

Augmenting Existing Virtual Environments with Plausible Naturalistic Clutter

R.A. Broughton and T.L.J. Howard
*Advanced Interfaces Group, School of Computer Science,
University of Manchester, United Kingdom, M13 9PL
russ@cs.man.ac.uk*

Abstract

We present an approach for enriching a Virtual Environment by augmenting it with the kind of naturalistic clutter often found in real-world scenes. Rather than manually specifying clutter distribution, our semi-automatic procedural approach seeks to add plausible clutter with minimal user-interaction, and few constraints on the format in which the Virtual Environment is specified. We suggest that the integration of suitable and plausible clutter to a Virtual Environment increases realism at low cost.

Keywords: Clutter, Procedural Augmentation, Virtual Environments, Genetic Algorithm, Rigid Body Dynamics

1. Introduction

The representation of a real scene by a Virtual Environment (VE) is almost always an idealized approximation. Despite being geometrically accurate, VEs are often conspicuously pristine in contrast to real-world scenes. This has arisen out of an historical need to streamline the computational and memory demands of a VE, but with a steady increase in processing power the limitation has now shifted to the labor involved in creating the geometry and textures in the scene. Recently, advances have been made in an attempt to diffuse this sense of sterility with techniques such as dirty-looking surfaces [1][2] but scenes are still commonly comprised of a relatively low number of key polygonal models and lack the clutter encountered in the real world.

Manually adding a large number of clutter models to a VE ranges from a chore in the small scale, to being totally unfeasible in the large scale. Our work seeks to

address this problem, and we present an approach for adding potentially massive numbers of models to an existing VE in a manner that minimizes user interaction, while placing the clutter as plausibly as possible within the context of the VE.

2. Clutter

Figure 1 shows an example of a real-world scene that includes the kind of extraneous material commonly considered as clutter – objects such as the boxes and papers on the desk, for example. Throughout this paper, we use the term ‘clutter’ to refer to the models with which we will augment an existing VE, in contrast to ‘furniture’ that refers to the models that already comprise the VE. The clutter models will invariably, although not necessarily, be relatively small in comparison to the furniture.




Figure 1: Clutter in a real-world scene

The composition and positions of clutter are not fixed, but change freely within a framework imposed on the VE that constrains it to appear sensible within the context of the type of furniture present.

2.1 Representing Clutter

A clutter model is defined as any 3D model represented in a format suitable for addition into the existing VE in which it is to be placed. In our prototype implementation, which is based on MAVERIK [3], clutter models can be natively defined in VRML97 [4] and AC3D [5].

The range of VEs to which our method can be applied is directly related to the range of suitable clutter that is made available within a repository of clutter items. These items are stored within a hierarchy, in which any member of a subtree below a node N is considered a subclass of N. Figure 2 illustrates this, with rectangles representing classes and ellipses representing the nodes at which specific clutter items are stored; ‘...’ indicates the omitted remainder of the class.

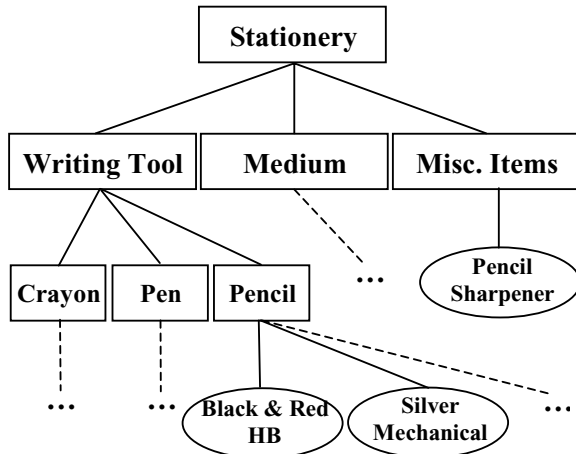


Figure 2: Stationery subsection of clutter

A problem with VEs is that they are conceptually all different, and no absolute scale exists across all possible VEs. To counter this, all clutter models within the repository adhere to a notional scale and therefore are proportional to each other. When adding clutter to a given VE, its scale can be compared to the notional one of the repository and the clutter items scaled accordingly as they are added.

3. Virtual Environment Preparation

Assumptions about the existing VE are kept to a minimum. All that is required is that the furniture polygons be individually accessible and can therefore be placed into a winged-edge data structure [6] (modified to allow holes and non-manifold surfaces)

for fast polygonal adjacency queries. This implies that all polygonal edges should have one immediate neighbor or be on a hole or boundary at the edge of a mesh.

