Centrum voor Wiskunde en Informatica
Centre for Mathematics and Computer Science

R.C. Veltkamp

The γ-neighborhood graph

Computer Science/Department of Interactive Systems Report CS-R9034 July

---

The Centre for Mathematics and Computer Science is a research institute of the Stichting Mathematisch Centrum, which was founded on February 11, 1946, as a nonprofit institution aiming at the promotion of mathematics, computer science, and their applications. It is sponsored by the Dutch Government through the Netherlands Organization for the Advancement of Research (N.W.O.).

The γ-Neighborhood Graph

REMCO C. VELTKAMP
CWI, Department of Interactive Systems
Kruislaan 413, 1098 SJ Amsterdam, The Netherlands; e-mail: remco@cwi.nl

Abstract

This report presents a novel two-parameter geometric graph, the γ-neighborhood graph. This graph unifies a number of geometric graphs such as the convex hull, the Delaunay triangulation, and in 2D also the Gabriel graph, and the circle-based β-skeleton, into a continuous spectrum of geometric graphs that ranges from the void, to the complete graph. The two parameters provide for a great flexibility in the analysis of a set of sites. For specific ranges of the parameters, the corresponding graph can be efficiently constructed.

1987 CR Categories:
F.2.2 [Analysis of Algorithms and Problem Complexity] Nonnumerical Algorithms and Problems
G.2 [Discrete Mathematics]
I.3.5 [Computer Graphics] Computational Geometry and Object Modeling
Key words and phrases: Computational Geometry, Neighborhood Graphs

This research was done while the author was working at Leiden University, The Netherlands.

1 Introduction

In the computational geometry discipline, old and new geometric techniques are brought together and unified. An example of this is the development in geometric graphs. A major unifying effect in computational geometry was brought about by the Delaunay triangulation [Del28, Del34], and its dual Voronoi diagram [Vor08]. Old geometric graphs such as the convex hull and the Euclidean minimum spanning tree, and new, parameterized graphs such as the α-shape [EKS83] and the β-skeleton [KR85] are intimately related to the Delaunay triangulation. An even more general graph is presented in this paper: the γ-neighborhood graph. It is a two-parameter graph, unifying the Delaunay triangulation, convex hull and the β-skeleton into a continuous spectrum of geometric graphs ranging from the void to the complete graph.

In [KR85] it is said that a geometric graph describes the internal shape of a set of sites, when it connects essential neighbors among the essential sites. The external shape is described when the graph connects essential neighbors among the essential extreme sites. In which way sites or pairs of sites are essential is determined by the definition of the graph or, when appropriate, the neighborhood. It will be shown that the γ-neighborhood graph can describe the internal, as well as the external shape.

Report CS-R9034
Centre for Mathematics and Computer Science
P.O. Box 4079, 1009 AB Amsterdam, The Netherlands

Before the γ-neighborhood graph is introduced in section 3, the next section gives an introductory overview of geometric graphs and some relations between them. Section 4 gives some examples of the γ-neighborhood graph for specific choices of the parameters. Section 5 deals with the computation of the γ-neighborhood graph, and the resulting computational complexities. Section 6 finally, shortly explains how the graph is used in computational morphology, specifically in finding an external shape for a set of sites.

2 Overview of geometric graphs

Some of the graphs mentioned in this section are truly proximity or neighborhood graphs. Such graphs join two sites by an edge when a certain neighborhood is empty. The neighborhood is called empty if no sites lie in its interior, except when an entire half-space is involved. A half-space with its boundary through $x_1, \ldots, x_k$, is called empty if no sites lie in the open half-space; sites may lie on the boundary of the half-space except inside or on the $(k-1)$D polygon through $x_1, \ldots, x_k$. For example in 2D, a half-plane through x_1 and x_2 is empty if the open half-plane is empty, and no sites lie on the line segment between x_1 and x_2 (but sites may lie on the boundary of the half-plane outside that segment). A sphere of infinite radius is also considered a half-space.

In the following, all distances are Euclidean, and N denotes the number of sites.

2.1 Definitions

Closest pairs (CP). The closest pairs of a set of sites are the pairs of sites that have the smallest distance to each other, among all pairs. Note that there can be more than one closest pair. CP is disconnected, except for $N = 2$, or when all sites are equidistant.

Nearest neighbors graph (NNG). In the nearest neighbor graph, each site is connected to the site that is nearest. Since all the pairs of sites that are each others nearest neighbor contain the pairs with the smallest distance of all, CP $\subseteq$ NNG. In general NNG is disconnected.

Euclidean minimum spanning tree (EMST). This graph is connected and has no cycles. EMST is spanning in the sense that it connects all sites, and minimum with respect to the sum of all Euclidean distances between connected sites. In EMST, each site is connected to its nearest neighbor, and thus NNG $\subseteq$ EMST. NNG actually is a minimum spanning forest, so in the special case that it is connected, it coincides with EMST.

Infinite strip graph (∞-SG). Two sites x and y are connected if and only if the infinite strip bounded by two parallel planes through x and y, that are perpendicular to $x-y$, is empty. The EMST must also connect such a pair of sites in order to be connected, except when two strips coincide for different pairs of sites. So in non-degenerate cases ∞-SG $\subseteq$ EMST.

Sphere of influence graph (SIG). The SIG is introduced by Toussaint [Tou88]. For each site x, let r_x be the distance to the closest site. The sphere of influence graph connects two sites x and y, if and only if the spheres with radii r_x and r_y, and centered at x and y respectively, intersect in more than one point. Clearly, each site is connected to its nearest neighbor, so that NNG $\subseteq$ SIG.

Relative neighborhood graph (RNG). Two sites x_i and x_j are defined to be relatively close if

$$d(x_i, x_j) \leq \max\{d(x_i, x_l), d(x_j, x_l)\}, \text{ for all } l \neq i, j.$$

For arbitrary x_l, $d(x_i, x_j) \leq \max\{d(x_i, x_l), d(x_j, x_l)\}$, if x_l does not lie in the interior of the intersection of the two spheres centered at x_i and x_j and with radius $d(x_i, x_j)$. The intersection is called the relative neighborhood, and two sites are relative neighbors, if their relative neighborhood is empty. The RNG connects all relative neighbors. In the original definition by [Lan69], the "$\leq$" is replaced by a "$<$", but the above definition has become common in computational geometry [Tou80], and corresponds to the notion of "empty neighborhood". It is shown in [Tou80] that EMST $\subseteq$ RNG.

Gabriel graph (GG). The Gabriel neighborhood (named after [GS69]) of two sites x and y is the smallest sphere through x and y, which has radius $d(x, y)/2$. Because the Gabriel neighborhood is contained in the relative neighborhood, it is empty when the latter is empty, and therefor RNG $\subseteq$ GG.

Convex hull (CH). The convex hull of a set of sites is the smallest polytope containing all the sites. Indeed this polytope is convex. The CH connects k sites with each other if a half-space with its boundary through these sites is empty in the sense stated at the beginning of this section. As a result, if $k + 1$ or more sites on the CH lie in a plane, the faces are kept $(k - 1)$-simplicial (a simplex or k-simplex is the k-dimensional analogue of the triangle in the plane and the tetrahedron in 3-space, having $k + 1$ vertices).

Delaunay triangulation (DT). In [Vor08] a partitioning of space into simplices L_i is defined, which is therefore called an L-subdivision or L-partition. A definition of the L-subdivision given by Delaunay [Del28, Del34], defines a simplex to be part of the L-subdivision if the sphere through its vertices, which are some sites from S, contains no other sites. The L-subdivision is therefore called Delaunay triangulation, or sometimes closest point Delaunay triangulation. In three-dimensional space, we can call this a Delaunay tetrahedralization, although in general k-space this subdivision is still called a triangulation.

In the case that more than $k + 1$ sites lie on a sphere, connecting all these sites with each other would generate overlapping simplices. Instead, the DT arbitrarily connects sites to generate non-overlapping simplices that fill the space enclosed by the convex hull of these sites. A degenerate Delaunay triangulation is therefore not unique.

Clearly an empty sphere passes through the end-points of each edge in a Delaunay triangulation. Conversely, if an empty sphere passes through two sites, then there is a largest possible empty sphere through these two sites. This sphere either passes through $k - 1$ other sites, or through $k - 2$ other sites and has an infinite radius. In the latter case, the "sphere" is a half-space, and the two end-points lie on the convex hull. In both cases the two sites form an edge in the Delaunay triangulation. It follows immediately that CH $\subseteq$ DT and GG $\subseteq$ DT.

