Geometric Modelling with α-Complexes

Bart H.M. Gerritsen
Netherlands Institute of Applied Geoscience TNO
P.O.Box 6012, NL-2600 JA, Delft, The Netherlands
E.mail: b.gerritsen@nitg.tno.nl

Klaas van der Werff
Dept. of Design, Engineering and Production
Delft University of Technology
Delft, The Netherlands

Remco C. Veltkamp
Dept. of Computer Science
Utrecht University
Utrecht, The Netherlands

Abstract

The shape of real objects can be so complicated, that only a sampling data point set can accurately represent them. Analytic descriptions are too complicated or impossible. Natural objects, for example, can be vague and rough with many holes. For this kind of modelling, α-complexes offer advantages over triangulations and hulls at little extra computational cost. Geometric and topological descriptions are well-formalised, with the flexibility to capture holes, up to a complete separation. Spatial distribution of the point set and the attachment of weights make "special modelling effects" possible. We explore in this paper the merits of geometric modelling with α-complexes, with the objective of evaluating their practical value. We discuss the α-complex as a model description and as a representation scheme. Varying the α-value is intuitive, but weighting can be tedious. We present a few strategies. We also show how to run FEM computations on α-complexes. α-Complexes form a useful addition to existing approaches and are applicable to a number of problems not (easily) handled by existing approaches.

Keywords

geometric modelling, α-complexes, representation scheme, weighting strategies, evaluation.

Nomenclature

α	=	alpha value in $[0..\infty)$
ρ	=	index into ordered α-family $\mathcal{A} = \{\mathcal{C}_\alpha^\rho(\tilde{S})\}$ of all α-complexes for a given set $\tilde{S}$
$\mathcal{C}_\alpha(\tilde{S})$	=	α-complex; $\mathcal{C}_\alpha(\tilde{S}) \in \mathcal{A}$ and $\mathcal{C}_\alpha(\tilde{S}) \subseteq \mathcal{T}(\tilde{S})$
$\mathcal{T}(\tilde{S})$	=	regular triangulation of point set $\tilde{S}$
S	=	finite sampling data point set, $S \subset E^d$
W	=	set of real weights to attach to S
$\tilde{S}$	=	weighted sampling data point set, $\tilde{S} = S \otimes W$
k-face	=	simplex of dimension k, $-1 \leq k \leq d$
E^d	=	d-dimensional euclidian space, $d \geq 0$
$\mathbf{x}$	=	point in E^d
$\tilde{\mathbf{x}} \cong (\mathbf{x}, w)$	=	weighted point in $E^d \times W$

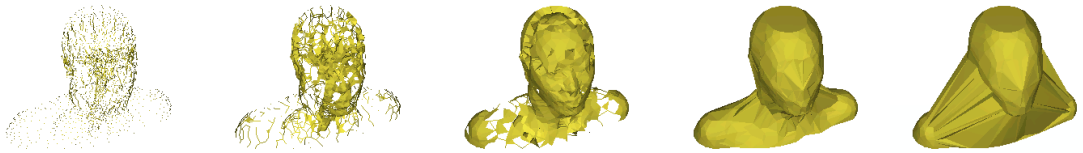

Figure 1: example of an evolving zero-weight α-complex: left: $\alpha = 3.096634e+01$, yielding a complex that is little more than just the point set with a few singular edges, developing into right: an overly fat ($\alpha = 1.805128e+03$) complex obscuring all the fine details of the face and the neck, almost grown into the convex hull. The best α for the details of the face is approx. $\alpha = 1.444991e+02$, the value of the centre picture (data set by Silicon Graphics, $n = 2780$).

1 Introduction

Over the past couple of decades, engineers and designers came 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 at *engineering* objects and assemblies, with regular mathematical properties. As soon as *natural* or *measured* 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 tolerance analysis. Also, many of today's modelling approaches are incapable of handling changing topologies. Finally, objects may get too complicated to find an analytical description, with the only remaining form of representation being a sampling data point set. *Faceted polyhedral objects* are typical representatives. But similar problems are encountered in complex *motion planning*, *assembly planning*, *visibility analysis* and a wealth of other geometry-dominated problems.

This is exactly the field where α-complexes came up-front as an alternative, capable of modelling typical engineering as well as natural objects. α-Complex-modelling offers a machinery to study various appearances of the "shape" of a set of points. Roughly speaking, it may be seen as a triangulation of that point set governed by a single geometric parameter α, giving rise to the term α-complex. The spatial occupancy by an α-complex is called an *α-shape* (see fig. 1). If the partitioning of the underlying space is irrelevant, α-shape can also be read where α-complex is written. The sample data points can be attributed a *weight*; a real-valued magnitude expressing the dominance of that point over other points. 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. Sometimes, a distinction is made between non-weighted and weighted α-complexes. Here, that distinction will not be made and α-complexes herein shall always be *weighted*. Of course, all weights can be chosen equal, possibly zero. Weighting, on top of variation of α, makes "special modelling effects" possible, and is in fact the key to practical applications.

This paper is organised as follows. After this introduction and a brief recapitulation of previous work, an informal introduction to α-complexes will be given, followed by the distinct steps of modelling with α-complexes. Next, weighting strategies will be introduced, both for "free form" modelling and for shape reconstruction. Evaluation criteria for representations of the models are presented next. Then, a few geometric modelling case studies are discussed: free form shapes, reconstructed polyhedral engineering and natural objects, and dynamic and numerical problems. Eventually, the merits of modelling with α-complexes are evaluated

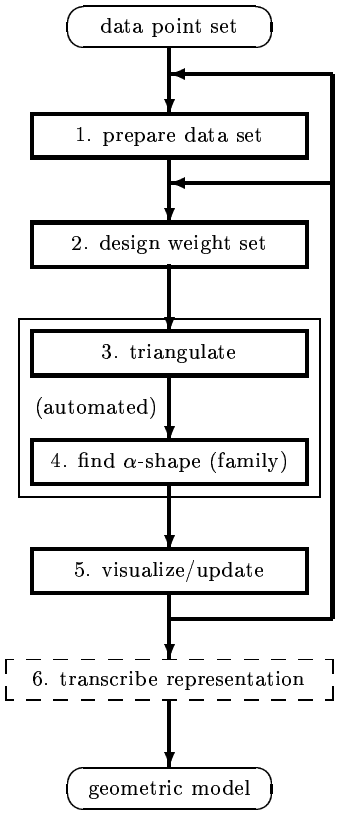

Figure 2: the sequence of basic steps involved in modelling with α-complexes, some steps may be void. Starting point is always a sampling data point set and the final result is always an α-complex, that may be transcribed into another representation. Dashed boxes represent optional steps. Triangulation and α-complex computation are steps with no or little human interaction. The modelling ingenuity is chiefly in the first steps.

against the earlier defined criteria, with regard to their practical value for the industrial and engineering community. The paper concludes with conclusions and suggestions for further research. An appendix reviews the mathematics underpinning the case studies carried out.

2 Previous Work

α-Complexes are mainly due to Edelsbrunner ([9];[10];[12]) and build on the results of regular triangulations of weighted point sets (e.g., [18];[8]). Studies on the geometry of spheres and balls are classical (e.g., [6];[5];[19]), with recent work in [25],[16] and [4]. Paoluzzi et al., in [22], pointed out the suitability of multi-dimensional simplicial complexes for geometric modelling. They did not encompass, however, non-regular complexes, like α-complexes. Edelsbrunner et al., in [11], described an application of α-shape-modelling to molecular geometric modelling and Gerritsen, in [13] described applications in earth sciences. Literature on weighting strategies is sparse; the theory of weighting is fairly well understood, particularly in conjunction to sphere-geometry (e.g., [2];[3];[9], [10]), but the use form a modellers' perspective is faintly addressed ([1]; [13]). The mathematical foundation of representation schemes is mainly due to Requicha and Voelcker ([23]). In more recent work, Kalay, in [17] further works out a number

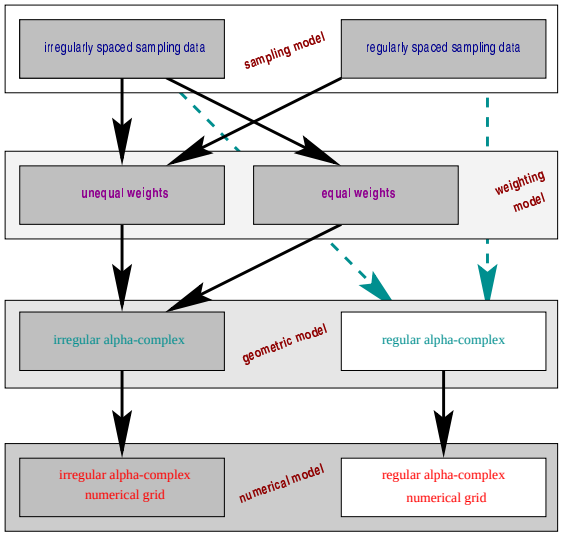

Figure 3: schematic view on the relationship between sampling data point set organisation, geometry and numerical model geometry. Gray-ed boxes show the path followed with α-complexes. Observe that both irregularly sampled data and regularly sampled data can be turned into an α-complex. For regularly spaced data samples (e.g., digital pictures), the use of unequal weights is a prerequisite.

