

Schelling Points on 3D Surface Meshes

Xiaobai Chen, Abulhair Saparov, Bill Pang, and Thomas Funkhouser
Princeton University

Abstract

This paper investigates “Schelling points” on 3D meshes, feature points selected by people in a pure coordination game due to their salience. To collect data for this investigation, we designed an on-line experiment that asked people to select points on 3D surfaces that they expect will be selected by other people. We then analyzed properties of the selected points, finding that: 1) Schelling point sets are usually highly symmetric, and 2) local curvature properties (e.g., Gauss curvature) are most helpful for identifying obvious Schelling points (tips of protrusions), but 3) global properties (e.g., segment centeredness, proximity to a symmetry axis, etc.) are required to explain more subtle features. Based on these observations, we use regression analysis to combine multiple properties into an analytical model that predicts where Schelling points are likely to be on new meshes. We find that this model benefits from a variety of surface properties, particularly when training data comes from examples in the same object class.

Keywords: 3D shape analysis, feature detection, shape matching

Links:  DL  PDF

1 Introduction

Detection of “salient” feature points on 3D surfaces is a fundamental problem in computer graphics with many applications in shape analysis and related fields, including object recognition [Johnson 2000], shape matching [Zhang et al. 2008], shape-based retrieval [Funkhouser and Shilane 2006], metamorphosis [Alexa 2000], cross-parameterization [Kraevoy and Sheffer 2004], texture mapping [Zhang et al. 2005], deformation transfer [Sumner and Popovic 2004], shape approximation [Lee et al. 2005], viewpoint selection [Lee et al. 2005], symmetry detection [Xu et al. 2009], and part-based segmentation [Katz et al. 2005; Zhou and Huang 2004].

Although many definitions have been proposed for what constitutes “salient” feature points in the computer graphics literature (e.g., maxima of average geodesic distance, maxima of Gauss curvature, maxima of mean curvature differences at increasing scales, etc.), none of them captures the social/psychological essence of salience as defined in the Oxford English Dictionary: “the quality or fact of being more prominent in a person’s awareness or in his memory of past experience” [Simpson 1989]. It is this definition that captures the semantic essence of a stable feature point, and therefore one that we believe is useful for applications in computer graphics.

ACM Reference Format

Chen, X., Saparov, A., Pang, B., Funkhouser, T. 2012. Schelling Points on 3D Surface Meshes. *ACM Trans. Graph.* 31 4, Article 29 (July 2012), 12 pages. DOI = 10.1145/2185520.2185525 <http://doi.acm.org/10.1145/2185520.2185525>.

Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, fax +1 (212) 869-0481, or permissions@acm.org.
© 2012 ACM 0730-0301/2012/08-ART29 \$15.00 DOI 10.1145/2185520.2185525 <http://doi.acm.org/10.1145/2185520.2185525>




Figure 1: Schelling points (red). Positions on a surface selected consistently by many people when trying to match each other without communication.

The goal of this paper is to develop a model of salience on 3D surface meshes based on this social/psychological definition. To achieve this goal, we leverage the concept of focal points as introduced by Schelling in his seminal paper on pure coordination games [Schelling 1960]. In game theory, “a focal point (also called a Schelling point) is a solution that people will tend to use in the absence of communication, because it seems natural, special or relevant to them” [Parker 2011]. To discover focal points, Schelling performed user studies in which he asked people to make selections that they expect will match other people’s selections. Since the people were not allowed to communicate in any way, they usually picked the most conspicuous point, often using semantic information not provided in the input. For example, if asked to select a time and place in New York City to meet someone without any prior communication, people tended to choose Grand Central Terminal at noon. This choice is clearly the result of prior semantic knowledge that makes that choice more prominent in a person’s awareness or in his memory of past experience (i.e., social/psychological salience).

Inspired by Schelling’s work, we designed a user study that asks people to select points on 3D surface meshes that they expect will be selected by other people and used the acquired data to build a model of mesh saliency. Our primary research contribution is the analysis of the collected data. We find that: 1) Schelling point sets are usually highly symmetric, and 2) local curvature properties (e.g., absolute value of minimum curvature) are most helpful for predicting obvious Schelling point features, but 3) global properties (e.g., symmetry and segment centeredness) are required to explain other features. Our secondary research contributions are the methods used for acquiring, analyzing, and predicting Schelling points on 3D meshes. They include: 1) the design of a user study, 2) the analysis of consistency and properties of Schelling points, 3) a learned model for predicting the distribution of Schelling points, and 4) an algorithm that uses the learned model to extract Schelling points on new meshes.

2 Related Work

Detection and matching of salient feature points are classical problems in computer vision, geometric modeling, computer-aided design, computer graphics, and several other fields [van Kaick et al. 2010]. In this section, we cover the most related work, focusing on methods designed for saliency estimation at a point.

Feature Point Detection: There has been lots of recent work on automatically extracting feature points on 3D meshes. Example methods include local maxima of average geodesic distance (AGD) to other points on the surface [Zhou and Huang 2004; Zhang et al. 2005; Zhang et al. 2008], local maxima of differences of Gaussians at multiple scales [Castellani et al. 2008], local maxima of Gaussian curvature [Lipman and Funkhouser 2009], properties of the Heat Kernel Signature [Sun et al. 2009], scale-space analysis of mean curvature flow [Zaharescu et al. 2009], other multi-scale analysis [Li and Guskov 2005; Li and Guskov 2007; Novotni et al. 2005; Schlattmann et al. 2008; Sonthi et al. 1997], points on the convex hull after MDS embedding [Katz et al. 2005], leaf nodes of a curve skeleton extraction [Hisada et al. 2002], points with unlikely local shape descriptors [Chua and Jarvis 1996; Johnson 2000], and points with local shape descriptors distinctive of an object class [Shilane and Funkhouser 2007]. Most methods are based on differential properties of the surface (e.g., [Castellani et al. 2008]), while others are based on global properties [Zhang et al. 2008] and/or shape descriptor statistics (e.g., [Johnson 2000]). Previous work has analyzed which of these feature points is most stable under various models of perturbation [Bronstein et al. 2010], but none as studied how they relate to semantic salience in the Schelling sense.

Saliency Estimation: In related work, methods have been proposed to define continuous measures of “saliency” across a surface for mesh processing applications. Motivated by perceptual criteria [Hoffman and Singh 1997], [Lee et al. 2005] used a center-surround filter of curvature across multiple scales to select salient regions for mesh simplification and viewpoint selection. [Gal and Cohen-Or 2006] computed the saliency of a region based on its size relative to the whole object, its curvature, the variance of curvature, and the number of curvature changes within the region. [Shilane and Funkhouser 2007] defined the “distinction” of a surface region based on its similarity to objects within the same class and difference to objects in other classes. This paper studies how these types of saliency measure relate to human-selected focal points and provides a new saliency measure learned from examples.

Saliency Evaluation and Comparison: [Kim et al. 2010] studied how eye fixations relate to mesh saliency [Lee et al. 2005], and other studies have evaluated and compared feature point detection algorithms for images in computer vision [Ko and Nam 2006; Moreels and Perona 2007; Privitera and Stark 2000; Schmid et al. 2000; Stark and Schiele 2007; Zuliani et al. 2004]. For example, [Schmid et al. 2000] and [Sebe and Lew 2003] compare interest point detectors on the basis of stability, repeatability, and information content. [Privitera and Stark 2000] compare region of interest algorithms based on how well they match eye fixations measured with an eye tracker. [Huang et al. 2009] compare algorithms with respect to points collected in a game (Photoshoot) that asks people to select points on images that they expect will match a partner’s selection, much like the ESP game proposed by [Von Ahn and Dabish 2008]. This work is similar to ours in that it also asks people to match point selections. However, the methodology of that study can produce bias in selected points due to training effects of immediate feedback (it is not a pure coordination game), it considers only 2D points in natural images, it does not provide any analysis of how image properties correlate with selected points, it does not suggest new properties correlated with the selected points (e.g., symmetry,

segmentations, etc.), and it does not provide a predictor of points learned from the collected examples.