3.1 Furniture Labeling

As clutter is so ubiquitous, it is easy to underestimate the implicit assumptions that are made about it. In addition to the assumption that clutter must appear to be in a state of physical equilibrium (assuming that it is stationary), the clutter selection must also be coherent with the theme of the VE to which it is added.

Although furniture models visibly represent real-world objects, they intrinsically share none of their other real-world properties, and so it is necessary to explicitly define a set of values for each furniture item that can be used to reason about what type of clutter should appear there. We refer to this process as ‘tagging’.

The main purpose of each tag is to associate a composition of clutter with furniture. This takes the form of a number of references to entries in the clutter hierarchy that can vary in resolution from the superclass of all clutter, to specific clutter items. Additionally, each clutter type C is given a fractional weight C_w to indicate how commonly it occurs relative to others. For convenience, we constrain the sum of the C_w to be 1. An example tag for an office desk might associate it with the following clutter composition: 0.7 stationary, 0.2 books, 0.05 snack food and 0.05 computer media.

An additional optional fractional coefficient F_A may also be assigned to a furniture item F, to represent the attraction that the furniture item exhibits to clutter. This is particularly useful if a furniture item is for some reason adverse to clutter, then F_A is set to a low fractional value to reflect this. The default value is set to 1, thus ensuring that tags for which this value is not specified are treated as normal.

When creating a tag set, a theme common to this work is highlighted: the trade-off between time spent preparing the VE and the quality of clutter produced. At one extreme, instead of tagging all the furniture in a scene, the user can simply specify a set of global clutter composition values that will be applied to every surface on which clutter will be added. At the other extreme, every furniture item could be given a specific tag,

custom-developed for the VE in question. It is envisaged that the approach taken by most users would fall somewhere between these two extremes, whereby a reasonably non-specific tag set is developed of commonly reoccurring furniture items, which can be used before the user fine-tunes the composition of any furniture items that have special requirements within the VE.

3.2 Identifying Clutter Regions

We currently limit clutter placement to horizontal flat areas on furniture items which we call ‘regions’. A region is defined as a set of adjacent polygons whose normal vectors are almost parallel (within a tolerance of angle ϕ) with the world up-vector. We identify regions using the following algorithm:

1. Establish three sets to contain references to the constituent polygons of a furniture item: *Model*, *Untested* and *Partial Region*. Only one reference to each polygon exists in the union of the sets.
2. *Model* initially contains all of the polygons of the furniture
3. *Untested* stores all polygons that border the current region and have not had their orientation checked, and is initially empty.
4. *Partial Region* is also initially empty and is used to accumulate groups of polygons that together form a coherent distinct region.
5. Let U represent the normalized world-up vector.
6. Move a random polygon P from *Model* to *Untested*.
7. Let A be a random polygon taken from *Untested*. Calculate N , the normalized normal vector of A .
8. If $(\text{acos}(N \bullet U) < \phi)$ then move A to *Partial Region* and any of its adjacent polygons still present in *Model* to *Untested*, otherwise discard A .
9. Repeat steps 7 and 8 until *Untested* is empty, at which point the polygons in *Partial Region* define a single complete region and should be stored with the other regions for later use in determining where to place clutter. Empty *Partial Region*.
10. Repeat steps 6 to 9 until *Model* is empty, at which point all of the regions on the furniture item have been found.

These regions represent the basic granularity of clutter property specification. Each region initially inherits the clutter properties of the furniture item to which it belongs. The user is then free to amend the properties of each region – in practice assigning lower granularity clutter as required. Note that no polygon is ever considered twice, making it a fast operation.

3.3 Weighting Clutter Regions

Of the number of potential regions identified, many will be unsuitable for clutter either because they may be too small, or because they are wholly or partially internal to a furniture item and so clutter placed on them would not be visible. Similarly, certain regions are more attractive to clutter than others and so should be allocated more clutter than others. To represent this, each region has a weighting allocated to it that represents how attractive it is to clutter relative to the other regions in a scene. This weighting is determined primarily by considering the area of a region, and its visibility. During the calculation of this weight, regions that fall below size and visibility thresholds are discarded.

