

Introducing 3D Venn and Euler Diagrams

Peter Rodgers¹, Jean Flower², and Gem Stapleton³

¹ University of Kent, UK

p.j.rodgers@kent.ac.uk

² Autodesk, UK

³ Visual Modelling Group, University of Brighton, UK

g.e.stapleton@brighton.ac.uk

Abstract. In 2D, Venn and Euler diagrams consist of labelled simple closed curves and have been widely studied. The advent of 3D display and interaction mechanisms means that extending these diagrams to 3D is now feasible. However, 3D versions of these diagrams have not yet been examined. Here, we begin the investigation into 3D Euler diagrams by defining them to comprise of labelled, orientable closed surfaces. As in 2D, these 3D Euler diagrams visually represent the set-theoretic notions of intersection, containment and disjointness. We extend the concept of wellformedness to the 3D case and compare it to wellformedness in the 2D case. In particular, we demonstrate that some data can be visualized with wellformed 3D diagrams that cannot be visualized with wellformed 2D diagrams. We also note that whilst there is only one topologically distinct embedding of wellformed Venn-3 in 2D, there are four such embeddings in 3D when the surfaces are topologically equivalent to spheres. Furthermore, we hypothesize that all data sets can be visualized with 3D Euler diagrams whereas this is not the case for 2D Euler diagrams, unless non-simple curves and/or duplicated labels are permitted. As this paper is the first to consider 3D Venn and Euler diagrams, we include a set of open problems and conjectures to stimulate further research.

1 Introduction

Euler diagrams represent intersection, containment and disjointness of sets. Currently, these diagrams are drawn in the plane and consist of labelled simple closed curves. These *2D Euler diagrams* have been widely studied over the last few years and much progress has been made on their theoretical underpinning and techniques for automatically drawing them.

Here we introduce the concept of *3D Euler diagrams*. We know of no other work defining this type of 3D representation and, thus, this paper focusses on setting the groundwork for discussing this new diagrammatic type. Furthermore, it provides a platform to engage the community in discussion about the various issues in 3D Euler diagram research. 3D Euler diagrams consist of labelled orientable closed surfaces drawn in \mathbb{R}^3 . An example of a 2D and a 3D Euler diagram representing the same information can be seen in figure 1. This 3D diagram, as well as all of the 3D Euler diagrams drawn in this paper, can be accessed from




Fig. 1. A 2D Euler diagram with an equivalent 3D Euler diagram.




Fig. 2. Four topologically distinct wellformed Venn-3s.

www.eulerdiagrams.com/3D/workshop/. Using the freely available Autodesk Design Review software, one can rotate and explore the 3D diagrams.

We define *3D Venn diagrams* as 3D Euler diagrams where all combinations of surface intersections are present. An interesting comparison between 2D and 3D is in the common Venn-3 case, i.e the Venn diagram representing exactly three sets. It is known that there is only one topologically distinct embedding of wellformed Venn-3 in 2D [9]. In 3D, there are infinitely many topologically distinct embeddings of wellformed Venn-3 when the surfaces are closed and orientable (i.e. connected sums of tori). When the surfaces are topologically equivalent to the sphere, there are at least four topologically distinct embeddings of wellformed 3D Venn-3, shown in figure 2.

Whilst 3D Venn and Euler diagrams are interesting in their own right, we believe that there are also solid practical motivations for examining them. Firstly, the recent advances in hardware available for 3D display and interaction (eg. 3D televisions and Microsoft Kinect) support 3D visualization. As Venn and Euler diagrams form an important aspect of 2D visualization, it is reasonable to expect that they will also be important for 3D visualization.

Secondly, there are intrinsic benefits to exploring 3D with respect to Euler diagrams. When 2D Euler diagrams are defined as consisting of (anything equivalent to) simple closed curves without duplicated labels (for instance [4] and [6]), not all data sets can be visualized. This is clearly a major limitation. Subsequently, the definition of a 2D Euler diagram was relaxed, permitting diagrams to have non-simple curves and duplicated curve labels [7, 10]. Under this new approach, all data sets can be visualized but potentially at the cost of significantly reduced usability [8]. However, as we note later in this paper, in the 3D case we conjecture that it is possible to draw all data sets (encapsulated by *diagram descriptions*) without duplicate labels and non-simple surfaces. Consequently, this major limitation on undrawability is overcome in the 3D case.




