Non-Realistic 3D Object Stylization

Julian Kratt^{1*} Ferdinand Eisenkeil¹ Sören Pirk¹ Andrei Sharf² Oliver Deussen¹
¹University of Konstanz, Germany ²Ben-Gurion University, Israel

Figure 1: Our method creates abstract stylized objects from a given input model (left). We analyze the shape and its geometry to guide the stylization and abstraction of the object. Essentially, the user makes a selection from a prioritized list of style operands and applies it on the object. The stylized versions of the input can be rendered in various ways using non-photorealistic rendering.

Abstract

In this paper we introduce the novel paradigm of non-realistic 3D stylization, where the expressiveness of a given 3D model is manifested in the 3D shape itself, rather than only in its rendering. We analyze the input model using abstraction, simplification, and symmetrization operators to determine important features that are later represented by new geometry. Doing so, we create a stylized and expressive representation of the input that can be rendered or might be printed using a 3D printer. We conducted a user study to verify the proposed stylizations and demonstrate the approach by using standard geometry of buildings as well as natural and technical objects.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems; I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques;

Keywords: non-realistic rendering, geometric modeling, solid modeling, digital surface processing, surface parametrization

1 Introduction

For a long time, one of the primary goals of computer graphics research has been the generation of realistic models and photorealistic images. With the evolution of 3D digital content and the emergence of novel technologies such as 3D printing and urban modeling, symbolic shape representations have also come to receive more attention. This shift from realistic rendering to stylized representations has been apparent in the field of non-photorealistic rendering

(NPR). Corresponding methods typically assume that the geometric model is given or render the model in a particular style, focusing on aesthetic or technical aspects to increase the clarity of the subject.

We introduce a novel paradigm of non-realistic 3D stylization, where the expressiveness of a model is manifested in the 3D shape itself, rather than only in its rendering. Since the input model is altered to express certain features and semantics, we call this approach non-realistic stylization. Essentially, 3D stylization provides alternate shape representations moving away from classical surface meshes towards high-level semantic models. One of the first algorithmic 3D stylizations to visualize phenomena such as hair, fire or glass was proposed by Perlin and Hoffert [Perlin and Hoffert 1989]. Among many others, abstraction, simplification, and symmetrization are means for modeling an object in an expressive style. These techniques serve as non-realistic representations that spark and inspire human perception of shape and are thus powerful visual communication tools for creating a variety of digital content.

Due to complexity and diversity of shapes, the stylization of 3D models is a challenging problem. While the abstraction and simplification of geometry is an interesting domain of research [Mehra et al. 2009a] our stylization does not simplify the object and may even increase its complexity in order to represent its essence (Figure 1). Our work is inspired by the method of Nan et al. [2011] who analyze the intricate relations in 2D urban drawings for their simplification. Since in 3D this complexity may increase exponentially, we take a semi-automatic approach. We analyze the 3D shape and provide a list of compatible stylization suggestions for the user. The stylizations are either based on clipping the fill patterns against the input or by growing them explicitly. Our method introduces a novel editing framework for stylization and abstraction of 3D structures by geometry replacement. While our work borrows from NPR where a wide variety of expressive styles and filters are applied to render raster 2D images, we develop high-level editing operands for semantic modeling. These operands provide a wide range of which amend 3D shape while accounting for specific geometrical, topological and volumetric features.

We performed an initial user study to clarify which geometric structures are commonly used to stylize models. Based on the findings our set of fill patterns was designed and adapted. While the presented pipeline allows generating visually pleasing stylizations for a large variety of input models, the findings of the user study help the system to suggest good initial fill patterns that later can be refined by the user.

*e-mail: julian.kratt@uni-konstanz.de

Figure 2: *Hand-drawn stylizations of an input model created by subjects of an initial user study.*

2 Related Work

Related work comes from the fields of non-photorealistic rendering, 3D abstraction as well as urban modeling techniques.

Non-Photorealistic Rendering: A major goal of non-photorealistic rendering (NPR) methods lies in highlighting, exaggerating and clarifying object characteristics and features [Gooch and Gooch 2001; Strothotte and Schlechtweg 2002]. Many styles have been explored in NPR literature to serve different artistic needs, such as pen and ink [Salisbury et al. 1997; Winkenbach and Salesin 1994], painting [Meier 1996], informal sketching [Raskar and Cohen 1999], and charcoal drawings [Cornish et al. 2001].

