

Schelling Meshes

Luther Power¹ and Manfred Lau²

¹Lancaster University, UK ²City University of Hong Kong




Figure 1: The shapes from the categories of chairs and tables are ranked according to their Schelling frequencies.

Abstract

The concept of “Schelling points” on 3D shapes has been explored for points on the surface of a 3D mesh. In this paper, we introduce the notion of “Schelling meshes” which extends the Schelling concept to 3D meshes as a whole themselves. We collect Schelling-based data for meshes where participants are given a group of shapes and asked to choose those with the aim of matching with what they expect others to choose. We analyze the data by computing the Schelling frequency of each shape and characterizing the qualitative features that make a shape “Schelling”. We show that the Schelling frequencies can be learned and demonstrate Schelling-guided shape applications.

CCS Concepts

• **Computing methodologies** → Perception; Mesh models;

1. Introduction

“Schelling points” are choices that are selected by people when they choose to match others’ selections with no communication beforehand [Sch81]. The notion of Schelling points on 3D meshes have been studied [CSPF12], where Schelling points are points on the surfaces of meshes that people expect will be selected by others. In this paper, we extend this notion from points on the surfaces of 3D meshes to the meshes themselves. Instead of selecting among points on a mesh, people will select meshes among a set of meshes. We use the term “Schelling meshes” to describe the meshes that are selected this way. We believe that the notion of Schelling meshes can be another tool for 3D shape analysis.

This paper explores the ideas introduced in the unpublished poster in this topic [PL17]. The contributions of this paper are:

- We introduce the Schelling concept for meshes as a whole.
- We collect data for the notion of Schelling meshes by applying the Schelling concept of asking people to choose answers that

they expect others will choose. For the case of 3D meshes, we provide participants with groups of 3D shapes and ask them to choose any number of shapes from each group with the goal of matching the selections made by other participants. We define the “Schelling frequency” of a shape to be how often it is chosen in a Schelling sense according to collected data.

- We explore the characteristics (i.e. qualitative features and quantitative shape descriptors) that make a shape more “Schelling”.
- We show that the notion of Schelling meshes can be learned and learn a function to predict Schelling scores for new meshes.
- We demonstrate applications with the Schelling-guided visualization and Schelling-guided search of 3D shapes.

2. Related Work

Mesh Saliency. There has been much work in computing the visual saliency of meshes, originating with “Mesh Saliency” [LVJ05]. A recent survey [LLS*16] describes visual mesh saliency methods

and their applications. In addition to visual saliency, a recent work introduces the idea of computing tactile salient information for a mesh [LDS*16]. In this paper, we also compute a kind of saliency in the Schelling aspect of meshes. The work that introduces the idea of computing Schelling points on mesh surfaces [CSPF12] has inspired our work. In contrast to considering the Schelling concept of mesh points, we consider the Schelling concept of the meshes as a whole themselves.

Crowdsourcing. There has been work in applying crowdsourcing to collect data from humans to solve various graphics problems. This approach has been used to compute a style similarity function for clip art [GAGH14], fonts [OLAH14], and 3D models [LHLF15, LKS15]. In this paper, we use crowdsourcing to collect data on how humans select Schelling meshes.

Shape Analysis. The analysis of 3D models is a well studied area and there has been work in many different problems [XKHK17]. For example, there is work in computing 3D shape distributions [OFCD01] and shape diameter functions [GSCO07]. This paper explores various aspects of the “Schelling Meshes” problem.

3. Collecting Schelling Meshes Data

In this section, we describe the processes of collecting data to study Schelling meshes. We collected 169 3D models from online sources (e.g. ShapeNet [CFG*15]). These include everyday objects: chairs, tables, lamps, and a variety of abstract shapes. To display a 3D shape on a 2D screen, we generated a repeatedly rotating view of each shape: each full rotation takes three seconds followed by a pause of half a second. We choose this representation rather than multiple images as we believe this makes it easier to visualize a 3D shape as a whole.