The approximate area of a region (not the exact area since the polygons may not be coplanar) is calculated as the sum of the areas of the polygons that comprise it, where the area of a single polygon is calculated as described in [7]. If this total area is below the threshold determined by a notional minimum area, scaled relative to the VE, then it is discarded. The weighting derived from region area is based on the region’s size as a fraction of the largest region within the region and therefore need not be scaled.

Depending on the geometry of the furniture model, some of the regions generated may be internal to it or only visible from such a restricted range of viewpoints as to not merit adding clutter there. To decide which regions merit keeping, each needs to have a visibility value determined for it, so that those whose that fall below a certain threshold can be discarded. Calculating an exact value is analogous to calculating a global illumination solution for all of the prospective regions and therefore computationally demanding. Instead, a simplified Monte Carlo-based approach was taken in the results illustrated in this work; we do not evaluate an exact visibility factor for each region, but compute an approximate one by firing rays from a number of evenly spaced points on the surface of a hemisphere centered over the region. These rays are fired at random points within the region and the visibility factor

is determined as the fraction of rays that reach the surface without intersecting another polygon first, as shown in Figure 3. In the figure, three rays are shot of which one intersects the region, which would allocate a visibility factor of 1/3 to the region. The hemisphere can be expanded to test against the area immediately above the region, to boundaries of the furniture item, or be set at the limits of the VE to encompass the whole scene.

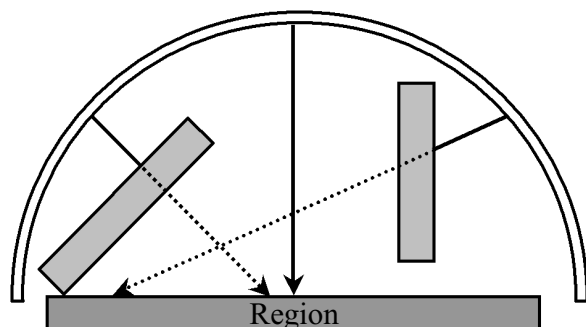


Figure 3: Calculating the region's visibility

'Evenly spacing' points on the surface of a hemisphere is a computationally complex task, and is slightly ambiguous in its definition. The approach we adopt is to space the points along a generalized parametric spiral [8] fitted to the surface of a hemisphere. Evenly spacing the points along the spiral gives a set of points that are approximately evenly distributed on the surface of a hemisphere without the computational overhead of spacing them 'exactly'.

The nature of the approach taken to calculating the region visibility means that it is best suited to a very low or very high number of rays. In Figure 4, the visibility factor was determined by only using a very low number of rays (order of 10), but was then only used as a basis to dismiss regions, rather than contributing to their overall attractiveness. Similarly, using a high number of rays (around 100) will give a good measure of the region's visibility that can be used in weighting the region. However, a number that lies in the middle ground is unsuitable, since it is overkill for just determining visibility, but still too coarse to be trusted as an overall fraction.

Once each region has been scored for attractiveness, this weighting is then multiplied by the clutter attractiveness fraction of the furniture item that it belongs to give an overall score relative to the other regions. The total amount of clutter to add will then be distributed in accordance to the ratio of the regions' values.

In Figure 4, five regions have been identified on a sofa model and colored, where the ratio of green to red represents the attractiveness of the region (green is high-attractiveness, red is low-attractiveness).

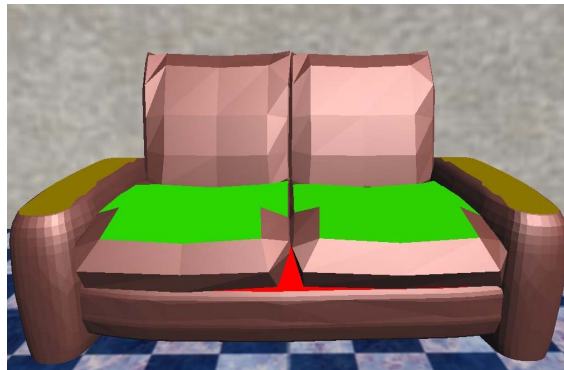


Figure 4: Clutter regions identified on a couch

4. Adding Clutter

We now describe the process of automatically selecting specific clutter items and placing them on regions. The total amount of clutter to add to a scene may be specified either as an absolute number of items or a total volume to add, where the volume of each item added would be approximated to that of its bounding box. This amount is then distributed between the regions present in accordance with the ratio of their respective clutter attractiveness values. Specific clutter items then need to be assigned that conform to each region's clutter composition.