Since low-level geometry does not provide natural prioritization of shape features, NPR techniques often focus on highlighting view-specific features to exaggerate or convey form [Cole et al. 2012]. A significant amount of research has been devoted to identifying object features. This ranges from silhouettes, ridges and valleys [Na et al. 2005] to contours and suggestive contours [DeCarlo et al. 2003]. 2D abstraction is a continuing line of research in non-photorealism on clarifying and exposing the essential structures in images [DeCarlo and Santella 2002], segmentation-based rendering [Kolliopoulos et al. 2006], enhanced representations of photographs [Orzan et al. 2007], video [Wang et al. 2004], [Winemöller et al. 2006], and line drawings [Barla et al. 2005; Jeong et al. 2005; Lee et al. 2007]. Bengtsson et al. [1991] obtained abstractions by studying contours at different scales and recently attempts have been made to learn abstractions using a set of exemplars [Fatih Demirci et al. 2009] or by performing abstraction through organizing shapes and analyzing local and global features of images [Mi et al. 2009].

Other avenues explored for abstraction include rule-based simplification [Brown et al. 1993], user-guided parametric models [Falcid et al. 1998], or topology-based inference [Biasotti et al. 2000]. These approaches aim to abstract the content of 2D imagery using segmentation, clustering, and scale-space processing to find regions or lines that can be omitted from imagery.

Geometry-based Stylization: A large amount of work addresses the abstraction and stylization of 2D line drawings. In the context of perceptual abstraction, Grabli et al. [2004] simplify line drawings using a complexity measure which accounts for stroke density and regularity variations. Barla et al. [2005; 2006], present algorithms for line drawing simplification and synthesis based on perceptual line grouping, accounting for proximity, color and continuation principles. Many of the existing methods analyze the geometry for grouping and simplification [Shesh and Chen 2008], the arrangement of patterns and the relations of neighboring elements [Ijiri et al. 2008] or the appearance and placement of stroke-based vec-

tors [Hurtut et al. 2009]. Wang et al. [2004] perform abstraction of video sequences by semi-automatic segmentation of semantical contiguous volumes.

While most of the existing work in non-realistic stylization focuses on 2D models, only a few methods address the stylization and generation of iconic representations of 3D models [Fleischer et al. 1995; Cutler et al. 2002]. In an early work, Akleman et al. [2004] generate 3D face caricatures to highlight artistic concepts in the modeling process. Gal et al. [2007] create 3D collages on top of target shapes by using a database of objects as primitive building blocks. Theobalt et al. [2007] extend this idea and propose a method for automatically transforming animated meshes into abstract representations. In a recent work, Mehra et al. [2009b] create envelope shapes for complex 3D objects to guide their visual simplification. McCrae [2011] introduced the use of planar section for creating shape abstractions and even more recently Vidimce et al. [2013] provide means for the synthesis of multi-material 3D printed objects.

Motivated by procedural modeling and constructive solid geometry, researchers have long proposed to approximate a given 3D model with parametric parts [Attene et al. 2006; Schmidt 2010]. Parametric descriptions enable the creation of different abstraction levels by direct manipulation, for instance by removing some of the parts while preserving others. However, most discrete digital models lack such semantic information and deducing regularity from 3D geometry poses a difficult problem [Pauly et al. 2008]. Nan et al. [2011] introduced a method for structural summarization and the abstraction of complex spatial arrangements found in architectural drawings. Their method is based on the well-known Gestalt rules, which summarize how forms, patterns, and semantics are perceived by humans from bits and pieces of geometric information.

Urban Abstraction: The abstraction and simplification of buildings and urban scenes has been of interest to researchers for improving, clarifying, and emphasizing visual representations of urban data sets. While Adabala et al. [2007] propose to create maps as a combination of two- and three-dimensional information to increase functionality and aesthetic appeal, Grabler et al. [2008] introduce a technique for the automated design of tourist maps with focus on emphasizing important cues. The approach of Loya et al. [2008] employs Fourier series to identify periodic features and to approximate facade structures from images. Glander et al. [2009] focus on the perception of city models by introducing a hierarchical abstraction of buildings and streets. Their work allows computing simplified representations of important landmarks and supports smooth transitions of varying levels of detail. Sidiropoulos and Vasilakos [2006] provide an overview of different techniques for modeling cities and discuss various symbolic and realistic visualization.

In contrast to previous techniques we focus on the stylization of 3D input models. Our method not necessarily produces simplified but complex, however, visually pleasing replicas that can also be printed on custom hardware.

