

Kinetic Visualization

A Technique for Illustrating 3D Shape and Structure

Eric B. Lum

Aleksander Stompel

Kwan-Liu Ma

Department of Computer Science
University of California at Davis

Abstract

Motion provides strong visual cues for the perception of shape and depth, as demonstrated by cognitive scientists and visual artists. This paper presents a novel visualization technique – kinetic visualization – using particle systems to add supplemental motion cues which can aid in the perception of shape and spatial relationships of static objects. Based on a set of rules following perceptual and physical principles, particles flowing over the surface of an object not only bring out, but also attract attention to essential shape information of the object that might not be readily visible with conventional rendering that uses lighting and view changes. Replacing still images with animations in this fashion, we demonstrate with both surface and volumetric models in the accompanying *mpeg* file that in many cases the resulting visualizations effectively enhance the perception of three-dimensional shape and structure. The results of a preliminary user study that we have conducted also show clear evidence that the supplemental motion cues helped.

Keywords: animation, visual perception, particle systems, scientific visualization, volume rendering

Note to Reviewers

Because of the nature of the techniques presented, while reading the paper the reviewers are advised to watch the accompanying videos in order to follow the exposition. A copy of the video can also be downloaded from:

<http://www.cs.ucdavis.edu/~ma/kinvis.mpg>

1 Introduction

Time varying sequences of images are widely used in visualization as a means to provide an extra dimension of information for perception to occur. This animation might be as simple as the changing of camera or object positions or can include animations resulting from time varying changes in the data itself. However, using motion that is independent of changes in viewing direction for conveying the shape information of *static objects* has been a rather unexplored area. In this paper, we describe a new visualization technique, which

we call *kinetic visualization*, creating animations that illustrate the shape of a static object in a perceptually intuitive manner.

This work is motivated by the observation that the flow of fast moving water over a rock, a dynamic flame from an open fire, or even a flock of birds exhibit motion that gives the perception of shape. Our technique is built on the inspirations we received from kinetic art [17], the studies done in cognitive science, specifically on structure-from-motion perception [1, 13], the ideas of particle systems [10], and the work of Interrante [8] on using texture to convey the shape of overlapping transparent surfaces. It is unique because we are able to apply motion as a supplemental cue to enhance perception of shape and structure, and because the motion is created not only according to the characteristics of the data but also using a set of rules based loosely on physics and biology.

1.1 Visual Cues

With traditional rendering methods lighting provides valuable spatial cues that assist spatial perception. Considering Lambertian surfaces, the illumination equation [4] accounting for both ambient light and diffuse light is

$$I = I_a k_a + I_p k_d (N \cdot L)$$

where I_p is the pointlight source's intensity. The dot product in this equation has the effect of transforming the three-dimensional surface normal into a one dimension light intensity seen by the viewer. This loss of dimensionality results in ambiguity in surface orientation since multiple normal orientations can map to the same light intensity. For example, under some conditions concave and convex shapes can have similar appearances, despite the surface orientations being entirely different.

The normal direction can be made less ambiguous with the addition of specular lighting. Phong illumination adds a $V \cdot R$ term, where V is the view direction and R is the reflected light direction, which has the effect of indicating shape using not only the normal vector, but also the derived reflectance vector. This vector is once again transformed into a one-dimensional quantity with a dot-product operation. By rotating an object, the viewer can better resolve the shape of an

object since rotation varies the direction of the normals with respect to the light and viewer. This helps disambiguate the loss of dimensionality from the dot-product operation. The changing of viewpoint can also aid in spatial perception by exposing different sets of silhouette edges on the object.

The focus of our work is to use motion from moving particles to even further disambiguate surface orientation. The technique introduced in this paper is not meant to be a replacement for traditional rendering techniques that use lighting and viewpoint changes to indicate shape, rather it can augment those methods for more intuitive and effective visualization.