Fig. 3. An non-wellformed 2D diagram and an equivalent wellformed 3D diagram.

Wellformedness properties are a key aspect of drawing of Euler diagrams. In 2D, they relate to how the curves intersect and to the properties of the regions present. In 3D, we generalize them to how the surfaces intersect and the properties of the solids to which the surfaces give rise. The 2D Euler diagram on the left of figure 3 is not wellformed because it has a triple point of intersection between the curves. By contrast, the same data can be represented in a wellformed manner in 3D, as shown in the righthand side of figure 3; we will demonstrate, in section 4, that any way of drawing a 2D Euler diagram representing the same data breaks a wellformedness property. Often, there exists wellformed 3D Euler diagrams for data that has no wellformed 2D representation.

The remainder of this paper is as follows. Section 2 formally defines 3D Euler diagrams and related concepts. Section 3 generalizes wellformedness properties of 2D Euler diagrams to the 3D context. Section 4 establishes that more data sets can be drawn wellformed with 3D Euler diagrams than with 2D Euler diagrams. We then go on to examine future work and propose open questions in section 5. Finally, section 6 concludes.

2 What is a 3D Euler Diagram?

3D Euler diagrams are formed from closed surfaces embedded in \mathbb{R}^3 rather than closed curves embedded in \mathbb{R}^2 . We refer the reader to [12] for a formal definition of a 2D Euler diagram and associated wellformedness properties. As with closed curves in 2D Euler diagrams, which are typically required to be simple, we choose not to use arbitrary surfaces in 3D Euler diagrams. This is because we want to be able to define certain properties of 3D Euler diagrams that require the surfaces to be ‘nice’. Choosing our surfaces to be orientable gives us a well-understood notion of what constitutes the interior. Hence, we define 3D Euler diagrams as follows, where \mathcal{L} is a set of labels that we use to label the surfaces:

Definition 1. A *3D Euler diagram* is a pair, $d = (\mathcal{S}, l)$, where

1. \mathcal{S} is a finite set of closed, orientable surfaces embedded in \mathbb{R}^3 , and
2. $l: \mathcal{S} \rightarrow \mathcal{L}$ is an injective function that labels each surface.

In 2D Euler diagrams, zones are sets of points in the plane that are inside all curves in a given set and outside the rest of the curves in the diagram. In figure 1, both diagrams have five zones. Zones are fundamentally important, since these correspond to the semantics of the diagram: between them, the present zones must represent all of the non-empty set intersections. We now generalize the notion of a zone to the 3D case:

Definition 2. A **zone** in a 3D Euler diagram, $d = (\mathcal{S}, l)$, is a set of points, z , in \mathbb{R}^3 for which there exists a subset, S , of \mathcal{S} such that

1. every point, p_{in} , in z is inside all of the surfaces in S and outside all of the surfaces in $\mathcal{S} - S$, and
2. z is maximal with this property.

Such a zone, z , is described by $des(z) = \{l(s) : s \in S\}$. The set of zones in d is denoted $\mathcal{Z}(d)$.

In the visualization process, one starts with a description of the to-be-drawn diagram. A diagram description is a list of the set intersections that must be present in the diagram, given the sets to be visualized, thus precisely encapsulating the categories in which data items lie. For example, suppose we wish to visualize the sets P , Q , and R , and the intersections we wish to visualize are $P \cap \overline{Q} \cap \overline{R}$, $Q \cap \overline{P} \cap \overline{R}$, $R \cap \overline{P} \cap \overline{Q}$ (i.e. the set intersections that comprise elements in exactly one of the three sets), along with $P \cap Q \cap \overline{R}$, $P \cap R \cap \overline{Q}$ and $Q \cap R \cap \overline{P}$ (i.e. the set intersections that comprise elements that are in exactly two of the sets). Further, we also must visualize the set intersection that comprises elements in none of the three sets, namely $\overline{P} \cap \overline{Q} \cap \overline{R}$. This is more succinctly represented as $\emptyset, P, Q, R, PQ, PR, QR$, listing the non-complemented sets from each specified intersection, and is visualized by both diagrams in figure 3. More formally these diagrams have description $\{\emptyset, \{P\}, \{Q\}, \{R\}, \{P, Q\}, \{P, R\}, \{Q, R\}\}$, but we will abuse notation as just illustrated.