3 Stylization Patterns Drawn in Initial Study

The number of possible primitives that could be used for stylization is rather limited since the entire process is subjective. However one can observe that several patterns are more suited than others. In order to find out which geometric structures are commonly used we performed a study with fifteen students from a design and architectural school in total that are skilled in drawing by hand. The subjects were introduced to our new paradigm of 3D model stylization and asked to draw stylized versions of four given input models: two buildings, a tree model (based on the geometric modeling of trees proposed by Livny et al. [Livny et al. 2011]) and a wooden toy. The examples are shown in Figure 1, Figure 10 (upper right) and Figure 11 (lower left). These objects were printed on papers with less opacity to guide the subjects. In addition, to help subjects to gain better spatial understanding, we provided videos showing each object rotating around its medial axis. The subjects were asked to apply geometrical shapes in order to provide a 3D abstraction on paper. Additionally they hatched the shapes to give a cue about the spatial impression.

Figure 2 shows some results of the study. The input model is represented on the very left side in Figure 1. Throughout the study some patterns were prominently utilized to highlight or exaggerate certain features of the models. For example to express the essence of buildings, such as floors or repetitive windows, most of the subjects drew structures like planes, crossing planes, tubes and parallel rings in the way that those characteristics were preserved. Especially for irregular structures, the participants suggested filling geometry such as spirals and splines.

Similar to 3D buildings with symmetric structures such as windows repetitive patterns in other arbitrary 3D objects are treated in the same way by the subjects. Another observation is that symmetric parts of the input model are stylized in the same way. Additionally, one important result of the study is that the patterns are very often arranged along the main axis of the selected segment.

Based on our findings we propose eight fill patterns (see Figure 3). Initially we had a larger set of fill patterns from which two were additionally implemented: 3D Voronoi cells and organic growth patterns.

4 Overview

A sequence of four general steps is performed to produce a stylized object (see Figure 4). We assume that the model is already divided into a number of segments. Segmentation can be performed automatically using common segmentation techniques ([Nan et al. 2011]) or defined by the user in a pre-processing step. All following steps can be computed within a few minutes, so the possibility for user interaction with fast feedback is preserved.

Shape Analysis and Fill Pattern Ranking: We analyze each segment individually to determine the location and orientation of filling primitives. To identify global and local features we make use of symmetry detection and global-to-local features such as repetition patterns and volume distribution. The goal of the analysis is to provide the user a ranking of fill patterns that preserves the characteristic features of the given segment. We may use different analysis of the shape such as local curvature, earth-movers distance [Rubner et al. 1997] to database objects, which are already abstracted, as

well as upright orientation [Fu et al. 2008]. In our work we use only a subset of possible alignment methods, RANSAC [Fischler and Bolles 1981] and PCA [Pearson 1901], in order to compute the best fitting representation and orientation of the fill patterns. After the analysis we provide the user with a ranking of possible patterns.

Figure 4: *System Overview: Given a pre-segmented input model, we initially analyze the geometry of each part and compatible stylization patterns are suggested to the user. Next, we apply the patterns as volumetric operands inside the shape. For an aesthetic appearance the generated abstraction model can be rendered in a number of different ways.*

Primitive Generation: Once we know what kind of fill pattern can be used to abstract the segment, the geometry of the fill pattern is generated. This pattern is later used to be clipped against the input model. The orientation of the PCA from the previous step as well as repetition and repetitive pattern analysis allows us to fit the fill patterns to the structures of the input model. Floor height or prominent facade elements are such primary and repetitive structures. The user is able to select multiple segments at a time to apply the same selected pattern to all of them.

Another possibility is to grow a fill pattern inside the segment. This has the advantage that no clipping is needed later and that the generated geometry can interact with the segment boundary. We direct the growth of such patterns by using distance fields from the segment boundary and predefined growth directions.

Filling and Representation: To fit the geometry of the fill patterns to the model we use the input mesh for clipping. The clipped geometry can be used for representing the model; it can be exported for 3D printing or subsequent rendering methods.

Stylization: In the last step we use the clipped and filled geometry to generate a stylized version of the input. At this step common 2D NPR approaches such as toon shading or hatching enhance the visual appearance of the abstraction.

5 Shape Analysis and Pattern Ranking

Proper stylization depends on the shape of the segments and of the combination of different elements to a stylized model. While the latter is a creative decision the former can be supported by a geometric analysis of the segment shapes. Thus we perform some tests to the geometry. By using RANSAC [Fischler and Bolles 1981] we determine the geometric primitive that is most similar to the segment, while PCA [Pearson 1901] determines the direction with the highest variances as main axes of the object and also the ratio between the dimensions of these axes.