When a clutter item of a specific type is to be selected, the tree is traversed downwards from the node of that type until a leaf node that identifies a clutter item is reached. During this traversal any decisions on which child node to visit are made by a random selection biased in accordance with the probabilistic weightings that each edge possesses. For example, in Figure 2, the edges leading from 'Stationary' to 'Writing Tool', 'Medium' and 'Misc. Items' might be manually weighted as 0.2, 0.7 and 0.1 respectively. These weights must be manually set. In their absence, weightings are calculated in proportion to the number of clutter items in each node's subtree.

4.1 The Dynamics Model

For clutter to appear realistic in an environment, it must appear to be obeying real-world physics, thus avoiding incongruities like floating objects and objects seemingly at rest in an imbalanced state. This implies that an underlying physical model must be applied to

every clutter item inserted into the scene to bring it to a state of rest. In other words, clutter is no longer added directly to a graphical world, but rather to a physical world whose state is represented by the graphical one.

The Open Dynamics Engine (ODE) [9] was chosen to represent the physical world. ODE is a library for simulating rigid body dynamics based on a Lagrange multiple velocity model that offers an integrated collision detection system (implemented using OPCODE [10]) and a choice of time-stepping methods.

ODE decouples the collision geometry from the dynamic body being simulated. In our application furniture items exist only as collision geometry and are therefore immovable objects (because we do not wish the addition of clutter to change the existing objects in the scene). Conversely, each clutter item has a dynamic body, but uses a bounding box for collision detection. The speed increase that this provides at the expense of accuracy was felt worthwhile. In future work we will replace this with a bounding box hierarchy [11] to further narrow the gap between performance and accuracy.

4.2 Adding Clutter

Having determined the composition of clutter to place across a surface and the order in which it will be added, one approach might be to simply scatter the clutter across the region in a completely random fashion. This would be achieved by dropping each clutter item in turn at a random position within the region (the process of identifying a random position is illustrated in Figure 5, where point P might arise from vertex weightings 0.1, 0.2 and 0.7 for W_1 to W_3 respectively) and with a random initial orientation and then allowing the physical model to bring the item to rest before inserting the next, such as might be encountered in a skip or rubbish dump. However, observation has shown that in practice, even the most seemingly messy scenes exhibit a semblance of order, in that reoccurring patterns of clutter can be seen within them. By identifying these patterns and using them to derive a metric that evaluates how closely a prospective clutter item's position/orientation conforms to them, it is possible to optimize the placement of each item.

This metric takes the form:

$$\frac{(A*\text{PatternType1}() + B*\text{PatternType2}() + C*...)}{(A+B+C+...)}$$

where $\text{PatternTypeN}()$ is a 'pattern evaluation function', which monitors how closely the current clutter configuration conforms to a single specific pattern and returns a number in the range [0,1]. Each function's result is multiplied by a coefficient used to weight the importance of that particular pattern within the context of the current clutter (shown as A, B and C in the equation). Hence, every clutter item needs to store a set of values for the coefficients that represents how important each of the factors is to it. Rather than specifying these values for every clutter item, they can be specified at any level within the clutter hierarchy and are subsequently inherited by subclasses that can optionally override them.




Figure 5: Finding a random point in a region

As an example of pattern usage and coefficient weighting, imagine a pattern type that scores the position of clutter based on how closely it appears to an electrical socket. This would be important in the case of adding clutter kettles in a kitchen VE, but not important at all in the case of cutlery clutter and so would have its coefficient set to zero.

4.2.1 Common Clutter Patterns

We consider the following kinds of patterns, derived from observation of real-world human-placed clutter:

- **Regional Containment.** To prevent clutter spilling out of the regions, a fractional score is awarded for the volume of the bounding box that lies within the volume defined by the region's polygons and the global up vector.
- **Positional Conformity.** This refers to the property that clutter tends to appear within

close proximity to other clutter of the same type. To model this, related clutter to be added to each region is grouped into units we call ‘batches.’ The fractional likelihood of a clutter item appearing as part of a batch is controlled by one of its parameters. Hence, positional conformity is measured by calculating the volume of the convex bounding hull of all of the bounding boxes of the items in the batch, as a fraction of their combined individual volumes. When applied loosely, this causes clustering of a batch’s items. When stringently enforced, it eventually causes clutter to organize itself in configurations that minimize volume and results in familiar configurations such as piles and, in the extreme, neat stacks.

