexes together in a new framework, and add a new weighting scheme for natural objects. A formal evaluation of the α -complex as a *representation* and the capacity to run FEM computations over an α -complex are seen as critical indicators for conversions, migration paths, and embedding of this approach in existing modelling environments.

1.2. Previous work

Alpha shapes were introduced in Ref. [8] and further elaborated upon in Refs. [9,10,13]. Teichmann and Capps [29] replaced the ordinary Euclidean metric by a local and anisotropic metric, depending on the point configuration in some neighbourhood. They did not consider weighted α -complexes, although their adapted metric bears close relationship with a modification of weights and the (Laguerre) weighted distance which is discussed below. An alternate definition of the shape of a point set was recently proposed by Melkemi and Djebali [20], coined A-shapes, and also allowing for a locally different coarseness of the shape. Weighted distances and weighted Voronoi diagrams (or: *power diagrams*) have been described by Aurenhammer [1]. Lee discusses regular triangulations in Ref. [18]. Edelsbrunner (e.g. [10,11]) showed the duality between the geometric realisation of the nerve of an (in)homogeneous ball union and the regular triangulation based α -complex. Edelsbrunner further showed that a deformation retraction can be defined, that takes the ball union to the nerve, vice versa and ‘holes’ in the ball union cover are homotopy equivalent to holes in the corresponding α -complex. Recent developments in stochastic geometry can be taken from



Fig. 1. Example of an evolving zero-weight α -complex of a bust. Left: a small α -value, yielding a complex that is little more than just the point set with a few singular edges, developing into right: high α -value, yielding an overly fat complex obscuring all the fine details of the face and the next that almost evolved into the convex hull (data set by Silicon Graphics, 2780 points).

(e.g. Refs. [3,27]. The theory of coverage processes can be found for example in Hall [16] and the theory of random Voronoi tessellations have been presented by Møller [21]. A discussion of the Stienen model can be found in Ref. [27]. A status of clustering techniques and data partitionings in general was recently presented by Bock [5]. The mathematical foundation of representation schemes is mainly due to Ref. [24]. In more recent work, Ref. [17] further works out a number of details of various representation schemes. Applications of α -complexes in a geometric context were discussed by Mücke [22]. Edelsbrunner et al. [12] (molecular modelling), Gerritsen [14] (subsurface modelling) and Gerritsen et al. [15] (engineering and natural objects).

1.3. Organisation of this paper

From this section onwards, this paper is organised as follows. Section 2 presents the formal background of α -complexes. To that extent, it starts with the notion of a sampling point pattern, or *point process*. A mathematical formulation of the point process facilitates a theoretical model to be fit to an empirical (sampling data) point set. Two related subdivisions, induced by the same point process, appear to encode how the α -complex will develop with changing weight and α : the *weighted Voronoi diagram* and the corresponding *regular triangulation*. They will be introduced in Section 2 as well. Having discussed these geometric building blocks, they will be plugged into a *modelling framework* in Section 3. The framework brings together the empirical and modelled point process and shows the steps that lead to the estimation of the weights and the iterative improvement of the α -complex representation of an object. A *weighting strategy* based on physical properties will be introduced in Section 4. Then, a few modelling case studies will be discussed in Section 5. Eventually, the merits of modelling with α -complexes are evaluated in Section 6, followed by conclusions and suggestions for further research.

2. α -complexes

2.1. The Poisson point process

An α -complex is always based on (or: *generated by*) a point set. Sampling an object yields an empirical point set. It may be convenient to fit a theoretical model to such a point set, in similar vein with fitting a regression line to a point cloud or a Gaussian curve to a histogram. In stochastic geometry, such model point sets are usually referred to as *point processes*. Based on a point set $\mathbb{S} \subset E^d$, say, in some arbitrary dimension $d \geq 0$, the point process $\Phi_{\mathbb{S}}(\bar{x})$ defines some spatial relation, empirical or in terms of probabilities, among the members of \mathbb{S} , and among arbitrary points $\bar{x} \in E^d$ and members of \mathbb{S} . For example, \mathbb{S} may be a randomly distributed set of points in E^d , and the point process may formalise the distribution of distances among the members

of \mathbb{S} and between an arbitrary \bar{x} and the nearest-by member of \mathbb{S} . In the context of point processes, the set of points \mathbb{S} is usually called the set of *landmarks*, the role of which will be further explained in Section 3. In this paper, a set of points \mathbb{S} is always finite and compact. None of its members coalesces with another member and the set is not empty. \mathbb{V} is identical to set \mathbb{S} , but each of its members has now been assigned a real-valued weight $w \in \mathbb{W} \subset \mathbb{R}$. \mathbb{V} so becomes the ordered Cartesian product $\mathbb{V} = \mathbb{S} \otimes \mathbb{W}$ with $\mathbb{V} \subset E^d \times \mathbb{R}$. A weighted point $\bar{v}_j = (\bar{s}_j, w_j)$ is a 2-tuple that can be geometrically represented by a d -ball \mathbf{B}_j of radius $\sqrt{w_j}$, centred in \bar{s}_j . Its *region-of-dominance* is bounded by a $(d-1)$ -sphere $\mathbf{S}_j = \partial\mathbf{B}_j$, called the *weight-sphere*.

One branch of stochastic geometry is concerned with the study of point patterns *generating* triangulations and tessellations [3,16,21,27]. This paper addresses this type of problem applied to the input data set for weighted α -complexes. More precisely, the design of a weight set for a weighted α -complex, such that the level of geometric detail (or: resolution) can be controlled locally. A point process is called a *marked* point process (an *mpp* for short) if every point carries some value drawn from some *marker set* \mathbb{M} embedded in *marker space* M . Often, the marker represents some geometric shape (typically, a disc, a rectangle, a sphere, etc.) of which the size is proportional to the marker value. This is the case in this paper too: the marker value $w \in \mathbb{W}$, the *weight*, is represented by a (d) -ball of radius \sqrt{w} , giving the mpp $\Phi_{\mathbb{V}}(\bar{x})$. The weights are generally not independent.

Among other things, we are interested in the *coverage* of space by the union of the interiors of these balls. Weight cannot be uncoupled from the underlying point process. As it turns out, weighting is basically a *coverage problem*. Such an mpp is called a *coverage process*, for obvious reasons. If only the total coverage or non-coverage of space is of interest, then coverage processes are commonly indicated as *Boolean models* or more generally *germ-grain models* [16,27]. A *Poisson point process* (a *ppp* for short) is a point process of which the points are spaced according to a Poisson-distribution with *intensity* λ , with λ defined as the number of events per unit of space. In 3D, for example, λ represents the number of points per unit volume. If sampling the process is invariant to shifting the sampling window, then the process is called *stationary* and if also invariant under rotation, the process is called *isotropic*. A point process that is stationary and isotopic is called a *homogeneous* process and *inhomogeneous* otherwise. Inhomogeneous ppps cannot be characterised by a scalar λ ; in that case λ is different everywhere in space: $\lambda = \lambda(\bar{x})$. Coverage processes can be studied as topological coverings, partitionings of space in non-overlapping partitions. See Ref. [23] for details. Voronoi tessellations and triangulations are such space partitionings.

2.2. The weighted Voronoi tessellation

In an ‘ordinary’ first-order, Euclidean nearest-site Voronoi




Fig. 2. Left: Laguerre distance $L(\bar{v}_1, \bar{x}) = \pi^2$ of \bar{x} with respect to (1)-sphere $S_{w_1}(\bar{v}_1) = S_{r^2}(\bar{v}_2)$ in E^2 . Right: similar configuration in E^3 ; \bar{x} has been assigned an orthogonal weight: $\bar{v}_2 = (\bar{x}, w_2) = (\bar{x}, \pi^2)$ intersects $\bar{v}_1 = (\bar{s}, w_1)$ exactly in \bar{p} and \bar{p}' and $L(\bar{v}_2, \bar{v}_1) = \langle \bar{x} - \bar{s}, \bar{x} - \bar{s} \rangle - (\pi^2 + r^2) = 0$.

diagram (e.g. [1,9]), the nearest-neighbour distance is based on Euclidean metric. In the presence of weights, Voronoi cells are defined by *weighted distances*. In this paper, weighted distance is always the Laguerre distance L . For $\bar{v}_i, \bar{v}_j \in \mathbb{V}$, with $\bar{v}_i = (\bar{s}_i, w_i)$ and $\bar{v}_j = (\bar{s}_j, w_j)$, $L(\bar{v}_i, \bar{v}_j)$, $\bar{s}_{i,j} \in \mathbb{S}, \bar{w}_{i,j} \in \mathbb{W}$, is defined as:

$$L(\bar{v}_i, \bar{v}_j) = \langle \bar{s}_i - \bar{s}_j, \bar{s}_i - \bar{s}_j \rangle - (w_i + w_j) \tag{1}$$

with $\langle \cdot, \cdot \rangle$ denoting the standard inner-product. Apparently, the Laguerre distance between two weighted points is equal to the squared Euclidean distance minus the sum of their weights, cf. Fig. 2. Spaces endowed with a Laguerre weighted distance are non-metric. Even if all weights are positive, some weighted distances may become negative. Moreover, if only one weight is chosen large enough, *all* weighted distances are negative.