Figure 3: Fill patterns used for stylization. These patterns are visualized in the user interface as buttons. The user selects a pattern that is then applied to the input model. Left: Clipping patterns that are derived from the user study. Right: Spiral and spline structures. The presented growing patterns (Voronoi and branching structures) are not proposed by our ranking because they are not based on the result of the user study.

Both values help us to select good fill patterns from our pre-defined library. It turns out that the regularity of the fill patterns is related to the regularity of the segment geometry. As our study shows, users prefer more regular patterns for box-like or cylinder-like segments and more organic patterns for irregular segment shapes.

Figure 5: Schematic diagram that shows the outcome of RANSAC and PCA. While the input mesh is given on the left, the center illustrates the shapes detected with RANSAC and the right-hand side points out the main directions of PCA for each segment.

5.1 Geometric features for parameter derivation

To determine the general shape of the segment, we use RANSAC. This helps us to limit the set of possible fill patterns that could be used for stylization. RANSAC is an iterative algorithm to estimate parameters of a mathematical model from a set of data points. In our case, the data points are sampled from the input mesh and fitted to implicit surfaces. We use RANSAC with six randomly chosen data points instead of solving the equation system exactly with every vertex of the mesh in order to increase computation speed. The algorithm converges either after a finite number of iterations or when a matching model is found. The other reason is that if we use all data points, the shape may not match the surface described by an implicit function completely. In our system we are able to identify whether a segment has a box-, ellipsoid- or a cone-like structure. This could easily be extended to more shapes.

In addition to RANSAC we perform PCA on the input. We take advantage of the characteristic of the PCA that the principal components are orthogonal to each other. The resulting orthogonal main axes help to direct our orthogonal filling structures like crossing planes. Taking the eigenvalues of the principal components we are able to produce a proposal of an optimal distribution of these primitives along the main directions. Both of the analysis methods are illustrated in Figure 5. It is necessary to provide both methods, for instance a box and a sphere can have the same principal components, while they are completely different forms.

5.2 Semantic Features

The stylization should preserve characteristic semantic features of the input model, such as repetitive patterns like windows on a facade. To identify such structures we exploit the help of the user. The user can draw simple lines on the object surface to mark prominent structures. Our system can subsequently detect the repetitive patterns of these structures with a very simple adaption of the Gestalt rules [Nan et al. 2011]. As illustrated in Figure 6, only two edges of a window in the facade of a building need to be marked in order to find all of the necessary edges that will give us information about the vertical distribution of planes in the segment.

Figure 6: The inset illustrates where the user draws lines (red) and which features are detected (blue). In the complete building repetitive patterns of these two selected edges are detected.

5.3 Ranking

We propose a ranking for the different filling primitives. The ranking provides an estimation about the quality of a fill pattern for a selected segment. The user has to decide on which method the ranking is based, whether PCA, information derived on selected repet-

itive patterns, or RANSAC. For instance if the three eigenvalues calculated by the PCA are identical, we propose crossed planes as fill pattern. Only one large eigenvalue would lead to parallel planes along the main principal component. If RANSAC estimates a cone as the basic structure our primary proposal is a sequence of rings. The ranking is based on our initial evaluation which is described in Section 3.

6 Clipping and Filling

The ranking that is produced on the basis of PCA and RANSAC provides us with different patterns for abstraction. These patterns are divided into two groups. The first group consists of parametric models, such as combination of planes and rings, and can be clipped against the input surface. The other group is composed of growing and fill patterns, which have a more organic appearance. These organic patterns can be further subdivided into explicitly given geometry and patterns, which are computed from given seed points within the volume. Using the distance field of the object, we can control this process. In contrast to discrete clipping, which is fast to compute, filling with growth patterns takes more computational time, but does not need any clipping.

6.1 Oriented Growth Patterns

The first type of filling primitive is any kind of spline-like structure. In its simplest form we sample points inside the volume and interpolate these by a spline curve. To give the splines thickness we use the *Parallel Transport Frame* (PTF) approach [Hanson and Ma 1995] to compute smoothly varying coordinate frames over the spline curve, which are used to define generalized cylinders. Besides the appearance of the spline curve the shape of the filling is mainly given by the sampling of the points. The distance field helps us to detect whether a given sample point is inside the volume or not. Further we can define a range ϵ around the zero-level set that goes inside the volume (see Figure 7). All our models are normalized to a unit cube and we initially suggest a value of 0.01 for ϵ . Thus the resulting structure is quite shape-aware, besides larger distances to the surface leads to a more general filling.