We have applied kinetic visualization to two different types of static data. One includes surface models represented as polygonal meshes, in which case particle motion is influenced by surface normal, principal curvature direction, and curvature magnitude. The other type of static data is regularly sampled scalar volumetric data where scalar value, gradient magnitude, gradient direction, principal curvature direction, and opacity are used in the calculation of particles motion.

2 Related Work

The perception of depth through motion, called "structure-from-motion", has long been studied in psychology [14]. Treue et al. [13] demonstrate that the movement of points on an object can give the perception of shape, using as stimulus a rotating cylinder with a random dot pattern on the surface. Their work shows a "building up" time is required for mental generation of a surface representation from the integration of point velocities. They also find that subjects were able to perform various tasks with peak performance when points had lifetimes of at least 125 milliseconds (ms) and that with lifetimes of less than 60 ms shape perception from motion did not occur.

Further work by Andersen and Bradley [1] demonstrates that structure-from-motion perception requires either a large number of dots, or fewer dots that appear in varying positions over time [1]. Their work also suggests that the middle temporal area (MT) of the brain is essential for the structure-from-motion perception.

Wanger, Ferwerda, and Greenberg [15] explored visual cues by conducting three psychophysical experiments in which the accuracy of interactive spatial manipulation performed by subjects was measured. Their study shows that different visual cues facilitate different tasks. Motion is found to have a substantial positive effect on performance accuracy of orienting tasks in which spatial location is less important but relative alignment information is needed.

Kinetic art incorporates real or apparent movement in a painting or sculpture. Kinetic artists often use various means to de-emphasize form and color in favor of movement. From studies in neurology, it is evident that an entire area of the brain is devoted to processing motion information. Zeki [17]

proposes that the same area is essential for appreciating kinetic art. In neurological terms, when activity in one area of the brain increases, activities in other areas would decrease. We need to take this into account when emphasizing motion.

Motion blur [3, 9] captures the effect of motion in stills and is also widely used in producing realistic animations. The work presented in this paper deals with the inverse problem, where instead of using a static image to represent a dynamic phenomenon, dynamic animations are generated for the visualization of static data.

Motion without movement [5] assigns perceptual motion to objects that remain in fixed positions by using oriented filtering. This technique can generate a continuous display of instantaneous motion. Line integral convolution [2], based on the same principle, low-pass filters a noise function along a vector field direction to create visualization of direction information.

Our work applies particle systems, which have been used to model a set of objects over time using a set of rules [10]. They have been applied to the modeling of a wide variety of phenomena, including smoke, fire, and trees, using a set of either deterministic or stochastic rules of motion [4]. These rules can be based on physics, for example gravity, or even biology, as is the case with flocking behaviors.

The shape, density, transparency and size of particles can have an impact on the visual appearance and resulting perception cues. Interrante [8] has done a comprehensive study on using opaque stroke texture to improve the perception of shape and depth information of overlapping transparent surfaces. Our work considers particle shape to some extent, although the focus of our work is particle motion, rather than shape.

Using particles as a representation of shape is also related to point based rendering. Point based rendering algorithms typically use reconstruction filters that disguise the appearance of the point representation [18]. In some ways our work can be thought of as a variation of point based rendering where the points move over time and are intentionally made visible.

In the volumetric case, our work is analogous to splatting [16] with a limited budget of splats. The location and size of each particle are not specified to represent the entire volume, but rather are positioned such that their location and movement create a dynamic representation of the static volume. In this way, our technique allows for the volume visualization of extremely large volumetric datasets with a limited rendering budget.

3 Motion Strategies

In this section we discuss the set of rules we apply to generate geometrically meaningful motion. The overall goal is to create rules resulting in particles that indicate shape by smoothly flowing over an object, with locally consistent directions, and a density distribution that does not let particles

"clump" together in regions of little interest. Many of the rules imposed on the particles are loosely based on biology or physics. It is our belief that these types of rules are desirable since they are similar to the types of stimulus the human visual system has been adapted to process.