A point that seeds a Voronoi cell is called a *nucleus*, to discriminate it from Voronoi vertices. In a Voronoi diagram generated by \mathbb{V} , all members of \mathbb{V} become a nucleus. The Voronoi cell defined by nucleus $\bar{v}_j = (\bar{s}_j, w_j)$ is composed of all those points \bar{x} for which \bar{v}_j is the nearest nucleus, in terms of weighted distance, see Fig. 3, left picture. A Voronoi diagram is built of k -faces: k -dimensional elements such as vertices, edges, etc. up to d -faces (or: d -cells). A k -face is denoted by $\sigma^{(k)}$. The bounding $(d - 1)$ -face separating one Voronoi cell from an adjacent cell is contained in the *radical plane* of the two nuclei and

Voronoi vertices are at the intersection points of radical planes, called the *radical centres*. A radical plane H_{ij}^α of two weighted points \bar{v}_i and \bar{v}_j is a hyperplane consisting of points \bar{x} at equal weighted distances to both weighted points: $L(\bar{x}, \bar{v}_i) = L(\bar{x}, \bar{v}_j)$, for all points \bar{x} in H_{ij}^α . Assume that nuclei have unequal weights but further grow at equal rate, e.g. proportional to α . Then radical centres are the last points of the Voronoi cell to be swept out by the growing nuclei (see Fig. 3). The associated growing nuclei reach their radical centre simultaneously.

Voronoi tessellations are often studied in combination with ball unions, as they are in this paper. After all, a Voronoi tessellation can be constructed by a growing ball union $\mathcal{B} = \cup \mathbf{B}_j$ centred in the nuclei \bar{v}_j from which all the overlap has been removed and alternately, by a ball union \mathcal{A} of α -balls in the radical centres. Refer to Fig. 3. In this paper, the focus will be on \mathcal{B} . What relates a family of growing balls to a corresponding family of α -exposed faces is the notion of a (geometric realisation of a) nerve $\mathcal{N}(\mathbb{V})$. Defining the balls as subsets of \mathbb{E}^d , $\mathcal{N}(\mathbb{V})$ contains a k -face connecting $k + 1$ weighted points if the $k + 1$ subsets have a common intersection. In other words, as soon as three balls intersect in a common point they become connected by a triangle, etc. The event of two balls touching, so forming an edge in the nerve happens to coincide with the event of that edge becoming α -exposed (dashed spheres in right picture Fig. 3). Edelsbrunner [11] showed this duality and also showed that a deformation–retraction can be formulated that takes the family of ball unions onto the α -family. Observe that *any* pair of intersecting balls also forms a Voronoi edge, as any pair of balls (intersecting or not) has a radical plane.

This duality relationship will be exploited in the weighting strategy described in Section 4. It reduces introducing or removing vacancies from α -complexes to a coverage-by-weight problem, also known as the *space filling problem* [10]. Changing weight *independently* brings about local changes in the ball of the nucleus and in the radical planes of the nucleus with its neighbours and, therefore, in the Voronoi cells of the nucleus and its neighbours. If the ball radius changes, then, consequently, the α -exposedness of the corresponding edge in the triangulation changes.




Fig. 3. Relationship between weighted Voronoi diagram (solid), regular triangulation (dashed) and α -exposedness. Left: $L(\bar{x}, \bar{v})$ to the closest nucleus, equal to the squared radius of the dashed sphere. Centre coverage \mathcal{B} by grey alpha-plus-weight balls on the vertices of \mathbb{V} and the associated coverage \mathcal{A} by dotted α -balls in the radical centres. Right: local radical α -balls $A_{\bar{x}}^\alpha$.

2.3. Regular triangulation

Associated to the weighted Voronoi diagram is the *regular triangulation*. A regular triangulation of a point process is determined by two criteria: firstly, *regularity* of each of its constituting k -faces and secondly, *orthogonality*. In fact, these two criteria come together in the following: for each k -face in a regular triangulation there exists an empty $(d - 1)$ -sphere, that orthogonally intersects the weight spheres of all vertices of the face, so that all other vertices not belonging to that face have non-negative weighted distance to each of the vertices of the face (regularity). Notice the conjunction of the orthogonality criterion with the α -exposedness: once α is high enough to erect such a sphere, the face mentioned above becomes α -exposed. For $\alpha \rightarrow \infty$, all such spheres can be created, and $\mathcal{N}(\mathbb{V})$ grows into triangulation $\mathcal{T}(\mathbb{V})$ of \mathbb{V} . Also, notice that by definition, such spheres must be located in the radical centres to obtain orthogonal intersection.

Delaunay triangulations can be regarded as zero-weight regular triangulations, which is easily understood if one realises that if weight decays to zero, the local Delaunay criterion of $d + 1$ vertices lying on a common $(d - 1)$ -sphere with no point lying in its interior, complies precisely to the regularity and orthogonality criterion. Fig. 4, right picture demonstrates this.

2.4. α -Filtration

The fact that the faces of the α -complex are also found in the triangulation of the same point set, suggests that an α -complex can be obtained from that triangulation by *filtration*, and, indeed, such approaches have been proposed [10,13,22]. An α -complex $\mathcal{C}_\alpha(\mathbb{S}) \subseteq \mathcal{C}$ is a sub-complex of the simplicial complex $\mathcal{C} \cong \mathcal{T}(\mathbb{S})$ formed by a triangulation \mathcal{T} of the point set $\mathbb{S} \subset \mathbb{E}^d$. \mathcal{C}_α is not just any sub-complex drawn from \mathcal{C} . Rather, \mathcal{C}_α can only result from an α -filtration: let $\mathbb{F} = \{\mathbf{F}^{(k)}\}$ be the set of k -faces (in this case: k -simplices) of $\mathcal{C} \cong \mathcal{T}$, with $0 \leq k \leq d$. Then, the faces $\mathbb{F}_\alpha = \{\mathbf{F}^{(k)} \mid \mathbf{F}^{(k)} \in \mathcal{C}_\alpha\}$ of \mathcal{C}_α are those faces of \mathbb{F} that are α -

exposed for a given value $\alpha \geq 0$. All α -exposed faces pass the filtration, those that are not α -exposed do not. Firstly, observe that apparently $\mathbb{F}_\alpha \subseteq \mathbb{F}$, in other words, the α -complex is completely erected by faces also found in the triangulation. Secondly, a more concise definition is needed of α -exposedness. Assume an α that monotonically increases from zero towards infinity. A face \mathbf{F} becomes α -exposed for some α -value if an empty $(d - 1)$ sphere exists of radius $\sqrt{\alpha}$, that intersects all weight-spheres \mathbf{S}_j of vertices \bar{v}_j of \mathbf{F} orthogonally (e.g. [4,7,11,19]). Notice that such an α -sphere is *not* located in a point $\bar{s} \in \mathbb{S}$. A *weighted α -complex* $\mathcal{C}_\alpha(\mathbb{V}) \subseteq \mathcal{C}$ is a sub-complex of the simplicial complex $\mathcal{C} \cong \mathcal{T}(\mathbb{V})$ formed by a regular triangulation \mathcal{T} of \mathbb{V} [18]. The spatial occupancy by an α -complex is called an α -shape and an α -complex is, in fact, a triangulated α -shape. If the partitioning of the underlying space is irrelevant, α -shape can also be read where α -complex is written. Varying α defines a finite ordered family of α -complexes, called α -family \mathbb{A} . The point set is the lower extreme member of that family, the triangulation the upper extreme α -complex and the corresponding convex hull the equivalent α -shape.

3. Modelling framework

3.1. Modelling steps

Modelling with α -complexes is a multi-step process. The first step is always the collection of a set of observed landmarks \mathbb{S} that sample the object to be modelled. Fig. 5 shows this set of observed landmarks at the top-right. With natural objects in mind, such points will typically result from some sampling measurement (MRI-scan, seismic, tomography, etc.) on a physical example. A proper sampling representation contains at least all *anatomic landmarks*: characteristic points without which, salient features of the object go undetected. *Pseudo-landmarks* may be added by modellers to enforce certain geometric details, or to imply topological constraints. The sampling data point set defines an *empirical point process* $\Phi_{\mathbb{S}}(\bar{x})$, to which a *model point process* $\hat{\Phi}_{\mathbb{S}}(\bar{x})$ can be fit. In the process of sampling physical objects, it is convenient to discriminate landmarks from the *properties* measured in the landmarks. Landmarks form the point process, value sets are input to the markers of the corresponding marked point process. A property set measured in the landmarks will be referred to as a *value set* $\mathbb{F}_{\mathbb{S}}$. Multiple such sets may exist for a single point process. Value sets, thus, play an important role in certain weighting strategies, to be discussed in Section 4.