Because the sphere through the vertices of a simplex in the DT is empty, the two spheres through k sites of any simplex, can be regarded as "the Delaunay neighborhood".

α-shape. The notion of a parameterized generalized disc is introduced in [EKS83]. A generalized disc of radius $1/\alpha$ is defined as a disc of radius $1/\alpha$ if $\alpha > 0$, the complement of a disc of radius $1/(-\alpha)$ if $\alpha < 0$, and a half-space if $\alpha = 0$. For an arbitrary real α and a set S of sites in the plane, the α-hull of S is defined as the intersection of all closed generalized discs of radius $1/\alpha$ that contain all the sites of S. Replacing circular arcs on the boundary of the α-hull by straight edges gives the α-shape. This definition is easily generalized to higher dimensions, replacing disks by spheres, and straight edges by flat triangles. The α-shape is a subgraph of the closest point Delaunay triangulation, if $\alpha \geq 0$, and a subgraph of the so-called furthest point Delaunay triangulation if $\alpha \leq 0$. The 0-shape coincides with the convex hull.

β-skeleton. The β-skeleton is a planar parameterized graph, introduced in a lune-based and a circle-based variant [KR85]. The following definition is a slightly modified version of the original, in order to normalize the parameter to lie between -1 and 1. The lune-based β-neighborhood for two sites x and y is

1. the intersection of two circles of radius $d(x,y)/2(1+\beta)$ that pass through x and y, if $\beta \in [-1,0]$,
2. the intersection of two circles of radius $d(x,y)/2(1-\beta)$ centered at the points $x+(y-x)\beta/2$ and $y+(x-y)\beta/2$, if $\beta \in [0,1]$.

The circle-based β-neighborhood for two sites x and y is

1. the intersection of two circles of radius $d(x,y)/2(1+\beta)$ that pass through x and y, if $\beta \in [-1,0]$,
2. the union of two circles of radius $d(x,y)/2(1-\beta)$, that pass through x and y, if $\beta \in [0,1]$.

The β-skeleton connects those sites whose β-neighborhood is empty. When $\beta = 0$, both the lune-based and the circle-based neighborhood coincide with the Gabriel neighborhood. Both β-neighborhoods contain the Gabriel neighborhood when $\beta < 0$, so that the corresponding skeletons are contained in the GG. When $\beta = 1/2$ the lune-based β-neighborhood reduces to the relative neighborhood. When $\beta = 1$, the lune-based skeleton reduces to ∞-SG, and the circle-based skeleton to the void graph. For $\beta = -1$, both skeletons reduce to the complete graph if no three sites are collinear. The spectrum of β-neighborhoods for the whole range of the parameter is illustrated in figure 1. The generalization of the lune-based β-neighborhood to higher dimensions is straightforward. Nothing is said in [KR85] about a higher dimensional circle-based β-neighborhood. The lune-based β-skeleton is used for the analysis of empirical networks.

The convex hull and its parameterized generalization, the α-shape, describe aspects of the external structure of a set of sites. All other geometric graphs mentioned here describe different aspects of the internal structure. The inclusion relations between all these graphs is depicted in figure 2.

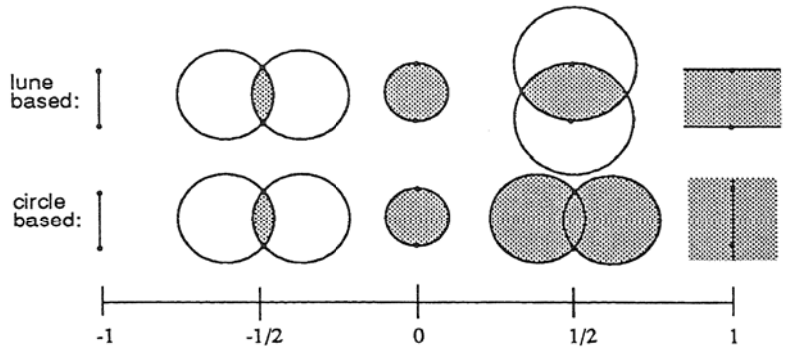

Figure 1: Overview of the spectrum of planar β-neighborhoods.

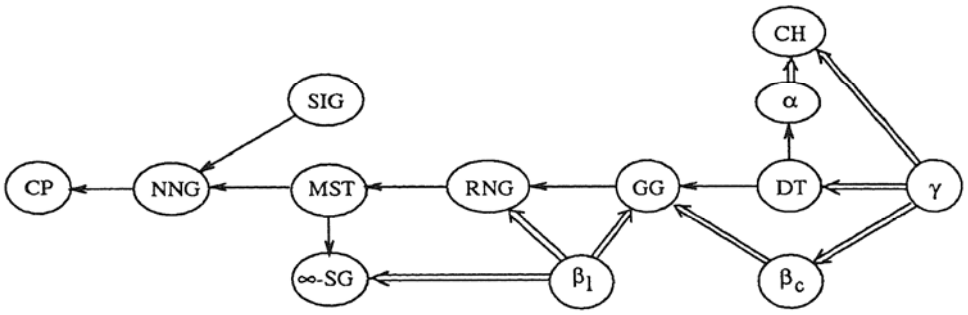

Figure 2: Dependencies between geometric graphs. α represents the α-shape, β_l the lune-based, β_c the circle-based β-skeleton, and $\gamma$the γ-graph. The void and the complete graph are omitted. Graph1 $\leftarrow$ graph2 denotes graph1 $\subseteq$ graph2, and graph1 $\Leftarrow$ graph2 indicates that the parameterized graph1 reduces to graph2 for specific parameter values.

2.2 Computational complexities

The following table lists the time complexities to compute the graphs, and the references where these results can be found.

	$2D$		kD	
	upper bound	reference	upper bound	reference
CP	$\Theta(N \log N)$	[BS76]	$\Theta(N \log N)$	[BS76]
NNG	$\Theta(N \log N)$	[BS76]	$O(N(\log N)^{k-1})$	[BS76]
EMST	$\Theta(N \log N)$	[SH75]	$O(N^2)$	[Pri57]
∞-SG	$O(N^2)$		$O(N^2)$	
SIG	$\Theta(N \log N)$	[Tou88]	$O(N^2)$	
RNG	$\Theta(N \log N)$	[Sup83]	$O(N^3)$	[Tou80]
GG	$\Theta(N \log N)$	[MS80]	$O(N^3)$	
DT	$\Theta(N \log N)$	[LS80]	$O(N^{1+\lceil k/2 \rceil})$	[Bro79] + [Sei81]
CH	$\Theta(N \log N)$	[PH77]	$O(N^{\lceil (k+1)/2 \rceil})$	[Sei81]
α-shape	$\Theta(N \log N)$	[EKS83]	$O(N^{1+\lceil k/2 \rceil})$	[EKS83]
β_l-skeleton	$O(N^2)$, $O(N^3)$	[KR85]	$O(N^3)$	
β_c-skeleton	$O(N \log N)$, $O(N^3)$	[KR85]		

I do not know a reference for the ∞-SG. But since ∞-SG $\subseteq$ EMST, one can examine each of the $N-1$ edges in the EMST, and check if any site lies in the infinite strip. This check takes $O(N)$ time, giving a total of $O(N^2)$.

The higher dimensional SIG can be constructed after computing the NNG, by examining each of the $O(N^2)$ pairs of sites in constant time, yielding $O(N^2)$ total time. The higher dimensional GG can be constructed by taking all $O(N^2)$ pairs of sites, and examine each neighborhood for inclusion of all $N-2$ other sites. This results in $O(N^3)$ time.

The planar lune-based β-skeleton (denoted by β_l in the table above) is a subgraph of the GG when $\beta \geq 0$. Examining all the $O(N)$ edges of the GG takes $O(N)$ time per edge, giving a total of $O(N^2)$. The planar circle-based β-skeleton (β_c) is also a subgraph of the GG when $\beta \geq 0$. Checking whether the Delaunay neighborhood of each edge in the GG contains the circle-based β-neighborhood takes constant time, giving order $O(N \log N)$. The $O(N^3)$ time applies to the case $\beta < 0$, and results from a brute force algorithm, just as the $O(N^3)$ for the higher dimensional lune-based β-skeleton (remember that the circle-based β-skeleton has no direct higher dimensional analogue).

3 The γ-neighborhood graph

In this section the γ-neighborhood graph is defined for arbitrary dimension k. In the following we will use 'γ-graph', '$\gamma(\gamma_0, \gamma_1)$-graph', or simply '$\gamma(\gamma_0, \gamma_1)$' and similar expressions, to denote the appropriate neighborhood graph.