3.1 Motion Along the Surfaces

Since we would like to better illustrate an object's shape, rules are imposed to constrain the motion of particles to be along a surface. The motion of particles along an object's surface can help improve shape perception, over time presenting the viewer with a set of vectors (trajectories) that run parallel to a surface. In the case of viewing a mesh, this rule is accomplished by simply constraining the particles to lie on the mesh.

In the case of volumetric data, the rules are applied to restrict motion along directions of high gradient. Movement in the direction of a particle is reduced along the gradient direction depending on gradient magnitude as described in the following equation:

$$\vec{v}_{n+1} = \vec{v}_n - (\vec{g} \cdot \vec{v}_n) \vec{g}$$

where \vec{v}_{n+1} is the new direction, and \vec{g} is the gradient direction. This results in particles that move along the surface that is of interest to us and don't leave that surface. The expected numerical instabilities which could be expected with such updates were unnoticeable for even the longest living particles, and thus ignored by us. After every iteration, velocities are normalized to have constant magnitude, or speed. A particle with reduced speed in projected screen space thus provides cues it is either moving in a direction near parallel to the view direction or is far from the viewer and thus has reduced speed on the screen as a result of perspective. If particle speed was allowed to vary, such depth and orientation cues would be lost.

3.2 Principal Curvature Direction

The principal curvature directions (PCDs) indicate the direction of minimum and maximum curvature on a surface. Interrante [7] describes how line integral convolution along the principal curvature directions can generate brush-like textures that create perceptually intuitive visualizations since the resulting textures "follow the shape" of the object being rendered. Similarly we use principal curvature directions to create particles that "follow the shape" of a surface. Particle velocities are adjusted so the particles flow in a direction that favors the primary principal curvature direction. This can be expressed as:

$$\vec{v}_{n+1} = (\vec{v}_n(1 - sc_m)) + (\vec{v}_c(sc_m))$$

where s is a scale factor for the principal curvature direction specified interactively by scientist, \vec{v}_c is the principal

curvature direction vector, and c_m is the magnitude of the principal curvature direction vector. Note that a curvature direction at any point is the same forward as backwards. When PCD is incorporated with the velocity, its orientation is adjusted so that it is most consistent with the current velocity of the particle. The PCD rules result in particles that smoothly flow over an object, although the particles are not guaranteed to move in the same direction.

3.3 Consistent Orientations

The motion of dots in opposite directions can suppress response to the middle temporal (MT) area of the brain and can give perceptual cues of differences in depth [1]. We therefore use a set of rules that move particles in directions consistent with their neighbors. This is particularly important since the PCD-based rule in the previous section results in particles that can follow a PCD in opposite directions. We use two different types of rules to enforce consistency.

The first method we use to give the particles more consistent orientations is to give the particles flock-like behavior. Flocks exhibit motion that are fluid, with each member still exhibiting individual behavior. Thus flocking can be used to add local uniformity to particle motion while still allowing particles to have motion shaped by outside forces like principal curvature direction. Reynolds [11] presents a method for creating flock-like particle systems using behaviors that include velocity matching, collision avoidance, and flock centering. We have found that adjusting particle velocity towards the direction of flock velocity to be an effective method in yielding more consistent particle velocities. This rule makes each particle attempt to match the velocity of its neighbors. We use a weighted function that computes the average flock velocity for each particle based on neighboring particles velocities and also on the distance of other particles from the primary one. The equation used to adjust velocity of particles to the flock velocity can be expressed as:

$$\vec{v}_{n+1} = ((1 - k)\vec{v}_n) + (k\vec{v}_f)$$

where k is a percentage of the contribution by flock vector to the particle velocity which can be specified interactively by scientist, and \vec{v}_f is the flock vector for this given particle. By not enforcing strict velocity matching, particles can still exhibit motion influenced by other rules, like principal curvature direction, while still adding consistency with respect to their neighbors. Collision avoidance is used to give particles a more uniform distribution and will be discussed in the next section. Flock centering is not used since it is not our intention to have the particles stay together as a coherent unit but rather to create particles exhibiting locally flock-like behavior. A simpler method for giving particles a consistent orientation is to simply define a "preferred" direction the particles must move, which can be expressed by the following equation:

$$\vec{v}_{n+1} = ((1-k)\vec{v}_n) + (k\vec{d})$$

where k is a percentage of the contribution by preferred direction which can be specified interactively by scientist, and \vec{d} is the preferred direction. The result is a flow of particles that move over a surface with an appearance similar to water flowing over an object. One drawback of this approach is that at the extreme ends of an object, where the particles flow from and flow into, the direction of the particles is not consistent; that is, the particles would move in opposite directions either to or from a point on the surface.

3.4 Particle Density

Treue et al. [13] demonstrate that if moving stimulus become too sparse, shape perception is diminished. Consideration must therefore be taken with regard to particle density. Since the number of particles has a direct influence on rendering time, it is desirable to have a set of rules that efficiently uses a limited budget of particles. In addition, rules regarding particle density are necessary since following principle curvature directions can result in particles accumulating in local minima.

By using a set of rules based on magnetic repulsion, more uniform particle densities can be achieved. Particles are modeled as having a magnetic charge of the same sign, and are repelled from their neighbors with a force inversely proportional to the square of the distance of the neighbors. This is similar to the rule Reynolds uses for flock collision avoidance [11]. In order to avoid numeric instability from particles that are too close, the total force is clamped. The adjustment to the position of each particle can be expressed by the following equation:

$$P_{n+1} = P_n c d_t$$

Where P is the position of particle, c is the percentage of the contribution by magnetic repulsion which can be specified interactively by scientist, and d_t is a delta time, which is needed since the contribution is made directly to the position and not the velocity vector. Using this technique results in more uniform particle densities. Use of this rule, however, must be limited, since it can result in particles that move with velocities contrary to principal curvature direction.

Another method for controlling particle density is to use particle lifetimes designed to prune particles from high density regions, and favor them in regions of low density. During each update iteration we calculate the density around each particle and upon it we decide if particle need to be removed or can stay, where we use a threshold value that is about twice as large as the average density for the whole system as the cut point. Then, we reinsert the same amount of particles that was removed where we require each new insertion to fall into a region below the average density value for

the whole system. The function that calculates the density for each particle can also be artificially manipulated to yield higher particle densities in regions of interest. For example, density can be "reported" lower at regions of high curvature, or in front-facing visible regions. In addition, in the volumetric case, density can be "reported" higher in highly transparent regions resulting in a higher density of particle in opaque regions. It is important, however, that particles are not allowed to die too quickly since particles shown for too short a period are ineffective for the perception of structure from motion. Therefore we require each particle to stay alive for certain amount of iterations before it can be removed. This required value can be specified interactively by scientist.

3.5 Particle Color

The color of each particle can be varied to provide additional information. Gooch et al. [6] describe how variation in color tone, or temperature, can be used to indicate shading, reserving variation in color intensity for silhouettes and highlights. Schussman et al. [12] use tone to indicate line direction when visualizing magnetic field lines. Either of these ideas can be incorporated as a particle system rule. The tone of each particle can be varied from cool to warm based on lighting. Particles can also have their color temperature varied depending on direction, with particles having a velocity toward the viewer being rendered in warmer colors than receding particles. Particle tone can also be used to indicate other scalar values, such as curvature magnitude or gradient magnitude for volumetric datasets.