Important is the *organisation* of the sample points, more in particular, the *randomness* of the spatial distribution, see Fig. 6. Many sample methods simply generate data on some *lattice*, e.g. the regular grid underlying a digital picture or




Fig. 4. Left: radical planes H^{δ} of three weighted points in \mathbb{E}^2 that meet in a radical centre C_{ijk}^{δ} . The real ball B_{α}^{δ} (dashed) is empty and the dashed triangle will be contained in the triangulation. Right: regular triangulation of three vertices decaying into a zero-weight Delaunay triangulation (for clarity, intermediate radical planes and α -balls are not shown).




Fig. 5. The sequence of basic steps involved in modelling with α -complexes, assembled in a modelling framework with empirical branch at the right and theoretical branch at the left. MoM refers to some measure-of-merit for selection of the most appropriate α -complex.

stacked seismic data and the slices of a CT scan, see for example Refs. [2,6,25,26]. Filtering these lattices *before-hand* leads to a situation comparable to irregularly spaced data, even though some regularity is being preserved. If no filtering takes place and all points on the lattice are included as sample points, \mathcal{S} and $\Phi_{\mathcal{S}}(\vec{x})$ inherit the regularity of the lattice. Perfectly regularly distributed points are uniformly distributed, though, but *not random* and events so forming a point process are *not independent*. If weights are chosen equal too, then the α -family will have only few members. The faces of the triangulation enter the α -complex at just a few discrete values of α , many at a time. It is clear that in that case the *variation* must come from unequal weighting. Fig. 6 depicts this.

Sometimes, the use of transformations on \mathcal{S} is desirable, to ‘precondition’ the landmarks. Preconditioning is typically aiming at an improvement of the spatial distribution of the points. Affine transformations T^* can be split in a translation $T_{\vec{a}}$ and a linear part T . Let T be the matrix associated with the transform T . Furthermore, we insert the coordinates of the n points $s_j \in \mathcal{S}$ into the rows of a *con-*

figuration matrix $X_{\mathcal{S}}$ of dimension $n \times d$. For an affine transformation, we may so write:

$$T^* = T_{\vec{a}} + X'_{\mathcal{S}}T \tag{2}$$

where T is the $d \times d$ transformation matrix. The Lebesgue measure can be understood as a generalisation of the length to general dimension, and the Lebesgue measure $v_{\mathcal{L}}^d$ can loosely be seen as a d -volume. The translation has no impact, neither on the development of the α -complex, nor on its Lebesgue measure. On the contrary, the linear part has generally impact on both. Furthermore, the effect of $-T$ is regarded similar to the effect of T in this context. Assume that $\ker T = \{0\}$. The impact, in the sense of the *magnification factor* ω_T on $v_{\mathcal{L}}^d$ has an upper bound given by the norm of the linear transformation $\|T\|$. More specifically, with $v_{\mathcal{L}}^d(T(\mathcal{S})) = \omega_T v_{\mathcal{L}}^d(\mathcal{S})$, one has that:

$$\omega_T \leq \|T\| = \max \left\{ \frac{T(\vec{s})}{\|\vec{s}\|} \right\} \tag{3}$$

After diagonalisation of matrix T , we find the eigenvalues




Fig. 6. Schematic view on the relationship between sampling data point set organisation, geometry and numerical model geometry. Grey-ed boxes show the path followed with α -complexes.

in $\text{diag}\mathbf{T}$ showing the scaling along specific coordinate directions and $\omega_{\mathbf{T}}$ is equal to their product, i.e. to $\det \mathbf{T}$. For the trivial cases of the identity transform ($\mathbf{T} = \mathbf{I}$) or an orthogonal transformation such as a rotation ($\mathbf{T}^t\mathbf{T} = \mathbf{I}$), $\omega_{\mathbf{T}} = 1$ and the Lebesgue measure remains invariant. In the case of *isotropic* scaling, $\mathbf{T} = \mu\mathbf{I}$, for some $\mu \in \mathbb{R}$, $0 \leq \mu < \infty$ we have that $\omega_{\mathbf{T}} = \mu^d$.

In practice, it may be convenient to span locations of the landmarks and value sets in a common hyperspace, particularly if strong clustering effects are expected in the property values. Normalisation steps (Fig. 5) can bring property values on an equal scale first, so as to avoid extreme spatial distortion. A single multi-dimensional α -complex can capture these clusters (to be explained in Section 4) along with their landmark locations. Projections and intersections can be used in the preliminary analysis of the shape of the object. An example will be given in Section 5. In practice, the dimensions of hyperspace are restricted, with an asymptotic run time and storage complexity for the creation of \mathcal{C}_α of $\mathcal{O}(n^{\lfloor d/2 \rfloor})$ and $\mathcal{O}(n^{\lfloor (d+1)/2 \rfloor})$, respectively. For clarity, hyper-spatial modelling is not explicitly indicated in Fig. 5, but all it takes is a merge of landmark set and value set.

The next step to undertake is the design of a proper *weight set* $\mathbb{W} \subset \mathbb{R}$. There is much to say to this often decisive step, and we postpone to Section 4 the discussion as to how to derive the weight values. Basically, after preparing \mathbb{S} and \mathbb{W} , the ordered Cartesian product $\mathbb{V} = \mathbb{S} \otimes \mathbb{W}$ is triangulated. Then, the α -family is determined. Triangulation and determining the α -family take little human input or intervention. The modelling ingenuity is mostly in the prepara-

tion of point set and weights. Visual inspection, or some *object and/or costs function* must reveal how well the current α -complex models the target object. Sometimes, this applies to geometric and/or topological features, sometimes to the cluster analysis or classification in property space, sometimes to physical constraints, such as volume, inertia or distortion of the gravity field. A numerical and crisp *measure of merits* is not always available, however, and a goodness-of-fit may be subjective. A final α -complex may be converted into alternate representations, as desired.

3.2. Updating the model

If further modification is needed, this may take place along two ways.

1. Modification of the weights.
2. Insertion of additional anatomic or pseudo-landmarks.

Minor (local) modification can commonly be inflicted by weight modification, if failing, new landmarks have to be inserted. A theoretical model point process is in support of both of these alternatives. For instance, by providing probability density functions and expected values. Interesting theoretical distributions and values that the point process $\Phi_{\mathbb{V}}$ may deliver are: $\mathbf{g}(\hat{\alpha}^{\mathbb{V}})$ for exposedness, expected coverage $\mathbf{E}[\hat{\kappa}]$ and the probability density function $\mathbf{g}(\hat{\mathcal{L}}(\vec{v}, \mathbf{NN}^{(k)}(\vec{v})))$ for the theoretical Euclidean and Laguerre distances to the k -th nearest neighbour (to be discussed in Section 4). The empirical *local coverage* κ in combination with the *nearest-neighbour* and *local-furthest neighbour*

graphs point out the trouble spots. More on this in the next section.

3.3. Geometrically coupled objects

In nature, many *geometrically coupled objects* or *complementary objects* are found. Sometimes, one is interested in one object, sometimes in the hole space or void space it captures, sometimes in the combined system of objects. Let $\mathcal{W} = |\mathcal{C}_\alpha|$ be the underlying space of an α -complex represented object. The total void space $\neg \mathcal{W}$ of one shape \mathcal{W} defines the coupled (dual) space $\bar{\mathcal{W}}$, v.v. Here, \neg denotes the hole space relative to the convex hull \mathcal{H} . The geometry duality of geometrically coupled objects is preserved under simplicial maps. Examples are fluid-saturated pore space, composite material and immiscible multi-phase fluid system. Modelling with α -complexes always yields an unambiguous definition of both \mathcal{W} and $\bar{\mathcal{W}}$.