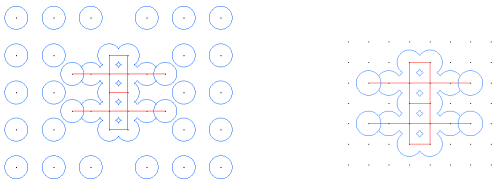

Figure 4: left: irregularly spaced but equally weighted points, right: regularly spaced, but unequally weighted points. Spheres represent weight; same size means same weight. Points with no spheres (right) have zero weight. The α-complex here is a 1-skeleton ("wire-frame"). Observe the similar and inter-changeable effect of α and weight: $\mathcal{C}_{\alpha=w}(S \times W|_{w_i=0}) = \mathcal{C}_{\alpha=0}(S \times W|_{w_i=\alpha})$

of details of various representation schemes. Further, refer to [24];[29].

3 α-Complexes

A formal description of α-complexes can be found in the Appendix. Less ponderously, given a sampling data point set S, *sampling* (or: probing) the target object, an α-complex based on that S, uniquely and unambiguously defines the object's interior and exterior, including holes. Here, the term *holes* is used for a variety of features; cf. fig. 18. Neither interior nor exterior need be connected and topologies may be manifold or non-manifold. The α-value reflects the distance over which neighbouring data points can connect. A low value of α yields an α-complex close to just the data point set, with no or few neighbours connected, a higher value of α results in an α-complex which matches or nearly matches the triangulation (see figure 1). With weights attached to the points, the evolvement of an α-complex is not only determined by *proximity* (as reflected by α), but also by *dominance* (weights). In a neighbourhood of high weights, points tend to connect more easily than in neighbourhoods with lower weights. The α-complex and the eventual regular triangulation share faces (vertices, edges, triangles, tetrahedra, $\cdots$). Finally, a triangulation is for example also a *cellular decomposition*, paving

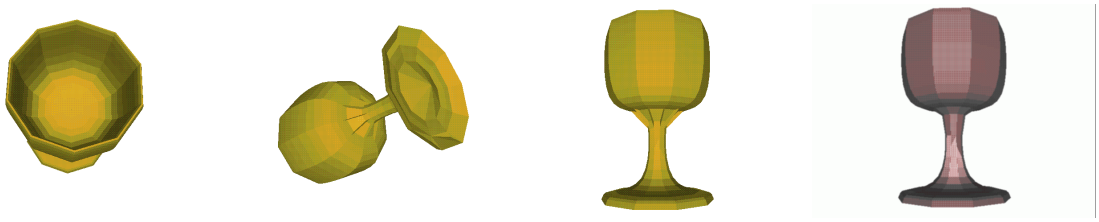

Figure 5: Left: leftmost three figures show the zero-weight α-complex of a goblet data set ($n = 520$). Close examination shows extraneous faces in the stem-to-cup transition. Applying a non-zero weight set (fig. 14) lifts this defect. The final result in shown at the right (data set obtained from GEOMVIEW DISTRIBUTION).

Figure 6: rabbit-like animal, shaped by spraying a coloured weight on a $400x400$ regular rectangular canvas. RGB-colours are translated to normalised weight $w \in [-1..1]$: $w = \pm\frac{1}{2}\bar{R} \pm \frac{1}{3}\bar{G} \pm \frac{1}{6}\bar{B}$, where $\bar{R}$,$\bar{G}$ and $\bar{B}$ denote normalised colour components (gray-levels) for red (heavy) weight, green (medium) weight and blue (light) weight, resp. Left: blue-coloured S (observe sparsely sampled spots). Centre: α-complex of the blue-mask rabbit. Right: rabbit overlaid with green weight mask; the 6 handles (holes) of the central figure have now vanished.

the way to various "transcriptions".

4 Modelling Steps

Modelling with α-complexes is a multi-step process, see fig. 2. Basically, after the preparation of the *sampling data point set* S and the *weight set* W, under optional transformations, the Cartesian product of S and W is triangulated. Then, sorted by α, the finite set of α-complexes is determined. As α grows, the α-complex grows into an ever greater sub-complex of the triangulation. Once identified, the appropriate α-complex may be converted into alternate representations. α-Complex modelling is best explained by making the following subdivision:

Problem 1; free form shaping: the α-complex representing the object is built from a modeller-generated, *synthetic* sampling data point set, thereby honouring geometric and topological constraints. Generally, the set of possible solutions to the problem is not limited to a single solution and changes to the sampling data point set S and/or the weight set W may be used to alter the shape.

Problem 2; shape reconstruction: the sampling data point set is now based on *observed* data. Generally, the observation will be subject to noise and errors. Although basically multiple solutions can be found, a single "best fit" can generally be identified according to optimisation of some cost function. The sampling data point set S is commonly

left untouched (you don't change observations) and control must therefore come from changes to weight set W.

Both classes of problems may be submitted to numerical analysis. Under certain conditions, the use of transformations of S can be justified, to "precondition" the data. Transformations cause the spacing to alter and, as shown in figure 4, the effect is complementary to the effect of changing weights.

5 Weighting Strategies

In the presence of regularly spaced data (figure 3), if not in most cases, the use of weight is essential to gain control over the geometry of the α-complex-model. With *weighting strategy*, we refer to the process of obtaining a weight set 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.

5.1 Unstructured Weighting

Weight can be assigned on a per-point basis. This form of weighting will be referred to as *unstructured weighting*. The use of unstructured weighting, applicable to both regularly and irregularly spaced data, is of limited practical value. It is typically used in combination with free-form shaping of simple shapes. To some extend, the painting (spraying) of weight can be regarded as unstructured weighting (cf. fig. 6). Unfortunately, such tools are still lacking for 3D. A best-practice solution would be the use of sliced data.

5.2 Masking

Consider the following weighting strategy for a regularly spaced sampling point sets S:

step 1 Create an empty canvas to contain the eventual weight set W of the size of the data set and initialise the canvas to zero-weight.

step 2a Instantiate a (library-based standard) weight mask by generating it with appropriate parameters.

step 2b Translate, rotate and scale the mask, as appropriate.

step 2c Map the instantiated mask onto the canvas. Mapping onto is essentially a "pixel"-based logical or mathematical operation.

step 3 Construct the weighted α-complex, based on $S \times W$.

Repeat steps 2 and 3, as appropriate.
This strategy is referred to as *masking*. Fig. 7 depicts this process, and fig. 9 shows an example using a mask that cuts out a folium-of-Descarte-shape. Standard masks can be collected in a library. Masks will be based on standard weight function, like the ones in table 1. These functions can be taken relative to a point, a line, a plane, etc. Arbitrary constrained weight functions can be compiled using for example Lagrange polynomials. Basically, masking is applicable to regularly shaped sampling data point sets of any model space dimension E^d, although tools become increasingly sparse for dimensions larger than 2.

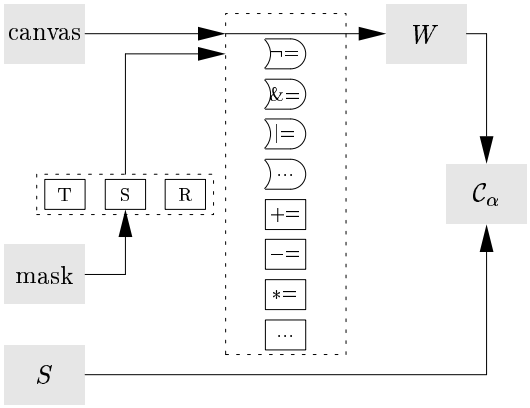

Figure 7: schematic overview of the masking technique; the aim is to compose a weight set W, that "highlights" the desired geometry ($\mathcal{C}_\alpha$) in the regularly spaced sampling data point set S. Starting point is an empty canvas (grid) of the size of S. Successively, masks (*partial, basic weight sets*) are added to the weight set W using transformations and one or more logical or mathematical operations. The Cartesian product $S \times W$ is then input to $\mathcal{C}_\alpha$.

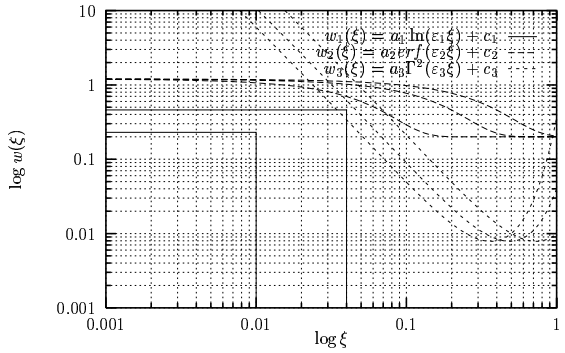

Figure 8: Weight functions with "cut-off" effect. $w_1(\xi)$ has upper-bound for $0 < \xi < \varepsilon^{-1}$. The rapid weight drop at $\xi = \varepsilon^{-1}$ can readily be seen from the figure. Vertices within this ε^{-1}-ball centred at $\xi = (0, 0, \ldots, 0)$ will be attracted to one another, those outside the ball will be distracted for equal values of α. Weight functions $w_2(\xi)$ and $w_3(\xi)$ have similar effects

5.3 Physical Properties

Shape reconstruction is relevant as far as object models are to be derived from real-live counterparts. For example when a measuring robot establishes the tolerance misfit of a cast-iron part. Commonly, when visualising observed sampling data, characteristic physical parameters are expressed in terms of colour. The target object to reconstruct the shape of, is determined by the value of these colours. The exact geometry depends on a number of things. Among others:

- The observation technique features.
- The "contrast" in the data.
- The threshold value chosen to discriminate "background" (embedding environment) from "foreground" (the object).
- The suite of properties observed.