Figure 7: We use distance fields to guide our growth patterns. Left: Geometry is distributed only within a small distance ϵ to the surface. Right: Spiral structures.

Spiral-like structures represent the next type of fill pattern. In order to resemble the given segment well we produce spirals at its border, which however, still remain within the volume. Starting from a base point p in the segment and a given direction n we compute the largest enclosed diameter for the spiral along the axes by utilizing the distance field (see Figure 7). Therefore we start with an initial radius r_{init} and check how many points on the corresponding circle lie inside and outside the segment. If the ratio between these two numbers is under a given threshold we increase the radius gradually until we find the optimal one. Once we have found the maximal possible radius we proceed with the next base point, sampled along the direction n , in the same way. Finally, we have a set of base

points and their corresponding radii, which are now interpolated to build the spiral.

6.2 Seed-Point-Based Growth Patterns

For more artificial stylizations without correlation to geometric features of the input model, organic structures may be wanted. In our system we use a space colonization algorithm to generate tree-like branching structures within the segment. The method was originally introduced to generate leaf venation patterns [Runions et al. 2005; Runions et al. 2007]. In tree modeling, the general idea is that branches compete for space while they are growing. To simulate this complex process seed points are distributed in space, where each point acts like an attractor. Starting from one point in space the tree skeleton is iteratively developed by adding new branches oriented by nearby attractors within a radius of influence. Once attractor points are reached they are excluded from further computations. Figure 8 illustrates the evolution of the space colonization pattern over time.

3D Voronoi volumes [Du et al. 1999] are also seed-point-based filling patterns. These structures are provided as an organic fill pattern for arbitrary segment shapes. While these Voronoi structures are very aesthetic, their calculation is very time and space consuming. The Voronoi structures are controlled by distributing the seed points in the same way as already described above. These structures as well as the space colonization filling are not derived by geometric features and only the user is able to decide whether they should be used for stylization.

Figure 8: Branching structures generated by the space colonization method. We use distance fields to control the sampling of seed points. Left: Input model. Center and right: Seed points are only generated in regions with small distance to the surface, which results in organic structures appearing at the boundary of the input.

7 Results

With support of the user interface, we applied our abstraction system to a wide range of different input models. Besides arbitrary objects, the meshes presented in [Mehra et al. 2009a] are used to generate abstracted geometry. In terms of the appearance of the final rendered result, the user is able to post-process the abstraction geometry. Beneath offline ray-tracing, it is also possible to apply different 2D non-photorealistic methods like suggestive contour detection or toon shading.

All models are abstracted in different ways with the primary proposal, that are generated by our system, as well as other abstraction

possibilities. Figure 9 and 10 show different input models and the offline rendered abstractions. Most of the input meshes model real towers, such as the Arc de Triomphe and the Petrona Towers. A more complex result of our system can be seen in Figure 13. Here the user selected growing structures to represent the castle.

Figure 9: Stylization of two building models in different ways of abstraction.

We are also able to apply our method to botanical tree models. For this we use the intermediate cluster representation for tree modeling proposed by Livny et al. [Livny et al. 2011]. In their work, a tree is represented by a skeleton graph, which represents the main branching structure and a set of lobes. The lobes represent the area of the tree’s foliage. The clusters define a closed volume that is filled with our primitives. All of these clusters can be treated as one segment or individually. Figure 11 shows two input trees, their intermediate representations and two stylizations per tree. 3D printing is also possible. As illustrated in Figure 12, we did this with the lower left tower model given in Figure 10.

Figure 12: Two 3D printed building stylizations.

7.1 Evaluation

In order to verify the ranking of the fill patterns as well as the pattern itself, proposed by our system, we performed a small study with the same fifteen users from the initial study. Given a segmented input model the users were asked to interactively stylize the object using our proposed framework but without having any suggestions from the system. All fill patterns were shown from our library. To represent characteristic features in architectural models such as windows and ceilings the major part of participants used variants

and combinations of planes and tubes. The users selected the type of primitive with respect to the spatial extension of a segment. For instance the spherical wheels of the elk model (see Figure 1) were represented in most cases by crossed planes. This reflects nicely the compact structure of the wheels without a prominent axis.

For organic models such as trees the users tended to select more irregular patterns (e.g. splines). This kind of selection matches the top three suggestions given by the ranking of our proposed system.