Special consideration must be taken into account with regard to particle color when combined with traditional rendering techniques. For example, if particles are to be drawn on top of a surface, the particles should not be the exact same color as the surface or they will not be visible. In addition it is often necessary to have particle color intensity vary based on shading parameters. This is particularly helpful when the particles are dense, since they can obscure the lighting cues provided by the underlying surface. If particles are lit, a different set of lighting parameters should be used for the particles in order to avoid their blending in with the surface and becoming difficult to see, especially when a particle is in a darker region. For example, if particle and surface are both rendered in extremely dark colors, it can be difficult to see the particles, even if they differ in hue from the surface. In our implementation we allow scientist to vary the particle color based on gradient magnitude, view vector and velocity vector. Each of these adjustments in color can be used alone or combined together. For gradient magnitude based adjustment we increase the hue component value of the color depending on the magnitude of the gradient. For view vector based adjustment we use the Z component of the view vector to shift the hue component of the color in each direction depending on the Z Value. And eventually, for color adjustment based on particle velocity we map each X,Y,Z component of the velocity vector into each primary color: Red, Green, and

Blue. To avoid rapid changes in particles colors, we use previous color value in computing the new color as well. Thus changes are smooth and pleasing to the human eye.

3.6 Particle Shape

The size and shape of each particle can also influence how it is perceived. For example, if particle size is varied based on density such that the gaps between particles are filled, more traditional point based rendering occurs. Since for our work individual particles must be visible for their motion to be perceived, particles are rendered small enough that overlap with neighbors is minimal.

There are a number of ways that particle size can be varied. Particles can be rendered in perspective such that closer particles appear larger than further particles, providing a visual cue of particle position. Particle size can be varied based on local density such that the gaps between particles is uniform, similar to splatting. Finally, particle size can simply be kept constant.

Interrante [8] found stroke length to be critically important in her work using strokes oriented along principle curvature directions. Since in our work direction is indicated by temporal means, the importance of particle shape is reduced. Nevertheless, particles can be given a stroke-like appearance as a temporal anti-aliasing mechanism, or to simply make particle direction more clear. Particles can be rendered as small "comets" with a tail that indicates the direction the particle came from in the previous frames. Particles can also be drawn as curves that indicate the position of the particle over some period of time.

In our implementation we allow the scientist to turn off and on the various rules and tune motion parameters until the desired visualization is achieved. The interactivity of this process allows the quick selection of parameters that are appropriate for emphasizing the regions of interest to the user.

4 Demonstration

We experimentally studied kinetic visualization on a PC with an AMD Athlon 1.4 Ghz processor and Geforce 3 graphics card. With this low-cost system we are able to render thousands to tens of thousands particles at 20 frames per second, depending on the type of particle system rules used. If more particles are used than can be rendered at interactive rates, animation can be generated in an offline batch mode.

To demonstrate kinetic visualization, several animation sequences have been made and included in the VHS tape accompanying this paper. Note that all rendering, including volume rendering, was done in hardware to achieve maximum interactivity. Consequently, the image quality, especially for the volumetric models, is not as good as what a software renderer could achieve. Figure 1 shows the four models used for the demonstration: a PET scan of a mouse brain ($256 \times 256 \times 256$), a CT tooth volume ($256 \times 256 \times 161$),




Figure 1: Models

a distorted ball model (15872 polygons), and a subdivided Venus model (5672 polygons).

The first video sequence shows the use of our technique in the visualization of a mouse brain PET volumetric dataset. The particles help to illustrate one of the function levels while direct volume rendering gives their motion context. The following still image shows the type of shape ambiguity that can exist with traditional rendering techniques. It is difficult to distinguish the concave and convex portions of the model. With the addition of the particles, the shape becomes immediately apparent. It is not the particles by themselves that clarify the shape, rather, it the extra shape cues they provide that work in addition to traditional rendering.

The "rules" portion of the video gives examples of each of the different rules we apply. Notice that with the absence of rules, the random motion of the particles on the Venus model does little to clarify shape. By having the particles follow the first principal curvature direction, the particles clearly "follow the shape" of the model.

The next sequence shows particles moving along the tooth dataset, but with inconsistent orientations. Although the particles seem to have a slight shape clarifying effect, their contrary motions are distracting and make them difficult to follow. With the addition of flocking, the particles still move along the shape of the tooth, but move in a much more locally consistent manner. In the following sequence, particles flow down the Venus model, in a manner similar to water. The downward tendency adds consistency to the motion, yet the particles still show some tendency toward following the first principal curvature direction.