3.4. Towards living objects

Also, the recording and replaying of ‘living’ objects becomes increasingly important. A suitable approach to accomplish this is by means of a continuous or finite-state *evolutionary model*. Time-dependent landmarks and value sets make the α -complex dynamic. The topology is implicitly dynamic. Some landmarks may even ‘die’ and new ones may be ‘born’. In practice, creating, transforming and destructing in- and external processes rule the evolution. In mathematical morphology, the set operations *erosion* (erosion by an eroder object) and *dilation* (growth by accretion) are particularly well defined and provide formalisations for these processes [3,27]. Transformations were discussed earlier. α -complexes, being based on point sets, lend themselves particularly well for evolutionary models subject to point set operations. We follow closely the approach given in Ref. [28]. Let $\{\mathbb{S}_{t_0}, \mathbb{S}_{t_0+\Delta t}, \mathbb{S}_{t_0+2\Delta t}, \dots, \mathbb{S}_\infty\}$ denote a time discrete *evolution* of a natural object represented by set \mathbb{S} . For ease of notation, take $t_0 = 0$, and let index j denote the state $j\Delta t$ units later. Time state \mathbb{S}_j is mapped into \mathbb{S}_{j+1} by map \mathbf{G} . In practice, map \mathbf{G} will often be *implicit*, i.e. for $\mathbf{G} : \mathbb{S}_{j+1} \times \mathbb{S}_j \mapsto \mathbb{S}_{j+1}$, we have that $\mathbb{S}_{j+1} = \mathbf{G}(\mathbb{S}_{j+1}, \mathbb{S}_j)$. Map \mathbf{G} will be represented by a matrix \mathbf{G}_j and constraints can be plugged into this system. Implicit forms are typically solved using a backward Euler or Runge–Kutta method. An *evolutionary map* Γ^n is an n -fold product sequence of \mathbf{G} , as follows: $\Gamma^n = \mathbf{G}_0 \circ \mathbf{G}_1 \circ \mathbf{G}_2 \cdots \mathbf{G}_n$. Assume $\|\mathbf{G}\| \leq |\gamma|$. Then, trivially, $\|\Gamma^n\| \leq |\gamma|^n$.

4. Weighting strategy

With *weighting strategy*, we refer to the process of obtaining a weight set \mathbb{W} such that for some $\alpha \in [0, \infty)$ the resulting α -complex is fulfilling all the geometric and topological constraints imposed upon the object’s model,

see Fig. 5; but what is a proper strategy to design weight sets accomplishing this? Of course, the answer to this question depends on the *problem* on the one hand and the *modelling objectives* on the other hand. In Ref. [15], a couple of strategies have been discussed. Here, weights will be derived from measured physical property values; e.g. the water content per unit volume when modelling a cloud or the spectral content of the emitted light when modelling a comet. Typically, *inhomogeneously distributed* weight induces the geometry and the topology of the shape, according to the measured physical properties $\mathbb{F}_\mathbb{S}$ in the landmarks \mathbb{S} (Fig. 6). This strategy forms a natural link between the observations and measurement of a *physical example* and the optimal fit of some *theoretical model*. Often with real-live objects, the *exact* or *most likely* shape depends on the physical property being measured and *how* measuring took place. An echo, an infrared image, a PET scan, they easily yield different geometries of the same object. We refer to such shapes as *property-ruled* shapes. Many natural objects ‘extracted’ from observed data are, in fact, property-ruled. See Section 5.1 for an example.

4.1. Neighbourhoods and nearest neighbours

The start of the weighting step is always a data analysis of the value sets $\mathbb{F}_\mathbb{S}$, generally one value set per property, see Section 3. To study different regions, some sense of neighbourhood must be defined and the definition must be convenient in the context of the regular triangulation. Recall that an α -complex is filtered off from a triangulation. The *closed star* $\text{St } \vec{v}$ of a vertex \vec{v} is the subset of the triangulation \mathcal{T} of all closed k -faces incident upon \vec{v} and all their sub-faces. Then, for a neighbourhood:

Definition 1 (neighbourhood). A *neighbourhood* $\mathcal{N}(\vec{v})$ is the space covered by the closed star $\text{St } \vec{v}$ of \vec{v} ●

Among the faces in this $\mathcal{N}(\vec{v})$, at least one 0-face can be denoted as the *nearest neighbour* $\text{NN}(\vec{v})$ and at least one 0-face as the *local furthest neighbour* $\text{FN}(\vec{v})$, closest-by and furthest-apart within $\mathcal{N}(\vec{v})$, respectively. Nearest- and local-furthest neighbours are not unique. The k -th nearest neighbour is the 0-face in $\mathbb{V} \cap \text{St } \vec{v}$ that is the k -th closest to \vec{v} . $\mathcal{N}(\vec{v})$ is convex and $\text{NN}(\vec{v})$ and $\text{FN}(\vec{v})$ are always contained in the boundary $\partial \mathcal{N}(\vec{v})$. Each $\mathcal{N}(\vec{v})$ has one or more associated radical centres. A radical centre is *not* necessarily located in $\text{Int } \mathcal{N}(\vec{v})$; it can also be contained in the $\text{Bd } \mathcal{N}(\vec{v})$ or the $\text{Ext } \mathcal{N}(\vec{v})$, which is typically the case with triangles with a ‘bad aspect-ratio’ (slivers, say), often occurring in neighbourhoods close to or at the boundary (see Fig. 4, left picture).

For the neighbourhood analysis underpinning the process of weight estimation, it is a convenient step to determine the nearest neighbour graph $\mathcal{NNG}(\mathbb{S})$ and the local furthest neighbour graph $\mathcal{LFNG}(\mathbb{S})$. $\mathcal{NNG}(\mathbb{S})$ and $\mathcal{LFNG}(\mathbb{S})$ reveal a first clue of problem spots for weighting. Another

indication of sparse neighbourhood in the ppp can be found by the classical *chi-square* test and a further confirmation can be attained by a tentative coverage-by-weight analysis. Assume that an initial weight set has been designed and attached to \mathbb{S} and assume Voronoi cells to be bounded, as appropriate, by convex hull \mathcal{H} . The local coverage $\kappa(\vec{s}_j)$ is defined as:

$$\kappa = \frac{v_{\mathcal{G}}^d(\mathbf{B}_{\epsilon}(\vec{s}_j) \cap V(\vec{s}_j))}{v_{\mathcal{G}}^d(V(\vec{s}_j))} \quad (4)$$

Further details will be discussed below. In practice, these steps usually point at the same ‘problem spots’.

Let m -vector $\vec{p}(\vec{s}) \in E^m$ be the vector of (scalar) properties measured in point \vec{s} . These property values may be organised in an $n \times m$ -property matrix $\mathbf{P}_{\mathbb{S}}$ in which the j -th row contains the property vector of \vec{s}_j . The ‘geometric’ space, represented by configuration matrix $\mathbf{X}_{\mathbb{S}}$, cf. Eq. (2), may now be augmented by the ‘property’ space, represented by $\mathbf{P}_{\mathbb{S}}$, in the following sense:

$$\mathbf{X}_{\mathbb{S}^*} = [\mathbf{X}_{\mathbb{S}} | \mathbf{P}_{\mathbb{S}}] \quad (5)$$

where $\mathbf{X}_{\mathbb{S}^*}$ denotes the augmented matrix resulting in a hyper-spatial representation discussed in Section 3 based on hyper-spatial landmark set \mathbb{S}^* .

Weighting based on physical properties can be regarded as a projection of m property coordinates onto l weight coordinates, i. e. $\mathbf{W} : E^m \mapsto R^l$. Basically, this projection is continuous. In this work, we always have $l = 1$, so that the set of weights \mathbb{W} can be represented by an n -vector \vec{w} :

$$\vec{w} = \mathbf{P}_{\mathbb{S}} \Lambda \quad (6)$$

with $\Lambda = [\mu_1, \mu_2, \dots, \mu_m]^t$, $\sum_{j=1}^m \mu_j = 1$, an $m \times 1$ coefficient matrix of the *relative contribution* μ_j of each of the properties to the weight. Not only these μ ’s are important, also the eventual *coverage* κ by weight markers. In Eq. (6), $\|\vec{w}\|$ depends on the magnitude of the properties, which is inconvenient. One solution is to normalise property values (Fig. 5), another solution is to normalise the resulting weight. To balance weight with distances and local coverage again, the normalised weight must be up-scaled using a weight transformation. Finding Λ and the weight transformation will be discussed in greater detail below.

4.2. Determining the relative contribution

Relative contribution vector Λ commonly results from a *discriminant* or *principal component* or *factor analysis*. The goal of such operations is generally to investigate or enhance *classifiability*, by which we understand the ability to subdivide data in clusters and classes. A recent overview on the status of cluster analysis and data classifications is given by Bock [5]. Clustering in this work is basically distance-based and clustering is, therefore, geometric clustering. We assume the ordinary Euclidean metric applicable to the property space E^m . More particularly, in this context, a *cluster* is defined as:

Definition 2 ((geometric realisation of a) cluster). Let $\Psi : E^m \mapsto R$, $\Psi(\vec{x}) = c$ be an implicit definition of closed d -volume in E^m . Then a (geometric realisation of a) cluster of geometric data values in a space E^m is defined as a finite compact region $C \subset E^m$ containing data that satisfy some $c_1 \leq \Psi(\vec{x}) \leq c_2$, with $c_1, c_2 \in R$ and $c_1 \leq c_2$. ●

The $\mathcal{NN}\mathcal{G}$ and $\mathcal{LFN}\mathcal{G}$ graph can also be extended over property space. Obviously, edges in such a nearest neighbour graph $\mathcal{NN}\mathcal{G}$ indicate pairs that tend to cluster earlier in a common cluster than pairs in the corresponding $\mathcal{LFN}\mathcal{G}$ as the mean length of edges in $\mathcal{LFN}\mathcal{G}$ is at least the mean length of edges in $\mathcal{NN}\mathcal{G}$. For a single property p_k , associate an $n \times n$ dissimilarity matrix Δ with property matrix $\mathbf{P}_{\mathbb{S}}$, for which:

$$\begin{cases} \text{(i)} & \delta_{ij} \geq 0 \text{ for } i \neq j \\ \text{(ii)} & \delta_{ji} = \delta_{ij} \text{ for all } 1 \leq i, j \leq n \\ \text{(iii)} & \delta_{ii} = 0 \text{ for all } 1 \leq i \leq n \end{cases} \quad (7)$$

A subdivision $\mathcal{Y} = \{Y_1, Y_2, \dots, Y_q\}$ of \mathbb{S}^* in q clusters is a sub-partitioning if $\forall j, 1 \leq j \leq q$:

$$\begin{cases} \text{(i)} & Y_j \subset \mathbb{S}^* \\ \text{(ii)} & Y_j \neq \emptyset \\ \text{(iii)} & Y_i \cap Y_j = \emptyset, i \neq j \end{cases} \quad (8)$$

If $\cup_j Y_j = \mathbb{S}^*$ then the sub-partitioning becomes a partitioning. Each partitioning according to this definition, is also a covering of \mathbb{S}^* .