Definition 3. A **diagram description**, D , is a subset of $\mathbb{P}\mathcal{L}$ that includes \emptyset . The **description** of a 3D Euler diagram, $d = (\mathcal{S}, l)$, is $\{des(z) : z \in \mathcal{Z}(d)\}$.

The classic drawing problem, generalized to 3D, is *given a diagram description, D , draw a 3D Euler diagram with description D* . In 2D, this problem is often subject to a range of extra constraints that typically relate to the wellformedness properties. For instance, we may wish to find a diagram that has no concurrency between surfaces. We generalize the wellformedness properties to 3D in the next section.

3 Wellformedness Properties of 3D Euler Diagrams

There are various wellformedness properties that can be applied to 2D Euler diagrams [12]. These are informally described in table 1, where we also present their generalizations to 3D. Examples of non-wellformed diagrams in both 2D

Property	2D Case	3D Case
Connected Zones	Every zone is a connected component of \mathbb{R}^2 .	Every zone is a connected component of \mathbb{R}^3 .
n -Point	Every point in \mathbb{R}^2 is passed through at most $n = 2$ times by the curves.	Every point in \mathbb{R}^3 is passed through at most $n = 3$ times by the surfaces.
Crossings	Whenever two curves intersect, they cross transversely.	Whenever two surfaces intersect, they cross transversely.
Line Concurrency	No two curves share a common line segment.	No three surfaces share a common line segment.
Surface Concurrency	N/A	No two surfaces share a common sub-surface.

Table 1. Wellformedness properties.

and 3D are shown in table 2. Some of the wellformedness properties in 3D are obvious generalizations of the 2D case, but others benefit from further discussion.

First, consider the n -points properties. For the 2D case, a diagram is non-wellformed if it contains a triple point (i.e. a 3-point). The reason that the presence of 2-points does not render a diagram non-wellformed is because whenever two curves intersect, a 2-point is formed. However, given three curves that pairwise intersect it need not be the case that a 3-point is formed. Thus, 3-points are avoidable in 2D. However, 3-points are not avoidable in 3D. This is illustrated in figure 4. Here, three spheres intersect to form Venn-3. The cross-section of the diagram shown on the right of the figure illustrates that a three 3-point is formed.

Now consider now line concurrency. In 2D, diagrams that have two curves running concurrently along a line segment are not wellformed. However, in 3D, such a property is unavoidable: two surfaces that intersect and share common interior points necessarily share a common line segment. For instance, the Venn-3 diagram drawn with spheres on the left of figure 4 has line concurrency.




Fig. 4. The necessity of 3-points in 3D.










Property	2D Case	3D Case
Connected Zones	 <p>The zone PQ is disconnected.</p>	 <p>Here, P is a sphere with a 'sausage', Q, through it. The zone inside Q and outside P is disconnected.</p>
n -point	 <p>The curves P, Q, and R form two 3-points.</p>	 <p>The spheres P, Q, and R form a 4-point where they all intersect with S.</p>
Crossings	 <p>The curves P and Q intersect at a point where they do not cross (as do R and S).</p>	 <p>The sphere R intersects with Q but does not cross Q; a cross-section is shown on the right.</p>
Line Concurrency	 <p>The two curves P and Q share a common line segment.</p>	 <p>The three tori share a common line segment; a cross-section is shown on the right.</p>
Surface Concurrency	N/A	 <p>The two 'squashed' spheres share a disc-like surface.</p>

Table 2. Examples of non-wellformed diagrams.