In the definition of the γ-graph we use the following notation: for $k \geq 2$, $r(x_1, \ldots, x_k)$ denotes the radius of the smallest sphere through sites $x_1, \ldots, x_k$ in kD space. Thus for $k = 2$, $r(x_1, x_2)$ equals $d(x_1, x_2)/2$.

The neighborhood graph $\gamma(\gamma_0, \gamma_1)$ is defined for $-1 \leq \gamma_0, \gamma_1 \leq 1$, and $|\gamma_0| \leq |\gamma_1|$. In kD space, the graph connects sites $x_1, \ldots, x_k$ pairwise ($k(k-1)/2$ edges) if an empty neighborhood $\mathrm{N}(\gamma_0, \gamma_1)$ is associated with these sites, that is defined by two kD spheres in the following way:

1. the spheres have radii $r(x_1,\ldots,x_k)/(1-|\gamma_0|)^{1/(k-1)}$ and $r(x_1,\ldots,x_k)/(1-|\gamma_1|)^{1/(k-1)}$,
2. if $\gamma_0\gamma_1 < 0$, the centers of the spheres lie on the same side of the plane through $x_1,\ldots,x_k$, if $\gamma_0\gamma_1 > 0$ the centers lie on both sides of that plane,
3. if $\gamma_1 \leq 0$, we take the intersection of the two spheres, if $\gamma_1 \geq 0$, we take the union.

Note that there can be two neighborhoods if γ_0 and γ_1 are both non-zero, and that the graph connects $x_1,\ldots,x_k$, as soon as one neighborhood is empty. Note further that this definition is valid for $\gamma_0 = \gamma_1 = 0$. In that case, the two spheres coincide, their common center lies in the plane through $x_1,\ldots,x_k$, and the intersection equals the union.

For $k = 2$ the definition involves two sites and two circles, and $r(x_1,x_2)$ is scaled by factors $1/(1-|\gamma_0|)$ and $1/(1-|\gamma_1)|$. The planar $\gamma(\gamma_0,\gamma_1)$ reduces to well known geometric graphs for special values of γ_0 and γ_1:

- $\gamma_0 = \gamma_1 = 0$. The resulting neighborhood $N(0,0)$ is the smallest circle through x_1 and x_2, which is the Gabriel neighborhood. $\gamma(0,0)$ is the Gabriel graph.
- $\gamma_0 = \gamma_1 = -1$. The intersection of the two half-planes yields the line through x_1 and x_2. If no three or more sites are collinear, then $\gamma(-1,-1)$ is the complete graph.
- $\gamma_0 = \gamma_1 = 1$. The union of the two half-planes gives the entire plane. If no three or more sites are collinear, $\gamma(1,1)$ is a void graph.
- $\gamma_0 = -1$, $\gamma_1 = 1$ and $\gamma_0 = 1$, $\gamma_1 = -1$. In both cases the two half-planes lie on the same side of the line through x_1 and x_2. They therefore coincide (more general, $\gamma(\gamma_0,-\gamma_0) = \gamma(-\gamma_0,\gamma_0)$). The neighborhood is empty if all other sites lie on one side of the line through x_1 and x_2, or on the line, but outside the segment from x_1 to x_2. That occurs only if x_1 and x_2 lie on the convex hull. Therefore, $\gamma(-1,1)$ and $\gamma(1,-1)$ are the convex hull of the set of sites.
- $\gamma_0 = \gamma_1$. The graph $\gamma(\gamma_0,\gamma_0)$ reduces to the circle-based β-neighborhood graph.

The relation of the planar γ-graph with other geometric graphs is depicted in figure 2. Figure 3 gives a graphical overview of the whole spectrum of planar neighborhoods.

The power $1/(k-1)$ in the denominator has no essential importance, and could be omitted without changing the concept of the γ-graph. The use of it is made clear in [Vel89].

In 3D space, the definition of the neighborhood involves three sites and two spheres, and $r(x_1,x_2,x_3)$ is scaled by factors $1/(1-|\gamma_0|)^{1/2}$ and $1/(1-|\gamma_1|)^{1/2}$. Because the Gabriel neighborhood is the smallest sphere trough two sites, whereas the kD $N(0,0)$ involves k sites, these neighborhoods do not coincide for $k > 2$. Further, it is not clear how to relate the higher dimensional $N(\gamma_0,\gamma_0)$ to the circle-based β-neighborhood, since Kirkpatrick and Radke [KR85] do not tell how to generalize it to higher dimensions. Generalization of the lune based neighborhood is straightforward, involving two sites in any dimension. The kD complete and the void graph however, result from the γ-graph if no $k+1$ or more sites lie in a $(k-1)$D plane. The convex hull equals $\gamma(-1,1)$ and $\gamma(1,-1)$ in any dimension. Again, $\gamma(\gamma_0,-\gamma_0) = \gamma(-\gamma_0,\gamma_0)$.

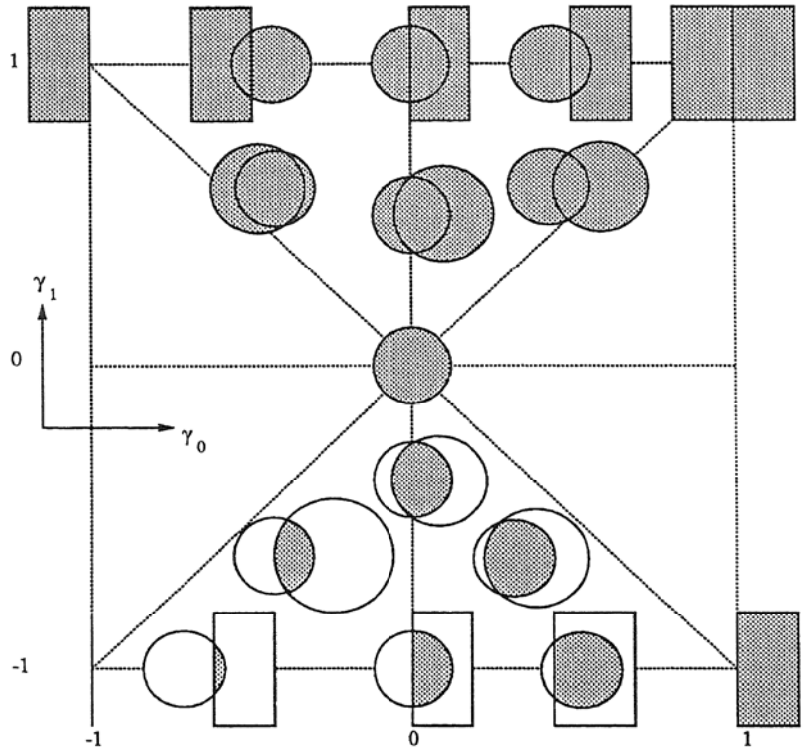

Figure 3: Overview of the spectrum of planar $\gamma(\gamma_0, \gamma_1)$-neighborhoods, $-1 \leq \gamma_0, \gamma_1 \leq 1$, and $-1 \leq \gamma_0/\gamma_1 \leq 1$ (rectangles denote half-spaces).

So far we have considered fixed values of the γ-parameters. We can also look at the largest values of the γ-parameters, for which the corresponding neighborhood is still empty. That is the value for which the sphere touches a $(k+1)$th site, or is either 1 or -1 if there is no such site. We define $\gamma([\gamma_0, \gamma_1], [\gamma_2, \gamma_3])$ to be the graph connecting sites $x_1, \ldots, x_k$ with each other, if the largest γ-parameter values for which the corresponding neighborhood is still empty, lie in $[\gamma_0, \gamma_1]$ and $[\gamma_2, \gamma_3]$ respectively.

The $\gamma([-1, 1], [0, 1])$-graph connects sites $x_1, \ldots, x_k$ in kD space if there are two spheres through these sites, of arbitrary radius, such that the union is empty. This is exactly a definition of the Delaunay triangulation, if no more than $k + 1$ sites are cospherical. If there are more than $k + 1$ cospherical sites, $\gamma([-1, 1], [0, 1])$ connects them all, whereas the Delaunay triangulation arbitrarily connects k sites, as long as the resulting $(k - 1)$D-faces do not intersect.

The γ-graph describes the internal structure of a set of sites. But it also describes aspects of the external structure. For example, the $\gamma(-1, 1)$-graph reduces to the convex hull. Also, the next section gives an example in which special γ-parameter values give a clear external structure. And section 6 will show how it is used to find a boundary through all sites.