FUNCTION	DEFINITION
gamma(ζ)	$\Gamma_\gamma(\zeta) = \int_0^\infty \zeta^{\gamma-1} e^{-\zeta} d\zeta,\ \gamma > 0$
-ln(ζ)	$-\ln(\zeta) = \int_0^\zeta \eta^{-1} d\eta,\ \zeta > 0$
error(ζ)	$\mathrm{erf}(\zeta) = \frac{2}{\sqrt{\pi}} \int_0^\zeta e^{-\eta^2} d\eta$
normal(ζ)	$p(\zeta) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{(\zeta-\mu)^2}{2\sigma^2}}$
sign(ζ)	$\mathrm{sign}(\zeta) = \begin{cases} -1 & \text{for } \zeta < 0 \\ 0 & \text{for } \zeta = 0 \\ 1 & \text{for } \zeta > 0 \end{cases}$
block(ζ)	$\mathrm{p}(\zeta) = \begin{cases} 0 & \text{for } \zeta \notin [a,b] \\ 1 & \text{for } \zeta \in [a,b] \end{cases}$

Table 1: natural weight functions exhibiting a "cut-off" effect. We may use the linear transform $\zeta = \beta_1 \xi + \beta_2$, with $\beta_1, \beta_2 \in R$, to allow shifting and scaling of the "cut-off" point.

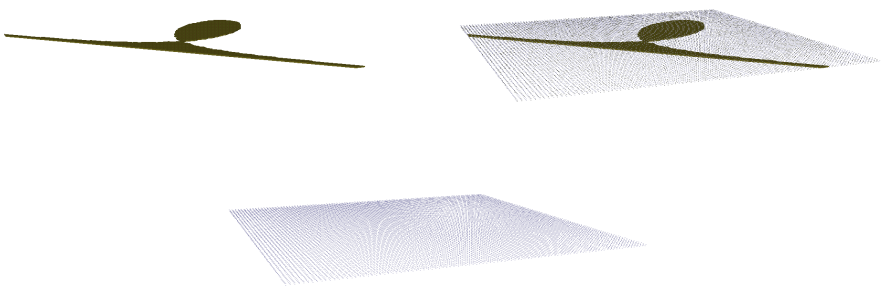

Figure 9: the masking process: the regularly gridded set S, forming the canvas (bottom), is weighted using predefined masks (middle); here a mask containing Descarte's folium. The resulting α-complex is shown at the top.

We refer to such shapes as *property-ruled* shapes. Many natural objects, "extracted" from observed data, are in fact property-ruled. An example may illustrate this: a surgeon removing a tumour will somehow rely on (a suite of) tissue properties to determine where to cut.

Often, observed sampling date get blurred by noise and other undesired effects. If such features hamper the finding of the right α-complex, filtering away remote areas outside the region of interest is easily accomplished with an AND-mask, i.e., a logical & =-operation. Another problem with observed data is the translation of vector- and tensor-properties to scalar, *omni-directional* weights.

5.4 Neighbourhood Weighting

In fig. 4, the complementarity between distance and weight has been shown. The weighting strategy based on nearest neighbour distances exploits exactly that. Fig. 10 shows the distribution of the nearest neighbour distances for the scapula-sampling data set of fig. 17.

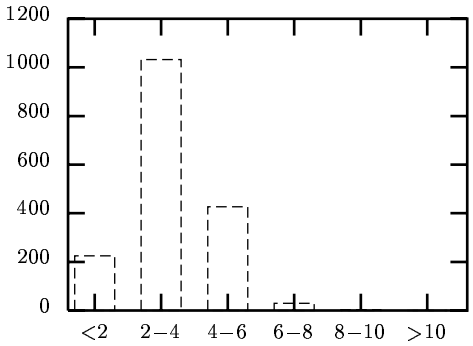

Figure 10: distribution of the nearest neighbour distances of the scapula-sampling data set of fig. 17.

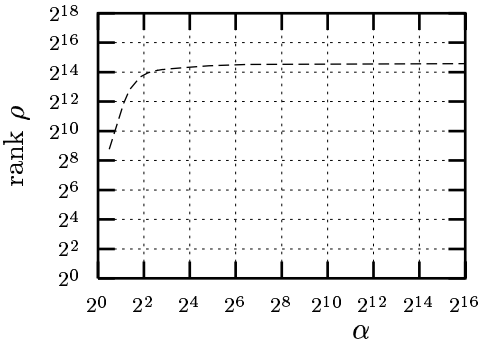

Figure 11: the increase of α vs. α-rank of the scapula-sampling data set of fig. 17.

The goal with nearest neighbour distances based weight assignment is to accomplish that the variance of the Laguerre distances (see App.) is minimised. This also affects the relation of α vs. α-rank, which is related to the accumulative distribution of the nearest neighbour distances (see fig. 4), and the curvature K vs. α-rank. The α-rank is the index into the finite α-family sequence, ordered by increasing α.

6 Evaluation Criteria

Computer models capturing real world objects and phenomena can be regarded as languages of representation composed of symbol structures and grammar rules to manipulate them, assigning a meaning to a collection of symbols and an interpretation in terms of the real world objects they represent([17]). For computer modelling purposes, different approaches in describing the geometry and topology of modelled objects exist ([17];[29];[27]). E.g.,

Volumetric description: The objects' geometries and topologies are described in terms of *underlying space* or *hyper-volume*.

Boundary description: The objects' boundaries are described, bounding the objects interiors and separating the interior from the exterior. When partitioned, the objects' internal boundaries, separating the internal cells are also part of this description.

Part-whole description: In this approach the objects' compositions are described in terms of *primitive* parts.

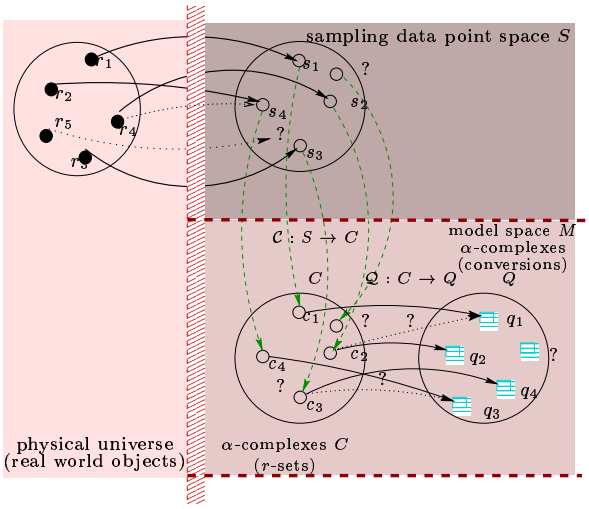

Figure 12: "chain" of representation-to-representation mappings in modelling with α-complexes. S is the sampling data point set, sampling the real world objects R. C is the set of α-complexes, Q the set of conversions. Dotted arrows and question mark-labelled elements refer to the questions raised in the text.

This subdivision suggests a much sharper distinction between these approaches than practical implementations permit. Object descriptions can be qualified by a number of criteria ([17];[29], chapter 7). 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* $Int\ \mathcal{P}$ and *exterior* $Ext\ \mathcal{P}$, separated by *boundary* $Bd\ \mathcal{P}$. Point location, for example, is feasible when this criterion is met.

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

Rigidity: the geometry of the modelled object shall not depend necessarily on its position, i.e., on its location or its orientation in space. This implies that a mapping from sample space S to model space A shall cause the α-complex to be *position invariant*.

Continuity: the represented object shall not be composed of disjunct (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 $d+1$ $(k-1)$-simplices. If so, k-faces can be represented as regularised set (r-sets).

Disjunct interiors: the interior $Int\ f_i$ and $Int\ f_j$ of any two k-faces f_i and f_j shall be disjunct, i.e., $Int\ f_i \cap Int\ f_j = \emptyset, \forall i, j$. This also excludes self-intersection.

Orientability: all k-faces shall be orientable. Notice that holes usually have a "negative" orientation, for area, volume, etc.

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

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. Criteria for representation schemes are:

Uniqueness: the cardinality of all representations that describe 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 in sect. 8. Question marks and dashed mapping arrows in fig. 12 will be discussed therein.

7 Experimental Cases

In this section, a number of α-complex-based geometric modelling cases, of both engineering objects and natural objects will be presented. Also, a few words will be spent on numerical modelling with α-complexes.

7.1 Engineering Objects

A goblet is a simple central axis circular symmetric object, as depicted in fig. 5. It is easily approximated by a zero-weight α-complex, with only few defects. A single-zone/single value weight model (cf. fig. 14) removes these defects, yielding a perfect goblet.