The Schelling concept as applied to 3D shapes requires humans to analyze a shape *relative to other shapes*. For the case of finding Schelling points on mesh surfaces [CSPF12], the participants selected some number of points on a mesh surface given many possible points to select from. Analogous to this but for meshes as a whole, we show participants many possible shapes to select from and ask them to select some number of them. Each participant makes this selection based on the Schelling concept: to try to match other participants’ responses based on what he/she thinks others will choose. We decided to let the participants choose from each group of shapes that belong to the same shape category (e.g. chairs). It is possible to mix the categories but separating them is already interesting for us to study the concept of Schelling meshes. For each group of shapes, the order of the shapes is randomized each time it appears as a question for a participant.

We use a crowdsourcing platform (Amazon Mechanical Turk) to collect data. Each question shows all of our shapes for a shape category. Each shape has a selection box for the participant to indicate choosing it or not, and each selection box is independent from the others. Participants were first given written instructions: “For each question, your task is to choose from a selection of shapes. Other participants will be given the same task. You should choose shapes that will most likely match with their selections. Note that you will not be able to communicate with other participants, and this is intentional.” They were also told to choose at least one shape per question. Each HIT (Human Intelligence Task or set of questions on Mechanical Turk) has four questions, one for each of our




Figure 2: The whole set of data is randomly sampled 10 times, each time with half of the participants. The y-axis includes each shape (from top to bottom) from the set of abstract shapes, chairs, lamps, and tables.

shape categories. We have 49 shapes of chairs, 33 tables, 49 lamps, and 38 “abstract” shapes. At the end of each HIT, we include an optional text box and asked participants to provide “a few words describing why you selected the shapes that you did.”

Since the user selections are subjective and there are no right or wrong answers, we decided to not have any control questions to filter out potentially bad users. In the instructions, we tried to encourage users to carefully work on the questions by specifying to users that: “If you randomly choose your answers, your HIT responses will not be taken, and you will not be paid.” Furthermore, the users are only allowed to work on our HITs if their acceptance rate of previous HITs they have done is at least 80%. A participant takes about 1 to 5 minutes for each HIT and we paid \$0.10 for each HIT. We collected data for 102 participants.

Checking for Data Consistency. We wish to see whether the collected data is consistent. Since the data is subjective, there is no right or wrong answer to compare against. Hence we check the consistency within the collected data. The main idea is to split the whole set of data into different groups and check whether the groups have similar distributions.

For each of the 169 shapes, the whole set of data consists of whether each of the 102 participants selected it. We randomly sample from this 10 times, where each time we randomly pick half (or 51) of the participant responses. As half of the responses still give us information about all shapes, this gives 10 vectors of 169 values. Figure 2 shows a visual representation. We can see that there is much correspondence in the horizontal rows, where some rows are mostly blue and light blue and some rows are mostly yellow and orange. This means that across the 10 vectors, the distributions of the values are similar. Quantitatively, we perform a two-sample Kolmogorov-Smirnov test for each pair (from 10) of 169 values, and find that the p-value is not less than 0.05 in each case. This provides evidence that these 10 vectors come from the same distribution and that there is consistency in the whole set of data.

4. Results and Analysis

Schelling Frequencies. We define the concept of *Schelling frequency* of a shape to be how likely it will be selected in a Schelling sense. Although the Schelling concept is a relative concept, we compute the Schelling frequency for each shape in order to give a score for every shape. For each shape, the Schelling frequency is the number of participants who selected it divided by the total




Figure 3: The shapes from the categories of lamps and abstract shapes are ranked according to their Schelling frequencies.

number of participants. Figures 1 and 3 show plots of all shapes we used in each shape category on a 1D line of Schelling frequency.

We describe the patterns that we observe in the plots. For chairs, the shapes that stand out more or look more unique have higher Schelling frequencies. In contrast, the more normal-looking chairs are all clustered into one big group near the left side of the 1D line. For tables and lamps, similarly, the shapes that stand out have higher Schelling frequencies, while the more normal-looking shapes are mostly clustered on the left side. For abstract shapes, the shapes with rings or holes and the statue shapes have higher Schelling frequencies, while those with lower Schelling frequencies tend to be blobs or just some unknown shapes.