This capability of external structure description somewhat contrasts to the lune based β-skeleton. In this graph, the neighborhood is located between the two sites. The resulting graph therefore emphasizes connections between sites, which makes it suitable for network analysis. With the γ-graph, the spheres are located aside the k involved sites. Especially

when the neighborhood is the union of these spheres, the γ-graph is more like (a part of) a tesselation.

4 Examples

The types of graphs that result from specific choices of the parameters is most clearly demonstrated with planar graphs. Figures 4, 5, and 6 show planar γ-graphs on the same set of 20 sites.

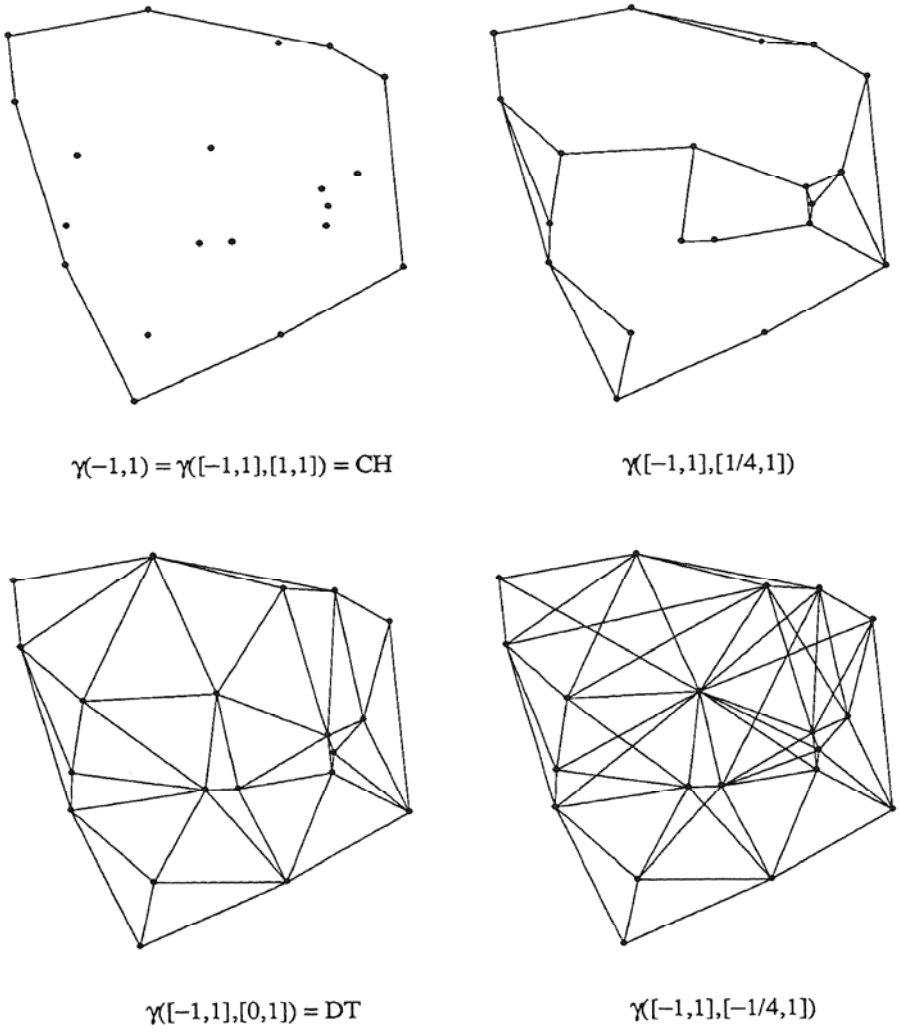

Figure 4: Planar γ-graphs on the same set of sites, containing more and more edges.

Figure 4 shows a sequence of $\gamma([-1,1],[\gamma_0,1])$-graphs. For $\gamma_0 = 1$, this yields the convex hull. Lowering γ_0 introduces more and more edges in the graph, until for $\gamma_0 = 0$ the graph coincides with the Delaunay triangulation. When γ_0 gets negative, edges cross each other. For $\gamma_0 = -1$, the graph would be complete (not shown).

The graphs in figure 5 all coincide with a circle-based β-skeleton. The neighborhoods $N(0.2, 0.2)$, $N(0.1, 0.1)$, $N(0, 0)$, and $N(-0.2, -0.2)$ get smaller and smaller. The emptiness requirement gets less restrictive, so that more pairs of sites are considered as neighbors. The $\gamma(0,0)$-graph equals the Gabriel graph.

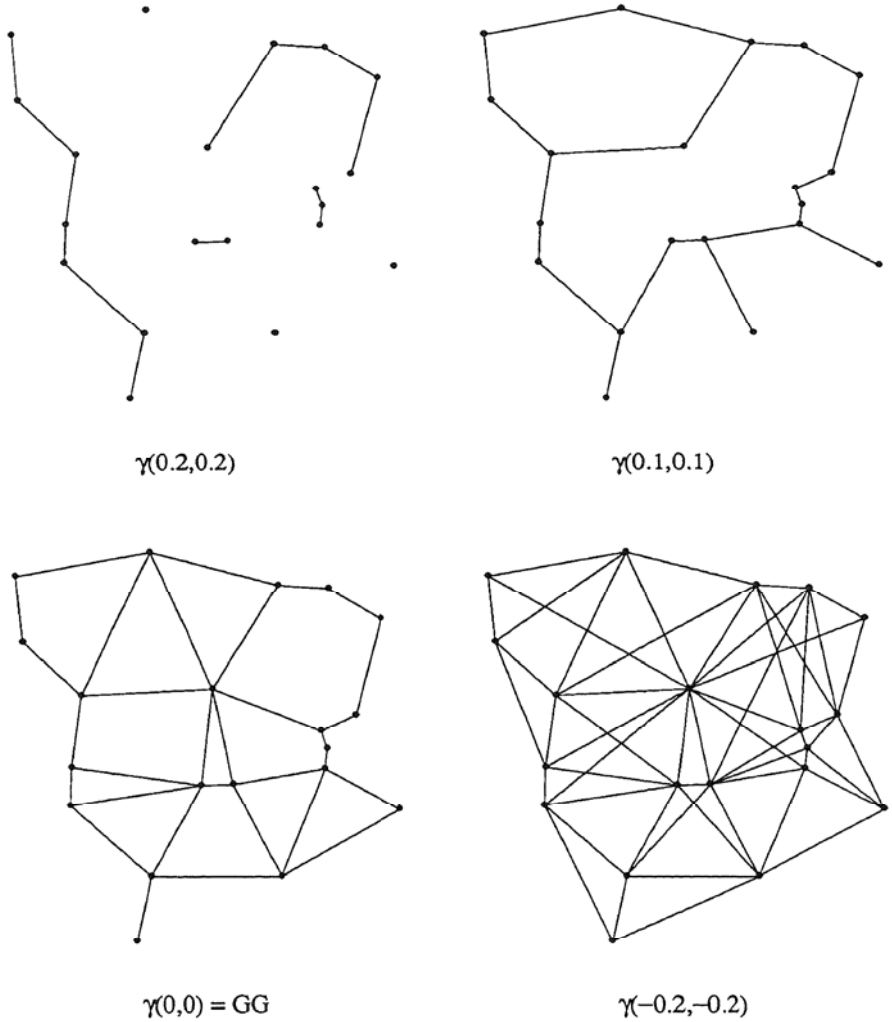

Figure 5: A sequence of planar γ-graphs on the same set of sites, reducing to β-skeletons.