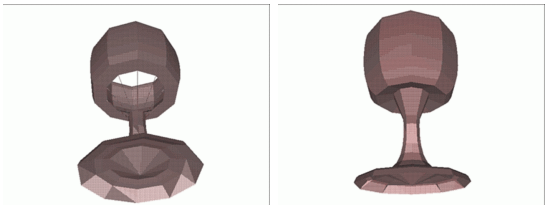

Figure 13: the effect of improper weight sets for the goblet α-complexmodel. Left: the negative stem-to-cup weight zone is too great. By the time the increasing α can override the negative weight, the entire zone will be bridged at once (observe the singular edges). Right: the negative weight zone is too small; the points at the inside bottom of the cup are able to connect, whereas the points from stem-to-cup are not. Compare to fig. 5

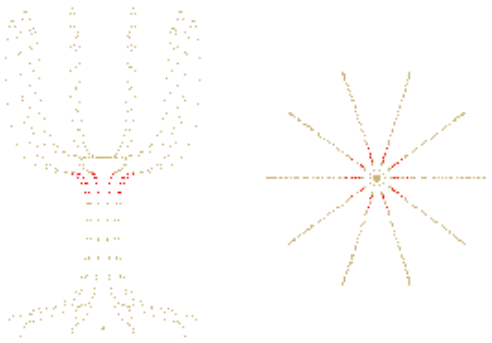

Figure 14: single zone/single value weights, obtained by the *block*(ζ)-function (table 1). Observe the bolder dots in the zone from stem-to-cup. Left: front view, right: top view

Figure 15: the comet West, as observed and produced by Laborde, March 1976. (image obtained from Univ. of Arizona Internet site)

7.2 Rabbit

A very simple example of a free-form natural object is given by the rabbit-like creature of fig. 6. The rabbit has been created by spraying coloured weight on a canvas. Sprayed weight can be homogenised for example using a "smearing"-filter. Approx. the same result could have been obtained (refer to fig. 3 and 4) by spraying properly-spaced zero-weight data points.

7.3 Comet West

A comet is a irregularly shaped natural object of frozen gas and rocky debris, orbiting around the sun. A comet has a focal bright approx. 10 km wide kernel, called the *nucleus*. When approaching the sun, a comet develops three tails: the *coma*, a large trailing cloud of diffuse material, the *ion tail* of ionised plasma, and the *hydrogen envelope*, with hydrogens that escaped the comet's gravity. Refer to figure 15. The comet West[1] has been observed by various observers during its bright appearance in 1976.

The model of fig. 16 was obtained by exploiting the effect demonstrated in fig. 4. Rather than turning physical properties (represented here by complementary colours yellow and

[1] named after the astronomer West, who first described his observation of the comet

Figure 16: α-complex of the comet West. Left: the initial model, by gray-scale weighting (threshold: 0.4), centre: nucleus plus ion tail, right: α-complex of the colour parameter space. X-axis represents red, Y-axis green and Z-axis blue. Multiple "clusters" are visible: a blue cluster (dark) representing the ion tail and the edges of the dust tail, a gray cluster along the main diagonal representing the dust tail and a rest cluster (nucleus, coma, dust tail).

blue) directly into weight, we augmented the 2-dimensional sampling point space S by a 3-dimensional parameter space $R \times G \times B$, thus obtaining a 5-dimensional hyper-space $\tilde{S}$. Generating an α-complex of the so obtained augmented set $\tilde{S}$ causes a "natural" clustering to appear, by colour value. Intersections and projections can be used to extract the part of the α-complex seeked for. This approach works best for strong clusterings. Notice that the colour determines the "direction", but the intensity distribution (contrast) determines the distances within and between clusters. The clustering in this case is weak, but nevertheless the approach works.

7.4 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 [26]. Fig. 17 shows an α-complex of the left and right shoulder blade (*scapula*). The picture shows a zero-weight α-complex, $\alpha = \mathcal{O}(10^4)$, $w = 0, \forall \tilde{s} \in S$, that captures the bones fairly well, except for relatively few singular triangles at the *glenoid cavity* (not visible). Weighting was required to turn the complex into a model good enough to be fed into an FEA-tool. This sampling data point set, with nearest neighbour distance distribution as in fig. 10, is a good candidate for neighbourhood weighting.

7.5 Numerical Modelling

An α-complex, in terms of numerical computing, is generally an *unstructured grid* and in practice, Finite Element Analysis (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 non-connected parts. 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. The framework of our methods is to turn the tetrahedra of the α-complex ("foreground material") into tetrahedra FEM-elements. The α-complex is embedded in the corresponding convex hull ("background material" or "bulk"). After applying boundary and/or initial conditions, computations can be started.

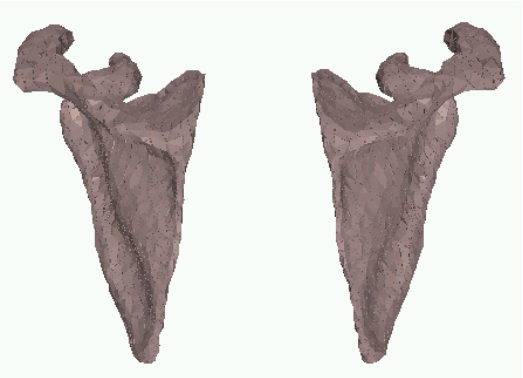

Figure 17: α-complex of the left and right shoulder blade (data set by courtesy of Delft Univ. of Techn., dep. Elec.Eng., $n = 1717$).

Our method exploits the fact that an α-complex and a convex hull of a point set share faces: every vertex,edge,triangle and tetrahedron of the α-complex also belongs to the corresponding triangulation. A convex hull, a cellular decomposition *continuum*, can always be input to any FEA-package, and *boundary conditions* can be attached to its boundary faces. Sometimes, FEA-codes run tests with respect to tetrahedra aspect-ratio's. Slivers will not pass such tests. Most slivers are developed when the α-complex approaches the convex hull. Therefore, rather than taking the convex hull itself, we lower α a bit and take an α-complex close to it (fig. 11). The rank of the α-complex to embed must be smaller, which is generally the case.
The approach is then:

step 1: Compute the α-complex family and select the embedding background α-complex a number of ranks before reaching the convex hull. The number of ranks to drop is determined by the presence of slivers. Assign background material properties to the elements of the selected α-complex.

step 2: Compute the normals of all the triangles in the border, in order to find out how (directional) boundary conditions need to be attached.

step 3: Assign foreground material properties to the target 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 replace the background material properties by foreground properties.

step 4: Run the FEA-analysis and optimise the α-complex as appropriate.

An application of this approach can be found in [13].

8 Evaluation

In this section, the α-complex will be evaluated against the criteria stated in section 6. Both as a model description and a representation scheme. There is not so much 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.
Modelling with α-complexes 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*).
- (Faceted) boundary description, built from $(d-1)$-faces (*facets*) of the α-complex.

The solidity-criterion becomes slightly more complicated if we relax the convex polytope to become an α-complex, with internal voids and possibly singular faces (fig. 18, table 2). Singular faces can be identified (see App.) and removed, and voids can be identified. Removal of singular faces ensures homogenous dimensionality, only if α is high enough to permit the forming of d-faces. It also ensures 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. This behaviour can be very beneficial in dynamic object description, like growing cracks, oxidation or eroding sand bodies, but complicates validity and consistency verifications on the model as well as further operations on the model. For many problems, the background-embedding approach of sect. 7.5 can be followed. For finite complexes, finite storage complexity can be shown to exist, as well as finite time complexity for operations upon them. To a great extend, characteristics are as with cellular decomposition. The faces of an α-complex are orientable. Voids usually have a "negative" orientation, for area, volume, etc. If $Bd\,\mathcal{C}_\alpha$ (i.e., the corresponding α-shape) can be triangulated (tessellated with triangles), such that the boundary is closed and connected, then these two properties ensure well-formedness. But 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 and dominance, not by position or orientation. Transformations that can be represented by a homeomorphic mapping do not change the order within the α-family. They only change α in the ordered pairs (ρ, α), in contrast with transformations that change the nearest neighbours, like scaling by different factors in different directions. The interior of any two k-faces is always disjunct, as they are faces in the underlying triangulation. This also excludes self-intersection.

In a sense, α-complexes can also be seen as a means of representation. First question to be raised (see fig. 12) concerns the domain: can every physical object be represented by a sampling data point set? One may argue as to whether an object can always be measured, or represented by the pixels of photograph? For the class of natural and engineering objects of interest in this paper, the answer is yes. How adequately this can be done, depends mainly on the observation/generation technique. This question cannot be addressed in the frame of this paper. 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. The step to account for discretisation with finite-precision computer internal number representations, like in [28], is not explicitly shown but assumed to be implicit in this step.

Once a data point set has been sampled, the next question is, can every point set be triangulated? Disregarding degenerate cases (that can be perturbed) all sets can be triangulated.

Some remarks have to be made, however. Points must be in *general position* (see App.). Duplicates must be dropped, and also *redundant* vertices, too close to dominant points. Singular (minimum) objects, sampled by a single point, result in a 0-triangulation of a single 0-face (vertex). The only α-complex in the α-family is identical to this triangulation. Generally, a proper spatial distribution of the sampling points is required, in order to obtain a good-quality triangulation. It can be shown ([12]) that a unique relation exists between the sampling data point set, the α-complex and its underlying space (fig. 1). Variations in the α-value, even though constrained, may instantiate a finite family of models (e.g., fig. 1) 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, like material properties, central moments, etc. The simplex, the basic building block, allows for relatively simple computational schemes.