In the general case, α -complex $\mathcal{C}_{\alpha}(\mathbb{S}^*)$ will have m interiors and l exteriors. Each separate sub-complex forms a *clique*. A *clique partitioning* is a partitioning in q clusters which minimises $\sum_{j=1}^q \sum_{k,l \in Y_j} \delta_{kl}$. In other words, a clique partitioning minimises the total dissimilarities among cluster members of the same cluster, for each cluster. Dissimilarity matrix Δ can be reorganised such that Δ_j is its $n_j \times m$ sub-matrix where n_j is the cardinality of clique Y_j . Sub-matrix Δ_j is minimised in the sense that $\sum_{k,l \in Y_j} \delta_{kl} = \min$. Now, let the geometric realisation of a clique partitioning be found by means of a general dimension α -complex in property space. Let $\alpha = 0$. Then, the clique partitioning found is the absolute minimum of 0 for $q = \text{card } \mathbb{S}$ clusters: each point a cluster. Each row in Δ is a clique sub-matrix Δ_j and all its columns are zero. Now, let α increase monotonically. We then have that if $\alpha_2 > \alpha_1 \Rightarrow q_2 \leq q_1$. Let $\vec{s}_l^* = \text{NN}(\vec{s}_k^*)$ at distance δ_{kl} . The points cease to exist as a separate clique as soon as α exceeds $\alpha = \delta_{kl}^2/4$. Let $E[\delta_{k,\text{NN}(k)}]$ and $E[\delta_{k,\text{FN}(k)}]$ denote the expected Euclidean distance (mean edge length) to the nearest and the local furthest neighbour of \vec{s}_k , respectively. Then, trivially, $E[\delta_{k,\text{NN}(k)}] \leq E[\delta_{k,\text{FN}(k)}]$. Furthermore, let $\vec{s}_k \in \mathcal{NN}\mathcal{G}$ and $\vec{s}_l \in \mathcal{LFN}\mathcal{G}$. Then, as a result, for any α and any clique Y_j , probability $\text{Pr}[\vec{s}_k \in Y_j] \geq \text{Pr}[\vec{s}_l \in Y_j]$, which re-establishes the indicative capacity of $\mathcal{NN}\mathcal{G}$ and $\mathcal{LFN}\mathcal{G}$.

4.3. Scaling the weight

Normalised weights still have no relation to the typical nearest-neighbour distances to abridge. A scaling (or: *weight transformation*) is needed to obtain sufficient coverage by weight to let a variation of α do the rest. In order not to destroy the distribution of Eq. (6), weight transformation must be equal for all points, i.e. an *isotropic scaling* $\varrho(\vec{v}) = c$. In that case, κ is not known in advance and can only be determined afterwards. For the modelling of natural objects, this kind of weight transformation suffices. Sometimes, however, there may be a priori knowledge that leads to the demand that (part of) the object must be hole-free. In that case, (part of) the result obtained after weight transformation so far may be further modified to remove any observed holes. Hole-free implies that κ must reach unity, by adjustment of individual weights, probably in combination with an increase of α . This, of course, violates isotropic scaling.

The first thing we are interested in is the weight distribution belonging to the moment that points in the α -complex are about to be connected by an edge. To understand when that happens, we introduce the *Stienen model*. The Stienen model (e.g. [27]) is a Voronoi tessellation and a ball union in which balls have maximum radius without overlapping. Weights of nuclei are such that their ball touches the nearest radical plane, i.e. the nearest boundary of their Voronoi cell. As a result, balls of the nuclei of two adjacent Voronoi cells touch at their common cell boundary if the two nuclei are nearest-neighbours. The nearest neighbour relationship is not symmetric and, as a result, not all balls actually touch their nearest neighbour ball, see Fig. 7. In the Stienen model, none of the radical centres can be covered and as a result, none of the cells can be completely swept out and all local coverage $\kappa_j < 1$. The geometric realisation of the nerve of the Stienen model is a sub-graph of $\mathcal{N}\mathcal{G}$. The Stienen model may serve as a guide to determine the weight of each of the individual points so that:

1. the weight has its local maximum value,
2. without enforcing any of the edges by weight only.

Honouring these two constraints, the Stienen model produces safe lower bound values for the weight set \mathbb{W} : all edges become α -exposed for some finite non-negative α , and all first edges in the star $St \vec{s}$ become α -exposed for minimum α . The latter does, of course, not imply that all edges become α -exposed for *the same* value of α . Increasing α causes the edges in the star to enter the α -complex one-by-one, until ultimately, the edge with the local-furthest neighbour enters.

Does the Stienen model have an equivalent that may serve as an upper bound for the weight? No such model could be found. As explained in Section 2.2, the radical centres (Voronoi vertices) are the last points of the Voronoi cell to be swept out by the growing balls in the nuclei. This

can be further made precise, in that the radical centre contained in the radical plane of a nucleus \vec{v} with its local furthest neighbour $FN(\vec{v})$ is the very last point being covered as α grows. If weight is increased such that the marker reaches this point, the entire Voronoi cell will be covered by the marker, i.e. $\kappa = 1$. This holds for each cell and full coverage will be obtained. Unfortunately, this strategy also introduces many redundant points. Also notice that if $\kappa = 1$ in each Voronoi cell, varying α is no longer relevant. In general, assigning weights to nuclei modifies the weighted Voronoi diagram, as it causes radical planes to move. See Fig. 4, right picture. In practice, the size of the weight marker touching the furthest Voronoi vertex (radical centre) is, therefore, multiplied by a factor $0 \leq \omega \leq 1$. Of course, ω should be chosen such that the weight is greater than the weight based on the Stienen model. In the absence of better terms, we to speak of the *nearest Stienen model* and the *local furthest Stienen model*.

If all Voronoi cells incident upon a radical centre have local coverage $\kappa = 1$, that radical centre cannot lie inside a hole. Since all balls grow in equal amounts as α changes, a hole shrinks so that the radical centre is always in its interior. On disappearance, the hole consists of *only* the radical centre. However, for an *arbitrary* radical centre to become covered, it is not necessary that *all* Voronoi cells incident upon a radical centre have $\kappa = 1$. A local coverage $\kappa = 1$ is only reached if also the radical centre with the *local-furthest* neighbour is covered, and not all radical centres are generated by a vertex contained in $\mathcal{L}\mathcal{F}\mathcal{N}\mathcal{G}$. Local coverage is generally not independent. In addition, the desired coverage is not necessarily obtained through weight alone. In practice, it is convenient to let some coverage be conquered by α . A *target* α should be chosen which is approximately of the magnitude of w . The scaling of weight relative to some *target* or *expected* α (Fig. 5) is best understood when their combined effect is examined. Call the combined effect ϵ ,




Fig. 7. Stienen model of a Poisson point process in E^2 . Two nuclei may touch at their common cell boundary iff the two are nearest-neighbours, but not all nearest neighbour balls do, due to asymmetry. Observe that none of the radical centres can be covered.




Fig. 8. Left: marker radius $r = \sqrt{\epsilon}$ versus α for various fixed weights w , depicted on a $^2\log$ -linear scale. Right: relative influence of α . The figures show the limited influence for $\alpha \ll w$ and $\alpha \gg w$. Observe the negative weight-curves and their local maximum for their derivative $dr/d\alpha$.