Perceptual Psychology: There have been many studies in perceptual psychology to understand how people assign importance to regions of images. Most have been based on visual attention [Koch and Ullman 1985; Milanese et al. 1994; Tsotsos et al. 1995; Itti et al. 1998; Rosenholtz 1999; Santella and DeCarlo 2004]). For example, [Koch and Ullman 1985] proposed a model that salient points would be ones that are different from their surroundings. Other perceptual psychology studies have studied which points are most important for approximation of a contour. For example, [Attneave 1954] showed 80 subjects a series of 16 shapes drawn as 2D contours and then asked them to select “a pattern of 10 dots which would resemble the shape as closely as possible.” His main finding was that most people select points where the “contour is most different from a straight line” – i.e., where the curvature has large magnitude. While these studies are related to ours, we aim to discover which points have semantic salience in 3D meshes.

Studying how people do X: There have very recently been studies in computer graphics aimed at understanding how people perform tasks of interest in computer graphics. For example, [Cole et al. 2008] studied where artists draw lines when make line drawings, and [Chen et al. 2009] analyzed how people decompose surfaces into parts. In both cases, machine learning algorithms were then developed to produce predictions for new surfaces based on training a classifier on examples collected from humans [Cole et al. 2008; Kalogerakis et al. 2010]. Our study follows this line of research, focusing on feature point detection. It is novel not only because it considers a new question: “where do people select feature points?” but also because it asks the question in a different way: “where do other people select feature points?” and because it includes new algorithms for analyzing and extracting feature points based on the collected data.

3 Approach

The goal of this paper is to develop a model of semantic salience for 3D surfaces. Our general approach is to gather a large collection of feature points from people and then to study what geometric properties distinguish them from others.

Although this general idea may sound straight-forward, it is surprisingly difficult to design a study to collect useful data on salience. We did not want to simply ask people to “please click on important points” because different people might have different ideas of what it means to be important (e.g., functional, structural, social, visual, etc.), and it would be difficult to ask the question in a way that reveals people’s intentions without leading them to an answer.

During a pilot study, we tried an approach based on [Attneave 1954], where we asked people to select points on a 3D surface from which another person could recognize the object class by just viewing those points. This approach was a resounding failure. We found that most people selected points only on a 2D silhouette curve of the surface as seen from a single canonical view (e.g., the outline of a fish as seen from the side), which does not seem to match a notion of semantic salience that is useful for 3D applications.

Ultimately, we arrived at an approach based on Schelling’s focal points: we ask people to select points they think will be selected by others. The concept of focal points was introduced in pure coordination game theory in the 1960’s [Schelling 1960] (page 57) – Schelling found that people asked to make selections that match other people’s selections amongst seemingly equivalent distinct options (segregated Nash equilibria) often make the same choices without any communication or feedback. For example, if people

were asked to choose “heads” or “tails” with no other information besides the goal of matching as many other people as possible, 86% chose heads [Schelling 1960]. If asked to match other people’s selection on a map, most people agreed on just a few points (e.g., prominent intersections). It was conjectured that the commonly selected focal points (later called Schelling points) arise from a strategy or outcome with properties of “prominence or conspicuousness” [Schelling 1960].

A decade later, Lewis used Schelling’s ideas to introduce the term “salience,” which he defined, first, as the property of an outcome of “standing out from the rest by its uniqueness in some conspicuous respect” and, second, as “being unique in some way everyone will notice, expect the others to notice” [Lewis 1969] (page 35). He used formal game models to characterize coordination strategies, hypothesizing that people use “common knowledge” to select amongst distinct Nash equilibria. It is this “common knowledge” about feature points on 3D surfaces that we aim to capture and leverage in our work.

Inspired by the ideas of Schelling and Lewis, we have designed a method to study semantic salience on 3D meshes. Specifically, we first acquire a large number of Schelling points by asking people to select points on 3D surface meshes that they expect will be selected by other people (Section 4). We then analyze properties of the collected point sets, asking questions like: “how consistent are Schelling points selected on different meshes within the same object class?” and “how are the locations of Schelling points associated with geometric properties of the surface?” (Section 5). Next, we train a regression model that predicts the likelihood that a point on a surface is a Schelling point based on its surface properties (Section 6). Finally, we provide an algorithm to predict a set of Schelling points for new meshes (Section 7).

4 Study Design

The first and most difficult issue faced in our investigation is how to design a study to acquire Schelling points from many people for many types of 3D surfaces. We would like to collect enough data to analyze how Schelling points relate to surface properties across a wide variety of meshes and to train a predictive model that can be used to estimate Schelling points for new 3D meshes. Of course, this is difficult because it requires recruiting and supervising many (possibly hundreds or thousands of) human subjects in a user study.

To address this issue, we performed our study on-line. Following the approach of [Chen et al. 2009] and [Cole et al. 2009], we recruited subjects for our study through Amazon’s Mechanical Turk (AMT) [Amazon 2009], an on-line platform that matches people willing to work with paying tasks. Alternatively, we could have designed an on-line game to acquire input (as in [Huang et al. 2009; Von Ahn and Dabbish 2008]), but that approach would have required attracting a player population, which is beyond the scope of this paper. Since tasks on the AMT are typically short in duration (a minute or two), inexpensive (around 10 cents), and accessible on-line (in a web page), it is well-suited for studies like ours that require lots of people to do simple, menial tasks (e.g., clicking points on a surface).

The challenge with any on-line study is to design a protocol that acquires useful information from a diverse population of subjects. Unlike a laboratory study, where a handful of screened subjects are trained, monitored, and employed for several hours, we have access to a much larger number of people, but less control over subject selection and less trust that every individual is motivated to do a good job. As such, our challenge is to design a study that motivates people to work responsibly and incorporates unbiased mechanisms to discard data from those that don’t.

We designed an easy-to-learn task in which subjects that provide “better” input get paid more, and those that do not do a “good” job get paid nothing. Specifically, we presented each user with a 3D mesh shown in an interactive viewer with a crystal ball camera control and a simple method for clicking on points with the mouse (one key for adding a point at the selected position on the mesh and another key for removing a previously added point). Then, we provided the following instructions: “select points on the surface of a 3D object likely to be selected by other people. We will ask many people to do the same task and see how your selection matches others.” Our reward structure ranked people according to a scoring function that provides positive credit for each point also selected by at least 25% of other people, negative credit for other points, and zero credit if less than ten points were selected. Based on the ranking, we paid the top scoring 30% of people X , the next 60% $X/2$, and the bottom 10% were not paid. This incentive structure was chosen based on the results of a pilot study, which suggested that it is useful to provide motivation for people to select more than just the very obvious Schelling points (e.g., the ends of limbs on animals) and that it is better to acquire too many points from people than too few (if a surface has fewer Schelling points, then “extra” points will be distributed somewhat randomly).

Data Filtering: Even though our pay structure encourages people to do a good job, we employ three filters to discard point sets we expect to provide bad data: i.e., when the user: 1) clicks too few points (less than ten), 2) maintains approximately the same camera viewpoint for the entire interactive session (the cumulative camera rotation is less than 36 degrees), and 3) clicks too hastily (the average time per click is less than one second). These filters were chosen on the basis of a pilot study, which showed empirically that a significant fraction of people provided careful data and that these simple filters were effective at conservatively discarding the careless data while retaining much of the good data. In this regard, we designed our filters to favor false-negatives (discarding good data) over false-positives (accepting careless data), since it is easy to collect data on the AMT.