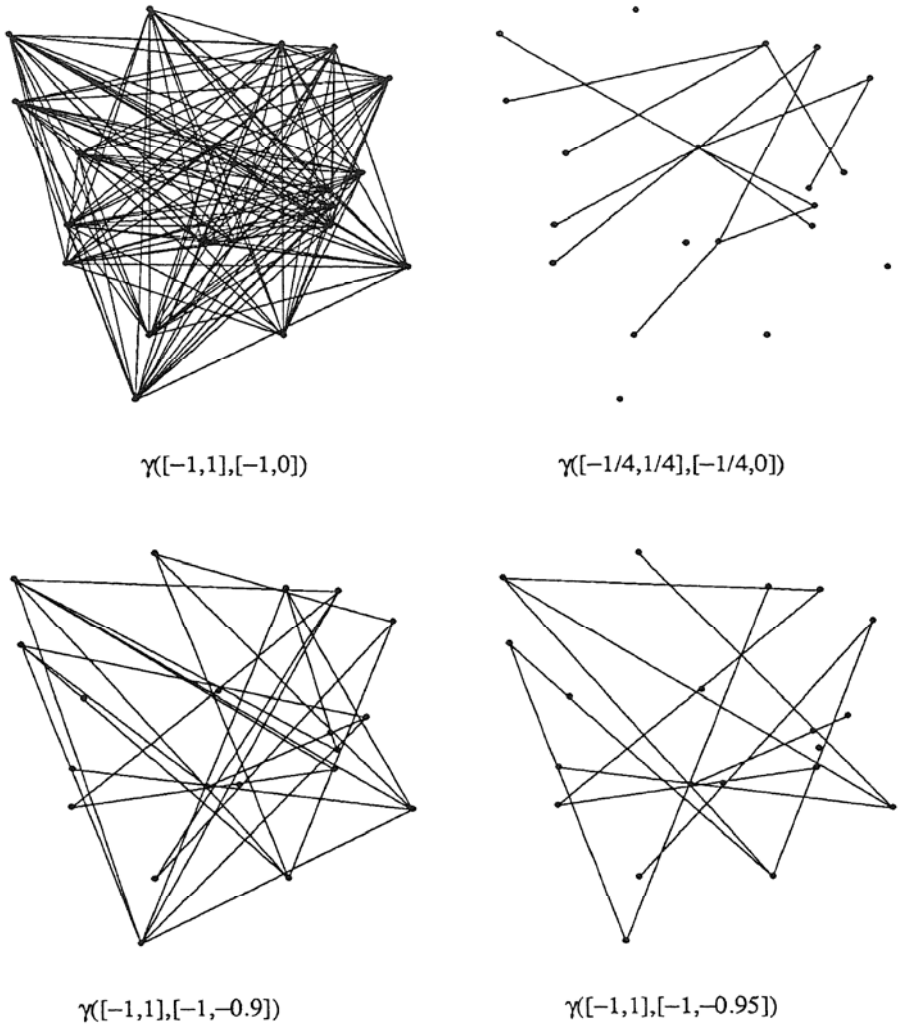

Figure 6: A sequence of planar γ-graphs on the same set of sites. The neighborhoods consist of the intersection of circles.

Figure 6 depicts graphs that result when only intersections of circles are allowed as neighborhood. The graph $\gamma([-1,1],[-1,0])$ connects all pairs of sites that have no empty circle through themselves. It it the complement of the Delaunay triangulation. In $\gamma([-1/4,1/4][,-1/4,0])$, the intersections are forced to have a certain fatness. On the other hand, $\gamma([-1,1],[-1,-0.9])$ allows only thin intersections. The edges now join sites only when there is another site close to the edge. In $\gamma([-1,1],[-1,-0.9])$ the neighborhoods are so thin, that only sites are connected if there is another site almost on the edge.

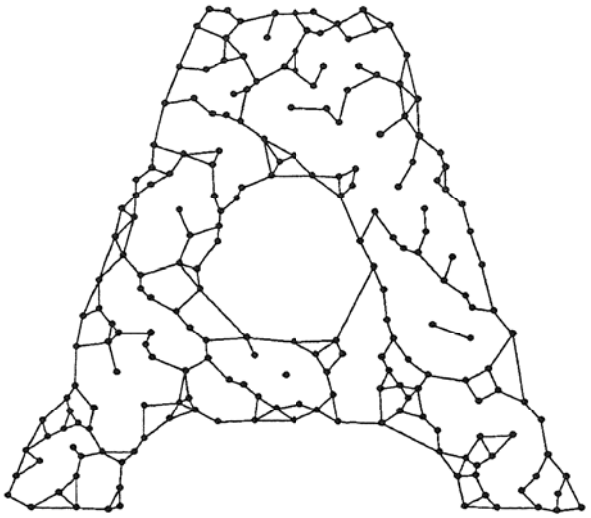

γ(–0.15,0.3)

Figure 7: The γ-graph describing an external structure on the planar set of sites after [EKS83].

Figure 7 shows a set of sites similar to the one used in [EKS83] to illustrate the α-shape. The $\gamma(-0.15, 0.3)$-graph turns out to give a clear boundary, although also internal sites are connected. The α-shape, designed to give the boundary of a cluster of sites, yields a single inner and outer contour. However, the two γ-parameters give more freedom for finding some external structure than the single parameter β-skeleton (see [KR85] for a β-skeleton on the set of sites from [EKS83]).

Figure 8 at last, shows projections of two stereo pairs of 3D γ-graphs on the same set of 30 sites. The 3D $\gamma(0, 0)$-graph connects three sites with each other if the smallest sphere through these sites is empty. Note that in 3D, this differs from the Gabriel graph, which connects two sites if the smallest sphere through these sites is empty. The picture of the $\gamma(0, 0)$-graph only slightly differs from a typical 3D Delaunay triangulation. This is because the triangles that belong to the $\gamma(0, 0)$ but not to the Delaunay triangulation, can have edges that are also edges of other triangles in the $\gamma(0, 0)$. All edges of a triangle can thus be displayed, while the triangle does not belong the graph. In this example, the Delaunay triangulation consists of 257 triangles (constituting 248 tetrahedra), and the $\gamma(0, 0)$ of only 150 triangles. The $\gamma([-1/4, 1/4], [1/2, 3/4])$-graph is disconnected, but shows that always three sites are connected to each other, if at all.

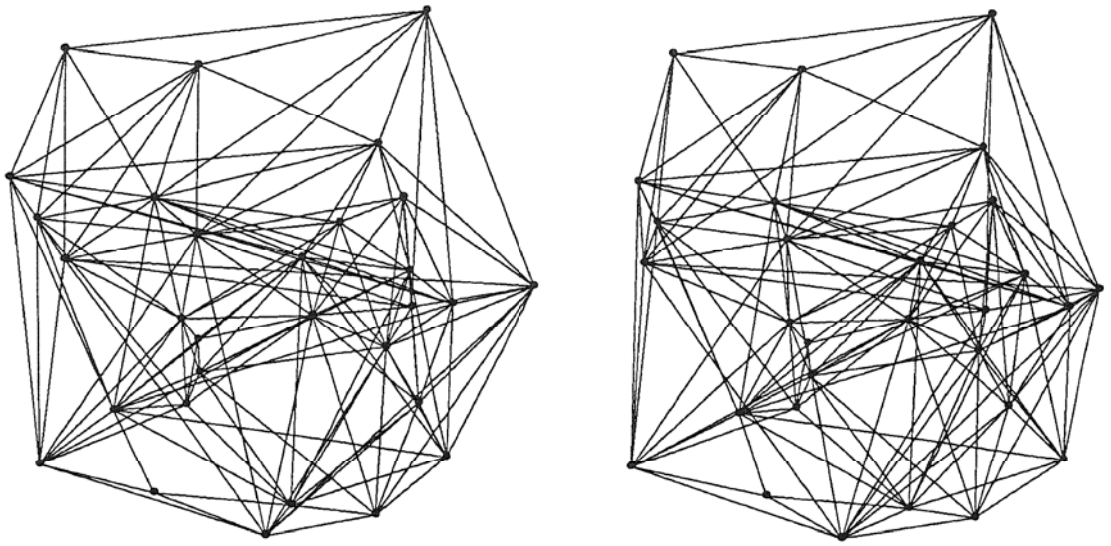

$\gamma(0,0)$

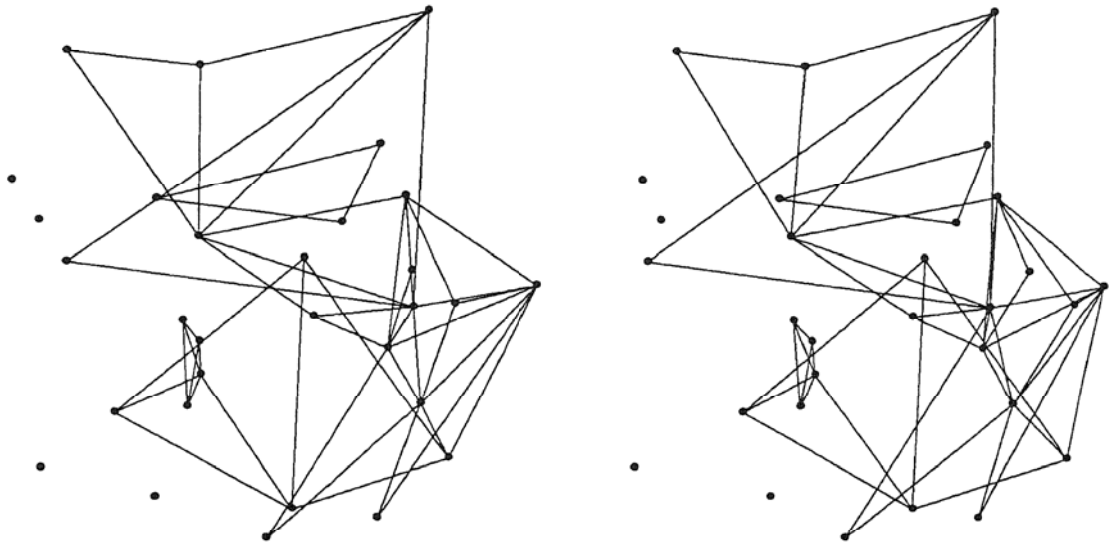

$\gamma([-1/4,1/4],[1/2,3/4])$

Figure 8: Two stereo-pairs of perspectively projected 3D γ-graphs.