Fig. 9. α -Complex of the comet West. Left: as observed. Centre: nucleus (grey-scale weighting, with threshold: 0.4). Right: ion tail, obtained by extraction of the blue cluster points shown in Fig. 10.

and define $\epsilon(\alpha)$ as:

$$\epsilon(\alpha) = \begin{cases} \alpha + w, & \text{for } : w + \alpha \geq 0 \\ i^2(\alpha + w), & \text{for } : w + \alpha < 0 \end{cases} \quad (9)$$

In Eq. (9), ϵ is equal to the square of the radius r of the weight marker. Then for the derivative $dr/d\alpha$:

$$\frac{dr}{d\alpha} = \begin{cases} \frac{1}{2\sqrt{\alpha + w}}, & \text{for } : w + \alpha \geq 0 \\ \frac{-i}{2\sqrt{\alpha + w}}, & \text{for } : w + \alpha < 0 \end{cases} \quad (10)$$

Then clearly, the relative influence of α is small as long as alpha is small compared to w , i.e. $\alpha \ll w$, strongly increasing when $\alpha \approx w$ and α completely wipes out the influence of w when $\alpha \gg w$. See Fig. 8. The scaling of the weight q may be chosen such that $E[\alpha]$ and w are approximately on equal scale, in order to experience the effects of both. Local variation depends on the distribution of the ratio between nearest and local furthest neighbour distances and the distribution of the local coverage κ . Variation is equalised to some extent by varying local values for w .

5. Experimental cases

In Refs. [14,15], various engineering and natural objects have been presented. In this section, a few models of natural objects will be presented. Also, a few words will be spent on numerical modelling with α -complexes.

5.1. The comet West

A comet is an irregularly shaped natural object of frozen gas and rocky debris, orbiting around the sun. A comet has a focal, bright, approximately 10 km wide kernel, called the *nucleus*. When approaching the sun, a comet develops three tails: the bright *coma*, a large trailing cloud of diffuse material, the pale-blue *ion tail* of ionised plasma, and the yellowish *hydrogen envelope*, with hydrogens that escaped the comet’s gravity. The comet West¹ was observed by various observers during its bright appearance in 1976. See Fig. 9, left picture.

The model of Fig. 9 was created by hyper-spatial modelling, with property space spanned by the spectral content of the light emitted by the comet. In many aspects, this case uncovers the limits of this approach. Spectral contrast is very weak and clustering is bad; Fig. 10 depicts the normalised property space in the unit cube, looking down the grey-scale diagonal. Nonetheless, three ‘clusters’ are visible: a blue cluster (dark) representing the ion tail and the edges of the dust tail, a grey cluster along the main diagonal representing the dust tail and a rest cluster (nucleus, coma, dust tail). To demonstrate the modelling capacity, the coma was in part reconstructed from various pictures, using ‘paintbrushes’ to spray landmarks on cross sections. Fig. 11 shows the results.

¹ Named after the astronomer West, who first described his observation of the comet.




Fig. 10. α -Complex of the colour spectrum property space. X -axis represents red, Y -axis green and Z -axis blue. Clusters represent different parts of the comet. The dark cluster represents the blue ion tail. The α -value increases left-to-right. Separation is weak and clusters concentrate around the grey-diagonal.




Fig. 11. 3D α -complex of the comet West, reconstructed from cross sections, that were edited using Fig. 9 as a template. Dust and ion tail have been extended and modified (the claim that the model fits the observations must be abandoned).

5.2. Medical modelling

Organ imaging is an important aid in clinical diagnosis. α -complexes can be used in (multi-dimensional) modelling of organs, see for example Ref. [26]. Fig. 12 shows an α -complex of the left and right shoulder blade (*scapula*), capturing the bones fairly well, except for relatively few singular triangles at the *glenoid cavity* (not visible). Weighting (nearest neighbour Stienen model) was required to turn the complex into a model good enough to be fed into a finite element analysis (FEA) tool. In this case, no physical properties based value sets were used. Weighting has been targeted at a hole-free scapula. The nearest- and local furthest neighbour graphs (Fig. 13) immediately revealed the problem spots. In practice, weight and volume will also be important constraining factors.

5.3. Numerical modelling

An α -complex, in terms of numerical computing, is generally an *unstructured grid* and in practice, FEA will be the analysis method to apply. The problem with α -complexes emerges from the fact that an α -complex may contain topological features that cannot be coped with by state of the art FEA codes, such as singular faces and separations. In addition, α -complexes often contain ‘slivers’: sharp tetrahedra with a very odd aspect ratio. We developed a new approach to run FEA over α -complexes, that overcomes most of these problems. We turn tetrahedral of the α -complex model (‘foreground material’) into tetrahedral

FEM elements. This α -complex of the object is embedded in the hosting regular triangulation (‘background material’ or ‘bulk’). After applying boundary and/or initial conditions, computations can be started.

Our method exploits the fact that an α -complex is a sub-complex of the triangulation. A triangulation, which is a cellular decomposition *continuum*, can always be input to an FEA package, and *boundary conditions* can be attached to its boundary faces. Sometimes, FEA codes run test with respect to tetrahedral aspect ratios. Slivers will not pass such tests. Most slivers are developed when the α -shape approaches the convex hull, i.e. for the highest α -values. Therefore, rather than taking the bulk at the maximum α , we lower α somewhat to get rid of these slivers.




Fig. 12. α -Complex of the left and right shoulder blade (data set by courtesy of Delft University of Technology, Department of Electrical Engineering, 3434 points).




Fig. 13. Nearest neighbour (left) and local-furthest neighbour graph (right) of the shoulder blade.

The approach is outlined below, an application can be found in Ref. [14].

- Step 1:** Compute the α -family and select an embedding background α -complex close to the triangulation. Assign background material properties to the elements of the embedding α -complex.
- Step 2:** Compute the normals of all the triangles in the border of this complex, in order to find out how (directional) boundary conditions need to be attached.
- Step 3:** Assign foreground material properties to the object model α -complex to analyse and ‘inject’ this complex into the embedding background complex, i.e. for each tetrahedron in the foreground α -complex, locate the corresponding tetrahedron in the embedding background complex, and flip the background material properties into foreground properties.
- Step 4:** Run the FEA analysis and optimise the α -complex as appropriate.

6. Evaluation

The introduction of α -complex modelling in existing environments is far from trivial. In practice, three main obstacles can be identified: visualisation; the use of existing numerical tools; and conversions, back and forth. What they have in common is the need for data or model conversion. Existing visualisation tools often need modification before they can cope with singular faces. This point is not further discussed here. The coupling with numerical tools, more specifically FEM codes, have been discussed above. There remain the model conversions in a broader sense, required for coupling or transcription to enterprise-wide CAD/CAM systems. In practice, *homeomorphic* conversions, mapping model elements in a one-to-one fashion, offer the best guarantees against loss of geometric and topological information. Therefore, homeomorphic conversions will be the goal.

For computer modelling purposes, different approaches in describing the geometry and topology of modelled objects exist and the quality of conversions depends mostly on the

target solid model description. A model description based on an α -complex carries some characteristics of a:

- part–whole description, with primitive instancing of k -simplices;
- volumetric description, with a cellular complex built from d -simplices (possibly, after *regularisation*); and
- (faceted) boundary description, built from $(d - 1)$ faces (*facets*) of the α -complex.

After describing the evaluation criteria, α -complexes will be evaluated against these three model description classes.

Loosely formulated, a *representation scheme* is a way to describe real world objects in a symbolic notation. Widely used schemes are: constructive solid geometry (CSG); boundary representation (B-rep); cell decomposition; spatial enumeration; and voxel representation. In a sense, α -complexes can also be seen as a means of representation. Therefore, α -complexes will also be evaluated as a representation scheme.

Not much is known about α -complexes in this regard. Some reasoning is possible, however, by comparing the α -complex model description to triangulations, simplicial complexes and cellular models, and the representation scheme to CSG-, B-rep and cellular decomposition representations. Table 2 summarises the principal findings of this comparison. We believe that it gives sufficient indication with regard to the use of α -complex modelling with existing tools and environments.

6.1. Evaluation criteria

Solid object model descriptions can be qualified by a number of criteria. Apart from well-formedness, generality and completeness, the following criteria are considered.

Solidity: a convex polytope \mathcal{P} (e.g. a simplex), divides space into two regions: *interior* $\text{Int } \mathcal{P}$ and *exterior* $\text{Ext } \mathcal{P}$, separated by *boundary* $\text{Bp } \mathcal{P}$. Point location, for example, is feasible when this criterion is met.

Homogeneous dimensionality: $\text{Bd } \mathcal{C}_\alpha$ shall be fully incident upon the interior $\text{Int } \mathcal{C}_\alpha$ and shall be composed of regular faces only.