Conversion mapping transforms an α-complex representation of an object onto another such representation Q (fig. 12). Its characteristics depend heavily on the target representation of the conversion. If the representation in $\mathcal{C}_\alpha$ and in Q are *homeomorphic*, then a bijective relation can be shown to exist ([21]) and conversion mapping is basically without loss of information. Topological relations existing in $\mathcal{C}_\alpha$ are preserved under conversion. Examples of homeomorphic representations in $\mathcal{C}_\alpha$ and Q are transcription of a non-degenerate regularised α-complexinto a cellular decomposition. An example of a non-homeomorphic representation is a projection from E^d onto E^{d-1}. 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.

Let us now contrast the α-complex description with B-rep and CSG. What sets out α-complex modelling is the capability of handling a large number of cells, singular faces and disjunct parts. For example, inserting massive amounts of cells in a CSG-model is not very convenient. 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 will fail to describe the d-volume of an arbitrary α-complex. The representation that comes closest to what is needed to accurately and unambiguously describe α-complexes would be a part-whole description, built up using instantiated simplicial primitives.

Finally, the handling of holes will be addressed in greater detail. Assume a connected interior for a while. Further, using less ponderous definitions (see fig. 18):

A cavity is a hole, or a *depression* lying in the boundary of the object, accessible via the (unbounded) exterior of the object.

A pocket is a cavity with a small "neck".

A void is an entirely internal hole. This is a special hole in that it creates an additional (bounded) exterior of the object.

A tunnel is a pathway from boundary to boundary, piercing the objects. It can be regarded as a void with two distinct connections to the single unbounded exterior of the object.

Zeid, in [29], distinguishes four classes of polyhedral objects. For example, a triangulation classifies as a class 1 object. Refer to fig. 12 and table 2. Every convex polyhedron can

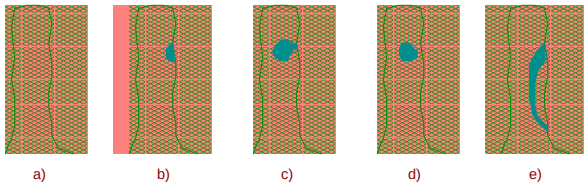

Figure 18: diapir with various forms of holes, to illustrate the terminology; left to right: a): diapir without a holes (a simple shape), b): diapir with cavity or boundary hole, c): diapir with pocket, d): diapir with internal hole, or void, e): diapir with tunnel or handle causing the genus to become nonzero. A void induces a separation of the exterior, while the others do not.

Class	k-faces	holes	tunnels	Description
1	x			simple polyhedra, 1 loop, 1 shell, simple k-shell
2	x			same, but with inner loops, non-conforming 2-faces
3	x	x		same, with cavities, pockets and voids
4	x		x	same, with handles (non-zero genus)

Table 2: topological classification of polyhedral objects, adapted, after Zeid ([29]).

be triangulated (e.g.,[21]). Triangulations of a point set in E^d have only 1 unbounded cell; the unbounded exterior. When a (triangulated) object contains one or more voids, m say, there will be m additional (bounded) exteriors, completely surrounded by the interior. Notice that disjoint exteriors imply multiply connected interiors, and vice versa. Starting out with a triangulation (i.e., $\mathcal{C}_\alpha \cong \mathcal{T}$), 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. Observe that during this process, the genus (number of handles) changes. We may even have that closure $Cl\ \mathcal{C}_\alpha$ is not a separation, but $Int\ \mathcal{C}_\alpha$ is. An α-shape (the underlying space of an α-complex) differs from the convex hull (the underlying space of the triangulation) by the total amount of space occupied by the various holes. In other words, for an arbitrary α-complex, the union of underlying space of the complex and the holes yields the convex hull. If the α-complex is identical to the triangulation, then the underlying space of the holes is a null-space.

The domain of the α-complex representation scheme is further expanded, if holes can be explicitely represented and treated as part in a bigger assembly. Therefore, we now introduce a *nil-object*, defined as follows. A nil-object in euclidian d-space E^d is an object represented by an α-complex containing the improper (-1)-dimensional simplicial face $\{\emptyset\}$ as its only face. This definition also holds for void space *within* an α-complex. Nil-objects, homeomorphic to an open d-ball, are geometrically and topologically conforming (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.

Overlooking all the above aspects, table 3 compares the α-complex representation with CSG, B-rep and cellular decompositions. A few supplementary aspects have been added. In daily practice, many hybrid representations exists, often tailored to particular purposes.

	FACETED B-REP	CSG	CELL DECOMP.	α-COMPLEX
solidity	yes	yes	yes	if regularized
continuity	yes	yes	yes	no, only part-by-part
rigidity	yes	yes	yes	yes
homo. dim.	yes	if regularized	yes	if regularized
closure	yes	if regularized	yes	if regularized
orientability	yes	yes	yes	yes
disjunct Int	yes	yes	yes	yes
finiteness	yes	yes	yes	yes
domain	rigid homogenous manifold complex shapes massive description	rigid homogenous (non-)manifold moderate complexity lean description	rigid homogenous manifold complex shapes massive description	rigid homogenous non-manifold very complex shapes massive description, reproducible from S
validity	euler rules, expensive	regularized, set-theoretic	validation expensive	euler + α+weighted metric, expensive
completeness	homogenous properties	homogenous properties	homo. prop. per cell	homo. prop. per face
uniqueness	no	no	yes	no

Table 3: α-Complexes compared to B-rep,CSG and cell decomposition. Top part: α-complex as model description, bottom part: α-complex as representation schemes

9 Conclusions

The α-complex description of an object can be compared, to some extend, with a part-whole description using k-simplices as primitive instances. For a high-enough α-value, it is also a volumetric description, with a cellular complex built from d-simplices (possibly, after regularisation) and a boundary description, by the $(d-1)$-facets. Compared with other model descriptions, the biggest problem is occurring in the event of disjunct parts and singular faces. As a result, the α-complex as an object description violates the criteria of continuity, solidity, closure and homogeneous dimensionality. Regularisation is required to repair these violations. In addition, the background-embedding approach (sect. 7.5) can be followed. The rigidity-criterion is met, like finite time and storage complexity. The faces of an α-complex are orientable. The boundary is generally not closed and connected, only after regularisation and on a part-by-part basis. The interior of any two k-faces is always disjunct, self-intersections are excluded. The α-complex as a representation scheme is unambiguous but not unique. Validation is generally expensive.

We have evaluated the use of α-complexes in a couple of applications, among others, on free-form shaping and polyhedral shape reconstruction, both of engineering as well as natural objects. We designed and experimented with various weight strategies. We also ran FEM computations on α-complexes. Relatively complex objects with many holes and separated parts 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. Considerable effort has gone into finding suitable weighting strategies. A few strategies have been developed and evaluated, but further research is still to be carried out on this subject. Occasionally, working with the accumulating effects of transformations and weight requires detailed knowledge of the underlying sampling data set and modelled object. Line-oriented data and data with strongly directional patterns were found to be ill-suited. Transformations can sometimes be used to precondition the spatial distribution. In shape reconstructions where the geometry depends on some physical phenomenon, finding a strong relation between properties and weight is usually the critical factor to obtain an intuitive form of modelling. With respect to natural objects, still more "realism" is 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 ($\mathcal{O}(10^2..10^3)$ sample points), results show that implementations can generally be made fast enough for interactive use. For more complex and bigger problems ($\mathcal{O}(10^4..10^6)$ sample points), storage becomes increasingly critical. α-Complexes can be uniquely reproduced from point sets, which makes an unevaluated description feasible. This reduces storage requirements. Finally, surface/surface and other intersections of faceted objects, though not addressed in this paper, remains a research area still suffering from great immaturity and concern (e.g, [29],chp. 7), that may seriously hinder further applications of α-complexes. Further research is targeted at improving weighting strategies, such as neighbourhood weighting and anisotropic weighting, and tools.

Acknowledgements

Authors wish to thank the Netherlands Institute of Applied Geoscience TNO – *National Geological Survey* for their financial support.