Next, the tooth is shown without density controlling rules.




Figure 2: Selected test images.

As the particles move over time, they tend to accumulate in ridges as a result of following the first principal curvature direction. With the absence of particles in some regions, the shape becomes less clear. With the addition of magnetic repulsion, the distribution of particles becomes much more uniform and the resulting video reveals more shape information.

The next sequence illustrates the effect of changing particle size. When particles are large, they can cover a surface much like spitting, but their motion becomes obscured. When particles are small, they can be difficult to see, and do little to improve perception. The last sequence shows kinetic visualization of the PET data with changing view direction.

5 User Study

To evaluate the effectiveness of kinetic visualization a user study was conducted. Height field data sets were generated with several randomly placed peaks and valleys, selected examples of which are seen in Figure 2. Subjects were asked to identify with a mouse pointer the closest and furthest (minimum and maximum) points on the surface. The height and depth of the tallest and shallowest peaks were at least twice as high or shallow as all others.

Twenty-two subjects were shown 14 different data sets, half of which were randomly rendered as either a static image with traditional Phong shading, or with kinetic visualization. Thus, the combined subject selected a total of 252 minimum and maximum points on surfaces rendered using each method. The results, summarized in Table 1, indicate that users could more accurately select both the minimum and maximum points on the surface with kinetic visualization. Although the scope of this user study was fairly limited, we feel the results are extremely encouraging. In particular since the same motion parameters were constant for all data sets, when we have found kinetic visualization to be most effective when the parameters are fine tuned to a particular data set.

Table 1: User study results.

Task	% Correct w/ KV	% Correct w/o KV
Find Max	65	49
Find Min	57	42
Combined	61	46

6 Conclusions

This paper shows a further step towards making perceptually effective visualizations by adding visually rich motion cues. While more work will be required, our current results are encouraging, demonstrating that it is feasible and desirable to capitalize on motion cues for the purpose of enhancing perception of 3D shape and spatial relationships.

We have shown that kinetic visualization nicely supplements conventional rendering for illustrating both volumetric and surface data models. We have also shown how the moving particles help reveal surface shape and orientation. By utilizing low-cost commodity hardware, the kinetic visualization system we have built is very affordable. The selective rendering based on particle budget ensures the critical interactivity required for kinetic visualization.

For some classes of data, however, some limitations in the effectiveness of our technique can be observed. In cases where the principal curvature direction is not well defined, for example flat or spherical regions, the effectiveness of having particles move along a principal curvature direction is limited. The use of optimization strategies, like that described by Hertzmann and Zorin [Hertzmann2000], could be used to add direction consistency in these regions, but consideration would also need to be given toward avoiding smoothing subtle features of interest.

Another limitation of our technique is it is not appropriate for visualizing time varying phenomena. Since the motion of particles in our work is based on the geometric properties of a data set, the motion can give the perception of motion that is contrary to that which is physically occurring. For example, our technique would not be appropriate for visualizing fluid flow since the motion of particles could give a misleading indication of flow direction. A final limitation of our technique, is that in some cases geometric structure is obvious or known a priori, thus the extra cues provided by our technique are not necessary and can even be distracting. However, scientific visualization, by its very nature, often involves attempting to gain insight into the shape and structure of unknown phenomena.

Despite the limitations listed above, we believe that the effectiveness of kinetic visualization to clarify structure and enhance understanding of ambiguous shapes makes it a useful technique that deserves further study. Further work includes experimenting with new rules for the particle movement to convey even more information. More in-depth user studies on kinetic visualization could provide valuable feedback into which types of rules should be improved. Ad-

ditional future work includes using improved methods for computing principal curvature directions, and accelerating the integrated rendering as much as possible to attain even higher interactivity. It is hoped the power of motion cues will be embraced by more people, helping them effectively perceive/illustrate complex or ambiguous object shape and spatial relationship.