4 Drawability of 3D Euler Diagrams

One of our key motivations for developing 3D Euler diagrams is that they allow more diagram descriptions to be drawn in a wellformed manner than is the case for 2D Euler diagrams. Recall that a 2D Euler diagram comprises a set of labelled simple closed curves such that no label is used on more than one curve.

Definition 4. *Let D be a diagram description. Then D can be **drawn well-formed** if there exists a Euler diagram with description D which satisfies all of the wellformedness properties of table 1.*

We now establish that every diagram description that can be drawn wellformed in 2D can be drawn wellformed in 3D. Our proof strategy is to convert a wellformed 2D diagram into a wellformed 3D diagram with the same description. To illustrate the approach, consider figure 5. Here, the wellformed 2D diagram is converted into a 3D diagram by rotating the 2D diagram around line that does not pass through any of the curves. Each resulting surface is a torus and the final 3D diagram is shown on the right.




Fig. 5. Converting a wellformed 2D diagram into a wellformed 3D diagram.

Theorem 1. *Let D be a diagram description. If D can be drawn wellformed in 2D then D can be drawn wellformed in 3D.*

Proof (Sketch). Suppose that D can be drawn wellformed in 2D. Choose any wellformed 2D diagram, d_2 , with description D . Draw a line, λ , that does not pass through any curve in d_2 . Rotate d_2 about λ by 2π to create a 3D Euler diagram, d_3 . Each closed curve, c_2 , in d_2 gives rise to a torus, t_3 , in d_3 and we label t_3 the same as c_2 . It can be shown that each zone, z_2 , in d_2 gives rise to a zone, z_3 , in d_3 with the same description and that no other zones appear in d_3 . That is, the description of d_3 is D . The wellformedness of d_3 can be trivially established using the wellformedness of d_2 .

We now demonstrate that there are diagram descriptions that cannot be drawn wellformed in 2D that can be drawn wellformed in 3D. Work by Flower and Howse [4] identified necessary and sufficient conditions for when a diagram description can be drawn wellformed in the 2D case. We will demonstrate that their conditions are not both necessary and sufficient in 3D, failing in multiple ways. Their approach starts by converting a diagram description into a graph,

called the *super-dual*, and looks at properties of this graph to establish drawability.

Definition 5. Given a diagram description, D , the **super-dual** of D is a graph, $G = (V, E)$, where $V = D$ is the set of vertices and, for every pair of vertices, v_1 and v_2 , there is an edge between v_1 and v_2 if and only if v_1 and v_2 differ by a single label (recall elements of D , i.e. the vertices, are sets of labels).

Assuming we have a super-dual that is planar, we can draw that graph in the plane without edges crossing. Given such an embedding of the super-dual, we can attempt to form the required Euler diagram. An illustration of the process is given in figure 6. Here, we start with description \emptyset, P, Q, PQ and turn it into the super-dual, shown on the left of figure 6. The curves of the 2D Euler diagram are constructed by enclosing the vertices appropriately. For example, to draw a curve labelled P we enclose the vertices that include P but no others. The curve labelled Q is similarly formed. Finally, we delete the super-dual and are left with the required 2D Euler diagram, shown on the right.

The preceding example is rather simple, but it is by examining the super-dual and, if necessary, its subgraphs that we can determine wellformed drawability in the 2D case. Now, the curves of a (wellformed) 2D Euler diagram are all simple which means that for each curve, c , the set of points inside c is a simply connected region. In terms of a super-dual, this implies that the maximal subgraph induced by the vertices that contain the label of c is connected and, moreover, that the subgraph induced by the vertices that do not contain the label of c is also connected. This key insight led Flower and Howse to define the *connectivity conditions* for graphs; these are used to establish properties of super-duals arising from diagram descriptions.

Definition 6 (Connectivity Conditions [4]). Let $G = (V, E)$ be a graph such that $V \subseteq \mathbb{P}\mathcal{L}$. The **connectivity conditions** for G are:

1. G is connected,
2. for each curve label, λ , in \mathcal{L} , the maximal subgraph of G whose vertices include λ is connected, and
3. for each curve label, λ , in \mathcal{L} , the maximal subgraph of G whose vertices do not include λ is connected.