Mesh Selection: In order to cover a wide variety of object categories, and to leverage data collected in previous studies, we chose to collect Schelling points for the 3D meshes from the Watertight Track of the 2007 SHREC Shape-based Retrieval Contest. This data set contains 400 meshes spread evenly among 20 object categories (human, cup, glasses, airplane, ant, chair, octopus, table, teddy, hand, plier, fish, bird, spring, armadillo, buste, mechanical part, bearing, vase, and four-legged animal), of which we use all but the “spring.” It forms an interesting data set for our study because it contains categories with different articulated poses (ant, octopus, bird, teddy, pliers, glasses), different local parts (human, hand, four-leg animals), and different global structures (fish, table, cup, mechanical part, bearing, chair, buste, and vase). Also, the data set is highly tessellated (an average of 10,223 vertices per mesh), and so we can limit selected points to vertices of the mesh. Finally, since it has been used for previous studies of mesh matching [Giorgi et al. 2007] and segmentation [Chen et al. 2009], it provides an opportunity to investigate how Schelling points correspond with other types of data.

Protocol Implementation: To reduce bias in the acquired data, we implemented a scheduling program that records the AMT identifier and IP address of every subject doing our task. That program ensures that no single AMT identifier or IP address can work on the same mesh twice, that meshes are distributed in randomized order, and that data is provided by approximately the same number of subjects for every mesh. The interactive Java Applet presented to the user for viewing meshes and selecting points is initialized with a random camera direction and two virtual lights, all that point at the object centroid from a distance relative to the object bounding box

size. As the user rotates the meshes, the lights move with the camera, staying along the equator at 90 degree angles to the camera. The Applet records the position, time stamp, and camera parameters for every point added or deleted from the data set, and sends the data back to our server. No feedback or payments were made to any subject until all data was collected to avoid training effects.

Data Collection: Using this protocol, we used the AMT to acquire 24,124 point sets from 1,696 unique AMT accounts. Of this raw data, 9,965 point sets (44%) from 1,060 unique AMT accounts passed all three data filters – 16% had too few points, 14% were entered with too little camera motion, and 28% were clicked too hastily. The 9,965 point sets in our final data set contain 201,304 points in total, an average of approximately 10,000 points per object category, 500 points per mesh, and 20 points per point set. Each of the 380 meshes was represented by data from 25 people, on average, and every mesh had at least 23 point sets.

Schelling Point Extraction: As a post-process, we extract a discrete set of Schelling Points from the collected data – i.e., the ones upon which many people agreed. Of course, since the data is discrete (on vertices of a mesh), and there is spatial noise σ in the point selection process, some aggregation is required to identify commonly selected points. To address this issue, we first construct a function on every mesh M that indicates the number of times every vertex V was selected by different people. Then, to reduce spatial noise, we blur that function geodesically with a Gaussian filter with maximal value 1 and $\sigma = 0.01R$, where R represents the radius of the mesh ($R = \sqrt{\text{SurfaceArea}(M)}$). The result is a smooth *Schelling distribution function*, $S_D(V)$, roughly estimating the probability that each vertex V will be selected by a person in our study (shown in red in Figures 1 and 2).

Then, we extract a discrete set of *Schelling points*, S , from local maxima of $S_D(V)$ and build an indicator function, $S_P(V)$, that tells whether vertex V is a Schelling point ($S_P(V)$ is one at Schelling points and zero otherwise). To form this set, we select every vertex V with $S_D(V)$ greater than 12.5% and spatial separation by more than $0.02R$ from any other vertex V_k with higher $S_D(V_k)$. These choices include at least the vertices that were selected with spatial error less than σ by at least 25% of the people in our study (the positive payment threshold in our study). Some examples point sets are shown in blue in Figure 2.




Figure 2: *Schelling points.* The saturation of red depicts the estimated fraction of people selecting a mesh vertex in our study. Extracted Schelling points are shown in blue.

5 Analysis of Schelling Points

This large data set provides an opportunity to investigate a number of questions of potential interest in computer graphics and perceptual psychology, including “How consistently do people select points on the same mesh?,” “How symmetric are the selected point sets?,” “How are the selected points distributed on the surface?,” and “What geometric properties of a 3D surface are prominent at selected points?.” This section takes steps towards addressing these questions.

A. How consistently do people select points on the same mesh? The first and most basic question is whether people select points consistently in our study.

To address this question, we compute the geodesic distances between points selected by different people on the same mesh. For the k -th point $P_{i,k}^M$ in every point set P_i^M selected on every mesh M , we compute the geodesic distance $d_M(P_{i,k}^M, P_j^M)$ from that point to the closest point in every other point set P_j^M collected on the same mesh, normalizing for scale by dividing by square root of the surface area. We then analyze the consistency of the point sets by plotting cumulative distributions of $d_M(P_{i,k}^M, P_j^M)$ indicating the fraction of point sets that are consistent with each selected point for a range of distance thresholds.

For example, Figure 3a shows cumulative distributions where the horizontal axis represents normalized distance thresholds ($d_M(P_{i,k}^M, P_j^M)$) and the vertical axis represents the fraction of points within that threshold of point sets selected by different people on the same mesh. The thick black curve represents the overall average of all points in the collected data set, aggregated first over all points within the same mesh, then over all meshes within the same object class, and finally over all classes in the data set to avoid over-weighting point sets or meshes with large numbers of points. The thin solid colored curves depict averages for different subsets of the points based on the order in which a user selected a point within his/her interactive session (e.g., the red curve shows distances from the first point selected in each session, the purple curve shows the same for the second and third point, the blue curve for points 4-7, etc. These curves tell a more detailed story of how consistency relates to the order that points were selected.

Looking at the curves in Figure 3a, we can readily make two observations. First, different people tend to pick points fairly consistently with one another on the same mesh. Using a normalized distance of $0.05R$ (approximately 3 inches on a human body) as a threshold to classify whether a point $P_{i,k}^M$ is “consistent” with a point set P_j^M , we find that 48.5% of the selected points are “consistent” with other point sets. This level of consistency is far greater than 12.8% that would be observed if points were selected randomly. Second, we observe that points selected earlier in an interactive session tend to be more consistent than ones selected later (there are higher densities near $d_M(p_i, P_j) = 0$ in the curves representing points selected earlier), which suggests that people choose the most “salient” points first, and others later. These results, combined with the visualizations in Figure 2, suggest that people selected points in a non-random, consistent way in our study.

B. How consistently do people select points on different meshes of the same object category? A second question of potential interest is whether people select semantically equivalent Schelling points on different meshes within the same object category. For example, if they select a particular point (e.g., the knee) when the object is in one pose (e.g., a running dog), do they also select that point when the same object is in a different pose (e.g., a




Figure 3: Consistency of points collected in our study. The plots show cumulative distributions of normalized distance from each point to the closest point in another point set. Different colored curves depict averages for different subsets points. The thick black curve represents the overall average of all points. The thin colored ones separate points based on the order in which they were selected during an interactive session: the consistency of the first point selected is shown in the top red curve, points 2-3 are shown in purple, points 4-7 in blue, etc. Note that the average normalized edge length for these meshes is 0.013, and the normalized distance threshold used to call a point “consistent” in our analysis is 0.05.

sitting dog)? And, do they select a semantically equivalent point on a different object within the same general category (e.g., all four-legged animals)?

To address these questions, we produced a semantic mapping between all pairs of the twenty meshes within each of the twenty object categories and used those mappings to compute cumulative distributions of normalized distances between selected points. To produce the mappings, an expert user selected a set of landmark points for every mesh that is semantically consistent with every other mesh in the same category. For example, in the four legged animal category, every mesh was marked with landmark points representing the “tip of the nose,” “tip of right ear,” “middle of the back,” etc. These landmark points provided a coarse point correspondence from which a dense one-way inter-surface mapping $A \rightarrow B$ was established from vertices of A to vertices of B using a simple procedure based on similarities of pairwise geodesic distances (following the strategy outlined in [Bronstein et al. 2006], except maintaining the explicit correspondences between landmark points).