The user study also shows that stylization is a very subjective process. When applying patterns to a segment the discrepancy of spatial parameters is quite large while the choice of the pattern itself is often similar. For instance, users applied planes to the same segment with individually different spacing.

Figure 13: Stylization of a more complex model.

8 Conclusion

By replacing input geometry with stylized geometric fill patterns, we provide a step towards high-level abstraction of geometry in addition to non-photorealistic rendering. This opens an additional space for abstraction and stylization and helps to create abstract models for 3D printing. We applied two powerful methods in a semi-automatic way to generate fill patterns. Based on geometric features we generate new primitives, which are clipped against the input mesh. The second approach is the use of distance fields: By using an implicit function, an explicit filling mesh is generated. For analysis of the input models, the current system is a very basic implementation, drawing on only a few analysis methods. In future, this will have to be extended. Furthermore, the user should be able to create style configurations as a proposal of the filling shape. Even if our approach results in printable 3D models there is still enough space for improvements. In future work we will focus on fill patterns optimized for 3D printing.

Acknowledgements

We thank the anonymous reviewers and Jonas Christen for their valuable suggestions. Also we thank the design students from Zurich University of the Arts for attending the study and all students that supported us. Additionally we want to mention aimatshape (<http://www.aimatshape.net>) for providing input models for our system.

References

- ADABALA, N., VARMA, M., AND TOYAMA, K. 2007. Computer aided generation of stylized maps: Research articles. *Comput. Animat. Virtual Worlds*, 133–140.
- AKLEMAN, E., AND REISCH, J. 2004. Modeling expressive 3d caricatures. In *ACM SIGGRAPH 2004 Sketches*, SIGGRAPH ’04, 61–.

Figure 10: Stylization of various building models in different ways of abstraction. For each group of buildings, the left model is the input model. The right ones are abstractions with different filling patterns and renderings.

Figure 11: Stylization of tree models. Left to right: input tree model, intermediate cluster-based representation and two stylizations rendered in different ways.