Fig. 6. Constructing a 2D Euler diagram from the super-dual.

Theorem 2 (2D Connectivity Test [4]). *Let D be a diagram description whose super-dual fails connectivity conditions. Then there is no 2D Euler diagram, d , with description D that is wellformed.*

The connectivity conditions are necessary for wellformed drawability in 3D:

Theorem 3 (3D Connectivity Test). *Let D be a diagram description whose super-dual fails connectivity conditions. Then there is no 3D Euler diagram, d , with description D that is wellformed.*

Flower and Howse further introduce the *face conditions*, which we now informally explain via an example; for full details we refer to [4]. Consider the diagram description $\emptyset, P, Q, PQ, R, PR, QR$. This has super-dual as shown on the left of figure 7. All three curves will pass through the face f , which will lead either to a 3-point (as shown in the figure), a disconnected zone, or an un-required zone (to create such a zone, nudge one of the curves to remove the triple point). The non-wellformedness of the diagram is determined by examining the edges around f . By traversing the simple cycle around f , each time we pass along an edge we write down the curve label that is in one of the incident vertices but not the other to form a word, say $w = RPQRPQ$. By examining alternations of letters in w , we can see which curves are required to cross. For instance, P and Q alternate, since $PQPQ$ is a scattered subword of w . This tells us that (the curves labelled) P and Q must cross in f . Similarly, P and R must cross and Q and R must cross. This indicates the possible presence of a triple point. In general, a combinatorial analysis of the words around faces in the graph is used to determine whether the plane embedding will give rise to a wellformed 2D Euler diagram. The face conditions for the graph, roughly speaking, identify whether too many crossings occur for wellformedness to be achieved. Of note is that our example is very simple and the actual details are more complex than we have illustrated. In any case, the super-dual in figure 7 is planar, passes the connectivity conditions, but fails the face conditions. Hence, this embedding of the super-dual cannot be used to draw a wellformed 2D Euler diagram with the specified description. Moreover, there is no different choice of embedding which passes the face conditions.




Fig. 7. Failure of the face conditions.

For some diagram descriptions, but not the one just considered, it is possible to remove edges from the super-dual whilst ensuring connectivity holds and produce a subgraph, G , that has an embedding which passes the face conditions. A further complication is the potential lack of planarity of the super-dual. Again, we may be able to remove edges to create a planar subgraph G with the properties just described. In either case, if such a G exists then D is drawable as a wellformed 2D Euler diagram, otherwise it is not. This key result is captured in the following theorem:

Theorem 4 (2D Drawability – Necessary and Sufficient Conditions [4]). *Let D be a diagram description with super-dual $SG(D)$. There exists a wellformed 2D Euler diagram that is a drawing of D iff there exists a planar subgraph, G , of $SG(D)$ obtained by removing edges from $SG(D)$, which passes the connectivity conditions and has a plane embedding that passes the face conditions.*

We can immediately generalize one side of this theorem to the 3D case:

Theorem 5 (3D Drawability – Sufficient Conditions). *Let D be a diagram description with super-dual $SG(D)$. If there exists a planar subgraph, G , of $SG(D)$ obtained by removing edges from $SG(D)$, which passes the connectivity conditions and has a plane embedding that passes the face conditions then there exists a wellformed 3D Euler diagram that is a drawing of D .*

Proof. By theorem 4, a wellformed 2D Euler diagram exists. The result then follows from theorem 1.

We now demonstrate that there are diagram descriptions that are not drawable wellformed in 2D (they fail one or more of the conditions in theorem 4) but that are drawable wellformed in 3D. The three examples below fail the conditions of Theorem 4 in different ways. This shows that there are more diagram descriptions that can be drawn wellformed in 3D than 2D and that the conditions to determine drawability in 2D are not useful for determining drawability in the context of 3D diagrams. For these three examples, the wellformedness of the 3D representations suggests that they are more readable than the 2D representations and the 3D diagrams display a pleasing symmetry.