5 Complexity issues

The following three lemmas tell how γ-graphs are related to each other. They give cues how to construct an arbitrary γ-graph. The two subsequent theorema tell how efficiently they can be computed.

Lemma 1. $\gamma(\gamma_0, \gamma_1) = \gamma([\gamma_0, 1], [\gamma_1, 1])$ if $\gamma_1 \geq 0$.

Proof. If $\gamma_1 \geq 0$, the neighborhood is defined by the union of two spheres. That neighborhood is contained in all neighborhoods that are the union of two larger or equal sized spheres, see figure 9. So when $N(\gamma_0, \gamma_1)$ is empty, the largest parameter values for which the neighborhood is still empty, are not less than γ_0 and γ_1 respectively. Thus $\gamma(\gamma_0, \gamma_1) \subseteq \gamma([\gamma_0, 1], [\gamma_1, 1])$. Conversely, look at the pairs of sites having parameter values for which the largest neighborhood is still empty, and are not less than γ_0 and γ_1 respectively. They also have an empty $N(\gamma_0, \gamma_1)$, so $\gamma(\gamma_0, \gamma_1) \supseteq \gamma([\gamma_0, 1], [\gamma_1, 1])$. ☺

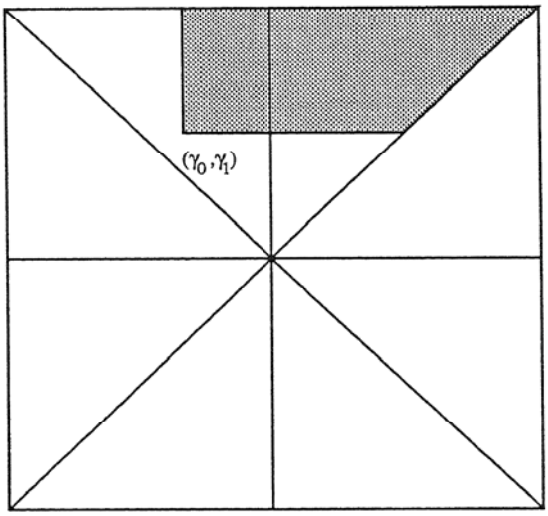

Figure 9: Shaded area denotes neighborhoods that contain $N(\gamma_0, \gamma_1)$ (lemma 1).

Lemma 2. $\gamma(\gamma_0, \gamma_1) = \gamma([\gamma_0, 1], [\gamma_1, 1]) \cup \gamma([\gamma_1, 1], [|\gamma_0|, 1])$ if $\gamma_1 \leq 0$.

Proof. If $\gamma_1 \leq 0$, the neighborhood is the intersection of two spheres. That neighborhood is contained in all neighborhoods that are the intersection or union of spheres defined by parameters larger than γ_0 and γ_1. But $N(\gamma_2, -\gamma_2)$ equals $N(-\gamma_2, \gamma_2)$, specifically for $\gamma_1 \leq \gamma_2 \leq \gamma_0$. $N(\gamma_0, \gamma_1)$ is therefore also contained in the neighborhood that are the union of spheres defined by parameters larger than γ_1 and $|\gamma_0|$ respectively, see figure 10. So when $N(\gamma_0, \gamma_1)$ is empty, the largest parameter values for which the neighborhood is still empty, are not less than γ_0 and γ_1, or γ_1 and $|\gamma_0|$. Thus $\gamma(\gamma_0, \gamma_1) \subseteq \gamma([\gamma_0, 1], [\gamma_1, 1]) \cup \gamma([\gamma_1, 1], [|\gamma_0|, 1])$. Conversely, look at the pairs of sites having parameter values for which the largest neighborhood is still empty, and are not less than γ_0 and γ_1 or γ_1 and $|\gamma_0|$. They also have an empty neighborhood defined by γ_0 and γ_1, so $\gamma(\gamma_0, \gamma_1) \supseteq \gamma([\gamma_0, 1], [\gamma_1, 1]) \cup \gamma([\gamma_1, 1], [|\gamma_0|, 1])$. ☺

Lemma 3. $\gamma([\gamma_0, \gamma_1], [\gamma_2, \gamma_3]) \subseteq \gamma([\gamma_4, \gamma_5], [\gamma_6, \gamma_7])$ if $[\gamma_0, \gamma_1] \subseteq [\gamma_4, \gamma_5]$ and $[\gamma_2, \gamma_3] \subseteq [\gamma_6, \gamma_7]$.

Proof. Consider the pairs of sites joined by an edge in $\gamma([\gamma_0, \gamma_1], [\gamma_2, \gamma_3])$. Their largest γ-parameters defining an empty neighborhood, lie in $[\gamma_0, \gamma_1]$ and $[\gamma_2, \gamma_3]$. Then they surely lie in $[\gamma_4, \gamma_5] \supseteq [\gamma_0, \gamma_1]$ and $[\gamma_6, \gamma_7] \supseteq [\gamma_2, \gamma_3]$. So all these sites are also joined by an edge in $\gamma([\gamma_4, \gamma_5], [\gamma_6, \gamma_7])$, and thus $\gamma([\gamma_0, \gamma_1], [\gamma_2, \gamma_3]) \subseteq \gamma([\gamma_4, \gamma_5], [\gamma_6, \gamma_7])$. ☺

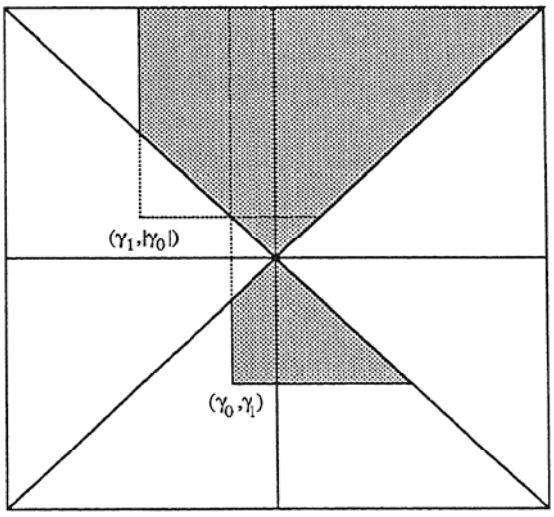

Figure 10: Shaded area denotes neighborhoods that contain $N(\gamma_0, \gamma_1)$ (lemma 2).

Lemmas 1, 2, and 3 are illustrated in figure 11. According to lemma 1, $\gamma(1/4, 1/2) = \gamma([1/4, 1], [1/2, 1])$ (say G1), and likewise $\gamma(-1/4, 1/4)$ is equal to $\gamma([-1/4, 1], [1/4, 1])$ (G2). Lemma 2 says that $\gamma(0, -1/4) = \gamma([0, 1], [-1/4, 1]) \cup \gamma([-1/4, 1], [0, 1])$ (G3), and $\gamma(0, -1/2) = \gamma([0, 1], [-1/2, 1]) \cup \gamma([-1/2, 1], [0, 1])$ (G4). And lemma 3 says G1 $\subseteq$ G2 $\subseteq$ G3 $\subseteq$ G4, as illustrated in the figure.