Using these inter-surface mappings, we can study the consistency of points selected by people on different meshes within the same object class. Specifically, for each point $P_{i,k}^M$ selected by a person on mesh M , we use the inter-surface mapping $M \rightarrow \tilde{M}$ to transfer it to the domain of every other mesh \tilde{M} in the same class to form $P_{i,k}^{\tilde{M}}$. Then, for every point set $P_j^{\tilde{M}}$ collected on mesh \tilde{M} , we compute $d_M(P_{i,k}^{\tilde{M}}, P_j^{\tilde{M}})$, the normalized geodesic distance from the point $P_{i,k}^{\tilde{M}}$ to the closest point in $P_j^{\tilde{M}}$, and add it to the cumulative distributions using the same procedure as described in the previous subsection. The results are shown in Figures 3b.

Interestingly, we find that the consistency of points selected by different people on different meshes in the same class is almost, but not quite, as high as the consistency of points selected by different people on the same mesh. Overall, 39.4% of the points selected on one mesh are “consistent” with point sets selected on different meshes of the same object class. Since the local geometry of meshes within the same class are sometimes very different, this result suggests that the criteria people used for selecting points is not based only on absolute geometric properties (e.g., Gauss curvature), but rather on relative properties (e.g., extrema of Gauss curvature) or on semantic features that are consistent across different instances within the same class.

C. How symmetric are the selected point sets? A third question of interest regards symmetry: if a person selects a point on a symmetric object, how often do they also select its symmetric correspondence(s)? For example, if a person selects the right ear on a head, how often do they also select the left ear?

To investigate this question, we manually produced a symmetry mapping from every mesh onto itself to establish dense symmetric point correspondences for each of the 319 meshes with intrinsic reflective symmetry. This was done by clicking on symmetric pairs of landmark points (e.g., right eye to left eye, left elbow to right elbow, etc.) and then interpolating those pairs to form a dense correspondence over the entire surface using an algorithm based on [Bronstein et al. 2006].

Given the symmetric mapping, we analyze the consistency of every point set with its symmetric correspondence using the methods described in the previous subsections – i.e., we build cumulative distributions of the normalized geodesic distance from every point to the closest point in the same point set after the symmetric mapping has been applied (Figure 3c). Our results reveal that point sets selected by people are highly symmetric – 76.0% of the points are “consistent” with the symmetric mapping. More specifically, comparing to the histograms in Figure 3, we find that point sets selected by people are far more consistent with their symmetric mapping than they are with point sets selected by different people. As such, we conclude that symmetry is an important cue for selecting “salient” points in our study.

D. How are the selected points distributed on a surface? A fourth question of interest is to characterize the spatial distribution of points selected by people. To address this question, we show the histogram of distances from each point to other points in the same set for meshes of different object class in Figure 4. The plot shows a separate curve for each of the 19 object classes (dotted lines), along with the overall average (thick red line).

Three interesting observations can be made from these histograms. First, people tend to spread points fairly evenly on a mesh – i.e., there are few points that are either very close to or very far from other points. Second, the spacing for different object classes is different. For example, points on busts (statue heads) tend to be closer to one another (dotted light blue curve on left), while points on octopi are more widely spaced (light blue curve on right). Third, the

spacing of points appears to be dictated more by the size of the mesh shown to the person than the physical size of the object in the real world. For example, the spacing of points on a pair of eye-glasses is the same as for a human body, and also the same as for an airplane. Similarly, the spacing of points on a human hand is different if shown alone or as part of an entire body. From these observations, we hypothesize that people think of the virtual objects at the scale they are shown and select the largest features available on the screen regardless of their size in the real world.




Figure 4: Histogram of spacings between points collected for different mesh categories. Distances are listed as fractions the square root of the mesh area. Red line shows average.

F. What geometric properties distinguish Schelling points?

Finally, it is interesting to ask whether it is possible to characterize common geometric properties of Schelling points. For example, are Schelling points commonly at extrema of Gauss curvature? Can Schelling points be predicted by analysis of average geodesic distance or other geometric properties commonly used for feature extraction in computer graphics? And, are there other geometric properties of meshes that are associated with Schelling points, but have not been used for point feature extraction before?

To investigate these questions, we compute a number of geometric properties for every vertex V of every mesh M in the data set, and then analyze how these properties explain the placement of Schelling points based on information theory statistics. The set of properties considered includes:

- **Curvatures:** the Gauss, mean, minimum, and maximum curvatures at V , as computed in [Rusinkiewicz 2004].
- **Mesh Saliency:** the mesh saliency at V (at scales 0.1, 0.3, 0.5, and 0.7), as computed with code provided by [Lee et al. 2005].
- **Geodesic Distance:** the average (AGD), median (MGD), standard deviation (SDGD), tenth percentile (10GD), ninetieth percentile (90GD), and maximum of geodesic distance (GGD) from V to all other vertices in M , as estimated by Dijkstra’s algorithm.
- **Shape Diameter Function (SDF):** the median length of rays traced from V through the interior of M as described in [Shapira et al. 2008] using the implementation provided by [Kalogerakis et al. 2010].
- **Heat Kernel Signature:** the amount of heat diffused from V to itself within time t , for five equally spaced time durations

ranging from very small (HKS1) to very large (HKS101), as computed with the implementation provided by [Sun et al. 2009].

- **Up:** the Z coordinate (ZPosition) and normal direction (ZNormal), assuming it is prescribed by standards in the modeling language.
- **Symmetry:** the Intrinsic Reflective Intrinsic Symmetry Axis function proposed by [Xu et al. 2009], per our implementation.
- **Segment Centeredness:** the centeredness of the point within its part, as computed by first decomposing the mesh into segments automatically with the algorithm in [Shapira et al. 2008] and then computing the distance from V to the closest segmentation boundary divided by the maximal distance to the closest segmentation boundary for all vertices in the segment containing V .

Most of these properties have been used before for characterizing saliency and/or generating point sets on meshes. However, we consider two new ones that are specifically motivated by visual inspection of Schelling point sets collected in our study: symmetry and segment centeredness. We observe that people often select points on symmetry axes and centers of large convex parts in the absence of other more distinguishing features nearby (e.g., the center of the lens of the glasses and the belly button of the teddy bear in Figure 2). Thus, we expect that extrema of these functions (e.g., segment centeredness = 1) will be correlated with locations of selected points. This hypothesis is corroborated by the histogram shown in Figure 5, which plots frequencies of selected points versus segment centeredness (for automatic SDF segmentations, as used throughout this paper: and for manual segmentations provided by the Benchmark of 3D Mesh Segmentation [Chen et al. 2009], which are included only for comparison sake). We find that selected points appear most often at the center of extract parts – i.e., the rightmost bin of the histogram contains 25% of the distribution.




Figure 5: Schelling points are most often at centers of segments (far right of plot).

We also consider functions derived from these basic geometric properties. Specifically, we include the difference between minimum and maximum curvature (CurvDiff), which is smallest at centers of local rotational symmetry. We include the absolute value of Gauss, mean, minimum, and maximum curvatures, which are largest at critical points, peaks/divots, peaks/divots, and ridges/valleys, respectively. We also include blurs at four levels ($\sigma = 0.01R$, $\sigma = 0.02R$, $\sigma = 0.04R$, and $\sigma = 0.08R$) for all

properties except HKS, CurvRing, and Geodesic Distance, which are already very smooth. Finally, for every property, p , and function of that property, $f(p)$, we include a percentile transformation, $\%(p)$, that encodes the percentage of vertices that have a smaller value within the same mesh.

To analyze how Schelling points are associated with these properties, we calculate information theory statistics and coverage plots to estimate how well property distributions explain the Schelling distribution. Results for a representative set of properties and functions on them are shown in Table 1 and Figure 6.