Example 1: Figure 7 shows a planar super-dual that passes the connectivity conditions but fails the face conditions, so the corresponding 2D representation shown in figure 7 is not well-formed (it has a triple-point). All plane embeddings




Fig. 8. Failure of planarity of the super-dual.




Fig. 9. Failure of planarity with no planar subgraph that passes connectivity.

of this graph fail the face conditions. Removing edges from this graph will never result in a graph that passes the connectivity and face conditions. Hence, there is no wellformed 2D Euler diagram with the given description. A wellformed 3D Euler diagram with this description can be seen in figure 7.

Example 2: Figure 8 shows a super-dual that is non-planar, since it is homeomorphic to $K_{5,5}$, so some edge removal is necessary to achieve planarity. We demonstrate that any way in which edges can be removed to achieve planarity whilst maintaining connectivity does not produce a graph which passes the face conditions. If we remove an edge from the super-dual that is not incident to \emptyset then we break the connectivity conditions. If we remove an edge which is incident to \emptyset then the graph becomes planar and connectivity is preserved, so we then look for an embedding which passes the face conditions. However, any plane embedding of the graphs resulting from the removal of exactly one of these edges, shown here with the edge between \emptyset and S removed, fails the face conditions. Continuing this kind of analysis, it can be demonstrated that there is no wellformed 2D Euler diagram with the given description. A wellformed 3D Euler diagram with the same description can be seen in figure 8.

Example 3: Figure 9 shows a super-dual that is non-planar, since it has a proper subgraph homeomorphic to $K_{5,5}$, so again some edge removal is necessary to achieve planarity. However, all planar subgraphs fail the connectivity conditions. Take one edge as an example, say the edge between T and RT . This edge is in the maximal subgraph of the super-dual whose vertices include T . The removal of the T - RT edge would disconnect this subgraph, breaking connectivity. The same argument prevents removal of any edge not incident to \emptyset . Thus, the only edges we can consider removing are those incident with \emptyset . However, the subgraph obtained by removing the vertex \emptyset is homeomorphic to $K_{5,5}$. This implies that we cannot obtain a planar subgraph by removing edges from the super-dual whilst preserving connectivity. Hence, there is no wellformed 2D Euler diagram with the given description. A wellformed 3D Euler diagram with the same description can be seen in figure 9.

Thus, more diagram descriptions are drawable wellformed in 3D than in 2D. In particular, the face conditions need not be passed in order for us to have wellformed drawability in 3D and we need not have planarity of the dual.




Fig. 10. Constructing 3D diagrams from descriptions.

5 Future Work and Open Problems

There are numerous open questions in 3D. Following the previous section:

Open Problem 1 *What are necessary and sufficient conditions for determining wellformed drawability in the 3D case?*

We have demonstrated that the connectivity conditions are necessary, but there is no obvious generalization of the face conditions to the 3D case (the notion of a face does not translate to 3D graphs). We conjecture that a different approach is needed and it is very possible that this could provide a new perspective on wellformed drawability in the 2D case as well.

An important question is how the definition of Euler diagrams needs to be relaxed in order to draw every diagram description. As discussed previously, in 2D we require either non-simple curves or duplicated label use. We believe that the definition given in this paper for the 3D case is sufficient for drawability in general, if we do not impose any wellformedness properties:

Conjecture 1 *For every diagram description there exists a 3D Euler diagram with that description.*

We are confident that this conjecture is true because we believe the following method for construction works in general. Given the diagram description $D = \{\emptyset, \{P\}, \{Q\}, \{P, Q\}$ for each element, z , in D create one sphere for each label in z and draw them concurrently. For each pair of elements, z_1 and z_2 , in D , if they share a non-empty set of labels, $L = z_1 \cap z_2$, then join the spheres with labels in L arising from z_1 and z_2 . This example can be seen in figure 10.