- ATTENE, M., FALCIDIENO, B., AND SPAGNUOLO, M. 2006. Hierarchical mesh segmentation based on fitting primitives. *Vis. Comput.* 22, 181–193.
- BARLA, P., THOLLOT, J., AND SILLION, F. X. 2005. Geometric clustering for line drawing simplification. In *ACM SIGGRAPH 2005 Sketches*, SIGGRAPH '05.
- BARLA, P., BRESLAV, S., THOLLOT, J., SILLION, F. X., AND MARKOSIAN, L. 2006. Stroke pattern analysis and synthesis. *Comput. Graph. Forum*, 663–671.
- BENGTSSON, A., AND EKLUNDH, J.-O. 1991. Shape representation by multiscale contour approximation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 85–93.
- BIASOTTI, S., FALCIDIENO, B., AND SPAGNUOLO, M., 2000. Shape abstraction using computational topology techniques.
- BROWN, G., FORTE, P., MALYAN, R., AND BARNWELL, P. 1993. A non-linear shape abstraction technique. In *Proceedings of the 5th International Conference on Computer Analysis of Images and Patterns*, CAIP '93, 223–230.
- COLE, F., GOLOVINSKIY, A., LIMPAEGER, A., BARROS, H. S., FINKELSTEIN, A., FUNKHOUSER, T., AND RUSINKIEWICZ, S. 2012. Where do people draw lines? *Commun. ACM* 55, 1 (Jan.), 107–115.
- CORNISH, D., ROWAN, A., AND LUEBKE, D. 2001. View-dependent particles for interactive non-photorealistic rendering. In *Graphics Interface 2001*, GRIN'01, 151–158.
- CUTLER, B., DORSEY, J., McMILLAN, L., MÜLLER, M., AND JAGNOW, R. 2002. A procedural approach to authoring solid models. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, NY, USA, SIGGRAPH '02, 302–311.
- DECARLO, D., AND SANTELLA, A. 2002. Stylization and abstraction of photographs. *ACM Trans. Graph.* 21, 3 (July), 769–776.
- DECARLO, D., FINKELSTEIN, A., RUSINKIEWICZ, S., AND SANTELLA, A. 2003. Suggestive contours for conveying shape. *ACM Trans. Graph.* 22, 3 (July), 848–855.
- DU, Q., FABER, V., AND GUNZBURGER, M. 1999. Centroidal voronoi tessellations: Applications and algorithms. *SIAM Review* 41, 4, 637–676.
- FALCID, B., FALCIDIENO, B., AND SPAGNUOLO, M. 1998. A shape abstraction paradigm for modeling geometry and semantics. In *Proceedings of the Computer Graphics International 1998*, CGI '98, 646–.
- FATIH DEMIRCI, M., SHOKOUFANDEH, A., AND DICKINSON, S. J. 2009. Skeletal shape abstraction from examples. *IEEE Trans. Pattern Anal. Mach. Intell.*, 944–952.
- FISCHLER, M. A., AND BOLLES, R. C. 1981. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* 24, 6 (June), 381–395.
- FLEISCHER, K. W., LAIDLAW, D. H., CURRIN, B. L., AND BARR, A. H. 1995. Cellular texture generation. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, NY, USA, SIGGRAPH '95, 239–248.
- FU, H., COHEN-OR, D., DROR, G., AND SHEFFER, A. 2008. Upright orientation of man-made objects. *ACM Trans. Graph.* 27, 3.
- GAL, R., SORKINE, O., POPA, T., SHEFFER, A., AND COHEN-OR, D. 2007. 3d collage: expressive non-realistic modeling. In *Proceedings of the 5th international symposium on Non-photorealistic animation and rendering*, NPAR '07, 7–14.
- GLANDER, T., AND DLLNER, J. 2009. Abstract representations for interactive visualization of virtual 3d city models. *Computers, Environment and Urban Systems*, 375–387.
- GOOCH, B., AND GOOCH, A. 2001. *Non-Photorealistic Rendering*. AK Peters Ltd.
- GRABLER, F., AGRAWALA, M., SUMNER, R. W., AND PAULY, M. 2008. Automatic generation of tourist maps. *ACM Trans. Graph.* 27, 3 (Aug.), 100:1–100:11.
- GRABLI, S., DURAND, F., AND SILLION, F. X. 2004. Density measure for line-drawing simplification. In *Proceedings of the Computer Graphics and Applications, 12th Pacific Conference*, PG '04, 309–318.
- HANSON, A. J., AND MA, H. 1995. Parallel transport approach to curve framing. Tech. rep.
- HURTUT, T., LANDES, P.-E., THOLLOT, J., GOUSSEAU, Y., DROUILLHET, R., AND COEURJOLLY, J.-F. 2009. Appearance-guided synthesis of element arrangements by example. In *Proceedings of the 7th International Symposium on Non-Photorealistic Animation and Rendering*, NPAR '09, 51–60.
- IJIRI, T., MECH, R., IGARASHI, T., AND MILLER, G. S. P. 2008. An example-based procedural system for element arrangement. *Comput. Graph. Forum*, 429–436.
- JEONG, K., NI, A., LEE, S., AND MARKOSIAN, L. 2005. Detail control in line drawings of 3d meshes. *The Visual Computer*, 698–706.
- KOLLIPOPOULOS, A., WANG, J. M., AND HERTZMANN, A. 2006. Segmentation-based 3d artistic rendering. In *Proceedings of the 17th Eurographics conference on Rendering Techniques*, EGSR'06, 361–370.
- LEE, Y., MARKOSIAN, L., LEE, S., AND HUGHES, J. F. 2007. Line drawings via abstracted shading. *ACM Trans. Graph.* 26, 3 (July).
- LIVNY, Y., PIRK, S., CHENG, Z., YAN, F., DEUSSEN, O., COHEN-OR, D., AND CHEN, B. 2011. Texture-lobes for tree modelling. *ACM Trans. Graph.* 30, 4 (July), 53:1–53:10.
- LOYA, A., ADABALA, N., DAS, A., AND MISHRA, P. 2008. A practical approach to image-guided building façade abstraction. In *Computer Graphics International*.
- MCCRAE, J., SINGH, K., AND MITRA, N. J. 2011. Slices: A shape-proxy based on planar sections. *ACM Trans. Graph.* 30, 6 (Dec.), 168:1–168:12.
- MEHRA, R., ZHOU, Q., LONG, J., SHEFFER, A., GOOCH, A., AND MITRA, N. J. 2009. Abstraction of man-made shapes. *ACM Trans. Graph.* 28, 5 (Dec.), 137:1–137:10.
- MEHRA, R., ZHOU, Q., LONG, J., SHEFFER, A., GOOCH, A., AND MITRA, N. J. 2009. Abstraction of man-made shapes. *ACM Trans. Graph.* 28, 5 (Dec.), 137:1–137:10.