Information Gain: The middle two columns of Table 1 show the Information Gain, $G(f)$, for functions f of several properties p . $G(f)$ is the difference between entropy of the Schelling point indicator function, $S_P(V)$ and its average entropy conditioned on discretized values of the property f – higher values of Information Gain indicate a stronger predictor of the Schelling point locations independent of other properties. From these results, we see that curvature is the strongest single cue for feature point selection – many of the Schelling points are at the tips of protrusions (fingers and toes), which have large positive minimum curvature, and a few are at conspicuous saddle points (e.g., between fingers), which have highly negative Gauss curvature. Other properties commonly used for feature point detection (e.g., Saliency, HKS, and AGD) also provide fairly good predictors.

Random Forest Importance: The rightmost two columns of Table 1 show the importance of f as computed by the Random Forests of [Breiman 2001], which estimate the importance of a feature in combinations with others by building a large number of decision trees trained with different sets of properties and measuring the differences between prediction errors with and without each property. It suggests that indeed minimum curvature and Gauss curvature are most useful properties for predicting Schelling points. It also suggests that the *added value* of AGD, HKS, mean curvature, and other common properties is limited, since they are redundant with minimum and Gauss curvatures. On the other hand, the importance of global properties, such as symmetry, SDF, and segment centeredness, is relatively high (84, 54, and 33, respectively), which suggests they provide cues that are independent of other properties.

We note that percentiles (columns 4 and 6) appear more useful for predicting Schelling points than raw geometric property values: e.g., vertices having curvature higher than others on the same surface are more likely to be Schelling points than ones within any specific range of values.

Coverage plots: Figure 6 shows stack plots representing the fraction of Schelling points in S “explained” by extrema of different mesh properties. To produce this result, we extracted the K largest local maxima of each property ($K = 14$, matching the average size of a point set collected in our study), with maxima spread by at least geodesic distance $0.02R$, using the algorithm described in Section 7. Then, for each Schelling point, $s_i \in S$, we considered the surface properties in the order from bottom to top, marking the property with a vote if it is the first to have one of its K largest local maxima within $0.05R$ of s_i . After all votes are cast (one for each Schelling point on each mesh), we average the votes per mesh, and then average the fractions per class (shown in the first 19 bars in Figure 6), and finally average those fractions to get a result for the entire data set (shown in the bar labeled “Overall” on the far right).

These results suggest that extrema of minimum curvature “explain” 56.7% of the Schelling points (bottom blue region on the rightmost bar). Of the remaining Schelling points, 6.7% are explained by Gauss curvature (red); and 9.1% of the ones still remaining are explained by symmetry (IRSA); and so on. Please note that there is variation in the results for different object classes:

Property (p)	Best Filter (f)	Info Gain		Importance	
		G(f)	G(%(f))	I(f)	I(%(f))
MinCurv	$ B(p, 1) $	195	203	216	801
GaussCurv	$ p $	201	211	59	267
Symmetry	$B(p, 4)$	12	25	18	84
SDF	p	17	26	31	54
SegCenter	$B(p, 3)$	24	43	15	33
HKS(101)	p	104	141	15	31
Saliency(0.3)	p	48	85	26	30
MaxCurv	$B(p , 4)$	53	99	19	30
ZPosition	p	20	51	29	27
ZNormal	p	11	14	24	25
MeanCurv	$B(p, 4)$	43	98	11	23
CurvDiff	p	29	46	22	21
AGD	p	25	85	18	18

Table 1: Information gain ($\times 1000$) and random forest importance of mesh properties (rows) for predicting the Schelling point distribution. p represents the property value, f is the best filter found for p , $B(\cdot, \sigma)$ is Gaussian blur, $\%(f)$ is its percentile within a mesh, $G(\cdot)$ is information gain, and $I(\cdot)$ is an estimate of importance by analysis of Random Forests.

curvature measures are almost completely sufficient for explaining Schelling points on airplanes, for example; but symmetry and segment centeredness are more helpful properties for explaining Schelling points on teddy bears.

Interestingly, approximately 11% of Schelling points are not described by the strongest K local maxima of any property considered. Examining those points visually, we find that they tend to reside either at semantic features off-center on smooth surfaces (e.g., the eyes on the teddy bears) and/or at geometric features that did not appear among the strongest maxima (e.g., a knee or elbow when the appendage is not bent). Perhaps other geometric properties could be discovered to explain these remaining points in future work.




Figure 6: Stackplot showing fractions of Schelling points first “explained” by extrema of each property considered in the order shown bottom to top.

6 Prediction of Schelling Distributions

In this section, we investigate whether it is possible to predict mesh saliency using the data collected in our study. That is, for a mesh M never seen before, we aim to provide a method that estimates the probability that a person would have selected each vertex V of the mesh if it were included in our study – i.e., the predicted Schelling point distribution, $\hat{S}_D(V)$.

Transfer: A simple method to approach this problem would be to transfer the Schelling points from other similar meshes for which AMT data has been collected. That is, given a new mesh M , we could select a set of K similar meshes M_i from the collected Schelling point data set, compute a map $m_i(M \rightarrow M_i)$ from each vertex V_j of M to the corresponding vertex $m_i(V_j)$ in each similar mesh M_i , and then predict the transferred Schelling point distribution $\hat{S}_T(V)$ for vertices of M by aggregating the Schelling points mapped from other meshes: $\hat{S}_T(V) = 1/K \sum_i^K S_P(m_i(V))$. This method is very simple. However, it will work well only when an accurate map can be computed, and when two meshes have the same number of Schelling points in the same exact arrangement, neither of which is often the case.

Regression: A second method is to build a regression model based on geometric properties of a surface. Given a “training” set of meshes M_i with collected Schelling points, we compute the geometric properties listed in Section 5 (Gauss curvature, etc.) for every vertex of every mesh M_i in a training set and then learn a model that relates those properties to the Schelling point indicator function $S_P(V)$. This model is simple to implement, since it considers each vertex of the mesh independently, and it can be applied to arbitrary new surfaces.

Our implementation for building the regression model is based on MSP regression trees as provided by Weka [Witten and Frank 2005]). This model builds a decision tree that takes in a training set of vertices labeled with the properties described in the previous section along with the observed value of the Schelling point distribution at that vertex. It analyzes this data and builds a binary tree that splits feature space into distinct regions and then fits a linear regression model for $\hat{S}(V)$ for each region independently. It was chosen for our study because it can fit non-linear relationships between input and output variables (piecewise linear), and it provides an explanation for how the model operates (the decision tree), which is much easier to decipher than most other regression models.

As an example, Figure 7 shows the top nodes of a MSP regression tree learned by our system when trained on the sampled data for teddy bears. Examining the tree, it is interesting to note that the top few levels of splits take into account different properties: minimum curvature, symmetry, segmentation centerness, shape diameter, saliency, etc. Also, different derived properties are utilized, with percentiles more often chosen for splits in the decision tree, and values or derivatives more often chosen for linear equations in the leaf nodes. Overall, this tree suggests that several types of properties can be combined to make better predictions than any one alone. Of course, nothing in our study is dependent on this particular choice of regression model, and we believe that several other alternatives could have been used just as effectively.

Results: To evaluate and compare how effectively the transfer and regression methods work, we performed a series of leave-one-out experiments where we utilized the Schelling points collected from people on all but one mesh to predict the Schelling point distribution on the one held out, and then we evaluated the quality of the prediction by comparison to the Schelling point data collected for that mesh. Evaluation was performed by computing the Information Gain between the predicted and actual point set distributions, as in Table 1.