We focus now on a specific class of Euler diagrams which is the widely known family of Venn diagrams [9]. In 2D, Venn diagrams are Euler diagrams where all 2^n possible intersections between n sets are represented by connected regions, that is there are 2^n zones each of which is connected. In 3D:

Definition 7. *A 3D Venn diagram, $d = (\mathcal{S}, l)$, is a 3D Euler diagram where there are $2^{|\mathcal{S}|}$ zones, each of which is connected.*

Four topologically distinct embeddings of Venn-3 are shown in the introduction, figure 2. To see that they are distinct, we make arguments about their zones. The first (leftmost) Venn-3 has only simply connected zones. The second Venn-3 has exactly two zones that not simply connected, namely P and PR . The third Venn-3 has exactly two zones that are not simply connected, namely

\emptyset and R . The fourth (rightmost) Venn-3 also has exactly two zones that are not simply connected, namely QR and PQR . The four diagrams are pairwise topologically distinct because the non-simply connected zones are contained by different numbers of surfaces.

Conjecture 2 *There are exactly four topologically distinct embeddings of well-formed 3D Venn-3 when the surfaces are all topologically equivalent to spheres.*

A variety of other open problems can be stated for 3D Venn diagrams, some of which have been answered for 2D Venn diagrams (see [9] for an excellent survey on results for 2D Venn diagrams). One such example is:

Open Problem 2 *How many topologically distinct embeddings of wellformed Venn- n exist when the surfaces are all topologically equivalent to spheres?*

Returning to the more general case of Euler diagrams, there has been considerable interest in drawing them with curves of particular shapes. For instance, Stapleton et al. [13] identified a class of diagram descriptions that could be drawn using only circles and Wilkinson devised a method that only drew Euler diagrams using circles [14]. Kestler et al. devised a method for drawing Euler diagrams with regular polygons [5] and others have considered drawing Venn diagrams where the curves have other geometric shapes, such as triangles [1]. Thus, curve shape is considered interesting and important in the 2D case. For 3D, this generalizes to surface shape (where we no longer mean ‘up to topological equivalence’). We pose the following two problems concerning surface shape:

Open Problem 3 *What class of diagram descriptions can be drawn when the surfaces are all some specified shape, such as spheres?*

Open Problem 4 *Can all diagram descriptions that can be drawn wellformed in 2D using only circles can be drawn wellformed in 3D using only spheres?*




Fig. 11. Converting from circles to spheres.

A naïve method for converting a diagram drawn with circles into one drawn with spheres is to use each circle to generate a sphere. However, this construction approach need not lead to the required diagram description or preserve wellformedness. This is demonstrated in figure 11: the leftmost Euler diagram is

drawn with circles, the 3D Euler diagram is obtained by converting the circles to spheres and the remaining diagrams show cross-sections of the 3D Euler diagram. Unfortunately, this created an extra zone, that inside only S and, moreover, this zone is disconnected. We conjecture that there does not exist a wellformed 3D diagram drawn with spheres with the same diagram description as figure 11. However, we believe that some classes of diagram descriptions drawable wellformed with circles can be drawn wellformed with spheres:

Conjecture 3 *The class of inductively pierced descriptions, introduced in [13] and generalized in [10], which can all be drawn wellformed with circles in 2D can be drawn wellformed with spheres in 3D.*

Finally, there has also been significant interest in drawing 2D Euler diagrams in a so-called *area-proportional* manner. In the area-proportional 2D case, the zones must have specified areas. Key publications on area-proportional Venn and Euler diagram drawing include Chow and Rodgers [2], Chow and Ruskey [3], Kestler et al. [5], Stapleton et al. [11], and Wilkinson [14]. These methods mostly consider drawing the diagrams where the curves have specific shapes, such as circles. In 3D, the area-proportional case generalizes to the zones having specified volumes, the *volume-proportional* case.

Definition 8. *A **volume specification** is a function, $v: \mathbb{P}\mathcal{L} - \{\emptyset\} \rightarrow \mathbb{R}^+ \cup \{0\}$. The diagram description **induced** by v is $\{z : v(z) \neq \emptyset \vee z = \emptyset\}$. A 3D Euler diagram **conforms** to a volume specification, v , if its description is induced by v and its zones have the volumes specified by v .*

Open Problem 5 *What class of volume specifications can be drawn in a wellformed manner?*

Open Problem 6 *What class of volume specifications can be drawn where the surfaces are all some specified shape?*

6 Conclusion

In this paper we have introduced the concept of 3D Euler diagrams, formally defining them as orientable closed surfaces which implies the surfaces are simple. We have compared them with 2D Euler diagrams and discovered that 3D Euler diagrams have some benefits over 2D Euler diagrams in terms of drawability when wellformedness is considered. In particular, we have shown that there are more diagram descriptions that can be drawn wellformed with 3D diagrams than 2D diagrams and we conjecture that all diagram descriptions can be drawn in 3D without allowing non-simple surfaces and duplicate label use, unlike in 2D.