Participant Descriptions of Selected Meshes. We try to understand the characteristics of Schelling Meshes from the user text descriptions. 46 total participants gave comments. 18 participants mentioned they made selections based on *appeal*, *aesthetics*, or *beauty*. One example user comment is: “I basically selected items that I liked and items which I thought other people would like as well.” 16 users said they selected shapes that *stand out*, *are different*, or *catchy*. For example, one user commented: “I selected the shapes that are unusual and different from others.” A few participants said they chose *familiar* shapes. For example, one user commented: “I hope that most of the other participants have also gone for similar designs as they are common and easy to remember.”

Correlation with Qualitative Features. For shapes that tend to be Schelling frequent, some characteristics that we observe that were also mentioned by participants are that they “stand out” from the other shapes, are unique, or are visually appealing. We collect data to test whether these features are related to being “Schelling.” We use the same setup on Amazon Mechanical Turk as described in Section 3, but with different participants. For each of the above three features, we ask users to provide a Likert score on a 1-5 scale. We collected data for 15 users and paid \$0.10 per HIT. Table 1 shows the results of correlating between the scores for each feature and the Schelling frequencies for the shapes in each category.

As expected, there are positive and significant correlations in most cases. The abstracts category has a small negative correlation for “stand-out”, as shapes that were selected more often in a Schelling sense are more normal for the abstracts category.

Correlations between	Schelling Frequency			
	Chairs	Tables	Lamps	Abstracts
Stand-out	0.584	0.635	0.302	-0.103
Unique	0.439	0.634	0.242	0.130
Visually Appealing	0.151	0.370	0.499	0.601

Table 1: Correlations between some qualitative features and Schelling frequencies. The values are Pearson correlation coefficients and bolded values indicate that the corresponding p-value is less than 0.05 which means the correlation is significant.

Comparison with 3D Shape Descriptors. We compute for each shape the histogram of some common 3D shape descriptors: D2 shape distribution [OFCD01], gaussian curvature and mean curvature [SMS*03], and shape diameter function [GSCO07]. We plot each of these versus the shapes in Schelling order, and we visually observe no clear correlation between each descriptor and Schelling frequency.

5. Learning and Predicting Schelling Scores

We learn a function to predict Schelling scores for new shapes, such that we do not need to collect Schelling-based data in general. We use the term “Schelling score” to denote a predicted Schelling value from the learned function, as opposed to “Schelling frequency” which is computed directly from the collected data. The function is a multi-layer neural network, with a 3D shape as input (or multiple views of depth images of the shape) and its Schelling score as output. We perform data augmentation by having 100 training samples per shape, by randomly picking the multiple depth images for each sample. For each shape category, we take all samples and perform 10-fold cross-validation to predict a Schelling score for each sample. We then correlate between the set of Schelling scores and

Schelling frequencies. The correlation coefficient for chairs is 0.82, tables is 0.77, lamps is 0.71, and abstract shapes is 0.91. These high values show that the notion of Schelling meshes can be learned.

Prediction with 3D Shape Descriptors. We also learn a function to take as input 3D shape descriptors and compute as output the shape’s Schelling score. The shape descriptors are the same as above: D2 shape distribution, gaussian curvature and mean curvature, and shape diameter function. Each of these gives us a histogram and they are concatenated into a single vector. The function is a multi-layer neural network and each shape provides only one training sample. We perform 10-fold cross validation as above. The correlation coefficients are small negative values for all four shape categories. As we do not intend that a simple set of geometric descriptors can predict Schelling scores, the results are as expected and show that the concept of Schelling meshes is complex.

6. Applications

We demonstrate the potential uses of the concept of Schelling meshes in some Schelling-based applications.

Schelling-based Visualization. The Schelling concept can be used to visualize a set of shapes. For example, the plots in Figures 1 and 3 show two major clusters in the furniture shapes, with shapes that stand out placed in the middle or right side of the spectrum and shapes that are more common placed in a large group on the left side. Note that it would otherwise be difficult to place the shapes this way, as for example common shape descriptors such as shape distribution or curvature would not have the same effect. An example setting where this kind of visualization would be useful is for outlier or anomaly detection.

Schelling-based Shape Search. The Schelling concept can be used for search and retrieval applications of 3D model datasets. The idea is to use the Schelling frequency as a distance metric such that the distance between two shapes is the difference between their Schelling frequencies. Figure 4 shows an example with a query chair that is a bit unusual. The first row shows the top-5 search results based on Schelling frequency and they are chairs that may be different in shape but are similar in their unusualness. The search results with the other 3D shape descriptors are very different, and mostly return chairs that are more similar in shape and more normal than the query.