Appendix: Mathematical Concepts

Throughout, let E^d denote a **Euclidean d-space** . Further, let $\alpha \in R$ be a non-negative real number on ray domain $[0, \infty)$, let weight w be a real number, $w \in R$, and let a finite set $S \subset E^d$ of sampling data points be augmented with weights. As a result, the representations for a weighted point $\tilde{\mathbf{s}}$ becomes a $d+1$-tuples over the **Cartesian product** $E^d \times R$:

$$\tilde{\mathbf{x}} = (x_1, x_2, \ldots, x_d, w) \cong (\mathbf{x}, w) \tag{1}$$

Only when confusion might arise, we will use the tilde over data point symbols again, to recall their weighted nature. We will drop the tilde-notation otherwise; in general, we will assume data points to be weighted. The points $\mathbf{s}$ of a set $S \subseteq E^d$ are said to be in **general position** in E^d if every subset Y_j, $Y_j \subseteq S$, $card(Y_j) \leq d+1$ is geometrically independent. A set $S \subseteq E^d$, with $S = \{\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_{d+1}\}$ is said to be **geometrically independent** if for all $i, 1 \leq i \leq d+1$:

$$\left.\begin{array}{rcl} \sum_{i=1}^{d+1} a_i \mathbf{s}_i & = & 0 \\ \sum_{i=1}^{d+1} a_i & = & 0 \end{array}\right\} \Rightarrow a_i = 0 \tag{2}$$

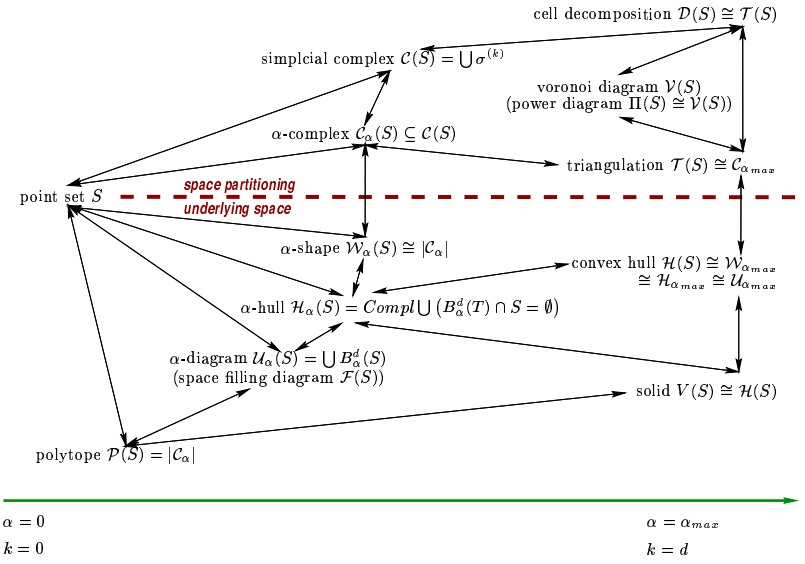

Figure 1: relationships between mathematical concepts discussed in this paper.

Let $d(\mathbf{p},\mathbf{q}) = \|\mathbf{p}-\mathbf{q}\| = \langle \mathbf{p}-\mathbf{q}, \mathbf{p}-\mathbf{q}\rangle^{\frac{1}{2}} = \sqrt{\sum_{k=1}^{d}(p_k - q_k)^2}$ denote the **Euclidean distance** between any two points $\mathbf{p},\mathbf{q} \in E^d$ and let $\varepsilon \in R$ denote a real number, with $\varepsilon \geq 0$. A $(d-1)$**-sphere** $S_\varepsilon^{d-1}(\mathbf{c})$ with radius ε and centred at $\mathbf{c}$ is the set of points $\mathbf{x} \in \mathrm{E^d}$, for which $S_\varepsilon^{d-1}(\mathbf{c}) = \{\mathbf{x} \in E^d \mid \|\mathbf{x}-\mathbf{c}\|^2 - \varepsilon^2 = 0\}$, or, in normalised form:

$$S_\varepsilon^{d-1}(\mathbf{c}) = \{\mathbf{x} \in E^d \mid \mathbf{x^T x} - 2\mathbf{c^T x} + \|\mathbf{c}\|^2 - \varepsilon^2 = 0\} \tag{3}$$

Let line L (figure 2) belong to a pencil through point $\mathbf{x}$ and be given by $\mathbf{y} = \mathbf{u}\lambda$, such that $\mathbf{u^T u} = 1$. Line L enters sphere S for $\lambda = \lambda_1$ and leaves sphere S at a point on L where $\lambda = \lambda_2$. The normalised equation for S (see equation (3)), takes the form:

$$A_2\lambda^2 + A_1\lambda + A_0 = 0 \tag{4}$$

where:

$$\begin{cases} A_2 = & \mathbf{u^T u} = 1 \\ A_1 = & 2(\mathbf{x^T} - \mathbf{c^T})\mathbf{u} \\ A_0 = & \mathbf{x^T x} - 2\mathbf{c^T x} + \|\mathbf{c}\|^2 - \varepsilon^2 \end{cases} \tag{5}$$

This is an ordinary quadratic equation in λ which has two real roots $\lambda_{1,2}$ if $D = A_1^2 - 4A_2A_0 > 0$, i.e., if $\varepsilon^2 > 0$. If $\varepsilon^2 < 0$, roots will be imaginary, associated with a negative weight $w = -\varpi = i^2\varpi$ where $\varpi > 0$ and $i^2 = -1$. A sphere with $\varepsilon^2 < 0$ is called an **i-sphere** .
The **power of a point x** (figure 2) has the following geometric interpretation ([2],[10],[4]). With $\mathbf{x}$ a point in E^d and $\bar{S}_\varepsilon^{d-1}(\mathbf{c})$ a closed $(d-1)$-sphere in E^d, the power $Pow_{S(c)}^{d-1}(\mathbf{x})$ of $\mathbf{x}$ with respect to $\bar{S}_\varepsilon^{d-1}(\mathbf{c})$ is the square of the distance $\pi = d(\mathbf{x},\mathbf{p}) = \|\mathbf{x}-\mathbf{p}\|$, with $\mathbf{p}$ being a contact point of a line through $\mathbf{x}$, touching the sphere. Refer to figure 2. $\mathbf{p'}$ is the other tangent point, and for $\mathbf{p}$ we may also read $\mathbf{p'}$, so $\pi = d(\mathbf{x},\mathbf{p}) = d(\mathbf{x},\mathbf{p'})$.

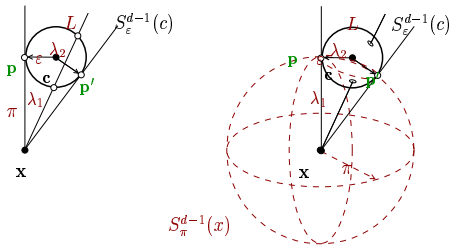

Figure 2: left: power $Pow_{S(c)}(\mathbf{x})$ of $\mathbf{x}$ with respect to $(d-1)$-sphere $\bar{S}_\varepsilon^{d-1}(\mathbf{c})$, here depicted in E^2 (left) and E^3 (right). Line L belongs to a pencil through point $\mathbf{x}$ and is given by $\mathbf{y} = \mathbf{u}\lambda$, such that $\mathbf{u}^\mathbf{T}\mathbf{u} = 1$. Power $Pow_{S(c)}(\mathbf{x}) = \lambda_1\lambda_2$. In $\mathbf{p}$ and $\mathbf{p}'$, $\lambda_1 = \lambda_2 = \lambda$ and $\lambda_1\lambda_2 = \lambda^2 = \pi^2$. Notice that if $\varepsilon \to 0$, then $\pi = \sqrt{Pow_{S(c)}(\mathbf{x})} = \|\mathbf{x} - \mathbf{c}\|$. Right: the geometric interpretation of two orthogonal weighted points $\tilde{\mathbf{x}} = (\mathbf{x}, w_x)$ and $\tilde{\mathbf{c}} = (\mathbf{c}, w_c)$ as $(d-1)$-spheres with radius $\pi = \sqrt{Pow_{S(c)}(\mathbf{x})}$ and $\varepsilon = \sqrt{Pow_{S(x)}(\mathbf{c})}$, resp. The two spheres are orthogonal (meet orthogonally) if they intersect exactly in $\mathbf{p}$ and $\mathbf{p}'$. In that case, we have that $\|\mathbf{x} - \mathbf{c}\|^2 - w_x - w_c = \|\mathbf{x} - \mathbf{c}\|^2 - \pi^2 - \varepsilon^2 = 0$.

If radius ε of $\bar{S}_\varepsilon^{d-1}(\mathbf{c})$ grows, such that a point $\mathbf{x}$ becomes contained in the border $Bd\ \bar{S}_\varepsilon^{d-1}(\mathbf{c})$ of the sphere or in the interior $Int\ \bar{S}_\varepsilon^{d-1}(\mathbf{c})$ of the sphere, then its power $Pow(\mathbf{x}, \mathbf{c})$ becomes zero, and negative, resp.

A **weighted closed d-ball** $\tilde{\bar{B}}^d_{\sqrt{w_c}}(\mathbf{c})$ can now be defined as a d-ball with radius $\sqrt{w_c}$, centred at point $\mathbf{c}$, of which the interior consists of points having a non-positive power with respect to B:

$$\tilde{\bar{B}}^d_{\sqrt{w_c}}(\mathbf{c}) = \{\mathbf{x} \in E^d |\ \|\mathbf{x} - \mathbf{c}\|^2 - \sqrt{\mathrm{w_c}}^2 \leq 0\} \tag{6}$$

For a negative weight $w_c = -\varpi = i^2\varpi < 0$, $\tilde{\mathbf{c}} = (\mathbf{c}, w_c) = (\mathbf{c}, -\varpi)$ can be interpreted as an i-sphere (d-ball) centred at $\mathbf{c}$ with radius $-i\sqrt{\varpi}$. Often in geometry, therefore, weighted points are treated as **d-balls**. If in the same equation, we define for the **weighted norm** :

$$\|\tilde{\mathbf{x}} - \tilde{\mathbf{c}}\| = \|\mathbf{x} - \mathbf{c}\|^2 - \sqrt{w_c}^2 \tag{7}$$

or, for its associated **(weighted) Laguerre distance**:

$$d_{B(c)}(\mathbf{x}) = d(\mathbf{x}, \mathbf{c})^2 - \sqrt{w_c}^2 \tag{8}$$

we can rewrite equation (6) into:

$$\tilde{\bar{B}}^d_{\sqrt{w_c}}(\mathbf{c}) = \{\mathbf{x} \in E^d \mid \|\tilde{\mathbf{x}} - \tilde{\mathbf{c}}\| = d_{B(c)}(\mathbf{x}) \leq 0\} \tag{9}$$

Notice that the norm $\|\tilde{\mathbf{x}} - \tilde{\mathbf{c}}\| = \|\mathbf{x} - \mathbf{c}\|^2 - w_c$ resembles $Pow^d_{B(c)}(\mathbf{x})$. Alternate notations: $\|\tilde{\mathbf{x}} - \tilde{\mathbf{c}}\| = d_{B(c)}(\mathbf{x}) = \tilde{d}(\mathbf{x}, \mathbf{c}) = Pow(\mathbf{x}, \mathbf{c}) = Pow^d_{B(c)}(\mathbf{x})$, some of which were already introduced.