We established that there are four topologically distinct wellformed embeddings of 3D Venn-3, whereas there is only one such embedding in 2D. This demonstrates that there is more choice in terms of how we layout diagrams in 3D over 2D, which is likely to be beneficial when more sets need to be represented. This gives further insight into why more diagram descriptions can be

drawn wellformed in 3D: we have greater control over which zones are topologically adjacent and topological adjacency impacts whether we can add a new surface (or curve in 2D) and maintain wellformedness.

By presenting a series of open questions and conjectures, we hope to stimulate research progress on 3D Euler diagrams. In some cases there is no obvious way to extend existing 2D results to the 3D case, such as Open Problem 1 concerning the drawability of wellformed diagrams. Hence, a different approach is likely to be required. It may be that results in the 3D case will allow more progress to be made in the 2D case.

With the advent of recent affordable 3D display, interaction and printing devices, 3D visualization has the potential to be commonplace. We expect that 3D Euler diagrams will form a useful component in this field.

Acknowledgement Gem Stapleton was partially supported by an Autodesk Education Grant.

References

1. J. Carroll. Drawing Venn triangles. Technical report, HP Labs HPL-2000-73, 2000.
2. S. Chow, P. Rodgers. Constructing area-proportional Venn and Euler diagrams with three circles. In *Euler Diagrams 2005*.
3. S. Chow, F. Ruskey. Towards a general solution to drawing area-proportional Euler diagrams. In *Euler Diagrams 2004, ENTCS*, pages 3–18, 2005.
4. J. Flower, J. Howse. Generating Euler diagrams. In *Diagrams 2002*, pp 61–75. Springer, 2002.
5. H. Kestler, A. Muller, T. Gress, M. Buchholz. Generalized Venn diagrams: A new method for visualizing complex genetic set relations. *Journal of Bioinformatics*, 21(8):1592–1595, 2005.
6. O. Lemon, I. Pratt. Spatial logic and the complexity of diagrammatic reasoning. *Machine GRAPHICS and VISION*, 6(1):89–108, 1997.
7. P. Rodgers, L. Zhang, A. Fish. General Euler diagram generation. In *Diagrams 2008*, pp 13–27. Springer, 2008.
8. P. Rodgers, L. Zhang, H. Purchase. Wellformedness properties in Euler diagrams: Which should be used? *accepted for IEEE Transactions on Visualization and Computer Graphics*, 2012.
9. F. Ruskey. A survey of Venn diagrams. *Electronic Journal of Combinatorics*, 1997. www.combinatorics.org/Surveys/ds5/VennEJC.html.
10. G. Stapleton, J. Flower, P. Rodgers, J. Howse. Automatically drawing Euler diagrams with circles. *Journal of Visual Languages and Computing*, accepted 2012.
11. G. Stapleton, P. Rodgers, J. Howse A general method for drawing area-proportional Euler diagrams. *Journal of Visual Languages and Computing*, 22(6):426–442, 2011.
12. G. Stapleton, P. Rodgers, J. Howse, J. Taylor. Properties of Euler diagrams. In *Layout of Software Engineering Diagrams*, pp 2–16. EASST, 2007.
13. G. Stapleton, L. Zhang, J. Howse, P. Rodgers. Drawing Euler diagrams with circles: The theory of piercings. *IEEE Transactions on Visualisation and Computer Graphics*, 17(7):1020-1032, 2011.
14. L. Wilkinson. Exact and approximate area-proportional circular Venn and Euler diagrams. *IEEE Transactions on Visualization and Computer Graphics*, available online, 2011.