7. Discussion

In this paper, we introduce the notion of Schelling meshes where the “Schelling points” are the “meshes” themselves and we study various aspects of this problem. We hope that this paper will inspire more work and more applications of this notion.

One limitation in the data collection is that we currently have four shape categories. Future work can include more categories and more shapes in each category. We can also mix the shapes in all the categories, although the abstract shapes are somewhat mixed already. Mixing the categories may make it more complex to study the characteristics of Schelling meshes, but it can be more general.

Furthermore, the context of the set of shapes in each question is important. The number of shapes and the variety of shapes within a set that users pick from may affect the computed Schelling frequencies. It would be a good future direction to study this dependency.




Figure 4: Schelling-based Shape Search. An example of searching with a query shape (shown on left). The rows correspond to searching with: Schelling frequency, histograms of gaussian curvature, mean curvature, D2 distribution, and shape diameter function. In each case, the top-5 closest shapes are shown.

References

- [CFG*15] CHANG A. X., FUNKHOUSER T. A., GUIBAS L. J., HAN-RAHAN P., HUANG Q., LI Z., SAVARESE S., SAVVA M., SONG S., SU H., XIAO J., YI L., YU F.: ShapeNet: An information-rich 3D model repository. *CoRR* (2015). 2
- [CSPF12] CHEN X., SAPAROV A., PANG B., FUNKHOUSER T.: Schelling points on 3D surface meshes. *ACM Trans. Graph.* 31, 4 (2012), 29. 1, 2
- [GAGH14] GARCES E., AGARWALA A., GUTIERREZ D., HERTZMANN A.: A Similarity Measure for Illustration Style. *ACM Trans. Graph.* 33, 4 (July 2014), 93:1–93:9. 2
- [GSCO07] GAL R., SHAMIR A., COHEN-OR D.: Pose-oblivious shape signature. *IEEE TVCG* 13, 2 (Mar. 2007), 261–271. 2, 3
- [LDS*16] LAU M., DEV K., SHI W., DORSEY J., RUSHMEIER H.: Tactile Mesh Saliency. *ACM Trans. Graph.* 35, 4 (2016), 52:1–52:11. 2
- [LHLF15] LIU T., HERTZMANN A., LI W., FUNKHOUSER T.: Style Compatibility for 3D Furniture Models. *ACM Trans. Graph.* 34, 4 (July 2015), 85:1–85:9. 2
- [LKS15] LUN Z., KALOGERAKIS E., SHEFFER A.: Elements of style: Learning perceptual shape style similarity. *ACM Trans. Graph.* 34, 4 (July 2015), 84:1–84:14. 2
- [LLS*16] LIU X., LIU L., SONG W., LIU Y., MA L.: Shape context based mesh saliency detection and its applications: A survey. *Computers & Graphics* 57 (2016), 12–30. 1
- [LVJ05] LEE C. H., VARSHNEY A., JACOBS D. W.: Mesh Saliency. *ACM Trans. Graph.* 24, 3 (2005), 659–666. 1
- [OFCD01] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Matching 3D Models with Shape Distributions. *SMI* (2001), 154–166. 2, 3
- [OLAH14] O’DONOVAN P., LIBEKS J., AGARWALA A., HERTZMANN A.: Exploratory Font Selection Using Crowdsourced Attributes. *ACM Trans. Graph.* 33, 4 (July 2014), 92:1–92:9. 2
- [PL17] POWER L., LAU M.: Schelling Meshes. *ACM Symposium on Applied Perception, Poster* (2017). 1
- [Sch81] SCHELLING T. C.: *The Strategy of Conflict*. Harvard University Press, 1981. 1
- [SMS*03] SURAZHSKY T., MAGID E., SOLDEA O., ELBER G., RIVLIN E.: A comparison of gaussian and mean curvatures estimation methods on triangular meshes. *ICRA* (2003). 3
- [XKHK17] XU K., KIM V. G., HUANG Q., KALOGERAKIS E.: Data-driven shape analysis and processing. *Computer Graphics Forum* 36, 1 (Jan. 2017), 101–132. 2