- **Default Orientation.** This metric considers the item’s default orientation to be its correct one and measures the difference between this and its present orientation within the VE.
- **Orientation Conformity.** Items in the same batch tend to have the same orientation (such as the piles of paper on the right-hand desk in Figure 1). To capture this, the mean orientation of all of the items in the batch is calculated and the fractional difference between this and the current item’s orientation is returned.
- **Boundary Alignment.** This captures the tendency of, particularly in neat environments, clutter being orientated so that the side (of the notional bounding box containing it) nearest a region edge is aligned to be parallel with it.
- **Disruption Caused.** When a new item is added, a score inversely proportional to the impact that it has on the existing items’ positions and orientations is determined. This can be used to ensure that existing clutter patterns are not disrupted.

Of these patterns, Regional Containment and Positional Conformity are generally most important and their coefficients should be weighted accordingly. More potential patterns can be identified, particularly within groups of identical clutter and so could be modeled by the user, but the patterns listed are considered the most widely applicable.

It should be noted that although the approach described is only concerned with general clutter specified at a regional resolution and not at the resolution of specific clutter batches or even items, user influence could be exerted over regional distribution via the introduction of user-specified patterns. For example, a pattern could be introduced that scores clutter on how closely it appears to a user-specified point or area, causing clutter to tend towards it by an amount proportional to its corresponding coefficient.

4.2.2 Evaluating Clutter

As a means of evaluating the metric for each clutter item to be added, a Genetic Algorithm (GA) was chosen, where the metric forms the fitness function of the GA. We use GALib [13]. The choice of a GA was made as it can quickly find approximate ‘good’ solutions (although not necessarily the best) and because it is good at avoiding the multitude of local optima that the pattern evaluation function will produce and therefore producing very neat clutter if the user is prepared to allocate the necessary processing time.

Figure 6 illustrates in 2D how a genome (in GALib parlance, sometimes called a chromosome) could be developed to describe potential clutter placement positions and orientations. No Y coordinate is specified, but rather it is determined as the minimum value necessary to elevate the item above whatever existing clutter falls within the volume that its cross-section defines in conjunction with the world up-vector (the Y direction in the figure).




Figure 6: 2D Genome composition example

The actual genome used comprises:

- A 2D position within the region's plane, where the 'height' (measured along the world up-vector) is determined as is shown in Figure 6. This position is initialized at random within the region, as shown in Figure 5.
- An orientation for the item's bounding box. This is represented as a random integer in the range 1 to 6, where each value corresponds to a different side of the box facing upwards.
- A rotation in $[0, 360^\circ]$ about the world up-vector that is applied after the orientation has been set.

The restriction on the range of initial orientations is specified as that of a die to limit the search space. The effect that this has on the range of expressible clutter orientations is surprisingly slight, since irrespective of the range of orientations that an item could initially be set to in space, the difference between these and one of the die orientations is generally cancelled out during the process of bringing the object to rest on the region.

No obvious terminating condition exists for the GA. The use of a minimum fitness function value that must be met is risky, since it is generally indeterminable in advance whether the maximum attainable fitness score for a clutter item will meet this. Instead it is better to allocate a fixed amount of processing time to each clutter item, this can be specified in terms of an absolute time, or in terms of the number of generations (and population size) up to which the GA will run. Long processing periods will give 'neat' results, but this is not always what is required, since the neatness of clutter can vary over clutter batches and regions within the VE. To model this, each region and clutter type needs to be assigned a neatness value F_N and C_N respectively. These are used in the calculation of an overall neatness, O_N , fraction for each batch, where \bar{C}_N is the mean of all of the C_N values in the batch:

$$O_N = \bar{C}_N \times F_N$$

For each VE minimum and maximum processing bounds can then be set for the placement of each clutter item and O_N used to determine as a fractional value between these the amount of processing required.

5. Results

Figure 7 shows the result of adding ten books to the end of a shelf with the clutter metrics weighted to cause the books to appear in a row (an additional metric was introduced that ensured the books' spines faced away from the wall). The current approach is ill suited to producing such a specific pattern since it requires the books to be placed exactly next to each other and, as is invariably the case with global optimization problems, the amount of processing required increases dramatically when an 'exact' solution is required. Hence, even though a relatively large population size of 35 was used over 30 generations the result still appears slightly messy. Although, arguably such a specific pattern requirement doesn't fall under our definition of clutter, since the user has a definite preconception of *exactly* how it should appear.