Two weighted points $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{x}}_j$ are said to be **orthogonal** (figure 2) if:

$$\|\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j\| = \|\mathbf{x}_i - \mathbf{x}_j\|^2 - w_i - w_j = 0 \tag{10}$$

Following Edelsbrunner ([10]), with $\tilde{\mathbf{x}}_i = (\mathbf{x}_i, w_i)$ and $\tilde{\mathbf{x}}_j = (\mathbf{x}_j, w_j)$ two weighted points, we also write:

$$\|\tilde{\mathbf{x}}_{i_{+\varpi_i}} - \tilde{\mathbf{x}}_{j_{+\varpi_j}}\| = \|\mathbf{x}_i - \mathbf{x}_j\|^2 - (w_i + \varpi_i) - (w_j + \varpi_j) \tag{11}$$

For example, in figure 2, we have that $\|\tilde{\mathbf{x}}_{-\pi} - \tilde{\mathbf{c}}\| = \|\mathbf{x} - \mathbf{c}\|^2 - (\pi - \pi)^2 - \varepsilon^2 = \|\tilde{\mathbf{y}} - \tilde{\mathbf{c}}\|$, with $\tilde{\mathbf{y}} = (\mathbf{x}, 0)$.

A **weighted alpha ball** is defined as follows:

$$\tilde{B}^d_{\sqrt{\alpha+w_c}}(\mathbf{c}) = \{\mathbf{x} \in E^d \mid \|\mathbf{x}-\mathbf{c}\|^2 - \sqrt{\alpha+w_c}^2 < 0\} \tag{12}$$

As special cases, α-ball $\tilde{B}^d_{\sqrt{\alpha+w_c}}(\mathbf{c})$ with $w_c = 0$ represents an unweighted point (unweighted α-ball) and merely a weighted point if $\alpha = 0$. With $\alpha \in [0..\infty)$, B has a "start"-radius of $\sqrt{w_c}$ and grows proportional to $\sqrt{\alpha}$ as α increases.

A **hyperplane** of d-space E^d is a $(d-1)$-dimensional bi-secting facet spanned by $d+1$ affinely independent points $\mathbf{x}_i \in E^d$. For a hyperplane:

$$H = \{\mathbf{x}, \mathbf{x} \in E^d \mid \ < \mathbf{x}, \mathbf{a} >= 0\} \tag{13}$$

for some $\mathbf{a} \neq \mathbf{0}$. Generally, $d+1$ hyperplanes suffice to span a **d-polyhedron** (a d-simplex in this case). A d-polyhedron can be seen as the subspace obtained by the intersection of all closed half-spaces created by $d+1$ or more bisecting hyperplanes. A **d-polytope** $\mathcal{P}^{(d)}(S)$ is a bounded d-polyhedron. If the polytope is convex, it coincides with the **convex hull** $\mathcal{H}(S)$, the smallest convex set that contains S, obtained by the intersection of all half-spaces containing S.

A **radical hyperplane** is the hyperplane of equal powers with respect to two $(d-1)$-spheres $S_1(\mathbf{c}_1)$ and $S_2(\mathbf{c}_2)$, $\mathbf{c}_1 \neq \mathbf{c}_2$. With normal equations:

$$\mathbf{x^T x} - 2\mathbf{c_1^T x} + \|\mathbf{c_1}\|^2 - \varepsilon_1^2 = \mathbf{x^T x} - 2\mathbf{c_2^T x} + \|\mathbf{c_2}\|^2 - \varepsilon_2^2 \tag{14}$$

we obtain for the radical hyperplane:

$$H = \{\mathbf{x} \in E^d \parallel 2(\mathbf{c}_2 - \mathbf{c}_1)^{\mathbf{T}}\mathbf{x} - (\|\mathbf{c_2}\|^2 - \varepsilon_2^2 - \|\mathbf{c_1}\|^2 + \varepsilon_1^2) = 0\} \tag{15}$$

Observe in figure 2 (right picture, spherical segment through $\mathbf{p}$ and $\mathbf{p}'$), that H is perpendicular to the line segment connecting both centres. Also notice that $S_1 \cap S_2 \neq \emptyset \Rightarrow H \cap S_1 \neq \emptyset$ and $H \cap S_1 \neq \emptyset \Rightarrow H \cap S_2 \neq \emptyset$. It follows ([2],[4]) that in E^d, three $(d-1)$-spheres in general position span radical hyperplanes that intersect in a single $(d-3)$-flat, called the **radical centre** .

A **finite simplicial k-complex** $\mathcal{C}(S)$, with $S = \bigcup Y_j$ is a connected or unconnected collection of simplexes (or: simplices), with $0 \leq k \leq d$ and $\sigma^{(-1)} = \emptyset$, all satisfying two properties ([20],[14],[21],[15]):

$$\begin{cases} \sigma^{(k)}_{Y_j} \in \mathcal{C} & \Rightarrow \sigma^{(k-1)}_{Y'_j} \in \mathcal{C}' \subseteq \mathcal{C} \\ \sigma^{(k)}_{Y_j} \in \mathcal{C},\ \sigma^{(k-1)}_{Y'_j} \in \mathcal{C}' & \Rightarrow \mathcal{C} \cap \mathcal{C}' = \emptyset \quad \vee \\ & \qquad\quad \mathcal{C} \cap \mathcal{C}' = \mathcal{C}' \end{cases} \tag{16}$$

Apparently, every sub-face of a simplex needs to be in $\mathcal{C}$ too, and every sub-complex $\mathcal{C}' \subseteq \mathcal{C}$. A simplicial complex may degenerate to a point set, a linear graph, etc., having some highest dimension $m < d$.

A **d-triangulation** $\mathcal{T}^{(d)}(S)$ based on a finite point set $S \subseteq E^d$, $\mathbf{s}_i$ not all in a common hyperplane, is a subdivision of E^d in d-simplices, such that ([9]):

$$\begin{cases} \text{Compl Conv } S \text{ is the only unbounded } d\text{-cell in } \mathcal{T} \\ \bigcup \sigma^{(0)}_{\mathcal{T}} \cong \{\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_n\} = S \\ \text{every bounded } d\text{-cell in } \mathcal{T} \text{ is a } d\text{-simplex} \end{cases} \tag{17}$$

where *Conv* S is the convex set, spanning the convex hull $\mathcal{H}(S)$. *Compl* denotes the complement of a set.

An **α-complex** $\mathcal{C}_\alpha$ is a one-parameter finite simplicial k-complex. As such, it is also a sub-complex of the regular triangulation $\mathcal{T}$, which is a d-simplicial complex (figure 1). Let $\sigma_{Y_j}^{(k)} \in \mathcal{C}_\alpha$ denote a k-simplex belonging to the complex, with $Y_j \subseteq S$. Assume ([10]) that for every such $\sigma_{Y_j}^{(k)}$ there exists a $(k-k)$-flat (a point) that coincides with the radical centre $\mathbf{r}$. Further, we assume all subsets to be affinely independent. A k-simplex $\sigma_{Y_j}^{(k)}$ is **α-exposed** if it has a radical centre $\mathbf{r}_Y$ such that:

$$\|\tilde{\mathbf{y}}_k - \tilde{\mathbf{r}}_Y\| = \begin{cases} 0 & ,\forall \mathbf{y}_k \in Y_j \quad \text{(i)} \\ > 0 & ,\forall \mathbf{y}_k \in S - Y_j \quad \text{(ii)} \end{cases} \tag{18}$$

Clearly, (i) states the orthogonality condition whereas (ii) states the regularity condition. A k-simplex, with vertex set Y_j, $Card(Y_j) = k+1$, is called **regular** if $\|\tilde{\mathbf{s}}_i - \tilde{\mathbf{y}}_k\| > 0,\ \forall \mathbf{s}_i \in S - Y_j, \mathbf{y}_k \in Y_j$. A **weighted alpha complex** $\mathcal{C}_\alpha(S)$ is a sub-complex of the regular triangulation $\mathcal{T}(S)$, containing all α-exposed k-faces of $\mathcal{T}$. Notice that equation (18) defines exactly those faces of the regular triangulation $\mathcal{T}$ that are α-exposed. Also notice that if $\sigma_{Y_j}^{(k)}$ is made up of zero-weight points, then d-ball $B_\alpha(\mathbf{r}_Y)$ reduces to a ball $B_\alpha(\mathbf{r}'_Y)$, with $\tilde{\mathbf{r}}' = (\mathbf{r}', 0)$ and the $k+1$ vertices $\mathbf{y}_k \in \sigma_{Y_j}^{(k)}$ will be contained in the boundary $Bd\ B_\alpha(\mathbf{r}'_Y)$ and $Int\ B_\alpha(\mathbf{r}'_Y) = \emptyset$. The following also holds:

$$\alpha_i < \alpha_j \quad \Rightarrow \quad \mathcal{C}_{\alpha_i} \subseteq \mathcal{C}_{\alpha_j} \tag{19}$$