Acknowledgements

This work has been sponsored by the National Science Foundation under contract ACI 9983641 (PECASE Award) and through the Large Scientific and Software Data Set Visualization (LSSDSV) program under contract ACI 9982251. The authors are particularly grateful to all of the people who participated in our user study.

References

- [1] ANDERSEN, R. A., AND BRADLEY, D. C. Perception of three-dimensional structure from motion. *Trends in Cognitive Science* 2, 6 (June 1998), 222–228.
- [2] CABRAL, B., AND LEEDOM, L. Imaging vector fields using line integral convolution. In *SIGGRAPH '93 Conference Proceedings* (August 1993), pp. 263–270.
- [3] COOK, R. L., PORTER, T., AND CARPENTER, L. Distributed ray tracing. In *SIGGRAPH '84 Conference Proceedings* (July 1984), pp. 137–145.
- [4] FOLEY, D. J., VAN DAM, A., FEINER, S. K., AND HUGHES, J. F. *Computer Graphics: Principles and Practice*. Addison Wesley, 1996.
- [5] FREEMAN, W. T., ADELSON, E. H., AND HEEGER, D. J. Motion without movement. In *SIGGRAPH '91 Conference Proceedings* (July 1991), pp. 27–30.
- [6] GOOCH, A., GOOCH, B., SHIRLEY, P., AND COHEN, E. A non-photorealistic lighting model for automatic technical illustration. In *SIGGRAPH '98 Conference Proceedings* (July 1998), pp. 447–452.
- [7] INTERRANTE, V. Illustrating surface shape in volume data via principal direction-driven 3d line integral convolution. In *SIGGRAPH '97 Conference Proceedings* (August 1997), pp. 109–116.
- [8] INTERRANTE, V., FUCHS, H., AND PIZER, S. Conveying the 3d shape of smoothly curving transparent surfaces via texture. *IEEE Transactions on Visualization and Computer Graphics* 3, 2 (April-June 1997), 98–117.
- [9] POTMESIL, M., AND CHAKRAVARTY, I. Modeling motion blur in computer-generated images. In *SIGGRAPH '83 Conference Proceedings* (July 1983), pp. 389–399.
- [10] REEVES, W. T. Particle systems—a technique for modeling a class of fuzzy objects. In *SIGGRAPH '83 Conference Proceedings* (July 1983), pp. 359–376.
- [11] REYNOLDS, C. W. Flocks, herds, and schools: A distributed behavioral model. In *SIGGRAPH '87 Conference Proceedings* (July 1987), pp. 25–34.
- [12] SCHUSSMAN, G., MA, K.-L., SCHISSEL, D., AND EVANS, T. Visualizing diiii-d tokamak magnetic field lines,. In *Proceedings of IEEE Visualization 2000 Conference* (October 2000), pp. 501–504.
- [13] TREUE, S., HUSAIN, M., AND ANDERSEN, R. A. Human perception of structure from motion. *Vision Research* 31 (1991), 59–75.
- [14] WALLACH, H., AND O'CONNELL, D. N. The kinetic depth effect. *Journal of Experimental Psychology* 45 (1953), 205–217.
- [15] WANGER, L. R., FERWERDA, J. A., AND GREENBERG, D. P. Perceiving spatial relationships in computer-generated images. *IEEE Computer Graphics and Applications* 20, 3 (May 1992), 44–58.
- [16] WESTOVER, L. Interactive volume rendering. In *Chapel Hill Workshop on Volume Visualization* (1989), pp. 9–16.
- [17] ZEKI, S. *Inner Vision*. Oxford University Press, 1999.
- [18] ZWICKER, M., PFISTER, H., VAN BAAR, J., AND GROSS, M. Surface splatting. In *SIGGRAPH 2001 Conference Proceedings* (August 2001), pp. 371–378.