- MEIER, B. J. 1996. Painterly rendering for animation. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, SIGGRAPH '96, 477–484.
- MI, X., DECARLO, D., AND STONE, M. 2009. Abstraction of 2d shapes in terms of parts. In *Proceedings of the 7th International Symposium on Non-Photorealistic Animation and Rendering*, NPAR '09, 15–24.
- NA, K. G., JUNG, M. R., LEE, J., AND SONG, C. 2005. Redeeming valleys and ridges for line-drawing. In *Proceedings of the 6th Pacific-Rim conference on Advances in Multimedia Information Processing - Volume Part I*, PCM'05, 327–338.
- NAN, L., SHARF, A., XIE, K., WONG, T.-T., DEUSSEN, O., COHEN-OR, D., AND CHEN, B. 2011. Conjoining gestalt rules for abstraction of architectural drawings. *ACM Trans. Graph.* 30, 6 (Dec.), 185:1–185:10.
- ORZAN, A., BOUSSEAU, A., BARLA, P., AND THOLLOT, J. 2007. Structure-preserving manipulation of photographs. In *Proceedings of the 5th international symposium on Non-photorealistic animation and rendering*, NPAR '07, 103–110.
- PAULY, M., MITRA, N. J., WALLNER, J., POTTMANN, H., AND GUIBAS, L. J. 2008. Discovering structural regularity in 3d geometry. *ACM Trans. Graph.* 27, 3 (Aug.), 43:1–43:11.
- PEARSON, K. 1901. *On Lines and Planes of Closest Fit to Systems of Points in Space*. University College.
- PERLIN, K., AND HOFFERT, E. M. 1989. Hypertexture. *SIGGRAPH Comput. Graph.* 23, 3 (July), 253–262.
- RASKAR, R., AND COHEN, M. 1999. Image precision silhouette edges. In *Proceedings of the 1999 symposium on Interactive 3D graphics*, I3D '99, 135–140.
- RUBNER, Y., GUIBAS, L., AND TOMASI, C. 1997. The earth movers distance, multi-dimensional scaling, and color-based image retrieval. In *in Proceedings of the ARPA Image Understanding Workshop*, 661–668.
- RUNIONS, A., FUHRER, M., LANE, B., FEDERL, P., ROLLAND-LAGAN, A.-G., AND PRUSINKIEWICZ, P. 2005. Modeling and visualization of leaf venation patterns. *ACM Trans. Graph.* 24, 3 (July), 702–711.
- RUNIONS, A., LANE, B., AND PRUSINKIEWICZ, P. 2007. Modeling trees with a space colonization algorithm. In *Proceedings of the Third Eurographics Conference on Natural Phenomena*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, NPH'07, 63–70.
- SALISBURY, M. P., WONG, M. T., HUGHES, J. F., AND SALESIN, D. H. 1997. Orientable textures for image-based pen-and-ink illustration. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '97, 401–406.
- SCHMIDT, R. 2010. *Part-Based Representation and Editing of 3D Surface Models*. PhD thesis, University of Toronto, Canada.
- SHESH, A., AND CHEN, B. 2008. Efficient and dynamic simplification of line drawings. *Comput. Graph. Forum*, 537–545.
- SIDIROPOULOS, G., AND VASILAKOS, A. 2006. Survey: Ultra-real or symbolic visualization? the case of the city through time. *Comput. Graph.*, 299–310.
- STROTHOTTE, T., AND SCHLECHTWEG, S. 2002. *Non-photorealistic computer graphics: modeling, rendering, and animation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- THEOBALT, C., RÖSSL, C., DE AGUIAR, E., AND SEIDEL, H.-P. 2007. Animation collage. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation*, SCA '07, 271–280.
- VIDIMČE, K., WANG, S.-P., RAGAN-KELLEY, J., AND MATUSIK, W. 2013. Openfab: A programmable pipeline for multi-material fabrication. *ACM Trans. Graph.* 32, 4 (July), 136:1–136:12.
- WANG, J., XU, Y., SHUM, H.-Y., AND COHEN, M. F. 2004. Video tooning. *ACM Trans. Graph.* 23, 3 (Aug.), 574–583.
- WINKENBACH, G., AND SALESIN, D. H. 1994. Computer-generated pen-and-ink illustration. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, SIGGRAPH '94, 91–100.
- WINNEMÖLLER, H., OLSEN, S. C., AND GOOCH, B. 2006. Real-time video abstraction. *ACM Trans. Graph.* 25, 3 (July), 1221–1226.