Rigidity: the geometry of the modelled object shall not

Table 2

α -complexes compared to B-rep and CSG and cell decomposition. Top part: α -complex as a model description, bottom part: α -complex as a representation scheme

	Faceted B-rep	CSG	Cell decomp.	α -Complex
Solidity	Yes	Yes	Yes	If regularised
Continuity	Yes	Yes	Yes	No, only part-by-part
Rigidity	Yes	Yes	Yes	Yes
Homo. dim.	Yes	If regularised	Yes	If regularised
Closure	Yes	If regularised	Yes	If regularised
Orientability	Yes	Yes	Yes	Yes
Disjunct Int	Yes	Yes	Yes	Yes
Finiteness	Yes	Yes	Yes	Yes
Domain	Rigid homogeneous Manifold Complex shapes Massive descriptions	Rigid homogeneous (Non)-manifold Moderate complexity Lean descriptions	Rigid homogeneous Manifold Complex shapes Massive description	Rigid homogeneous Non-manifold Very complex shapes Massive description, reproducible from \mathbb{V}
Validity	Euler rules, expensive	Regularised, set-theoretic	Validation expensive	Euler + α + weight expensive
Completeness	Homogeneous properties	Homogeneous properties	Homo. prop. per cell	Homo. prop. per face
Uniqueness	No	No	Yes	No

depend necessarily on its position, i.e. on its location or its orientation in space.

Continuity: the represented object shall not be composed of unconnected parts.

Closure: k -faces shall have no incidences to singular faces. For example, in the case of a triangulated object, every k -simplex shall be incident upon exactly $d + 1$ ($k - 1$)-simplicies. If so, k -faces can be represented as regularised sets (r -sets).

Disjunct interiors: the interior $\text{Int } F_i$ and $\text{Int } F_j$ of any two k -faces F_i and F_j shall be disjunct.

Orientability: all k -faces shall be orientable.

Finite time and storage complexity: supporting finite time complexity of operations and finite storage complexity of the results.

Criteria for representation schemes are the following.

Uniqueness: the uniqueness of a representation that describes the modelled object under this representation scheme.

Completeness: the richness of description, in support of operations, analysis and conversions.

Domain: the class of physical objects that can be represented and the class of valid representations that the scheme can produce.

Validity: the validity of the objects produced under this scheme.

These criteria will be applied to the evaluation of α -complexes below, both as a model description and a representation scheme.

6.2. Model description and representation

For α -complexes, with their internal voids and possibly

singular faces, the solidity criterion is generally not met. Provided that α is high enough to permit the forming of d -faces, removal of singular faces ensures homogeneous dimensionality and closure. The exterior can be composed of disjunct parts, all but one bounded and one unbounded. Continuity of the interior cannot be guaranteed, moreover, α -complexes can scatter into many disjunct parts. For finite complexes, finite storage complexity can be shown to exist, as well as finite time complexity for operations upon them. To a great extent, characteristics are as with cellular decomposition. Faces of an α -complex are orientable. Voids usually have a 'negative' orientation, for area, volume, etc. If $\text{Bd } \mathcal{C}_\alpha$ can be triangulated such that the boundary is closed and connected, then these two properties ensure well-formedness. However, generally, continuity is not guaranteed and boundary is not connected, only on a part-by-part basis. The rigidity criterion is obviously met: the α -complex is determined by distance versus α and weight, not by position or orientation. The interior of any two k -faces is always disjunct, as they are faces in the underlying triangulation, a non-overlapping covering. This also excludes self-intersection.

6.3. Representation

The first question concerns the domain: can every physical object be represented by a sampling data point set? For the class of objects discussed in this paper, the answer to this question is assumed yes. A sampling data point set does not uniquely represent a physical object: many such sets can sample the same object, and one sampling point set can represent multiple physical objects. Data may also be added subset by subset, like with sliced data.

Once a data point set has been sampled, the next question is, can every point set be triangulated? Generally,




Fig. 14. Hierarchy of holes related to the value of α . The position of a separation in the hierarchy may be vulnerable for dispute. To reflect this, it is set off from the rest.

homogeneous ppps are good enough to obtain a good-quality triangulation. It can be shown [13] that a unique relation exists between the sampling data point set, the α -complex and its underlying space. Variations in the α -value, even though constrained, may instantiate an α -family of models that cannot be thought of as equivalent classes, due to different topologies. α -complexes cannot be traced back uniquely to sampling data point sets, but they can be uniquely reproduced from them, which makes an *unevaluated* description feasible, thus reducing storage requirements. With regard to the completeness; the data structure storing the α -complex can be easily augmented with auxiliary parameters, such as material properties, central moments, etc. The simplex, the basic building block, allows for relatively simple computational schemes.

6.4. Conversions

Conversion mapping transforms an α -complex representation of an object onto another such representation. If *homeomorphic* (e.g. [23]), a bijective relation exists between source and target and conversion mapping is basically without loss of topological information. Multiple conversions may exist for a single α -complex and, depending on the type of the conversion, generally, multiple α -complexes may lead to the same conversion. Inserting massive amounts of cells in a CSG model is not very convenient, and unless regularised, singular k -faces will cause a CSG representation to fail on an arbitrary α -complex. For example, when determining the union of a cellular volume with a singular point. An α -complex is not necessarily convex. As a result, boundary representations fail to describe the d -volume of an arbitrary α -complex. The remedy to rule out singular faces is regularisation. The handling of holes will be addressed in greater detail next. First, topological relationships of holes will be examined. Next, a *nil-object* will be introduced to represent a hole in conversion, with the aim to achieve homeomorphism in conversion.

6.5. The representation of holes

In the discussion of holes, it is often more convenient to consider α -shape \mathcal{W}_α rather than the corresponding α -complex \mathcal{C}_α . An α -shape differs from the convex hull by the total amount of space occupied by the various holes, i.e. $\mathcal{W}_\alpha = \mathcal{H} - \neg \mathcal{W}_\alpha$. If $\mathcal{W}_\alpha \cong \mathcal{T}$, then $\neg \mathcal{W}_\alpha = |0|$. The complement $\neg \mathcal{W}_\alpha$ is referred to as void space and void space is always relative to the convex hull. Starting out

from the hull (i.e. $\mathcal{W}_\alpha \cong \mathcal{H}$), and assuming a monotonically decreasing value of α , voids may grow into pockets, pockets into cavities and/or handles and handles may grow into separations of the α -complex, cf. Fig. 14. Observe that during this process, the genus (number of handles) changes. Consequences for border and interior may differ: we may even have that closure $\text{Cl } \mathcal{C}_\alpha$ is not a separation, whereas $\text{Int } \mathcal{C}_\alpha$ is.

In general, an α -shape will have m interiors $\mathcal{I}_l \subseteq \mathcal{W}_\alpha$ and l exteriors $\mathcal{E}_k \subseteq \neg \mathcal{W}_\alpha$. Holes are open sets and only their closures may intersect, i.e. they may meet, but not overlap, i.e. for an arbitrary pair of holes \mathcal{E}_i and \mathcal{E}_j :

$$\begin{cases} \text{(i)} & \text{Bd Cl } \mathcal{E}_i \cap \text{Bd Cl } \mathcal{E}_j \neq \emptyset \\ \text{(ii)} & \mathcal{E}_i \cap \mathcal{E}_j = \emptyset \end{cases} \quad (11)$$

where $i \neq j$, $\text{Bd Cl } \mathcal{E} = \partial \bar{\mathcal{E}}$ and the interior of an open set is equal to the open set itself. Holes can be intersected by singular faces. Such singular faces appear as a ‘spike’ intruding the hole. Despite the spike, that is contained in the surrounding α -complex, the hole remains topologically homeomorphic with an open disc. A single hole may also be bisected so that two adjacent holes are formed, the closures of which meet. Fig. 15 illustrates this.

From a topological point of view, as long as the interiors are not separated by a closed boundary, two adjacent holes are topologically one. However, from a representation point of view, every not yet α -exposed d -simplex in the (triangulated) void space is to be considered a separate hole.

6.6. Nil-objects

The domain of the α -complex representation scheme is further expanded, if holes can be explicitly represented and treated as part in a bigger assembly. Therefore, we define a *nil-object* as follows: a nil-object in Euclidean d -space E^d is an object represented by an α -complex containing the




Fig. 15. Left: \mathcal{E}_i and \mathcal{E}_j meet. Right: as long as the interiors of \mathcal{E}_i and \mathcal{E}_j are not separated and the bisecting internal boundary separating them is not closed, \mathcal{E}_i and \mathcal{E}_j represent a single hole.

Table 3

Originating, transforming and terminating processes that model the birth-life-death life cycle by means of nil-objects. The * denotes that model space C has been augmented by the nil object

Type of process	Physical process (example)	Morpho-dynamic process	Map
Originating	Deposition, sorption, reaction	Dilation	$T : nil \mapsto *C$
Transforming	Transformation	Linear evolution map	$\Gamma^n : *C \mapsto *C$
Terminating	Erosion, reaction, solution	Erosion	$\Gamma^n : *C \mapsto nil$

improper (-1) -dimensional simplicial face $\{0\}$ as its only face (Fig. 14). This definition also holds for void space *within* an α -complex. Nil-objects, homeomorphic to an open d -ball, are geometrically and topologically confirming (fitting) to any neighbouring object. As a consequence, nil-objects can be inserted in between any two adjacent objects. Nil-objects map to background elements in a conversion mapping onto FEM models. With this definition of a nil-object, consistent *spatial occupancy* of model space E^d and well-formedness of the modelled object are more easily obtained.