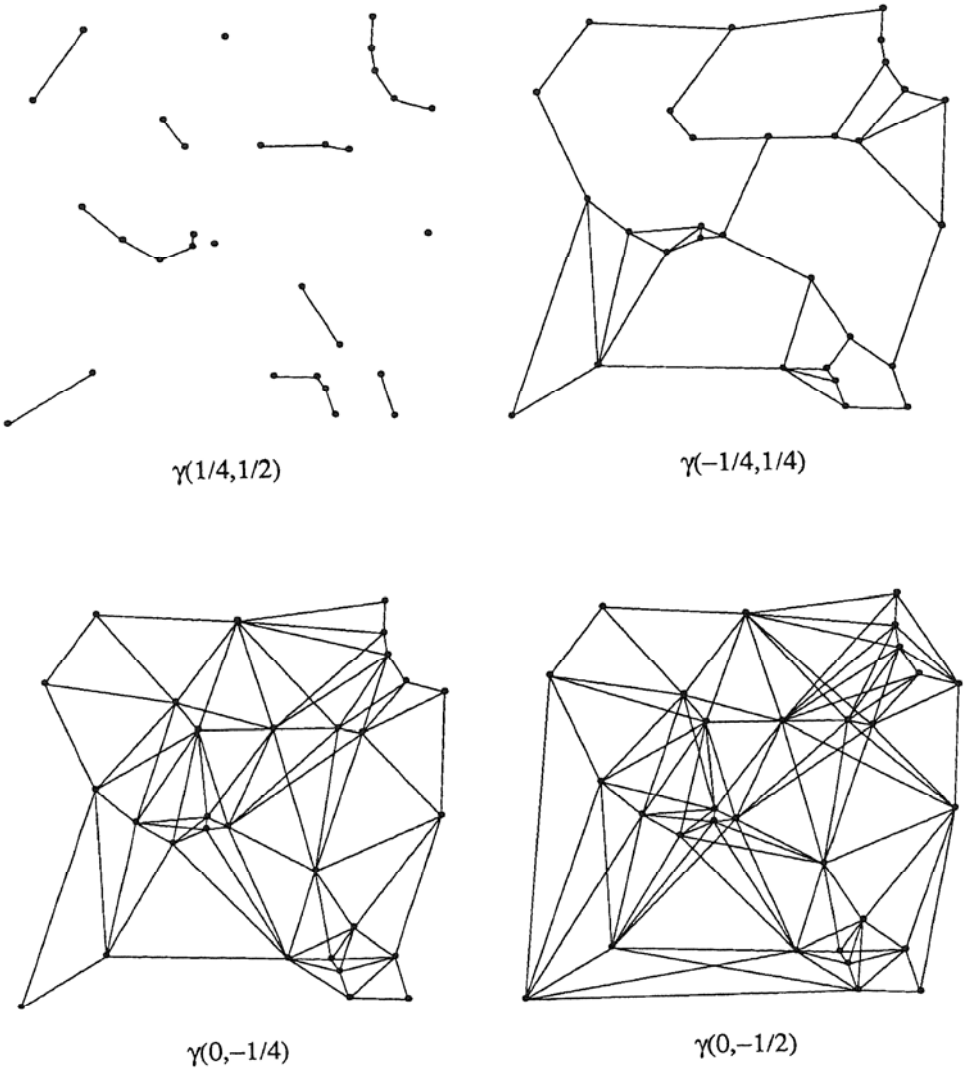

Figure 11: Four γ-graphs on the same set of 30 sites.
$\gamma(1/4, 1/2) \subseteq \gamma(-1/4, 1/4) \subseteq \gamma(0, -1/4) \subseteq \gamma(0, -1/2)$.

Theorem 1. Any k-dimensional γ-graph can be computed in $O(N^{k+1})$ time.
Proof. A brute force algorithm takes all the $\binom{N}{k}$ possible combinations of k sites, and checks whether any of the $N-k$ other sites lie in the neighborhood defined by the γ-parameters. This amounts to $O\left((N-k)\binom{N}{k}\right) = O(N^{k+1})$ time. ☺

Theorem 2. For non-degenerate sets of sites, the planar $\gamma([\gamma_0,\gamma_1],[\gamma_2,\gamma_3])$ can be computed in $O(N \log N)$ time, and the k-dimensional one in $O(N^{1+\lceil k/2 \rceil})$ time, provided that $[\gamma_0,\gamma_1] \subseteq [-1,1]$ and $[\gamma_2,\gamma_3] \subseteq [0,1]$.
Proof. When $[\gamma_0,\gamma_1] \subseteq [-1,1]$ and $[\gamma_2,\gamma_3] \subseteq [0,1]$, $\gamma([\gamma_0,\gamma_1],[\gamma_2,\gamma_3])$ is a subgraph of $\gamma([-1,1],[0,1])$ according to lemma 3. When the position of the sites is non-degenerate, $\gamma([-1,1],[0,1])$ is the Delaunay triangulation. After computing the DT, we can check whether the γ-parameter values of each (k-1)-simplex lie in the allowable range in constant time. The upper bounds to compute the Delaunay triangulation thus carry over to the γ-graph. These are $O(N \log N)$ for 2D, and $O(N^{1+\lceil k/2 \rceil})$ for kD (see section 2). ☺

For the Delaunay triangulation, $O(N \log N)$ is optimal. Whether this is optimal for the γ-graph, depends on the parameter values. It is clearly not optimal when the γ-graph reduces to the void graph.

Both $O(N \log N)$ and $O(N^{1+\lceil k/2 \rceil})$ only apply to non-degenerate cases. Because in the degenerate case that all sites lie on a k-dimensional sphere, the size of the output is already $O\left(\binom{N}{k}\right) = O(N^k)$.

6 Application

The γ-graph can be used to (re)construct a boundary of a set of sites [Vel89]. The sites are thought to be measured from the surface of a 2-dimensional or 3-dimensional object. The problem is to find a simple closed polygon or polyhedron, passing through all the sites. This is also called a Hamilton polygon, or Hamilton polyhedron.

In order to find a Hamilton polygon or polyhedron, we take $\gamma([-1,1],[0,1])$, and successively remove boundary simplices from the hull (initially the convex hull). We go on shrinking the hull, until all vertices are included in the boundary. However, by removing a simplex we may not introduce an isolated vertex, dangling boundary segments, or a self-intersecting boundary. The following short description for the planar case indicates the way simplices are selected for deletion.

A value is associated to all current boundary edges, based on the γ-values of the boundary segments. However, we keep the sign of the γ-values of the boundary edges consistent with the following rule: if $-1 \leq \gamma < 0$, the center of the associated circle lies on the side of the boundary segment that is outside the current boundary, and if $0 < \gamma \leq 1$, the center lies on the side of the boundary segment that is inside the current boundary. The selection of the triangle to be removed is based on the attempt to, informally speaking, change the shape of the current boundary not too much, relative to the size of the triangle. Formally, we choose the triangle with the largest interior angle at the vertex opposite the boundary edge.

Let us call the radius of the circle through the vertices of the triangle that we consider, R, the γ-value of the boundary edge corresponding to that triangle, γ, and the two vertices on the boundary, x_1 and x_2. We abbreviate $r(x_1,x_2)$ to r. If $\gamma \geq 0$, the angle ϕ at the interior vertex, increases when r/R increases. If $\gamma \leq 0$, ϕ increases when $2 - r/R$ increases.

The exact relation is given by the sine rule: $r/R = \sin\phi$. By definition, r/R equals $1 - |\gamma|$, which is $1 - \gamma$ for a non-negative γ. Similarly, $2 - r/R$ expands to $1 + |\gamma|$, which equals $1 - \gamma$ for a non-positive γ. This results in the following selection rule:

> among all removable triangles, delete the one whose boundary edge has the largest value for $1 - \gamma$ (or equivalently, the smallest value for γ).

A more detailed description, including the 3-dimensional case, is given in [Vel89].

The method described so far does not always succeed. In the first place, the shrinking operation can get locked, although the initial graph does contain a Hamilton polygon. This happens when there are no more removable edges, while not all the vertices are included in the boundary yet. Secondly, there exist non-degenerate non-Hamiltonian Delaunay triangulations [Dil87], and thus $\gamma([-1,1],[0,1])$-graphs. Methods based on shrinking from the Delaunay triangulation therefore guarantee no success.

In both cases, the solution is to shrink from a $\gamma([-1,1],[\gamma_0,1])$-graph, for some $\gamma_0 < 0$. Such a graph contains more edges. The extra triangles will have smaller interior angles at the vertex opposite to the boundary edge, than overlapping triangles from $\gamma([-1,1],[0,1])$. They offer more choice in selecting a boundary edge for deletion. For a γ_0 small enough, $\gamma([-1,1],[\gamma_0,1])$ will be Hamiltonian (the complete graph $\gamma([-1,1],[-1,1]) = \gamma(-1,-1)$ always contains a Hamilton polygon), and locking of the shrinking process will not occur.