For the first tests, we used vertex-to-vertex maps to transfer Schelling points to the held out mesh M from other meshes in the same object class. Maps were constructed with two different methods: 1) manually by interpolating human-selected landmark corre-

```

% (Blur (MinCurv, 4) <= 65.56
| Symmetry > 35.669
| | SDF > 0.528
| | | Symmetry <= 101.366
| | | | % (GeodesicTenPercentile) <= 44.399
| | | | Symmetry > 101.366
| | | | ZNormal <= 0.287
% (Blur (MinCurv4) > 65.56
| SegCenter <= 0.772
| | % (MaxCurv) <= 37.769
| | | % (Saliency0.7) <= 89.877
| | | | % (|Blur (MaxCurve)|, 3) <= 33.576
| | | | % (Saliency0.7) > 89.877
| | | | % (Blur (|MaxCurv|), 3) <= 77.394
| | | % (MaxCurv) > 37.769
| | | % (SDF) <= 43.417
| | | | Blur (Symmetry, 4) <= 59.815
| | SegCenter > 0.772
| | MeanCurv > 5.973
| | | % (ZNormal) <= 92.829
| | | | Blur (Symmetry, 4) <= 63.682

```

Figure 7: Top nodes of an MSP regression tree learned from Schelling points on teddy bears.

spondences (TrueMap),¹ and 2) automatically using Blended Intrinsic Maps [Kim et al. 2011] (BlendedMap). In both cases, we considered cases where points are transferred from just the most similar mesh as computed with a geodesic D2 shap descriptor (Closest), and where they are transferred and aggregated from all meshes in the same object class (InClass). Results of these tests, averaged over all 380 meshes, are shown in the top four rows of Table 2. They suggest that transferring points does not work well, in general. For many classes of objects (e.g., fish, vase, chair, airplane, bird, etc.), different meshes have different arrangements of parts, accounting for the failures of even TrueMap. For other classes (e.g., ant, hand, etc.), computing semantically correct maps between meshes is difficult, accounting for the worse performance of BlendedMaps.

For the second tests, we used regression trees learned from Schelling points on a subset of meshes to predict $\hat{S}_D(V)$ for others. We executed two experiments that differ in how the training sets were chosen. In the first experiment (InClass), regression trees were built in leave-one-out style using training data from meshes of the same class – i.e., the model was trained on 19 out of the 20 meshes in each object class and then tested on the 20th (as in [Kalogerakis et al. 2010]). In the second experiment (OutClass), the regression tree for each mesh was trained only using data from meshes of other object classes. Results of this study are shown in Table 2. They suggest that regression significantly outperforms transfer algorithms on this data set, even when training is performed only on examples from different object classes. They also show that the regression model combining many surface properties outperforms any single property – i.e., the highest Information Gain achieved for any single property (GaussCurvature) is 0.211 (second row of Table 1), in comparison to 0.257 (OutClass) and 0.341 (InClass) achieved with regression.

Visualizations of Schelling point distributions predicted with a regression model trained InClass are shown in Figure 8. The meshes and color scheme are the same as in Figure 1 to facilitate direct visual comparison. From these images (and others provided in supplemental materials), we see that these learned models are not as precise as the actual Schelling distributions. However, they predict tips of protrusions well (corners in the airplane) as well as some subtle features (e.g., centers of eyeglasses, eyes of the bust, belly button of the teddy bear, etc.) that are difficult to recognize with a threshold on any single surface property.

¹TrueMaps are considered only for didactic purposes. Since they are entered manually, they could not be used in an automatic prediction system.

Prediction Method (p)	Source Data	Best Filter (f)	Info Gain
PredictMap	Closest	$\%(Blur(p, 2))$	70
	InClass	$\%(Blur(p, 2))$	105
TrueMap	Closest	$\%(Blur(p, 2))$	160
	InClass	$\%(Blur(p, 2))$	236
Regression	OutClass	$Blur(p, 1)$	257
	InClass	$Blur(p, 1)$	341

Table 2: Information gain ($\times 1000$) of Schelling point predictions.




Figure 8: Visualizations of Schelling point distributions predicted by our algorithm after training on properties of different meshes in the same object class (InClass).

7 Prediction of Schelling Points

For many applications, it is important to not only estimate mesh saliency, but also to extract a discrete set of salient feature points. For example, surface matching algorithms often start by detecting a set of features and then searching for correspondences between them. To be successful, these algorithms require a set of feature points that have many of the properties of Schelling points: stability, spacing, symmetry, etc.

Many algorithms are possible to address this problem, including ones that perform a combinatorial optimization over possible candidate sets, choosing the one that maximizes some objective function. However, for the purposes of this study, we choose a simple greedy algorithm that has been used in several other saliency experiments (e.g., [Shilane and Funkhouser 2007]). The algorithm aims to select a set of vertices $\{V_i\}$ that maximizes the sum of $\hat{S}(V_i)$ at selected vertices, subject to the constraint that no two vertices in $\{V_i\}$ are too close to one another. It does so by sorting vertices from highest $\hat{S}(V)$ to lowest and then repeatedly selecting the next vertex remaining whose position is not closer than a normalized geodesic distance threshold, D , to any previously selected vertex ($D = 0.05$). This process avoids selecting many points near one another on the mesh and provides an easy way to reduce the point set size by increasing the distance threshold D or terminating the greedy search after a given number has been selected.

Results: To investigate how effectively an algorithm can predict Schelling points, we perform an analysis of the similarities of the automatically generated point sets to the ones collected from people. For this analysis, we extract $k = 14$ points for each mesh (the even number closest to the median of point sets collected from

people) and then follow the general methodology described in Section 5A. That is, we measure the distance from points in our automatically extracted point sets to the Schelling points S_P on the same mesh and plot the cumulative distribution with respect to increasing normalized distance thresholds.

For comparison sake, we consider point sets extracted with the same algorithm from maxima of every surface property considered in Section 5, and we consider point sets transferred from Schelling points on other meshes within the same class using the mapping algorithms described in Section 6.

Figure 9 shows the results. As in Figure 3a, the horizontal axis contains increasing distance thresholds, D , and the vertical axis shows consistency using the cumulative distribution of point pairs whose normalized geodesic distance is less than D . The red curve on top represents the consistency between Schelling points S_P collected from people with point sets extracted automatically from the distribution predicted with regression using InClass training. The lower curves represent the best point sets generated with other methods. These results suggest that the points selected by our algorithm from the distribution learned with regression using InClass training are closer to the Schelling points collected in our study than those predicted with other methods. We also find that they are slightly more consistent with the Schelling points than the average point set collected from people, suggesting that the regression and point extraction algorithms for predicting Schelling point sets are effectively relating surface properties to Schelling point locations.

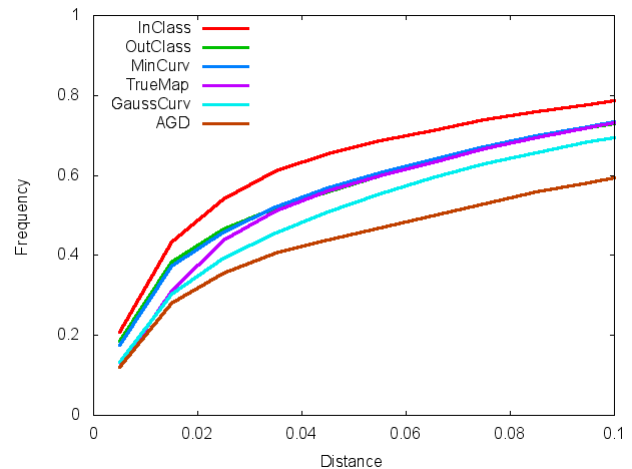


Figure 9: Consistency between Schelling points and point sets extracted with different algorithms.

8 Validation with Controlled User Study

As a final test, we investigated the impact of collecting data via the Amazon Mechanical Turk (AMT) on the main conclusions of our study. Although the AMT is becoming a common platform for perceptual experiments in computer graphics (e.g., [Cole et al. 2009]) and there are several studies validating AMT studies in the laboratory (e.g., [Heer and Bostock 2010]), it is valuable to ask whether the data we gathered from the AMT is representative of that which would be collected from a controlled group of participants in a laboratory environment.