Not only conversion mappings benefit from the nil-object, evolutionary maps also do so. Living objects represented by α -complexes can be created by the use of evolutionary maps that act upon the landmarks, see Section 3. A nil-object may aid in the definition of homeomorphic evolutionary maps, as indicated in Table 3.

7. Conclusions and further research

Relatively complex objects with many holes and separated parts, like natural objects, can be modelled conveniently using α -complexes. No tedious description is required; just a sampling point set will do. Objects to be reconstructed can be sampled by a camera or measuring robot, for example, offering a suitable entry point for reverse engineering. Varying the α -value is intuitive, but the design of a proper weight set, an essential step in practice, can be cumbersome. Weighting has an *omni-directional* effect, ignoring tensor-like and vector-like phenomena. Further research has, therefore, been initiated, into the representation of vectorial and tensorial weight. For natural objects, modelled by example, weights can be derived from a well-chosen combination of discriminating observed physical properties. Theoretical support by a coverage process was found to provide valuable insight in the initial design and incremental improvement of the weights. Transformations can sometimes be used to improve the spatial distribution of such a point process. Landmarks with strongly directional patterns were found to be ill suited. Accumulating effects of transformations and weight require detailed knowledge of the underlying sampling data set and modelled object. When applying transformations, it is important to understand the effect on the Lebesgue measure and on the development of the α -complex with varying α .

Dynamic modelling of α -complex models is conveni-

ently supported by morphological set operations upon the generating landmark set. The result is an evolutionary map, driven by the morphological process. The ability to handle singular faces and disjunct parts in object descriptions based on α -complexes is both a strength and a complication, particularly in conversions, numerical modelling and the evolutionary map. Holes and geometrically coupled complementary objects, such as fluids in pore space, are inherently defined, along with the object itself. An α -complex may violate the criteria of continuity, solidity, closure and homogeneous dimensionality, but regularisation and the background-embedding approach (Section 5.3) can be followed to remedy this. The introduction of the related nil-object greatly aids in achieving homeomorphic maps. The α -complex as a representation scheme is unambiguous, but not unique. Validation is generally expensive.

Creating complex natural-like objects such as prostheses frequently takes a combined free form design and shape reconstruction from an example, for which α -complexes are natural candidates. It offers good possibilities for knowledge-based modelling and variational geometry. Crisp measures-of-merits and cost function are mostly lacking for natural objects, and topological constraints are often undefined. With the help of physical constraints such as mass and volume, fairly realistic geometries and topologies can be obtained, following the approach of Fig. 5. We paid no attention to visual realism, but further improvements may be anticipated by the use of α -complex-based object geometries with natural object texture maps. Generally, detailed models require vast amounts of data. For simple models ($O(10^2)$ – $O(10^3)$ sample points), results show that implementations can generally be made fast enough for interactive use. For more complex and bigger problems ($O(10^4)$ – $O(10^6)$ sample points), storage becomes increasingly critical. The largest model created had almost 4×10^5 points and took more than 1 Gb internal memory. α -Complexes can be uniquely reproduced from point sets, which makes an unevaluated description feasible. This reduces persistent storage requirements.

Acknowledgements

The authors wish to thank The Netherlands Institute of

Applied Geoscience TNO—National Geological Survey for their financial support.

References

- [1] Aurenhammer F. Power diagrams: properties, algorithms and applications. *SIAM J Comput* 1987;16(1):78–96.
- [2] Badley ME. Practical seismic interpretation. Boston: International Human Resources Development Corporation, 1985.
- [3] Barndorff-Nielsen OE, Kendall WS, Van Lieshout MNM. Number 80 in Monographs on statistics and applied probability (Cox et al., series editors). Stochastic geometry; likelihood and computation. Boca Raton: Chapman & Hall/CRC, 1999.
- [4] Bix R. Topics in geometry. San Diego, CA: Academic Press, 1994.
- [5] Bock HH. Classification and clustering: problems for the future. In: Diday E, Lechevallier Y, Schader M, Bertrand P, Burtschy B, editors. New approaches in classification and data analysis. Studies in classification, data analysis and knowledge organization, Berlin: Springer-Verlag, 1996.
- [6] Brown AR. Number 42 in AAPG Memoir. Interpretation of three-dimensional seismic data. 4th ed. Tulsa: AAPG, 1996.
- [7] Coolidge JT. A treatise on the circle and the sphere. London: Oxford at the Clarendon Press, 1916.
- [8] Edelsbrunner H, Kirkpatrick DG, Seidel R. On the shape of a set of points in the plane. *IEEE Trans Inform Theory* 1983;IT-29:551–9.
- [9] Edelsbrunner H. Algorithms in combinatorial geometry. Berlin-Heidelberg: Springer-Verlag, 1987.
- [10] Edelsbrunner H. Weighted alpha shapes. Technical Report UIUCDCS-R-92-1760 (UILU-ENG-92-1740), University of Illinois at Urbana-Champaign, Department of Computer Science, Urbana, Illinois, July 1992.
- [11] Edelsbrunner H. The union of balls and its dual shape. In: Proc Annu ACM Sympos Comput Geom. ACM, 1993. p. 218–31.
- [12] Edelsbrunner H, Facello M, Liang J. On the definition and the construction of pockets in macromolecules. *Discrete Applied Mathematics* 1998;88:83–102.
- [13] Edelsbrunner H, Muecke EP. Three-dimensional alpha shapes. *ACM Transactions on Graphics* 1994;13:43–72.
- [14] Gerritsen BHM. On the use of alpha complexes in subsurface modeling. In: Proc ENSG Conference on 3D Modeling of Natural Objects, a Challenge for the 2000s, ENSG, Nancy, France, 4–5 June, 1998.
- [15] Gerritsen BHM, Van der Werff K, Veltkamp RC. Geometric modeling with α -complexes. In: Proc of the Third Symp on Tools and Methods for Competitive Engineering, The Delft University of Technology, Delft, The Netherlands, 18–21 April, 2000. Delft University Press. p. 117–30.
- [16] Hall P. Introduction to the theory of coverage processes, Wiley series in probability and mathematical statistics. Chichester, UK: Wiley, 1988.
- [17] Kalay YE. (series editor Kalay YE). Modeling objects and environments. Principles of computer-aided design. New York: Wiley, 1989.
- [18] Lee CW. Regular triangulations of convex polytopes. Applied geometry and discrete mathematics: the Victor Klee Festschrift, DIMACS series in discrete mathematics and theoretical computer science, vol. 4. Providence, RI, AMS Press, 1991. p. 443–56.
- [19] Maxwell EA. Elementary coordinate geometry. London: Oxford at the Clarendon Press, 1952.
- [20] Melkemi M, Djebali M. Computing the shape of a planar point set. *Pattern Recognition* 2000;33:1423–36.
- [21] Møller J. Lectures on random Voronoi tessellations, Number 87 in Lecture notes in statistics. New York: Springer-Verlag, 1994.
- [22] Muecke EP. Shapes and implementations in three-dimensional geometry. PhD thesis, University of Illinois at Urbana-Champaign, 1993.
- [23] Munkres JH. Elements of algebraic topology. Redwood, CA: Addison-Wesley, 1984.
- [24] Requicha AAG. Representation of rigid solid objects. In: Encarnação J, editor. Computer Aided Design: Modelling, Systems Engineering, CAD Systems, CREST Advanced Course, Darmstadt, September 1980, Lecture notes in computer science (series editors Goos G, Hartmanis J) Berlin-Heidelberg: Springer-Verlag, 1980.
- [25] Robertson JD, Fisher DA. Complex seismic trace attributes. *Geophysics: the Leading Edge of Exploration*, 1988, Vol. 7, No. 6, p. 22–6.
- [26] Smets G, de Groof M, Nuyts J, van der Meulen D, Suetens P, Marchal G, Oosterlinck A. Interpretation of 3D medical scenes. In: Freeman H, editor. Machine vision for three-dimensional scenes, Boston: Academic Press, 1990. p. 163–93.
- [27] Stoyan D., Kendall W.S., Mecke J. Stochastic geometry and its applications, Wiley series in probability and statistics. 2nd ed. Chichester, UK: Wiley, 1995.
- [28] Stuart AM, Humphries AR. Dynamical systems and numerical analysis, Number 2 in Cambridge monographs on applied and computational mathematics (Gialet PG, Iserles A, Kohn RV, Wrights MH, series editors). Cambridge: Cambridge University Press, 1996.
- [29] Teichmann M, Capps M. Surface reconstruction with anisotropic density-scaled alpha shapes. In: IEEE Visualization '98. IEEE, 1998. p. 67–72.