Recall that for sufficiently low a value of α, e.g, $\alpha \to 0$, the α-complex $\mathcal{C}_\alpha = S^2$. This complex is growing, with increasing value of α, into triangulation $\mathcal{T}(S)$, and the **underlying space** $|\mathcal{C}_\alpha|$ into the convex hull $Conv\ S = \mathcal{H}(S) = |\mathcal{C}_{\alpha\to\infty}|$.
The set A of all ordered pairs $(\varrho, \mathcal{C}_\alpha(S))$ denotes the finite set of α-complexes of set S, called the **α-family**, and a 2-tuple $\mathcal{A} = (\varrho_i, \mathcal{C}_{\alpha_i}(S)) \in A$ denotes a member of that set. Here, ϱ is a unique index into the α-family, such that:

$$\varrho_i < \varrho_j \quad \Rightarrow \quad \alpha_i < \alpha_j \tag{20}$$

and:

$$\varrho_i < \varrho_j \quad \Leftrightarrow \quad \mathcal{C}_{\alpha_i} \subset \mathcal{C}_{\alpha_j} \tag{21}$$

yields a proper sub-complex, whereas in general, equation (19) doesn't. α only allows for a partial ordering of the α-family, but ϱ allows for a strict partial ordering using a linear order relation "<". Refer to [20] for further details. Notice also that this definition allows us to topologically sort the set A, according to the index ϱ.
A sub-complex $\mathcal{C}' \subseteq \mathcal{C}$, with k-faces $k \leq l$, is called an **l-skeleton** $\mathcal{C}^{(l)}$. Generally, an α-complex contains l-skeletons, with $0 \leq l \leq d$. As a result, k-faces contained in α-complexes may be singular, regular or interior for a specific value of α. Let $0 \leq j < m = j+1 \leq d$ and let $\sigma^{(j)}, \sigma^{(m)} \in \mathcal{C}_\alpha$. Now we can tell singular k-faces apart from regular ones, as follows [12]:

$$\begin{cases} \sigma_{Y_j} \in Bd\ \mathcal{C}_\alpha \ \wedge\ \sigma_{Y_j} \notin \sigma_{Y_m} & \to \text{singular} \\ \sigma_{Y_j} \in Bd\ \mathcal{C}_\alpha \ \wedge\ \sigma_{Y_j} \in \sigma_{Y_m} & \to \text{regular} \end{cases} \tag{22}$$

[2] redundant vertices may have been dropped

A face in the interior of $\mathcal{C}_\alpha$, i.e., $\sigma_{Y_r} \in Int\ \mathcal{C}_\alpha$, is classified as *interior face*.

For simplicial complexes, (specifically in E^d) finding expressions for the number of k-faces is not trivial. However, a fully triangulated single d-skeleton serves as an upper bound: with $\alpha \to \infty \Rightarrow O(\mathcal{C}_\alpha) \to O(\mathcal{T})$, where $\mathcal{T}$ is a triangulation of point set S. Refer to Edelsbrunner, [10],[12], Delfinado and Edelsbrunner, [7], Gross and Tucker, [15].

Weights can be generated on a per point basis, they can also be generated as a function $f : E^d \to R$. Weights are called **unstructured weights** if two arbitrary functions values (function images) $f(\mathbf{x}_1)$ and $f(\mathbf{x}_2)$, $\mathbf{x}_1, \mathbf{x}_2 \in E^d$ have auto-covariance $cov_{xx}(\mathbf{x}_1, \mathbf{x}_2) \to 0$.

References

[1] Oswin Aichholzer, Franz Aurenhammer, D.T. Lee, Danny Z. Chen, and Evanthia Papadopoulou. Skew voronoi diagrams. *to appear in: International Journal of Computational Geometry & Applications*, 1999.

[2] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM J. Comput.*, 16(1):78–96, February 1987.

[3] F. Aurenhammer and H. Imai. Geometric relations among Voronoi diagrams. In *Proc. 4th Sympos. Theoret. Aspects Comput. Sci.*, volume 247 of *Lecture Notes Comput. Sci.*, pages 53–65. Springer-Verlag, 1987.

[4] Wolfgang Boehm and Hartmut Prautzsch. *Geometric Concepts for Geometric Design.* A. K. Peters, Ltd., Wellesley, MA, 1994.

[5] Louis Brand. *Vector and Tensor Analysis.* John Wiley & Sons, Inc., London, ninth printing edition, 1947.

[6] Julian L. Coolidge. *A treatise on the circle and the sphere.* Oxford at the Clarendon Press, London, 1916.

[7] C. J. A. Delfinado and H. Edelsbrunner. An incremental algorithm for betti numbers of simplicial complexes. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 232–239, 1993.

[8] H. Edelsbrunner and N. R. Shah. Incremental topological flipping works for regular triangulations. *Algorithmica*, 15:223–241, 1996.

[9] Herbert Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer-Verlag, Berlin-Heidelberg, 1987.

[10] Herbert Edelsbrunner. 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] Herbert Edelsbrunner, Michael Facello, and Jie Liang. On the definition and the construction of pockets in macromolecules. *Discrete Applied Mathematics*, 88:83–102, 1998.

[12] Herbert Edelsbrunner and Ernst P. Muecke. Three-dimensional alpha shapes. *ACM Transactions on Graphics*, 13:43–72, 1994.

[13] Bart H. M. Gerritsen. On the use of alpha complexes in subsurface modelling. In *Proc. ENSG Conference on 3D Modeling of Natural Objects, A Challenge for the 2000's.* ENSG, Nancy, France, June 4–5 1998.

[14] P.J. Giblin. *Graphs, Surfaces and Homology; An Introduction to Algebraic Topology.* Chapman and Hall Mathematics Series. Chapman and Hall, London, 1977.

[15] Jonathan L. Gross and Thomas W. Tucker. *Topological Graph Theory.* John Wiley & Sons, New York, 1987.

[16] Liang-shin Hahn. *Complex Numbers & Geometry.* Spectrum. The Mathematical Association of America, Washington DC, 1994.

[17] Yehuda E. Kalay. *Modeling Objects and Environments.* Principles of Computer-Aided Design (Series Ed. Yehuda E. Kalay). John Wiley & Sons, New York, 1989.

[18] C. W. Lee. Regular triangulations of convex polytopes. In *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, volume 4 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 443–456. AMS Press, 1991.

[19] E.A. Maxwell. *Elementary Coordinate Geometry.* Oxford at the Clarendon Press, 1952.

[20] James R. Munkres. *Topology, a first Course.* Prentice-Hall, Inc., Englewood Cliffs, New Jersey, third edition, 1975.

[21] James R. Munkres. *Elements of Algebraic Topology.* Addison-Wesley, Redwood, CA, 1984.

[22] A. Paoluzzi, F. Bernardini, C. Cattani, and V. Ferrucci. Dimension-independent modelling with simplicial complexes. *ACM Transactions on Graphics*, 12(1):56–102, January 1993.

[23] A.A.G. Requicha. Representation of rigid solid objects. In Jose Encarnacao, editor, *Computer Aided Design; Modelling, systems engineering, CAD systems. CREST Advanced Coarse, Darmstadt, September 1980*, Lecture Notes in Computer Science (series eds. G. Goos and J. Hartmanis). Springer-Verlag, Berlin-Heidelberg, 1980.

[24] Aristides A. G. Requicha. Toward a theory of geometric tolerancing. *Int. J. Robot. Res.*, 2(4):45–60, Winter 1983.

[25] Hans Schwerdtfeger. *Geometry of complex numbers; Circle geometry, Moebius transformation, non-euclidean geometry.* Dover Publications, Inc., New York, dover edition, 1979.

[26] G. Smets, M. de Groof, J. Nuyts, D. van der Meulen, P. Suetens, G. Marchal, and A. Oosterlinck. Interpretation of $3d$ medical scenes. In Herbert Freeman, editor, *Machine Vision for Three-Dimensional Scenes*, pages 163–193. Academic Press, Inc., Boston, 1990.

[27] Dean L. Taylor. *Computer-Aided Design.* Addison-Wesley series in Mechanical Engineering. Addison-Wesley Publishing Company, Inc., Reading, Massachusetts, 1992.

[28] Michael F. Worboys. A generic model for planar geographical objects. *International Journal of Geographical Information Systems*, 6(5):353–372, 1992.

[29] Ibrahim Zeid. *CAD/CAM Theory and Practice.* McGraw-Hill Series in Mechanical Engineering (series eds. Jack P. Holman and John R. Lloyd). McGraw-Hill, Inc., New York, 1991.