Figure 12 shows a Delaunay triangulation from [Dil87], that contains no boundary through all the sites. The three other γ-graphs make it feasible to find such a boundary.

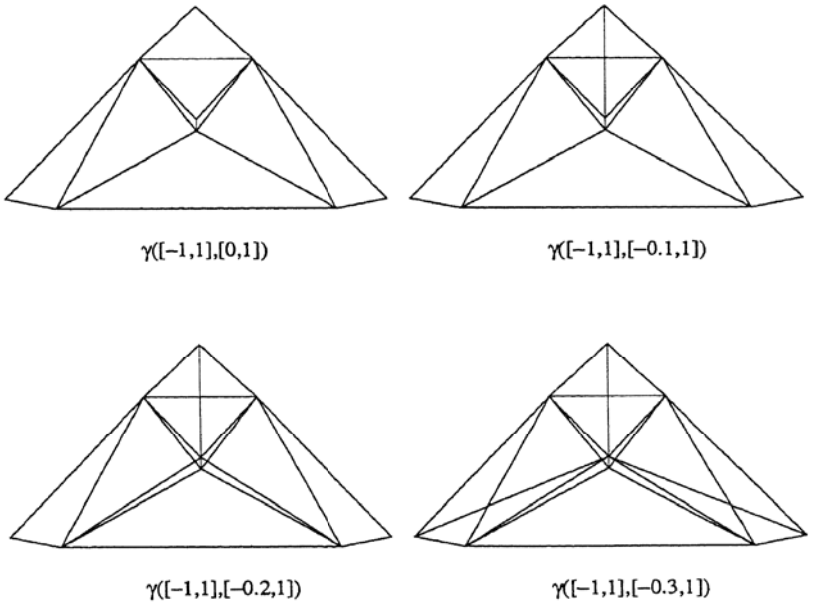

Figure 12: Planar non-Hamiltonian Delaunay triangulation (after [Dil87]), and three Hamiltonian γ-graphs.

7 Conclusions

In this paper we have introduced the γ-neighborhood graph. The γ-graph describes the internal structure of a set of sites. It has been shown briefly that it can also be used to find the external structure, specifically a simple boundary through all sites. The inclusion hierarchy CP $\subseteq$ NNG $\subseteq$ MST $\subseteq$ RNG $\subseteq$ GG $\subseteq$ DT has been extended: DT $\subseteq$ γ-graph. The γ-graph provides for a general framework in describing neighborhood graphs. It unifies the convex hull, the Delaunay triangulation, and in 2D also the Gabriel graph and the circle-based β-skeleton, into a continuous spectrum ranging from the void to the complete graph.

The neighborhood $N(\gamma_0, \gamma_1)$ is defined only for $\gamma_0, \gamma_1 \in [-1, 1]$, and $|\gamma_0| \leq |\gamma_1|$ For k sites and specific parameters γ_0 and γ_1, there can be two neighborhoods, since the spheres can be interchanged. The sites are connected if at least one of the two neighborhoods is empty. We could also use the parameters $|\gamma_0| \geq |\gamma_1|$ and completely specify the position of the spheres, for example γ_0 specifies 'the left', and γ_1 'the right' sphere. We feel no the urge to do so, because there is in general no need to specify a preference for one direction.

We have shown that the $\gamma([\gamma_0, \gamma_1], [\gamma_2, \gamma_3])$ can be constructed efficiently if it is a subgraph of the Delaunay triangulation, that is, when $[\gamma_0, \gamma_1] \subseteq [-1, 1]$, and $[\gamma_2, \gamma_3] \subseteq [0, 1]$, and the Delaunay triangulation is non-degenerate. Programs have been developed that construct the γ-graphs. They have been written in the programming language C, on a UNIX workstation. All the example graphs in this paper have been generated by these programs.

There are several directions for further research. The most urgent is the development of output sensitive algorithms. Of course for $\gamma < 0$, the worst case size of the γ-graph is $O(N^3)$, but an algorithm having a time complexity that depends on the size of the output can probably do better than $O(N^3)$. Also for $\gamma > 0$ an output sensitive algorithm can be profitable, since the size of the γ-graph may be sub-linear in N.

Little is known from stochastic geometry about probabilistic properties of geometric graphs (some results are known about the DT [Mil70], the GG, and the RNG [Dev88]). Insight in the expected number of edges in the γ-graph may lead to the development of efficient algorithms for the average case.

A final research suggestion is the construction of γ-graphs on sets of weighted sites.

References

[Bro79] K. Q. Brown. Voronoi diagrams from convex hulls. *Information Processing Letters*, 9(5):223 – 228, December 1979.

[BS76] J. L. Bentley and M. I. Shamos. Divide-and-conquer in multidimensional space. In *Proceedings of the 8th Annual Symposium on Theory of Computing*, pages 220 – 230, 1976.

[Del28] B. Delaunay. Sur la sphère vide. In *Proceedings of the International Congress on Mathematics (Toronto 1924)*, volume 1, pages 695 – 700, Toronto, 1928. University of Toronto Press.

[Del34] B. Delaunay. Sur la sphère vide. *Izvestija Akademii Nauk S.S.S.R. Otdelenie Matematiceskich i Estestvennych Nauk (Bulletin de l'Académie des Sciences de*

l'URSS, VII Série, Classe des Sciences Mathématiques et Naturelles), pages 793 – 800, 1934.

[Dev88] L. Devroye. The expected size of some graphs in computational geometry. *Computers & Mathematics with Applications*, 15(1):53 –64, 1988.

[Dil87] M. B. Dillencourt. A non-Hamiltonian, nondegenerate Delaunay triangulation. *Information Processing Letters*, 25(3):149 – 151, May 1987.

[EKS83] H. Edelsbrunner, D. G. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. *IEEE Transactions on Information Theory*, IT-29(4):551 – 559, July 1983.

[GS69] K. R. Gabriel and R. R. Sokal. A new statistical approach to geographic variation analysis. *Systematic Zoology*, 18:259 – 278, 1969.

[KR85] D. G. Kirkpatrick and J. D. Radke. *A Framework for Computational Morphology*, pages 217 – 248. Elsevier Science Publishers, 1985.

[Lan69] P. M. Lankford. Regionalization: Theory and alternative algorithms. *Geographical Analysis*, 1(2):169 – 212, April 1969.

[LS80] D. T. Lee and B. J. Schachter. Two algorithms for constructing the delaunay triangulation. *International Journal of Computers and Information Science*, 9(3):219 – 242, 1980.

[Mil70] R. E. Miles. On the homogeneous planar Poisson point process. *Mathematical Biosciences*, 6:85 – 127, 1970.

[MS80] D. W. Matula and R. R. Sokal. Properties of gabriel graphs relevant to geographic variation research and the clustering of points in the plane. *Geographical Analysis*, 12:205 – 222, July 1980.

[PH77] F. P. Preparata and S. J. Hong. Convex hulls of finite sets of points in two and three dimensions. *Communications of the ACM*, 20(2):87 – 93, February 1977.

[Pri57] R. C. Prim. Shortest connection networks and some generalizations. *Bell Systems Technical Journal*, 36:1389 – 1401, 1957.

[Sei81] R. Seidel. A convex hull algorithm optimal for point sets in even dimensions. Technical Report 81 – 14, University of British Columbia (MSc thesis), Vancouver, September 1981.

[SH75] M. I. Shamos and D. Hoey. Closest point problems. In *Proceedings of the 16th Annual IEEE Symposium on Foundations of Computer Science*, pages 151 – 162, New York, October 1975. IEEE.

[Sup83] K. J. Supowit. The relative neighbourhood graph with an application to minimum spanning trees. *Journal of the ACM*, 30(3):428 – 447, July 1983.

[Tou80] G. T. Toussaint. The relative neighbourhood graph of a finite planar set. *Pattern Recognition*, 12(4):261 – 268, 1980.

[Tou88] G. T. Toussaint. *A Graph Theoretical Primal Sketch*, volume 6 of *Machine Intelligence and Pattern Recognition*, pages 229 – 260. Noth-Holland, 1988.

[Vel89] R. C. Veltkamp. 2D and 3D computational morphology on the γ-neighborhood graph. *Acta Stereologica*, 8(2/2):595 – 600, November 1989.

[Vor08] G. Voronoï. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième mémoire — Recherche sur les parralléloèdres primitifs, Introduction et première partie. *Journal für die reine und angewandte Mathematik*, 134:198 – 287, 1908.