Of course, it is not practical to duplicate the AMT experiment exactly in the laboratory, since it would take hundreds (or possibly thousands) of hours to collect an equivalent amount of data. In-

stead, we ran a small study in which we recruited 30 volunteers through email inquiries to acquaintances and students unfamiliar with our project. Of these volunteers, 20 were male and 10 were female. They span a range of ages: 6 were ≤ 20 years old, 16 (21-30), 2 (31-40), 2 (41-50), 3 (51-60), and 1 (>60). All except four self-evaluate as having at least “plenty” of experience with computers, but only 9 report having “plenty” of experience with graphics. Every participant performed the study on a Windows computer with a 3-button mouse.

We asked each of the participants to select points on 19 meshes (one selected randomly from each object category) with the exact same instructions and user interface provided on the AMT (substituting the incentive to gain “points” for selections consistent with others’ rather than higher monetary payment). After completion, participants completed an exit survey, in which 25 responded “yes” or “mostly” to a question asking whether they completed the task to their satisfaction. We rejected data from the other users, leaving 474 point sets from 25 participants. Of those, 354 (74.7%) passed all three data filters – 66 (13.9%) had too few points, 23 (4.9%) were entered with too little camera motion, and 31 (6.5%) were clicked too hastily. This is a significantly lower rejection rate than for the data collected on the AMT.

Analyzing the point sets passing all filters, we find that they closely matches the characteristics of data collected on the AMT. In particular, the distribution of distances indicating the consistency of points selected by different people on the same mesh (Figure 10a) and the reflective symmetry of points selected by the same user on the same mesh are almost identical to those collected on the AMT for the same set of meshes (Figure 10b). Of course, there is not enough data to study consistency across different meshes in the same class, or to study variations within specific subsets of objects. However, all statistics computed for this small, controlled data set closely match those of the larger AMT data set.




Figure 10: Comparison of consistency and symmetry of point sets collected in a small, controlled study (left) versus ones collected from people on the Amazon Mechanical Turk (right). (a) The top row shows the consistency of points selected by different people on the same mesh (like Figure 3a). (b) The bottom row shows the consistency of points selected by the same person on symmetric sides of a mesh (like Figure 3c).




Figure 11: The predicted Schelling distribution \hat{S}_D (left) can be used to preserve salient detail during mesh simplification.

9 Conclusion

In this paper, we have described a study of Schelling points on 3D meshes – i.e., points that people expect to be selected by other people. With an on-line experiment, we gathered 9,965 points sets containing a total of 201,304 points. During both qualitative and quantitative analysis, we find that these points appear mainly at “semantically stable” positions on a 3D mesh – e.g., extrema of curvature, on axes of symmetry, centers of segments, etc. We also find that they are selected fairly consistently by different people on the same mesh, and slightly less so on different meshes of the same object class. However, they are selected very consistently on symmetric parts within the same mesh. We have used this data to train an algorithm for predicting the distribution of Schelling point on new meshes and used it in algorithms for feature point detection.

The saliency measures produced with our methods could be used in a variety of applications in computer graphics (e.g., [Alexa 2000; Funkhouser and Shilane 2006; Johnson 2000; Katz et al. 2005; Lee et al. 2005; Zhang et al. 2005]). For example, Figure 11 shows how a predicted Schelling distribution can guide a mesh simplification algorithm to preserve features with higher expected semantic salience (as in [Lee et al. 2005]). In this case, quadric errors used by QSLim [Garland and Heckbert 1997] were scaled by \hat{S}_D^2 learned with regression from InClass examples. Note how detail is better preserved in salient areas.

Our study is just a first step and thus has several limitations that suggest topics for future work. First, it considers only watertight meshes with relatively smooth features, and therefore some of the findings regarding local geometric features may not generalize to polygon-soup models commonly found in computer graphics repositories. Second, it considers only point features: studying line and region features would also be valuable. Third, it studies geometric surface properties of the collected data: future work could study other aspects of the data, including the time-dependent strategy people used to select feature points. Finally, it focuses only on applications in computer graphics (mesh saliency and feature point detection): future work might consider questions in perceptual psychology.

Acknowledgments

We thank Doug DeCarlo for introducing us to the work of Thomas Schelling and Maneesh Singh for useful discussions during the early phases of the project. We also thank Daniela Giorgi and AIM@SHAPE for access to the SHREC 2007 Watertight Models and Evangelos Kalogerakis for use of shape feature extraction code. Finally, we thank NSF (CCF-0937139, CNS-0831374), Intel (ISTC-VC), Adobe, and Google for partial funding of this project.

References

- ALEXA, M. 2000. Merging polyhedral shapes with scattered features. *Visual Computer* 16, 26–37.
- AMAZON, 2009. Mechanical turk. <http://www.mturk.com>.
- ATTNEAVE, F. 1954. Some informational aspects of visual perception. *Psychological Review* 61, 3.
- BREIMAN, L. 2001. Random forests. *Machine Learning* 45, 1, 5–32.
- BRONSTEIN, A., BRONSTEIN, M., AND KIMMEL, R. 2006. Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Science*, 1168–1172.
- BRONSTEIN, A., BRONSTEIN, M., BUSTOS, B., CASTELLANI, U., CRISANI, M., FALCIDIENO, B., GUIBAS, L., KOKKINOS, I., MURINO, V., OVSJANIKOV, M., PATANE, G., SIPIRAN, I., SPAGNUOLO, M., AND SUN, J. 2010. SHREC 2011: robust feature detection and description benchmark. In *Eurographics Workshop on 3D Object Retrieval*.
- CASTELLANI, U., CRISTANI, M., FANTONI, S., AND MURINO, V. 2008. Sparse points matching by combining 3d mesh saliency with statistical descriptors. *Computer Graphics Forum* 27, 2, 643–652.
- CHEN, X., GOLOVINSKIY, A., AND FUNKHOUSER, T. 2009. A benchmark for 3D mesh segmentation. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28, 3 (Aug.).
- CHUA, C., AND JARVIS, R. 1996. Point signatures: A new representation for 3D object recognition. *International Journal of Computer Vision* 25, 1, 63–85.
- COLE, F., GOLOVINSKIY, A., LIMPAECHER, A., BARROS, H. S., FINKELSTEIN, A., FUNKHOUSER, T., AND RUSINKIEWICZ, S. 2008. Where do people draw lines? *ACM Transactions on Graphics (Proc. SIGGRAPH)* 27, 3 (Aug.).
- COLE, F., SANIK, K., DECARLO, D., FINKELSTEIN, A., FUNKHOUSER, T., AND AN D MANISH SINGH, S. R. 2009. How well do line drawings depict shape? *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28, 3 (Aug.).
- FUNKHOUSER, T., AND SHILANE, P. 2006. Partial matching of 3d shapes with priority-driven search. In *Symposium on Geometry Processing*.
- GAL, R., AND COHEN-OR, D. 2006. Salient geometric features for partial shape matching and similarity. *ACM Transaction on Graphics* (January).
- GARLAND, M., AND HECKBERT, P. S. 1997. Surface simplification using quadric error metrics. In *Proceedings of SIGGRAPH 1997, Computer Graphics Proceedings, Annual Conference Series*, 209–216.
- GIORGI, D., BIASOTTI, S., AND PARABOSCHI, L., 2007. SHREC:SHape REtrieval Contest: Watertight models track, <http://watertight.ge.imati.cnr.it/>.
- HEER, J., AND BOSTOCK, M. 2010. Crowdsourcing graphical perception: Using Mechanical Turk to assess visualization design. In *ACM Human Factors in Computing Systems (CHI)*, 203–212.
- HISADA, M., BELYAEV, A., AND KUNII, T. 2002. A skeleton-based approach for detection of perceptually salient features on polygonal surfaces. *Computer Graphics Forum* 21, 4, 689–700.
- HOFFMAN, D. D., AND SINGH, M. 1997. Saliency of visual parts. vol. 63.
- HUANG, T., CHENG, K., AND CHUANG, Y. 2009. A collaborative benchmark for region of interest detection algorithms. 296–303.
- ITTI, L., KOCH, C., AND NEIBUR, E. 1998. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, 11, 1254–1259.
- JOHNSON, A. 2000. Surface landmark selection and matching in natural terrain. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2, 413–420. Using saliency in choosing spin images.
- KALOGERAKIS, E., HERTZMANN, A., AND SINGH, K. 2010. Learning 3d mesh segmentation and labeling. *ACM Transactions on Graphics (proc. SIGGRAPH)* 29, 3.
- KATZ, S., LEIFMAN, G., AND TAL, A. 2005. Mesh segmentation using feature point and core extraction. *Visual Computer* (September).
- KIM, Y., VARSHNEY, A., AND ADN FRANÇOIS GUIMBRETIERE, D. J. 2010. Mesh saliency and human eye fixations. *ACM Transactions on Applied Perception* 7, 2 (February).
- KIM, V., LIPMAN, Y., AND FUNKHOUSER, T. 2011. Blended intrinsic maps. *ACM Transactions on Graphics (SIGGRAPH 2011)* (jul).