Figure 7: Clutter books on a shelf

Conversely, Figure 8 shows the result of a relatively low amount of computation (an average population size of 12 over 12 generations was used). Despite the relative loose clustering of the book piles they still appear plausible since no exact configuration is expected of them.




Figure 8: Desktop clutter

6. Conclusion & Future Work

The approach outlined in this paper is capable of semi-automatically producing plausible results and is extensible to a wide range of VEs and clutter models. It provides very significant timesavings over manual modeling and incorporation of clutter items for a VE, although the user must still invest time in devising suitable furniture and clutter parameters and also ensuring that a range of suitable clutter models is available. We are presently investigating methods of model re-use, with techniques such as procedural generation of clutter models presenting promising avenues of research [14], to further minimize the work required on the user's part.

The suitability of the clutter patterns used in the placement stage appears to be the predominant factor in determining how realistic the resultant clutter appears. In light of this, current work is focusing on refining and extending the general set of patterns described.

Although the approach outlined is intended to produce clutter that appears human-placed, it could potentially be extended to model natural phenomena, such as leaves around a tree. Care must be taken though, since the metric necessary to evaluate leaf placement may well be more complex and slower than directly physically simulating the falling of leaves within the VE. Future work will investigate the possibility of introducing a hybrid method that combines procedural methods that explicitly place clutter in certain preset patterns, thereby helping to counter problems such as Figure 7 highlights.

As graphics technology advances, the limiting factor for creating realistic Virtual Environments is moving from what can be displayed by the computer to what can be created by the user. In light of this, we propose that continuing research into techniques that simplify the creation of convincing virtual environments is essential, and an exciting topic.

7. Acknowledgements

It is a pleasure to acknowledge our colleagues in the Advanced Interfaces Group, whose comments, suggestions, support and friendship have been invaluable.

8. References

- [1] J. Dorsey, H.K. Pedersen, P. Hanrahan, "Flow and Changes in Appearance", in *ACM Trans. Graph. (Proc. SIGGRAPH)*, pages 411-420, 1996.
- [2] J. Dorsey, A. Edelman, H.W. Jensen, J. Legakis et al, "Modeling and rendering of weathered stone", in *ACM Trans. Graph. (Proc. SIGGRAPH)*, pages 225-234, 1999.
- [3] R.J. Hubbard, J. Cook et al., "GNU/MAVERIK: A Micro-Kernel for Large-Scale Virtual Environments", *Presence: Teleoperators and Virtual Environments*, **10**, pages 22-34, February 2001, ISSN 1054-7460.
- [4] VRML97 International Standard ISO/IEC 14772-1, www.web3d.org/x3d/specifications/vrml/, 1997.
- [5] "AC3D", Lancaster University, England, <http://www.ac3d.org>.
- [6] B. Baumgart, "A polyhedral representation for computer vision", *Proceedings of the AFIPS National Computer Conference*, pages 589-596, 1975.
- [7] J. O'Rourke, "Area of a Polygon", *Computational Geometry in C* (2nd Edition), Section 1.3, Cambridge University Press, ISBN 0-52164-0105 1998.
- [8] E.B. Saff and A. B. J. Kuijlaars, "Distributing Many Points On a Sphere", *The Mathematical Intelligencer* **19**(1), pages 5-11, 1997, ISSN 0343-6993.
- [9] R. Smith, "The Open Dynamics Engine", <http://ode.org/>, 2001.
- [10] P. Terdiman, "OPCODE User Manuel", <http://www.codercorner.com/Opcode.htm>, 2002.
- [11] M.C. Lin, S. Gottschalk, D. Manocha, "OBBTree: A hierarchical structure for rapid interference detection", in *ACM Trans. Graph. (Proc. SIGGRAPH)*, pages 171-180, 1996.
- [12] G. Turk, "Generating Random Points in Triangles", *Graphics Gems I* (Andrew S. Glassner, Ed.), pages 24-28, Morgan Kaufmann, ISBN 0-12286-1663, 1990.
- [13] M. Wall, "GALib: A C++ Library of Genetic Algorithm Components", <http://lancet.mit.edu/ga/>, 1996.
- [14] B. Cutler, J. Dorsey, L. McMillan et al, "A procedural approach to authoring solid models", *ACM Transactions on Graphics*, **21**(3), pages 302-311, July 2002.