- KO, B., AND NAM, J. 2006. Object-of-interest image segmentation based on human attention and semantic region clustering. *J Opt Soc Am A Opt Image Sci Vis* 23, 10 (October), 2462–2470.
- KOCH, C., AND ULLMAN, S. 1985. Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology* 4, 219–227.
- KRAEVOY, V., AND SHEFFER, A. 2004. Cross-parameterization and compatible remeshing of 3d models. *ACM Transactions on Graphics (Proc SIGGRAPH)* 23, 3, 861–869.
- LEE, C. H., VARSHNEY, A., AND JACOBS, D. W. 2005. Mesh saliency. *ACM Transactions on Graphics (SIGGRAPH 2005)* (aug).
- LEWIS, D. 1969. *Convention: A Philosophical Study*. Harvard University Press.
- LI, X., AND GUSKOV, I. 2005. Multi-scale features for approximate alignment of point-based surfaces. In *Symposium on Geometry Processing*.
- LI, X., AND GUSKOV, I. 2007. 3d object recognition from range images using pyramid matching. In *Workshop on 3D Representation for Recognition (3dRR)*.
- LIPMAN, Y., AND FUNKHOUSER, T. 2009. Mobius voting for surface correspondence. *ACM Transactions on Graphics (SIGGRAPH 2009)* (August).
- MILANES, R., WECHSLER, H., GIL, S., BOST, J., AND PUN, T. 1994. Integration of bottom-up and top-down cues for visual attention using non-linear relaxation. *IEEE Computer Vision and Pattern Recognition*, 781–785.
- MOREELS, P., AND PERONA, P. 2007. Evaluation of features detectors and descriptors based on 3d objects. *IJCV* 73, 3 (July), 263–284.
- NOVOTNI, M., DEGENER, P., AND KLEIN, R. 2005. Correspondence generation and matching of 3d shape subparts. Tech. Rep. CG-2005-2, Universität Bonn, June.
- PARKER, P. 2011. *Webster's On-line Dictionary: The Rosetta Edition*. <http://www.websters-online-dictionary.org>.
- PRIVITERA, C., AND STARK, L. 2000. Algorithms for defining visual regions-of-interest: Comparison with eye fixations. *PAMI* 22, 9 (September), 970–982.
- ROSENHOLTZ, R. 1999. A simple saliency model predicts a number of motion popout phenomena. *Vision Research* 39, 19, 3157–3163.
- RUSINKIEWICZ, S. 2004. Estimating curvatures and their derivatives on triangle meshes. In *Symposium on 3D Data Processing, Visualization, and Transmission*.
- SANTELLA, A., AND DECARLO, D. 2004. Robust clustering of eye movement recordings for quantification of visual interest. In *Eye Tracking Research and Applications (ETRA)*, 27–34.
- SCHELLING, T. 1960. *The Strategy of Conflict*. Harvard University Press.
- SCHLATTMANN, M., DEGENER, P., AND KLEIN, R. 2008. Scale space based feature point detection on surfaces. *Journal of WSCG* 16 (February).
- SCHMID, C., MOHR, R., AND BAUCKHAGE, C. 2000. Evaluation of interest point detectors. *IJCV* 37, 2 (June), 151–172.
- SEBE, N., AND LEW, M. 2003. Comparing salient point detectors. *Pattern Recognition Letters* 24, 1-3 (January), 89–96.
- SHAPIRA, L., SHAMIR, A., AND COHEN-OR, D. 2008. Consistent mesh partitioning and skeletonisation using the shape diameter function. *Vis. Comput.* 24, 4, 249–259.
- SHILANE, P., AND FUNKHOUSER, T. 2007. Distinctive regions of 3d surfaces. *ACM Transactions on Graphics* 26, 2 (June).
- SIMPSON, J. 1989. *Oxford English Dictionary, Second Edition*. Oxford University Press. <http://dictionary.oed.com>.
- SONTHI, R., KUNJUR, G., AND GADH, R. 1997. Shape feature determination using the curvature region representation. In *Proc. Solid Modeling*, ACM.
- STARK, M., AND SCHIELE, B. 2007. How good are local features for classes of geometric objects. 1–8.
- SUMNER, R., AND POPOVIC, J. 2004. Deformation transfer for triangle meshes. *ACM Transactions on Graphics (Proc SIGGRAPH)* 23, 3, 399–405.
- SUN, J., OVSJANIKOV, M., AND GUIBAS, L. 2009. A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. In *Computer Graphics Forum*, vol. 28, Wiley Online Library, 1383–1392.
- TSOTSOS, J., CULHANE, S., WAI, W., LAI, Y., DAVIS, N., AND NUFLO, F. 1995. Modeling visual-attention via selective tuning. *Artificial Intelligence* 78, 1-2, 507–545.
- VAN KAICK, O., ZHANG, H., HAMARNEH, G., AND COHEN-OR, D. 2010. A survey on shape correspondence. In *Proc. Eurographics State-of-the-art Report*.
- VON AHN, L., AND DABBISH, L. 2008. Designing games with a purpose. *Communications of the ACM* 51, 8, 58–67.
- WITTEN, I. H., AND FRANK, E. 2005. *Data mining: Practical machine learning tools and techniques*, 2nd edition.
- XU, K., ZHANG, H., TAGLIASACCHI, A., LIU, L., LI, G., MENG, M., AND XIONG, Y. 2009. Partial intrinsic reflectional symmetry of 3d shapes. *ACM Transactions on Graphics, (Proceedings SIGGRAPH Asia 2009)* 28, 5, to appear.
- ZAHARESCU, A., BOYER, E., VARANASI, K., AND HORAUD, R. 2009. Surface feature detection and description with applications to mesh matching. In *CVPR*.
- ZHANG, E., MISCHAIKOW, K., AND TURK, G. 2005. Feature-based surface parameterization and texture mapping. *ACM Transactions on Graphics* 24, 1.
- ZHANG, H., SHEFFER, A., COHEN-OR, D., ZHOU, Q., VAN KAICK, O., AND TAGLIASACCHI, A. 2008. Deformation-driven shape correspondence. *Comput. Graph. Forum* 27, 5, 1431–1439.
- ZHOU, Y., AND HUANG, Z. 2004. Decomposing polygon meshes by means of critical points. In *MMM*, 187–195.
- ZULIANI, M., KENNEY, C., AND MANJUNATH, B. 2004. A mathematical comparison of point detectors. In *Computer Vision and Pattern Recognition Workshop*, 172.